

13th International Conference on Learning Representations

(ICLR 2025)

24-28 April 2025, Singapore

ISBN: 979-8-3313-2085-0

Title and Authors	Page#
<i>DataEnvGym: Data Generation Agents in Teacher Environments with Student Feedback</i> <i>Zaid Khan, Elias Stengel-Eskin, Jaemin Cho, Mohit Bansal</i>	1
<i>Dynamic Neural Fortresses: An Adaptive Shield for Model Extraction Defense</i> <i>Siyu Luan, Zhenyi Wang, Li Shen, Zonghua Gu, Chao Wu, Dacheng Tao</i>	29
<i>VILA-U: A Unified Foundation Model Integrating Visual Understanding and Generation</i> <i>Yecheng Wu, Zhuoyang Zhang, Junyu Chen, Haotian Tang, Dacheng Li, Yunhao Fang, Ligeng Zhu, Enze Xie, Hongxu Yin, Li Yi, Song Han, Yao Lu</i>	52
<i>Port-Hamiltonian Architectural Bias for Long-Range Propagation in Deep Graph Networks</i> <i>Simon Heilig, Alessio Gravina, Alessandro Trenta, Claudio Gallicchio, Davide Bacciu</i>	71
<i>Adaptive Gradient Clipping for Robust Federated Learning</i> <i>Youssef Allouah, Rachid Guerraoui, Nirupam Gupta, Ahmed Jellouli, Geovani Rizk, John Stephan</i>	103
<i>Holistically Evaluating the Environmental Impact of Creating Language Models</i> <i>Jacob Morrison, Clara Na, Jared Fernandez, Tim Dettmers, Emma Strubell, Jesse Dodge</i>	148
<i>Synthetic Continued Pretraining</i> <i>Zitong Yang, Neil Band, Shuangping Li, Emmanuel Candes, Tatsunori Hashimoto</i>	164
<i>Weakly-Supervised Affordance Grounding Guided by Part-Level Semantic Priors</i> <i>Peiran Xu, Yadong Mu</i>	207
<i>Divide and Translate: Compositional First-Order Logic Translation and Verification for Complex Logical Reasoning</i> <i>Hyun Ryu, Gyeongman Kim, Hyemin S. Lee, Eunho Yang</i>	236
<i>Multi-Modal Agent Tuning: Building a VLM-Driven Agent for Efficient Tool Usage</i> <i>Zhi Gao, Bofei Zhang, Pengxiang Li, Xiaojian Ma, Tao Yuan, Yue Fan, Yuwei Wu, Yunde Jia, Song-Chun Zhu, Qing Li</i>	266
<i>AdvWave: Stealthy Adversarial Jailbreak Attack Against Large Audio-Language Models</i> <i>Mintong Kang, Chejian Xu, Bo Li</i>	298
<i>Neural Eulerian Scene Flow Fields</i> <i>Kyle Vedder, Neehar Peri, Ishan Khatri, Siyi Li, Eric Eaton, Mehmet Kemal Kocamaz, Yue Wang, Zhiding Yu, Deva Ramanan, Joachim Peherl</i>	315

<p>ProtComposer: Compositional Protein Structure Generation with 3D Ellipsoids</p> <p>Hannes Stark, Bowen Jing, Tomas Geffner, Jason Yim, Tommi Jaakkola, Arash Vahdat, Karsten Kreis</p>	334
<p>Multi-Modal Brain Encoding Models for Multi-Modal Stimuli</p> <p>Subba Reddy Oota, Khushbu Pahwa, Mounika Marreddy, Maneesh Kumar Singh, Manish Gupta, Bapi Raju Surampudi</p>	363
<p>Personality Alignment of Large Language Models</p> <p>Minjun Zhu, Yixuan Weng, Linyi Yang, Yue Zhang</p>	389
<p>PN-GAIL: Leveraging Non-Optimal Information from Imperfect Demonstrations</p> <p>Qiang Liu, Huiqiao Fu, Kaiqiang Tang, Chunlin Chen, Daoyi Dong</p>	439
<p>Mufu: Multilingual Fused Learning for Low-Resource Translation with LLM</p> <p>Zheng Wei Lim, Nitish Gupta, Honglin Yu, Trevor Cohn</p>	465
<p>Improved Sampling of Diffusion Models in Fluid Dynamics with Tweedie's Formula</p> <p>Youssef Shehata, Benjamin Holzs Schuh, Nils Thuerey</p>	496
<p>Emergence of a High-Dimensional Abstraction Phase in Language Transformers</p> <p>Emily Cheng, Diego Doimo, Corentin Kervadec, Iuri Macocco, Lei Yu, Alessandro Laio, Marco Baroni</p>	525
<p>PETRA: Parallel End-To-End Training with Reversible Architectures</p> <p>Stephane Rivaud, Louis Fournier, Thomas Pumir, Eugene Belilovsky, Michael Eickenberg, Edouard Oyallon</p>	551
<p>ThunderKittens: Simple, Fast, and $\textit{Adorable}$ Kernels</p> <p>Benjamin Frederick Spector, Simran Arora, Aaryan Singhal, Arjun Parthasarathy, Daniel Y Fu, Christopher Re</p>	568
<p>RTDiff: Reverse Trajectory Synthesis Via Diffusion for Offline Reinforcement Learning</p> <p>Qianlan Yang, Yu-Xiong Wang</p>	602
<p>Physics-Informed Deep Inverse Operator Networks for Solving PDE Inverse Problems</p> <p>Sung Woong Cho, Hwijae Son</p>	621
<p>Probing the Latent Hierarchical Structure of Data Via Diffusion Models</p> <p>Antonio Sclocchi, Alessandro Favero, Noam Itzhak Levi, Matthieu Wyart</p>	645
<p>Universal Generalization Guarantees for Wasserstein Distributionally Robust Models</p> <p>Tam Le, Jerome Malick</p>	675
<p>Comparing Targeting Strategies for Maximizing Social Welfare with Limited Resources</p> <p>Vibhhu Sharma, Bryan Wilder</p>	710
<p>Planning Anything with Rigor: General-Purpose Zero-Shot Planning with LLM-Based Formalized Programming</p> <p>Yilun Hao, Yang Zhang, Chuchu Fan</p>	731

<p>Brain-Inspired L_p-Convolution Benefits Large Kernels and Aligns Better with Visual Cortex</p> <p><i>Jea Kwon, Sungjun Lim, Kyungwoo Song, C. Justin Lee</i></p>	788
<p>Incremental Causal Effect for Time to Treatment Initialization</p> <p><i>Andrew Ying, Zhichen Zhao, Ronghui Xu</i></p>	819
<p>VD3D: Taming Large Video Diffusion Transformers for 3D Camera Control</p> <p><i>Sherwin Bahmani, Ivan Skorokhodov, Aliaksandr Siarohin, Willi Menapace, Guocheng Qian, Michael Vasilkovsky, Hsin-Ying Lee, Chaoyang Wang, Jiaxu Zou, Andrea Tagliasacchi, David B. Lindell, Sergey Tulyakov</i></p>	836
<p>Variational Bayesian Pseudo-Coreset</p> <p><i>Hyungi Lee, Seungyoo Lee, Juho Lee</i></p>	862
<p>PEAR: Primitive Enabled Adaptive Relabeling for Boosting Hierarchical Reinforcement Learning</p> <p><i>Utsav Singh, Vinay P. Namboodiri</i></p>	898
<p>Going Beyond Feature Similarity: Effective Dataset Distillation Based on Class-Aware Conditional Mutual Information</p> <p><i>Xinhao Zhong, Bin Chen, Hao Fang, Xulin Gu, Shu-Tao Xia, En-Hui Yang</i></p>	922
<p>Rethinking Artistic Copyright Infringements in the Era of Text-To-Image Generative Models</p> <p><i>Mazda Moayeri, Sriram Balasubramanian, Samyadeep Basu, Priyatham Kattakinda, Atoosa Chegini, Robert Brauneis, Soheil Feizi</i></p>	940
<p>CO-MOT: Boosting End-To-End Transformer-Based Multi-Object Tracking Via Coopetition Label Assignment and Shadow Sets</p> <p><i>Feng Yan, Weixin Luo, Yujie Zhong, Yiyang Gan, Lin Ma</i></p>	974
<p>On Bits and Bandits: Quantifying the Regret-Information Trade-Off</p> <p><i>Itai Shufaro, Nadav Merlis, Nir Weinberger, Shie Mannor</i></p>	987
<p>A Percolation Model of Emergence: Analyzing Transformers Trained on a Formal Language</p> <p><i>Ekdeep Singh Lubana, Kyogo Kawaguchi, Robert P. Dick, Hidenori Tanaka</i></p>	1012
<p>MIM-Refiner: A Contrastive Learning Boost from Intermediate Pre-Trained Masked Image Modeling Representations</p> <p><i>Benedikt Alkin, Lukas Miklautz, Sepp Hochreiter, Johannes Brandstetter</i></p>	1071
<p>Unsupervised Zero-Shot Reinforcement Learning Via Dual-Value Forward-Backward Representation</p> <p><i>Jingbo Sun, Songjun Tu, Qichao Zhang, Haoran Li, Xin Liu, Yaran Chen, Ke Chen, Dongbin Zhao</i></p>	1108
<p>Regret-Optimal List Replicable Bandit Learning: Matching Upper and Lower Bounds</p> <p><i>Michael Chen, A. Pavan, N. V. Vinodchandran, Ruosong Wang, Lin Yang</i></p>	1138

<p>BOND: Aligning LLMs with Best-Of-N Distillation</p> <p><i>Pier Giuseppe Sessa, Robert Dadashi-Tazehozhi, Leonard Hussenot, Johan Ferret, Nino Vieillard, Alexandre Rame, Bobak Shahriari, Sarah Perrin, Abram L. Friesen, Geoffrey Cideron, Sertan Girgin, Piotr Stanczyk, Andrea Michi, Danila Sinopalnikov, Sabela Ramos Garea, Amélie Héliou, Aliaksei Severyn, Matthew Hoffman, Nikola Momchev, Olivier Bachem</i></p>	1154
<p>SAVA: Scalable Learning-Agnostic Data Valuation</p> <p><i>Samuel Kessler, Tam Le, Vu Nguyen</i></p>	1175
<p>Edge-Aware Image Smoothing with Relative Wavelet Domain Representation</p> <p><i>Huiqing Qi, Xiaoliu Luo, Tingting Li, Fang Li</i></p>	1202
<p>ORSO: Accelerating Reward Design Via Online Reward Selection and Policy Optimization</p> <p><i>Chen Bo Calvin Zhang, Zhang-Wei Hong, Aldo Pacchiano, Pulkit Agrawal</i></p>	1224
<p>Toward Exploratory Inverse Constraint Inference with Generative Diffusion Verifiers</p> <p><i>Runyi Zhao, Sheng Xu, Bo Yue, Guiliang Liu</i></p>	1268
<p>Influence-Guided Diffusion for Dataset Distillation</p> <p><i>Mingyang Chen, Jiawei Du, Bo Huang, Yi Wang, Xiaobo Zhang, Wei Wang</i></p>	1287
<p>Methods for Convex $\\$(L_0, L_1)\\$-Smooth Optimization: Clipping, Acceleration, and Adaptivity</p> <p><i>Eduard Gorbunov, Nazarii Tupitsa, Sayantan Choudhury, Alen Aliev, Peter Richtárik, Samuel Horváth, Martin Takác</i></p>	1309
<p>Uni-Sign: Toward Unified Sign Language Understanding at Scale</p> <p><i>Zecheng Li, Wengang Zhou, Weichao Zhao, Kepeng Wu, Hezhen Hu, Houqiang Li</i></p>	1354
<p>Benchmarking Vision Language Model Unlearning Via Fictitious Facial Identity Dataset</p> <p><i>Yingzi Ma, Jiong Xiao Wang, Fei Wang, Siyuan Ma, Jiazhao Li, Jinsheng Pan, Xiujun Li, Furong Huang, Lichao Sun, Bo Li, Yejin Choi, Muhao Chen, Chaowei Xiao</i></p>	1374
<p>Monitoring Latent World States in Language Models with Propositional Probes</p> <p><i>Jiahai Feng, Stuart Russell, Jacob Steinhardt</i></p>	1396
<p>Factor Graph-Based Interpretable Neural Networks</p> <p><i>Yicong Li, Kuanjiu Zhou, Shuo Yu, Qiang Zhang, Renqiang Luo, Xiaodong Li, Feng Xia</i></p>	1419
<p>In Vivo Cell-Type and Brain Region Classification Via Multimodal Contrastive Learning</p> <p><i>Han Yu, Hanrui Lyu, Yixun Xu, Charlie Windolf, Eric Kenji Lee, Fan Yang, Andrew M Shelton, Olivier Winter, International Brain Laboratory, Eva L Dyer, Chandramouli Chandrasekaran, Nicholas A. Steinmetz, Liam Paninski, Cole Lincoln Hurwitz</i></p>	1444
<p>OmniRe: Omni Urban Scene Reconstruction</p> <p><i>Ziyu Chen, Jiawei Yang, Jiahui Huang, Riccardo De Lutio, Janick Martinez Esturo, Boris Ivanovic, Or Litany, Zan Gojcic, Sanja Fidler, Marco Pavone, Li Song, Yue Wang</i></p>	1486
<p>TRACE: Temporal Grounding Video LLM Via Causal Event Modeling</p> <p><i>Yongxin Guo, Jingyu Liu, Mingda Li, Qingbin Liu, Xi Chen, Xiaoying Tang</i></p>	1506

<p>Transformer Meets Twicing: Harnessing Unattended Residual Information</p> <p>Laziz Abdullaev, Tan Minh Nguyen</p>	1531
<p>Robust Watermarking Using Generative Priors Against Image Editing: From Benchmarking to Advances</p> <p>Shilin Lu, Zihan Zhou, Jiayou Lu, Yuanzhi Zhu, Adams Wai-Kin Kong</p>	1555
<p>Open-Vocabulary Customization from CLIP Via Data-Free Knowledge Distillation</p> <p>Yongxian Wei, Zixuan Hu, Li Shen, Zhenyi Wang, Chun Yuan, Dacheng Tao</p>	1590
<p>Chain-Of-Action: Faithful and Multimodal Question Answering Through Large Language Models</p> <p>Zhenyu Pan, Haozheng Luo, Manling Li, Han Liu</p>	1614
<p>6D Object Pose Tracking in Internet Videos for Robotic Manipulation</p> <p>Georgy Ponimatkin, Martin Cifka, Tomas Soucek, M�d�ric Fourmy, Yann Labb�, Vladimir Petrik, Josef Sivic</p>	1634
<p>TimeMixer++: A General Time Series Pattern Machine for Universal Predictive Analysis</p> <p>Shiyu Wang, Jiawei Li, Xiaoming Shi, Zhou Ye, Baichuan Mo, Wenzhe Lin, Ju Shengtong, Zhixuan Chu, Ming Jin</p>	1662
<p>Towards Fast, Specialized Machine Learning Force Fields: Distilling Foundation Models Via Energy Hessians</p> <p>Ishan Amin, Sanjeev Raja, Aditi S. Krishnapriyan</p>	1699
<p>Improved Diffusion-Based Generative Model with Better Adversarial Robustness</p> <p>Zekun Wang, Mingyang Yi, Shuchen Xue, Zhenguo Li, Ming Liu, Bing Qin, Zhi-Ming Ma</p>	1722
<p>Feast Your Eyes: Mixture-Of-Resolution Adaptation for Multimodal Large Language Models</p> <p>Gen Luo, Yiyi Zhou, Yuxin Zhang, Xiawu Zheng, Xiaoshuai Sun, Rongrong Ji</p>	1752
<p>Long Context Compression with Activation Beacon</p> <p>Peitian Zhang, Zheng Liu, Shitao Xiao, Ninglu Shao, Qiwei Ye, Zhicheng Dou</p>	1768
<p>Towards Robust and Parameter-Efficient Knowledge Unlearning for LLMs</p> <p>Sungmin Cha, Sungjun Cho, Dasol Hwang, Moontae Lee</p>	1785
<p>Advancing Out-Of-Distribution Detection Via Local Neuroplasticity</p> <p>Alessandro Canevaro, Julian Schmidt, Mohammad Sajad Marvi, Hang Yu, Georg Martius, Julian Jordan</p>	1808
<p>Perplexed by Perplexity: Perplexity-Based Data Pruning with Small Reference Models</p> <p>Zachary Ankner, Cody Blakeney, Kartik Sreenivasan, Max Marion, Matthew L Leavitt, Mansheej Paul</p>	1833
<p>Optimizing Neural Network Representations of Boolean Networks</p> <p>Joshua Russell, Ignacio Gavier, Devdhar Patel, Edward Rietman, Hava T Siegelmann</p>	1853

<p>Wide Neural Networks Trained with Weight Decay Provably Exhibit Neural Collapse</p> <p><i>Arthur Jacot, Peter Šúkeník, Zihan Wang, Marco Mondelli</i></p>	1905
<p>Procedural Knowledge in Pretraining Drives Reasoning in Large Language Models</p> <p><i>Laura Ruis, Maximilian Mozes, Juhan Bae, Siddhartha Rao Kamalakara, Dwaraknath Gnaneshwar, Acyr Locatelli, Robert Kirk, Tim Rocktäschel, Edward Grefenstette, Max Bartolo</i></p>	1932
<p>Conformal Generative Modeling with Improved Sample Efficiency Through Sequential Greedy Filtering</p> <p><i>Klaus-Rudolf Kladny, Bernhard Schölkopf, Michael Muehlebach</i></p>	1995
<p>Rapid Selection and Ordering of In-Context Demonstrations Via Prompt Embedding Clustering</p> <p><i>Kha Pham, Hung Le, Man Ngo, Truyen Tran</i></p>	2025
<p>Fast Uncovering of Protein Sequence Diversity from Structure</p> <p><i>Luca Alessandro Silva, Barthelemy Meynard-Piganeau, Carlo Lucibello, Christoph Feinauer</i></p>	2042
<p>MathCoder2: Better Math Reasoning from Continued Pretraining on Model-Translated Mathematical Code</p> <p><i>Zimu Lu, Aojun Zhou, Ke Wang, Houxing Ren, Weikang Shi, Junting Pan, Mingjie Zhan, Hongsheng Li</i></p>	2070
<p>Animate-X: Universal Character Image Animation with Enhanced Motion Representation</p> <p><i>Shuai Tan, Biao Gong, Xiang Wang, Shiwei Zhang, Dandan Zheng, Ruobing Zheng, Kecheng Zheng, Jingdong Chen, Ming Yang</i></p>	2091
<p>Bayesian Optimization Via Continual Variational Last Layer Training</p> <p><i>Paul Brunzema, Mikkel Jordahn, John Willes, Sebastian Trimpe, Jasper Snoek, James Harrison</i></p>	2120
<p>PAL: Sample-Efficient Personalized Reward Modeling for Pluralistic Alignment</p> <p><i>Daiwei Chen, Yi Chen, Aniket Rege, Zhi Wang, Ramya Korlakai Vinayak</i></p>	2150
<p>Century: A Framework and Dataset for Evaluating Historical Contextualisation of Sensitive Images</p> <p><i>Canfer Akbulut, Kevin Robinson, Maribeth Rauh, Isabela Albuquerque, Olivia Wiles, Laura Weidinger, Verena Rieser, Yana Hasson, Nahema Marchal, Iason Gabriel, William Isaac, Lisa Anne Hendricks</i></p>	2195
<p>Test-Time Adaptation for Combating Missing Modalities in Egocentric Videos</p> <p><i>Merey Ramazanova, Alejandro Pardo, Bernard Ghanem, Motasem Alfarra</i></p>	2246
<p>Learning Diverse Attacks on Large Language Models for Robust Red-Teaming and Safety Tuning</p> <p><i>Seanie Lee, Minsu Kim, Lynn Cherif, David Dobre, Juho Lee, Sung Ju Hwang, Kenji Kawaguchi, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, Moksh Jain</i></p>	2266
<p>Revisiting Mode Connectivity in Neural Networks with Bezier Surface</p> <p><i>Jie Ren, Pin-Yu Chen, Ren Wang</i></p>	2289
<p>Towards Principled Evaluations of Sparse Autoencoders for Interpretability and Control</p> <p><i>Aleksandar Makelov, Georg Lange, Neel Nanda</i></p>	2307

<p>TexTailor: Customized Text-Aligned Texturing Via Effective Resampling</p> <p>Suin Lee, Daeshik Kim</p>	2356
<p>Monet: Mixture of Monosemantic Experts for Transformers</p> <p>Jungwoo Park, Ahn Young Jin, Kee-Eung Kim, Jaewoo Kang</p>	2373
<p>GenSE: Generative Speech Enhancement Via Language Models Using Hierarchical Modeling</p> <p>Jixun Yao, Hexin Liu, Chen Chen, Yuchen Hu, Engsiang Chng, Lei Xie</p>	2405
<p>Bridging Jensen Gap for Max-Min Group Fairness Optimization in Recommendation</p> <p>Chen Xu, Yuxin Li, Wenjie Wang, Liang Pang, Jun Xu, Tat-Seng Chua</p>	2426
<p>RMP-SAM: Towards Real-Time Multi-Purpose Segment Anything</p> <p>Shilin Xu, Haobo Yuan, Qingyu Shi, Lu Qi, Jingbo Wang, Yibo Yang, Yining Li, Kai Chen, Yunhai Tong, Bernard Ghanem, Xiangtai Li, Ming-Hsuan Yang</p>	2454
<p>Generation and Comprehension Hand-In-Hand: Vision-Guided Expression Diffusion for Boosting Referring Expression Generation and Comprehension</p> <p>Jingcheng Ke, Jun-Cheng Chen, I-Hong Jhuo, Chia-Wen Lin, Yen-Yu Lin</p>	2477
<p>Aligning Language Models with Demonstrated Feedback</p> <p>Omar Shaikh, Michelle S. Lam, Joey Hejna, Yijia Shao, Hyundong Justin Cho, Michael S. Bernstein, Diyi Yang</p>	2493
<p>Coreset Spectral Clustering</p> <p>Ben Jourdan, Gregory Schwartzman, Peter Macgregor, He Sun</p>	2521
<p>NetMoE: Accelerating MoE Training Through Dynamic Sample Placement</p> <p>Xinyi Liu, Yujie Wang, Fangcheng Fu, Xupeng Miao, Shenhan Zhu, Xiaonan Nie, Bin Cui</p>	2537
<p>Fully-Inductive Node Classification on Arbitrary Graphs</p> <p>Jianan Zhao, Zhaocheng Zhu, Mikhail Galkin, Hesham Mostafa, Michael M. Bronstein, Jian Tang</p>	2556
<p>More Experts than Galaxies: Conditionally-Overlapping Experts with Biologically-Inspired Fixed Routing</p> <p>Sagi Shaier, Francisco Pereira, Katharina Von Der Wense, Lawrence Hunter, Matt Jones</p>	2575
<p>SysCaps: Language Interfaces for Simulation Surrogates of Complex Systems</p> <p>Patrick Emami, Zhaonan Li, Saumya Sinha, Truc Nguyen</p>	2605
<p>Aligned Better, Listen Better for Audio-Visual Large Language Models</p> <p>Yuxin Guo, Shuailei Ma, Shijie Ma, Xiaoyi Bao, Chen-Wei Xie, Kecheng Zheng, Tingyu Weng, Siyang Sun, Yun Zheng, Wei Zou</p>	2628
<p>RevisEval: Improving LLM-As-A-Judge Via Response-Adapted References</p> <p>Qiyuan Zhang, Yufei Wang, Tiezheng Yu, Yuxin Jiang, Chuhan Wu, Liangyou Li, Yasheng Wang, Xin Jiang, Lifeng Shang, Ruiming Tang, Fuyuan Lyu, Chen Ma</p>	2653

<p>GenXD: Generating Any 3D and 4D Scenes</p> <p>Yuyang Zhao, Chung-Ching Lin, Kevin Lin, Zhiwen Yan, Linjie Li, Zhengyuan Yang, Jianfeng Wang, Gim Hee Lee, Lijuan Wang</p>	2681
<p>MoS: Unleashing Parameter Efficiency of Low-Rank Adaptation with Mixture of Shards</p> <p>Sheng Wang, Liheng Chen, Pengan Chen, Jingwei Dong, Boyang Xue, Jiyue Jiang, Lingpeng Kong, Chuan Wu</p>	2702
<p>MAP: Low-Compute Model Merging with Amortized Pareto Fronts Via Quadratic Approximation</p> <p>Lu Li, Tianyu Zhang, Zhiqi Bu, Suyuchen Wang, Huan He, Jie Fu, Yonghui Wu, Jiang Bian, Yong Chen, Yoshua Bengio</p>	2719
<p>Cross-Attention Head Position Patterns Can Align with Human Visual Concepts in Text-To-Image Generative Models</p> <p>Jungwon Park, Jungmin Ko, Dongnam Byun, Jangwon Suh, Wonjong Rhee</p>	2752
<p>Variational Search Distributions</p> <p>Daniel M. Steinberg, Rafael Oliveira, Cheng Soon Ong, Edwin V. Bonilla</p>	2812
<p>Mastering Task Arithmetic: τ as a Key Indicator for Weight Disentanglement</p> <p>Kotaro Yoshida, Yuji Naraki, Takafumi Horie, Ryosuke Yamaki, Ryotaro Shimizu, Yuki Saito, Julian McAuley, Hiroki Naganuma</p>	2852
<p>SurFhead: Affine Rig Blending for Geometrically Accurate 2D Gaussian Surfel Head Avatars</p> <p>Jaeseong Lee, Taewoong Kang, Marcel Buehler, Min-Jung Kim, Sungwon Hwang, Junha Hyung, Hyojin Jang, Jaegul Choo</p>	2877
<p>Discrete GCBF Proximal Policy Optimization for Multi-Agent Safe Optimal Control</p> <p>Songyuan Zhang, Oswin So, Mitchell Black, Chuchu Fan</p>	2902
<p>Are Transformers Able to Reason by Connecting Separated Knowledge in Training Data?</p> <p>Yutong Yin, Zhaoran Wang</p>	2933
<p>High-Dimensional Analysis of Knowledge Distillation: Weak-To-Strong Generalization and Scaling Laws</p> <p>Muhammed Emrullah Ildiz, Halil Alperen Gozeten, Ege Onur Taga, Marco Mondelli, Samet Oymak</p>	2967
<p>Lambda-Skip Connections: The Architectural Component that Prevents Rank Collapse</p> <p>Federico Arangath Joseph, Jerome Sieber, Melanie Zeilinger, Carmen Amo Alonso</p>	3007
<p>Wicked Oddities: Selectively Poisoning for Effective Clean-Label Backdoor Attacks</p> <p>Nguyen Hung-Quang, Ngoc-Hieu Nguyen, The-Anh Ta, Thanh Nguyen-Tang, Kok-Seng Wong, Hoang Thanh-Tung, Khoa D Doan</p>	3033
<p>Differentially Private Learners for Heterogeneous Treatment Effects</p> <p>Maresa Schröder, Valentyn Melnychuk, Stefan Feuerriegel</p>	3056

<p>BiGR: Harnessing Binary Latent Codes for Image Generation and Improved Visual Representation Capabilities</p> <p>Shaozhe Hao, Xuantong Liu, Xianbiao Qi, Shihao Zhao, Bojia Zi, Rong Xiao, Kai Han, Kwan-Yee K. Wong</p>	3085
<p>Neuroplastic Expansion in Deep Reinforcement Learning</p> <p>Jiashun Liu, Johan Samir Obando Ceron, Aaron Courville, Ling Pan</p>	3114
<p>Balanced Ranking with Relative Centrality: A Multi-Core Periphery Perspective</p> <p>Chandra Sekhar Mukherjee, Jiapeng Zhang</p>	3137
<p>A General Framework for Producing Interpretable Semantic Text Embeddings</p> <p>Yiqun Sun, Qiang Huang, Yixuan Tang, Anthony Kum Hoe Tung, Jun Yu</p>	3175
<p>Explore Theory of Mind: Program-Guided Adversarial Data Generation for Theory of Mind Reasoning</p> <p>Melanie Sclar, Jane Dwivedi-Yu, Maryam Fazel-Zarandi, Yulia Tsvetkov, Yonatan Bisk, Yejin Choi, Asli Celikyilmaz</p>	3199
<p>A Deep Inverse-Mapping Model for a Flapping Robotic Wing</p> <p>Hadar Sharvit, Raz Karl, Tsevi Beatus</p>	3225
<p>How Gradient Descent Balances Features: A Dynamical Analysis for Two-Layer Neural Networks</p> <p>Zhenyu Zhu, Fanghui Liu, Volkan Cevher</p>	3244
<p>Towards Understanding Why FixMatch Generalizes Better than Supervised Learning</p> <p>Jingyang Li, Jiachun Pan, Vincent Y. F. Tan, Kim-Chuan Toh, Pan Zhou</p>	3288
<p>HQGS: High-Quality Novel View Synthesis with Gaussian Splatting in Degraded Scenes</p> <p>Xin Lin, Shi Luo, Xiaojun Shan, Xiaoyu Zhou, Chao Ren, Lu Qi, Ming-Hsuan Yang, Nuno Vasconcelos</p>	3326
<p>PharmacoMatch: Efficient 3D Pharmacophore Screening Via Neural Subgraph Matching</p> <p>Daniel Rose, Oliver Wieder, Thomas Seidel, Thierry Langer</p>	3343
<p>Make Haste Slowly: A Theory of Emergent Structured Mixed Selectivity in Feature Learning ReLU Networks</p> <p>Devon Jarvis, Richard Klein, Benjamin Rosman, Andrew M Saxe</p>	3367
<p>CREIMBO: Cross-Regional Ensemble Interactions in Multi-View Brain Observations</p> <p>Noga Mudrik, Ryan Ly, Oliver Ruebel, Adam Shabti Charles</p>	3402
<p>Beyond Graphs: Can Large Language Models Comprehend Hypergraphs?</p> <p>Yifan Feng, Chengwu Yang, Xingliang Hou, Shaoyi Du, Shihui Ying, Zongze Wu, Yue Gao</p>	3445
<p>HELMET: How to Evaluate Long-Context Models Effectively and Thoroughly</p> <p>Howard Yen, Tianyu Gao, Minmin Hou, Ke Ding, Daniel Fleischer, Peter Izsak, Moshe Wasserblat, Danqi Chen</p>	3473

<p>Conformal Structured Prediction</p> <p><i>Botong Zhang, Shuo Li, Osbert Bastani</i></p>	3525
<p>Learning from Weak Labelers as Constraints</p> <p><i>Vishwajeet Agrawal, Rattana Pukdee, Maria Florina Balcan, Pradeep Kumar Ravikumar</i></p>	3546
<p>DeFT: Decoding with Flash Tree-Attention for Efficient Tree-Structured LLM Inference</p> <p><i>Jinwei Yao, Kaiqi Chen, Kexun Zhang, Jiaxuan You, Binhang Yuan, Zeke Wang, Tao Lin</i></p>	3587
<p>SelectFormer in Data Markets: Privacy-Preserving and Efficient Data Selection for Transformers with Multi-Party Computation</p> <p><i>Xu Ouyang, Felix Xiaozhu Lin, Yangfeng Ji</i></p>	3619
<p>Be More Diverse than the Most Diverse: Optimal Mixtures of Generative Models Via Mixture-UCB Bandit Algorithms</p> <p><i>Parham Rezaei, Farzan Farnia, Cheuk Ting Li</i></p>	3634
<p>Talking Turns: Benchmarking Audio Foundation Models on Turn-Taking Dynamics</p> <p><i>Siddhant Arora, Zhiyun Lu, Chung-Cheng Chiu, Ruoming Pang, Shinji Watanabe</i></p>	3663
<p>Advancing LLM Reasoning Generalists with Preference Trees</p> <p><i>Lifan Yuan, Ganqu Cui, Hanbin Wang, Ning Ding, Xingyao Wang, Boji Shan, Zeyuan Liu, Jia Deng, Huimin Chen, Ruobing Xie, Yankai Lin, Zhenghao Liu, Bowen Zhou, Hao Peng, Zhiyuan Liu, Maosong Sun</i></p>	3691
<p>Looped Transformers for Length Generalization</p> <p><i>Ying Fan, Yilun Du, Kannan Ramchandran, Kangwook Lee</i></p>	3714
<p>On the Identification of Temporal Causal Representation with Instantaneous Dependence</p> <p><i>Zijian Li, Yifan Shen, Kaitao Zheng, Ruichu Cai, Xiangchen Song, Mingming Gong, Guangyi Chen, Kun Zhang</i></p>	3733
<p>HiLo: A Learning Framework for Generalized Category Discovery Robust to Domain Shifts</p> <p><i>Hongjun Wang, Sagar Vaze, Kai Han</i></p>	3773
<p>Few-Class Arena: A Benchmark for Efficient Selection of Vision Models and Dataset Difficulty Measurement</p> <p><i>Bryan Bo Cao, Lawrence O’Gorman, Michael Coss, Shubham Jain</i></p>	3815
<p>Mixture of In-Context Prompters for Tabular PFNs</p> <p><i>Derek Qiang Xu, F Olcay Cirit, Reza Asadi, Yizhou Sun, Wei Wang</i></p>	3847
<p>From Commands to Prompts: LLM-Based Semantic File System for AIOS</p> <p><i>Zeru Shi, Kai Mei, Mingyu Jin, Yongye Su, Chaoji Zuo, Wenyue Hua, Wujiang Xu, Yujie Ren, Zirui Liu, Mengnan Du, Dong Deng, Yongfeng Zhang</i></p>	3883
<p>I Can Hear You: Selective Robust Training for Deepfake Audio Detection</p> <p><i>Zirui Zhang, Wei Hao, Aroon Sankoh, William Lin, Emanuel Mendiola-Ortiz, Junfeng Yang, Chengzhi Mao</i></p>	3907

Algorithmic Stability Based Generalization Bounds for Adversarial Training Runzhi Tian, Yongyi Mao	3926
DeepRTL: Bridging Verilog Understanding and Generation with a Unified Representation Model Yi Liu, Changran Xu, Yunhao Zhou, Zeju Li, Qiang Xu	3948
CLIPDrag: Combining Text-Based and Drag-Based Instructions for Image Editing Ziqi Jiang, Zhen Wang, Long Chen	3971
Discovering Temporally Compositional Neural Manifolds with Switching Infinite GPFA Changmin Yu, Maneesh Sahani, Máté Lengyel	3988
Online Reward-Weighted Fine-Tuning of Flow Matching with Wasserstein Regularization Jiajun Fan, Shuaike Shen, Chaoran Cheng, Yuxin Chen, Chumeng Liang, Ge Liu	4010
DOPL: Direct Online Preference Learning for Restless Bandits with Preference Feedback Guojun Xiong, Ujwal Dinesha, Debajoy Mukherjee, Jian Li, Srinivas Shakkottai	4071
Towards Universality: Studying Mechanistic Similarity Across Language Model Architectures Junxuan Wang, Xuyang Ge, Wentao Shu, Qiong Tang, Yunhua Zhou, Zhengfu He, Xipeng Qiu	4103
GLoRa: A Benchmark to Evaluate the Ability to Learn Long-Range Dependencies in Graphs Dongzhuoran Zhou, Evgeny Kharlamov, Egor V. Kostylev	4130
ParaSolver: A Hierarchical Parallel Integral Solver for Diffusion Models Jianrong Lu, Zhiyu Zhu, Junhui Hou	4151
Towards Hyperparameter-Free Optimization with Differential Privacy Ruixuan Liu, Zhiqi Bu	4178
Descent with Misaligned Gradients and Applications to Hidden Convexity Aditya Bhaskara, Ashok Cutkosky, Ravi Kumar, Manish Purohit	4194
Spectro-Riemannian Graph Neural Networks Karish Grover, Haiyang Yu, Xiang Song, Qi Zhu, Han Xie, Vassilis N. Ioannidis, Christos Faloutsos	4211
T-Stitch: Accelerating Sampling in Pre-Trained Diffusion Models with Trajectory Stitching Zizheng Pan, Bohan Zhuang, De-An Huang, Weili Nie, Zhiding Yu, Chaowei Xiao, Jianfei Cai, Anima Anandkumar	4238
Automated Proof Generation for Rust Code Via Self-Evolution Tianyu Chen, Shuai Lu, Shan Lu, Yeyun Gong, Chenyuan Yang, Xuheng Li, Md Rakib Hossain Misu, Hao Yu, Nan Duan, Peng Cheng, Fan Yang, Shuvendu K Lahiri, Tao Xie, Lidong Zhou	4273
The Superposition of Diffusion Models Using the Itô Density Estimator Marta Skreta, Lazar Atanackovic, Joey Bose, Alexander Tong, Kirill Neklyudov	4301
ZIP: An Efficient Zeroth-Order Prompt Tuning for Black-Box Vision-Language Models Seonghwan Park, Jaehyeon Jeong, Yongjun Kim, Jaeho Lee, Namhoon Lee	4355

<p>General Scene Adaptation for Vision-And-Language Navigation</p> <p>Haodong Hong, Yanyuan Qiao, Sen Wang, Jiajun Liu, Qi Wu</p>	4389
<p>Flow Matching Achieves Almost Minimax Optimal Convergence</p> <p>Kenji Fukumizu, Taiji Suzuki, Noboru Isobe, Kazusato Oko, Masanori Koyama</p>	4417
<p>Language-Assisted Feature Transformation for Anomaly Detection</p> <p>Eunggu Yun, Heonjin Ha, Yeongwoo Nam, Bryan Dongik Lee</p>	4450
<p>Preference Elicitation for Offline Reinforcement Learning</p> <p>Alizée Pace, Bernhard Schölkopf, Gunnar Ratsch, Giorgia Ramponi</p>	4471
<p>UniINF: Best-Of-Both-Worlds Algorithm for Parameter-Free Heavy-Tailed MABs</p> <p>Yu Chen, Jiatai Huang, Yan Dai, Longbo Huang</p>	4500
<p>Global Convergence of Policy Gradient in Average Reward MDPs</p> <p>Navdeep Kumar, Yashaswini Murthy, Itai Shufaro, Kfir Yehuda Levy, R. Srikant, Shie Mannor</p>	4537
<p>MoDGS: Dynamic Gaussian Splatting from Casually-Captured Monocular Videos with Depth Priors</p> <p>Qingming Liu, Yuan Liu, Jiepeng Wang, Xianqiang Lyu, Peng Wang, Wenping Wang, Junhui Hou</p>	4568
<p>CollabEdit: Towards Non-Destructive Collaborative Knowledge Editing</p> <p>Jiamu Zheng, Jinghuai Zhang, Tianyu Du, Xuhong Zhang, Jianwei Yin, Tao Lin</p>	4595
<p>Gradient Correlation is a Key Ingredient to Accelerate SGD with Momentum</p> <p>Julien Hermant, Marien Renaud, Jean-François Aujol, Charles Dossal, Aude Rondepierre</p>	4616
<p>On a Connection Between Imitation Learning and RLHF</p> <p>Teng Xiao, Yige Yuan, Mingxiao Li, Zhengyu Chen, Vasant G Honavar</p>	4670
<p>PersonalLLM: Tailoring LLMs to Individual Preferences</p> <p>Thomas P Zollo, Andrew Wei Tung Siah, Naimeng Ye, Ang Li, Hongseok Namkoong</p>	4686
<p>ELBOing Stein: Variational Bayes with Stein Mixture Inference</p> <p>Ola Rønning, Eric Nalisnick, Christophe Ley, Padhraic Smyth, Thomas Hamelryck</p>	4709
<p>Boost Self-Supervised Dataset Distillation Via Parameterization, Predefined Augmentation, and Approximation</p> <p>Sheng-Feng Yu, Jia-Jiun Yao, Wei-Chen Chiu</p>	4734
<p>SynQ: Accurate Zero-Shot Quantization by Synthesis-Aware Fine-Tuning</p> <p>Minjun Kim, Jongjin Kim, U Kang</p>	4754
<p>MEGA-Bench: Scaling Multimodal Evaluation to Over 500 Real-World Tasks</p> <p>Jiacheng Chen, Tianhao Liang, Sherman Siu, Zhengqing Wang, Kai Wang, Yubo Wang, Yuansheng Ni, Ziyang Jiang, Wang Zhu, Bohan Lyu, Dongfu Jiang, Xuan He, Yuan Liu, Hexiang Hu, Xiang Yue, Wenhui Chen</p>	4780

<p>VisualAgentBench: Towards Large Multimodal Models as Visual Foundation Agents</p> <p>Xiao Liu, Tianjie Zhang, Yu Gu, Jiat Long long, Song Xixuan, Yifan Xu, Shudan Zhang, Hanyu Lai, Jiadai Sun, Xinyue Yang, Yu Yang, Zehan Qi, Shuntian Yao, Xueqiao Sun, Siyi Cheng, Qinkai Zheng, Hao Yu, Hanchen Zhang, Wenyi Hong, Ming Ding, Lihang Pan, Xiaotao Gu, Aohan Zeng, Zhengxiao Du, Chan Hee Song, Yu Su, Yuxiao Dong, Jie Tang</p>	4891
<p>InvestESG: A Multi-Agent Reinforcement Learning Benchmark for Studying Climate Investment as a Social Dilemma</p> <p>Xiaoxuan Hou, Jiayi Yuan, Joel Z Leibo, Natasha Jaques</p>	4949
<p>Sparse Autoencoders Reveal Temporal Difference Learning in Large Language Models</p> <p>Can Demircan, Tankred Saanum, Akshay Kumar Jagadish, Marcel Binz, Eric Schulz</p>	4972
<p>DICE: Data Influence Cascade in Decentralized Learning</p> <p>Tongtian Zhu, Wenhao Li, Can Wang, Fengxiang He</p>	4998
<p>Long-Context Linear System Identification</p> <p>Oguz Kaan Yüksel, Mathieu Even, Nicolas Flammarion</p>	5034
<p>Has the Deep Neural Network Learned the Stochastic Process? an Evaluation Viewpoint</p> <p>Harshit Kumar, Beomseok Kang, Biswadeep Chakraborty, Saibal Mukhopadhyay</p>	5069
<p>Null Counterfactual Factor Interactions for Goal-Conditioned Reinforcement Learning</p> <p>Caleb Chuck, Fan Feng, Carl Qi, Chang Shi, Siddhant Agarwal, Amy Zhang, Scott Niekum</p>	5102
<p>Robust Gymnasium: A Unified Modular Benchmark for Robust Reinforcement Learning</p> <p>Shangding Gu, Laixi Shi, Muning Wen, Ming Jin, Eric Mazumdar, Yuejie Chi, Adam Wierman, Costas Spanos</p>	5127
<p>Residual Kernel Policy Network: Enhancing Stability and Robustness in RKHS-Based Reinforcement Learning</p> <p>Yixian Zhang, Huaze Tang, Huijing Lin, Wenbo Ding</p>	5155
<p>Revisit the Open Nature of Open Vocabulary Semantic Segmentation</p> <p>Qiming Huang, Han Hu, Jianbo Jiao</p>	5181
<p>Unlocking Global Optimality in Bilevel Optimization: A Pilot Study</p> <p>Quan Xiao, Tianyi Chen</p>	5200
<p>Distribution Backtracking Builds a Faster Convergence Trajectory for Diffusion Distillation</p> <p>Shengyuan Zhang, Ling Yang, Zejian Li, An Zhao, Chenye Meng, Changyuan Yang, Guang Yang, Zhiyuan Yang, Lingyun Sun</p>	5254
<p>Trajectory Attention for Fine-Grained Video Motion Control</p> <p>Zeqi Xiao, Wenqi Ouyang, Yifan Zhou, Shuai Yang, Lei Yang, Jianlou Si, Xingang Pan</p>	5283
<p>Improving Long-Text Alignment for Text-To-Image Diffusion Models</p> <p>Luping Liu, Chao Du, Tianyu Pang, Zehan Wang, Chongxuan Li, Dong Xu</p>	5298

<p>Qinco2: Vector Compression and Search with Improved Implicit Neural Codebooks</p> <p><i>Théophane Vallaëys, Matthew J. Muckley, Jakob Verbeek, Matthijs Douze</i></p>	5328
<p>NL-Eye: Abductive NLI for Images</p> <p><i>Mor Ventura, Michael Toker, Nitay Calderon, Zorik Gekhman, Yonatan Bitton, Roi Reichart</i></p>	5346
<p>Seq-VCR: Preventing Collapse in Intermediate Transformer Representations for Enhanced Reasoning</p> <p><i>Md Rifat Arefin, Gopeshh Subbaraj, Nicolas Gontier, Yann Lecun, Irina Rish, Ravid Shwartz-Ziv, Christopher Pal</i></p>	5379
<p>From Tokens to Words: On the Inner Lexicon of LLMs</p> <p><i>Guy Kaplan, Matanel Oren, Yuval Reif, Roy Schwartz</i></p>	5401
<p>Learning Spatiotemporal Dynamical Systems from Point Process Observations</p> <p><i>Valerii Iakovlev, Harri Lähdesmäki</i></p>	5426
<p>Revisiting a Design Choice in Gradient Temporal Difference Learning</p> <p><i>Xiaochi Qian, Shangtong Zhang</i></p>	5447
<p>Refining CLIP's Spatial Awareness: A Visual-Centric Perspective</p> <p><i>Congpei Qiu, Yanhao Wu, Wei Ke, Xiuxiu Bai, Tong Zhang</i></p>	5479
<p>LongGenBench: Benchmarking Long-Form Generation in Long Context LLMs</p> <p><i>Yuhao Wu, Ming Shan Hee, Zhiqiang Hu, Roy Ka-Wei Lee</i></p>	5506
<p>LeFusion: Controllable Pathology Synthesis Via Lesion-Focused Diffusion Models</p> <p><i>Hantao Zhang, Yuhe Liu, Jiancheng Yang, Shouhong Wan, Xinyuan Wang, Wei Peng, Pascal Fua</i></p>	5528
<p>UniCoTT: A Unified Framework for Structural Chain-Of-Thought Distillation</p> <p><i>Xianwei Zhuang, Zhihong Zhu, Zhichang Wang, Xuxin Cheng, Yuexian Zou</i></p>	5550
<p>Video Action Differencing</p> <p><i>James Burgess, Xiaohan Wang, Yuhui Zhang, Anita Rau, Alejandro Lozano, Lisa Dunlap, Trevor Darrell, Serena Yeung-Levy</i></p>	5575
<p>Not All Prompts Are Made Equal: Prompt-Based Pruning of Text-To-Image Diffusion Models</p> <p><i>Alireza Ganjdanesh, Reza Shirkavand, Shangqian Gao, Heng Huang</i></p>	5608
<p>Broaden Your SCOPE! Efficient Multi-Turn Conversation Planning for LLMs with Semantic Space</p> <p><i>Zhiliang Chen, Xinyuan Niu, Chuan-Sheng Foo, Bryan Kian Hsiang Low</i></p>	5639
<p>On Speeding Up Language Model Evaluation</p> <p><i>Jin Peng Zhou, Christian K Belardi, Ruihan Wu, Travis Zhang, Carla P Gomes, Wen Sun, Kilian Q Weinberger</i></p>	5664
<p>What Does it Mean to Be a Transformer? Insights from a Theoretical Hessian Analysis</p> <p><i>Weronika Ormaniec, Felix Dangel, Sidak Pal Singh</i></p>	5684

<p>Recite, Reconstruct, Recollect: Memorization in LMs as a Multifaceted Phenomenon</p> <p><i>Usvsn Sai Prashanth, Alvin Deng, Kyle O'Brien, Jyothir S V, Mohammad Aflah Khan, Jaydeep Borkar, Christopher A. Choquette-Choo, Jacob Ray Fuehne, Stella Biderman, Tracy Ke, Katherine Lee, Naomi Saphra</i></p>	5720
<p>Model-Agnostic Knowledge Guided Correction for Improved Neural Surrogate Rollout</p> <p><i>Bharat Srikishan, Daniel O'Malley, Mohamed Mehana, Nicholas Lubbers, Nikhil Muralidhar</i></p>	5741
<p>Fast Training and Sampling of Restricted Boltzmann Machines</p> <p><i>Nicolas Bereux, Aurélien Decelle, Cyril Furtlehner, Lorenzo Rosset, Beatriz Seoane</i></p>	5760
<p>Interpretable Causal Representation Learning for Biological Data in the Pathway Space</p> <p><i>Jesus De La Fuente Cedeño, Robert Lehmann, Carlos Ruiz-Arenas, Jan Voges, Irene Marín-Goñi, Xabier Martinez De Morentin, David Gomez-Cabrero, Idoia Ochoa, Jesper Tegnér, Vincenzo Lagani, Mikel Hernaez</i></p>	5794
<p>DiffPuter: Empowering Diffusion Models for Missing Data Imputation</p> <p><i>Hengrui Zhang, Liancheng Fang, Qitian Wu, Philip S. Yu</i></p>	5822
<p>CtrlLoRA: An Extensible and Efficient Framework for Controllable Image Generation</p> <p><i>Yifeng Xu, Zhenliang He, Shiguang Shan, Xilin Chen</i></p>	5844
<p>Justice Or Prejudice? Quantifying Biases in LLM-As-A-Judge</p> <p><i>Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, Nitesh V Chawla, Xiangliang Zhang</i></p>	5867
<p>Learning to Contextualize Web Pages for Enhanced Decision Making by LLM Agents</p> <p><i>Dongjun Lee, Juyong Lee, Kyuyoung Kim, Jihoon Tack, Jinwoo Shin, Yee Whye Teh, Kimin Lee</i></p>	5907
<p>WorkflowLLM: Enhancing Workflow Orchestration Capability of Large Language Models</p> <p><i>Shengda Fan, Xin Cong, Yuepeng Fu, Zhong Zhang, Shuyan Zhang, Yuanwei Liu, Yesai Wu, Yankai Lin, Zhiyuan Liu, Maosong Sun</i></p>	5941
<p>Compositional Entailment Learning for Hyperbolic Vision-Language Models</p> <p><i>Avik Pal, Max Van Spengler, Guido Maria D'Amely Di Melendugno, Alessandro Flaborea, Fabio Galasso, Pascal Mettes</i></p>	5969
<p>OptionZero: Planning with Learned Options</p> <p><i>Po-Wei Huang, Pei-Chiun Peng, Hung Guei, Ti-Rong Wu</i></p>	5998
<p>Group Distributionally Robust Dataset Distillation with Risk Minimization</p> <p><i>Saeed Vahidian, Mingyu Wang, Jianyang Gu, Vyacheslav Kungurtsev, Wei Jiang, Yiran Chen</i></p>	6024
<p>Iterative Substructure Extraction for Molecular Relational Learning with Interactive Graph Information Bottleneck</p> <p><i>Shuai Zhang, Junfeng Fang, Xuqiang Li, Hongxin Xiang, Alan Xia, Ye Wei, Wenjie Du, Yang Wang</i></p>	6048

<p>CR2PQ: Continuous Relative Rotary Positional Query for Dense Visual Representation Learning</p> <p>Shaofeng Zhang, Qiang Zhou, Sitong Wu, Haoru Tan, Zhibin Wang, Jinfa Huang, Junchi Yan</p>	6077
<p>ConMix: Contrastive Mixup at Representation Level for Long-Tailed Deep Clustering</p> <p>Zhixin Li, Yuheng Jia</p>	6091
<p>Multi-Modal and Multi-Attribute Generation of Single Cells with CFGen</p> <p>Alessandro Palma, Till Richter, Hanyi Zhang, Manuel Lubetzki, Alexander Tong, Andrea Dittadi, Fabian J Theis</p>	6111
<p>Simulating Human-Like Daily Activities with Desire-Driven Autonomy</p> <p>Yiding Wang, Yuxuan Chen, Fangwei Zhong, Long Ma, Yizhou Wang</p>	6152
<p>An Information Criterion for Controlled Disentanglement of Multimodal Data</p> <p>Chenyu Wang, Sharut Gupta, Xinyi Zhang, Sana Tonekaboni, Stefanie Jegelka, Tommi Jaakkola, Caroline Uhler</p>	6198
<p>Extendable and Iterative Structure Learning Strategy for Bayesian Networks</p> <p>Hamid Kalantari, Russell Greiner, Pouria Ramazi</p>	6228
<p>Automatic Curriculum Expert Iteration for Reliable LLM Reasoning</p> <p>Zirui Zhao, Hanze Dong, Amrita Saha, Caiming Xiong, Doyen Sahoo</p>	6241
<p>RelitLRM: Generative Relightable Radiance for Large Reconstruction Models</p> <p>Tianyuan Zhang, Zhengfei Kuang, Haian Jin, Zexiang Xu, Sai Bi, Hao Tan, He Zhang, Yiwei Hu, Milos Hasan, William T. Freeman, Kai Zhang, Fujun Luan</p>	6263
<p>Smaller, Weaker, Yet Better: Training LLM Reasoners Via Compute-Optimal Sampling</p> <p>Hritik Bansal, Arian Hosseini, Rishabh Agarwal, Vinh Q. Tran, Mehran Kazemi</p>	6286
<p>Multi-Field Adaptive Retrieval</p> <p>Millicent Li, Tongfei Chen, Benjamin Van Durme, Patrick Xia</p>	6311
<p>IPDreamer: Appearance-Controllable 3D Object Generation with Complex Image Prompts</p> <p>Bohan Zeng, Shanglin Li, Yutang Feng, Ling Yang, Juan Zhang, Hong Li, Jiaming Liu, Conghui He, Wentao Zhang, Jianzhuang Liu, Baochang Zhang, Shuicheng Yan</p>	6333
<p>Visual Description Grounding Reduces Hallucinations and Boosts Reasoning in LVLMs</p> <p>Sreyan Ghosh, Chandra Kiran Reddy Evuru, Sonal Kumar, Utkarsh Tyagi, Oriol Nieto, Zeyu Jin, Dinesh Manocha</p>	6356
<p>Input Space Mode Connectivity in Deep Neural Networks</p> <p>Jakub Vrabel, Ori Shem-Ur, Yaron Oz, David Krueger</p>	6394
<p>An Intelligent Agentic System for Complex Image Restoration Problems</p> <p>Kaiwen Zhu, Jinjin Gu, Zhiyuan You, Yu Qiao, Chao Dong</p>	6419

<p><i>Dream to Manipulate: Compositional World Models Empowering Robot Imitation Learning with Imagination</i></p> <p>Leonardo Barcellona, Andrii Zadaianchuk, Davide Allegro, Samuele Papa, Stefano Ghidoni, Efstratios Gavves</p>	6448
<p><i>From Pixels to Tokens: Byte-Pair Encoding on Quantized Visual Modalities</i></p> <p>Wanpeng Zhang, Zilong Xie, Yicheng Feng, Yijiang Li, Xingrun Xing, Sipeng Zheng, Zongqing Lu</p>	6483
<p><i>Scaling Speech-Text Pre-Training with Synthetic Interleaved Data</i></p> <p>Aohan Zeng, Zhengxiao Du, Mingdao Liu, Lei Zhang, Shengmin Jiang, Yuxiao Dong, Jie Tang</p>	6507
<p><i>CREMA: Generalizable and Efficient Video-Language Reasoning Via Multimodal Modular Fusion</i></p> <p>Shoubin Yu, Jaehong Yoon, Mohit Bansal</p>	6531
<p><i>Learn-By-Interact: A Data-Centric Framework for Self-Adaptive Agents in Realistic Environments</i></p> <p>Hongjin Su, Ruoxi Sun, Jinsung Yoon, Pengcheng Yin, Tao Yu, Sercan O Arik</p>	6556
<p><i>BRAID: Input-Driven Nonlinear Dynamical Modeling of Neural-Behavioral Data</i></p> <p>Parsa Vahidi, Omid G. Sani, Maryam Shanechi</p>	6610
<p><i>Machine Unlearning Via Simulated Oracle Matching</i></p> <p>Kristian Georgiev, Roy Rinberg, Sung Min Park, Shivam Garg, Andrew Ilyas, Aleksander Madry, Seth Neel</p>	6640
<p><i>Enhancing Federated Domain Adaptation with Multi-Domain Prototype-Based Federated Fine-Tuning</i></p> <p>Jingyuan Zhang, Yiyang Duan, Shuaicheng Niu, Yang Cao, Wei Yang Bryan Lim</p>	6679
<p><i>Multi-Label Node Classification with Label Influence Propagation</i></p> <p>Yifei Sun, Zemin Liu, Bryan Hooi, Yang Yang, Rizal Fathony, Jia Chen, Bingsheng He</p>	6702
<p><i>AI2TALE: An Innovative Information Theory-Based Approach for Learning to Localize Phishing Attacks</i></p> <p>Van Nguyen, Tingmin Wu, Xingliang Yuan, Marthie Grobler, Surya Nepal, Carsten Rudolph</p>	6727
<p><i>BrainOOD: Out-Of-Distribution Generalizable Brain Network Analysis</i></p> <p>Jiaxing Xu, Yongqiang Chen, Xia Dong, Mengcheng Lan, Tiancheng Huang, Qingtian Bian, James Cheng, Yiping Ke</p>	6751
<p><i>Tailoring Mixup to Data for Calibration</i></p> <p>Quentin Bouniot, Pavlo Mozharovskiy, Florence D'Alché-Buc</p>	6766
<p><i>Selective Label Enhancement Learning for Test-Time Adaptation</i></p> <p>Yihao Hu, Congyu Qiao, Xin Geng, Ning Xu</p>	6797

<p>RankSHAP: Shapley Value Based Feature Attributions for Learning to Rank</p> <p>Tanya Chowdhury, Yair Zick, James Allan</p>	6820
<p>Efficient Automated Circuit Discovery in Transformers Using Contextual Decomposition</p> <p>Aliyah R. Hsu, Georgia Zhou, Yeshwanth Cherapanamjeri, Yaxuan Huang, Anobel Odisho, Peter R. Carroll, Bin Yu</p>	6850
<p>Durable Quantization Conditioned Misalignment Attack on Large Language Models</p> <p>Peiran Dong, Haowei Li, Song Guo</p>	6869
<p>Learning a Fast Mixing Exogenous Block MDP Using a Single Trajectory</p> <p>Alexander Levine, Peter Stone, Amy Zhang</p>	6882
<p>Demystifying Online Clustering of Bandits: Enhanced Exploration Under Stochastic and Smoothed Adversarial Contexts</p> <p>Zhuohua Li, Maoli Liu, Xiangxiang Dai, John C. S. Lui</p>	6929
<p>On Quantizing Neural Representation for Variable-Rate Video Coding</p> <p>Junqi Shi, Zhuji Chen, Hanfei Li, Qi Zhao, Ming Lu, Tong Chen, Zhan Ma</p>	6964
<p>Test of Time: A Benchmark for Evaluating LLMs on Temporal Reasoning</p> <p>Bahare Fatemi, Mehran Kazemi, Anton Tsitsulin, Karishma Malkan, Jinyeong Yim, John Palowitch, Sungyong Seo, Jonathan Halcrow, Bryan Perozzi</p>	6988
<p>FedTMOS: Efficient One-Shot Federated Learning with Tsetlin Machine</p> <p>Shannon How Shi Qi, Jagmohan Chauhan, Geoff V. Merrett, Jonathon Hare</p>	7010
<p>Instruct-SkillMix: A Powerful Pipeline for LLM Instruction Tuning</p> <p>Simran Kaur, Simon Park, Anirudh Goyal, Sanjeev Arora</p>	7030
<p>Cross-Modal Safety Mechanism Transfer in Large Vision-Language Models</p> <p>Shicheng Xu, Liang Pang, Yunchang Zhu, Huawei Shen, Xueqi Cheng</p>	7172
<p>Single-Agent Poisoning Attacks Suffice to Ruin Multi-Agent Learning</p> <p>Fan Yao, Yuwei Cheng, Ermin Wei, Haifeng Xu</p>	7194
<p>Planning in Natural Language Improves LLM Search for Code Generation</p> <p>Evan Z Wang, Federico Cassano, Catherine Wu, Yunfeng Bai, William Song, Vaskar Nath, Ziwen Han, Sean M. Hendryx, Summer Yue, Hugh Zhang</p>	7230
<p>Structuring Benchmark into Knowledge Graphs to Assist Large Language Models in Retrieving and Designing Models</p> <p>Hanmo Liu, Shimin Di, Jialiang Wang, Zhili Wang, Jiachuan Wang, Xiaofang Zhou, Lei Chen</p>	7277
<p>KNN Attention Demystified: A Theoretical Exploration for Scalable Transformers</p> <p>Themistoklis Haris</p>	7301

<p>FreDF: Learning to Forecast in the Frequency Domain</p> <p>Hao Wang, Lichen Pan, Yuan Shen, Zhichao Chen, Degui Yang, Yifei Yang, Sen Zhang, Xinggao Liu, Haoxuan Li, Dacheng Tao</p>	7329
<p>Unlocking Point Processes Through Point Set Diffusion</p> <p>David Lüdke, Enric Rabasseda Raventós, Marcel Kollovich, Stephan Günnemann</p>	7359
<p>Training-Free Message Passing for Learning on Hypergraphs</p> <p>Bohan Tang, Zexi Liu, Keyue Jiang, Siheng Chen, Xiaowen Dong</p>	7378
<p>CONGO: Compressive Online Gradient Optimization</p> <p>Jeremy Carleton, Prathik Vijaykumar, Divyanshu Saxena, Dheeraj Narasimha, Srinivas Shakkottai, Aditya Akella</p>	7407
<p>ReMoE: Fully Differentiable Mixture-Of-Experts with ReLU Routing</p> <p>Ziteng Wang, Jun Zhu, Jianfei Chen</p>	7441
<p>Chemistry-Inspired Diffusion with Non-Differentiable Guidance</p> <p>Yuchen Shen, Chenhao Zhang, Sijie Fu, Chenghui Zhou, Newell Washburn, Barnabas Póczos</p>	7463
<p>O(d/T) Convergence Theory for Diffusion Probabilistic Models Under Minimal Assumptions</p> <p>Gen Li, Yuling Yan</p>	7498
<p>Compute-Constrained Data Selection</p> <p>Junjie Yin, Alexander M Rush</p>	7522
<p>Tamper-Resistant Safeguards for Open-Weight LLMs</p> <p>Rishub Tamirisa, Bhrgu Bharathi, Long Phan, Andy Zhou, Alice Gatti, Tarun Suresh, Maxwell Lin, Justin Wang, Rowan Wang, Ron Arel, Andy Zou, Dawn Song, Bo Li, Dan Hendrycks, Mantas Mazeika</p>	7544
<p>Learning from Negative Feedback, Or Positive Feedback Or Both</p> <p>Abbas Abdolmaleki, Bilal Piot, Bobak Shahriari, Jost Tobias Springenberg, Tim Hertweck, Michael Bloesch, Rishabh Joshi, Thomas Lampe, Junhyuk Oh, Nicolas Heess, Jonas Buchli, Martin Riedmiller</p>	7572
<p>Scaling LLM Test-Time Compute Optimally Can Be More Effective than Scaling Parameters for Reasoning</p> <p>Charlie Victor Snell, Jaehoon Lee, Kelvin Xu, Aviral Kumar</p>	7595
<p>Accelerating Goal-Conditioned Reinforcement Learning Algorithms and Research</p> <p>Michal Borkiewicz, Wladyslaw Palucki, Vivek Myers, Tadeusz Dziarmaga, Tomasz Arczewski, Lukasz Kucinski, Benjamin Eysenbach</p>	7630
<p>DreamBench++: A Human-Aligned Benchmark for Personalized Image Generation</p> <p>Yuang Peng, Yuxin Cui, Haomiao Tang, Zekun Qi, Runpei Dong, Jing Bai, Chunrui Han, Zheng Ge, Xiangyu Zhang, Shu-Tao Xia</p>	7653
<p>AdaGrad Under Anisotropic Smoothness</p> <p>Yuxing Liu, Rui Pan, Tong Zhang</p>	7676

<p>Fourier Head: Helping Large Language Models Learn Complex Probability Distributions</p> <p>Nate Gillman, Daksh Aggarwal, Michael Freeman, Chen Sun</p>	7711
<p>ND-SDF: Learning Normal Deflection Fields for High-Fidelity Indoor Reconstruction</p> <p>Ziyu Tang, Weicai Ye, Yifan Wang, Di Huang, Hujun Bao, Tong He, Guofeng Zhang</p>	7746
<p>Cached Multi-Lora Composition for Multi-Concept Image Generation</p> <p>Xiandong Zou, Mingzhu Shen, Christos-Savvas Bouganis, Yiren Zhao</p>	7776
<p>Transformers Learn Low Sensitivity Functions: Investigations and Implications</p> <p>Bhavya Vasudeva, Deqing Fu, Tianyi Zhou, Elliott Kau, Youqi Huang, Vatsal Sharan</p>	7805
<p>REFINE: Inversion-Free Backdoor Defense Via Model Reprogramming</p> <p>Yukun Chen, Shuo Shao, Enhao Huang, Yiming Li, Pin-Yu Chen, Zhan Qin, Kui Ren</p>	7835
<p>Free Hunch: Denoiser Covariance Estimation for Diffusion Models Without Extra Costs</p> <p>Severi Rissanen, Markus Heinonen, Arno Solin</p>	7863
<p>Multi-Objective Antibody Design with Constrained Preference Optimization</p> <p>Milong Ren, Zaikai He, Haicang Zhang</p>	7897
<p>Revealing and Mitigating Over-Attention in Knowledge Editing</p> <p>Pinzheng Wang, Zecheng Tang, Keyan Zhou, Juntao Li, Qiaoming Zhu, Min Zhang</p>	7926
<p>Democratic Training Against Universal Adversarial Perturbations</p> <p>Bing Sun, Jun Sun, Wei Zhao</p>	7953
<p>A-Bench: Are LMMs Masters at Evaluating AI-Generated Images?</p> <p>Zicheng Zhang, Haoning Wu, Chunyi Li, Yingjie Zhou, Wei Sun, Xiongkuo Min, Zijian Chen, Xiaohong Liu, Weisi Lin, Guangtao Zhai</p>	7975
<p>Global Identifiability of Overcomplete Dictionary Learning Via L1 and Volume Minimization</p> <p>Yuchen Sun, Kejun Huang</p>	7997
<p>LoCA: Location-Aware Cosine Adaptation for Parameter-Efficient Fine-Tuning</p> <p>Zhekai Du, Yinjie Min, Jingjing Li, Ke Lu, Changliang Zou, Lihua Peng, Tingjin Chu, Mingming Gong</p>	8015
<p>Beyond Squared Error: Exploring Loss Design for Enhanced Training of Generative Flow Networks</p> <p>Rui Hu, Yifan Zhang, Zhuoran Li, Longbo Huang</p>	8056
<p>Evidential Learning-Based Certainty Estimation for Robust Dense Feature Matching</p> <p>Lile Cai, Chuan-Sheng Foo, Xun Xu, Zaiwang Gu, Jun Cheng, Xulei Yang</p>	8080
<p>On the Self-Verification Limitations of Large Language Models on Reasoning and Planning Tasks</p> <p>Kaya Stechly, Karthik Valmeekam, Subbarao Kambhampati</p>	8097

<p>Flat Reward in Policy Parameter Space Implies Robust Reinforcement Learning</p> <p>Hyun Kyu Lee, Sung Whan Yoon</p>	8151
<p>Graph Assisted Offline-Online Deep Reinforcement Learning for Dynamic Workflow Scheduling</p> <p>Yifan Yang, Gang Chen, Hui Ma, Cong Zhang, Zhiguang Cao, Mengjie Zhang</p>	8174
<p>Rethinking Light Decoder-Based Solvers for Vehicle Routing Problems</p> <p>Ziwei Huang, Jianan Zhou, Zhiguang Cao, Yixin Xu</p>	8203
<p>Self-Evolving Multi-Agent Collaboration Networks for Software Development</p> <p>Yue Hu, Yuzhu Cai, Yaxin Du, Xinyu Zhu, Xiangrui Liu, Zijie Yu, Yuchen Hou, Shuo Tang, Siheng Chen</p>	8224
<p>Let Me Grok for You: Accelerating Grokking Via Embedding Transfer from a Weaker Model</p> <p>Zhiwei Xu, Zhiyu Ni, Yixin Wang, Wei Hu</p>	8257
<p>Pushing the Limits of All-Atom Geometric Graph Neural Networks: Pre-Training, Scaling, and Zero-Shot Transfer</p> <p>Zihan Pengmei, Zhengyuan Shen, Zichen Wang, Marcus D. Collins, Huzefa Rangwala</p>	8286
<p>What Do You See in Common? Learning Hierarchical Prototypes Over Tree-Of-Life to Discover Evolutionary Traits</p> <p>Harish Babu Manogaran, M. Maruf, Arka Daw, Kazi Sajeed Mehrab, Caleb Patrick Charpentier, Josef Uyeda, Wasila M Dahdul, Matthew J Thompson, Elizabeth G Campolongo, Kaiya L Provost, Wei-Lun Chao, Tanya Berger-Wolf, Paula Mabee, Hilmar Lapp, Anuj Karpatne</p>	8318
<p>Neural Multi-Objective Combinatorial Optimization Via Graph-Image Multimodal Fusion</p> <p>Jinbiao Chen, Jiahai Wang, Zhiguang Cao, Yaoxin Wu</p>	8355
<p>Metabench - A Sparse Benchmark of Reasoning and Knowledge in Large Language Models</p> <p>Alex Kipnis, Konstantinos Voudouris, Luca M. Schulze Buschoff, Eric Schulz</p>	8375
<p>Local Patterns Generalize Better for Novel Anomalies</p> <p>Yalong Jiang</p>	8412
<p>Walk the Talk? Measuring the Faithfulness of Large Language Model Explanations</p> <p>Katie Matton, Robert Ness, John Gutttag, Emre Kiciman</p>	8439
<p>REVISITING MULTI-PERMUTATION EQUIVARIANCE THROUGH THE LENS OF IRREDUCIBLE REPRESENTATIONS</p> <p>Yonatan Sverdlov, Ido Springer, Nadav Dym</p>	8505
<p>DUALFormer: Dual Graph Transformer</p> <p>Jiaming Zhuo, Yuwei Liu, Yintong Lu, Ziyi Ma, Kun Fu, Chuan Wang, Yuanfang Guo, Zhen Wang, Xiaochun Cao, Liang Yang</p>	8528
<p>STRAP: Robot Sub-Trajectory Retrieval for Augmented Policy Learning</p> <p>Marius Memmel, Jacob Berg, Bingqing Chen, Abhishek Gupta, Jonathan Francis</p>	8551

<p>Provable Weak-To-Strong Generalization Via Benign Overfitting</p> <p><i>David Xing Wu, Anant Sahai</i></p>	8571
<p>IterComp: Iterative Composition-Aware Feedback Learning from Model Gallery for Text-To-Image Generation</p> <p><i>Xinchen Zhang, Ling Yang, Guohao Li, Yaqi Cai, Xie Jiake, Yong Tang, Yuju Yang, Mengdi Wang, Bin Cui</i></p>	8608
<p>Test-Time Ensemble Via Linear Mode Connectivity: A Path to Better Adaptation</p> <p><i>Byungjai Kim, Chanho Ahn, Wissam J. Baddar, Kikyung Kim, Huijin Lee, Saehyun Ahn, Seungju Han, Sungjoo Suh, Eunho Yang</i></p>	8629
<p>Can Transformers Do Enumerative Geometry?</p> <p><i>Baran Hashemi, Roderic Guigo Corominas, Alessandro Giacchetto</i></p>	8655
<p>Exploring the Loss Landscape of Regularized Neural Networks Via Convex Duality</p> <p><i>Sungyoon Kim, Aaron Mishkin, Mert Pilanci</i></p>	8676
<p>Scaling Transformers for Low-Bitrate High-Quality Speech Coding</p> <p><i>Julian D Parker, Anton Smirnov, Jordi Pons, Cj Carr, Zack Zukowski, Zach Evans, Xubo Liu</i></p>	8730
<p>IFORMER: INTEGRATING CONVNET AND TRANSFORMER FOR MOBILE APPLICATION</p> <p><i>Chuanyang Zheng</i></p>	8755
<p>Nova: Generative Language Models for Assembly Code with Hierarchical Attention and Contrastive Learning</p> <p><i>Nan Jiang, Chengxiao Wang, Kevin Liu, Xiangzhe Xu, Lin Tan, Xiangyu Zhang, Petr Babkin</i></p>	8770
<p>HyperFace: Generating Synthetic Face Recognition Datasets by Exploring Face Embedding Hypersphere</p> <p><i>Hatef Otroshi Shahreza, Sébastien Marcel</i></p>	8792
<p>MLLM Can See? Dynamic Correction Decoding for Hallucination Mitigation</p> <p><i>Chenxi Wang, Xiang Chen, Ningyu Zhang, Bozhong Tian, Haoming Xu, Shumin Deng, Huajun Chen</i></p>	8811
<p>Solving Hidden Monotone Variational Inequalities with Surrogate Losses</p> <p><i>Ryan D'Orazio, Danilo Vucetic, Zichu Liu, Junhyung Lyle Kim, Ioannis Mitliagkas, Gauthier Gidel</i></p>	8836
<p>COME: Test-Time Adaption by Conservatively Minimizing Entropy</p> <p><i>Qingyang Zhang, Yatao Bian, Xinke Kong, Peilin Zhao, Changqing Zhang</i></p>	8865
<p>Bayesian Analysis of Combinatorial Gaussian Process Bandits</p> <p><i>Jack Sandberg, Niklas Åkerblom, Morteza Haghiri Chehreghani</i></p>	8895
<p>A Probabilistic Perspective on Unlearning and Alignment for Large Language Models</p> <p><i>Yan Scholten, Stephan Günnemann, Leo Schwinn</i></p>	8929

<p>Deep Networks Learn Features from Local Discontinuities in the Label Function</p> <p><i>Prithaj Banerjee, Harish Guruprasad Ramaswamy, Mahesh Lorik Yadav, Chandra Shekar Lakshminarayanan</i></p>	8946
<p>On the Completeness of Invariant Geometric Deep Learning Models</p> <p><i>Zian Li, Xiyuan Wang, Shijia Kang, Muhan Zhang</i></p>	8969
<p>Point Cluster: A Compact Message Unit for Communication-Efficient Collaborative Perception</p> <p><i>Zihan Ding, Jiahui Fu, Si Liu, Hongyu Li, Siheng Chen, Hongsheng Li, Shifeng Zhang, Xu Zhou</i></p>	9012
<p>Elucidating the Preconditioning in Consistency Distillation</p> <p><i>Kaiwen Zheng, Guande He, Jianfei Chen, Fan Bao, Jun Zhu</i></p>	9029
<p>Binary Losses for Density Ratio Estimation</p> <p><i>Werner Zellinger</i></p>	9046
<p>Graph-Guided Scene Reconstruction from Images with 3D Gaussian Splatting</p> <p><i>Chong Cheng, Gaochao Song, Yiyang Yao, Qinzheng Zhou, Gangjian Zhang, Hao Wang</i></p>	9079
<p>How to Probe: Simple Yet Effective Techniques for Improving Post-Hoc Explanations</p> <p><i>Siddhartha Gairola, Moritz Böhle, Francesco Locatello, Bernt Schiele</i></p>	9099
<p>Neural Fluid Simulation on Geometric Surfaces</p> <p><i>Haoxiang Wang, Tao Yu, Hui Qiao, Qionghai Dai</i></p>	9154
<p>Measuring Non-Adversarial Reproduction of Training Data in Large Language Models</p> <p><i>Michael Aerni, Javier Rando, Edoardo Debenedetti, Nicholas Carlini, Daphne Ippolito, Florian Tramèr</i></p>	9182
<p>PnP-Flow: Plug-And-Play Image Restoration with Flow Matching</p> <p><i>Ségolène Tiffany Martin, Anne Gagneux, Paul Hagemann, Gabriele Steidl</i></p>	9225
<p>HR-Extreme: A High-Resolution Dataset for Extreme Weather Forecasting</p> <p><i>Nian Ran, Peng Xiao, Yue Wang, Wesley Shi, Jianxin Lin, Qi Meng, Richard Allmendinger</i></p>	9252
<p>RegMix: Data Mixture as Regression for Language Model Pre-Training</p> <p><i>Qian Liu, Xiaosen Zheng, Niklas Muennighoff, Guangtao Zeng, Longxu Dou, Tianyu Pang, Jing Jiang, Min Lin</i></p>	9269
<p>HASARD: A Benchmark for Vision-Based Safe Reinforcement Learning in Embodied Agents</p> <p><i>Tristan Tomilin, Meng Fang, Mykola Pechenizkiy</i></p>	9304
<p>Enhancing Compositional Text-To-Image Generation with Reliable Random Seeds</p> <p><i>Shuangqi Li, Hieu Le, Jingyi Xu, Mathieu Salzmann</i></p>	9337
<p>LoR-VP: Low-Rank Visual Prompting for Efficient Vision Model Adaptation</p> <p><i>Can Jin, Ying Li, Mingyu Zhao, Shiyu Zhao, Zhenting Wang, Xiaoxiao He, Ligong Han, Tong Che, Dimitris N. Metaxas</i></p>	9359

<p>Tool Decoding: A Plug-And-Play Approach to Enhancing Language Models for Tool Usage</p> <p>Chenheng Zhang, Mingqing Xiao, Yifei Wang, Lizhe Fang, Haihan Zhang, Zhouchen Lin</p>	9377
<p>Can One Modality Model Synergize Training of Other Modality Models?</p> <p>Jae-Jun Lee, Sung Whan Yoon</p>	9398
<p>MathGAP: Out-Of-Distribution Evaluation on Problems with Arbitrarily Complex Proofs</p> <p>Andreas Opedal, Haruki Shirakami, Bernhard Schölkopf, Abulhair Saparov, Mrinmaya Sachan</p>	9421
<p>Asynchronous Federated Reinforcement Learning with Policy Gradient Updates: Algorithm Design and Convergence Analysis</p> <p>Guangchen Lan, Dong-Jun Han, Abolfazl Hashemi, Vaneet Aggarwal, Christopher Brinton</p>	9444
<p>MAGNet: Motif-Agnostic Generation of Molecules from Scaffolds</p> <p>Leon Hetzel, Johanna Sommer, Bastian Rieck, Fabian J Theis, Stephan Günemann</p>	9475
<p>Intent3D: 3D Object Detection in RGB-D Scans Based on Human Intention</p> <p>Weitai Kang, Mengxue Qu, Jyoti Kini, Yunchao Wei, Mubarak Shah, Yan Yan</p>	9500
<p>Affine Steerable Equivariant Layer for Canonicalization of Neural Networks</p> <p>Yikang Li, Yeqing Qiu, Yuxuan Chen, Zhouchen Lin</p>	9522
<p>Prioritized Generative Replay</p> <p>Renhao Wang, Kevin Frans, Pieter Abbeel, Sergey Levine, Alexei A Efros</p>	9543
<p>Everything, Everywhere, All at Once: Is Mechanistic Interpretability Identifiable?</p> <p>Maxime Méloux, Silviu Maniu, François Portet, Maxime Peyrard</p>	9562
<p>Quantized Spike-Driven Transformer</p> <p>Xuerui Qiu, Malu Zhang, Jieyuan Zhang, Wenjie Wei, Honglin Cao, Junsheng Guo, Rui-Jie Zhu, Yimeng Shan, Yang Yang, Haizhou Li</p>	9586
<p>Energy-Based Backdoor Defense Against Federated Graph Learning</p> <p>Guancheng Wan, Zitong Shi, Wenke Huang, Guibin Zhang, Dacheng Tao, Mang Ye</p>	9611
<p>A Stochastic Approach to the Subset Selection Problem Via Mirror Descent</p> <p>Dan Greenstein, Elazar Gershuni, Ilan Ben-Bassat, Yaroslav Fyodorov, Ran Moshe, Fiana Raiber, Alex Shtoff, Oren Somekh, Nadav Hallak</p>	9641
<p>Solving Differential Equations with Constrained Learning</p> <p>Viggo Moro, Luiz F. O. Chamon</p>	9685
<p>Holistic Reasoning with Long-Context LMs: A Benchmark for Database Operations on Massive Textual Data</p> <p>Seiji Maekawa, Hayate Iso, Nikita Bhutani</p>	9728
<p>DEBORA: Efficient Bilevel Optimization-Based low-Rank Adaptation</p> <p>Emanuele Zangrando, Sara Venturini, Francesco Rinaldi, Francesco Tudisco</p>	9753

Learning Under Temporal Label Noise <i>Sujay Nagaraj, Walter Gerych, Sana Tonekaboni, Anna Goldenberg, Berk Ustun, Thomas Hartvigsen</i>	9773
ComaDICE: Offline Cooperative Multi-Agent Reinforcement Learning with Stationary Distribution Shift Regularization <i>The Viet Bui, Thanh Hong Nguyen, Tien Anh Mai</i>	9800
KoonPro: A Variance-Aware Koopman Probabilistic Model Enhanced by Neural Process for Time Series Forecasting <i>Ronghua Zheng, Hanru Bai, Weiyang Ding</i>	9815
Charting the Design Space of Neural Graph Representations for Subgraph Matching <i>Vaibhav Raj, Indradyumna Roy, Ashwin Ramachandran, Soumen Chakrabarti, Abir De</i>	9840
Score-Based Free-Form Architectures for High-Dimensional Fokker-Planck Equations <i>Feng Liu, Faguo Wu, Xiao Zhang</i>	9888
Breaking Neural Network Scaling Laws with Modularity <i>Akhilan Boopathy, Sunshine Jiang, William Yue, Jaedong Hwang, Abhiram Iyer, Ila R Fiete</i>	9917
GraphEval: A Lightweight Graph-Based LLM Framework for Idea Evaluation <i>Tao Feng, Yihang Sun, Jiaxuan You</i>	9946
You Only Prune Once: Designing Calibration-Free Model Compression with Policy Learning <i>Ayan Sengupta, Siddhant Chaudhary, Tanmoy Chakraborty</i>	9975
TTVD: Towards a Geometric Framework for Test-Time Adaptation Based on Voronoi Diagram <i>Mingxi Lei, Chunwei Ma, Meng Ding, Yufan Zhou, Ziyun Huang, Jinhui Xu</i>	9999
PhiNets: Brain-Inspired Non-Contrastive Learning Based on Temporal Prediction Hypothesis <i>Satoki Ishikawa, Makoto Yamada, Han Bao, Yuki Takezawa</i>	10022
SD-LoRA: Scalable Decoupled Low-Rank Adaptation for Class Incremental Learning <i>Yichen Wu, Hongming Piao, Long-Kai Huang, Renzhen Wang, Wanhua Li, Hanspeter Pfister, Deyu Meng, Kede Ma, Ying Wei</i>	10049
NeuralPlane: Structured 3D Reconstruction in Planar Primitives with Neural Fields <i>Hanqiao Ye, Yuzhou Liu, Yangdong Liu, Shuhan Shen</i>	10066
Captured by Captions: On Memorization and Its Mitigation in CLIP Models <i>Wenhao Wang, Adam Dziedzic, Grace C. Kim, Michael Backes, Franziska Boenisch</i>	10090
Contextualizing Biological Perturbation Experiments Through Language <i>Menghua Wu, Russell Littman, Jacob Levine, Lin Qiu, Tommaso Biancalani, David Richmond, Jan-Christian Huetter</i>	10114
Satisficing Regret Minimization in Bandits <i>Qing Feng, Tianyi Ma, Ruihao Zhu</i>	10141

GotenNet: Rethinking Efficient 3D Equivariant Graph Neural Networks Sarp Aykent, Tian Xia	10153
Steering Large Language Models Between Code Execution and Textual Reasoning Yongchao Chen, Harsh Jhamtani, Srinagesh Sharma, Chuchu Fan, Chi Wang	10178
Bonsai: Gradient-Free Graph Condensation for Node Classification Mridul Gupta, Samyak Jain, Vansh Ramani, Hariprasad Kodamana, Sayan Ranu	10210
Manifold Constraint Reduces Exposure Bias in Accelerated Diffusion Sampling Yuzhe Yao, Jun Chen, Zeyi Huang, Haonan Lin, Mengmeng Wang, Guang Dai, Jingdong Wang	10231
Precise Localization of Memories: A Fine-Grained Neuron-Level Knowledge Editing Technique for LLMs Haowen Pan, Xiaozhi Wang, Yixin Cao, Zenglin Shi, Xun Yang, Juanzi Li, Meng Wang	10268
VVC-Gym: A Fixed-Wing UAV Reinforcement Learning Environment for Multi-Goal Long-Horizon Problems Xudong Gong, Feng Dawei, Kele Xu, Weijia Wang, Zhangjun Sun, Xing Zhou, Si Zheng, Bo Ding, Huaimin Wang	10289
Expressivity of Neural Networks with Random Weights and Learned Biases Ezekiel Williams, Alexandre Payeur, Avery Hee-Woon Ryoo, Thomas Jiralerspong, Matthew G Perich, Luca Mazzucato, Guillaume Lajoie	10313
Unlearn and Burn: Adversarial Machine Unlearning Requests Destroy Model Accuracy Yangsibo Huang, Daogao Liu, Lynn Chua, Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, Milad Nasr, Amer Sinha, Chiyuan Zhang	10339
Adversarial Policy Optimization for Offline Preference-Based Reinforcement Learning Hyungkyu Kang, Min-Hwan Oh	10360
Learning Task Belief Similarity with Latent Dynamics for Meta-Reinforcement Learning Menglong Zhang, Fuyuan Qian, Quanying Liu	10395
Time After Time: Deep-Q Effect Estimation for Interventions on When and What to Do Yoav Wald, Mark Goldstein, Yonathan Efroni, Wouter A. C. Van Amsterdam, Rajesh Ranganath	10422
DPLM-2: A Multimodal Diffusion Protein Language Model Xinyou Wang, Zaixiang Zheng, Fei Ye, Dongyu Xue, Shujian Huang, Quanquan Gu	10436
COPER: Correlation-Based Permutations for Multi-View Clustering Ran Eisenberg, Jonathan Svirsky, Ofir Lindenbaum	10464
Doubly Optimal Policy Evaluation for Reinforcement Learning Shuze Liu, Claire Chen, Shangtong Zhang	10490
BlendRL: A Framework for Merging Symbolic and Neural Policy Learning Hikaru Shindo, Quentin Delfosse, Devendra Singh Dhami, Kristian Kersting	10518

<p>Metamizer: A Versatile Neural Optimizer for Fast and Accurate Physics Simulations</p> <p>Nils Wandel, Stefan Schulz, Reinhard Klein</p>	10550
<p>Privacy Auditing of Large Language Models</p> <p>Ashwinee Panda, Xinyu Tang, Christopher A. Choquette-Choo, Milad Nasr, Prateek Mittal</p>	10573
<p>Training Free Guided Flow-Matching with Optimal Control</p> <p>Luran Wang, Chaoran Cheng, Yizhen Liao, Yanru Qu, Ge Liu</p>	10590
<p>Not-So-Optimal Transport Flows for 3D Point Cloud Generation</p> <p>Ka-Hei Hui, Chao Liu, Xiaohui Zeng, Chi-Wing Fu, Arash Vahdat</p>	10623
<p>Painting with Words: Elevating Detailed Image Captioning with Benchmark and Alignment Learning</p> <p>Qinghao Ye, Xianhan Zeng, Fu Li, Chunyuan Li, Haoqi Fan</p>	10647
<p>Pyramidal Flow Matching for Efficient Video Generative Modeling</p> <p>Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song, Yadong Mu, Zhouchen Lin</p>	10672
<p>MMEgo: Towards Building Egocentric Multimodal LLMs for Video QA</p> <p>Hanrong Ye, Haotian Zhang, Erik Daxberger, Lin Chen, Zongyu Lin, Yanghao Li, Bowen Zhang, Haoxuan You, Dan Xu, Zhe Gan, Jiasen Lu, Yinfei Yang</p>	10697
<p>Model Merging with SVD to Tie the Knots</p> <p>George Stoica, Pratik Ramesh, Boglarka Ecsedi, Leshem Choshen, Judy Hoffman</p>	10716
<p>Scaling Up the Banded Matrix Factorization Mechanism for Large Scale Differentially Private ML</p> <p>Ryan McKenna</p>	10735
<p>Mutual Reasoning Makes Smaller LLMs Stronger Problem-Solver</p> <p>Zhenting Qi, Mingyuan Ma, Jiahang Xu, Li Lyna Zhang, Fan Yang, Mao Yang</p>	10757
<p>Diverse Policies Recovering Via Pointwise Mutual Information Weighted Imitation Learning</p> <p>Hanlin Yang, Jian Yao, Weiming Liu, Qing Wang, Hanmin Qin, Kong Hansheng, Kirk Tang, Jiechao Xiong, Chao Yu, Kai Li, Junliang Xing, Hongwu Chen, Juchao Zhuo, Qiang Fu, Yang Wei, Haobo Fu</p>	10777
<p>How Discrete and Continuous Diffusion Meet: Comprehensive Analysis of Discrete Diffusion Models Via a Stochastic Integral Framework</p> <p>Yinuo Ren, Haoxuan Chen, Grant M. Rotskoff, Lexing Ying</p>	10796
<p>Interference Among First-Price Pricing Equilibria: A Bias and Variance Analysis</p> <p>Luofeng Liao, Christian Kroer, Sergei Leonenkov, Okke Schrijvers, Liang Shi, Nicolas Stier Moses, Congshan Zhang</p>	10834
<p>ManiSkill-HAB: A Benchmark for Low-Level Manipulation in Home Rearrangement Tasks</p> <p>Arth Shukla, Stone Tao, Hao Su</p>	10856

<p>LoRA-X: Bridging Foundation Models with Training-Free Cross-Model Adaptation</p> <p>Farzad Farhadzadeh, Debasmit Das, Shubhankar Borse, Fatih Porikli</p>	10886
<p>Curriculum-Aware Training for Discriminating Molecular Property Prediction Models</p> <p>Hansi Yang, Quanming Yao, James Kwok</p>	10910
<p>LLM Unlearning Via Loss Adjustment with Only Forget Data</p> <p>Yaxuan Wang, Jiaheng Wei, Chris Yuhao Liu, Jinlong Pang, Quan Liu, Ankit Shah, Yujia Bao, Yang Liu, Wei Wei</p>	10929
<p>Variational Diffusion Posterior Sampling with Midpoint Guidance</p> <p>Badr Moufad, Yazid Janati, Lisa Bedin, Alain Oliviero Durmus, Randal Douc, Eric Moulines, Jimmy Olsson</p>	10958
<p>Towards Hierarchical Rectified Flow</p> <p>Yichi Zhang, Yici Yan, Alex Schwing, Zhizhen Zhao</p>	11003
<p>HyperDAS: Towards Automating Mechanistic Interpretability with Hypernetworks</p> <p>Jiuding Sun, Jing Huang, Sidharth Baskaran, Karel D'Oosterlinck, Christopher Potts, Michael Sklar, Atticus Geiger</p>	11025
<p>Preference Diffusion for Recommendation</p> <p>Shuo Liu, An Zhang, Guoqing Hu, Hong Qian, Tat-Seng Chua</p>	11043
<p>MTU-Bench: A Multi-Granularity Tool-Use Benchmark for Large Language Models</p> <p>Pei Wang, Yanan Wu, Noah Wang, Jiaheng Liu, Xiaoshuai Song, Z. Y. Peng, Ken Deng, Chenchen Zhang, Jiakaiwang, Junran Peng, Ge Zhang, Hangyu Guo, Zhaoxiang Zhang, Wenbo Su, Bo Zheng</p>	11081
<p>MotherNet: Fast Training and Inference Via Hyper-Network Transformers</p> <p>Andreas C Mueller, Carlo A Curino, Raghu Ramakrishnan</p>	11120
<p>Token-Supervised Value Models for Enhancing Mathematical Problem-Solving Capabilities of Large Language Models</p> <p>Jung Hyun Lee, June Yong Yang, Byeongho Heo, Dongyoon Han, Kyungsu Kim, Eunho Yang, Kang Min Yoo</p>	11141
<p>Reading Your Heart: Learning ECG Words and Sentences Via Pre-Training ECG Language Model</p> <p>Jiarui Jin, Haoyu Wang, Hongyan Li, Jun Li, Jiahui Pan, Shenda Hong</p>	11161
<p>BrainUICL: An Unsupervised Individual Continual Learning Framework for EEG Applications</p> <p>Yangxuan Zhou, Sha Zhao, Jiquan Wang, Haiteng Jiang, Shijian Li, Tao Li, Gang Pan</p>	11182
<p>Benign Overfitting in Out-Of-Distribution Generalization of Linear Models</p> <p>Shange Tang, Jiayun Wu, Jianqing Fan, Chi Jin</p>	11204
<p>Hummingbird: High Fidelity Image Generation Via Multimodal Context Alignment</p> <p>Minh-Quan Le, Gaurav Mittal, Tianjian Meng, A S M Iftekhar, Vishwas Suryanarayanan, Barun Patra, Dimitris Samaras, Mei Chen</p>	11258

<p>HD-Painter: High-Resolution and Prompt-Faithful Text-Guided Image Inpainting with Diffusion Models</p> <p>Hayk Manukyan, Andranik Sargsyan, Barsegh Atanyan, Zhangyang Wang, Shant Navasardyan, Humphrey Shi</p>	11283
<p>Data Taggants: Dataset Ownership Verification Via Harmless Targeted Data Poisoning</p> <p>Wassim Bouaziz, Nicolas Usunier, El-Mahdi El-Mhamdi</p>	11313
<p>NutriBench: A Dataset for Evaluating Large Language Models in Nutrition Estimation from Meal Descriptions</p> <p>Mehak Preet Dhaliwal, Andong Hua, Laya Pullela, Ryan Burke, Yao Qin</p>	11332
<p>Linear Transformer Topological Masking with Graph Random Features</p> <p>Isaac Reid, Kumar Avinava Dubey, Deepali Jain, William F Whitney, Amr Ahmed, Joshua Ainslie, Alex Bewley, Mithun George Jacob, Aranyak Mehta, David Rendleman, Connor Schenck, Richard E. Turner, René Wagner, Adrian Weller, Krzysztof Marcin Choromanski</p>	11356
<p>CSA: Data-Efficient Mapping of Unimodal Features to Multimodal Features</p> <p>Po-Han Li, Sandeep P. Chinchali, Ufuk Topcu</p>	11381
<p>Safety Alignment Should Be Made More than Just a Few Tokens Deep</p> <p>Xiangyu Qi, Ashwinee Panda, Kaifeng Lyu, Xiao Ma, Subhrajit Roy, Ahmad Beirami, Prateek Mittal, Peter Henderson</p>	11401
<p>MTSAM: Multi-Task Fine-Tuning for Segment Anything Model</p> <p>Xuehao Wang, Zhan Zhuang, Feiyang Ye, Yu Zhang</p>	11432
<p>STAR: Stability-Inducing Weight Perturbation for Continual Learning</p> <p>Masih Eskandar, Tooba Imtiaz, Davin Hill, Zifeng Wang, Jennifer Dy</p>	11454
<p>Unsupervised Multiple Kernel Learning for Graphs Via Ordinality Preservation</p> <p>Yan Sun, Stanley Kok</p>	11470
<p>Answer, Assemble, Ace: Understanding How LMs Answer Multiple Choice Questions</p> <p>Sarah Wiegrefe, Oyvind Tafjord, Yonatan Belinkov, Hannaneh Hajishirzi, Ashish Sabharwal</p>	11512
<p>Physiome-ODE: A Benchmark for Irregularly Sampled Multivariate Time-Series Forecasting Based on Biological ODEs</p> <p>Christian Klötergens, Vijaya Krishna Yalavarthi, Randolph Scholz, Maximilian Stubbemann, Stefan Born, Lars Schmidt-Thieme</p>	11538
<p>Towards Understanding Why Label Smoothing Degrades Selective Classification and How to Fix it</p> <p>Guoxuan Xia, Olivier Laurent, Gianni Franchi, Christos-Savvas Bouganis</p>	11566
<p>Dynamic Negative Guidance of Diffusion Models</p> <p>Felix Koulischer, Johannes Deleu, Gabriel Raya, Thomas Demeester, Luca Ambrogioni</p>	11600

<p>Identification of Intermittent Temporal Latent Process</p> <p>Yuke Li, Yujia Zheng, Guangyi Chen, Kun Zhang, Heng Huang</p>	11630
<p>Learning How Hard to Think: Input-Adaptive Allocation of LM Computation</p> <p>Mehul Damani, Idan Shenfeld, Andi Peng, Andreea Bobu, Jacob Andreas</p>	11650
<p>GeoX: Geometric Problem Solving Through Unified Formalized Vision-Language Pre-Training</p> <p>Renqiu Xia, Mingsheng Li, Hancheng Ye, Wenjie Wu, Hongbin Zhou, Jiakang Yuan, Tianshuo Peng, Xinyu Cai, Xiangchao Yan, Bin Wang, Conghui He, Botian Shi, Tao Chen, Junchi Yan, Bo Zhang</p>	11670
<p>FedLWS: Federated Learning with Adaptive Layer-Wise Weight Shrinking</p> <p>Changlong Shi, Jinmeng Li, He Zhao, Dan Dan Guo, Yi Chang</p>	11689
<p>MAD-TD: Model-Augmented Data Stabilizes High Update Ratio RL</p> <p>Claas A Voelcker, Marcel Hussing, Eric Eaton, Amir-Massoud Farahmand, Igor Gilitschenski</p>	11714
<p>Transformers Are Universal In-Context Learners</p> <p>Takashi Furuya, Maarten V. De Hoop, Gabriel Peyré</p>	11745
<p>MLE-Bench: Evaluating Machine Learning Agents on Machine Learning Engineering</p> <p>Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, Aleksander Madry, Lilian Weng</p>	11771
<p>High-Dimension Prototype is a Better Incremental Object Detection Learner</p> <p>Yanjie Wang, Liqun Chen, Tianming Zhao, Tao Zhang, Guodong Wang, Luxin Yan, Sheng Zhong, Jiahuan Zhou, Xu Zou</p>	11800
<p>SPA: 3D Spatial-Awareness Enables Effective Embodied Representation</p> <p>Haoyi Zhu, Honghui Yang, Yating Wang, Jiange Yang, Limin Wang, Tong He</p>	11818
<p>Gap-Dependent Bounds for Q-Learning Using Reference-Advantage Decomposition</p> <p>Zhong Zheng, Haochen Zhang, Lingzhou Xue</p>	11849
<p>UniGS: Unified Language-Image-3D Pretraining with Gaussian Splatting</p> <p>Haoyuan Li, Zhou Yanpeng, Tao Tang, Jifei Song, Yihan Zeng, Michael Kampffmeyer, Hang Xu, Xiaodan Liang</p>	11902
<p>POGEMA: A Benchmark Platform for Cooperative Multi-Agent Pathfinding</p> <p>Alexey Skrynnik, Anton Andreychuk, Anatolii Borzilov, Alexander Chernyavskiy, Konstantin Yakovlev, Aleksandr Panov</p>	11923
<p>Inference Optimal VLMs Need Fewer Visual Tokens and More Parameters</p> <p>Kevin Li, Sachin Goyal, João D. Semedo, J Zico Kolter</p>	11954
<p>Compositional 4D Dynamic Scenes Understanding with Physics Priors for Video Question Answering</p> <p>Xingrui Wang, Wufei Ma, Angtian Wang, Shuo Chen, Adam Kortylewski, Alan Yuille</p>	11972

<p>Divergence-Enhanced Knowledge-Guided Context Optimization for Visual-Language Prompt Tuning</p> <p><i>Yilun Li, Miaomiao Cheng, Xu Han, Wei Song</i></p>	11985
<p>A Multiscale Frequency Domain Causal Framework for Enhanced Pathological Analysis</p> <p><i>Xiaoyu Cui, Weixing Chen, Jiandong Su</i></p>	11997
<p>STBLLM: Breaking the 1-Bit Barrier with Structured Binary LLMs</p> <p><i>Peijie Dong, Lujun Li, Yuedong Zhong, Dayou Du, Ruibo Fan, Yuhan Chen, Zhenheng Tang, Qiang Wang, Wei Xue, Yike Guo, Xiaowen Chu</i></p>	12014
<p>Episodic Memories Generation and Evaluation Benchmark for Large Language Models</p> <p><i>Alexis Huet, Zied Ben Houidi, Dario Rossi</i></p>	12040
<p>Towards Foundation Models for Mixed Integer Linear Programming</p> <p><i>Sirui Li, Janardhan Kulkarni, Ishai Menache, Cathy Wu, Beibin Li</i></p>	12089
<p>ScienceAgentBench: Toward Rigorous Assessment of Language Agents for Data-Driven Scientific Discovery</p> <p><i>Ziru Chen, Shijie Chen, Yuting Ning, Qianheng Zhang, Boshi Wang, Botao Yu, Yifei Li, Zeyi Liao, Chen Wei, Zitong Lu, Vishal Dey, Mingyi Xue, Frazier N. Baker, Benjamin Burns, Daniel Adu-Ampratwum, Xuhui Huang, Xia Ning, Song Gao, Yu Su, Huan Sun</i></p>	12138
<p>TabWak: A Watermark for Tabular Diffusion Models</p> <p><i>Chaoyi Zhu, Jiayi Tang, Jeroen M. Galjaard, Pin-Yu Chen, Robert Birke, Cornelis Bos, Lydia Y. Chen</i></p>	12195
<p>Mini-Monkey: Alleviating the Semantic Sawtooth Effect for Lightweight MLLMs Via Complementary Image Pyramid</p> <p><i>Mingxin Huang, Yuliang Liu, Dingkang Liang, Lianwen Jin, Xiang Bai</i></p>	12223
<p>LaGeM: A Large Geometry Model for 3D Representation Learning and Diffusion</p> <p><i>Biao Zhang, Peter Wonka</i></p>	12243
<p>Wayward Concepts in Multimodal Models</p> <p><i>Brandon Trabucco, Max A Gurinas, Kyle Doherty, Russ Salakhutdinov</i></p>	12261
<p>Multi-Label Test-Time Adaptation with Bound Entropy Minimization</p> <p><i>Xiangyu Wu, Feng Yu, Yang Yang, Qing-Guo Chen, Jianfeng Lu</i></p>	12287
<p>Inference-Aware Fine-Tuning for Best-Of-N Sampling in Large Language Models</p> <p><i>Yinlam Chow, Guy Tennenholtz, Izzeddin Gur, Vincent Zhuang, Bo Dai, Aviral Kumar, Rishabh Agarwal, Sridhar Thiagarajan, Craig Boutilier, Aleksandra Faust</i></p>	12304
<p>On the Adversarial Risk of Test Time Adaptation: An Investigation into Realistic Test-Time Data Poisoning</p> <p><i>Yongyi Su, Yushu Li, Nanqing Liu, Kui Jia, Xulei Yang, Chuan-Sheng Foo, Xun Xu</i></p>	12328
<p>When Attention Sink Emerges in Language Models: An Empirical View</p> <p><i>Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, Min Lin</i></p>	12353

MADGEN: Mass-Spec Attends to De Novo Molecular Generation <i>Yinkai Wang, Xiaohui Chen, Liping Liu, Soha Hassoun</i>	12384
Bad-PFL: Exploiting Backdoor Attacks Against Personalized Federated Learning <i>Mingyuan Fan, Zhanyi Hu, Fuyi Wang, Cen Chen</i>	12400
Small-To-Large Generalization: Training Data Influences Models Consistently Across Scale <i>Alaa Khaddaj, Logan Engstrom, Aleksander Madry</i>	12428
Graph Sparsification Via Mixture of Graphs <i>Guibin Zhang, Xiangguo Sun, Yanwei Yue, Chonghe Jiang, Kun Wang, Tianlong Chen, Shirui Pan</i>	12468
Graph Neural Ricci Flow: Evolving Feature from a Curvature Perspective <i>Jialong Chen, Bowen Deng, Zhen Wang, Chuan Chen, Zibin Zheng</i>	12497
Regretful Decisions Under Label Noise <i>Sujay Nagaraj, Yang Liu, Flavio Calmon, Berk Ustun</i>	12526
GRAIN: Exact Graph Reconstruction from Gradients <i>Maria Drencheva, Ivo Petrov, Maximilian Baader, Dimitar Iliev Dimitrov, Martin Vechev</i>	12546
Training One-Dimensional Graph Neural Networks is NP-Hard <i>Robert Galian, Mathis Rocton, Simon Wietheger</i>	12571
Geometry-Aware RL for Manipulation of Varying Shapes and Deformable Objects <i>Tai Hoang, Huy Le, Philipp Becker, Vien Anh Ngo, Gerhard Neumann</i>	12588
What Makes a Good Diffusion Planner for Decision Making? <i>Haofei Lu, Dongqi Han, Yifei Shen, Dongsheng Li</i>	12618
Object-Centric Pretraining Via Target Encoder Bootstrapping <i>Nikola Đukic, Tim Lebailly, Tinne Tuytelaars</i>	12648
PuzzleFusion++: Auto-Agglomerative 3D Fracture Assembly by Denoise and Verify <i>Zhengqing Wang, Jiacheng Chen, Yasutaka Furukawa</i>	12662
MIA-Bench: Towards Better Instruction Following Evaluation of Multimodal LLMs <i>Yusu Qian, Hanrong Ye, Jean-Philippe Fauconnier, Peter Grasch, Yinfei Yang, Zhe Gan</i>	12686
Lawma: The Power of Specialization for Legal Annotation <i>Ricardo Dominguez-Olmedo, Vedant Nanda, Rediet Abebe, Stefan Bechtold, Christoph Engel, Jens Frankenreiter, Krishna P. Gummadi, Moritz Hardt, Michael Livermore</i>	12707
Bandit Learning in Matching Markets with Indifference <i>Fang Kong, Jingqi Tang, Mingzhu Li, Pinyan Lu, John C. S. Lui, Shuai Li</i>	12752
WardropNet: Traffic Flow Predictions Via Equilibrium-Augmented Learning <i>Kai Jungel, Dario Paccagnan, Axel Parmentier, Maximilian Schiffer</i>	12771

<p>Manifold Induced Biases for Zero-Shot and Few-Shot Detection of Generated Images</p> <p><i>Jonathan Brokman, Amit Giloni, Omer Hofman, Roman Vainshtein, Hisashi Kojima, Guy Gilboa</i></p>	12803
<p>Prompting Fairness: Integrating Causality to Debias Large Language Models</p> <p><i>Jingling Li, Zeyu Tang, Xiaoyu Liu, Peter Spirtes, Kun Zhang, Liu Leqi, Yang Liu</i></p>	12829
<p>Visually Consistent Hierarchical Image Classification</p> <p><i>Seulki Park, Youren Zhang, Stella X. Yu, Sara Beery, Jonathan Huang</i></p>	12853
<p>SC-OmniGS: Self-Calibrating Omnidirectional Gaussian Splatting</p> <p><i>Huajian Huang, Yingshu Chen, Longwei Li, Hui Cheng, Tristan Braud, Yajie Zhao, Sai-Kit Yeung</i></p>	12876
<p>CoTFormer: A Chain of Thought Driven Architecture with Budget-Adaptive Computation Cost at Inference</p> <p><i>Amirkeivan Mohtashami, Matteo Pagliardini, Martin Jaggi</i></p>	12897
<p>Simplifying Deep Temporal Difference Learning</p> <p><i>Matteo Gallici, Mattie Fellows, Benjamin Ellis, Bartomeu Pou, Ivan Masmitja, Jakob Nicolaus Foerster, Mario Martin</i></p>	12915
<p>DON'T STOP ME NOW: EMBEDDING BASED SCHEDULING FOR LLMS</p> <p><i>Rana Shahout, Eran Malach, Chunwei Liu, Weifan Jiang, Minlan Yu, Michael Mitzenmacher</i></p>	12958
<p>3DGS-Drag: Dragging Gaussians for Intuitive Point-Based 3D Editing</p> <p><i>Jiahua Dong, Yu-Xiong Wang</i></p>	12982
<p>A Black Swan Hypothesis: The Role of Human Irrationality in AI Safety</p> <p><i>Hyunin Lee, Chanwoo Park, David Abel, Ming Jin</i></p>	12999
<p>Examining Alignment of Large Language Models Through Representative Heuristics: The Case of Political Stereotypes</p> <p><i>Sullam Jeoung, Yubin Ge, Haohan Wang, Jana Diesner</i></p>	13033
<p>Sketching for Convex and Nonconvex Regularized Least Squares with Sharp Guarantees</p> <p><i>Yingzhen Yang, Ping Li</i></p>	13055
<p>CHiP: Cross-Modal Hierarchical Direct Preference Optimization for Multimodal LLMs</p> <p><i>Jinlan Fu, Huangfushenzhen, Hao Fei, Xiaoyu Shen, Bryan Hooi, Xipeng Qiu, See-Kiong Ng</i></p>	13083
<p>Generalization Through Variance: How Noise Shapes Inductive Biases in Diffusion Models</p> <p><i>John Vastola</i></p>	13104
<p>LASeR: Towards Diversified and Generalizable Robot Design with Large Language Models</p> <p><i>Junru Song, Yang Yang, Huan Xiao, Wei Peng, Wen Yao, Feifei Wang</i></p>	13143
<p>Alchemy: Amplifying Theorem-Proving Capability Through Symbolic Mutation</p> <p><i>Shaonan Wu, Shuai Lu, Yeyun Gong, Nan Duan, Ping Wei</i></p>	13173

<p>CityAnchor: City-Scale 3D Visual Grounding with Multi-Modality LLMs</p> <p><i>Jinpeng Li, Haiping Wang, Jiabin Chen, Yuan Liu, Zhiyang Dou, Yuexin Ma, Sibe Yang, Yuan Li, Wenping Wang, Zhen Dong, Bisheng Yang</i></p>	13204
<p>LiFT: Learning to Fine-Tune Via Bayesian Parameter Efficient Meta Fine-Tuning</p> <p><i>Minyoung Kim, Timothy Hospedales</i></p>	13228
<p>TestGenEval: A Real World Unit Test Generation and Test Completion Benchmark</p> <p><i>Kush Jain, Gabriel Synnaeve, Baptiste Roziere</i></p>	13251
<p>Collab: Controlled Decoding Using Mixture of Agents for LLM Alignment</p> <p><i>Souradip Chakraborty, Sujay Bhatt, Udari Madhushani Sehwag, Soumya Suvra Ghosal, Jiahao Qiu, Mengdi Wang, Dinesh Manocha, Furong Huang, Alec Koppel, Sumitra Ganesh</i></p>	13304
<p>Proximal Mapping Loss: Understanding Loss Functions in Crowd Counting & Localization</p> <p><i>Wei Lin, Jia Wan, Antoni B. Chan</i></p>	13326
<p>On the Importance of Language-Driven Representation Learning for Heterogeneous Federated Learning</p> <p><i>Yunlu Yan, Chun-Mei Feng, Wangmeng Zuo, Salman Khan, Yong Liu, Lei Zhu</i></p>	13349
<p>Logically Consistent Language Models Via Neuro-Symbolic Integration</p> <p><i>Diego Calanzone, Stefano Teso, Antonio Vergari</i></p>	13373
<p>Lie Algebra Canonicalization: Equivariant Neural Operators Under Arbitrary Lie Groups</p> <p><i>Zakhar Shumaylov, Peter Zaika, James Rowbottom, Ferdia Sherry, Melanie Weber, Carola-Bibiane Schönlieb</i></p>	13401
<p>Exploring a Principled Framework for Deep Subspace Clustering</p> <p><i>Xianghan Meng, Zhiyuan Huang, Wei He, Xianbiao Qi, Rong Xiao, Chun-Guang Li</i></p>	13450
<p>AI Sandbagging: Language Models Can Strategically Underperform on Evaluations</p> <p><i>Teun Van Der Weij, Felix Hofstätter, Oliver Jaffe, Samuel F. Brown, Francis Rhys Ward</i></p>	13484
<p>Entropy-Based Activation Function Optimization: A Method on Searching Better Activation Functions</p> <p><i>Haoyuan Sun, Zihao Wu, Bo Xia, Pu Chang, Zibin Dong, Yifu Yuan, Yongzhe Chang, Xueqian Wang</i></p>	13522
<p>See What You Are Told: Visual Attention Sink in Large Multimodal Models</p> <p><i>Seil Kang, Jinyeong Kim, Junhyeok Kim, Seong Jae Hwang</i></p>	13547
<p>Modality-Specialized Synergizers for Interleaved Vision-Language Generalists</p> <p><i>Zhiyang Xu, Minqian Liu, Ying Shen, Joy Rimchala, Jiaxin Zhang, Qifan Wang, Yu Cheng, Lifu Huang</i></p>	13575
<p>XFinder: Large Language Models as Automated Evaluators for Reliable Evaluation</p> <p><i>Qingchen Yu, Zifan Zheng, Shichao Song, Zhiyu Li, Feiyu Xiong, Bo Tang, Ding Chen</i></p>	13601
<p>Self-Boosting Large Language Models with Synthetic Preference Data</p> <p><i>Qingxiu Dong, Li Dong, Xingxing Zhang, Zhifang Sui, Furu Wei</i></p>	13644

<p>AnoLLM: Large Language Models for Tabular Anomaly Detection</p> <p><i>Che-Ping Tsai, Ganyu Teng, Phillip Wallis, Wei Ding</i></p>	13668
<p>NextBestPath: Efficient 3D Mapping of Unseen Environments</p> <p><i>Shiyao Li, Antoine Guedon, Clémentin Boittiaux, Shizhe Chen, Vincent Lepetit</i></p>	13699
<p>Hierarchical World Models as Visual Whole-Body Humanoid Controllers</p> <p><i>Nicklas Hansen, Jyothir S V, Vlad Sobal, Yann Lecun, Xiaolong Wang, Hao Su</i></p>	13718
<p>Better Instruction-Following Through Minimum Bayes Risk</p> <p><i>Ian Wu, Patrick Fernandes, Amanda Bertsch, Seungone Kim, Sina Khoshfetrat Pakazad, Graham Neubig</i></p>	13739
<p>Theory on Mixture-Of-Experts in Continual Learning</p> <p><i>Hongbo Li, Sen Lin, Lingjie Duan, Yingbin Liang, Ness Shroff</i></p>	13773
<p>Drama: Mamba-Enabled Model-Based Reinforcement Learning is Sample and Parameter Efficient</p> <p><i>Wenlong Wang, Ivana Dusparic, Yucheng Shi, Ke Zhang, Vinny Cahill</i></p>	13811
<p>Forte : Finding Outliers with Representation Typicality Estimation</p> <p><i>Debargha Ganguly, Warren Richard Morningstar, Andrew Seohwan Yu, Vipin Chaudhary</i></p>	13836
<p>Can Reinforcement Learning Solve Asymmetric Combinatorial-Continuous Zero-Sum Games?</p> <p><i>Yuheng Li, Wang Panpan, Haipeng Chen</i></p>	13874
<p>Do WGANs Succeed Because They Minimize the Wasserstein Distance? Lessons from Discrete Generators</p> <p><i>Ariel Elnekave, Yair Weiss</i></p>	13903
<p>Block-Attention for Efficient Prefilling</p> <p><i>Dongyang Ma, Yan Wang, Tian Lan</i></p>	13925
<p>PPT: Patch Order Do Matters in Time Series Pretext Task</p> <p><i>Jaeho Kim, Kwangryeol Park, Sukmin Yun, Seulki Lee</i></p>	13940
<p>EdgeRunner: Auto-Regressive Auto-Encoder for Artistic Mesh Generation</p> <p><i>Jiaxiang Tang, Zhaoshuo Li, Zekun Hao, Xian Liu, Gang Zeng, Ming-Yu Liu, Qinsheng Zhang</i></p>	13973
<p>Language Models Are Advanced Anonymizers</p> <p><i>Robin Staab, Mark Vero, Mislav Balunovic, Martin Vechev</i></p>	13995
<p>Do Vision-Language Models Represent Space and How? Evaluating Spatial Frame of Reference Under Ambiguities</p> <p><i>Zheyuan Zhang, Fengyuan Hu, Jayjun Lee, Freda Shi, Parisa Kordjamshidi, Joyce Chai, Ziqiao Ma</i></p>	14036
<p>Stiefel Flow Matching for Moment-Constrained Structure Elucidation</p> <p><i>Austin Henry Cheng, Alston Lo, Kin Long Kelvin Lee, Santiago Miret, Alan Aspuru-Guzik</i></p>	14063

<p>Towards Unbiased Learning in Semi-Supervised Semantic Segmentation</p> <p>Rui Sun, Huayu Mai, Wangkai Li, Tianzhu Zhang</p>	14094
<p>HALL-E: Hierarchical Neural Codec Language Model for Minute-Long Zero-Shot Text-To-Speech Synthesis</p> <p>Yuto Nishimura, Takumi Hirose, Masanari Ohi, Hideki Nakayama, Nakamasa Inoue</p>	14126
<p>RRM: Robust Reward Model Training Mitigates Reward Hacking</p> <p>Tianqi Liu, Wei Xiong, Jie Ren, Lichang Chen, Junru Wu, Rishabh Joshi, Yang Gao, Jiaming Shen, Zhen Qin, Tianhe Yu, Daniel Sohn, Anastasia Makarova, Jeremiah Zhe Liu, Yuan Liu, Bilal Piot, Abe Ittycheriah, Aviral Kumar, Mohammad Saleh</p>	14150
<p>Less is More: Masking Elements in Image Condition Features Avoids Content Leakages in Style Transfer Diffusion Models</p> <p>Lin Zhu, Xinbing Wang, Chenghu Zhou, Qinying Gu, Nanyang Ye</p>	14169
<p>Optimal Brain Apoptosis</p> <p>Mingyuan Sun, Zheng Fang, Jiayu Wang, Junjie Jiang, Delei Kong, Chenming Hu, Yuetong Fang, Renjing Xu</p>	14202
<p>PICASO: Permutation-Invariant Context Composition with State Space Models</p> <p>Tian Yu Liu, Alessandro Achille, Matthew Trager, Aditya Golatkar, Luca Zancato, Stefano Soatto</p>	14221
<p>ViSAGe: Video-To-Spatial Audio Generation</p> <p>Jaeyeon Kim, Heeseung Yun, Gunhee Kim</p>	14239
<p>Quality Measures for Dynamic Graph Generative Models</p> <p>Ryien Hosseini, Filippo Simini, Venkatram Vishwanath, Rebecca Willett, Henry Hoffmann</p>	14260
<p>Synthesizing Programmatic Reinforcement Learning Policies with Large Language Model Guided Search</p> <p>Max Liu, Chan-Hung Yu, Wei-Hsu Lee, Cheng-Wei Hung, Yen-Chun Chen, Shao-Hua Sun</p>	14279
<p>SciLitLLM: How to Adapt LLMs for Scientific Literature Understanding</p> <p>Sihang Li, Jin Huang, Jiayi Zhuang, Yaorui Shi, Xiaochen Cai, Mingjun Xu, Xiang Wang, Linfeng Zhang, Guolin Ke, Hengxing Cai</p>	14335
<p>TGB-Seq Benchmark: Challenging Temporal GNNs with Complex Sequential Dynamics</p> <p>Lu Yi, Jie Peng, Yanping Zheng, Fengran Mo, Zhewei Wei, Yuhang Ye, Yue Zixuan, Zengfeng Huang</p>	14359
<p>ToolACE: Winning the Points of LLM Function Calling</p> <p>Weiwen Liu, Xu Huang, Xingshan Zeng, Xinlong Hao, Shuai Yu, Dexun Li, Shuai Wang, Weinan Gan, Zhengying Liu, Yuanqing Yu, Zezhong Wang, Yuxian Wang, Wu Ning, Yutai Hou, Bin Wang, Chuhan Wu, Wang Xinzhi, Yong Liu, Yasheng Wang, Duyu Tang, Dandan Tu, Lifeng Shang, Xin Jiang, Ruiming Tang, Defu Lian, Qun Liu, Enhong Chen</p>	14380
<p>MoDeGPT: Modular Decomposition for Large Language Model Compression</p> <p>Chi-Heng Lin, Shangqian Gao, James Seale Smith, Abhishek Patel, Shikhar Tuli, Yilin Shen, Hongxia Jin, Yen-Chang Hsu</p>	14403

<p>Dynamic Contrastive Skill Learning with State-Transition Based Skill Clustering and Dynamic Length Adjustment</p> <p><i>Jinwoo Choi, Seung-Woo Seo</i></p>	14439
<p>Optimal Strong Regret and Violation in Constrained MDPs Via Policy Optimization</p> <p><i>Francesco Emanuele Stradi, Matteo Castiglioni, Alberto Marchesi, Nicola Gatti</i></p>	14462
<p>TetSphere Splatting: Representing High-Quality Geometry with Lagrangian Volumetric Meshes</p> <p><i>Minghao Guo, Bohan Wang, Kaiming He, Wojciech Matusik</i></p>	14488
<p>Can LLMs Separate Instructions from Data? and What Do We Even Mean by That?</p> <p><i>Egor Zverev, Sahar Abdelnabi, Soroush Tabesh, Mario Fritz, Christoph H. Lampert</i></p>	14513
<p>Reasoning-Enhanced Healthcare Predictions with Knowledge Graph Community Retrieval</p> <p><i>Pengcheng Jiang, Cao Xiao, Minhao Jiang, Parminder Bhatia, Taha Kass-Hout, Jimeng Sun, Jiawei Han</i></p>	14546
<p>GDrag: Towards General-Purpose Interactive Editing with Anti-Ambiguity Point Diffusion</p> <p><i>Xiaojian Lin, Hanhui Li, Yuhao Cheng, Yiqiang Yan, Xiaodan Liang</i></p>	14592
<p>Controllable Generation Via Locally Constrained Resampling</p> <p><i>Kareem Ahmed, Kai-Wei Chang, Guy Van Den Broeck</i></p>	14615
<p>Streaming Video Question-Answering with In-Context Video KV-Cache Retrieval</p> <p><i>Shangzhe Di, Zhelun Yu, Guanghao Zhang, Haoyuan Li, Taozhong, Hao Cheng, Bolin Li, Wanggui He, Fangxun Shu, Hao Jiang</i></p>	14631
<p>C-Adapter: Adapting Deep Classifiers for Efficient Conformal Prediction Sets</p> <p><i>Kangdao Liu, Hao Zeng, Jianguo Huang, Huiping Zhuang, Chi Man Vong, Hongxin Wei</i></p>	14644
<p>SEPARATE: A Simple Low-Rank Projection for Gradient Compression in Modern Large-Scale Model Training Process</p> <p><i>Hanzhen Zhao, Xingyu Xie, Cong Fang, Zhouchen Lin</i></p>	14668
<p>EgoExo-Gen: Ego-Centric Video Prediction by Watching Exo-Centric Videos</p> <p><i>Jilan Xu, Yifei Huang, Baoqi Pei, Junlin Hou, Qingqiu Li, Guo Chen, Yuejie Zhang, Rui Feng, Weidi Xie</i></p>	14694
<p>Longhorn: State Space Models Are Amortized Online Learners</p> <p><i>Bo Liu, Rui Wang, Lemeng Wu, Yihao Feng, Peter Stone, Qiang Liu</i></p>	14712
<p>Automated Filtering of Human Feedback Data for Aligning Text-To-Image Diffusion Models</p> <p><i>Yongjin Yang, Sihyeon Kim, Hojung Jung, Sangmin Bae, Sangmook Kim, Se-Young Yun, Kimin Lee</i></p>	14728
<p>CraftRTL: High-Quality Synthetic Data Generation for Verilog Code Models with Correct-By-Construction Non-Textual Representations and Targeted Code Repair</p> <p><i>Mingjie Liu, Yun-Da Tsai, Wenfei Zhou, Haoxing Ren</i></p>	14757

<p>From Artificial Needles to Real Haystacks: Improving Retrieval Capabilities in LLMs by Finetuning on Synthetic Data</p> <p><i>Zheyang Xiong, Vasilis Papageorgiou, Kangwook Lee, Dimitris Papailiopoulos</i></p>	14803
<p>Causal Discovery Via Bayesian Optimization</p> <p><i>Bao Duong, Sunil Gupta, Thin Nguyen</i></p>	14819
<p>STAMP: Scalable Task- And Model-Agnostic Collaborative Perception</p> <p><i>Xiangbo Gao, Runsheng Xu, Jiachen Li, Ziran Wang, Zhiwen Fan, Zhengzhong Tu</i></p>	14853
<p>GenEx: Generating an Explorable World</p> <p><i>Taiming Lu, Tianmin Shu, Alan Yuille, Daniel Khashabi, Jieneng Chen</i></p>	14874
<p>Explain Yourself, Briefly! Self-Explaining Neural Networks with Concise Sufficient Reasons</p> <p><i>Shahaf Bassan, Ron Eliav, Shlomit Gur</i></p>	14900
<p>Don't Flatten, Tokenize! Unlocking the Key to SoftMoE's Efficacy in Deep RL</p> <p><i>Ghada Sokar, Johan Samir Obando Ceron, Aaron Courville, Hugo Larochelle, Pablo Samuel Castro</i></p>	14938
<p>DenseMatcher: Learning 3D Semantic Correspondence for Category-Level Manipulation from a Single Demo</p> <p><i>Junzhe Zhu, Yuanchen Ju, Junyi Zhang, Muhan Wang, Zhecheng Yuan, Kaizhe Hu, Huazhe Xu</i></p>	14960
<p>ARLON: Boosting Diffusion Transformers with Autoregressive Models for Long Video Generation</p> <p><i>Zongyi Li, Shujie Hu, Shujie Liu, Long Zhou, Jeongsoo Choi, Lingwei Meng, Xun Guo, Jinyu Li, Hefei Ling, Furu Wei</i></p>	14982
<p>Reconstructive Visual Instruction Tuning</p> <p><i>Haochen Wang, Anlin Zheng, Yucheng Zhao, Tiancai Wang, Zheng Ge, Xiangyu Zhang, Zhaoxiang Zhang</i></p>	15001
<p>PFGuard: A Generative Framework with Privacy and Fairness Safeguards</p> <p><i>Soyeon Kim, Yuji Roh, Geon Heo, Steven Euijong Whang</i></p>	15027
<p>Forking Paths in Neural Text Generation</p> <p><i>Eric J Bigelow, Ari Holtzman, Hidenori Tanaka, Tomer Ullman</i></p>	15055
<p>Guaranteed Generation from Large Language Models</p> <p><i>Minbeom Kim, Thibaut Thonet, Jos Rozen, Hwaran Lee, Kyomin Jung, Marc Dymetman</i></p>	15089
<p>Breach by a Thousand Leaks: Unsafe Information Leakage in 'Safe' AI Responses</p> <p><i>David Glukhov, Ziwen Han, Iliia Shumailov, Vardan Papyan, Nicolas Papernot</i></p>	15117
<p>Sharpness-Aware Minimization: General Analysis and Improved Rates</p> <p><i>Dimitris Oikonomou, Nicolas Loizou</i></p>	15149
<p>CONDA: Adaptive Concept Bottleneck for Foundation Models Under Distribution Shifts</p> <p><i>Jihye Choi, Jayaram Raghuram, Yixuan Li, Somesh Jha</i></p>	15176

<p>RouteLLM: Learning to Route LLMs from Preference Data</p> <p><i>Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E. Gonzalez, M Waleed Kadous, Ion Stoica</i></p>	15202
<p>Learn Hybrid Prototypes for Multivariate Time Series Anomaly Detection</p> <p><i>Ke-Yuan Shen</i></p>	15218
<p>Federated Domain Generalization with Data-Free On-Server Matching Gradient</p> <p><i>Trong Binh Nguyen, Duong Minh Nguyen, Jinsun Park, Viet Quoc Pham, Won-Joo Hwang</i></p>	15241
<p>Direct Post-Training Preference Alignment for Multi-Agent Motion Generation Model Using Implicit Feedback from Pre-Training Demonstrations</p> <p><i>Thomas Tian, Kratarth Goel</i></p>	15267
<p>GameGen-X: Interactive Open-World Game Video Generation</p> <p><i>Haoxuan Che, Xuanhua He, Quande Liu, Cheng Jin, Hao Chen</i></p>	15280
<p>LoLCATs: On Low-Rank Linearizing of Large Language Models</p> <p><i>Michael Zhang, Simran Arora, Rahul Chalamala, Benjamin Frederick Spector, Alan Wu, Krithik Ramesh, Aaryan Singhal, Christopher Re</i></p>	15328
<p>Neural Context Flows for Meta-Learning of Dynamical Systems</p> <p><i>Roussel Desmond Nzoyem, David A. W. Barton, Tom Deakin</i></p>	15382
<p>Generalizability of Neural Networks Minimizing Empirical Risk Based on Expressive Power</p> <p><i>Lijia Yu, Yibo Miao, Yifan Zhu, Xiao-Shan Gao, Lijun Zhang</i></p>	15422
<p>Semantics-Adaptive Activation Intervention for LLMs Via Dynamic Steering Vectors</p> <p><i>Weixuan Wang, Jingyuan Yang, Wei Peng</i></p>	15455
<p>FreqPrior: Improving Video Diffusion Models with Frequency Filtering Gaussian Noise</p> <p><i>Yunlong Yuan, Yuanfan Guo, Chunwei Wang, Wei Zhang, Hang Xu, Li Zhang</i></p>	15473
<p>Shapley-Guided Utility Learning for Effective Graph Inference Data Valuation</p> <p><i>Hongliang Chi, Qiong Wu, Zhengyi Zhou, Yao Ma</i></p>	15495
<p>Towards a Unified and Verified Understanding of Group-Operation Networks</p> <p><i>Wilson Wu, Louis Jaburi, Jacob Drori, Jason Gross</i></p>	15528
<p>ParFam -- (Neural Guided) Symbolic Regression Via Continuous Global Optimization</p> <p><i>Philipp Scholl, Katharina Bieker, Hillary Hauger, Gitta Kutyniok</i></p>	15560
<p>Amortized Control of Continuous State Space Feynman-Kac Model for Irregular Time Series</p> <p><i>Byoungwoo Park, Hyungi Lee, Juho Lee</i></p>	15599
<p>Order-Aware Interactive Segmentation</p> <p><i>Bin Wang, Anwesa Choudhuri, Meng Zheng, Zhongpai Gao, Benjamin Planche, Andong Deng, Qin Liu, Terrence Chen, Ulas Bagci, Ziyang Wu</i></p>	15634

<p>Do LLMs Have Consistent Values?</p> <p><i>Naama Rozen, Liat Bezalel, Gal Elidan, Amir Globerson, Ella Daniel</i></p>	15659
<p>Joint Fine-Tuning and Conversion of Pretrained Speech and Language Models Towards Linear Complexity</p> <p><i>Mutian He, Philip N. Garner</i></p>	15686
<p>LevAttention: Time, Space and Streaming Efficient Algorithm for Heavy Attentions</p> <p><i>Ravindran Kannan, Chiranjib Bhattacharyya, Praneeth Kacham, David Woodruff</i></p>	15704
<p>Edge Prompt Tuning for Graph Neural Networks</p> <p><i>Xingbo Fu, Yinhan He, Jundong Li</i></p>	15721
<p>Data Pruning by Information Maximization</p> <p><i>Haoru Tan, Sitong Wu, Wei Huang, Shizhen Zhao, Xiaojuan Qi</i></p>	15743
<p>Interpreting and Editing Vision-Language Representations to Mitigate Hallucinations</p> <p><i>Nicholas Jiang, Anish Kachinthaya, Suzanne Petryk, Yossi Gandelsman</i></p>	15766
<p>Causal Identification for Complex Functional Longitudinal Studies</p> <p><i>Andrew Ying</i></p>	15790
<p>Denoising Task Difficulty-Based Curriculum for Training Diffusion Models</p> <p><i>Jin-Young Kim, Hyojun Go, Soonwoo Kwon, Hyun-Gyoon Kim</i></p>	15813
<p>SiMHand: Mining Similar Hands for Large-Scale 3D Hand Pose Pre-Training</p> <p><i>Nie Lin, Takehiko Ohkawa, Yifei Huang, Mingfang Zhang, Minjie Cai, Ming Li, Ryosuke Furuta, Yoichi Sato</i></p>	15827
<p>Semi-Supervised CLIP Adaptation by Enforcing Semantic and Trapezoidal Consistency</p> <p><i>Kai Gan, Bo Ye, Min-Ling Zhang, Tong Wei</i></p>	15846
<p>On the Optimization and Generalization of Two-Layer Transformers with Sign Gradient Descent</p> <p><i>Bingrui Li, Wei Huang, Andi Han, Zhanpeng Zhou, Taiji Suzuki, Jun Zhu, Jianfei Chen</i></p>	15868
<p>LANTERN: Accelerating Visual Autoregressive Models with Relaxed Speculative Decoding</p> <p><i>Doohyuk Jang, Sihwan Park, June Yong Yang, Yeonsung Jung, Jihun Yun, Souvik Kundu, Sung-Yub Kim, Eunho Yang</i></p>	15947
<p>Sports-Traj: A Unified Trajectory Generation Model for Multi-Agent Movement in Sports</p> <p><i>Yi Xu, Yun Fu</i></p>	15977
<p>Vision and Language Synergy for Rehearsal Free Continual Learning</p> <p><i>Muhammad Anwar Ma'Sum, Mahardhika Pratama, Savitha Ramasamy, Lin Liu, H Habibullah, Ryszard Kowalczyk</i></p>	15997
<p>DebGCD: Debiased Learning with Distribution Guidance for Generalized Category Discovery</p> <p><i>Yuanpei Liu, Kai Han</i></p>	16028

<p>Enhancing Zeroth-Order Fine-Tuning for Language Models with Low-Rank Structures</p> <p>Yiming Chen, Yuan Zhang, Liyuan Cao, Kun Yuan, Zaiwen Wen</p>	16042
<p>Adaptive Retention & Correction: Test-Time Training for Continual Learning</p> <p>Haoran Chen, Micah Goldblum, Zuxuan Wu, Yu-Gang Jiang</p>	16069
<p>A CLIP-Powered Framework for Robust and Generalizable Data Selection</p> <p>Suorong Yang, Peng Ye, Wanli Ouyang, Dongzhan Zhou, Furao Shen</p>	16082
<p>Verifying Properties of Binary Neural Networks Using Sparse Polynomial Optimization</p> <p>Jianting Yang, Srecko Durasinovic, Jean B. Lasserre, Victor Magron, Jun Zhao</p>	16101
<p>Sparse Autoencoders Do Not Find Canonical Units of Analysis</p> <p>Patrick Leask, Bart Bussmann, Michael T Pearce, Joseph Isaac Bloom, Curt Tigges, Noura Al Moubayed, Lee Sharkey, Neel Nanda</p>	16126
<p>SPaR: Self-Play with Tree-Search Refinement to Improve Instruction-Following in Large Language Models</p> <p>Jiale Cheng, Xiao Liu, Cunxiang Wang, Xiaotao Gu, Yida Lu, Dan Zhang, Yuxiao Dong, Jie Tang, Hongning Wang, Minlie Huang</p>	16152
<p>Transformers Struggle to Learn to Search</p> <p>Abulhair Saparov, Srushti Ajay Pawar, Shreyas Pimpalgaonkar, Nitish Joshi, Richard Yuanzhe Pang, Vishakh Padmakumar, Mehran Kazemi, Najoung Kim, He He</p>	16172
<p>Ada-K Routing: Boosting the Efficiency of MoE-Based LLMs</p> <p>Tongtian Yue, Longteng Guo, Jie Cheng, Xuange Gao, Hua Huang, Jing Liu</p>	16198
<p>VideoPhy: Evaluating Physical Commonsense for Video Generation</p> <p>Hritik Bansal, Zongyu Lin, Tianyi Xie, Zeshun Zong, Michal Yarom, Yonatan Bitton, Chenfanfu Jiang, Yizhou Sun, Kai-Wei Chang, Aditya Grover</p>	16215
<p>Optimization by Parallel Quasi-Quantum Annealing with Gradient-Based Sampling</p> <p>Yuma Ichikawa, Yamato Arai</p>	16262
<p>AutoCGP: Closed-Loop Concept-Guided Policies from Unlabeled Demonstrations</p> <p>Pei Zhou, Ruizhe Liu, Qian Luo, Fan Wang, Yibing Song, Yanchao Yang</p>	16283
<p>An Efficient Framework for Crediting Data Contributors of Diffusion Models</p> <p>Mingyu Lu, Chris Lin, Chanwoo Kim, Su-In Lee</p>	16311
<p>PRDP: Progressively Refined Differentiable Physics</p> <p>Kanishk Bhatia, Felix Koehler, Nils Thuerey</p>	16345
<p>Hyper-Connections</p> <p>Defa Zhu, Hongzhi Huang, Zihao Huang, Yutao Zeng, Yunyao Mao, Banggu Wu, Qiyang Min, Xun Zhou</p>	16385

<p>Inverse Constitutional AI: Compressing Preferences into Principles</p> <p>Arduin Findeis, Timo Kaufmann, Eyke Hüllermeier, Samuel Albanie, Robert D. Mullins</p>	16422
<p>When Do GFlowNets Learn the Right Distribution?</p> <p>Tiago Silva, Rodrigo Barreto Alves, Eliezer De Souza Da Silva, Amauri H Souza, Vikas Garg, Samuel Kaski, Diego Mesquita</p>	16466
<p>Debiasing Federated Learning with Correlated Client Participation</p> <p>Zhenyu Sun, Ziyang Zhang, Zheng Xu, Gauri Joshi, Pranay Sharma, Ermin Wei</p>	16492
<p>SqueezeAttention: 2D Management of KV-Cache in LLM Inference Via Layer-Wise Optimal Budget</p> <p>Zihao Wang, Bin Cui, Shaoduo Gan</p>	16524
<p>Temporal Flexibility in Spiking Neural Networks: Towards Generalization Across Time Steps and Deployment Friendliness</p> <p>Kangrui Du, Yuhang Wu, Shikuang Deng, Shi Gu</p>	16538
<p>Beyond Sequence: Impact of Geometric Context for RNA Property Prediction</p> <p>Junjie Xu, Artem Moskalev, Tommaso Mansi, Mangal Prakash, Rui Liao</p>	16558
<p>3D-Properties: Identifying Challenges in DPO and Charting a Path Forward</p> <p>Yuzi Yan, Yibo Miao, Jialian Li, Yipin Zhang, Jian Xie, Zhijie Deng, Dong Yan</p>	16588
<p>OmniPhysGS: 3D Constitutive Gaussians for General Physics-Based Dynamics Generation</p> <p>Yuchen Lin, Chenguo Lin, Jianjin Xu, Yadong Mu</p>	16611
<p>Reward Learning from Multiple Feedback Types</p> <p>Yannick Metz, Andras Geiszl, Raphaël Baur, Mennatallah El-Assady</p>	16639
<p>Visual Haystacks: A Vision-Centric Needle-In-A-Haystack Benchmark</p> <p>Tsung-Han Wu, Giscard Biamby, Jerome Quenum, Ritwik Gupta, Joseph E. Gonzalez, Trevor Darrell, David Chan</p>	16690
<p>An Optimal Discriminator Weighted Imitation Perspective for Reinforcement Learning</p> <p>Haoran Xu, Shuozhe Li, Harshit Sikchi, Scott Niekum, Amy Zhang</p>	16714
<p>Causal Order: The Key to Leveraging Imperfect Experts in Causal Inference</p> <p>Aniket Vashishtha, Abbavaram Gowtham Reddy, Abhinav Kumar, Saketh Bachu, Vineeth N. Balasubramanian, Amit Sharma</p>	16735
<p>SIMPL: Scalable and Hassle-Free Optimisation of Neural Representations from Behaviour</p> <p>Tom George, Pierre Glaser, Kim Stachenfeld, Caswell Barry, Claudia Clopath</p>	16777
<p>ALBAR: Adversarial Learning Approach to Mitigate Biases in Action Recognition</p> <p>Joseph Fiorese, Ishan Rajendrakumar Dave, Mubarak Shah</p>	16802
<p>AttriBoT: A Bag of Tricks for Efficiently Approximating Leave-One-Out Context Attribution</p> <p>Fengyuan Liu, Nikhil Kandpal, Colin Raffel</p>	16821

<p>Towards Optimal Multi-Draft Speculative Decoding</p> <p>Zhengmian Hu, Tong Zheng, Vignesh Viswanathan, Ziyi Chen, Ryan A. Rossi, Yihan Wu, Dinesh Manocha, Heng Huang</p>	16845
<p>IRIS: LLM-Assisted Static Analysis for Detecting Security Vulnerabilities</p> <p>Ziyang Li, Saikat Dutta, Mayur Naik</p>	16868
<p>LocoVR: Multiuser Indoor Locomotion Dataset in Virtual Reality</p> <p>Kojiro Takeyama, Yimeng Liu, Misha Sra</p>	16892
<p>Multi-Objective Differentiable Neural Architecture Search</p> <p>Rhea Sanjay Sukthanker, Arber Zela, Benedikt Staffler, Samuel Dooley, Josif Grabocka, Frank Hutter</p>	16917
<p>SplatFormer: Point Transformer for Robust 3D Gaussian Splatting</p> <p>Yutong Chen, Marko Mihajlovic, Xiyi Chen, Yiming Wang, Sergey Prokudin, Siyu Tang</p>	16961
<p>Hierarchically Encapsulated Representation for Protocol Design in Self-Driving Labs</p> <p>Yu-Zhe Shi, Mingchen Liu, Fanxu Meng, Qiao Xu, Zhangqian Bi, Kun He, Lecheng Ruan, Qining Wang</p>	16988
<p>DiffusionGuard: A Robust Defense Against Malicious Diffusion-Based Image Editing</p> <p>June Suk Choi, Kyungmin Lee, Jongheon Jeong, Saining Xie, Jinwoo Shin, Kimin Lee</p>	17038
<p>Controllable Unlearning for Image-To-Image Generative Models Via ϵ-Constrained Optimization</p> <p>Xiaohua Feng, Yuyuan Li, Chaochao Chen, Li Zhang, Longfei Li, Jun Zhou, Xiaolin Zheng</p>	17085
<p>Deconstructing Denoising Diffusion Models for Self-Supervised Learning</p> <p>Xinlei Chen, Zhuang Liu, Saining Xie, Kaiming He</p>	17129
<p>LLMOPT: Learning to Define and Solve General Optimization Problems from Scratch</p> <p>Caigao Jiang, Xiang Shu, Hong Qian, Xingyu Lu, Jun Zhou, Aimin Zhou, Yang Yu</p>	17144
<p>DeepLTL: Learning to Efficiently Satisfy Complex LTL Specifications for Multi-Task RL</p> <p>Mathias Jackermeier, Alessandro Abate</p>	17171
<p>MetaOOD: Automatic Selection of OOD Detection Models</p> <p>Yuehan Qin, Yichi Zhang, Yi Nian, Xueying Ding, Yue Zhao</p>	17200
<p>Infilling Score: A Pretraining Data Detection Algorithm for Large Language Models</p> <p>Negin Raoof, Litu Rout, Giannis Daras, Sujay Sanghavi, Constantine Caramanis, Sanjay Shakkottai, Alex Dimakis</p>	17222
<p>Integrating Protein Dynamics into Structure-Based Drug Design Via Full-Atom Stochastic Flows</p> <p>Xiangxin Zhou, Yi Xiao, Haowei Lin, Xinheng He, Jiaqi Guan, Yang Wang, Qiang Liu, Feng Zhou, Liang Wang, Jianzhu Ma</p>	17239

<p>Montessori-Instruct: Generate Influential Training Data Tailored for Student Learning</p> <p>Xiaochuan Li, Zichun Yu, Chenyan Xiong</p>	17268
<p>Multi-Reward as Condition for Instruction-Based Image Editing</p> <p>Xin Gu, Ming Li, Libo Zhang, Fan Chen, Longyin Wen, Tiejian Luo, Sijie Zhu</p>	17294
<p>Zero-Shot Whole-Body Humanoid Control Via Behavioral Foundation Models</p> <p>Andrea Tirinzoni, Ahmed Touati, Jesse Farebrother, Mateusz Guzek, Anssi Kanervisto, Yingchen Xu, Alessandro Lazaric, Matteo Pirota</p>	17314
<p>Meta Flow Matching: Integrating Vector Fields on the Wasserstein Manifold</p> <p>Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio, Alexander Tong, Kirill Neklyudov</p>	17370
<p>Differentiable and Learnable Wireless Simulation with Geometric Transformers</p> <p>Thomas Hehn, Markus Peschl, Tribhuvanesh Orekondy, Arash Behboodi, Johann Brehmer</p>	17395
<p>Exact Computation of Any-Order Shapley Interactions for Graph Neural Networks</p> <p>Maximilian Muschalik, Fabian Fumagalli, Paolo Frazzetto, Janine Strotherm, Luca Hermes, Alessandro Sperduti, Eyke Hüllermeier, Barbara Hammer</p>	17416
<p>Conditional Diffusion with Ordinal Regression: Longitudinal Data Generation for Neurodegenerative Disease Studies</p> <p>Hyuna Cho, Ziquan Wei, Seungjoo Lee, Tingting Dan, Guorong Wu, Won Hwa Kim</p>	17459
<p>The Complexity of Two-Team Polymatrix Games with Independent Adversaries</p> <p>Alexandros Hollender, Gilbert Maystre, Sai Ganesh Nagarajan</p>	17481
<p>R-Sparse: Rank-Aware Activation Sparsity for Efficient LLM Inference</p> <p>Zhenyu Zhang, Zechun Liu, Yuandong Tian, Harshit Khaitan, Zhangyang Wang, Steven Li</p>	17497
<p>Doubly Robust Identification of Treatment Effects from Multiple Environments</p> <p>Piersilvio De Bartolomeis, Julia Kostin, Javier Abad, Yixin Wang, Fanny Yang</p>	17512
<p>Sample Then Identify: A General Framework for Risk Control and Assessment in Multimodal Large Language Models</p> <p>Qingni Wang, Tiantian Geng, Zhiyuan Wang, Teng Wang, Bo Fu, Feng Zheng</p>	17543
<p>Mixture of Parrots: Experts Improve Memorization More than Reasoning</p> <p>Samy Jelassi, Clara Mohri, David Brandfonbrener, Alex Gu, Nikhil Vyas, Nikhil Anand, David Alvarez-Melis, Yuanzhi Li, Sham M. Kakade, Eran Malach</p>	17561
<p>RobuRCDet: Enhancing Robustness of Radar-Camera Fusion in Bird's Eye View for 3D Object Detection</p> <p>Jingtong Yue, Zhiwei Lin, Xin Lin, Xiaoyu Zhou, Xiangtai Li, Lu Qi, Yongtao Wang, Ming-Hsuan Yang</p>	17588
<p>Boltzmann Semantic Score: A Semantic Metric for Evaluating Large Vision Models Using Large Language Models</p> <p>Ali Khajegili Mirabadi, Katherine Rich, Hossein Farahani, Ali Bashashati</p>	17604

<p>OSDA Agent: Leveraging Large Language Models for De Novo Design of Organic Structure Directing Agents</p> <p>Zhaolin Hu, Yixiao Zhou, Zhongan Wang, Xin Li, Weimin Yang, Hehe Fan, Yi Yang</p>	17630
<p>CheapNet: Cross-Attention on Hierarchical Representations for Efficient Protein-Ligand Binding Affinity Prediction</p> <p>Hyukjun Lim, Sun Kim, Sangseon Lee</p>	17663
<p>Hymba: A Hybrid-Head Architecture for Small Language Models</p> <p>Xin Dong, Yonggan Fu, Shizhe Diao, Wonmin Byeon, Zijia Chen, Ameya Sunil Mahabaleshwarkar, Shih-Yang Liu, Matthijs Van Keirsbilck, Min-Hung Chen, Yoshi Suhara, Yingyan Celine Lin, Jan Kautz, Pavlo Molchanov</p>	17697
<p>Open-CK: A Large Multi-Physics Fields Coupling Benchmarks in Combustion Kinetics</p> <p>Zaige Fei, Fan Xu, Junyuan Mao, Yuxuan Liang, Qingsong Wen, Kun Wang, Hao Wu, Yang Wang</p>	17724
<p>Pangea: A Fully Open Multilingual Multimodal LLM for 39 Languages</p> <p>Xiang Yue, Yueqi Song, Akari Asai, Seungone Kim, Jean De Dieu Nyandwi, Simran Khanuja, Anjali Kantharuban, Lintang Sutawika, Sathyanarayanan Ramamoorthy, Graham Neubig</p>	17751
<p>Self-Play Preference Optimization for Language Model Alignment</p> <p>Yue Wu, Zhiqing Sun, Huizhuo Yuan, Kaixuan Ji, Yiming Yang, Quanquan Gu</p>	17805
<p>CityGaussianV2: Efficient and Geometrically Accurate Reconstruction for Large-Scale Scenes</p> <p>Yang Liu, Chuanchen Luo, Zhongkai Mao, Junran Peng, Zhaoxiang Zhang</p>	17830
<p>The Hidden Cost of Waiting for Accurate Predictions</p> <p>Ali Shirali, Ariel D. Procaccia, Rediet Abebe</p>	17848
<p>Tuning-Free Bilevel Optimization: New Algorithms and Convergence Analysis</p> <p>Yifan Yang, Hao Ban, Minhui Huang, Shiqian Ma, Kaiyi Ji</p>	17883
<p>Grounding Continuous Representations in Geometry: Equivariant Neural Fields</p> <p>David Wessels, David M Knigge, Riccardo Valperga, Samuele Papa, Sharvaree Vadgama, Efstratios Gavves, Erik J Bekkers</p>	17934
<p>Consistent Flow Distillation for Text-To-3D Generation</p> <p>Runjie Yan, Yinbo Chen, Xiaolong Wang</p>	17955
<p>Retrieval Augmented Diffusion Model for Structure-Informed Antibody Design and Optimization</p> <p>Zichen Wang, Yaokun Ji, Jianing Tian, Shuangjia Zheng</p>	17987
<p>Rewarding Progress: Scaling Automated Process Verifiers for LLM Reasoning</p> <p>Amrith Setlur, Chirag Nagpal, Adam Fisch, Xinyang Geng, Jacob Eisenstein, Rishabh Agarwal, Alekh Agarwal, Jonathan Berant, Aviral Kumar</p>	18006
<p>Investigating Pattern Neurons in Urban Time Series Forecasting</p> <p>Chengxin Wang, Yiran Zhao, Shaofeng Cai, Gary Tan</p>	18037

<p>Implicit Search Via Discrete Diffusion: A Study on Chess</p> <p><i>Jiacheng Ye, Zhenyu Wu, Jiahui Gao, Zhiyong Wu, Xin Jiang, Zhenguo Li, Lingpeng Kong</i></p>	18055
<p>Generating Graphs Via Spectral Diffusion</p> <p><i>Giorgia Minello, Alessandro Biciato, Luca Rossi, Andrea Torsello, Luca Cosmo</i></p>	18077
<p>SaMer: A Scenario-Aware Multi-Dimensional Evaluator for Large Language Models</p> <p><i>Kehua Feng, Keyan Ding, Jing Yu, Yiwen Qu, Zhiwen Chen, Chengfei Lv, Gang Yu, Qiang Zhang, Huajun Chen</i></p>	18097
<p>AugKD: Ingenious Augmentations Empower Knowledge Distillation for Image Super-Resolution</p> <p><i>Yun Zhang, Wei Li, Simiao Li, Hanting Chen, Zhijun Tu, Bingyi Jing, Shaohui Lin, Jie Hu, Wenjia Wang</i></p>	18119
<p>AgentHarm: A Benchmark for Measuring Harmfulness of LLM Agents</p> <p><i>Maksym Andriushchenko, Alexandra Souly, Mateusz Dziemian, Derek Duenas, Maxwell Lin, Justin Wang, Dan Hendrycks, Andy Zou, J Zico Kolter, Matt Fredrikson, Yarin Gal, Xander Davies</i></p>	18136
<p>IterGen: Iterative Semantic-Aware Structured LLM Generation with Backtracking</p> <p><i>Shubham Ugare, Rohan Gumaste, Tarun Suresh, Gagandeep Singh, Sasa Misailovic</i></p>	18172
<p>I2VControl-Camera: Precise Video Camera Control with Adjustable Motion Strength</p> <p><i>Wanquan Feng, Jiawei Liu, Pengqi Tu, Tianhao Qi, Mingzhen Sun, Tianxiang Ma, Songtao Zhao, Siyu Zhou, Qian He</i></p>	18198
<p>Self-Updatable Large Language Models by Integrating Context into Model Parameters</p> <p><i>Yu Wang, Xinshuang Liu, Xiusi Chen, Sean O'Brien, Junda Wu, Julian McAuley</i></p>	18211
<p>Efficient Inference for Large Language Model-Based Generative Recommendation</p> <p><i>Xinyu Lin, Chaoqun Yang, Wenjie Wang, Yongqi Li, Cunxiao Du, Fuli Feng, See-Kiong Ng, Tat-Seng Chua</i></p>	18230
<p>DeLLMa: Decision Making Under Uncertainty with Large Language Models</p> <p><i>Ollie Liu, Deqing Fu, Dani Yogatama, Willie Neiswanger</i></p>	18256
<p>Uni²Det: Unified and Universal Framework for Prompt-Guided Multi-Dataset 3D Detection</p> <p><i>Yubin Wang, Zhikang Zou, Xiaoqing Ye, Xiao Tan, Errui Ding, Cairong Zhao</i></p>	18293
<p>VibeCheck: Discover and Quantify Qualitative Differences in Large Language Models</p> <p><i>Lisa Dunlap, Krishna Mandal, Trevor Darrell, Jacob Steinhardt, Joseph E. Gonzalez</i></p>	18307
<p>Sharpness-Aware Minimization Efficiently Selects Flatter Minima Late in Training</p> <p><i>Zhanpeng Zhou, Mingze Wang, Yuchen Mao, Bingrui Li, Junchi Yan</i></p>	18336
<p>XAlguiFormer: Explainable Artificial Intelligence Guided Transformer for Brain Disorder Identification</p> <p><i>Hanning Guo, Farah Abdellatif, Yu Fu, N. Jon Shah, Abigail Morrison, Jürgen Dammers</i></p>	18368

<p>Approximation Algorithms for Combinatorial Optimization with Predictions</p> <p><i>Antonios Antoniadis, Marek Elias, Adam Polak, Moritz Venzin</i></p>	18387
<p>HG-Adapter: Improving Pre-Trained Heterogeneous Graph Neural Networks with Dual Adapters</p> <p><i>Yujie Mo, Runpeng Yu, Xiaofeng Zhu, Xinchao Wang</i></p>	18401
<p>MMDisCo: Multi-Modal Discriminator-Guided Cooperative Diffusion for Joint Audio and Video Generation</p> <p><i>Akio Hayakawa, Masato Ishii, Takashi Shibuya, Yuki Mitsufuji</i></p>	18427
<p>Targeted Attack Improves Protection Against Unauthorized Diffusion Customization</p> <p><i>Boyang Zheng, Chumeng Liang, Xiaoyu Wu</i></p>	18453
<p>Efficient Action-Constrained Reinforcement Learning Via Acceptance-Rejection Method and Augmented MDPs</p> <p><i>Wei Hung, Shao-Hua Sun, Ping-Chun Hsieh</i></p>	18484
<p>DiffGAD: A Diffusion-Based Unsupervised Graph Anomaly Detector</p> <p><i>Jinghan Li, Yuan Gao, Jinda Lu, Junfeng Fang, Congcong Wen, Hui Lin, Xiang Wang</i></p>	18507
<p>Repetition Improves Language Model Embeddings</p> <p><i>Jacob Mitchell Springer, Suhas Kotha, Daniel Fried, Graham Neubig, Aditi Raghunathan</i></p>	18530
<p>Diff3DS: Generating View-Consistent 3D Sketch Via Differentiable Curve Rendering</p> <p><i>Yibo Zhang, Lihong Wang, Changqing Zou, Tieru Wu, Rui Ma</i></p>	18567
<p>Wasserstein-Regularized Conformal Prediction Under General Distribution Shift</p> <p><i>Rui Xu, Chao Chen, Yue Sun, Parvathinathan Venkatasubramaniam, Sihong Xie</i></p>	18593
<p>Decentralized Optimization with Coupled Constraints</p> <p><i>Demyan Yarmoshik, Alexander Rogozin, Nikita Kiselev, Daniil Dorin, Alexander Gasnikov, Dmitry Kovalev</i></p>	18619
<p>Conformal Language Model Reasoning with Coherent Factuality</p> <p><i>Maxon Rubin-Toles, Maya Gambhir, Keshav Ramji, Aaron Roth, Surbhi Goel</i></p>	18641
<p>DexTrack: Towards Generalizable Neural Tracking Control for Dexterous Manipulation from Human References</p> <p><i>Xueyi Liu, Jianbieke Adalibieke, Qianwei Han, Yuzhe Qin, Li Yi</i></p>	18670
<p>Does Refusal Training in LLMs Generalize to the Past Tense?</p> <p><i>Maksym Andriushchenko, Nicolas Flammarion</i></p>	18695
<p>Learning Geometric Reasoning Networks for Robot Task and Motion Planning</p> <p><i>Smail Ait Bouhsain, Rachid Alami, Thierry Simeon</i></p>	18710

<p>GSM-Symbolic: Understanding the Limitations of Mathematical Reasoning in Large Language Models</p> <p>Seyed Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, Mehrdad Farajtabar</p>	18734
<p>Beyond Random Augmentations: Pretraining with Hard Views</p> <p>Fabio Ferreira, Ivo Rapant, Jörg K. H. Franke, Frank Hutter</p>	18757
<p>Periodic Materials Generation Using Text-Guided Joint Diffusion Model</p> <p>Kishalay Das, Subhojyoti Khastagir, Pawan Goyal, Seung-Cheol Lee, Satadeep Bhattacharjee, Niloy Ganguly</p>	18779
<p>Towards a General Time Series Anomaly Detector with Adaptive Bottlenecks and Dual Adversarial Decoders</p> <p>Qichao Shentu, Beibu Li, Kai Zhao, Yang Shu, Zhongwen Rao, Lujia Pan, Bin Yang, Chenjuan Guo</p>	18810
<p>Toward Understanding In-Context Vs. In-Weight Learning</p> <p>Bryan Chan, Xinyi Chen, András György, Dale Schuurmans</p>	18834
<p>Robust Feature Learning for Multi-Index Models in High Dimensions</p> <p>Alireza Mousavi-Hosseini, Adel Javanmard, Murat A Erdogdu</p>	18869
<p>ActSafe: Active Exploration with Safety Constraints for Reinforcement Learning</p> <p>Yarden As, Bhavya Sukhija, Lenart Treven, Carmelo Sferrazza, Stelian Coros, Andreas Krause</p>	18908
<p>Sequential Stochastic Combinatorial Optimization Using Hierarchical Reinforcement Learning</p> <p>Xinsong Feng, Zihan Yu, Yanhai Xiong, Haipeng Chen</p>	18933
<p>KBLaM: Knowledge Base Augmented Language Model</p> <p>Xi Wang, Taketomo Isazawa, Liana Mikaelyan, James Hensman</p>	18960
<p>Learning-Guided Rolling Horizon Optimization for Long-Horizon Flexible Job-Shop Scheduling</p> <p>Sirui Li, Wenbin Ouyang, Yining Ma, Cathy Wu</p>	18990
<p>MagicPIG: LSH Sampling for Efficient LLM Generation</p> <p>Zhuoming Chen, Ranajoy Sadhukhan, Zihao Ye, Yang Zhou, Jianyu Zhang, Niklas Nolte, Yuandong Tian, Matthijs Douze, Leon Bottou, Zhihao Jia, Beidi Chen</p>	19035
<p>Can LLMs Really Learn to Translate a Low-Resource Language from One Grammar Book?</p> <p>Seth Aycock, David Stap, Di Wu, Christof Monz, Khalil Sima'An</p>	19057
<p>Denoising Autoregressive Transformers for Scalable Text-To-Image Generation</p> <p>Jiatao Gu, Yuyang Wang, Yizhe Zhang, Qihang Zhang, Dinghui Zhang, Navdeep Jaitly, Joshua M. Susskind, Shuangfei Zhai</p>	19081
<p>From Sparse Dependence to Sparse Attention: Unveiling How Chain-Of-Thought Enhances Transformer Sample Efficiency</p> <p>Kaiyue Wen, Huaqing Zhang, Hongzhou Lin, Jingzhao Zhang</p>	19103

<p>Looking Backward: Streaming Video-To-Video Translation with Feature Banks</p> <p><i>Feng Liang, Akio Kodaira, Chenfeng Xu, Masayoshi Tomizuka, Kurt Keutzer, Diana Marculescu</i></p>	19148
<p>HADAMRNN: BINARY AND SPARSE TERNARY ORTHOGONAL RNNS</p> <p><i>Armand Foucault, Francois Malgouyres, Franck Mamalet</i></p>	19169
<p>Glimpse: Enabling White-Box Methods to Use Proprietary Models for Zero-Shot LLM-Generated Text Detection</p> <p><i>Guangsheng Bao, Yanbin Zhao, Juncai He, Yue Zhang</i></p>	19199
<p>Stabilized Neural Prediction of Potential Outcomes in Continuous Time</p> <p><i>Konstantin Hess, Stefan Feuerriegel</i></p>	19221
<p>MIRACLE 3D: Memory-Efficient Integrated Robust Approach for Continual Learning on 3D Point Clouds Via Shape Model Construction</p> <p><i>Hossein Resani, Behrooz Nasihatkon</i></p>	19251
<p>Learning Structured Representations by Embedding Class Hierarchy with Fast Optimal Transport</p> <p><i>Siqi Zeng, Sixian Du, Makoto Yamada, Han Zhao</i></p>	19265
<p>Vevo: Controllable Zero-Shot Voice Imitation with Self-Supervised Disentanglement</p> <p><i>Xueyao Zhang, Xiaohui Zhang, Kainan Peng, Zhenyu Tang, Vimal Manohar, Yingru Liu, Jeff Hwang, Dangna Li, Yuhao Wang, Julian Chan, Yuan Huang, Zhizheng Wu, Mingbo Ma</i></p>	19296
<p>A Theoretical Analysis of Self-Supervised Learning for Vision Transformers</p> <p><i>Yu Huang, Zixin Wen, Yuejie Chi, Yingbin Liang</i></p>	19320
<p>L-WISE: Boosting Human Visual Category Learning Through Model-Based Image Selection and Enhancement</p> <p><i>Morgan Bruce Talbot, Gabriel Kreiman, James J. Dicarlo, Guy Gaziv</i></p>	19383
<p>Exploiting Structure in Offline Multi-Agent RL: The Benefits of Low Interaction Rank</p> <p><i>Wenhao Zhan, Scott Fujimoto, Zheqing Zhu, Jason D. Lee, Daniel Jiang, Yonathan Efroni</i></p>	19416
<p>Global Convergence in Neural ODEs: Impact of Activation Functions</p> <p><i>Tianxiang Gao, Siyuan Sun, Hailiang Liu, Hongyang Gao</i></p>	19451
<p>Geometry of Neural Reinforcement Learning in Continuous State and Action Spaces</p> <p><i>Saket Tiwari, Omer Gottesman, George Konidaris</i></p>	19505
<p>Intelligent Go-Explore: Standing on the Shoulders of Giant Foundation Models</p> <p><i>Cong Lu, Shengran Hu, Jeff Clune</i></p>	19544
<p>FairDen: Fair Density-Based Clustering</p> <p><i>Lena Krieger, Anna Beer, Pernille Matthews, Anneka Myrup Thiesson, Ira Assent</i></p>	19570

<p>ConFIG: Towards Conflict-Free Training of Physics Informed Neural Networks</p> <p><i>Qiang Liu, Mengyu Chu, Nils Thuerey</i></p>	19590
<p>Mitigate the Gap: Improving Cross-Modal Alignment in CLIP</p> <p><i>Sedigheh Eslami, Gerard De Melo</i></p>	19626
<p>Training Language Models on Synthetic Edit Sequences Improves Code Synthesis</p> <p><i>Ulyana Piterbarg, Lerrel Pinto, Rob Fergus</i></p>	19645
<p>Group-Robust Sample Reweighting for Subpopulation Shifts Via Influence Functions</p> <p><i>Rui Qiao, Zhaoxuan Wu, Jingtian Wang, Pang Wei Koh, Bryan Kian Hsiang Low</i></p>	19670
<p>Towards Multiple Character Image Animation Through Enhancing Implicit Decoupling</p> <p><i>Jingyun Xue, Hongfa Wang, Qi Tian, Yue Ma, Andong Wang, Zhiyuan Zhao, Shaobo Min, Wenzhe Zhao, Kaihao Zhang, Heung-Yeung Shum, Wei Liu, Mengyang Liu, Wenhan Luo</i></p>	19692
<p>Adaptive Data Optimization: Dynamic Sample Selection with Scaling Laws</p> <p><i>Yiding Jiang, Allan Zhou, Zhili Feng, Sadhika Malladi, J Zico Kolter</i></p>	19718
<p>High-Quality Joint Image and Video Tokenization with Causal VAE</p> <p><i>Dawit Mureja Argaw, Xian Liu, Qinsheng Zhang, Joon Son Chung, Ming-Yu Liu, Fitsum Reda</i></p>	19742
<p>Robust Representation Consistency Model Via Contrastive Denoising</p> <p><i>Jiachen Lei, Julius Berner, Jiong Xiao Wang, Zhongzhu Chen, Chaowei Xiao, Zhongjie Ba, Kui Ren, Jun Zhu, Anima Anandkumar</i></p>	19759
<p>AtomSurf: Surface Representation for Learning on Protein Structures</p> <p><i>Vincent Mallet, Yangyang Miao, Souhaib Attaiki, Bruno Correia, Maks Ovsjanikov</i></p>	19779
<p>Ask, and it Shall Be Given: On the Turing Completeness of Prompting</p> <p><i>Ruizhong Qiu, Zhe Xu, Wenxuan Bao, Hanghang Tong</i></p>	19805
<p>Prediction Risk and Estimation Risk of the Ridgeless Least Squares Estimator Under General Assumptions on Regression Errors</p> <p><i>Sungyoon Lee, Sokbae Lee</i></p>	19829
<p>LR0.FM: Low-Res Benchmark and Improving Robustness for Zero-Shot Classification in Foundation Models</p> <p><i>Priyank Pathak, Shyam Marjit, Shruti Vyas, Yogesh S Rawat</i></p>	19847
<p>Is Factuality Enhancement a Free Lunch for LLMs? Better Factuality Can Lead to Worse Context-Faithfulness</p> <p><i>Baolong Bi, Shenghua Liu, Yiwei Wang, Lingrui Mei, Junfeng Fang, Hongcheng Gao, Shiyu Ni, Xueqi Cheng</i></p>	19864
<p>Understanding and Enhancing the Transferability of Jailbreaking Attacks</p> <p><i>Runqi Lin, Bo Han, Fengwang Li, Tongliang Liu</i></p>	19896

Injecting Universal Jailbreak Backdoors into LLMs in Minutes Zhuowei Chen, Qiannan Zhang, Shichao Pei	19920
Toward Generalizing Visual Brain Decoding to Unseen Subjects Xiangtao Kong, Kexin Huang, Ping Li, Lei Zhang	19937
KAA: Kolmogorov-Arnold Attention for Enhancing Attentive Graph Neural Networks Taoran Fang, Tianhong Gao, Chunping Wang, Yihaoshang, Wei Chow, Lei Chen, Yang Yang	19951
Confidence Elicitation: A New Attack Vector for Large Language Models Brian Formento, Chuan-Sheng Foo, See-Kiong Ng	19976
Transformers Provably Learn Two-Mixture of Linear Classification Via Gradient Flow Hongru Yang, Zhangyang Wang, Jason D. Lee, Yingbin Liang	20001
MatExpert: Decomposing Materials Discovery by Mimicking Human Experts Qianggang Ding, Santiago Miret, Bang Liu	20059
RecDreamer: Consistent Text-To-3D Generation Via Uniform Score Distillation Chenxi Zheng, Yihong Lin, Bangzhen Liu, Xuemiao Xu, Yongwei Nie, Shengfeng He	20079
Multi-Robot Motion Planning with Diffusion Models Yorai Shaoul, Itamar Mishani, Shivam Vats, Jiaoyang Li, Maxim Likhachev	20116
Depth Pro: Sharp Monocular Metric Depth in Less than a Second Alexey Bochkovskiy, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan Richter, Vladlen Koltun	20137
DisPose: Disentangling Pose Guidance for Controllable Human Image Animation Hongxiang Li, Yaowei Li, Yuhang Yang, Junjie Cao, Zhihong Zhu, Xuxin Cheng, Long Chen	20173
QMP: Q-Switch Mixture of Policies for Multi-Task Behavior Sharing Grace Zhang, Ayush Jain, Injune Hwang, Shao-Hua Sun, Joseph J Lim	20192
Dense Video Object Captioning from Disjoint Supervision Xingyi Zhou, Anurag Arnab, Chen Sun, Cordelia Schmid	20221
Mitigating Modality Prior-Induced Hallucinations in Multimodal Large Language Models Via Deciphering Attention Causality Guanyu Zhou, Yibo Yan, Xin Zou, Kun Wang, Aiwei Liu, Xuming Hu	20241
Searching for Optimal Solutions with LLMs Via Bayesian Optimization Dhruv Agarwal, Manoj Ghuhan Arivazhagan, Rajarshi Das, Sandesh Swamy, Sopan Khosla, Rashmi Gangadharaiah	20266
Scalable Decentralized Learning with Teleportation Yuki Takezawa, Sebastian U Stich	20288

<p>Motion-Agent: A Conversational Framework for Human Motion Generation with LLMs</p> <p>Qi Wu, Yubo Zhao, Yifan Wang, Xinhang Liu, Yu-Wing Tai, Chi-Keung Tang</p>	20314
<p>DARE the Extreme: Revisiting Delta-Parameter Pruning for Fine-Tuned Models</p> <p>Wenlong Deng, Yize Zhao, Vala Vakilian, Minghui Chen, Xiaoxiao Li, Christos Thrampoulidis</p>	20335
<p>Learning Efficient Positional Encodings with Graph Neural Networks</p> <p>Charilaos Kanatsoulis, Evelyn Choi, Stefanie Jegelka, Jure Leskovec, Alejandro Ribeiro</p>	20370
<p>Training Neural Networks as Recognizers of Formal Languages</p> <p>Alexandra Butoi, Ghazal Khalighinejad, Anej Svete, Josef Valvoda, Ryan Cotterell, Brian Dussell</p>	20398
<p>Generating Freeform Endoskeletal Robots</p> <p>Muhan Li, Lingji Kong, Sam Kriegman</p>	20442
<p>Efficient Evolutionary Search Over Chemical Space with Large Language Models</p> <p>Haorui Wang, Marta Skreta, Cher Tian Ser, Wenhao Gao, Lingkai Kong, Felix Strieth-Kalthoff, Chenru Duan, Yuchen Zhuang, Yue Yu, Yanqiao Zhu, Yuanqi Du, Alan Aspuru-Guzik, Kirill Neklyudov, Chao Zhang</p>	20467
<p>TopoLM: Brain-Like Spatio-Functional Organization in a Topographic Language Model</p> <p>Neil Rathi, Johannes Mehrer, Badr Alkhamissi, Taha Osama A Binhuraib, Nicholas Blauch, Martin Schrimpf</p>	20501
<p>Inner Information Analysis Algorithm for Deep Neural Network Based on Community</p> <p>Guipeng Lan, Shuai Xiao, Meng Xi, Jiabao Wen, Jiachen Yang</p>	20527
<p>Revealing the 3D Cosmic Web Through Gravitationally Constrained Neural Fields</p> <p>Brandon Zhao, Aviad Levis, Liam Connor, Pratul P. Srinivasan, Katherine Bouman</p>	20548
<p>E(n) Equivariant Topological Neural Networks</p> <p>Claudio Battiloro, Ege Karaismailoglu, Mauricio Tec, George Dasoulas, Michelle Audirac, Francesca Dominici</p>	20563
<p>Recovering Manifold Structure Using Ollivier Ricci Curvature</p> <p>Tristan Luca Saidi, Abigail Hickok, Andrew J. Blumberg</p>	20606
<p>AgentStudio: A Toolkit for Building General Virtual Agents</p> <p>Longtao Zheng, Zhiyuan Huang, Zhenghai Xue, Xinrun Wang, Bo An, Shuicheng Yan</p>	20666
<p>Standardizing Structural Causal Models</p> <p>Weronika Ormaniec, Scott Sussex, Lars Lorch, Bernhard Schölkopf, Andreas Krause</p>	20708
<p>InCoDe: Interpretable Compressed Descriptions for Image Generation</p> <p>Armand Comas, Aditya Chattopadhyay, Feliu Formosa, Changyu Liu, Octavia Camps, Rene Vidal</p>	20751
<p>MotionClone: Training-Free Motion Cloning for Controllable Video Generation</p> <p>Pengyang Ling, Jiazi Bu, Pan Zhang, Xiaoyi Dong, Yuhang Zang, Tong Wu, Huaian Chen, Jiaqi Wang, Yi Jin</p>	20783

<p>LLaMaFlex: Many-In-One LLMs Via Generalized Pruning and Weight Sharing</p> <p>Ruisi Cai, Saurav Muralidharan, Hongxu Yin, Zhangyang Wang, Jan Kautz, Pavlo Molchanov</p>	20806
<p>Efficient Diversity-Preserving Diffusion Alignment Via Gradient-Informed GFlowNets</p> <p>Zhen Liu, Tim Z. Xiao, Weiyang Liu, Yoshua Bengio, Dinghui Zhang</p>	20820
<p>A Geometric Framework for Understanding Memorization in Generative Models</p> <p>Brendan Leigh Ross, Hamidreza Kamkari, Tongzi Wu, Rasa Hosseinzadeh, Zhaoyan Liu, George Stein, Jesse C. Cresswell, Gabriel Loaiza-Ganem</p>	20856
<p>Boosting Multiple Views for Pretrained-Based Continual Learning</p> <p>Quyên Tran, Tung Lam Tran, Khanh Doan, Toan Tran, Dinh Phung, Khoat Than, Trung Le</p>	20889
<p>Revisiting Convolution Architecture in the Realm of DNA Foundation Models</p> <p>Yu Bo, Weian Mao, Yanjun Shao, Weiqiang Bai, Peng Ye, Xinzhu Ma, Junbo Zhao, Hao Chen, Chunhua Shen</p>	20918
<p>Risk-Sensitive Diffusion: Robustly Optimizing Diffusion Models with Noisy Samples</p> <p>Yangming Li, Max Ruiz Luyten, Mihaela Van Der Schaar</p>	20941
<p>DeepGate4: Efficient and Effective Representation Learning for Circuit Design at Scale</p> <p>Ziyang Zheng, Shan Huang, Jianyuan Zhong, Zhengyuan Shi, Guohao Dai, Ningyi Xu, Qiang Xu</p>	20972
<p>TraceVLA: Visual Trace Prompting Enhances Spatial-Temporal Awareness for Generalist Robotic Policies</p> <p>Ruijie Zheng, Yongyuan Liang, Shuaiyi Huang, Jianfeng Gao, Hal Daumé III, Andrey Kolobov, Furong Huang, Jianwei Yang</p>	20991
<p>BigDocs: An Open Dataset for Training Multimodal Models on Document and Code Tasks</p> <p>Juan A. Rodriguez, Xiangru Jian, Siba Smarak Panigrahi, Tianyu Zhang, Aarash Feizi, Abhay Puri, Akshay Kalkunte Suresh, François Savard, Ahmed Masry, Shravan Nayak, Rabiul Awal, Mahsa Massoud, Amirhossein Abaskohi, Zichao Li, Suyuchen Wang, Pierre-Andre Noel, Mats Leon Richter, Saverio Vadacchino, Shubham Agarwal, Sanket Biswas, Sara Shanian, Ying Zhang, Sathwik Tejaswi Madhusudhan, Joao Monteiro, Krishnamurthy Dj Dvijotham, Torsten Scholak, Nicolas Chapados, Sepideh Kharaghani, Sean Hughes, M. Özsu, Siva Reddy, Marco Pedersoli, Yoshua Bengio, Christopher Pal, Issam H. Laradji, Spandana Gella, Perouz Taslakian, David Vazquez, Sai Rajeswar</p>	21011
<p>Adapting Multi-Modal Large Language Model to Concept Drift from Pre-Training Onwards</p> <p>Xiaoyu Yang, Jie Lu, En Yu</p>	21068
<p>Rethinking Shapley Value for Negative Interactions in Non-Convex Games</p> <p>Wonjoon Chang, Myeongjin Lee, Jaesik Choi</p>	21091
<p>Fugatto 1: Foundational Generative Audio Transformer Opus 1</p> <p>Rafael Valle, Rohan Badlani, Zhifeng Kong, Sang-Gil Lee, Arushi Goel, Sungwon Kim, Joao Felipe Santos, Shuqi Dai, Siddharth Gururani, Aya Aljafari, Alexander H. Liu, Kevin J. Shih, Ryan Prenger, Wei Ping, Chao-Han Huck Yang, Bryan Catanzaro</p>	21111

<p>FLIP: Flow-Centric Generative Planning as General-Purpose Manipulation World Model</p> <p>Chongkai Gao, Haozhuo Zhang, Zhixuan Xu, Cai Zhehao, Lin Shao</p>	21140
<p>Understanding Constraint Inference in Safety-Critical Inverse Reinforcement Learning</p> <p>Bo Yue, Shufan Wang, Ashish Gaurav, Jian Li, Pascal Poupart, Guiliang Liu</p>	21162
<p>No Training, No Problem: Rethinking Classifier-Free Guidance for Diffusion Models</p> <p>Seyedmorteza Sadat, Manuel Kansy, Otmar Hilliges, Romann M. Weber</p>	21190
<p>X-Fi: A Modality-Invariant Foundation Model for Multimodal Human Sensing</p> <p>Xinyan Chen, Jianfei Yang</p>	21216
<p>The Case for Cleaner Biosignals: High-Fidelity Neural Compressor Enables Transfer from Cleaner iEEG to Noisier EEG</p> <p>Francesco S. Carzaniga, Gary Tom Hoppeler, Michael Hersche, Kaspar Schindler, Abbas Rahimi</p>	21236
<p>The Foundations of Tokenization: Statistical and Computational Concerns</p> <p>Juan Luis Gastaldi, John Terilla, Luca Malagutti, Brian Duseell, Tim Vieira, Ryan Cotterell</p>	21284
<p>TopoGaussian: Inferring Internal Topology Structures from Visual Clues</p> <p>Xiaoyu Xiong, Changyu Hu, Chunru Lin, Pingchuan Ma, Chuang Gan, Tao Du</p>	21302
<p>FormalAlign: Automated Alignment Evaluation for Autoformalization</p> <p>Jianqiao Lu, Yingjia Wan, Yinya Huang, Jing Xiong, Zhengying Liu, Zhijiang Guo</p>	21332
<p>SleepSMC: Ubiquitous Sleep Staging Via Supervised Multimodal Coordination</p> <p>Shuo Ma, Yingwei Zhang, Yiqiang Chen, Hualei Wang, Yuan Jin, Wei Zhang, Ziyu Jia</p>	21360
<p>Hot-Pluggable Federated Learning: Bridging General and Personalized FL Via Dynamic Selection</p> <p>Lei Shen, Zhenheng Tang, Lijun Wu, Yonggang Zhang, Xiaowen Chu, Tao Qin, Bo Han</p>	21382
<p>PRISM: Privacy-Preserving Improved Stochastic Masking for Federated Generative Models</p> <p>Kyeongkook Seo, Dong-Jun Han, Jaejun Yoo</p>	21419
<p>Linear Partial Gromov-Wasserstein Embedding</p> <p>Yikun Bai, Abihith Kothapalli, Hengrong Du, Rocio Diaz Martin, Soheil Kolouri</p>	21443
<p>MuHBoost: Multi-Label Boosting for Practical Longitudinal Human Behavior Modeling</p> <p>Nguyen T Thach, Patrick Habecker, Anika R. Eisenbraun, W. Alex Mason, Kimberly A. Tyler, Bilal Khan, Hau Chan</p>	21497
<p>Mini-Batch Coresets for Memory-Efficient Language Model Training on Data Mixtures</p> <p>Dang Nguyen, Wenhan Yang, Rathul Anand, Yu Yang, Baharan Mirzasoleiman</p>	21526
<p>ROUTE: Robust Multitask Tuning and Collaboration for Text-To-SQL</p> <p>Yang Qin, Chao Chen, Zhihang Fu, Ze Chen, Dezhong Peng, Peng Hu, Jieping Ye</p>	21547

<p>A Quantum Circuit-Based Compression Perspective for Parameter-Efficient Learning</p> <p><i>Chen-Yu Liu, Chao-Han Huck Yang, Hsi-Sheng Goan, Min-Hsiu Hsieh</i></p>	21571
<p>I^2AM: Interpreting Image-To-Image Latent Diffusion Models Via Bi-Attribution Maps</p> <p><i>Junseo Park, Hyeryung Jang</i></p>	21596
<p>NarrativeBridge: Enhancing Video Captioning with Causal-Temporal Narrative</p> <p><i>Asmar Nadeem, Faegheh Sardari, Robert Dawes, Syed Sameed Husain, Adrian Hilton, Armin Mustafa</i></p>	21619
<p>Learning Molecular Representation in a Cell</p> <p><i>Gang Liu, Srijit Seal, John Arevalo, Zhenwen Liang, Anne E Carpenter, Meng Jiang, Shantanu Singh</i></p>	21648
<p>Human-Aligned Chess with a Bit of Search</p> <p><i>Yiming Zhang, Athul Paul Jacob, Vivian Lai, Daniel Fried, Daphne Ippolito</i></p>	21671
<p>Generative Adapter: Contextualizing Language Models in Parameters with a Single Forward Pass</p> <p><i>Tong Chen, Hao Fang, Patrick Xia, Xiaodong Liu, Benjamin Van Durme, Luke Zettlemoyer, Jianfeng Gao, Hao Cheng</i></p>	21693
<p>Generative Representational Instruction Tuning</p> <p><i>Niklas Muennighoff, Hongjin Su, Liang Wang, Nan Yang, Furu Wei, Tao Yu, Amanpreet Singh, Douwe Kiela</i></p>	21714
<p>Kolmogorov-Arnold Transformer</p> <p><i>Xingyi Yang, Xinchao Wang</i></p>	21784
<p>Mix-LN: Unleashing the Power of Deeper Layers by Combining Pre-LN and Post-LN</p> <p><i>Pengxiang Li, Lu Yin, Shiwei Liu</i></p>	21808
<p>Human Simulacra: Benchmarking the Personification of Large Language Models</p> <p><i>Qiuqie Xie, Qiming Feng, Tianqi Zhang, Qingqiu Li, Linyi Yang, Yuejie Zhang, Rui Feng, Liang He, Shang Gao, Yue Zhang</i></p>	21823
<p>NetFormer: An Interpretable Model for Recovering Dynamical Connectivity in Neuronal Population Dynamics</p> <p><i>Ziyu Lu, Wuwei Zhang, Trung Le, Hao Wang, Uygur Sümbül, Eric Todd Sheabrown, Lu Mi</i></p>	21866
<p>Can Generative AI Solve Your In-Context Learning Problem? a Martingale Perspective</p> <p><i>Andrew Jesson, Nicolas Beltran-Velez, David Blei</i></p>	21895
<p>Fourier Sliced-Wasserstein Embedding for Multisets and Measures</p> <p><i>Tal Amir, Nadav Dym</i></p>	21918
<p>Adaptive Teachers for Amortized Samplers</p> <p><i>Minsu Kim, Sanghyeok Choi, Taeyoung Yun, Emmanuel Bengio, Leo Feng, Jarrid Rector-Brooks, Sungsoo Ahn, Jinkyoo Park, Nikolay Malkin, Yoshua Bengio</i></p>	21958

<p>Latent Safety-Constrained Policy Approach for Safe Offline Reinforcement Learning</p> <p><i>Prajwal Koirala, Zhanhong Jiang, Soumik Sarkar, Cody Fleming</i></p>	21985
<p>Flavors of Margin: Implicit Bias of Steepest Descent in Homogeneous Neural Networks</p> <p><i>Nikolaos Tsilivis, Gal Vardi, Julia Kempe</i></p>	22012
<p>Making Transformer Decoders Better Differentiable Indexers</p> <p><i>Wuchao Li, Kai Zheng, Defu Lian, Qi Liu, Wentian Bao, Yun En Yu, Yang Song, Han Li, Kun Gai</i></p>	22039
<p>Second-Order Fine-Tuning Without Pain for LLMs: A Hessian Informed Zeroth-Order Optimizer</p> <p><i>Yanjun Zhao, Sizhe Dang, Haishan Ye, Guang Dai, Yi Qian, Ivor Tsang</i></p>	22055
<p>Draw-And-Understand: Leveraging Visual Prompts to Enable MLLMs to Comprehend What You Want</p> <p><i>Weifeng Lin, Xinyu Wei, Ruichuan An, Peng Gao, Bocheng Zou, Yulin Luo, Siyuan Huang, Shanghang Zhang, Hongsheng Li</i></p>	22080
<p>MeshMask: Physics-Based Simulations with Masked Graph Neural Networks</p> <p><i>Paul Garnier, Vincent Lannelongue, Jonathan Viquerat, Elie Hachem</i></p>	22110
<p>Perturbation-Restrained Sequential Model Editing</p> <p><i>Jun-Yu Ma, Hong Wang, Hao-Xiang Xu, Zhen-Hua Ling, Jia-Chen Gu</i></p>	22127
<p>On the Linear Speedup of Personalized Federated Reinforcement Learning with Shared Representations</p> <p><i>Guojun Xiong, Shufan Wang, Daniel Jiang, Jian Li</i></p>	22155
<p>Distilling Reinforcement Learning Algorithms for In-Context Model-Based Planning</p> <p><i>Jaehyeon Son, Soochan Lee, Gunhee Kim</i></p>	22198
<p>No Preference Left Behind: Group Distributional Preference Optimization</p> <p><i>Binwei Yao, Zefan Cai, Yun-Shiuan Chuang, Shanglin Yang, Ming Jiang, Diyi Yang, Junjie Hu</i></p>	22216
<p>Rare-To-Frequent: Unlocking Compositional Generation Power of Diffusion Models on Rare Concepts with LLM Guidance</p> <p><i>Dongmin Park, Sebin Kim, Taehong Moon, Minkyu Kim, Kangwook Lee, Jaewoong Cho</i></p>	22235
<p>What Matters When Repurposing Diffusion Models for General Dense Perception Tasks?</p> <p><i>Guangkai Xu, Yongtao Ge, Mingyu Liu, Chengxiang Fan, Kangyang Xie, Zhiyue Zhao, Hao Chen, Chunhua Shen</i></p>	22262
<p>Horizon Generalization in Reinforcement Learning</p> <p><i>Vivek Myers, Catherine Ji, Benjamin Eysenbach</i></p>	22276
<p>Constructing Confidence Intervals for Average Treatment Effects from Multiple Datasets</p> <p><i>Yuxin Wang, Maresa Schröder, Dennis Frauen, Jonas Schweisthal, Konstantin Hess, Stefan Feuerriegel</i></p>	22301

<p>AutoDAN-Turbo: A Lifelong Agent for Strategy Self-Exploration to Jailbreak LLMs</p> <p>Xiaogeng Liu, Peiran Li, G. Edward Suh, Yevgeniy Vorobeychik, Zhuoqing Mao, Somesh Jha, Patrick McDaniel, Huan Sun, Bo Li, Chaowei Xiao</p>	22337
<p>Diffusion Transformers for Tabular Data Time Series Generation</p> <p>Fabrizio Garuti, Enver Sangineto, Simone Luetto, Lorenzo Forni, Rita Cucchiara</p>	22385
<p>Human-Inspired Episodic Memory for Infinite Context LLMs</p> <p>Zafeirios Fountas, Martin Benfeghoul, Adnan Oomerjee, Fenia Christopoulou, Gerasimos Lampouras, Haitham Bou Ammar, Jun Wang</p>	22411
<p>Gradient Descent with Generalized Newton's Method</p> <p>Zhiqi Bu, Shiyun Xu</p>	22450
<p>Samba: Simple Hybrid State Space Models for Efficient Unlimited Context Language Modeling</p> <p>Liliang Ren, Yang Liu, Yadong Lu, Yelong Shen, Chen Liang, Weizhu Chen</p>	22475
<p>CycleResearcher: Improving Automated Research Via Automated Review</p> <p>Yixuan Weng, Minjun Zhu, Guangsheng Bao, Hongbo Zhang, Jindong Wang, Yue Zhang, Linyi Yang</p>	22500
<p>Understanding and Mitigating Hallucination in Large Vision-Language Models Via Modular Attribution and Intervention</p> <p>Tianyun Yang, Ziniu Li, Juan Cao, Chang Xu</p>	22541
<p>Linear SCM Identification in the Presence of Confounders and Gaussian Noise</p> <p>Vahideh Sanjaroonpouri, Pouria Ramazi</p>	22564
<p>Severing Spurious Correlations with Data Pruning</p> <p>Varun Mulchandani, Jung-Eun Kim</p>	22582
<p>S4M: S4 for Multivariate Time Series Forecasting with Missing Values</p> <p>Peng Jing, Meiqi Yang, Qiong Zhang, Xiaoxiao Li</p>	22598
<p>CBMA: Improving Conformal Prediction Through Bayesian Model Averaging</p> <p>Pankaj Bhagwat, Linglong Kong, Bei Jiang</p>	22625
<p>ProteinBench: A Holistic Evaluation of Protein Foundation Models</p> <p>Fei Ye, Zaixiang Zheng, Dongyu Xue, Yuning Shen, Lihao Wang, Yiming Ma, Yan Wang, Xinyou Wang, Xiangxin Zhou, Quanquan Gu</p>	22644
<p>Scaling Instruction-Tuned LLMs to Million-Token Contexts Via Hierarchical Synthetic Data Generation</p> <p>Linda He, Jue Wang, Maurice Weber, Shang Zhu, Ben Athiwaratkun, Ce Zhang</p>	22679
<p>Leveraging Variable Sparsity to Refine Pareto Stationarity in Multi-Objective Optimization</p> <p>Zeou Hu, Yaoliang Yu</p>	22705
<p>Lightweight Neural App Control</p> <p>Filippos Christianos, Georgios Papoudakis, Thomas Coste, Jianye Hao, Jun Wang, Kun Shao</p>	22736

<p>Non-Equilibrium Dynamics of Hybrid Continuous-Discrete Ground-State Sampling</p> <p><i>Timothee Leleu, Sam Reifenstein</i></p>	22754
<p>Generating CAD Code with Vision-Language Models for 3D Designs</p> <p><i>Kamel Alrashedy, Pradyumna Tambwekar, Zulfiqar Haider Zaidi, Megan Langwasser, Wei Xu, Matthew Gombolay</i></p>	22779
<p>LucidPPN: Unambiguous Prototypical Parts Network for User-Centric Interpretable Computer Vision</p> <p><i>Mateusz Pach, Koryna Lewandowska, Jacek Tabor, Bartosz Michal Zielinski, Dawid Damian Rymarczyk</i></p>	22806
<p>Dualformer: Controllable Fast and Slow Thinking by Learning with Randomized Reasoning Traces</p> <p><i>Dijia Su, Sainbayar Sukhbaatar, Michael Rabbat, Yuandong Tian, Qinqing Zheng</i></p>	22903
<p>Harnessing Diversity for Important Data Selection in Pretraining Large Language Models</p> <p><i>Chi Zhang, Huaping Zhong, Kuan Zhang, Chengliang Chai, Rui Wang, Xinlin Zhuang, Tianyi Bai, Qiu Jiantao, Lei Cao, Ju Fan, Ye Yuan, Guoren Wang, Conghui He</i></p>	22941
<p>Test-Time Adaptation for Cross-Modal Retrieval with Query Shift</p> <p><i>Haobin Li, Peng Hu, Qianjun Zhang, Xi Peng, Xitingliu, Mouxing Yang</i></p>	22965
<p>PolyPythias: Stability and Outliers Across Fifty Language Model Pre-Training Runs</p> <p><i>Oskar Van Der Wal, Pietro Lesci, Max Müller-Eberstein, Naomi Saphra, Hailey Schoelkopf, Willem Zuidema, Stella Biderman</i></p>	22993
<p>Towards Generalizable Reinforcement Learning Via Causality-Guided Self-Adaptive Representations</p> <p><i>Yupei Yang, Biwei Huang, Fan Feng, Xinyue Wang, Shikui Tu, Lei Xu</i></p>	23018
<p>SmartPretrain: Model-Agnostic and Dataset-Agnostic Representation Learning for Motion Prediction</p> <p><i>Yang Zhou, Hao Shao, Letian Wang, Steven L. Waslander, Hongsheng Li, Yu Liu</i></p>	23050
<p>RB-Modulation: Training-Free Stylization Using Reference-Based Modulation</p> <p><i>Litu Rout, Yujia Chen, Nataniel Ruiz, Abhishek Kumar, Constantine Caramanis, Sanjay Shakkottai, Wen-Sheng Chu</i></p>	23071
<p>Selective Induction Heads: How Transformers Select Causal Structures in Context</p> <p><i>Francesco D'Angelo, Francesco Croce, Nicolas Flammarion</i></p>	23107
<p>Bridging the Gap Between Variational Inference and Stochastic Gradient MCMC in Function Space</p> <p><i>Mengjing Wu, Junyu Xuan, Jie Lu</i></p>	23146
<p>Backtracking Improves Generation Safety</p> <p><i>Yiming Zhang, Jianfeng Chi, Hailey Nguyen, Kartikeya Upasani, Daniel M. Bikel, Jason E Weston, Eric Michael Smith</i></p>	23174

<p>Towards Bridging Generalization and Expressivity of Graph Neural Networks</p> <p>Shouheng Li, Floris Geerts, Dongwoo Kim, Qing Wang</p>	23192
<p>Differentiable Causal Discovery for Latent Hierarchical Causal Models</p> <p>Parjanya Prajakta Prashant, Ignavier Ng, Kun Zhang, Biwei Huang</p>	23212
<p>Grounding by Trying: LLMs with Reinforcement Learning-Enhanced Retrieval</p> <p>Sheryl Hsu, Omar Khattab, Chelsea Finn, Archit Sharma</p>	23238
<p>Animate Your Thoughts: Reconstruction of Dynamic Natural Vision from Human Brain Activity</p> <p>Yizhuo Lu, Changde Du, Chong Wang, Xuanliu Zhu, Liuyun Jiang, Xujin Li, Huiguang He</p>	23255
<p>Spread Preference Annotation: Direct Preference Judgment for Efficient LLM Alignment</p> <p>Dongyoung Kim, Kimin Lee, Jinwoo Shin, Jaehyung Kim</p>	23294
<p>RuAG: Learned-Rule-Augmented Generation for Large Language Models</p> <p>Yudi Zhang, Pei Xiao, Lu Wang, Chaoyun Zhang, Meng Fang, Yali Du, Yevgeniy Puzyrev, Randolph Yao, Si Qin, Qingwei Lin, Mykola Pechenizkiy, Dongmei Zhang, Saravan Rajmohan, Qi Zhang</p>	23316
<p>ACE: All-Round Creator and Editor Following Instructions Via Diffusion Transformer</p> <p>Zhen Han, Zeyinzi Jiang, Yulin Pan, Jingfeng Zhang, Chaojie Mao, Chen-Wei Xie, Yu Liu, Jingren Zhou</p>	23340
<p>PolyhedronNet: Representation Learning for Polyhedra with Surface-Attributed Graph</p> <p>Dazhou Yu, Genpei Zhang, Liang Zhao</p>	23356
<p>SOO-Bench: Benchmarks for Evaluating the Stability of Offline Black-Box Optimization</p> <p>Hong Qian, Yiyi Zhu, Xiang Shu, Shuo Liu, Yaolin Wen, Xin An, Huakang Lu, Aimin Zhou, Ke Tang, Yang Yu</p>	23373
<p>Multimodal Unsupervised Domain Generalization by Retrieving Across the Modality Gap</p> <p>Christopher Liao, Christian So, Theodoros Tsiligkaridis, Brian Kulis</p>	23438
<p>RGB-Event ISP: The Dataset and Benchmark</p> <p>Yunfan Lu, Yanlin Qian, Ziyang Rao, Junren Xiao, Liming Chen, Hui Xiong</p>	23468
<p>LoRanPAC: Low-Rank Random Features and Pre-Trained Models for Bridging Theory and Practice in Continual Learning</p> <p>Liangzu Peng, Juan Elenter, Joshua Agterberg, Alejandro Ribeiro, Rene Vidal</p>	23483
<p>MatryoshkaKV: Adaptive KV Compression Via Trainable Orthogonal Projection</p> <p>Bokai Lin, Zihao Zeng, Zipeng Xiao, Siqi Kou, Tianqi Hou, Xiaofeng Gao, Hao Zhang, Zhijie Deng</p>	23530
<p>Synthio: Augmenting Small-Scale Audio Classification Datasets with Synthetic Data</p> <p>Sreyan Ghosh, Sonal Kumar, Zhifeng Kong, Rafael Valle, Bryan Catanzaro, Dinesh Manocha</p>	23552
<p>MR-GSM8K: A Meta-Reasoning Benchmark for Large Language Model Evaluation</p> <p>Zhongshen Zeng, Pengguang Chen, Shu Liu, Haiyun Jiang, Jiaya Jia</p>	23575

<p>ImDy: Human Inverse Dynamics from Imitated Observations</p> <p>Xinpeng Liu, Junxuan Liang, Zili Lin, Haowen Hou, Yong-Lu Li, Cewu Lu</p>	23597
<p>CoMRes: Semi-Supervised Time Series Forecasting Utilizing Consensus Promotion of Multi-Resolution</p> <p>Yunju Cho, Jay-Yoon Lee</p>	23618
<p>Prototype Antithesis for Biological Few-Shot Class-Incremental Learning</p> <p>Binghao Liu, Han Yang, Fang Wan, Fei Gu</p>	23642
<p>GeoLoRA: Geometric Integration for Parameter Efficient Fine-Tuning</p> <p>Steffen Schotthöfer, Emanuele Zangrando, Gianluca Ceruti, Francesco Tudisco, Jonas Kusch</p>	23656
<p>Test-Time Adaptation for Image Compression with Distribution Regularization</p> <p>Kecheng Chen, Pingping Zhang, Tiexin Qin, Shiqi Wang, Hong Yan, Haoliang Li</p>	23686
<p>On the Transfer of Object-Centric Representation Learning</p> <p>Aniket Rajiv Didolkar, Andrii Zadaianchuk, Anirudh Goyal, Michael Curtis Mozer, Yoshua Bengio, Georg Martius, Maximilian Seitzer</p>	23704
<p>Leave-One-Out Stable Conformal Prediction</p> <p>Kiljae Lee, Yuan Zhang</p>	23741
<p>QNBO: quasi-Newton Meets Bilevel Optimization</p> <p>Sheng Fang, Yongjin Liu, Wei Yao, Chengming Yu, Jin Zhang</p>	23766
<p>Dysca: A Dynamic and Scalable Benchmark for Evaluating Perception Ability of LVLMS</p> <p>Jie Zhang, Zhongqi Wang, Mengqi Lei, Zheng Yuan, Bei Yan, Shiguang Shan, Xilin Chen</p>	23807
<p>Efficient Residual Learning with Mixture-Of-Experts for Universal Dexterous Grasping</p> <p>Ziye Huang, Haoqi Yuan, Yuhui Fu, Zongqing Lu</p>	23823
<p>Breaking the $\sqrt{\log(1/\Delta_2)}$ Barrier: Better Batched Best Arm Identification with Adaptive Grids</p> <p>Tianyuan Jin, Qin Zhang, Dongruo Zhou</p>	23842
<p>The Geometry of Categorical and Hierarchical Concepts in Large Language Models</p> <p>Kiho Park, Yo Joong Choe, Yibo Jiang, Victor Veitch</p>	23863
<p>MotionAura: Generating High-Quality and Motion Consistent Videos Using Discrete Diffusion</p> <p>Onkar Kishor Susladkar, Jishu Sen Gupta, Chirag Sehgal, Sparsh Mittal, Rekha Singhal</p>	23886
<p>ReMatching Dynamic Reconstruction Flow</p> <p>Sara Oblak, Despoina Paschalidou, Sanja Fidler, Matan Atzmon</p>	23920
<p>Repulsive Latent Score Distillation for Solving Inverse Problems</p> <p>Nicolas Zilberstein, Morteza Mardani, Santiago Segarra</p>	23947

<p>Neural Networks on Symmetric Spaces of Noncompact Type</p> <p>Xuan Son Nguyen, Shuo Yang, Aymeric Histace</p>	23991
<p>Interactive Speculative Planning: Enhance Agent Efficiency Through Co-Design of System and User Interface</p> <p>Wenyue Hua, Mengting Wan, Jagannath Shashank Subramanya Sai Vadrevu, Ryan Nadel, Yongfeng Zhang, Chi Wang</p>	24028
<p>Number Cookbook: Number Understanding of Language Models and How to Improve it</p> <p>Haotong Yang, Yi Hu, Shijia Kang, Zhouchen Lin, Muhan Zhang</p>	24056
<p>ControlAR: Controllable Image Generation with Autoregressive Models</p> <p>Zongming Li, Tianheng Cheng, Shoufa Chen, Peize Sun, Haocheng Shen, Longjin Ran, Xiaoxin Chen, Wenyu Liu, Xinggang Wang</p>	24095
<p>Range, Not Independence, Drives Modularity in Biologically Inspired Representations</p> <p>Will Dorrell, Kyle Hsu, Luke Hollingsworth, Jin Hwa Lee, Jiajun Wu, Chelsea Finn, Peter E. Latham, Timothy Edward John Behrens, James C. R. Whittington</p>	24120
<p>Intermediate Layer Classifiers for OOD Generalization</p> <p>Arnas Uselis, Seong Joon Oh</p>	24157
<p>2-Rectifications Are Enough for Straight Flows: A Theoretical Insight into Wasserstein Convergence</p> <p>Saptarshi Roy, Vansh Bansal, Purnamrita Sarkar, Alessandro Rinaldo</p>	24187
<p>ADAM Optimization with Adaptive Batch Selection</p> <p>Gyu Yeol Kim, Min-Hwan Oh</p>	24216
<p>Flow Distillation Sampling: Regularizing 3D Gaussians with Pre-Trained Matching Priors</p> <p>Lin-Zhuo Chen, Kangjie Liu, Youtian Lin, Zhihao Li, Siyu Zhu, Xun Cao, Yao Yao</p>	24258
<p>T2V-Turbo-V2: Enhancing Video Model Post-Training Through Data, Reward, and Conditional Guidance Design</p> <p>Jiachen Li, Qian Long, Jian Zheng, Xiaofeng Gao, Robinson Piramuthu, Wenhui Chen, William Yang Wang</p>	24271
<p>Efficient and Trustworthy Causal Discovery with Latent Variables and Complex Relations</p> <p>Xiu-Chuan Li, Tongliang Liu</p>	24298
<p>A Large-Scale Training Paradigm for Graph Generative Models</p> <p>Yu Wang, Ryan A. Rossi, Namyong Park, Huiyuan Chen, Nesreen K. Ahmed, Puja Trivedi, Franck Dernoncourt, Danai Koutra, Tyler Derr</p>	24333
<p>Fitting Networks with a Cancellation Trick</p> <p>Jiashun Jin, Jingming Wang</p>	24362

<p>\mathbb{X}-Sample Contrastive Loss: Improving Contrastive Learning with Sample Similarity Graphs</p> <p>Vlad Sobal, Mark Ibrahim, Randall Balestriero, Vivien Cabannes, Diane Bouchacourt, Pietro Astolfi, Kyunghyun Cho, Yann Lecun</p>	24374
<p>Swift4D: Adaptive Divide-And-Conquer Gaussian Splatting for Compact and Efficient Reconstruction of Dynamic Scene</p> <p>Jiahao Wu, Rui Peng, Zhiyan Wang, Lu Xiao, Luyang Tang, Jinbo Yan, Kaiqiang Xiong, Ronggang Wang</p>	24396
<p>Large Language Models Meet Symbolic Provers for Logical Reasoning Evaluation</p> <p>Chengwen Qi, Ren Ma, Bowen Li, He Du, Binyuan Hui, Jinwang Wu, Yuanjun Laili, Conghui He</p>	24414
<p>Hessian-Free Online Certified Unlearning</p> <p>Xinbao Qiao, Meng Zhang, Ming Tang, Ermin Wei</p>	24434
<p>The KoLMogorov Test: Compression by Code Generation</p> <p>Ori Yoran, Kunhao Zheng, Fabian Gloeckle, Jonas Gehring, Gabriel Synnaeve, Taco Cohen</p>	24471
<p>Are Large Vision Language Models Good Game Players?</p> <p>Xinyu Wang, Bohan Zhuang, Qi Wu</p>	24502
<p>Out-Of-Distribution Generalization for Total Variation Based Invariant Risk Minimization</p> <p>Yuanchao Wang, Zhao-Rong Lai, Tianqi Zhong</p>	24540
<p>On Statistical Rates of Conditional Diffusion Transformers: Approximation, Estimation and Minimax Optimality</p> <p>Jerry Yao-Chieh Hu, Weimin Wu, Yi-Chen Lee, Yu-Chao Huang, Minshuo Chen, Han Liu</p>	24557
<p>Dynamical Diffusion: Learning Temporal Dynamics with Diffusion Models</p> <p>Xingzhuo Guo, Yu Zhang, Baixu Chen, Haoran Xu, Jianmin Wang, Mingsheng Long</p>	24642
<p>Dataset Distillation Via Knowledge Distillation: Towards Efficient Self-Supervised Pre-Training of Deep Networks</p> <p>Siddharth Joshi, Jiayi Ni, Baharan Mirzasoleiman</p>	24664
<p>A New Perspective on Shampoo's Preconditioner</p> <p>Depen Morwani, Itai Shapira, Nikhil Vyas, Eran Malach, Sham M. Kakade, Lucas Janson</p>	24691
<p>EqNIO: Subequivariant Neural Inertial Odometry</p> <p>Royina Karegoudra Jayanth, Yinshuang Xu, Ziyun Wang, Evangelos Chatzipantazis, Kostas Daniilidis, Daniel Gehrig</p>	24712
<p>Does Training with Synthetic Data Truly Protect Privacy?</p> <p>Yunpeng Zhao, Jie Zhang</p>	24749
<p>Can LLMs Solve Longer Math Word Problems Better?</p> <p>Xin Xu, Tong Xiao, Zitong Chao, Zhenya Huang, Can Yang, Yang Wang</p>	24767

<p>Diff-PIC: Revolutionizing Particle-In-Cell Nuclear Fusion Simulation with Diffusion Models</p> <p>Chuan Liu, Chunshu Wu, Shihui Cao, Mingkai Chen, James Chenhao Liang, Ang Li, Michael Huang, Chuang Ren, Ying Nian Wu, Dongfang Liu, Tong Geng</p>	24799
<p>Utilitarian Algorithm Configuration for Infinite Parameter Spaces</p> <p>Devon R. Graham, Kevin Leyton-Brown</p>	24813
<p>Foundation Models Secretly Understand Neural Network Weights: Enhancing Hypernetwork Architectures with Foundation Models</p> <p>Jeffrey Gu, Serena Yeung-Levy</p>	24825
<p>CodeMMLU: A Multi-Task Benchmark for Assessing Code Understanding & Reasoning Capabilities of CodeLLMs</p> <p>Dung Manh Nguyen, Thang Chau Phan, Nam Le Hai, Tien-Thong Doan, Nam V. Nguyen, Quang Pham, Nghi D. Q. Bui</p>	24838
<p>Real2Code: Reconstruct Articulated Objects Via Code Generation</p> <p>Zhao Mandi, Yijia Weng, Dominik Bauer, Shuran Song</p>	24897
<p>Towards Robust Alignment of Language Models: Distributionally Robustifying Direct Preference Optimization</p> <p>Junkang Wu, Yuexiang Xie, Zhengyi Yang, Jiancan Wu, Jiawei Chen, Jinyang Gao, Bolin Ding, Xiang Wang, Xiangnan He</p>	24916
<p>Polynomial Composition Activations: Unleashing the Dynamics of Large Language Models</p> <p>Zhijian Zhuo, Ya Wang, Yutao Zeng, Xiaoqing Li, Xun Zhou, Jinwen Ma</p>	24944
<p>How to Evaluate Reward Models for RLHF</p> <p>Evan Frick, Tianle Li, Connor Chen, Wei-Lin Chiang, Anastasios Nikolas Angelopoulos, Jiantao Jiao, Banghua Zhu, Joseph E. Gonzalez, Ion Stoica</p>	24971
<p>Beyond FVD: An Enhanced Evaluation Metrics for Video Generation Distribution Quality</p> <p>Ge Ya Luo, Gian Mario Favero, Zhihao Luo, Alexia Jolicoeur-Martineau, Christopher Pal</p>	25007
<p>Interaction Asymmetry: A General Principle for Learning Composable Abstractions</p> <p>Jack Brady, Julius Von Kügelgen, Sebastien Lachapelle, Simon Buchholz, Thomas Kipf, Wieland Brendel</p>	25042
<p>CoInD: Enabling Logical Compositions in Diffusion Models</p> <p>Sachit Gaudi, Gautam Sreekumar, Vishnu Boddeti</p>	25097
<p>Generative Verifiers: Reward Modeling as Next-Token Prediction</p> <p>Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, Rishabh Agarwal</p>	25130
<p>One-Prompt-One-Story: Free-Lunch Consistent Text-To-Image Generation Using a Single Prompt</p> <p>Tao Liu, Kai Wang, Senmao Li, Joost Van De Weijer, Fahad Shahbaz Khan, Shiqi Yang, Yaxing Wang, Jian Yang, Ming-Ming Cheng</p>	25160

<p>Certified Robustness Under Bounded Levenshtein Distance</p> <p><i>Eliás Abad Rocamora, Grigorios Chrysos, Volkan Cevher</i></p>	25188
<p>Restyling Unsupervised Concept Based Interpretable Networks with Generative Models</p> <p><i>Jayneel Parekh, Quentin Bouniot, Pavlo Mozharovskyi, Alasdair Newson, Florence D'Alché-Buc</i></p>	25211
<p>GeoLP: A Synthetic Dataset to Guide Large-Scale Rule Induction</p> <p><i>Si Chen, Richong Zhang, Xu Zhang</i></p>	25246
<p>Online Preference Alignment for Language Models Via Count-Based Exploration</p> <p><i>Chenjia Bai, Yang Zhang, Shuang Qiu, Qiaosheng Zhang, Kang Xu, Xuelong Li</i></p>	25267
<p>On-The-Fly Preference Alignment Via Principle-Guided Decoding</p> <p><i>Mingye Zhu, Yi Liu, Lei Zhang, Junbo Guo, Zhendong Mao</i></p>	25290
<p>DuoAttention: Efficient Long-Context LLM Inference with Retrieval and Streaming Heads</p> <p><i>Guangxuan Xiao, Jiaming Tang, Jingwei Zuo, Junxian Guo, Shang Yang, Haotian Tang, Yao Fu, Song Han</i></p>	25314
<p>Improving Neural Optimal Transport Via Displacement Interpolation</p> <p><i>Jaemoo Choi, Yongxin Chen, Jaewoong Choi</i></p>	25340
<p>Innovative Thinking, Infinite Humor: Humor Research of Large Language Models Through Structured Thought Leaps</p> <p><i>Han Wang, Yilin Zhao, Dian Li, Xiaohan Wang, Sinbadliu, Xuguang Lan, Hui Wang</i></p>	25365
<p>BrainACTIV: Identifying Visuo-Semantic Properties Driving Cortical Selectivity Using Diffusion-Based Image Manipulation</p> <p><i>Diego Garcia Cerdas, Christina Sartzetaki, Magnus Petersen, Gemma Roig, Pascal Mettes, Iris Groen</i></p>	25390
<p>Subgraph Federated Learning for Local Generalization</p> <p><i>Sungwon Kim, Yoonho Lee, Yunhak Oh, Namkyeong Lee, Sukwon Yun, Junseok Lee, Sein Kim, Carl Yang, Chanyoung Park</i></p>	25426
<p>Towards Interpreting Visual Information Processing in Vision-Language Models</p> <p><i>Clement Neo, Luke Ong, Philip Torr, Mor Geva, David Krueger, Fazl Barez</i></p>	25461
<p>LiveCodeBench: Holistic and Contamination Free Evaluation of Large Language Models for Code</p> <p><i>Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, Ion Stoica</i></p>	25479
<p>ELICIT: LLM Augmentation Via External In-Context Capability</p> <p><i>Futing Wang, Jianhao Yan, Yue Zhang, Tao Lin</i></p>	25520
<p>Robust Barycenter Estimation Using Semi-Unbalanced Neural Optimal Transport</p> <p><i>Milena Gazdieva, Jaemoo Choi, Alexander Kolesov, Jaewoong Choi, Petr Mokrov, Alexander Korotin</i></p>	25554

<p>CR-CTC: Consistency Regularization on CTC for Improved Speech Recognition</p> <p>Zengwei Yao, Wei Kang, Xiaoyu Yang, Fangjun Kuang, Liyong Guo, Han Zhu, Zengrui Jin, Zhaoqing Li, Long Lin, Daniel Povey</p>	25584
<p>Distributed Speculative Inference (DSI): Speculation Parallelism for Provably Faster Lossless Language Model Inference</p> <p>Nadav Timor, Jonathan Mamou, Daniel Korat, Moshe Berchansky, Oren Pereg, Moshe Wasserblat, Tomer Galanti, Michal Gordon-Kiwkowitz, David Harel</p>	25603
<p>Digi-Q: Learning VLM Q-Value Functions for Training Device-Control Agents</p> <p>Hao Bai, Yifei Zhou, Li Erran Li, Sergey Levine, Aviral Kumar</p>	25628
<p>Occlusion-Aware Non-Rigid Point Cloud Registration Via Unsupervised Neural Deformation Correntropy</p> <p>Mingyang Zhao, Gaofeng Meng, Dong-Ming Yan</p>	25649
<p>DRL: Decomposed Representation Learning for Tabular Anomaly Detection</p> <p>Hangting Ye, He Zhao, Wei Fan, Mingyuan Zhou, Dan Dan Guo, Yi Chang</p>	25674
<p>REvolve: Reward Evolution with Large Language Models Using Human Feedback</p> <p>Rishi Hazra, Alkis Sygkounas, Andreas Persson, Amy Loutfi, Pedro Zuidberg Dos Martires</p>	25710
<p>Fine-Grained Verifiers: Preference Modeling as Next-Token Prediction in Vision-Language Alignment</p> <p>Chenhang Cui, An Zhang, Yiyang Zhou, Zhaorun Chen, Gelei Deng, Huaxiu Yao, Tat-Seng Chua</p>	25752
<p>Training Language Models to Self-Correct Via Reinforcement Learning</p> <p>Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, Lei M Zhang, Kay McKinney, Disha Shrivastava, Cosmin Paduraru, George Tucker, Doina Precup, Feryal Behbahani, Aleksandra Faust</p>	25776
<p>HyPoGen: Optimization-Biased Hypernetworks for Generalizable Policy Generation</p> <p>Hanxiang Ren, Li Sun, Xulong Wang, Pei Zhou, Zewen Wu, Siyan Dong, Difan Zou, Youyi Zheng, Yanchao Yang</p>	25803
<p>MrSteve: Instruction-Following Agents in Minecraft with What-Where-When Memory</p> <p>Junyeong Park, Junmo Cho, Sungjin Ahn</p>	25840
<p>High-Dynamic Radar Sequence Prediction for Weather Nowcasting Using Spatiotemporal Coherent Gaussian Representation</p> <p>Ziye Wang, Yiran Qin, Lin Zeng, Ruimao Zhang</p>	25879
<p>\$R^2\$-Guard: Robust Reasoning Enabled LLM Guardrail Via Knowledge-Enhanced Logical Reasoning</p> <p>Mintong Kang, Bo Li</p>	25899

<p>How Do Large Language Models Understand Graph Patterns? a Benchmark for Graph Pattern Comprehension</p> <p>Xinnan Dai, Haohao Qu, Yifei Shen, Bohang Zhang, Qihao Wen, Wenqi Fan, Dongsheng Li, Jiliang Tang, Caihua Shan</p>	25917
<p>Diversity Empowers Intelligence: Integrating Expertise of Software Engineering Agents</p> <p>Kexun Zhang, Weiran Yao, Zuxin Liu, Yihao Feng, Zhiwei Liu, Rithesh R N, Tian Lan, Lei Li, Renze Lou, Jiacheng Xu, Bo Pang, Yingbo Zhou, Shelby Heinecke, Silvio Savarese, Huan Wang, Caiming Xiong</p>	25948
<p>When Narrower is Better: The Narrow Width Limit of Bayesian Parallel Branching Neural Networks</p> <p>Zechen Zhang, Haim Sompolinsky</p>	25972
<p>A Deep Generative Learning Approach for Two-Stage Adaptive Robust Optimization</p> <p>Aron Brenner, Rahman Khorramfar, Jennifer Z Sun, Saurabh Amin</p>	25996
<p>Exponential Topology-Enabled Scalable Communication in Multi-Agent Reinforcement Learning</p> <p>Xinran Li, Xiaolu Wang, Chenjia Bai, Jun Zhang</p>	26014
<p>What Should a Neuron Aim For? Designing Local Objective Functions Based on Information Theory</p> <p>Andreas Christian Schneider, Valentin Neuhaus, David Alexander Ehrlich, Abdullah Makkeh, Alexander S Ecker, Viola Priesemann, Michael Wibral</p>	26035
<p>Find a Winning Sign: Sign is All We Need to Win the Lottery</p> <p>Junghun Oh, Sungyong Baik, Kyoung Mu Lee</p>	26059
<p>Feedback Favors the Generalization of Neural ODEs</p> <p>Jindou Jia, Zihan Yang, Meng Wang, Kexin Guo, Jianfei Yang, Xiang Yu, Lei Guo</p>	26074
<p>Towards Faster Decentralized Stochastic Optimization with Communication Compression</p> <p>Rustem Islamov, Yuan Gao, Sebastian U Stich</p>	26101
<p>Class Distribution-Induced Attention Map for Open-Vocabulary Semantic Segmentations</p> <p>Dong Un Kang, Hayeon Kim, Se Young Chun</p>	26135
<p>Geometric Inductive Biases of Deep Networks: The Role of Data and Architecture</p> <p>Sajad Movahedi, Antonio Orvieto, Seyed-Mohsen Moosavi-Dezfooli</p>	26159
<p>Zeroth-Order Policy Gradient for Reinforcement Learning from Human Feedback Without Reward Inference</p> <p>Qining Zhang, Lei Ying</p>	26193
<p>Adversarial Latent Feature Augmentation for Fairness</p> <p>Hoin Jung, Junyi Chai, Xiaoqian Wang</p>	26208

<p>Wasserstein Distances, Neuronal Entanglement, and Sparsity</p> <p><i>Shashata Sawmya, Linghao Kong, Iliia Markov, Dan Alistarh, Nir N Shavit</i></p>	26244
<p>Comparing Noisy Neural Population Dynamics Using Optimal Transport Distances</p> <p><i>Amin Nejatbakhsh, Victor Geada, Alex H Williams, David Lipshutz</i></p>	26275
<p>Multiview Equivariance Improves 3D Correspondence Understanding with Minimal Feature Finetuning</p> <p><i>Yang You, Yixin Li, Congyue Deng, Yue Wang, Leonidas Guibas</i></p>	26302
<p>Rethinking Reward Model Evaluation: Are We Barking Up the Wrong Tree?</p> <p><i>Xueru Wen, Jie Lou, Yaojie Lu, Hongyu Lin, Xingyu, Xinyu Lu, Ben He, Xianpei Han, Debing Zhang, Le Sun</i></p>	26316
<p>SPDIM: Source-Free Unsupervised Conditional and Label Shift Adaptation in EEG</p> <p><i>Shanglin Li, Motoaki Kawanabe, Reinmar J Kobler</i></p>	26338
<p>Towards Generalization Bounds of GCNs for Adversarially Robust Node Classification</p> <p><i>Wen Wen, Han Li, Tieliang Gong, Hong Chen</i></p>	26363
<p>ChatQA 2: Bridging the Gap to Proprietary LLMs in Long Context and RAG Capabilities</p> <p><i>Peng Xu, Wei Ping, Xianchao Wu, Chejian Xu, Zihan Liu, Mohammad Shoeybi, Bryan Catanzaro</i></p>	26405
<p>VL-ICL Bench: The Devil in the Details of Multimodal In-Context Learning</p> <p><i>Yongshuo Zong, Ondrej Bohdal, Timothy Hospedales</i></p>	26420
<p>Bridging Context Gaps: Leveraging Coreference Resolution for Long Contextual Understanding</p> <p><i>Yanming Liu, Xinyue Peng, Jiannan Cao, Shi Bo, Yanxin Shen, Tianyu Du, Sheng Cheng, Xun Wang, Jianwei Yin, Xuhong Zhang</i></p>	26463
<p>OPTAMI: Global Superlinear Convergence of High-Order Methods</p> <p><i>Dmitry Kamzolov, Artem Agafonov, Dmitry Pasechnyuk, Alexander Gasnikov, Martin Takác</i></p>	26480
<p>CertainlyUncertain: A Benchmark and Metric for Multimodal Epistemic and Aleatoric Awareness</p> <p><i>Khyathi Chandu, Linjie Li, Anas Awadalla, Ximing Lu, Jae Sung Park, Jack Hessel, Lijuan Wang, Yejin Choi</i></p>	26510
<p>TAID: Temporally Adaptive Interpolated Distillation for Efficient Knowledge Transfer in Language Models</p> <p><i>Makoto Shing, Kou Misaki, Han Bao, Sho Yokoi, Takuya Akiba</i></p>	26540
<p>Scrutinize What We Ignore: Reining in Task Representation Shift of Context-Based Offline Meta Reinforcement Learning</p> <p><i>Hai Zhang, Boyuan Zheng, Tianying Ji, Jinhang Liu, Anqi Guo, Junqiao Zhao, Lanqing Li</i></p>	26564

<p>EQMARL: Entangled Quantum Multi-Agent Reinforcement Learning for Distributed Cooperation Over Quantum Channels</p> <p>Alexander Derieux, Walid Saad</p>	26586
<p>Open-YOLO 3D: Towards Fast and Accurate Open-Vocabulary 3D Instance Segmentation</p> <p>Mohamed El Amine Boudjoghra, Angela Dai, Jean Lahoud, Hisham Cholakkal, Rao Muhammad Anwer, Salman Khan, Fahad Shahbaz Khan</p>	26618
<p>CViT: Continuous Vision Transformer for Operator Learning</p> <p>Sifan Wang, Jacob H Seidman, Shyam Sankaran, Hanwen Wang, George J. Pappas, Paris Perdikaris</p>	26632
<p>Self-Play with Execution Feedback: Improving Instruction-Following Capabilities of Large Language Models</p> <p>Guanting Dong, Keming Lu, Chengpeng Li, Tingyu Xia, Bowen Yu, Chang Zhou, Jingren Zhou</p>	26668
<p>MagicDec: Breaking the Latency-Throughput Tradeoff for Long Context Generation with Speculative Decoding</p> <p>Ranajoy Sadhukhan, Jian Chen, Zhuoming Chen, Vashisth Tiwari, Ruihang Lai, Jinyuan Shi, Ian En-Hsu Yen, Avner May, Tianqi Chen, Beidi Chen</p>	26696
<p>HexGen-2: Disaggregated Generative Inference of LLMs in Heterogeneous Environment</p> <p>Youhe Jiang, Ran Yan, Binhang Yuan</p>	26712
<p>X-ALMA: Plug & Play Modules and Adaptive Rejection for Quality Translation at Scale</p> <p>Haoran Xu, Kenton Murray, Philipp Koehn, Hieu Hoang, Akiko Eriguchi, Huda Khayrallah</p>	26730
<p>Bias Mitigation in Graph Diffusion Models</p> <p>Meng Yu, Kun Zhan</p>	26757
<p>CFD: Learning Generalized Molecular Representation Via Concept-Enhanced Feedback Disentanglement</p> <p>Aming Wu, Cheng Deng</p>	26775
<p>Masked Diffusion Models Are Secretly Time-Agnostic Masked Models and Exploit Inaccurate Categorical Sampling</p> <p>Kaiwen Zheng, Yongxin Chen, Hanzi Mao, Ming-Yu Liu, Jun Zhu, Qinsheng Zhang</p>	26796
<p>Accelerating Neural Network Training: An Analysis of the AlgoPerf Competition</p> <p>Priya Kasimbeg, Frank Schneider, Runa Eschenhagen, Juhan Bae, Chandramouli Shama Sastry, Mark Saroufim, Boyuan Feng, Less Wright, Edward Z. Yang, Zachary Nado, Sourabh Medapati, Philipp Hennig, Michael Rabbat, George E. Dahl</p>	26838
<p>BodyGen: Advancing Towards Efficient Embodiment Co-Design</p> <p>Haofei Lu, Zhe Wu, Junliang Xing, Jianshu Li, Ruoyu Li, Zhe Li, Yuanchun Shi</p>	26861
<p>OCCAM: Towards Cost-Efficient and Accuracy-Aware Classification Inference</p> <p>Dujian Ding, Bicheng Xu, Laks V. S. Lakshmanan</p>	26883

<p>Accelerating Training with Neuron Interaction and Nowcasting Networks</p> <p><i>Boris Knyazev, Abhinav Moudgil, Guillaume Lajoie, Eugene Belilovsky, Simon Lacoste-Julien</i></p>	26907
<p>Continuous Autoregressive Modeling with Stochastic Monotonic Alignment for Speech Synthesis</p> <p><i>Weiwei Lin, Chenhang He</i></p>	26927
<p>Tight Time Complexities in Parallel Stochastic Optimization with Arbitrary Computation Dynamics</p> <p><i>Alexander Tyurin</i></p>	26945
<p>Credal Wrapper of Model Averaging for Uncertainty Estimation in Classification</p> <p><i>Kaizheng Wang, Fabio Cuzzolin, Keivan Shariatmadar, David Moens, Hans Hallez</i></p>	26988
<p>CHASE-SQL: Multi-Path Reasoning and Preference Optimized Candidate Selection in Text-To-SQL</p> <p><i>Mohammadreza Pourreza, Hailong Li, Ruoxi Sun, Yeounoh Chung, Shayan Talaei, Gaurav Tarlok Kakkar, Yu Gan, Amin Saberi, Fatma Ozcan, Sercan O Arik</i></p>	27017
<p>Strength Estimation and Human-Like Strength Adjustment in Games</p> <p><i>Chun Jung Chen, Chung-Chin Shih, Ti-Rong Wu</i></p>	27048
<p>Beyond Single Concept Vector: Modeling Concept Subspace in LLMs with Gaussian Distribution</p> <p><i>Haiyan Zhao, Heng Zhao, Bo Shen, Ali Payani, Fan Yang, Mengnan Du</i></p>	27066
<p>Regressing the Relative Future: Efficient Policy Optimization for Multi-Turn RLHF</p> <p><i>Zhaolin Gao, Wenhao Zhan, Jonathan Daniel Chang, Gokul Swamy, Kianté Brantley, Jason D. Lee, Wen Sun</i></p>	27095
<p>Towards Homogeneous Lexical Tone Decoding from Heterogeneous Intracranial Recordings</p> <p><i>Di Wu, Siyuan Li, Chen Feng, Lu Cao, Yue Zhang, Jie Yang, Mohamad Sawan</i></p>	27135
<p>From Risk to Uncertainty: Generating Predictive Uncertainty Measures Via Bayesian Estimation</p> <p><i>Nikita Kotelevskii, Vladimir Kondratyev, Martin Takác, Eric Moulines, Maxim Panov</i></p>	27157
<p>Efficient Learning with Sine-Activated Low-Rank Matrices</p> <p><i>Yiping Ji, Hemanth Saratchandran, Cameron Gordon, Zeyu Zhang, Simon Lucey</i></p>	27194
<p>Knowledge Distillation with Multi-Granularity Mixture of Priors for Image Super-Resolution</p> <p><i>Simiao Li, Yun Zhang, Wei Li, Hanting Chen, Wenjia Wang, Bingyi Jing, Shaohui Lin, Jie Hu</i></p>	27216
<p>Modeling Future Conversation Turns to Teach LLMs to Ask Clarifying Questions</p> <p><i>Michael Jq Zhang, W. Bradley Knox, Eunsol Choi</i></p>	27233
<p>Progressive Compression with Universally Quantized Diffusion Models</p> <p><i>Yibo Yang, Justus Will, Stephan Mandt</i></p>	27254

<p>Can Textual Gradient Work in Federated Learning?</p> <p>Minghui Chen, Ruinan Jin, Wenlong Deng, Yuanyuan Chen, Zhi Huang, Han Yu, Xiaoxiao Li</p>	27276
<p>Overcoming Lower-Level Constraints in Bilevel Optimization: A Novel Approach with Regularized Gap Functions</p> <p>Wei Yao, Haian Yin, Shangzhi Zeng, Jin Zhang</p>	27298
<p>Decentralized Sporadic Federated Learning: A Unified Algorithmic Framework with Convergence Guarantees</p> <p>Shahryar Zehtabi, Dong-Jun Han, Rohit Parasnis, Seyyedali Hosseinalipour, Christopher Brinton</p>	27332
<p>Flash Inference: Near Linear Time Inference for Long Convolution Sequence Models and Beyond</p> <p>Costin-Andrei Oncescu, Sanket Purandare, Stratos Idreos, Sham M. Kakade</p>	27402
<p>Physics-Aligned Field Reconstruction with Diffusion Bridge</p> <p>Zeyu Li, Hongkun Dou, Shen Fang, Wang Han, Yue Deng, Lijun Yang</p>	27428
<p>Speech Robust Bench: A Robustness Benchmark for Speech Recognition</p> <p>Muhammad A Shah, David Solans Noguero, Mikko A. Heikkilä, Bhiksha Raj, Nicolas Kourtellis</p>	27465
<p>Exploiting Distribution Constraints for Scalable and Efficient Image Retrieval</p> <p>Mohammad Omama, Po-Han Li, Sandeep P. Chinchali</p>	27492
<p>Efficient and Context-Aware Label Propagation for Zero-/Few-Shot Training-Free Adaptation of Vision-Language Model</p> <p>Yushu Li, Yongyi Su, Adam Goodge, Kui Jia, Xun Xu</p>	27508
<p>Identifying Latent State Transitions in Non-Linear Dynamical Systems</p> <p>Çaglar Hizli, Çagatay Yildiz, Matthias Bethge, S. T. John, Pekka Marttinen</p>	27529
<p>Gyrogrouop Batch Normalization</p> <p>Ziheng Chen, Yue Song, Xiaojun Wu, Nicu Sebe</p>	27559
<p>Imputation for Prediction: Beware of Diminishing Returns.</p> <p>Marine Le Morvan, Gael Varoquaux</p>	27597
<p>MotionDreamer: One-To-Many Motion Synthesis with Localized Generative Masked Transformer</p> <p>Yilin Wang, Chuan Guo, Yuxuan Mu, Muhammad Gohar Javed, Xinxin Zuo, Juwei Lu, Hai Jiang, Li Cheng</p>	27637
<p>Mask-DPO: Generalizable Fine-Grained Factuality Alignment of LLMs</p> <p>Yuzhe Gu, Wenwei Zhang, Chengqi Lyu, Dahua Lin, Kai Chen</p>	27654
<p>QA-Calibration of Language Model Confidence Scores</p> <p>Putra Manggala, Atalanti A. Mastakouri, Elke Kirschbaum, Shiva Kasiviswanathan, Aaditya Ramdas</p>	27673

<p>Controlling Space and Time with Diffusion Models</p> <p><i>Daniel Watson, Saurabh Saxena, Lala Li, Andrea Tagliasacchi, David J. Fleet</i></p>	27697
<p>Denoising with a Joint-Embedding Predictive Architecture</p> <p><i>Dengsheng Chen, Jie Hu, Xiaoming Wei, Enhua Wu</i></p>	27721
<p>Efficient Training of Neural Stochastic Differential Equations by Matching Finite Dimensional Distributions</p> <p><i>Jianxin Zhang, Josh Viktorov, Doosan Jung, Emily Pitler</i></p>	27760
<p>Learning Shape-Independent Transformation Via Spherical Representations for Category-Level Object Pose Estimation</p> <p><i>Huan Ren, Wenfei Yang, Xiang Liu, Shifeng Zhang, Tianzhu Zhang</i></p>	27773
<p>GSE: Group-Wise Sparse and Explainable Adversarial Attacks</p> <p><i>Shpresim Sadiku, Moritz Wagner, Sebastian Pokutta</i></p>	27794
<p>CLIBD: Bridging Vision and Genomics for Biodiversity Monitoring at Scale</p> <p><i>Zeming Gong, Austin Wang, Xiaoliang Huo, Joakim Bruslund Haurum, Scott C. Lowe, Graham W. Taylor, Angel X Chang</i></p>	27811
<p>Probabilistic Language-Image Pre-Training</p> <p><i>Sanghyuk Chun, Wonjae Kim, Song Park, Sangdoon Yun</i></p>	27842
<p>Not All Language Model Features Are One-Dimensionally Linear</p> <p><i>Joshua Engels, Eric J Michaud, Isaac Liao, Wes Gurnee, Max Tegmark</i></p>	27874
<p>MMFakeBench: A Mixed-Source Multimodal Misinformation Detection Benchmark for LVLMS</p> <p><i>Xuannan Liu, Zekun Li, Pei Pei Li, Huaibo Huang, Shuhan Xia, Xing Cui, Linzhi Huang, Weihong Deng, Zhaofeng He</i></p>	27906
<p>Diffusing to the Top: Boost Graph Neural Networks with Minimal Hyperparameter Tuning</p> <p><i>Lequan Lin, Dai Shi, Andi Han, Zhiyong Wang, Junbin Gao</i></p>	27932
<p>Hybrid Regularization Improves Diffusion-Based Inverse Problem Solving</p> <p><i>Hongkun Dou, Zeyu Li, Jinyang Du, Lijun Yang, Wen Yao, Yue Deng</i></p>	27967
<p>Morphing Tokens Draw Strong Masked Image Models</p> <p><i>Taekyung Kim, Byeongho Heo, Dongyoon Han</i></p>	28002
<p>MallowsPO: Fine-Tune Your LLM with Preference Dispersions</p> <p><i>Haoxian Chen, Hanyang Zhao, Henry Lam, David Yao, Wenpin Tang</i></p>	28026
<p>Policy Design in Long-Run Welfare Dynamics</p> <p><i>Jiduan Wu, Rediet Abebe, Moritz Hardt, Ana-Andreea Stoica</i></p>	28059
<p>Exact Certification of (Graph) Neural Networks Against Label Poisoning</p> <p><i>Mahalakshmi Sabanayagam, Lukas Gosch, Stephan Günnemann, Debarghya Ghoshdastidar</i></p>	28088

<p>Refine-By-Align: Reference-Guided Artifacts Refinement Through Semantic Alignment</p> <p>Yizhi Song, Liu He, Zhifei Zhang, Soo Ye Kim, He Zhang, Wei Xiong, Zhe Lin, Brian L. Price, Scott Cohen, Jianming Zhang, Daniel Aliaga</p>	28117
<p>DELTA: DENSE EFFICIENT LONG-RANGE 3D TRACKING FOR ANY VIDEO</p> <p>Tuan Duc Ngo, Peiye Zhuang, Evangelos Kalogerakis, Chuang Gan, Sergey Tulyakov, Hsin-Ying Lee, Chaoyang Wang</p>	28132
<p>TIPS: Text-Image Pretraining with Spatial Awareness</p> <p>Kevis-Kokitsi Maninis, Kaifeng Chen, Soham Ghosh, Arjun Karpur, Koert Chen, Ye Xia, Bingyi Cao, Daniel Salz, Guangxing Han, Jan Dlabal, Dan Gnanapragasam, Mojtaba Seyedhosseini, Howard Zhou, Andre Araujo</p>	28154
<p>Gaussian-Based Instance-Adaptive Intensity Modeling for Point-Supervised Facial Expression Spotting</p> <p>Yicheng Deng, Hideaki Hayashi, Hajime Nagahara</p>	28178
<p>Large Language Models Assume People Are More Rational than We Really Are</p> <p>Ryan Liu, Jiayi Geng, Joshua Peterson, Ilya Sucholutsky, Thomas L. Griffiths</p>	28193
<p>Dynamic Sparse Training Versus Dense Training: The Unexpected Winner in Image Corruption Robustness</p> <p>Boqian Wu, Qiao Xiao, Shunxin Wang, Nicola Strisciuglio, Mykola Pechenizkiy, Maurice Van Keulen, Decebal Constantin Mocanu, Elena Mocanu</p>	28220
<p>FOSP: Fine-Tuning Offline Safe Policy Through World Models</p> <p>Chenyang Cao, Yucheng Xin, Silang Wu, Longxiang He, Zichen Yan, Junbo Tan, Xueqian Wang</p>	28247
<p>Estimating the Probabilities of Rare Outputs in Language Models</p> <p>Gabriel Wu, Jacob Hilton</p>	28279
<p>Fast Feedforward 3D Gaussian Splatting Compression</p> <p>Yihang Chen, Qianyi Wu, Mengyao Li, Weiyao Lin, Mehrtash Harandi, Jianfei Cai</p>	28308
<p>In-Context Time Series Predictor</p> <p>Jiecheng Lu, Yan Sun, Shihao Yang</p>	28322
<p>CodePlan: Unlocking Reasoning Potential in Large Language Models by Scaling Code-Form Planning</p> <p>Jiaxin Wen, Jian Guan, Hongning Wang, Wei Wu, Minlie Huang</p>	28349
<p>Interactive Adjustment for Human Trajectory Prediction with Individual Feedback</p> <p>Jianhua Sun, Yuxuan Li, Liang Chai, Cewu Lu</p>	28372
<p>ADMM for Structured Fractional Minimization</p> <p>Ganzhao Yuan</p>	28392

<p>Exploring the Effectiveness of Object-Centric Representations in Visual Question Answering: Comparative Insights with Foundation Models</p> <p>Amir Mohammad Karimi Mamaghan, Samuele Papa, Karl Henrik Johansson, Stefan Bauer, Andrea Dittadi</p>	28433
<p>Faster Algorithms for Structured Linear and Kernel Support Vector Machines</p> <p>Yuzhou Gu, Zhao Song, Lichen Zhang</p>	28469
<p>Beyond Next Token Prediction: Patch-Level Training for Large Language Models</p> <p>Chenze Shao, Fandong Meng, Jie Zhou</p>	28538
<p>PooDLe: Pooled and Dense Self-Supervised Learning from Naturalistic Videos</p> <p>Alex N Wang, Christopher Hoang, Yuwen Xiong, Yann Lecun, Mengye Ren</p>	28557
<p>Efficient Policy Evaluation with Safety Constraint for Reinforcement Learning</p> <p>Claire Chen, Shuze Liu, Shangtong Zhang</p>	28580
<p>How Feature Learning Can Improve Neural Scaling Laws</p> <p>Blake Bordelon, Alexander Atanasov, Cengiz Pehlevan</p>	28602
<p>Clique Number Estimation Via Differentiable Functions of Adjacency Matrix Permutations</p> <p>Indradyumna Roy, Eeshaan Jain, Soumen Chakrabarti, Abir De</p>	28633
<p>MLLMs Know Where to Look: Training-Free Perception of Small Visual Details with Multimodal LLMs</p> <p>Jiarui Zhang, Mahyar Khayatkhoei, Prateek Chhikara, Filip Ilievski</p>	28669
<p>Quantum-PEFT: Ultra Parameter-Efficient Fine-Tuning</p> <p>Toshiaki Koike-Akino, Francesco Tonin, Yongtao Wu, Zhengqing Wu, Leyla Naz Candogan, Volkan Cevher</p>	28689
<p>SyllableLM: Learning Coarse Semantic Units for Speech Language Models</p> <p>Alan Baade, Puyuan Peng, David Harwath</p>	28711
<p>Training-Free Activation Sparsity in Large Language Models</p> <p>James Liu, Pragaash Ponnusamy, Tianle Cai, Han Guo, Yoon Kim, Ben Athiwaratkun</p>	28733
<p>Transformer-Squared: Self-Adaptive LLMs</p> <p>Qi Sun, Edoardo Cetin, Yujin Tang</p>	28754
<p>Provable Unlearning in Topic Modeling and Downstream Tasks</p> <p>Stanley Wei, Sadhika Malladi, Sanjeev Arora, Amartya Sanyal</p>	28772
<p>Data Selection Via Optimal Control for Language Models</p> <p>Yuxian Gu, Li Dong, Hongning Wang, Yaru Hao, Qingxiu Dong, Furu Wei, Minlie Huang</p>	28805
<p>Limits of Deep Learning: Sequence Modeling Through the Lens of Complexity Theory</p> <p>Nikola Zubic, Federico Soldà, Aurelio Sulser, Davide Scaramuzza</p>	28833

<p>Reinforcement Learning with Combinatorial Actions for Coupled Restless Bandits</p> <p><i>Lily Xu, Bryan Wilder, Elias Boutros Khalil, Milind Tambe</i></p>	28864
<p>Decoupling Layout from Glyph in Online Chinese Handwriting Generation</p> <p><i>Minsi Ren, Yan-Ming Zhang, Yi Chen</i></p>	28886
<p>Locality-Aware Gaussian Compression for Fast and High-Quality Rendering</p> <p><i>Seungjoo Shin, Jaesik Park, Sunghyun Cho</i></p>	28904
<p>Competing Large Language Models in Multi-Agent Gaming Environments</p> <p><i>Jen-Tse Huang, Eric John Li, Man Ho Lam, Tian Liang, Wenxuan Wang, Youliang Yuan, Wenxiang Jiao, Xing Wang, Zhaopeng Tu, Michael Lyu</i></p>	28932
<p>MixMax: Distributional Robustness in Function Space Via Optimal Data Mixtures</p> <p><i>Anvith Thudi, Chris J. Maddison</i></p>	28975
<p>Sequential Controlled Langevin Diffusions</p> <p><i>Junhua Chen, Lorenz Richter, Julius Berner, Denis Blessing, Gerhard Neumann, Anima Anandkumar</i></p>	28996
<p>Reasoning with Latent Thoughts: On the Power of Looped Transformers</p> <p><i>Nikunj Saunshi, Nishanth Dikkala, Zhiyuan Li, Sanjiv Kumar, Sashank J. Reddi</i></p>	29031
<p>Fine-Tuning Attention Modules Only: Enhancing Weight Disentanglement in Task Arithmetic</p> <p><i>Ruochen Jin, Bojian Hou, Jiancong Xiao, Weijie J Su, Li Shen</i></p>	29058
<p>Enhancing Document Understanding with Group Position Embedding: A Novel Approach to Incorporate Layout Information</p> <p><i>Yuke Zhu, Yue Zhang, Dongdong Liu, Chi Xie, Zihua Xiong, Bo Zheng, Sheng Guo</i></p>	29077
<p>Representation Alignment for Generation: Training Diffusion Transformers is Easier than You Think</p> <p><i>Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, Saining Xie</i></p>	29100
<p>Cocoon: Robust Multi-Modal Perception with Uncertainty-Aware Sensor Fusion</p> <p><i>Minkyung Cho, Yulong Cao, Jiachen Sun, Qingzhao Zhang, Marco Pavone, Jeong Joon Park, Heng Yang, Zhuoqing Mao</i></p>	29143
<p>Improving Data Efficiency Via Curating LLM-Driven Rating Systems</p> <p><i>Jinlong Pang, Jiaheng Wei, Ankit Shah, Zhaowei Zhu, Yaxuan Wang, Chen Qian, Yang Liu, Yujia Bao, Wei Wei</i></p>	29167
<p>Improving Deep Regression with Tightness</p> <p><i>Shihao Zhang, Yuguang Yan, Angela Yao</i></p>	29198
<p>Gaussian Splatting Lucas-Kanade</p> <p><i>Liuyue Xie, Joel Julin, Koichiro Niinuma, Laszlo Attila Jeni</i></p>	29215

<p>OmniSep: Unified Omni-Modality Sound Separation with Query-Mixup</p> <p>Xize Cheng, Siqi Zheng, Zehan Wang, Minghui Fang, Ziang Zhang, Rongjie Huang, Shengpeng Ji, Jialong Zuo, Tao Jin, Zhou Zhao</p>	29235
<p>Continual Slow-And-Fast Adaptation of Latent Neural Dynamics (CoSFan): Meta-Learning What-How & When to Adapt</p> <p>Ryan Missel, Linwei Wang</p>	29253
<p>Balancing Act: Diversity and Consistency in Large Language Model Ensembles</p> <p>Ahmed Abdulaal, Chen Jin, Nina Montaña-Brown, Aryo Pradipta Gema, Daniel C. Castro, Daniel C. Alexander, Philip Alexander Teare, Tom Diethel, Dino Oglie, Amrutha Saseendran</p>	29287
<p>Noise-Conditioned Energy-Based Annealed Rewards (NEAR): A Generative Framework for Imitation Learning from Observation</p> <p>Anish Abhijit Diwan, Julen Urain, Jens Kober, Jan Peters</p>	29320
<p>OATS: Outlier-Aware Pruning Through Sparse and Low Rank Decomposition</p> <p>Stephen Zhang, Vardan Papyan</p>	29342
<p>Bootstrapping Language Models with DPO Implicit Rewards</p> <p>Changyu Chen, Zichen Liu, Chao Du, Tianyu Pang, Qian Liu, Arunesh Sinha, Pradeep Varakantham, Min Lin</p>	29368
<p>FACTS: A Factored State-Space Framework for World Modelling</p> <p>Li Nanbo, Firas Laakom, Yucheng Xu, Wenyi Wang, Jürgen Schmidhuber</p>	29386
<p>Single Teacher, Multiple Perspectives: Teacher Knowledge Augmentation for Enhanced Knowledge Distillation</p> <p>Md Imtiaz Hossain, Sharmen Akhter, Choong Seon Hong, Eui-Nam Huh</p>	29415
<p>Learning Structured Universe Graph with Outlier OOD Detection for Partial Matching</p> <p>Zetian Jiang, Jiaxin Lu, Haizhao Fan, Tianzhe Wang, Junchi Yan</p>	29427
<p>Progressive Token Length Scaling in Transformer Encoders for Efficient Universal Segmentation</p> <p>Abhishek Aich, Yumin Suh, Samuel Schuler, Manmohan Chandraker</p>	29443
<p>MOFFlow: Flow Matching for Structure Prediction of Metal-Organic Frameworks</p> <p>Nayoung Kim, Seongsu Kim, Minsu Kim, Jinkyoo Park, Sungsoo Ahn</p>	29461
<p>Calibrating Expressions of Certainty</p> <p>Peiqi Wang, Barbara D. Lam, Yingcheng Liu, Ameneh Asgari-Targhi, Rameswar Panda, William M Wells, Tina Kapur, Polina Golland</p>	29482
<p>Budgeted Online Continual Learning by Adaptive Layer Freezing and Frequency-Based Sampling</p> <p>Minhyuk Seo, Hyunseo Koh, Jonghyun Choi</p>	29513

<p>Aligned Datasets Improve Detection of Latent Diffusion-Generated Images</p> <p>Anirudh Sundara Rajan, Utkarsh Ojha, Jedidiah Schloesser, Yong Jae Lee</p>	29548
<p>Dynamic Assortment Selection and Pricing with Censored Preference Feedback</p> <p>Jung-Hun Kim, Min-Hwan Oh</p>	29567
<p>DEPfold: RNA Secondary Structure Prediction as Dependency Parsing.</p> <p>Ke Wang, Shay B Cohen</p>	29601
<p>Operator Deep Smoothing for Implied Volatility</p> <p>Ruben Wiedemann, Antoine Jacquier, Lukas Gonon</p>	29623
<p>Sensitivity-Constrained Fourier Neural Operators for Forward and Inverse Problems in Parametric Differential Equations</p> <p>Abdolmehdi Behroozi, Chaopeng Shen, Daniel Kifer</p>	29663
<p>Generalizable Human Gaussians from Single-View Image</p> <p>Jinnan Chen, Chen Li, Jianfeng Zhang, Lingting Zhu, Buzhen Huang, Hanlin Chen, Gim Hee Lee</p>	29685
<p>Mitigating the Backdoor Effect for Multi-Task Model Merging Via Safety-Aware Subspace</p> <p>Jinluan Yang, Anke Tang, Didi Zhu, Zhengyu Chen, Li Shen, Fei Wu</p>	29701
<p>Learning Diagrams: A Graphical Language for Compositional Training Regimes</p> <p>Mason Lary, Richard Samuelson, Alexander Wilentz, Alina Zare, Matthew Klawonn, James Fairbanks</p>	29727
<p>Stabilizing Reinforcement Learning in Differentiable Multiphysics Simulation</p> <p>Eliot Xing, Vernon Luk, Jean Oh</p>	29740
<p>DeepTAGE: Deep Temporal-Aligned Gradient Enhancement for Optimizing Spiking Neural Networks</p> <p>Wei Liu, Li Yang, Mingxuan Zhao, Shuxun Wang, Jin Gao, Wenjuan Li, Bing Li, Weiming Hu</p>	29774
<p>BirdSet: A Large-Scale Dataset for Audio Classification in Avian Bioacoustics</p> <p>Lukas Rauch, Raphael Schwinger, Moritz Wirth, René Heinrich, Denis Huseljic, Marek Herde, Jonas Lange, Stefan Kahl, Bernhard Sick, Sven Tomforde, Christoph Scholz</p>	29786
<p>Tool-Planner: Task Planning with Clusters Across Multiple Tools</p> <p>Yanming Liu, Xinyue Peng, Jiannan Cao, Shi Bo, Yuwei Zhang, Xuhong Zhang, Sheng Cheng, Xun Wang, Jianwei Yin, Tianyu Du</p>	29825
<p>Reinforcement Learning for Control of Non-Markovian Cellular Population Dynamics</p> <p>Josiah C Kratz, Jacob Adamczyk</p>	29874
<p>Towards a Learning Theory of Representation Alignment</p> <p>Francesco Insulla, Shuo Huang, Lorenzo Rosasco</p>	29889
<p>The "Law" of the Unconscious Contrastive Learner: Probabilistic Alignment of Unpaired Modalities</p> <p>Yongwei Che, Benjamin Eysenbach</p>	29904

<p>Training Free Exponential Context Extension Via Cascading KV Cache</p> <p>Jeffrey Willette, Heejun Lee, Youngwan Lee, Myeongjae Jeon, Sung Ju Hwang</p>	29924
<p>TC-MoE: Augmenting Mixture of Experts with Ternary Expert Choice</p> <p>Shen Yan, Xingyan Bin, Sijun Zhang, Yisen Wang, Zhouchen Lin</p>	29944
<p>DSBench: How Far Are Data Science Agents from Becoming Data Science Experts?</p> <p>Liqiang Jing, Zhehui Huang, Xiaoyang Wang, Wenlin Yao, Wenhao Yu, Kaixin Ma, Hongming Zhang, Xinya Du, Dong Yu</p>	29960
<p>Leveraging Sub-Optimal Data for Human-In-The-Loop Reinforcement Learning</p> <p>Calarina Muslimani, Matthew E. Taylor</p>	30013
<p>Trivialized Momentum Facilitates Diffusion Generative Modeling on Lie Groups</p> <p>Yuchen Zhu, Tianrong Chen, Lingkai Kong, Evangelos Theodorou, Molei Tao</p>	30049
<p>Gaussian-Det: Learning Closed-Surface Gaussians for 3D Object Detection</p> <p>Hongru Yan, Yu Zheng, Yueqi Duan</p>	30079
<p>FreeVS: Generative View Synthesis on Free Driving Trajectory</p> <p>Qitai Wang, Lue Fan, Yuqi Wang, Yuntao Chen, Zhaoxiang Zhang</p>	30098
<p>SCOPE: A Self-Supervised Framework for Improving Faithfulness in Conditional Text Generation</p> <p>Song Duong, Florian Le Bronnec, Alexandre Allauzen, Vincent Guigue, Alberto Lumbreras, Laure Soulier, Patrick Gallinari</p>	30121
<p>Offline Hierarchical Reinforcement Learning Via Inverse Optimization</p> <p>Carolin Schmidt, Daniele Gammelli, James Harrison, Marco Pavone, Filipe Rodrigues</p>	30152
<p>Convergence and Implicit Bias of Gradient Descent on Continual Linear Classification</p> <p>Hyunji Jung, Hanseul Cho, Chulhee Yun</p>	30183
<p>Unposed Sparse Views Room Layout Reconstruction in the Age of Pretrain Model</p> <p>Yaxuan Huang, Xili Dai, Jianan Wang, Xianbiao Qi, Yixing Yuan, Xiangyu Yue</p>	30250
<p>When Graph Neural Networks Meet Dynamic Mode Decomposition</p> <p>Dai Shi, Lequan Lin, Andi Han, Zhiyong Wang, Yi Guo, Junbin Gao</p>	30268
<p>Controlled LLM Decoding Via Discrete Auto-Regressive Biasing</p> <p>Patrick Pynadath, Ruqi Zhang</p>	30297
<p>Attention Layers Provably Solve Single-Location Regression</p> <p>Pierre Marion, Raphaël Berthier, Gérard Biau, Claire Boyer</p>	30319
<p>RepoGraph: Enhancing AI Software Engineering with Repository-Level Code Graph</p> <p>Siru Ouyang, Wenhao Yu, Kaixin Ma, Zilin Xiao, Zhihan Zhang, Mengzhao Jia, Jiawei Han, Hongming Zhang, Dong Yu</p>	30361

<p>Two Sparse Matrices Are Better than One: Sparsifying Neural Networks with Double Sparse Factorization</p> <p>Vladimír Boža, Vladimír Macko</p>	30385
<p>The Last Iterate Advantage: Empirical Auditing and Principled Heuristic Analysis of Differentially Private SGD</p> <p>Milad Nasr, Thomas Steinke, Borja Balle, Christopher A. Choquette-Choo, Arun Ganesh, Matthew Jagielski, Jamie Hayes, Abhradeep Guha Thakurta, Adam Smith, Andreas Terzis</p>	30404
<p>ImProver: Agent-Based Automated Proof Optimization</p> <p>Riyaz Ahuja, Jeremy Avigad, Prasad Tetali, Sean Welleck</p>	30423
<p>Large Scale Knowledge Washing</p> <p>Yu Wang, Ruihan Wu, Zexue He, Xiusi Chen, Julian McAuley</p>	30446
<p>Learning to Generate Diverse Pedestrian Movements from Web Videos with Noisy Labels</p> <p>Zhizheng Liu, Joe Lin, Wayne Wu, Bolei Zhou</p>	30467
<p>Boosting Methods for Interval-Censored Data with Regression and Classification</p> <p>Yuan Bian, Grace Yi, Wenqing He</p>	30483
<p>Provable Convergence and Limitations of Geometric Tempering for Langevin Dynamics</p> <p>Omar Chehab, Anna Korba, Austin J Stromme, Adrien Vacher</p>	30515
<p>Interpreting Emergent Planning in Model-Free Reinforcement Learning</p> <p>Thomas Bush, Stephen Chung, Usman Anwar, Adrià Garriga-Alonso, David Krueger</p>	30558
<p>Student-Informed Teacher Training</p> <p>Nico Messikommer, Jiaxu Xing, Elie Aljalbout, Davide Scaramuzza</p>	30641
<p>Federated Residual Low-Rank Adaptation of Large Language Models</p> <p>Yunlu Yan, Chun-Mei Feng, Wangmeng Zuo, Rick Siow Mong Goh, Yong Liu, Lei Zhu</p>	30657
<p>Convex Formulations for Training Two-Layer ReLU Neural Networks</p> <p>Karthik Prakhya, Tolga Birdal, Alp Yurtsever</p>	30682
<p>Scaling FP8 Training to Trillion-Token LLMs</p> <p>Maxim Fishman, Brian Chmiel, Ron Banner, Daniel Soudry</p>	30705
<p>KinPFN: Bayesian Approximation of RNA Folding Kinetics Using Prior-Data Fitted Networks</p> <p>Dominik Scheuer, Frederic Runge, Jörg K. H. Franke, Michael T. Wolfinger, Christoph Flamm, Frank Hutter</p>	30719
<p>ViDiT-Q: Efficient and Accurate Quantization of Diffusion Transformers for Image and Video Generation</p> <p>Tianchen Zhao, Tongcheng Fang, Haofeng Huang, Rui Wan, Widyadewi Soedarmadji, Enshu Liu, Shiyao Li, Zinan Lin, Guohao Dai, Shengen Yan, Huazhong Yang, Xuefei Ning, Yu Wang</p>	30754

<p>Time-MoE: Billion-Scale Time Series Foundation Models with Mixture of Experts</p> <p>Xiaoming Shi, Shiyu Wang, Yuqi Nie, Dianqi Li, Zhou Ye, Qingsong Wen, Ming Jin</p>	30785
<p>Efficient Reinforcement Learning with Large Language Model Priors</p> <p>Xue Yan, Yan Song, Xidong Feng, Mengyue Yang, Haifeng Zhang, Haitham Bou Ammar, Jun Wang</p>	30818
<p>Eliminating Oversaturation and Artifacts of High Guidance Scales in Diffusion Models</p> <p>Seyedmorteza Sadat, Otmar Hilliges, Romann M. Weber</p>	30843
<p>Have the VLMs Lost Confidence? a Study of Sycophancy in VLMs</p> <p>Shuo Li, Tao Ji, Xiaoran Fan, Linsheng Lu, Leyi Yang, Yuming Yang, Zhiheng Xi, Rui Zheng, Yuran Wang, Xh. Zhao, Tao Gui, Qi Zhang, Xuanjing Huang</p>	30870
<p>Ensembling Diffusion Models Via Adaptive Feature Aggregation</p> <p>Cong Wang, Kuan Tian, Yonghang Guan, Fei Shen, Zhiwei Jiang, Qing Gu, Jun Zhang</p>	30891
<p>Evaluating Large Language Models Through Role-Guide and Self-Reflection: A Comparative Study</p> <p>Lili Zhao, Yang Wang, Qi Liu, Mengyun Wang, Wei Chen, Zhichao Sheng, Shijin Wang</p>	30914
<p>Bayesian Optimization of Antibodies Informed by a Generative Model of Evolving Sequences</p> <p>Alan Nawzad Amin, Nate Gruver, Yilun Kuang, Yucen Lily Li, Hunter Elliott, Calvin McCarter, Aniruddh Raghu, Peyton Greenside, Andrew Gordon Wilson</p>	30938
<p>Cut Your Losses in Large-Vocabulary Language Models</p> <p>Erik Wijmans, Brody Huval, Alexander Hertzberg, Vladlen Koltun, Philipp Kraehenbuehl</p>	30962
<p>Black-Box Detection of Language Model Watermarks</p> <p>Thibaud Gloaguen, Nikola Jovanovic, Robin Staab, Martin Vechev</p>	30982
<p>Policy Decorator: Model-Agnostic Online Refinement for Large Policy Model</p> <p>Xiu Yuan, Tongzhou Mu, Stone Tao, Yunhao Fang, Mengke Zhang, Hao Su</p>	31006
<p>Dreamweaver: Learning Compositional World Models from Pixels</p> <p>Junyeob Baek, Yi-Fu Wu, Gautam Singh, Sungjin Ahn</p>	31042
<p>Large Convolutional Model Tuning Via Filter Subspace</p> <p>Wei Chen, Zichen Miao, Qiang Qiu</p>	31067
<p>CFG++: Manifold-Constrained Classifier Free Guidance for Diffusion Models</p> <p>Hyungjin Chung, Jeongsol Kim, Geon Yeong Park, Hyelin Nam, Jong Chul Ye</p>	31096
<p>PAD: Personalized Alignment of LLMs at Decoding-Time</p> <p>Ruizhe Chen, Xiaotian Zhang, Meng Luo, Wenhao Chai, Zuozhu Liu</p>	31123
<p>Beyond Correlation: The Impact of Human Uncertainty in Measuring the Effectiveness of Automatic Evaluation and LLM-As-A-Judge</p> <p>Aparna Elangovan, Lei Xu, Jongwoo Ko, Mahsa Elyasi, Ling Liu, Sravan Babu Bodapati, Dan Roth</p>	31149

<p>Beyond Worst-Case Dimensionality Reduction for Sparse Vectors</p> <p>Sandeep Silwal, David Woodruff, Qiuyi Zhang</p>	31179
<p>MGMapNet: Multi-Granularity Representation Learning for End-To-End Vectorized HD Map Construction</p> <p>Jing Yang, Minyue Jiang, Sen Yang, Xiao Tan, Yingying Li, Errui Ding, Jingdong Wang, Hanli Wang</p>	31213
<p>Improved Techniques for Optimization-Based Jailbreaking on Large Language Models</p> <p>Xiaojun Jia, Tianyu Pang, Chao Du, Yihao Huang, Jindong Gu, Yang Liu, Xiaochun Cao, Min Lin</p>	31225
<p>DiffSplat: Repurposing Image Diffusion Models for Scalable Gaussian Splat Generation</p> <p>Chenguo Lin, Panwang Pan, Bangbang Yang, Zeming Li, Yadong Mu</p>	31247
<p>Mitigating Parameter Interference in Model Merging Via Sharpness-Aware Fine-Tuning</p> <p>Yeoreum Lee, Jinwook Jung, Sungyong Baik</p>	31270
<p>Looking Inward: Language Models Can Learn About Themselves by Introspection</p> <p>Felix Jedidja Binder, James Chua, Tomek Korbak, Henry Sleight, John Hughes, Robert Long, Ethan Perez, Miles Turpin, Owain Evans</p>	31293
<p>Pairwise Elimination with Instance-Dependent Guarantees for Bandits with Cost Subsidy</p> <p>Ishank Juneja, Carlee Joe-Wong, Osman Yagan</p>	31340
<p>ReGen: Generative Robot Simulation Via Inverse Design</p> <p>Phat Tan Nguyen, Tsun-Hsuan Wang, Zhang-Wei Hong, Erfan Aasi, Andrew Silva, Guy Rosman, Sertac Karaman, Daniela Rus</p>	31398
<p>Emergence of Meta-Stable Clustering in Mean-Field Transformer Models</p> <p>Giuseppe Bruno, Federico Pasqualotto, Andrea Agazzi</p>	31425
<p>Lumina-T2X: Scalable Flow-Based Large Diffusion Transformer for Flexible Resolution Generation</p> <p>Peng Gao, Le Zhuo, Dongyang Liu, Ruoyi Du, Xu Luo, Longtian Qiu, Yuhang Zhang, Rongjie Huang, Shijie Geng, Renrui Zhang, Junlin Xie, Wenqi Shao, Zhengkai Jiang, Tianshuo Yang, Weicai Ye, Tong He, Jingwen He, Junjun He, Yu Qiao, Hongsheng Li</p>	31456
<p>Probe Before You Talk: Towards Black-Box Defense Against Backdoor Unalignment for Large Language Models</p> <p>Biao Yi, Tiansheng Huang, Sishuo Chen, Tong Li, Zheli Liu, Zhixuan Chu, Yiming Li</p>	31495
<p>BadJudge: Backdoor Vulnerabilities of LLM-As-A-Judge</p> <p>Terry Tong, Fei Wang, Zhe Zhao, Muhao Chen</p>	31522
<p>Ctrl-U: Robust Conditional Image Generation Via Uncertainty-Aware Reward Modeling</p> <p>Guiyu Zhang, Huan-Ang Gao, Zijian Jiang, Hao Zhao, Zhedong Zheng</p>	31548
<p>Filtered Not Mixed: Filtering-Based Online Gating for Mixture of Large Language Models</p> <p>Raeid Saqur, Anastasis Kratsios, Florian Krach, Yannick Limmer, Blanka Horvath, Frank Rudzicz</p>	31568

<p><i>The Effectiveness of Curvature-Based Rewiring and the Role of Hyperparameters in GNNs Revisited</i></p> <p><i>Floriano Tori, Vincent Holst, Vincent Ginis</i></p>	31602
<p><i>Learning Neural Networks with Distribution Shift: Efficiently Certifiable Guarantees</i></p> <p><i>Gautam Chandrasekaran, Adam Klivans, Lin Lin Lee, Konstantinos Stavropoulos</i></p>	31624
<p><i>Long-Horizon Visual Instruction Generation with Logic and Attribute Self-Reflection</i></p> <p><i>Yucheng Suo, Fan Ma, Kaixin Shen, Linchao Zhu, Yi Yang</i></p>	31658
<p><i>Selective Task Group Updates for Multi-Task Optimization</i></p> <p><i>Wooseong Jeong, Kuk-Jin Yoon</i></p>	31681
<p><i>On Linear Representations and Pretraining Data Frequency in Language Models</i></p> <p><i>Jack Merullo, Noah A. Smith, Sarah Wiegrefe, Yanai Elazar</i></p>	31715
<p><i>TweedieMix: Improving Multi-Concept Fusion for Diffusion-Based Image/Video Generation</i></p> <p><i>Gihyun Kwon, Jong Chul Ye</i></p>	31740
<p><i>Learning Clustering-Based Prototypes for Compositional Zero-Shot Learning</i></p> <p><i>Hongyu Qu, Jianan Wei, Xiangbo Shu, Wenguan Wang</i></p>	31762
<p><i>AgentTrek: Agent Trajectory Synthesis Via Guiding Replay with Web Tutorials</i></p> <p><i>Yiheng Xu, Dunjie Lu, Zhennan Shen, Junli Wang, Zekun Wang, Yuchen Mao, Caiming Xiong, Tao Yu</i></p>	31784
<p><i>Neural Exploratory Landscape Analysis for Meta-Black-Box-Optimization</i></p> <p><i>Zeyuan Ma, Jiacheng Chen, Hongshu Guo, Yue-Jiao Gong</i></p>	31806
<p><i>A Meta-Learning Approach to Bayesian Causal Discovery</i></p> <p><i>Anish Dhir, Matthew Ashman, James Requeima, Mark Van Der Wilk</i></p>	31827
<p><i>Unveiling the Secret Recipe: A Guide for Supervised Fine-Tuning Small LLMs</i></p> <p><i>Aldo Pareja, Nikhil Shivakumar Nayak, Hao Wang, Krishnateja Killamsetty, Shivchander Sudalairaj, Wenlong Zhao, Seungwook Han, Abhishek Bhandwalder, Guangxuan Xu, Kai Xu, Ligong Han, Luke Inglis, Akash Srivastava</i></p>	31848
<p><i>Swiss Army Knife: Synergizing Biases in Knowledge from Vision Foundation Models for Multi-Task Learning</i></p> <p><i>Yuxiang Lu, Shengcao Cao, Yu-Xiong Wang</i></p>	31881
<p><i>The Adaptive Complexity of Parallelized Log-Concave Sampling</i></p> <p><i>Huanjian Zhou, Baoxiang Wang, Masashi Sugiyama</i></p>	31918
<p><i>Your Mixture-Of-Experts LLM is Secretly an Embedding Model for Free</i></p> <p><i>Ziyue Li, Tianyi Zhou</i></p>	31944
<p><i>DynaPrompt: Dynamic Test-Time Prompt Tuning</i></p> <p><i>Zehao Xiao, Shilin Yan, Jack Hong, Jiayin Cai, Xiaolong Jiang, Yao Hu, Jiayi Shen, Cheems Wang, Cees G. M. Snoek</i></p>	31962

<p>A Graph Enhanced Symbolic Discovery Framework for Efficient Logic Optimization</p> <p>Yinqi Bai, Jie Wang, Lei Chen, Zhihai Wang, Yufei Kuang, Mingxuan Yuan, Jianye Hao, Feng Wu</p>	31979
<p>Diffusion Bridge Implicit Models</p> <p>Kaiwen Zheng, Guande He, Jianfei Chen, Fan Bao, Jun Zhu</p>	32011
<p>Speculative Knowledge Distillation: Bridging the Teacher-Student Gap Through Interleaved Sampling</p> <p>Wenda Xu, Rujun Han, Zifeng Wang, Long Le, Dhruv Madeka, Lei Li, William Yang Wang, Rishabh Agarwal, Chen-Yu Lee, Tomas Pfister</p>	32039
<p>3D-MoIT5: Leveraging Discrete Structural Information for Molecule-Text Modeling</p> <p>Qizhi Pei, Rui Yan, Kaiyuan Gao, Jinhua Zhu, Lijun Wu</p>	32070
<p>PIN: Prolate Spheroidal Wave Function-Based Implicit Neural Representations</p> <p>Dhananjaya Jayasundara, Heng Zhao, Demetrio Labate, Vishal M. Patel</p>	32096
<p>Knowledge Entropy Decay During Language Model Pretraining Hinders New Knowledge Acquisition</p> <p>Jiyeon Kim, Hyunji Lee, Hyowon Cho, Joel Jang, Hyeonbin Hwang, Seungpil Won, Youbin Ahn, Dohaeng Lee, Minjoon Seo</p>	32118
<p>Matrix Product Sketching Via Coordinated Sampling</p> <p>Majid Daliri, Juliana Freire, Danrong Li, Christopher Musco</p>	32142
<p>Disentangled Representation Learning with the Gromov-Monge Gap</p> <p>Théo Uscidda, Luca Eyring, Karsten Roth, Fabian J Theis, Zeynep Akata, Marco Cuturi</p>	32160
<p>BadRobot: Jailbreaking Embodied LLM Agents in the Physical World</p> <p>Hangtao Zhang, Chenyu Zhu, Xianlong Wang, Ziqi Zhou, Changgan Yin, Minghui Li, Lulu Xue, Yichen Wang, Shengshan Hu, Aishan Liu, Peijin Guo, Leo Yu Zhang</p>	32192
<p>Look Before You Leap: Universal Emergent Mechanism for Retrieval in Language Models</p> <p>Alexandre Variengien, Eric Winsor</p>	32232
<p>Arithmetic Transformers Can Length-Generalize in Both Operand Length and Count</p> <p>Hanseul Cho, Jaeyoung Cha, Srinadh Bhojanapalli, Chulhee Yun</p>	32275
<p>Language Representations Can Be What Recommenders Need: Findings and Potentials</p> <p>Leheng Sheng, An Zhang, Yi Zhang, Yuxin Chen, Xiang Wang, Tat-Seng Chua</p>	32318
<p>Persistent Pre-Training Poisoning of LLMs</p> <p>Yiming Zhang, Javier Rando, Ivan Evtimov, Jianfeng Chi, Eric Michael Smith, Nicholas Carlini, Florian Tramèr, Daphne Ippolito</p>	32345
<p>Rethinking Self-Distillation: Label Averaging and Enhanced Soft Label Refinement with Partial Labels</p> <p>Hyeonsu Jeong, Hye Won Chung</p>	32363

MIND Over Body: Adaptive Thinking Using Dynamic Computation <i>Mrinal Mathur, Barak A. Pearlmutter, Sergey M. Plis</i>	32414
TidalDecode: Fast and Accurate LLM Decoding with Position Persistent Sparse Attention <i>Lijie Yang, Zhihao Zhang, Zhuofu Chen, Zikun Li, Zhihao Jia</i>	32448
SWIFT: On-The-Fly Self-Speculative Decoding for LLM Inference Acceleration <i>Heming Xia, Yongqi Li, Jun Zhang, Cunxiao Du, Wenjie Li</i>	32470
UniDetox: Universal Detoxification of Large Language Models Via Dataset Distillation <i>Huimin Lu, Masaru Isonuma, Junichiro Mori, Ichiro Sakata</i>	32494
Value-Aligned Behavior Cloning for Offline Reinforcement Learning Via Bi-Level Optimization <i>Xingyu Jiang, Ning Gao, Xiuhui Zhang, Hongkun Dou, Yue Deng</i>	32518
Learning Spatial-Semantic Features for Robust Video Object Segmentation <i>Xin Li, Deshui Miao, Zhenyu He, Yaowei Wang, Huchuan Lu, Ming-Hsuan Yang</i>	32543
Vision CNNs Trained to Estimate Spatial Latents Learned Similar Ventral-Stream-Aligned Representations <i>Yudi Xie, Weichen Huang, Esther Alter, Jeremy Schwartz, Joshua B. Tenenbaum, James J. Dicarlo</i>	32563
ToVE: Efficient Vision-Language Learning Via Knowledge Transfer from Vision Experts <i>Yuanchen Wu, Junlong Du, Ke Yan, Shouhong Ding, Xiaoqiang Li</i>	32593
DataMan: Data Manager for Pre-Training Large Language Models <i>Ru Peng, Kexin Yang, Yawen Zeng, Junyang Lin, Dayiheng Liu, Junbo Zhao</i>	32612
Generating Physical Dynamics Under Priors <i>Zihan Zhou, Xiaoxue Wang, Tianshu Yu</i>	32665
Preserving Deep Representations in One-Shot Pruning: A Hessian-Free Second-Order Optimization Framework <i>Ryan Lucas, Rahul Mazumder</i>	32700
Segment Any 3D Object with Language <i>Seungjun Lee, Yuyang Zhao, Gim Hee Lee</i>	32725
SymmetricDiffusers: Learning Discrete Diffusion on Finite Symmetric Groups <i>Yongxing Zhang, Donglin Yang, Renjie Liao</i>	32748
Why Does the Effective Context Length of LLMs Fall Short? <i>Chenxin An, Jun Zhang, Ming Zhong, Lei Li, Shansan Gong, Yao Luo, Jingjing Xu, Lingpeng Kong</i>	32774
Continuous Ensemble Weather Forecasting with Diffusion Models <i>Martin Andrae, Tomas Landelius, Joel Oskarsson, Fredrik Lindsten</i>	32794
Matcha: Mitigating Graph Structure Shifts with Test-Time Adaptation <i>Wenxuan Bao, Zhichen Zeng, Zhining Liu, Hanghang Tong, Jingrui He</i>	32819

<p>Improved Approximation Algorithms for k-Submodular Maximization Via Multilinear Extension</p> <p>Huanjian Zhou, Lingxiao Huang, Baoxiang Wang</p>	32858
<p>Rethinking Audio-Visual Adversarial Vulnerability from Temporal and Modality Perspectives</p> <p>Zeliang Zhang, Susan Liang, Daiki Shimada, Chenliang Xu</p>	32888
<p>Efficient Multi-Agent Offline Coordination Via Diffusion-Based Trajectory Stitching</p> <p>Lei Yuan, Yuqi Bian, Lihe Li, Ziqian Zhang, Cong Guan, Yang Yu</p>	32905
<p>Agent-Oriented Planning in Multi-Agent Systems</p> <p>Ao Li, Yuexiang Xie, Songze Li, Fugee Tsung, Bolin Ding, Yaliang Li</p>	32933
<p>CAKE: Cascading and Adaptive KV Cache Eviction with Layer Preferences</p> <p>Ziran Qin, Yuchen Cao, Mingbao Lin, Wen Hu, Shixuan Fan, Ke Cheng, Weiyao Lin, Jianguo Li</p>	32956
<p>Privacy-Preserving Personalized Federated Prompt Learning for Multimodal Large Language Models</p> <p>Linh Tran, Wei Sun, Stacy Patterson, Ana Milanova</p>	32989
<p>Controllable Safety Alignment: Inference-Time Adaptation to Diverse Safety Requirements</p> <p>Jingyu Zhang, Ahmed Elgohary, Ahmed Magooda, Daniel Khashabi, Benjamin Van Durme</p>	33008
<p>AutoUAD: Hyper-Parameter Optimization for Unsupervised Anomaly Detection</p> <p>Wei Dai, Jicong Fan</p>	33040
<p>GPUdrive: Data-Driven, Multi-Agent Driving Simulation at 1 Million FPS</p> <p>Saman Kazemkhani, Aarav Pandya, Daphne Cornelisse, Brennan Shacklett, Eugene Vinitsky</p>	33069
<p>TD-Paint: Faster Diffusion Inpainting Through Time-Aware Pixel Conditioning</p> <p>Tsiry Mayet, Pourya Shamsolmoali, Simon Bernard, Eric Granger, Romain Hérault, Clement Chatelain</p>	33086
<p>Robust Weight Initialization for Tanh Neural Networks with Fixed Point Analysis</p> <p>Hyun Woo Lee, Hayoung Choi, Hyunju Kim</p>	33115
<p>On Large Language Model Continual Unlearning</p> <p>Chongyang Gao, Lixu Wang, Kaize Ding, Chenkai Weng, Xiao Wang, Qi Zhu</p>	33136
<p>Influence Functions for Scalable Data Attribution in Diffusion Models</p> <p>Bruno Kacper Mlodozieniec, Runa Eschenhagen, Juhan Bae, Alexander Immer, David Krueger, Richard E. Turner</p>	33166
<p>Strong Model Collapse</p> <p>Elvis Dohmatob, Yunzhen Feng, Arjun Subramonian, Julia Kempe</p>	33203

<p>What Secrets Do Your Manifolds Hold? Understanding the Local Geometry of Generative Models</p> <p>Ahmed Imtiaz Humayun, Ibtihel Amara, Cristina Nader Vasconcelos, Deepak Ramachandran, Candice Schumann, Junfeng He, Katherine A Heller, Golnoosh Farnadi, Negar Rostamzadeh, Mohammad Havaei</p>	33239
<p>GraphRouter: A Graph-Based Router for LLM Selections</p> <p>Tao Feng, Yanzhen Shen, Jiaxuan You</p>	33265
<p>On Minimizing Adversarial Counterfactual Error in Adversarial Reinforcement Learning</p> <p>Roman Belaire, Arunesh Sinha, Pradeep Varakantham</p>	33283
<p>R2Det: Exploring Relaxed Rotation Equivariance in 2D Object Detection</p> <p>Zhiqiang Wu, Yingjie Liu, Hanlin Dong, Xuan Tang, Jian Yang, Bo Jin, Mingsong Chen, Xian Wei</p>	33308
<p>How Much of My Dataset Did You Use? Quantitative Data Usage Inference in Machine Learning</p> <p>Yao Tong, Jiayuan Ye, Sajjad Zarifzadeh, Reza Shokri</p>	33329
<p>EMMA: Empowering Multi-Modal Mamba with Structural and Hierarchical Alignment</p> <p>Yifei Xing, Xiangyuan Lan, Ruiping Wang, Dongmei Jiang, Wenjun Huang, Zheng Qingfang, Yaowei Wang</p>	33369
<p>Reveal Object in Lensless Photography Via Region Gaze and Amplification</p> <p>Yin Xiangjun, Huihui Yue</p>	33398
<p>Local-Prompt: Extensible Local Prompts for Few-Shot Out-Of-Distribution Detection</p> <p>Fanhu Zeng, Zhen Cheng, Fei Zhu, Hongxin Wei, Xu-Yao Zhang</p>	33416
<p>CBQ: Cross-Block Quantization for Large Language Models</p> <p>Xin Ding, Xiaoyu Liu, Zhijun Tu, Yun Zhang, Wei Li, Jie Hu, Hanting Chen, Yehui Tang, Zhiwei Xiong, Baoqun Yin, Yunhe Wang</p>	33434
<p>Offline RL in Regular Decision Processes: Sample Efficiency Via Language Metrics</p> <p>Ahana Deb, Roberto Cipollone, Anders Jonsson, Alessandro Ronca, Mohammad Sadegh Talebi</p>	33454
<p>Adapt-\$\infty\$: Scalable Continual Multimodal Instruction Tuning Via Dynamic Data Selection</p> <p>Adyasha Maharana, Jaehong Yoon, Tianlong Chen, Mohit Bansal</p>	33479
<p>Layerwise Recurrent Router for Mixture-Of-Experts</p> <p>Zihan Qiu, Zeyu Huang, Shuang Cheng, Yizhi Zhou, Zili Wang, Ivan Titov, Jie Fu</p>	33500
<p>Multi-Resolution Decomposable Diffusion Model for Non-Stationary Time Series Anomaly Detection</p> <p>Guojin Zhong, Pan Wang, Jin Yuan, Zhiyong Li, Long Chen</p>	33527
<p>How Does Vision-Language Adaptation Impact the Safety of Vision Language Models?</p> <p>Seongyun Lee, Geewook Kim, Jiyeon Kim, Hyunji Lee, Hoyeon Chang, Sue Hyun Park, Minjoon Seo</p>	33554

<p>Parameter Expanded Stochastic Gradient Markov Chain Monte Carlo</p> <p><i>Hyunsu Kim, Giung Nam, Chulhee Yun, Hongseok Yang, Juho Lee</i></p>	33593
<p>INFER: A Neural-Symbolic Model for Extrapolation Reasoning on Temporal Knowledge Graph</p> <p><i>Ningyuan Li, Haihong E, Tianyu Yao, Tianyi Hu, Yuhan Li, Haoran Luo, Meina Song, Yifan Zhu</i></p>	33622
<p>Poison-Splat: Computation Cost Attack on 3D Gaussian Splatting</p> <p><i>Jiahao Lu, Yifan Zhang, QiuHong Shen, Xinchao Wang, Shuicheng Yan</i></p>	33638
<p>MaskGCT: Zero-Shot Text-To-Speech with Masked Generative Codec Transformer</p> <p><i>Yuancheng Wang, Haoyue Zhan, Liwei Liu, Ruihong Zeng, Haotian Guo, Jiachen Zheng, Qiang Zhang, Xueyao Zhang, Shunsi Zhang, Zhizheng Wu</i></p>	33670
<p>Offline RL with Smooth OOD Generalization in Convex Hull and Its Neighborhood</p> <p><i>Qingmao Yao, Zhichao Lei, Tianyuan Chen, Ziyue Yuan, Xuefan Chen, Jianxiang Liu, Faguo Wu, Xiao Zhang</i></p>	33694
<p>EMOS: Embodiment-Aware Heterogeneous Multi-Robot Operating System with LLM Agents</p> <p><i>Junting Chen, Checheng Yu, Xunzhe Zhou, Tianqi Xu, Yao Mu, Mengkang Hu, Wenqi Shao, Yikai Wang, Guohao Li, Lin Shao</i></p>	33721
<p>The Ramanujan Library - Automated Discovery on the Hypergraph of Integer Relations</p> <p><i>Itay Beit Halachmi, Ido Kaminer</i></p>	33742
<p>Retrieval Head Mechanistically Explains Long-Context Factuality</p> <p><i>Wenhao Wu, Yizhong Wang, Guangxuan Xiao, Hao Peng, Yao Fu</i></p>	33762
<p>Topological Blindspots: Understanding and Extending Topological Deep Learning Through the Lens of Expressivity</p> <p><i>Yam Eitan, Yoav Gelberg, Guy Bar-Shalom, Fabrizio Frasca, Michael M. Bronstein, Haggai Maron</i></p>	33776
<p>The Rise and Down of Babel Tower: Investigating the Evolution Process of Multilingual Code Large Language Model</p> <p><i>Jiawei Chen, Wentao Chen, Jing Su, Jingjing Xu, Hongyu Lin, Mengjie Ren, Yaojie Lu, Xianpei Han, Le Sun</i></p>	33816
<p>BEEM: Boosting Performance of Early Exit DNNs Using Multi-Exit Classifiers as Experts</p> <p><i>Divya Jyoti Bajpai, Manjesh Kumar Hanawal</i></p>	33832
<p>Bisimulation Metric for Model Predictive Control</p> <p><i>Yutaka Shimizu, Masayoshi Tomizuka</i></p>	33848
<p>Action Sequence Augmentation for Action Anticipation</p> <p><i>Yihui Qiu, Deepu Rajan</i></p>	33870
<p>Lasso Bandit with Compatibility Condition on Optimal Arm</p> <p><i>Harin Lee, Taehyun Hwang, Min-Hwan Oh</i></p>	33888

<p>Safety-Prioritizing Curricula for Constrained Reinforcement Learning</p> <p><i>Cevahir Koprulu, Thiago D. Simão, Nils Jansen, Ufuk Topcu</i></p>	33946
<p>Combating Dimensional Collapse in LLM Pre-Training Data Via Submodular File Selection</p> <p><i>Ziqing Fan, Siyuan Du, Shengchao Hu, Pingjie Wang, Li Shen, Ya Zhang, Dacheng Tao, Yanfeng Wang</i></p>	33975
<p>Regulatory DNA Sequence Design with Reinforcement Learning</p> <p><i>Zhao Yang, Bing Su, Chuan Cao, Ji-Rong Wen</i></p>	34000
<p>Differentially Private Optimization for Non-Decomposable Objective Functions</p> <p><i>Weiwei Kong, Andres Munoz Medina, Mónica Ribero</i></p>	34026
<p>Efficient and Robust Neural Combinatorial Optimization Via Wasserstein-Based Coresets</p> <p><i>Xu Wang, Fuyou Miao, Wenjie Liu, Yan Xiong</i></p>	34045
<p>LOIRE: LifeLong Learning on Incremental Data Via Pre-Trained Language Model gRowth Efficiently</p> <p><i>Xue Han, Yitong Wang, Junlan Feng, Wenchun. Gao, Qian Hu, Chao Deng</i></p>	34080
<p>DataGen: Unified Synthetic Dataset Generation Via Large Language Models</p> <p><i>Yue Huang, Siyuan Wu, Chujie Gao, Dongping Chen, Qihui Zhang, Yao Wan, Tianyi Zhou, Chaowei Xiao, Jianfeng Gao, Lichao Sun, Xiangliang Zhang</i></p>	34101
<p>Shh, Don't Say That! Domain Certification in LLMs</p> <p><i>Cornelius Emde, Alasdair Paren, Preetham Arvind, Maxime Guillaume Kayser, Tom Rainforth, Thomas Lukasiewicz, Philip Torr, Adel Bibi</i></p>	34136
<p>Federated Continual Learning Goes Online: Uncertainty-Aware Memory Management for Vision Tasks and Beyond</p> <p><i>Giuseppe Serra, Florian Buettner</i></p>	34173
<p>Capability Localization: Capabilities Can Be Localized Rather than Individual Knowledge</p> <p><i>Xiusheng Huang, Jiaxiang Liu, Yequan Wang, Jun Zhao, Kang Liu</i></p>	34194
<p>The Crucial Role of Samplers in Online Direct Preference Optimization</p> <p><i>Ruizhe Shi, Runlong Zhou, Simon Shaolei Du</i></p>	34211
<p>Weak-To-Strong Preference Optimization: Stealing Reward from Weak Aligned Model</p> <p><i>Wenhong Zhu, Zhiwei He, Xiaofeng Wang, Pengfei Liu, Rui Wang</i></p>	34242
<p>API Pack: A Massive Multi-Programming Language Dataset for API Call Generation</p> <p><i>Zhen Guo, Adriana Meza Soria, Wei Sun, Yikang Shen, Rameswar Panda</i></p>	34271
<p>MIA-DPO: Multi-Image Augmented Direct Preference Optimization for Large Vision-Language Models</p> <p><i>Ziyu Liu, Yuhang Zang, Xiaoyi Dong, Pan Zhang, Yuhang Cao, Haodong Duan, Conghui He, Yuanjun Xiong, Dahua Lin, Jiaqi Wang</i></p>	34299

<p>Controllable Satellite-To-Street-View Synthesis with Precise Pose Alignment and Zero-Shot Environmental Control</p> <p>Xianghui Ze, Zhenbo Song, Qiwei Wang, Jianfeng Lu, Yujiao Shi</p>	34327
<p>Amulet: ReAlignment During Test Time for Personalized Preference Adaptation of LLMs</p> <p>Zhaowei Zhang, Fengshuo Bai, Qizhi Chen, Chengdong Ma, Mingzhi Wang, Haoran Sun, Zilong Zheng, Yaodong Yang</p>	34349
<p>DreamCatalyst: Fast and High-Quality 3D Editing Via Controlling Editability and Identity Preservation</p> <p>Jiwook Kim, Seonho Lee, Jaeyo Shin, Jiho Choi, Hyunjung Shim</p>	34380
<p>BIRD: A Trustworthy Bayesian Inference Framework for Large Language Models</p> <p>Yu Feng, Ben Zhou, Weidong Lin, Dan Roth</p>	34407
<p>Statistical Advantages of Perturbing Cosine Router in Mixture of Experts</p> <p>Huy Nguyen, Pedram Akbarian, Huyen Trang Pham, Thien Trang Nguyen Vu, Shujian Zhang, Nhat Ho</p>	34436
<p>STAFF: Speculative Coreset Selection for Task-Specific Fine-Tuning</p> <p>Xiaoyu Zhang, Juan Zhai, Shiqing Ma, Chao Shen, Tianlin Li, Weipeng Jiang, Yang Liu</p>	34469
<p>Can We Ignore Labels in Out of Distribution Detection?</p> <p>Hong Yang, Qi Yu, Travis Desell</p>	34495
<p>Efficient Sparse PCA Via Block-Diagonalization</p> <p>Alberto Del Pia, Dekun Zhou, Yinglun Zhu</p>	34520
<p>TAU-106K: A New Dataset for Comprehensive Understanding of Traffic Accident</p> <p>Yixuan Zhou, Long Bai, Sijia Cai, Bing Deng, Xing Xu, Heng Tao Shen</p>	34548
<p>Fast Unsupervised Ground Metric Learning with tree-Wasserstein Distance</p> <p>Kira Michaela Dusterwald, Samo Hromadka, Makoto Yamada</p>	34574
<p>Turning Up the Heat: Min-P Sampling for Creative and Coherent LLM Outputs</p> <p>Nguyen Nhat Minh, Andrew Baker, Clement Neo, Allen G Roush, Andreas Kirsch, Ravid Shwartz-Ziv</p>	34593
<p>Optimality of Matrix Mechanism on ℓ_p-Metric</p> <p>Zongrui Zou, Jingcheng Liu, Jalaj Upadhyay</p>	34627
<p>Event-Driven Online Vertical Federated Learning</p> <p>Ganyu Wang, Boyu Wang, Bin Gu, Charles Ling</p>	34654
<p>TOMATO: Assessing Visual Temporal Reasoning Capabilities in Multimodal Foundation Models</p> <p>Ziyao Shangguan, Chuhan Li, Yuxuan Ding, Yanan Zheng, Yilun Zhao, Tesca Fitzgerald, Arman Cohan</p>	34685
<p>On Stochastic Contextual Bandits with Knapsacks in Small Budget Regime</p> <p>Hengquan Guo, Xin Liu</p>	34827

<p>Enhancing End-To-End Autonomous Driving with Latent World Model</p> <p>Yingyan Li, Lue Fan, Jiawei He, Yuqi Wang, Yuntao Chen, Zhaoxiang Zhang, Tieniu Tan</p>	34849
<p>SV-RAG: LoRA-Contextualizing Adaptation of MLLMs for Long Document Understanding</p> <p>Jian Chen, Ruiyi Zhang, Yufan Zhou, Tong Yu, Franck Dernoncourt, Jiuxiang Gu, Ryan A. Rossi, Changyou Chen, Tong Sun</p>	34867
<p>Efficient Discovery of Pareto Front for Multi-Objective Reinforcement Learning</p> <p>Ruohong Liu, Yuxin Pan, Linjie Xu, Lei Song, Pengcheng You, Yize Chen, Jiang Bian</p>	34883
<p>AgentRefine: Enhancing Agent Generalization Through Refinement Tuning</p> <p>Dayuan Fu, Keqing He, Yejie Wang, Wentao Hong, Zhuoma Gongque, Weihao Zeng, Wei Wang, Jingang Wang, Xunliang Cai, Weiran Xu</p>	34910
<p>Determine-Then-Ensemble: Necessity of Top-K Union for Large Language Model Ensembling</p> <p>Yuxuan Yao, Han Wu, Mingyang Liu, Sichun Luo, Xiongwei Han, Jie Liu, Zhijiang Guo, Linqi Song</p>	34930
<p>Diffusion²: Dynamic 3D Content Generation Via Score Composition of Video and Multi-View Diffusion Models</p> <p>Zeyu Yang, Zijie Pan, Chun Gu, Li Zhang</p>	34949
<p>Empowering Users in Digital Privacy Management Through Interactive LLM-Based Agents</p> <p>Bolun Sun, Yifan Zhou, Haiyun Jiang</p>	34971
<p>MAESTRO: Masked Encoding Set Transformer with Self-Distillation</p> <p>Matthew Eric Lee, Jaesik Kim, Matei Ionita, Jonghyun Lee, Michelle L. McKeague, Yonghyun Nam, Irene Khavin, Yidi Huang, Victoria Fang, Sokratis Apostolidis, Divij Mathew, Shwetank, Ajinkya Pattekar, Zahabia Rangwala, Amit Bar-Or, Benjamin A Fensterheim, Benjamin A. Abramoff, Rennie L. Rhee, Damian Maseda, Allison R Greenplate, John Wherry, Dokyoon Kim</p>	34992
<p>Teaching Human Behavior Improves Content Understanding Abilities of VLMs</p> <p>Somesh Kumar Singh, Harini S I, Yaman Kumar Singla, Changyou Chen, Rajiv Ratn Shah, Veeky Baths, Balaji Krishnamurthy</p>	35019
<p>CogCoM: A Visual Language Model with Chain-Of-Manipulations Reasoning</p> <p>Ji Qi, Ming Ding, Weihang Wang, Yushi Bai, Qingsong Lv, Wenyi Hong, Bin Xu, Lei Hou, Juanzi Li, Yuxiao Dong, Jie Tang</p>	35051
<p>Generalization Guarantees for Representation Learning Via Data-Dependent Gaussian Mixture Priors</p> <p>Milad Sefidgaran, Abdellatif Zaidi, Piotr Krasnowski</p>	35072
<p>Recovery of Causal Graph Involving Latent Variables Via Homologous Surrogates</p> <p>Xiu-Chuan Li, Jun Wang, Tongliang Liu</p>	35107
<p>OpenPRM: Building Open-Domain Process-Based Reward Models with Preference Trees</p> <p>Kaiyan Zhang, Jiayuan Zhang, Haoxin Li, Xuekai Zhu, Ermo Hua, Xingtai Lv, Ning Ding, Biqing Qi, Bowen Zhou</p>	35131

<p>SPARTUN3D: Situated Spatial Understanding of 3D World in Large Language Model</p> <p>Yue Zhang, Zhiyang Xu, Ying Shen, Parisa Kordjamshidi, Lifu Huang</p>	35154
<p>MMRole: A Comprehensive Framework for Developing and Evaluating Multimodal Role-Playing Agents</p> <p>Yanqi Dai, Huanran Hu, Lei Wang, Shengjie Jin, Xu Chen, Zhiwu Lu</p>	35173
<p>Linear Spherical Sliced Optimal Transport: A Fast Metric for Comparing Spherical Data</p> <p>Xinran Liu, Yikun Bai, Rocio Diaz Martin, Kaiwen Shi, Ashkan Shahbazi, Bennett Allan Landman, Catie Chang, Soheil Kolouri</p>	35195
<p>ϕ-Update: A Class of Policy Update Methods with Policy Convergence Guarantee</p> <p>Wenye Li, Jiakai Liu, Ke Wei</p>	35238
<p>Learning Equivariant Non-Local Electron Density Functionals</p> <p>Nicholas Gao, Eike Eberhard, Stephan Günnemann</p>	35272
<p>Asynchronous RLHF: Faster and More Efficient Off-Policy RL for Language Models</p> <p>Michael Noukhovitch, Shengyi Huang, Sophie Xhonneux, Arian Hosseini, Rishabh Agarwal, Aaron Courville</p>	35297
<p>Scalable Bayesian Learning with Posteriors</p> <p>Samuel Duffield, Kaelan Donatella, Johnathan Chiu, Phoebe Klett, Daniel Simpson</p>	35324
<p>Proxy Denoising for Source-Free Domain Adaptation</p> <p>Song Tang, Wenxin Su, Yan Gan, Mao Ye, Jianwei Dr. Zhang, Xiatian Zhu</p>	35348
<p>Proving Olympiad Inequalities by Synergizing LLMs and Symbolic Reasoning</p> <p>Zenan Li, Zhaoyu Li, Wen Tang, Xian Zhang, Yuan Yao, Xujie Si, Fan Yang, Kaiyu Yang, Xiaoxing Ma</p>	35370
<p>CS-Bench: A Comprehensive Benchmark for Large Language Models Towards Computer Science Mastery</p> <p>Xiaoshuai Song, Muxi Diao, Guanting Dong, Zhengyang Wang, Yujia Fu, Runqi Qiao, Zhexu Wang, Dayuan Fu, Huangxuan Wu, Bin Liang, Weihao Zeng, Yejie Wang, Zhuoma Gongque, Jianing Yu, Qiuna Tan, Weiran Xu</p>	35396
<p>Not All Heads Matter: A Head-Level KV Cache Compression Method with Integrated Retrieval and Reasoning</p> <p>Yu Fu, Zefan Cai, Abedelkadir Asi, Wayne Xiong, Yue Dong, Wen Xiao</p>	35441
<p>Exploring the Forgetting in Adversarial Training: A Novel Method for Enhancing Robustness</p> <p>Xianglu Wang, Hu Ding</p>	35463
<p>Do Large Language Models Truly Understand Geometric Structures?</p> <p>Xiaofeng Wang, Yiming Wang, Wenhong Zhu, Rui Wang</p>	35487
<p>Provable Convergence Bounds for Hybrid Dynamical Sampling and Optimization</p> <p>Matthew X. Burns, Qingyuan Hou, Michael Huang</p>	35520

<p>Towards Scalable Topological Regularizers</p> <p>Wong Hiu Tung, Darrick Lee, Hong Yan</p>	35553
<p>In Search of Forgotten Domain Generalization</p> <p>Prasanna Mayilvahanan, Roland S. Zimmermann, Thaddäus Wiedemer, Evgenia Rusak, Attila Juhas, Matthias Bethge, Wieland Brendel</p>	35584
<p>What is Wrong with Perplexity for Long-Context Language Modeling?</p> <p>Lizhe Fang, Yifei Wang, Zhaoyang Liu, Chenheng Zhang, Stefanie Jegelka, Jinyang Gao, Bolin Ding, Yisen Wang</p>	35628
<p>LongMamba: Enhancing Mamba's Long-Context Capabilities Via Training-Free Receptive Field Enlargement</p> <p>Zhifan Ye, Kejing Xia, Yonggan Fu, Xin Dong, Jihoon Hong, Xiangchi Yuan, Shizhe Diao, Jan Kautz, Pavlo Molchanov, Yingyan Celine Lin</p>	35651
<p>Learned Reference-Based Diffusion Sampler for Multi-Modal Distributions</p> <p>Maxence Noble, Louis Grenioux, Marylou Gabri�, Alain Oliviero Durmus</p>	35670
<p>Unlearning Or Obfuscating? Jogging the Memory of Unlearned LLMs Via Benign Relearning</p> <p>Shengyuan Hu, Yiwei Fu, Steven Wu, Virginia Smith</p>	35725
<p>Tight Lower Bounds Under Asymmetric High-Order H�lder Smoothness and Uniform Convexity</p> <p>Site Bai, Brian Bullins</p>	35757
<p>Efficient Imitation Under Misspecification</p> <p>Nicolas Espinosa-Dice, Sanjiban Choudhury, Wen Sun, Gokul Swamy</p>	35784
<p>Exploring Learning Complexity for Efficient Downstream Dataset Pruning</p> <p>Wenyu Jiang, Zhenlong Liu, Zejian Xie, Songxin Zhang, Bingyi Jing, Hongxin Wei</p>	35823
<p>When LLMs Play the Telephone Game: Cultural Attractors as Conceptual Tools to Evaluate LLMs in Multi-Turn Settings</p> <p>J�r�my Perez, Grgur Kovac, Corentin L�ger, C�dric Colas, Gaia Molinaro, Maxime Derex, Pierre-Yves Oudeyer, Cl�ment Moulin-Frier</p>	35844
<p>Looking Backward: Retrospective Backward Synthesis for Goal-Conditioned GFlowNets</p> <p>Haoran He, Can Chang, Huazhe Xu, Ling Pan</p>	35884
<p>GIFT: Unlocking Full Potential of Labels in Distilled Dataset at Near-Zero Cost</p> <p>Xinyi Shang, Peng Sun, Tao Lin</p>	35906
<p>Causal Graph Transformer for Treatment Effect Estimation Under Unknown Interference</p> <p>Anpeng Wu, Haiyi Qiu, Zhengming Chen, Zijian Li, Ruoxuan Xiong, Fei Wu, Kun Zhang</p>	35931
<p>Federated Q\$-Learning with Reference-Advantage Decomposition: Almost Optimal Regret and Logarithmic Communication Cost</p> <p>Zhong Zheng, Haochen Zhang, Lingzhou Xue</p>	35961

<p>UGMathBench: A Diverse and Dynamic Benchmark for Undergraduate-Level Mathematical Reasoning with Large Language Models</p> <p>Xin Xu, Jiaxin Zhang, Tianhao Chen, Zitong Chao, Jishan Hu, Can Yang</p>	36017
<p>BALROG: Benchmarking Agentic LLM and VLM Reasoning on Games</p> <p>Davide Paglieri, Bartłomiej Cupiał, Samuel Coward, Ulyana Piterbarg, Maciej Wolczyk, Akbir Khan, Eduardo Pignatelli, Lukasz Kucinski, Lerrel Pinto, Rob Fergus, Jakob Nicolaus Foerster, Jack Parker-Holder, Tim Rocktäschel</p>	36061
<p>When Does Compositional Structure Yield Compositional Generalization? a Kernel Theory.</p> <p>Samuel Lippl, Kim Stachenfeld</p>	36098
<p>Differentiable Integer Linear Programming</p> <p>Zijie Geng, Jie Wang, Xijun Li, Fangzhou Zhu, Jianye Hao, Bin Li, Feng Wu</p>	36138
<p>More RLHF, More Trust? on the Impact of Preference Alignment on Trustworthiness</p> <p>Aaron Jiaxun Li, Satyapriya Krishna, Himabindu Lakkaraju</p>	36165
<p>Spherical Tree-Sliced Wasserstein Distance</p> <p>Hoang V. Tran, Thanh Chu, Minh-Khoi Nguyen-Nhat, Huyen Trang Pham, Tam Le, Tan Minh Nguyen</p>	36186
<p>Simple ReFlow: Improved Techniques for Fast Flow Models</p> <p>Beomsu Kim, Yu-Guan Hsieh, Michal Klein, Marco Cuturi, Jong Chul Ye, Bahjat Kawar, James Thornton</p>	36219
<p>Weighted-Reward Preference Optimization for Implicit Model Fusion</p> <p>Ziyi Yang, Fanqi Wan, Longguang Zhong, Tianyuan Shi, Xiaojun Quan</p>	36246
<p>Wavelet Diffusion Neural Operator</p> <p>Peiyan Hu, Rui Wang, Xiang Zheng, Tao Zhang, Haodong Feng, Ruiqi Feng, Long Wei, Yue Wang, Zhi-Ming Ma, Tailin Wu</p>	36265
<p>The Semantic Hub Hypothesis: Language Models Share Semantic Representations Across Languages and Modalities</p> <p>Zhaofeng Wu, Xinyan Velocity Yu, Dani Yogatama, Jiasen Lu, Yoon Kim</p>	36308
<p>Vision Language Models Are In-Context Value Learners</p> <p>Yecheng Jason Ma, Joey Hejna, Chuyuan Fu, Dhruv Shah, Jacky Liang, Zhuo Xu, Sean Kirmani, Peng Xu, Danny Driess, Ted Xiao, Osbert Bastani, Dinesh Jayaraman, Wenhao Yu, Tingnan Zhang, Dorsa Sadigh, Fei Xia</p>	36327
<p>Node Similarities Under Random Projections: Limits and Pathological Cases</p> <p>Tvrtko Tadic, Cassiano O Becker, Jennifer Neville</p>	36353
<p>Block Verification Accelerates Speculative Decoding</p> <p>Ziteng Sun, Uri Mendlovic, Yaniv Leviathan, Asaf Aharoni, Jae Hun Ro, Ahmad Beirami, Ananda Theertha Suresh</p>	36393

<p>Expand and Compress: Exploring Tuning Principles for Continual Spatio-Temporal Graph Forecasting</p> <p>Wei Chen, Yuxuan Liang</p>	36423
<p>Teaching LLMs How to Learn with Contextual Fine-Tuning</p> <p>Younwoo Choi, Muhammad Adil Asif, Ziwen Han, John Willes, Rahul Krishnan</p>	36449
<p>FIG: Flow with Interpolant Guidance for Linear Inverse Problems</p> <p>Yici Yan, Yichi Zhang, Xiangming Meng, Zhizhen Zhao</p>	36478
<p>EmbedLLM: Learning Compact Representations of Large Language Models</p> <p>Richard Zhuang, Tianhao Wu, Zhaojin Wen, Andrew Li, Jiantao Jiao, Kannan Ramchandran</p>	36503
<p>OptiBench Meets ReSocratic: Measure and Improve LLMs for Optimization Modeling</p> <p>Zhicheng Yang, Yiwei Wang, Yinya Huang, Zhijiang Guo, Wei Shi, Xiongwei Han, Liang Feng, Linqi Song, Xiaodan Liang, Jing Tang</p>	36517
<p>Inference Scaling for Long-Context Retrieval Augmented Generation</p> <p>Zhenrui Yue, Honglei Zhuang, Aijun Bai, Kai Hui, Rolf Jagerman, Hansi Zeng, Zhen Qin, Dong Wang, Xuanhui Wang, Michael Bendersky</p>	36551
<p>ADePT: Adaptive Decomposed Prompt Tuning for Parameter-Efficient Fine-Tuning</p> <p>Pengwei Tang, Xiaolin Hu, Yong Liu</p>	36576
<p>As Simple as Fine-Tuning: LLM Alignment Via Bidirectional Negative Feedback Loss</p> <p>Xin Mao, Huimin Xu, Feng-Lin Li, Ziqi Jin, Wang Chen, Wei Zhang, Anh Tuan Luu</p>	36599
<p>Robust-PIFu: Robust Pixel-Aligned Implicit Function for 3D Human Digitalization from a Single Image</p> <p>Kennard Chan, Fayao Liu, Guosheng Lin, Chuan-Sheng Foo, Weisi Lin</p>	36621
<p>Gramian Multimodal Representation Learning and Alignment</p> <p>Giordano Cicchetti, Eleonora Grassucci, Luigi Sigillo, Danilo Comminiello</p>	36648
<p>Fantastic Copyrighted Beasts and How (Not) to Generate Them</p> <p>Luxi He, Yangsibo Huang, Weijia Shi, Tinghao Xie, Haotian Liu, Yue Wang, Luke Zettlemoyer, Chiyuan Zhang, Danqi Chen, Peter Henderson</p>	36670
<p>Diffusion Generative Modeling for Spatially Resolved Gene Expression Inference from Histology Images</p> <p>Sichen Zhu, Yuchen Zhu, Molei Tao, Peng Qiu</p>	36698
<p>Realistic Evaluation of Deep Partial-Label Learning Algorithms</p> <p>Wei Wang, Dong-Dong Wu, Jindong Wang, Gang Niu, Min-Ling Zhang, Masashi Sugiyama</p>	36723
<p>DELIFT: Data Efficient Language Model Instruction Fine-Tuning</p> <p>Ishika Agarwal, Krishnateja Killamsetty, Lucian Popa, Marina Danilevsky</p>	36748

<p>Graph-Based Document Structure Analysis</p> <p>Yufan Chen, Ruiping Liu, Junwei Zheng, Di Wen, Kunyu Peng, Jiaming Zhang, Rainer Stiefelhagen</p>	36770
<p>PhyMPGN: Physics-Encoded Message Passing Graph Network for Spatiotemporal PDE Systems</p> <p>Bocheng Zeng, Qi Wang, Mengtao Yan, Yang Liu, Ruizhi Chengze, Yi Zhang, Hongsheng Liu, Zidong Wang, Hao Sun</p>	36794
<p>Identifiability for Gaussian Processes with Holomorphic Kernels</p> <p>Ameer Qaqish, Didong Li</p>	36822
<p>Transformer Learns Optimal Variable Selection in Group-Sparse Classification</p> <p>Chenyang Zhang, Xuran Meng, Yuan Cao</p>	36851
<p>GridMix: Exploring Spatial Modulation for Neural Fields in PDE Modeling</p> <p>Honghui Wang, Shiji Song, Gao Huang</p>	36904
<p>Improving Probabilistic Diffusion Models with Optimal Diagonal Covariance Matching</p> <p>Zijing Ou, Mingtian Zhang, Andi Zhang, Tim Z. Xiao, Yingzhen Li, David Barber</p>	36922
<p>A Unified Theory of Quantum Neural Network Loss Landscapes</p> <p>Eric Ricardo Anschuetz</p>	36948
<p>CL-MFAP: A Contrastive Learning-Based Multimodal Foundation Model for Molecular Property Prediction and Antibiotic Screening</p> <p>Gen Zhou, Sugitha Janarthanan, Yutong Lu, Pingzhao Hu</p>	37008
<p>Sitcom-Crafter: A Plot-Driven Human Motion Generation System in 3D Scenes</p> <p>Jianqi Chen, Panwen Hu, Xiaojun Chang, Zhenwei Shi, Michael Kampffmeyer, Xiaodan Liang</p>	37029
<p>Beyond-Expert Performance with Limited Demonstrations: Efficient Imitation Learning with Double Exploration</p> <p>Heyang Zhao, Xingrui Yu, David Mark Bossens, Ivor Tsang, Quanquan Gu</p>	37063
<p>Eliminating Position Bias of Language Models: A Mechanistic Approach</p> <p>Ziqi Wang, Hanlin Zhang, Xiner Li, Kuan-Hao Huang, Chi Han, Shuiwang Ji, Sham M. Kakade, Hao Peng, Heng Ji</p>	37090
<p>Langevin Soft Actor-Critic: Efficient Exploration Through Uncertainty-Driven Critic Learning</p> <p>Haque Ishfaq, Guangyuan Wang, Sami Nur Islam, Doina Precup</p>	37118
<p>Population Transformer: Learning Population-Level Representations of Neural Activity</p> <p>Geeling Chau, Christopher Wang, Sabera J Talukder, Vighnesh Subramaniam, Saraswati Soedarmadji, Yisong Yue, Boris Katz, Andrei Barbu</p>	37145
<p>GReaTer: Gradients Over Reasoning Makes Smaller Language Models Strong Prompt Optimizers</p> <p>Sarkar Snigdha Sarathi Das, Ryo Kamoi, Bo Pang, Yusen Zhang, Caiming Xiong, Rui Zhang</p>	37167

<p>Enabling Realtime Reinforcement Learning at Scale with Staggered Asynchronous Inference</p> <p>Matthew Riemer, Gopeshh Subbaraj, Glen Berseth, Irina Rish</p>	37199
<p>On Evaluating the Durability of Safeguards for Open-Weight LLMs</p> <p>Xiangyu Qi, Boyi Wei, Nicholas Carlini, Yangsibo Huang, Tinghao Xie, Luxi He, Matthew Jagielski, Milad Nasr, Prateek Mittal, Peter Henderson</p>	37225
<p>Tuning Timestep-Distilled Diffusion Model Using Pairwise Sample Optimization</p> <p>Zichen Miao, Zhengyuan Yang, Kevin Lin, Ze Wang, Zicheng Liu, Lijuan Wang, Qiang Qiu</p>	37256
<p>Physics of Language Models: Part 3.3, Knowledge Capacity Scaling Laws</p> <p>Zeyuan Allen-Zhu, Yuanzhi Li</p>	37278
<p>DRoP: Distributionally Robust Data Pruning</p> <p>Artem M Vysogorets, Kartik Ahuja, Julia Kempe</p>	37288
<p>Discrete Copula Diffusion</p> <p>Anji Liu, Oliver Broadrick, Mathias Niepert, Guy Van Den Broeck</p>	37313
<p>Sylber: Syllabic Embedding Representation of Speech from Raw Audio</p> <p>Cheol Jun Cho, Nicholas Lee, Akshat Gupta, Dhruv Agarwal, Ethan Chen, Alan Black, Gopala Anumanchipalli</p>	37340
<p>Conformal Prediction Sets Can Cause Disparate Impact</p> <p>Jesse C. Cresswell, Bhargava Kumar, Yi Sui, Mouloud Belbahri</p>	37367
<p>Shot2Story: A New Benchmark for Comprehensive Understanding of Multi-Shot Videos</p> <p>Mingfei Han, Linjie Yang, Xiaojun Chang, Lina Yao, Heng Wang</p>	37387
<p>JudgeBench: A Benchmark for Evaluating LLM-Based Judges</p> <p>Sijun Tan, Siyuan Zhuang, Kyle Montgomery, William Yuan Tang, Alejandro Cuadron, Chenguang Wang, Raluca Popa, Ion Stoica</p>	37400
<p>ADBMs: Adversarial Diffusion Bridge Model for Reliable Adversarial Purification</p> <p>Xiao Li, Wenxuan Sun, Huanran Chen, Qiongxiu Li, Yingzhe He, Jie Shi, Xiaolin Hu</p>	37427
<p>InstaRevive: One-Step Image Enhancement Via Dynamic Score Matching</p> <p>Yixuan Zhu, Haolin Wang, Ao Li, Wenliang Zhao, Yansong Tang, Jingxuan Niu, Lei Chen, Jie Zhou, Jiwen Lu</p>	37450
<p>Black Sheep in the Herd: Playing with Spuriously Correlated Attributes for Vision-Language Recognition</p> <p>Xinyu Tian, Shu Zou, Zhaoyuan Yang, Mengqi He, Jing Zhang</p>	37474
<p>TeaserGen: Generating Teasers for Long Documentaries</p> <p>Weihan Xu, Paul Pu Liang, Haven Kim, Julian McAuley, Taylor Berg-Kirkpatrick, Hao-Wen Dong</p>	37508

<p>Synergy Between Sufficient Changes and Sparse Mixing Procedure for Disentangled Representation Learning</p> <p>Zijian Li, Shunxing Fan, Yujia Zheng, Ignavier Ng, Shaoan Xie, Guangyi Chen, Xinshuai Dong, Ruichu Cai, Kun Zhang</p>	37529
<p>Fine-Tuning Discrete Diffusion Models Via Reward Optimization with Applications to DNA and Protein Design</p> <p>Chenyu Wang, Masatoshi Uehara, Yichun He, Amy Wang, Avantika Lal, Tommi Jaakkola, Sergey Levine, Aviv Regev, Hanchen, Tommaso Biancalani</p>	37554
<p>Natural Language Inference Improves Compositionality in Vision-Language Models</p> <p>Paola Cascante-Bonilla, Yu Hou, Yang Trista Cao, Hal Daumé III, Rachel Rudinger</p>	37583
<p>Supervised and Semi-Supervised Diffusion Maps with Label-Driven Diffusion</p> <p>Harel Mendelman, Ronen Talmon</p>	37605
<p>PhysPDE: Rethinking PDE Discovery and a Physical Hypothesis Selection Benchmark</p> <p>Mingquan Feng, Yixin Huang, Yizhou Liu, Bofang Jiang, Junchi Yan</p>	37640
<p>Multi-Domain Distribution Learning for De Novo Drug Design</p> <p>Arne Schneuing, Iliia Igashov, Adrian W. Dobbelstein, Thomas Castiglione, Michael M. Bronstein, Bruno Correia</p>	37663
<p>A Computational Framework for Modeling Emergence of Color Vision in the Human Brain</p> <p>Atsunobu Kotani, Ren Ng</p>	37700
<p>SLMRec: Distilling Large Language Models into Small for Sequential Recommendation</p> <p>Wujiang Xu, Qitian Wu, Zujie Liang, Jiaojiao Han, Xuying Ning, Yunxiao Shi, Wenfang Lin, Yongfeng Zhang</p>	37733
<p>Bridging the Data Provenance Gap Across Text, Speech, and Video</p> <p>Shayne Longpre, Nikhil Singh, Manuel Cherep, Kushagra Tiwary, Joanna Materzynska, William Brannon, Robert Mahari, Naana Obeng-Marnu, Manan Dey, Mohammed Hamdy, Nayan Saxena, Ahmad Mustafa Anis, Emad A. Alghamdi, Vu Minh Chien, Da Yin, Kun Qian, Yizhi Li, Minnie Liang, An Dinh, Shrestha Mohanty, Deividas Mataciunas, Tobin South, Jianguo Zhang, Ariel N. Lee, Campbell S. Lund, Christopher Klamm, Damien Sileo, Diganta Misra, Enrico Shippole, Kevin Klyman, Lester James Validad Miranda, Niklas Muennighoff, Seonghyeon Ye, Seungone Kim, Vipul Gupta, Vivek Sharma, Xuhui Zhou, Caiming Xiong, Luis Villa, Stella Biderman, Alex Pentland, Sara Hooker, Jad Kabbara</p>	37755
<p>Near-Optimal Policy Identification in Robust Constrained Markov Decision Processes Via Epigraph Form</p> <p>Toshinori Kitamura, Tadashi Kozuno, Wataru Kumagai, Kenta Hoshino, Yohei Hosoe, Kazumi Kasaura, Masashi Hamaya, Paavo Parmas, Yutaka Matsuo</p>	37834
<p>3D StreetUnveiler with Semantic-Aware 2DGS - A Simple Baseline</p> <p>Jingwei Xu, Yikai Wang, Yiqun Zhao, Yanwei Fu, Shenghua Gao</p>	37871
<p>Grounding Video Models to Actions Through Goal Conditioned Exploration</p> <p>Yunhao Luo, Yilun Du</p>	37905

<p>Beyond Interpretability: The Gains of Feature Monosemanticity on Model Robustness</p> <p><i>Qi Zhang, Yifei Wang, Jingyi Cui, Xiang Pan, Qi Lei, Stefanie Jegelka, Yisen Wang</i></p>	37938
<p>Local Loss Optimization in the Infinite Width: Stable Parameterization of Predictive Coding Networks and Target Propagation</p> <p><i>Satoki Ishikawa, Rio Yokota, Ryo Karakida</i></p>	37971
<p>GNNs Getting ComFy: Community and Feature Similarity Guided Rewiring</p> <p><i>Celia Rubio-Madrigal, Adarsh Jamadandi, Rebekka Burkholz</i></p>	38011
<p>SWE-Search: Enhancing Software Agents with Monte Carlo Tree Search and Iterative Refinement</p> <p><i>Antonis Antoniadis, Albert Örwall, Kexun Zhang, Yuxi Xie, Anirudh Goyal, William Yang Wang</i></p>	38045
<p>Implicit In-Context Learning</p> <p><i>Zhuwei Li, Zihao Xu, Ligong Han, Yunhe Gao, Song Wen, Di Liu, Hao Wang, Dimitris N. Metaxas</i></p>	38076
<p>Self-Normalized Resets for Plasticity in Continual Learning</p> <p><i>Vivek Farias, Adam Daniel Jozefiak</i></p>	38099
<p>Singular Subspace Perturbation Bounds Via Rectangular Random Matrix Diffusions</p> <p><i>Peiyao Lai, Oren Mangoubi</i></p>	38134
<p>Unifying Unsupervised Graph-Level Anomaly Detection and Out-Of-Distribution Detection: A Benchmark</p> <p><i>Yili Wang, Yixin Liu, Xu Shen, Chenyu Li, Rui Miao, Kaize Ding, Ying Wang, Shirui Pan, Xin Wang</i></p>	38164
<p>Ferret-UI 2: Mastering Universal User Interface Understanding Across Platforms</p> <p><i>Zhangheng Li, Keen You, Haotian Zhang, Di Feng, Harsh Agrawal, Xiujun Li, Mohana Prasad Sathya Moorthy, Jeffrey Nichols, Yinfei Yang, Zhe Gan</i></p>	38191
<p>SMT: Fine-Tuning Large Language Models with Sparse Matrices</p> <p><i>Haoze He, Juncheng B Li, Xuan Jiang, Heather Miller</i></p>	38215
<p>ExACT: Teaching AI Agents to Explore with Reflective-MCTS and Exploratory Learning</p> <p><i>Xiao Yu, Baolin Peng, Vineeth Vajipey, Hao Cheng, Michel Galley, Jianfeng Gao, Zhou Yu</i></p>	38240
<p>Unlocking the Power of Function Vectors for Characterizing and Mitigating Catastrophic Forgetting in Continual Instruction Tuning</p> <p><i>Gangwei Jiang, Caigao Jiang, Zhaoyi Li, Siqiao Xue, Jun Zhou, Linqi Song, Defu Lian, Ying Wei</i></p>	38268
<p>Adversarial Mixup Unlearning</p> <p><i>Zhuoyi Peng, Yixuan Tang, Yi Yang</i></p>	38294
<p>Neural Functions for Learning Periodic Signal</p> <p><i>Woojin Cho, Minju Jo, Kookjin Lee, Noseong Park</i></p>	38314

<p>Improving the Sparse Structure Learning of Spiking Neural Networks from the View of Compression Efficiency</p> <p><i>Jiangrong Shen, Qi Xu, Gang Pan, Badong Chen</i></p>	38355
<p>Rethinking Diffusion Posterior Sampling: From Conditional Score Estimator to Maximizing a Posterior</p> <p><i>Tongda Xu, Xiyan Cai, Xinjie Zhang, Xingtong Ge, Dailan He, Ming Sun, Jingjing Liu, Ya-Qin Zhang, Jian Li, Yan Wang</i></p>	38368
<p>NeurFlow: Interpreting Neural Networks Through Neuron Groups and Functional Interactions</p> <p><i>Tue Minh Cao, Nhat Hoang-Xuan, Hieu Pham, Phi Le Nguyen, My T. Thai</i></p>	38392
<p>Exploring Channel Distinguishability in Local Neighborhoods of the Model Space in Quantum Neural Networks</p> <p><i>Sabrina Herbst, Sandeep Suresh Cranganore, Vincenzo De Maio, Ivona Brandic</i></p>	38416
<p>Precise Parameter Localization for Textual Generation in Diffusion Models</p> <p><i>Lukasz Staniszewski, Bartosz Cywinski, Franziska Boenisch, Kamil Deja, Adam Dziedzic</i></p>	38437
<p>OVTR: End-To-End Open-Vocabulary Multiple Object Tracking with Transformer</p> <p><i>Jinyang Li, En Yu, Sijia Chen, Wenbing Tao</i></p>	38461
<p>RaSA: Rank-Sharing Low-Rank Adaptation</p> <p><i>Zhiwei He, Zhaopeng Tu, Xing Wang, Xingyu Chen, Zhijie Wang, Jiahao Xu, Tian Liang, Wenxiang Jiao, Zhuosheng Zhang, Rui Wang</i></p>	38482
<p>How Learnable Grids Recover Fine Detail in Low Dimensions: A Neural Tangent Kernel Analysis of Multigrid Parametric Encodings</p> <p><i>Samuel Audia, Soheil Feizi, Matthias Zwicker, Dinesh Manocha</i></p>	38497
<p>Rational Decision-Making Agent with Learning Internal Utility Judgment</p> <p><i>Yining Ye, Xin Cong, Shizuo Tian, Yujia Qin, Chong Liu, Yankai Lin, Zhiyuan Liu, Maosong Sun</i></p>	38518
<p>Taming Transformer Without Using Learning Rate Warmup</p> <p><i>Xianbiao Qi, Yelin He, Jiaquan Ye, Chun-Guang Li, Bojia Zi, Xili Dai, Qin Zou, Rong Xiao</i></p>	38538
<p>Unhackable Temporal Reward for Scalable Video MLLMs</p> <p><i>En Yu, Kangheng Lin, Liang Zhao, Yana Wei, Zining Zhu, Haoran Wei, Jianjian Sun, Zheng Ge, Xiangyu Zhang, Jingyu Wang, Wenbing Tao</i></p>	38571
<p>Efficient Active Imitation Learning with Random Network Distillation</p> <p><i>Emilien Biré, Anthony Kobanda, Ludovic Denoyer, Rémy Portelas</i></p>	38594
<p>Strategist: Self-Improvement of LLM Decision Making Via Bi-Level Tree Search</p> <p><i>Jonathan Light, Min Cai, Weiqin Chen, Guanzhi Wang, Xiusi Chen, Wei Cheng, Yisong Yue, Ziniu Hu</i></p>	38613
<p>CrossMPT: Cross-Attention Message-Passing Transformer for Error Correcting Codes</p> <p><i>Seong-Joon Park, Hee-Youl Kwak, Sang-Hyo Kim, Yongjune Kim, Jong-Seon No</i></p>	38681

<p>How to Verify Any (Reasonable) Distribution Property: Computationally Sound Argument Systems for Distributions</p> <p>Tal Herman, Guy N. Rothblum</p>	38700
<p>Build-A-Scene: Interactive 3D Layout Control for Diffusion-Based Image Generation</p> <p>Abdelrahman Eldesokey, Peter Wonka</p>	38726
<p>MMQA: Evaluating LLMs with Multi-Table Multi-Hop Complex Questions</p> <p>Jian Wu, Linyi Yang, Dongyuan Li, Yuliang Ji, Manabu Okumura, Yue Zhang</p>	38744
<p>Integrative Decoding: Improving Factuality Via Implicit Self-Consistency</p> <p>Yi Cheng, Xiao Liang, Yeyun Gong, Wen Xiao, Song Wang, Yuji Zhang, Wenjun Hou, Kaishuai Xu, Wenge Liu, Wenjie Li, Jian Jiao, Qi Chen, Peng Cheng, Wayne Xiong</p>	38762
<p>StructRAG: Boosting Knowledge Intensive Reasoning of LLMs Via Inference-Time Hybrid Information Structurization</p> <p>Zhuoqun Li, Xuanang Chen, Haiyang Yu, Hongyu Lin, Yaojie Lu, Qiaoyu Tang, Fei Huang, Xianpei Han, Le Sun, Yongbin Li</p>	38808
<p>Composing Unbalanced Flows for Flexible Docking and Relaxation</p> <p>Gabriele Corso, Vignesh Ram Somnath, Noah Getz, Regina Barzilay, Tommi Jaakkola, Andreas Krause</p>	38826
<p>Bilinear MLPs Enable Weight-Based Mechanistic Interpretability</p> <p>Michael T Pearce, Thomas Dooms, Alice Rigg, Jose Oramas, Lee Sharkey</p>	38850
<p>Presto! Distilling Steps and Layers for Accelerating Music Generation</p> <p>Zachary Novack, Ge Zhu, Jonah Casebeer, Julian McAuley, Taylor Berg-Kirkpatrick, Nicholas J. Bryan</p>	38878
<p>Enhancing Pre-Trained Representation Classifiability Can Boost Its Interpretability</p> <p>Shufan Shen, Zhaobo Qi, Junshu Sun, Qingming Huang, Qi Tian, Shuhui Wang</p>	38903
<p>Safety Representations for Safer Policy Learning</p> <p>Kaustubh Mani, Vincent Mai, Charlie Gauthier, Annie S Chen, Samer B. Nashed, Liam Paull</p>	38928
<p>Bi-Factorial Preference Optimization: Balancing Safety-Helpfulness in Language Models</p> <p>Wenxuan Zhang, Philip Torr, Mohamed Elhoseiny, Adel Bibi</p>	38950
<p>GraphBridge: Towards Arbitrary Transfer Learning in GNNs</p> <p>Li Ju, Xingyi Yang, Qi Li, Xinchao Wang</p>	38976
<p>Meissonic: Revitalizing Masked Generative Transformers for Efficient High-Resolution Text-To-Image Synthesis</p> <p>Jinbin Bai, Tian Ye, Wei Chow, Enxin Song, Qing-Guo Chen, Xiangtai Li, Zhen Dong, Lei Zhu, Shuicheng Yan</p>	39001
<p>Score Forgetting Distillation: A Swift, Data-Free Method for Machine Unlearning in Diffusion Models</p> <p>Tianqi Chen, Shujian Zhang, Mingyuan Zhou</p>	39059

<p>Towards Auto-Regressive Next-Token Prediction: In-Context Learning Emerges from Generalization</p> <p>Zixuan Gong, Xiaolin Hu, Huayi Tang, Yong Liu</p>	39084
<p>TFG-Flow: Training-Free Guidance in Multimodal Generative Flow</p> <p>Haowei Lin, Shanda Li, Haotian Ye, Yiming Yang, Stefano Ermon, Yitao Liang, Jianzhu Ma</p>	39130
<p>ADMM for Nonconvex Optimization Under Minimal Continuity Assumption</p> <p>Ganzhao Yuan</p>	39166
<p>SCBench: A KV Cache-Centric Analysis of Long-Context Methods</p> <p>Yucheng Li, Huiqiang Jiang, Qianhui Wu, Xufang Luo, Surin Ahn, Chengruidong Zhang, Amir H. Abdi, Dongsheng Li, Jianfeng Gao, Yuqing Yang, Lili Qiu</p>	39209
<p>Multi-Agent Cooperation Through Learning-Aware Policy Gradients</p> <p>Alexander Meulemans, Seijin Kobayashi, Johannes Von Oswald, Nino Scherrer, Eric Elmoznino, Blake Aaron Richards, Guillaume Lajoie, Blaise Aguera Y Arcas, Joao Sacramento</p>	39240
<p>Scalable Influence and Fact Tracing for Large Language Model Pretraining</p> <p>Tyler A. Chang, Dheeraj Rajagopal, Tolga Bolukbasi, Lucas Dixon, Ian Tenney</p>	39272
<p>QPM: Discrete Optimization for Globally Interpretable Image Classification</p> <p>Thomas Norrenbrock, Timo Kaiser, Sovan Biswas, Ramesh Manuvinakurike, Bodo Rosenhahn</p>	39294
<p>Geometry Image Diffusion: Fast and Data-Efficient Text-To-3D with Image-Based Surface Representation</p> <p>Slava Elizarov, Ciara Rowles, Simon Donné</p>	39339
<p>Skill Expansion and Composition in Parameter Space</p> <p>Tenglong Liu, Jianxiong Li, Yinan Zheng, Haoyi Niu, Yixing Lan, Xin Xu, Xianyuan Zhan</p>	39359
<p>Rethinking Neural Multi-Objective Combinatorial Optimization Via Neat Weight Embedding</p> <p>Jinbiao Chen, Zhiguang Cao, Jiahai Wang, Yaoxin Wu, Hanzhang Qin, Zizhen Zhang, Yue-Jiao Gong</p>	39396
<p>Restructuring Vector Quantization with the Rotation Trick</p> <p>Christopher Fifty, Ronald Guenther Junkins, Dennis Duan, Aniketh Iyengar, Jerry Weihong Liu, Ehsan Amid, Sebastian Thrun, Christopher Re</p>	39416
<p>DocMIA: Document-Level Membership Inference Attacks Against DocVQA Models</p> <p>Khanh Nguyen, Raouf Kerkouche, Mario Fritz, Dimosthenis Karatzas</p>	39452
<p>Directional Gradient Projection for Robust Fine-Tuning of Foundation Models</p> <p>Chengyue Huang, Junjiao Tian, Brisa Manechotesuwan, Shivang Chopra, Zsolt Kira</p>	39476
<p>ClimaQA: An Automated Evaluation Framework for Climate Question Answering Models</p> <p>Veeramakali Vignesh Manivannan, Yasaman Jafari, Srikar Eranky, Spencer Ho, Rose Yu, Duncan Watson-Parris, Yian Ma, Leon Bergen, Taylor Berg-Kirkpatrick</p>	39497

<p>SaLoRA: Safety-Alignment Preserved Low-Rank Adaptation</p> <p>Mingjie Li, Wai Man Si, Michael Backes, Yang Zhang, Yisen Wang</p>	39523
<p>Basis Sharing: Cross-Layer Parameter Sharing for Large Language Model Compression</p> <p>Jingcun Wang, Yu-Guang Chen, Ing-Chao Lin, Bing Li, Grace Li Zhang</p>	39540
<p>Interpreting the Second-Order Effects of Neurons in CLIP</p> <p>Yossi Gandelsman, Alexei A Efros, Jacob Steinhardt</p>	39559
<p>The Crystal Ball Hypothesis in Diffusion Models: Anticipating Object Positions from Initial Noise</p> <p>Yuanhao Ban, Ruochen Wang, Tianyi Zhou, Boqing Gong, Cho-Jui Hsieh, Minhao Cheng</p>	39579
<p>Towards Counterfactual Fairness Through Auxiliary Variables</p> <p>Bowei Tian, Ziyao Wang, Shwai He, Wanghao Ye, Guoheng Sun, Yucong Dai, Yongkai Wu, Ang Li</p>	39604
<p>Optimizing $\\$(L_0, L_1)\\$-Smooth Functions by Gradient Methods</p> <p>Daniil Vankov, Anton Rodomanov, Angelia Nedich, Lalitha Sankar, Sebastian U Stich</p>	39615
<p>MuseGNN: Forming Scalable, Convergent GNN Layers that Minimize a Sampling-Based Energy</p> <p>Haitian Jiang, Renjie Liu, Zengfeng Huang, Yichuan Wang, Xiao Yan, Zhenkun Cai, Minjie Wang, David Wipf</p>	39642
<p>A Conditional Independence Test in the Presence of Discretization</p> <p>Boyang Sun, Yu Yao, Guang-Yuan Hao, Yumou Qiu, Kun Zhang</p>	39668
<p>Rethinking Visual Counterfactual Explanations Through Region Constraint</p> <p>Bartłomiej Sobieski, Jakub Grzywaczewski, Bartłomiej Sadlej, Matthew Tivnan, Przemysław Biecek</p>	39701
<p>Weakly Supervised Video Scene Graph Generation Via Natural Language Supervision</p> <p>Kibum Kim, Kanghoon Yoon, Yeonjun In, Jaehyeong Jeon, Jinyoung Moon, Donghyun Kim, Chanyoung Park</p>	39754
<p>Learning Stochastic Dynamics from Snapshots Through Regularized Unbalanced Optimal Transport</p> <p>Zhenyi Zhang, Tiejun Li, Peijie Zhou</p>	39777
<p>Rare Event Modeling with Self-Regularized Normalizing Flows: What Can We Learn from a Single Failure?</p> <p>Charles Dawson, Van Tran, Max Z. Li, Chuchu Fan</p>	39809
<p>MAPS: Advancing Multi-Modal Reasoning in Expert-Level Physical Science</p> <p>Erle Zhu, Yadi Liu, Zhe Zhang, Xujun Li, Jinzhou, Xinjie Yu, Minlie Huang, Hongning Wang</p>	39831
<p>Adversarially Robust Out-Of-Distribution Detection Using Lyapunov-Stabilized Embeddings</p> <p>Hossein Mirzaei, Mackenzie W Mathis</p>	39857

<p>Minimal Variance Model Aggregation: A Principled, Non-Intrusive, and Versatile Integration of Black Box Models</p> <p>Theo Bourdais, Houman Owhadi</p>	39875
<p>Oscillatory State-Space Models</p> <p>T. Konstantin Rusch, Daniela Rus</p>	39901
<p>RFWave: Multi-Band Rectified Flow for Audio Waveform Reconstruction</p> <p>Peng Liu, Dongyang Dai, Zhiyong Wu</p>	39921
<p>Unify ML4TSP: Drawing Methodological Principles for TSP and Beyond from Streamlined Design Space of Learning and Search</p> <p>Yang Li, Jiale Ma, Wenzheng Pan, Runzhong Wang, Haoyu Geng, Nianzu Yang, Junchi Yan</p>	39954
<p>Robust Simulation-Based Inference Under Missing Data Via Neural Processes</p> <p>Yogesh Verma, Ayush Bharti, Vikas Garg</p>	39991
<p>Beyond Surface Structure: A Causal Assessment of LLMs' Comprehension Ability</p> <p>Yujin Han, Lei Xu, Sirui Chen, Difan Zou, Chaochao Lu</p>	40014
<p>Bridging Compressed Image Latents and Multimodal Large Language Models</p> <p>Chia-Hao Kao, Cheng Chien, Yu-Jen Tseng, Yi-Hsin Chen, Alessandro Gnutti, Shao-Yuan Lo, Wen-Hsiao Peng, Riccardo Leonardi</p>	40042
<p>Mutual Effort for Efficiency: A Similarity-Based Token Pruning for Vision Transformers in Self-Supervised Learning</p> <p>Sheng Li, Qitao Tan, Yue Dai, Zhenglun Kong, Tianyu Wang, Jun Liu, Ao Li, Ninghao Liu, Yufei Ding, Xulong Tang, Geng Yuan</p>	40063
<p>Learning Gain Map for Inverse Tone Mapping</p> <p>Yinuo Liao, Yuanshen Guan, Ruikang Xu, Jiacheng Li, Shida Sun, Zhiwei Xiong</p>	40081
<p>3D-AffordanceLLM: Harnessing Large Language Models for Open-Vocabulary Affordance Detection in 3D Worlds</p> <p>Hengshuo Chu, Xiang Deng, Qi Lv, Xiaoyang Chen, Yinchuan Li, Jianye Hao, Liqiang Nie</p>	40099
<p>Advancing Mathematical Reasoning in Language Models: The Impact of Problem-Solving Data, Data Synthesis Methods, and Training Stages</p> <p>Zui Chen, Tianqiao Liu, Tongqing, Mi Tian, Weiqi Luo, Zitao Liu</p>	40116
<p>Round and Round We Go! What Makes Rotary Positional Encodings Useful?</p> <p>Federico Barbero, Alex Vitvitskyi, Christos Perivolaropoulos, Razvan Pascanu, Petar Velickovic</p>	40136
<p>LoRA-Pro: Are Low-Rank Adapters Properly Optimized?</p> <p>Zhengbo Wang, Jian Liang, Ran He, Zilei Wang, Tieniu Tan</p>	40167

<p>Dynamic Loss-Based Sample Reweighting for Improved Large Language Model Pretraining</p> <p>Daouda Sow, Herbert Woisetschläger, Saikiran Bulusu, Shiqiang Wang, Hans Arno Jacobsen, Yingbin Liang</p>	40189
<p>Remove Symmetries to Control Model Expressivity and Improve Optimization</p> <p>Liu Ziyin, Yizhou Xu, Isaac L. Chuang</p>	40220
<p>Efficient Biological Data Acquisition Through Inference Set Design</p> <p>Ihor Neporozhnyi, Julien Roy, Emmanuel Bengio, Jason Hartford</p>	40243
<p>Residual-MPPI: Online Policy Customization for Continuous Control</p> <p>Pengcheng Wang, Chenran Li, Catherine Weaver, Kenta Kawamoto, Masayoshi Tomizuka, Chen Tang, Wei Zhan</p>	40265
<p>Representative Guidance: Diffusion Model Sampling with Coherence</p> <p>Anh-Dung Dinh, Daochang Liu, Chang Xu</p>	40295
<p>Depth Any Video with Scalable Synthetic Data</p> <p>Honghui Yang, Di Huang, Wei Yin, Chunhua Shen, Haifeng Liu, Xiaofei He, Binbin Lin, Wanli Ouyang, Tong He</p>	40316
<p>A Truncated Newton Method for Optimal Transport</p> <p>Mete Kemertas, Amir-Massoud Farahmand, Allan Douglas Jepson</p>	40331
<p>3DTrajMaster: Mastering 3D Trajectory for Multi-Entity Motion in Video Generation</p> <p>Xiao Fu, Xian Liu, Xintao Wang, Sida Peng, Menghan Xia, Xiaoyu Shi, Ziyang Yuan, Pengfei Wan, Di Zhang, Dahua Lin</p>	40365
<p>Drop-Upcycling: Training Sparse Mixture of Experts with Partial Re-Initialization</p> <p>Taishi Nakamura, Takuya Akiba, Kazuki Fujii, Yusuke Oda, Rio Yokota, Jun Suzuki</p>	40379
<p>T-JEPA: Augmentation-Free Self-Supervised Learning for Tabular Data</p> <p>Hugo Thimonier, José Lucas De Melo Costa, Fabrice Popineau, Arpad Rimmel, Bich-Liên Doan</p>	40407
<p>MAI: A Multi-Turn Aggregation-Iteration Model for Composed Image Retrieval</p> <p>Yanzhe Chen, Zhiwen Yang, Jinglin Xu, Yuxin Peng</p>	40433
<p>LIFe-GoM: Generalizable Human Rendering with Learned Iterative Feedback Over Multi-Resolution Gaussians-On-Mesh</p> <p>Jing Wen, Alex Schwing, Shenlong Wang</p>	40453
<p>Uncovering Gaps in How Humans and LLMs Interpret Subjective Language</p> <p>Erik Jones, Arjun Patrawala, Jacob Steinhardt</p>	40473
<p>PortLLM: Personalizing Evolving Large Language Models with Training-Free and Portable Model Patches</p> <p>Rana Shahroz, Pingzhi Li, Sukwon Yun, Zhenyu Wang, Shahriar Nirjon, Chau-Wai Wong, Tianlong Chen</p>	40513

<p>Geometry of Long-Tailed Representation Learning: Rebalancing Features for Skewed Distributions</p> <p>Lingjie Yi, Jiachen Yao, Weimin Lyu, Haibin Ling, Raphael Douady, Chao Chen</p>	40531
<p>Learning to Select Nodes in Branch and Bound with Sufficient Tree Representation</p> <p>Sijia Zhang, Shuli Zeng, Shaoang Li, Feng Wu, Xiangyang Li</p>	40564
<p>ImpScore: A Learnable Metric for Quantifying the Implicitness Level of Sentences</p> <p>Yuxin Wang, Xiaomeng Zhu, Weimin Lyu, Saeed Hassanpour, Soroush Vosoughi</p>	40587
<p>Rethinking the Role of Frames for SE(3)-Invariant Crystal Structure Modeling</p> <p>Yusei Ito, Tatsunori Tanai, Ryo Igarashi, Yoshitaka Ushiku, Kanta Ono</p>	40619
<p>On the Role of Attention Heads in Large Language Model Safety</p> <p>Zhenhong Zhou, Haiyang Yu, Xinghua Zhang, Rongwu Xu, Fei Huang, Kun Wang, Yang Liu, Junfeng Fang, Yongbin Li</p>	40641
<p>An Auditing Test to Detect Behavioral Shift in Language Models</p> <p>Leo Richter, Xuanli He, Pasquale Minervini, Matt Kusner</p>	40671
<p>Learn Your Reference Model for Real Good Alignment</p> <p>Alexey Gorbatovski, Boris Shaposhnikov, Alexey Malakhov, Nikita Surnachev, Yaroslav Aksenov, Ian Maksimov, Nikita Balagansky, Daniil Gavrilov</p>	40698
<p>Sensitivity Verification for Additive Decision Tree Ensembles</p> <p>Arhaan Ahmad, Tanay Vineet Tayal, Ashutosh Gupta, S. Akshay</p>	40736
<p>Mixture-Of-Agents Enhances Large Language Model Capabilities</p> <p>Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, James Zou</p>	40752
<p>Mix-CPT: A Domain Adaptation Framework Via Decoupling Knowledge Learning and Format Alignment</p> <p>Jinhao Jiang, Junyi Li, Xin Zhao, Yang Song, Tao Zhang, Ji-Rong Wen</p>	40772
<p>Atlas Gaussians Diffusion for 3D Generation</p> <p>Haitao Yang, Yuan Dong, Hanwen Jiang, Dejia Xu, Georgios Pavlakos, Qixing Huang</p>	40791
<p>RNNs Are Not Transformers (Yet): The Key Bottleneck on In-Context Retrieval</p> <p>Kaiyue Wen, Xingyu Dang, Kaifeng Lyu</p>	40811
<p>SecureGS: Boosting the Security and Fidelity of 3D Gaussian Splatting Steganography</p> <p>Xuanyu Zhang, Jiarui Meng, Zhipei Xu, Shuzhou Yang, Yanmin Wu, Ronggang Wang, Jian Zhang</p>	40855
<p>SeIKD: Selective Knowledge Distillation Via Optimal Transport Perspective</p> <p>Liangliang Shi, Zhengyan Shi, Junchi Yan</p>	40875
<p>BANGS: Game-Theoretic Node Selection for Graph Self-Training</p> <p>Fangxin Wang, Kay Liu, Sourav Medya, Philip S. Yu</p>	40889

<p>Model Risk-Sensitive Offline Reinforcement Learning</p> <p>Gwangpyo Yoo, Honguk Woo</p>	40917
<p>Second Order Bounds for Contextual Bandits with Function Approximation</p> <p>Aldo Pacchiano</p>	40937
<p>HAMSTER: Hierarchical Action Models for Open-World Robot Manipulation</p> <p>Yi Li, Yuquan Deng, Jesse Zhang, Joel Jang, Marius Memmel, Caelan Reed Garrett, Fabio Ramos, Dieter Fox, Anqi Li, Abhishek Gupta, Ankit Goyal</p>	40972
<p>Sharpness-Aware Black-Box Optimization</p> <p>Feiyang Ye, Yueming Lyu, Xuehao Wang, Masashi Sugiyama, Yu Zhang, Ivor Tsang</p>	41001
<p>Sharper Guarantees for Learning Neural Network Classifiers with Gradient Methods</p> <p>Hossein Taheri, Christos Thrampoulidis, Arya Mazumdar</p>	41032
<p>Improving Complex Reasoning with Dynamic Prompt Corruption: A Soft Prompt Optimization Approach</p> <p>Sinan Fan, Liang Xie, Chen Shen, Ge Teng, Xiaosong Yuan, Xiaofeng Zhang, Chenxi Huang, Wenxiao Wang, Xiaofei He, Jieping Ye</p>	41044
<p>Partially Observed Trajectory Inference Using Optimal Transport and a Dynamics Prior</p> <p>Anming Gu, Edward Chien, Kristjan Greenewald</p>	41060
<p>Bayesian Experimental Design Via Contrastive Diffusions</p> <p>Jacopo Lollo, Christophe Heinkelé, Pierre Alliez, Florence Forbes</p>	41095
<p>MediConfusion: Can You Trust Your AI Radiologist? Probing the Reliability of Multimodal Medical Foundation Models</p> <p>Mohammad Shahab Sepehri, Zalan Fabian, Maryam Soltanolkotabi, Mahdi Soltanolkotabi</p>	41122
<p>Energy-Weighted Flow Matching for Offline Reinforcement Learning</p> <p>Shiyuan Zhang, Weitong Zhang, Quanquan Gu</p>	41147
<p>SAM 2: Segment Anything in Images and Videos</p> <p>Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollar, Christoph Feichtenhofer</p>	41175
<p>Causal Graphical Models for Vision-Language Compositional Understanding</p> <p>Fiorenzo Parascandolo, Nicholas Moratelli, Enver Sangineto, Lorenzo Baraldi, Rita Cucchiara</p>	41219
<p>BioDiscoveryAgent: An AI Agent for Designing Genetic Perturbation Experiments</p> <p>Yusuf H Roohani, Andrew H. Lee, Qian Huang, Jian Vora, Zachary Steinhart, Kexin Huang, Alexander Marson, Percy Liang, Jure Leskovec</p>	41245
<p>Machine Unlearning Fails to Remove Data Poisoning Attacks</p> <p>Martin Pawelczyk, Jimmy Z. Di, Yiwei Lu, Gautam Kamath, Ayush Sekhari, Seth Neel</p>	41295

<p>Diffusion Bridge AutoEncoders for Unsupervised Representation Learning</p> <p>Yeongmin Kim, Kwanghyeon Lee, Minsang Park, Byeonghu Na, Il-Chul Moon</p>	41327
<p>Learning Continually by Spectral Regularization</p> <p>Alex Lewandowski, Michal Borkiewicz, Saurabh Kumar, András György, Dale Schuurmans, Mateusz Ostaszewski, Marlos C. Machado</p>	41369
<p>Data Shapley in One Training Run</p> <p>Jiachen T. Wang, Prateek Mittal, Dawn Song, Ruoxi Jia</p>	41396
<p>Real-Time Video Generation with Pyramid Attention Broadcast</p> <p>Xuanlei Zhao, Xiaolong Jin, Kai Wang, Yang You</p>	41434
<p>Diffusion-Based Decoupled Deterministic and Uncertain Framework for Probabilistic Multivariate Time Series Forecasting</p> <p>Qi Li, Zhenyu Zhang, Lei Yao, Zhaoxia Li, Tianyi Zhong, Yong Zhang</p>	41458
<p>Adaptive Camera Sensor for Vision Models</p> <p>Eunsu Baek, Sung-Hwan Han, Taesik Gong, Hyung-Sin Kim</p>	41487
<p>DLEFT-MKC: Dynamic Late Fusion Multiple Kernel Clustering with Robust Tensor Learning Via Min-Max Optimization</p> <p>Yi Zhang, Siwei Wang, Jiyuan Liu, Shengju Yu, Zhibin Dong, Suyuan Liu, Xinwang Liu, En Zhu</p>	41513
<p>Accelerating Inference of Retrieval-Augmented Generation Via Sparse Context Selection</p> <p>Yun Zhu, Jia-Chen Gu, Caitlin Sikora, Ho Ko, Yinxiao Liu, Chu-Cheng Lin, Lei Shu, Liangchen Luo, Lei Meng, Bang Liu, Jindong Chen</p>	41536
<p>LASER: A Neuro-Symbolic Framework for Learning Spatio-Temporal Scene Graphs with Weak Supervision</p> <p>Jiani Huang, Ziyang Li, Mayur Naik, Ser-Nam Lim</p>	41553
<p>ADAPT: Attentive Self-Distillation and Dual-Decoder Prediction Fusion for Continual Panoptic Segmentation</p> <p>Ze Yang, Shichao Dong, Ruibo Li, Nan Song, Guosheng Lin</p>	41576
<p>Agents' Room: Narrative Generation Through Multi-Step Collaboration</p> <p>Fantine Huot, Reinald Kim Amplayo, Jennimaria Palomaki, Alice Shoshana Jakobovits, Elizabeth Clark, Mirella Lapata</p>	41593
<p>Pareto Prompt Optimization</p> <p>Guang Zhao, Byung-Jun Yoon, Gilchan Park, Shantenu Jha, Shinjae Yoo, Xiaoning Qian</p>	41627
<p>SAFREE: Training-Free and Adaptive Guard for Safe Text-To-Image and Video Generation</p> <p>Jaehong Yoon, Shoubin Yu, Vaidehi Patil, Huaxiu Yao, Mohit Bansal</p>	41643
<p>Watch Less, Do More: Implicit Skill Discovery for Video-Conditioned Policy</p> <p>Jiangxing Wang, Zongqing Lu</p>	41670

<p>Efficient Exploration and Discriminative World Model Learning with an Object-Centric Abstraction</p> <p>Anthony Gx-Chen, Kenneth Marino, Rob Fergus</p>	41688
<p>Towards Improving Exploration Through Sibling Augmented GFlowNets</p> <p>Kanika Madan, Alex Lamb, Emmanuel Bengio, Glen Berseth, Yoshua Bengio</p>	41726
<p>Mixture Compressor for Mixture-Of-Experts LLMs Gains More</p> <p>Wei Huang, Yue Liao, Jianhui Liu, Ruifei He, Haoru Tan, Shiming Zhang, Hongsheng Li, Si Liu, Xiaojuan Qi</p>	41745
<p>Probabilistic Neural Pruning Via Sparsity Evolutionary Fokker-Planck-Kolmogorov Equation</p> <p>Zhanfeng Mo, Haosen Shi, Sinno Jialin Pan</p>	41766
<p>Towards Robust Multimodal Open-Set Test-Time Adaptation Via Adaptive Entropy-Aware Optimization</p> <p>Hao Dong, Eleni Chatzi, Olga Fink</p>	41799
<p>GI-GS: Global Illumination Decomposition on Gaussian Splatting for Inverse Rendering</p> <p>Hongze Chen, Zehong Lin, Jun Zhang</p>	41825
<p>Causal Representation Learning from Multimodal Biomedical Observations</p> <p>Yuewen Sun, Lingjing Kong, Guangyi Chen, Loka Li, Gongxu Luo, Zijian Li, Yixuan Zhang, Yujia Zheng, Mengyue Yang, Petar Stojanov, Eran Segal, Eric P. Xing, Kun Zhang</p>	41855
<p>NatureLM-Audio: An Audio-Language Foundation Model for Bioacoustics</p> <p>David Robinson, Marius Miron, Masato Hagiwara, Olivier Pietquin</p>	41888
<p>Adversarial Search Engine Optimization for Large Language Models</p> <p>Fredrik Nestaas, Edoardo Debenedetti, Florian Tramèr</p>	41909
<p>OBI-Bench: Can LMMs Aid in Study of Ancient Script on Oracle Bones?</p> <p>Zijian Chen, Tingzhu Chen, Wenjun Zhang, Guangtao Zhai</p>	41941
<p>OmniEdit: Building Image Editing Generalist Models Through Specialist Supervision</p> <p>Cong Wei, Zheyang Xiong, Weiming Ren, Xeron Du, Ge Zhang, Wenhui Chen</p>	41972
<p>OvercookedV2: Rethinking Overcooked for Zero-Shot Coordination</p> <p>Tobias Gessler, Tin Dizdarevic, Ani Calinescu, Benjamin Ellis, Andrei Lupu, Jakob Nicolaus Foerster</p>	41985
<p>SAGEPhos: Sage Bio-Coupled and Augmented Fusion for Phosphorylation Site Detection</p> <p>Jingjie Zhang, Hanqun Cao, Zijun Gao, Xiaorui Wang, Chunbin Gu</p>	42018
<p>Can Knowledge Editing Really Correct Hallucinations?</p> <p>Baixiang Huang, Canyu Chen, Xiong Xiao Xu, Ali Payani, Kai Shu</p>	42035
<p>VL-Cache: Sparsity and Modality-Aware KV Cache Compression for Vision-Language Model Inference Acceleration</p> <p>Dezhan Tu, Danylo Vashchilenko, Yuzhe Lu, Panpan Xu</p>	42069

<p>Guided Score Identity Distillation for Data-Free One-Step Text-To-Image Generation</p> <p>Mingyuan Zhou, Zhendong Wang, Huangjie Zheng, Hai Huang</p>	42090
<p>Efficient Online Reinforcement Learning Fine-Tuning Need Not Retain Offline Data</p> <p>Zhiyuan Zhou, Andy Peng, Qiyang Li, Sergey Levine, Aviral Kumar</p>	42117
<p>RandLoRA: Full Rank Parameter-Efficient Fine-Tuning of Large Models</p> <p>Paul Albert, Frederic Z. Zhang, Hemanth Saratchandran, Cristian Rodriguez-Opazo, Anton Van Den Hengel, Ehsan Abbasnejad</p>	42143
<p>World Model on Million-Length Video and Language with Blockwise RingAttention</p> <p>Hao Liu, Wilson Yan, Matei Zaharia, Pieter Abbeel</p>	42156
<p>MMIE: Massive Multimodal Interleaved Comprehension Benchmark for Large Vision-Language Models</p> <p>Peng Xia, Siwei Han, Shi Qiu, Yiyang Zhou, Zhaoyang Wang, Wenhao Zheng, Zhaorun Chen, Chenhang Cui, Mingyu Ding, Linjie Li, Lijuan Wang, Huaxiu Yao</p>	42181
<p>MamKO: Mamba-Based Koopman Operator for Modeling and Predictive Control</p> <p>Zhaoyang Li, Minghao Han, Xunyuan Yin</p>	42215
<p>Scale-Aware Contrastive Reverse Distillation for Unsupervised Medical Anomaly Detection</p> <p>Chunlei Li, Yilei Shi, Jingliang Hu, Xiao Xiang Zhu, Lichao Mou</p>	42245
<p>Masked Temporal Interpolation Diffusion for Procedure Planning in Instructional Videos</p> <p>Yufan Zhou, Zhaobo Qi, Lingshuai Lin, Junqi Jing, Tingting Chai, Beichen Zhang, Shuhui Wang, Weigang Zhang</p>	42260
<p>Mind Control Through Causal Inference: Predicting Clean Images from Poisoned Data</p> <p>Mengxuan Hu, Zihan Guan, Yi Zeng, Junfeng Guo, Zhongliang Zhou, Jielu Zhang, Ruoxi Jia, Anil Kumar Vullikanti, Sheng Li</p>	42281
<p>PWM: Policy Learning with Multi-Task World Models</p> <p>Ignat Georgiev, Varun Giridhar, Nicklas Hansen, Animesh Garg</p>	42309
<p>AutoG: Towards Automatic Graph Construction from Tabular Data</p> <p>Zhikai Chen, Han Xie, Jian Zhang, Xiang Song, Jiliang Tang, Huzefa Rangwala, George Karypis</p>	42330
<p>DeITA: An Online Document-Level Translation Agent Based on Multi-Level Memory</p> <p>Yutong Wang, Jiali Zeng, Xuebo Liu, Derek F. Wong, Fandong Meng, Jie Zhou, Min Zhang</p>	42364
<p>MixEval-X: Any-To-Any Evaluations from Real-World Data Mixture</p> <p>Jinjie Ni, Yifan Song, Deepanway Ghosal, Bo Li, David Junhao Zhang, Xiang Yue, Fuzhao Xue, Yuntian Deng, Zian Zheng, Kaichen Zhang, Mahir Shah, Kabir Jain, Yang You, Michael Shieh</p>	42388
<p>TDDBench: A Benchmark for Training Data Detection</p> <p>Zhihao Zhu, Yi Yang, Defu Lian</p>	42446

<p>Posterior-Mean Rectified Flow: Towards Minimum MSE Photo-Realistic Image Restoration</p> <p>Guy Ohayon, Tomer Michaeli, Michael Elad</p>	42473
<p>Uncertainty-Aware Decoding with Minimum Bayes Risk</p> <p>Nico Daheim, Clara Meister, Thomas Möllenhoff, Iryna Gurevych</p>	42513
<p>BingoGuard: LLM Content Moderation Tools with Risk Levels</p> <p>Fan Yin, Philippe Laban, Xiangyu Peng, Yilun Zhou, Yixin Mao, Vaibhav Vats, Linnea Ross, Divyansh Agarwal, Caiming Xiong, Chien-Sheng Wu</p>	42536
<p>Rethinking Evaluation of Sparse Autoencoders Through the Representation of Polysemous Words</p> <p>Gouki Minegishi, Hiroki Furuta, Yusuke Iwasawa, Yutaka Matsuo</p>	42569
<p>GEVRM: Goal-Expressive Video Generation Model for Robust Visual Manipulation</p> <p>Hongyin Zhang, Pengxiang Ding, Shangke Lyu, Ying Peng, Donglin Wang</p>	42598
<p>Certifying Counterfactual Bias in LLMs</p> <p>Isha Chaudhary, Qian Hu, Manoj Kumar, Morteza Ziyadi, Rahul Gupta, Gagandeep Singh</p>	42621
<p>Refine Knowledge of Large Language Models Via Adaptive Contrastive Learning</p> <p>Yinghui Li, Haojing Huang, Jiayi Kuang, Yangning Li, Shu-Yu Guo, Chao Qu, Xiaoyu Tan, Hai-Tao Zheng, Ying Shen, Philip S. Yu</p>	42647
<p>Union-Over-Intersections: Object Detection Beyond Winner-Takes-All</p> <p>Aritra Bhowmik, Pascal Mettes, Martin R. Oswald, Cees G. M. Snoek</p>	42664
<p>DiTTo-TTS: Diffusion Transformers for Scalable Text-To-Speech Without Domain-Specific Factors</p> <p>Keon Lee, Dong Won Kim, Jaehyeon Kim, Seungjun Chung, Jaewoong Cho</p>	42681
<p>Learning Generalizable Skills from Offline Multi-Task Data for Multi-Agent Cooperation</p> <p>Sicong Liu, Yang Shu, Chenjuan Guo, Bin Yang</p>	42715
<p>Unified Convergence Analysis for Score-Based Diffusion Models with Deterministic Samplers</p> <p>Runjia Li, Qiwei Di, Quanquan Gu</p>	42742
<p>Straight to Zero: Why Linearly Decaying the Learning Rate to Zero Works Best for LLMs</p> <p>Shane Bergsma, Nolan Simran Dey, Gurpreet Gosal, Gavia Gray, Daria Soboleva, Joel Hestness</p>	42799
<p>Cross-Entropy is All You Need to Invert the Data Generating Process</p> <p>Patrik Reizinger, Alice Bizeul, Attila Juhas, Julia E Vogt, Randall Balestriero, Wieland Brendel, David Klindt</p>	42831
<p>Bridging and Modeling Correlations in Pairwise Data for Direct Preference Optimization</p> <p>Yuxin Jiang, Bo Huang, Yufei Wang, Xingshan Zeng, Liangyou Li, Yasheng Wang, Xin Jiang, Lifeng Shang, Ruiming Tang, Wei Wang</p>	42862

<p>STAR: Synthesis of Tailored Architectures</p> <p><i>Armin W Thomas, Rom Parnichkun, Alexander Amini, Stefano Massaroli, Michael Poli</i></p>	42882
<p>Improving Uncertainty Estimation Through Semantically Diverse Language Generation</p> <p><i>Lukas Aichberger, Kajetan Schweighofer, Mykyta Ielanskyi, Sepp Hochreiter</i></p>	42935
<p>No Location Left Behind: Measuring and Improving the Fairness of Implicit Representations for Earth Data</p> <p><i>Daniel Cai, Randall Balestriero</i></p>	42960
<p>Gaussian Head & Shoulders: High Fidelity Neural Upper Body Avatars with Anchor Gaussian Guided Texture Warping</p> <p><i>Tianhao Walter Wu, Jing Yang, Zhilin Guo, Jingyi Wan, Fangcheng Zhong, Cengiz Oztireli</i></p>	42987
<p>Investigating the Pre-Training Dynamics of In-Context Learning: Task Recognition Vs. Task Learning</p> <p><i>Xiaolei Wang, Xinyu Tang, Junyi Li, Xin Zhao, Ji-Rong Wen</i></p>	43001
<p>Large Language Models Are Interpretable Learners</p> <p><i>Ruo Chen Wang, Si Si, Felix Yu, Dorothea Wiesmann Rothuizen, Cho-Jui Hsieh, Inderjit S Dhillon</i></p>	43018
<p>GSBA^{AK}: Stop-\$K\$ Geometric Score-Based Black-Box Attack</p> <p><i>Md Farhamdur Reza, Richeng Jin, Tianfu Wu, Huaiyu Dai</i></p>	43042
<p>Semantic Image Inversion and Editing Using Rectified Stochastic Differential Equations</p> <p><i>Litu Rout, Yujia Chen, Nataniel Ruiz, Constantine Caramanis, Sanjay Shakkottai, Wen-Sheng Chu</i></p>	43074
<p>Rethinking and Improving Autoformalization: Towards a Faithful Metric and a Dependency Retrieval-Based Approach</p> <p><i>Qi Liu, Xinhao Zheng, Xudong Lu, Qinxiang Cao, Junchi Yan</i></p>	43110
<p>Mining Your Own Secrets: Diffusion Classifier Scores for Continual Personalization of Text-To-Image Diffusion Models</p> <p><i>Saurav Jha, Shiqi Yang, Masato Ishii, Mengjie Zhao, Christian Simon, Muhammad Jehanzeb Mirza, Dong Gong, Lina Yao, Shusuke Takahashi, Yuki Mitsufuji</i></p>	43139
<p>Rethinking LLM Unlearning Objectives: A Gradient Perspective and Go Beyond</p> <p><i>Qizhou Wang, Jin Peng Zhou, Zhanke Zhou, Saebyeol Shin, Bo Han, Kilian Q Weinberger</i></p>	43169
<p>Improving Pretraining Data Using Perplexity Correlations</p> <p><i>Tristan Thrush, Christopher Potts, Tatsunori Hashimoto</i></p>	43204
<p>The Utility and Complexity of In- And Out-Of-Distribution Machine Unlearning</p> <p><i>Youssef Allouah, Joshua Kazdan, Rachid Guerraoui, Sanmi Koyejo</i></p>	43239
<p>DPaI: Differentiable Pruning at Initialization with Node-Path Balance Principle</p> <p><i>Lichuan Xiang, Quan Nguyen-Tri, Lan-Cuong Nguyen, Hoang Pham, Khoat Than, Long Tran-Thanh, Hongkai Wen</i></p>	43265

<p>AlphaEdit: Null-Space Constrained Knowledge Editing for Language Models</p> <p>Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, Tat-Seng Chua</p>	43287
<p>Privately Counting Partially Ordered Data</p> <p>Matthew Joseph, Mónica Ribero, Alexander Yu</p>	43318
<p>MM1.5: Methods, Analysis & Insights from Multimodal LLM Fine-Tuning</p> <p>Haotian Zhang, Mingfei Gao, Zhe Gan, Philipp Dufter, Nina Wenzel, Forrest Huang, Dhruvi Shah, Xianzhi Du, Bowen Zhang, Yanghao Li, Sam Dodge, Keen You, Zhen Yang, Aleksei Timofeev, Mingze Xu, Hong-You Chen, Jean-Philippe Fauconnier, Zhengfeng Lai, Haoxuan You, Zirui Wang, Afshin Dehghan, Peter Grasch, Yinfei Yang</p>	43333
<p>Immunogenicity Prediction with Dual Attention Enables Vaccine Target Selection</p> <p>Song Li, Yang Tan, Song Ke, Liang Hong, Bingxin Zhou</p>	43379
<p>Modeling Complex System Dynamics with Flow Matching Across Time and Conditions</p> <p>Martin Rohbeck, Edward De Brouwer, Charlotte Bunne, Jan-Christian Huetter, Anne Biton, Kelvin Y. Chen, Aviv Regev, Romain Lopez</p>	43392
<p>Understanding Factual Recall in Transformers Via Associative Memories</p> <p>Eshaan Nichani, Jason D. Lee, Alberto Bietti</p>	43417
<p>SonicSim: A Customizable Simulation Platform for Speech Processing in Moving Sound Source Scenarios</p> <p>Kai Li, Wendi Sang, Chang Zeng, Runxuan Yang, Guo Chen, Xiaolin Hu</p>	43459
<p>Jailbreaking Leading Safety-Aligned LLMs with Simple Adaptive Attacks</p> <p>Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion</p>	43486
<p>Super(ficial)-Alignment: Strong Models May Deceive Weak Models in Weak-To-Strong Generalization</p> <p>Wenkai Yang, Shiqi Shen, Guangyao Shen, Wei Yao, Yong Liu, Gong Zhi, Yankai Lin, Ji-Rong Wen</p>	43514
<p>Correcting the Mythos of KL-Regularization: Direct Alignment Without Overoptimization Via Chi-Squared Preference Optimization</p> <p>Audrey Huang, Wenhao Zhan, Tengyang Xie, Jason D. Lee, Wen Sun, Akshay Krishnamurthy, Dylan J Foster</p>	43539
<p>Neural Stochastic Differential Equations for Uncertainty-Aware Offline RL</p> <p>Cevahir Koprulu, Franck Djeumou, Ufuk Topcu</p>	43594
<p>Computationally Efficient RL Under Linear Bellman Completeness for Deterministic Dynamics</p> <p>Runzhe Wu, Ayush Sekhari, Akshay Krishnamurthy, Wen Sun</p>	43619
<p>Glauber Generative Model: Discrete Diffusion Models Via Binary Classification</p> <p>Harshit Varma, Dheeraj Mysore Nagaraj, Karthikeyan Shanmugam</p>	43661

<p>SVBench: A Benchmark with Temporal Multi-Turn Dialogues for Streaming Video Understanding</p> <p>Zhenyu Yang, Yuhang Hu, Zemin Du, Dizhan Xue, Shengsheng Qian, Jiahong Wu, Fan Yang, Weiming Dong, Changsheng Xu</p>	43694
<p>D_2O: Dynamic Discriminative Operations for Efficient Long-Context Inference of Large Language Models</p> <p>Zhongwei Wan, Xinjian Wu, Yu Zhang, Yi Xin, Chaofan Tao, Zhihong Zhu, Xin Wang, Siqi Luo, Jing Xiong, Longyue Wang, Mi Zhang</p>	43729
<p>Generalizing Weisfeiler-Lehman Kernels to Subgraphs</p> <p>Dongkwan Kim, Alice Oh</p>	43755
<p>Deep Distributed Optimization for Large-Scale Quadratic Programming</p> <p>Augustinos D Saravanos, Hunter Kuperman, Alex Oshin, Arshiya Taj Abdul, Vincent Pacelli, Evangelos Theodorou</p>	43770
<p>Learning LLM-As-A-Judge for Preference Alignment</p> <p>Ziyi Ye, Xiangsheng Li, Qiuchi Li, Qingyao Ai, Yujia Zhou, Wei Shen, Dong Yan, Yiqun Liu</p>	43810
<p>Dynamic-LLaVA: Efficient Multimodal Large Language Models Via Dynamic Vision-Language Context Sparsification</p> <p>Wenxuan Huang, Zijie Zhai, Yunhang Shen, Shaosheng Cao, Fei Zhao, Xiangfeng Xu, Zheyu Ye, Shaohui Lin</p>	43838
<p>Fast and Slow Streams for Online Time Series Forecasting Without Information Leakage</p> <p>Ying-Yee Ava Lau, Zhiwen Shao, Dit-Yan Yeung</p>	43867
<p>Parameter and Memory Efficient Pretraining Via Low-Rank Riemannian Optimization</p> <p>Zhanfeng Mo, Long-Kai Huang, Sinno Jialin Pan</p>	43894
<p>SymDiff: Equivariant Diffusion Via Stochastic Symmetrisation</p> <p>Leo Zhang, Kianoosh Ashouritaklimi, Yee Whye Teh, Rob Cornish</p>	43924
<p>ECHOPulse: ECG Controlled Echocardiogram Video Generation</p> <p>Yiwei Li, Sekeun Kim, Zihao Wu, Hanqi Jiang, Yi Pan, Pengfei Jin, Sifan Song, Yucheng Shi, Xiaowei Yu, Tianze Yang, Tianming Liu, Quanzheng Li, Xiang Li</p>	43947
<p>TASAR: Transfer-Based Attack on Skeletal Action Recognition</p> <p>Yunfeng Diao, Baiqi Wu, Ruixuan Zhang, Ajian Liu, Xiaoshuai Hao, Xingxing Wei, Meng Wang, He Wang</p>	43968
<p>Language Guided Skill Discovery</p> <p>Seungeun Rho, Laura Smith, Tianyu Li, Sergey Levine, Xue Bin Peng, Sehoon Ha</p>	43989
<p>Valid Conformal Prediction for Dynamic GNNs</p> <p>Ed Davis, Ian Gallagher, Daniel John Lawson, Patrick Rubin-Delanchy</p>	44011

<p>MM-EMBED: UNIVERSAL MULTIMODAL RETRIEVAL WITH MULTIMODAL LLMS</p> <p>Sheng-Chieh Lin, Chankyu Lee, Mohammad Shoeybi, Jimmy Lin, Bryan Catanzaro, Wei Ping</p>	44036
<p>Sparse Feature Circuits: Discovering and Editing Interpretable Causal Graphs in Language Models</p> <p>Samuel Marks, Can Rager, Eric J Michaud, Yonatan Belinkov, David Bau, Aaron Mueller</p>	44056
<p>ConcreTizer: Model Inversion Attack Via Occupancy Classification and Dispersion Control for 3D Point Cloud Restoration</p> <p>Youngseok Kim, Sunwook Hwang, Hyung-Sin Kim, Saewoong Bahk</p>	44092
<p>How Far Are We from True Unlearnability?</p> <p>Kai Ye, Liangcai Su, Chenxiong Qian</p>	44106
<p>DeepSeek-Prover-V1.5: Harnessing Proof Assistant Feedback for Reinforcement Learning and Monte-Carlo Tree Search</p> <p>Huajian Xin, Z. Z. Ren, Junxiao Song, Zhihong Shao, Wanjia Zhao, Haocheng Wang, Bo Liu, Liyue Zhang, Xuan Lu, Qiushi Du, Wenjun Gao, Haowei Zhang, Qihao Zhu, Dejian Yang, Zhibin Gou, Z. F. Wu, Fuli Luo, Chong Ruan</p>	44123
<p>Simple Guidance Mechanisms for Discrete Diffusion Models</p> <p>Yair Schiff, Subham Sekhar Sahoo, Hao Phung, Guanghan Wang, Sam Boshar, Hugo Dalla-Torre, Bernardo P De Almeida, Alexander M Rush, Thomas Pierrot, Volodymyr Kuleshov</p>	44153
<p>Optimizing Posterior Samples for Bayesian Optimization Via Rootfinding</p> <p>Taiwo Adebiyi, Bach Do, Ruda Zhang</p>	44199
<p>Episodic Novelty Through Temporal Distance</p> <p>Yuhua Jiang, Qihan Liu, Yiqin Yang, Xiaoteng Ma, Dianyu Zhong, Hao Hu, Jun Yang, Bin Liang, Bo Xu, Chongjie Zhang, Qianchuan Zhao</p>	44228
<p>Disentangling 3D Animal Pose Dynamics with Scrubbed Conditional Latent Variables</p> <p>Joshua Huang Wu, Hari Koneru, James Russell Ravenel, Anshuman Sabath, James Michael Roach, Shaun Sze-Xian Lim, Michael R Tadross, Alex H Williams, Timothy W Dunn</p>	44255
<p>Bootstrapped Model Predictive Control</p> <p>Yuhang Wang, Hanwei Guo, Sizhe Wang, Long Qian, Xuguang Lan</p>	44283
<p>Lipschitz Bandits in Optimal Space</p> <p>Xiaoyi Zhu, Zengfeng Huang</p>	44302
<p>Sail into the Headwind: Alignment Via Robust Rewards and Dynamic Labels Against Reward Hacking</p> <p>Paria Rashidinejad, Yuandong Tian</p>	44314
<p>Near-Optimal Online Learning for Multi-Agent Submodular Coordination: Tight Approximation and Communication Efficiency</p> <p>Qixin Zhang, Zongqi Wan, Yu Yang, Li Shen, Dacheng Tao</p>	44359

<p>Interpreting Language Reward Models Via Contrastive Explanations</p> <p><i>Junqi Jiang, Tom Bewley, Saumitra Mishra, Freddy Lecue, Manuela Veloso</i></p>	44392
<p>Minimax Optimal Reinforcement Learning with Quasi-Optimism</p> <p><i>Harin Lee, Min-Hwan Oh</i></p>	44412
<p>Residual Connections and Normalization Can Provably Prevent Oversmoothing in GNNs</p> <p><i>Michael Scholkemper, Xinyi Wu, Ali Jadbabaie, Michael T Schaub</i></p>	44460
<p>Multimodal Situational Safety</p> <p><i>Kaiwen Zhou, Chengzhi Liu, Xuandong Zhao, Anderson Compalas, Dawn Song, Xin Eric Wang</i></p>	44490
<p>Maximizing the Potential of Synthetic Data: Insights from Random Matrix Theory</p> <p><i>Aymane El Firdoussi, Mohamed El Amine Seddik, Soufiane Hayou, Reda Alami, Ahmed Alzubaidi, Hakim Hacid</i></p>	44521
<p>Let Your Features Tell the Differences: Understanding Graph Convolution by Feature Splitting</p> <p><i>Yilun Zheng, Xiang Li, Sitao Luan, Xiaojiang Peng, Lihui Chen</i></p>	44557
<p>MuPT: A Generative Symbolic Music Pretrained Transformer</p> <p><i>Xingwei Qu, Yuelin Bai, Yinghao Ma, Ziya Zhou, Ka Man Lo, Jiaheng Liu, Ruibin Yuan, Lejun Min, Xueling Liu, Tianyu Zhang, Xeron Du, Shuyue Guo, Yiming Liang, Yizhi Li, Shangda Wu, Junting Zhou, Tianyu Zheng, Ziyang Ma, Fengze Han, Wei Xue, Gus Xia, Emmanouil Benetos, Xiang Yue, Chenghua Lin, Xu Tan, Wenhao Huang, Jie Fu, Ge Zhang</i></p>	44591
<p>Improving Graph Neural Networks by Learning Continuous Edge Directions</p> <p><i>Seong Ho Pahng, Sahand Hormoz</i></p>	44618
<p>Uncertainty and Influence Aware Reward Model Refinement for Reinforcement Learning from Human Feedback</p> <p><i>Zexu Sun, Yiju Guo, Yankai Lin, Xu Chen, Qi Qi, Xing Tang, Xiuqiang He, Ji-Rong Wen</i></p>	44654
<p>Efficient Model-Based Reinforcement Learning Through Optimistic Thompson Sampling</p> <p><i>Jasmine Bayrooti, Carl Henrik Ek, Amanda Prorok</i></p>	44676
<p>THE ROBUSTNESS OF DIFFERENTIABLE CAUSAL DISCOVERY IN MISSPECIFIED SCENARIOS</p> <p><i>Huiyang Yi, Yanyan He, Duxin Chen, Mingyu Kang, He Wang, Wenwu Yu</i></p>	44697
<p>Adam-Mini: Use Fewer Learning Rates to Gain More</p> <p><i>Yushun Zhang, Congliang Chen, Ziniu Li, Tian Ding, Chenwei Wu, Diederik P Kingma, Yinyu Ye, Zhi-Quan Luo, Ruoyu Sun</i></p>	44736
<p>Streamlining Redundant Layers to Compress Large Language Models</p> <p><i>Xiaodong Chen, Yuxuan Hu, Jing Zhang, Yanling Wang, Cuiping Li, Hong Chen</i></p>	44767
<p>Pareto Low-Rank Adapters: Efficient Multi-Task Learning with Preferences</p> <p><i>Nikolaos Dimitriadis, Pascal Frossard, François Fleuret</i></p>	44789

<p>Improving Language Model Distillation Through Hidden State Matching</p> <p>Sayantan Dasgupta, Trevor Cohn</p>	44813
<p>Optimizing 4D Gaussians for Dynamic Scene Video from Single Landscape Images</p> <p>In-Hwan Jin, Haesoo Choo, Seong-Hun Jeong, Park Heemoon, Junghwan Kim, Oh-Joon Kwon, Kyeongbo Kong</p>	44828
<p>DynAlign: Unsupervised Dynamic Taxonomy Alignment for Cross-Domain Segmentation</p> <p>Han Sun, Rui Gong, Ismail Nejjar, Olga Fink</p>	44849
<p>EcoFace: Audio-Visual Emotional Co-Disentanglement Speech-Driven 3D Talking Face Generation</p> <p>Jiajian Xie, Shengyu Zhang, Mengze Li, Chengfei Lv, Zhou Zhao, Fei Wu</p>	44873
<p>Spatial-Mamba: Effective Visual State Space Models Via Structure-Aware State Fusion</p> <p>Chaodong Xiao, Minghan Li, Zhengqiang Zhang, Deyu Meng, Lei Zhang</p>	44892
<p>Self-MoE: Towards Compositional Large Language Models with Self-Specialized Experts</p> <p>Junmo Kang, Leonid Karlinsky, Hongyin Luo, Zhen Wang, Jacob A Hansen, James R. Glass, David Daniel Cox, Rameswar Panda, Rogerio Feris, Alan Ritter</p>	44911
<p>SOAP: Improving and Stabilizing Shampoo Using Adam for Language Modeling</p> <p>Nikhil Vyas, Depen Morwani, Rosie Zhao, Itai Shapira, David Brandfonbrener, Lucas Janson, Sham M. Kakade</p>	44935
<p>The Optimization Landscape of SGD Across the Feature Learning Strength</p> <p>Alexander Atanasov, Alexandru Meterez, James B Simon, Cengiz Pehlevan</p>	44957
<p>X-Drive: Cross-Modality Consistent Multi-Sensor Data Synthesis for Driving Scenarios</p> <p>Yichen Xie, Chenfeng Xu, Chensheng Peng, Shuqi Zhao, Nhat Ho, Alexander T. Pham, Mingyu Ding, Masayoshi Tomizuka, Wei Zhan</p>	44997
<p>Intelligence at the Edge of Chaos</p> <p>Shiyang Zhang, Aakash Patel, Syed A Rizvi, Nianchen Liu, Sizhuang He, Amin Karbasi, Emanuele Zappala, David Van Dijk</p>	45014
<p>HGM³: Hierarchical Generative Masked Motion Modeling with Hard Token Mining</p> <p>Minjae Jeong, Yechan Hwang, Jaejin Lee, Sungyoon Jung, Won Hwa Kim</p>	45031
<p>Fast and Accurate Blind Flexible Docking</p> <p>Zizhuo Zhang, Lijun Wu, Kaiyuan Gao, Jiangchao Yao, Tao Qin, Bo Han</p>	45049
<p>Near-Optimal Active Regression of Single-Index Models</p> <p>Yi Li, Wai Ming Tai</p>	45074
<p>Implicit Bias of Mirror Flow for Shallow Neural Networks in Univariate Regression</p> <p>Shuang Liang, Guido Montufar</p>	45116

<p>Provably Robust Explainable Graph Neural Networks Against Graph Perturbation Attacks</p> <p><i>Jiate Li, Meng Pang, Yun Dong, Jinyuan Jia, Binghui Wang</i></p>	45172
<p>SafeDiffuser: Safe Planning with Diffusion Probabilistic Models</p> <p><i>Wei Xiao, Tsun-Hsuan Wang, Chuang Gan, Ramin Hasani, Mathias Lechner, Daniela Rus</i></p>	45192
<p>DenseGrounding: Improving Dense Language-Vision Semantics for Ego-Centric 3D Visual Grounding</p> <p><i>Henry Zheng, Hao Shi, Qihang Peng, Yong Xien Chng, Rui Huang, Yepeng Weng, Zhongchao Shi, Gao Huang</i></p>	45214
<p>Duoduo CLIP: Efficient 3D Understanding with Multi-View Images</p> <p><i>Han-Hung Lee, Yiming Zhang, Angel X Chang</i></p>	45231
<p>Controllable Context Sensitivity and the Knob Behind it</p> <p><i>Julian Minder, Kevin Du, Niklas Stoehr, Giovanni Monea, Chris Wendler, Robert West, Ryan Cotterell</i></p>	45253
<p>Representational Similarity Via Interpretable Visual Concepts</p> <p><i>Neehar Kondapaneni, Oisin Mac Aodha, Pietro Perona</i></p>	45283
<p>Do LLMs Estimate Uncertainty Well in Instruction-Following?</p> <p><i>Juyeon Heo, Miao Xiong, Christina Heinze-Deml, Jaya Narain</i></p>	45315
<p>Start Smart: Leveraging Gradients for Enhancing Mask-Based XAI Methods</p> <p><i>Buelent Uendes, Shujian Yu, Mark Hoogendoorn</i></p>	45339
<p>Multiplicative Logit Adjustment Approximates Neural-Collapse-Aware Decision Boundary Adjustment</p> <p><i>Naoya Hasegawa, Issei Sato</i></p>	45400
<p>Memory Mosaics</p> <p><i>Jianyu Zhang, Niklas Nolte, Ranajoy Sadhukhan, Beidi Chen, Leon Bottou</i></p>	45433
<p>Harnessing Webpage UIs for Text-Rich Visual Understanding</p> <p><i>Junpeng Liu, Tianyue Ou, Yifan Song, Yuxiao Qu, Wai Lam, Chenyan Xiong, Wenhui Chen, Graham Neubig, Xiang Yue</i></p>	45455
<p>Rodimus*: Breaking the Accuracy-Efficiency Trade-Off with Efficient Attentions</p> <p><i>Zhihao He, Hang Yu, Zi Gong, Shizhan Liu, Jianguo Li, Weiyao Lin</i></p>	45490
<p>The Hyperfitting Phenomenon: Sharpening and Stabilizing LLMs for Open-Ended Text Generation</p> <p><i>Fredrik Carlsson, Fangyu Liu, Daniel Ward, Murathan Kurfali, Joakim Nivre</i></p>	45526
<p>Second-Order Min-Max Optimization with Lazy Hessians</p> <p><i>Lesi Chen, Chengchang Liu, Jingzhao Zhang</i></p>	45552

<p>Bio-XLSTM: Generative Modeling, Representation and In-Context Learning of Biological and Chemical Sequences</p> <p><i>Niklas Schmidinger, Lisa Schneckenreiter, Philipp Seidl, Johannes Schimunek, Pieter-Jan Hoedt, Johannes Brandstetter, Andreas Mayr, Sohvi Luukkonen, Sepp Hochreiter, Günter Klambauer</i></p>	45566
<p>From Attention to Activation: Unraveling the Enigmas of Large Language Models</p> <p><i>Prannay Kaul, Chengcheng Ma, Ismail Elezi, Jiankang Deng</i></p>	45608
<p>Diffusion-NPO: Negative Preference Optimization for Better Preference Aligned Generation of Diffusion Models</p> <p><i>Fu-Yun Wang, Yunhao Shui, Jingtian Piao, Keqiang Sun, Hongsheng Li</i></p>	45660
<p>Tell Me About Yourself: LLMs Are Aware of Their Learned Behaviors</p> <p><i>Jan Betley, Xuchan Bao, Martín Soto, Anna Sztyber-Betley, James Chua, Owain Evans</i></p>	45683
<p>3DitScene: Editing Any Scene Via Language-Guided Disentangled Gaussian Splatting</p> <p><i>Qihang Zhang, Yinghao Xu, Chaoyang Wang, Hsin-Ying Lee, Gordon Wetzstein, Bolei Zhou, Ceyuan Yang</i></p>	45736
<p>Adaptive Batch Size for Privately Finding Second-Order Stationary Points</p> <p><i>Daogao Liu, Kunal Talwar</i></p>	45752
<p>The Value of Sensory Information to a Robot</p> <p><i>Arjun Krishna, Edward S. Hu, Dinesh Jayaraman</i></p>	45771
<p>Watermark Anything with Localized Messages</p> <p><i>Tom Sander, Pierre Fernandez, Alain Oliviero Durmus, Teddy Furon, Matthijs Douze</i></p>	45800
<p>AndroidWorld: A Dynamic Benchmarking Environment for Autonomous Agents</p> <p><i>Christopher Rawles, Sarah Clinckemaillie, Yifan Chang, Jonathan Waltz, Gabrielle Lau, Marybeth Fair, Alice Li, William E Bishop, Wei Li, Folawiyi Campbell-Ajala, Daniel Kenji Toyama, Robert James Berry, Divya Tyamagundlu, Timothy P Lillicrap, Oriana Riva</i></p>	45831
<p>AI as Humanity's Saliery: Quantifying Linguistic Creativity of Language Models Via Systematic Attribution of Machine Text Against Web Text</p> <p><i>Ximing Lu, Melanie Sclar, Skyler Hallinan, Niloofar Mireshghallah, Jiacheng Liu, Seungju Han, Allyson Ettinger, Liwei Jiang, Khyathi Chandu, Nouha Dziri, Yejin Choi</i></p>	45867
<p>Why In-Context Learning Models Are Good Few-Shot Learners?</p> <p><i>Shiguang Wu, Yaqing Wang, Quanming Yao</i></p>	45919
<p>Revisiting Text-To-Image Evaluation with Gecko: On Metrics, Prompts, and Human Rating</p> <p><i>Olivia Wiles, Chuhan Zhang, Isabela Albuquerque, Ivana Kajic, Su Wang, Emanuele Bugliarello, Yasumasa Onoe, Pinelopi Papalampidi, Ira Ktena, Christopher Knutsen, Cyrus Rashtchian, Anant Nawalgaria, Jordi Pont-Tuset, Aida Nematzadeh</i></p>	45942
<p>Stochastic Variance-Reduced Gaussian Variational Inference on the Bures-Wasserstein Manifold</p> <p><i>Hoang Phuc Hau Luu, Hanlin Yu, Bernardo Williams, Marcelo Hartmann, Arto Klami</i></p>	45958

<p>IV-Mixed Sampler: Leveraging Image Diffusion Models for Enhanced Video Synthesis</p> <p>Shitong Shao, Zikai Zhou, Bai Lichen, Haoyi Xiong, Zeke Xie</p>	45983
<p>Sparse Autoencoders Reveal Selective Remapping of Visual Concepts During Adaptation</p> <p>Hyesu Lim, Jinho Choi, Jaegul Choo, Steffen Schneider</p>	46012
<p>Can In-Context Learning Really Generalize to Out-Of-Distribution Tasks?</p> <p>Qixun Wang, Yifei Wang, Xianghua Ying, Yisen Wang</p>	46038
<p>Adversarial Training Can Provably Improve Robustness: Theoretical Analysis of Feature Learning Process Under Structured Data</p> <p>Binghui Li, Yuanzhi Li</p>	46060
<p>Fast Summation of Radial Kernels Via QMC Slicing</p> <p>Johannes Hertrich, Tim Jahn, Michael Quellmalz</p>	46096
<p>Towards Explaining the Power of Constant-Depth Graph Neural Networks for Structured Linear Programming</p> <p>Qian Li, Minghui Ouyang, Tian Ding, Yuyi Wang, Qingjiang Shi, Ruoyu Sun</p>	46132
<p>M³PC: Test-Time Model Predictive Control Using Pretrained Masked Trajectory Model</p> <p>Kehan Wen, Yutong Hu, Yao Mu, Lei Ke</p>	46148
<p>Narrowing Information Bottleneck Theory for Multimodal Image-Text Representations Interpretability</p> <p>Zhiyu Zhu, Zhibo Jin, Jiayu Zhang, Nan Yang, Jiahao Huang, Jianlong Zhou, Fang Chen</p>	46171
<p>Differential Learning Kinetics Govern the Transition from Memorization to Generalization During In-Context Learning</p> <p>Alex Nguyen, Gautam Reddy</p>	46189
<p>Proactive Privacy Amnesia for Large Language Models: Safeguarding PII with Negligible Impact on Model Utility</p> <p>Martin Kuo, Jingyang Zhang, Jianyi Zhang, Minxue Tang, Louis Divalentin, Aolin Ding, Jingwei Sun, William Chen, Amin Hass, Tianlong Chen, Yiran Chen, Hai Li</p>	46210
<p>NeuroLM: A Universal Multi-Task Foundation Model for Bridging the Gap Between Language and EEG Signals</p> <p>Weibang Jiang, Yansen Wang, Bao-Liang Lu, Dongsheng Li</p>	46232
<p>Training Large Language Models for Retrieval-Augmented Question Answering Through Backtracking Correction</p> <p>Huawen Feng, Zekun Yao, Junhao Zheng, Qianli Ma</p>	46254
<p>Utility-Directed Conformal Prediction: A Decision-Aware Framework for Actionable Uncertainty Quantification</p> <p>Santiago Cortes-Gomez, Carlos Miguel Patiño, Yewon Byun, Steven Wu, Eric Horvitz, Bryan Wilder</p>	46273

<p>Conditional Testing Based on Localized Conformal ρ-Values</p> <p><i>Xiaoyang Wu, Lin Lu, Zhaojun Wang, Changliang Zou</i></p>	46293
<p>Connecting Federated ADMM to Bayes</p> <p><i>Siddharth Swaroop, Mohammad Emtiyaz Khan, Finale Doshi-Velez</i></p>	46326
<p>Sparse Components Distinguish Visual Pathways & Their Alignment to Neural Networks</p> <p><i>Ammar I Marvi, Nancy Kanwisher, Meenakshi Khosla</i></p>	46346
<p>On Designing General and Expressive Quantum Graph Neural Networks with Applications to MILP Instance Representation</p> <p><i>Xinyu Ye, Hao Xiong, Jianhao Huang, Ziang Chen, Jia Wang, Junchi Yan</i></p>	46364
<p>Generalization V.s. Memorization: Tracing Language Models' Capabilities Back to Pretraining Data</p> <p><i>Xinyi Wang, Antonis Antoniadis, Yanai Elazar, Alfonso Amayuelas, Alon Albalak, Kexun Zhang, William Yang Wang</i></p>	46390
<p>Leveraging Submodule Linearity Enhances Task Arithmetic Performance in LLMs</p> <p><i>Rui Dai, Sile Hu, Xu Shen, Yonggang Zhang, Xinmei Tian, Jieping Ye</i></p>	46411
<p>Risk-Sensitive Variational Actor-Critic: A Model-Based Approach</p> <p><i>Alonso Granados, Reza Ebrahimi, Jason Pacheco</i></p>	46439
<p>Can Neural Networks Achieve Optimal Computational-Statistical Tradeoff? an Analysis on Single-Index Model</p> <p><i>Siyu Chen, Beining Wu, Miao Lu, Zhuoran Yang, Tianhao Wang</i></p>	46461
<p>Action Abstractions for Amortized Sampling</p> <p><i>Oussama Boussif, Lena Nehale Ezzine, Joseph D Viviano, Michal Koziarski, Moksh Jain, Nikolay Malkin, Emmanuel Bengio, Rim Assouel, Yoshua Bengio</i></p>	46546
<p>LeanQuant: Accurate and Scalable Large Language Model Quantization with Loss-Error-Aware Grid</p> <p><i>Tianyi Zhang, Anshumali Shrivastava</i></p>	46576
<p>Improving Reasoning Performance in Large Language Models Via Representation Engineering</p> <p><i>Bertram Højer, Oliver Simon Jarvis, Stefan Heinrich</i></p>	46600
<p>ClassDiffusion: More Aligned Personalization Tuning with Explicit Class Guidance</p> <p><i>Jiannan Huang, Jun Hao Liew, Hanshu Yan, Yuyang Yin, Yao Zhao, Humphrey Shi, Yunchao Wei</i></p>	46618
<p>PseDet: Revisiting the Power of Pseudo Label in Incremental Object Detection</p> <p><i>Qiuchen Wang, Zehui Chen, Chenhongyi Yang, Jiaming Liu, Zhenyu Li, Feng Zhao</i></p>	46641
<p>Post-Hoc Reward Calibration: A Case Study on Length Bias</p> <p><i>Zeyu Huang, Zihan Qiu, Zili Wang, Edoardo Ponti, Ivan Titov</i></p>	46658

<p>CEB: Compositional Evaluation Benchmark for Fairness in Large Language Models</p> <p>Song Wang, Peng Wang, Tong Zhou, Yushun Dong, Zhen Tan, Jundong Li</p>	46685
<p>Multi-Session, Multi-Task Neural Decoding from Distinct Cell-Types and Brain Regions</p> <p>Mehdi Azabou, Krystal Xuejing Pan, Vinam Arora, Ian Jarratt Knight, Eva L Dyer, Blake Aaron Richards</p>	46727
<p>Recognize Any Surgical Object: Unleashing the Power of Weakly-Supervised Data</p> <p>Jiajie Li, Brian R Quaranto, Chenhui Xu, Ishan Mishra, Ruiyang Qin, Dancheng Liu, Peter C W Kim, Jinjun Xiong</p>	46751
<p>Autonomous Evaluation of LLMs for Truth Maintenance and Reasoning Tasks</p> <p>Rushang Karia, Daniel Richard Bramblett, Daksh Dobhal, Siddharth Srivastava</p>	46764
<p>Can We Talk Models into Seeing the World Differently?</p> <p>Paul Gavrikov, Jovita Lukasik, Steffen Jung, Robert Geirhos, Muhammad Jehanzeb Mirza, Margret Keuper, Janis Keuper</p>	46810
<p>LLaRA: Supercharging Robot Learning Data for Vision-Language Policy</p> <p>Xiang Li, Cristina Mata, Jongwoo Park, Kumara Kahatapitiya, Yoo Sung Jang, Jinghuan Shang, Kanchana Ranasinghe, Ryan D Burgert, Mu Cai, Yong Jae Lee, Michael S Ryoo</p>	46842
<p>MedTrinity-25M: A Large-Scale Multimodal Dataset with Multigranular Annotations for Medicine</p> <p>Yunfei Xie, Ce Zhou, Lang Gao, Juncheng Wu, Xianhang Li, Hong-Yu Zhou, Sheng Liu, Lei Xing, James Zou, Cihang Xie, Yuyin Zhou</p>	46889
<p>Enhancing Learning with Label Differential Privacy by Vector Approximation</p> <p>Puning Zhao, Jiafei Wu, Zhe Liu, Li Shen, Zhikun Zhang, Rongfei Fan, Le Sun, Qingming Li</p>	46914
<p>DeciMamba: Exploring the Length Extrapolation Potential of Mamba</p> <p>Assaf Ben-Kish, Itamar Zimerman, Shady Abu-Hussein, Nadav Cohen, Amir Globerson, Lior Wolf, Raja Giryes</p>	46932
<p>Building Interactable Replicas of Complex Articulated Objects Via Gaussian Splatting</p> <p>Yu Liu, Baoxiong Jia, Ruijie Lu, Junfeng Ni, Song-Chun Zhu, Siyuan Huang</p>	46955
<p>Selective Aggregation for Low-Rank Adaptation in Federated Learning</p> <p>Pengxin Guo, Shuang Zeng, Yanran Wang, Huijie Fan, Feifei Wang, Liangqiong Qu</p>	46979
<p>REMEDY: Recipe Merging Dynamics in Large Vision-Language Models</p> <p>Didi Zhu, Yibing Song, Tao Shen, Ziyu Zhao, Jinluan Yang, Min Zhang, Chao Wu</p>	47004
<p>End-To-End Learning of Gaussian Mixture Priors for Diffusion Sampler</p> <p>Denis Blessing, Xiaogang Jia, Gerhard Neumann</p>	47023
<p>Mind the GAP: Glimpse-Based Active Perception Improves Generalization and Sample Efficiency of Visual Reasoning</p> <p>Oleh Kolner, Thomas Ortner, Stanislaw Wozniak, Angeliki Pantazi</p>	47050

<p>SePer: Measure Retrieval Utility Through the Lens of Semantic Perplexity Reduction</p> <p><i>Lu Dai, Yijie Xu, Jinhui Ye, Hao Liu, Hui Xiong</i></p>	47073
<p>A Simple Yet Effective $\Delta\Delta$ Predictor is an Unsupervised Antibody Optimizer and Explainer</p> <p><i>Lirong Wu, Yunfan Liu, Haitao Lin, Yufei Huang, Guojiang Zhao, Zhifeng Gao, Stan Z. Li</i></p>	47104
<p>Learning Color Equivariant Representations</p> <p><i>Yulong Yang, Felix O'Mahony, Christine Allen-Blanchette</i></p>	47118
<p>CoRNStack: High-Quality Contrastive Data for Better Code Retrieval and Reranking</p> <p><i>Tarun Suresh, Revanth Gangi Reddy, Yifei Xu, Zach Nussbaum, Andriy Mulyar, Brandon Duderstadt, Heng Ji</i></p>	47145
<p>PaCA: Partial Connection Adaptation for Efficient Fine-Tuning</p> <p><i>Sunghyeon Woo, Sol Namkung, Sunwoo Lee, Inho Jeong, Beomseok Kim, Dongsuk Jeon</i></p>	47163
<p>Noisy Test-Time Adaptation in Vision-Language Models</p> <p><i>Chentao Cao, Zhun Zhong, Zhanke Zhou, Tongliang Liu, Yang Liu, Kun Zhang, Bo Han</i></p>	47183
<p>Measuring and Enhancing Trustworthiness of LLMs in RAG Through Grounded Attributions and Learning to Refuse</p> <p><i>Maojia Song, Shang Hong Sim, Rishabh Bhardwaj, Hai Leong Chieu, Navonil Majumder, Soujanya Poria</i></p>	47215
<p>What Makes a Maze Look Like a Maze?</p> <p><i>Joy Hsu, Jiayuan Mao, Joshua B. Tenenbaum, Noah Goodman, Jiajun Wu</i></p>	47256
<p>What's New in My Data? Novelty Exploration Via Contrastive Generation</p> <p><i>Masaru Isonuma, Ivan Titov</i></p>	47278
<p>Language Models Scale Reliably with Over-Training and on Downstream Tasks</p> <p><i>Samir Yitzhak Gadre, Georgios Smyrnis, Vaishaal Shankar, Suchin Gururangan, Mitchell Wortsman, Rulin Shao, Jean Mercat, Alex Fang, Jeffrey Li, Sedrick Keh, Rui Xin, Marianna Nezhurina, Igor Vasiljevic, Luca Soldaini, Jenia Jitsev, Alex Dimakis, Gabriel Ilharco, Pang Wei Koh, Shuran Song, Thomas Kollar, Yair Carmon, Achal Dave, Reinhard Heckel, Niklas Muennighoff, Ludwig Schmidt</i></p>	47297
<p>WeatherGFM: Learning a Weather Generalist Foundation Model Via In-Context Learning</p> <p><i>Xiangyu Zhao, Zhiwang Zhou, Zhangwenlong, Yihao Liu, Xiangyu Chen, Junchao Gong, Hao Chen, Ben Fei, Shiqi Chen, Wanli Ouyang, Xiao-Ming Wu, Lei Bai</i></p>	47319
<p>Uncertainty Modeling for Fine-Tuned Implicit Functions</p> <p><i>Anna Susmelj, Mael Macuglia, Natasa Tagasovska, Reto Sutter, Sebastiano Caprara, Jean-Philippe Thiran, Ender Konukoglu</i></p>	47340
<p>Learning to Steer Markovian Agents Under Model Uncertainty</p> <p><i>Jiawei Huang, Vinzenz Thoma, Zebang Shen, Heinrich H. Nax, Niao He</i></p>	47363

<p>GenARM: Reward Guided Generation with Autoregressive Reward Model for Test-Time Alignment</p> <p>Yuancheng Xu, Udari Madhushani Sehwag, Alec Koppel, Sicheng Zhu, Bang An, Furong Huang, Sumitra Ganesh</p>	47398
<p>ToolDial: Multi-Turn Dialogue Generation Method for Tool-Augmented Language Models</p> <p>Jeonghoon Shim, Gyuhyeon Seo, Cheongsu Lim, Yohan Jo</p>	47419
<p>Scaling Diffusion Language Models Via Adaptation from Autoregressive Models</p> <p>Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An, Peilin Zhao, Wei Bi, Jiawei Han, Hao Peng, Lingpeng Kong</p>	47445
<p>MMSearch: Unveiling the Potential of Large Models as Multi-Modal Search Engines</p> <p>Dongzhi Jiang, Renrui Zhang, Ziyu Guo, Yanmin Wu, Jiayi Lei, Pengshuo Qiu, Pan Lu, Zehui Chen, Guanglu Song, Peng Gao, Yu Liu, Chunyuan Li, Hongsheng Li</p>	47473
<p>Temporal Difference Learning: Why it Can Be Fast and How it Will Be Faster</p> <p>Patrick Schnell, Luca Guastoni, Nils Thuerey</p>	47511
<p>RECAST: Reparameterized, Compact Weight Adaptation for Sequential Tasks</p> <p>Nazia Tasnim, Bryan A. Plummer</p>	47535
<p>Learning Dynamics of Deep Matrix Factorization Beyond the Edge of Stability</p> <p>Avrajit Ghosh, Soo Min Kwon, Rongrong Wang, Saiprasad Ravishankar, Qing Qu</p>	47556
<p>Optimizing Importance Weighting in the Presence of Sub-Population Shifts</p> <p>Floris Holstegel, Bram Wouters, Noud Van Giersbergen, Cees Diks</p>	47604
<p>PerturboLLaVA: Reducing Multimodal Hallucinations with Perturbative Visual Training</p> <p>Cong Chen, Mingyu Liu, Chenchen Jing, Yizhou Zhou, Fengyun Rao, Hao Chen, Bo Zhang, Chunhua Shen</p>	47632
<p>An Empirical Analysis of Uncertainty in Large Language Model Evaluations</p> <p>Qiuqie Xie, Qingqiu Li, Zhuohao Yu, Yuejie Zhang, Yue Zhang, Linyi Yang</p>	47668
<p>LiNeS: Post-Training Layer Scaling Prevents Forgetting and Enhances Model Merging</p> <p>Ke Wang, Nikolaos Dimitriadis, Alessandro Favero, Guillermo Ortiz-Jimenez, François Fleuret, Pascal Frossard</p>	47700
<p>RetroInText: A Multimodal Large Language Model Enhanced Framework for Retrosynthetic Planning Via In-Context Representation Learning</p> <p>Chenglong Kang, Xiaoyi Liu, Fei Guo</p>	47734
<p>Merging LoRAs Like Playing LEGO: Pushing the Modularity of LoRA to Extremes Through Rank-Wise Clustering</p> <p>Ziyu Zhao, Tao Shen, Didi Zhu, Zexi Li, Jing Su, Xuwu Wang, Fei Wu</p>	47752

<p>Rethinking the Generalization of Drug Target Affinity Prediction Algorithms Via Similarity Aware Evaluation</p> <p>Chenbin Zhang, Zhiqiang Hu, Jiang Chuchu, Wen Chen, Jie Xu, Shaoting Zhang</p>	47770
<p>OpenVid-1M: A Large-Scale High-Quality Dataset for Text-To-Video Generation</p> <p>Kepan Nan, Rui Xie, Penghao Zhou, Tiehan Fan, Zhenheng Yang, Zhijie Chen, Xiang Li, Jian Yang, Ying Tai</p>	47797
<p>Exploiting Hidden Symmetry to Improve Objective Perturbation for DP Linear Learners with a Nonsmooth L1-Norm</p> <p>Du Chen, Geoffrey A. Chua</p>	47817
<p>Minimalistic Predictions for Online Class Constraint Scheduling</p> <p>Dorian Guyot, Alexandra Anna Lassota</p>	47847
<p>Diffusion-Based Neural Network Weights Generation</p> <p>Bedionita Soro, Bruno Andreis, Hayeon Lee, Wonyong Jeong, Song Chong, Frank Hutter, Sung Ju Hwang</p>	47860
<p>PostEdit: Posterior Sampling for Efficient Zero-Shot Image Editing</p> <p>Feng Tian, Yixuan Li, Yichao Yan, Shanyan Guan, Yanhao Ge, Xiaokang Yang</p>	47892
<p>Attribute-Based Visual Reprogramming for Vision-Language Models</p> <p>Chengyi Cai, Zesheng Ye, Lei Feng, Jianzhong Qi, Feng Liu</p>	47923
<p>Words in Motion: Extracting Interpretable Control Vectors for Motion Transformers</p> <p>Omer Sahin Tas, Royden Wagner</p>	47950
<p>Discretization-Invariance? on the Discretization Mismatch Errors in Neural Operators</p> <p>Wenhan Gao, Ruichen Xu, Yuefan Deng, Yi Liu</p>	47974
<p>Systems with Switching Causal Relations: A Meta-Causal Perspective</p> <p>Moritz Willig, Tim Tobiasch, Florian Peter Busch, Jonas Seng, Devendra Singh Dhami, Kristian Kersting</p>	47999
<p>ZETA: Leveraging \mathcal{Z}-Order Curves for Efficient Top-k Attention</p> <p>Qiu hao Zeng, Jerry Huang, Peng Lu, Gezheng Xu, Boxing Chen, Charles Ling, Boyu Wang</p>	48020
<p>GeSubNet: Gene Interaction Inference for Disease Subtype Network Generation</p> <p>Ziwei Yang, Zheng Chen, Xin Liu, Rikuto Kotoge, Peng Chen, Yasuko Matsubara, Yasushi Sakurai, Jimeng Sun</p>	48045
<p>Active Task Disambiguation with LLMs</p> <p>Kasia Kobalcyk, Nicolás Astorga, Tennison Liu, Mihaela Van Der Schaar</p>	48071
<p>Streaming Video Understanding and Multi-Round Interaction with Memory-Enhanced Knowledge</p> <p>Haomiao Xiong, Zongxin Yang, Jiazuo Yu, Yunzhi Zhuge, Lu Zhang, Jiawen Zhu, Huchuan Lu</p>	48096

<p>An Exploration with Entropy Constrained 3D Gaussians for 2D Video Compression</p> <p><i>Xiang Liu, Bin Chen, Zimo Liu, Yaowei Wang, Shu-Tao Xia</i></p>	48116
<p>From Few to Many: Self-Improving Many-Shot Reasoners Through Iterative Optimization and Generation</p> <p><i>Xingchen Wan, Han Zhou, Ruoxi Sun, Sercan O Arik</i></p>	48142
<p>Competitive Fair Scheduling with Predictions</p> <p><i>Tianming Zhao, Chunqiu Xia, Xiaomin Chang, Chunhao Li, Wei Li, Albert Zomaya</i></p>	48171
<p>EFFICIENT JAILBREAK ATTACK SEQUENCES ON LARGE LANGUAGE MODELS VIA MULTI-ARMED BANDIT-BASED CONTEXT SWITCHING</p> <p><i>Aditya Ramesh, Shivam Bhardwaj, Aditya Saibewar, Manohar Kaul</i></p>	48209
<p>How Does Critical Batch Size Scale in Pre-Training?</p> <p><i>Hanlin Zhang, Depen Morwani, Nikhil Vyas, Jingfeng Wu, Difan Zou, Udaya Ghai, Dean Foster, Sham M. Kakade</i></p>	48254
<p>AssembleFlow: Rigid Flow Matching with Inertial Frames for Molecular Assembly</p> <p><i>Hongyu Guo, Yoshua Bengio, Shengchao Liu</i></p>	48281
<p>AnalogGenie: A Generative Engine for Automatic Discovery of Analog Circuit Topologies</p> <p><i>Jian Gao, Weidong Cao, Junyi Yang, Xuan Zhang</i></p>	48314
<p>The Computational Complexity of Positive Non-Clashing Teaching in Graphs</p> <p><i>Robert Galian, Liana Khazaliya, Fionn Mc Inerney, Mathis Rocton</i></p>	48333
<p>MMAD: A Comprehensive Benchmark for Multimodal Large Language Models in Industrial Anomaly Detection</p> <p><i>Xi Jiang, Jian Li, Hanqiu Deng, Yong Liu, Bin-Bin Gao, Yifeng Zhou, Jialin Li, Chengjie Wang, Feng Zheng</i></p>	48346
<p>Accelerated Training Through Iterative Gradient Propagation Along the Residual Path</p> <p><i>Erwan Fagnou, Paul Caillon, Blaise Delattre, Alexandre Allauzen</i></p>	48369
<p>Fundamental Limits of Prompt Tuning Transformers: Universality, Capacity and Efficiency</p> <p><i>Jerry Yao-Chieh Hu, Wei-Po Wang, Ammar Gilani, Chenyang Li, Zhao Song, Han Liu</i></p>	48389
<p>Dual Process Learning: Controlling Use of In-Context Vs. In-Weights Strategies with Weight Forgetting</p> <p><i>Suraj Anand, Michael A. Lepori, Jack Merullo, Ellie Pavlick</i></p>	48442
<p>Efficient Off-Policy Learning for High-Dimensional Action Spaces</p> <p><i>Fabian Otto, Philipp Becker, Vien Anh Ngo, Gerhard Neumann</i></p>	48469
<p>Instant Policy: In-Context Imitation Learning Via Graph Diffusion</p> <p><i>Vitalis Vosylius, Edward Johns</i></p>	48488

<p>BLEND: Behavior-Guided Neural Population Dynamics Modeling Via Privileged Knowledge Distillation</p> <p>Zhengrui Guo, Fangxu Zhou, Wei Wu, Qichen Sun, Lishuang Feng, Jinzhuo Wang, Hao Chen</p>	48508
<p>Autoregressive Video Generation Without Vector Quantization</p> <p>Haoge Deng, Ting Pan, Haiwen Diao, Zhengxiong Luo, Yufeng Cui, Huchuan Lu, Shiguang Shan, Yonggang Qi, Xinlong Wang</p>	48533
<p>From Decoupling to Adaptive Transformation: A Wider Optimization Space for PTQ</p> <p>Zhaojing Wen, Qiulin Zhang, Yuan Zhang, Rudan Chen, Xichao Yang, Di Xie, Jiang Zhu</p>	48549
<p>Self-Attention-Based Contextual Modulation Improves Neural System Identification</p> <p>Isaac Lin, Tianye Wang, Shang Gao, Tang Shiming, Tai Sing Lee</p>	48563
<p>VTDexManip: A Dataset and Benchmark for Visual-Tactile Pretraining and Dexterous Manipulation with Reinforcement Learning</p> <p>Qingtao Liu, Yu Cui, Zhengnan Sun, Gaofeng Li, Jiming Chen, Qi Ye</p>	48585
<p>ComLoRA: A Competitive Learning Approach for Enhancing LoRA</p> <p>Qiushi Huang, Tom Ko, Lilian Tang, Yu Zhang</p>	48611
<p>Jamba: Hybrid Transformer-Mamba Language Models</p> <p>Barak Lenz, Opher Lieber, Alan Arazi, Amir Bergman, Avshalom Manevich, Barak Peleg, Ben Aviram, Chen Almagor, Clara Fridman, Dan Padnos, Daniel Gissin, Daniel Jannai, Dor Muhlgay, Dor Zimberg, Edden M. Gerber, Elad Dolev, Eran Krakovsky, Erez Safahi, Erez Schwartz, Gal Cohen, Gal Shachaf, Haim Rozenblum, Hofit Bata, Ido Blass, Inbal Magar, Itay Dalmedigos, Jhonathan Osin, Julie Fadlon, Maria Rozman, Matan Danos, Michael Gokhman, Mor Zusman, Naama Gidron, Nir Ratner, Noam Gat, Noam Rozen, Oded Fried, Ohad Leshno, Omer Antverg, Omri Abend, Or Dagan, Orit Cohavi, Raz Alon, Ro'el Belson, Roi Cohen, Rom Gilad, Roman Glozman, Shahar Lev, Shai Shalev-Shwartz, Shaked Haim Meirum, Tal Delbari, Tal Ness, Tomer Asida, Tom Ben Gal, Tom Braude, Uriya Pumerantz, Josh Cohen, Yonatan Belinkov, Yuval Globerson, Yuval Peleg Levy, Yoav Shoham</p>	48624
<p>SimPER: A Minimalist Approach to Preference Alignment Without Hyperparameters</p> <p>Teng Xiao, Yige Yuan, Zhengyu Chen, Mingxiao Li, Shangsong Liang, Zhaochun Ren, Vasant G Honavar</p>	48650
<p>The AdEMAMix Optimizer: Better, Faster, Older</p> <p>Matteo Pagliardini, Pierre Ablin, David Grangier</p>	48670
<p>Data Mixing Laws: Optimizing Data Mixtures by Predicting Language Modeling Performance</p> <p>Jiasheng Ye, Peiju Liu, Tianxiang Sun, Jun Zhan, Yunhua Zhou, Xipeng Qiu</p>	48713
<p>On the Feature Learning in Diffusion Models</p> <p>Andi Han, Wei Huang, Yuan Cao, Difan Zou</p>	48738
<p>U-Shaped and Inverted-U Scaling Behind Emergent Abilities of Large Language Models</p> <p>Tung-Yu Wu, Melody Lo</p>	48806

<p>Sufficient Context: A New Lens on Retrieval Augmented Generation Systems</p> <p>Hailey Joren, Jianyi Zhang, Chun-Sung Ferng, Da-Cheng Juan, Ankur Taly, Cyrus Rashtchian</p>	48839
<p>On the Convergence of No-Regret Dynamics in Information Retrieval Games with Proportional Ranking Functions</p> <p>Omer Madmon, Idan Pipano, Itamar Reinman, Moshe Tennenholtz</p>	48864
<p>Towards Realistic Data Generation for Real-World Super-Resolution</p> <p>Long Peng, Wenbo Li, Renjing Pei, Jingjing Ren, Jiaqi Xu, Yang Wang, Yang Cao, Zheng-Jun Zha</p>	48888
<p>Positive-Unlabeled Diffusion Models for Preventing Sensitive Data Generation</p> <p>Hiroshi Takahashi, Tomoharu Iwata, Atsutoshi Kumagai, Yuuki Yamanaka, Tomoya Yamashita</p>	48913
<p>OmnixR: Evaluating Omni-Modality Language Models on Reasoning Across Modalities</p> <p>Lichang Chen, Hexiang Hu, Mingda Zhang, Yiwen Chen, Zifeng Wang, Yandong Li, Pranav Shyam, Tianyi Zhou, Heng Huang, Ming-Hsuan Yang, Boqing Gong</p>	48926
<p>Long-Sequence Recommendation Models Need Decoupled Embeddings</p> <p>Ningya Feng, Junwei Pan, Jialong Wu, Baixu Chen, Ximei Wang, Qianli, Xian Hu, Jie Jiang, Mingsheng Long</p>	48945
<p>Preference Optimization for Reasoning with Pseudo Feedback</p> <p>Fangkai Jiao, Geyang Guo, Xingxing Zhang, Nancy F. Chen, Shafiq Joty, Furu Wei</p>	48967
<p>SAMRefiner: Taming Segment Anything Model for Universal Mask Refinement</p> <p>Yuqi Lin, Hengjia Li, Wenqi Shao, Zheng Yang, Jun Zhao, Xiaofei He, Ping Luo, Kaipeng Zhang</p>	48995
<p>An Undetectable Watermark for Generative Image Models</p> <p>Sam Gunn, Xuandong Zhao, Dawn Song</p>	49020
<p>The OMG Dataset: An Open MetaGenomic Corpus for Mixed-Modality Genomic Language Modeling</p> <p>Andre Cornman, Jacob West-Roberts, Antonio Pedro Camargo, Simon Roux, Martin Beracochea, Milot Mirdita, Sergey Ovchinnikov, Yunha Hwang</p>	49046
<p>To Clip Or Not to Clip: The Dynamics of SGD with Gradient Clipping in High-Dimensions</p> <p>Noah Marshall, Ke Liang Xiao, Atish Agarwala, Elliot Paquette</p>	49071
<p>ℓ_0-Zero: Gradient-Based Optimization of ℓ_0-Norm Adversarial Examples</p> <p>Antonio Emanuele Cinà, Francesco Villani, Maura Pintor, Lea Schönherr, Battista Biggio, Marcello Pelillo</p>	49108
<p>Training on the Test Task Confounds Evaluation and Emergence</p> <p>Ricardo Dominguez-Olmedo, Florian E. Dorner, Moritz Hardt</p>	49121
<p>Unsupervised Meta-Learning Via In-Context Learning</p> <p>Anna Vettoruzzo, Lorenzo Braccaioli, Joaquin Vanschoren, Marlena Nowaczyk</p>	49152

<p>Quantifying Generalization Complexity for Large Language Models</p> <p>Zhenting Qi, Hongyin Luo, Xuliang Huang, Zhuokai Zhao, Yibo Jiang, Xiangjun Fan, Himabindu Lakkaraju, James R. Glass</p>	49178
<p>Hotspot-Driven Peptide Design Via Multi-Fragment Autoregressive Extension</p> <p>Jiahan Li, Tong Chen, Shitong Luo, Chaoran Cheng, Jiaqi Guan, Ruihan Guo, Sheng Wang, Ge Liu, Jian Peng, Jianzhu Ma</p>	49201
<p>Fluid: Scaling Autoregressive Text-To-Image Generative Models with Continuous Tokens</p> <p>Lijie Fan, Tianhong Li, Siyang Qin, Yuanzhen Li, Chen Sun, Michael Rubinstein, Deqing Sun, Kaiming He, Yonglong Tian</p>	49234
<p>Private Mechanism Design Via Quantile Estimation</p> <p>Yuanyuan Yang, Tao Xiao, Bhuvesh Kumar, Jamie Heather Morgenstern</p>	49248
<p>Conformalized Survival Analysis for General Right-Censored Data</p> <p>Hen Davidov, Shai Feldman, Gil Shamai, Ron Kimmel, Yaniv Romano</p>	49262
<p>Learning a Neural Solver for Parametric PDEs to Enhance Physics-Informed Methods</p> <p>Lise Le Boudec, Emmanuel De Bezenac, Louis Serrano, Ramon Daniel Regueiro-Espino, Yuan Yin, Patrick Gallinari</p>	49291
<p>Boosting Perturbed Gradient Ascent for Last-Iterate Convergence in Games</p> <p>Kenshi Abe, Mitsuki Sakamoto, Kaito Ariu, Atsushi Iwasaki</p>	49343
<p>Efficiently Democratizing Medical LLMs for 50 Languages Via a Mixture of Language Family Experts</p> <p>Guorui Zheng, Xidong Wang, Juhao Liang, Nuo Chen, Yuping Zheng, Benyou Wang</p>	49374
<p>Denoising as Adaptation: Noise-Space Domain Adaptation for Image Restoration</p> <p>Kang Liao, Zongsheng Yue, Zhouxia Wang, Chen Change Loy</p>	49403
<p>Empowering LLM Agents with Zero-Shot Optimal Decision-Making Through Q-Learning</p> <p>Jiajun Chai, Sicheng Li, Yuqian Fu, Dongbin Zhao, Yuanheng Zhu</p>	49431
<p>ToddlerDiffusion: Interactive Structured Image Generation with Cascaded Schrödinger Bridge</p> <p>Eslam Mohamed Bakr, Liangbing Zhao, Vincent Tao Hu, Matthieu Cord, Patrick Perez, Mohamed Elhoseiny</p>	49478
<p>CatVTON: Concatenation is All You Need for Virtual Try-On with Diffusion Models</p> <p>Zheng Chong, Xiao Dong, Haoxiang Li, Shiyue Zhang, Wenqing Zhang, Hanqing Zhao, Xujie Zhang, Dongmei Jiang, Xiaodan Liang</p>	49513
<p>AVHBench: A Cross-Modal Hallucination Benchmark for Audio-Visual Large Language Models</p> <p>Kim Sung-Bin, Oh Hyun-Bin, Jungmok Lee, Arda Senocak, Joon Son Chung, Tae-Hyun Oh</p>	49529
<p>Multiagent Finetuning: Self Improvement with Diverse Reasoning Chains</p> <p>Vighnesh Subramaniam, Yilun Du, Joshua B. Tenenbaum, Antonio Torralba, Shuang Li, Igor Mordatch</p>	49557

<p>Reinforcement Learning from Imperfect Corrective Actions and Proxy Rewards</p> <p>Zhaohui Jiang, Xuening Feng, Paul Weng, Yifei Zhu, Yan Song, Tianze Zhou, Yujing Hu, Tangjie Lv, Changjie Fan</p>	49580
<p>Cyclic Contrastive Knowledge Transfer for Open-Vocabulary Object Detection</p> <p>Chuhan Zhang, Chaoyang Zhu, Pingcheng Dong, Long Chen, Dong Zhang</p>	49610
<p>Indirect Gradient Matching for Adversarial Robust Distillation</p> <p>Hongsin Lee, Seungju Cho, Changick Kim</p>	49625
<p>DAMO: Decoding by Accumulating Activations Momentum for Mitigating Hallucinations in Vision-Language Models</p> <p>Kaishen Wang, Hengrui Gu, Meijun Gao, Kaixiong Zhou</p>	49647
<p>HeadMap: Locating and Enhancing Knowledge Circuits in LLMs</p> <p>Xuehao Wang, Liyuan Wang, Binghuai Lin, Yu Zhang</p>	49664
<p>Neural Phylogeny: Fine-Tuning Relationship Detection Among Neural Networks</p> <p>Runpeng Yu, Xinchao Wang</p>	49682
<p>From Models to Microtheories: Distilling a Model's Topical Knowledge for Grounded Question-Answering</p> <p>Nathaniel Weir, Bhavana Dalvi Mishra, Orion Weller, Oyvind Tafjord, Sam Hornstein, Alexander Sabol, Peter Jansen, Benjamin Van Durme, Peter Clark</p>	49706
<p>UniDrive: Towards Universal Driving Perception Across Camera Configurations</p> <p>Ye Li, Wenzhao Zheng, Xiaonan Huang, Kurt Keutzer</p>	49732
<p>Towards Calibrated Deep Clustering Network</p> <p>Yuheng Jia, Jianhong Cheng, Hui Liu, Junhui Hou</p>	49747
<p>AHA: A Vision-Language-Model for Detecting and Reasoning Over Failures in Robotic Manipulation</p> <p>Jiafei Duan, Wilbert Pumacay, Nishanth Kumar, Yi Ru Wang, Shulin Tian, Wentao Yuan, Ranjay Krishna, Dieter Fox, Ajay Mandlekar, Yijie Guo</p>	49772
<p>Simple is Effective: The Roles of Graphs and Large Language Models in Knowledge-Graph-Based Retrieval-Augmented Generation</p> <p>Mufei Li, Siqi Miao, Pan Li</p>	49797
<p>CURIE: Evaluating LLMs on Multitask Scientific Long-Context Understanding and Reasoning</p> <p>Hao Cui, Zahra Shamsi, Gowoon Cheon, Xuejian Ma, Shutong Li, Maria Tikhonovskaya, Peter Christian Norgaard, Nayantara Mudur, Martyna Beata Plomecka, Paul Raccuglia, Yasaman Bahri, Victor V. Albert, Pranesh Srinivasan, Haining Pan, Philippe Faist, Brian A Rohr, Michael J. Statt, Dan Morris, Drew Purves, Elise Kleeman, Ruth Alcantara, Matthew Abraham, Muqthar Mohammad, Ean Phing Vanlee, Chenfei Jiang, Elizabeth Dorfman, Eun-Ah Kim, Michael Brenner, Sameera S Ponda, Subhashini Venugopalan</p>	49826

<p>Personalized Representation from Personalized Generation</p> <p><i>Shobhita Sundaram, Julia Chae, Yonglong Tian, Sara Beery, Phillip Isola</i></p>	49874
<p>Trained Transformer Classifiers Generalize and Exhibit Benign Overfitting In-Context</p> <p><i>Spencer Frei, Gal Vardi</i></p>	49909
<p>Residual Deep Gaussian Processes on Manifolds</p> <p><i>Kacper Wyrwal, Andreas Krause, Viacheslav Borovitskiy</i></p>	49939
<p>TaskGalaxy: Scaling Multi-Modal Instruction Fine-Tuning with Tens of Thousands Vision Task Types</p> <p><i>Jiankang Chen, Tianke Zhang, Changyi Liu, Haojie Ding, Yaya Shi, Cheng. Feng, Huihui Xiao, Bin Wen, Fan Yang, Tingting Gao, Di Zhang</i></p>	49965
<p>BaB-ND: Long-Horizon Motion Planning with Branch-And-Bound and Neural Dynamics</p> <p><i>Keyi Shen, Jiangwei Yu, Jose Barreiros, Huan Zhang, Yunzhu Li</i></p>	50016
<p>SimBa: Simplicity Bias for Scaling Up Parameters in Deep Reinforcement Learning</p> <p><i>Hojoon Lee, Dongyoon Hwang, Donghu Kim, Hyunseung Kim, Jun Jet Tai, Kaushik Subramanian, Peter R. Wurman, Jaegul Choo, Peter Stone, Takuma Seno</i></p>	50050
<p>DUET: Decentralized Bilevel Optimization Without Lower-Level Strong Convexity</p> <p><i>Zhen Qin, Zhuqing Liu, Songtao Lu, Yingbin Liang, Jia Liu</i></p>	50083
<p>Latent Space Chain-Of-Embedding Enables Output-Free LLM Self-Evaluation</p> <p><i>Yiming Wang, Pei Zhang, Baosong Yang, Derek F. Wong, Rui Wang</i></p>	50130
<p>Multimodality Helps Few-Shot 3D Point Cloud Semantic Segmentation</p> <p><i>Zhaochong An, Guolei Sun, Yun Liu, Runjia Li, Min Wu, Ming-Ming Cheng, Ender Konukoglu, Serge Belongie</i></p>	50163
<p>Large Language Models Can Become Strong Self-Detoxifiers</p> <p><i>Ching-Yun Ko, Pin-Yu Chen, Payel Das, Youssef Mroueh, Soham Dan, Georgios Kollias, Subhjit Chaudhury, Tejaswini Pedapati, Luca Daniel</i></p>	50185
<p>The Directionality of Optimization Trajectories in Neural Networks</p> <p><i>Sidak Pal Singh, Bobby He, Thomas Hofmann, Bernhard Schölkopf</i></p>	50205
<p>Injective Flows for Star-Like Manifolds</p> <p><i>Marcello Massimo Negri, Jonathan Aellen, Volker Roth</i></p>	50266
<p>Robustness of Quantum Algorithms for Nonconvex Optimization</p> <p><i>Weiyuan Gong, Chenyi Zhang, Tongyang Li</i></p>	50288
<p>Revisiting Nearest Neighbor for Tabular Data: A Deep Tabular Baseline Two Decades Later</p> <p><i>Han-Jia Ye, Huai-Hong Yin, De-Chuan Zhan, Wei-Lun Chao</i></p>	50302

<p>Specialized Foundation Models Struggle to Beat Supervised Baselines</p> <p>Zongzhe Xu, Ritvik Gupta, Wenduo Cheng, Alexander Shen, Junhong Shen, Ameet Talwalkar, Mikhail Khodak</p>	50322
<p>Video-STaR: Self-Training Enables Video Instruction Tuning with Any Supervision</p> <p>Orr Zohar, Xiaohan Wang, Yonatan Bitton, Idan Szpektor, Serena Yeung-Levy</p>	50348
<p>Revisiting Large-Scale Non-Convex Distributionally Robust Optimization</p> <p>Qi Zhang, Yi Zhou, Simon Khan, Ashley Prater-Bennette, Lixin Shen, Shaofeng Zou</p>	50363
<p>Any-Step Dynamics Model Improves Future Predictions for Online and Offline Reinforcement Learning</p> <p>Haoxin Lin, Yu-Yan Xu, Yihao Sun, Zhilong Zhang, Yi-Chen Li, Chengxing Jia, Junyin Ye, Jiayi Zhang, Yang Yu</p>	50390
<p>Storybooth: Training-Free Multi-Subject Consistency for Improved Visual Storytelling</p> <p>Jaskirat Singh, Junshen K Chen, Jonas K Kohler, Michael F Cohen</p>	50414
<p>Learning Interleaved Image-Text Comprehension in Vision-Language Large Models</p> <p>Chenyu Zhou, Mengdan Zhang, Peixian Chen, Chaoyou Fu, Yunhang Shen, Xiewu Zheng, Xing Sun, Rongrong Ji</p>	50427
<p>Extending Mercer's Expansion to Indefinite and Asymmetric Kernels</p> <p>Sungwoo Jeong, Alex Townsend</p>	50451
<p>Identifiable Exchangeable Mechanisms for Causal Structure and Representation Learning</p> <p>Patrik Reizinger, Siyuan Guo, Ferenc Huszár, Bernhard Schölkopf, Wieland Brendel</p>	50469
<p>To Trust Or Not to Trust? Enhancing Large Language Models' Situated Faithfulness to External Contexts</p> <p>Yukun Huang, Sanxing Chen, Hongyi Cai, Bhuwan Dhingra</p>	50497
<p>Gaussian Differentially Private Human Faces Under a Face Radial Curve Representation</p> <p>Carlos J Soto, Matthew Reimherr, Aleksandra Slavkovic, Mark Shriver</p>	50527
<p>RelCon: Relative Contrastive Learning for a Motion Foundation Model for Wearable Data</p> <p>Maxwell A Xu, Jaya Narain, Gregory Darnell, Haraldur T Hallgrímsson, Hyewon Jeong, Darren Forde, Richard Andres Fineman, Karthik Jayaraman Raghuram, James Matthew Rehg, Shirley You Ren</p>	50546
<p>Efficient Dictionary Learning with Switch Sparse Autoencoders</p> <p>Anish Mudide, Joshua Engels, Eric J Michaud, Max Tegmark, Christian Schroeder De Witt</p>	50565
<p>Root Cause Analysis of Anomalies in Multivariate Time Series Through Granger Causal Discovery</p> <p>Xiao Han, Saima Absar, Lu Zhang, Shuhan Yuan</p>	50580

<p>INCLUDE: Evaluating Multilingual Language Understanding with Regional Knowledge</p> <p>Angelika Romanou, Negar Foroutan, Anna Sotnikova, Sree Harsha Nelaturu, Shivalika Singh, Rishabh Maheshwary, Micol Altomare, Zeming Chen, Mohamed A. Haggag, Snegha A, Alfonso Amayuelas, Azril Hafizi Amirudin, Danylo Boiko, Michael Chang, Jenny Chim, Gal Cohen, Aditya Kumar Dalmia, Abraham Diress, Sharad Duwal, Daniil Dzenhaliou, Daniel Fernando Erazo Florez, Fabian Farestam, Joseph Marvin Imperial, Shayekh Bin Islam, Perttu Isotalo, Maral Jabbarishiviari, Börje F. Karlsson, Eldar Khalilov, Christopher Klamm, Fajri Koto, Dominik Krzeminski, Gabriel Adriano De Melo, Syrielle Montariol, Yiyang Nan, Joel Niklaus, Jekaterina Novikova, Johan Samir Obando Ceron, Debjit Paul, Esther Ploeger, Jebish Purbey, Swati Rajwal, Selvan Sunitha Ravi, Sara Rydell, Roshan Santhosh, Drishti Sharma, Marjana Prifti Skenduli, Arshia Soltani Moakhar, Bardia Soltani Moakhar, Ayush Kumar Tarun, Azmine Toushik Wasi, Thenuka Ovin Weerasinghe, Serhan Yilmaz, Mike Zhang, Imanol Schlag, Marzieh Fadaee, Sara Hooker, Antoine Bosselut</p>	50597
<p>LLMs Can Plan Only if We Tell Them</p> <p>Bilgehan Sel, Ruoxi Jia, Ming Jin</p>	50629
<p>Scaling Large Language Model-Based Multi-Agent Collaboration</p> <p>Chen Qian, Zihao Xie, Yifei Wang, Wei Liu, Kunlun Zhu, Hanchen Xia, Yufan Dang, Zhuoyun Du, Weize Chen, Cheng Yang, Zhiyuan Liu, Maosong Sun</p>	50670
<p>Feedback Schrödinger Bridge Matching</p> <p>Panagiotis Theodoropoulos, Nikolaos Komianos, Vincent Pacelli, Guan-Horng Liu, Evangelos Theodorou</p>	50688
<p>Predictive Uncertainty Quantification for Bird's Eye View Segmentation: A Benchmark and Novel Loss Function</p> <p>Linlin Yu, Bowen Yang, Tianhao Wang, Kangshuo Li, Feng Chen</p>	50717
<p>VLAS: Vision-Language-Action Model with Speech Instructions for Customized Robot Manipulation</p> <p>Wei Zhao, Pengxiang Ding, Zhang Min, Zhefei Gong, Shuanghao Bai, Han Zhao, Donglin Wang</p>	50757
<p>MME-RealWorld: Could Your Multimodal LLM Challenge High-Resolution Real-World Scenarios that Are Difficult for Humans?</p> <p>Yifan Zhang, Huanyu Zhang, Haochen Tian, Chaoyou Fu, Shuangqing Zhang, Junfei Wu, Feng Li, Kun Wang, Qingsong Wen, Zhang Zhang, Liang Wang, Rong Jin</p>	50775
<p>MLLM as Retriever: Interactively Learning Multimodal Retrieval for Embodied Agents</p> <p>Junpeng Yue, Xinrun Xu, Börje F. Karlsson, Zongqing Lu</p>	50822
<p>Tracking the Copyright of Large Vision-Language Models Through Parameter Learning Adversarial Images</p> <p>Yubo Wang, Jianting Tang, Chaohu Liu, Linli Xu</p>	50852
<p>Prevalence of Negative Transfer in Continual Reinforcement Learning: Analyses and a Simple Baseline</p> <p>Hongjoon Ahn, Jinu Hyeon, Youngmin Oh, Bosun Hwang, Taesup Moon</p>	50872

<p>Language Imbalance Driven Rewarding for Multilingual Self-Improving</p> <p>Wen Yang, Junhong Wu, Chen Wang, Chengqing Zong, Jiajun Zhang</p>	50897
<p>LayerDAG: A Layerwise Autoregressive Diffusion Model for Directed Acyclic Graph Generation</p> <p>Mufei Li, Viraj Shitole, Eli Chien, Changhai Man, Zhaodong Wang, Srinivas, Ying Zhang, Tushar Krishna, Pan Li</p>	50923
<p>VSTAR: Generative Temporal Nursing for Longer Dynamic Video Synthesis</p> <p>Yumeng Li, William H. Beluch, Margret Keuper, Dan Zhang, Anna Khoreva</p>	50941
<p>Ranking-Aware Adapter for Text-Driven Image Ordering with CLIP</p> <p>Wei-Hsiang Yu, Yen-Yu Lin, Ming-Hsuan Yang, Yi-Hsuan Tsai</p>	50965
<p>Neural Wave Equation for Irregularly Sampled Sequence Data</p> <p>Arkaprava Majumdar, M Anand Krishna, P. K. Srijith</p>	50985
<p>Learning and Aligning Single-Neuron Invariance Manifolds in Visual Cortex</p> <p>Mohammad Bashiri, Luca Baroni, Ján Antolík, Fabian H. Sinz</p>	51007
<p>No Equations Needed: Learning System Dynamics Without Relying on Closed-Form ODEs</p> <p>Krzysztof Kacprzyk, Mihaela Van Der Schaar</p>	51033
<p>Re-Evaluating Open-Ended Evaluation of Large Language Models</p> <p>Siqi Liu, Ian Gemp, Luke Marris, Georgios Piliouras, Nicolas Heess, Marc Lanctot</p>	51077
<p>Divergence-Regularized Discounted Aggregation: Equilibrium Finding in Multiplayer Partially Observable Stochastic Games</p> <p>Runyu Lu, Yuanheng Zhu, Dongbin Zhao</p>	51105
<p>ReAttention: Training-Free Infinite Context with Finite Attention Scope</p> <p>Xiaoran Liu, Ruixiao Li, Zhigeng Liu, Qipeng Guo, Yuerong Song, Kai Lv, Hang Yan, Linlin Li, Qun Liu, Xipeng Qiu</p>	51137
<p>A Non-Contrastive Learning Framework for Sequential Recommendation with Preference-Preserving Profile Generation</p> <p>Huimin Zeng, Xiaojie Wang, Anoop Jain, Zhicheng Dou, Dong Wang</p>	51158
<p>Locally Connected Echo State Networks for Time Series Forecasting</p> <p>Filip Matzner, František Mráz</p>	51172
<p>Adaptive Deployment of Untrusted LLMs Reduces Distributed Threats</p> <p>Jiaxin Wen, Vivek Hebbar, Caleb Larson, Aryan Bhatt, Ansh Radhakrishnan, Mrinank Sharma, Henry Sleight, Shi Feng, He He, Ethan Perez, Buck Shlegeris, Akbir Khan</p>	51194
<p>From GNNs to Trees: Multi-Granular Interpretability for Graph Neural Networks</p> <p>Jie Yang, Yuwen Wang, Kaixuan Chen, Tongya Zheng, Yihe Zhou, Zhenbang Xiao, Ji Cao, Mingli Song, Shunyu Liu</p>	51238

Inverse Rendering Using Multi-Bounce Path Tracing and Reservoir Sampling Yuxin Dai, Qi Wang, Jingsen Zhu, Dianbing Xi, Yuchi Huo, Chen Qian, Ying He	51262
Accelerating Task Generalisation with Multi-Level Skill Hierarchies Thomas P Cannon, Özgür Simsek	51284
MetaUrban: An Embodied AI Simulation Platform for Urban Micromobility Wayne Wu, Honglin He, Jack He, Yiran Wang, Chenda Duan, Zhizheng Liu, Quanyi Li, Bolei Zhou	51327
Toward Guidance-Free AR Visual Generation Via Condition Contrastive Alignment Huayu Chen, Hang Su, Peize Sun, Jun Zhu	51374
MeshAnything: Artist-Created Mesh Generation with Autoregressive Transformers Yiwen Chen, Tong He, Di Huang, Weicai Ye, Sijin Chen, Jiayang Tang, Zhongang Cai, Lei Yang, Gang Yu, Guosheng Lin, Chi Zhang	51394
EVA: Geometric Inverse Design for Fast Protein Motif-Scaffolding with Coupled Flow Yufei Huang, Yunshu Liu, Lirong Wu, Haitao Lin, Cheng Tan, Odin Zhang, Zhangyang Gao, Siyuan Li, Zicheng Liu, Yunfan Liu, Tailin Wu, Stan Z. Li	51415
Boosting the Visual Interpretability of CLIP Via Adversarial Fine-Tuning Shizhan Gong, Haoyu Lei, Qi Dou, Farzan Farnia	51435
Training-Free Camera Control for Video Generation Chen Hou, Zhibo Chen	51464
MambaQuant: Quantizing the Mamba Family with Variance Aligned Rotation Methods Zukang Xu, Yuxuan Yue, Xing Hu, Dawei Yang, Zhihang Yuan, Zixu Jiang, Zhixuan Chen, Jiangyongyu, Xuchen, Sifan Zhou	51481
BenTo: Benchmark Reduction with In-Context Transferability Hongyu Zhao, Ming Li, Lichao Sun, Tianyi Zhou	51501
MiniCTX: Neural Theorem Proving with (Long-)Contexts Jiewen Hu, Thomas Zhu, Sean Welleck	51518
Simple Yet Effective Incomplete Multi-View Clustering: Similarity-Level Imputation and Intra-View Hybrid-Group Prototype Construction Shengju Yu, Zhibin Dong, Siwei Wang, Pei Zhang, Yi Zhang, Xinwang Liu, Naiyang Guan, Tiejun Li, Yiu-Ming Cheung	51543
A Differentiable Rank-Based Objective for Better Feature Learning Krunoslav Lehman Pavasovic, Giulio Biroli, Levent Sagun	51570
Reconsidering Faithfulness in Regular, Self-Explainable and Domain Invariant GNNs Steve Azzolin, Antonio Longa, Stefano Teso, Andrea Passerini	51601
Learning Long Range Dependencies on Graphs Via Random Walks Dexiong Chen, Till Hendrik Schulz, Karsten Borgwardt	51635

<p>Enhancing Prediction Performance Through Influence Measure</p> <p>Shuguang Yu, Wenqian Xu, Xinyi Zhou, Xuechun Wang, Hongtu Zhu, Fan Zhou</p>	51662
<p>Convergent Privacy Loss of Noisy-SGD Without Convexity and Smoothness</p> <p>Eli Chien, Pan Li</p>	51686
<p>ShortcutsBench: A Large-Scale Real-World Benchmark for API-Based Agents</p> <p>Haiyang Shen, Yue Li, Desong Meng, Dongqi Cai, Sheng Qi, Li Zhang, Mengwei Xu, Yun Ma</p>	51709
<p>Improving Semantic Understanding in Speech Language Models Via Brain-Tuning</p> <p>Omer Moussa, Dietrich Klakow, Mariya Toneva</p>	51735
<p>CtD: Composition Through Decomposition in Emergent Communication</p> <p>Boaz Carmeli, Ron Meir, Yonatan Belinkov</p>	51762
<p>Manifolds, Random Matrices and Spectral Gaps: The Geometric Phases of Generative Diffusion</p> <p>Enrico Ventura, Beatrice Achilli, Gianluigi Silvestri, Carlo Lucibello, Luca Ambrogioni</p>	51796
<p>MA²E: Addressing Partial Observability in Multi-Agent Reinforcement Learning with Masked Auto-Encoder</p> <p>Sehyeok Kang, Yongsik Lee, Gahee Kim, Song Chong, Se-Young Yun</p>	51818
<p>Data-Centric Prediction Explanation Via Kernelized Stein Discrepancy</p> <p>Mahtab Sarvmaili, Hassan Sajjad, Ga Wu</p>	51839
<p>Shallow Diffusion Networks Provably Learn Hidden Low-Dimensional Structure</p> <p>Nicholas Matthew Boffi, Arthur Jacot, Stephen Tu, Ingvar Ziemann</p>	51857
<p>Timer-XL: Long-Context Transformers for Unified Time Series Forecasting</p> <p>Yong Liu, Guo Qin, Xiangdong Huang, Jianmin Wang, Mingsheng Long</p>	51892
<p>RMB: Comprehensively Benchmarking Reward Models in LLM Alignment</p> <p>Enyu Zhou, Guodong Zheng, Binghai Wang, Zhiheng Xi, Shihan Dou, Rong Bao, Wei Shen, Limao Xiong, Jessica Fan, Yurong Mou, Rui Zheng, Tao Gui, Qi Zhang, Xuanjing Huang</p>	51917
<p>Learning to Plan Before Answering: Self-Teaching LLMs to Learn Abstract Plans for Problem Solving</p> <p>Jin Zhang, Flood Sung, Zhilin Yang, Yang Gao, Chongjie Zhang</p>	51964
<p>Detecting Backdoor Samples in Contrastive Language Image Pretraining</p> <p>Hanxun Huang, Sarah Monazam Erfani, Yige Li, Xingjun Ma, James Bailey</p>	51980
<p>Unveiling the Magic of Code Reasoning Through Hypothesis Decomposition and Amendment</p> <p>Yuze Zhao, Tianyun Ji, Wenjun Feng, Zhenya Huang, Qi Liu, Zhiding Liu, Yixiao Ma, Kai Zhang, Enhong Chen</p>	52008
<p>PolaFormer: Polarity-Aware Linear Attention for Vision Transformers</p> <p>Weikang Meng, Yadan Luo, Xin Li, Dongmei Jiang, Zheng Zhang</p>	52032

<p>Unearthing Skill-Level Insights for Understanding Trade-Offs of Foundation Models</p> <p>Mazda Moayeri, Vidhisha Balachandran, Varun Chandrasekaran, Safoora Yousefi, Thomas Fel, Soheil Feizi, Besmira Nushi, Neel Joshi, Vibhav Vineet</p>	52050
<p>A Multi-Power Law for Loss Curve Prediction Across Learning Rate Schedules</p> <p>Kairong Luo, Haodong Wen, Shengding Hu, Zhenbo Sun, Zhiyuan Liu, Maosong Sun, Kaifeng Lyu, Wenguang Chen</p>	52078
<p>Semantic Loss Guided Data Efficient Supervised Fine Tuning for Safe Responses in LLMs</p> <p>Yuxiao Lu, Arunesh Sinha, Pradeep Varakantham</p>	52123
<p>Air Quality Prediction with Physics-Guided Dual Neural ODEs in Open Systems</p> <p>Jindong Tian, Yuxuan Liang, Ronghui Xu, Peng Chen, Chenjuan Guo, Aoying Zhou, Lujia Pan, Zhongwen Rao, Bin Yang</p>	52136
<p>DeeperForward: Enhanced Forward-Forward Training for Deeper and Better Performance</p> <p>Liang Sun, Yang Zhang, Weizhao He, Jiajun Wen, Linlin Shen, Weicheng Xie</p>	52159
<p>CHAMP: Conformalized 3D Human Multi-Hypothesis Pose Estimators</p> <p>Harry Zhang, Luca Carlone</p>	52183
<p>Streaming Algorithms for ℓ_p Flows and ℓ_p Regression</p> <p>Amit Chakrabarti, Jeffrey Jiang, David Woodruff, Taisuke Yasuda</p>	52208
<p>Circuit Transformer: A Transformer that Preserves Logical Equivalence</p> <p>Xihan Li, Xing Li, Lei Chen, Xing Zhang, Mingxuan Yuan, Jun Wang</p>	52227
<p>Learning Local Equivariant Representations for Quantum Operators</p> <p>Zhanghao Zhouyin, Zixi Gan, Shishir Kumar Pandey, Linfeng Zhang, Qiangqiang Gu</p>	52248
<p>LongWriter: Unleashing 10,000+ Word Generation from Long Context LLMs</p> <p>Yushi Bai, Jiajie Zhang, Xin Lv, Linzhi Zheng, Siqi Zhu, Lei Hou, Yuxiao Dong, Jie Tang, Juanzi Li</p>	52268
<p>AstroCompress: A Benchmark Dataset for Multi-Purpose Compression of Astronomical Data</p> <p>Tuan Truong, Rithwik Sudharsan, Yibo Yang, Peter Xiangyuan Ma, Ruihan Yang, Stephan Mandt, Joshua S. Bloom</p>	52287
<p>Diffusion State-Guided Projected Gradient for Inverse Problems</p> <p>Rayhan Zirvi, Bahareh Tolooshams, Anima Anandkumar</p>	52312
<p>SoundCTM: Unifying Score-Based and Consistency Models for Full-Band Text-To-Sound Generation</p> <p>Koichi Saito, Dongjun Kim, Takashi Shibuya, Chieh-Hsin Lai, Zhi Zhong, Yuhta Takida, Yuki Mitsufuji</p>	52338
<p>Can Watermarks Be Used to Detect LLM IP Infringement for Free?</p> <p>Zhengyue Zhao, Xiaogeng Liu, Somesh Jha, Patrick McDaniel, Bo Li, Chaowei Xiao</p>	52368

<p>LLMs Know More than They Show: On the Intrinsic Representation of LLM Hallucinations</p> <p>Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, Yonatan Belinkov</p>	52385
<p>Copyright-Protected Language Generation Via Adaptive Model Fusion</p> <p>Javier Abad, Konstantin Donhauser, Francesco Pinto, Fanny Yang</p>	52419
<p>CaPo: Cooperative Plan Optimization for Efficient Embodied Multi-Agent Cooperation</p> <p>Jie Liu, Pan Zhou, Yingjun Du, Ah-Hwee Tan, Cees G. M. Snoek, Jan-Jakob Sonke, Efstratios Gavves</p>	52464
<p>Uncovering Latent Memories in Large Language Models</p> <p>Sunny Duan, Mikail Khona, Abhram Iyer, Rylan Schaeffer, Ila R Fiete</p>	52489
<p>ConceptPrune: Concept Editing in Diffusion Models Via Skilled Neuron Pruning</p> <p>Ruchika Chavhan, Da Li, Timothy Hospedales</p>	52512
<p>Beyond Single Neurons: Population Response Geometry in Digital Twins of Mouse Visual Cortex</p> <p>Dario Liscai, Emanuele Luconi, Alessandro Marin Vargas, Alessandro Sanzeni</p>	52536
<p>ShEPHERD: Diffusing Shape, Electrostatics, and Pharmacophores for Bioisosteric Drug Design</p> <p>Keir Adams, Kento Abeywardane, Jenna Fromer, Connor W. Coley</p>	52565
<p>Open-Set Graph Anomaly Detection Via Normal Structure Regularisation</p> <p>Qizhou Wang, Guansong Pang, Mahsa Salehi, Xiaokun Xia, Christopher Leckie</p>	52628
<p>Federated Granger Causality Learning for Interdependent Clients with State Space Representation</p> <p>Ayush Mohanty, Nazal Mohamed, Paritosh Ramanan, Nagi Gebraeel</p>	52652
<p>StringLLM: Understanding the String Processing Capability of Large Language Models</p> <p>Xilong Wang, Hao Fu, Jindong Wang, Neil Zhenqiang Gong</p>	52681
<p>Safety Layers in Aligned Large Language Models: The Key to LLM Security</p> <p>Shen Li, Liuyi Yao, Lan Zhang, Yaliang Li</p>	52697
<p>ChemAgent: Self-Updating Memories in Large Language Models Improves Chemical Reasoning</p> <p>Xiangru Tang, Tianyu Hu, Muyang Ye, Yanjun Shao, Xunjian Yin, Siru Ouyang, Wangchunshu Zhou, Pan Lu, Zhuosheng Zhang, Yilun Zhao, Arman Cohan, Mark Gerstein</p>	52724
<p>Diffusion Attribution Score: Evaluating Training Data Influence in Diffusion Models</p> <p>Jinxu Lin, Linwei Tao, Minjing Dong, Chang Xu</p>	52761
<p>Addressing Label Shift in Distributed Learning Via Entropy Regularization</p> <p>Zhiyuan Wu, Changkyu Choi, Xiangcheng Cao, Volkan Cevher, Ali Ramezani-Kebrya</p>	52789
<p>Sketch2Diagram: Generating Vector Diagrams from Hand-Drawn Sketches</p> <p>Itsumi Saito, Haruto Yoshida, Keisuke Sakaguchi</p>	52825

<p>Agent Skill Acquisition for Large Language Models Via CycleQD</p> <p><i>So Kuroki, Taishi Nakamura, Takuya Akiba, Yujin Tang</i></p>	52848
<p>Transformer Block Coupling and Its Correlation with Generalization in LLMs</p> <p><i>Murdock Aubry, Haoming Meng, Anton Sugolov, Vardan Papyan</i></p>	52870
<p>Node-Time Conditional Prompt Learning in Dynamic Graphs</p> <p><i>Xingtong Yu, Zhenghao Liu, Xinming Zhang, Yuan Fang</i></p>	52905
<p>Ward: Provable RAG Dataset Inference Via LLM Watermarks</p> <p><i>Nikola Jovanovic, Robin Staab, Maximilian Baader, Martin Vechev</i></p>	52925
<p>SMITE: Segment Me in Time</p> <p><i>Amirhossein Alimohammadi, Sauradip Nag, Saeid Asgari, Andrea Tagliasacchi, Ghassan Hamarneh, Ali Mahdavi Amiri</i></p>	52952
<p>To Tackle Adversarial Transferability: A Novel Ensemble Training Method with Fourier Transformation</p> <p><i>Wanlin Zhang, Weichen Lin, Ruomin Huang, Shihong Song, Hu Ding</i></p>	52978
<p>Provably Safeguarding a Classifier from OOD and Adversarial Samples</p> <p><i>Nicolas Atienza, Johanne Cohen, Christophe Labreuche, Michele Sebag</i></p>	53001
<p>Computing Circuits Optimization Via Model-Based Circuit Genetic Evolution</p> <p><i>Zhihai Wang, Jie Wang, Xilin Xia, Dongsheng Zuo, Lei Chen, Yuzhe Ma, Jianye Hao, Mingxuan Yuan, Feng Wu</i></p>	53022
<p>An Asynchronous Bundle Method for Distributed Learning Problems</p> <p><i>Daniel Cederberg, Xuyang Wu, Stephen P. Boyd, Mikael Johansson</i></p>	53054
<p>OmniCorpus: A Unified Multimodal Corpus of 10 Billion-Level Images Interleaved with Text</p> <p><i>Qingyun Li, Zhe Chen, Weiyun Wang, Wenhai Wang, Shenglong Ye, Zhenjiang Jin, Guanzhou Chen, Yanan He, Zhangwei Gao, Erfei Cui, Jiashuo Yu, Hao Tian, Jiasheng Zhou, Chao Xu, Bin Wang, Xingjian Wei, Wei Li, Wenjian Zhang, Bo Zhang, Pinlong Cai, Licheng Wen, Xiangchao Yan, Pei Chu, Yi Wang, Min Dou, Changyao Tian, Xizhou Zhu, Lewei Lu, Yushi Chen, Junjun He, Tong Lu, Yali Wang, Limin Wang, Dahua Lin, Yu Qiao, Botian Shi, Conghui He, Jifeng Dai</i></p>	53080
<p>Diffusing States and Matching Scores: A New Framework for Imitation Learning</p> <p><i>Runzhe Wu, Yiding Chen, Gokul Swamy, Kianté Brantley, Wen Sun</i></p>	53123
<p>Dobi-SVD: Differentiable SVD for LLM Compression and Some New Perspectives</p> <p><i>Wang Qinsi, Jinghan Ke, Masayoshi Tomizuka, Kurt Keutzer, Chenfeng Xu</i></p>	53149
<p>Standard Gaussian Process is All You Need for High-Dimensional Bayesian Optimization</p> <p><i>Zhitong Xu, Haitao Wang, Jeff M. Phillips, Shandian Zhe</i></p>	53179
<p>Finally Rank-Breaking Conquers MNL Bandits: Optimal and Efficient Algorithms for MNL Assortment</p> <p><i>Aadirupa Saha, Pierre Gaillard</i></p>	53200

<p>INS: Interaction-Aware Synthesis to Enhance Offline Multi-Agent Reinforcement Learning</p> <p>Yuqian Fu, Yuanheng Zhu, Jian Zhao, Jiajun Chai, Dongbin Zhao</p>	53222
<p>YOLO-RD: Introducing Relevant and Compact Explicit Knowledge to YOLO by Retriever-Dictionary</p> <p>Hao-Tang Tsui, Chien-Yao Wang, Hong-Yuan Mark Liao</p>	53247
<p>Reti-Diff: Illumination Degradation Image Restoration with Retinex-Based Latent Diffusion Model</p> <p>Chunming He, Chengyu Fang, Yulun Zhang, Longxiang Tang, Jinfa Huang, Kai Li, Zhenhua Guo, Xiu Li, Sina Farsiu</p>	53270
<p>Navigating the Digital World as Humans Do: Universal Visual Grounding for GUI Agents</p> <p>Boyu Gou, Ruohan Wang, Boyuan Zheng, Yanan Xie, Cheng Chang, Yiheng Shu, Huan Sun, Yu Su</p>	53291
<p>Adapters for Altering LLM Vocabularies: What Languages Benefit the Most?</p> <p>Hyojung Han, Akiko Eriguchi, Haoran Xu, Hieu Hoang, Marine Carpuat, Huda Khayrallah</p>	53324
<p>Distilling Structural Representations into Protein Sequence Models</p> <p>Jeffrey Ouyang-Zhang, Chengyue Gong, Yue Zhao, Philipp Kraehenbuehl, Adam Klivans, Daniel Jesus Diaz</p>	53351
<p>InstaSHAP: Interpretable Additive Models Explain Shapley Values Instantly</p> <p>James Enouen, Yan Liu</p>	53364
<p>Collapsed Language Models Promote Fairness</p> <p>Jingxuan Xu, Wuyang Chen, Linyi Li, Yao Zhao, Yunchao Wei</p>	53410
<p>GALA: Geometry-Aware Local Adaptive Grids for Detailed 3D Generation</p> <p>Dingdong Yang, Yizhi Wang, Konrad Schindler, Ali Mahdavi Amiri, Hao Zhang</p>	53431
<p>PaPaGei: Open Foundation Models for Optical Physiological Signals</p> <p>Arvind Pillai, Dimitris Spathis, Fahim Kawsar, Mohammad Malekzadeh</p>	53455
<p>One for All and All for One: Efficient Computation of Partial Wasserstein Distances on the Line</p> <p>Laetitia Chapel, Romain Tavenard</p>	53487
<p>Distribution-Specific Agnostic Conditional Classification with Halfspaces</p> <p>Jizhou Huang, Brendan Juba</p>	53510
<p>PaRa: Personalizing Text-To-Image Diffusion Via Parameter Rank Reduction</p> <p>Shangyu Chen, Zizheng Pan, Jianfei Cai, Dinh Phung</p>	53523
<p>Selective Unlearning Via Representation Erasure Using Domain Adversarial Training</p> <p>Nazanin Mohammadi Sepahvand, Eleni Triantafillou, Hugo Larochelle, Doina Precup, James J. Clark, Daniel M. Roy, Gintare Karolina Dziugaite</p>	53548
<p>Cauchy-Schwarz Regularizers</p> <p>Sueda Taner, Ziyi Wang, Christoph Studer</p>	53567

<p>Slot-Guided Adaptation of Pre-Trained Diffusion Models for Object-Centric Learning and Compositional Generation</p> <p><i>Kaan Akan, Yu cel Yemez</i></p>	53594
<p>SysBench: Can LLMs Follow System Message?</p> <p><i>Yanzhao Qin, Tao Zhang, Tao Zhang, Yanjun Shen, Wenjing Luo, Sunhaoze, Yan Zhang, Yujing Qiao, Weipeng Chen, Zenan Zhou, Wentao Zhang, Bin Cui</i></p>	53608
<p>Finding Shared Decodable Concepts and Their Negations in the Brain</p> <p><i>Cory Daniel Efird, Alex Murphy, Joel Zylberberg, Alona Fyshe</i></p>	53633
<p>Projection Head is Secretly an Information Bottleneck</p> <p><i>Zhuo Ouyang, Kaiwen Hu, Qi Zhang, Yifei Wang, Yisen Wang</i></p>	53653
<p>Semi-Parametric Retrieval Via Binary Bag-Of-Tokens Index</p> <p><i>Jiawei Zhou, Li Dong, Furu Wei, Lei Chen</i></p>	53673
<p>Ensembles of Low-Rank Expert Adapters</p> <p><i>Yinghao Li, Vianne R. Gao, Chao Zhang, Mohamadali Torkamani</i></p>	53692
<p>Prompt as Knowledge Bank: Boost Vision-Language Model Via Structural Representation for Zero-Shot Medical Detection</p> <p><i>Yuguang Yang, Tongfei Chen, Haoyu Huang, Linlin Yang, Chunyu Xie, Dawei Leng, Xianbin Cao, Baochang Zhang</i></p>	53721
<p>Taming Overconfidence in LLMs: Reward Calibration in RLHF</p> <p><i>Jixuan Leng, Chengsong Huang, Banghua Zhu, Ji axin Huang</i></p>	53744
<p>FlashRNN: I/O-Aware Optimization of Traditional RNNs on Modern Hardware</p> <p><i>Korbinian Pöppel, Maximilian Beck, Sepp Hochreiter</i></p>	53778
<p>Robust Root Cause Diagnosis Using In-Distribution Interventions</p> <p><i>Lokesh Nagalapatti, Ashutosh Srivastava, Sunita Sarawagi, Amit Sharma</i></p>	53803
<p>TabReD: Analyzing Pitfalls and Filling the Gaps in Tabular Deep Learning Benchmarks</p> <p><i>Ivan Rubachev, Nikolay Kartashev, Yury Gorishniy, Artem Babenko</i></p>	53831
<p>Learning to Engineer Protein Flexibility</p> <p><i>Petr Kouba, Joan Planas-Iglesias, Jiri Damborsky, Jiri Sedlar, Stanislav Mazurenko, Josef Sivic</i></p>	53868
<p>OmniBind: Large-Scale Omni Multimodal Representation Via Binding Spaces</p> <p><i>Zehan Wang, Ziang Zhang, Minjie Hong, Hang Zhang, Luping Liu, Rongjie Huang, Xize Cheng, Shengpeng Ji, Tao Jin, Hengshuang Zhao, Zhou Zhao</i></p>	53891
<p>Controlling Language and Diffusion Models by Transporting Activations</p> <p><i>Pau Rodriguez, Arno Blaas, Michal Klein, Luca Zappella, Nicholas Apostoloff, Marco Cuturi, Xavier Suau</i></p>	53912

<p>HiBug2: Efficient and Interpretable Error Slice Discovery for Comprehensive Model Debugging</p> <p>Muxi Chen, Chenchen Zhao, Qiang Xu</p>	53956
<p>Enhancing Graph of Thought: Enhancing Prompts with LLM Rationales and Dynamic Temperature Control</p> <p>Sunguk Shin, Youngjoon Kim</p>	53981
<p>ADIFF: Explaining Audio Difference Using Natural Language</p> <p>Soham Deshmukh, Shuo Han, Rita Singh, Bhiksha Raj</p>	53999
<p>RobustKV: Defending Large Language Models Against Jailbreak Attacks Via KV Eviction</p> <p>Tanqiu Jiang, Zian Wang, Jiacheng Liang, Changjiang Li, Yuhui Wang, Ting Wang</p>	54033
<p>Discrete Latent Plans Via Semantic Skill Abstractions</p> <p>Haobin Jiang, Jiangxing Wang, Zongqing Lu</p>	54051
<p>OASIS Uncovers: High-Quality T2I Models, Same Old Stereotypes</p> <p>Sepehr Dehdashtian, Gautam Sreekumar, Vishnu Boddeti</p>	54076
<p>Online Reinforcement Learning in Non-Stationary Context-Driven Environments</p> <p>Pouya Hamadani, Arash Nasr-Esfahany, Malte Schwarzkopf, Siddhartha Sen, Mohammad Alizadeh</p>	54098
<p>DaWin: Training-Free Dynamic Weight Interpolation for Robust Adaptation</p> <p>Changdae Oh, Yixuan Li, Kyungwoo Song, Sangdoon Yun, Dongyoon Han</p>	54133
<p>CarbonSense: A Multimodal Dataset and Baseline for Carbon Flux Modelling</p> <p>Matthew Fortier, Mats Leon Richter, Oliver Sonnentag, Christopher Pal</p>	54162
<p>SPAM: Spike-Aware Adam with Momentum Reset for Stable LLM Training</p> <p>Tianjin Huang, Ziquan Zhu, Gaojie Jin, Lu Liu, Zhangyang Wang, Shiwei Liu</p>	54195
<p>QERA: An Analytical Framework for Quantization Error Reconstruction</p> <p>Cheng Zhang, Jeffrey T. H. Wong, Can Xiao, George Anthony Constantinides, Yiren Zhao</p>	54216
<p>UniGEM: A Unified Approach to Generation and Property Prediction for Molecules</p> <p>Shikun Feng, Yuyan Ni, Lu Yan, Zhi-Ming Ma, Wei-Ying Ma, Yanyan Lan</p>	54246
<p>Gaussian Mixture Counterfactual Generator</p> <p>Jong-Hoon Ahn, Akshay Vashist</p>	54272
<p>TANGO: Co-Speech Gesture Video Reenactment with Hierarchical Audio Motion Embedding and Diffusion Interpolation</p> <p>Haiyang Liu, Xingchao Yang, Tomoya Akiyama, Yuantian Huang, Qiaoge Li, Shigeru Kuriyama, Takafumi Taketomi</p>	54298
<p>Nonlinear Multiregion Neural Dynamics with Parametric Impulse Response Communication Channels</p> <p>Matthew Dowling, Cristina Savin</p>	54314

Reducing Hallucinations in Large Vision-Language Models Via Latent Space Steering Sheng Liu, Haotian Ye, James Zou	54333
Rethinking Graph Neural Networks from a Geometric Perspective of Node Features Feng Ji, Yanan Zhao, Kai Zhao, Hanyang Meng, Jielong Yang, Wee Peng Tay	54351
BitStack: Any-Size Compression of Large Language Models in Variable Memory Environments Xinghao Wang, Pengyu Wang, Bo Wang, Dong Zhang, Yunhua Zhou, Xipeng Qiu	54383
Regularizing Energy Among Training Samples for Out-Of-Distribution Generalization Yiting Chen, Qitian Wu, Junchi Yan	54406
Accessing Vision Foundation Models Via ImageNet-1K Yitian Zhang, Xu Ma, Yue Bai, Huan Wang, Yun Fu	54425
Do Vision & Language Decoders Use Images and Text Equally? How Self-Consistent Are Their Explanations? Letitia Parcalabescu, Anette Frank	54445
For Better Or for Worse? Learning Minimum Variance Features with Label Augmentation Muthu Chidambaram, Rong Ge	54475
VEDIT: Latent Prediction Architecture for Procedural Video Representation Learning Han Lin, Tushar Nagarajan, Nicolas Ballas, Mido Assran, Mojtaba Komeili, Mohit Bansal, Koustuv Sinha	54500
Diffusion Actor-Critic: Formulating Constrained Policy Iteration as Diffusion Noise Regression for Offline Reinforcement Learning Linjiajie Fang, Ruoxue Liu, Jing Zhang, Wenjia Wang, Bingyi Jing	54521
CG-Bench: Clue-Grounded Question Answering Benchmark for Long Video Understanding Guo Chen, Yicheng Liu, Yifei Huang, Baoqi Pei, Jilan Xu, Yuping He, Tong Lu, Yali Wang, Limin Wang	54542
Adaptive Q\$-Network: On-The-Fly Target Selection for Deep Reinforcement Learning Théo Vincent, Fabian Wahren, Jan Peters, Boris Belousov, Carlo D'Eramo	54578
Computational Limits of Low-Rank Adaptation (LoRA) Fine-Tuning for Transformer Models Jerry Yao-Chieh Hu, Maojiang Su, En-Jui Kuo, Zhao Song, Han Liu	54606
Policy Optimization Under Imperfect Human Interactions with Agent-Gated Shared Autonomy Zhenghai Xue, Bo An, Shuicheng Yan	54646
Diff-Prompt: Diffusion-Driven Prompt Generator with Mask Supervision Weicai Yan, Wang Lin, Zirun Guo, Ye Wang, Fangming Feng, Xiaoda Yang, Zehan Wang, Tao Jin	54666
IDInit: A Universal and Stable Initialization Method for Neural Network Training Yu Pan, Chaozheng Wang, Zekai Wu, Qifan Wang, Min Zhang, Zenglin Xu	54688

<p>ForecastBench: A Dynamic Benchmark of AI Forecasting Capabilities</p> <p><i>Ezra Karger, Houtan Bastani, Chen Yueh-Han, Zachary Jacobs, Danny Halawi, Fred Zhang, Philip Tetlock</i></p>	54716
<p>Discrete Codebook World Models for Continuous Control</p> <p><i>Aidan Scannell, Mohammadreza Nakhaeinezhadfar, Kalle Kujanpää, Yi Zhao, Kevin Sebastian Luck, Arno Solin, Joni Pajarinen</i></p>	54754
<p>DiSK: Differentially Private Optimizer with Simplified Kalman Filter for Noise Reduction</p> <p><i>Xinwei Zhang, Zhiqi Bu, Borja Balle, Mingyi Hong, Meisam Razaviyayn, Vahab Mirrokni</i></p>	54792
<p>RFMamba: Frequency-Aware State Space Model for RF-Based Human-Centric Perception</p> <p><i>Rui Zhang, Ruixu Geng, Yadong Li, Ruiyuan Song, Hanqin Gong, Dongheng Zhang, Yang Hu, Yan Chen</i></p>	54820
<p>Can LLMs Understand Time Series Anomalies?</p> <p><i>Zihao Zhou, Rose Yu</i></p>	54837
<p>NV-Embed: Improved Techniques for Training LLMs as Generalist Embedding Models</p> <p><i>Chankyu Lee, Rajarshi Roy, Mengyao Xu, Jonathan Raiman, Mohammad Shoeybi, Bryan Catanzaro, Wei Ping</i></p>	54876
<p>MambaExtend: A Training-Free Approach to Improve Long Context Extension of Mamba</p> <p><i>Seyedarmin Azizi, Souvik Kundu, Mohammad Erfan Sadeghi, Massoud Pedram</i></p>	54900
<p>Enhancing Cognition and Explainability of Multimodal Foundation Models with Self-Synthesized Data</p> <p><i>Yucheng Shi, Quanzheng Li, Jin Sun, Xiang Li, Ninghao Liu</i></p>	54916
<p>Catastrophic Failure of LLM Unlearning Via Quantization</p> <p><i>Zhiwei Zhang, Fali Wang, Xiaomin Li, Zongyu Wu, Xianfeng Tang, Hui Liu, Qi He, Wenpeng Yin, Suhang Wang</i></p>	54940
<p>CPSample: Classifier Protected Sampling for Guarding Training Data During Diffusion</p> <p><i>Joshua Kazdan, Hao Sun, Jiaqi Han, Felix Petersen, Frederick Vu, Stefano Ermon</i></p>	54964
<p>Contractive Dynamical Imitation Policies for Efficient Out-Of-Sample Recovery</p> <p><i>Amin Abyaneh, Mahrokh Ghoddousi Boroujeni, Hsiu-Chin Lin, Giancarlo Ferrari-Trecate</i></p>	54987
<p>Efficient and Accurate Explanation Estimation with Distribution Compression</p> <p><i>Hubert Baniecki, Giuseppe Casalicchio, Bernd Bischl, Przemyslaw Biecek</i></p>	55020
<p>Agent S: An Open Agentic Framework that Uses Computers Like a Human</p> <p><i>Saaket Agashe, Jiuzhou Han, Shuyu Gan, Jiachen Yang, Ang Li, Xin Eric Wang</i></p>	55060
<p>Diversity-Rewarded CFG Distillation</p> <p><i>Geoffrey Cideron, Andrea Agostinelli, Johan Ferret, Sertan Girgin, Romuald Elie, Olivier Bachem, Sarah Perrin, Alexandre Rame</i></p>	55083

<p>Symbolic Regression Via MDLformer-Guided Search: From Minimizing Prediction Error to Minimizing Description Length</p> <p>Zihan Yu, Jingtao Ding, Yong Li, Depeng Jin</p>	55106
<p>MonST3R: A Simple Approach for Estimating Geometry in the Presence of Motion</p> <p>Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, Ming-Hsuan Yang</p>	55134
<p>Unifying Causal Representation Learning with the Invariance Principle</p> <p>Dingling Yao, Dario Rancati, Riccardo Cadei, Marco Fumero, Francesco Locatello</p>	55158
<p>Cut the Crap: An Economical Communication Pipeline for LLM-Based Multi-Agent Systems</p> <p>Guibin Zhang, Yanwei Yue, Zhixun Li, Sukwon Yun, Guancheng Wan, Kun Wang, Dawei Cheng, Jeffrey Xu Yu, Tianlong Chen</p>	55202
<p>Hidden in the Noise: Two-Stage Robust Watermarking for Images</p> <p>Kasra Arabi, Benjamin Feuer, R. Teal Witter, Chinmay Hegde, Niv Cohen</p>	55242
<p>Learning Chaos in a Linear Way</p> <p>Xiaoyuan Cheng, Yi He, Yiming Yang, Xiao Xue, Sibao Cheng, Daniel Giles, Xiaohang Tang, Yukun Hu</p>	55276
<p>Differentially Private Steering for Large Language Model Alignment</p> <p>Anmol Goel, Yaxi Hu, Iryna Gurevych, Amartya Sanyal</p>	55319
<p>Point-Based Instance Completion with Scene Constraints</p> <p>Wesley Khademi, Fuxin Li</p>	55341
<p>Leveraging Driver Field-Of-View for Multimodal Ego-Trajectory Prediction</p> <p>M. Eren Akbiyik, Nedko Savov, Danda Pani Paudel, Nikola Popovic, Christian Vater, Otmar Hilliges, Luc Van Gool, Xi Wang</p>	55372
<p>One Hundred Neural Networks and Brains Watching Videos: Lessons from Alignment</p> <p>Christina Sartzetaki, Gemma Roig, Cees G. M. Snoek, Iris Groen</p>	55401
<p>Causal Concept Graph Models: Beyond Causal Opacity in Deep Learning</p> <p>Gabriele Dominici, Pietro Barbiero, Mateo Espinosa Zarlenga, Alberto Termine, Martin Gjoreski, Giuseppe Marra, Marc Langheinrich</p>	55428
<p>Towards Domain Adaptive Neural Contextual Bandits</p> <p>Ziyan Wang, Xiaoming Huo, Hao Wang</p>	55458
<p>Frame-Voyager: Learning to Query Frames for Video Large Language Models</p> <p>Sicheng Yu, Chengkai Jin, Huanyu Wang, Zhenghao Chen, Sheng Jin, Zhongrong Zuo, Xu Xiaolei, Zhenbang Sun, Bingni Zhang, Jiawei Wu, Hao Zhang, Qianru Sun</p>	55482
<p>SVD-LLM: Truncation-Aware Singular Value Decomposition for Large Language Model Compression</p> <p>Xin Wang, Yu Zheng, Zhongwei Wan, Mi Zhang</p>	55508

<p>From Promise to Practice: Realizing High-Performance Decentralized Training</p> <p>Zesen Wang, Jiaojiao Zhang, Xuyang Wu, Mikael Johansson</p>	55529
<p>Does Safety Training of LLMs Generalize to Semantically Related Natural Prompts?</p> <p>Sravanti Addepalli, Yerram Varun, Arun Suggala, Karthikeyan Shanmugam, Prateek Jain</p>	55575
<p>Multi-Dimensional Conformal Prediction</p> <p>Yam Tawachi, Bracha Laufer-Goldshtein</p>	55596
<p>Efficient Top-M Data Values Identification for Data Selection</p> <p>Xiaoqiang Lin, Xinyi Xu, See-Kiong Ng, Bryan Kian Hsiang Low</p>	55626
<p>Model Editing as a Robust and Denoised Variant of DPO: A Case Study on Toxicity</p> <p>Rheeya Uppaal, Apratim Dey, Yiting He, Yiqiao Zhong, Junjie Hu</p>	55653
<p>Framer: Interactive Frame Interpolation</p> <p>Wen Wang, Qiuyu Wang, Kecheng Zheng, Hao Ouyang, Zhekai Chen, Biao Gong, Hao Chen, Yujun Shen, Chunhua Shen</p>	55685
<p>DistRL: An Asynchronous Distributed Reinforcement Learning Framework for On-Device Control Agent</p> <p>Taiyi Wang, Zhihao Wu, Jianheng Liu, Jianye Hao, Jun Wang, Kun Shao</p>	55713
<p>DCT-CryptoNets: Scaling Private Inference in the Frequency Domain</p> <p>Arjun Roy, Kaushik Roy</p>	55739
<p>Complementary Label Learning with Positive Label Guessing and Negative Label Enhancement</p> <p>Yuhang Li, Zhuying Li, Yuheng Jia</p>	55757
<p>What Matters in Learning from Large-Scale Datasets for Robot Manipulation</p> <p>Vaibhav Saxena, Matthew Bronars, Nadun Ranawaka Arachchige, Kuancheng Wang, Woo Chul Shin, Soroush Nasiriany, Ajay Mandlekar, Danfei Xu</p>	55785
<p>SLoPe: Double-Pruned Sparse Plus Lazy Low-Rank Adapter Pretraining of LLMs</p> <p>Mohammad Mozaffari, Amir Yazdanbakhsh, Zhao Zhang, Maryam Mehri Dehnavi</p>	55814
<p>RESuM: A Rare Event Surrogate Model for Physics Detector Design</p> <p>Ann-Kathrin Schuetz, A. W. P. Poon, Aobo Li</p>	55844
<p>Generalized Principal-Agent Problem with a Learning Agent</p> <p>Tao Lin, Yiling Chen</p>	55865
<p>CogVideoX: Text-To-Video Diffusion Models with an Expert Transformer</p> <p>Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, Da Yin, Yuxuan. Zhang, Weihang Wang, Yean Cheng, Bin Xu, Xiaotao Gu, Yuxiao Dong, Jie Tang</p>	55897
<p>HAINAN: Fast and Accurate Transducer for Hybrid-Autoregressive ASR</p> <p>Hainan Xu, Travis M. Bartley, Vladimir Bataev, Boris Ginsburg</p>	55927

<p>Cross-Domain Offline Policy Adaptation with Optimal Transport and Dataset Constraint</p> <p><i>Jiafei Lyu, Mengbei Yan, Zhongjian Qiao, Runze Liu, Xiaoteng Ma, Deheng Ye, Jing-Wen Yang, Zongqing Lu, Xiu Li</i></p>	55943
<p>Adversarial Score Identity Distillation: Rapidly Surpassing the Teacher in One Step</p> <p><i>Mingyuan Zhou, Huangjie Zheng, Yi Gu, Zhendong Wang, Hai Huang</i></p>	55973
<p>LoRA3D: Low-Rank Self-Calibration of 3D Geometric Foundation Models</p> <p><i>Ziqi Lu, Heng Yang, Danfei Xu, Boyi Li, Boris Ivanovic, Marco Pavone, Yue Wang</i></p>	56007
<p>Is Uniform Expressivity Too Restrictive? Towards Efficient Expressivity of GNNs</p> <p><i>Sammy Khalife, Josué Tonelli-Cueto</i></p>	56030
<p>Bringing NeRFs to the Latent Space: Inverse Graphics Autoencoder</p> <p><i>Antoine Schnepf, Karim Kassab, Jean-Yves Franceschi, Laurent Caraffa, Flavian Vasile, Jeremie Mary, Andrew I. Comport, Valerie Gouet-Brunet</i></p>	56051
<p>HMoRA: Making LLMs More Effective with Hierarchical Mixture of LoRA Experts</p> <p><i>Mengqi Liao, Wei Chen, Junfeng Shen, Shengnan Guo, Huaiyu Wan</i></p>	56079
<p>PT-T2I/V: An Efficient Proxy-Tokenized Diffusion Transformer for Text-To-Image/Video-Task</p> <p><i>Jing Wang, Ao Ma, Jiasong Feng, Dawei Leng, Yuhui Yin, Xiaodan Liang</i></p>	56101
<p>Robust Conformal Prediction with a Single Binary Certificate</p> <p><i>Soroush H. Zargarbashi, Aleksandar Bojchevski</i></p>	56118
<p>Dynamic Gaussians Mesh: Consistent Mesh Reconstruction from Dynamic Scenes</p> <p><i>Isabella Liu, Hao Su, Xiaolong Wang</i></p>	56144
<p>AdaManip: Adaptive Articulated Object Manipulation Environments and Policy Learning</p> <p><i>Yuanfei Wang, Xiaojie Zhang, Ruihai Wu, Yu Li, Yan Shen, Mingdong Wu, Zhaofeng He, Yizhou Wang, Hao Dong</i></p>	56180
<p>Behavioral Entropy-Guided Dataset Generation for Offline Reinforcement Learning</p> <p><i>Wesley A. Suttle, Aamodh Suresh, Carlos Nieto-Granda</i></p>	56196
<p>Unsupervised Disentanglement of Content and Style Via Variance-Invariance Constraints</p> <p><i>Yuxuan Wu, Ziyu Wang, Bhiksha Raj, Gus Xia</i></p>	56219
<p>Towards Out-Of-Modal Generalization Without Instance-Level Modal Correspondence</p> <p><i>Zhuo Huang, Gang Niu, Bo Han, Masashi Sugiyama, Tongliang Liu</i></p>	56245
<p>Eliciting Human Preferences with Language Models</p> <p><i>Belinda Z. Li, Alex Tamkin, Noah Goodman, Jacob Andreas</i></p>	56265
<p>AdaRankGrad: Adaptive Gradient Rank and Moments for Memory-Efficient LLMs Training and Fine-Tuning</p> <p><i>Yehonathan Refael, Jonathan Svirsky, Boris Shustin, Wasim Huleihel, Ofir Lindenbaum</i></p>	56295

<p>TempMe: Video Temporal Token Merging for Efficient Text-Video Retrieval</p> <p>Leqi Shen, Tianxiang Hao, Tao He, Sicheng Zhao, Yifeng Zhang, Pengzhang Liu, Yongjun Bao, Guiguang Ding</p>	56317
<p>Non-Adversarial Inverse Reinforcement Learning Via Successor Feature Matching</p> <p>Arnav Kumar Jain, Harley Wiltzer, Jesse Farebrother, Irina Rish, Glen Berseth, Sanjiban Choudhury</p>	56339
<p>On the Fourier Analysis in the SO(3) Space : The EquiLoPO Network</p> <p>Dmitrii Zhemchuzhnikov, Sergei Grudin</p>	56367
<p>GRNAde: Geometric Deep Learning for 3D RNA Inverse Design</p> <p>Chaitanya K. Joshi, Arian Rokkum Jamasb, Ramon Viñas Torné, Charles Harris, Simon V Mathis, Alex Morehead, Rishabh Anand, Pietro Lio</p>	56390
<p>Enhancing Robust Fairness Via Confusional Spectral Regularization</p> <p>Gaojie Jin, Sihao Wu, Jiayu Liu, Tianjin Huang, Ronghui Mu</p>	56416
<p>Pursuing Better Decision Boundaries for Long-Tailed Object Detection Via Category Information Amount</p> <p>Yanbiao Ma, Wei Dai, Jiayi Chen</p>	56438
<p>Last-Iterate Convergence Properties of Regret-Matching Algorithms in Games</p> <p>Yang Cai, Gabriele Farina, Julien Grand-Clément, Christian Kroer, Chung-Wei Lee, Haipeng Luo, Weiqiang Zheng</p>	56454
<p>Palu: KV-Cache Compression with Low-Rank Projection</p> <p>Chi-Chih Chang, Wei-Cheng Lin, Chien-Yu Lin, Chong-Yan Chen, Yu-Fang Hu, Pei-Shuo Wang, Ning-Chi Huang, Luis Ceze, Mohamed S. Abdelfattah, Kai-Chiang Wu</p>	56482
<p>Do as I Do (Safely): Mitigating Task-Specific Fine-Tuning Risks in Large Language Models</p> <p>Francisco Eiras, Aleksandar Petrov, Philip Torr, M. Pawan Kumar, Adel Bibi</p>	56510
<p>POTEC: Off-Policy Contextual Bandits for Large Action Spaces Via Policy Decomposition</p> <p>Yuta Saito, Jihan Yao, Thorsten Joachims</p>	56529
<p>Token Statistics Transformer: Linear-Time Attention Via Variational Rate Reduction</p> <p>Ziyang Wu, Tianjiao Ding, Yifu Lu, Druv Pai, Jingyuan Zhang, Weida Wang, Yaodong Yu, Yi Ma, Benjamin David Haeffele</p>	56554
<p>Predicate Hierarchies Improve Few-Shot State Classification</p> <p>Emily Jin, Joy Hsu, Jiajun Wu</p>	56579
<p>LaMP: Language-Motion Pretraining for Motion Generation, Retrieval, and Captioning</p> <p>Zhe Li, Weihao Yuan, Yisheng He, Lingteng Qiu, Shenhao Zhu, Xiaodong Gu, Weichao Shen, Yuan Dong, Zilong Dong, Laurence Tianruo Yang</p>	56604
<p>Local Steps Speed Up Local GD for Heterogeneous Distributed Logistic Regression</p> <p>Michael Crawshaw, Blake Woodworth, Mingrui Liu</p>	56617

<p>Neuron Based Personality Trait Induction in Large Language Models</p> <p><i>Jia Deng, Tianyi Tang, Yanbin Yin, Wenhao Yang, Xin Zhao, Ji-Rong Wen</i></p>	56679
<p>Simplifying, Stabilizing and Scaling Continuous-Time Consistency Models</p> <p><i>Cheng Lu, Yang Song</i></p>	56704
<p>Analysis of Linear Mode Connectivity Via Permutation-Based Weight Matching: With Insights into Other Permutation Search Methods</p> <p><i>Akira Ito, Masanori Yamada, Atsutoshi Kumagai</i></p>	56743
<p>A Periodic Bayesian Flow for Material Generation</p> <p><i>Hanlin Wu, Yuxuan Song, Jingjing Gong, Ziyao Cao, Yawen Ouyang, Jianbing Zhang, Hao Zhou, Wei-Ying Ma, Jingjing Liu</i></p>	56780
<p>Boltzmann-Aligned Inverse Folding Model as a Predictor of Mutational Effects on Protein-Protein Interactions</p> <p><i>Xiaoran Jiao, Weian Mao, Wengong Jin, Peiyuan Yang, Hao Chen, Chunhua Shen</i></p>	56810
<p>Accelerating Auto-Regressive Text-To-Image Generation with Training-Free Speculative Jacobi Decoding</p> <p><i>Yao Teng, Han Shi, Xian Liu, Xuefei Ning, Guohao Dai, Yu Wang, Zhenguo Li, Xihui Liu</i></p>	56828
<p>Task Descriptors Help Transformers Learn Linear Models In-Context</p> <p><i>Ruomin Huang, Rong Ge</i></p>	56853
<p>CATCH: Channel-Aware Multivariate Time Series Anomaly Detection Via Frequency Patching</p> <p><i>Xingjian Wu, Xiangfei Qiu, Zhengyu Li, Yihang Wang, Jilin Hu, Chenjuan Guo, Hui Xiong, Bin Yang</i></p>	56894
<p>Can LLMs Generate Novel Research Ideas? a Large-Scale Human Study with 100+ NLP Researchers</p> <p><i>Chenglei Si, Diyi Yang, Tatsunori Hashimoto</i></p>	56923
<p>SMI-Editor: Edit-Based SMILES Language Model with Fragment-Level Supervision</p> <p><i>Kangjie Zheng, Siyue Liang, Junwei Yang, Bin Feng, Zequn Liu, Wei Ju, Zhiping Xiao, Ming Zhang</i></p>	57013
<p>Tracking Objects that Change in Appearance with Phase Synchrony</p> <p><i>Sabine Muzellec, Drew Linsley, Alekh Karkada Ashok, Ennio Mingolla, Girik Malik, Rufin Vanrullen, Thomas Serre</i></p>	57043
<p>STORM: Spatio-Temporal Reconstruction Model for Large-Scale Outdoor Scenes</p> <p><i>Jiawei Yang, Jiahui Huang, Boris Ivanovic, Yuxiao Chen, Yan Wang, Boyi Li, Yurong You, Apoorva Sharma, Maximilian Igl, Peter Karkus, Danfei Xu, Yue Wang, Marco Pavone</i></p>	57082
<p>LLM-SR: Scientific Equation Discovery Via Programming with Large Language Models</p> <p><i>Parshin Shojaee, Kazem Meidani, Shashank Gupta, Amir Barati Farimani, Chandan K. Reddy</i></p>	57102
<p>CubeDiff: Repurposing Diffusion-Based Image Models for Panorama Generation</p> <p><i>Nikolai Kalischek, Michael Oechsle, Fabian Manhardt, Philipp Henzler, Konrad Schindler, Federico Tombari</i></p>	57134

<p>Trusted Multi-View Classification Via Evolutionary Multi-View Fusion</p> <p>Xinyan Liang, Pinhan Fu, Yuhua Qian, Qian Guo, Guoqing Liu</p>	57147
<p>CAT-3DGS: A Context-Adaptive Triplane Approach to Rate-Distortion-Optimized 3DGS Compression</p> <p>Yu-Ting Zhan, Cheng-Yuan Ho, Hebi Yang, Yi-Hsin Chen, Jui Chiu Chiang, Yu-Lun Liu, Wen-Hsiao Peng</p>	57173
<p>DriveTransformer: Unified Transformer for Scalable End-To-End Autonomous Driving</p> <p>Xiaosong Jia, Junqi You, Zhiyuan Zhang, Junchi Yan</p>	57196
<p>An Engorgio Prompt Makes Large Language Model Babble on</p> <p>Jianshuo Dong, Ziyuan Zhang, Qingjie Zhang, Tianwei Zhang, Hao Wang, Hewu Li, Qi Li, Chao Zhang, Ke Xu, Han Qiu</p>	57213
<p>OpenRCA: Can Large Language Models Locate the Root Cause of Software Failures?</p> <p>Junjielong Xu, Qinan Zhang, Zhiqing Zhong, Shilin He, Chaoyun Zhang, Qingwei Lin, Dan Pei, Pinjia He, Dongmei Zhang, Qi Zhang</p>	57241
<p>Understanding Warmup-Stable-Decay Learning Rates: A River Valley Loss Landscape View</p> <p>Kaiyue Wen, Zhiyuan Li, Jason S. Wang, David Leo Wright Hall, Percy Liang, Tengyu Ma</p>	57270
<p>Bridging the Gap Between F-Divergences and Bayes Hilbert Spaces</p> <p>Linus Lach, Alexander Willi Fottner, Yarema Okhrin</p>	57316
<p>AdalR: Adaptive All-In-One Image Restoration Via Frequency Mining and Modulation</p> <p>Yuning Cui, Syed Waqas Zamir, Salman Khan, Alois Knoll, Mubarak Shah, Fahad Shahbaz Khan</p>	57335
<p>Chain-Of-Region: Visual Language Models Need Details for Diagram Analysis</p> <p>Xue Li, Yiyu Sun, Wei Cheng, Yinglun Zhu, Haifeng Chen</p>	57357
<p>AdvPaint: Protecting Images from Inpainting Manipulation Via Adversarial Attention Disruption</p> <p>Joonsung Jeon, Woo Jae Kim, Suhyeon Ha, Sooel Son, Sung-Eui Yoon</p>	57374
<p>DynamicCity: Large-Scale 4D Occupancy Generation from Dynamic Scenes</p> <p>Hengwei Bian, Lingdong Kong, Haozhe Xie, Liang Pan, Yu Qiao, Ziwei Liu</p>	57388
<p>Do Egocentric Video-Language Models Truly Understand Hand-Object Interactions?</p> <p>Boshen Xu, Ziheng Wang, Yang Du, Zhinan Song, Sipeng Zheng, Qin Jin</p>	57423
<p>Following the Human Thread in Social Navigation</p> <p>Luca Scofano, Alessio Sampieri, Tommaso Campari, Valentino Sacco, Indro Spinelli, Lamberto Ballan, Fabio Galasso</p>	57451
<p>SynCamMaster: Synchronizing Multi-Camera Video Generation from Diverse Viewpoints</p> <p>Jianhong Bai, Menghan Xia, Xintao Wang, Ziyang Yuan, Zuozhu Liu, Haoji Hu, Pengfei Wan, Di Zhang</p>	57472
<p>Is Your Video Language Model a Reliable Judge?</p> <p>Ming Liu, Wensheng Zhang</p>	57495

<p>OGBench: Benchmarking Offline Goal-Conditioned RL</p> <p>Seohong Park, Kevin Frans, Benjamin Eysenbach, Sergey Levine</p>	57515
<p>Grounding Multimodal Large Language Model in GUI World</p> <p>Weixian Lei, Difei Gao, Mike Zheng Shou</p>	57561
<p>Overcoming False Illusions in Real-World Face Restoration with Multi-Modal Guided Diffusion Model</p> <p>Keda Tao, Jinjin Gu, Yulun Zhang, Xiucheng Wang, Nan Cheng</p>	57585
<p>Rethinking Invariance Regularization in Adversarial Training to Improve Robustness-Accuracy Trade-Off</p> <p>Futa Kai Waseda, Ching-Chun Chang, Isao Echizen</p>	57610
<p>MQuAKE-Remastered: Multi-Hop Knowledge Editing Can Only Be Advanced with Reliable Evaluations</p> <p>Shaochen Zhong, Yifan Lu, Lize Shao, Bhargav Bhushanam, Xiaocong Du, Yixin Wan, Yucheng Shi, Daochen Zha, Yiwei Wang, Ninghao Liu, Kaixiong Zhou, Shuai Xu, Kai-Wei Chang, Louis Feng, Vipin Chaudhary, Xia Hu</p>	57634
<p>3DIS: Depth-Driven Decoupled Image Synthesis for Universal Multi-Instance Generation</p> <p>Dewei Zhou, Ji Xie, Zongxin Yang, Yi Yang</p>	57663
<p>Efficient Distribution Matching of Representations Via Noise-Injected Deep InfoMax</p> <p>Ivan Butakov, Alexander Semenenko, Alexander Tolmachev, Andrey Gladkov, Marina Munkhoeva, Alexey Frolov</p>	57676
<p>Coreset Selection Via Reducible Loss in Continual Learning</p> <p>Ruilin Tong, Yuhang Liu, Javen Qinfeng Shi, Dong Gong</p>	57701
<p>Adaptive Length Image Tokenization Via Recurrent Allocation</p> <p>Shivam Duggal, Phillip Isola, Antonio Torralba, William T. Freeman</p>	57737
<p>URLOST: Unsupervised Representation Learning Without Stationarity Or Topology</p> <p>Zeyu Yun, Juexiao Zhang, Yann Lecun, Yubei Chen</p>	57764
<p>Diffusion Transformer Captures Spatial-Temporal Dependencies: A Theory for Gaussian Process Data</p> <p>Hengyu Fu, Zehao Dou, Jiawei Guo, Mengdi Wang, Minshuo Chen</p>	57789
<p>MLPs Learn In-Context on Regression and Classification Tasks</p> <p>William Lingxiao Tong, Cengiz Pehlevan</p>	57840
<p>SFS: Smarter Code Space Search Improves LLM Inference Scaling</p> <p>Jonathan Light, Yue Wu, Yiyu Sun, Wenchao Yu, Yanchi Liu, Xujiang Zhao, Ziniu Hu, Haifeng Chen, Wei Cheng</p>	57870
<p>From Tokens to Lattices: Emergent Lattice Structures in Language Models</p> <p>Bo Xiong, Steffen Staab</p>	57907

<p>Image Watermarks Are Removable Using Controllable Regeneration from Clean Noise</p> <p>Yepeng Liu, Yiren Song, Hai Ci, Yu Zhang, Haofan Wang, Mike Zheng Shou, Yuheng Bu</p>	57922
<p>Mechanistic Permutability: Match Features Across Layers</p> <p>Nikita Balagansky, Ian Maksimov, Daniil Gavrilov</p>	57940
<p>Moral Alignment for LLM Agents</p> <p>Elizaveta Tennant, Stephen Hailes, Mirco Musolesi</p>	57958
<p>Preble: Efficient Distributed Prompt Scheduling for LLM Serving</p> <p>Vikranth Srivatsa, Zijian He, Reyna Abhyankar, Dongming Li, Yiyang Zhang</p>	57992
<p>Diffusion Policy Policy Optimization</p> <p>Allen Z. Ren, Justin Lidard, Lars Lien Ankile, Anthony Simeonov, Pulkit Agrawal, Anirudha Majumdar, Benjamin Burchfiel, Hongkai Dai, Max Simchowitz</p>	58018
<p>Predictive Inverse Dynamics Models Are Scalable Learners for Robotic Manipulation</p> <p>Yang Tian, Sizhe Yang, Jia Zeng, Ping Wang, Dahua Lin, Hao Dong, Jiangmiao Pang</p>	58060
<p>Compute-Optimal LLMs Provably Generalize Better with Scale</p> <p>Marc Anton Finzi, Sanyam Kapoor, Diego Granzio, Anming Gu, Christopher De Sa, J Zico Kolter, Andrew Gordon Wilson</p>	58080
<p>Learning to Discover Regulatory Elements for Gene Expression Prediction</p> <p>Xingyu Su, Haiyang Yu, Degui Zhi, Shuiwang Ji</p>	58104
<p>Apollo-MILP: An Alternating Prediction-Correction Neural Solving Framework for Mixed-Integer Linear Programming</p> <p>Haoyang Liu, Jie Wang, Zijie Geng, Xijun Li, Yuxuan Zong, Fangzhou Zhu, Jianye Hao, Feng Wu</p>	58120
<p>D-FINE: Redefine Regression Task of DETRs as Fine-Grained Distribution Refinement</p> <p>Yansong Peng, Hebei Li, Peixi Wu, Yueyi Zhang, Xiaoyan Sun, Feng Wu</p>	58148
<p>Mitigating Memorization in Language Models</p> <p>Mansi Sakarvadia, Aswathy Ajith, Arham Mushtaq Khan, Nathaniel C Hudson, Caleb Geniesse, Kyle Chard, Yaoqing Yang, Ian Foster, Michael W. Mahoney</p>	58165
<p>GOFA: A Generative One-For-All Model for Joint Graph Language Modeling</p> <p>Lecheng Kong, Jiarui Feng, Hao Liu, Chengsong Huang, Jiaxin Huang, Yixin Chen, Muhan Zhang</p>	58209
<p>QP-SNN: Quantized and Pruned Spiking Neural Networks</p> <p>Wenjie Wei, Malu Zhang, Zijian Zhou, Ammar Belatreche, Yimeng Shan, Yu Liang, Honglin Cao, Jieyuan Zhang, Yang Yang</p>	58241
<p>LICORICE: Label-Efficient Concept-Based Interpretable Reinforcement Learning</p> <p>Zhuorui Ye, Stephanie Milani, Geoffrey J. Gordon, Fei Fang</p>	58267

<p>Think While You Generate: Discrete Diffusion with Planned Denoising</p> <p>Sulin Liu, Juno Nam, Andrew Campbell, Hannes Stark, Yilun Xu, Tommi Jaakkola, Rafael Gomez-Bombarelli</p>	58296
<p>Asymmetric Factorized Bilinear Operation for Vision Transformer</p> <p>Junjie Wu, Qilong Wang, Jiangtao Xie, Pengfei Zhu, Qinghua Hu</p>	58344
<p>Probabilistic Geometric Principal Component Analysis with Application to Neural Data</p> <p>Han-Lin Hsieh, Maryam Shanechi</p>	58365
<p>WildBench: Benchmarking LLMs with Challenging Tasks from Real Users in the Wild</p> <p>Bill Yuchen Lin, Yuntian Deng, Khyathi Chandu, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, Yejin Choi</p>	58384
<p>Associative Memory and Dead Neurons</p> <p>Vladimir Fanaskov, Ivan Oseledets</p>	58403
<p>Spectral-Refiner: Accurate Fine-Tuning of Spatiotemporal Fourier Neural Operator for Turbulent Flows</p> <p>Shuhao Cao, Francesco Brarda, Ruipeng Li, Yuanzhe Xi</p>	58421
<p>SimXRD-4M: Big Simulated X-Ray Diffraction Data and Crystal Symmetry Classification Benchmark</p> <p>Bin Cao, Yang Liu, Zinan Zheng, Ruifeng Tan, Jia Li, Tong-Yi Zhang</p>	58458
<p>Zigzag Diffusion Sampling: Diffusion Models Can Self-Improve Via Self-Reflection</p> <p>Bai Lichen, Shitong Shao, Zikai Zhou, Zipeng Qi, Zhiqiang Xu, Haoyi Xiong, Zeke Xie</p>	58483
<p>X-NeMo: Expressive Neural Motion Reenactment Via Disentangled Latent Attention</p> <p>Xiaochen Zhao, Hongyi Xu, Guoxian Song, You Xie, Chenxu Zhang, Xiu Li, Linjie Luo, Jinli Suo, Yebin Liu</p>	58514
<p>ParetoFlow: Guided Flows in Multi-Objective Optimization</p> <p>Ye Yuan, Can Chen, Christopher Pal, Xue Liu</p>	58533
<p>Ready-To-React: Online Reaction Policy for Two-Character Interaction Generation</p> <p>Zhi Cen, Huaijin Pi, Sida Peng, Qing Shuai, Yujun Shen, Hujun Bao, Xiaowei Zhou, Ruizhen Hu</p>	58560
<p>HELM: Hierarchical Encoding for mRNA Language Modeling</p> <p>Mehdi Yazdani-Jahromi, Mangal Prakash, Tommaso Mansi, Artem Moskalev, Rui Liao</p>	58580
<p>Fragment and Geometry Aware Tokenization of Molecules for Structure-Based Drug Design Using Language Models</p> <p>Cong Fu, Xiner Li, Blake Olson, Heng Ji, Shuiwang Ji</p>	58604
<p>WizardMath: Empowering Mathematical Reasoning for Large Language Models Via Reinforced Evol-Instruct</p> <p>Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jian-Guang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, Yansong Tang, Dongmei Zhang</p>	58629

<p>Commit0: Library Generation from Scratch</p> <p>Wenting Zhao, Nan Jiang, Celine Lee, Justin T Chiu, Claire Cardie, Matthias Gallé, Alexander M Rush</p>	58666
<p>Continuous Exposure Learning for Low-Light Image Enhancement Using Neural ODEs</p> <p>Donggoo Jung, Daehyun Kim, Tae Hyun Kim</p>	58682
<p>HelpSteer2-Preference: Complementing Ratings with Preferences</p> <p>Zhilin Wang, Alexander Bukharin, Olivier Delalleau, Daniel Egert, Gerald Shen, Jiaqi Zeng, Oleksii Kuchaiev, Yi Dong</p>	58708
<p>MAVIS: Mathematical Visual Instruction Tuning with an Automatic Data Engine</p> <p>Renrui Zhang, Xinyu Wei, Dongzhi Jiang, Ziyu Guo, Yichi Zhang, Chengzhuo Tong, Jiaming Liu, Aojun Zhou, Shanghang Zhang, Peng Gao, Hongsheng Li</p>	58736
<p>Active Learning for Continual Learning: Keeping the Past Alive in the Present</p> <p>Jaehyun Park, Dongmin Park, Jae-Gil Lee</p>	58771
<p>Quamba: A Post-Training Quantization Recipe for Selective State Space Models</p> <p>Hung-Yueh Chiang, Chi-Chih Chang, Natalia Frumkin, Kai-Chiang Wu, Diana Marculescu</p>	58791
<p>DRESSing Up LLM: Efficient Stylized Question-Answering Via Style Subspace Editing</p> <p>Xinyu Ma, Yifeng Xu, Yang Lin, Tianlong Wang, Xu Chu, Xin Gao, Junfeng Zhao, Yasha Wang</p>	58818
<p>CBGBench: Fill in the Blank of Protein-Molecule Complex Binding Graph</p> <p>Haitao Lin, Guojiang Zhao, Odin Zhang, Yufei Huang, Lirong Wu, Cheng Tan, Zicheng Liu, Zhifeng Gao, Stan Z. Li</p>	58842
<p>Web Agents with World Models: Learning and Leveraging Environment Dynamics in Web Navigation</p> <p>Hyungjoo Chae, Namyoung Kim, Kai Tzu-lunn Ong, Minju Gwak, Gwanwoo Song, Jihoon Kim, Sunghwan Kim, Dongha Lee, Jinyoung Yeo</p>	58869
<p>From Isolated Conversations to Hierarchical Schemas: Dynamic Tree Memory Representation for LLMs</p> <p>Alireza Rezazadeh, Zichao Li, Wei Wei, Yujia Bao</p>	58901
<p>GS-CPR: Efficient Camera Pose Refinement Via 3D Gaussian Splatting</p> <p>Changkun Liu, Shuai Chen, Yash Sanjay Bhalgat, Siyan Hu, Ming Cheng, Zirui Wang, Victor Adrian Prisacariu, Tristan Braud</p>	58935
<p>AgentSquare: Automatic LLM Agent Search in Modular Design Space</p> <p>Yu Shang, Yu Li, Keyu Zhao, Likai Ma, Jiahe Liu, Fengli Xu, Yong Li</p>	58955
<p>Re-Aligning Language to Visual Objects with an Agentic Workflow</p> <p>Yuming Chen, Jiangyan Feng, Haodong Zhang, Lijun Gong, Feng Zhu, Rui Zhao, Qibin Hou, Ming-Ming Cheng, Yibing Song</p>	58980
<p>Visually Guided Decoding: Gradient-Free Hard Prompt Inversion with Language Models</p> <p>Donghoon Kim, Minji Bae, Kyuhong Shim, Byonghyo Shim</p>	59013

<p>Subtask-Aware Visual Reward Learning from Segmented Demonstrations</p> <p>Changyeon Kim, Minho Heo, Doohyun Lee, Honglak Lee, Jinwoo Shin, Joseph J Lim, Kimin Lee</p>	59039
<p>Towards Federated RLHF with Aggregated Client Preference for LLMs</p> <p>Feijie Wu, Xiaoze Liu, Haoyu Wang, Xingchen Wang, Lu Su, Jing Gao</p>	59063
<p>Beyond Model Collapse: Scaling Up with Synthesized Data Requires Verification</p> <p>Yunzhen Feng, Elvis Dohmatob, Pu Yang, Francois Charton, Julia Kempe</p>	59087
<p>TIGeR: Unifying Text-To-Image Generation and Retrieval with Large Multimodal Models</p> <p>Leigang Qu, Haochuan Li, Tan Wang, Wenjie Wang, Yongqi Li, Liqiang Nie, Tat-Seng Chua</p>	59116
<p>Efficient Perplexity Bound and Ratio Matching in Discrete Diffusion Language Models</p> <p>Etrit Haxholli, Yeti Z. Gurbuz, Ogul Can, Eli Waxman</p>	59146
<p>FlickerFusion: Intra-Trajectory Domain Generalizing Multi-Agent Reinforcement Learning</p> <p>Woosung Koh, Wonbeen Oh, Siyeol Kim, Suhin Shin, Hyeongjin Kim, Jaein Jang, Junghyun Lee, Se-Young Yun</p>	59197
<p>ALLaM: Large Language Models for Arabic and English</p> <p>M Saiful Bari, Yazeed Alnumay, Norah A. Alzahrani, Nouf M. Alotaibi, Hisham Abdullah Alyahya, Sultan Alrashed, Faisal Abdulrahman Mirza, Shaykhah Z. Alsubaie, Hassan A. Alahmed, Ghadah Alabduljabbar, Raghad Alkhathran, Yousef Almushayqih, Raneem Alnajim, Salman Alsubaihi, Maryam Al Mansour, Saad Amin Hassan, Dr. Majed Alrubaian, Ali Alammari, Zaki Alawami, Abdulmohsen Al-Thubaity, Ahmed Abdelali, Jeril Kuriakose, Abdalghani Abujabal, Nora Al-Twairesh, Areeb Alowisheq, Haidar Khan</p>	59235
<p>From Layers to States: A State Space Model Perspective to Deep Neural Network Layer Dynamics</p> <p>Qinshuo Liu, Weiqin Zhao, Wei Huang, Yanwen Fang, Lequan Yu, Guodong Li</p>	59271
<p>Correlated Proxies: A New Definition and Improved Mitigation for Reward Hacking</p> <p>Cassidy Laidlaw, Shivam Singhal, Anca Dragan</p>	59291
<p>Last Iterate Convergence of Incremental Methods as a Model of Forgetting</p> <p>Xufeng Cai, Jelena Diakonikolas</p>	59334
<p>Faster Diffusion Sampling with Randomized Midpoints: Sequential and Parallel</p> <p>Shivam Gupta, Linda Cai, Sitan Chen</p>	59369
<p>OpenMathInstruct-2: Accelerating AI for Math with Massive Open-Source Instruction Data</p> <p>Shubham Toshniwal, Wei Du, Ivan Moshkov, Branislav Kisacanicin, Alexan Ayrapetyan, Igor Gitman</p>	59405
<p>Efficient Online Pruning and Abstraction for Imperfect Information Extensive-Form Games</p> <p>Boning Li, Longbo Huang</p>	59438
<p>Anti-Exposure Bias in Diffusion Models</p> <p>Junyu Zhang, Daochang Liu, Eunbyung Park, Shichao Zhang, Chang Xu</p>	59458

<p>Mind the Gap: Examining the Self-Improvement Capabilities of Large Language Models</p> <p>Yuda Song, Hanlin Zhang, Carson Eisenach, Sham M. Kakade, Dean Foster, Udaya Ghai</p>	59479
<p>Judge Decoding: Faster Speculative Sampling Requires Going Beyond Model Alignment</p> <p>Gregor Bachmann, Sotiris Anagnostidis, Albert Pumarola, Markos Georgopoulos, Artsiom Sanakoyeu, Yuming Du, Edgar Schönfeld, Ali Thabet, Jonas K Kohler</p>	59517
<p>BP-Modified Local Loss for Efficient Training of Deep Neural Networks</p> <p>Ren Lianhai, Qianxiao Li</p>	59538
<p>A General Framework for Off-Policy Learning with Partially-Observed Reward</p> <p>Rikiya Takehi, Masahiro Asami, Kosuke Kawakami, Yuta Saito</p>	59556
<p>CipherPrune: Efficient and Scalable Private Transformer Inference</p> <p>Yancheng Zhang, Jiaqi Xue, Mengxin Zheng, Mimi Xie, Mingzhe Zhang, Lei Jiang, Qian Lou</p>	59578
<p>Atomas: Hierarchical Adaptive Alignment on Molecule-Text for Unified Molecule Understanding and Generation</p> <p>Yikun Zhang, Geyan Ye, Chaohao Yuan, Bo Han, Long-Kai Huang, Jianhua Yao, Wei Liu, Yu Rong</p>	59597
<p>Pursuing Feature Separation Based on Neural Collapse for Out-Of-Distribution Detection</p> <p>Yingwen Wu, Ruiji Yu, Xinwen Cheng, Zhengbao He, Xiaolin Huang</p>	59626
<p>MetaDesigner: Advancing Artistic Typography Through AI-Driven, User-Centric, and Multilingual WordArt Synthesis</p> <p>Jun-Yan He, Zhi-Qi Cheng, Chenyang Li, Jingdong Sun, Qi He, Wangmeng Xiang, Hanyuan Chen, Jin-Peng Lan, Xianhui Lin, Kang Zhu, Bin Luo, Yifeng Geng, Xuansong Xie, Alexander G Hauptmann</p>	59646
<p>MamBEV: Enabling State Space Models to Learn Birds-Eye-View Representations</p> <p>Hongyu Ke, Jack Morris, Kentaro Oguchi, Xiaofei Cao, Yongkang Liu, Haoxin Wang, Yi Ding</p>	59670
<p>Motion Control of High-Dimensional Musculoskeletal Systems with Hierarchical Model-Based Planning</p> <p>Yunyue Wei, Shanning Zhuang, Vincent Zhuang, Yanan Sui</p>	59689
<p>Continuity-Preserving Convolutional Autoencoders for Learning Continuous Latent Dynamical Models from Images</p> <p>Aiqing Zhu, Yuting Pan, Qianxiao Li</p>	59707
<p>NNsight and NDIF: Democratizing Access to Open-Weight Foundation Model Internals</p> <p>Jaden Fried Fiotto-Kaufman, Alexander Russell Loftus, Eric Todd, Jannik Brinkmann, Koyena Pal, Dmitrii Troitskii, Michael Ripa, Adam Belfki, Can Rager, Caden Juang, Aaron Mueller, Samuel Marks, Arnab Sen Sharma, Francesca Lucchetti, Nikhil Prakash, Carla E. Brodley, Arjun Guha, Jonathan Bell, Byron C Wallace, David Bau</p>	59731
<p>E(3)-Equivariant Models Cannot Learn Chirality: Field-Based Molecular Generation</p> <p>Alexandru Dumitrescu, Dani Korpela, Markus Heinonen, Yogesh Verma, Valerii Iakovlev, Vikas Garg, Harri Lähdesmäki</p>	59765

Discovering Clone Negatives Via Adaptive Contrastive Learning for Image-Text Matching Renjie Pan, Jihao Dong, Hua Yang	59800
Topological Zigzag Spaghetti for Diffusion-Based Generation and Prediction on Graphs Yuzhou Chen, Yulia Gel	59827
Layout-Your-3D: Controllable and Precise 3D Generation with 2D Blueprint Junwei Zhou, Xueting Li, Lu Qi, Ming-Hsuan Yang	59845
SEBRA : Debiasing Through Self-Guided Bias Ranking Adarsh Kappiyath, Abhra Chaudhuri, Ajay Kumar Jaiswal, Ziquan Liu, Yunpeng Li, Xiatian Zhu, Lu Yin	59868
NUDGE: Lightweight Non-Parametric Fine-Tuning of Embeddings for Retrieval Sepanta Zeighami, Zac Wellmer, Aditya Parameswaran	59891
Zeroth-Order Fine-Tuning of LLMs with Transferable Static Sparsity Wentao Guo, Jikai Long, Yimeng Zeng, Zirui Liu, Xinyu Yang, Yide Ran, Jacob R. Gardner, Osbert Bastani, Christopher De Sa, Xiaodong Yu, Beidi Chen, Zhaozhuo Xu	59924
TVNet: A Novel Time Series Analysis Method Based on Dynamic Convolution and 3D-Variation Chenghan Li, Mingchen Li, Ruisheng Diao	59965
Style Outweighs Substance: Failure Modes of LLM Judges in Alignment Benchmarking Benjamin Feuer, Micah Goldblum, Teresa Datta, Sanjana Nambiar, Raz Besaleli, Samuel Dooley, Max Cembalest, John P Dickerson	60004
MMR: A Large-Scale Benchmark Dataset for Multi-Target and Multi-Granularity Reasoning Segmentation Donggon Jang, Yucheol Cho, Suin Lee, Taehyeon Kim, Daeshik Kim	60035
HQ-Edit: A High-Quality Dataset for Instruction-Based Image Editing Mude Hui, Siwei Yang, Bingchen Zhao, Yichun Shi, Heng Wang, Peng Wang, Cihang Xie, Yuyin Zhou	60061
On the Adversarial Vulnerability of Label-Free Test-Time Adaptation Shahriar Rifat, Jonathan Ashdown, Michael J. De Lucia, Ananthram Swami, Francesco Restuccia	60079
Generalization Bounds for Canonicalization: A Comparative Study with Group Averaging Behrooz Tahmasebi, Stefanie Jegelka	60095
ThinK: Thinner Key Cache by Query-Driven Pruning Yuhui Xu, Zhanming Jie, Hanze Dong, Lei Wang, Xudong Lu, Aojun Zhou, Amrita Saha, Caiming Xiong, Doyen Sahoo	60115
Multi-Draft Speculative Sampling: Canonical Decomposition and Theoretical Limits Ashish J Khisti, Mohammadreza Ebrahimi, Hassan Dbouk, Arash Behboodi, Roland Memisevic, Christos Louizos	60134
Weak to Strong Generalization for Large Language Models with Multi-Capabilities Yucheng Zhou, Jianbing Shen, Yu Cheng	60175

<p>Transformers Provably Solve Parity Efficiently with Chain of Thought</p> <p>Juno Kim, Taiji Suzuki</p>	60205
<p>Shedding Light on Time Series Classification Using Interpretability Gated Networks</p> <p>Yunshi Wen, Tengfei Ma, Ronny Luss, Debarun Bhattacharjya, Achille Fokoue, Anak Agung Julius</p>	60232
<p>A Tight Convergence Analysis of Inexact Stochastic Proximal Point Algorithm for Stochastic Composite Optimization Problems</p> <p>Shulan Zhu, Chenglong Bao, Defeng Sun, Yancheng Yuan</p>	60268
<p>TOP-ERL: Transformer-Based Off-Policy Episodic Reinforcement Learning</p> <p>Ge Li, Dong Tian, Hongyi Zhou, Xinkai Jiang, Rudolf Lioutikov, Gerhard Neumann</p>	60311
<p>Learning-Augmented Search Data Structures</p> <p>Chunkai Fu, Brandon G. Nguyen, Jung Hoon Seo, Ryan S. Zesch, Samson Zhou</p>	60341
<p>Learning Mirror Maps in Policy Mirror Descent</p> <p>Carlo Alfano, Sebastian Rene Towers, Silvia Sapora, Chris Lu, Patrick Rebeschini</p>	60369
<p>Fiddler: CPU-GPU Orchestration for Fast Inference of Mixture-Of-Experts Models</p> <p>Keisuke Kamahori, Tian Tang, Yile Gu, Kan Zhu, Baris Kasikci</p>	60388
<p>Linear Combinations of Latents in Generative Models: Subspaces and Beyond</p> <p>Erik Bodin, Alexandru I. Stere, Dragos D Margineantu, Carl Henrik Ek, Henry Moss</p>	60405
<p>IDA-VLM: Towards Movie Understanding Via ID-Aware Large Vision-Language Model</p> <p>Yatai Ji, Shilong Zhang, Jie Wu, Peize Sun, Weifeng Chen, Xuefeng Xiao, Sidi Yang, Yujiu Yang, Ping Luo</p>	60435
<p>Towards Semantic Equivalence of Tokenization in Multimodal LLM</p> <p>Shengqiong Wu, Hao Fei, Xiangtai Li, Jiayi Ji, Hanwang Zhang, Tat-Seng Chua, Shuicheng Yan</p>	60449
<p>Outlier Synthesis Via Hamiltonian Monte Carlo for Out-Of-Distribution Detection</p> <p>Hengzhuang Li, Teng Zhang</p>	60474
<p>Chain-Of-Thought Provably Enables Learning the (Otherwise) Unlearnable</p> <p>Chenxiao Yang, Zhiyuan Li, David Wipf</p>	60514
<p>Training Nonlinear Transformers for Chain-Of-Thought Inference: A Theoretical Generalization Analysis</p> <p>Hongkang Li, Songtao Lu, Pin-Yu Chen, Xiaodong Cui, Meng Wang</p>	60542
<p>PQMass: Probabilistic Assessment of the Quality of Generative Models Using Probability Mass Estimation</p> <p>Pablo Lemos, Sammy Nasser Sharief, Nikolay Malkin, Salma Salhi, Connor Stone, Laurence Perreault-Levasseur, Yashar Hezaveh</p>	60581

<p>Relax and Merge: A Simple Yet Effective Framework for Solving Fair ℓ_1-Means and ℓ_1-Sparse Wasserstein Barycenter Problems</p> <p>Shihong Song, Guanlin Mo, Hu Ding</p>	60608
<p>CryoSPHERE: Single-Particle HETerogeneous REconstruction from Cryo EM</p> <p>Gabriel Ducrocq, Lukas Grunewald, Sebastian Westenhoff, Fredrik Lindsten</p>	60634
<p>SANA: Efficient High-Resolution Text-To-Image Synthesis with Linear Diffusion Transformers</p> <p>Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Haotian Tang, Yujun Lin, Zhekai Zhang, Muyang Li, Ligeng Zhu, Yao Lu, Song Han</p>	60648
<p>On the Almost Sure Convergence of the Stochastic Three Points Algorithm</p> <p>Taha El Bakkali El Kadi, Omar Saadi</p>	60673
<p>OS-ATLAS: Foundation Action Model for Generalist GUI Agents</p> <p>Zhiyong Wu, Zhenyu Wu, Fangzhi Xu, Yan Wang, Qiushi Sun, Chengyou Jia, Kanzhi Cheng, Zichen Ding, Liheng Chen, Paul Pu Liang, Yu Qiao</p>	60692
<p>Online Clustering with Nearly Optimal Consistency</p> <p>T-H. Hubert Chan, Shaofeng H.-C. Jiang, Tianyi Wu, Mengshi Zhao</p>	60711
<p>Balanced Neural ODEs: Nonlinear Model Order Reduction and Koopman Operator Approximations</p> <p>Julius Aka, Johannes Brunnemann, Jörg Eiden, Arne Speerforck, Lars Mikelsons</p>	60741
<p>TimeSuite: Improving MLLMs for Long Video Understanding Via Grounded Tuning</p> <p>Xiangyu Zeng, Kunchang Li, Chenting Wang, Xinhao Li, Tianxiang Jiang, Ziang Yan, Songze Li, Yansong Shi, Zhengrong Yue, Yi Wang, Yali Wang, Yu Qiao, Limin Wang</p>	60771
<p>Text2PDE: Latent Diffusion Models for Accessible Physics Simulation</p> <p>Anthony Zhou, Zijie Li, Michael Schneier, John R Buchanan Jr, Amir Barati Farimani</p>	60796
<p>Visual Agents as Fast and Slow Thinkers</p> <p>Guangyan Sun, Mingyu Jin, Zhenting Wang, Cheng-Long Wang, Siqi Ma, Qifan Wang, Tong Geng, Ying Nian Wu, Yongfeng Zhang, Dongfang Liu</p>	60826
<p>Learning to Help in Multi-Class Settings</p> <p>Yu Wu, Yansong Li, Zeyu Dong, Nitya Sathyavageeswaran, Anand D. Sarwate</p>	60852
<p>The Same but Different: Structural Similarities and Differences in Multilingual Language Modeling</p> <p>Ruochen Zhang, Qinan Yu, Matianyu Zang, Carsten Eickhoff, Ellie Pavlick</p>	60882
<p>Distilling Dataset into Neural Field</p> <p>Donghyeok Shin, Heesun Bae, Gyuwon Sim, Wanmo Kang, Il-Chul Moon</p>	60903
<p>RAG-SR: Retrieval-Augmented Generation for Neural Symbolic Regression</p> <p>Hengzhe Zhang, Qi Chen, Bing Xue, Wolfgang Banzhaf, Mengjie Zhang</p>	60935

<p>Causal Information Prioritization for Efficient Reinforcement Learning</p> <p>Hongye Cao, Fan Feng, Tianpei Yang, Jing Huo, Yang Gao</p>	60964
<p>Weighted Multi-Prompt Learning with Description-Free Large Language Model Distillation</p> <p>Sua Lee, Kyubum Shin, Jung Ho Park</p>	61001
<p>Efficient Diffusion Transformer Policies with Mixture of Expert Denoisers for Multitask Learning</p> <p>Moritz Reuss, Jyothish Pari, Pulkit Agrawal, Rudolf Lioutikov</p>	61017
<p>HARDMath: A Benchmark Dataset for Challenging Problems in Applied Mathematics</p> <p>Jingxuan Fan, Sarah Martinson, Erik Y. Wang, Kaylie Hausknecht, Jonah Brenner, Danxian Liu, Nianli Peng, Corey Wang, Michael Brenner</p>	61046
<p>Is Your Model Really a Good Math Reasoner? Evaluating Mathematical Reasoning with Checklist</p> <p>Zihao Zhou, Shudong Liu, Maizhen Ning, Wei Liu, Jindong Wang, Derek F. Wong, Xiaowei Huang, Qiufeng Wang, Kaizhu Huang</p>	61080
<p>Many-Objective Multi-Solution Transport</p> <p>Ziyue Li, Tian Li, Virginia Smith, Jeff Bilmes, Tianyi Zhou</p>	61124
<p>Rectified Diffusion: Straightness is Not Your Need in Rectified Flow</p> <p>Fu-Yun Wang, Ling Yang, Zhaoyang Huang, Mengdi Wang, Hongsheng Li</p>	61150
<p>AdaWM: Adaptive World Model Based Planning for Autonomous Driving</p> <p>Hang Wang, Xin Ye, Feng Tao, Chenbin Pan, Abhirup Mallik, Burhaneddin Yaman, Liu Ren, Junshan Zhang</p>	61176
<p>Measuring and Improving Persuasiveness of Large Language Models</p> <p>Somesh Kumar Singh, Yaman Kumar Singla, Harini S I, Balaji Krishnamurthy</p>	61201
<p>Scale-Free Graph-Language Models</p> <p>Jianglin Lu, Yixuan Liu, Yitian Zhang, Yun Fu</p>	61257
<p>Probabilistic Conformal Prediction with Approximate Conditional Validity</p> <p>Vincent Plassier, Alexander Fishkov, Mohsen Guizani, Maxim Panov, Eric Moulines</p>	61272
<p>ACC-Collab: An Actor-Critic Approach to Multi-Agent LLM Collaboration</p> <p>Andrew Estornell, Jean-Francois Ton, Yuanshun Yao, Yang Liu</p>	61305
<p>YouTube-SL-25: A Large-Scale, Open-Domain Multilingual Sign Language Parallel Corpus</p> <p>Garrett Tanzer, Biao Zhang</p>	61329
<p>Generalization in VAE and Diffusion Models: A Unified Information-Theoretic Analysis</p> <p>Qi Chen, Jierui Zhu, Florian Shkurti</p>	61343
<p>Vision-RWKV: Efficient and Scalable Visual Perception with RWKV-Like Architectures</p> <p>Yuchen Duan, Weiyun Wang, Zhe Chen, Xizhou Zhu, Lewei Lu, Tong Lu, Yu Qiao, Hongsheng Li, Jifeng Dai, Wenhai Wang</p>	61381

<p>How New Data Permeates LLM Knowledge and How to Dilute it</p> <p><i>Chen Sun, Renat Aksitov, Andrey Zhmoginov, Nolan Andrew Miller, Max Vladymyrov, Ulrich Rueckert, Been Kim, Mark Sandler</i></p>	61398
<p>The Unreasonable Ineffectiveness of the Deeper Layers</p> <p><i>Andrey Gromov, Kushal Tirumala, Hassan Shapourian, Paolo Glorioso, Dan Roberts</i></p>	61431
<p>On the Learn-To-Optimize Capabilities of Transformers in In-Context Sparse Recovery</p> <p><i>Renpu Liu, Ruida Zhou, Cong Shen, Jing Yang</i></p>	61446
<p>Long-Short Decision Transformer: Bridging Global and Local Dependencies for Generalized Decision-Making</p> <p><i>Jincheng Wang, Penny Karanasou, Pengyuan Wei, Elia Gatti, Diego Martinez Plasencia, Dimitrios Kanoulas</i></p>	61481
<p>KinFormer: Generalizable Dynamical Symbolic Regression for Catalytic Organic Reaction Kinetics</p> <p><i>Jindou Chen, Jidong Tian, Liang Wu, Chenxinwei, Xiaokang Yang, Yaohui Jin, Yanyan Xu</i></p>	61506
<p>It Helps to Take a Second Opinion: Teaching Smaller LLMs to Deliberate Mutually Via Selective Rationale Optimisation</p> <p><i>Sohan Patnaik, Milan Aggarwal, Sumit Bhatia, Balaji Krishnamurthy</i></p>	61529
<p>Shifting the Paradigm: A Diffeomorphism Between Time Series Data Manifolds for Achieving Shift-Invariancy in Deep Learning</p> <p><i>Berken Utku Demirel, Christian Holz</i></p>	61554
<p>Misspecified \mathcal{Q}^*-Learning with Sparse Linear Function Approximation: Tight Bounds on Approximation Error</p> <p><i>Ally Yalei Du, Lin Yang, Ruosong Wang</i></p>	61594
<p>Plastic Learning with Deep Fourier Features</p> <p><i>Alex Lewandowski, Dale Schuurmans, Marlos C. Machado</i></p>	61614
<p>RefactorBench: Evaluating Stateful Reasoning in Language Agents Through Code</p> <p><i>Dhruv Gautam, Spandan Garg, Jinu Jang, Neel Sundaresan, Roshanak Zilouchian Moghaddam</i></p>	61640
<p>NeRAF: 3D Scene Infused Neural Radiance and Acoustic Fields</p> <p><i>Amandine Brunetto, Sascha Hornauer, Fabien Moutarde</i></p>	61672
<p>Formation of Representations in Neural Networks</p> <p><i>Liu Ziyin, Isaac L. Chuang, Tomer Galanti, Tomaso A Poggio</i></p>	61696
<p>Dynamic Modeling of Patients, Modalities and Tasks Via Multi-Modal Multi-Task Mixture of Experts</p> <p><i>Chenwei Wu, Zitao Shuai, Zhengxu Tang, Luning Wang, Liyue Shen</i></p>	61725

<p>Hyperbolic Genome Embeddings</p> <p>Raiyan R. Khan, Philippe Chlenski, Itsik Pe'Er</p>	61740
<p>Nonconvex Stochastic Optimization Under Heavy-Tailed Noises: Optimal Convergence Without Gradient Clipping</p> <p>Zijian Liu, Zhengyuan Zhou</p>	61770
<p>A Watermark for Order-Agnostic Language Models</p> <p>Ruibo Chen, Yihan Wu, Yanshuo Chen, Chenxi Liu, Junfeng Guo, Heng Huang</p>	61796
<p>Conditional Diffusion Models Are Minimax-Optimal and Manifold-Adaptive for Conditional Distribution Estimation</p> <p>Rong Tang, Lizhen Lin, Yun Yang</p>	61813
<p>Understanding Virtual Nodes: Oversquashing and Node Heterogeneity</p> <p>Joshua Southern, Francesco Di Giovanni, Michael M. Bronstein, Johannes F. Lutzeyer</p>	61826
<p>Neural Causal Graph for Interpretable and Intervenable Classification</p> <p>Jiawei Wang, Shaofei Lu, Da Cao, Dongyu Wang, Yuquan Le, Zhe Quan, Tat-Seng Chua</p>	61855
<p>MAP: Multi-Human-Value Alignment Palette</p> <p>Xinran Wang, Qi Le, Ammar Ahmed, Enmao Diao, Yi Zhou, Nathalie Baracaldo, Jie Ding, Ali Anwar</p>	61876
<p>ONLINE EPSILON NET & PIERCING SET FOR GEOMETRIC CONCEPTS</p> <p>Sujoy Bhore, Devdan Dey, Satyam Singh</p>	61906
<p>ViBiDSampler: Enhancing Video Interpolation Using Bidirectional Diffusion Sampler</p> <p>Serin Yang, Taesung Kwon, Jong Chul Ye</p>	61926
<p>Limits to Scalable Evaluation at the Frontier: LLM as Judge Won't Beat Twice the Data</p> <p>Florian E. Dörner, Vivian Yvonne Nastl, Moritz Hardt</p>	61945
<p>Multi-Scale Fusion for Object Representation</p> <p>Rongzhen Zhao, Vivienne Huiling Wang, Juho Kannala, Joni Pajarinen</p>	61970
<p>Learning Mask Invariant Mutual Information for Masked Image Modeling</p> <p>Tao Huang, Yanxiang Ma, Shan You, Chang Xu</p>	61987
<p>Chain-Of-Focus Prompting: Leveraging Sequential Visual Cues to Prompt Large Autoregressive Vision Models</p> <p>Jiyang Zheng, Jialiang Shen, Yu Yao, Min Wang, Yang Yang, Dadong Wang, Tongliang Liu</p>	62006
<p>MeToken: Uniform Micro-Environment Token Boosts Post-Translational Modification Prediction</p> <p>Cheng Tan, Zhenxiao Cao, Zhangyang Gao, Lirong Wu, Siyuan Li, Yufei Huang, Jun Xia, Bozhen Hu, Stan Z. Li</p>	62026
<p>CBraMod: A Criss-Cross Brain Foundation Model for EEG Decoding</p> <p>Jiquan Wang, Sha Zhao, Zhiling Luo, Yangxuan Zhou, Haiteng Jiang, Shijian Li, Tao Li, Gang Pan</p>	62056

<p>Zero-Shot Imputation with Foundation Inference Models for Dynamical Systems</p> <p><i>Patrick Seifner, Kostadin Cvejoski, Antonia Körner, Ramses J Sanchez</i></p>	62093
<p>ProtPainter: Draw Or Drag Protein Via Topology-Guided Diffusion</p> <p><i>Zhengxi Lu, Shizhuo Cheng, Tintin Jiang, Yan Zhang, Min Zhang</i></p>	62140
<p>Preserving Diversity in Supervised Fine-Tuning of Large Language Models</p> <p><i>Ziniu Li, Congliang Chen, Tian Xu, Zeyu Qin, Jiancong Xiao, Zhi-Quan Luo, Ruoyu Sun</i></p>	62165
<p>Towards Marginal Fairness Sliced Wasserstein Barycenter</p> <p><i>Khai Nguyen, Hai Nguyen, Nhat Ho</i></p>	62193
<p>Constraint-Conditioned Actor-Critic for Offline Safe Reinforcement Learning</p> <p><i>Zijian Guo, Weichao Zhou, Shengao Wang, Wenchao Li</i></p>	62222
<p>Chunk-Distilled Language Modeling</p> <p><i>Yanhong Li, Karen Livescu, Jiawei Zhou</i></p>	62244
<p>Beyond Autoregression: Discrete Diffusion for Complex Reasoning and Planning</p> <p><i>Jiacheng Ye, Jiahui Gao, Shansan Gong, Lin Zheng, Xin Jiang, Zhenguo Li, Lingpeng Kong</i></p>	62282
<p>Efficiently Learning at Test-Time: Active Fine-Tuning of LLMs</p> <p><i>Jonas Hübötter, Sascha Bongni, Ido Hakimi, Andreas Krause</i></p>	62306
<p>DynFr: An Efficient Framework for Machine Unlearning in Random Forest</p> <p><i>Shurong Wang, Zhuoyang Shen, Xinbao Qiao, Tongning Zhang, Meng Zhang</i></p>	62364
<p>Aligned LLMs Are Not Aligned Browser Agents</p> <p><i>Priyanshu Kumar, Elaine Lau, Saranya Vijayakumar, Tu Trinh, Elaine T Chang, Vaughn Robinson, Shuyan Zhou, Matt Fredrikson, Sean M. Hendryx, Summer Yue, Zifan Wang</i></p>	62386
<p>EffoVPR: Effective Foundation Model Utilization for Visual Place Recognition</p> <p><i>Issar Tzachor, Boaz Lerner, Matan Levy, Michael Green, Tal Berkovitz Shalev, Gavriel Habib, Dvir Samuel, Noam Korngut Zailer, Or Shimshi, Nir Darshan, Rami Ben-Ari</i></p>	62408
<p>CAMEx: Curvature-Aware Merging of Experts</p> <p><i>Viet Dung Nguyen, Minh Nguyen Hoang, Luc Nguyen, Rachel Teo, Tan Minh Nguyen, Linh Duy Tran</i></p>	62431
<p>Neuron-Based Multifractal Analysis of Neuron Interaction Dynamics in Large Models</p> <p><i>Xiongye Xiao, Heng Ping, Chenyu Zhou, Defu Cao, Yaxing Li, Yi-Zhuo Zhou, Shixuan Li, Nikos Kanakaris, Paul Bogdan</i></p>	62457
<p>Towards Automated Knowledge Integration from Human-Interpretable Representations</p> <p><i>Kasia Kobalczyk, Mihaela Van Der Schaar</i></p>	62505
<p>Optimal Flow Transport and Its Entropic Regularization: A GPU-Friendly Matrix Iterative Algorithm for Flow Balance Satisfaction</p> <p><i>Liangliang Shi, Yufeng Li, Kaipeng Zeng, Yihui Tu, Junchi Yan</i></p>	62541

<p>From Search to Sampling: Generative Models for Robust Algorithmic Recourse</p> <p><i>Prateek Garg, Lokesh Nagalapatti, Sunita Sarawagi</i></p>	62560
<p>ActionReasoningBench: Reasoning About Actions with and Without Ramification Constraints</p> <p><i>Divij Handa, Pavel Dolin, Shrinidhi Kumbhar, Tran Cao Son, Chitta Baral</i></p>	62579
<p>Equivariant Masked Position Prediction for Efficient Molecular Representation</p> <p><i>Junyi An, Chao Qu, Yun-Fei Shi, Xinhao Liu, Qianwei Tang, Fenglei Cao, Yuan Qi</i></p>	62619
<p>Where Am I and What Will I See: An Auto-Regressive Model for Spatial Localization and View Prediction</p> <p><i>Junyi Chen, Di Huang, Weicai Ye, Wanli Ouyang, Tong He</i></p>	62643
<p>Decoupled Graph Energy-Based Model for Node Out-Of-Distribution Detection on Heterophilic Graphs</p> <p><i>Yuhan Chen, Yihong Luo, Yifan Song, Pengwen Dai, Jing Tang, Xiaochun Cao</i></p>	62656
<p>Stochastic Polyak Step-Sizes and Momentum: Convergence Guarantees and Practical Performance</p> <p><i>Dimitris Oikonomou, Nicolas Loizou</i></p>	62679
<p>TRENDy: Temporal Regression of Effective Nonlinear Dynamics</p> <p><i>Matt Ricci, Guy Pelc, Zoe Piran, Noa Moriel, Mor Nitzan</i></p>	62719
<p>Direct Distributional Optimization for Provable Alignment of Diffusion Models</p> <p><i>Ryotaro Kawata, Kazusato Oko, Atsushi Nitanda, Taiji Suzuki</i></p>	62743
<p>Evaluating Semantic Variation in Text-To-Image Synthesis: A Causal Perspective</p> <p><i>Xiangru Zhu, Penglei Sun, Yaoxian Song, Yanghua Xiao, Zhixu Li, Chengyu Wang, Jun Huang, Bei Yang, Xiaoxiao Xu</i></p>	62792
<p>Graph Neural Networks Are More than Filters: Revisiting and Benchmarking from a Spectral Perspective</p> <p><i>Yushun Dong, Patrick Soga, Yinhan He, Song Wang, Jundong Li</i></p>	62841
<p>Artificial Kuramoto Oscillatory Neurons</p> <p><i>Takeru Miyato, Sindy Löwe, Andreas Geiger, Max Welling</i></p>	62863
<p>Theory on Score-Mismatched Diffusion Models and Zero-Shot Conditional Samplers</p> <p><i>Yuchen Liang, Peizhong Ju, Yingbin Liang, Ness Shroff</i></p>	62908
<p>Signature Kernel Conditional Independence Tests in Causal Discovery for Stochastic Processes</p> <p><i>Georg Manten, Cecilia Casolo, Emilio Ferrucci, Søren Wengel Mogensen, Cristopher Salvi, Niki Kilbertus</i></p>	62970
<p>Revisiting Source-Free Domain Adaptation: A New Perspective Via Uncertainty Control</p> <p><i>Gezheng Xu, Hui Guo, Li Yi, Charles Ling, Boyu Wang, Grace Yi</i></p>	63007

<p>REGENT: A Retrieval-Augmented Generalist Agent that Can Act In-Context in New Environments</p> <p><i>Kaustubh Sridhar, Souradeep Dutta, Dinesh Jayaraman, Insup Lee</i></p>	63047
<p>On the Expressive Power of Sparse Geometric MPNNs</p> <p><i>Yonatan Sverdlov, Nadav Dym</i></p>	63077
<p>Ctrl-Adapter: An Efficient and Versatile Framework for Adapting Diverse Controls to Any Diffusion Model</p> <p><i>Han Lin, Jaemin Cho, Abhay Zala, Mohit Bansal</i></p>	63105
<p>Tree-Wasserstein Distance for High Dimensional Data with a Latent Feature Hierarchy</p> <p><i>Ya-Wei Eileen Lin, Ronald R. Coifman, Gal Mishne, Ronen Talmon</i></p>	63150
<p>Large-Scale and Fine-Grained Vision-Language Pre-Training for Enhanced CT Image Understanding</p> <p><i>Zhongyi Shui, Jianpeng Zhang, Weiwei Cao, Sinuo Wang, Ruizhe Guo, Le Lu, Lin Yang, Xianghua Ye, Tingbo Liang, Qi Zhang, Ling Zhang</i></p>	63188
<p>InversionGNN: A Dual Path Network for Multi-Property Molecular Optimization</p> <p><i>Yifan Niu, Ziqi Gao, Tingyang Xu, Yang Liu, Yatao Bian, Yu Rong, Junzhou Huang, Jia Li</i></p>	63202
<p>An Effective Manifold-Based Optimization Method for Distributionally Robust Classification</p> <p><i>Jiawei Huang, Hu Ding</i></p>	63225
<p>ContextGNN: Beyond Two-Tower Recommendation Systems</p> <p><i>Yiwen Yuan, Zecheng Zhang, Xinwei He, Akihiro Nitta, Weihua Hu, Manan Shah, Blaž Stojanovic, Shenyang Huang, Jan Eric Lenssen, Jure Leskovec, Matthias Fey</i></p>	63264
<p>TPO: Aligning Large Language Models with Multi-Branch & Multi-Step Preference Trees</p> <p><i>Weibin Liao, Xu Chu, Yasha Wang</i></p>	63277
<p>Internet of Agents: Weaving a Web of Heterogeneous Agents for Collaborative Intelligence</p> <p><i>Weize Chen, Ziming You, Ran Li, Yitong Guan, Chen Qian, Chenyang Zhao, Cheng Yang, Ruobing Xie, Zhiyuan Liu, Maosong Sun</i></p>	63300
<p>ANaGRAM: A Natural Gradient Relative to Adapted Model for Efficient PINNs Learning</p> <p><i>Nilo Schwencke, Cyril Furtlehner</i></p>	63338
<p>CAX: Cellular Automata Accelerated in JAX</p> <p><i>Maxence Faldor, Antoine Cully</i></p>	63383
<p>EC-Diffuser: Multi-Object Manipulation Via Entity-Centric Behavior Generation</p> <p><i>Carl Qi, Dan Haramati, Tal Daniel, Aviv Tamar, Amy Zhang</i></p>	63397
<p>Few for Many: Tchebycheff Set Scalarization for Many-Objective Optimization</p> <p><i>Xi Lin, Yilu Liu, Xiaoyuan Zhang, Fei Liu, Zhenkun Wang, Qingfu Zhang</i></p>	63421

<p>ChartMoE: Mixture of Diversely Aligned Expert Connector for Chart Understanding</p> <p>Zhengzhuo Xu, Bowen Qu, Yiyan Qi, Sinan Du, Chengjin Xu, Chun Yuan, Jian Guo</p>	63452
<p>Strategic Classification with Externalities</p> <p>Safwan Hossain, Evi Micha, Yiling Chen, Ariel D. Procaccia</p>	63475
<p>Show-O: One Single Transformer to Unify Multimodal Understanding and Generation</p> <p>Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, Mike Zheng Shou</p>	63499
<p>Bayesian Treatment of the Spectrum of the Empirical Kernel in (Sub)Linear-Width Neural Networks</p> <p>Ouns El Harzli, Bernardo Cuenca Grau</p>	63524
<p>Looking into User's Long-Term Interests Through the Lens of Conservative Evidential Learning</p> <p>Dingrong Wang, Krishna Prasad Neupane, Ervine Zheng, Qi Yu</p>	63538
<p>On the Benefits of Memory for Modeling Time-Dependent PDEs</p> <p>Ricardo Buitrago, Tanya Marwah, Albert Gu, Andrej Risteski</p>	63567
<p>Arithmetic Without Algorithms: Language Models Solve Math with a Bag of Heuristics</p> <p>Yaniv Nikankin, Anja Reusch, Aaron Mueller, Yonatan Belinkov</p>	63598
<p>Maintaining Structural Integrity in Parameter Spaces for Parameter Efficient Fine-Tuning</p> <p>Chongjie Si, Xuehui Wang, Xue Yang, Zhengqin Xu, Qingyun Li, Jifeng Dai, Yu Qiao, Xiaokang Yang, Wei Shen</p>	63625
<p>SimpleTM: A Simple Baseline for Multivariate Time Series Forecasting</p> <p>Hui Chen, Viet Luong, Lopamudra Mukherjee, Vikas Singh</p>	63646
<p>Inverse Attention Agents for Multi-Agent Systems</p> <p>Qian Long, Ruoyan Li, Minglu Zhao, Tao Gao, Demetri Terzopoulos</p>	63672
<p>Model-Based Offline Reinforcement Learning with Lower Expectile Q-Learning</p> <p>Kwanyoung Park, Youngwoon Lee</p>	63690
<p>Scalable Mechanistic Neural Networks</p> <p>Jiale Chen, Dingling Yao, Adeel Pervez, Dan Alistarh, Francesco Locatello</p>	63716
<p>Erasing Concept Combination from Text-To-Image Diffusion Model</p> <p>Hongyi Nie, Quanming Yao, Yang Liu, Zhen Wang, Yatao Bian</p>	63738
<p>SparsyFed: Sparse Adaptive Federated Learning</p> <p>Adriano Guastella, Lorenzo Sani, Alex Iacob, Alessio Mora, Paolo Bellavista, Nicholas Donald Lane</p>	63759
<p>Minimax Optimal Two-Stage Algorithm for Moment Estimation Under Covariate Shift</p> <p>Zhen Zhang, Xin Liu, Shaoli Wang, Jiaye Teng</p>	63792

<p>SmartRAG: Jointly Learn RAG-Related Tasks from the Environment Feedback</p> <p><i>Jingsheng Gao, Linxu Li, Ke Ji, Weiyuan Li, Yixin Lian, Yuzhuo Fu, Bin Dai</i></p>	63821
<p>Analytic DAG Constraints for Differentiable DAG Learning</p> <p><i>Zhen Zhang, Ignavier Ng, Dong Gong, Yuhang Liu, Mingming Gong, Biwei Huang, Kun Zhang, Anton Van Den Hengel, Javen Qinfeng Shi</i></p>	63845
<p>ZAPBench: A Benchmark for Whole-Brain Activity Prediction in Zebrafish</p> <p><i>Jan-Matthis Lueckmann, Alexander Immer, Alex Bo-Yuan Chen, Peter H. Li, Mariela D Petkova, Nirmla A Iyer, Luuk Willem Hesselink, Aparna Dev, Gudrun Ihrke, Woohyun Park, Alyson Petruncio, Aubrey Weigel, Wyatt Korff, Florian Engert, Jeff Lichtman, Misha Ahrens, Michal Januszewski, Viren Jain</i></p>	63871
<p>Meta-Continual Learning of Neural Fields</p> <p><i>Seungyoon Woo, Junhyeog Yun, Gunhee Kim</i></p>	63909
<p>Generative Classifiers Avoid Shortcut Solutions</p> <p><i>Alexander Cong Li, Ananya Kumar, Deepak Pathak</i></p>	63932
<p>Deep Random Features for Scalable Interpolation of Spatiotemporal Data</p> <p><i>Weibin Chen, Azhir Mahmood, Michel Tsamados, So Takao</i></p>	63953
<p>Physics of Language Models: Part 3.2, Knowledge Manipulation</p> <p><i>Zeyuan Allen-Zhu, Yuanzhi Li</i></p>	63984
<p>Oryx MLLM: On-Demand Spatial-Temporal Understanding at Arbitrary Resolution</p> <p><i>Zuyan Liu, Yuhao Dong, Ziwei Liu, Winston Hu, Jiwen Lu, Yongming Rao</i></p>	63993
<p>DarkBench: Benchmarking Dark Patterns in Large Language Models</p> <p><i>Esben Kran, Hieu Minh Nguyen, Akash Kundu, Sami Jawhar, Jinsuk Park, Mateusz Maria Jurewicz</i></p>	64016
<p>SINGAPO: Single Image Controlled Generation of Articulated Parts in Objects</p> <p><i>Jiayi Liu, Denys Iliash, Angel X Chang, Manolis Savva, Ali Mahdavi Amiri</i></p>	64037
<p>Moner: Motion Correction in Undersampled Radial MRI with Unsupervised Neural Representation</p> <p><i>Qing Wu, Chenhe Du, Xuanyu Tian, Jingyi Yu, Yuyao Zhang, Hongjiang Wei</i></p>	64059
<p>A Sanity Check for AI-Generated Image Detection</p> <p><i>Shilin Yan, Ouxiang Li, Jiayin Cai, Yanbin Hao, Xiaolong Jiang, Yao Hu, Weidi Xie</i></p>	64081
<p>Nonasymptotic Analysis of Stochastic Gradient Descent with the Richardson–Romberg Extrapolation</p> <p><i>Marina Sheshukova, Denis Belomestny, Alain Oliviero Durmus, Eric Moulines, Alexey Naumov, Sergey Samsonov</i></p>	64100
<p>SSLAM: Enhancing Self-Supervised Models with Audio Mixtures for Polyphonic Soundscapes</p> <p><i>Tony Alex, Sara Atito, Armin Mustafa, Muhammad Awais, Philip J B Jackson</i></p>	64131

<p>Promptriever: Instruction-Trained Retrievers Can Be Prompted Like Language Models</p> <p>Orion Weller, Benjamin Van Durme, Dawn Lawrie, Ashwin Paranjape, Yuhao Zhang, Jack Hessel</p>	64150
<p>Towards Scalable Exact Machine Unlearning Using Parameter-Efficient Fine-Tuning</p> <p>Somnath Basu Roy Chowdhury, Krzysztof Marcin Choromanski, Arijit Sehanobish, Kumar Avinava Dubey, Snigdha Chaturvedi</p>	64174
<p>Geometry-Aware Approaches for Balancing Performance and Theoretical Guarantees in Linear Bandits</p> <p>Yuwei Luo, Mohsen Bayati</p>	64201
<p>Samba: Synchronized Set-Of-Sequences Modeling for Multiple Object Tracking</p> <p>Mattia Segu, Luigi Piccinelli, Siyuan Li, Yung-Hsu Yang, Luc Van Gool, Bernt Schiele</p>	64213
<p>Counterfactual Generative Modeling with Variational Causal Inference</p> <p>Yulun Wu, Louis McConnell, Claudia Iriondo</p>	64227
<p>Rethinking Classifier Re-Training in Long-Tailed Recognition: Label Over-Smooth Can Balance</p> <p>Siyu Sun, Han Lu, Jiangtong Li, Yichen Xie, Tianjiao Li, Xiaokang Yang, Liqing Zhang, Junchi Yan</p>	64262
<p>OMG: Opacity Matters in Material Modeling with Gaussian Splatting</p> <p>Silong Yong, Venkata Nagarjun Pudureddyur Manivannan, Bernhard Kerbl, Zifu Wan, Simon Stepputtis, Katia P. Sycara, Yaqi Xie</p>	64279
<p>TIS-DPO: Token-Level Importance Sampling for Direct Preference Optimization with Estimated Weights</p> <p>Aiwei Liu, Haoping Bai, Zhiyun Lu, Yanchao Sun, Xiang Kong, Xiaoming Simon Wang, Jiulong Shan, Albin Madappally Jose, Xiaojiang Liu, Lijie Wen, Philip S. Yu, Meng Cao</p>	64298
<p>On the Byzantine-Resilience of Distillation-Based Federated Learning</p> <p>Christophe Roux, Max Zimmer, Sebastian Pokutta</p>	64328
<p>Think-On-Graph 2.0: Deep and Faithful Large Language Model Reasoning with Knowledge-Guided Retrieval Augmented Generation</p> <p>Shengjie Ma, Chengjin Xu, Xuhui Jiang, Muzhi Li, Huaren Qu, Cehao Yang, Jiabin Mao, Jian Guo</p>	64354
<p>FlexPrefill: A Context-Aware Sparse Attention Mechanism for Efficient Long-Sequence Inference</p> <p>Xunhao Lai, Jianqiao Lu, Yao Luo, Yiyuan Ma, Xun Zhou</p>	64379
<p>Resolution Attack: Exploiting Image Compression to Deceive Deep Neural Networks</p> <p>Wangjia Yu, Xiaomeng Fu, Qiao Li, Jizhong Han, Xiaodan Zhang</p>	64406
<p>Provably Reliable Conformal Prediction Sets in the Presence of Data Poisoning</p> <p>Yan Scholten, Stephan Günnemann</p>	64437
<p>Procedural Synthesis of Synthesizable Molecules</p> <p>Michael Sun, Alston Lo, Minghao Guo, Jie Chen, Connor W. Coley, Wojciech Matusik</p>	64462

<p>ColPali: Efficient Document Retrieval with Vision Language Models</p> <p>Manuel Faysse, Hugues Sibille, Tony Wu, Bilel Omrani, Gautier Viaud, Celine Hudelot, Pierre Colombo</p>	64477
<p>Loss Landscape of Shallow ReLU-Like Neural Networks: Stationary Points, Saddle Escape, and Network Embedding</p> <p>Zhengqing Wu, Berfin Simsek, François Gaston Ged</p>	64503
<p>SpinQuant: LLM Quantization with Learned Rotations</p> <p>Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, Tijmen Blankevoort</p>	64546
<p>Conflict-Averse Gradient Aggregation for Constrained Multi-Objective Reinforcement Learning</p> <p>Dohyeong Kim, Mineui Hong, Jeongho Park, Songhwa Oh</p>	64570
<p>Towards Unified Human Motion-Language Understanding Via Sparse Interpretable Characterization</p> <p>Guangtao Lyu, Chenghao Xu, Jiexi Yan, Muli Yang, Cheng Deng</p>	64600
<p>Wavelet-Based Positional Representation for Long Context</p> <p>Yui Oka, Taku Hasegawa, Kyosuke Nishida, Kuniko Saito</p>	64625
<p>ET-SEED: EFFICIENT TRAJECTORY-LEVEL SE(3) EQUIVARIANT DIFFUSION POLICY</p> <p>Chenrui Tie, Yue Chen, Ruihai Wu, Boxuan Dong, Zeyi Li, Chongkai Gao, Hao Dong</p>	64653
<p>Large (Vision) Language Models Are Unsupervised In-Context Learners</p> <p>Artyom Gadetsky, Andrei Atanov, Yulun Jiang, Zhitong Gao, Ghazal Hosseini Mighan, Amir Zamir, Maria Brbic</p>	64672
<p>Robotouille: An Asynchronous Planning Benchmark for LLM Agents</p> <p>Gonzalo Gonzalez-Pumariega, Leong Su Yean, Neha Sunkara, Sanjiban Choudhury</p>	64697
<p>Programming Refusal with Conditional Activation Steering</p> <p>Bruce W. Lee, Inkit Padhi, Karthikeyan Natesan Ramamurthy, Erik Miehl, Pierre Dognin, Manish Nagireddy, Amit Dhurandhar</p>	64733
<p>Not All LLM-Generated Data Are Equal: Rethinking Data Weighting in Text Classification</p> <p>Hsun-Yu Kuo, Yin-Hsiang Liao, Yu-Chieh Chao, Wei-Yun Ma, Pu-Jen Cheng</p>	64759
<p>Distance-Based Tree-Sliced Wasserstein Distance</p> <p>Hoang V. Tran, Minh-Khoi Nguyen-Nhat, Huyen Trang Pham, Thanh Chu, Tam Le, Tan Minh Nguyen</p>	64784
<p>A Theoretically-Principled Sparse, Connected, and Rigid Graph Representation of Molecules</p> <p>Shih-Hsin Wang, Yuhao Huang, Justin M. Baker, Yuan-En Sun, Qi Tang, Bao Wang</p>	64814
<p>FaceShot: Bring Any Character into Life</p> <p>Junyao Gao, Yanan Sun, Fei Shen, Xin Jiang, Zhening Xing, Kai Chen, Cairong Zhao</p>	64840

<p>A Theoretical Framework for Partially-Observed Reward States in RLHF</p> <p><i>Chinmaya Kausik, Mirco Mutti, Aldo Pacchiano, Ambuj Tewari</i></p>	64859
<p>OpenHands: An Open Platform for AI Software Developers as Generalist Agents</p> <p><i>Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan, Hao Peng, Heng Ji, Graham Neubig</i></p>	64910
<p>Adversarial Attacks on Data Attribution</p> <p><i>Xinhe Wang, Pingbang Hu, Junwei Deng, Jiaqi W. Ma</i></p>	64948
<p>Variance-Reducing Couplings for Random Features</p> <p><i>Isaac Reid, Stratis Markou, Krzysztof Marcin Choromanski, Richard E. Turner, Adrian Weller</i></p>	64970
<p>SIM: Surface-Based fMRI Analysis for Inter-Subject Multimodal Decoding from Movie-Watching Experiments</p> <p><i>Simon Dahan, Gabriel Bénédict, Logan Zane John Williams, Yourong Guo, Daniel Rueckert, Robert Leech, Emma Claire Robinson</i></p>	65010
<p>SageAttention: Accurate 8-Bit Attention for Plug-And-Play Inference Acceleration</p> <p><i>Jintao Zhang, Jia Wei, Pengle Zhang, Jun Zhu, Jianfei Chen</i></p>	65038
<p>Beyond Content Relevance: Evaluating Instruction Following in Retrieval Models</p> <p><i>Jianqun Zhou, Yuanlei Zheng, Wei Chen, Qianqian Zheng, Shang Zeyuan, Wei Zhang, Rui Meng, Xiaoyu Shen</i></p>	65058
<p>One Step Diffusion Via Shortcut Models</p> <p><i>Kevin Frans, Danijar Hafner, Sergey Levine, Pieter Abbeel</i></p>	65090
<p>Steering Masked Discrete Diffusion Models Via Discrete Denoising Posterior Prediction</p> <p><i>Jarrid Rector-Brooks, Mohsin Hasan, Zhangzhi Peng, Cheng-Hao Liu, Sarthak Mittal, Nouha Dziri, Michael M. Bronstein, Pranam Chatterjee, Alexander Tong, Joey Bose</i></p>	65107
<p>Knowledge Graph Finetuning Enhances Knowledge Manipulation in Large Language Models</p> <p><i>Hanzhu Chen, Xu Shen, Jie Wang, Zehao Wang, Qitan Lv, Junjie He, Rong Wu, Feng Wu, Jieping Ye</i></p>	65139
<p>Fast Direct: Query-Efficient Online Black-Box Guidance for Diffusion-Model Target Generation</p> <p><i>Kim Yong Tan, Yueming Lyu, Ivor Tsang, Yew-Soon Ong</i></p>	65162
<p>Achieving Dimension-Free Communication in Federated Learning Via Zeroth-Order Optimization</p> <p><i>Zhe Li, Bicheng Ying, Zidong Liu, Chaosheng Dong, Haibo Yang</i></p>	65200
<p>Self-Supervised Contrastive Learning Performs Non-Linear System Identification</p> <p><i>Rodrigo González Laiz, Tobias Schmidt, Steffen Schneider</i></p>	65235
<p>Equivariant Denoisers Cannot Copy Graphs: Align Your Graph Diffusion Models</p> <p><i>Najwa Laabid, Severi Rissanen, Markus Heinonen, Arno Solin, Vikas Garg</i></p>	65273

<p>Non-Myopic Generation of Language Models for Reasoning and Planning</p> <p>Chang Ma, Haiteng Zhao, Junlei Zhang, Junxian He, Lingpeng Kong</p>	65312
<p>Stable Segment Anything Model</p> <p>Qi Fan, Xin Tao, Lei Ke, Mingqiao Ye, Di Zhang, Pengfei Wan, Yu-Wing Tai, Chi-Keung Tang</p>	65340
<p>Kernel-Based Optimally Weighted Conformal Time-Series Prediction</p> <p>Jonghyeok Lee, Chen Xu, Yao Xie</p>	65354
<p>DS-LLM: Leveraging Dynamical Systems to Enhance Both Training and Inference of Large Language Models</p> <p>Ruibing Song, Chuan Liu, Chunshu Wu, Ang Li, Dongfang Liu, Ying Nian Wu, Tong Geng</p>	65379
<p>TokenFormer: Rethinking Transformer Scaling with Tokenized Model Parameters</p> <p>Haiyang Wang, Yue Fan, Muhammad Ferjad Naeem, Yongqin Xian, Jan Eric Lenssen, Liwei Wang, Federico Tombari, Bernt Schiele</p>	65395
<p>DreamDistribution: Learning Prompt Distribution for Diverse In-Distribution Generation</p> <p>Brian Nlong Zhao, Yuhang Xiao, Jiashu Xu, Xinyang Jiang, Yifan Yang, Dongsheng Li, Laurent Itti, Vibhav Vineet, Yunhao Ge</p>	65413
<p>KaSA: Knowledge-Aware Singular-Value Adaptation of Large Language Models</p> <p>Fan Wang, Juyong Jiang, Chansung Park, Sunghun Kim, Jing Tang</p>	65444
<p>MaestroMotif: Skill Design from Artificial Intelligence Feedback</p> <p>Martin Klissarov, Mikael Henaff, Roberta Raileanu, Shagun Sodhani, Pascal Vincent, Amy Zhang, Pierre-Luc Bacon, Doina Precup, Marlos C. Machado, Pierluca D'Oro</p>	65483
<p>Track-On: Transformer-Based Online Point Tracking with Memory</p> <p>Görkay Aydemir, Xiongyi Cai, Weidi Xie, Fatma Guney</p>	65512
<p>Earlier Tokens Contribute More: Learning Direct Preference Optimization from Temporal Decay Perspective</p> <p>Ruichen Shao, Bei Li, Gangao Liu, Yang Chen, Zhouxiang, Jingang Wang, Xunliang Cai, Peng Li</p>	65535
<p>LLaVA-Interleave: Tackling Multi-Image, Video, and 3D in Large Multimodal Models</p> <p>Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, Chunyuan Li</p>	65558
<p>Unlocking the Potential of Model Calibration in Federated Learning</p> <p>Yun-Wei Chu, Dong-Jun Han, Seyyedali Hosseinalipour, Christopher Brinton</p>	65576
<p>Revealing and Reducing Gender Biases in Vision and Language Assistants (VLAs)</p> <p>Leander Girkbach, Stephan Alaniz, Yiran Huang, Trevor Darrell, Zeynep Akata</p>	65604
<p>DoF: A Diffusion Factorization Framework for Offline Multi-Agent Reinforcement Learning</p> <p>Chao Li, Ziwei Deng, Chenxing Lin, Wenqi Chen, Yongquan Fu, Weiquan Liu, Chenglu Wen, Cheng Wang, Siqi Shen</p>	65642

<p>Hypothetical Minds: Scaffolding Theory of Mind for Multi-Agent Tasks with Large Language Models</p> <p>Logan Cross, Violet Xiang, Agam Bhatia, Daniel Lk Yamins, Nick Haber</p>	65681
<p>Long-Context LLMs Meet RAG: Overcoming Challenges for Long Inputs in RAG</p> <p>Bowen Jin, Jinsung Yoon, Jiawei Han, Sercan O Arik</p>	65721
<p>Endowing Visual Reprogramming with Adversarial Robustness</p> <p>Shengjie Zhou, Xin Cheng, Haiyang Xu, Ming Yan, Tao Xiang, Feng Liu, Lei Feng</p>	65760
<p>Learning High-Degree Parities: The Crucial Role of the Initialization</p> <p>Emmanuel Abbe, Elisabetta Cornacchia, Jan Hazla, Donald Kougang-Yombi</p>	65785
<p>ADAM: An Embodied Causal Agent in Open-World Environments</p> <p>Shu Yu, Chaochao Lu</p>	65832
<p>Bootstrapping Language-Guided Navigation Learning with Self-Refining Data Flywheel</p> <p>Zun Wang, Jialu Li, Yicong Hong, Songze Li, Kunchang Li, Shoubin Yu, Yi Wang, Yu Qiao, Yali Wang, Mohit Bansal, Limin Wang</p>	65852
<p>WebRL: Training LLM Web Agents Via Self-Evolving Online Curriculum Reinforcement Learning</p> <p>Zehan Qi, Xiao Liu, Jat Long long, Hanyu Lai, Xueqiao Sun, Jiadai Sun, Xinyue Yang, Yu Yang, Shuntian Yao, Wei Xu, Jie Tang, Yuxiao Dong</p>	65879
<p>Differential Transformer</p> <p>Tianzhu Ye, Li Dong, Yuqing Xia, Yutao Sun, Yi Zhu, Gao Huang, Furu Wei</p>	65910
<p>An Online Learning Theory of Trading-Volume Maximization</p> <p>Tommaso Cesari, Roberto Colomboni</p>	65931
<p>Progressive Mixed-Precision Decoding for Efficient LLM Inference</p> <p>Hao Mark Chen, Fuwen Tan, Alexandros Kouris, Royson Lee, Hongxiang Fan, Stylianos Venieris</p>	65949
<p>AgentOccam: A Simple Yet Strong Baseline for LLM-Based Web Agents</p> <p>Ke Yang, Yao Liu, Sapana Chaudhary, Rasool Fakoor, Pratik Chaudhari, George Karypis, Huzefa Rangwala</p>	65966
<p>GlycanML: A Multi-Task and Multi-Structure Benchmark for Glycan Machine Learning</p> <p>Minghao Xu, Yunteng Geng, Yihang Zhang, Ling Yang, Jian Tang, Wentao Zhang</p>	65999
<p>No Free Lunch: Fundamental Limits of Learning Non-Hallucinating Generative Models</p> <p>Changlong Wu, Ananth Grama, Wojciech Szpankowski</p>	66015
<p>Facilitating Multi-Turn Function Calling for LLMs Via Compositional Instruction Tuning</p> <p>Mingyang Chen, Sunhaoze, Tianpeng Li, Fan Yang, Hao Liang, Keerlu, Bin Cui, Wentao Zhang, Zenan Zhou, Weipeng Chen</p>	66028

<p>Learning Vector Fields of Differential Equations on Manifolds with Geometrically Constrained Operator-Valued Kernels</p> <p><i>Daning Huang, Hanyang He, John Harlim, Yan Li</i></p>	66069
<p>Matérn Kernels for Tunable Implicit Surface Reconstruction</p> <p><i>Maximilian Weiherer, Bernhard Egger</i></p>	66102
<p>Dynamic Low-Rank Sparse Adaptation for Large Language Models</p> <p><i>Weizhong Huang, Yuxin Zhang, Xiawu Zheng, Liuyang, Jing Lin, Yiwu Yao, Rongrong Ji</i></p>	66124
<p>Vertical Federated Learning with Missing Features During Training and Inference</p> <p><i>Pedro Valdeira, Shiqiang Wang, Yuejie Chi</i></p>	66150
<p>Understanding Long Videos with Multimodal Language Models</p> <p><i>Kanchana Ranasinghe, Xiang Li, Kumara Kahatapitiya, Michael S Ryoo</i></p>	66170
<p>GTR: Improving Large 3D Reconstruction Models Through Geometry and Texture Refinement</p> <p><i>Peiye Zhuang, Songfang Han, Chaoyang Wang, Aliaksandr Siarohin, Jiaxu Zou, Michael Vasilkovsky, Vladislav Shakhrai, Sergei Korolev, Sergey Tulyakov, Hsin-Ying Lee</i></p>	66196
<p>Do Mice Grok? Glimpses of Hidden Progress in Sensory Cortex</p> <p><i>Tanishq Kumar, Blake Bordelon, Cengiz Pehlevan, Venkatesh N Murthy, Samuel J. Gershman</i></p>	66213
<p>Learning Graph Quantized Tokenizers</p> <p><i>Limei Wang, Kaveh Hassani, Si Zhang, Dongqi Fu, Baichuan Yuan, Weilin Cong, Zhigang Hua, Hao Wu, Ning Yao, Bo Long</i></p>	66233
<p>\$\$-Exponential Family for Policy Optimization</p> <p><i>Lingwei Zhu, Haseeb Shah, Han Wang, Yukie Nagai, Martha White</i></p>	66255
<p>SPA-BENCH: A COMPREHENSIVE BENCHMARK FOR SMARTPHONE AGENT EVALUATION</p> <p><i>Jingxuan Chen, Derek Yuen, Bin Xie, Yuhao Yang, Gongwei Chen, Zhihao Wu, Li Yixing, Xurui Zhou, Weiwen Liu, Shuai Wang, Kaiwen Zhou, Rui Shao, Liqiang Nie, Yasheng Wang, Jianye Hao, Jun Wang, Kun Shao</i></p>	66283
<p>Support is All You Need for Certified VAE Training</p> <p><i>Changming Xu, Debangshu Banerjee, Deepak Vasisht, Gagandeep Singh</i></p>	66321
<p>KAN: Kolmogorov–Arnold Networks</p> <p><i>Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljacic, Thomas Y. Hou, Max Tegmark</i></p>	66342
<p>Structure Language Models for Protein Conformation Generation</p> <p><i>Jiarui Lu, Xiaoyin Chen, Stephen Zhewen Lu, Chence Shi, Hongyu Guo, Yoshua Bengio, Jian Tang</i></p>	66389
<p>A Second-Order Perspective on Model Compositionality and Incremental Learning</p> <p><i>Angelo Porrello, Lorenzo Bonicelli, Pietro Buzzega, Monica Millunzi, Simone Calderara, Rita Cucchiara</i></p>	66423

<p>Learning Causal Alignment for Reliable Disease Diagnosis</p> <p>Mingzhou Liu, Ching-Wen Lee, Xinwei Sun, Xueqing Yu, Yu Qiao, Yizhou Wang</p>	66454
<p>Solving New Tasks by Adapting Internet Video Knowledge</p> <p>Calvin Luo, Zilai Zeng, Yilun Du, Chen Sun</p>	66473
<p>Stealthy Shield Defense: A Conditional Mutual Information-Based Approach Against Black-Box Model Inversion Attacks</p> <p>Tianqu Zhuang, Hongyao Yu, Yixiang Qiu, Hao Fang, Bin Chen, Shu-Tao Xia</p>	66492
<p>InstructRAG: Instructing Retrieval-Augmented Generation Via Self-Synthesized Rationales</p> <p>Zhepei Wei, Wei-Lin Chen, Yu Meng</p>	66510
<p>Spreading Out-Of-Distribution Detection on Graphs</p> <p>Daeho Um, Jongin Lim, Sunoh Kim, Yuneil Yeo, Yoonho Jung</p>	66534
<p>A Benchmark for Semantic Sensitive Information in LLMs Outputs</p> <p>Qingjie Zhang, Han Qiu, Di Wang, Yiming Li, Tianwei Zhang, Wenyu Zhu, Haiqin Weng, Liu Yan, Chao Zhang</p>	66565
<p>Analyzing and Boosting the Power of Fine-Grained Visual Recognition for Multi-Modal Large Language Models</p> <p>Hulingxiao He, Geng Li, Zijun Geng, Jinglin Xu, Yuxin Peng</p>	66593
<p>Instance-Dependent Early Stopping</p> <p>Suqin Yuan, Runqi Lin, Lei Feng, Bo Han, Tongliang Liu</p>	66611
<p>Exploring Local Memorization in Diffusion Models Via Bright Ending Attention</p> <p>Chen Chen, Daochang Liu, Mubarak Shah, Chang Xu</p>	66636
<p>GaussianAnything: Interactive Point Cloud Flow Matching for 3D Generation</p> <p>Yushi Lan, Shangchen Zhou, Zhaoyang Lyu, Fangzhou Hong, Shuai Yang, Bo Dai, Xingang Pan, Chen Change Loy</p>	66651
<p>No Pose, No Problem: Surprisingly Simple 3D Gaussian Splats from Sparse Unposed Images</p> <p>Botao Ye, Sifei Liu, Haofei Xu, Xueting Li, Marc Pollefeys, Ming-Hsuan Yang, Songyou Peng</p>	66676
<p>ACES: Automatic Cohort Extraction System for Event-Stream Datasets</p> <p>Justin Xu, Jack Gallifant, Alistair Johnson, Matthew B. A. McDermott</p>	66701
<p>Content-Style Learning from Unaligned Domains: Identifiability Under Unknown Latent Dimensions</p> <p>Sagar Shrestha, Xiao Fu</p>	66717
<p>NExT-Mol: 3D Diffusion Meets 1D Language Modeling for 3D Molecule Generation</p> <p>Zhiyuan Liu, Yanchen Luo, Han Huang, Enzhi Zhang, Sihang Li, Junfeng Fang, Yaorui Shi, Xiang Wang, Kenji Kawaguchi, Tat-Seng Chua</p>	66755

B-STaR: Monitoring and Balancing Exploration and Exploitation in Self-Taught Reasoners Weihao Zeng, Yuzhen Huang, Lulu Zhao, Yijun Wang, Zifei Shan, Junxian He	66787
Modeling Fine-Grained Hand-Object Dynamics for Egocentric Video Representation Learning Baoqi Pei, Yifei Huang, Jilan Xu, Guo Chen, Yuping He, Lijin Yang, Yali Wang, Weidi Xie, Yu Qiao, Fei Wu, Limin Wang	66805
Provable Benefit of Annealed Langevin Monte Carlo for Non-Log-Concave Sampling Wei Guo, Molei Tao, Yongxin Chen	66827
Task-Adaptive Pretrained Language Models Via Clustered-Importance Sampling David Grangier, Simin Fan, Skyler Seto, Pierre Ablin	66850
Varying Shades of Wrong: Aligning LLMs with Wrong Answers Only Jihan Yao, Wenxuan Ding, Shangbin Feng, Lucy Lu Wang, Yulia Tsvetkov	66873
QuaDiM: A Conditional Diffusion Model for Quantum State Property Estimation Yehui Tang, Mabiao Long, Junchi Yan	66910
On the Hölder Stability of Multiset and Graph Neural Networks Yair Davidson, Nadav Dym	66932
U-μSρ: The Unit-Scaled Maximal Update Parametrization Charlie Blake, Constantin Eichenberg, Josef Dean, Lukas Balles, Luke Yuri Prince, Björn Deiseroth, Andres Felipe Cruz-Salinas, Carlo Luschi, Samuel Weinbach, Douglas Orr	66975
Expected Sliced Transport Plans Xinran Liu, Rocio Diaz Martin, Yikun Bai, Ashkan Shahbazi, Matthew Thorpe, Akram Aldroubi, Soheil Kolouri	67024
Swift Hydra: Self-Reinforcing Generative Framework for Anomaly Detection with Multiple Mamba Models Nguyen Hoang Khoi Do, Truc Nguyen, Malik Hassanaly, Raed Alharbi, Jung Taek Seo, My T. Thai	67048
Diffusion Models Are Real-Time Game Engines Dani Valevski, Yaniv Leviathan, Moab Arar, Shlomi Fruchter	67092
XLand-100B: A Large-Scale Multi-Task Dataset for In-Context Reinforcement Learning Alexander Nikulin, Ilya Zisman, Alexey Zemtsov, Vladislav Kurenkov	67115
Can Video LLMs Refuse to Answer? Alignment for Answerability in Video Large Language Models Eunseop Yoon, Hee Suk Yoon, Mark A. Hasegawa-Johnson, Chang D. Yoo	67139
Universal Image Restoration Pre-Training Via Degradation Classification Jiakui Hu, Lujia Jin, Zhengjian Yao, Yanye Lu	67165
Revisit Micro-Batch Clipping: Adaptive Data Pruning Via Gradient Manipulation Lun Wang	67187

<p>Magnetic Preference Optimization: Achieving Last-Iterate Convergence for Language Model Alignment</p> <p>Mingzhi Wang, Chengdong Ma, Qizhi Chen, Linjian Meng, Yang Han, Jiancong Xiao, Zhaowei Zhang, Jing Huo, Weijie J Su, Yaodong Yang</p>	67580
<p>Learning Robust Representations with Long-Term Information for Generalization in Visual Reinforcement Learning</p> <p>Rui Yang, Jie Wang, Qijie Peng, Ruibo Guo, Guoping Wu, Bin Li</p>	67612
<p>What to Align in Multimodal Contrastive Learning?</p> <p>Benoit Dufumier, Javiera Castillo Navarro, Devis Tuia, Jean-Philippe Thiran</p>	67641
<p>A Transfer Learning Framework for Weak to Strong Generalization</p> <p>Seamus Somerstep, Felipe Maia Polo, Moulinath Banerjee, Yaacov Ritov, Mikhail Yurochkin, Yuekai Sun</p>	67666
<p>Scalable Discrete Diffusion Samplers: Combinatorial Optimization and Statistical Physics</p> <p>Sebastian Sanokowski, Wilhelm Franz Berghammer, Haoyu Peter Wang, Martin Ennemoser, Sepp Hochreiter, Sebastian Lehner</p>	67701
<p>Adaptive Shrinkage Estimation for Personalized Deep Kernel Regression in Modeling Brain Trajectories</p> <p>Vasiliki Tassopoulou, Haochang Shou, Christos Davatzikos</p>	67731
<p>Deep MMD Gradient Flow Without Adversarial Training</p> <p>Alexandre Galashov, Valentin De Bortoli, Arthur Gretton</p>	67760
<p>Daily Dilemmas: Revealing Value Preferences of LLMs with Quandaries of Daily Life</p> <p>Yu Ying Chiu, Liwei Jiang, Yejin Choi</p>	67795
<p>LLM-Wrapper: Black-Box Semantic-Aware Adaptation of Vision-Language Models for Referring Expression Comprehension</p> <p>Amaia Cardiel, Eloi Zablocki, Elias Ramzi, Oriane Siméoni, Matthieu Cord</p>	67833
<p>Multimodal Lego: Model Merging and Fine-Tuning Across Topologies and Modalities in Biomedicine</p> <p>Konstantin Hemker, Nikola Simidjievski, Mateja Jamnik</p>	67853
<p>Tackling Data Corruption in Offline Reinforcement Learning Via Sequence Modeling</p> <p>Jiawei Xu, Rui Yang, Shuang Qiu, Feng Luo, Meng Fang, Baoxiang Wang, Lei Han</p>	67875
<p>Bayesian WeakS-To-Strong from Text Classification to Generation</p> <p>Ziyun Cui, Ziyang Zhang, Guangzhi Sun, Wen Wu, Chao Zhang</p>	67904
<p>RTop-K: Ultra-Fast Row-Wise Top-K Selection for Neural Network Acceleration on GPUs</p> <p>Xi Xie, Yuebo Luo, Hongwu Peng, Caiwen Ding</p>	67922

<p>No Need to Talk: Asynchronous Mixture of Language Models</p> <p>Anastasiia Filippova, Angelos Katharopoulos, David Grangier, Ronan Collobert</p>	67937
<p>CL-DiffPhyCon: Closed-Loop Diffusion Control of Complex Physical Systems</p> <p>Long Wei, Haodong Feng, Yuchen Yang, Ruiqi Feng, Peiyan Hu, Xiang Zheng, Tao Zhang, Dixia Fan, Tailin Wu</p>	67960
<p>Humanizing the Machine: Proxy Attacks to Mislead LLM Detectors</p> <p>Tianchun Wang, Yuanzhou Chen, Zichuan Liu, Zhanwen Chen, Haifeng Chen, Xiang Zhang, Wei Cheng</p>	67997
<p>Data Scaling Laws in Imitation Learning for Robotic Manipulation</p> <p>Fanqi Lin, Yingdong Hu, Pingyue Sheng, Chuan Wen, Jiacheng You, Yang Gao</p>	68026
<p>Unlocking Efficient, Scalable, and Continual Knowledge Editing with Basis-Level Representation Fine-Tuning</p> <p>Tianci Liu, Ruirui Li, Yunzhe Qi, Hui Liu, Xianfeng Tang, Tianqi Zheng, Qingyu Yin, Monica Xiao Cheng, Jun Huan, Haoyu Wang, Jing Gao</p>	68060
<p>SoftCVI: Contrastive Variational Inference with Self-Generated Soft Labels</p> <p>Daniel Ward, Mark Beaumont, Matteo Fasiolo</p>	68081
<p>Transformers Can Learn Temporal Difference Methods for In-Context Reinforcement Learning</p> <p>Jiuqi Wang, Ethan Blaser, Hadi Daneshmand, Shangtong Zhang</p>	68107
<p>InfoGS: Efficient Structure-Aware 3D Gaussians Via Lightweight Information Shaping</p> <p>Yunchao Zhang, Guandao Yang, Leonidas Guibas, Yanchao Yang</p>	68160
<p>Should VLMs Be Pre-Trained with Image Data?</p> <p>Sedrick Keh, Jean Mercat, Samir Yitzhak Gadre, Kushal Arora, Igor Vasiljevic, Benjamin Burchfiel, Shuran Song, Russ Tedrake, Thomas Kollar, Ludwig Schmidt, Achal Dave</p>	68180
<p>Generalization and Distributed Learning of GFlowNets</p> <p>Tiago Silva, Amauri H Souza, Omar Rivasplata, Vikas Garg, Samuel Kaski, Diego Mesquita</p>	68205
<p>MS-Diffusion: Multi-Subject Zero-Shot Image Personalization with Layout Guidance</p> <p>Xierui Wang, Siming Fu, Qihan Huang, Wanggui He, Hao Jiang</p>	68248
<p>Let SSMS Be ConvNets: State-Space Modeling with Optimal Tensor Contractions</p> <p>Yan Ru Pei</p>	68277
<p>ESE: Espresso Sentence Embeddings</p> <p>Xianming Li, Zongxi Li, Jing Li, Haoran Xie, Qing Li</p>	68302
<p>Logicbreaks: A Framework for Understanding Subversion of Rule-Based Inference</p> <p>Anton Xue, Avishree Khare, Rajeev Alur, Surbhi Goel, Eric Wong</p>	68316

<p>Gaussian Ensemble Belief Propagation for Efficient Inference in High-Dimensional, Black-Box Systems</p> <p>Dan Mackinlay, Russell Tsuchida, Daniel Edward Pagendam, Petra Kuhnert</p>	68346
<p>SegLLM: Multi-Round Reasoning Segmentation with Large Language Models</p> <p>Xudong Wang, Shaolun Zhang, Shufan Li, Kehan Li, Konstantinos Kallidromitis, Yusuke Kato, Kazuki Kozuka, Trevor Darrell</p>	68378
<p>Gap Preserving Distillation by Building Bidirectional Mappings with a Dynamic Teacher</p> <p>Yong Guo, Shulian Zhang, Haolin Pan, Jing Liu, Yulun Zhang, Jian Chen</p>	68400
<p>Magpie: Alignment Data Synthesis from Scratch by Prompting Aligned LLMs with Nothing</p> <p>Zhangchen Xu, Fengqing Jiang, Luyao Niu, Yuntian Deng, Radha Poovendran, Yejin Choi, Bill Yuchen Lin</p>	68422
<p>RAG-DDR: Optimizing Retrieval-Augmented Generation Using Differentiable Data Rewards</p> <p>Xinze Li, Sen Mei, Zhenghao Liu, Yukun Yan, Shuo Wang, Shi Yu, Zheni Zeng, Hao Chen, Ge Yu, Zhiyuan Liu, Maosong Sun, Chenyan Xiong</p>	68459
<p>Mitigating Reward Over-Optimization in RLHF Via Behavior-Supported Regularization</p> <p>Juntao Dai, Taiye Chen, Yaodong Yang, Qian Zheng, Gang Pan</p>	68481
<p>CONTRA: Conformal Prediction Region Via Normalizing Flow Transformation</p> <p>Zhenhan Fang, Aixin Tan, Jian Huang</p>	68516
<p>Diverse Preference Learning for Capabilities and Alignment</p> <p>Stewart Slocum, Asher Parker-Sartori, Dylan Hadfield-Menell</p>	68534
<p>Domain Guidance: A Simple Transfer Approach for a Pre-Trained Diffusion Model</p> <p>Jincheng Zhong, Xiangcheng Zhang, Jianmin Wang, Mingsheng Long</p>	68565
<p>Control-Oriented Clustering of Visual Latent Representation</p> <p>Han Qi, Haocheng Yin, Heng Yang</p>	68584
<p>Scalable and Certifiable Graph Unlearning: Overcoming the Approximation Error Barrier</p> <p>Lu Yi, Zhewei Wei</p>	68617
<p>LoCoDL: Communication-Efficient Distributed Learning with Local Training and Compression</p> <p>Laurent Condat, Arto Maranjyan, Peter Richtárik</p>	68655
<p>Convergence of Score-Based Discrete Diffusion Models: A Discrete-Time Analysis</p> <p>Zikun Zhang, Zixiang Chen, Quanquan Gu</p>	68676
<p>UV-Attack: Physical-World Adversarial Attacks on Person Detection Via Dynamic-NeRF-Based UV Mapping</p> <p>Yanjie Li, Kaisheng Liang, Bin Xiao</p>	68702
<p>Improved Training Technique for Latent Consistency Models</p> <p>Quan Dao, Khanh Doan, Di Liu, Trung Le, Dimitris N. Metaxas</p>	68725

Towards a Complete Logical Framework for GNN Expressiveness <i>Tuo Xu</i>	68743
Autoregressive Pretraining with Mamba in Vision <i>Sucheng Ren, Xianhang Li, Haoqin Tu, Feng Wang, Fangxun Shu, Lei Zhang, Jieru Mei, Linjie Yang, Peng Wang, Heng Wang, Alan Yuille, Cihang Xie</i>	68783
On Scaling Up 3D Gaussian Splatting Training <i>Hexu Zhao, Haoyang Weng, Daohan Lu, Ang Li, Jinyang Li, Aurojit Panda, Saining Xie</i>	68798
Scalable Decision-Making in Stochastic Environments Through Learned Temporal Abstraction <i>Baiting Luo, Ava Pettet, Aron Laszka, Abhishek Dubey, Ayan Mukhopadhyay</i>	68818
MPLUG-Owl3: Towards Long Image-Sequence Understanding in Multi-Modal Large Language Models <i>Jiabo Ye, Haiyang Xu, Haowei Liu, Anwen Hu, Ming Yan, Qi Qian, Ji Zhang, Fei Huang, Jingren Zhou</i>	68839
Boltzmann Priors for Implicit Transfer Operators <i>Juan Viguera Diez, Mathias Jacob Schreiner, Ola Engkvist, Simon Olsson</i>	68862
ReSi: A Comprehensive Benchmark for Representational Similarity Measures <i>Max Klabunde, Tassilo Wald, Tobias Schumacher, Klaus Maier-Hein, Markus Strohmaier, Florian Lemmerich</i>	68891
Leveraging Flatness to Improve Information-Theoretic Generalization Bounds for SGD <i>Ze Peng, Jian Zhang, Yisen Wang, Lei Qi, Yinghuan Shi, Yang Gao</i>	68931
Steering Protein Family Design Through Profile Bayesian Flow <i>Jingjing Gong, Yu Pei, Siyu Long, Yuxuan Song, Zhe Zhang, Wenhao Huang, Ziyao Cao, Shuyi Zhang, Hao Zhou, Wei-Ying Ma</i>	68980
MorphoDiff: Cellular Morphology Painting with Diffusion Models <i>Zeinab Navidi, Jun Ma, Esteban Miglietta, Le Liu, Anne E Carpenter, Beth A Cimini, Benjamin Haibe-Kains, Bo Wang</i>	69013
A Training-Free Sub-Quadratic Cost Transformer Model Serving Framework with Hierarchically Pruned Attention <i>Heejun Lee, Geon Park, Youngwan Lee, Jaduk Suh, Jina Kim, Wonyong Jeong, Bumsik Kim, Hyemin Lee, Myeongjae Jeon, Sung Ju Hwang</i>	69038
TEASER: Token Enhanced Spatial Modeling for Expressions Reconstruction <i>Yunfei Liu, Lei Zhu, Lijian Lin, Ye Zhu, Ailing Zhang, Yu Li</i>	69084
Learning from Imperfect Human Feedback: A Tale from Corruption-Robust Dueling <i>Yuwei Cheng, Fan Yao, Xuefeng Liu, Haifeng Xu</i>	69104
Tight Clusters Make Specialized Experts <i>Stefan Nielsen, Rachel Teo, Laziz Abdullaev, Tan Minh Nguyen</i>	69136

<p>Think Thrice Before You Act: Progressive Thought Refinement in Large Language Models</p> <p>Chengyu Du, Jinyi Han, Yizhou Ying, Aili Chen, Qianyu He, Haokun Zhao, Haoran Guo, Sirui Xia, Jiaqing Liang, Zulong Chen, Liangyue Li, Yanghua Xiao</p>	69163
<p>Iterative Nash Policy Optimization: Aligning LLMs with General Preferences Via No-Regret Learning</p> <p>Yuheng Zhang, Dian Yu, Baolin Peng, Linfeng Song, Ye Tian, Mingyue Huo, Nan Jiang, Haitao Mi, Dong Yu</p>	69194
<p>Adam Exploits ℓ_∞-Geometry of Loss Landscape Via Coordinate-Wise Adaptivity</p> <p>Shuo Xie, Mohamad Amin Mohamadi, Zhiyuan Li</p>	69211
<p>AdaFisher: Adaptive Second Order Optimization Via Fisher Information</p> <p>Damien Martins Gomes, Yanlei Zhang, Eugene Belilovsky, Guy Wolf, Mahdi S. Hosseini</p>	69240
<p>Interpretable Unsupervised Joint Denoising and Enhancement for Real-World Low-Light Scenarios</p> <p>Li Huaqiu, Huxiaowan, Haoqian Wang</p>	69282
<p>Streamlining Prediction in Bayesian Deep Learning</p> <p>Rui Li, Marcus Klasson, Arno Solin, Martin Trapp</p>	69295
<p>Beyond Random Masking: When Dropout Meets Graph Convolutional Networks</p> <p>Yuankai Luo, Xiao-Ming Wu, Hao Zhu</p>	69326
<p>G-LLaVA: Solving Geometric Problem with Multi-Modal Large Language Model</p> <p>Jiahui Gao, Renjie Pi, Jipeng Zhang, Jiacheng Ye, Wanjun Zhong, Yufei Wang, Lanqing Hong, Jianhua Han, Hang Xu, Zhenguo Li, Lingpeng Kong</p>	69346
<p>On the Optimization Landscape of Low Rank Adaptation Methods for Large Language Models</p> <p>Xu-Hui Liu, Yali Du, Jun Wang, Yang Yu</p>	69368
<p>EC-DIT: Scaling Diffusion Transformers with Adaptive Expert-Choice Routing</p> <p>Haotian Sun, Tao Lei, Bowen Zhang, Yanghao Li, Haoshuo Huang, Ruoming Pang, Bo Dai, Nan Du</p>	69394
<p>ICLR: In-Context Learning of Representations</p> <p>Core Francisco Park, Andrew Lee, Ekdeep Singh Lubana, Yongyi Yang, Maya Okawa, Kento Nishi, Martin Wattenberg, Hidenori Tanaka</p>	69413
<p>Explanations of GNN on Evolving Graphs Via Axiomatic Layer Edges</p> <p>Yazheng Liu, Sihong Xie</p>	69440
<p>Predicting the Energy Landscape of Stochastic Dynamical System Via Physics-Informed Self-Supervised Learning</p> <p>Ruikun Li, Huandong Wang, Qingmin Liao, Yong Li</p>	69465
<p>Execution-Guided Within-Prompt Search for Programming-By-Example</p> <p>Gust Verbruggen, Ashish Tiwari, Mukul Singh, Vu Le, Sumit Gulwani</p>	69489

<p>SONICS: Synthetic Or Not - Identifying Counterfeit Songs</p> <p><i>Md Awsafur Rahman, Zaber Ibn Abdul Hakim, Najibul Haque Sarker, Bishmoy Paul, Shaikh Anowarul Fattah</i></p>	69501
<p>SuperCorrect: Advancing Small LLM Reasoning with Thought Template Distillation and Self-Correction</p> <p><i>Ling Yang, Zhaochen Yu, Tianjun Zhang, Minkai Xu, Joseph E. Gonzalez, Bin Cui, Shuicheng Yan</i></p>	69532
<p>LLaMA-Omni: Seamless Speech Interaction with Large Language Models</p> <p><i>Qingkai Fang, Shoutao Guo, Yan Zhou, Zhengrui Ma, Shaolei Zhang, Yang Feng</i></p>	69565
<p>Understanding and Mitigating Bottlenecks of State Space Models Through the Lens of Recency and Over-Smoothing</p> <p><i>Peihao Wang, Ruisi Cai, Yuehao Wang, Jiajun Zhu, Pragya Srivastava, Zhangyang Wang, Pan Li</i></p>	69583
<p>Scalable Benchmarking and Robust Learning for Noise-Free Ego-Motion and 3D Reconstruction from Noisy Video</p> <p><i>Xiaohao Xu, Tianyi Zhang, Shibo Zhao, Xiang Li, Sibao Wang, Yongqi Chen, Ye Li, Bhiksha Raj, Matthew Johnson-Roberson, Sebastian Scherer, Xiaonan Huang</i></p>	69612
<p>CLoSD: Closing the Loop Between Simulation and Diffusion for Multi-Task Character Control</p> <p><i>Guy Tevet, Sigal Raab, Setareh Cohan, Daniele Reda, Zhengyi Luo, Xue Bin Peng, Amit Haim Bermano, Michiel Van De Panne</i></p>	69704
<p>LongMemEval: Benchmarking Chat Assistants on Long-Term Interactive Memory</p> <p><i>Di Wu, Hongwei Wang, Wenhao Yu, Yuwei Zhang, Kai-Wei Chang, Dong Yu</i></p>	69719
<p>Connectome Mapping: Shape-Memory Network Via Interpretation of Contextual Semantic Information</p> <p><i>Kyungsu Lee, Haeyun Lee, Jae Youn Hwang</i></p>	69747
<p>TEOChat: A Large Vision-Language Assistant for Temporal Earth Observation Data</p> <p><i>Jeremy Andrew Irvin, Emily Ruoyu Liu, Joyce C. Chen, Ines Dormoy, Jinyoung Kim, Samar Khanna, Zhuo Zheng, Stefano Ermon</i></p>	69776
<p>A Robust Method to Discover Causal Or Anticausal Relation</p> <p><i>Yu Yao, Yang Zhou, Bo Han, Mingming Gong, Kun Zhang, Tongliang Liu</i></p>	69805
<p>Debiasing Mini-Batch Quadratics for Applications in Deep Learning</p> <p><i>Lukas Tatzel, Bálint Mucsányi, Osane Hackel, Philipp Hennig</i></p>	69830
<p>Topograph: An Efficient Graph-Based Framework for Strictly Topology Preserving Image Segmentation</p> <p><i>Laurin Lux, Alexander H Berger, Alexander Weers, Nico Stucki, Daniel Rueckert, Ulrich Bauer, Johannes C. Paetzold</i></p>	69870
<p>Spectral Compressive Imaging Via Unmixing-Driven Subspace Diffusion Refinement</p> <p><i>Haijin Zeng, Benteng Sun, Yongyong Chen, Jingyong Su, Yong Xu</i></p>	69895

<p>Towards Self-Supervised Covariance Estimation in Deep Heteroscedastic Regression</p> <p>Megh Shukla, Aziz Shameem, Mathieu Salzmann, Alexandre Alahi</p>	69923
<p>A Closer Look at Machine Unlearning for Large Language Models</p> <p>Xiaojian Yuan, Tianyu Pang, Chao Du, Kejiang Chen, Weiming Zhang, Min Lin</p>	69946
<p>Underdamped Diffusion Bridges with Applications to Sampling</p> <p>Denis Blessing, Julius Berner, Lorenz Richter, Gerhard Neumann</p>	69972
<p>Understanding Matrix Function Normalizations in Covariance Pooling Through the Lens of Riemannian Geometry</p> <p>Ziheng Chen, Yue Song, Xiaojun Wu, Gaowen Liu, Nicu Sebe</p>	70005
<p>Rethinking Invariance in In-Context Learning</p> <p>Lizhe Fang, Yifei Wang, Khashayar Gatmiry, Lei Fang, Yisen Wang</p>	70032
<p>Do Contemporary Causal Inference Models Capture Real-World Heterogeneity? Findings from a Large-Scale Benchmark</p> <p>Haining Yu, Yizhou Sun</p>	70055
<p>CofCA: A STEP-WISE Counterfactual Multi-Hop QA Benchmark</p> <p>Jian Wu, Linyi Yang, Zhen Wang, Manabu Okumura, Yue Zhang</p>	70110
<p>Forgetting Transformer: Softmax Attention with a Forget Gate</p> <p>Zhixuan Lin, Evgenii Nikishin, Xu He, Aaron Courville</p>	70129
<p>Retri3D: 3D Neural Graphics Representation Retrieval</p> <p>Yushi Guan, Daniel Kwan, Jean Sebastien Dandurand, Xi Yan, Ruofan Liang, Yuxuan Zhang, Nilesh Jain, Nilesh Ahuja, Selvakumar Panneer, Nandita Vijaykumar</p>	70164
<p>γ-MoD: Exploring Mixture-Of-Depth Adaptation for Multimodal Large Language Models</p> <p>Yaxin Luo, Gen Luo, Jiayi Ji, Yiyi Zhou, Xiaoshuai Sun, Zhiqiang Shen, Rongrong Ji</p>	70207
<p>Iterative Label Refinement Matters More than Preference Optimization Under Weak Supervision</p> <p>Yaowen Ye, Cassidy Laidlaw, Jacob Steinhardt</p>	70221
<p>VOILA: Evaluation of MLLMs for Perceptual Understanding and Analogical Reasoning</p> <p>Nilay Yilmaz, Maitreya Patel, Yiran Lawrence Luo, Tejas Gokhale, Chitta Baral, Suren Jayasuriya, Yezhou Yang</p>	70243
<p>Efficient Source-Free Time-Series Adaptation Via Parameter Subspace Disentanglement</p> <p>Gaurav Patel, Christopher Michael Sandino, Behrooz Mahasseni, Ellen L. Zippi, Erdrin Azemi, Ali Moin, Juri Minxha</p>	70269
<p>HART: Efficient Visual Generation with Hybrid Autoregressive Transformer</p> <p>Haotian Tang, Yecheng Wu, Shang Yang, Enze Xie, Junsong Chen, Junyu Chen, Zhuoyang Zhang, Han Cai, Yao Lu, Song Han</p>	70311

<p>Generalization Bounds and Model Complexity for Kolmogorov–Arnold Networks</p> <p><i>Xianyang Zhang, Huijuan Zhou</i></p>	70331
<p>PhysBench: Benchmarking and Enhancing Vision-Language Models for Physical World Understanding</p> <p><i>Wei Chow, Jiageng Mao, Boyi Li, Daniel Seita, Vitor Campagnolo Guizilini, Yue Wang</i></p>	70354
<p>SoftMatcha: A Soft and Fast Pattern Matcher for Billion-Scale Corpus Searches</p> <p><i>Hiroyuki Deguchi, Go Kamoda, Yusuke Matsushita, Chihiro Taguchi, Kohei Suenaga, Masaki Waga, Sho Yokoi</i></p>	70504
<p>How Many Samples Are Needed to Train a Deep Neural Network?</p> <p><i>Pegah Golestaneh, Mahsa Taheri, Johannes Lederer</i></p>	70530
<p>Generalized Behavior Learning from Diverse Demonstrations</p> <p><i>Varshith Sreeramdas, Rohan R Paleja, Letian Chen, Sanne Van Waveren, Matthew Gombolay</i></p>	70556
<p>SFESS: Score Function Estimators for k-Subset Sampling</p> <p><i>Klas Wijk, Ricardo Vinuesa, Hossein Azizpour</i></p>	70583
<p>Navigating Neural Space: Revisiting Concept Activation Vectors to Overcome Directional Divergence</p> <p><i>Frederik Pahde, Maximilian Dreyer, Moritz Weckbecker, Leander Weber, Christopher J. Anders, Thomas Wiegand, Wojciech Samek, Sebastian Lapuschkin</i></p>	70598
<p>From an LLM Swarm to a PDDL-Empowered Hive: Planning Self-Executed Instructions in a Multi-Modal Jungle</p> <p><i>Kaustubh Vyas, Damien Graux, Yijun Yang, Sebastien Montella, Chenxin Diao, Wendi Zhou, Pavlos Vougiouklis, Ruofei Lai, Yang Ren, Keshuang Li, Jeff Z. Pan</i></p>	70642
<p>GUI-World: A Video Benchmark and Dataset for Multimodal GUI-Oriented Understanding</p> <p><i>Dongping Chen, Yue Huang, Siyuan Wu, Jingyu Tang, Huichi Zhou, Qihui Zhang, Zhigang He, Yilin Bai, Chujie Gao, Liuyi Chen, Yiqiang Li, Chenlong Wang, Yue Yu, Tianshuo Zhou, Zhen Li, Yi Gui, Yao Wan, Pan Zhou, Jianfeng Gao, Lichao Sun</i></p>	70666
<p>Size-Generalizable RNA Structure Evaluation by Exploring Hierarchical Geometries</p> <p><i>Zongzhao Li, Jiacheng Cen, Wenbing Huang, Taifeng Wang, Le Song</i></p>	70719
<p>Demystifying Topological Message-Passing with Relational Structures: A Case Study on Oversquashing in Simplicial Message-Passing</p> <p><i>Diaaeldin Taha, James Chapman, Marzieh Eidi, Karel Devriendt, Guido Montufar</i></p>	70745
<p>Model Equality Testing: Which Model is This API Serving?</p> <p><i>Irena Gao, Percy Liang, Carlos Guestrin</i></p>	70795
<p>Generalized Video Moment Retrieval</p> <p><i>You Qin, Qilong Wu, Yicong Li, Wei Ji, Li Li, Pengcheng Cai, Lina Wei, Roger Zimmermann</i></p>	70809

<p>Causal Effect Estimation with Mixed Latent Confounders and Post-Treatment Variables</p> <p>Yaochen Zhu, Jing Ma, Liang Wu, Qi Guo, Liangjie Hong, Jundong Li</p>	70826
<p>ImageFolder: Autoregressive Image Generation with Folded Tokens</p> <p>Xiang Li, Kai Qiu, Hao Chen, Jason Kuen, Jiuxiang Gu, Bhiksha Raj, Zhe Lin</p>	70848
<p>RM-Bench: Benchmarking Reward Models of Language Models with Subtlety and Style</p> <p>Yantao Liu, Zijun Yao, Rui Min, Yixin Cao, Lei Hou, Juanzi Li</p>	70868
<p>Ambient Diffusion Posterior Sampling: Solving Inverse Problems with Diffusion Models Trained on Corrupted Data</p> <p>Asad Aali, Giannis Daras, Brett Levac, Sidharth Kumar, Alex Dimakis, Jon Tamir</p>	70901
<p>State Space Models Are Provably Comparable to Transformers in Dynamic Token Selection</p> <p>Naoki Nishikawa, Taiji Suzuki</p>	70929
<p>Transformers Handle Endogeneity in In-Context Linear Regression</p> <p>Haodong Liang, Krishna Balasubramanian, Lifeng Lai</p>	70972
<p>Advantage Alignment Algorithms</p> <p>Juan Agustin Duque, Milad Aghajohari, Tim Cooijmans, Razvan Ciuca, Tianyu Zhang, Gauthier Gidel, Aaron Courville</p>	71011
<p>Homomorphism Counts as Structural Encodings for Graph Learning</p> <p>Linus Bao, Emily Jin, Michael M. Bronstein, Ismail Ilkan Ceylan, Matthias Lanzinger</p>	71037
<p>A New Framework for Evaluating Model Out-Of-Distribution Generalisation for the Biochemical Domain</p> <p>Raul Fernandez-Diaz, Hoang Thanh Lam, Vanessa López, Denis C. Shields</p>	71066
<p>Multi-Perspective Data Augmentation for Few-Shot Object Detection</p> <p>Anh Khoa Nguyen Vu, Truong Quoc Truong, Vinh-Tiep Nguyen, Thanh Duc Ngo, Thanh-Toan Do, Tam V. Nguyen</p>	71092
<p>Unleashing the Potential of Vision-Language Pre-Training for 3D Zero-Shot Lesion Segmentation Via Mask-Attribute Alignment</p> <p>Yankai Jiang, Wenhui Lei, Xiaofan Zhang, Shaoting Zhang</p>	71111
<p>Model-Agnostic Meta-Learners for Estimating Heterogeneous Treatment Effects Over Time</p> <p>Dennis Frauen, Konstantin Hess, Stefan Feuerriegel</p>	71137
<p>GETS: Ensemble Temperature Scaling for Calibration in Graph Neural Networks</p> <p>Dingyi Zhuang, Chonghe Jiang, Yunhan Zheng, Shenhao Wang, Jinhua Zhao</p>	71167
<p>InstaTrain: Adaptive Training Via Ultra-Fast Natural Annealing Within Dynamical Systems</p> <p>Chuan Liu, Ruibing Song, Chunshu Wu, Pouya Haghi, Tong Geng</p>	71190
<p>MovieDreamer: Hierarchical Generation for Coherent Long Visual Sequences</p> <p>Canyu Zhao, Mingyu Liu, Wen Wang, Weihua Chen, Fan Wang, Hao Chen, Bo Zhang, Chunhua Shen</p>	71203

<p>Addax: Utilizing Zeroth-Order Gradients to Improve Memory Efficiency and Performance of SGD for Fine-Tuning Language Models</p> <p>Zeman Li, Xinwei Zhang, Peilin Zhong, Yuan Deng, Meisam Razaviyayn, Vahab Mirrokni</p>	71234
<p>MMDT: Decoding the Trustworthiness and Safety of Multimodal Foundation Models</p> <p>Chejian Xu, Jiawei Zhang, Zhaorun Chen, Chulin Xie, Mintong Kang, Yujin Potter, Zhun Wang, Zhuowen Yuan, Alexander Xiong, Zidi Xiong, Chenhui Zhang, Lingzhi Yuan, Yi Zeng, Peiyang Xu, Chengquan Guo, Andy Zhou, Jeffrey Ziwei Tan, Xuandong Zhao, Francesco Pinto, Zhen Xiang, Yu Gai, Zinan Lin, Dan Hendrycks, Bo Li, Dawn Song</p>	71272
<p>ACTIVE: Offline Reinforcement Learning Via Adaptive Imitation and In-Sample V^{π}-Ensemble</p> <p>Tianyuan Chen, Ronglong Cai, Faguo Wu, Xiao Zhang</p>	71369
<p>Do LLMs "know" Internally When They Follow Instructions?</p> <p>Juyeon Heo, Christina Heinze-Deml, Oussama Elachqar, Kwan Ho Ryan Chan, Shirley You Ren, Andrew Miller, Udhyakumar Nallasamy, Jaya Narain</p>	71388
<p>Fréchet Wavelet Distance: A Domain-Agnostic Metric for Image Generation</p> <p>Lokesh Veeramacheni, Moritz Wolter, Hilde Kuehne, Juergen Gall</p>	71407
<p>Scale-Aware Recognition in Satellite Images Under Resource Constraints</p> <p>Shreelekha Revankar, Cheng Perng Phoo, Utkarsh Mall, Bharath Hariharan, Kavita Bala</p>	71438
<p>Intrinsic Dimension Correlation: Uncovering Nonlinear Connections in Multimodal Representations</p> <p>Lorenzo Basile, Santiago Acevedo, Luca Bortolussi, Fabio Anselmi, Alex Rodriguez</p>	71451
<p>Statistical Tractability of Off-Policy Evaluation of History-Dependent Policies in POMDPs</p> <p>Yuheng Zhang, Nan Jiang</p>	71477
<p>Do Deep Neural Network Solutions Form a Star Domain?</p> <p>Ankit Sonthalia, Alexander Rubinstein, Ehsan Abbasnejad, Seong Joon Oh</p>	71503
<p>Learning from End User Data with Shuffled Differential Privacy Over Kernel Densities</p> <p>Tal Wagner</p>	71532
<p>Revisiting Prefix-Tuning: Statistical Benefits of Reparameterization Among Prompts</p> <p>Minh Le, Chau Nguyen, Huy Nguyen, Quyen Tran, Trung Le, Nhat Ho</p>	71564
<p>PIORF: Physics-Informed Ollivier-Ricci Flow for Long-Range Interactions in Mesh Graph Neural Networks</p> <p>Youn-Yeol Yu, Jeongwhan Choi, Jaehyeon Park, Kookjin Lee, Noseong Park</p>	71606
<p>From Exploration to Mastery: Enabling LLMs to Master Tools Via Self-Driven Interactions</p> <p>Changle Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Jun Xu, Ji-Rong Wen</p>	71631
<p>Generalization, Expressivity, and Universality of Graph Neural Networks on Attributed Graphs</p> <p>Levi Rauchwerger, Stefanie Jegelka, Ron Levie</p>	71654

<p>CoMotion: Concurrent Multi-Person 3D Motion</p> <p><i>Alejandro Newell, Peiyun Hu, Lahav Lipson, Stephan Richter, Vladlen Koltun</i></p>	71715
<p>Improving Generalization and Robustness in SNNs Through Signed Rate Encoding and Sparse Encoding Attacks</p> <p><i>Bhaskar Mukhoty, Hilal Alquabeh, Bin Gu</i></p>	71739
<p>Bridging the Semantic Gap Between Text and Table: A Case Study on NL2SQL</p> <p><i>Lin Long, Xijun Gu, Xinjie Sun, Wentao Ye, Haobo Wang, Sai Wu, Gang Chen, Junbo Zhao</i></p>	71756
<p>Microcanonical Langevin Ensembles: Advancing the Sampling of Bayesian Neural Networks</p> <p><i>Emanuel Sommer, Jakob Robnik, Giorgi Nozadze, Uros Seljak, David Rügamer</i></p>	71779
<p>Do LLM Agents Have Regret? a Case Study in Online Learning and Games</p> <p><i>Chanwoo Park, Xiangyu Liu, Asuman E. Ozdaglar, Kaiqing Zhang</i></p>	71801
<p>ILLUSION: Unveiling Truth with a Comprehensive Multi-Modal, Multi-Lingual Deepfake Dataset</p> <p><i>Kartik Thakral, Rishabh Ranjan, Akanksha Singh, Akshat Jain, Mayank Vatsa, Richa Singh</i></p>	71822
<p>Systematic Relational Reasoning with Epistemic Graph Neural Networks</p> <p><i>Irtaza Khalid, Steven Schockaert</i></p>	71847
<p>Laplace Sample Information: Data Informativeness Through a Bayesian Lens</p> <p><i>Johannes Kaiser, Kristian Schwethelm, Daniel Rueckert, Georgios Kaissis</i></p>	71886
<p>ETA: Evaluating Then Aligning Safety of Vision Language Models at Inference Time</p> <p><i>Yi Ding, Bolian Li, Ruqi Zhang</i></p>	71923
<p>VisualPredicator: Learning Abstract World Models with Neuro-Symbolic Predicates for Robot Planning</p> <p><i>Yichao Liang, Nishanth Kumar, Hao Tang, Adrian Weller, Joshua B. Tenenbaum, Tom Silver, Joao F. Henriques, Kevin Ellis</i></p>	71952
<p>Highly Efficient Self-Adaptive Reward Shaping for Reinforcement Learning</p> <p><i>Haozhe Ma, Zhengding Luo, Thanh Vinh Vo, Kuankuan Sima, Tze-Yun Leong</i></p>	71981
<p>The Computational Complexity of Circuit Discovery for Inner Interpretability</p> <p><i>Federico Adolfi, Martina G. Vilas, Todd Wareham</i></p>	72003
<p>Aligning Human Motion Generation with Human Perceptions</p> <p><i>Haoru Wang, Wentao Zhu, Luyi Miao, Yishu Xu, Feng Gao, Qi Tian, Yizhou Wang</i></p>	72019
<p>Linear Combination of Saved Checkpoints Makes Consistency and Diffusion Models Better</p> <p><i>Enshu Liu, Junyi Zhu, Zinan Lin, Xuefei Ning, Shuaiqi Wang, Matthew B. Blaschko, Sergey Yekhanin, Shengen Yan, Guohao Dai, Huazhong Yang, Yu Wang</i></p>	72050
<p>PEARL: Parallel Speculative Decoding with Adaptive Draft Length</p> <p><i>Tianyu Liu, Yun Li, Qitan Lv, Kai Liu, Jianchen Zhu, Winston Hu, Xiao Sun</i></p>	72087

<p>Continuous Diffusion for Mixed-Type Tabular Data</p> <p>Markus Mueller, Kathrin Gruber, Dennis Fok</p>	72107
<p>Both Ears Wide Open: Towards Language-Driven Spatial Audio Generation</p> <p>Peiwen Sun, Sitong Cheng, Xiangtai Li, Zhen Ye, Huadai Liu, Honggang Zhang, Wei Xue, Yike Guo</p>	72147
<p>Mitigating Information Loss in Tree-Based Reinforcement Learning Via Direct Optimization</p> <p>Sascha Marton, Tim Grams, Florian Vogt, Stefan Lüdtkke, Christian Bartelt, Heiner Stuckenschmidt</p>	72191
<p>GravMAD: Grounded Spatial Value Maps Guided Action Diffusion for Generalized 3D Manipulation</p> <p>Yangtao Chen, Zixuan Chen, Junhui Yin, Jing Huo, Pinzhuo Tian, Jieqi Shi, Yang Gao</p>	72224
<p>LVSM: A Large View Synthesis Model with Minimal 3D Inductive Bias</p> <p>Haian Jin, Hanwen Jiang, Hao Tan, Kai Zhang, Sai Bi, Tianyuan Zhang, Fujun Luan, Noah Snively, Zexiang Xu</p>	72257
<p>A3D: Does Diffusion Dream About 3D Alignment?</p> <p>Savva Victorovich Ignatyev, Nina Konovalova, Daniil Selikhanovych, Oleg Voynov, Nikolay Patakin, Ilya Olkov, Dmitry Senushkin, Alexey Artemov, Anton Konushin, Alexander Filippov, Peter Wonka, Evgeny Burnaev</p>	72278
<p>Near, Far: Patch-Ordering Enhances Vision Foundation Models' Scene Understanding</p> <p>Valentin Pariza, Mohammadreza Salehi, Gertjan J. Burghouts, Francesco Locatello, Yuki M Asano</p>	72303
<p>Is Your Multimodal Language Model Oversensitive to Safe Queries?</p> <p>Xirui Li, Hengguang Zhou, Ruochen Wang, Tianyi Zhou, Minhao Cheng, Cho-Jui Hsieh</p>	72331
<p>EconSG: Efficient and Multi-View Consistent Open-Vocabulary 3D Semantic Gaussians</p> <p>Can Zhang, Gim Hee Lee</p>	72383
<p>Locality Alignment Improves Vision-Language Models</p> <p>Ian Connick Covert, Tony Sun, James Zou, Tatsunori Hashimoto</p>	72401
<p>LongPO: Long Context Self-Evolution of Large Language Models Through Short-To-Long Preference Optimization</p> <p>Guanzheng Chen, Xin Li, Michael Shieh, Lidong Bing</p>	72440
<p>Demystifying the Token Dynamics of Deep Selective State Space Models</p> <p>Thieu Vo, Duy-Tung Pham, Xin T. Tong, Tan Minh Nguyen</p>	72457
<p>DEEM: Diffusion Models Serve as the Eyes of Large Language Models for Image Perception</p> <p>Run Luo, Yunshui Li, Longze Chen, Wanwei He, Ting-En Lin, Ziqiang Liu, Lei Zhang, Zikai Song, Hamid Alinejad-Rokny, Xiaobo Xia, Tongliang Liu, Binyuan Hui, Min Yang</p>	72484
<p>Unsupervised Model Tree Heritage Recovery</p> <p>Eliahu Horwitz, Asaf Shul, Yedid Hoshen</p>	72512

<p>Do LLMs Recognize Your Preferences? Evaluating Personalized Preference Following in LLMs</p> <p>Siyao Zhao, Mingyi Hong, Yang Liu, Devamanyu Hazarika, Kaixiang Lin</p>	72531
<p>4K4DGen: Panoramic 4D Generation at 4K Resolution</p> <p>Renjie Li, Panwang Pan, Bangbang Yang, Dejie Xu, Shijie Zhou, Xuanyang Zhang, Zeming Li, Achuta Kadambi, Zhangyang Wang, Zhengzhong Tu, Zhiwen Fan</p>	72575
<p>Exploratory Preference Optimization: Harnessing Implicit Q*-Approximation for Sample-Efficient RLHF</p> <p>Tengyang Xie, Dylan J Foster, Akshay Krishnamurthy, Corby Rosset, Ahmed Hassan Awadallah, Alexander Rakhlin</p>	72597
<p>Balancing Bias in Two-Sided Markets for Fair Stable Matchings</p> <p>Siyuan Wu, Leong Hou U, Panagiotis Karras</p>	72635
<p>Improving Unsupervised Constituency Parsing Via Maximizing Semantic Information</p> <p>Junjie Chen, Xiangheng He, Yusuke Miyao, Danushka Bollegala</p>	72654
<p>Model-Free Offline Reinforcement Learning with Enhanced Robustness</p> <p>Chi Zhang, Zain Ulabedeen Farhat, George K. Atia, Yue Wang</p>	72680
<p>SGD with Memory: Fundamental Properties and Stochastic Acceleration</p> <p>Dmitry Yarotsky, Maksim Velikanov</p>	72725
<p>Decoupled Finetuning for Domain Generalizable Semantic Segmentation</p> <p>Jaehyun Park, Donghyeon Kwon, Seong Joon Oh, Suha Kwak</p>	72768
<p>Bidirectional Decoding: Improving Action Chunking Via Guided Test-Time Sampling</p> <p>Yuejiang Liu, Jubayer Ibn Hamid, Annie Xie, Yoonho Lee, Max Du, Chelsea Finn</p>	72787
<p>A Simple Framework for Open-Vocabulary Zero-Shot Segmentation</p> <p>Thomas Stegmüller, Tim Lebailly, Nikola Đukic, Behzad Bozorgtabar, Tinne Tuytelaars, Jean-Philippe Thiran</p>	72821
<p>How Much is a Noisy Image Worth? Data Scaling Laws for Ambient Diffusion.</p> <p>Giannis Daras, Yeshwanth Cherapanamjeri, Constantinos Costis Daskalakis</p>	72843
<p>Spiking Vision Transformer with Saccadic Attention</p> <p>Shuai Wang, Malu Zhang, Dehao Zhang, Ammar Belatreche, Yichen Xiao, Yu Liang, Yimeng Shan, Qian Sun, Enqi Zhang, Yang Yang</p>	72872
<p>Breaking the Reclustering Barrier in Centroid-Based Deep Clustering</p> <p>Lukas Miklautz, Timo Klein, Kevin Sidak, Collin Leiber, Thomas Lang, Andrii Shkabrii, Sebastian Tschatschek, Claudia Plant</p>	72894
<p>What's the Move? Hybrid Imitation Learning Via Salient Points</p> <p>Priya Sundaresan, Hengyuan Hu, Quan Vuong, Jeannette Bohg, Dorsa Sadigh</p>	72935

<p>Towards General-Purpose Model-Free Reinforcement Learning</p> <p>Scott Fujimoto, Pierluca D'Oro, Amy Zhang, Yuandong Tian, Michael Rabbat</p>	72951
<p>Monte Carlo Planning with Large Language Model for Text-Based Game Agents</p> <p>Zijing Shi, Meng Fang, Ling Chen</p>	72995
<p>PointOBB-V2: Towards Simpler, Faster, and Stronger Single Point Supervised Oriented Object Detection</p> <p>Botao Ren, Xue Yang, Yi Yu, Junwei Luo, Zhidong Deng</p>	73016
<p>Neural Interactive Proofs</p> <p>Lewis Hammond, Sam Adam-Day</p>	73031
<p>Neuralized Markov Random Field for Interaction-Aware Stochastic Human Trajectory Prediction</p> <p>Zilin Fang, David Hsu, Gim Hee Lee</p>	73082
<p>Global Well-Posedness and Convergence Analysis of Score-Based Generative Models Via Sharp Lipschitz Estimates</p> <p>Connor Mooney, Zhongjian Wang, Jack Xin, Yifeng Yu</p>	73098
<p>Transformers Learn to Implement Multi-Step Gradient Descent with Chain of Thought</p> <p>Jianhao Huang, Zixuan Wang, Jason D. Lee</p>	73125
<p>Reflexive Guidance: Improving OoDD in Vision-Language Models Via Self-Guided Image-Adaptive Concept Generation</p> <p>Jihyo Kim, Seulbi Lee, Sangheum Hwang</p>	73186
<p>MaxInfoRL: Boosting Exploration in Reinforcement Learning Through Information Gain Maximization</p> <p>Bhavya Sukhija, Stelian Coros, Andreas Krause, Pieter Abbeel, Carmelo Sferrazza</p>	73223
<p>Consistency Checks for Language Model Forecasters</p> <p>Daniel Paleka, Abhimanyu Pallavi Sudhir, Alejandro Alvarez, Vineeth Bhat, Adam Shen, Evan Wang, Florian Tramèr</p>	73252
<p>Gated Delta Networks: Improving Mamba2 with Delta Rule</p> <p>Songlin Yang, Jan Kautz, Ali Hatamizadeh</p>	73311
<p>Scaling Stick-Breaking Attention: An Efficient Implementation and In-Depth Study</p> <p>Shawn Tan, Songlin Yang, Aaron Courville, Rameswar Panda, Yikang Shen</p>	73332
<p>TULIP: Token-Length Upgraded CLIP</p> <p>Ivona Najdenkoska, Mohammad Mahdi Derakhshani, Yuki M Asano, Nanne Van Noord, Marcel Worring, Cees G. M. Snoek</p>	73349

<p>OSTQuant: Refining Large Language Model Quantization with Orthogonal and Scaling Transformations for Better Distribution Fitting</p> <p>Xing Hu, Yuan Cheng, Dawei Yang, Zhixuan Chen, Zukang Xu, Jiangyongyu, Xuchen, Zhihang Yuan, Zhe Jiang, Sifan Zhou</p>	73373
<p>Almost Optimal Batch-Regret Tradeoff for Batch Linear Contextual Bandits</p> <p>Zihan Zhang, Xiangyang Ji, Yuan Zhou</p>	73399
<p>A Statistical Framework for Ranking LLM-Based Chatbots</p> <p>Siavash Ameli, Siyuan Zhuang, Ion Stoica, Michael W. Mahoney</p>	73430
<p>Easing Training Process of Rectified Flow Models Via Lengthening Inter-Path Distance</p> <p>Xu Shifeng, Yanzhu Liu, Adams Wai-Kin Kong</p>	73466
<p>Lossy Compression with Pretrained Diffusion Models</p> <p>Jeremy Vonderfecht, Feng Liu</p>	73484
<p>Shape as Line Segments: Accurate and Flexible Implicit Surface Representation</p> <p>Siyu Ren, Junhui Hou</p>	73500
<p>Separation Power of Equivariant Neural Networks</p> <p>Marco Pacini, Xiaowen Dong, Bruno Lepri, Gabriele Santin</p>	73521
<p>CircuitFusion: Multimodal Circuit Representation Learning for Agile Chip Design</p> <p>Wenji Fang, Shang Liu, Jing Wang, Zhiyao Xie</p>	73554
<p>Concept Bottleneck Large Language Models</p> <p>Chung-En Sun, Tuomas Oikarinen, Berk Ustun, Tsui-Wei Weng</p>	73577
<p>Enhanced Diffusion Sampling Via Extrapolation with Multiple ODE Solutions</p> <p>Jinyoung Choi, Junoh Kang, Bohyung Han</p>	73618
<p>Semi-Supervised Vision-Centric 3D Occupancy World Model for Autonomous Driving</p> <p>Xiang Li, Pengfei Li, Yupeng Zheng, Wei Sun, Yan Wang, Yilun Chen</p>	73643
<p>Error-Quantified Conformal Inference for Time Series</p> <p>Junxi Wu, Dongjian Hu, Yajie Bao, Shu-Tao Xia, Changliang Zou</p>	73661
<p>Joint Gradient Balancing for Data Ordering in Finite-Sum Multi-Objective Optimization</p> <p>Hansi Yang, James Kwok</p>	73689
<p>TS-LIF: A Temporal Segment Spiking Neuron Network for Time Series Forecasting</p> <p>Feng Shibo, Wanjin Feng, Xingyu Gao, Peilin Zhao, Zhiqi Shen</p>	73712
<p>Exposure Bracketing is All You Need for a High-Quality Image</p> <p>Zhilu Zhang, Shuohao Zhang, Renlong Wu, Zifei Yan, Wangmeng Zuo</p>	73733

<p>Interleaved Scene Graphs for Interleaved Text-And-Image Generation Assessment</p> <p>Dongping Chen, Ruoxi Chen, Shu Pu, Zhaoyi Liu, Yanru Wu, Caixi Chen, Benlin Liu, Yue Huang, Yao Wan, Pan Zhou, Ranjay Krishna</p>	73759
<p>Warm Diffusion: Recipe for Blur-Noise Mixture Diffusion Models</p> <p>Hao-Chien Hsueh, Wen-Hsiao Peng, Ching-Chun Huang</p>	73823
<p>Homomorphism Expressivity of Spectral Invariant Graph Neural Networks</p> <p>Jingchu Gai, Yiheng Du, Bohang Zhang, Haggai Maron, Liwei Wang</p>	73841
<p>Halton Scheduler for Masked Generative Image Transformer</p> <p>Victor Besnier, Mickael Chen, David Hurych, Eduardo Valle, Matthieu Cord</p>	73883
<p>Optimality and Adaptivity of Deep Neural Features for Instrumental Variable Regression</p> <p>Juno Kim, Dimitri Meunier, Arthur Gretton, Taiji Suzuki, Zhu Li</p>	73904
<p>Greener GRASS: Enhancing GNNs with Encoding, Rewiring, and Attention</p> <p>Tongzhou Liao, Barnabas Poczos</p>	73948
<p>Broadening Target Distributions for Accelerated Diffusion Models Via a Novel Analysis Approach</p> <p>Yuchen Liang, Peizhong Ju, Yingbin Liang, Ness Shroff</p>	73966
<p>Rethinking Reward Modeling in Preference-Based Large Language Model Alignment</p> <p>Hao Sun, Yunyi Shen, Jean-Francois Ton</p>	74029
<p>High-Quality Text-To-3D Character Generation with SparseCubes and Sparse Transformers.</p> <p>Jiachen Qian, Hongye Yang, Shuang Wu, Jingxi Xu, Feihu Zhang</p>	74067
<p>PathGen-1.6M: 1.6 Million Pathology Image-Text Pairs Generation Through Multi-Agent Collaboration</p> <p>Yuxuan Sun, Yunlong Zhang, Yixuan Si, Chenglu Zhu, Kai Zhang, Zhongyi Shui, Jingxiong Li, Xuan Gong, Xinheng Lyu, Tao Lin, Lin Yang</p>	74086
<p>Progress Or Regress? Self-Improvement Reversal in Post-Training</p> <p>Ting Wu, Xuefeng Li, Pengfei Liu</p>	74129
<p>DICE: End-To-End Deformation Capture of Hand-Face Interactions from a Single Image</p> <p>Qingxuan Wu, Zhiyang Dou, Sirui Xu, Soshi Shimada, Chen Wang, Zhengming Yu, Yuan Liu, Cheng Lin, Zeyu Cao, Taku Komura, Vladislav Golyanik, Christian Theobalt, Wenping Wang, Lingjie Liu</p>	74150
<p>3DMolFormer: A Dual-Channel Framework for Structure-Based Drug Discovery</p> <p>Xiuyuan Hu, Guoqing Liu, Can Chen, Yang Zhao, Hao Zhang, Xue Liu</p>	74174
<p>Metric-Driven Attributions for Vision Transformers</p> <p>Chase Walker, Sumit Kumar Jha, Rickard Ewetz</p>	74196
<p>Generating Likely Counterfactuals Using Sum-Product Networks</p> <p>Jirí Nemecek, Tomáš Pevný, Jakub Marecek</p>	74233

<p>Optimal Protocols for Continual Learning Via Statistical Physics and Control Theory</p> <p><i>Francesco Mori, Stefano Sarao Mannelli, Francesca Mignacco</i></p>	74265
<p>UIFace: Unleashing Inherent Model Capabilities to Enhance Intra-Class Diversity in Synthetic Face Recognition</p> <p><i>Xiao Lin, Yuge Huang, Jianqing Xu, Yuxi Mi, Shuigeng Zhou, Shouhong Ding</i></p>	74288
<p>SRSA: Skill Retrieval and Adaptation for Robotic Assembly Tasks</p> <p><i>Yijie Guo, Bingjie Tang, Iretiayo Akinola, Dieter Fox, Abhishek Gupta, Yashraj Narang</i></p>	74302
<p>Diffusion Models as Cartoonists: The Curious Case of High Density Regions</p> <p><i>Rafal Karczewski, Markus Heinonen, Vikas Garg</i></p>	74332
<p>SWE-Bench Multimodal: Do AI Systems Generalize to Visual Software Domains?</p> <p><i>John Yang, Carlos E Jimenez, Alex L Zhang, Kilian Lieret, Joyce Yang, Xindi Wu, Ori Press, Niklas Muennighoff, Gabriel Synnaeve, Karthik R Narasimhan, Diyi Yang, Sida Wang, Ofir Press</i></p>	74366
<p>SEMDICE: Off-Policy State Entropy Maximization Via Stationary Distribution Correction Estimation</p> <p><i>Jongmin Lee, Meiqi Sun, Pieter Abbeel</i></p>	74402
<p>Learning System Dynamics Without Forgetting</p> <p><i>Xikun Zhang, Dongjin Song, Yushan Jiang, Yixin Chen, Dacheng Tao</i></p>	74428
<p>Protecting Against Simultaneous Data Poisoning Attacks</p> <p><i>Neel Alex, Shoaib Ahmed Siddiqui, Amartya Sanyal, David Krueger</i></p>	74451
<p>A Distributional Approach to Uncertainty-Aware Preference Alignment Using Offline Demonstrations</p> <p><i>Sheng Xu, Bo Yue, Hongyuan Zha, Guiliang Liu</i></p>	74476
<p>Hallo2: Long-Duration and High-Resolution Audio-Driven Portrait Image Animation</p> <p><i>Jiahao Cui, Hui Li, Yao Yao, Hao Zhu, Hanlin Shang, Kaihui Cheng, Hang Zhou, Siyu Zhu, Jingdong Wang</i></p>	74502
<p>DiffPC: Diffusion-Based High Perceptual Fidelity Image Compression with Semantic Refinement</p> <p><i>Yichong Xia, Yimin Zhou, Jinpeng Wang, Baoyi An, Haoqian Wang, Yaowei Wang, Bin Chen</i></p>	74515
<p>OCEAN: Offline Chain-Of-Thought Evaluation and Alignment in Large Language Models</p> <p><i>Junda Wu, Xintong Li, Ruoyu Wang, Yu Xia, Yuxin Xiong, Jianing Wang, Tong Yu, Xiang Chen, Branislav Kveton, Lina Yao, Jingbo Shang, Julian McAuley</i></p>	74542
<p>N-ForGOT: Towards Not-Forgetting and Generalization of Open Temporal Graph Learning</p> <p><i>Liping Wang, Xujia Li, Jingshu Peng, Yue Wang, Chen Zhang, Yan Zhou, Lei Chen</i></p>	74562
<p>Systematic Outliers in Large Language Models</p> <p><i>Yongqi An, Xu Zhao, Tao Yu, Ming Tang, Jinqiao Wang</i></p>	74581

<p>GS-LiDAR: Generating Realistic LiDAR Point Clouds with Panoramic Gaussian Splatting</p> <p>Junzhe Jiang, Chun Gu, Yurui Chen, Li Zhang</p>	74603
<p>Aligning Generative Denoising with Discriminative Objectives Unleashes Diffusion for Visual Perception</p> <p>Ziqi Pang, Xin Xu, Yu-Xiong Wang</p>	74618
<p>Beyond Mere Token Analysis: A Hypergraph Metric Space Framework for Defending Against Socially Engineered LLM Attacks</p> <p>Manohar Kaul, Aditya Saibewar, Sadbhavana Babar</p>	74641
<p>Sensor-Invariant Tactile Representation</p> <p>Harsh Gupta, Yuchen Mo, Shengmiao Jin, Wenzhen Yuan</p>	74682
<p>Bayesian Image Regression with Soft-Thresholded Conditional Autoregressive Prior</p> <p>Yuliang Xu, Jian Kang</p>	74706
<p>CLDyB: Towards Dynamic Benchmarking for Continual Learning with Pre-Trained Models</p> <p>Shengzhuang Chen, Yikai Liao, Xiaoxiao Sun, Kede Ma, Ying Wei</p>	74731
<p>Smoothing the Shift: Towards Stable Test-Time Adaptation Under Complex Multimodal Noises</p> <p>Zirun Guo, Tao Jin</p>	74764
<p>A Riemannian Framework for Learning Reduced-Order Lagrangian Dynamics</p> <p>Katharina Friedl, Noémie Jaquier, Jens Lundell, Tamim Asfour, Danica Kragic</p>	74782
<p>Vec2Face: Scaling Face Dataset Generation with Loosely Constrained Vectors</p> <p>Haiyu Wu, Jaskirat Singh, Sicong Tian, Liang Zheng, Kevin Bowyer</p>	74810
<p>$\{\tau\}$-Bench: A Benchmark for <u>T</u>ool-<u>A</u>gent-<u>U</u>ser Interaction in Real-World Domains</p> <p>Shunyu Yao, Noah Shinn, Pedram Razavi, Karthik R Narasimhan</p>	74824
<p>Closed-Form Merging of Parameter-Efficient Modules for Federated Continual Learning</p> <p>Riccardo Salami, Pietro Buzzega, Matteo Mosconi, Jacopo Bonato, Luigi Sabetta, Simone Calderara</p>	74877
<p>Local Convergence of Simultaneous Min-Max Algorithms to Differential Equilibrium on Riemannian Manifold</p> <p>Sixin Zhang</p>	74903
<p>A Causal Lens for Learning Long-Term Fair Policies</p> <p>Jacob Lear, Lu Zhang</p>	74937
<p>ConvCodeWorld: Benchmarking Conversational Code Generation in Reproducible Feedback Environments</p> <p>Hojae Han, Seung-Won Hwang, Rajhans Samdani, Yuxiong He</p>	74953

<p>UniCBE: An Uniformity-Driven Comparing Based Evaluation Framework with Unified Multi-Objective Optimization</p> <p>Peiwen Yuan, Shaoxiong Feng, Yiwei Li, Xinglin Wang, Yueqi Zhang, Jiayi Shi, Chuyi Tan, Boyuan Pan, Yao Hu, Kan Li</p>	74993
<p>Multimodal Large Language Models for Inverse Molecular Design with Retrosynthetic Planning</p> <p>Gang Liu, Michael Sun, Wojciech Matusik, Meng Jiang, Jie Chen</p>	75014
<p>Can a Large Language Model Be a Gaslighter?</p> <p>Wei Li, Luyao Zhu, Yang Song, Ruixi Lin, Rui Mao, Yang You</p>	75042
<p>Biologically Plausible Brain Graph Transformer</p> <p>Ciyuan Peng, Yuelong Huang, Qichao Dong, Shuo Yu, Feng Xia, Chengqi Zhang, Yaochu Jin</p>	75071
<p>A Theory of Initialisation's Impact on Specialisation</p> <p>Devon Jarvis, Sebastian Lee, Clémentine Carla Juliette Dominé, Andrew M Saxe, Stefano Sarao Mannelli</p>	75098
<p>FairMT-Bench: Benchmarking Fairness for Multi-Turn Dialogue in Conversational LLMs</p> <p>Zhiting Fan, Ruizhe Chen, Tianxiang Hu, Zuozhu Liu</p>	75127
<p>Faster Inference of Flow-Based Generative Models Via Improved Data-Noise Coupling</p> <p>Aram Davtyan, Leello Tadesse Dadi, Volkan Cevher, Paolo Favaro</p>	75156
<p>Self-Introspective Decoding: Alleviating Hallucinations for Large Vision-Language Models</p> <p>Fushuo Huo, Wenchao Xu, Zhong Zhang, Haozhao Wang, Zhicheng Chen, Peilin Zhao</p>	75182
<p>Reasoning of Large Language Models Over Knowledge Graphs with Super-Relations</p> <p>Song Wang, Junhong Lin, Xiaojie Guo, Julian Shun, Jundong Li, Yada Zhu</p>	75206
<p>Large Language Models Often Say One Thing and Do Another</p> <p>Ruoxi Xu, Hongyu Lin, Xianpei Han, Jia Zheng, Weixiang Zhou, Le Sun, Yingfei Sun</p>	75225
<p>PhyloLM: Inferring the Phylogeny of Large Language Models and Predicting Their Performances in Benchmarks</p> <p>Nicolas Yax, Pierre-Yves Oudeyer, Stefano Palminteri</p>	75242
<p>Mechanism and Emergence of Stacked Attention Heads in Multi-Layer Transformers</p> <p>Tiberiu Mu?at</p>	75295
<p>Kronecker Mask and Interpretive Prompts Are Language-Action Video Learners</p> <p>Yang Jingyi, Zitong Yu, Nixiuming, He Jia, Hui Li</p>	75316
<p>Generator Matching: Generative Modeling with Arbitrary Markov Processes</p> <p>Peter Holderrieth, Marton Havasi, Jason Yim, Neta Shaul, Itai Gat, Tommi Jaakkola, Brian Karrer, Ricky T. Q. Chen, Yaron Lipman</p>	75345
<p>Online-To-Offline RL for Agent Alignment</p> <p>Xu Liu, Haobo Fu, Stefano V Albrecht, Qiang Fu, Shuai Li</p>	75412

<p>Towards Realistic UAV Vision-Language Navigation: Platform, Benchmark, and Methodology</p> <p>Xiangyu Wang, Donglin Yang, Ziqin Wang, Hohin Kwan, Jinyu Chen, Wenjun Wu, Hongsheng Li, Yue Liao, Si Liu</p>	75433
<p>\$InterLCM\$: Low-Quality Images as Intermediate States of Latent Consistency Models for Effective Blind Face Restoration</p> <p>Senmao Li, Kai Wang, Joost Van De Weijer, Fahad Shahbaz Khan, Chun-Le Guo, Shiqi Yang, Yaxing Wang, Jian Yang, Ming-Ming Cheng</p>	75452
<p>Efficient Alternating Minimization with Applications to Weighted Low Rank Approximation</p> <p>Zhao Song, Mingquan Ye, Junze Yin, Lichen Zhang</p>	75485
<p>From Probability to Counterfactuals: The Increasing Complexity of Satisfiability in Pearl's Causal Hierarchy</p> <p>Julian Dörfler, Benito Van Der Zander, Markus Bläser, Maciej Liskiewicz</p>	75536
<p>Robust Transfer of Safety-Constrained Reinforcement Learning Agents</p> <p>Markel Zubia, Thiago D. Simão, Nils Jansen</p>	75558
<p>Training-Free Dataset Pruning for Instance Segmentation</p> <p>Yalun Dai, Lingao Xiao, Ivor Tsang, Yang He</p>	75588
<p>Brain Bandit: A Biologically Grounded Neural Network for Efficient Control of Exploration</p> <p>Chen Jiang, Jiahui An, Yating Liu, Ni Ji</p>	75612
<p>PaLD: Detection of Text Partially Written by Large Language Models</p> <p>Eric Lei, Hsiang Hsu, Chun-Fu Chen</p>	75644
<p>Graph Transformers Dream of Electric Flow</p> <p>Xiang Cheng, Lawrence Carin, Suvrit Sra</p>	75663
<p>Linear Representations of Political Perspective Emerge in Large Language Models</p> <p>Junsol Kim, James Evans, Aaron Schein</p>	75685
<p>MoLEx: Mixture of Layer Experts for Fine-Tuning with Sparse Upcycling</p> <p>Rachel Teo, Tan Minh Nguyen</p>	75717
<p>Beyond the Convexity Assumption: Realistic Tabular Data Generation Under Quantifier-Free Real Linear Constraints</p> <p>Mihaela C. Stoian, Eleonora Giunchiglia</p>	75743
<p>PianoMotion10M: Dataset and Benchmark for Hand Motion Generation in Piano Performance</p> <p>Qijun Gan, Song Wang, Shengtao Wu, Jianke Zhu</p>	75773
<p>Unleashing the Power of Task-Specific Directions in Parameter Efficient Fine-Tuning</p> <p>Chongjie Si, Zhiyi Shi, Shifan Zhang, Xiaokang Yang, Hanspeter Pfister, Wei Shen</p>	75791

<p>Emergent Orientation Maps — Mechanisms, Coding Efficiency and Robustness</p> <p>Haixin Zhong, Haoyu Wang, Wei P Dai, Yuchao Huang, Mingyi Huang, Rubin Wang, Anna Wang Roe, Yuguo Yu</p>	75824
<p>Reframing Structure-Based Drug Design Model Evaluation Via Metrics Correlated to Practical Needs</p> <p>Bowen Gao, Haichuan Tan, Yanwen Huang, Minsi Ren, Xiao Huang, Wei-Ying Ma, Ya-Qin Zhang, Yanyan Lan</p>	75848
<p>Mixture of Attentions for Speculative Decoding</p> <p>Matthieu Zimmer, Milan Gritta, Gerasimos Lampouras, Haitham Bou Ammar, Jun Wang</p>	75873
<p>Medium-Difficulty Samples Constitute Smoothed Decision Boundary for Knowledge Distillation on Pruned Datasets</p> <p>Yudong Chen, Xuwei Xu, Frank De Hoog, Jiajun Liu, Sen Wang</p>	75892
<p>Minimal Impact ControlNet: Advancing Multi-ControlNet Integration</p> <p>Shikun Sun, Min Zhou, Zixuan Wang, Xubin Li, Tiezheng Ge, Zijie Ye, Xiaoyu Qin, Junliang Xing, Bo Zheng, Jia Jia</p>	75910
<p>Concept-ROT: Poisoning Concepts in Large Language Models with Model Editing</p> <p>Keltin Grimes, Marco Christiani, David Shriver, Marissa Catherine Connor</p>	75933
<p>HOPE for a Robust Parameterization of Long-Memory State Space Models</p> <p>Annan Yu, Michael W. Mahoney, N. Benjamin Erichson</p>	75959
<p>TIGER: Time-Frequency Interleaved Gain Extraction and Reconstruction for Efficient Speech Separation</p> <p>Mohan Xu, Kai Li, Guo Chen, Xiaolin Hu</p>	75986
<p>VCR: A Task for Pixel-Level Complex Reasoning in Vision Language Models Via Restoring Occluded Text</p> <p>Tianyu Zhang, Suyuchen Wang, Lu Li, Ge Zhang, Perouz Taslakian, Sai Rajeswar, Jie Fu, Bang Liu, Yoshua Bengio</p>	76004
<p>On Discriminative Probabilistic Modeling for Self-Supervised Representation Learning</p> <p>Bokun Wang, Yunwen Lei, Yiming Ying, Tianbao Yang</p>	76033
<p>Can We Trust Embodied Agents? Exploring Backdoor Attacks Against Embodied LLM-Based Decision-Making Systems</p> <p>Ruo Chen Jiao, Shaoyuan Xie, Justin Yue, Takami Sato, Lixu Wang, Yixuan Wang, Qi Alfred Chen, Qi Zhu</p>	76064
<p>An Evolved Universal Transformer Memory</p> <p>Edoardo Ceting, Qi Sun, Tianyu Zhao, Yujin Tang</p>	76095
<p>Swing-By Dynamics in Concept Learning and Compositional Generalization</p> <p>Yongyi Yang, Core Francisco Park, Ekdeep Singh Lubana, Maya Okawa, Wei Hu, Hidenori Tanaka</p>	76130

<p>Jailbreak Antidote: Runtime Safety-Utility Balance Via Sparse Representation Adjustment in Large Language Models</p> <p>Guobin Shen, Dongcheng Zhao, Yiting Dong, Xiang He, Yi Zeng</p>	76162
<p>Articulate-Anything: Automatic Modeling of Articulated Objects Via a Vision-Language Foundation Model</p> <p>Long Le, Jason Xie, William Liang, Hung-Ju Wang, Yue Yang, Yecheng Jason Ma, Kyle Vedder, Arjun Krishna, Dinesh Jayaraman, Eric Eaton</p>	76190
<p>Language Models Need Inductive Biases to Count Inductively</p> <p>Yingshan Chang, Yonatan Bisk</p>	76215
<p>How to Find the Exact Pareto Front for Multi-Objective MDPs?</p> <p>Yining Li, Peizhong Ju, Ness Shroff</p>	76239
<p>Deep Incomplete Multi-View Learning Via Cyclic Permutation of VAEs</p> <p>Xin Gao, Jian Pu</p>	76270
<p>What Are Good Positional Encodings for Directed Graphs?</p> <p>Yinan Huang, Haoyu Peter Wang, Pan Li</p>	76305
<p>MMed-RAG: Versatile Multimodal RAG System for Medical Vision Language Models</p> <p>Peng Xia, Kangyu Zhu, Haoran Li, Tianze Wang, Weijia Shi, Sheng Wang, Linjun Zhang, James Zou, Huaxiu Yao</p>	76330
<p>Robust LLM Safeguarding Via Refusal Feature Adversarial Training</p> <p>Lei Yu, Virginie Do, Karen Hambardzumyan, Nicola Cancedda</p>	76360
<p>ODE-Based Smoothing Neural Network for Reinforcement Learning Tasks</p> <p>Yinuo Wang, Wenxuan Wang, Xujie Song, Tong Liu, Yuming Yin, Liangfa Chen, Likun Wang, Jingliang Duan, Shengbo Eben Li</p>	76384
<p>Enhancing Multilingual Reasoning in LLMs: Insights from Cross-Linguistic Correlations and Optimal Data Proportions</p> <p>Jiangkuo Wang, Suyv Ma, Mingpeng Wei</p>	76404
<p>Cafe-Talk: Generating 3D Talking Face Animation with Multimodal Coarse- And Fine-Grained Control</p> <p>Hejia Chen, Haoxian Zhang, Shoulong Zhang, Xiaoqiang Liu, Sisi Zhuang, Zhangyuan, Pengfei Wan, Di Zhang, Shuai Li</p>	76419
<p>Train Small, Infer Large: Memory-Efficient LoRA Training for Large Language Models</p> <p>Jun Zhang, Jue Wang, Huan Li, Lidan Shou, Ke Chen, Yang You, Guiming Xie, Xuejian Gong, Kunlong Zhou</p>	76440

<p>Dynamic-SUPERB Phase-2: A Collaboratively Expanding Benchmark for Measuring the Capabilities of Spoken Language Models with 180 Tasks</p> <p>Chien-Yu Huang, Wei-Chih Chen, Shu-Wen Yang, Andy T. Liu, Chen-An Li, Yu-Xiang Lin, Wei-Cheng Tseng, Anuj Diwan, Yi-Jen Shih, Jiatong Shi, William Chen, Chih-Kai Yang, Xuanjun Chen, Chi-Yuan Hsiao, Puyuan Peng, Shih-Heng Wang, Chun-Yi Kuan, Ke-Han Lu, Kai-Wei Chang, Fabian Alejandro Ritter Gutierrez, Huang Kuan-Po, Siddhant Arora, You-Kuan Lin, Chuang Ming To, Eunjung Yeo, Calvin Chang, Chung-Ming Chien, Kwanghee Choi, Cheng-Hsiu Hsieh, Yi-Cheng Lin, Chee-En Yu, I-Hsiang Chiu, Heitor Guimarães, Jionghao Han, Tzu-Quan Lin, Tzu-Yuan Lin, Homu Chang, Ting-Wu Chang, Chun Wei Chen, Shou-Jen Chen, Yu-Hua Chen, Hsi-Chun Cheng, Kunal Dhawan, Jia-Lin Fang, Shi-Xin Fang, Kuan Yu Fang Chiang, Chi An Fu, Hsien-Fu Hsiao, Ching Yu Hsu, Shao-Syuan Huang, Lee Chen Wei, Hsi-Che Lin, Hsuan-Hao Lin, Hsuan-Ting Lin, Jian-Ren Lin, Ting-Chun Liu, Li-Chun Lu, Tsung-Min Pai, Ankita Pasad, Shih-Yun Shan Kuan, Suwon Shon, Yuxun Tang, Yun-Shao Tsai, Wei Jui Chiang, Tzu-Chieh Wei, Chengxi Wu, Dien-Ruei Wu, Chao-Han Huck Yang, Chieh-Chi Yang, Jia Qi Yip, Shao-Xiang Yuan, Haibin Wu, Karen Livescu, David Harwath, Shinji Watanabe, Hung-Yi Lee</p>	76470
<p>Progressive Compositionality in Text-To-Image Generative Models</p> <p>Xu Han, Linghao Jin, Xiaofeng Liu, Paul Pu Liang</p>	76536
<p>Redefining the Task of Bioactivity Prediction</p> <p>Yanwen Huang, Bowen Gao, Yinjun Jia, Hongbo Ma, Wei-Ying Ma, Ya-Qin Zhang, Yanyan Lan</p>	76559
<p>GROOT-2: Weakly Supervised Multimodal Instruction Following Agents</p> <p>Shaofei Cai, Bowei Zhang, Zihao Wang, Haowei Lin, Xiaojian Ma, Anji Liu, Yitao Liang</p>	76580
<p>DRoC: Elevating Large Language Models for Complex Vehicle Routing Via Decomposed Retrieval of Constraints</p> <p>Xia Jiang, Yaoxin Wu, Chenhao Zhang, Yingqian Zhang</p>	76603
<p>Benchmarking Predictive Coding Networks -- Made Simple</p> <p>Luca Pinchetti, Chang Qi, Oleh Lokshyn, Cornelius Emde, Amine M'Charrak, Mufeng Tang, Simon Frieder, Bayar Menzat, Gaspard Oliviers, Rafal Bogacz, Thomas Lukasiewicz, Tommaso Salvatori</p>	76625
<p>Personalized Visual Instruction Tuning</p> <p>Renjie Pi, Jianshu Zhang, Tianyang Han, Jipeng Zhang, Rui Pan, Tong Zhang</p>	76659
<p>Quest: Query-Centric Data Synthesis Approach for Long-Context Scaling of Large Language Model</p> <p>Chaochen Gao, Xing W, Qi Fu, Songlin Hu</p>	76679
<p>Offline Model-Based Optimization by Learning to Rank</p> <p>Rong-Xi Tan, Ke Xue, Shen-Huan Lyu, Haopu Shang, Yaowang, Yaoyuan Wang, Fu Sheng, Chao Qian</p>	76700
<p>C-CLIP: Multimodal Continual Learning for Vision-Language Model</p> <p>Wenzhuo Liu, Fei Zhu, Longhui Wei, Qi Tian</p>	76728

<p>ECD: A Machine Learning Benchmark for Predicting Enhanced-Precision Electronic Charge Density in Crystalline Inorganic Materials</p> <p><i>Pin Chen, Zexin Xu, Qing Mo, Hongjin Zhong, Fengyang Xu, Yutong Lu</i></p>	76745
<p>Improved Finite-Particle Convergence Rates for Stein Variational Gradient Descent</p> <p><i>Sayan Banerjee, Krishna Balasubramanian, Promit Ghosal</i></p>	76764
<p>HiSplat: Hierarchical 3D Gaussian Splatting for Generalizable Sparse-View Reconstruction</p> <p><i>Shengji Tang, Weicai Ye, Peng Ye, Weihao Lin, Yang Zhou, Tao Chen, Wanli Ouyang</i></p>	76786
<p>Surgical, Cheap, and Flexible: Mitigating False Refusal in Language Models Via Single Vector Ablation</p> <p><i>Xinpeng Wang, Chengzhi Hu, Paul Röttger, Barbara Plank</i></p>	76800
<p>Partial Gromov-Wasserstein Metric</p> <p><i>Yikun Bai, Rocio Diaz Martin, Abihith Kothapalli, Hengrong Du, Xinran Liu, Soheil Kolouri</i></p>	76820
<p>Spurious Forgetting in Continual Learning of Language Models</p> <p><i>Junhao Zheng, Xidi Cai, Shengjie Qiu, Qianli Ma</i></p>	76881
<p>Bundle Neural Network for Message Diffusion on Graphs</p> <p><i>Jacob Bamberger, Federico Barbero, Xiaowen Dong, Michael M. Bronstein</i></p>	76947
<p>CapeX: Category-Agnostic Pose Estimation from Textual Point Explanation</p> <p><i>Matan Rusanovsky, Or Hirschorn, Shai Avidan</i></p>	76973
<p>PMoE: Prompting Diverse Experts Together Wins More in Visual Adaptation</p> <p><i>Shentong Mo, Xufang Luo, Dongsheng Li</i></p>	76991
<p>A Decade's Battle on Dataset Bias: Are We There Yet?</p> <p><i>Zhuang Liu, Kaiming He</i></p>	77010
<p>UNSURE: Self-Supervised Learning with Unknown Noise Level and Stein's Unbiased Risk Estimate</p> <p><i>Julián Tachella, Mike Davies, Laurent Jacques</i></p>	77027
<p>TabM: Advancing Tabular Deep Learning with Parameter-Efficient Ensembling</p> <p><i>Yury Gorishniy, Akim Kotelnikov, Artem Babenko</i></p>	77048
<p>Collaborative Discrete-Continuous Black-Box Prompt Learning for Language Models</p> <p><i>Hualin Zhang, Haozhen Zhang, Zhekai Liu, Bin Gu, Yi Chang</i></p>	77085
<p>NGPT: Normalized Transformer with Representation Learning on the Hypersphere</p> <p><i>Ilya Loshchilov, Cheng-Ping Hsieh, Simeng Sun, Boris Ginsburg</i></p>	77109
<p>Herald: A Natural Language Annotated Lean 4 Dataset</p> <p><i>Guoxiong Gao, Yutong Wang, Jiedong Jiang, Qi Gao, Zihan Qin, Tianyi Xu, Bin Dong</i></p>	77134

<p>Adaptive Energy Alignment for Accelerating Test-Time Adaptation</p> <p>Wonjeong Choi, Do-Yeon Kim, Jungwuk Park, Jungmoon Lee, Younghyun Park, Dong-Jun Han, Jaekyun Moon</p>	77177
<p>Density Estimation with LLMs: A Geometric Investigation of In-Context Learning Trajectories</p> <p>Toni J. B. Liu, Nicolas Boulle, Raphaël Sarfati, Christopher Earls</p>	77200
<p>GameArena: Evaluating LLM Reasoning Through Live Computer Games</p> <p>Lanxiang Hu, Qiyu Li, Anze Xie, Nan Jiang, Ion Stoica, Haojian Jin, Hao Zhang</p>	77235
<p>P-SPIKESSM: HARNESSING PROBABILISTIC SPIKING STATE SPACE MODELS FOR LONG-RANGE DEPENDENCY TASKS</p> <p>Malyaban Bal, Abhronil Sengupta</p>	77267
<p>FreeCG: Free the Design Space of Clebsch-Gordan Transform for Machine Learning Force Fields</p> <p>Shihao Shao, Haoran Geng, Zun Wang, Qinghua Cui</p>	77283
<p>TorchTitan: One-Stop PyTorch Native Solution for Production Ready LLM Pretraining</p> <p>Wanchao Liang, Tianyu Liu, Less Wright, Will Constable, Andrew Gu, Chien-Chin Huang, Iris Zhang, Wei Feng, Howard Huang, Junjie Wang, Sanket Purandare, Gokul Nadathur, Stratos Idreos</p>	77311
<p>NExUME: Adaptive Training and Inference for DNNs Under Intermittent Power Environments</p> <p>Cyan Subhra Mishra, Deeksha Chaudhary, Jack Sampson, Mahmut Kandemir, Chita R. Das</p>	77332
<p>PiCO: Peer Review in LLMs Based on Consistency Optimization</p> <p>Kun-Peng Ning, Shuo Yang, Yuyang Liu, Jia-Yu Yao, Zhenhui Liu, Yonghong Tian, Yibing Song, Li Yuan</p>	77363
<p>Revisiting Random Walks for Learning on Graphs</p> <p>Jinwoo Kim, Olga Zaghen, Ayhan Suleymanzade, Youngmin Ryou, Seunghoon Hong</p>	77389
<p>JetFormer: An Autoregressive Generative Model of Raw Images and Text</p> <p>Michael Tschannen, André Susano Pinto, Alexander Kolesnikov</p>	77440
<p>Beyond Linear Approximations: A Novel Pruning Approach for Attention Matrix</p> <p>Yingyu Liang, Jiangxuan Long, Zhenmei Shi, Zhao Song, Yufa Zhou</p>	77462
<p>ChartMimic: Evaluating LMM's Cross-Modal Reasoning Capability Via Chart-To-Code Generation</p> <p>Cheng Yang, Chufan Shi, Yaxin Liu, Bo Shui, Junjie Wang, Mohan Jing, Linran Xu, Xinyu Zhu, Siheng Li, Yuxiang Zhang, Gongye Liu, Xiaomei Nie, Deng Cai, Yujiu Yang</p>	77512
<p>Your Weak LLM is Secretly a Strong Teacher for Alignment</p> <p>Leitian Tao, Yixuan Li</p>	77569
<p>Exploring the Camera Bias of Person Re-Identification</p> <p>Myungseo Song, Jin-Woo Park, Jong-Seok Lee</p>	77592

<p>CTSyn: A Foundation Model for Cross Tabular Data Generation</p> <p>Xiaofeng Lin, Chenheng Xu, Matthew Yang, Guang Cheng</p>	77612
<p>Temporal Reasoning Transfer from Text to Video</p> <p>Lei Li, Yuanxin Liu, Linli Yao, Peiyuan Zhang, Chenxin An, Lean Wang, Xu Sun, Lingpeng Kong, Qi Liu</p>	77634
<p>Towards Understanding the Robustness of Diffusion-Based Purification: A Stochastic Perspective</p> <p>Yiming Liu, Kezhao Liu, Yao Xiao, Ziyi Dong, Xiaogang Xu, Pengxu Wei, Liang Lin</p>	77667
<p>Transfusion: Predict the Next Token and Diffuse Images with One Multi-Modal Model</p> <p>Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob Kahn, Xuezhe Ma, Luke Zettlemoyer, Omer Levy</p>	77694
<p>Semantix: An Energy-Guided Sampler for Semantic Style Transfer</p> <p>Huiang He, Minghui Hu, Chuanxia Zheng, Chaoyue Wang, Tat-Jen Cham</p>	77718
<p>Understanding Optimization in Deep Learning with Central Flows</p> <p>Jeremy Cohen, Alex Damian, Ameet Talwalkar, J Zico Kolter, Jason D. Lee</p>	77746
<p>Learning to Clarify: Multi-Turn Conversations with Action-Based Contrastive Self-Training</p> <p>Maximillian Chen, Ruoxi Sun, Tomas Pfister, Sercan O Arik</p>	77861
<p>Vision-LSTM: xLSTM as Generic Vision Backbone</p> <p>Benedikt Alkin, Maximilian Beck, Korbinian Pöppel, Sepp Hochreiter, Johannes Brandstetter</p>	77897
<p>BTBS-LNS: Binarized-Tightening, Branch and Search on Learning LNS Policies for MIP</p> <p>Hao Yuan, Wenli Ouyang, Changwen Zhang, Yong Sun, Liming Gong, Junchi Yan</p>	77920
<p>Logical Consistency of Large Language Models in Fact-Checking</p> <p>Bishwamitra Ghosh, Sarah Hasan, Naheed Anjum Arafat, Arijit Khan</p>	77945
<p>High-Dimensional Bayesian Optimisation with Gaussian Process Prior Variational Autoencoders</p> <p>Siddharth Ramchandran, Manuel Haussmann, Harri Lähdesmäki</p>	77982
<p>Bridging the Gap Between Database Search and $\text{emph}\{De Novo\}$ Peptide Sequencing with SearchNovo</p> <p>Jun Xia, Sizhe Liu, Jingbo Zhou, Shaorong Chen, Hongxin Xiang, Zicheng Liu, Yue Liu, Stan Z. Li</p>	78009
<p>Higher-Order Graphon Neural Networks: Approximation and Cut Distance</p> <p>Daniel Herbst, Stefanie Jegelka</p>	78025
<p>Instructional Segment Embedding: Improving LLM Safety with Instruction Hierarchy</p> <p>Tong Wu, Shujian Zhang, Kaiqiang Song, Silei Xu, Sanqiang Zhao, Ravi Agrawal, Sathish Reddy Indurthi, Chong Xiang, Prateek Mittal, Wenxuan Zhou</p>	78078
<p>Neural Spacetimes for DAG Representation Learning</p> <p>Haitz Sáez De Ocariz Borde, Anastasis Kratsios, Marc T. Law, Xiaowen Dong, Michael M. Bronstein</p>	78109

<p>Towards Understanding Text Hallucination of Diffusion Models Via Local Generation Bias</p> <p>Rui Lu, Runzhe Wang, Kaifeng Lyu, Xitai Jiang, Gao Huang, Mengdi Wang</p>	78149
<p>LiveBench: A Challenging, Contamination-Limited LLM Benchmark</p> <p>Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Benjamin Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Sreemanti Dey, Shubh-Agrawal, Sandeep Singh Sandha, Siddartha Venkat Naidu, Chinmay Hegde, Yann Lecun, Tom Goldstein, Willie Neiswanger, Micah Goldblum</p>	78177
<p>Energy-Based Diffusion Language Models for Text Generation</p> <p>Minkai Xu, Tomas Geffner, Karsten Kreis, Weili Nie, Yilun Xu, Jure Leskovec, Stefano Ermon, Arash Vahdat</p>	78214
<p>Flow: Modularized Agentic Workflow Automation</p> <p>Boye Niu, Yiliao Song, Kai Lian, Yifan Shen, Yu Yao, Kun Zhang, Tongliang Liu</p>	78235
<p>MetaMetrics: Calibrating Metrics for Generation Tasks Using Human Preferences</p> <p>Genta Indra Winata, David Anugraha, Lucky Susanto, Garry Kuwanto, Derry Tanti Wijaya</p>	78264
<p>Language Models Are Implicitly Continuous</p> <p>Samuele Marro, Davide Evangelista, X. Angelo Huang, Emanuele La Malfa, Michele Lombardi, Michael J. Wooldridge</p>	78300
<p>Your Absorbing Discrete Diffusion Secretly Models the Conditional Distributions of Clean Data</p> <p>Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, Chongxuan Li</p>	78341
<p>Learnable Expansion of Graph Operators for Multi-Modal Feature Fusion</p> <p>Dexuan Ding, Lei Wang, Liyun Zhu, Tom Gedeon, Piotr Koniusz</p>	78379
<p>REEF: Representation Encoding Fingerprints for Large Language Models</p> <p>Jie Zhang, Dongrui Liu, Chen Qian, Linfeng Zhang, Yong Liu, Yu Qiao, Jing Shao</p>	78402
<p>Towards Learning High-Precision Least Squares Algorithms with Sequence Models</p> <p>Jerry Weihong Liu, Jessica Grogan, Owen M Dugan, Ashish Rao, Simran Arora, Atri Rudra, Christopher Re</p>	78428
<p>Improved Convergence Rate for Diffusion Probabilistic Models</p> <p>Gen Li, Yuchen Jiao</p>	78503
<p>Group Downsampling with Equivariant Anti-Aliasing</p> <p>Md Ashiqur Rahman, Raymond A. Yeh</p>	78515
<p>ComPC: Completing a 3D Point Cloud with 2D Diffusion Priors</p> <p>Tianxin Huang, Zhiwen Yan, Yuyang Zhao, Gim Hee Lee</p>	78546
<p>Lean-STaR: Learning to Interleave Thinking and Proving</p> <p>Haohan Lin, Zhiqing Sun, Sean Welleck, Yiming Yang</p>	78566

<p>Flow-Based Variational Mutual Information: Fast and Flexible Approximations</p> <p>Caleb Dahlke, Jason Pacheco</p>	78588
<p>Context-Parametric Inversion: Why Instruction Finetuning May Not Actually Improve Context Reliance</p> <p>Sachin Goyal, Christina Baek, J Zico Kolter, Aditi Raghunathan</p>	78611
<p>MAST: Model-Agnostic Sparsified Training</p> <p>Yury Demidovich, Grigory Malinovsky, Egor Shulgin, Peter Richtárik</p>	78644
<p>Provable Robust Overfitting Mitigation in Wasserstein Distributionally Robust Optimization</p> <p>Shuang Liu, Yihan Wang, Yifan Zhu, Yibo Miao, Xiao-Shan Gao</p>	78696
<p>Value-Incentivized Preference Optimization: A Unified Approach to Online and Offline RLHF</p> <p>Shicong Cen, Jincheng Mei, Katayoon Goshvadi, Hanjun Dai, Tong Yang, Sherry Yang, Dale Schuurmans, Yuejie Chi, Bo Dai</p>	78719
<p>When Prompt Engineering Meets Software Engineering: CNL-P as Natural and Robust "APIs" for Human-AI Interaction</p> <p>Zhenchang Xing, Yang Liu, Zhuo Cheng, Qing Huang, Dehai Zhao, Daniel Sun, Chenhua Liu</p>	78745
<p>Attention with Markov: A Curious Case of Single-Layer Transformers</p> <p>Ashok Vardhan Makkuva, Marco Bondaschi, Adway Girish, Alliot Nagle, Martin Jaggi, Hyeji Kim, Michael Gastpar</p>	78773
<p>Fast Training of Sinusoidal Neural Fields Via Scaling Initialization</p> <p>Taesun Yeom, Sangyoon Lee, Jaeho Lee</p>	78800
<p>Learning the Optimal Stopping for Early Classification Within Finite Horizons Via Sequential Probability Ratio Test</p> <p>Akinori F. Ebihara, Taiki Miyagawa, Kazuyuki Sakurai, Hitoshi Imaoka</p>	78843
<p>Adaptive Backtracking Line Search</p> <p>Joao V. Cavalcanti, Laurent Lessard, Ashia C. Wilson</p>	78905
<p>Proactive Agent: Shifting LLM Agents from Reactive Responses to Active Assistance</p> <p>Yaxi Lu, Shenzhi Yang, Cheng Qian, Guirong Chen, Qinyu Luo, Yesai Wu, Huadong Wang, Xin Cong, Zhong Zhang, Yankai Lin, Weiwen Liu, Yasheng Wang, Zhiyuan Liu, Fangming Liu, Maosong Sun</p>	78952
<p>Fat-To-Thin Policy Optimization: Offline Reinforcement Learning with Sparse Policies</p> <p>Lingwei Zhu, Han Wang, Yukie Nagai</p>	78979
<p>Meta-Dynamical State Space Models for Integrative Neural Data Analysis</p> <p>Ayesha Vermani, Josue Nassar, Hyungju Jeon, Matthew Dowling, Il Memming Park</p>	78996
<p>3D Vision-Language Gaussian Splatting</p> <p>Qucheng Peng, Benjamin Planche, Zhongpai Gao, Meng Zheng, Anwesha Choudhuri, Terrence Chen, Chen Chen, Ziyang Wu</p>	79020

Reward Dimension Reduction for Scalable Multi-Objective Reinforcement Learning Giseung Park, Youngchul Sung	79044
VideoGrain: Modulating Space-Time Attention for Multi-Grained Video Editing Xiangpeng Yang, Linchao Zhu, Hehe Fan, Yi Yang	79063
Do as We Do, Not as You Think: The Conformity of Large Language Models Zhiyuan Weng, Guikun Chen, Wenguan Wang	79082
Better than Your Teacher: LLM Agents that Learn from Privileged AI Feedback Sanjiban Choudhury, Paloma Sodhi	79121
Is In-Context Learning Sufficient for Instruction Following in LLMs? Hao Zhao, Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion	79158
Conservative Contextual Bandits: Beyond Linear Representations Rohan Deb, Mohammad Ghavamzadeh, Arindam Banerjee	79182
Lotus: Diffusion-Based Visual Foundation Model for High-Quality Dense Prediction Jing He, Haodong Li, Wei Yin, Yixun Liang, Leheng Li, Kaiqiang Zhou, Hongbo Zhang, Bingbing Liu, Ying-Cong Chen	79224
Tractable Multi-Agent Reinforcement Learning Through Behavioral Economics Eric Mazumdar, Kishan Panaganti, Laixi Shi	79238
Differentiation and Specialization of Attention Heads Via the Refined Local Learning Coefficient George Wang, Jesse Hoogland, Stan Van Wingerden, Zach Furman, Daniel Mufet	79253
Robustness Reprogramming for Representation Learning Zhichao Hou, Mohamadali Torkamani, Hamid Krim, Xiaorui Liu	79299
Data Unlearning in Diffusion Models Silas Alberti, Kenan Hasanaliyev, Manav Shah, Stefano Ermon	79322
One Model Transfer to All: On Robust Jailbreak Prompts Generation Against LLMs Linbao Li, Yannan Liu, Daojing He, Yu Li	79339
LiveXiv - A Multi-Modal Live Benchmark Based on Arxiv Papers Content Nimrod Shabtay, Felipe Maia Polo, Sivan Doveh, Wei Lin, Muhammad Jehanzeb Mirza, Leshem Choshen, Mikhail Yurochkin, Yuekai Sun, Assaf Arbelle, Leonid Karlinsky, Raja Giryes	79356
6DGS: Enhanced Direction-Aware Gaussian Splatting for Volumetric Rendering Zhongpai Gao, Benjamin Planche, Meng Zheng, Anwesa Choudhuri, Terrence Chen, Ziyang Wu	79389
How Efficient is LLM-Generated Code? a Rigorous & High-Standard Benchmark Ruizhong Qiu, Weiliang Will Zeng, James Ezick, Christopher Lott, Hanghang Tong	79408
Data-Adaptive Differentially Private Prompt Synthesis for In-Context Learning Fengyu Gao, Ruida Zhou, Tianhao Wang, Cong Shen, Jing Yang	79437

<p>Adding Conditional Control to Diffusion Models with Reinforcement Learning</p> <p>Yulai Zhao, Masatoshi Uehara, Gabriele Scalia, Sunyuan Kung, Tommaso Biancalani, Sergey Levine, Ehsan Hajiramezani</p>	79466
<p>KOR-Bench: Benchmarking Language Models on Knowledge-Orthogonal Reasoning Tasks</p> <p>Kaijing Ma, Xeron Du, Yunran Wang, Haoran Zhang, Zhoufutuwen, Xingwei Qu, Jian Yang, Jiaheng Liu, Minghao Liu, Xiang Yue, Wenhao Huang, Ge Zhang</p>	79491
<p>Locality Sensitive Avatars from Video</p> <p>Chunjin Song, Zhijie Wu, Shih-Yang Su, Bastian Wandt, Leonid Sigal, Helge Rhodin</p>	79591
<p>Accelerated Over-Relaxation Heavy-Ball Method: Achieving Global Accelerated Convergence with Broad Generalization</p> <p>Jingrong Wei, Long Chen</p>	79621
<p>TabDiff: A Mixed-Type Diffusion Model for Tabular Data Generation</p> <p>Juntong Shi, Minkai Xu, Harper Hua, Hengrui Zhang, Stefano Ermon, Jure Leskovec</p>	79649
<p>Adversarial Machine Unlearning</p> <p>Zonglin Di, Sixie Yu, Yevgeniy Vorobeychik, Yang Liu</p>	79672
<p>SVG: 3D Stereoscopic Video Generation Via Denoising Frame Matrix</p> <p>Peng Dai, Feitong Tan, Qiangeng Xu, David Futschik, Ruofei Du, Sean Fanello, Xiaojuan Qi, Yinda Zhang</p>	79694
<p>Test-Time Adaptation for Regression by Subspace Alignment</p> <p>Kazuki Adachi, Shin'ya Yamaguchi, Atsutoshi Kumagai, Tomoki Hamagami</p>	79719
<p>U-Nets as Belief Propagation: Efficient Classification, Denoising, and Diffusion in Generative Hierarchical Models</p> <p>Song Mei</p>	79744
<p>Efficient Masked AutoEncoder for Video Object Counting and a Large-Scale Benchmark</p> <p>Bing Cao, Quanhao Lu, Jiekang Feng, Qilong Wang, Pengfei Zhu, Qinghua Hu</p>	79779
<p>Implicit Neural Surface Deformation with Explicit Velocity Fields</p> <p>Lu Sang, Zehranaz Canfes, Dongliang Cao, Florian Bernard, Daniel Cremers</p>	79801
<p>Context-Alignment: Activating and Enhancing LLMs Capabilities in Time Series</p> <p>Yuxiao Hu, Qian Li, Dongxiao Zhang, Jinyue Yan, Yuntian Chen</p>	79826
<p>Heavy-Tailed Diffusion with Denoising Levy Probabilistic Models</p> <p>Dario Shariatian, Umut Simsekli, Alain Oliviero Durmus</p>	79853
<p>Release the Powers of Prompt Tuning: Cross-Modality Prompt Transfer</p> <p>Ningyuan Zhang, Jie Lu, Keqiyun Li, Zhen Fang, Guangquan Zhang</p>	79887

<p>Emerging Safety Attack and Defense in Federated Instruction Tuning of Large Language Models</p> <p>Rui Ye, Jingyi Chai, Xiangrui Liu, Yaodong Yang, Yanfeng Wang, Siheng Chen</p>	79915
<p>SpaceGNN: Multi-Space Graph Neural Network for Node Anomaly Detection with Extremely Limited Labels</p> <p>Xiangyu Dong, Xingyi Zhang, Lei Chen, Mingxuan Yuan, Sibow Wang</p>	79934
<p>Cheating Automatic LLM Benchmarks: Null Models Achieve High Win Rates</p> <p>Xiaosen Zheng, Tianyu Pang, Chao Du, Qian Liu, Jing Jiang, Min Lin</p>	79958
<p>Deep Learning Alternatives of the Kolmogorov Superposition Theorem</p> <p>Leonardo Ferreira Guilhoto, Paris Perdikaris</p>	79992
<p>Aioli: A Unified Optimization Framework for Language Model Data Mixing</p> <p>Mayee F Chen, Michael Y. Hu, Nicholas Lourie, Kyunghyun Cho, Christopher Re</p>	80021
<p>Autocorrelation Matters: Understanding the Role of Initialization Schemes for State Space Models</p> <p>Fusheng Liu, Qianxiao Li</p>	80065
<p>Text-To-Image Rectified Flow as Plug-And-Play Priors</p> <p>Xiaofeng Yang, Chen Cheng, Xulei Yang, Fayao Liu, Guosheng Lin</p>	80086
<p>Graph Neural Networks Can (Often) Count Substructures</p> <p>Paolo Pellizzoni, Till Hendrik Schulz, Karsten Borgwardt</p>	80111
<p>Effective Post-Training Embedding Compression Via Temperature Control in Contrastive Training</p> <p>Georgiana Dinu, Corey D Barrett, Yi Xiang, Miguel Romero Calvo, Anna Currey, Xing Niu</p>	80138
<p>Enhance Multi-View Classification Through Multi-Scale Alignment and Expanded Boundary</p> <p>Yuena Lin, Yiyuan Wang, Gengyu Lyu, Yongjian Deng, Haichun Cai, Huibin Lin, Haobo Wang, Zhen Yang</p>	80155
<p>Scaling Offline Model-Based RL Via Jointly-Optimized World-Action Model Pretraining</p> <p>Jie Cheng, Ruixi Qiao, Yingwei Ma, Binhua Li, Gang Xiong, Qinghai Miao, Yongbin Li, Yisheng Lv</p>	80174
<p>Dynamic Mixture of Experts: An Auto-Tuning Approach for Efficient Transformer Models</p> <p>Yongxin Guo, Zhenglin Cheng, Xiaoying Tang, Zhaopeng Tu, Tao Lin</p>	80197
<p>Fundamental Limitations on Subquadratic Alternatives to Transformers</p> <p>Josh Alman, Hantao Yu</p>	80227
<p>On the Benefits of Attribute-Driven Graph Domain Adaptation</p> <p>Ruiyi Fang, Bingheng Li, Zhao Kang, Qiuhaio Zeng, Nima Hosseini Dashtbayaz, Ruizhi Pu, Boyu Wang, Charles Ling</p>	80251

<p>ZooProbe: A Data Engine for Evaluating, Exploring, and Evolving Large-Scale Training Data for Multimodal LLMs</p> <p><i>Yi-Kai Zhang, Shiyin Lu, Qing-Guo Chen, De-Chuan Zhan, Han-Jia Ye</i></p>	80272
<p>CryoFM: A Flow-Based Foundation Model for Cryo-EM Densities</p> <p><i>Yi Zhou, Yilai Li, Jing Yuan, Quanquan Gu</i></p>	80298
<p>PARTNR: A Benchmark for Planning and Reasoning in Embodied Multi-Agent Tasks</p> <p><i>Matthew Chang, Gunjan Chhablani, Alexander Clegg, Mikael Dallaire Cote, Ruta Desai, Michal Hlavac, Vladimir Karashchuk, Jacob Krantz, Roozbeh Mottaghi, Priyam Parashar, Siddharth Patki, Ishita Prasad, Xavier Puig, Akshara Rai, Ram Ramrakhya, Daniel Tran, Joanne Truong, John M Turner, Eric Undersander, Tsung-Yen Yang</i></p>	80327
<p>Breaking Mental Set to Improve Reasoning Through Diverse Multi-Agent Debate</p> <p><i>Yexiang Liu, Jie Cao, Zekun Li, Ran He, Tieniu Tan</i></p>	80391
<p>You Only Sample Once: Taming One-Step Text-To-Image Synthesis by Self-Cooperative Diffusion GANs</p> <p><i>Yihong Luo, Xiaolong Chen, Xinghua Qu, Tianyang Hu, Jing Tang</i></p>	80451
<p>MoE++: Accelerating Mixture-Of-Experts Methods with Zero-Computation Experts</p> <p><i>Peng Jin, Bo Zhu, Li Yuan, Shuicheng Yan</i></p>	80473
<p>Classic but Everlasting: Traditional Gradient-Based Algorithms Converge Fast Even in Time-Varying Multi-Player Games</p> <p><i>Yanzheng Chen, Jun Yu</i></p>	80498
<p>Adversarially Robust Anomaly Detection Through Spurious Negative Pair Mitigation</p> <p><i>Hossein Mirzaei, Mojtaba Nafez, Jafar Habibi, Mohammad Sabokrou, Mohammad Hossein Rohban</i></p>	80535
<p>Sparse Learning for State Space Models on Mobile</p> <p><i>Xuan Shen, Hangyu Zheng, Yifan Gong, Zhenglun Kong, Changdi Yang, Zheng Zhan, Yushu Wu, Xue Lin, Yanzhi Wang, Pu Zhao, Wei Niu</i></p>	80563
<p>Uncovering Overfitting in Large Language Model Editing</p> <p><i>Mengqi Zhang, Xiaotian Ye, Qiang Liu, Shu Wu, Pengjie Ren, Zhumin Chen</i></p>	80578
<p>Lightning-Fast Image Inversion and Editing for Text-To-Image Diffusion Models</p> <p><i>Dvir Samuel, Barak Meiri, Haggai Maron, Yoad Tewel, Nir Darshan, Shai Avidan, Gal Chechik, Rami Ben-Ari</i></p>	80597
<p>Node Identifiers: Compact, Discrete Representations for Efficient Graph Learning</p> <p><i>Yuankai Luo, Hongkang Li, Qijiong Liu, Lei Shi, Xiao-Ming Wu</i></p>	80618
<p>Automated Design of Agentic Systems</p> <p><i>Shengran Hu, Cong Lu, Jeff Clune</i></p>	80648
<p>Learning Harmonized Representations for Speculative Sampling</p> <p><i>Lefan Zhang, Xiaodan Wang, Yanhua Huang, Ruiwen Xu</i></p>	80682

<p>A Simple Approach to Unifying Diffusion-Based Conditional Generation</p> <p>Xirui Li, Charles Herrmann, Kelvin C. K. Chan, Yinxiao Li, Deqing Sun, Chao Ma, Ming-Hsuan Yang</p>	80704
<p>Dynamic Diffusion Transformer</p> <p>Wangbo Zhao, Yizeng Han, Jiasheng Tang, Kai Wang, Yibing Song, Gao Huang, Fan Wang, Yang You</p>	80725
<p>MUSE: Machine Unlearning Six-Way Evaluation for Language Models</p> <p>Weijia Shi, Jaechan Lee, Yangsibo Huang, Sadhika Malladi, Jieyu Zhao, Ari Holtzman, Daogao Liu, Luke Zettlemoyer, Noah A. Smith, Chiyuan Zhang</p>	80758
<p>HShare: Fast LLM Decoding by Hierarchical Key-Value Sharing</p> <p>Huaijin Wu, Lianqiang Li, Hantao Huang, Tu Yi, Jihang Zhang, Minghui Yu, Junchi Yan</p>	80780
<p>Epistemic Monte Carlo Tree Search</p> <p>Yaniv Oren, Viliam Vavoc, Matthijs T. J. Spaan, Wendelin Boehmer</p>	80797
<p>Adversarial Generative Flow Network for Solving Vehicle Routing Problems</p> <p>Ni Zhang, Jingfeng Yang, Zhiguang Cao, Xu Chi</p>	80824
<p>Boosting Neural Combinatorial Optimization for Large-Scale Vehicle Routing Problems</p> <p>Fu Luo, Xi Lin, Yaoxin Wu, Zhenkun Wang, Tong Xialiang, Mingxuan Yuan, Qingfu Zhang</p>	80839
<p>A Theory for Token-Level Harmonization in Retrieval-Augmented Generation</p> <p>Shicheng Xu, Liang Pang, Huawei Shen, Xueqi Cheng</p>	80866
<p>Cybench: A Framework for Evaluating Cybersecurity Capabilities and Risks of Language Models</p> <p>Andy K Zhang, Neil Perry, Riya Dulepet, Joey Ji, Celeste Menders, Justin W Lin, Eliot Jones, Gashon Hussein, Samantha Liu, Donovan Julian Jasper, Pura Peetathawatchai, Ari Glenn, Vikram Sivashankar, Daniel Zamoshchin, Leo Glikbarg, Derek Askaryar, Haoxiang Yang, Aolin Zhang, Rishi Alluri, Nathan Tran, Rinnara Sangpisit, Kenny O Oseleononmen, Dan Boneh, Daniel E. Ho, Percy Liang</p>	80893
<p>Scaling and Evaluating Sparse Autoencoders</p> <p>Leo Gao, Tom Dupre La Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya Sutskever, Jan Leike, Jeffrey Wu</p>	81043
<p>Flow Matching with General Discrete Paths: A Kinetic-Optimal Perspective</p> <p>Neta Shaul, Itai Gat, Marton Havasi, Daniel Severo, Anuroop Sriram, Peter Holderrieth, Brian Karrer, Yaron Lipman, Ricky T. Q. Chen</p>	81077
<p>Learning to Adapt Frozen CLIP for Few-Shot Test-Time Domain Adaptation</p> <p>Zhixiang Chi, Li Gu, Huan Liu, Ziqiang Wang, Yanan Wu, Yang Wang, Konstantinos N Plataniotis</p>	81110
<p>Quantum (Inspired) D^2-Sampling with Applications</p> <p>Poojan Chetan Shah, Ragesh Jaiswal</p>	81132

<p>On the Relation Between Trainability and Dequantization of Variational Quantum Learning Models</p> <p><i>Elies Gil-Fuster, Casper Gyurik, Adrian Perez-Salinas, Vedran Dunjko</i></p>	81167
<p>Provence: Efficient and Robust Context Pruning for Retrieval-Augmented Generation</p> <p><i>Nadezhda Chirkova, Thibault Formal, Vassilina Nikoulina, Stéphane Clinchant</i></p>	81192
<p>Efficient Model Editing with Task-Localized Sparse Fine-Tuning</p> <p><i>Leonardo Iurada, Marco Ciccone, Tatiana Tommasi</i></p>	81217
<p>VLM2Vec: Training Vision-Language Models for Massive Multimodal Embedding Tasks</p> <p><i>Ziyan Jiang, Rui Meng, Xinyi Yang, Semih Yavuz, Yingbo Zhou, Wenhu Chen</i></p>	81242
<p>JPEG Inspired Deep Learning</p> <p><i>Ahmed H. Salamah, Kaixiang Zheng, Yiwen Liu, En-Hui Yang</i></p>	81267
<p>Gradient-Free Generation for Hard-Constrained Systems</p> <p><i>Chaoran Cheng, Boran Han, Danielle C. Maddix, Abdul Fatir Ansari, Andrew Stuart, Michael W. Mahoney, Bernie Wang</i></p>	81291
<p>Phidias: A Generative Model for Creating 3D Content from Text, Image, and 3D Conditions with Reference-Augmented Diffusion</p> <p><i>Zhenwei Wang, Tengfei Wang, Zexin He, Gerhard Petrus Hancke, Ziwei Liu, Rynson W. H. Lau</i></p>	81321
<p>Improving Neural Network Accuracy by Concurrently Training with a Twin Network</p> <p><i>Benjamin Vandersmissen, Lucas Deckers, Jose Oramas</i></p>	81341
<p>Scaling Long Context Training Data by Long-Distance Referrals</p> <p><i>Yonghao Zhuang, Lanxiang Hu, Longfei Yun, Souvik Kundu, Zhengzhong Liu, Eric P. Xing, Hao Zhang</i></p>	81355
<p>Studying the Interplay Between the Actor and Critic Representations in Reinforcement Learning</p> <p><i>Samuel Garcin, Trevor McInroe, Pablo Samuel Castro, Christopher G. Lucas, David Abel, Prakash Panangaden, Stefano V Albrecht</i></p>	81373
<p>MMAU: A Massive Multi-Task Audio Understanding and Reasoning Benchmark</p> <p><i>S Sakshi, Utkarsh Tyagi, Sonal Kumar, Ashish Seth, Ramaneswaran Selvakumar, Oriol Nieto, Ramani Duraiswami, Sreyan Ghosh, Dinesh Manocha</i></p>	81400
<p>ThinkBot: Embodied Instruction Following with Thought Chain Reasoning</p> <p><i>Guanxing Lu, Ziwei Wang, Changliu Liu, Jiwen Lu, Yansong Tang</i></p>	81436
<p>Training-Free Diffusion Model Alignment with Sampling Demons</p> <p><i>Po-Hung Yeh, Kuang-Huei Lee, Jun-Cheng Chen</i></p>	81455
<p>Stochastic Semi-Gradient Descent for Learning Mean Field Games with Population-Aware Function Approximation</p> <p><i>Chenyu Zhang, Xu Chen, Xuan Di</i></p>	81492

<p>ElasticTok: Adaptive Tokenization for Image and Video</p> <p><i>Wilson Yan, Volodymyr Mnih, Aleksandra Faust, Matei Zaharia, Pieter Abbeel, Hao Liu</i></p>	81540
<p>Knowledge Localization: Mission Not Accomplished? Enter Query Localization!</p> <p><i>Yuheng Chen, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao</i></p>	81561
<p>Measuring Memorization in RLHF for Code Completion</p> <p><i>Jamie Hayes, Ilya Shumailov, William P. Porter, Aneesh Pappu</i></p>	81586
<p>Holographic Node Representations: Pre-Training Task-Agnostic Node Embeddings</p> <p><i>Beatrice Bevilacqua, Joshua Robinson, Jure Leskovec, Bruno Ribeiro</i></p>	81615
<p>MGCFFN: A Neural MultiGrid Solver with Novel Fourier Neural Network for High Wave Number Helmholtz Equations</p> <p><i>Yan Xie, Minrui Lv, Chen-Song Zhang</i></p>	81651
<p>TopoNets: High Performing Vision and Language Models with Brain-Like Topography</p> <p><i>Mayukh Deb, Mainak Deb, Apurva Ratan Murty</i></p>	81669
<p>The Belief State Transformer</p> <p><i>Edward S. Hu, Kwangjun Ahn, Qinghua Liu, Haoran Xu, Manan Tomar, Ada Langford, Dinesh Jayaraman, Alex Lamb, John Langford</i></p>	81690
<p>Rationalizing and Augmenting Dynamic Graph Neural Networks</p> <p><i>Guibin Zhang, Yiyang Qi, Ziyang Cheng, Yanwei Yue, Dawei Cheng, Jian Guo</i></p>	81711
<p>Boosting Ray Search Procedure of Hard-Label Attacks with Transfer-Based Priors</p> <p><i>Chen Ma, Xinjie Xu, Shuyu Cheng, Qi Xuan</i></p>	81732
<p>Provable Uncertainty Decomposition Via Higher-Order Calibration</p> <p><i>Gustaf Ahndritz, Aravind Gollakota, Parikshit Gopalan, Charlotte Peale, Udi Wieder</i></p>	81780
<p>BAMDP Shaping: A Unified Framework for Intrinsic Motivation and Reward Shaping</p> <p><i>Aly Lidayan, Michael D Dennis, Stuart Russell</i></p>	81822
<p>State Space Model Meets Transformer: A New Paradigm for 3D Object Detection</p> <p><i>Chuxin Wang, Wenfei Yang, Xiang Liu, Tianzhu Zhang</i></p>	81852
<p>Temporal Heterogeneous Graph Generation with Privacy, Utility, and Efficiency</p> <p><i>Xinyu He, Dongqi Fu, Hanghang Tong, Ross Maciejewski, Jingrui He</i></p>	81876
<p>MiniPLM: Knowledge Distillation for Pre-Training Language Models</p> <p><i>Yuxian Gu, Hao Zhou, Fandong Meng, Jie Zhou, Minlie Huang</i></p>	81897
<p>An Image is Worth More than 16x16 Patches: Exploring Transformers on Individual Pixels</p> <p><i>Duy Kien Nguyen, Mido Assran, Unnat Jain, Martin R. Oswald, Cees G. M. Snoek, Xinlei Chen</i></p>	81924
<p>Surprising Effectiveness of Pretraining Ternary Language Model at Scale</p> <p><i>Ayush Kaushal, Tejas Vaidhya, Arnab Kumar Mondal, Tejas Pandey, Aaryan Bhagat, Irina Rish</i></p>	81945

SV4D: Dynamic 3D Content Generation with Multi-Frame and Multi-View Consistency Yiming Xie, Chun-Han Yao, Vikram Voleti, Huaizu Jiang, Varun Jampani	81993
Discriminator-Guided Embodied Planning for LLM Agent Haofu Qian, Chenjia Bai, Jiatao Zhang, Fei Wu, Wei Song, Xuelong Li	82014
FIRING-Net: A Filtered Feature Recycling Network for Speech Enhancement Xinmeng Xu, Yiqun Zhang, Jizhen Li, Yuhong Yang, Yong Luo, Weiping Tu	82043
UTILITY: Utilizing Explainable Reinforcement Learning to Improve Reinforcement Learning Shicheng Liu, Minghui Zhu	82070
RazorAttention: Efficient KV Cache Compression Through Retrieval Heads Hanlin Tang, Yang Lin, Jing Lin, Qingsen Han, Danning Ke, Shikuan Hong, Yiwu Yao, Gongyi Wang	82100
Graph Neural Preconditioners for Iterative Solutions of Sparse Linear Systems Jie Chen	82115
PolyNet: Learning Diverse Solution Strategies for Neural Combinatorial Optimization André Hottung, Mridul Mahajan, Kevin Tierney	82141
First-Person Fairness in Chatbots Tyna Eloundou, Alex Beutel, David G. Robinson, Keren Gu, Anna-Luisa Brakman, Pamela Mishkin, Meghan Shah, Johannes Heidecke, Lilian Weng, Adam Tauman Kalai	82160
Diffusion Feedback Helps CLIP See Better Wenxuan Wang, Quan Sun, Fan Zhang, Yepeng Tang, Jing Liu, Xinlong Wang	82195
Law of the Weakest Link: Cross Capabilities of Large Language Models Ming Zhong, Aston Zhang, Xuewei Wang, Rui Hou, Wenhan Xiong, Chenguang Zhu, Zhengxing Chen, Liang Tan, Chloe Bi, Mike Lewis, Sravya Popuri, Sharan Narang, Melanie Kambadur, Dhruv Mahajan, Sergey Edunov, Jiawei Han, Laurens Van Der Maaten	82214
Measuring and Improving Engagement of Text-To-Image Generation Models Varun Khurana, Yaman Kumar Singla, Jayakumar Subramanian, Changyou Chen, Rajiv Ratn Shah, Zhiqiang Xu, Balaji Krishnamurthy	82278
SplineGS: Learning Smooth Trajectories in Gaussian Splatting for Dynamic Scene Reconstruction Jihwan Yoon, Sangbeom Han, Jaeseok Oh, Minsik Lee	82310
GLOMA: Global Video Text Spotting with Morphological Association Han Wang, Yanjie Wang, Yang Li, Can Huang	82328
Learning the Complexity of Weakly Noisy Quantum States Yusen Wu, Bujiao Wu, Yanqi Song, Xiao Yuan, Jingbo Wang	82341
DOTS: Learning to Reason Dynamically in LLMs Via Optimal Reasoning Trajectories Search Murong Yue, Wenlin Yao, Haitao Mi, Dian Yu, Ziyu Yao, Dong Yu	82371

<p>Physics of Language Models: Part 2.1, Grade-School Math and the Hidden Reasoning Process</p> <p>Tian Ye, Zicheng Xu, Yuanzhi Li, Zeyuan Allen-Zhu</p>	82393
<p>Bridging Information Asymmetry in Text-Video Retrieval: A Data-Centric Approach</p> <p>Zechen Bai, Tianjun Xiao, Tong He, Pichao Wang, Zheng Zhang, Thomas Brox, Mike Zheng Shou</p>	82404
<p>Language Models Trained to Do Arithmetic Predict Human Risky and Intertemporal Choice</p> <p>Jian-Qiao Zhu, Haijiang Yan, Thomas L. Griffiths</p>	82427
<p>KGAREvion: An AI Agent for Knowledge-Intensive Biomedical QA</p> <p>Xiaorui Su, Yibo Wang, Shanghua Gao, Xiaolong Liu, Valentina Giunchiglia, Djork-Arné Clevert, Marinka Zitnik</p>	82446
<p>Asymptotic Analysis of Two-Layer Neural Networks After One Gradient Step Under Gaussian Mixtures Data with Structure</p> <p>Samet Demir, Zafer Dogan</p>	82474
<p>Noise Separation Guided Candidate Label Reconstruction for Noisy Partial Label Learning</p> <p>Xiaorui Peng, Yuheng Jia, Fuchao Yang, Ran Wang, Min-Ling Zhang</p>	82501
<p>Physics-Informed Temporal Difference Metric Learning for Robot Motion Planning</p> <p>Ruiqi Ni, Zherong Pan, Ahmed H Qureshi</p>	82522
<p>Rethinking Multiple-Instance Learning from Feature Space to Probability Space</p> <p>Zhaolong Du, Shasha Mao, Xuequan Lu, Mengnan Qi, Yimeng Zhang, Jing Gu, Licheng Jiao</p>	82540
<p>ClawMachine: Learning to Fetch Visual Tokens for Referential Comprehension</p> <p>Tianren Ma, Lingxi Xie, Yunjie Tian, Boyu Yang, Qixiang Ye</p>	82557
<p>Heavy-Tailed Diffusion Models</p> <p>Kushagra Pandey, Jaideep Pathak, Yilun Xu, Stephan Mandt, Michael Pritchard, Arash Vahdat, Morteza Mardani</p>	82578
<p>Hydra-SGG: Hybrid Relation Assignment for One-Stage Scene Graph Generation</p> <p>Minghan Chen, Guikun Chen, Wenguan Wang, Yi Yang</p>	82633
<p>Improving Large Language Model Planning with Action Sequence Similarity</p> <p>Xinran Zhao, Hanie Sedghi, Bernd Bohnet, Dale Schuurmans, Azade Nova</p>	82651
<p>ProAdvPrompter: A Two-Stage Journey to Effective Adversarial Prompting for LLMs</p> <p>Hao Di, Tong He, Haishan Ye, Yinghui Huang, Xiangyu Chang, Guang Dai, Ivor Tsang</p>	82676
<p>Real-Time Design of Architectural Structures with Differentiable Mechanics and Neural Networks</p> <p>Rafael Pastrana, Eder Medina, Isabel M. De Oliveira, Sigrid Adriaenssens, Ryan P Adams</p>	82701
<p>Learning Dynamics of LLM Finetuning</p> <p>Yi Ren, Danica J. Sutherland</p>	82725

<p>Physics-Informed Diffusion Models</p> <p><i>Jan-Hendrik Bastek, Waiching Sun, Dennis Kochmann</i></p>	82766
<p>Relation-Aware Diffusion for Heterogeneous Graphs with Partially Observed Features</p> <p><i>Daeho Um, Yoonji Lee, Jiwoong Park, Seulki Park, Yuneil Yeo, Seong Jin Ahn</i></p>	82792
<p>PeriodWave: Multi-Period Flow Matching for High-Fidelity Waveform Generation</p> <p><i>Sang-Hoon Lee, Ha-Yeong Choi, Seong-Whan Lee</i></p>	82819
<p>CLIPure: Purification in Latent Space Via CLIP for Adversarially Robust Zero-Shot Classification</p> <p><i>Mingkun Zhang, Keping Bi, Wei Chen, Jiafeng Guo, Xueqi Cheng</i></p>	82850
<p>Estimation of Single-Cell and Tissue Perturbation Effect in Spatial Transcriptomics Via Spatial Causal Disentanglement</p> <p><i>Stathis Megas, Daniel G. Chen, Krzysztof Polanski, Moshe Eliasof, Carola-Bibiane Schönlieb, Sarah A Teichmann</i></p>	82875
<p>Learning Video-Conditioned Policy on Unlabelled Data with Joint Embedding Predictive Transformer</p> <p><i>Hao Luo, Zongqing Lu</i></p>	82901
<p>Discrete Diffusion Schrödinger Bridge Matching for Graph Transformation</p> <p><i>Jun Hyeong Kim, Seonghwan Kim, Seokhyun Moon, Hyeongwoo Kim, Jeheon Woo, Woo Youn Kim</i></p>	82925
<p>Zero-Shot Forecasting of Chaotic Systems</p> <p><i>Yuanzhao Zhang, William Gilpin</i></p>	82972
<p>AvatarGO: Zero-Shot 4D Human-Object Interaction Generation and Animation</p> <p><i>Yukang Cao, Liang Pan, Kai Han, Kwan-Yee K. Wong, Ziwei Liu</i></p>	82999
<p>Interpretable Vision-Language Survival Analysis with Ordinal Inductive Bias for Computational Pathology</p> <p><i>Pei Liu, Luping Ji, Jiaxiang Gou, Bo Fu, Mao Ye</i></p>	83016
<p>Methods with Local Steps and Random Reshuffling for Generally Smooth Non-Convex Federated Optimization</p> <p><i>Yury Demidovich, Petr Ostroukhov, Grigory Malinovsky, Samuel Horváth, Martin Takác, Peter Richtárik, Eduard Gorbunov</i></p>	83045
<p>RainbowPO: A Unified Framework for Combining Improvements in Preference Optimization</p> <p><i>Hanyang Zhao, Genta Indra Winata, Anirban Das, Shi-Xiong Zhang, David Yao, Wenpin Tang, Sambit Sahu</i></p>	83113
<p>Robustness Inspired Graph Backdoor Defense</p> <p><i>Zhiwei Zhang, Minhua Lin, Junjie Xu, Zongyu Wu, Enyan Dai, Suhang Wang</i></p>	83136

<p>A Little Goes a Long Way: Efficient Long Context Training and Inference with Partial Contexts</p> <p>Suyu Ge, Xihui Lin, Yunan Zhang, Jiawei Han, Hao Peng</p>	83163
<p>MMWorld: Towards Multi-Discipline Multi-Faceted World Model Evaluation in Videos</p> <p>Xuehai He, Weixi Feng, Kaizhi Zheng, Yujie Lu, Wanrong Zhu, Jiachen Li, Yue Fan, Jianfeng Wang, Linjie Li, Zhengyuan Yang, Kevin Lin, William Yang Wang, Lijuan Wang, Xin Eric Wang</p>	83177
<p>Accelerating Neural ODEs: A Variational Formulation-Based Approach</p> <p>Hongjue Zhao, Yuchen Wang, Hairong Qi, Zijie Huang, Han Zhao, Lui Sha, Huajie Shao</p>	83211
<p>MuirBench: A Comprehensive Benchmark for Robust Multi-Image Understanding</p> <p>Fei Wang, Xingyu Fu, James Y. Huang, Zekun Li, Qin Liu, Xiaogeng Liu, Mingyu Derek Ma, Nan Xu, Wenxuan Zhou, Kai Zhang, Tianyi Lorena Yan, Wenjie Jacky Mo, Hsiang-Hui Liu, Pan Lu, Chunyuan Li, Chaowei Xiao, Kai-Wei Chang, Dan Roth, Sheng Zhang, Hoifung Poon, Muhao Chen</p>	83241
<p>Solving Video Inverse Problems Using Image Diffusion Models</p> <p>Taesung Kwon, Jong Chul Ye</p>	83268
<p>FreCaS: Efficient Higher-Resolution Image Generation Via Frequency-Aware Cascaded Sampling</p> <p>Zhengqiang Zhang, Ruihuang Li, Lei Zhang</p>	83293
<p>Self-Correcting Decoding with Generative Feedback for Mitigating Hallucinations in Large Vision-Language Models</p> <p>Ce Zhang, Zifu Wan, Zhehan Kan, Martin Q. Ma, Simon Stepputtis, Deva Ramanan, Russ Salakhutdinov, Louis-Philippe Morency, Katia P. Sycara, Yaqi Xie</p>	83306
<p>AuroraCap: Efficient, Performant Video Detailed Captioning and a New Benchmark</p> <p>Wenhao Chai, Enxin Song, Yilun Du, Chenlin Meng, Vashisht Madhavan, Omer Bar-Tal, Jenq-Neng Hwang, Saining Xie, Christopher D Manning</p>	83337
<p>T2V2: A Unified Non-Autoregressive Model for Speech Recognition and Synthesis Via Multitask Learning</p> <p>Nabarun Goswami, Hanqin Wang, Tatsuya Harada</p>	83382
<p>Booster: Tackling Harmful Fine-Tuning for Large Language Models Via Attenuating Harmful Perturbation</p> <p>Tiansheng Huang, Sihao Hu, Fatih Ilhan, Selim Furkan Tekin, Ling Liu</p>	83397
<p>Bounds on L_p Errors in Density Ratio Estimation Via f-Divergence Loss Functions</p> <p>Yoshiaki Kitazawa</p>	83422
<p>Regularization by Texts for Latent Diffusion Inverse Solvers</p> <p>Jeongsol Kim, Geon Yeong Park, Hyungjin Chung, Jong Chul Ye</p>	83470
<p>Hierarchical Autoregressive Transformers: Combining Byte- And Word-Level Processing for Robust, Adaptable Language Models</p> <p>Pit Neitemeier, Björn Deiseroth, Constantin Eichenberg, Lukas Balles</p>	83497

<p>Composable Interventions for Language Models</p> <p><i>Arinbjörn Kolbeinsson, Kyle O'Brien, Tianjin Huang, Shanghua Gao, Shiwei Liu, Jonathan Richard Schwarz, Anurag Jayant Vaidya, Faisal Mahmood, Marinka Zitnik, Tianlong Chen, Thomas Hartvigsen</i></p>	83521
<p>Gumbel Counterfactual Generation from Language Models</p> <p><i>Shauli Ravfogel, Anej Svete, Vésteinn Snæbjarnarson, Ryan Cotterell</i></p>	83555
<p>On Generalization Across Environments in Multi-Objective Reinforcement Learning</p> <p><i>Jayden Teoh, Pradeep Varakantham, Peter Vamplew</i></p>	83582
<p>Metalic: Meta-Learning In-Context with Protein Language Models</p> <p><i>Jacob Beck, Shikha Surana, Manus McAuliffe, Oliver Bent, Thomas D Barrett, Juan Jose Garau-Luis, Paul Duckworth</i></p>	83610
<p>MIND: Math Informed syNthetic Dialogues for Pretraining LLMs</p> <p><i>Syeda Nahida Akter, Shrimai Prabhumoye, John Kamalu, Sanjeev Satheesh, Eric Nyberg, Mostofa Patwary, Mohammad Shoeybi, Bryan Catanzaro</i></p>	83627
<p>Improving Convergence Guarantees of Random Subspace Second-Order Algorithm for Nonconvex Optimization</p> <p><i>Rei Higuchi, Pierre-Louis Poirion, Akiko Takeda</i></p>	83660
<p>Neural Sampling from Boltzmann Densities: Fisher-Rao Curves in the Wasserstein Geometry</p> <p><i>Jannis Chemseddine, Christian Wald, Richard Duong, Gabriele Steidl</i></p>	83703
<p>Approaching Rate-Distortion Limits in Neural Compression with Lattice Transform Coding</p> <p><i>Eric Lei, Hamed Hassani, Shirin Saeedi Bidokhti</i></p>	83728
<p>Grokking at the Edge of Numerical Stability</p> <p><i>Lucas Prieto, Melih Barsbey, Pedro A. M. Mediano, Tolga Birdal</i></p>	83746
<p>MELODI: Exploring Memory Compression for Long Contexts</p> <p><i>Yinpeng Chen, Delesley Hutchins, Aren Jansen, Andrey Zhmoginov, David Racz, Jesper Sparre Andersen</i></p>	83764
<p>Proteina: Scaling Flow-Based Protein Structure Generative Models</p> <p><i>Tomas Geffner, Kieran Didi, Zuobai Zhang, Danny Reidenbach, Zhonglin Cao, Jason Yim, Mario Geiger, Christian Dallago, Emine Kucukbenli, Arash Vahdat, Karsten Kreis</i></p>	83787
<p>Precedence-Constrained Winter Value for Effective Graph Data Valuation</p> <p><i>Hongliang Chi, Wei Jin, Charu C. Aggarwal, Yao Ma</i></p>	83836
<p>Enhancing the Scalability and Applicability of Kohn-Sham Hamiltonians for Molecular Systems</p> <p><i>Yunyang Li, Zaishuo Xia, Lin Huang, Xinran Wei, Samuel Harshe, Han Yang, Erpai Luo, Zun Wang, Jia Zhang, Chang Liu, Bin Shao, Mark Gerstein</i></p>	83868
<p>Cross-Embodiment Dexterous Grasping with Reinforcement Learning</p> <p><i>Haoqi Yuan, Bohan Zhou, Yuhui Fu, Zongqing Lu</i></p>	83913

<p>HiRA: Parameter-Efficient Hadamard High-Rank Adaptation for Large Language Models</p> <p>Qiushi Huang, Tom Ko, Zhan Zhuang, Lillian Tang, Yu Zhang</p>	83935
<p>Revisit Large-Scale Image-Caption Data in Pre-Training Multimodal Foundation Models</p> <p>Zhengfeng Lai, Vasileios Saveris, Chen Chen, Hong-You Chen, Haotian Zhang, Bowen Zhang, Wenze Hu, Juan Lao Tebar, Zhe Gan, Peter Grasch, Meng Cao, Yinfei Yang</p>	83955
<p>DECO: Unleashing the Potential of ConvNets for Query-Based Detection and Segmentation</p> <p>Xinghao Chen, Siwei Li, Yijing Yang, Yunhe Wang</p>	83975
<p>A Coefficient Makes SVRG Effective</p> <p>Yida Yin, Zhiqiu Xu, Zhiyuan Li, Trevor Darrell, Zhuang Liu</p>	83989
<p>Model-Based RL as a Minimalist Approach to Horizon-Free and Second-Order Bounds</p> <p>Zhiyong Wang, Dongruo Zhou, John C. S. Lui, Wen Sun</p>	84013
<p>Reliable and Diverse Evaluation of LLM Medical Knowledge Mastery</p> <p>Yuxuan Zhou, Xien Liu, Chen Ning, Xiao Zhang, Ji Wu</p>	84044
<p>PEARL: Towards Permutation-Resilient LLMs</p> <p>Liang Chen, Li Shen, Yang Deng, Xiaoyan Zhao, Bin Liang, Kam-Fai Wong</p>	84068
<p>COFlowNet: Conservative Constraints on Flows Enable High-Quality Candidate Generation</p> <p>Yudong Zhang, Xuan Yu, Xu Wang, Zhaoyang Sun, Chen Zhang, Pengkun Wang, Yang Wang</p>	84088
<p>Near-Exact Privacy Amplification for Matrix Mechanisms</p> <p>Christopher A. Choquette-Choo, Arun Ganesh, Saminul Haque, Thomas Steinke, Abhradeep Guha Thakurta</p>	84103
<p>RAPID: Retrieval Augmented Training of Differentially Private Diffusion Models</p> <p>Tanqiu Jiang, Changjiang Li, Fenglong Ma, Ting Wang</p>	84134
<p>Newton Meets Marchenko-Pastur: Massively Parallel Second-Order Optimization with Hessian Sketching and Debiasing</p> <p>Elad Romanov, Fangzhao Zhang, Mert Pilanci</p>	84152
<p>Block Diffusion: Interpolating Between Autoregressive and Diffusion Language Models</p> <p>Marianne Arriola, Aaron Gokaslan, Justin T Chiu, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Subham Sekhar Sahoo, Volodymyr Kuleshov</p>	84192
<p>HyperPLR: Hypergraph Generation Through Projection, Learning, and Reconstruction</p> <p>Weihuang Wen, Tianshu Yu</p>	84220
<p>Uncertainty Modeling in Graph Neural Networks Via Stochastic Differential Equations</p> <p>Richard Bergna, Sergio Calvo Ordoñez, Felix Opolka, Pietro Lio, José Miguel Hernández-Lobato</p>	84237
<p>Discovering Group Structures Via Unitary Representation Learning</p> <p>Dongsung Huh</p>	84254

<p>MIRAGE: Evaluating and Explaining Inductive Reasoning Process in Language Models</p> <p><i>Jiachun Li, Pengfei Cao, Zhuoran Jin, Yubo Chen, Kang Liu, Jun Zhao</i></p>	84279
<p>Fantastic Targets for Concept Erasure in Diffusion Models and Where to Find Them</p> <p><i>Anh Tuan Bui, Thuy-Trang Vu, Long Tung Vuong, Trung Le, Paul Montague, Tamas Abraham, Junae Kim, Dinh Phung</i></p>	84305
<p>Multi-Task Dense Predictions Via Unleashing the Power of Diffusion</p> <p><i>Yuqi Yang, Peng-Tao Jiang, Qibin Hou, Hao Zhang, Jinwei Chen, Bo Li</i></p>	84348
<p>On the Performance Analysis of Momentum Method: A Frequency Domain Perspective</p> <p><i>Xianliang Li, Jun Luo, Zhiwei Zheng, Hanxiao Wang, Li Luo, Lingkun Wen, Linlong Wu, Sheng Xu</i></p>	84366
<p>Backdooring Vision-Language Models with Out-Of-Distribution Data</p> <p><i>Weimin Lyu, Jiachen Yao, Saumya Gupta, Lu Pang, Tao Sun, Lingjie Yi, Lijie Hu, Haibin Ling, Chao Chen</i></p>	84388
<p>GPS: A Probabilistic Distributional Similarity with Gumbel Priors for Set-To-Set Matching</p> <p><i>Ziming Zhang, Fangzhou Lin, Haotian Liu, Jose Morales, Haichong Zhang, Kazunori Yamada, Vijaya B Kolachalama, Venkatesh Saligrama</i></p>	84407
<p>Scaling In-The-Wild Training for Diffusion-Based Illumination Harmonization and Editing by Imposing Consistent Light Transport</p> <p><i>Lvmin Zhang, Anyi Rao, Maneesh Agrawala</i></p>	84422
<p>Poisson-Dirac Neural Networks for Modeling Coupled Dynamical Systems Across Domains</p> <p><i>Razmik Arman Khosrovian, Takaharu Yaguchi, Hiroaki Yoshimura, Takashi Matsubara</i></p>	84440
<p>Perplexity Trap: PLM-Based Retrievers Overrate Low Perplexity Documents</p> <p><i>Haoyu Wang, Sunhao Dai, Haiyuan Zhao, Liang Pang, Xiao Zhang, Gang Wang, Zhenhua Dong, Jun Xu, Ji-Rong Wen</i></p>	84469
<p>Context-Aware Dynamic Pruning for Speech Foundation Models</p> <p><i>Masao Someki, Yifan Peng, Siddhant Arora, Markus Müller, Athanasios Mouchtaris, Grant Strimel, Jing Liu, Shinji Watanabe</i></p>	84489
<p>InverseBench: Benchmarking Plug-And-Play Diffusion Priors for Inverse Problems in Physical Sciences</p> <p><i>Hongkai Zheng, Wenda Chu, Bingliang Zhang, Zihui Wu, Austin Wang, Berthy Feng, Caifeng Zou, Yu Sun, Nikola Borislavov Kovachki, Zachary E Ross, Katherine Bouman, Yisong Yue</i></p>	84521
<p>SeedLM: Compressing LLM Weights into Seeds of Pseudo-Random Generators</p> <p><i>Rasoul Shafipour, David Harrison, Maxwell Horton, Jeffrey Marker, Houman Bedayat, Sachin Mehta, Mohammad Rastegari, Mahyar Najibi, Saman Naderiparizi</i></p>	84550
<p>Audio Large Language Models Can Be Descriptive Speech Quality Evaluators</p> <p><i>Chen Chen, Yuchen Hu, Siyin Wang, Helin Wang, Zhehuai Chen, Chao Zhang, Chao-Han Huck Yang, Engsiong Chng</i></p>	84567

<p>Mask in the Mirror: Implicit Sparsification</p> <p>Tom Jacobs, Rebekka Burkholz</p>	84582
<p>A Formal Framework for Understanding Length Generalization in Transformers</p> <p>Xinting Huang, Andy Yang, Satwik Bhattamishra, Yash Sarrof, Andreas Krebs, Hattie Zhou, Preetum Nakkiran, Michael Hahn</p>	84607
<p>Deriving Causal Order from Single-Variable Interventions: Guarantees & Algorithm</p> <p>Mathieu Chevalley, Patrick Schwab, Arash Mehrjou</p>	84692
<p>On Disentangled Training for Nonlinear Transform in Learned Image Compression</p> <p>Han Li, Shaohui Li, Wenrui Dai, Maida Cao, Nuowen Kan, Chenglin Li, Junni Zou, Hongkai Xiong</p>	84718
<p>Nonlinear Sequence Embedding by Monotone Variational Inequality</p> <p>Jonathan Yuyang Zhou, Yao Xie</p>	84739
<p>Exploring Prosocial Irrationality for LLM Agents: A Social Cognition View</p> <p>Xuan Liu, Jie Zhang, Haoyang Shang, Song Guo, Yang Chengxu, Quanyan Zhu</p>	84769
<p>Adversaries with Incentives: A Strategic Alternative to Adversarial Robustness</p> <p>Maayan Ehrenberg, Roy Ganz, Nir Rosenfeld</p>	84799
<p>Data Distillation for Extrapolative Protein Design Through Exact Preference Optimization</p> <p>Mostafa Karimi, Sharmi Banerjee, Tommi Jaakkola, Bella Dubrov, Shang Shang, Ron Benson</p>	84821
<p>Two Effects, One Trigger: On the Modality Gap, Object Bias, and Information Imbalance in Contrastive Vision-Language Models</p> <p>Simon Schrodli, David T. Hoffmann, Max Argus, Volker Fischer, Thomas Brox</p>	84854
<p>MambaPEFT: Exploring Parameter-Efficient Fine-Tuning for Mamba</p> <p>Masakazu Yoshimura, Teruaki Hayashi, Yota Maeda</p>	84883
<p>Unintentional Unalignment: Likelihood Displacement in Direct Preference Optimization</p> <p>Noam Razin, Sadhika Malladi, Adithya Bhaskar, Danqi Chen, Sanjeev Arora, Boris Hanin</p>	84908
<p>Trajectory-LLM: A Language-Based Data Generator for Trajectory Prediction in Autonomous Driving</p> <p>Kairui Yang, Zihao Guo, Gengjie Lin, Haotian Dong, Zhao Huang, Yipeng Wu, Die Zuo, Jibin Peng, Ziyuan Zhong, Xin Wang, Qing Guo, Xiaosong Jia, Junchi Yan, Di Lin</p>	84952
<p>Adaptive Rank Allocation: Speeding Up Modern Transformers with RaNA Adapters</p> <p>Roberto Garcia, Jerry Weihong Liu, Daniel Sorvisto, Sabri Eyuboglu</p>	84965
<p>Equivariant Neural Functional Networks for Transformers</p> <p>Hoang V. Tran, Thieu Vo, An Nguyen The, Tho Tran Huu, Minh-Khoi Nguyen-Nhat, Thanh Tran, Duy-Tung Pham, Tan Minh Nguyen</p>	84981
<p>Jailbreaking as a Reward Misspecification Problem</p> <p>Zhihui Xie, Jiahui Gao, Lei Li, Zhenguo Li, Qi Liu, Lingpeng Kong</p>	85053

<p>A Transfer Attack to Image Watermarks</p> <p>Yuepeng Hu, Zhengyuan Jiang, Moyang Guo, Neil Zhenqiang Gong</p>	85076
<p>Compositional Simulation-Based Inference for Time Series</p> <p>Manuel Gloeckler, Shoji Toyota, Kenji Fukumizu, Jakob H. Macke</p>	85103
<p>Towards Neural Scaling Laws for Time Series Foundation Models</p> <p>Qingren Yao, Chao-Han Huck Yang, Renhe Jiang, Yuxuan Liang, Ming Jin, Shirui Pan</p>	85137
<p>The Journey Matters: Average Parameter Count Over Pre-Training Unifies Sparse and Dense Scaling Laws</p> <p>Tian Jin, Ahmed Imtiaz Humayun, Utku Evci, Suvinay Subramanian, Amir Yazdanbakhsh, Dan Alistarh, Gintare Karolina Dziugaite</p>	85165
<p>Learning Partial Graph Matching Via Optimal Partial Transport</p> <p>Gathika Ratnayaka, James Nichols, Qing Wang</p>	85182
<p>Multiple Heads Are Better than One: Mixture of Modality Knowledge Experts for Entity Representation Learning</p> <p>Yichi Zhang, Zhuo Chen, Lingbing Guo, Yajing Xu, Binbin Hu, Ziqi Liu, Wen Zhang, Huajun Chen</p>	85203
<p>Benchmarking LLMs' Judgments with No Gold Standard</p> <p>Shengwei Xu, Yuxuan Lu, Grant Schoenebeck, Yuqing Kong</p>	85221
<p>Efficient Cross-Episode Meta-RL</p> <p>Gresa Shala, André Biedenkapp, Pierre Krack, Florian Walter, Josif Grabocka</p>	85254
<p>FaithEval: Can Your Language Model Stay Faithful to Context, Even if "The Moon is Made of Marshmallows"</p> <p>Yifei Ming, Senthil Purushwalkam, Shrey Pandit, Zixuan Ke, Xuan-Phi Nguyen, Caiming Xiong, Shafiq Joty</p>	85296
<p>InstantSwap: Fast Customized Concept Swapping Across Sharp Shape Differences</p> <p>Chenyang Zhu, Kai Li, Yue Ma, Longxiang Tang, Chengyu Fang, Chubin Chen, Qifeng Chen, Xiu Li</p>	85323
<p>Dimension Agnostic Neural Processes</p> <p>Hyungi Lee, Chaeyun Jang, Dong Bok Lee, Juho Lee</p>	85358
<p>Uncertainty Herding: One Active Learning Method for All Label Budgets</p> <p>Wonho Bae, Danica J. Sutherland, Gabriel L. Oliveira</p>	85392
<p>Information Theoretic Text-To-Image Alignment</p> <p>Chao Wang, Giulio Franzese, Alessandro Finamore, Massimo Gallo, Pietro Michiardi</p>	85410
<p>On the Optimal Memorization Capacity of Transformers</p> <p>Tokio Kajitsuka, Issei Sato</p>	85447
<p>Discriminating Image Representations with Principal Distortions</p> <p>Jenelle Feather, David Lipshutz, Sarah E Harvey, Alex H Williams, Eero P Simoncelli</p>	85489

<p>ScImage: How Good Are Multimodal Large Language Models at Scientific Text-To-Image Generation?</p> <p><i>Leixin Zhang, Steffen Eger, Yinjie Cheng, Weihe Zhai, Jonas Belouadi, Fahimeh Moafian, Zhixue Zhao</i></p>	85516
<p>On Rollouts in Model-Based Reinforcement Learning</p> <p><i>Bernd Frauenknecht, Devdutt Subhasish, Friedrich Solowjow, Sebastian Trimpe</i></p>	85542
<p>Functional Homotopy: Smoothing Discrete Optimization Via Continuous Parameters for LLM Jailbreak Attacks</p> <p><i>Zi Wang, Divyam Anshumaan, Ashish Hooda, Yudong Chen, Somesh Jha</i></p>	85568
<p>Matryoshka Multimodal Models</p> <p><i>Mu Cai, Jianwei Yang, Jianfeng Gao, Yong Jae Lee</i></p>	85588
<p>Capturing the Temporal Dependence of Training Data Influence</p> <p><i>Jiachen T. Wang, Dawn Song, James Zou, Prateek Mittal, Ruoxi Jia</i></p>	85607
<p>Trust Or Escalate: LLM Judges with Provable Guarantees for Human Agreement</p> <p><i>Jaehun Jung, Faeze Brahman, Yejin Choi</i></p>	85645
<p>SAM-CP: Marrying SAM with Composable Prompts for Versatile Segmentation</p> <p><i>Pengfei Chen, Lingxi Xie, Xinyue Huo, Xuehui Yu, Xiaopeng Zhang, Yingfei Sun, Zhenjun Han, Qi Tian</i></p>	85670
<p>The 3D-PC: A Benchmark for Visual Perspective Taking in Humans and Machines</p> <p><i>Drew Linsley, Peisen Zhou, Alekh Karkada Ashok, Akash Nagaraj, Gaurav Gaonkar, Francis E Lewis, Zygmunt Pizlo, Thomas Serre</i></p>	85693
<p>SeRA: Self-Reviewing and Alignment of LLMs Using Implicit Reward Margins</p> <p><i>Jongwoo Ko, Saket Dingliwal, Bhavana Ganesh, Sailik Sengupta, Sravan Babu Bodapati, Aram Galstyan</i></p>	85716
<p>Can Watermarked LLMs Be Identified by Users Via Crafted Prompts?</p> <p><i>Aiwei Liu, Sheng Guan, Yiming Liu, Leyi Pan, Yifei Zhang, Liancheng Fang, Lijie Wen, Philip S. Yu, Xuming Hu</i></p>	85743
<p>UniCon: Unidirectional Information Flow for Effective Control of Large-Scale Diffusion Models</p> <p><i>Fanghua Yu, Jinjin Gu, Jinfan Hu, Zheyuan Li, Chao Dong</i></p>	85771
<p>Learning Distributions of Complex Fluid Simulations with Diffusion Graph Networks</p> <p><i>Mario Lino Valencia, Tobias Pfaff, Nils Thuerey</i></p>	85797
<p>Efficiently Parameterized Neural Metriplectic Systems</p> <p><i>Anthony Gruber, Kookjin Lee, Haksoo Lim, Noseong Park, Nathaniel Trask</i></p>	85828
<p>SlowFast-VGen: Slow-Fast Learning for Action-Driven Long Video Generation</p> <p><i>Yining Hong, Beide Liu, Maxine Wu, Yuanhao Zhai, Kai-Wei Chang, Linjie Li, Kevin Lin, Chung-Ching Lin, Jianfeng Wang, Zhengyuan Yang, Ying Nian Wu, Lijuan Wang</i></p>	85853

<p>OmniKV: Dynamic Context Selection for Efficient Long-Context LLMs</p> <p><i>Jitai Hao, Yuke Zhu, Tian Wang, Jun Yu, Xin Xin, Bo Zheng, Zhaochun Ren, Sheng Guo</i></p>	85881
<p>L3Ms — Lagrange Large Language Models</p> <p><i>Guneet S. Dhillon, Xingjian Shi, Yee Whye Teh, Alex Smola</i></p>	85903
<p>Fair Submodular Cover</p> <p><i>Wenjing Chen, Shuo Xing, Samson Zhou, Victoria G. Crawford</i></p>	85918
<p>Adversarial Training for Defense Against Label Poisoning Attacks</p> <p><i>Melis Ilayda Bal, Volkan Cevher, Michael Muehlebach</i></p>	85946
<p>Combining Induction and Transduction for Abstract Reasoning</p> <p><i>Wen-Ding Li, Keya Hu, Carter Larsen, Yuqing Wu, Simon Alford, Caleb Woo, Spencer M. Dunn, Hao Tang, Wei-Long Zheng, Yewen Pu, Kevin Ellis</i></p>	85981
<p>NovelQA: Benchmarking Question Answering on Documents Exceeding 200K Tokens</p> <p><i>Cunxiang Wang, Ruoxi Ning, Boqi Pan, Tonghui Wu, Qipeng Guo, Cheng Deng, Guangsheng Bao, Xiangkun Hu, Zheng Zhang, Qian Wang, Yue Zhang</i></p>	86025
<p>Exploring the Design Space of Visual Context Representation in Video MLLMs</p> <p><i>Yifan Du, Yuqi Huo, Kun Zhou, Zijia Zhao, Haoyu Lu, Han Huang, Xin Zhao, Bingning Wang, Weipeng Chen, Ji-Rong Wen</i></p>	86048
<p>BOFormer: Learning to Solve Multi-Objective Bayesian Optimization Via Non-Markovian RL</p> <p><i>Yu Heng Hung, Kai-Jie Lin, Yu-Heng Lin, Chien-Yi Wang, Cheng Sun, Ping-Chun Hsieh</i></p>	86067
<p>Score-Based Self-Supervised MRI Denoising</p> <p><i>Jiachen Tu, Yaokun Shi, Fan Lam</i></p>	86095
<p>VideoWebArena: Evaluating Long Context Multimodal Agents with Video Understanding Web Tasks</p> <p><i>Lawrence Keunho Jang, Yinheng Li, Dan Zhao, Charles Ding, Justin Lin, Paul Pu Liang, Rogerio Bonatti, Kazuhito Koishida</i></p>	86127
<p>Lift Your Molecules: Molecular Graph Generation in Latent Euclidean Space</p> <p><i>Mohamed Amine Ketata, Nicholas Gao, Johanna Sommer, Tom Wollschläger, Stephan Günemann</i></p>	86152
<p>LeanAgent: Lifelong Learning for Formal Theorem Proving</p> <p><i>Adarsh Kumarappan, Mo Tiwari, Peiyang Song, Robert Joseph George, Chaowei Xiao, Anima Anandkumar</i></p>	86184
<p>CryoGEN: Generative Energy-Based Models for Cryogenic Electron Tomography Reconstruction</p> <p><i>Yunfei Teng, Yuxuan Ren, Kai Chen, Xi Chen, Zhaoming Chen, Qiwei Ye</i></p>	86224
<p>Weak-To-Strong Generalization Through the Data-Centric Lens</p> <p><i>Changho Shin, John Cooper, Frederic Sala</i></p>	86245

<p>Learning on One Mode: Addressing Multi-Modality in Offline Reinforcement Learning</p> <p>Mianchu Wang, Yue Jin, Giovanni Montana</p>	86284
<p>Strong Preferences Affect the Robustness of Preference Models and Value Alignment</p> <p>Ziwei Xu, Mohan Kankanhalli</p>	86302
<p>The Impact of Allocation Strategies in Subset Learning on the Expressive Power of Neural Networks</p> <p>Ofir Schliesselberg, Ran Darshan</p>	86323
<p>VICtoR: Learning Hierarchical Vision-Instruction Correlation Rewards for Long-Horizon Manipulation</p> <p>Kuo-Han Hung, Pang-Chi Lo, Jia-Fong Yeh, Han-Yuan Hsu, Yi-Ting Chen, Winston H. Hsu</p>	86351
<p>Self-Supervised Diffusion Models for Electron-Aware Molecular Representation Learning</p> <p>Gyoung S. Na, Chanyoung Park</p>	86379
<p>EG4D: Explicit Generation of 4D Object Without Score Distillation</p> <p>Qi Sun, Zhiyang Guo, Ziyu Wan, Jing Nathan Yan, Shengming Yin, Wengang Zhou, Jing Liao, Houqiang Li</p>	86402
<p>Trajectory-Class-Aware Multi-Agent Reinforcement Learning</p> <p>Hyungho Na, Kwanghyeon Lee, Sumin Lee, Il-Chul Moon</p>	86428
<p>LLaVA-Mini: Efficient Image and Video Large Multimodal Models with One Vision Token</p> <p>Shaolei Zhang, Qingkai Fang, Zhe Yang, Yang Feng</p>	86455
<p>SG-I2V: Self-Guided Trajectory Control in Image-To-Video Generation</p> <p>Koichi Namekata, Sherwin Bahmani, Ziyi Wu, Yash Kant, Igor Gilitschenski, David B. Lindell</p>	86481
<p>ReNovo: Retrieval-Based <i>De Novo</i> Mass Spectrometry Peptide Sequencing</p> <p>Shaorong Chen, Jun Xia, Jingbo Zhou, Lecheng Zhang, Zhangyang Gao, Bozhen Hu, Cheng Tan, Wenjie Du, Stan Z. Li</p>	86504
<p>Accurate and Scalable Graph Neural Networks Via Message Invariance</p> <p>Zhihao Shi, Jie Wang, Zhiwei Zhuang, Xize Liang, Bin Li, Feng Wu</p>	86520
<p>How Much is Unseen Depends Chiefly on Information About the Seen</p> <p>Seongmin Lee, Marcel Boehme</p>	86552
<p>Linear Mode Connectivity in Differentiable Tree Ensembles</p> <p>Ryuichi Kanoh, Mahito Sugiyama</p>	86576
<p>Latent-EnSF: A Latent Ensemble Score Filter for High-Dimensional Data Assimilation with Sparse Observation Data</p> <p>Phillip Si, Peng Chen</p>	86601
<p>On the Expressiveness of Rational ReLU Neural Networks with Bounded Depth</p> <p>Gennadiy Averkov, Christopher Hojny, Maximilian Merkert</p>	86618

<p>The Breakdown of Gaussian Universality in Classification of High-Dimensional Linear Factor Mixtures</p> <p>Xiaoyi Mai, Zhenyu Liao</p>	86633
<p>Polyrating: A Cost-Effective and Bias-Aware Rating System for LLM Evaluation</p> <p>Jasper Dekoninck, Maximilian Baader, Martin Vechev</p>	86665
<p>Circuit Representation Learning with Masked Gate Modeling and Verilog-AIG Alignment</p> <p>Haoyuan Wu, Haisheng Zheng, Yuan Pu, Bei Yu</p>	86691
<p>Deep Kernel Posterior Learning Under Infinite Variance Prior Weights</p> <p>Jorge Loria, Anindya Bhadra</p>	86708
<p>Certifying Language Model Robustness with Fuzzed Randomized Smoothing: An Efficient Defense Against Backdoor Attacks</p> <p>Bowei He, Lihao Yin, Huiling Zhen, Jianping Zhang, Lanqing Hong, Mingxuan Yuan, Chen Ma</p>	86731
<p>MRAG-Bench: Vision-Centric Evaluation for Retrieval-Augmented Multimodal Models</p> <p>Wenbo Hu, Jia-Chen Gu, Zi-Yi Dou, Mohsen Fayyaz, Pan Lu, Kai-Wei Chang, Nanyun Peng</p>	86765
<p>Mitigating Spurious Correlations in Zero-Shot Multimodal Models</p> <p>Shenyu Lu, Junyi Chai, Xiaoqian Wang</p>	86789
<p>Privacy-Aware Lifelong Learning</p> <p>Ozan Ozdenizci, Elmar Rueckert, Robert Legenstein</p>	86816
<p>Learning Graph Invariance by Harnessing Spuriousity</p> <p>Tianjun Yao, Yongqiang Chen, Kai Hu, Tongliang Liu, Kun Zhang, Zhiqiang Shen</p>	86837
<p>Weighted Point Set Embedding for Multimodal Contrastive Learning Toward Optimal Similarity Metric</p> <p>Toshimitsu Uesaka, Taiji Suzuki, Yuhta Takida, Chieh-Hsin Lai, Naoki Murata, Yuki Mitsufuji</p>	86874
<p>TODO: Enhancing LLM Alignment with Ternary Preferences</p> <p>Yuxiang Guo, Lu Yin, Bo Jiang, Jiaqi Zhang</p>	86896
<p>IGL-Bench: Establishing the Comprehensive Benchmark for Imbalanced Graph Learning</p> <p>Jiawen Qin, Haonan Yuan, Qingyun Sun, Lyujin Xu, Jiaqi Yuan, Pengfeng Huang, Zhaonan Wang, Xingcheng Fu, Hao Peng, Jianxin Li, Philip S. Yu</p>	86918
<p>VAE-Var: Variational Autoencoder-Enhanced Variational Methods for Data Assimilation in Meteorology</p> <p>Yi Xiao, Qilong Jia, Kun Chen, Lei Bai, Wei Xue</p>	86976
<p>IDArb: Intrinsic Decomposition for Arbitrary Number of Input Views and Illuminations</p> <p>Zhibing Li, Tong Wu, Jing Tan, Mengchen Zhang, Jiaqi Wang, Dahua Lin</p>	87005

<p>McEval: Massively Multilingual Code Evaluation</p> <p>Linzheng Chai, Shukai Liu, Jian Yang, Yuwei Yin, Jinke, Jiaheng Liu, Tao Sun, Ge Zhang, Changyu Ren, Hongcheng Guo, Noah Wang, Boyang Wang, Xianjie Wu, Bing Wang, Tongliang Li, Liqun Yang, Sufeng Duan, Zhaoxiang Zhang, Zhoujun Li</p>	87030
<p>Towards Certification of Uncertainty Calibration Under Adversarial Attacks</p> <p>Cornelius Emde, Francesco Pinto, Thomas Lukasiewicz, Philip Torr, Adel Bibi</p>	87057
<p>Counterfactual Realizability</p> <p>Arvind Raghavan, Elias Bareinboim</p>	87098
<p>Learning Splitting Heuristics in Divide-And-Conquer SAT Solvers with Reinforcement Learning</p> <p>Shumao Zhai, Ning Ge</p>	87115
<p>Learning Randomized Algorithms with Transformers</p> <p>Johannes Von Oswald, Seijin Kobayashi, Yassir Akram, Angelika Steger</p>	87137
<p>Biologically Constrained Barrel Cortex Model Integrates Whisker Inputs and Replicates Key Brain Network Dynamics</p> <p>Tianfang Zhu, Dongli Hu, Jiandong Zhou, Kai Du, Anan Li</p>	87158
<p>SynFlowNet: Design of Diverse and Novel Molecules with Synthesis Constraints</p> <p>Miruna Cretu, Charles Harris, Ilia Igashov, Arne Schneuing, Marwin Segler, Bruno Correia, Julien Roy, Emmanuel Bengio, Pietro Lio</p>	87172
<p>Gnothi Seauton: Empowering Faithful Self-Interpretability in Black-Box Transformers</p> <p>Shaobo Wang, Hongxuan Tang, Mingyang Wang, Hongrui Zhang, Xuyang Liu, Weiya Li, Xuming Hu, Linfeng Zhang</p>	87198
<p>AIR-BENCH 2024: A Safety Benchmark Based on Regulation and Policies Specified Risk Categories</p> <p>Yi Zeng, Yu Yang, Andy Zhou, Jeffrey Ziwei Tan, Yuheng Tu, Yifan Mai, Kevin Klyman, Minzhou Pan, Ruoxi Jia, Dawn Song, Percy Liang, Bo Li</p>	87227
<p>Protein Language Model Fitness is a Matter of Preference</p> <p>Cade W Gordon, Amy X. Lu, Pieter Abbeel</p>	87262
<p>How DNNs Break the Curse of Dimensionality: Compositionality and Symmetry Learning</p> <p>Arthur Jacot, Seok Hoan Choi, Yuxiao Wen</p>	87279
<p>Unlocking State-Tracking in Linear RNNs Through Negative Eigenvalues</p> <p>Riccardo Grazzi, Julien Siems, Arber Zela, Jörg K. H. Franke, Frank Hutter, Massimiliano Pontil</p>	87314
<p>Triples as the Key: Structuring Makes Decomposition and Verification Easier in LLM-Based TableQA</p> <p>Zhen Yang, Ziwei Du, Minghan Zhang, Wei Du, Jie Chen, Zhen Duan, Shu Zhao</p>	87347
<p>Efficient Causal Decision Making with One-Sided Feedback</p> <p>Jianing Chu, Shu Yang, Wenbin Lu, Pulak Ghosh</p>	87372

<p>LLaVA-MoD: Making LLaVA Tiny Via MoE-Knowledge Distillation</p> <p>Fangxun Shu, Yue Liao, Lei Zhang, Le Zhuo, Chenning Xu, Guanghao Zhang, Haonan Shi, Long Chan, Taozhong, Zhelun Yu, Wanggui He, Siming Fu, Haoyuan Li, Si Liu, Hongsheng Li, Hao Jiang</p>	87394
<p>FlowDec: A Flow-Based Full-Band General Audio Codec with High Perceptual Quality</p> <p>Simon Welker, Matthew Le, Ricky T. Q. Chen, Wei-Ning Hsu, Timo Gerkmann, Alexander Richard, Yi-Chiao Wu</p>	87413
<p>Towards a Theoretical Understanding of Synthetic Data in LLM Post-Training: A Reverse-Bottleneck Perspective</p> <p>Zeyu Gan, Yong Liu</p>	87441
<p>Flow Matching with Gaussian Process Priors for Probabilistic Time Series Forecasting</p> <p>Marcel Kollovich, Marten Lienen, David Lüdke, Leo Schwinn, Stephan Günemann</p>	87465
<p>Think Then React: Towards Unconstrained Action-To-Reaction Motion Generation</p> <p>Wenhui Tan, Boyuan Li, Chuhao Jin, Wenbing Huang, Xiting Wang, Ruihua Song</p>	87488
<p>Unbounded: A Generative Infinite Game of Character Life Simulation</p> <p>Jialu Li, Yuanzhen Li, Neal Wadhwa, Yael Pritch, David E. Jacobs, Michael Rubinstein, Mohit Bansal, Nataniel Ruiz</p>	87502
<p>Reconciling Model Multiplicity for Downstream Decision Making</p> <p>Ally Yalei Du, Dung Daniel Ngo, Steven Wu</p>	87529
<p>Generalizable Motion Planning Via Operator Learning</p> <p>Sharath Matada, Luke Bhan, Yuanyuan Shi, Nikolay Atanasov</p>	87561
<p>Decision Tree Induction Through LLMs Via Semantically-Aware Evolution</p> <p>Tennison Liu, Nicolas Huynh, Mihaela Van Der Schaar</p>	87586
<p>Fictitious Synthetic Data Can Improve LLM Factuality Via Prerequisite Learning</p> <p>Yujian Liu, Shiyu Chang, Tommi Jaakkola, Yang Zhang</p>	87620
<p>Beyond Autoregression: Fast LLMs Via Self-Distillation Through Time</p> <p>Justin Deschenaux, Caglar Gulcehre</p>	87644
<p>Learning Hierarchical Polynomials of Multiple Nonlinear Features</p> <p>Hengyu Fu, Zihao Wang, Eshaan Nichani, Jason D. Lee</p>	87683
<p>Zero-Shot Model-Based Reinforcement Learning Using Large Language Models</p> <p>Abdelhakim Benechehab, Youssef Attia El Hili, Ambroise Odonnat, Oussama Zekri, Albert Thomas, Giuseppe Paolo, Maurizio Filippone, Ievgen Redko, Balázs Kégl</p>	87738
<p>Seeing Eye to AI: Human Alignment Via Gaze-Based Response Rewards for Large Language Models</p> <p>Ángela López-Cardona, Carlos Segura, Alexandros Karatzoglou, Sergi Abadal, Ioannis Arapakis</p>	87768

<p>VideoShield: Regulating Diffusion-Based Video Generation Models Via Watermarking</p> <p>Runyi Hu, Jie Zhang, Yiming Li, Jiwei Li, Qing Guo, Han Qiu, Tianwei Zhang</p>	87793
<p>Selective Attention Improves Transformer</p> <p>Yaniv Leviathan, Matan Kalman, Yossi Matias</p>	87824
<p>Schur's Positive-Definite Network: Deep Learning in the SPD Cone with Structure</p> <p>Can Pouliquen, Mathurin Massias, Titouan Vayer</p>	87850
<p>AniSDF: Fused-Granularity Neural Surfaces with Anisotropic Encoding for High-Fidelity 3D Reconstruction</p> <p>Jingnan Gao, Zhuo Chen, Xiaokang Yang, Yichao Yan</p>	87866
<p>Decoupled Subgraph Federated Learning</p> <p>Javad Aliakbari, Johan Östman, Alexandre Graell I Amat</p>	87888
<p>PostCast: Generalizable Postprocessing for Precipitation Nowcasting Via Unsupervised Blurriness Modeling</p> <p>Junchao Gong, Siwei Tu, Weidong Yang, Ben Fei, Kun Chen, Zhangwenlong, Xiaokang Yang, Wanli Ouyang, Lei Bai</p>	87920
<p>RA-TTA: Retrieval-Augmented Test-Time Adaptation for Vision-Language Models</p> <p>Youngjun Lee, Doyoung Kim, Junhyeok Kang, Jihwan Bang, Hwanjun Song, Jae-Gil Lee</p>	87942
<p>Attention as a Hypernetwork</p> <p>Simon Schug, Seijin Kobayashi, Yassir Akram, Joao Sacramento, Razvan Pascanu</p>	87967
<p>Q-SFT: Q-Learning for Language Models Via Supervised Fine-Tuning</p> <p>Joey Hong, Anca Dragan, Sergey Levine</p>	87994
<p>Agent Security Bench (ASB): Formalizing and Benchmarking Attacks and Defenses in LLM-Based Agents</p> <p>Hanrong Zhang, Jingyuan Huang, Kai Mei, Yifei Yao, Zhenting Wang, Chenlu Zhan, Hongwei Wang, Yongfeng Zhang</p>	88011
<p>Learning to Search from Demonstration Sequences</p> <p>Dixant Mittal, Liwei Kang, Wee Sun Lee</p>	88047
<p>Robustness Auditing for Linear Regression: To Singularity and Beyond</p> <p>Ittai Rubinstein, Samuel B. Hopkins</p>	88073
<p>Does SGD Really Happen in Tiny Subspaces?</p> <p>Minhak Song, Kwangjun Ahn, Chulhee Yun</p>	88137
<p>Advancing Prompt-Based Methods for Replay-Independent General Continual Learning</p> <p>Zhiqi Kang, Liyuan Wang, Xingxing Zhang, Karteek Alahari</p>	88172

<p>How Low Can You Go? Searching for the Intrinsic Dimensionality of Complex Networks Using Metric Node Embeddings</p> <p><i>Nikolaos Nakis, Niels Raunkjær Holm, Andreas Lyhne Fiehn, Morten Mørup</i></p>	88191
<p>Multimodal Quantitative Language for Generative Recommendation</p> <p><i>Jiayang Zhai, Zi-Feng Mai, Chang-Dong Wang, Feidiao Yang, Xiawu Zheng, Hui Li, Yonghong Tian</i></p>	88215
<p>MMKE-Bench: A Multimodal Editing Benchmark for Diverse Visual Knowledge</p> <p><i>Yuntao Du., Kailin Jiang, Zhi Gao, Chenrui Shi, Zilong Zheng, Siyuan Qi, Qing Li</i></p>	88233
<p>Visual-O1: Understanding Ambiguous Instructions Via Multi-Modal Multi-Turn Chain-Of-Thoughts Reasoning</p> <p><i>Minheng Ni, Yutao Fan, Lei Zhang, Wangmeng Zuo</i></p>	88263
<p>UniMatch: Universal Matching from Atom to Task for Few-Shot Drug Discovery</p> <p><i>Ruifeng Li, Mingqian Li, Wei Liu, Yuhua Zhou, Xiangxin Zhou, Yuan Yao, Qiang Zhang, Hongyang Chen</i></p>	88290
<p>CyberHost: A One-Stage Diffusion Framework for Audio-Driven Talking Body Generation</p> <p><i>Gaojie Lin, Jianwen Jiang, Chao Liang, Tianyun Zhong, Jiaqi Yang, Zerong Zheng, Yanbo Zheng</i></p>	88317
<p>Learning to Explore and Exploit with GNNs for Unsupervised Combinatorial Optimization</p> <p><i>Utku Umur Acikalin, Aaron M Ferber, Carla P Gomes</i></p>	88340
<p>Co³Gesture: Towards Coherent Concurrent Co-Speech 3D Gesture Generation with Interactive Diffusion</p> <p><i>Xingqun Qi, Yatian Wang, Hengyuan Zhang, Jiahao Pan, Wei Xue, Shanghang Zhang, Wenhan Luo, Qifeng Liu, Yike Guo</i></p>	88359
<p>PINP: Physics-Informed Neural Predictor with Latent Estimation of Fluid Flows</p> <p><i>Huaguan Chen, Yang Liu, Hao Sun</i></p>	88378
<p>VoxDialogue: Can Spoken Dialogue Systems Understand Information Beyond Words?</p> <p><i>Xize Cheng, Ruofan Hu, Xiaoda Yang, Jingyu Lu, Dongjie Fu, Zehan Wang, Shengpeng Ji, Rongjie Huang, Boyang Zhang, Tao Jin, Zhou Zhao</i></p>	88409
<p>OccProphet: Pushing the Efficiency Frontier of Camera-Only 4D Occupancy Forecasting with an Observer-Forecaster-Refiner Framework</p> <p><i>Junliang Chen, Huaiyuan Xu, Yi Wang, Lap-Pui Chau</i></p>	88425
<p>Joint Reward and Policy Learning with Demonstrations and Human Feedback Improves Alignment</p> <p><i>Chenliang Li, Siliang Zeng, Zeyi Liao, Jiayang Li, Dongyeop Kang, Alfredo Garcia, Mingyi Hong</i></p>	88441
<p>Approximating Full Conformal Prediction for Neural Network Regression with Gauss-Newton Influence</p> <p><i>Dharmesh Tailor, Alvaro Correia, Eric Nalisnick, Christos Louizos</i></p>	88478

Regret Bounds for Episodic Risk-Sensitive Linear Quadratic Regulator Wenhao Xu, Xuefeng Gao, Xuedong He	88508
Beyond Canonicalization: How Tensorial Messages Improve Equivariant Message Passing Peter Lippmann, Gerrit Gerhartz, Roman Remme, Fred A. Hamprecht	88556
Divergence of Neural Tangent Kernel in Classification Problems Zixiong Yu, Songtao Tian, Guhan Chen	88577
Neural Dueling Bandits: Preference-Based Optimization with Human Feedback Arun Verma, Zhongxiang Dai, Xiaoqiang Lin, Patrick Jaillet, Bryan Kian Hsiang Low	88596
Learning to Solve Differential Equation Constrained Optimization Problems Vincenzo Di Vito Francesco, Mostafa Mohammadian, Kyri Baker, Ferdinando Fioretto	88634
Language Model Alignment in Multilingual Trolley Problems Zhijing Jin, Max Kleiman-Weiner, Giorgio Piatti, Sydney Levine, Jiarui Liu, Fernando Gonzalez Adauto, Francesco Ortu, András Strausz, Mrinmaya Sachan, Rada Mihalcea, Yejin Choi, Bernhard Schölkopf	88657
DEPT: Decoupled Embeddings for Pre-Training Language Models Alex Jacob, Lorenzo Sani, Meghdad Kurmanji, William F. Shen, Xinchu Qiu, Dongqi Cai, Yan Gao, Nicholas Donald Lane	88687
CREAM: Consistency Regularized Self-Rewarding Language Models Zhaoyang Wang, Weilei He, Zhiyuan Liang, Xuchao Zhang, Chetan Bansal, Ying Wei, Weitong Zhang, Huaxiu Yao	88717
Neuron Platonic Intrinsic Representation from Dynamics Using Contrastive Learning Wei Wu, Can Liao, Zizhen Deng, Zhengrui Guo, Jinzhuo Wang	88740
Fair Clustering in the Sliding Window Model Vincent Cohen-Addad, Shaofeng H.-C. Jiang, Qiaoyuan Yang, Yubo Zhang, Samson Zhou	88757
MindSimulator: Exploring Brain Concept Localization Via Synthetic fMRI Guangyin Bao, Qi Zhang, Zixuan Gong, Zhuojia Wu, Duoqian Miao	88785
Accelerating 3D Molecule Generation Via Jointly Geometric Optimal Transport Haokai Hong, Wanyu Lin, Kc Tan	88808
Towards Empowerment Gain Through Causal Structure Learning in Model-Based Reinforcement Learning Hongye Cao, Fan Feng, Meng Fang, Shaokang Dong, Tianpei Yang, Jing Huo, Yang Gao	88829
Neural Approximate Mirror Maps for Constrained Diffusion Models Berthy Feng, Ricardo Baptista, Katherine Bouman	88864
High-Precision Dichotomous Image Segmentation Via Probing Diffusion Capacity Qian Yu, Peng-Tao Jiang, Hao Zhang, Jinwei Chen, Bo Li, Lihe Zhang, Huchuan Lu	88888

<p>SEAL: Safety-Enhanced Aligned LLM Fine-Tuning Via Bilevel Data Selection</p> <p><i>Han Shen, Pin-Yu Chen, Payel Das, Tianyi Chen</i></p>	88902
<p>SWEb: A Large Web Dataset for the Scandinavian Languages</p> <p><i>Tobias Norlund, Tim Isbister, Amaru Cuba Gyllensten, Paul Gabriel Dos Santos, Danila Petrelli, Ariel Ekgren, Magnus Sahlgren</i></p>	88924
<p>Test-Time Alignment of Diffusion Models Without Reward Over-Optimization</p> <p><i>Sunwoo Kim, Minkyu Kim, Dongmin Park</i></p>	88948
<p>Correlation and Navigation in the Vocabulary Key Representation Space of Language Models</p> <p><i>Letian Peng, Chenyang An, Jingbo Shang</i></p>	88983
<p>TypedThinker: Diversify Large Language Model Reasoning with Typed Thinking</p> <p><i>Danqing Wang, Jianxin Ma, Fei Fang, Lei Li</i></p>	89005
<p>The Labyrinth of Links: Navigating the Associative Maze of Multi-Modal LLMs</p> <p><i>Hong Li, Nanxi Li, Yuanjie Chen, Jianbin Zhu, Qinlu Guo, Cewu Lu, Yong-Lu Li</i></p>	89027
<p>Scalable Extraction of Training Data from Aligned, Production Language Models</p> <p><i>Milad Nasr, Javier Rando, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne Ippolito, Christopher A. Choquette-Choo, Florian Tramèr, Katherine Lee</i></p>	89055
<p>Logic-Logit: A Logic-Based Approach to Choice Modeling</p> <p><i>Shuhan Zhang, Wendi Ren, Shuang Li</i></p>	89128
<p>DAWN: Dynamic Frame Avatar with Non-Autoregressive Diffusion Framework for Talking Head Video Generation</p> <p><i>Hanbo Cheng, Limin Lin, Chenyu Liu, Pengcheng Xia, Pengfei Hu, Jiefeng Ma, Jun Du, Jia Pan</i></p>	89149
<p>Memory Efficient Transformer Adapter for Dense Predictions</p> <p><i>Dong Zhang, Rui Yan, Pingcheng Dong, Kwang-Ting Cheng</i></p>	89168
<p>Learning View-Invariant World Models for Visual Robotic Manipulation</p> <p><i>Jing-Cheng Pang, Nan Tang, Kaiyuan Li, Yuting Tang, Xin-Qiang Cai, Zhen-Yu Zhang, Gang Niu, Masashi Sugiyama, Yang Yu</i></p>	89182
<p>Text4Seg: Reimagining Image Segmentation as Text Generation</p> <p><i>Mengcheng Lan, Chaofeng Chen, Yue Zhou, Jiaying Xu, Yiping Ke, Xinjiang Wang, Litong Feng, Wayne Zhang</i></p>	89206
<p>Linear Multistep Solver Distillation for Fast Sampling of Diffusion Models</p> <p><i>Yuchen Liang, Xiangzhong Fang, Hanting Chen, Yunhe Wang</i></p>	89234
<p>IMDPrompter: Adapting SAM to Image Manipulation Detection by Cross-View Automated Prompt Learning</p> <p><i>Quan Zhang, Yuxin Qi, Xi Tang, Jinwei Fang, Xi Lin, Ke Zhang, Chun Yuan</i></p>	89253

<p>Latent Radiance Fields with 3D-Aware 2D Representations</p> <p>Chaoyi Zhou, Xi Liu, Feng Luo, Siyu Huang</p>	89278
<p>FA3Set: Towards Analyzing Fast, Frequent, and Fine-Grained Events from Videos</p> <p>Zhaoyu Liu, Kan Jiang, Murong Ma, Zhe Hou, Yun Lin, Jin Song Dong</p>	89298
<p>Long-Tailed Adversarial Training with Self-Distillation</p> <p>Seungju Cho, Hongsin Lee, Changick Kim</p>	89313
<p>Frequency-Guided Masking for Enhanced Vision Self-Supervised Learning</p> <p>Amin Karimi Monsefi, Mengxi Zhou, Nastaran Karimi Monsefi, Ser-Nam Lim, Wei-Lun Chao, Rajiv Ramnath</p>	89332
<p>Revolutionizing EMCCD Denoising Through a Novel Physics-Based Learning Framework for Noise Modeling</p> <p>Haiyang Jiang, Tetsuichi Wazawa, Imari Sato, Takeharu Nagai, Yinqiang Zheng</p>	89361
<p>Designing Mechanical Meta-Materials by Learning Equivariant Flows</p> <p>Mehran Mirramezani, Anne S. Meeussen, Katia Bertoldi, Peter Orbanz, Ryan P Adams</p>	89379
<p>MCNC: Manifold-Constrained Reparameterization for Neural Compression</p> <p>Chayne Thrash, Reed Andreas, Ali Abbasi, Parsa Nooralinejad, Soroush Abbasi Koohpayegani, Hamed Pirsiavash, Soheil Kolouri</p>	89397
<p>KiVA: Kid-Inspired Visual Analogies for Testing Large Multimodal Models</p> <p>Eunice Yiu, Maan Qraitem, Anisa Noor Majhi, Charlie Wong, Yutong Bai, Shiry Ginosar, Alison Gopnik, Kate Saenko</p>	89417
<p>Inference Scaling Laws: An Empirical Analysis of Compute-Optimal Inference for LLM Problem-Solving</p> <p>Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, Yiming Yang</p>	89437
<p>Deep Weight Factorization: Sparse Learning Through the Lens of Artificial Symmetries</p> <p>Chris Kolb, Tobias Weber, Bernd Bischl, David Rügamer</p>	89459
<p>Convergence of Distributed Adaptive Optimization with Local Updates</p> <p>Ziheng Cheng, Margalit Glasgow</p>	89499
<p>K-HALU: Multiple Answer Korean Hallucination Benchmark for Large Language Models</p> <p>Jaehyung Seo, Heuseok Lim</p>	89557
<p>Activation Gradient Based Poisoned Sample Detection Against Backdoor Attacks</p> <p>Danni Yuan, Mingda Zhang, Shaokui Wei, Li Liu, Baoyuan Wu</p>	89588
<p>Training-Free LLM-Generated Text Detection by Mining Token Probability Sequences</p> <p>Yihuai Xu, Yongwei Wang, Yifei Bi, Huangsen Cao, Zhouhan Lin, Yu Zhao, Fei Wu</p>	89616

<p>Faster Cascades Via Speculative Decoding</p> <p><i>Harikrishna Narasimhan, Wittawat Jitkrittum, Ankit Singh Rawat, Seungyeon Kim, Neha Gupta, Aditya Krishna Menon, Sanjiv Kumar</i></p>	89643
<p>DynaMath: A Dynamic Visual Benchmark for Evaluating Mathematical Reasoning Robustness of Vision Language Models</p> <p><i>Chengke Zou, Xingang Guo, Rui Yang, Junyu Zhang, Bin Hu, Huan Zhang</i></p>	89682
<p>NeSyC: A Neuro-Symbolic Continual Learner for Complex Embodied Tasks in Open Domains</p> <p><i>Wonje Choi, Jinwoo Park, Sanghyun Ahn, Daehee Lee, Honguk Woo</i></p>	89729
<p>On the Modeling Capabilities of Large Language Models for Sequential Decision Making</p> <p><i>Martin Klissarov, R Devon Hjelm, Alexander T Toshev, Bogdan Mazouze</i></p>	89766
<p>Stochastic Bandits Robust to Adversarial Attacks</p> <p><i>Xuchuang Wang, Maoli Liu, Jinhang Zuo, Xutong Liu, John C. S. Lui, Mohammad Hajiesmaili</i></p>	89797
<p>An Effective Theory of Bias Amplification</p> <p><i>Arjun Subramonian, Samuel Bell, Levent Sagun, Elvis Dohmatob</i></p>	89817
<p>Bayesian Regularization of Latent Representation</p> <p><i>Chukwudi Paul Obite, Zhi Chang, Keyan Wu, Shiwei Lan</i></p>	89882
<p>Learning Interpretable Hierarchical Dynamical Systems Models from Time Series Data</p> <p><i>Manuel Brenner, Elias Weber, Georgia Koppe, Daniel Durstewitz</i></p>	89901
<p>Scaling Laws for Downstream Task Performance in Machine Translation</p> <p><i>Berivan Isik, Natalia Ponomareva, Hussein Hazimeh, Dimitris Paparas, Sergei Vassilvitskii, Sanmi Koyejo</i></p>	89938
<p>LoRA Done RITE: Robust Invariant Transformation Equilibration for LoRA Optimization</p> <p><i>Jui-Nan Yen, Si Si, Zhao Meng, Felix Yu, Sai Surya Duvvuri, Inderjit S Dhillon, Cho-Jui Hsieh, Sanjiv Kumar</i></p>	89960
<p>ParamΔ for Direct Mixing: Post-Train Large Language Model at Zero Cost</p> <p><i>Sheng Cao, Mingrui Wu, Karthik Prasad, Yuandong Tian, Zechun Liu</i></p>	89980
<p>ImagineNav: Prompting Vision-Language Models as Embodied Navigator Through Scene Imagination</p> <p><i>Xinxin Zhao, Wenzhe Cai, Likun Tang, Teng Wang</i></p>	90000
<p>Layer Swapping for Zero-Shot Cross-Lingual Transfer in Large Language Models</p> <p><i>Lucas Bandarkar, Benjamin Muller, Pritish Yuvraj, Rui Hou, Nayan Singhal, Hongjiang Lv, Bing Liu</i></p>	90015
<p>CARTS: Advancing Neural Theorem Proving with Diversified Tactic Calibration and Bias-Resistant Tree Search</p> <p><i>Xiao-Wen Yang, Zhi Zhou, Haiming Wang, Aoxue Li, Wen-Da Wei, Hui Jin, Zhenguo Li, Yu-Feng Li</i></p>	90039

<p>DisEnvisioner: Disentangled and Enriched Visual Prompt for Customized Image Generation</p> <p>Jing He, Haodong Li, Huyongzhe, Guibao Shen, Yingjie Cai, Weichao Qiu, Ying-Cong Chen</p>	90057
<p>Difference-Of-Submodular Bregman Divergence</p> <p>Masanari Kimura, Takahiro Kawashima, Tasuku Soma, Hideitsu Hino</p>	90070
<p>When is Task Vector Provably Effective for Model Editing? a Generalization Analysis of Nonlinear Transformers</p> <p>Hongkang Li, Yihua Zhang, Shuai Zhang, Pin-Yu Chen, Sijia Liu, Meng Wang</p>	90088
<p>Fengbo: A Clifford Neural Operator Pipeline for 3D PDEs in Computational Fluid Dynamics</p> <p>Alberto Pepe, Mattia Montanari, Joan Lasenby</p>	90121
<p>Routing Experts: Learning to Route Dynamic Experts in Existing Multi-Modal Large Language Models</p> <p>Qiong Wu, Zhaoxi Ke, Yiyi Zhou, Xiaoshuai Sun, Rongrong Ji</p>	90143
<p>GoodDrag: Towards Good Practices for Drag Editing with Diffusion Models</p> <p>Zewei Zhang, Huan Liu, Jun Chen, Xiangyu Xu</p>	90158
<p>MAGE: Model-Level Graph Neural Networks Explanations Via Motif-Based Graph Generation</p> <p>Zhaoning Yu, Hongyang Gao</p>	90182
<p>Benchmarking Agentic Workflow Generation</p> <p>Shuofei Qiao, Runnan Fang, Zhisong Qiu, Xiaobin Wang, Ningyu Zhang, Yong Jiang, Pengjun Xie, Fei Huang, Huajun Chen</p>	90201
<p>OSCAR: Operating System Control Via State-Aware Reasoning and Re-Planning</p> <p>Xiaoqiang Wang, Bang Liu</p>	90226
<p>Navigation-Guided Sparse Scene Representation for End-To-End Autonomous Driving</p> <p>Peidong Li, Dixiao Cui</p>	90249
<p>Decomposition Polyhedra of Piecewise Linear Functions</p> <p>Marie-Charlotte Brandenburg, Moritz Leo Grillo, Christoph Hertrich</p>	90266
<p>Benchmarking Multimodal Retrieval Augmented Generation with Dynamic VQA Dataset and Self-Adaptive Planning Agent</p> <p>Yangning Li, Yinghui Li, Xinyu Wang, Yong Jiang, Zhen Zhang, Xinran Zheng, Hui Wang, Hai-Tao Zheng, Fei Huang, Jingren Zhou, Philip S. Yu</p>	90296
<p>RecFlow: An Industrial Full Flow Recommendation Dataset</p> <p>Qi Liu, Kai Zheng, Rui Huang, Wuchao Li, Kuo Cai, Yuan Chai, Yanan Niu, Yiqun Hui, Bing Han, Na Mou, Hongning Wang, Wentian Bao, Yun En Yu, Guorui Zhou, Han Li, Yang Song, Defu Lian, Kun Gai</p>	90319
<p>The Pitfalls of Memorization: When Memorization Hurts Generalization</p> <p>Reza Bayat, Mohammad Pezeshki, Elvis Dohmatob, David Lopez-Paz, Pascal Vincent</p>	90337

<p>Youku Dense Caption: A Large-Scale Chinese Video Dense Caption Dataset and Benchmarks</p> <p>Zixuan Xiong, Guangwei Xu, Wenkai Zhang, Yuan Miao, Xuan Wu, Linhai, Ruijie Guo, Hai-Tao Zheng</p>	90364
<p>Reasoning Elicitation in Language Models Via Counterfactual Feedback</p> <p>Alihan Hüyük, Xinnuo Xu, Jacqueline R. M. A. Maasch, Aditya V. Nori, Javier Gonzalez</p>	90386
<p>Transformer Encoder Satisfiability: Complexity and Impact on Formal Reasoning</p> <p>Marco Sälzer, Eric Alsmann, Martin Lange</p>	90415
<p>Cross the Gap: Exposing the Intra-Modal Misalignment in CLIP Via Modality Inversion</p> <p>Marco Mistretta, Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, Andrew D. Bagdanov</p>	90437
<p>Tracing Representation Progression: Analyzing and Enhancing Layer-Wise Similarity</p> <p>Jiachen Jiang, Jinxin Zhou, Zhihui Zhu</p>	90459
<p>Image-Level Memorization Detection Via Inversion-Based Inference Perturbation</p> <p>Yue Jiang, Haokun Lin, Yang Bai, Bo Peng, Zhili Liu, Yueming Lyu, Yong Yang, Xingzheng, Jing Dong</p>	90485
<p>HaDeMiF: Hallucination Detection and Mitigation in Large Language Models</p> <p>Xiaoling Zhou, Mingjie Zhang, Zhemg Lee, Wei Ye, Shikun Zhang</p>	90505
<p>SVDQuant: Absorbing Outliers by Low-Rank Component for 4-Bit Diffusion Models</p> <p>Muyang Li, Yujun Lin, Zhekai Zhang, Tianle Cai, Xiuyu Li, Junxian Guo, Enze Xie, Chenlin Meng, Jun-Yan Zhu, Song Han</p>	90530
<p>Differentiable Optimization of Similarity Scores Between Models and Brains</p> <p>Nathan Cloos, Moufan Li, Markus Siegel, Scott L Brincat, Earl K Miller, Guangyu Robert Yang, Christopher J Cueva</p>	90558
<p>Solving Token Gradient Conflict in Mixture-Of-Experts for Large Vision-Language Model</p> <p>Longrong Yang, Dong Shen, Chaoxiang Cai, Fan Yang, Tingting Gao, Di Zhang, Xi Li</p>	90578
<p>DiscoveryBench: Towards Data-Driven Discovery with Large Language Models</p> <p>Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv Agarwal, Bhavana Dalvi Mishra, Abhijeetsingh Meena, Aryan Prakhar, Tirth Vora, Tushar Khot, Ashish Sabharwal, Peter Clark</p>	90605
<p>Latent Action Pretraining from Videos</p> <p>Seonghyeon Ye, Joel Jang, Byeongguk Jeon, Se June Joo, Jianwei Yang, Baolin Peng, Ajay Mandlekar, Reuben Tan, Yu-Wei Chao, Bill Yuchen Lin, Lars Liden, Kimin Lee, Jianfeng Gao, Luke Zettlemoyer, Dieter Fox, Minjoon Seo</p>	90629
<p>ObscuraCoder: Powering Efficient Code LM Pre-Training Via Obfuscation Grounding</p> <p>Indraneil Paul, Haoyi Yang, Goran Glavaš, Kristian Kersting, Iryna Gurevych</p>	90656
<p>MrT5: Dynamic Token Merging for Efficient Byte-Level Language Models</p> <p>Julie Kallini, Shikhar Murty, Christopher D Manning, Christopher Potts, Róbert Csordás</p>	90695
<p>TimeInf: Time Series Data Contribution Via Influence Functions</p> <p>Yizi Zhang, Jingyan Shen, Xiaoxue Xiong, Yongchan Kwon</p>	90719

<p>Open-World Reinforcement Learning Over Long Short-Term Imagination</p> <p><i>Jiajian Li, Qi Wang, Yunbo Wang, Xin Jin, Yang Li, Wenjun Zeng, Xiaokang Yang</i></p>	90744
<p>Universal Sharpness Dynamics in Neural Network Training: Fixed Point Analysis, Edge of Stability, and Route to Chaos</p> <p><i>Dayal Singh Kalra, Tianyu He, Maissam Barkeshli</i></p>	90767
<p>DistillHGNN: A Knowledge Distillation Approach for High-Speed Hypergraph Neural Networks</p> <p><i>Saman Forouzandeh, Parham Moradi Dw, Mahdi Jalili</i></p>	90802
<p>Differentially Private Federated Learning with Time-Adaptive Privacy Spending</p> <p><i>Shahrazad Kiani, Nupur Kulkarni, Adam Dziedzic, Stark Draper, Franziska Boenisch</i></p>	90824
<p>Decision Information Meets Large Language Models: The Future of Explainable Operations Research</p> <p><i>Yansen Zhang, Qingcan Kang, Wing Yin Yu, Haileigong, Xiaojin Fu, Xiongwei Han, Tao Zhong, Chen Ma</i></p>	90849
<p>Forget the Data and Fine-Tuning! Just Fold the Network to Compress</p> <p><i>Dong Wang, Haris Šikic, Lothar Thiele, Olga Saukh</i></p>	90874
<p>Overcoming Slow Decision Frequencies in Continuous Control: Model-Based Sequence Reinforcement Learning for Model-Free Control</p> <p><i>Devdhar Patel, Hava T Siegelmann</i></p>	90908
<p>FasterCache: Training-Free Video Diffusion Model Acceleration with High Quality</p> <p><i>Zhengyao Lv, Chenyang Si, Junhao Song, Zhenyu Yang, Yu Qiao, Ziwei Liu, Kwan-Yee K. Wong</i></p>	90940
<p>Growth Inhibitors for Suppressing Inappropriate Image Concepts in Diffusion Models</p> <p><i>Die Chen, Zhiwen Li, Mingyuan Fan, Cen Chen, Wenmeng Zhou, Yanhao Wang, Yaliang Li</i></p>	90965
<p>Auto-GDA: Automatic Domain Adaptation for Efficient Grounding Verification in Retrieval-Augmented Generation</p> <p><i>Tobias Leemann, Periklis Petridis, Giuseppe Vietri, Dionysis Manousakas, Aaron Roth, Sergul Aydore</i></p>	90986
<p>To CoT Or Not to CoT? Chain-Of-Thought Helps Mainly on Math and Symbolic Reasoning</p> <p><i>Zayne Rea Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, Greg Durrett</i></p>	91018
<p>In-Context Editing: Learning Knowledge from Self-Induced Distributions</p> <p><i>Siyuan Qi, Bangcheng Yang, Kailin Jiang, Xiaobo Wang, Jiaqi Li, Yifan Zhong, Yaodong Yang, Zilong Zheng</i></p>	91063
<p>Don't Take Things Out of Context: Attention Intervention for Enhancing Chain-Of-Thought Reasoning in Large Language Models</p> <p><i>Shaotian Yan, Chen Shen, Wenxiao Wang, Liang Xie, Junjie Liu, Jieping Ye</i></p>	91086

<p>Field-DiT: Diffusion Transformer on Unified Video, 3D, and Game Field Generation</p> <p><i>Kangfu Mei, Mo Zhou, Vishal M. Patel</i></p>	91108
<p>PIED: Physics-Informed Experimental Design for Inverse Problems</p> <p><i>Apivich Hemachandra, Gregory Kang Ruey Lau, See-Kiong Ng, Bryan Kian Hsiang Low</i></p>	91122
<p>Counterfactual Concept Bottleneck Models</p> <p><i>Gabriele Dominici, Pietro Barbiero, Francesco Giannini, Martin Gjoreski, Giuseppe Marra, Marc Langheinrich</i></p>	91153
<p>Adaptive Transformer Programs: Bridging the Gap Between Performance and Interpretability in Transformers</p> <p><i>Quoc-Vinh Lai-Dang, Taemin Kang, Seungah Son</i></p>	91177
<p>Hierarchical Uncertainty Estimation for Learning-Based Registration in Neuroimaging</p> <p><i>Xiaoling Hu, Karthik Gopinath, Peirong Liu, Malte Hoffmann, Koen Van Leemput, Oula Puonti, Juan Eugenio Iglesias</i></p>	91200
<p>Data Center Cooling System Optimization Using Offline Reinforcement Learning</p> <p><i>Xianyuan Zhan, Xiangyu Zhu, Peng Cheng, Xiao Hu, Ziteng He, Hanfei Geng, Jichao Leng, Huiwen Zheng, Chenhui Liu, Tianshun Hong, Yan Liang, Yunxin Liu, Feng Zhao</i></p>	91217
<p>Variational Best-Of-N Alignment</p> <p><i>Afra Amini, Tim Vieira, Elliott Ash, Ryan Cotterell</i></p>	91240
<p>Efficient Low-Bit Quantization with Adaptive Scales for Multi-Task Co-Training</p> <p><i>Boyu Liu, Haoyu Huang, Linlin Yang, Yanjing Li, Guodong Guo, Xianbin Cao, Baochang Zhang</i></p>	91261
<p>Adaptive Pruning of Pretrained Transformer Via Differential Inclusions</p> <p><i>Yizhuo Ding, Ke Fan, Yikai Wang, Xinwei Sun, Yanwei Fu</i></p>	91282
<p>PivotMesh: Generic 3D Mesh Generation Via Pivot Vertices Guidance</p> <p><i>Haohan Weng, Yikai Wang, Tong Zhang, C. L. Philip Chen, Jun Zhu</i></p>	91306
<p>Modeling Dynamic Social Vision Highlights Gaps Between Deep Learning and Humans</p> <p><i>Kathy Garcia, Emalie McMahon, Colin Conwell, Michael Bonner, Leyla Isik</i></p>	91326
<p>Salvage: Shapley-Distribution Approximation Learning Via Attribution Guided Exploration for Explainable Image Classification</p> <p><i>Mehdi Naouar, Hanne Raum, Jens Rahmfeld, Yannick Vogt, Joschka Boedecker, Gabriel Kalweit, Maria Kalweit</i></p>	91352
<p>Is Large-Scale Pretraining the Secret to Good Domain Generalization?</p> <p><i>Piotr Teterwak, Kuniaki Saito, Theodoros Tsiligkaridis, Bryan A. Plummer, Kate Saenko</i></p>	91371
<p>Do I Know This Entity? Knowledge Awareness and Hallucinations in Language Models</p> <p><i>Javier Ferrando, Oscar Balcells Obeso, Senthoooran Rajamanoharan, Neel Nanda</i></p>	91414

<p>Distributional Associations Vs In-Context Reasoning: A Study of Feed-Forward and Attention Layers</p> <p><i>Lei Chen, Joan Bruna, Alberto Bietti</i></p>	91450
<p>LongVILA: Scaling Long-Context Visual Language Models for Long Videos</p> <p><i>Yukang Chen, Fuzhao Xue, Dacheng Li, Qinghao Hu, Ligeng Zhu, Xiuyu Li, Yunhao Fang, Haotian Tang, Shang Yang, Zhijian Liu, Yihui He, Hongxu Yin, Pavlo Molchanov, Jan Kautz, Linxi Fan, Yuke Zhu, Yao Lu, Song Han</i></p>	91515
<p>Learning-Augmented Frequent Directions</p> <p><i>Anders Aamand, Justin Y. Chen, Siddharth Gollapudi, Sandeep Silwal, Hao Wu</i></p>	91535
<p>Intrinsic User-Centric Interpretability Through Global Mixture of Experts</p> <p><i>Vinitra Swamy, Syrielle Montariol, Julian Blackwell, Jibril Frej, Martin Jaggi, Tanja Käser</i></p>	91567
<p>Stable Hadamard Memory: Revitalizing Memory-Augmented Agents for Reinforcement Learning</p> <p><i>Hung Le, Dung Nguyen, Kien Do, Sunil Gupta, Svetha Venkatesh</i></p>	91582
<p>A Large-Scale Dataset and Benchmark for Commuting Origin-Destination Flow Generation</p> <p><i>Can Rong, Jingtao Ding, Yan Liu, Yong Li</i></p>	91600
<p>Loopy: Taming Audio-Driven Portrait Avatar with Long-Term Motion Dependency</p> <p><i>Jianwen Jiang, Chao Liang, Jiaqi Yang, Gaojie Lin, Tianyun Zhong, Yanbo Zheng</i></p>	91626
<p>Multi-Task Corrupted Prediction for Learning Robust Audio-Visual Speech Representation</p> <p><i>Sungnyun Kim, Sungwoo Cho, Sangmin Bae, Kangwook Jang, Se-Young Yun</i></p>	91645
<p>On Targeted Manipulation and Deception When Optimizing LLMs for User Feedback</p> <p><i>Marcus Williams, Micah Carroll, Adhyyan Narang, Constantin Weisser, Brendan Murphy, Anca Dragan</i></p>	91667
<p>I-Con: A Unifying Framework for Representation Learning</p> <p><i>Shaden Naif Alshammari, John R. Hershey, Axel Feldmann, William T. Freeman, Mark Hamilton</i></p>	91719
<p>Analyzing Neural Scaling Laws in Two-Layer Networks with Power-Law Data Spectra</p> <p><i>Roman Worschech, Bernd Rosenow</i></p>	91750
<p>Expected Return Symmetries</p> <p><i>Darius Muglich, Johannes Forkel, Elise Van Der Pol, Jakob Nicolaus Foerster</i></p>	91798
<p>Making Text Embedders Few-Shot Learners</p> <p><i>Chaofan Li, Minghao Qin, Shitao Xiao, Jianlyu Chen, Kun Luo, Defu Lian, Yingxia Shao, Zheng Liu</i></p>	91823
<p>Tree of Attributes Prompt Learning for Vision-Language Models</p> <p><i>Tong Ding, Wanhua Li, Zhongqi Miao, Hanspeter Pfister</i></p>	91841
<p>Decoding Game: On Minimax Optimality of Heuristic Text Generation Strategies</p> <p><i>Sijin Chen, Omar Hagrass, Jason Matthew Klusowski</i></p>	91861

<p>Graph Neural Networks Gone Hogwild</p> <p><i>Olga Solodova, Nick Richardson, Deniz Oktay, Ryan P Adams</i></p>	91881
<p>Scaling Laws for Precision</p> <p><i>Tanishq Kumar, Zachary Ankner, Benjamin Frederick Spector, Blake Bordelon, Niklas Muennighoff, Mansheej Paul, Cengiz Pehlevan, Christopher Re, Aditi Raghunathan</i></p>	91904
<p>Provably Accurate Shapley Value Estimation Via Leverage Score Sampling</p> <p><i>Christopher Musco, R. Teal Witter</i></p>	91936
<p>Rotated Runtime Smooth: Training-Free Activation Smoother for Accurate INT4 Inference</p> <p><i>Ke Yi, Zengke Liu, Jianwei Zhang, Chengyuan Li, Tong Zhang, Junyang Lin, Jingren Zhou</i></p>	91964
<p>Learning Fine-Grained Representations Through Textual Token Disentanglement in Composed Video Retrieval</p> <p><i>Yue Wu, Zhaobo Qi, Yiling Wu, Junshu Sun, Yaowei Wang, Shuhui Wang</i></p>	91981
<p>MGDA Converges Under Generalized Smoothness, Provably</p> <p><i>Qi Zhang, Peiyao Xiao, Shaofeng Zou, Kaiyi Ji</i></p>	92004
<p>Aligning Visual Contrastive Learning Models Via Preference Optimization</p> <p><i>Amirabbas Afzali, Borna Khodabandeh, Ali Rasekh, Mahyar Jafarinodeh, Sepehr Kazemi Ranjbar, Simon Gottschalk</i></p>	92036
<p>SaRA: High-Efficient Diffusion Model Fine-Tuning with Progressive Sparse Low-Rank Adaptation</p> <p><i>Teng Hu, Jiangning Zhang, Ran Yi, Hongrui Huang, Yabiao Wang, Lizhuang Ma</i></p>	92066
<p>Three Mechanisms of Feature Learning in a Linear Network</p> <p><i>Yizhou Xu, Liu Ziyin</i></p>	92079
<p>Deep Compression Autoencoder for Efficient High-Resolution Diffusion Models</p> <p><i>Junyu Chen, Han Cai, Junsong Chen, Enze Xie, Shang Yang, Haotian Tang, Muyang Li, Song Han</i></p>	92102
<p>ChroKnowledge: Unveiling Chronological Knowledge of Language Models in Multiple Domains</p> <p><i>Yein Park, Chanwoong Yoon, Jungwoo Park, Donghyeon Lee, Minbyul Jeong, Jaewoo Kang</i></p>	92124
<p>VLMaterial: Procedural Material Generation with Large Vision-Language Models</p> <p><i>Beichen Li, Rundi Wu, Armando Solar-Lezama, Changxi Zheng, Liang Shi, Bernd Bickel, Wojciech Matusik</i></p>	92164
<p>Learning Multi-Index Models with Neural Networks Via Mean-Field Langevin Dynamics</p> <p><i>Alireza Mousavi-Hosseini, Denny Wu, Murat A Erdogdu</i></p>	92185
<p>Needle Threading: Can LLMs Follow Threads Through Near-Million-Scale Haystacks?</p> <p><i>Jonathan Roberts, Kai Han, Samuel Albanie</i></p>	92219

<p>Self-Improvement in Language Models: The Sharpening Mechanism</p> <p>Audrey Huang, Adam Block, Dylan J Foster, Dhruv Rohatgi, Cyril Zhang, Max Simchowitz, Jordan T. Ash, Akshay Krishnamurthy</p>	92241
<p>DyCAST: Learning Dynamic Causal Structure from Time Series</p> <p>Yue Cheng, Bochen Lyu, Weiwei Xing, Zhanxing Zhu</p>	92294
<p>Building Math Agents with Multi-Turn Iterative Preference Learning</p> <p>Wei Xiong, Chengshuai Shi, Jiaming Shen, Aviv Rosenberg, Zhen Qin, Daniele Calandriello, Misha Khalman, Rishabh Joshi, Bilal Piot, Mohammad Saleh, Chi Jin, Tong Zhang, Tianqi Liu</p>	92319
<p>Effective Interplay Between Sparsity and Quantization: From Theory to Practice</p> <p>Simla Burcu Harma, Ayan Chakraborty, Elizaveta Kostenok, Danila Mishin, Dongho Ha, Babak Falsafi, Martin Jaggi, Ming Liu, Yunho Oh, Suvinay Subramanian, Amir Yazdanbakhsh</p>	92351
<p>Does Spatial Cognition Emerge in Frontier Models?</p> <p>Santhosh Kumar Ramakrishnan, Erik Wijmans, Philipp Kraehenbuehl, Vladlen Koltun</p>	92390
<p>Video In-Context Learning: Autoregressive Transformers Are Zero-Shot Video Imitators</p> <p>Wentao Zhang, Junliang Guo, Tianyu He, Li Zhao, Linli Xu, Jiang Bian</p>	92450
<p>Perm: A Parametric Representation for Multi-Style 3D Hair Modeling</p> <p>Chengan He, Xin Sun, Zhixin Shu, Fujun Luan, Soren Pirk, Jorge Alejandro Amador Herrera, Dominik Michels, Tuanfeng Yang Wang, Meng Zhang, Holly Rushmeier, Yi Zhou</p>	92471
<p>Tuning Frequency Bias of State Space Models</p> <p>Annan Yu, Dongwei Lyu, Soon Hoe Lim, Michael W. Mahoney, N. Benjamin Erichson</p>	92498
<p>Training Robust Ensembles Requires Rethinking Lipschitz Continuity</p> <p>Ali Ebrahimpour-Borojeny, Hari Sundaram, Varun Chandrasekaran</p>	92522
<p>RESfM: Robust Deep Equivariant Structure from Motion</p> <p>Fadi Khatib, Yoni Kasten, Dror Moran, Meirav Galun, Ronen Basri</p>	92548
<p>Optimal Learning of Kernel Logistic Regression for Complex Classification Scenarios</p> <p>Hongwei Wen, Annika Betken, Hanyuan Hang</p>	92567
<p>LancBiO: Dynamic Lanczos-Aided Bilevel Optimization Via Krylov Subspace</p> <p>Yan Yang, Bin Gao, Ya-Xiang Yuan</p>	92608
<p>Improved Algorithms for Kernel Matrix-Vector Multiplication Under Sparsity Assumptions</p> <p>Piotr Indyk, Michael Kapralov, Kshiteej Sheth, Tal Wagner</p>	92647
<p>Q-Adapter: Customizing Pre-Trained LLMs to New Preferences with Forgetting Mitigation</p> <p>Yi-Chen Li, Fuxiang Zhang, Wenjie Qiu, Lei Yuan, Chengxing Jia, Zongzhang Zhang, Yang Yu, Bo An</p>	92663
<p>Diffusion-Based Planning for Autonomous Driving with Flexible Guidance</p> <p>Yinan Zheng, Ruiming Liang, Kexin Zheng, Jinliang Zheng, Liyuan Mao, Jianxiong Li, Weihao Gu, Rui Ai, Shengbo Eben Li, Xianyuan Zhan, Jingjing Liu</p>	92688

<p>Learning to Communicate Through Implicit Communication Channels</p> <p><i>Han Wang, Binbin Chen, Tieying Zhang, Baoxiang Wang</i></p>	92709
<p>Credit-Based Self Organizing Maps: Training Deep Topographic Networks with Minimal Performance Degradation</p> <p><i>Amirozhan Dehghani, Xinyu Qian, Asa Farahani, Pouya Bashivan</i></p>	92726
<p>PFDiff: Training-Free Acceleration of Diffusion Models Combining Past and Future Scores</p> <p><i>Guangyi Wang, Yuren Cai, Lijiang Li, Wei Peng, Song-Zhi Su</i></p>	92745
<p>CausalRivers - Scaling Up Benchmarking of Causal Discovery for Real-World Time-Series</p> <p><i>Gideon Stein, Maha Shadaydeh, Jan Blunk, Niklas Penzel, Joachim Denzler</i></p>	92778
<p>Diffusion on Syntax Trees for Program Synthesis</p> <p><i>Shreyas Kapur, Erik Jenner, Stuart Russell</i></p>	92794
<p>GMValuator: Similarity-Based Data Valuation for Generative Models</p> <p><i>Jiaxi Yang, Wenlong Deng, Benlin Liu, Yangsibo Huang, James Zou, Xiaoxiao Li</i></p>	92813
<p>REBIND: Enhancing Ground-State Molecular Conformation Prediction Via Force-Based Graph Rewiring</p> <p><i>Taewon Kim, Hyunjin Seo, Sungsoo Ahn, Eunho Yang</i></p>	92840
<p>LLM-Based Typed Hyperresolution for Commonsense Reasoning with Knowledge Bases</p> <p><i>Armin Toroghi, Ali Pesaranghader, Tanmana Sadhu, Scott Sanner</i></p>	92858
<p>Explaining Modern Gated-Linear RNNs Via a Unified Implicit Attention Formulation</p> <p><i>Itamar Zimmerman, Ameen Ali Ali, Lior Wolf</i></p>	92894
<p>Scaling Up Masked Diffusion Models on Text</p> <p><i>Shen Nie, Fengqi Zhu, Chao Du, Tianyu Pang, Qian Liu, Guangtao Zeng, Min Lin, Chongxuan Li</i></p>	92913
<p>GenDataAgent: On-The-Fly Dataset Augmentation with Synthetic Data</p> <p><i>Zhiteng Li, Lele Chen, Jerone Andrews, Yunhao Ba, Yulun Zhang, Alice Xiang</i></p>	92937
<p>Scaling Autonomous Agents Via Automatic Reward Modeling and Planning</p> <p><i>Zhenfang Chen, Delin Chen, Rui Sun, Wenjun Liu, Chuang Gan</i></p>	92958
<p>Small Models Are LLM Knowledge Triggers for Medical Tabular Prediction</p> <p><i>Jiahuan Yan, Jintai Chen, Chaowen Hu, Bo Zheng, Yaojun Hu, Jimeng Sun, Jian Wu</i></p>	92990
<p>Probe Pruning: Accelerating LLMs Through Dynamic Pruning Via Model-Probing</p> <p><i>Qi Le, Enmao Diao, Ziyang Wang, Xinran Wang, Jie Ding, Li Yang, Ali Anwar</i></p>	93006
<p>Knowing Your Target: Target-Aware Transformer Makes Better Spatio-Temporal Video Grounding</p> <p><i>Xin Gu, Yaojie Shen, Chenxi Luo, Tiejian Luo, Yan Huang, Yuewei Lin, Heng Fan, Libo Zhang</i></p>	93026

<p>Improving Instruction-Following in Language Models Through Activation Steering</p> <p>Alessandro Stolfo, Vidhisha Balachandran, Safoora Yousefi, Eric Horvitz, Besmira Nushi</p>	93047
<p>Progressive Distillation Induces an Implicit Curriculum</p> <p>Abhishek Panigrahi, Bingbin Liu, Sadhika Malladi, Andrej Risteski, Surbhi Goel</p>	93081
<p>Re-Evaluating the Impact of Unseen-Class Unlabeled Data on Semi-Supervised Learning Model</p> <p>Rundong He, Yicong Dong, Lan-Zhe Guo, Yilong Yin, Tailin Wu</p>	93130
<p>Anyprefer: An Agentic Framework for Preference Data Synthesis</p> <p>Yiyang Zhou, Zhaoyang Wang, Tianle Wang, Shangyu Xing, Peng Xia, Bo Li, Kaiyuan Zheng, Zijian Zhang, Zhaorun Chen, Wenhao Zheng, Xuchao Zhang, Chetan Bansal, Weitong Zhang, Ying Wei, Mohit Bansal, Huaxiu Yao</p>	93150
<p>Process Reward Model with Q-Value Rankings</p> <p>Wendi Li, Yixuan Li</p>	93174
<p>NRGBoost: Energy-Based Generative Boosted Trees</p> <p>João Bravo</p>	93193
<p>Discovering Influential Neuron Path in Vision Transformers</p> <p>Yifan Wang, Yifei Liu, Yingdong Shi, Changming Li, Anqi Pang, Sibeil Yang, Jingyi Yu, Kan Ren</p>	93222
<p>Contextual Document Embeddings</p> <p>John Xavier Morris, Alexander M Rush</p>	93251
<p>Transition Path Sampling with Improved Off-Policy Training of Diffusion Path Samplers</p> <p>Kiyoung Seong, Seonghyun Park, Seonghwan Kim, Woo Youn Kim, Sungsoo Ahn</p>	93264
<p>Towards Synergistic Path-Based Explanations for Knowledge Graph Completion: Exploration and Evaluation</p> <p>Tengfei Ma, Xiang Song, Wen Tao, Mufei Li, Jiani Zhang, Xiaoqin Pan, Yijun Wang, Bosheng Song, Xiangxiang Zeng</p>	93287
<p>LARP: Tokenizing Videos with a Learned Autoregressive Generative Prior</p> <p>Hanyu Wang, Saksham Suri, Yixuan Ren, Hao Chen, Abhinav Shrivastava</p>	93308
<p>GaussianBlock: Building Part-Aware Compositional and Editable 3D Scene by Primitives and Gaussians</p> <p>Shuyi Jiang, Qihao Zhao, Hossein Rahmani, De Wen Soh, Jun Liu, Na Zhao</p>	93335
<p>A Spark of Vision-Language Intelligence: 2-Dimensional Autoregressive Transformer for Efficient Finegrained Image Generation</p> <p>Liang Chen, Sinan Tan, Zefan Cai, Weichu Xie, Haozhe Zhao, Yichi Zhang, Junyang Lin, Jinze Bai, Tianyu Liu, Baobao Chang</p>	93352
<p>The Power of LLM-Generated Synthetic Data for Stance Detection in Online Political Discussions</p> <p>Stefan Sylvius Wagner, Maike Behrendt, Marc Ziegele, Stefan Harmeling</p>	93377

<p>MMIU: Multimodal Multi-Image Understanding for Evaluating Large Vision-Language Models</p> <p>Fanqing Meng, Jin Wang, Chuanhao Li, Quanfeng Lu, Hao Tian, Tianshuo Yang, Jiaqi Liao, Xizhou Zhu, Jifeng Dai, Yu Qiao, Ping Luo, Kaipeng Zhang, Wenqi Shao</p>	93408
<p>SBSC: Step-By-Step Coding for Improving Mathematical Olympiad Performance</p> <p>Kunal Singh, Ankan Biswas, Sayandeep Bhowmick, Pradeep Moturi, Siva Kishore Gollapalli</p>	93457
<p>Feature Responsiveness Scores: Model-Agnostic Explanations for Recourse</p> <p>Seung Hyun Cheon, Anneke Wernerfelt, Sorelle Friedler, Berk Ustun</p>	93508
<p>TimeKAN: KAN-Based Frequency Decomposition Learning Architecture for Long-Term Time Series Forecasting</p> <p>Songtao Huang, Zhen Zhao, Can Li, Lei Bai</p>	93540
<p>Learning Evolving Tools for Large Language Models</p> <p>Guoxin Chen, Zhong Zhang, Xin Cong, Fangda Guo, Yesai Wu, Yankai Lin, Wenzheng Feng, Yasheng Wang</p>	93556
<p>A Theoretical Perspective: How to Prevent Model Collapse in Self-Consuming Training Loops</p> <p>Shi Fu, Yingjie Wang, Yuzhu Chen, Xinmei Tian, Dacheng Tao</p>	93586
<p>On Calibration of LLM-Based Guard Models for Reliable Content Moderation</p> <p>Hongfu Liu, Hengguan Huang, Xiangming Gu, Hao Wang, Ye Wang</p>	93620
<p>Towards Effective Evaluations and Comparisons for LLM Unlearning Methods</p> <p>Qizhou Wang, Bo Han, Puning Yang, Jianing Zhu, Tongliang Liu, Masashi Sugiyama</p>	93642
<p>Optimal Non-Asymptotic Rates of Value Iteration for Average-Reward Markov Decision Processes</p> <p>Jongmin Lee, Ernest K. Ryu</p>	93671
<p>FlashMask: Efficient and Rich Mask Extension of FlashAttention</p> <p>Guoxia Wang, Jinle Zeng, Xiyuan Xiao, Siming Wu, Jiabin Yang, Lujing Zheng, Zeyu Chen, Jiang Bian, Dianhai Yu, Haifeng Wang</p>	93714
<p>SINGER: Stochastic Network Graph Evolving Operator for High Dimensional PDEs</p> <p>Mingquan Feng, Yixin Huang, Weixin Liao, Yuhong Liu, Yizhou Liu, Junchi Yan</p>	93735
<p>Failures to Find Transferable Image Jailbreaks Between Vision-Language Models</p> <p>Rylan Schaeffer, Dan Valentine, Luke Bailey, James Chua, Cristobal Eyzaguirre, Zane Durante, Joe Benton, Brando Miranda, Henry Sleight, Tony Tong Wang, John Hughes, Rajashree Agrawal, Mrinank Sharma, Scott Emmons, Sanmi Koyejo, Ethan Perez</p>	93755
<p>Controllable Blur Data Augmentation Using 3D-Aware Motion Estimation</p> <p>Insoo Kim, Hana Lee, Hyong-Euk Lee, Jinwoo Shin</p>	93791
<p>When GNNs Meet Symmetry in ILPs: An Orbit-Based Feature Augmentation Approach</p> <p>Qian Chen, Lei Li, Qian Li, Jianghua Wu, Akang Wang, Ruoyu Sun, Xiaodong Luo, Tsung-Hui Chang, Qingjiang Shi</p>	93823

<p>Adaptive Methods Through the Lens of SDEs: Theoretical Insights on the Role of Noise</p> <p>Enea Monzio Compagnoni, Tianlin Liu, Rustem Islamov, Frank Norbert Proske, Antonio Orvieto, Aurelien Lucchi</p>	93839
<p>GANDALF: Generative Attention Based Data Augmentation and Predictive modeling Framework for Personalized Cancer Treatment</p> <p>Aishwarya Jayagopal, Yanrong Zhang, Robert John Walsh, Tuan Zea Tan, Anand D Jeyasekharan, Vaibhav Rajan</p>	93897
<p>Inspection and Control of Self-Generated-Text Recognition Ability in Llama3-8b-Instruct</p> <p>Christopher Ackerman, Nina Panickssery</p>	93918
<p>Relaxed Recursive Transformers: Effective Parameter Sharing with Layer-Wise LoRA</p> <p>Sangmin Bae, Adam Fisch, Hrayr Harutyunyan, Ziwei Ji, Seungyeon Kim, Tal Schuster</p>	93950
<p>BlueSuffix: Reinforced Blue Teaming for Vision-Language Models Against Jailbreak Attacks</p> <p>Yunhan Zhao, Xiang Zheng, Lin Luo, Yige Li, Xingjun Ma, Yu-Gang Jiang</p>	93996
<p>Semantic Aware Representation Learning for Lifelong Learning</p> <p>Fahad Sarfraz, Elahe Arani, Bahram Zonooz</p>	94016
<p>MA-RLHF: Reinforcement Learning from Human Feedback with Macro Actions</p> <p>Yekun Chai, Haoran Sun, Huang Fang, Shuohuan Wang, Yu Sun, Hua Wu</p>	94036
<p>Self-Supervised Diffusion MRI Denoising Via Iterative and Stable Refinement</p> <p>Chenxu Wu, Qingpeng Kong, Zihang Jiang, S Kevin Zhou</p>	94070
<p>GrabS: Generative Embodied Agent for 3D Object Segmentation Without Scene Supervision</p> <p>Zihui Zhang, Yafei Yang, Hongtao Wen, Bo Yang</p>	94110
<p>ZeroDiff: Solidified Visual-Semantic Correlation in Zero-Shot Learning</p> <p>Zihan Ye, Shreyank N Gowda, Shiming Chen, Xiaowei Huang, Haotian Xu, Fahad Shahbaz Khan, Yaochu Jin, Kaizhu Huang, Xiaobo Jin</p>	94134
<p>Scalable Universal T-Cell Receptor Embeddings from Adaptive Immune Repertoires</p> <p>Paidamoyo Chapfuwa, Ilker Demirel, Lorenzo Pisani, Javier Zazo, Elon Portugaly, H. Jabran Zahid, Julia Greissl</p>	94155
<p>Learning Successor Features with Distributed Hebbian Temporal Memory</p> <p>Evgenii Aleksandrovich Dzhivelikian, Petr Kuderov, Aleksandr Panov</p>	94176
<p>Scaling Optimal LR Across Token Horizons</p> <p>Johan Bjorck, Alon Benhaim, Vishrav Chaudhary, Furu Wei, Xia Song</p>	94201
<p>ST-GCond: Self-Supervised and Transferable Graph Dataset Condensation</p> <p>Beining Yang, Qingyun Sun, Cheng Ji, Xingcheng Fu, Jianxin Li</p>	94219
<p>A Solvable Attention for Neural Scaling Laws</p> <p>Bochen Lyu, Di Wang, Zhanxing Zhu</p>	94238

<p>Breaking Free from MMI: A New Frontier in Rationalization by Probing Input Utilization</p> <p>Wei Liu, Zhiying Deng, Zhongyu Niu, Jun Wang, Haozhao Wang, Zhigang Zeng, Ruixuan Li</p>	94278
<p>Topological Schrödinger Bridge Matching</p> <p>Maosheng Yang</p>	94302
<p>Breaking Class Barriers: Efficient Dataset Distillation Via Inter-Class Feature Compensator</p> <p>Xin Zhang, Jiawei Du, Ping Liu, Joey Tianyi Zhou</p>	94344
<p>Reassessing How to Compare and Improve the Calibration of Machine Learning Models</p> <p>Muthu Chidambaram, Rong Ge</p>	94357
<p>F-Fidelity: A Robust Framework for Faithfulness Evaluation of Explainable AI</p> <p>Xu Zheng, Farhad Shirani, Zhuomin Chen, Chaozhao Lin, Wei Cheng, Wenbo Guo, Dongsheng Luo</p>	94386
<p>Advancing Graph Generation Through Beta Diffusion</p> <p>Xinyang Liu, Yilin He, Bo Chen, Mingyuan Zhou</p>	94419
<p>Dynamic Multimodal Evaluation with Flexible Complexity by Vision-Language Bootstrapping</p> <p>Yue Yang, Shuibo Zhang, Kaipeng Zhang, Yi Bin, Yu Wang, Ping Luo, Wenqi Shao</p>	94445
<p>MACPO: Weak-To-Strong Alignment Via Multi-Agent Contrastive Preference Optimization</p> <p>Yougang Lyu, Lingyong Yan, Zihan Wang, Dawei Yin, Pengjie Ren, Maarten De Rijke, Zhaochun Ren</p>	94467
<p>Decoupling Angles and Strength in Low-Rank Adaptation</p> <p>Massimo Bini, Leander Gırrbach, Zeynep Akata</p>	94494
<p>SPORTU: A Comprehensive Sports Understanding Benchmark for Multimodal Large Language Models</p> <p>Haotian Xia, Zhengbang Yang, Junbo Zou, Rhys Tracy, Yuqing Wang, Chi Lu, Christopher Lai, Yanjun He, Xun Shao, Zhuoqing Xie, Yuan-Fang Wang, Weining Shen, Hanjie Chen</p>	94512
<p>Multi-Level Certified Defense Against Poisoning Attacks in Offline Reinforcement Learning</p> <p>Shijie Liu, Andrew Craig Cullen, Paul Montague, Sarah Monazam Erfani, Benjamin I. P. Rubinstein</p>	94571
<p>Which Tasks Should Be Compressed Together? a Causal Discovery Approach for Efficient Multi-Task Representation Compression</p> <p>Sha Guo, Jing Chen, Zixuan Hu, Zhuo Chen, Wenhan Yang, Yu Lin, Xing Jiang, Lingyu Duan</p>	94592
<p>Active Learning for Neural PDE Solvers</p> <p>Daniel Musekamp, Marimuthu Kalimuthu, David Holzmüller, Makoto Takamoto, Mathias Niepert</p>	94616
<p>Aria-MIDI: A Dataset of Piano MIDI Files for Symbolic Music Modeling</p> <p>Louis Bradshaw, Simon Colton</p>	94665
<p>SANER: Annotation-Free Societal Attribute Neutralizer for Debiasing CLIP</p> <p>Yusuke Hirota, Min-Hung Chen, Chien-Yi Wang, Yuta Nakashima, Yu-Chiang Frank Wang, Ryo Hachiuma</p>	94682

<p>Everything is Editable: Extend Knowledge Editing to Unstructured Data in Large Language Models</p> <p><i>Jingcheng Deng, Zihao Wei, Liang Pang, Hanxing Ding, Huawei Shen, Xueqi Cheng</i></p>	94701
<p>Zero-Shot Natural Language Explanations</p> <p><i>Fawaz Sammani, Nikos Deligiannis</i></p>	94726
<p>MANTRA: The Manifold Triangulations Assemblage</p> <p><i>Rubén Ballester, Ernst Röell, Daniel Bin Schmid, Mathieu Alain, Sergio Escalera, Carles Casacuberta, Bastian Rieck</i></p>	94750
<p>Long-Time Asymptotics of Noisy SVGD Outside the Population Limit</p> <p><i>Victor Priser, Pascal Bianchi, Adil Salim</i></p>	94780
<p>Beware of Calibration Data for Pruning Large Language Models</p> <p><i>Yixin Ji, Yang Xiang, Juntao Li, Qingrong Xia, Ping Li, Xinyu Duan, Zhefeng Wang, Min Zhang</i></p>	94812
<p>Fine-Tuning Can Help Detect Pretraining Data from Large Language Models</p> <p><i>Hengxiang Zhang, Songxin Zhang, Bingyi Jing, Hongxin Wei</i></p>	94829
<p>Centrality-Guided Pre-Training for Graph</p> <p><i>Bin Liang, Shiwei Chen, Lin Gui, Hui Wang, Yue Yu, Ruifeng Xu, Kam-Fai Wong</i></p>	94849
<p>MOOSE-Chem: Large Language Models for Rediscovering Unseen Chemistry Scientific Hypotheses</p> <p><i>Zonglin Yang, Wanhao Liu, Ben Gao, Tong Xie, Yuqiang Li, Wanli Ouyang, Soujanya Poria, Erik Cambria, Dongzhan Zhou</i></p>	94875
<p>Residual Stream Analysis with Multi-Layer SAEs</p> <p><i>Tim Lawson, Lucy Farnik, Conor Houghton, Laurence Aitchison</i></p>	94902
<p>Computational Explorations of Total Variation Distance</p> <p><i>Arnab Bhattacharyya, Sutanu Gayen, Kuldeep S. Meel, Dimitrios Myrisiotis, A. Pavan, N. V. Vinodchandran</i></p>	94947
<p>A Statistical Approach for Controlled Training Data Detection</p> <p><i>Zirui Hu, Yingjie Wang, Zheng Zhang, Hong Chen, Dacheng Tao</i></p>	94961
<p>Stem-OB: Generalizable Visual Imitation Learning with Stem-Like Convergent Observation Through Diffusion Inversion</p> <p><i>Kaizhe Hu, Zihang Rui, Yao He, Yuyao Liu, Pu Hua, Huazhe Xu</i></p>	94981
<p>Deep Signature: Characterization of Large-Scale Molecular Dynamics</p> <p><i>Tiexin Qin, Mengxu Zhu, Chunyang Li, Terry Lyons, Hong Yan, Haoliang Li</i></p>	95009
<p>SimulPL: Aligning Human Preferences in Simultaneous Machine Translation</p> <p><i>Donglei Yu, Yang Zhao, Jie Zhu, Yangyifan Xu, Yu Zhou, Chengqing Zong</i></p>	95027

DOCS: Quantifying Weight Similarity for Deeper Insights into Large Language Models Zeping Min, Xinshang Wang	95050
See it from My Perspective: How Language Affects Cultural Bias in Image Understanding Amith Ananthram, Elias Stengel-Eskin, Mohit Bansal, Kathleen McKeown	95081
Isometric Regularization for Manifolds of Functional Data Hyeongjun Heo, Seonghun Oh, Jae Yong Lee, Young Min Kim, Yonghyeon Lee	95103
When Selection Meets Intervention: Additional Complexities in Causal Discovery Haoyue Dai, Ignavier Ng, Jianle Sun, Zeyu Tang, Gongxu Luo, Xinshuai Dong, Peter Spirtes, Kun Zhang	95129
A Single Goal is All You Need: Skills and Exploration Emerge from Contrastive RL Without Rewards, Demonstrations, Or Subgoals Grace Liu, Michael Tang, Benjamin Eysenbach	95156
Learning to Discretize Denoising Diffusion ODEs Vinh Tong, Dung Trung Hoang, Anji Liu, Guy Van Den Broeck, Mathias Niepert	95179
Self-Supervised Monocular Depth Estimation Robust to Reflective Surface Leveraged by Triplet Mining Wonhyeok Choi, Kyumin Hwang, Wei Peng, Minwoo Choi, Sunghoon Im	95218
COAT: Compressing Optimizer States and Activations for Memory-Efficient FP8 Training Haocheng Xi, Han Cai, Ligeng Zhu, Yao Lu, Kurt Keutzer, Jianfei Chen, Song Han	95237
EmbodiedSAM: Online Segment Any 3D Thing in Real Time Xiuwei Xu, Huangxing Chen, Linqing Zhao, Ziwei Wang, Jie Zhou, Jiwen Lu	95260
Competition Dynamics Shape Algorithmic Phases of In-Context Learning Core Francisco Park, Ekdeep Singh Lubana, Hidenori Tanaka	95276
Speculative RAG: Enhancing Retrieval Augmented Generation Through Drafting Zilong Wang, Zifeng Wang, Long Le, Steven Zheng, Swaroop Mishra, Vincent Perot, Yuwei Zhang, Anush Mattapalli, Ankur Taly, Jingbo Shang, Chen-Yu Lee, Tomas Pfister	95329
Better Autoregressive Regression with LLMs Via Regression-Aware Fine-Tuning Michal Lukasik, Zhao Meng, Harikrishna Narasimhan, Yin-Wen Chang, Aditya Krishna Menon, Felix Yu, Sanjiv Kumar	95352
MindSearch: Mimicking Human Minds Elicits Deep AI Searcher Zehui Chen, Kuikun Liu, Qiuchen Wang, Jiangning Liu, Wenwei Zhang, Kai Chen, Feng Zhao	95383
ProtoSnap: Prototype Alignment for Cuneiform Signs Rachel Mikulinsky, Morris Alper, Shai Gordin, Enrique Jiménez, Yoram Cohen, Hadar Averbuch-Elor	95406
(Mis)Fitting Scaling Laws: A Survey of Scaling Law Fitting Techniques in Deep Learning Margaret Li, Sneha Kudugunta, Luke Zettlemoyer	95426

<p>Mixture of Experts Made Personalized: Federated Prompt Learning for Vision-Language Models</p> <p>Jun Luo, Chen Chen, Shandong Wu</p>	95458
<p>Vector-ICL: In-Context Learning with Continuous Vector Representations</p> <p>Yufan Zhuang, Chandan Singh, Liyuan Liu, Jingbo Shang, Jianfeng Gao</p>	95478
<p>A Generic Framework for Conformal Fairness</p> <p>Aditya T. Vadlamani, Anutam Srinivasan, Pranav Maneriker, Ali Payani, Srinivasan Parthasarathy</p>	95501
<p>Building, Reusing, and Generalizing Abstract Representations from Concrete Sequences</p> <p>Shuchen Wu, Mirko Thalmann, Peter Dayan, Zeynep Akata, Eric Schulz</p>	95534
<p>Optimistic Games for Combinatorial Bayesian Optimization with Application to Protein Design</p> <p>Melis Ilayda Bal, Pier Giuseppe Sessa, Mojmír Mutný, Andreas Krause</p>	95558
<p>Revisiting In-Context Learning Inference Circuit in Large Language Models</p> <p>Hakaze Cho, Mariko Kato, Yoshihiro Sakai, Naoya Inoue</p>	95589
<p>Optimizing Backward Policies in GFlowNets Via Trajectory Likelihood Maximization</p> <p>Timofei Gritsaev, Nikita Morozov, Sergey Samsonov, Daniil Tiapkin</p>	95626
<p>MoISpectra: Pre-Training 3D Molecular Representation with Multi-Modal Energy Spectra</p> <p>Liang Wang, Shaozhen Liu, Yu Rong, Deli Zhao, Qiang Liu, Shu Wu, Liang Wang</p>	95647
<p>SafeWatch: An Efficient Safety-Policy Following Video Guardrail Model with Transparent Explanations</p> <p>Zhaorun Chen, Francesco Pinto, Minzhou Pan, Bo Li</p>	95666
<p>Language Models Learn to Mislead Humans Via RLHF</p> <p>Jiaxin Wen, Ruiqi Zhong, Akbir Khan, Ethan Perez, Jacob Steinhardt, Minlie Huang, Samuel R. Bowman, He He, Shi Feng</p>	95709
<p>Enhancing Uncertainty Estimation and Interpretability with Bayesian Non-Negative Decision Layer</p> <p>Xinyue Hu, Zhibin Duan, Bo Chen, Mingyuan Zhou</p>	95732
<p>SeCom: On Memory Construction and Retrieval for Personalized Conversational Agents</p> <p>Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Xufang Luo, Hao Cheng, Dongsheng Li, Yuqing Yang, Chin-Yew Lin, H. Vicky Zhao, Lili Qiu, Jianfeng Gao</p>	95755
<p>Correlating Instruction-Tuning (in Multimodal Models) with Vision-Language Processing (in the Brain)</p> <p>Subba Reddy Oota, Akshett Rai Jindal, Ishani Mondal, Khushbu Pahwa, Satya Sai Srinath Namburi Gnvv, Manish Shrivastava, Maneesh Kumar Singh, Bapi Raju Surampudi, Manish Gupta</p>	95790
<p>Query-Based Knowledge Transfer for Heterogeneous Learning Environments</p> <p>Norah Alballa, Wenxuan Zhang, Ziquan Liu, Ahmed M. Abdelmoniem, Mohamed Elhoseiny, Marco Canini</p>	95820

<p>MaxCutPool: Differentiable Feature-Aware Maxcut for Pooling in Graph Neural Networks</p> <p>Carlo Abate, Filippo Maria Bianchi</p>	95852
<p>ToolGen: Unified Tool Retrieval and Calling Via Generation</p> <p>Renxi Wang, Xudong Han, Lei Ji, Shu Wang, Timothy Baldwin, Haonan Li</p>	95880
<p>Modeling Unseen Environments with Language-Guided Composable Causal Components in Reinforcement Learning</p> <p>Xinyue Wang, Biwei Huang</p>	95906
<p>TSC-Net: Prediction of Pedestrian Trajectories by Trajectory-Scene-Cell Classification</p> <p>Bo Hu, Tat-Jen Cham</p>	95935
<p>ContraDiff: Planning Towards High Return States Via Contrastive Learning</p> <p>Yixiang Shan, Zhengbang Zhu, Ting Long, Liang Qifan, Yi Chang, Weinan Zhang, Liang Yin</p>	95949
<p>EIA: ENVIRONMENTAL INJECTION ATTACK ON GENERALIST WEB AGENTS FOR PRIVACY LEAKAGE</p> <p>Zeyi Liao, Lingbo Mo, Chejian Xu, Mintong Kang, Jiawei Zhang, Chaowei Xiao, Yuan Tian, Bo Li, Huan Sun</p>	95971
<p>Spider 2.0: Evaluating Language Models on Real-World Enterprise Text-To-SQL Workflows</p> <p>Fangyu Lei, Jixuan Chen, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin Su, Zhaoqing Suo, Hongcheng Gao, Wenjing Hu, Pengcheng Yin, Victor Zhong, Caiming Xiong, Ruoxi Sun, Qian Liu, Sida Wang, Tao Yu</p>	96003
<p>DartControl: A Diffusion-Based Autoregressive Motion Model for Real-Time Text-Driven Motion Control</p> <p>Kaifeng Zhao, Gen Li, Siyu Tang</p>	96048
<p>Neural ODE Transformers: Analyzing Internal Dynamics and Adaptive Fine-Tuning</p> <p>Anh Tong, Thanh Nguyen-Tang, Dongeun Lee, Duc Nguyen, Toan Tran, David Leo Wright Hall, Cheongwoong Kang, Jaesik Choi</p>	96072
<p>Feature-Based Online Bilateral Trade</p> <p>Solenne Gaucher, Martino Bernasconi, Matteo Castiglioni, Andrea Celli, Vianney Perchet</p>	96106
<p>Discrete Distribution Networks</p> <p>Lei Yang</p>	96143
<p>SymmCD: Symmetry-Preserving Crystal Generation with Diffusion Models</p> <p>Daniel Levy, Siba Smarak Panigrahi, Sékou-Oumar Kaba, Qiang Zhu, Kin Long Kelvin Lee, Mikhail Galkin, Santiago Miret, Siamak Ravanbakhsh</p>	96167
<p>Beyond Circuit Connections: A Non-Message Passing Graph Transformer Approach for Quantum Error Mitigation</p> <p>Tianyi Bao, Xinyu Ye, Hang Ruan, Chang Liu, Wenjie Wu, Junchi Yan</p>	96191

<p>On Conformal Isometry of Grid Cells: Learning Distance-Preserving Position Embedding</p> <p><i>Dehong Xu, Ruiqi Gao, Wenhao Zhang, Xue-Xin Wei, Ying Nian Wu</i></p>	96216
<p>Can a MISL Fly? Analysis and Ingredients for Mutual Information Skill Learning</p> <p><i>Chongyi Zheng, Jens Tuyls, Joanne Peng, Benjamin Eysenbach</i></p>	96239
<p>Learning General-Purpose Biomedical Volume Representations Using Randomized Synthesis</p> <p><i>Neel Dey, Benjamin Billot, Hallee E. Wong, Clinton Wang, Mengwei Ren, Ellen Grant, Adrian V Dalca, Polina Golland</i></p>	96267
<p>IntersectionZoo: Eco-Driving for Benchmarking Multi-Agent Contextual Reinforcement Learning</p> <p><i>Vindula Jayawardana, Baptiste Freydt, Ao Qu, Cameron Hickert, Zhongxia Yan, Cathy Wu</i></p>	96299
<p>Syntactic and Semantic Control of Large Language Models Via Sequential Monte Carlo</p> <p><i>João Loula, Benjamin Lebrun, Li Du, Ben Lipkin, Clemente Pasti, Gabriel Grand, Tianyu Liu, Yahya Emara, Marjorie Freedman, Jason Eisner, Ryan Cotterell, Vikash Mansinghka, Alexander K. Lew, Tim Vieira, Timothy J. O'Donnell</i></p>	96326
<p>Deep Linear Probe Generators for Weight Space Learning</p> <p><i>Jonathan Kahana, Eliahu Horwitz, Imri Shuval, Yedid Hoshen</i></p>	96360
<p>Endless Jailbreaks with Bijection Learning</p> <p><i>Brian R. Y. Huang, Maximilian Li, Leonard Tang</i></p>	96378
<p>StochSync: Stochastic Diffusion Synchronization for Image Generation in Arbitrary Spaces</p> <p><i>Kyeongmin Yeo, Jaihoon Kim, Minhyuk Sung</i></p>	96399
<p>UniRestore3D: A Scalable Framework for General Shape Restoration</p> <p><i>Yuang Wang, Yujian Zhang, Sida Peng, Xingyi He, Haoyu Guo, Yujun Shen, Hujun Bao, Xiaowei Zhou</i></p>	96433
<p>Optimal Transport for Time Series Imputation</p> <p><i>Hao Wang, Zhengnan Li, Haoxuan Li, Xu Chen, Mingming Gong, Binchen, Zhichao Chen</i></p>	96455
<p>Reflective Gaussian Splatting</p> <p><i>Yuxuan Yao, Zixuan Zeng, Chun Gu, Xiatian Zhu, Li Zhang</i></p>	96480
<p>UNIP: Rethinking Pre-Trained Attention Patterns for Infrared Semantic Segmentation</p> <p><i>Tao Zhang, Jinyong Wen, Zhen Chen, Kun Ding, Shiming Xiang, Chunhong Pan</i></p>	96497
<p>Adjoint Matching: Fine-Tuning Flow and Diffusion Generative Models with Memoryless Stochastic Optimal Control</p> <p><i>Carles Domingo-Enrich, Michal Drozdal, Brian Karrer, Ricky T. Q. Chen</i></p>	96524
<p>Context Steering: Controllable Personalization at Inference Time</p> <p><i>Jerry Zhi-Yang He, Sashrika Pandey, Mariah L Schrum, Anca Dragan</i></p>	96580
<p>Going Beyond Static: Understanding Shifts with Time-Series Attribution</p> <p><i>Jiashuo Liu, Nabeel Seedat, Peng Cui, Mihaela Van Der Schaar</i></p>	96613

Consistency Models Made Easy <i>Zhengyang Geng, Ashwini Pokle, Weijian Luo, Justin Lin, J Zico Kolter</i>	96638
Attributing Culture-Conditioned Generations to Pretraining Corpora <i>Huihan Li, Arnav Goel, Keyu He, Xiang Ren</i>	96666
Contrastive Learning from Synthetic Audio Doppelgänger <i>Manuel Cherep, Nikhil Singh</i>	96692
JudgeLM: Fine-Tuned Large Language Models Are Scalable Judges <i>Lianghui Zhu, Xinggang Wang, Xinlong Wang</i>	96711
Unlocking Guidance for Discrete State-Space Diffusion and Flow Models <i>Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, Jennifer Listgarten</i>	96751
Advantage-Guided Distillation for Preference Alignment in Small Language Models <i>Shiping Gao, Fanqi Wan, Jiajian Guo, Xiaojun Quan, Qifan Wang</i>	96806
EDiT: A Local-SGD-Based Efficient Distributed Training Method for Large Language Models <i>Jialiang Cheng, Ning Gao, Yun Yue, Zhiling Ye, Jiadi Jiang, Jian Sha</i>	96829
AnyTouch: Learning Unified Static-Dynamic Representation Across Multiple Visuo-Tactile Sensors <i>Ruoxuan Feng, Jiangyu Hu, Wenke Xia, Tiancigao, Ao Shen, Yuhao Sun, Bin Fang, Di Hu</i>	96851
Geometry of Lightning Self-Attention: Identifiability and Dimension <i>Nathan W. Henry, Giovanni Luca Marchetti, Kathlén Kohn</i>	96872
Structural-Entropy-Based Sample Selection for Efficient and Effective Learning <i>Tianchi Xie, Jiangning Zhu, Guozu Ma, Minzhi Lin, Wei Chen, Weikai Yang, Shixia Liu</i>	96889
PixWizard: Versatile Image-To-Image Visual Assistant with Open-Language Instructions <i>Weifeng Lin, Xinyu Wei, Renrui Zhang, Le Zhuo, Shitian Zhao, Siyuan Huang, Junlin Xie, Peng Gao, Hongsheng Li</i>	96920
New Algorithms for the Learning-Augmented K-Means Problem <i>Junyu Huang, Qilong Feng, Ziyun Huang, Zhen Zhang, Jinhui Xu, Jianxin Wang</i>	96952
Diffusion Models Are Evolutionary Algorithms <i>Yanbo Zhang, Benedikt Hartl, Hananel Hazan, Michael Levin</i>	96985
InstantSplamp: Fast and Generalizable Stenography Framework for Generative Gaussian Splatting <i>Chenxin Li, Hengyu Liu, Zhiwen Fan, Wuyang Li, Yifan Liu, Panwang Pan, Yixuan Yuan</i>	97007
MP-Mat: A 3D-And-Instance-Aware Human Matting and Editing Framework with Multiplane Representation <i>Siyi Jiao, Wenzheng Zeng, Yerong Li, Huayu Zhang, Changxin Gao, Nong Sang, Mike Zheng Shou</i>	97021

<p>Graph Neural Networks for Edge Signals: Orientation Equivariance and Invariance</p> <p><i>Dominik Fuchsgruber, Tim Postuvan, Stephan Günnemann, Simon Geisler</i></p>	97044
<p>Rapidly Adapting Policies to the Real-World Via Simulation-Guided Fine-Tuning</p> <p><i>Patrick Yin, Tyler Westenbroek, Ching-An Cheng, Andrey Kolobov, Abhishek Gupta</i></p>	97078
<p>Improved Sampling Algorithms for Lévy-Itô Diffusion Models</p> <p><i>Vadim Popov, Assel Yermekova, Tasnima Sadekova, Artem Khrapov, Mikhail Sergeevich Kudinov</i></p>	97100
<p>OLMoE: Open Mixture-Of-Experts Language Models</p> <p><i>Niklas Muennighoff, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Jacob Morrison, Sewon Min, Weijia Shi, Evan Pete Walsh, Oyvind Tafjord, Nathan Lambert, Yuling Gu, Shane Arora, Akshita Bhagia, Dustin Schwenk, David Wadden, Alexander Wettig, Binyuan Hui, Tim Dettmers, Douwe Kiela, Ali Farhadi, Noah A. Smith, Pang Wei Koh, Amanpreet Singh, Hannaneh Hajishirzi</i></p>	97116
<p>Incorporating Visual Correspondence into Diffusion Model for Virtual Try-On</p> <p><i>Siqi Wan, Jingwen Chen, Yingwei Pan, Ting Yao, Tao Mei</i></p>	97177
<p>3D-SPATIAL MULTIMODAL MEMORY</p> <p><i>Xueyan Zou, Yuchen Song, Ri-Zhao Qiu, Xuanbin Peng, Jianglong Ye, Sifei Liu, Xiaolong Wang</i></p>	97199
<p>DSPO: Direct Score Preference Optimization for Diffusion Model Alignment</p> <p><i>Huaisheng Zhu, Teng Xiao, Vasant G Honavar</i></p>	97212
<p>Agree to Disagree: Demystifying Homogeneous Deep Ensembles Through Distributional Equivalence</p> <p><i>Yipei Wang, Xiaoqian Wang</i></p>	97234
<p>Tighter Privacy Auditing of DP-SGD in the Hidden State Threat Model</p> <p><i>Tudor Ioan Cebere, Aurélien Bellet, Nicolas Papernot</i></p>	97266
<p>PALMBENCH: A COMPREHENSIVE BENCHMARK OF COMPRESSED LARGE LANGUAGE MODELS ON MOBILE PLATFORMS</p> <p><i>Yilong Li, Jingyu Liu, Hao Zhang, M Badri Narayanan, Utkarsh Sharma, Shuai Zhang, Yijing Zeng, Jayaram Raghuram, Suman Banerjee</i></p>	97289
<p>Efficient Neuron Segmentation in Electron Microscopy by Affinity-Guided Queries</p> <p><i>Hang Chen, Chufeng Tang, Xiao Li, Xiaolin Hu</i></p>	97308
<p>GraphArena: Evaluating and Exploring Large Language Models on Graph Computation</p> <p><i>Jianheng Tang, Qifan Zhang, Yuhan Li, Nuo Chen, Jia Li</i></p>	97329
<p>OMNI-EPIC: Open-Endedness Via Models of Human Notions of Interestingness with Environments Programmed in Code</p> <p><i>Maxence Faldor, Jenny Zhang, Antoine Cully, Jeff Clune</i></p>	97357
<p>EditRoom: LLM-Parameterized Graph Diffusion for Composable 3D Room Layout Editing</p> <p><i>Kaizhi Zheng, Xiaotong Chen, Xuehai He, Jing Gu, Linjie Li, Zhengyuan Yang, Kevin Lin, Jianfeng Wang, Lijuan Wang, Xin Eric Wang</i></p>	97483

<p>Eagle: Exploring the Design Space for Multimodal LLMs with Mixture of Encoders</p> <p>Min Shi, Fuxiao Liu, Shihao Wang, Shijia Liao, Subhashree Radhakrishnan, Yilin Zhao, De-An Huang, Hongxu Yin, Karan Sapra, Yaser Yacoub, Humphrey Shi, Bryan Catanzaro, Andrew Tao, Jan Kautz, Zhiding Yu, Guilin Liu</p>	97501
<p>HarmAug: Effective Data Augmentation for Knowledge Distillation of Safety Guard Models</p> <p>Seanie Lee, Haebin Seong, Dong Bok Lee, Minki Kang, Xiaoyin Chen, Dominik Wagner, Yoshua Bengio, Juho Lee, Sung Ju Hwang</p>	97520
<p>Follow My Instruction and Spill the Beans: Scalable Data Extraction from Retrieval-Augmented Generation Systems</p> <p>Zhenting Qi, Hanlin Zhang, Eric P. Xing, Sham M. Kakade, Himabindu Lakkaraju</p>	97540
<p>Boosting Latent Diffusion with Perceptual Objectives</p> <p>Tariq Berrada, Pietro Astolfi, Melissa Hall, Marton Havasi, Yohann Benchetrit, Adriana Romero-Soriano, Karteek Alahari, Michal Drozdal, Jakob Verbeek</p>	97563
<p>PIG: Physics-Informed Gaussians as Adaptive Parametric Mesh Representations</p> <p>Namgyu Kang, Jaemin Oh, Youngjoon Hong, Eunbyung Park</p>	97577
<p>GOLD: Graph Out-Of-Distribution Detection Via Implicit Adversarial Latent Generation</p> <p>Danny Wang, Ruihong Qiu, Guangdong Bai, Zi Huang</p>	97602
<p>Efficient Stagewise Pretraining Via Progressive Subnetworks</p> <p>Abhishek Panigrahi, Nikunj Saunshi, Kaifeng Lyu, Sobhan Miryoosefi, Sashank J. Reddi, Satyen Kale, Sanjiv Kumar</p>	97632
<p>MOS: Model Synergy for Test-Time Adaptation on LiDAR-Based 3D Object Detection</p> <p>Zhuoxiao Chen, Junjie Meng, Mahsa Baktashmotlagh, Yonggang Zhang, Zi Huang, Yadan Luo</p>	97680
<p>Infinite-Resolution Integral Noise Warping for Diffusion Models</p> <p>Yitong Deng, Winnie Lin, Lingxiao Li, Dmitry Smirnov, Ryan D Burgert, Ning Yu, Vincent Dedun, Mohammad H. Taghavi</p>	97705
<p>Efficient Interpolation Between Extragradient and Proximal Methods for Weak MVIs</p> <p>Thomas Pethick, Ioannis Mavrothalassitis, Volkan Cevher</p>	97724
<p>Agent-To-Sim: Learning Interactive Behavior Models from Casual Longitudinal Videos</p> <p>Gengshan Yang, Andrea Bajcsy, Shunsuke Saito, Angjoo Kanazawa</p>	97752
<p>SigDiffusions: Score-Based Diffusion Models for Time Series Via Log-Signature Embeddings</p> <p>Barbora Barancikova, Zhuoyue Huang, Cristopher Salvi</p>	97775
<p>Reconstruction-Guided Policy: Enhancing Decision-Making Through Agent-Wise State Consistency</p> <p>Liang Qifan, Yixiang Shan, Haipeng Liu, Zhengbang Zhu, Ting Long, Weinan Zhang, Yuan Tian</p>	97811

<p>Sort-Free Gaussian Splatting Via Weighted Sum Rendering</p> <p><i>Qiqi Hou, Randall Rauwendaal, Zifeng Li, Hoang Le, Farzad Farhadzadeh, Fatih Porikli, Alexei Bourd, Amir Said</i></p>	97829
<p>Language Agents Meet Causality -- Bridging LLMs and Causal World Models</p> <p><i>John Gkountouras, Matthias Lindemann, Phillip Lippe, Efstratios Gavves, Ivan Titov</i></p>	97846
<p>Quality Over Quantity in Attention Layers: When Adding More Heads Hurts</p> <p><i>Noah Amsel, Gilad Yehudai, Joan Bruna</i></p>	97872
<p>A Common Pitfall of Margin-Based Language Model Alignment: Gradient Entanglement</p> <p><i>Hui Yuan, Yifan Zeng, Yue Wu, Huazheng Wang, Mengdi Wang, Liu Leqi</i></p>	97925
<p>BinaryDM: Accurate Weight Binarization for Efficient Diffusion Models</p> <p><i>Xingyu Zheng, Xianglong Liu, Haotong Qin, Xudong Ma, Mingyuan Zhang, Haojie Hao, Jiakai Wang, Zixiang Zhao, Jinyang Guo, Michele Magno</i></p>	97945
<p>NoVo: Norm Voting off Hallucinations with Attention Heads in Large Language Models</p> <p><i>Zheng Yi Ho, Siyuan Liang, Sen Zhang, Yibing Zhan, Dacheng Tao</i></p>	97967
<p>Language-Image Models with 3D Understanding</p> <p><i>Jang Hyun Cho, Boris Ivanovic, Yulong Cao, Edward Schmerling, Yue Wang, Xinshuo Weng, Boyi Li, Yurong You, Philipp Kraehenbuehl, Yan Wang, Marco Pavone</i></p>	97991
<p>Omni-MATH: A Universal Olympiad Level Mathematic Benchmark for Large Language Models</p> <p><i>Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma, Liang Chen, Runxin Xu, Zhengyang Tang, Benyou Wang, Daoguang Zan, Shanghaoran Quan, Ge Zhang, Lei Sha, Yichang Zhang, Xuancheng Ren, Tianyu Liu, Baobao Chang</i></p>	98023
<p>Dissecting Adversarial Robustness of Multimodal LM Agents</p> <p><i>Chen Henry Wu, Rishi Rajesh Shah, Jing Yu Koh, Russ Salakhutdinov, Daniel Fried, Aditi Raghunathan</i></p>	98053
<p>RDT-1B: A Diffusion Foundation Model for Bimanual Manipulation</p> <p><i>Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, Jun Zhu</i></p>	98075
<p>Scaling Wearable Foundation Models</p> <p><i>Girish Narayanswamy, Xin Liu, Kumar Ayush, Yuzhe Yang, Xuhai Xu, Shun Liao, Jake Garrison, Shyam A. Tailor, Jacob Sunshine, Yun Liu, Tim Althoff, Shrikanth Narayanan, Pushmeet Kohli, Jiening Zhan, Mark Malhotra, Shwetak Patel, Samy Abdel-Ghaffar, Daniel McDuff</i></p>	98103
<p>AIMS.au: A Dataset for the Analysis of Modern Slavery Countermeasures in Corporate Statements</p> <p><i>Adriana Eufrosina Bora, Pierre-Luc St-Charles, Mirko Bronzi, Arsene Fansi Tchango, Bruno Rousseau, Kerrie Mengersen</i></p>	98135
<p>ThermalGaussian: Thermal 3D Gaussian Splatting</p> <p><i>Rongfeng Lu, Hangyu Chen, Zunjie Zhu, Yuhang Qin, Ming Lu, Le Zhang, Chenggang Yan, Anke Xue</i></p>	98168

<p>WavTokenizer: An Efficient Acoustic Discrete Codec Tokenizer for Audio Language Modeling</p> <p>Shengpeng Ji, Ziyue Jiang, Wen Wang, Yifu Chen, Minghui Fang, Jialong Zuo, Qian Yang, Xize Cheng, Zehan Wang, Ruiqi Li, Ziang Zhang, Xiaoda Yang, Rongjie Huang, Yidi Jiang, Qian Chen, Siqi Zheng, Zhou Zhao</p>	98181
<p>GOttack: Universal Adversarial Attacks on Graph Neural Networks Via Graph Orbits Learning</p> <p>Zulfikar Alom, Tran Gia Bao Ngo, Murat Kantarcioglu, Cuneyt Gurcan Akcora</p>	98199
<p>Calibrating LLMs with Information-Theoretic Evidential Deep Learning</p> <p>Yawei Li, David Rügamer, Bernd Bischl, Mina Rezaei</p>	98213
<p>GPromptShield: Elevating Resilience in Graph Prompt Tuning Against Adversarial Attacks</p> <p>Shuhan Song, Ping Li, Ming Dun, Maolei Huang, Huawei Cao, Xiaochun Ye</p>	98240
<p>Multilevel Generative Samplers for Investigating Critical Phenomena</p> <p>Ankur Singha, Elia Cellini, Kim Andrea Nicoli, Karl Jansen, Stefan Kühn, Shinichi Nakajima</p>	98261
<p>On the Expressiveness and Spectral Bias of KANs</p> <p>Yixuan Wang, Jonathan W. Siegel, Ziming Liu, Thomas Y. Hou</p>	98295
<p>Federated Class-Incremental Learning: A Hybrid Approach Using Latent Exemplars and Data-Free Techniques to Address Local and Global Forgetting</p> <p>Milad Khademi Nori, Il Min Kim, Guanghui Wang</p>	98315
<p>Grammar Reinforcement Learning: Path and Cycle Counting in Graphs with a Context-Free Grammar and Transformer Approach</p> <p>Jason Piquenot, Maxime Berar, Romain Raveaux, Pierre Héroux, Jean-Yves Ramel, Sébastien Adam</p>	98329
<p>Dist Loss: Enhancing Regression in Few-Shot Region Through Distribution Distance Constraint</p> <p>Guangkun Nie, Gongzheng Tang, Shenda Hong</p>	98362
<p>UniCO: On Unified Combinatorial Optimization Via Problem Reduction to Matrix-Encoded General TSP</p> <p>Wenzheng Pan, Hao Xiong, Jiale Ma, Wentao Zhao, Yang Li, Junchi Yan</p>	98381
<p>MDSGen: Fast and Efficient Masked Diffusion Temporal-Aware Transformers for Open-Domain Sound Generation</p> <p>Trung X. Pham, Tri Ton, Chang D. Yoo</p>	98422
<p>Simple, Good, Fast: Self-Supervised World Models Free of Baggage</p> <p>Jan Robine, Marc Höftmann, Stefan Harmeling</p>	98448
<p>SORRY-Bench: Systematically Evaluating Large Language Model Safety Refusal</p> <p>Tinghao Xie, Xiangyu Qi, Yi Zeng, Yangsibo Huang, Udari Madhushani Sehwag, Kaixuan Huang, Luxi He, Boyi Wei, Dacheng Li, Ying Sheng, Ruoxi Jia, Bo Li, Kai Li, Danqi Chen, Peter Henderson, Prateek Mittal</p>	98469

<p>Adversarial Perturbations Cannot Reliably Protect Artists from Generative AI</p> <p>Robert Hönig, Javier Rando, Nicholas Carlini, Florian Tramèr</p>	98506
<p>PADRe: A Unifying Polynomial Attention Drop-In Replacement for Efficient Vision Transformer</p> <p>Pierre-David Letourneau, Manish Kumar Singh, Hsin-Pai Cheng, Shizhong Han, Yunxiao Shi, Dalton Jones, Matthew Harper Langston, Hong Cai, Fatih Porikli</p>	98547
<p>Mitigating Object Hallucination in MLLMs Via Data-Augmented Phrase-Level Alignment</p> <p>Pritam Sarkar, Sayna Ebrahimi, Ali Etemad, Ahmad Beirami, Sercan O Arik, Tomas Pfister</p>	98566
<p>Image and Video Tokenization with Binary Spherical Quantization</p> <p>Yue Zhao, Yuanjun Xiong, Philipp Kraehenbuehl</p>	98617
<p>BoneMet: An Open Large-Scale Multi-Modal Murine Dataset for Breast Cancer Bone Metastasis Diagnosis and Prognosis</p> <p>Tiankuo Chu, Fudong Lin, Shubo Wang, Jason Jiang, Wiley Jia-Wei Gong, Xu Yuan, Liyun Wang</p>	98642
<p>Effective and Efficient Time-Varying Counterfactual Prediction with State-Space Models</p> <p>Haotian Wang, Haoxuan Li, Hao Zou, Haoang Chi, Long Lan, Wanrong Huang, Wenjing Yang</p>	98663
<p>PABBO: Preferential Amortized Black-Box Optimization</p> <p>Xinyu Zhang, Daolang Huang, Samuel Kaski, Julien Martinelli</p>	98689
<p>Contextual Self-Paced Learning for Weakly Supervised Spatio-Temporal Video Grounding</p> <p>Akash Kumar, Zsolt Kira, Yogesh S Rawat</p>	98714
<p>A Unified Framework for Forward and Inverse Problems in Subsurface Imaging Using Latent Space Translations</p> <p>Naveen Gupta, Medha Sawhney, Arka Daw, Youzuo Lin, Anuj Karpatne</p>	98738
<p>Forewarned is Forearmed: Harnessing LLMs for Data Synthesis Via Failure-Induced Exploration</p> <p>Qintong Li, Jiahui Gao, Sheng Wang, Renjie Pi, Xueliang Zhao, Chuan Wu, Xin Jiang, Zhenguo Li, Lingpeng Kong</p>	98783
<p>UniWav: Towards Unified Pre-Training for Speech Representation Learning and Generation</p> <p>Alexander H. Liu, Sang-Gil Lee, Chao-Han Huck Yang, Yuan Gong, Yu-Chiang Frank Wang, James R. Glass, Rafael Valle, Bryan Catanzaro</p>	98805
<p>Brain Mapping with Dense Features: Grounding Cortical Semantic Selectivity in Natural Images with Vision Transformers</p> <p>Andrew Luo, Jacob Yeung, Rushikesh Zawar, Shaurya Rajat Dewan, Margaret Marie Henderson, Leila Wehbe, Michael J. Tarr</p>	98823
<p>Can Large Language Models Understand Symbolic Graphics Programs?</p> <p>Zeju Qiu, Weiyang Liu, Haiwen Feng, Zhen Liu, Tim Z. Xiao, Katherine M. Collins, Joshua B. Tenenbaum, Adrian Weller, Michael J. Black, Bernhard Schölkopf</p>	98874

<p>Learning Transformer-Based World Models with Contrastive Predictive Coding</p> <p>Maxime Burchi, Radu Timofte</p>	98921
<p>Generative Inbetweening: Adapting Image-To-Video Models for Keyframe Interpolation</p> <p>Xiaojuan Wang, Boyang Zhou, Brian Curless, Ira Kemelmacher-Shlizerman, Aleksander Holynski, Steve Seitz</p>	98943
<p>ELFS: Label-Free Coreset Selection with Proxy Training Dynamics</p> <p>Haizhong Zheng, Elisa Tsai, Yifu Lu, Jiachen Sun, Brian R. Bartoldson, Bhavya Kailkhura, Atul Prakash</p>	98959
<p>W-PCA Based Gradient-Free Proxy for Efficient Search of Lightweight Language Models</p> <p>Shang Wang</p>	98979
<p>BRIGHT: A Realistic and Challenging Benchmark for Reasoning-Intensive Retrieval</p> <p>Hongjin Su, Howard Yen, Mengzhou Xia, Weijia Shi, Niklas Muennighoff, Han-Yu Wang, Liu Haisu, Quan Shi, Zachary S Siegel, Michael Tang, Ruoxi Sun, Jinsung Yoon, Sercan O Arik, Danqi Chen, Tao Yu</p>	99001
<p>On the Price of Differential Privacy for Hierarchical Clustering</p> <p>Chengyuan Deng, Jie Gao, Jalaj Upadhyay, Chen Wang, Samson Zhou</p>	99052
<p>LaMPlace: Learning to Optimize Cross-Stage Metrics in Macro Placement</p> <p>Zijie Geng, Jie Wang, Ziyang Liu, Siyuan Xu, Zhentao Tang, Shixiong Kai, Mingxuan Yuan, Jianye Hao, Feng Wu</p>	99084
<p>Learning 3D Perception from Others' Predictions</p> <p>Jinsu Yoo, Zhenyang Feng, Tai-Yu Pan, Yihong Sun, Cheng Perng Phoo, Xiangyu Chen, Mark Campbell, Kilian Q Weinberger, Bharath Hariharan, Wei-Lun Chao</p>	99112
<p>Conformalized Interactive Imitation Learning: Handling Expert Shift and Intermittent Feedback</p> <p>Michelle D Zhao, Henny Admoni, Reid Simmons, Aaditya Ramdas, Andrea Bajcsy</p>	99133
<p>DenoiseVAE: Learning Molecule-Adaptive Noise Distributions for Denoising-Based 3D Molecular Pre-Training</p> <p>Yurou Liu, Jiahao Chen, Rui Jiao, Jiangmeng Li, Wenbing Huang, Bing Su</p>	99147
<p>HERO: Human-Feedback Efficient Reinforcement Learning for Online Diffusion Model Finetuning</p> <p>Ayano Hiranaka, Shang-Fu Chen, Chieh-Hsin Lai, Dongjun Kim, Naoki Murata, Takashi Shibuya, Wei-Hsiang Liao, Shao-Hua Sun, Yuki Mitsufuji</p>	99171
<p>AutoBencher: Towards Declarative Benchmark Construction</p> <p>Xiang Lisa Li, Farzaan Kaiyom, Evan Zheran Liu, Yifan Mai, Percy Liang, Tatsunori Hashimoto</p>	99201
<p>Progressive Parameter Efficient Transfer Learning for Semantic Segmentation</p> <p>Nan Zhou, Huiqun Wang, Yaoyan Zheng, Di Huang</p>	99231

<p>Handling Delay in Real-Time Reinforcement Learning</p> <p><i>Ivan Anokhin, Rishav Rishav, Matthew Riemer, Stephen Chung, Irina Rish, Samira Ebrahimi Kahou</i></p>	99247
<p>Understanding the Generalization of In-Context Learning in Transformers: An Empirical Study</p> <p><i>Xingxuan Zhang, Haoran Wang, Jiansheng Li, Yuan Xue, Shikai Guan, Renzhe Xu, Hao Zou, Han Yu, Peng Cui</i></p>	99271
<p>MeteoRA: Multiple-Tasks Embedded LoRA for Large Language Models</p> <p><i>Jingwei Xu, Junyu Lai, Yunpeng Huang</i></p>	99303
<p>Pedestrian Motion Reconstruction: A Large-Scale Benchmark Via Mixed Reality Rendering with Multiple Perspectives and Modalities</p> <p><i>Yichen Wang, Yiyi Zhang, Xinhao Hu, Li Niu, Jianfu Zhang, Yasushi Makihara, Yasushi Yagi, Pai Peng, Wenlong Liao, Tao He, Junchi Yan, Liqing Zhang</i></p>	99331
<p>ReCogLab: A Framework Testing Relational Reasoning & Cognitive Hypotheses on LLMs</p> <p><i>Andrew Liu, Henry Prior, Gargi Balasubramaniam, Rivka Moroshko, Amir Zait, Ilia Labzovsky, Danny Karmon, Ishita Dasgupta, Kim Stachenfeld, Kenneth Marino</i></p>	99346
<p>SiReRAG: Indexing Similar and Related Information for Multihop Reasoning</p> <p><i>Nan Zhang, Prafulla Kumar Choubey, Alexander Fabbri, Gabriel Bernadett-Shapiro, Rui Zhang, Prasenjit Mitra, Caiming Xiong, Chien-Sheng Wu</i></p>	99376
<p>Pacmann: Efficient Private Approximate Nearest Neighbor Search</p> <p><i>Mingxun Zhou, Elaine Shi, Giulia Fanti</i></p>	99396
<p>MarS: A Financial Market Simulation Engine Powered by Generative Foundation Model</p> <p><i>Junjie Li, Yang Liu, Weiqing Liu, Shikai Fang, Lewen Wang, Chang Xu, Jiang Bian</i></p>	99417
<p>Understanding and Enhancing Safety Mechanisms of LLMs Via Safety-Specific Neuron</p> <p><i>Yiran Zhao, Wenxuan Zhang, Yuxi Xie, Anirudh Goyal, Kenji Kawaguchi, Michael Shieh</i></p>	99452
<p>Causally Motivated Sycophancy Mitigation for Large Language Models</p> <p><i>Haoxi Li, Xueyang Tang, Jie Zhang, Song Guo, Sikai Bai, Peiran Dong, Yue Yu</i></p>	99467
<p>BigCodeBench: Benchmarking Code Generation with Diverse Function Calls and Complex Instructions</p> <p><i>Terry Yue Zhuo, Vu Minh Chien, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen Gong, James Hoang, Armel Randy Zebaze, Xiaoheng Hong, Wen-Ding Li, Jean Kaddour, Ming Xu, Zhihan Zhang, Prateek Yadav, Naman Jain, Alex Gu, Zhoujun Cheng, Jiawei Liu, Qian Liu, Zijian Wang, Binyuan Hui, Niklas Muennighoff, David Lo, Daniel Fried, Xiaoning Du, Harm De Vries, Leandro Von Werra</i></p>	99488
<p>CirT: Global Subseasonal-To-Seasonal Forecasting with Geometry-Inspired Transformer</p> <p><i>Yang Liu, Zinan Zheng, Jiashun Cheng, Fugee Tsung, Deli Zhao, Yu Rong, Jia Li</i></p>	99543
<p>Concept Bottleneck Language Models for Protein Design</p> <p><i>Aya Abdelsalam Ismail, Tuomas Oikarinen, Amy Wang, Julius Adebayo, Samuel Don Stanton, Hector Corrada Bravo, Kyunghyun Cho, Nathan C. Frey</i></p>	99565

<p>On the Crucial Role of Initialization for Matrix Factorization</p> <p>Bingcong Li, Liang Zhang, Aryan Mokhtari, Niao He</p>	99598
<p>Robots Pre-Train Robots: Manipulation-Centric Robotic Representation from Large-Scale Robot Datasets</p> <p>Guangqi Jiang, Yifei Sun, Tao Huang, Huanyu Li, Yongyuan Liang, Huazhe Xu</p>	99634
<p>LICO: Large Language Models for In-Context Molecular Optimization</p> <p>Tung Nguyen, Aditya Grover</p>	99655
<p>APE: Faster and Longer Context-Augmented Generation Via Adaptive Parallel Encoding</p> <p>Xinyu Yang, Tianqi Chen, Beidi Chen</p>	99677
<p>Quantitative Approximation for Neural Operators in Nonlinear Parabolic Equations</p> <p>Takashi Furuya, Koichi Taniguchi, Satoshi Okuda</p>	99698
<p>Interpretable Bilingual Multimodal Large Language Model for Diverse Biomedical Tasks</p> <p>Lehan Wang, Haonan Wang, Honglong Yang, Jiaji Mao, Zehong Yang, Jun Shen, Xiaomeng Li</p>	99727
<p>ReGenesis: LLMs Can Grow into Reasoning Generalists Via Self-Improvement</p> <p>Xiangyu Peng, Congying Xia, Xinyi Yang, Caiming Xiong, Chien-Sheng Wu, Chen Xing</p>	99760
<p>MaRS: A Fast Sampler for Mean Reverting Diffusion Based on ODE and SDE Solvers</p> <p>Ao Li, Wei Fang, Hongbo Zhao, Le Lu, Ge Yang, Minfeng Xu</p>	99796
<p>Disentangling Representations Through Multi-Task Learning</p> <p>Pantelis Vafidis, Aman Bhargava, Antonio Rangel</p>	99824
<p>Toward Efficient Multi-Agent Exploration with Trajectory Entropy Maximization</p> <p>Tianxu Li, Kun Zhu</p>	99867
<p>Robust Function-Calling for On-Device Language Model Via Function Masking</p> <p>Qiqiang Lin, Muning Wen, Qiuying Peng, Guanyu Nie, Junwei Liao, Jun Wang, Xiaoyun Mo, Jiamu Zhou, Cheng Cheng, Yin Zhao, Jun Wang, Weinan Zhang</p>	99887
<p>Revisiting Zeroth-Order Optimization: Minimum-Variance Two-Point Estimators and Directionally Aligned Perturbations</p> <p>Shaocong Ma, Heng Huang</p>	99908
<p>Nesterov Acceleration in Benignly Non-Convex Landscapes</p> <p>Kanan Gupta, Stephan Wojtowytsch</p>	99935
<p>Towards Understanding the Universality of Transformers for Next-Token Prediction</p> <p>Michael Eli Sander, Gabriel Peyré</p>	99968
<p>Elliptic Loss Regularization</p> <p>Ali Hasan, Haoming Yang, Yuting Ng, Vahid Tarokh</p>	99988

<p>Point-SAM: Promptable 3D Segmentation Model for Point Clouds</p> <p>Yuchen Zhou, Jiayuan Gu, Tung Yen Chiang, Fanbo Xiang, Hao Su</p>	100013
<p>COMBO: Compositional World Models for Embodied Multi-Agent Cooperation</p> <p>Hongxin Zhang, Zeyuan Wang, Qiushi Lyu, Zheyuan Zhang, Sunli Chen, Tianmin Shu, Behzad Dariush, Kwonjoon Lee, Yilun Du, Chuang Gan</p>	100033
<p>Permute-And-Flip: An Optimally Stable and Watermarkable Decoder for LLMs</p> <p>Xuandong Zhao, Lei Li, Yu-Xiang Wang</p>	100057
<p>Accelerating Diffusion Transformers with Token-Wise Feature Caching</p> <p>Chang Zou, Xuyang Liu, Ting Liu, Siteng Huang, Linfeng Zhang</p>	100086
<p>Generative Monoculture in Large Language Models</p> <p>Fan Wu, Emily Black, Varun Chandrasekaran</p>	100108
<p>Attention in Large Language Models Yields Efficient Zero-Shot Re-Rankers</p> <p>Shijie Chen, Bernal Jimenez Gutierrez, Yu Su</p>	100148
<p>Scaling Laws for Adversarial Attacks on Language Model Activations and Tokens</p> <p>Stanislav Fort</p>	100171
<p>FlexCAD: Unified and Versatile Controllable CAD Generation with Fine-Tuned Large Language Models</p> <p>Zhanwei Zhang, Shizhao Sun, Wenxiao Wang, Deng Cai, Jiang Bian</p>	100189
<p>Once-For-All: Controllable Generative Image Compression with Dynamic Granularity Adaptation</p> <p>Anqi Li, Feng Li, Yuxi Liu, Runmin Cong, Yao Zhao, Huihui Bai</p>	100213
<p>FLOPS: Forward Learning with OPTimal Sampling</p> <p>Tao Ren, Zishi Zhang, Jinyang Jiang, Guanghao Li, Zeliang Zhang, Mingqian Feng, Yijie Peng</p>	100234
<p>Integral Performance Approximation for Continuous-Time Reinforcement Learning Control</p> <p>Brent A. Wallace, Jennie Si</p>	100253
<p>GOAL: A Generalist Combinatorial Optimization Agent Learner</p> <p>Darko Drakulic, Sofia Michel, Jean-Marc Andreoli</p>	100295
<p>CameraCtrl: Enabling Camera Control for Video Diffusion Models</p> <p>Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, Ceyuan Yang</p>	100319
<p>AFlow: Automating Agentic Workflow Generation</p> <p>Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xiong-Hui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, Bingnan Zheng, Bang Liu, Yuyu Luo, Chenglin Wu</p>	100351
<p>Cross-Domain Off-Policy Evaluation and Learning for Contextual Bandits</p> <p>Yuta Natsubori, Masataka Ushiku, Yuta Saito</p>	100389

<p>Random is All You Need: Random Noise Injection on Feature Statistics for Generalizable Deep Image Denoising</p> <p>Zhengwei Yin, Hongjun Wang, Guixu Lin, Weihang Ran, Yinqiang Zheng</p>	100410
<p>NEAR: A Training-Free Pre-Estimator of Machine Learning Model Performance</p> <p>Raphael T. Husstein, Markus Reiher, Marco Eckhoff</p>	100434
<p>LOKI: A Comprehensive Synthetic Data Detection Benchmark Using Large Multimodal Models</p> <p>Junyan Ye, Baichuan Zhou, Zilong Huang, Junan Zhang, Tianyi Bai, Hengrui Kang, Jun He, Honglin Lin, Zihao Wang, Tong Wu, Zhizheng Wu, Yiping Chen, Dahua Lin, Conghui He, Weijia Li</p>	100455
<p>PhyloVAE: Unsupervised Learning of Phylogenetic Trees Via Variational Autoencoders</p> <p>Tianyu Xie, Harry Richman, Jiansi Gao, Frederick A Matsen, Cheng Zhang</p>	100538
<p>Deep Kernel Relative Test for Machine-Generated Text Detection</p> <p>Yiliao Song, Zhenqiao Yuan, Shuhai Zhang, Zhen Fang, Jun Yu, Feng Liu</p>	100560
<p>Actions Speak Louder than Words: Rate-Reward Trade-Off in Markov Decision Processes</p> <p>Haotian Wu, Gongpu Chen, Deniz Gunduz</p>	100587
<p>SpikeLLM: Scaling Up Spiking Neural Network to Large Language Models Via Saliency-Based Spiking</p> <p>Xingrun Xing, Boyan Gao, Zheng Liu, David A. Clifton, Shitao Xiao, Wanpeng Zhang, Li Du, Zheng Zhang, Guoqi Li, Jiajun Zhang</p>	100615
<p>EgoSim: Egocentric Exploration in Virtual Worlds with Multi-Modal Conditioning</p> <p>Wei Yu, Songheng Yin, Steve Easterbrook, Animesh Garg</p>	100635
<p>InterMask: 3D Human Interaction Generation Via Collaborative Masked Modeling</p> <p>Muhammad Gohar Javed, Chuan Guo, Li Cheng, Xingyu Li</p>	100650
<p>Joint Graph Rewiring and Feature Denoising Via Spectral Resonance</p> <p>Jonas Linkerhägner, Cheng Shi, Ivan Dokmanic</p>	100672
<p>SSOLE: Rethinking Orthogonal Low-Rank Embedding for Self-Supervised Learning</p> <p>Lun Huang, Qiang Qiu, Guillermo Sapiro</p>	100712
<p>ASTrA: Adversarial Self-Supervised Training with Adaptive-Attacks</p> <p>Prakash Chandra Chhipa, Gautam Vashishtha, Settur Jithamanyu, Rajkumar Saini, Mubarak Shah, Marcus Liwicki</p>	100735
<p>Semialgebraic Neural Networks: From Roots to Representations</p> <p>S David Mis, Matti Lassas, Maarten V. De Hoop</p>	100758
<p>Latent Bayesian Optimization Via Autoregressive Normalizing Flows</p> <p>Seunghun Lee, Jinyoung Park, Jaewon Chu, Minseo Yoon, Hyunwoo J. Kim</p>	100792

<p>Time-To-Event Pretraining for 3D Medical Imaging</p> <p>Zepeng Frazier Huo, Jason Alan Fries, Alejandro Lozano, Jeya Maria Jose Valanarasu, Ethan Steinberg, Louis Blankemeier, Akshay S Chaudhari, Curtis Langlotz, Nigam Shah</p>	100815
<p>Kinetix: Investigating the Training of General Agents Through Open-Ended Physics-Based Control Tasks</p> <p>Michael Matthews, Michael Beukman, Chris Lu, Jakob Nicolaus Foerster</p>	100852
<p>Do Stochastic, Feel Noiseless: Stable Stochastic Optimization Via a Double Momentum Mechanism</p> <p>Tehila Dahan, Kfir Yehuda Levy</p>	100902
<p>System 1.x: Learning to Balance Fast and Slow Planning with Language Models</p> <p>Swarnadeep Saha, Archiki Prasad, Justin Chen, Peter Hase, Elias Stengel-Eskin, Mohit Bansal</p>	100937
<p>Group Ligands Docking to Protein Pockets</p> <p>Jiaqi Guan, Jiahua Li, Xiangxin Zhou, Xingang Peng, Sheng Wang, Yunan Luo, Jian Peng, Jianzhu Ma</p>	100965
<p>NVS-Solver: Video Diffusion Model as Zero-Shot Novel View Synthesizer</p> <p>Meng You, Zhiyu Zhu, Hui Liu, Junhui Hou</p>	100983
<p>Differentiable Rule Induction from Raw Sequence Inputs</p> <p>Kun Gao, Katsumi Inoue, Yongzhi Cao, Hanpin Wang, Yang Feng</p>	101012
<p>Step-By-Step Reasoning for Math Problems Via Twisted Sequential Monte Carlo</p> <p>Shengyu Feng, Xiang Kong, Shuang Ma, Aonan Zhang, Dong Yin, Chong Wang, Ruoming Pang, Yiming Yang</p>	101026
<p>Improving Equivariant Networks with Probabilistic Symmetry Breaking</p> <p>Hannah Lawrence, Vasco Portilheiro, Yan Zhang, Sékou-Oumar Kaba</p>	101050
<p>Dataset Ownership Verification in Contrastive Pre-Trained Models</p> <p>Yuechen Xie, Jie Song, Mengqi Xue, Haofei Zhang, Xingen Wang, Bingde Hu, Genlang Chen, Mingli Song</p>	101078
<p>Add-It: Training-Free Object Insertion in Images with Pretrained Diffusion Models</p> <p>Yoad Tewel, Rinon Gal, Dvir Samuel, Yuval Atzmon, Lior Wolf, Gal Chechik</p>	101101
<p>KLay: Accelerating Arithmetic Circuits for Neurosymbolic AI</p> <p>Jaron Maene, Vincent Derkinderen, Pedro Zuidberg Dos Martires</p>	101124
<p>Deconstructing What Makes a Good Optimizer for Autoregressive Language Models</p> <p>Rosie Zhao, Depen Morwani, David Brandfonbrener, Nikhil Vyas, Sham M. Kakade</p>	101142
<p>Glad: A Streaming Scene Generator for Autonomous Driving</p> <p>Bin Xie, Yingfei Liu, Tiancai Wang, Jiale Cao, Xiangyu Zhang</p>	101163

<p>Context Clues: Evaluating Long Context Models for Clinical Prediction Tasks on EHR Data</p> <p>Michael Wornow, Suhana Bedi, Miguel Angel Fuentes Hernandez, Ethan Steinberg, Jason Alan Fries, Christopher Re, Sanmi Koyejo, Nigam Shah</p>	101181
<p>VisRAG: Vision-Based Retrieval-Augmented Generation on Multi-Modality Documents</p> <p>Shi Yu, Chaoyue Tang, Bokai Xu, Junbo Cui, Junhao Ran, Yukun Yan, Zhenghao Liu, Shuo Wang, Xu Han, Zhiyuan Liu, Maosong Sun</p>	101219
<p>Intervening Anchor Token: Decoding Strategy in Alleviating Hallucinations for MLLMs</p> <p>Feilong Tang, Zile Huang, Chengzhi Liu, Qiang Sun, Harry Yang, Ser-Nam Lim</p>	101244
<p>Exact Byte-Level Probabilities from Tokenized Language Models for FIM-Tasks and Model Ensembles</p> <p>Buu Phan, Brandon Amos, Itai Gat, Marton Havasi, Matthew J. Muckley, Karen Ullrich</p>	101276
<p>Min-K%+-: Improved Baseline for Pre-Training Data Detection from Large Language Models</p> <p>Jingyang Zhang, Jingwei Sun, Eric Yeats, Yang Ouyang, Martin Kuo, Jianyi Zhang, Hao Frank Yang, Hai Li</p>	101298
<p>InsightBench: Evaluating Business Analytics Agents Through Multi-Step Insight Generation</p> <p>Gaurav Sahu, Abhay Puri, Juan A. Rodriguez, Amirhossein Abaskohi, Mohammad Chegini, Alexandre Drouin, Perouz Taslakian, Valentina Zantedeschi, Alexandre Lacoste, David Vazquez, Nicolas Chapados, Christopher Pal, Sai Rajeswar, Issam H. Laradji</p>	101316
<p>A Skewness-Based Criterion for Addressing Heteroscedastic Noise in Causal Discovery</p> <p>Yingyu Lin, Yuxing Huang, Wenqin Liu, Haoran Deng, Ignavier Ng, Kun Zhang, Mingming Gong, Yian Ma, Biwei Huang</p>	101349
<p>Theory, Analysis, and Best Practices for Sigmoid Self-Attention</p> <p>Jason Ramapuram, Federico Danieli, Eeshan Gunesh Dhekane, Floris Weers, Dan Busbridge, Pierre Ablin, Tatiana Likhomanenko, Jagrit Digani, Zijin Gu, Amitis Shidani, Russell Webb</p>	101377
<p>Zero-Cost Proxy for Adversarial Robustness Evaluation</p> <p>Yuqi Feng, Yuwei Ou, Jiahao Fan, Yanan Sun</p>	101433
<p>Exact Community Recovery Under Side Information: Optimality of Spectral Algorithms</p> <p>Julia Gaudio, Nirmal Joshi</p>	101455
<p>Optimized Multi-Token Joint Decoding with Auxiliary Model for LLM Inference</p> <p>Zongyue Qin, Ziniu Hu, Zifan He, Neha Prakriya, Jason Cong, Yizhou Sun</p>	101501
<p>Federated Few-Shot Class-Incremental Learning</p> <p>Muhammad Anwar Ma'Sum, Mahardhika Pratama, Lin Liu, H Habibullah, Ryszard Kowalczyk</p>	101524
<p>Do You Keep an Eye on What I Ask? Mitigating Multimodal Hallucination Via Attention-Guided Ensemble Decoding</p> <p>Yeongjae Cho, Keonwoo Kim, Taebaek Hwang, Sungzoon Cho</p>	101554

<p>Lines of Thought in Large Language Models</p> <p>Raphaël Sarfati, Toni J. B. Liu, Nicolas Boulle, Christopher Earls</p>	101575
<p>Ultra-Sparse Memory Network</p> <p>Zihao Huang, Qiyang Min, Hongzhi Huang, Yutao Zeng, Defa Zhu, Ran Guo, Zhou Xun</p>	101595
<p>Simulating Training Dynamics to Reconstruct Training Data from Deep Neural Networks</p> <p>Hanling Tian, Yuhang Liu, Mingzhen He, Zhengbao He, Zhehao Huang, Ruikai Yang, Xiaolin Huang</p>	101614
<p>RocketEval: Efficient Automated LLM Evaluation Via Grading Checklist</p> <p>Tianjun Wei, Wei Wen, Ruizhi Qiao, Xing Sun, Jianghong Ma</p>	101641
<p>Needle in a Video Haystack: A Scalable Synthetic Evaluator for Video MLLMs</p> <p>Zijia Zhao, Haoyu Lu, Yuqi Huo, Yifan Du, Tongtian Yue, Longteng Guo, Bingning Wang, Weipeng Chen, Jing Liu</p>	101668
<p>Complexity Lower Bounds of Adaptive Gradient Algorithms for Non-Convex Stochastic Optimization Under Relaxed Smoothness</p> <p>Michael Crawshaw, Mingrui Liu</p>	101701
<p>Generalized Consistency Trajectory Models for Image Manipulation</p> <p>Beomsu Kim, Jaemin Kim, Jeongsol Kim, Jong Chul Ye</p>	101753
<p>TopoDiffusionNet: A Topology-Aware Diffusion Model</p> <p>Saumya Gupta, Dimitris Samaras, Chao Chen</p>	101776
<p>What Makes Large Language Models Reason in (Multi-Turn) Code Generation?</p> <p>Kunhao Zheng, Juliette Decugis, Jonas Gehring, Taco Cohen, Benjamin Negrevergne, Gabriel Synnaeve</p>	101791
<p>InstantPortrait: One-Step Portrait Editing Via Diffusion Multi-Objective Distillation</p> <p>Zhixin Lai, Keqiang Sun, Fu-Yun Wang, Dhritiman Sagar, Erli Ding</p>	101829
<p>Distilled Decoding 1: One-Step Sampling of Image Auto-Regressive Models with Flow Matching</p> <p>Enshu Liu, Xuefei Ning, Yu Wang, Zinan Lin</p>	101853
<p>LDAdam: Adaptive Optimization from Low-Dimensional Gradient Statistics</p> <p>Thomas Robert, Mher Safaryan, Ionut-Vlad Modoranu, Dan Alistarh</p>	101877
<p>Synergy and Diversity in CLIP: Enhancing Performance Through Adaptive Backbone Ensembling</p> <p>Cristian Rodriguez-Opazo, Ehsan Abbasnejad, Damien Teney, Hamed Damirchi, Edison Marrese-Taylor, Anton Van Den Hengel</p>	101914
<p>EvA: Erasing Spurious Correlations with Activations</p> <p>Qiyuan He, Kai Xu, Angela Yao</p>	101959

<p>Probabilistic Learning to Defer: Handling Missing Expert Annotations and Controlling Workload Distribution</p> <p><i>Cuong C. Nguyen, Thanh-Toan Do, Gustavo Carneiro</i></p>	101983
<p>MMTEB: Massive Multilingual Text Embedding Benchmark</p> <p><i>Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Márton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Sibli, Dominik Krzeminski, Genta Indra Winata, Saba Sturua, Saiteja Utpala, Mathieu Ciancone, Marion Schaeffer, Diganta Misra, Shreeya Dhakal, Jonathan Rystrom, Roman Solomatin, Ömer Veysel Çagatan, Akash Kundu, Martin Bernstorff, Shitao Xiao, Akshita Sukhlecha, Bhavish Pahwa, Rafal Poswiata, Kranthi Kiran Gv, Shawon Ashraf, Daniel Auras, Björn Plüster, Jan Philipp Harries, Loïc Magne, Isabelle Mohr, Dawei Zhu, Hippolyte Gisserot-Boukhlef, Tom Aarsen, Jan Kostkan, Konrad Wojtasik, Taemin Lee, Marek Suppa, Crystina Zhang, Roberta Rocca, Mohammed Hamdy, Andrianos Michail, John Yang, Manuel Faysse, Aleksei Vatolin, Nandan Thakur, Manan Dey, Dipam Vasani, Pranjal A Chitale, Simone Tedeschi, Nguyen Tai, Artem Snegirev, Mariya Hendriksen, Michael Günther, Mengzhou Xia, Weijia Shi, Xing Han Lü, Jordan Clive, Gayatri K, Maksimova Anna, Silvan Wehrli, Maria Tikhonova, Henil Shalin Panchal, Aleksandr Abramov, Malte Ostendorff, Zheng Liu, Simon Clematide, Lester James Validad Miranda, Alena Fenogenova, Guangyu Song, Ruqiya Bin Safi, Wen-Ding Li, Alessia Borghini, Federico Cassano, Lasse Hansen, Sara Hooker, Chenghao Xiao, Vaibhav Adlakha, Orion Weller, Siva Reddy, Niklas Muennighoff</i></p>	102004
<p>FreSh: Frequency Shifting for Accelerated Neural Representation Learning</p> <p><i>Adam Kania, Marko Mihajlovic, Sergey Prokudin, Jacek Tabor, Przemyslaw Spurek</i></p>	102061
<p>IgGM: A Generative Model for Functional Antibody and Nanobody Design</p> <p><i>Rubo Wang, Fandi Wu, Xingyu Gao, Jiaxiang Wu, Peilin Zhao, Jianhua Yao</i></p>	102085
<p>Understanding the Stability-Based Generalization of Personalized Federated Learning</p> <p><i>Yingqi Liu, Qinglun Li, Jie Tan, Yifan Shi, Li Shen, Xiaochun Cao</i></p>	102112
<p>Robust System Identification: Finite-Sample Guarantees and Connection to Regularization</p> <p><i>Hyuk Park, Grani A. Hanasusanto, Yingying Li</i></p>	102144
<p>SELF-EVOLVED REWARD LEARNING FOR LLMs</p> <p><i>Chenghua Huang, Zhizhen Fan, Lu Wang, Fangkai Yang, Pu Zhao, Zeqi Lin, Qingwei Lin, Dongmei Zhang, Saravan Rajmohan, Qi Zhang</i></p>	102167
<p>Unified Parameter-Efficient Unlearning for LLMs</p> <p><i>Chenlu Ding, Jiancan Wu, Yancheng Yuan, Jinda Lu, Kai Zhang, Alex Su, Xiang Wang, Xiangnan He</i></p>	102190
<p>Physics of Language Models: Part 2.2, How to Learn from Mistakes on Grade-School Math Problems</p> <p><i>Tian Ye, Zicheng Xu, Yuanzhi Li, Zeyuan Allen-Zhu</i></p>	102214
<p>Enhancing Clustered Federated Learning: Integration of Strategies and Improved Methodologies</p> <p><i>Yongxin Guo, Xiaoying Tang, Tao Lin</i></p>	102226
<p>Generalizing Reasoning Problems to Longer Lengths</p> <p><i>Changnan Xiao, Bing Liu</i></p>	102271

<p>Feature Averaging: An Implicit Bias of Gradient Descent Leading to Non-Robustness in Neural Networks</p> <p>Binghui Li, Zhixuan Pan, Kaifeng Lyu, Jian Li</p>	102297
<p>Fewer May Be Better: Enhancing Offline Reinforcement Learning with Reduced Dataset</p> <p>Yiqin Yang, Quanwei Wang, Chenghao Li, Hao Hu, Chengjie Wu, Yuhua Jiang, Dianyu Zhong, Ziyou Zhang, Qianchuan Zhao, Chongjie Zhang, Bo Xu</p>	102369
<p>Let the Code LLM Edit Itself When You Edit the Code</p> <p>Zhenyu He, Jun Zhang, Shengjie Luo, Jingjing Xu, Zhi Zhang, Di He</p>	102392
<p>Concept Pinpoint Eraser for Text-To-Image Diffusion Models Via Residual Attention Gate</p> <p>Byung Hyun Lee, Sungjin Lim, Seunggyu Lee, Dong Un Kang, Se Young Chun</p>	102409
<p>Oracle Efficient Truncated Statistics</p> <p>Konstantinos Karatapanis, Vasilis Kontonis, Christos Tzamos</p>	102449
<p>Self-Improving Robust Preference Optimization</p> <p>Eugene Choi, Arash Ahmadian, Matthieu Geist, Olivier Pietquin, Mohammad Gheshlaghi Azar</p>	102470
<p>To Code Or Not to Code? Exploring Impact of Code in Pre-Training</p> <p>Viraat Aryabumi, Yixuan Su, Raymond Ma, Adrien Morisot, Ivan Zhang, Acyr Locatelli, Marzieh Fadaee, Ahmet Üstün, Sara Hooker</p>	102492
<p>Enhancing Language Model Agents Using Diversity of Thoughts</p> <p>Vijay Lingam, Behrooz Omidvar Tehrani, Sujay Sanghavi, Gaurav Gupta, Sayan Ghosh, Linbo Liu, Jun Huan, Anoop Deoras</p>	102519
<p>PvNeXt: Rethinking Network Design and Temporal Motion for Point Cloud Video Recognition</p> <p>Jie Wang, Tingfa Xu, Lihe Ding, Xinjie Zhang, Long Bai, Jianan Li</p>	102538
<p>Radar: Fast Long-Context Decoding for Any Transformer</p> <p>Yongchang Hao, Mengyao Zhai, Hossein Hajimirsadeghi, Sepidehsadat Hosseini, Frederick Tung</p>	102555
<p>ReDeEP: Detecting Hallucination in Retrieval-Augmented Generation Via Mechanistic Interpretability</p> <p>Zhongxiang Sun, Xiaoxue Zang, Kai Zheng, Jun Xu, Xiao Zhang, Weijie Yu, Yang Song, Han Li</p>	102578
<p>MVTokenFlow: High-Quality 4D Content Generation Using Multiview Token Flow</p> <p>Hanzhuo Huang, Yuan Liu, Ge Zheng, Jiepeng Wang, Zhiyang Dou, Sibe Yang</p>	102608
<p>ARB-LLM: Alternating Refined Binarizations for Large Language Models</p> <p>Zhiteng Li, Xianglong Yan, Tianao Zhang, Haotong Qin, Dong Xie, Jiang Tian, Zhongchao Shi, Linghe Kong, Yulun Zhang, Xiaokang Yang</p>	102626
<p>Problem-Parameter-Free Federated Learning</p> <p>Wenjing Yan, Kai Zhang, Xiaolu Wang, Xuanyu Cao</p>	102639

<p><i>Fine-Tuning with Reserved Majority for Noise Reduction</i></p> <p><i>Shuyang Jiang, Yusheng Liao, Ya Zhang, Yanfeng Wang, Yu Wang</i></p>	102667
<p><i>Designing Concise ConvNets with Columnar Stages</i></p> <p><i>Ashish Kumar, Jaesik Park</i></p>	102687
<p><i>A Policy-Gradient Approach to Solving Imperfect-Information Games with Best-Iterate Convergence</i></p> <p><i>Mingyang Liu, Gabriele Farina, Asuman E. Ozdaglar</i></p>	102705
<p><i>From Lazy to Rich: Exact Learning Dynamics in Deep Linear Networks</i></p> <p><i>Clémentine Carla Juliette Dominé, Nicolas Anguita, Alexandra Maria Proca, Lukas Braun, Daniel Kunin, Pedro A. M. Mediano, Andrew M Saxe</i></p>	102747
<p><i>Shared-AE: Automatic Identification of Shared Subspaces in High-Dimensional Neural and Behavioral Activity</i></p> <p><i>Daiyao Yi, Hao Dong, Michael James Higley, Anne Churchland, Shreya Saxena</i></p>	102799
<p><i>Re-Imagining Multimodal Instruction Tuning: A Representation View</i></p> <p><i>Yiyang Liu, James Chenhao Liang, Ruixiang Tang, Yugyung Lee, Majid Rabbani, Sohail Dianat, Raghuvveer Rao, Lifu Huang, Dongfang Liu, Qifan Wang, Cheng Han</i></p>	102827
<p><i>Inverse Decision-Making Using Neural Amortized Bayesian Actors</i></p> <p><i>Dominik Straub, Tobias F. Niehues, Jan Peters, Constantin A. Rothkopf</i></p>	102851
<p><i>TLDR: Token-Level Detective Reward Model for Large Vision Language Models</i></p> <p><i>Deqing Fu, Tong Xiao, Rui Wang, Wang Zhu, Pengchuan Zhang, Guan Pang, Robin Jia, Lawrence Chen</i></p>	102873
<p><i>Semantic Temporal Abstraction Via Vision-Language Model Guidance for Efficient Reinforcement Learning</i></p> <p><i>Tian-Shuo Liu, Xu-Hui Liu, Ruifeng Chen, Lixuan Jin, Pengyuan Wang, Zhilong Zhang, Yang Yu</i></p>	102905
<p><i>Diff-2-In-1: Bridging Generation and Dense Perception with Diffusion Models</i></p> <p><i>Shuhong Zheng, Zhipeng Bao, Ruoyu Zhao, Martial Hebert, Yu-Xiong Wang</i></p>	102933
<p><i>Truncated Consistency Models</i></p> <p><i>Sangyun Lee, Yilun Xu, Tomas Geffner, Giulia Fanti, Karsten Kreis, Arash Vahdat, Weili Nie</i></p>	102961
<p><i>Rethinking Spiking Neural Networks from an Ensemble Learning Perspective</i></p> <p><i>Yongqi Ding, Lin Zuo, Mengmeng Jing, Pei He, Hanpu Deng</i></p>	102985
<p><i>DGQ: Distribution-Aware Group Quantization for Text-To-Image Diffusion Models</i></p> <p><i>Hyogon Ryu, Nahyeon Park, Hyunjung Shim</i></p>	103012
<p><i>Synthesizing Realistic fMRI: A Physiological Dynamics-Driven Hierarchical Diffusion Model for Efficient fMRI Acquisition</i></p> <p><i>Yufan Hu, Yujiang, Wuyang Li, Yixuan Yuan</i></p>	103035

<i>H4rm3l: A Language for Composable Jailbreak Attack Synthesis</i> <i>Moussa Koulako Bala Doumbouya, Ananjan Nandi, Gabriel Poesia, Davide Ghilardi, Anna Goldie, Federico Bianchi, Dan Jurafsky, Christopher D Manning</i>	103049
<i>One-For-All Few-Shot Anomaly Detection Via Instance-Induced Prompt Learning</i> <i>Wenxi Lv, Qinliang Su, Wenchao Xu</i>	103083