

International Conference on Machine Learning

ICML 2025

13-19 July 2025, Vancouver, British Columbia, Canada

ISBN: 979-8-3313-3135-1

Title and Authors	Page#
Semi-Supervised Blind Quality Assessment with Confidence-quantifiable Pseudo-label Learning for Authentic Images <i>Yan Zhong, Chenxi Yang, Suyuan Zhao, Tingting Jiang</i>	1
MindCustomer: Multi-Context Image Generation Blended with Brain Signal <i>Muzhou Yu, Shuyun Lin, Lei Ma, Bo Lei, Kaisheng Ma</i>	19
MixBridge: Heterogeneous Image-to-Image Backdoor Attack through Mixture of Schrödinger Bridges <i>Shixi Qin, Zhiyong Yang, Shilong Bao, Shi Wang, Qianqian Xu, Qingming Huang</i>	43
Conditioning Diffusions Using Malliavin Calculus <i>Jakiw Pidstrigach, Elizabeth Louise Baker, Carles Domingo-Enrich, George Deligiannidis, Nikolas Nüsken</i>	81
Leveraging Per-Instance Privacy for Machine Unlearning <i>Nazanin Mohammadi Sepahvand, Anvith Thudi, Berivan Isik, Ashmita Bhattacharyya, Nicolas Papernot, Eleni Triantafillou, Daniel M. Roy, Gintare Karolina Dziugaite</i>	105
TransPL: VQ-Code Transition Matrices for Pseudo-Labeling of Time Series Unsupervised Domain Adaptation <i>Jaeho Kim, Seulki Lee</i>	122
COKE: Core Kernel for More Efficient Approximation of Kernel Weights in Multiple Kernel Clustering <i>Weixuan Liang, Xinwang Liu, KE LIANG, Jiyuan Liu, En Zhu</i>	140
Analytical Construction on Geometric Architectures: Transitioning from Static to Temporal Link Prediction <i>Yadong Sun, Xiaofeng Cao, Ivor Tsang, Heng Tao Shen</i>	164
Test-time Correlation Alignment <i>Linjing You, Jiabao Lu, Xiayuan Huang</i>	183
Unbiased Recommender Learning from Implicit Feedback via Weakly Supervised Learning <i>Hao Wang, Zhichao Chen, Haotian Wang, Yanchao Tan, Licheng Pan, Tianqiao Liu, Xu Chen, Haoxuan Li, Zhouchen Lin</i>	213

MoEQuant: Enhancing Quantization for Mixture-of-Experts Large Language Models via Expert-Balanced Sampling and Affinity Guidance <i>Zhixuan Chen, Xing Hu, Dawei Yang, Zukang Xu, XUCHEN, Zhihang Yuan, Sifan Zhou, JiangyongYu</i>	234
Think Smarter not Harder: Adaptive Reasoning with Inference Aware Optimization <i>Zishun Yu, Tengyu Xu, Di Jin, Karthik Abinav Sankararaman, Yun He, Wenxuan Zhou, Zhouhao Zeng, Eryk Helenowski, Chen Zhu, Sinong Wang, Hao Ma, Han Fang</i>	250
Kernel-based Unsupervised Embedding Alignment for Enhanced Visual Representation in Vision-language Models <i>Shizhan Gong, Yankai Jiang, Qi Dou, Farzan Farnia</i>	272
Exploring Vision Semantic Prompt for Efficient Point Cloud Understanding <i>Yixin Zha, Chuxin Wang, Wenfei Yang, Tianzhu Zhang, Feng Wu</i>	292
SDE Matching: Scalable and Simulation-Free Training of Latent Stochastic Differential Equations <i>Grigory Bartosh, Dmitry Vetrov, Christian A. Naesseth</i>	309
Linear convergence of Sinkhorn’s algorithm for generalized static Schrödinger bridge <i>Rahul Choudhary, Hanbaek Lyu</i>	326
Positional Attention: Expressivity and Learnability of Algorithmic Computation <i>Artur Back de Luca, George Giapitzakis, Shenghao Yang, Petar Veličković, Kimon Fountoulakis</i>	353
MITIGATING OVER-EXPLORATION IN LATENT SPACE OPTIMIZATION USING LES <i>Omer Ronen, Ahmed Imtiaz Humayun, Richard Baraniuk, Randall Balestriero, Bin Yu</i>	409
Understanding Nonlinear Implicit Bias via Region Counts in Input Space <i>Jingwei Li, Jing Xu, Zifan Wang, Huishuai Zhang, Jingzhao Zhang</i>	438
Improving the Scaling Laws of Synthetic Data with Deliberate Practice <i>Reyhane Askari-Hemmat, Mohammad Pezeshki, Elvis Dohmatob, Florian Bordes, Pietro Astolfi, Melissa Hall, Jakob Verbeek, Michal Drozdal, Adriana Romero-Soriano</i>	455
sciLaMA: A Single-Cell Representation Learning Framework to Leverage Prior Knowledge from Large Language Models <i>Hongru Hu, Shuwen Zhang, Yongin Choi, Venkat S. Malladi, Gerald Quon</i>	486
CPCF: A Cross-Prompt Contrastive Framework for Referring Multimodal Large Language Models <i>Lanyun Zhu, Deyi Ji, Tianrun Chen, Haiyang Wu, De Wen Soh, Jun Liu</i>	505
Gap-Dependent Bounds for Federated Q-Learning <i>Haochen Zhang, Zhong Zheng, Lingzhou Xue</i>	519
Gumiho: A Hybrid Architecture to Prioritize Early Tokens in Speculative Decoding <i>Jinze Li, Yixing Xu, Haiduo Huang, Xuanwu Yin, Dong Li, Edith C. H. Ngai, Emad Barsoum</i>	564

NoLiMa: Long-Context Evaluation Beyond Literal Matching <i>Ali Modarressi, Hanieh Deilamsalehy, Franck Deroncourt, Trung Bui, Ryan A. Rossi, Seunghyun Yoon, Hinrich Schuetze</i>	589
Learning multivariate Gaussians with imperfect advice <i>Arnab Bhattacharyya, Davin Choo, Philips George John, Themis Gouleakis</i>	606
Widening the Network Mitigates the Impact of Data Heterogeneity on FedAvg <i>Like Jian, Dong Liu</i>	649
Density Ratio Estimation-based Bayesian Optimization with Semi-Supervised Learning <i>Jungtaek Kim</i>	676
Federated Disentangled Tuning with Textual Prior Decoupling and Visual Dynamic Adaptation <i>Yihao Yang, Wenke Huang, Guancheng Wan, Bin Yang, Mang Ye</i>	699
Enhancing Graph Invariant Learning from a Negative Inference Perspective <i>Kuo Yang, Zhengyang Zhou, Qihe Huang, Wenjie Du, Limin Li, Wu Jiang, Yang Wang</i>	710
Overcoming Spurious Solutions in Semi-Dual Neural Optimal Transport: A Smoothing Approach for Learning the Optimal Transport Plan <i>Jaemoo Choi, Jaewoong Choi, Dohyun Kwon</i>	729
OOD-Chameleon: Is Algorithm Selection for OOD Generalization Learnable? <i>Liangze Jiang, Damien Teney</i>	751
Raising the Bar: Investigating the Values of Large Language Models via Generative Evolving Testing <i>Han Jiang, Xiaoyuan Yi, Zhihua Wei, Ziang Xiao, Shu Wang, Xing Xie</i>	776
Improved Off-policy Reinforcement Learning in Biological Sequence Design <i>Hyeonah Kim, Minsu Kim, Taeyoung Yun, Sanghyeok Choi, Emmanuel Bengio, Alex Hernández-García, Jinkyoo Park</i>	824
Complete-Tree Space Favors Data-Efficient Link Prediction <i>Chi Gao, Lukai Li, Yancheng Zhou, Shangqi Guo</i>	850
DOLPHIN: A Programmable Framework for Scalable Neurosymbolic Learning <i>Aaditya Naik, Jason Liu, Claire Wang, Amish Sethi, Saikat Dutta, Mayur Naik, Eric Wong</i>	870
TabICL: A Tabular Foundation Model for In-Context Learning on Large Data <i>Jingang QU, David Holzmüller, Gaël Varoquaux, Marine Le Morvan</i>	893
Socialized Coevolution: Advancing a Better World through Cross-Task Collaboration <i>Xinjie Yao, Yu Wang, Pengfei Zhu, Wanyu Lin, Ruipu Zhao, Zhoupeng Guo, Weihao Li, Qinghua Hu</i>	924
MimicMotion: High-Quality Human Motion Video Generation with Confidence-aware Pose Guidance <i>Yuang Zhang, Jiayi Gu, Li-Wen Wang, Han Wang, JunqiCheng, Yuefeng Zhu, FangYuan Zou</i>	942

Joint Metric Space Embedding by Unbalanced Optimal Transport with Gromov–Wasserstein Marginal Penalization <i>Florian Beier, Moritz Piening, Robert Beinert, Gabriele Steidl</i>	957
Beyond Self-Repellent Kernels: History-Driven Target Towards Efficient Nonlinear MCMC on General Graphs <i>Jie Hu, Yi-Ting Ma, Do Young Eun</i>	975
DLP: Dynamic Layerwise Pruning in Large Language Models <i>Yuli Chen, Bo Cheng, Jiale Han, Yingying Zhang, Yingting Li, Shuhao Zhang</i>	1005
Revisiting Differentially Private Algorithms for Decentralized Online Learning <i>Xiaoyu Wang, Wenhao Yang, Chang Yao, Mingli Song, Yuanyu Wan</i>	1028
Optimal Task Order for Continual Learning of Multiple Tasks <i>Ziyan Li, Naoki Hiratani</i>	1051
Attention Mechanisms Perspective: Exploring LLM Processing of Graph-Structured Data <i>Zhong Guan, Likang Wu, Hongke Zhao, Ming He, Jianping Fan</i>	1077
Data Mixing Optimization for Supervised Fine-Tuning of Large Language Models <i>Yuan Li, Zhengzhong Liu, Eric P. Xing</i>	1105
Navigating the Social Welfare Frontier: Portfolios for Multi-objective Reinforcement Learning <i>Cheol Woo Kim, Jai Moondra, Shresth Verma, Madeleine Pollack, Ling kai Kong, Milind Tambe, Swati Gupta</i>	1124
La RoSA: Enhancing LLM Efficiency via Layerwise Rotated Sparse Activation <i>Kai Liu, Bowen Xu, Shaoyu Wu, Xin Chen, Hao Zhou, Yongliang Tao, lulu hu</i>	1147
Mitigating Local Cohesion and Global Sparseness in Graph Contrastive Learning with Fuzzy Boundaries <i>Yuena Lin, Haichun Cai, Jun-Yi Hang, Haobo Wang, Zhen Yang, Gengyu Lyu</i>	1166
ATA: Adaptive Task Allocation for Efficient Resource Management in Distributed Machine Learning <i>Arto Maranjyan, El Mehdi Saad, Peter Richtárik, Francesco Orabona</i>	1183
Gradient Aligned Regression via Pairwise Losses <i>Dixian Zhu, Tianbao Yang, Livnat Jerby</i>	1214
Scaling Video-Language Models to 10K Frames via Hierarchical Differential Distillation <i>Chuanqi Cheng, Jian Guan, Wei Wu, Rui Yan</i>	1237
Balancing Model Efficiency and Performance: Adaptive Pruner for Long-tailed Data <i>Zhe Zhao, HaiBin Wen, Pengkun Wang, ShuangWang, Zhenkun Wang, Qingfu Zhang, Yang Wang</i>	1255

Understanding High-Dimensional Bayesian Optimization <i>Leonard Papenmeier, Matthias Poloczek, Luigi Nardi</i>	1272
Towards a Unified Framework of Clustering-based Anomaly Detection <i>Zeyu Fang, Ming Gu, Sheng Zhou, Jiawei Chen, Qiaoyu Tan, Haishuai Wang, Jiajun Bu</i>	1294
ArrayDPS: Unsupervised Blind Speech Separation with a Diffusion Prior <i>Zhongweiyang Xu, Xulin Fan, Zhong-Qiu Wang, Xilin Jiang, Romit Roy Choudhury</i>	1311
Rethinking Score Distilling Sampling for 3D Editing and Generation <i>Xingyu Miao, Haoran Duan, Yang Long, Jungong Han</i>	1340
Do Not Mimic My Voice : Speaker Identity Unlearning for Zero-Shot Text-to-Speech <i>Taesoo Kim, Jinju Kim, Dong Chan Kim, Jong Hwan Ko, Gyeong-Moon Park</i>	1357
Learning Configurations for Data-Driven Multi-Objective Optimization <i>Zhiyang Chen, Hailong Yao, Xia Yin</i>	1380
Wolfpack Adversarial Attack for Robust Multi-Agent Reinforcement Learning <i>Sunwoo Lee, Jaebak Hwang, Yonghyeon Jo, Seungyul Han</i>	1402
A Likelihood Based Approach to Distribution Regression Using Conditional Deep Generative Models <i>Shivam Kumar, Yun Yang, Lizhen Lin</i>	1434
Test-Time Multimodal Backdoor Detection by Contrastive Prompting <i>Yuwei Niu, Shuo He, Qi Wei, Zongyu Wu, Feng Liu, Lei Feng</i>	1461
Mutual Learning for SAM Adaptation: A Dual Collaborative Network Framework for Source-Free Domain Transfer <i>Yabo Liu, Waikeng Wong, Chengliang Liu, Xiaoling Luo, Yong Xu, Jinghua Wang</i>	1481
Discriminative Finetuning of Generative Large Language Models without Reward Models and Human Preference Data <i>Siqi Guo, Ilgee Hong, Vicente Balmaseda, Changlong Yu, Liang Qiu, Xin Liu, Haoming Jiang, Tuo Zhao, Tianbao Yang</i>	1493
Reinforced Lifelong Editing for Language Models <i>Zherui Li, Houcheng Jiang, Hao Chen, Baolong Bi, Zhenhong Zhou, Fei Sun, Junfeng Fang, Xiang Wang</i>	1511
Provable Zero-Shot Generalization in Offline Reinforcement Learning <i>Zhiyong Wang, Chen Yang, John C.S. Lui, Dongruo Zhou</i>	1534
End-to-End Learning Framework for Solving Non-Markovian Optimal Control <i>Xiaole Zhang, Peiyu Zhang, Xiongye Xiao, Shixuan Li, Vasileios Tzoumas, Vijay Gupta, Paul Bogdan</i>	1556
Actor-Critics Can Achieve Optimal Sample Efficiency <i>Kevin Tan, Wei Fan, Yuting Wei</i>	1588

Causal Logistic Bandits with Counterfactual Fairness Constraints <i>Jiajun Chen, Jin Tian, Christopher John Quinn</i>	1634
Distillation Scaling Laws <i>Dan Busbridge, Amitis Shidani, Floris Weers, Jason Ramapuram, Etai Littwin, Russell Webb</i>	1664
E-LDA: Toward Interpretable LDA Topic Models with Strong Guarantees in Logarithmic Parallel Time <i>Adam Breuer</i>	1733
How Compositional Generalization and Creativity Improve as Diffusion Models are Trained <i>Alessandro Favero, Antonio Sclocchi, Francesco Cagnetta, Pascal Frossard, Matthieu Wyart</i>	1760
Liger: Linearizing Large Language Models to Gated Recurrent Structures <i>Disen Lan, Weigao Sun, Jiaxi Hu, Jusen Du, Yu Cheng</i>	1781
PEINR: A Physics-enhanced Implicit Neural Representation for High-Fidelity Flow Field Reconstruction <i>Liming Shen, Liang Deng, Chongke Bi, Yu Wang, Xinhai Chen, Yueqing Wang, Jie Liu</i>	1796
Learning to Reuse Policies in State Evolvable Environments <i>Ziqian Zhang, Bohan Yang, Lihe Li, Yuqi Bian, Ruiqi Xue, Feng Chen, Yi-Chen Li, Lei Yuan, Yang Yu</i>	1813
Optimistic Algorithms for Adaptive Estimation of the Average Treatment Effect <i>Ojash Neopane, Aaditya Ramdas, Aarti Singh</i>	1839
Massive Values in Self-Attention Modules are the Key to Contextual Knowledge Understanding <i>Mingyu Jin, Kai Mei, Wujiang Xu, Mingjie Sun, Ruixiang Tang, Mengnan Du, Zirui Liu, Yongfeng Zhang</i>	1855
Iterative Vectors: In-Context Gradient Steering without Backpropagation <i>Yiting Liu, Zhi-Hong Deng</i>	1889
Kona: An Efficient Privacy-Preservation Framework for KNN Classification by Communication Optimization <i>Guopeng Lin, Ruisheng Zhou, Shuyu Chen, Weili Han, Jin Tan, Wenjing Fang, Lei Wang, Tao Wei</i>	1912
MindAligner: Explicit Brain Functional Alignment for Cross-Subject Visual Decoding from Limited fMRI Data <i>Yiqin Dai, Zhouheng Yao, Chunfeng Song, Qihao Zheng, Weijian Mai, Kunyu Peng, Shuai Lu, Wanli Ouyang, Jian Yang, Jiamin Wu</i>	1929
Optimal Auction Design in the Joint Advertising <i>Yang Li, Yuchao Ma, Qi Qi</i>	1944
Positional Encoding meets Persistent Homology on Graphs <i>Yogesh Verma, Amauri H Souza, Vikas K Garg</i>	1961

Understanding Overadaptation in Supervised Fine-Tuning: The Role of Ensemble Methods <i>Yifan HAO, Xingyuan Pan, Hanning Zhang, Chenlu Ye, Rui Pan, Tong Zhang</i>	1982
Reinforcement Learning with Segment Feedback <i>Yihan Du, Anna Winnicki, Gal Dalal, Shie Mannor, R. Srikant</i>	2015
Hierarchical Reinforcement Learning with Uncertainty-Guided Diffusional Subgoals <i>Vivienne Huiling Wang, Tinghuai Wang, Joni Pajarinen</i>	2065
Understanding Fixed Predictions via Confined Regions <i>Connor Lawless, Tsui-Wei Weng, Berk Ustun, Madeleine Udell</i>	2088
PoisonBench: Assessing Language Model Vulnerability to Poisoned Preference Data <i>Tingchen Fu, Mrinank Sharma, Philip Torr, Shay B Cohen, David Krueger, Fazl Barez</i>	2114
Zebra: In-Context Generative Pretraining for Solving Parametric PDEs <i>Louis Serrano, Armand Kassai Koupai, Thomas X Wang, Pierre ERBACHER, Patrick Gallinari</i>	2140
Doubly Robust Fusion of Many Treatments for Policy Learning <i>Ke Zhu, Jianing Chu, Ilya Lipkovich, Wenyu Ye, Shu Yang</i>	2189
Understanding Bias Reinforcement in LLM Agents Debate <i>Jihwan Oh, Minchan Jeong, Jongwoo Ko, Se-Young Yun</i>	2207
SPD: Sync-Point Drop for Efficient Tensor Parallelism of Large Language Models <i>Han-Byul Kim, Duc N.M Hoang, Arnav Kundu, Mohammad Samragh, Minsik Cho</i>	2239
The Value of Prediction in Identifying the Worst-Off <i>Unai Fischer-Abaigar, Christoph Kern, Juan Carlos Perdomo</i>	2252
SANA 1.5: Efficient Scaling of Training-Time and Inference-Time Compute in Linear Diffusion Transformer <i>Enze Xie, Junsong Chen, Yuyang Zhao, Jincheng YU, Ligeng Zhu, Yujun Lin, Zhekai Zhang, Muyang Li, Junyu Chen, Han Cai, Bingchen Liu, Daquan Zhou, Song Han</i>	2275
A Geometric Approach to Personalized Recommendation with Set-Theoretic Constraints Using Box Embeddings <i>Shib Sankar Dasgupta, Michael Boratko, Andrew McCallum</i>	2296
AKRMap: Adaptive Kernel Regression for Trustworthy Visualization of Cross-Modal Embeddings <i>Yilin Ye, Junchao Huang, Xingchen ZENG, Jiazhi Xia, Wei Zeng</i>	2311
SITCOM: Step-wise Triple-Consistent Diffusion Sampling For Inverse Problems <i>Ismail Alkhouri, Shijun Liang, Cheng-Han Huang, Jimmy Dai, Qing Qu, Saiprasad Ravishankar, Rongrong Wang</i>	2330
Polynomial Time Learning Augmented Algorithms for NP-hard Permutation Problems <i>Evrpidis Bampis, Bruno Escoffier, Dimitris Fotakis, Panagiotis Patsilinakos, Michalis Xeferis</i>	2361

KoopSTD: Reliable Similarity Analysis between Dynamical Systems via Approximating Koopman Spectrum with Timescale Decoupling <i>Shimin Zhang, Ziyuan Ye, Yinsong Yan, Zeyang Song, Yujie Wu, Jibin Wu</i>	2372
Commute Graph Neural Networks <i>Wei Zhuo, Han Yu, Guang Tan, Xiaoxiao Li</i>	2392
Going Deeper into Locally Differentially Private Graph Neural Networks <i>Longzhu He, Chaozhuo Li, Peng Tang, Sen Su</i>	2409
Looking Beyond the Top-1: Transformers Determine Top Tokens in Order <i>Daria Lioubashevski, Tomer M. Schlank, Gabriel Stanovsky, Ariel Goldstein</i>	2427
NMA-tune: Generating Highly Designable and Dynamics Aware Protein Backbones <i>Urszula Julia Komorowska, Francisco Vargas, Alessandro Rondina, Pietro Lio, Mateja Jamnik</i>	2447
Score-Based Diffusion Policy Compatible with Reinforcement Learning via Optimal Transport <i>Mingyang Sun, Pengxiang Ding, Weinan Zhang, Donglin Wang</i>	2470
On Measuring Long-Range Interactions in Graph Neural Networks <i>Jacob Bamberger, Benjamin Gutteridge, Scott le Roux, Michael M. Bronstein, Xiaowen Dong</i>	2487
Pre-training Auto-regressive Robotic Models with 4D Representations <i>Dantong Niu, Yuvan Sharma, Haoru Xue, Giscard Biamby, Junyi Zhang, Ziteng Ji, Trevor Darrell, Roei Herzig</i>	2507
Improving Reward Model Generalization from Adversarial Process Enhanced Preferences <i>Zhilong Zhang, Tian Xu, Xinghao Du, Xingchen Cao, Yihao Sun, Yang Yu</i>	2526
Latent Action Learning Requires Supervision in the Presence of Distractors <i>Alexander Nikulin, Ilya Zisman, Denis Tarasov, Lyubaykin Nikita, Andrei Polubarov, Igor Kiselev, Vladislav Kurenkov</i>	2548
Supervised Contrastive Learning from Weakly-Labeled Audio Segments for Musical Version Matching <i>Joan Serrà, R. Oguz Araz, Dmitry Bogdanov, Yuki Mitsufuji</i>	2569
The Berkeley Function Calling Leaderboard (BFCL): From Tool Use to Agentic Evaluation of Large Language Models <i>Shishir G Patil, Huanzhi Mao, Fanjia Yan, Charlie Cheng-Jie Ji, Vishnu Suresh, Ion Stoica, Joseph E. Gonzalez</i>	2586
Implicit Regularization for Tubal Tensor Factorizations via Gradient Descent <i>Santhosh Karnik, Anna Veselovska, Mark Iwen, Felix Krahmer</i>	2608
Skip the Equations: Learning Behavior of Personalized Dynamical Systems Directly From Data <i>Krzysztof Kacprzyk, Julianna Piskorz, Mihaela van der Schaar</i>	2665

Weight matrices compression based on PDB model in deep neural networks <i>Xiaoling Wu, Junpeng Zhu, Zeng Li</i>	2690
Improving the Diffusability of Autoencoders <i>Ivan Skorokhodov, Sharath Girish, Benran Hu, Willi Menapace, Yanyu Li, Rameen Abdal, Sergey Tulyakov, Aliaksandr Siarohin</i>	2706
Metastable Dynamics of Chain-of-Thought Reasoning: Provable Benefits of Search, RL and Distillation <i>Juno Kim, Denny Wu, Jason D. Lee, Taiji Suzuki</i>	2736
Reward Modeling with Ordinal Feedback: Wisdom of the Crowd <i>Shang Liu, Yu Pan, Guanting Chen, Xiaocheng Li</i>	2771
Mastering Multiple-Expert Routing: Realizable H-Consistency and Strong Guarantees for Learning to Defer <i>Anqi Mao, Mehryar Mohri, Yutao Zhong</i>	2800
IMTS is Worth Time × Channel Patches: Visual Masked Autoencoders for Irregular Multivariate Time Series Prediction <i>Zhangyi Hu, Jiemin Wu, Hua XU, Mingqian Liao, Ninghui Feng, Bo Gao, Songning Lai, Yutao Yue</i>	2832
Interchangeable Token Embeddings for Extendable Vocabulary and Alpha-Equivalence <i>İlker Işık, Ramazan Gokberk Cinbis, Ebru Aydin Gol</i>	2850
GuardAgent: Safeguard LLM Agents via Knowledge-Enabled Reasoning <i>Zhen Xiang, Linzhi Zheng, Yanjie Li, Junyuan Hong, Qinbin Li, Han Xie, Jiawei Zhang, Zidi Xiong, Chulin Xie, Carl Yang, Dawn Song, Bo Li</i>	2869
Since Faithfulness Fails: The Performance Limits of Neural Causal Discovery <i>Mateusz Olko, Mateusz Gajewski, Joanna Wojciechowska, Mikołaj Morzy, Piotr Sankowski, Piotr Miłoś</i>	2896
Understanding and Improving Length Generalization in Recurrent Models <i>Ricardo Buitrago, Albert Gu</i>	2917
Measuring Diversity: Axioms and Challenges <i>Mikhail Mironov, Liudmila Prokhorenkova</i>	2936
The Polynomial Stein Discrepancy for Assessing Moment Convergence <i>Narayan Srinivasan, Matthew Sutton, Christopher Drovandi, Leah F South</i>	2952
Understanding the Kronecker Matrix-Vector Complexity of Linear Algebra <i>Raphael A Meyer, William Joseph Swartworth, David Woodruff</i>	2975
M³HF: Multi-agent Reinforcement Learning from Multi-phase Human Feedback of Mixed Quality <i>Ziyan Wang, Zhicheng Zhang, Fei Fang, Yali Du</i>	3000

Distributed Differentially Private Data Analytics via Secure Sketching <i>Jakob Burkhardt, Hannah Keller, Claudio Orlandi, Chris Schwiegelshohn</i>	3020
Structure Is All You Need: Structural Representation Learning on Hyper-Relational Knowledge Graphs <i>Jaejun Lee, Joyce Jiyoung Whang</i>	3050
Normalizing Flows are Capable Generative Models <i>Shuangfei Zhai, Ruixiang ZHANG, Preetum Nakkiran, David Berthelot, Jiatao Gu, Huangjie Zheng, Tianrong Chen, Miguel Ángel Bautista, Navdeep Jaitly, Joshua M. Susskind</i>	3069
I Think, Therefore I Diffuse: Enabling Multimodal In-Context Reasoning in Diffusion Models <i>Zhenxing Mi, Kuan-Chieh Wang, Guocheng Qian, Hanrong Ye, Runtao Liu, Sergey Tulyakov, Kfir Aberman, Dan Xu</i>	3091
L-Diffusion: Laplace Diffusion for Efficient Pathology Image Segmentation <i>Weihan Li, Linyun Zhou, YangJian, Shengxuming Zhang, Xiangtong Du, Xiuming Zhang, Jing Zhang, ChaoqingXu, Mingli Song, Zunlei Feng</i>	3111
NICE Data Selection for Instruction Tuning in LLMs with Non-differentiable Evaluation Metric <i>Jingtian Wang, Xiaoqiang Lin, Rui Qiao, Pang Wei Koh, Chuan-Sheng Foo, Bryan Kian Hsiang Low</i>	3131
WGFormer: An SE(3)-Transformer Driven by Wasserstein Gradient Flows for Molecular Ground-State Conformation Prediction <i>Fanmeng Wang, Minjie Cheng, Hongteng Xu</i>	3159
SlimLLM: Accurate Structured Pruning for Large Language Models <i>Jialong Guo, Xinghao Chen, Yehui Tang, Yunhe Wang</i>	3179
Mixture of Experts Provably Detect and Learn the Latent Cluster Structure in Gradient-Based Learning <i>Ryotaro Kawata, Kohsei Matsutani, Yuri Kinoshita, Naoki Nishikawa, Taiji Suzuki</i>	3190
Does One-shot Give the Best Shot? Mitigating Model Inconsistency in One-shot Federated Learning <i>Hui Zeng, Wenke Huang, Tongqing Zhou, Xinyi Wu, Guancheng Wan, Yingwen Chen, Zhiping Cai</i>	3249
Scaffold with Stochastic Gradients: New Analysis with Linear Speed-Up <i>Paul Mangold, Alain Oliviero Durmus, Aymeric Dieuleveut, Eric Moulines</i>	3267
Geometric Generative Modeling with Noise-Conditioned Graph Networks <i>Peter Pao-Huang, Mitchell Black, Xiaojie Qiu</i>	3312
Agent-Centric Actor-Critic for Asynchronous Multi-Agent Reinforcement Learning <i>Whiyoung Jung, Sunghoon Hong, Deunsol Yoon, Kanghoon Lee, Woohyung Lim</i>	3335
Adaptive Exploration for Multi-Reward Multi-Policy Evaluation <i>Alessio Russo, Aldo Pacchiano</i>	3357

Autoformulation of Mathematical Optimization Models Using LLMs <i>Nicolás Astorga, Tennison Liu, Yuanzhang Xiao, Mihaela van der Schaar</i>	3397
Emergence in non-neural models: grokking modular arithmetic via average gradient outer product <i>Neil Rohit Mallinar, Daniel Beaglehole, Libin Zhu, Adityanarayanan Radhakrishnan, Parthe Pandit, Mikhail Belkin</i>	3420
The Empirical Mean is Minimax Optimal for Local Glivenko-Cantelli <i>Doron Cohen, Aryeh Kontorovich, Roi Weiss</i>	3443
On the Statistical Mechanisms of Distributional Compositional Generalization <i>Jingwen Fu, Nanning Zheng</i>	3455
Federated Node-Level Clustering Network with Cross-Subgraph Link Mending <i>Jingxin Liu, Renda Han, Wenxuan Tu, Haotian Wang, Junlong Wu, Jieren Cheng</i>	3474
Causal Effect Identification in IvLiNGAM from Higher-Order Cumulants <i>Daniele Tramontano, Yaroslav Kivva, Saber Salehkaleybar, Negar Kiyavash, Mathias Drton</i>	3491
Extreme Value Policy Optimization for Safe Reinforcement Learning <i>Shiqing Gao, Yihang Zhou, Shuai Shao, Haoyu Luo, Yiheng Bing, Jiaxin Ding, Luoyi Fu, Xinbing Wang</i>	3512
NextCoder: Robust Adaptation of Code LMs to Diverse Code Edits <i>Tushar Aggarwal, Swayam Singh, Abhijeet Awasthi, Aditya Kanade, Nagarajan Natarajan</i>	3534
PDE-Transformer: Efficient and Versatile Transformers for Physics Simulations <i>Benjamin Holzs Schuh, Qiang Liu, Georg Kohl, Nils Thuerey</i>	3557
Constrained Pareto Set Identification with Bandit Feedback <i>Cyrille Kone, Emilie Kaufmann, Laura Richert</i>	3598
DyCodeEval: Dynamic Benchmarking of Reasoning Capabilities in Code Large Language Models Under Data Contamination <i>Simin Chen, Pranav Pusarla, Baishakhi Ray</i>	3635
LLM Data Selection and Utilization via Dynamic Bi-level Optimization <i>Yang Yu, Kai Han, Hang Zhou, Yehui Tang, Kaiqi Huang, Yunhe Wang, Dacheng Tao</i>	3655
MAS-GPT: Training LLMs to Build LLM-based Multi-Agent Systems <i>Rui Ye, Shuo Tang, Rui Ge, Yaxin Du, Zhenfei Yin, Siheng Chen, Jing Shao</i>	3669
Optimal Survey Design for Private Mean Estimation <i>Yu-Wei Chen, Raghu Pasupathy, Jordan Awan</i>	3697
Algorithm Development in Neural Networks: Insights from the Streaming Parity Task <i>Loek van Rossem, Andrew M Saxe</i>	3712

A Peer-review Look on Multi-modal Clustering: An Information Bottleneck Realization Method <i>Zhengzheng Lou, Hang Xue, Chaoyang Zhang, Shizhe Hu</i>	3740
Improving LLM Video Understanding with 16 Frames Per Second <i>Yixuan Li, Changli Tang, Jimin Zhuang, Yudong Yang, Guangzhi Sun, Wei Li, Zejun MA, Chao Zhang</i>	3756
Improving Parallel Program Performance with LLM Optimizers via Agent-System Interfaces <i>Anjiang Wei, Allen Nie, Thiago S. F. X. Teixeira, Rohan Yadav, Wonchan Lee, Ke Wang, Alex Aiken</i>	3771
Design Considerations in Offline Preference-based RL <i>Alekh Agarwal, Christoph Dann, Teodor Vanislavov Marinov</i>	3794
SHARP-Distill: A 68× Faster Recommender System with Hypergraph Neural Networks and Language Models <i>Saman Forouzandeh, Parham Moradi, Mahdi Jalili</i>	3808
Near Optimal Best Arm Identification for Clustered Bandits <i>Yash, Avishek Ghosh, Nikhil Karamchandani</i>	3845
Leveraging Randomness in Model and Data Partitioning for Privacy Amplification <i>Andy Dong, Wei-Ning Chen, Ayfer Ozgur</i>	3871
Learning without Isolation: Pathway Protection for Continual Learning <i>Zhikang Chen, Abudukelimu Wuerkaixi, Sen Cui, Haoxuan Li, Ding Li, Jingfeng Zhang, Bo Han, Gang Niu, Houfang Liu, Yi Yang, Sifan YANG, Changshui Zhang, Tianling Ren</i>	3896
HaploVL: A Single-Transformer Baseline for Multi-Modal Understanding <i>Rui Yang, Lin Song, Yicheng Xiao, Runhui Huang, Yixiao Ge, Ying Shan, Hengshuang Zhao</i>	3919
CellFlux: Simulating Cellular Morphology Changes via Flow Matching <i>Yuhui Zhang, Yuchang Su, Chenyu Wang, Tianhong Li, Zoe Wefers, Jeffrey J Nirschl, James Burgess, Daisy Ding, Alejandro Lozano, Emma Lundberg, Serena Yeung-Levy</i>	3938
Adversarial Perturbations Are Formed by Iteratively Learning Linear Combinations of the Right Singular Vectors of the Adversarial Jacobian <i>Thomas Paniagua, Chinmay Savadikar, Tianfu Wu</i>	3966
Flat-LoRA: Low-Rank Adaptation over a Flat Loss Landscape <i>Tao Li, Zhengbao He, Yujun Li, Yasheng Wang, Lifeng Shang, Xiaolin Huang</i>	3986
A First-order Generative Bilevel Optimization Framework for Diffusion Models <i>Quan Xiao, Hui Yuan, A F M Saif, Gaowen Liu, Ramana Rao Kompella, Mengdi Wang, Tianyi Chen</i>	4001
Certified Unlearning for Neural Networks <i>Anastasia Koloskova, Youssef Allouah, Animesh Jha, Rachid Guerraoui, Sanmi Koyejo</i>	4025
AlphaQCM: Alpha Discovery in Finance with Distributional Reinforcement Learning <i>Zhoufan Zhu, Ke Zhu</i>	4049

In-Context Linear Regression Demystified: Training Dynamics and Mechanistic Interpretability of Multi-Head Softmax Attention <i>Jianliang He, Xintian Pan, Siyu Chen, Zhuoran Yang</i>	4066
Off-Policy Actor-Critic for Adversarial Observation Robustness: Virtual Alternative Training via Symmetric Policy Evaluation <i>Kosuke Nakanishi, Akihiro Kubo, Yuji Yasui, Shin Ishii</i>	4123
Efficient LiDAR Reflectance Compression via Scanning Serialization <i>Jiahao Zhu, Kang You, Dandan Ding, Zhan Ma</i>	4162
Generalized additive models via direct optimization of regularized decision stump forests <i>Magzhan Gabidolla, Miguel Á. Carreira-Perpiñán</i>	4179
MF-LAL: Drug Compound Generation Using Multi-Fidelity Latent Space Active Learning <i>Peter Eckmann, Dongxia Wu, Germano Heinzelmann, Michael K Gilson, Rose Yu</i>	4194
Compress then Serve: Serving Thousands of LoRA Adapters with Little Overhead <i>Rickard Brühl Gabrielsson, Jiacheng Zhu, Onkar Bhardwaj, Leshem Choshen, Kristjan Greenewald, Mikhail Yurochkin, Justin Solomon</i>	4211
Prediction-Powered Adaptive Shrinkage Estimation <i>Sida Li, Nikolaos Ignatiadis</i>	4245
Deep Neural Cellular Potts Models <i>Koen Minartz, Tim D'Hondt, Leon Hillmann, Jörn Starruß, Lutz Brusch, Vlado Menkovski</i>	4285
Multi-Armed Bandits with Interference: Bridging Causal Inference and Adversarial Bandits <i>Su Jia, Peter I. Frazier, Nathan Kallus</i>	4306
Scalable First-order Method for Certifying Optimal k-Sparse GLMs <i>Jiachang Liu, Soroosh Shafiee, Andrea Lodi</i>	4327
Rejecting Hallucinated State Targets during Planning <i>Harry Zhao, Tristan Sylvain, Romain Laroche, Doina Precup, Yoshua Bengio</i>	4354
A Dynamical Systems-Inspired Pruning Strategy for Addressing Oversmoothing in Graph Attention Networks <i>Biswadeep Chakraborty, Harshit Kumar, Saibal Mukhopadhyay</i>	4380
Rank-One Modified Value Iteration <i>Arman Sharifi Kolarijani, Tolga Ok, Peyman Mohajerin Esfahani, Mohamad Amin Sharifi Kolarijani</i>	4406
Multi-Domain Graph Foundation Models: Robust Knowledge Transfer via Topology Alignment <i>Shuo Wang, Bokui Wang, Zhixiang Shen, Boyan Deng, zhao kang</i>	4426
Towards Understanding Fine-Tuning Mechanisms of LLMs via Circuit Analysis <i>Xu Wang, Yan Hu, Wenyu Du, Reynold Cheng, Benyou Wang, Difan Zou</i>	4442

HyperTree Planning: Enhancing LLM Reasoning via Hierarchical Thinking <i>Runquan Gui, Zhihai Wang, Jie Wang, Chi Ma, Huiling Zhen, Mingxuan Yuan, Jianye HAO, Defu Lian, Enhong Chen, Feng Wu</i>	4467
Origin Identification for Text-Guided Image-to-Image Diffusion Models <i>Wenhao Wang, Yifan Sun, Zongxin Yang, Zhentao Tan, Zhengdong Hu, Yi Yang</i>	4496
Statistical Collusion by Collectives on Learning Platforms <i>Etienne Gauthier, Francis Bach, Michael I. Jordan</i>	4512
Low-distortion and GPU-compatible Tree Embeddings in Hyperbolic Space <i>Max van Spengler, Pascal Mettes</i>	4535
TimeFilter: Patch-Specific Spatial-Temporal Graph Filtration for Time Series Forecasting <i>Yifan Hu, Guibin Zhang, Peiyuan Liu, Disen Lan, Naiqi Li, Dawei Cheng, Tao Dai, Shu-Tao Xia, Shirui Pan</i>	4559
Preference Controllable Reinforcement Learning with Advanced Multi-Objective Optimization <i>Yucheng Yang, Tianyi Zhou, Mykola Pechenizkiy, Meng Fang</i>	4578
Topology-Aware Dynamic Reweighting for Distribution Shifts on Graph <i>Wei Huang Zheng, Jiashuo Liu, Jiaxing Li, Jiayun Wu, Peng Cui, Youyong Kong</i>	4614
Long-Form Speech Generation with Spoken Language Models <i>Se Jin Park, Julian Salazar, Aren Jansen, Keisuke Kinoshita, Yong Man Ro, RJ Skerry-Ryan</i>	4630
Compositional Flows for 3D Molecule and Synthesis Pathway Co-design <i>Tony Shen, Seonghwan Seo, Ross Irwin, Kieran Didi, Simon Olsson, Woo Youn Kim, Martin Ester</i>	4647
Subspace Optimization for Large Language Models with Convergence Guarantees <i>Yutong He, Pengrui Li, Yipeng Hu, Chuyan Chen, Kun Yuan</i>	4676
Do We Need to Verify Step by Step? Rethinking Process Supervision from a Theoretical Perspective <i>Zeyu Jia, Alexander Rakhlin, Tengyang Xie</i>	4731
Random Registers for Cross-Domain Few-Shot Learning <i>Shuai Yi, Yixiong Zou, Yuhua Li, Ruixuan Li</i>	4757
Sampling Binary Data by Denoising through Score Functions <i>Francis Bach, Saeed Saremi</i>	4774
Statistical Test for Feature Selection Pipelines by Selective Inference <i>Tomohiro Shiraishi, Tatsuya Matsukawa, Shuichi Nishino, Ichiro Takeuchi</i>	4794
Control and Realism: Best of Both Worlds in Layout-to-Image without Training <i>Bonan Li, Yinhan Hu, Songhua Liu, Xinchao Wang</i>	4814
A Trichotomy for List Transductive Online Learning <i>Steve Hanneke, Amirreza Shaeiri</i>	4826

Contextures: Representations from Contexts <i>Runtian Zhai, Kai Yang, Burak Varici, Che-Ping Tsai, J Zico Kolter, Pradeep Kumar Ravikumar</i>	4840
Offline-to-Online Reinforcement Learning with Classifier-Free Diffusion Generation <i>Xiao Huang, Xu Liu, Enze Zhang, Tong Yu, Shuai Li</i>	4870
Learning to Incentivize in Repeated Principal-Agent Problems with Adversarial Agent Arrivals <i>Junyan Liu, Arnab Maiti, Artin Tajdini, Kevin Jamieson, Lillian J. Ratliff</i>	4882
FedClean: A General Robust Label Noise Correction for Federated Learning <i>Xiaoqian Jiang, Jing Zhang</i>	4903
Neural Guided Diffusion Bridges <i>Gefan Yang, Frank van der Meulen, Stefan Sommer</i>	4924
Learning Likelihood-Free Reference Priors <i>Nicholas George Bishop, Daniel Jarne Ornia, Joel Dyer, Ani Calinescu, Michael J. Wooldridge</i>	4945
Understanding the Emergence of Multimodal Representation Alignment <i>Megan Tjandrasuwita, Chanakya Ekbote, Liu Ziyin, Paul Pu Liang</i>	4965
Decision Mixer: Integrating Long-term and Local Dependencies via Dynamic Token Selection for Decision-Making <i>Hongling Zheng, Li Shen, Yong Luo, Deheng Ye, Bo Du, Jialie Shen, Dacheng Tao</i>	5003
Attention-Level Speculation <i>Jack Cai, Ammar Vora, Randolph Zhang, Mark O'Connor, Mark C. Jeffrey</i>	5021
NestQuant: nested lattice quantization for matrix products and LLMs <i>Semyon Savkin, Eitan Porat, Or Ordentlich, Yury Polyanskiy</i>	5049
Winner-takes-all for Multivariate Probabilistic Time Series Forecasting <i>Adrien Cortes, Remi Rehm, Victor Letzelter</i>	5070
Efficient Skill Discovery via Regret-Aware Optimization <i>He Zhang, Ming Zhou, Shaopeng Zhai, Ying Sun, Hui Xiong</i>	5095
Sparse Autoencoders for Hypothesis Generation <i>Rajiv Movva, Kenny Peng, Nikhil Garg, Jon Kleinberg, Emma Pierson</i>	5112
GraphGPT: Generative Pre-trained Graph Eulerian Transformer <i>Qifang Zhao, Weidong Ren, Tianyu Li, Hong Liu, Xingsheng He, Xiaoxiao Xu</i>	5139
Ultra-Resolution Adaptation with Ease <i>Ruonan Yu, Songhua Liu, Zhenxiong Tan, Xinchao Wang</i>	5166
GAPrompt: Geometry-Aware Point Cloud Prompt for 3D Vision Model <i>Zixiang Ai, Zichen Liu, Yuanhang Lei, Zhenyu Cui, Xu Zou, Jiahuan Zhou</i>	5187

Understanding the Logic of Direct Preference Alignment through Logic <i>Kyle Richardson, Vivek Srikumar, Ashish Sabharwal</i>	5200
What Limits Virtual Agent Application? OmniBench: A Scalable Multi-Dimensional Benchmark for Essential Virtual Agent Capabilities <i>Wendong Bu, Yang Wu, Qifan Yu, Minghe Gao, Bingchen Miao, Zhenkui Zhang, Kaihang Pan, liyunfei, Mengze Li, Wei Ji, Juncheng Li, Siliang Tang, Yueting Zhuang</i>	5219
Premise-Augmented Reasoning Chains Improve Error Identification in Math reasoning with LLMs <i>Sagnik Mukherjee, Abhinav Chinta, Takyoun Kim, Tarun Anoop Sharma, Dilek Hakkani Tur</i>	5243
Banyan: Improved Representation Learning with Explicit Structure <i>Mattia Oppè, Siddharth N</i>	5263
RAGGED: Towards Informed Design of Scalable and Stable RAG Systems <i>Jennifer Hsia, Afreen Shaikh, Zora Zhiruo Wang, Graham Neubig</i>	5278
SLiM: One-shot Quantization and Sparsity with Low-rank Approximation for LLM Weight Compression <i>Mohammad Mozaffari, Amir Yazdanbakhsh, Maryam Mehri Dehnavi</i>	5295
Geometry Informed Tokenization of Molecules for Language Model Generation <i>Xiner Li, Limei Wang, Youzhi Luo, Carl Edwards, Shurui Gui, Yuchao Lin, Heng Ji, Shuiwang Ji</i>	5321
A Checks-and-Balances Framework for Context-Aware Ethical AI Alignment <i>Edward Y Chang</i>	5354
Revolve: Optimizing AI Systems by Tracking Response Evolution in Textual Optimization <i>Peiyan Zhang, Haibo Jin, Leyang Hu, Xinnuo Li, Liying Kang, Man Luo, Yangqiu Song, Haohan Wang</i>	5373
Prices, Bids, Values: One ML-Powered Combinatorial Auction to Rule Them All <i>Ermis Soumalias, Jakob Heiss, Jakob Weissteiner, Sven Seuken</i>	5391
Adversarial Robust Generalization of Graph Neural Networks <i>Chang Cao, Han Li, Yulong Wang, Rui Wu, Hong Chen</i>	5436
Inference-Time Decomposition of Activations (ITDA): A Scalable Approach to Interpreting Large Language Models <i>Patrick Leask, Neel Nanda, Noura Al Moubayed</i>	5474
Human Body Restoration with One-Step Diffusion Model and A New Benchmark <i>Jue Gong, Jingkai Wang, Zheng Chen, Xing Liu, Hong Gu, Yulun Zhang, Xiaokang Yang</i>	5501
DVI:A Derivative-based Vision Network for INR <i>Runzhao Yang, Xiaolong Wu, Zhihong Zhang, Fabian Zhang, Tingxiong Xiao, Zongren Li, Kunlun He, Jinli Suo</i>	5512

The Missing Alignment Link of In-context Learning on Sequences <i>Harshvardhan Agarwal, Sunita Sarawagi</i>	5532
Open Your Eyes: Vision Enhances Message Passing Neural Networks in Link Prediction <i>Yanbin Wei, Xuehao Wang, Zhan Zhuang, Yang Chen, Shuhao Chen, Yulong Zhang, James Kwok, Yu Zhang</i>	5548
Demystifying Singular Defects in Large Language Models <i>Haoqi Wang, Tong Zhang, Mathieu Salzmann</i>	5569
Embedding Safety into RL: A New Take on Trust Region Methods <i>Nikola Milosevic, Johannes Müller, Nico Scherf</i>	5589
ToMA: Token Merge with Attention for Diffusion Models <i>Wenbo Lu, Shaoyi Zheng, Yuxuan Xia, Shengjie Wang</i>	5615
Enhancing Diversity In Parallel Agents: A Maximum State Entropy Exploration Story <i>Vincenzo de paola, Riccardo Zamboni, Mirco Mutti, Marcello Restelli</i>	5637
Restoring Calibration for Aligned Large Language Models: A Calibration-Aware Fine-Tuning Approach <i>Jiancong Xiao, Bojian Hou, Zhanliang Wang, Ruochen Jin, Qi Long, Weijie J Su, Li Shen</i>	5658
Hierarchical Overlapping Clustering on Graphs: Cost Function, Algorithm and Scalability <i>Yicheng Pan, Renjie Chen, Pengyu Long, Bingchen Fan</i>	5685
Efficient Personalized Adaptation for Physiological Signal Foundation Model <i>Chenrui Wu, Haishuai Wang, Xiang Zhang, Chengqi Zhang, Jiajun Bu</i>	5710
ELITE: Enhanced Language-Image Toxicity Evaluation for Safety <i>Wonjun Lee, Doehyeon Lee, Eugene Choi, Sangyoon Yu, Ashkan Yousefpour, Haon Park, Bumsub Ham, Suhyun Kim</i>	5729
Online Pre-Training for Offline-to-Online Reinforcement Learning <i>Yongjae Shin, Jeonghye Kim, Whiyoung Jung, Sunghoon Hong, Deunsol Yoon, Youngsoo Jang, Geon-Hyeong Kim, Jongseong Chae, Youngchul Sung, Kanghoon Lee, Woohyung Lim</i>	5753
Can Transformers Reason Logically? A Study in SAT Solving <i>Leyan Pan, Vijay Ganesh, Jacob Abernethy, Chris Esposito, Wenke Lee</i>	5776
Improving Your Model Ranking on Chatbot Arena by Vote Rigging <i>Rui Min, Tianyu Pang, Chao Du, Qian Liu, Minhao Cheng, Min Lin</i>	5816
OrthoRank: Token Selection via Sink Token Orthogonality for Efficient LLM inference <i>Seungjun Shin, Jaehoon Oh, Dokwan Oh</i>	5836
Learning Extrapolative Sequence Transformations from Markov Chains <i>Sophia Hager, Aleem Khan, Andrew Wang, Nicholas Andrews</i>	5859

TabPFN Unleashed: A Scalable and Effective Solution to Tabular Classification Problems <i>Siyang Liu, Han-Jia Ye</i>	5880
De-mark: Watermark Removal in Large Language Models <i>Ruibo Chen, Yihan Wu, Junfeng Guo, Heng Huang</i>	5906
Sorbet: A Neuromorphic Hardware-Compatible Transformer-Based Spiking Language Model <i>Kaiwen Tang, Zhanglu Yan, Weng-Fai Wong</i>	5924
VisionTS: Visual Masked Autoencoders Are Free-Lunch Zero-Shot Time Series Forecasters <i>Mouxiang Chen, Lefei Shen, Zhuo Li, Xiaoyun Joy Wang, Jianling Sun, Chenghao Liu</i>	5940
Prune 'n Predict: Optimizing LLM Decision-making with Conformal Prediction <i>Harit Vishwakarma, Alan Mishler, Thomas Cook, Niccolo Dalmasso, Natraj Raman, Sumitra Ganesh</i>	5969
NTPP: Generative Speech Language Modeling for Dual-Channel Spoken Dialogue via Next-Token-Pair Prediction <i>Qichao Wang, Ziqiao Meng, Wenqian Cui, Yifei Zhang, Pengcheng Wu, Bingzhe Wu, Irwin King, Liang Chen, Peilin Zhao</i>	6003
Uniform Mean Estimation for Heavy-Tailed Distributions via Median-of-Means <i>Mikael Møller Høgsgaard, Andrea Paudice</i>	6020
SpikF: Spiking Fourier Network for Efficient Long-term Prediction <i>Wenjie Wu, Dexuan Huo, Hong Chen</i>	6045
Compelling ReLU Networks to Exhibit Exponentially Many Linear Regions at Initialization and During Training <i>Max Milkert, David Hyde, Forrest John Laine</i>	6072
Cowpox: Towards the Immunity of VLM-based Multi-Agent Systems <i>YUTONG WU, Jie Zhang, Yiming Li, Chao Zhang, Qing Guo, Han Qiu, Nils Lukas, Tianwei Zhang</i>	6096
Understanding the Limits of Deep Tabular Methods with Temporal Shift <i>Haorun Cai, Han-Jia Ye</i>	6117
Interpreting CLIP with Hierarchical Sparse Autoencoders <i>Vladimir Zaigrajew, Hubert Baniecki, Przemyslaw Biecek</i>	6138
Differentially Private Analysis for Binary Response Models: Optimality, Estimation, and Inference <i>Ce Zhang, Yixin Han, Yafei Wang, Xiaodong Yan, Linglong Kong, Ting Li, Bei Jiang</i>	6177
CAT: Contrastive Adversarial Training for Evaluating the Robustness of Protective Perturbations in Latent Diffusion Models <i>Sen Peng, Mingyue Wang, Jianfei He, Jijia Yang, Xiaohua Jia</i>	6192

Pivoting Factorization: A Compact Meta Low-Rank Representation of Sparsity for Efficient Inference in Large Language Models <i>Jialin Zhao, Yingtao Zhang, Carlo Vittorio Cannistraci</i>	6206
AsymRnR: Video Diffusion Transformers Acceleration with Asymmetric Reduction and Restoration <i>Wenhao Sun, Rong-Cheng Tu, Jingyi Liao, Zhao Jin, Dacheng Tao</i>	6228
Improved Approximations for Hard Graph Problems using Predictions <i>Anders Aamand, Justin Y. Chen, Siddharth Gollapudi, Sandeep Silwal, Hao WU</i>	6246
Prior Knowledge Guided Neural Architecture Generation <i>Jingrong Xie, Han Ji, Yanan Sun</i>	6275
UI-Vision: A Desktop-centric GUI Benchmark for Visual Perception and Interaction <i>Shravan Nayak, Xiangru Jian, Kevin Qinghong Lin, Juan A. Rodriguez, Montek Kalsi, Nicolas Chapados, M. Tamer Özsu, Aishwarya Agrawal, David Vazquez, Christopher Pal, Perouz Taslakian, Spandana Gella, Sai Rajeswar</i>	6291
EGPlace: An Efficient Macro Placement Method via Evolutionary Search with Greedy Repositioning Guided Mutation <i>ji deng, Zhao Li, Ji Zhang, Jun Gao</i>	6326
Conservative Offline Goal-Conditioned Implicit V-Learning <i>Kaiqiang Ke, Qian Lin, Zongkai Liu, Shenghong He, Chao Yu</i>	6345
TimeStacker: A Novel Framework with Multilevel Observation for Capturing Nonstationary Patterns in Time Series Forecasting <i>Qinglong Liu, Cong Xu, Wenhao Jiang, Kaixuan Wang, Lin Ma, Haifeng Li</i>	6362
Online Learning in Risk Sensitive constrained MDP <i>Arnob Ghosh, Mehrdad Moharrami</i>	6380
Learn to Vaccinate: Combining Structure Learning and Effective Vaccination for Epidemic and Outbreak Control <i>Sepehr Elahi, Paula Mürmann, Patrick Thiran</i>	6400
Learning Optimal Multimodal Information Bottleneck Representations <i>Qilong Wu, Yiyang Shao, Jun Wang, Xiaobo Sun</i>	6433
Riemannian Diffusion Adaptation for Distributed Optimization on Manifolds <i>Xiuheng Wang, Ricardo Augusto Borsoi, Cédric Richard, Ali H. Sayed</i>	6472
A Simple Model of Inference Scaling Laws <i>Noam Itzhak Levi</i>	6494

Fast and Robust: Task Sampling with Posterior and Diversity Synergies for Adaptive Decision-Makers in Randomized Environments <i>Yun Qu, Cheems Wang, Yixiu Mao, Yiqin Lv, Xiangyang Ji</i>	6509
ROPO: Robust Preference Optimization for Large Language Models <i>Xize Liang, Chao Chen, Shuang Qiu, Jie Wang, Yue Wu, Zhihang Fu, Hanzhu Chen, Feng Wu, Jieping Ye</i>	6537
Graph Attention is Not Always Beneficial: A Theoretical Analysis of Graph Attention Mechanisms via Contextual Stochastic Block Models <i>Zhongtian Ma, Qiaosheng Zhang, Bocheng Zhou, Yexin Zhang, Shuyue Hu, Zhen Wang</i>	6568
Scalable Approximation Algorithms for p-Wasserstein Distance and Its Variants <i>Nathaniel Lahn, Sharath Raghvendra, Emma Saarinen, Pouyan Shirzadian</i>	6609
Graph Inverse Style Transfer for Counterfactual Explainability <i>Bardh Prenkaj, Efstratios Zaradoukas, Gjergji Kasneci</i>	6627
To Each Metric Its Decoding: Post-Hoc Optimal Decision Rules of Probabilistic Hierarchical Classifiers <i>Roman Plaud, Alexandre Perez-Lebel, Matthieu Labeau, Antoine Saillenfest, Thomas Bonald</i>	6654
rStar-Math: Small LLMs Can Master Math Reasoning with Self-Evolved Deep Thinking <i>Xinyu Guan, Li Lyna Zhang, Yifei Liu, Ning Shang, Youran Sun, Yi Zhu, Fan Yang, Mao Yang</i>	6698
EduLLM: Leveraging Large Language Models and Framelet-Based Signed Hypergraph Neural Networks for Student Performance Prediction <i>Ming Li, Yukang Cheng, Lu Bai, Feilong Cao, Ke Lv, Jiye Liang, Pietro Lio</i>	6720
Parameter-Efficient Fine-Tuning of State Space Models <i>Kevin Galim, Wonjun Kang, Yuchen Zeng, Hyung Il Koo, Kangwook Lee</i>	6745
Efficient Online Reinforcement Learning for Diffusion Policy <i>Haitong Ma, Tianyi Chen, Kai Wang, Na Li, Bo Dai</i>	6781
Pruning for GNNs: Lower Complexity with Comparable Expressiveness <i>Dun Ma, Jianguo Chen, Wenguo Yang, Suixiang Gao, Shengminjie Chen</i>	6798
Approximately Correct Label Distribution Learning <i>Weiwei Li, Haitao Wu, Yunan Lu, Xiuyi Jia</i>	6834
Average Sensitivity of Hierarchical k-Median Clustering <i>Shijie Li, Weiqiang He, Ruobing Bai, Pan Peng</i>	6846
Cradle: Empowering Foundation Agents towards General Computer Control <i>Weihao Tan, Wentao Zhang, Xinrun Xu, Haochong Xia, Ziluo Ding, Boyu Li, Bohan Zhou, Junpeng Yue, Jiechuan Jiang, Yewen Li, Ruyi An, Molei Qin, Chuqiao Zong, Longtao Zheng, Yulie Wu, Xiaoqiang Chai, Yifei Bi, Tianbao Xie, Pengjie Gu, Xiyun Li, Ceyao Zhang, Long Tian, Chaojie Wang, Xinrun Wang, Börje F. Karlsson, Bo An, Shuicheng YAN, Zongqing Lu</i>	6867

Partition First, Embed Later: Laplacian-Based Feature Partitioning for Refined Embedding and Visualization of High-Dimensional Data <i>Erez Peterfreund, Ofir Lindenbaum, Yuval Kluger, Boris Landa</i>	6935
SHIELD: Multi-task Multi-distribution Vehicle Routing Solver with Sparsity and Hierarchy <i>Yong Liang Goh, Zhiguang Cao, Yining Ma, Jianan Zhou, Mohammed Haroon Dupty, Wee Sun Lee</i>	7003
Adjoint Sampling: Highly Scalable Diffusion Samplers via Adjoint Matching <i>Aaron J Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram, Daniel S. Levine, Brandon M Wood, Bin Hu, Brandon Amos, Brian Karrer, Xiang Fu, Guan-Hong Liu, Ricky T. Q. Chen</i>	7036
DISCO: learning to DISCover an evolution Operator for multi-physics-agnostic prediction <i>Rudy Morel, Jiequn Han, Edouard Oyallon</i>	7070
Test-Time Selective Adaptation for Uni-Modal Distribution Shift in Multi-Modal Data <i>MingCai Chen, Baoming Zhang, Zongbo Han, Wenyu Jiang, Yanmeng Wang, Shuai Feng, Yuntao Du., Bingkun BAO</i>	7095
Cannot See the Forest for the Trees: Invoking Heuristics and Biases to Elicit Irrational Choices of LLMs <i>Haoming Yang, Ke Ma, Xiaojun Jia, Yingfei Sun, Qianqian Xu, Qingming Huang</i>	7112
Geometric and Physical Constraints Synergistically Enhance Neural PDE Surrogates <i>Yunfei Huang, David S. Greenberg</i>	7138
Peripheral Memory for LLMs: Integration of Sequential Memory Banks with Adaptive Querying <i>Songlin Zhai, Yuan Meng, Yongrui Chen, Yiwei Wang, Guilin Qi</i>	7184
Understanding Complexity in VideoQA via Visual Program Generation <i>Cristobal Eyzaguirre, Igor Vasiljevic, Achal Dave, Jiajun Wu, Rares Andrei Ambrus, Thomas Kollar, Juan Carlos Niebles, Pavel Tokmakov</i>	7196
Earley-Driven Dynamic Pruning for Efficient Structured Decoding <i>Xintong Sun, Chi Wei, Minghao Tian, Shiwen Ni</i>	7220
Efficient Length-Generalizable Attention via Causal Retrieval for Long-Context Language Modeling <i>Xiang Hu, Zhihao Teng, Jun Zhao, Wei Wu, Kewei Tu</i>	7237
Towards a General Time Series Forecasting Model with Unified Representation and Adaptive Transfer <i>Yihang Wang, Yuying Qiu, Peng Chen, Kai Zhao, Yang Shu, Zhongwen Rao, Lujia Pan, Bin Yang, Chenjuan Guo</i>	7254
Reinforce LLM Reasoning through Multi-Agent Reflection <i>Yurun Yuan, Tengyang Xie</i>	7279

Sounding that Object: Interactive Object-Aware Image to Audio Generation <i>Tingle Li, Baihe Huang, Xiaobin Zhuang, Dongya Jia, Jiawei Chen, Yuping Wang, Zhuo Chen, Gopala Anumanchipalli, Yuxuan Wang</i>	7310
Beyond Minimax Rates in Group Distributionally Robust Optimization via a Novel Notion of Sparsity <i>Quan M. Nguyen, Nishant A Mehta, Cristóbal A Guzmán</i>	7331
HybridGS: High-Efficiency Gaussian Splatting Data Compression using Dual-Channel Sparse Representation and Point Cloud Encoder <i>Qi Yang, Le Yang, Geert Van der Auwera, Zhu Li</i>	7374
Efficient Distributed Optimization under Heavy-Tailed Noise <i>Su Hyeong Lee, Manzil Zaheer, Tian Li</i>	7395
SAeUron: Interpretable Concept Unlearning in Diffusion Models with Sparse Autoencoders <i>Bartosz Cywiński, Kamil Deja</i>	7445
Sample-specific Noise Injection for Diffusion-based Adversarial Purification <i>Yuhao Sun, Jiacheng Zhang, Zesheng Ye, Chaowei Xiao, Feng Liu</i>	7483
Targeted Unlearning with Single Layer Unlearning Gradient <i>Zikui Cai, Yaoteng Tan, M. Salman Asif</i>	7506
Global Optimization with a Power-Transformed Objective and Gaussian Smoothing <i>Chen Xu</i>	7540
CostFilter-AD: Enhancing Anomaly Detection through Matching Cost Filtering <i>Zhe Zhang, Mingxiu Cai, Hanxiao Wang, Gaochang Wu, Tianyou Chai, Xiatian Zhu</i>	7568
Cross-Modal Alignment via Variational Copula Modelling <i>Feng Wu, Tsai Hor Chan, Fuying Wang, Guosheng Yin, Lequan Yu</i>	7593
Targeted control of fast prototyping through domain-specific interface <i>Yu-Zhe Shi, Mingchen Liu, Hanlu Ma, Qiao Xu, Huamin Qu, Kun He, Lecheng Ruan, Qining Wang</i>	7612
Mixture of Experts Made Intrinsically Interpretable <i>Xingyi Yang, Constantin Venhoff, Ashkan Khakzar, Christian Schroeder de Witt, Puneet K. Dokania, Adel Bibi, Philip Torr</i>	7634
Logarithmic Regret for Online KL-Regularized Reinforcement Learning <i>Heyang Zhao, Chenlu Ye, Wei Xiong, Quanquan Gu, Tong Zhang</i>	7652
Probabilistic Interactive 3D Segmentation with Hierarchical Neural Processes <i>Jie Liu, Pan Zhou, Zehao Xiao, Jiayi Shen, Wenzhe Yin, Jan-Jakob Sonke, Efstratios Gavves</i>	7673
Efficient Multi-modal Long Context Learning for Training-free Adaptation <i>Zehong Ma, Shiliang Zhang, Longhui Wei, Qi Tian</i>	7693

Enhancing Foundation Models with Federated Domain Knowledge Infusion <i>Jiaqi Wang, Jingtao Li, Weiming Zhuang, Chen Chen, Lingjuan Lyu, Fenglong Ma</i>	7709
PoisonedEye: Knowledge Poisoning Attack on Retrieval-Augmented Generation based Large Vision-Language Models <i>Chenyang Zhang, Xiaoyu Zhang, Jian Lou, Kai Wu, Zilong Wang, Xiaofeng Chen</i>	7724
Flex3D: Feed-Forward 3D Generation with Flexible Reconstruction Model and Input View Curation <i>Junlin Han, Jianyuan Wang, Andrea Vedaldi, Philip Torr, Filippos Kokkinos</i>	7744
SyncMind: Measuring Agent Out-of-Sync Recovery in Collaborative Software Engineering <i>Xuehang Guo, Xingyao Wang, Yangyi Chen, Sha Li, Chi Han, Manling Li, Heng Ji</i>	7762
Exploiting Presentative Feature Distributions for Parameter-Efficient Continual Learning of Large Language Models <i>Xin Cheng, Jiabo Ye, Haiyang Xu, Ming Yan, Ji Zhang, Feng Liu, Fei Huang, Lei Feng</i>	7836
Understanding the Forgetting of (Replay-based) Continual Learning via Feature Learning: Angle Matters <i>Hongyi Wan, Shiyuan Ren, Wei Huang, Miao Zhang, Xiang Deng, Yixin Bao, Liqiang Nie</i>	7859
Efficient Molecular Conformer Generation with SO(3)-Averaged Flow Matching and Reflow <i>Zhonglin Cao, Mario Geiger, Allan Dos Santos Costa, Danny Reidenbach, Karsten Kreis, Tomas Geffner, Franco Pellegrini, Guoqing Zhou, Emine Kucukbenli</i>	7923
Imagine While Reasoning in Space: Multimodal Visualization-of-Thought <i>Chengzu Li, Wenshan Wu, Huanyu Zhang, Yan Xia, Shaoguang Mao, Li Dong, Ivan Vulić, Furu Wei</i>	7944
Stochastic Layer-Wise Shuffle for Improving Vision Mamba Training <i>Zizheng Huang, Haoxing Chen, Jiaqi Li, Jun Lan, Huijia Zhu, Weiqiang Wang, Limin Wang</i>	7969
Large Language Models to Diffusion Finetuning <i>Edoardo Cetin, Tianyu Zhao, Yujin Tang</i>	7983
PENCIL: Long Thoughts with Short Memory <i>Chenxiao Yang, Nathan Srebro, David McAllester, Zhiyuan Li</i>	8006
Multinoulli Extension: A Lossless Yet Effective Probabilistic Framework for Subset Selection over Partition Constraints <i>Qixin Zhang, Wei Huang, Can Jin, Puning Zhao, Yao Shu, Li Shen, Dacheng Tao</i>	8071
FDGen: A Fairness-Aware Graph Generation Model <i>Zichong Wang, Wenbin Zhang</i>	8108
Provably Near-Optimal Federated Ensemble Distillation with Negligible Overhead <i>Won-Jun Jang, Hyeon-Seo Park, Si-Hyeon Lee</i>	8125

Joint MoE Scaling Laws: Mixture of Experts Can Be Memory Efficient <i>Jan Ludziejewski, Maciej Pióro, Jakub Krajewski, Maciej Stefaniak, Michał Krutul, Jan Małaśnicki, Marek Cygan, Piotr Sankowski, Kamil Adamczewski, Piotr Miłoś, Sebastian Jaszczur</i>	8154
Polynomial-Delay MAG Listing with Novel Locally Complete Orientation Rules <i>Tian-Zuo Wang, Wen-Bo Du, Zhi-Hua Zhou</i>	8172
Relative Error Fair Clustering in the Weak-Strong Oracle Model <i>Vladimir Braverman, Prathamesh Dharangutte, Shaofeng H.-C. Jiang, Hoai-An Nguyen, Chen Wang, Yubo Zhang, Samson Zhou</i>	8193
Graph-Assisted Stitching for Offline Hierarchical Reinforcement Learning <i>Seungho Baek, taegeon park, Jongchan Park, Seungjun Oh, Yusung Kim</i>	8229
Mask-Enhanced Autoregressive Prediction: Pay Less Attention to Learn More <i>Xialie Zhuang, Zhikai Jia, Jianjin Li, Zhenyu Zhang, Li Shen, Zheng Cao, Shiwei Liu</i>	8247
Thermalizer: Stable autoregressive neural emulation of spatiotemporal chaos <i>Christian Pedersen, Laure Zanna, Joan Bruna</i>	8264
SpargeAttention: Accurate and Training-free Sparse Attention Accelerating Any Model Inference <i>Jintao Zhang, Chendong Xiang, Haofeng Huang, Jia wei, Haocheng Xi, Jun Zhu, Jianfei Chen</i>	8291
Vision-Language Models Create Cross-Modal Task Representations <i>Grace Luo, Trevor Darrell, Amir Bar</i>	8308
FAB-PPI: Frequentist, Assisted by Bayes, Prediction-Powered Inference <i>Stefano Cortinovis, Francois Caron</i>	8328
Sassha: Sharpness-aware Adaptive Second-order Optimization with Stable Hessian Approximation <i>Dahun Shin, Dongyeop Lee, Jinseok Chung, Namhoon Lee</i>	8357
The Hidden Life of Tokens: Reducing Hallucination of Large Vision-Language Models Via Visual Information Steering <i>Zhuowei Li, Haizhou Shi, Yunhe Gao, Di Liu, Zhenting Wang, Yuxiao Chen, Ting Liu, Long Zhao, Hao Wang, Dimitris N. Metaxas</i>	8384
Offline Learning for Combinatorial Multi-armed Bandits <i>Xutong Liu, Xiangxiang Dai, Jinhang Zuo, Siwei Wang, Carlee Joe-Wong, John C.S. Lui, Wei Chen</i>	8405
A Manifold Perspective on the Statistical Generalization of Graph Neural Networks <i>Zhiyang Wang, Juan Cervino, Alejandro Ribeiro</i>	8444
On Understanding Attention-Based In-Context Learning for Categorical Data <i>Aaron T Wang, William Convertino, Xiang Cheng, Ricardo Henao, Lawrence Carin</i>	8482

Confidence Difference Reflects Various Supervised Signals in Confidence-Difference Classification <i>Yuanchao Dai, Ximing Li, Changchun Li</i>	8510
CSV-Occ: Fusing Multi-frame Alignment for Occupancy Prediction with Temporal Cross State Space Model and Central Voting Mechanism <i>Ziming Zhu, Yu Zhu, Jiahao Chen, LingXiaofeng, Huanlei Chen, Lihua Sun</i>	8525
A New Approach to Backtracking Counterfactual Explanations: A Unified Causal Framework for Efficient Model Interpretability <i>Pouria Fatemi, Ehsan Sharifian, Mohammad Hossein Yassaee</i>	8544
Weak-to-Strong Jailbreaking on Large Language Models <i>Xuandong Zhao, Xianjun Yang, Tianyu Pang, Chao Du, Lei Li, Yu-Xiang Wang, William Yang Wang</i>	8557
PDE-Controller: LLMs for Autoformalization and Reasoning of PDEs <i>Mauricio Soroco, Jialin Song, Mengzhou Xia, Kye Emond, Weiran Sun, Wuyang Chen</i>	8579
Weisfeiler and Leman Go Gambling: Why Expressive Lottery Tickets Win <i>Lorenz Kummer, Samir Moustafa, Anatol Ehrlich, Franka Bause, Nikolaus Suess, Wilfried N. Gansterer, Nils Morten Kriege</i>	8611
Event-Customized Image Generation <i>Zhen Wang, Yilei JIANG, Dong Zheng, Jun Xiao, Long Chen</i>	8631
Tool Unlearning for Tool-Augmented LLMs <i>Jiali Cheng, Hadi Amiri</i>	8652
FLAM: Frame-Wise Language-Audio Modeling <i>Yusong Wu, Christos Tsirigotis, Ke Chen, Cheng-Zhi Anna Huang, Aaron Courville, Oriol Nieto, Prem Seetharaman, Justin Salamon</i>	8669
CSG-ODE: ControlSynth Graph ODE For Modeling Complex Evolution of Dynamic Graphs <i>Zhiqiang Wang, Xiaoyi Wang, Jianqing Liang</i>	8691
Kinetic Langevin Diffusion for Crystalline Materials Generation <i>François R J Cornet, Federico Bergamin, Arghya Bhowmik, Juan Maria Garcia-Lastra, Jes Frellsen, Mikkel N. Schmidt</i>	8711
Fairness on Principal Stratum: A New Perspective on Counterfactual Fairness <i>Haoxuan Li, Zeyu Tang, Zhichao Jiang, Zhuangyan Fang, Yue Liu, Zhi Geng, Kun Zhang</i>	8741
Avoiding Leakage Poisoning: Concept Interventions Under Distribution Shifts <i>Mateo Espinosa Zarlenga, Gabriele Dominici, Pietro Barbiero, Zohreh Shams, Mateja Jamnik</i>	8754
Breaking the $n^{1.5}$ Additive Error Barrier for Private and Efficient Graph Sparsification via Private Expander Decomposition <i>Anders Aamand, Justin Y. Chen, Mina Dalirrooyfard, Slobodan Mitrović, Yuriy Nevmyvaka, Sandeep Silwal, Yinzhan Xu</i>	8786

Meta-Black-Box-Optimization through Offline Q-function Learning <i>Zeyuan Ma, Zhiguang Cao, Zhou Jiang, Hongshu Guo, Yue-Jiao Gong</i>	8800
GPEN: Global Position Encoding Network for Enhanced Subgraph Representation Learning <i>Nannan Wu, Yuming Huang, Yiming Zhao, Jie Chen, Wenjun Wang</i>	8820
SeedLoRA: A Fusion Approach to Efficient LLM Fine-Tuning <i>Yong Liu, Di Fu, Shenggan Cheng, Zirui Zhu, Yang Luo, Minhao Cheng, Cho-Jui Hsieh, Yang You</i>	8836
Federated Incomplete Multi-view Clustering with Globally Fused Graph Guidance <i>Guoqing Chao, Zhenghao Zhang, Lei Meng, Jie Wen, Dianhui Chu</i>	8849
Decomposition of Graphic Design with Unified Multimodal Model <i>Hui Nie, Zhao Zhang, Yutao Cheng, Maoke Yang, Gonglei Shi, Qingsong Xie, Jie Shao, Xinglong Wu</i>	8862
QuantSpec: Self-Speculative Decoding with Hierarchical Quantized KV Cache <i>Rishabh Tiwari, Haocheng Xi, Aditya Tomar, Coleman Richard Charles Hooper, Sehoon Kim, Maxwell Horton, Mahyar Najibi, Michael W. Mahoney, Kurt Keutzer, Amir Gholami</i>	8874
Privacy Attacks on Image AutoRegressive Models <i>Antoni Kowalczyk, Jan Dubiński, Franziska Boenisch, Adam Dziedzic</i>	8893
TruthFlow: Truthful LLM Generation via Representation Flow Correction <i>Hanyu Wang, Bochuan Cao, Yuanpu Cao, Jinghui Chen</i>	8922
Neural Discovery in Mathematics: Do Machines Dream of Colored Planes? <i>Konrad Munding, Max Zimmer, Aldo Kiem, Christoph Spiegel, Sebastian Pokutta</i>	8944
SafeArena: Evaluating the Safety of Autonomous Web Agents <i>Ada Defne Tur, Nicholas Meade, Xing Han Lù, Alejandra Zambrano, Arkil Patel, Esin DURMUS, Spandana Gella, Karolina Stanczak, Siva Reddy</i>	8964
LineFlow: A Framework to Learn Active Control of Production Lines <i>Kai Müller, Martin Wenzel, Tobias Windisch</i>	9002
SAFER: A Calibrated Risk-Aware Multimodal Recommendation Model for Dynamic Treatment Regimes <i>Yishan Shen, Yuyang Ye, Hui Xiong, Yong Chen</i>	9026
A Lens into Interpretable Transformer Mistakes via Semantic Dependency <i>Ruo-Jing Dong, Yu Yao, Bo Han, Tongliang Liu</i>	9047
Continuous-Time Analysis of Heavy Ball Momentum in Min-Max Games <i>Yi Feng, Kaito Fujii, Stratis Skoulakis, Xiao Wang, Volkan Cevher</i>	9072
Omni-Angle Assault: An Invisible and Powerful Physical Adversarial Attack on Face Recognition <i>Shuai Yuan, Hongwei Li, Rui Zhang, Hangcheng Cao, Wenbo Jiang, Tao Ni, Wenshu Fan, Qingchuan Zhao, Guowen Xu</i>	9113

Prompt-to-Leaderboard: Prompt-Adaptive LLM Evaluations <i>Evan Frick, Connor Chen, Joseph Tennyson, Tianle Li, Wei-Lin Chiang, Anastasios Nikolas Angelopoulos, Ion Stoica</i>	9125
Mechanistic PDE Networks for Discovery of Governing Equations <i>Adeel Pervez, Efstratios Gavves, Francesco Locatello</i>	9143
Statistical and Computational Guarantees of Kernel Max-Sliced Wasserstein Distances <i>Jie Wang, March Boedihardjo, Yao Xie</i>	9155
Eigenspectrum Analysis of Neural Networks without Aspect Ratio Bias <i>Yuanzhe Hu, Kinshuk Goel, Vlad Killiakov, Yaoqing Yang</i>	9183
Confounder-Free Continual Learning via Recursive Feature Normalization <i>Yash Shah, Camila Gonzalez, Mohammad H. Abbasi, Qingyu Zhao, Kilian M. Pohl, Ehsan Adeli</i>	9207
SparseVLM: Visual Token Sparsification for Efficient Vision-Language Model Inference <i>Yuan Zhang, Chun-Kai Fan, Junpeng Ma, Wenzhao Zheng, Tao Huang, Kuan Cheng, Denis A Gudovskiy, Tomoyuki Okuno, Yohei Nakata, Kurt Keutzer, Shanghang Zhang</i>	9238
The Geometry of Refusal in Large Language Models: Concept Cones and Representational Independence <i>Tom Wollschläger, Jannes Elstner, Simon Geisler, Vincent Cohen-Addad, Stephan Günnemann, Johannes Gasteiger</i>	9256
One Arrow, Two Hawks: Sharpness-aware Minimization for Federated Learning via Global Model Trajectory <i>Yuhang Li, Tong Liu, Yangguang Cui, Ming Hu, Xiaoqiang Li</i>	9282
Test-Time Graph Neural Dataset Search With Generative Projection <i>Xin Zheng, Wei Huang, Chuan Zhou, Ming Li, Shirui Pan</i>	9307
The Energy Loss Phenomenon in RLHF: A New Perspective on Mitigating Reward Hacking <i>Yuchun Miao, Sen Zhang, Liang Ding, Yuqi Zhang, Lefei Zhang, Dacheng Tao</i>	9321
Scaling Laws for Floating-Point Quantization Training <i>Xingwu Sun, Shuaipeng Li, Ruobing Xie, Weidong Han, Kan Wu, Zhen Yang, Yixing Li, An Wang, SHUAI LI, Jinbao Xue, Yu Cheng, Yangyu Tao, Zhanhui Kang, Cheng-zhong Xu, Di Wang, Jie Jiang</i>	9351
Disentangled Graph Spectral Domain Adaptation <i>Liang Yang, Xin Chen, Jiaming Zhuo, Di Jin, Chuan Wang, Xiaochun Cao, Zhen Wang, Yuanfang Guo</i>	9378
OmiAD: One-Step Adaptive Masked Diffusion Model for Multi-class Anomaly Detection via Adversarial Distillation <i>Yaoxuan Feng, Wenchao Chen, Yuxin Li, Bo Chen, Yubiao Wang, Zixuan Zhao, Hongwei Liu, Mingyuan Zhou</i>	9395
Fully Heteroscedastic Count Regression with Deep Double Poisson Networks <i>Spencer Young, Porter Jenkins, Longchao Da, Jeff Dotson, Hua Wei</i>	9425

Flexibility-conditioned protein structure design with flow matching <i>Vsevolod Viliuga, Leif Seute, Nicolas Wolf, Simon Wagner, Arne Elofsson, Jan Stühmer, Frauke Gräter</i>	9458
All-atom Diffusion Transformers: Unified generative modelling of molecules and materials <i>Chaitanya K. Joshi, Xiang Fu, Yi-Lun Liao, Vahe Gharakhanyan, Benjamin Kurt Miller, Anuroop Sriram, Zachary Ward Ulissi</i>	9479
LMAct: A Benchmark for In-Context Imitation Learning with Long Multimodal Demonstrations <i>Anian Ruoss, Fabio Pardo, Harris Chan, Bonnie Li, Volodymyr Mnih, Tim Genewein</i>	9504
Controlling Neural Collapse Enhances Out-of-Distribution Detection and Transfer Learning <i>Md Yousuf Harun, Jhair Gallardo, Christopher Kanan</i>	9546
EffiCoder: Enhancing Code Generation in Large Language Models through Efficiency-Aware Fine-tuning <i>Dong HUANG, Guangtao Zeng, Jianbo Dai, Meng Luo, Han Weng, Yuhao QING, Heming Cui, Zhijiang Guo, Jie Zhang</i>	9574
It's My Data Too: Private ML for Datasets with Multi-User Training Examples <i>Arun Ganesh, Ryan McKenna, Hugh Brendan McMahan, Adam Smith, Fan Wu</i>	9593
Autonomy-of-Experts Models <i>Ang Lv, Ruobing Xie, Yining Qian, Songhao Wu, Xingwu Sun, Zhanhui Kang, Di Wang, Rui Yan</i>	9610
Decision-aware Training of Spatiotemporal Forecasting Models to Select a Top-K Subset of Sites for Intervention <i>Kyle Heuton, Frederick Muench, Shikhar Shrestha, Thomas J. Stopka, Michael C Hughes</i>	9625
HetSSNet: Spatial-Spectral Heterogeneous Graph Learning Network for Panchromatic and Multispectral Images Fusion <i>Mengting Ma, Yizhen Jiang, Mengjiao Zhao, Jiaxin Li, Wei Zhang</i>	9644
Compute Optimal Inference and Provable Amortisation Gap in Sparse Autoencoders <i>Charles O'Neill, Alim Gumran, David Klindt</i>	9660
NBDI: A Simple and Effective Termination Condition for Skill Extraction from Task-Agnostic Demonstrations <i>Myunsoo Kim, Hayeong Lee, Seong-Woong Shim, JunHo Seo, Byung-Jun Lee</i>	9680
A Certified Unlearning Approach without Access to Source Data <i>Umit Yigit Basaran, Sk Miraj Ahmed, Amit Roy-Chowdhury, Basak Guler</i>	9705
Improving Out-of-Distribution Detection with Markov Logic Networks <i>Konstantin Kirchheim, Frank Ortmeier</i>	9725
Anytime-Constrained Equilibria in Polynomial Time <i>Jeremy McMahan</i>	9736

GRADEO: Towards Human-Like Evaluation for Text-to-Video Generation via Multi-Step Reasoning <i>Zhun Mou, Bin Xia, Zhengchao Huang, Wenming Yang, Jiaya Jia</i>	9754
Maximum Total Correlation Reinforcement Learning <i>Bang You, Puze Liu, Huaping Liu, Jan Peters, Oleg Arenz</i>	9780
Label Distribution Propagation-based Label Completion for Crowdsourcing <i>Tong Wu, Liangxiao Jiang, Wenjun Zhang, Chaoqun Li</i>	9803
Multiple-policy Evaluation via Density Estimation <i>Yilei Chen, Aldo Pacchiano, Ioannis Paschalidis</i>	9816
CodeSync: Synchronizing Large Language Models with Dynamic Code Evolution at Scale <i>Chenlong Wang, Zhaoyang Chu, Zhengxiang Cheng, Xuyi Yang, Kaiyue Qiu, Yao Wan, Zhou Zhao, Xuanhua Shi, Hai Jin, Dongping Chen</i>	9843
PARQ: Piecewise-Affine Regularized Quantization <i>Lisa Jin, Jianhao Ma, Zechun Liu, Andrey Gromov, Aaron Defazio, Lin Xiao</i>	9872
TOPLOC: A Locality Sensitive Hashing Scheme for Trustless Verifiable Inference <i>Jack Min Ong, Matthew Di Ferrante, Aaron Pazdera, Ryan Garner, Sami Jaghouar, Manveer Basra, Max Ryabinin, Johannes Hagemann</i>	9891
Progressively Label Enhancement for Large Language Model Alignment <i>Biao Liu, Ning Xu, Xin Geng</i>	9907
BILBO: BILevel Bayesian Optimization <i>Ruth Wan Theng Chew, Quoc Phong Nguyen, Bryan Kian Hsiang Low</i>	9921
(How) Do Language Models Track State? <i>Belinda Z. Li, Zifan Carl Guo, Jacob Andreas</i>	9941
From Spectrum-free towards Baseline-view-free: Double-track Proximity Driven Multi-view Clustering <i>Shengju Yu, Zhibin Dong, Siwei Wang, Suyuan Liu, KE LIANG, Xinwang Liu, Yue Liu, Yi Zhang</i>	9965
Free Process Rewards without Process Labels <i>Lifan Yuan, Wendi Li, Huayu Chen, Ganqu Cui, Ning Ding, Kaiyan Zhang, Bowen Zhou, Zhiyuan Liu, Hao Peng</i>	9995
MPO: An Efficient Post-Processing Framework for Mixing Diverse Preference Alignment <i>Tianze Wang, Dongnan Gui, Yifan Hu, Shuhang Lin, Linjun Zhang</i>	10010
All-atom inverse protein folding through discrete flow matching <i>Kai Yi, Kiarash Jamali, Sjors HW Scheres</i>	10032

DiTAR: Diffusion Transformer Autoregressive Modeling for Speech Generation <i>Dongya Jia, Zhuo Chen, Jiawei Chen, Chenpeng Du, Jian Wu, Jian Cong, Xiaobin Zhuang, Chumin Li, Zhen Wei, Yuping Wang, Yuxuan Wang</i>	10046
Can RLHF be More Efficient with Imperfect Reward Models? A Policy Coverage Perspective <i>Jiawei Huang, Bingcong Li, Christoph Dann, Niao He</i>	10062
Topology-aware Neural Flux Prediction Guided by Physics <i>Haoyang Jiang, Jindong Wang, Xingquan Zhu, Yi He</i>	10098
Evaluating Neuron Explanations: A Unified Framework with Sanity Checks <i>Tuomas Oikarinen, Ge Yan, Tsui-Wei Weng</i>	10113
Beyond Zero Initialization: Investigating the Impact of Non-Zero Initialization on LoRA Fine-Tuning Dynamics <i>Shiwei Li, Xiandi Luo, Xing Tang, Haozhao Wang, Hao Chen, weihongluo, Yuhua Li, xiuqiang He, Ruixuan Li</i>	10155
X-Transfer Attacks: Towards Super Transferable Adversarial Attacks on CLIP <i>Hanxun Huang, Sarah Monazam Erfani, Yige Li, Xingjun Ma, James Bailey</i>	10172
Outsourced Diffusion Sampling: Efficient Posterior Inference in Latent Spaces of Generative Models <i>Siddarth Venkatraman, Mohsin Hasan, Minsu Kim, Luca Scimeca, Marcin Sendera, Yoshua Bengio, Glen Berseth, Nikolay Malkin</i>	10203
Self-Consistency Preference Optimization <i>Archiki Prasad, Weizhe Yuan, Richard Yuanzhe Pang, Jing Xu, Maryam Fazel-Zarandi, Mohit Bansal, Sainbayar Sukhbaatar, Jason E Weston, Jane Yu</i>	10231
Ehrenfeucht-Haussler Rank and Chain of Thought <i>Pablo Barcelo, Alexander Kozachinskiy, Tomasz Steifer</i>	10246
C2IQL: Constraint-Conditioned Implicit Q-learning for Safe Offline Reinforcement Learning <i>Zifan LIU, Xinran Li, Jun Zhang</i>	10256
BinauralFlow: A Causal and Streamable Approach for High-Quality Binaural Speech Synthesis with Flow Matching Models <i>Susan Liang, Dejan Markovic, Israel D. Gebru, Steven Krenn, Todd Keebler, Jacob Sandakly, Frank Yu, Samuel Hassel, Chenliang Xu, Alexander Richard</i>	10271
Clustering Items through Bandit Feedback: Finding the Right Feature out of Many <i>Maximilian Graf, Thuot Victor, Nicolas Verzelen</i>	10289
Balanced Learning for Domain Adaptive Semantic Segmentation <i>Wangkai Li, Rui Sun, Bohao Liao, Zhaoyang Li, Tianzhu Zhang</i>	10319
One-dimensional Path Convolution <i>Xuanshu Luo, Martin Werner</i>	10345

Heavy-Tailed Linear Bandits: Huber Regression with One-Pass Update <i>Jing Wang, Yu-Jie Zhang, Peng Zhao, Zhi-Hua Zhou</i>	10361
Improving Memory Efficiency for Training KANs via Meta Learning <i>Zhangchi Zhao, Jun Shu, Deyu Meng, Zongben Xu</i>	10384
GIVE: Structured Reasoning of Large Language Models with Knowledge Graph Inspired Veracity Extrapolation <i>Jiashu He, Mingyu Derek Ma, Jinxuan Fan, Dan Roth, Wei Wang, Alejandro Ribeiro</i>	10404
Copilot Arena: A Platform for Code LLM Evaluation in the Wild <i>Wayne Chi, Valerie Chen, Anastasios Nikolas Angelopoulos, Wei-Lin Chiang, Aditya Mittal, Naman Jain, Tianjun Zhang, Ion Stoica, Chris Donahue, Ameet Talwalkar</i>	10426
Segment Anyword: Mask Prompt Inversion for Open-Set Grounded Segmentation <i>Zhihua Liu, Amrutha Saseendran, Lei Tong, Xilin He, Fariba Yousefi, Nikolay Burlutskiy, Dino Oglic, Tom Diethe, Philip Alexander Teare, Huiyu Zhou, Chen Jin</i>	10455
Towards the Causal Complete Cause of Multi-Modal Representation Learning <i>Jingyao Wang, Siyu Zhao, Wenwen Qiang, Jiangmeng Li, Changwen Zheng, Fuchun Sun, Hui Xiong</i>	10490
Nesterov Method for Asynchronous Pipeline Parallel Optimization <i>Thalaiyasingam Ajanthan, Sameera Ramasinghe, Yan Zuo, Gil Avraham, Alexander Long</i>	10528
TCP-Diffusion: A Multi-modal Diffusion Model for Global Tropical Cyclone Precipitation Forecasting with Change Awareness <i>Cheng Huang, Pan Mu, Cong Bai, Peter AG Watson</i>	10544
MaskTwins: Dual-form Complementary Masking for Domain-Adaptive Image Segmentation <i>Jiawen Wang, Yinda Chen, Xiaoyu Liu, Che Liu, Dong Liu, Jianqing Gao, Zhiwei Xiong</i>	10564
SERENA: A Unified Stochastic Recursive Variance Reduced Gradient Framework for Riemannian Non-Convex Optimization <i>Yan Liu, Mingjie Chen, Chaojie Ji, Hao Zhang, Ruxin Wang</i>	10590
Safe-EF: Error Feedback for Non-smooth Constrained Optimization <i>Rustem Islamov, Yarden As, Ilyas Fatkhullin</i>	10615
Field Matching: an Electrostatic Paradigm to Generate and Transfer Data <i>Alexander Kolesov, S. I. Manukhov, Vladimir Vladimirovich Palyulin, Alexander Korotin</i>	10659
Convergence of Policy Mirror Descent Beyond Compatible Function Approximation <i>Uri Sherman, Tomer Koren, Yishay Mansour</i>	10680
Ab Initio Nonparametric Variable Selection for Scalable Symbolic Regression with Large p <i>Shengbin Ye, Meng Li</i>	10719
Zero-Shot Offline Imitation Learning via Optimal Transport <i>Thomas Rumpf, Marco Bagatella, Nico Gürtler, Jonas Frey, Georg Martius</i>	10741

Can Transformers Learn Full Bayesian Inference in Context? <i>Arik Reuter, Tim G. J. Rudner, Vincent Fortuin, David Rügamer</i>	10778
Efficient Optimization with Orthogonality Constraint: a Randomized Riemannian Submanifold Method <i>Andi Han, Pierre-Louis Poirion, Akiko Takeda</i>	10830
Flow-based Domain Randomization for Learning and Sequencing Robotic Skills <i>Aidan Curtis, Eric Li, Michael Noseworthy, Nishad Gothoskar, Sachin Chitta, Hui Li, Leslie Pack Kaelbling, Nicole E Carey</i>	10860
Compressing tree ensembles through Level-wise Optimization and Pruning <i>Laurens Devos, Timo Martens, Deniz Can Oruc, Wannes Meert, Hendrik Blockeel, Jesse Davis</i>	10878
Pairwise Maximum Likelihood For Multi-Class Logistic Regression Model With Multiple Rare Classes <i>Xuetong Li, Danyang Huang, Hansheng Wang</i>	10903
Kernel Quantile Embeddings and Associated Probability Metrics <i>Masha Naslidnyk, Siu Lun Chau, Francois-Xavier Briol, Krikamol Muandet</i>	10920
Learning from Suboptimal Data in Continuous Control via Auto-Regressive Soft Q-Network <i>Jijia Liu, Feng Gao, Qingmin Liao, Chao Yu, Yu Wang</i>	10951
Efficient ANN-SNN Conversion with Error Compensation Learning <i>Chang Liu, Jiangrong Shen, Xuming Ran, Mingkun Xu, Qi Xu, Yi Xu, Gang Pan</i>	10973
Feature-Mapping Topology Optimization with Neural Heaviside Signed Distance Functions <i>Aleksandr Kolomeitsev, ANH-HUY PHAN</i>	10987
TimeStep Master: Asymmetrical Mixture of Timestep LoRA Experts for Versatile and Efficient Diffusion Models in Vision <i>Shaobin Zhuang, Yiwei Guo, Yanbo Ding, Kunchang Li, Xinyuan Chen, Yaohui Wang, Fangyikang Wang, Ying Zhang, Chen Li, Yali Wang</i>	11013
TeLoGraF: Temporal Logic Planning via Graph-encoded Flow Matching <i>Yue Meng, Chuchu Fan</i>	11030
A Generalizable Physics-Enhanced State Space Model for Long-Term Dynamics Forecasting in Complex Environments <i>Yuchen Wang, Hongjue Zhao, Haohong Lin, Enze Xu, Lifang He, Huajie Shao</i>	11057
WikiBigEdit: Understanding the Limits of Lifelong Knowledge Editing in LLMs <i>Lukas Thede, Karsten Roth, Matthias Bethge, Zeynep Akata, Thomas Hartvigsen</i>	11087
S4S: Solving for a Fast Diffusion Model Solver <i>Eric Frankel, Sitan Chen, Jerry Li, Pang Wei Koh, Lillian J. Ratliff, Sewoong Oh</i>	11116

On the Query Complexity of Verifier-Assisted Language Generation <i>Edoardo Botta, Yuchen Li, Aashay Mehta, Jordan T. Ash, Cyril Zhang, Andrej Risteski</i>	11149
Action Dubber: Timing Audible Actions via Inflectional Flow <i>Wenlong Wan, Weiying Zheng, Tianyi Xiang, Guiqing Li, Shengfeng He</i>	11181
OptMATH: A Scalable Bidirectional Data Synthesis Framework for Optimization Modeling <i>Hongliang Lu, Zhonglin Xie, Yaoyu Wu, Can Ren, Yuxuan Chen, Zaiwen Wen</i>	11198
The Diffusion Duality <i>Subham Sekhar Sahoo, Justin Deschenaux, Aaron Gokaslan, Guanghan Wang, Justin T Chiu, Volodymyr Kuleshov</i>	11232
Accelerating Unbiased LLM Evaluation via Synthetic Feedback <i>Zhaoyi Zhou, Yuda Song, Andrea Zanette</i>	11268
Reward-Guided Iterative Refinement in Diffusion Models at Test-Time with Applications to Protein and DNA Design <i>Masatoshi Uehara, Xingyu Su, Yulai Zhao, Xiner Li, Aviv Regev, Shuiwang Ji, Sergey Levine, Tommaso Biancalani</i>	11286
Cut out and Replay: A Simple yet Versatile Strategy for Multi-Label Online Continual Learning <i>Xinrui Wang, Shao-Yuan Li, Jiaqiang Zhang, Songcan Chen</i>	11301
ProDiff: Prototype-Guided Diffusion for Minimal Information Trajectory Imputation <i>Tianci Bu, Le Zhou, Wenchuan Yang, Jianhong Mou, Kang Yang, Suoyi Tan, Feng Yao, Jingyuan Wang, Xin Lu</i>	11320
Phase and Amplitude-aware Prompting for Enhancing Adversarial Robustness <i>Yibo Xu, Dawei Zhou, Decheng Liu, Nannan Wang</i>	11338
Lightweight Dataset Pruning without Full Training via Example Difficulty and Prediction Uncertainty <i>Yeseul Cho, Baekrok Shin, Changmin Kang, Chulhee Yun</i>	11353
An End-to-End Model for Logits-Based Large Language Models Watermarking <i>KA HIM WONG, Jicheng Zhou, Jiantao Zhou, Yain-Whar Si</i>	11395
SEAD: Unsupervised Ensemble of Streaming Anomaly Detectors <i>Saumya Gaurang Shah, Abishek Sankararaman, Balakrishnan Murali Narayanaswamy, Vikramank Singh</i>	11416
Robust Consensus Anchor Learning for Efficient Multi-view Subspace Clustering <i>Yalan Qin, Nan Pu, Guorui Feng, Nicu Sebe</i>	11435
RZ-NAS: Enhancing LLM-guided Neural Architecture Search via Reflective Zero-Cost Strategy <i>Zipeng Ji, Guanghui Zhu, Chunfeng Yuan, Yihua Huang</i>	11446
Generalization Analysis for Controllable Learning <i>Yi-Fan Zhang, Xiao Zhang, Min-Ling Zhang</i>	11464

Clustering via Self-Supervised Diffusion <i>Roy Uziel, Irit Chelly, Oren Freifeld, Ari Pakman</i>	11491
Pareto-frontier Entropy Search with Variational Lower Bound Maximization <i>Masanori Ishikura, Masayuki Karasuyama</i>	11507
Archetypal SAE: Adaptive and Stable Dictionary Learning for Concept Extraction in Large Vision Models <i>Thomas Fel, Ekdeep Singh Lubana, Jacob S. Prince, Matthew Kowal, Victor Boutin, Isabel Papadimitriou, Binxu Wang, Martin Wattenberg, Demba E. Ba, Talia Konkle</i>	11540
Learning Distribution-wise Control in Representation Space for Language Models <i>Chunyuan Deng, Ruidi Chang, Hanjie Chen</i>	11570
Expert Race: A Flexible Routing Strategy for Scaling Diffusion Transformer with Mixture of Experts <i>Yike Yuan, Ziyu Wang, Zihao Huang, Defa Zhu, Xun Zhou, Jingyi Yu, Qiyang Min</i>	11595
OneForecast: A Universal Framework for Global and Regional Weather Forecasting <i>Yuan Gao, Hao Wu, Ruiqi Shu, huanshuo dong, Fan Xu, Rui Ray Chen, Yibo Yan, Qingsong Wen, Xuming Hu, Kun Wang, Jiahao Wu, Li Qing, Hui Xiong, Xiaomeng Huang</i>	11607
Representation Surgery in Model Merging with Probabilistic Modeling <i>Qi Wei, Shuo He, Enneng Yang, Tingcong Liu, Haobo Wang, Lei Feng, Bo An</i>	11647
One Example Shown, Many Concepts Known! Counterexample-Driven Conceptual Reasoning in Mathematical LLMs <i>Yinghui Li, Jiayi Kuang, Haojing Huang, Zhikun Xu, Xinnian Liang, Yi Yu, Wenlian Lu, Yangning Li, Xiaoyu Tan, Chao Qu, Ying Shen, Hai-Tao Zheng, Philip S. Yu</i>	11665
The Limits of Predicting Agents from Behaviour <i>Alexis Bellot, Jonathan Richens, Tom Everitt</i>	11683
Unlocking Post-hoc Dataset Inference with Synthetic Data <i>Bihe Zhao, Pratyush Maini, Franziska Boenisch, Adam Dziedzic</i>	11719
Test-Time Adaptation with Binary Feedback <i>Taeckyung Lee, Sorn Chottanaturak, Junsu Kim, Jinwoo Shin, Taesik Gong, Sung-Ju Lee</i>	11741
IMPACT: Iterative Mask-based Parallel Decoding for Text-to-Audio Generation with Diffusion Modeling <i>Kuan-Po Huang, Shu-wen Yang, HUY PHAN, Bo-Ru Lu, Byeonggeun Kim, Sashank Macha, Qingming Tang, Shalini Ghosh, Hung-yi Lee, Chieh-Chi Kao, Chao Wang</i>	11761
TMetaNet: Topological Meta-Learning Framework for Dynamic Link Prediction <i>Hao Li, Hao Wan, Yuzhou Chen, Dongsheng Ye, Yulia Gel, Hao Jiang</i>	11779
DPCore: Dynamic Prompt Coreset for Continual Test-Time Adaptation <i>Yunbei Zhang, Akshay Mehra, Shuaicheng Niu, Jihun Hamm</i>	11802

ReFocus: Visual Editing as a Chain of Thought for Structured Image Understanding <i>Xingyu Fu, Minqian Liu, Zhengyuan Yang, John Richard Corring, Yijuan Lu, Jianwei Yang, Dan Roth, Dinei Florencio, Cha Zhang</i>	11824
Harmonizing Geometry and Uncertainty: Diffusion with Hyperspheres <i>Muskan Dosi, Chiranjeev Chiranjeev, Kartik Thakral, Mayank Vatsa, Richa Singh</i>	11847
Can We Predict Performance of Large Models across Vision-Language Tasks? <i>Qinyu Zhao, Ming Xu, Kartik Gupta, Akshay Asthana, Liang Zheng, Stephen Gould</i>	11873
Diffusion Adversarial Post-Training for One-Step Video Generation <i>Shanchuan Lin, Xin Xia, Yuxi Ren, Ceyuan Yang, Xuefeng Xiao, Lu Jiang</i>	11899
Verification Learning: Make Unsupervised Neuro-Symbolic System Feasible <i>Lin-Han Jia, Wen-Chao Hu, Jie-Jing Shao, Lan-Zhe Guo, Yu-Feng Li</i>	11915
A Unified Approach to Routing and Cascading for LLMs <i>Jasper Dekoninck, Maximilian Baader, Martin Vechev</i>	11930
TabFSBench: Tabular Benchmark for Feature Shifts in Open Environments <i>Zi-Jian Cheng, Ziyi Jia, Zhi Zhou, Yu-Feng Li, Lan-Zhe Guo</i>	11954
Merge-Friendly Post-Training Quantization for Multi-Target Domain Adaptation <i>Juncheol Shin, Minsang Seok, Seonggon Kim, Eunhyeok Park</i>	12019
MoE-SVD: Structured Mixture-of-Experts LLMs Compression via Singular Value Decomposition <i>Wei Li, Lujun Li, Hao Gu, You-Liang Huang, Mark G. Lee, Shengjie Sun, Wei Xue, Yike Guo</i>	12029
Near-Optimal Decision Trees in a SPLIT Second <i>Varun Babbar, Hayden McTavish, Cynthia Rudin, Margo Seltzer</i>	12051
Efficient Motion Prompt Learning for Robust Visual Tracking <i>Jie Zhao, Xin Chen, Yongsheng Yuan, Michael Felsberg, Dong Wang, Huchuan Lu</i>	12113
3D-LMVIC: Learning-based Multi-View Image Compression with 3D Gaussian Geometric Priors <i>Yujun Huang, Bin Chen, Niu Lian, Xin Wang, Baoyi An, Tao Dai, Shu-Tao Xia</i>	12131
Multi-View Graph Clustering via Node-Guided Contrastive Encoding <i>Yazhou Ren, Junlong Ke, Zichen Wen, Tianyi Wu, Yang Yang, Xiaorong Pu, Lifang He</i>	12149
Towards characterizing the value of edge embeddings in Graph Neural Networks <i>Dhruv Rohatgi, Tanya Marwah, Zachary Chase Lipton, Jianfeng Lu, Ankur Moitra, Andrej Risteski</i>	12164
Selective Prompt Anchoring for Code Generation <i>Yuan Tian, Tianyi Zhang</i>	12183
Overcoming the Curse of Dimensionality in Reinforcement Learning Through Approximate Factorization <i>Chenbei Lu, Laixi Shi, Zaiwei Chen, Chenye Wu, Adam Wierman</i>	12207

∞-Video: A Training-Free Approach to Long Video Understanding via Continuous-Time Memory Consolidation <i>Saul Santos, António Farinhas, Daniel C McNamee, Andre Martins</i>	12258
Enhancing the Influence of Labels on Unlabeled Nodes in Graph Convolutional Networks <i>Jincheng Huang, Yujie Mo, Xiaoshuang Shi, Lei Feng, Xiaofeng Zhu</i>	12275
An All-Atom Generative Model for Designing Protein Complexes <i>Ruizhe Chen, Dongyu Xue, Xiangxin Zhou, Zaixiang Zheng, xiangxiang Zeng, Quanquan Gu</i>	12296
Lightweight-Mark: Rethinking Deep Learning-Based Watermarking <i>Yupeng Qiu, Han Fang, Ee-Chien Chang</i>	12326
MoMa: Modulating Mamba for Adapting Image Foundation Models to Video Recognition <i>Yuhuan Yang, Chaofan Ma, Zhenjie Mao, Jiangchao Yao, Ya Zhang, Yanfeng Wang</i>	12348
ROME is Forged in Adversity: Robust Distilled Datasets via Information Bottleneck <i>Zheng Zhou, Wenquan Feng, Qiaosheng Zhang, Shuchang Lyu, Qi Zhao, Guangliang Cheng</i>	12361
Teaching Transformers Causal Reasoning through Axiomatic Training <i>Aniket Vashishtha, Abhinav Kumar, Atharva Pandey, Abbavaram Gowtham Reddy, Kabir Ahuja, Vineeth N. Balasubramanian, Amit Sharma</i>	12388
Near-Optimal Consistency-Robustness Trade-Offs for Learning-Augmented Online Knapsack Problems <i>Mohammadreza Daneshvaramoli, Helia Karisani, Adam Lechowicz, Bo Sun, Cameron N Musco, Mohammad Hajiesmaili</i>	12408
STAIR: Improving Safety Alignment with Introspective Reasoning <i>Yichi Zhang, Siyuan Zhang, Yao Huang, Zeyu Xia, Zhengwei Fang, Xiao Yang, Ranjie Duan, Dong Yan, Yinpeng Dong, Jun Zhu</i>	12439
Long-Term TalkingFace Generation via Motion-Prior Conditional Diffusion Model <i>Fei Shen, Cong Wang, Junyao Gao, Qin Guo, Jisheng Dang, Jinhui Tang, Tat-Seng Chua</i>	12463
Closed-form Solutions: A New Perspective on Solving Differential Equations <i>Shu Wei, Yanjie Li, Lina Yu, Weijun Li, Min Wu, Linjun Sun, Jingyi Liu, Hong Qin, Yusong Deng, Jufeng Han, Yan Pang</i>	12479
Diffuse Everything: Multimodal Diffusion Models on Arbitrary State Spaces <i>Kevin Rojas, Yuchen Zhu, Sichen Zhu, Felix X-F. Ye, Molei Tao</i>	12501
AnyEdit: Edit Any Knowledge Encoded in Language Models <i>Houcheng Jiang, Junfeng Fang, Ningyu Zhang, Mingyang Wan, Guojun Ma, Xiang Wang, Xiangnan He, Tat-Seng Chua</i>	12534
Score-based Pullback Riemannian Geometry: Extracting the Data Manifold Geometry using Anisotropic Flows <i>Willem Diepeveen, Georgios Batzolis, Zakhar Shumaylov, Carola-Bibiane Schönlieb</i>	12558

David and Goliath: Small One-step Model Beats Large Diffusion with Score Post-training <i>Weijian Luo, colin zhang, Debing Zhang, Zhengyang Geng</i>	12586
Programming Every Example: Lifting Pre-training Data Quality Like Experts at Scale <i>Fan Zhou, Zengzhi Wang, Qian Liu, Junlong Li, Pengfei Liu</i>	12606
KABB: Knowledge-Aware Bayesian Bandits for Dynamic Expert Coordination in Multi-Agent Systems <i>Jusheng Zhang, Zimeng Huang, Yijia Fan, Ningyuan Liu, Mingyan Li, Zhuojie Yang, Jiawei Yao, Jian Wang, Keze Wang</i>	12653
Learning Progress Driven Multi-Agent Curriculum <i>Wenshuai Zhao, Zhiyuan Li, Joni Pajarinen</i>	12684
AtlasD: Automatic Local Symmetry Discovery <i>Manu Bhat, Jonghyun Park, Jianke Yang, Nima Dehmamy, Robin Walters, Rose Yu</i>	12701
Transolver+ ++: An Accurate Neural Solver for PDEs on Million-Scale Geometries <i>Huakun Luo, Haixu Wu, Hang Zhou, Lanxiang Xing, Yichen Di, Jianmin Wang, Mingsheng Long</i>	12720
Safety Certificate against Latent Variables with Partially Unidentifiable Dynamics <i>Haoming Jing, Yorie Nakahira</i>	12738
EgoPrivacy: What Your First-Person Camera Says About You? <i>Yijiang Li, Genpei Zhang, Jiacheng Cheng, Yi Li, Xiaojun Shan, Dashan Gao, Jiancheng Lyu, Yuan Li, Ning Bi, Nuno Vasconcelos</i>	12754
SAM2Act: Integrating Visual Foundation Model with A Memory Architecture for Robotic Manipulation <i>Haoquan Fang, Markus Grotz, Wilbert Pumacay, Yi Ru Wang, Dieter Fox, Ranjay Krishna, Jiafei Duan</i>	12770
Emergent Misalignment: Narrow finetuning can produce broadly misaligned LLMs <i>Jan Betley, Daniel Chee Hian Tan, Niels Warncke, Anna Sztyber-Betley, Xuchan Bao, Martín Soto, Nathan Labenz, Owain Evans</i>	12788
BCE vs. CE in Deep Feature Learning <i>Qiufu Li, Huibin Xiao, Linlin Shen</i>	12814
The Panaceas for Improving Low-Rank Decomposition in Communication-Efficient Federated Learning <i>Shiwei Li, Xiandi Luo, Haozhao Wang, Xing Tang, Shijie Xu, weihongluo, Yuhua Li, xiuqiang He, Ruixuan Li</i>	12849
IBCircuit: Towards Holistic Circuit Discovery with Information Bottleneck <i>Tian Bian, Yifan Niu, Chaohao Yuan, Chengzhi Piao, Bingzhe Wu, Long-Kai Huang, Yu Rong, Tingyang Xu, Hong Cheng, Jia Li</i>	12875
Differentiable Quadratic Optimization For the Maximum Independent Set Problem <i>Ismail Alkhouri, Cedric Le Denmat, Yingjie Li, CUNXI YU, Jia Liu, Rongrong Wang, Alvaro Velasquez</i>	12889

Gradual Transition from Bellman Optimality Operator to Bellman Operator in Online Reinforcement Learning <i>Motoki Omura, Kazuki Ota, Takayuki Osa, Yusuke Mukuta, Tatsuya Harada</i>	12915
Resolving Lexical Bias in Model Editing <i>Hammad Rizwan, Domenic Rosati, Ga Wu, Hassan Sajjad</i>	12935
Learning Minimum-Size BDDs: Towards Efficient Exact Algorithms <i>Christian Komusiewicz, Andre Schidler, Frank Sommer, Manuel Sorge, Luca Pascal Staus</i>	12958
FEAT-KD: Learning Concise Representations for Single and Multi-Target Regression via TabNet Knowledge Distillation <i>Kei Sen Fong, Mehul Motani</i>	12978
Discriminative Policy Optimization for Token-Level Reward Models <i>Hongzhan Chen, Tao Yang, Shiping Gao, Ruijun Chen, Xiaojun Quan, Hongtao Tian, Ting Yao</i>	12992
Time to Spike? Understanding the Representational Power of Spiking Neural Networks in Discrete Time <i>Duc Anh Nguyen, Ernesto Araya, Adalbert Fono, Gitta Kutyniok</i>	13012
Antidote: Post-fine-tuning Safety Alignment for Large Language Models against Harmful Fine-tuning Attack <i>Tiansheng Huang, Gautam Bhattacharya, Pratik Joshi, Joshua Kimball, Ling Liu</i>	13046
Test-Time Preference Optimization: On-the-Fly Alignment via Iterative Textual Feedback <i>Yafu Li, Xuyang Hu, Xiaoye Qu, Linjie Li, Yu Cheng</i>	13062
On The Concurrence of Layer-wise Preconditioning Methods and Provable Feature Learning <i>Thomas TCK Zhang, Behrad Moniri, Ansh Nagwekar, Faraz Rahman, Anton Xue, Hamed Hassani, Nikolai Matni</i>	13106
(How) Can Transformers Predict Pseudo-Random Numbers? <i>Tao Tao, Darshil Doshi, Dayal Singh Kalra, Tianyu He, Maissam Barkeshli</i>	13147
The Surprising Effectiveness of Test-Time Training for Few-Shot Learning <i>Ekin Akyürek, Mehul Damani, Adam Zweiger, Linlu Qiu, Han Guo, Jyothish Pari, Yoon Kim, Jacob Andreas</i>	13176
Aligning Protein Conformation Ensemble Generation with Physical Feedback <i>Jiarui Lu, Xiaoyin Chen, Stephen Zhewen Lu, Aurelie Lozano, Vijil Chenthamarakshan, Payel Das, Jian Tang</i>	13198
Transformative or Conservative? Conservation laws for ResNets and Transformers <i>Sibylle Marcotte, Rémi Gribonval, Gabriel Peyré</i>	13214
Fine-Grained Captioning of Long Videos through Scene Graph Consolidation <i>Sanghyeok Chu, Seonguk Seo, Bohyung Han</i>	13251

Unveiling AI's Blind Spots: An Oracle for In-Domain, Out-of-Domain, and Adversarial Errors <i>Shuangpeng Han, Mengmi Zhang</i>	13268
Nearly Optimal Algorithms for Contextual Dueling Bandits from Adversarial Feedback <i>Qiwei Di, Jiafan He, Quanquan Gu</i>	13289
Memory Layers at Scale <i>Vincent-Pierre Berges, Barlas Oguz, Daniel HAZIZA, Wen-tau Yih, Luke Zettlemoyer, Gargi Ghosh</i>	13322
Be a Goldfish: Forgetting Bad Conditioning in Sparse Linear Regression via Variational Autoencoders <i>Kuheli Pratihar, Debdeep Mukhopadhyay</i>	13334
RLTHF: Targeted Human Feedback for LLM Alignment <i>Yifei Xu, Tusher Chakraborty, Emre Kiciman, Bibek Aryal, Srinagesh Sharma, Songwu Lu, Ranveer Chandra</i>	13354
SongGen: A Single Stage Auto-regressive Transformer for Text-to-Song Generation <i>Zihan Liu, Shuangrui Ding, Zhixiong Zhang, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Yuhang Cao, Dahua Lin, Jiaqi Wang</i>	13374
Generalization of noisy SGD in unbounded non-convex settings <i>Leello Tadesse Dadi, Volkan Cevher</i>	13388
On Zero-Initialized Attention: Optimal Prompt and Gating Factor Estimation <i>Nghiem Tuong Diep, Huy Nguyen, Chau Nguyen, Minh Le, Duy Minh Ho Nguyen, Daniel Sonntag, Mathias Niepert, Nhat Ho</i>	13410
Residual TPP: A Unified Lightweight Approach for Event Stream Data Analysis <i>Ruoxin Yuan, Guanhua Fang</i>	13443
Eliciting Language Model Behaviors with Investigator Agents <i>Xiang Lisa Li, Neil Chowdhury, Daniel D. Johnson, Tatsunori Hashimoto, Percy Liang, Sarah Schwettmann, Jacob Steinhardt</i>	13466
Model Swarms: Collaborative Search to Adapt LLM Experts via Swarm Intelligence <i>Shangbin Feng, Zifeng Wang, Yike Wang, Sayna Ebrahimi, Hamid Palangi, Lesly Miculicich, Achin Kulshrestha, Nathalie Rauschmayr, Yejin Choi, Yulia Tsvetkov, Chen-Yu Lee, Tomas Pfister</i>	13486
Learning-Augmented Algorithms for MTS with Bandit Access to Multiple Predictors <i>Matei Gabriel Cosa, Marek Elias</i>	13513
Learning Adaptive Lighting via Channel-Aware Guidance <i>Qirui Yang, Peng-Tao Jiang, Hao Zhang, Jinwei Chen, Bo Li, Huanjing Yue, Jingyu Yang</i>	13535
Reward-Guided Speculative Decoding for Efficient LLM Reasoning <i>Baohao Liao, Yuhui Xu, Hanze Dong, Junnan Li, Christof Monz, Silvio Savarese, Doyen Sahoo, Caiming Xiong</i>	13552

Provable Length Generalization in Sequence Prediction via Spectral Filtering <i>Annie Marsden, Evan Dogariu, Naman Agarwal, Xinyi Chen, Daniel Suo, Elad Hazan</i>	13570
Continuous Visual Autoregressive Generation via Score Maximization <i>Chenze Shao, Fandong Meng, Jie Zhou</i>	13595
DexScale: Automating Data Scaling for Sim2Real Generalizable Robot Control <i>Guiliang Liu, Yueci Deng, Runyi Zhao, Huayi Zhou, Jian Chen, Jietao Chen, Ruiyan Xu, Yunxin Tai, Kui Jia</i>	13610
How Transformers Learn Structured Data: Insights From Hierarchical Filtering <i>Jerome Garnier-Brun, Marc Mezard, Emanuele Moscato, Luca Saglietti</i>	13629
DIME: Diffusion-Based Maximum Entropy Reinforcement Learning <i>Onur Celik, Zechu Li, Denis Blessing, Ge Li, Daniel Palenicek, Jan Peters, Georgia Chalvatzaki, Gerhard Neumann</i>	13646
Adversarial Reasoning at Jailbreaking Time <i>Mahdi Sabbaghi, Paul Kassianik, George J. Pappas, Amin Karbasi, Hamed Hassani</i>	13666
Steerable Transformers for Volumetric Data <i>Soumyabrata Kundu, Risi Kondor</i>	13692
Geometric Resampling in Nearly Linear Time for Follow-the-Perturbed-Leader with Best-of-Both-Worlds Guarantee in Bandit Problems <i>Botao Chen, Jongyeong Lee, Junya Honda</i>	13713
Quantifying Prediction Consistency Under Fine-tuning Multiplicity in Tabular LLMs <i>Faisal Hamman, Pasan Dissanayake, Saumitra Mishra, Freddy Lecue, Sanghamitra Dutta</i>	13737
Compositional Risk Minimization <i>Divyat Mahajan, Mohammad Pezeshki, Charles Arnal, Ioannis Mitliagkas, Kartik Ahuja, Pascal Vincent</i>	13760
Explicit Exploration for High-Welfare Equilibria in Game-Theoretic Multiagent Reinforcement Learning <i>Austin A. Nguyen, Anri Gu, Michael P. Wellman</i>	13811
Compositional Condition Question Answering in Tabular Understanding <i>Jun-Peng Jiang, Tao Zhou, De-Chuan Zhan, Han-Jia Ye</i>	13831
Teaching Physical Awareness to LLMs through Sounds <i>Weiguo Wang, Andy Nie, Wenrui Zhou, Yi Kai, Chengchen Hu</i>	13851
Distributed Nonparametric Estimation: from Sparse to Dense Samples per Terminal <i>Deheng Yuan, Tao Guo, Zhongyi Huang</i>	13871
Calibrating Video Watch-time Predictions with Credible Prototype Alignment <i>Chao Cui, Shisong Tang, Fan Li, Jiechao Gao, Hechang Chen</i>	13904

GenZSL: Generative Zero-Shot Learning Via Inductive Variational Autoencoder <i>Shiming Chen, Dingjie Fu, Salman Khan, Fahad Shahbaz Khan</i>	13926
Sortformer: A Novel Approach for Permutation-Resolved Speaker Supervision in Speech-to-Text Systems <i>Taejin Park, Ivan Medennikov, Kunal Dhawan, Weiqing Wang, He Huang, Nithin Rao Koluguri, Krishna C Puvvada, Jagadeesh Balam, Boris Ginsburg</i>	13941
MA-LoT: Model-Collaboration Lean-based Long Chain-of-Thought Reasoning enhances Formal Theorem Proving <i>Ruida WANG, Rui Pan, Yuxin Li, Jipeng Zhang, Yizhen Jia, Shizhe Diao, Renjie Pi, Junjie Hu, Tong Zhang</i>	13958
Learning Imbalanced Data with Beneficial Label Noise <i>Guangzheng Hu, Feng Liu, Mingming Gong, Guanghui Wang, Lihua Peng</i>	13991
CTBench: A Library and Benchmark for Certified Training <i>Yuhao Mao, Stefan Balaucu, Martin Vechev</i>	14026
Optimization over Sparse Support-Preserving Sets: Two-Step Projection with Global Optimality Guarantees <i>William de Vazelhes, Xiaotong Yuan, Bin Gu</i>	14053
Otter: Generating Tests from Issues to Validate SWE Patches <i>Toufique Ahmed, Jatin Ganhotra, Rangeet Pan, Avraham Shinnar, Saurabh Sinha, Martin Hirzel</i>	14114
The Devil Is in the Details: Tackling Unimodal Spurious Correlations for Generalizable Multimodal Reward Models <i>Zichao Li, Xueru Wen, Jie Lou, Yuqiu Ji, Yaojie Lu, Xianpei Han, Debing Zhang, Le Sun</i>	14134
Splitting & Integrating: Out-of-Distribution Detection via Adversarial Gradient Attribution <i>Jiayu Zhang, Xinyi Wang, Zhibo Jin, Zhiyu Zhu, Jianlong Zhou, Fang Chen, Huaming Chen</i>	14151
IRBridge: Solving Image Restoration Bridge with Pre-trained Generative Diffusion Models <i>Hanting Wang, Tao Jin, Wang Lin, Shulei Wang, Hai Huang, Shengpeng Ji, Zhou Zhao</i>	14163
Efficiently Access Diffusion Fisher: Within the Outer Product Span Space <i>Fangyikang Wang, Hubery Yin, Shaobin Zhuang, Huminhao Zhu, Yinan Li, Lei Qian, Chao Zhang, Hanbin Zhao, Hui Qian, Chen Li</i>	14187
Knowledge Swapping via Learning and Unlearning <i>Mingyu Xing, Lechao Cheng, Shengeng Tang, Yaxiong Wang, Zhun Zhong, Meng Wang</i>	14212
FRUGAL: Memory-Efficient Optimization by Reducing State Overhead for Scalable Training <i>Philip Zmushko, Aleksandr Beznosikov, Martin Takáč, Samuel Horváth</i>	14223
Time-VLM: Exploring Multimodal Vision-Language Models for Augmented Time Series Forecasting <i>Siru Zhong, Weilin Ruan, Ming Jin, Huan Li, Qingsong Wen, Yuxuan Liang</i>	14255

The Global Convergence Time of Stochastic Gradient Descent in Non-Convex Landscapes: Sharp Estimates via Large Deviations <i>Waïss Azizian, Franck Iutzeler, Jerome Malick, Panayotis Mertikopoulos</i>	14275
A Reduction Framework for Distributionally Robust Reinforcement Learning under Average Reward <i>Zachary Andrew Roch, George K. Atia, Yue Wang</i>	14338
Piloting Structure-Based Drug Design via Modality-Specific Optimal Schedule <i>Keyue Qiu, Yuxuan Song, Zhehuan Fan, Peidong Liu, Zhe Zhang, Mingyue Zheng, Hao Zhou, Wei-Ying Ma</i>	14380
Enhancing Foundation Models for Time Series Forecasting via Wavelet-based Tokenization <i>Luca Masserano, Abdul Fatir Ansari, Boran Han, Xiyuan Zhang, Christos Faloutsos, Michael W. Mahoney, Andrew Gordon Wilson, Youngsuk Park, Syama Sundar Rangapuram, Danielle C. Maddix, Bernie Wang</i>	14406
DragLoRA: Online Optimization of LoRA Adapters for Drag-based Image Editing in Diffusion Model <i>Siwei Xia, Li Sun, Tiantian Sun, Qingli Li</i>	14434
The Surprising Agreement Between Convex Optimization Theory and Learning-Rate Scheduling for Large Model Training <i>Fabian Schaipp, Alexander Hägele, Adrien Taylor, Umut Simsekli, Francis Bach</i>	14449
Learning Survival Distributions with the Asymmetric Laplace Distribution <i>Deming Sheng, Ricardo Henao</i>	14477
AssistanceZero: Scalably Solving Assistance Games <i>Cassidy Laidlaw, Eli Bronstein, Timothy Guo, Dylan Feng, Lukas Berglund, Justin Svegliato, Stuart Russell, Anca Dragan</i>	14515
Flexible, Efficient, and Stable Adversarial Attacks on Machine Unlearning <i>Zihan Zhou, Yang Zhou, Zijie Zhang, Lingjuan Lyu, Da Yan, Ruoming Jin, Dejing Dou</i>	14543
Dynamic Sparse Training of Diagonally Sparse Networks <i>Abhishek Tyagi, Arjun Iyer, William H Renninger, Christopher Kanan, Yuhao Zhu</i>	14592
X-Hacking: The Threat of Misguided AutoML <i>Rahul Sharma, Sumantrak Mukherjee, Andrea Sipka, Eyke Hüllermeier, Sebastian Josef Vollmer, Sergey Redyuk, David Antony Selby</i>	14616
Regularized Langevin Dynamics for Combinatorial Optimization <i>Shengyu Feng, Yiming Yang</i>	14654
L3A: Label-Augmented Analytic Adaptation for Multi-Label Class Incremental Learning <i>Xiang Zhang, Run He, Chen Jiao, Di Fang, Ming Li, Ziqian Zeng, Cen Chen, Huiping Zhuang</i>	14670

Is Your Model Fairly Certain? Uncertainty-Aware Fairness Evaluation for LLMs <i>Yinong Oliver Wang, Nivedha Sivakumar, Falaah Arif Khan, Katherine Metcalf, Adam Golinski, Natalie Mackraz, Barry-John Theobald, Luca Zappella, Nicholas Apostoloff</i>	14682
Polynomial-Time Approximability of Constrained Reinforcement Learning <i>Jeremy McMahan</i>	14706
Visual Graph Arena: Evaluating Visual Conceptualization of Vision and Multimodal Large Language Models <i>Zahra Babaiee, Peyman Kiasari, Daniela Rus, Radu Grosu</i>	14729
Constrain Alignment with Sparse Autoencoders <i>Qingyu Yin, Chak Tou Leong, Hongbo Zhang, Minjun Zhu, Hanqi Yan, Qiang Zhang, Yulan He, Wenjie Li, Jun Wang, Yue Zhang, Linyi Yang</i>	14762
EAGLES: Towards Effective, Efficient, and Economical Federated Graph Learning via Unified Sparsification <i>Zitong Shi, Guancheng Wan, Wenke Huang, Guibin Zhang, He Li, Carl Yang, Mang Ye</i>	14777
Learning from others' mistakes: Finetuning machine translation models with span-level error annotations <i>Lily H Zhang, Hamid Dadkhahi, Mara Finkelstein, Firas Trabelsi, Jiaming Luo, Markus Freitag</i>	14796
Automated Red Teaming with GOAT: the Generative Offensive Agent Tester <i>Maya Pavlova, Erik Brinkman, Krithika Iyer, Vitor Albiero, Joanna Bitton, Hailey Nguyen, Cristian Canton Ferrer, Ivan Evtimov, Aaron Grattafiori</i>	14813
Quantum Speedups in Regret Analysis of Infinite Horizon Average-Reward Markov Decision Processes <i>Bhargav Ganguly, Yang Xu, Vaneet Aggarwal</i>	14831
DCBM: Data-Efficient Visual Concept Bottleneck Models <i>Katharina Prasse, Patrick Knab, Sascha Marton, Christian Bartelt, Margret Keuper</i>	14851
Improved Last-Iterate Convergence of Shuffling Gradient Methods for Nonsmooth Convex Optimization <i>Zijian Liu, Zhengyuan Zhou</i>	14882
Preference-CFR: Beyond Nash Equilibrium for Better Game Strategies <i>Qi Ju, Thomas Tellier, Meng Sun, Zhemei Fang, YunFeng Luo</i>	14924
RBench: Graduate-level Multi-disciplinary Benchmarks for LLM & MLLM Complex Reasoning Evaluation <i>Meng-Hao Guo, Jiajun Xu, Yi Zhang, Jiayi Song, Haoyang Peng, Yi-Xuan Deng, Xinzhi Dong, Kiyohiro Nakayama, Zhengyang Geng, Chen Wang, Bolin Ni, Guo-Wei Yang, Yongming Rao, Houwen Peng, Han Hu, Gordon Wetzstein, Shi-min Hu</i>	14942

Language Models as Implicit Tree Search <i>Ziliang Chen, Zhao-Rong Lai, Yufeng Yang, Liangda Fang, ZHANFU YANG, Liang Lin</i>	14960
DSP: Dynamic Sequence Parallelism for Multi-Dimensional Transformers <i>Xuanlei Zhao, Shenggan Cheng, Chang Chen, Zangwei Zheng, Ziming Liu, Zheming Yang, Yang You</i>	14982
Asymmetric Decision-Making in Online Knowledge Distillation: Unifying Consensus and Divergence <i>Zhaowei Chen, Borui Zhao, Yuchen Ge, Yuhao Chen, Renjie Song, Jiajun Liang</i>	14997
Multi-Session Budget Optimization for Forward Auction-based Federated Learning <i>Xiaoli Tang, Han Yu, Zengxiang Li, Xiaoxiao Li</i>	15008
PAC Learning with Improvements <i>Idan Attias, Avrim Blum, Keziah Naggita, Donya Saleh, Dravyansh Sharma, Matthew Walter</i>	15020
DyPolySeg: Taylor Series-Inspired Dynamic Polynomial Fitting Network for Few-shot Point Cloud Semantic Segmentation <i>Changshuo Wang, Xiang Fang, Prayag Tiwari</i>	15052
Predictive Performance of Deep Quantum Data Re-uploading Models <i>Xin Wang, Hanxiao Tao, Rebing Wu</i>	15064
Generalizing Causal Effects from Randomized Controlled Trials to Target Populations across Diverse Environments <i>Baohong Li, Yingrong Wang, Anpeng Wu, Ming Ma, Ruoxuan Xiong, Kun Kuang</i>	15098
Minimum Width for Universal Approximation using Squashable Activation Functions <i>Jonghyun Shin, Namjun Kim, Geonho Hwang, Sejun Park</i>	15120
Probing Visual Language Priors in VLMs <i>Tiange Luo, Ang Cao, Gunhee Lee, Justin Johnson, Honglak Lee</i>	15146
Safety Reasoning with Guidelines <i>Haoyu Wang, Zeyu Qin, Li Shen, Xueqian Wang, Dacheng Tao, Minhao Cheng</i>	15183
Be Confident: Uncovering Overfitting in MLLM Multi-Task Tuning <i>Wenke Huang, Jian Liang, Guancheng Wan, Didi Zhu, He Li, Jiawei Shao, Mang Ye, Bo Du, Dacheng Tao</i>	15208
Discrete Markov Probabilistic Models: An Improved Discrete Score-Based Framework with sharp convergence bounds under minimal assumptions <i>Le-Tuyet-Nhi PHAM, Dario Shariatian, Antonio Ocello, Giovanni Conforti, Alain Olivier Durmus</i>	15224
Runtime Analysis of Evolutionary NAS for Multiclass Classification <i>Zeqiong Lv, Chao Qian, Yun Liu, Jiahao Fan, Yanan Sun</i>	15288
On the Dynamic Regret of Following the Regularized Leader: Optimism with History Pruning <i>Naram Mhaisen, George Iosifidis</i>	15307

ParallelComp: Parallel Long-Context Compressor for Length Extrapolation <i>Jing Xiong, Jianghan Shen, Chuanyang Zheng, Zhongwei Wan, Chenyang Zhao, Chiwun Yang, Fanghua Ye, Hongxia Yang, Lingpeng Kong, Ngai Wong</i>	15334
Benchmarking Abstract and Reasoning Abilities Through A Theoretical Perspective <i>Qingchuan Ma, Yuhang Wu, Xiawu Zheng, Rongrong Ji</i>	15353
The Price of Linear Time: Error Analysis of Structured Kernel Interpolation <i>Alexander Moreno, Justin Xiao, Jonathan Mei</i>	15380
What Do Learning Dynamics Reveal About Generalization in LLM Mathematical Reasoning? <i>Katie Kang, Amrith Setlur, Dibya Ghosh, Jacob Steinhardt, Claire Tomlin, Sergey Levine, Aviral Kumar</i>	15409
Representative Language Generation <i>Charlotte Peale, Vinod Raman, Omer Reingold</i>	15423
OpenworldAUC: Towards Unified Evaluation and Optimization for Open-world Prompt Tuning <i>Cong Hua, Qianqian Xu, Zhiyong Yang, Zitai Wang, Shilong Bao, Qingming Huang</i>	15447
Generative Data Mining with Longtail-Guided Diffusion <i>David S Hayden, Mao Ye, Timur Garipov, Gregory P. Meyer, Carl Vondrick, Zhao Chen, Yuning Chai, Eric M Wolff, Siddhartha Srinivasa</i>	15493
How Much Can Transfer? BRIDGE: Bounded Multi-Domain Graph Foundation Model with Generalization Guarantees <i>Haonan Yuan, Qingyun Sun, Junhua Shi, Xingcheng Fu, Bryan Hooi, Jianxin Li, Philip S. Yu</i>	15513
Ladder-Residual: Parallelism-Aware Architecture for Accelerating Large Model Inference with Communication Overlapping <i>Muru Zhang, Mayank Mishra, Zhongzhu Zhou, William Brandon, Jue WANG, Yoon Kim, Jonathan Ragan-Kelley, Shuaiwen Leon Song, Ben Athiwaratkun, Tri Dao</i>	15554
On the Resilience of LLM-Based Multi-Agent Collaboration with Faulty Agents <i>Jen-tse Huang, Jiaxu Zhou, Tailin Jin, Xuhui Zhou, Zixi Chen, Wenxuan Wang, Youliang Yuan, Michael Lyu, Maarten Sap</i>	15568
Beyond Message Passing: Neural Graph Pattern Machine <i>Zehong Wang, Zheyuan Zhang, Tianyi Ma, Nitesh V Chawla, Chuxu Zhang, Yanfang Ye</i>	15593
A Selective Learning Method for Temporal Graph Continual Learning <i>Hanmo Liu, Shimin Di, Haoyang Li, Xun Jian, Yue Wang, Lei Chen</i>	15615
Representation Shattering in Transformers: A Synthetic Study with Knowledge Editing <i>Kento Nishi, Rahul Ramesh, Maya Okawa, Mikail Khona, Hidenori Tanaka, Ekdeep Singh Lubana</i>	15634
Near-optimal Sketchy Natural Gradients for Physics-Informed Neural Networks <i>Maricela Best Mckay, Avleen Kaur, Chen Greif, Brian Wetton</i>	15663

Function-to-Style Guidance of LLMs for Code Translation <i>Longhui Zhang, Bin Wang, Jiahao Wang, Xiaofeng Zhao, Min Zhang, Hao Yang, Meishan Zhang, YU LI, Jing Li, Jun Yu, Min Zhang</i>	15678
Test-Time Training Provably Improves Transformers as In-context Learners <i>Halil Alperen Gozeten, Muhammed Emrullah Ildiz, Xuechen Zhang, Mahdi Soltanolkotabi, Marco Mondelli, Samet Oymak</i>	15694
Forest-of-Thought: Scaling Test-Time Compute for Enhancing LLM Reasoning <i>Zhenni Bi, Kai Han, Chuanjian Liu, Yehui Tang, Yunhe Wang</i>	15724
When Can Proxies Improve the Sample Complexity of Preference Learning? <i>Yuchen Zhu, Daniel Augusto de Souza, Zhengyan Shi, Mengyue Yang, Pasquale Minervini, Matt Kusner, Alexander Nicholas D'Amour</i>	15739
Sparse Training from Random Initialization: Aligning Lottery Ticket Masks using Weight Symmetry <i>Mohammed Adnan, Rohan Jain, Ekansh Sharma, Rahul Krishnan, Yani Ioannou</i>	15764
Robust Multi-Agent Reinforcement Learning with Stochastic Adversary <i>Ziyuan Zhou, Guanjun Liu, Mengchu Zhou, Weiran Guo</i>	15786
An Online Statistical Framework for Out-of-Distribution Detection <i>Xinsong Ma, Xin Zou, Weiwei Liu</i>	15810
CoSER: Coordinating LLM-Based Persona Simulation of Established Roles <i>Xintao Wang, Heng Wang, Yifei Zhang, Xinfeng Yuan, Rui Xu, Jen-tse Huang, Siyu Yuan, Haoran Guo, Jiangjie Chen, Shuchang Zhou, Wei Wang, Yanghua Xiao</i>	15826
Organize the Web: Constructing Domains Enhances Pre-Training Data Curation <i>Alexander Wettig, Kyle Lo, Sewon Min, Hannaneh Hajishirzi, Danqi Chen, Luca Soldaini</i>	15863
An in depth look at the Procrustes-Wasserstein distance: properties and barycenters <i>Davide Adamo, Marco Corneli, Manon Vuillien, Emmanuelle Vila</i>	15891
Pixel2Feature Attack (P2FA): Rethinking the Perturbed Space to Enhance Adversarial Transferability <i>Renpu Liu, Hao Wu, Jiawei Zhang, Xin Cheng, Xiangyang Luo, Bin Ma, Jinwei Wang</i>	15907
In-Context Deep Learning via Transformer Models <i>Weimin Wu, Maojiang Su, Jerry Yao-Chieh Hu, Zhao Song, Han Liu</i>	15925
ADDQ: Adaptive distributional double Q-learning <i>Leif Döring, Benedikt Wille, Maximilian Birr, Mihail Bîrsan, Martin Slowik</i>	15974
Wait-Less Offline Tuning and Re-solving for Online Decision Making <i>Jingruo Sun, Wenzhi Gao, Ellen Vitercik, Yinyu Ye</i>	16021

NeuralCohort: Cohort-aware Neural Representation Learning for Healthcare Analytics <i>Changshuo Liu, Lingze Zeng, Kaiping Zheng, Shaofeng Cai, Beng Chin Ooi, James Wei Luen Yip</i>	16052
Efficient Noise Calculation in Deep Learning-based MRI Reconstructions <i>Onat Dalmaz, Arjun D Desai, Reinhard Heckel, Tolga Cukur, Akshay S Chaudhari, Brian Hargreaves</i>	16071
Demystifying the Paradox of Importance Sampling with an Estimated History-Dependent Behavior Policy in Off-Policy Evaluation <i>Hongyi Zhou, Josiah P. Hanna, Jin Zhu, Ying Yang, Chengchun Shi</i>	16105
Intersectional Fairness in Reinforcement Learning with Large State and Constraint Spaces <i>Eric Eaton, Marcel Hussing, Michael Kearns, Aaron Roth, Sikata Bela Sengupta, Jessica Sorrell</i>	16132
Towards Graph Foundation Models: Learning Generalities Across Graphs via Task-Trees <i>Zehong Wang, Zheyuan Zhang, Tianyi Ma, Nitesh V Chawla, Chuxu Zhang, Yanfang Ye</i>	16148
InfoSEM: A Deep Generative Model with Informative Priors for Gene Regulatory Network Inference <i>Tianyu Cui, Song-Jun Xu, Artem Moskalev, Shuwei Li, Tommaso Mansi, Mangal Prakash, Rui Liao</i>	16186
PolyConf: Unlocking Polymer Conformation Generation through Hierarchical Generative Models <i>Fanmeng Wang, Wentao Guo, Qi Ou, Hongshuai Wang, Haitao Lin, Hongteng Xu, Zhifeng Gao</i>	16202
A Model of Place Field Reorganization During Reward Maximization <i>M Ganesh Kumar, Blake Bordelon, Jacob A Zavatore-Veth, Cengiz Pehlevan</i>	16216
Accelerated Diffusion Models via Speculative Sampling <i>Valentin De Bortoli, Alexandre Galashov, Arthur Gretton, Arnaud Doucet</i>	16254
Context Matters: Query-aware Dynamic Long Sequence Modeling of Gigapixel Images <i>Zhengrui Guo, Qichen Sun, Jiabo MA, Lishuang Feng, Jinzhuo Wang, Hao Chen</i>	16296
TimePoint: Accelerated Time Series Alignment via Self-Supervised Keypoint and Descriptor Learning <i>Ron Shapira Weber, Shahar benishay, Andrey Lavrinenko, Shahaf E. Finder, Oren Freifeld</i>	16319
Topological Signatures of Adversaries in Multimodal Alignments <i>Minh N. Vu, Geigh Zollicoffer, Huy Mai, Ben Nebgen, Boian Alexandrov, Manish Bhattarai</i>	16344
A General Representation-Based Approach to Multi-Source Domain Adaptation <i>Ignavier Ng, Yan Li, Zijian Li, Yujia Zheng, Guangyi Chen, Kun Zhang</i>	16368
MetricEmbedding: Accelerate Metric Nearness by Tropical Inner Product <i>Muyang Cao, Jiajun Yu, Xin Du, Gang Pan, Wei Wang</i>	16391
Collapse or Thrive: Perils and Promises of Synthetic Data in a Self-Generating World <i>Joshua Kazdan, Rylan Schaeffer, Apratim Dey, Matthias Gerstgrasser, Rafael Rafailov, David L. Donoho, Sanmi Koyejo</i>	16415

Multi-Stage Manipulation with Demonstration-Augmented Reward, Policy, and World Model Learning <i>Adrià López Escoriza, Nicklas Hansen, Stone Tao, Tongzhou Mu, Hao Su</i>	16441
Vision-Language Model Selection and Reuse for Downstream Adaptation <i>Hao-Zhe Tan, Zhi Zhou, Yu-feng Li, Lan-Zhe Guo</i>	16463
A Machine Learning Approach to Duality in Statistical Physics <i>Prateek Gupta, Andrea E. V. Ferrari, Nabil Iqbal</i>	16476
Tilted Sharpness-Aware Minimization <i>Tian Li, Tianyi Zhou, Jeff Bilmes</i>	16496
AdvAgent: Controllable Blackbox Red-teaming on Web Agents <i>Chejian Xu, Mintong Kang, Jiawei Zhang, Zeyi Liao, Lingbo Mo, Mengqi Yuan, Huan Sun, Bo Li</i>	16516
Strategic A/B testing via Maximum Probability-driven Two-armed Bandit <i>Yu Zhang, Shanshan Zhao, Bokui Wan, Jinjuan Wang, Xiaodong Yan</i>	16529
CoCoA-Mix: Confusion-and-Confidence-Aware Mixture Model for Context Optimization <i>Dasol Hong, Wooju Lee, Hyun Myung</i>	16550
Tensor-Var: Efficient Four-Dimensional Variational Data Assimilation <i>Yiming Yang, Xiaoyuan Cheng, Daniel Giles, Sibo Cheng, Yi He, Xiao Xue, Boli Chen, Yukun Hu</i>	16572
A Meta-learner for Heterogeneous Effects in Difference-in-Differences <i>Hui Lan, Haoge Chang, Eleanor Wiske Dillon, Vasilis Syrgkanis</i>	16603
Diffusion models for Gaussian distributions: Exact solutions and Wasserstein errors <i>Emile Pierret, Bruno Galerne</i>	16635
Heterogeneous Label Shift: Theory and Algorithm <i>Chao Xu, Xijia Tang, Chenping Hou</i>	16662
iN2V: Bringing Transductive Node Embeddings to Inductive Graphs <i>Nicolas Lell, Ansgar Scherp</i>	16682
Speculative Prefill: Turbocharging TTFT with Lightweight and Training-Free Token Importance Estimation <i>Jingyu Liu, Beidi Chen, Ce Zhang</i>	16705
Revisiting Noise Resilience Strategies in Gesture Recognition: Short-Term Enhancement in sEMG Analysis <i>Weiye Guo, Ziyue Qiao, Ying Sun, Yijie Xu, Hui Xiong</i>	16727
Right Now, Wrong Then: Non-Stationary Direct Preference Optimization under Preference Drift <i>Seongho Son, William Bankes, Sayak Ray Chowdhury, Brooks Paige, Ilija Bogunovic</i>	16745

Arrow: Accelerator for Time Series Causal Discovery with Time Weaving <i>Yuanyuan Yao, Yuan Dong, Lu Chen, Kun Kuang, Ziquan Fang, Cheng Long, Yunjun Gao, TIANYI LI</i>	16776
LLaVA-ReID: Selective Multi-image Questioner for Interactive Person Re-Identification <i>Yiding Lu, Mouxing Yang, Dezhong Peng, Peng Hu, Yijie Lin, Xi Peng</i>	16795
Defending LVLMS Against Vision Attacks Through Partial-Perception Supervision <i>Qi Zhou, Dongxia Wang, Tianlin Li, Yun Lin, Yang Liu, Jin Song Dong, Qing Guo</i>	16815
Partially Observable Reinforcement Learning with Memory Traces <i>Onno Eberhard, Michael Muehlebach, Claire Vernade</i>	16842
Detecting Strategic Deception with Linear Probes <i>Nicholas Goldowsky-Dill, Bilal Chughtai, Stefan Heimersheim, Marius Hobbhahn</i>	16858
Compression via Pre-trained Transformers: A Study on Byte-Level Multimodal Data <i>David Heurtel-Depeiges, Anian Ruoss, Joel Veness, Tim Genewein</i>	16890
Clients Collaborate: Flexible Differentially Private Federated Learning with Guaranteed Improvement of Utility-Privacy Trade-off <i>Yuecheng Li, Lele Fu, Tong Wang, Jian Lou, Bin Chen, Lei Yang, Jian Shen, Zibin Zheng, Chuan Chen</i>	16906
Which Attention Heads Matter for In-Context Learning? <i>Kayo Yin, Jacob Steinhardt</i>	16928
Devil is in the Details: Density Guidance for Detail-Aware Generation with Flow Models <i>Rafal Karczewski, Markus Heinonen, Vikas K Garg</i>	16962
Best Subset Selection: Optimal Pursuit for Feature Selection and Elimination <i>Zhihan Zhu, Yanhao Zhang, Yong Xia</i>	16992
Sable: a Performant, Efficient and Scalable Sequence Model for MARL <i>Omayma Mahjoub, Sasha Abramowitz, Ruan John de Kock, Wiem Khelifi, Simon Verster Du Toit, Jemma Daniel, Louay Ben Nessir, Louise Beyers, Juan Claude Formanek, Liam Clark, Arnu Pretorius</i>	17021
Optimal Sensor Scheduling and Selection for Continuous-Discrete Kalman Filtering with Auxiliary Dynamics <i>Mohamad Al Ahdab, John Leth, Zheng-Hua Tan</i>	17057
The Illusion of Role Separation: Hidden Shortcuts in LLM Role Learning (and How to Fix Them) <i>Zihao Wang, Yibo Jiang, Jiahao Yu, Heqing Huang</i>	17084
Bellman Unbiasedness: Toward Provably Efficient Distributional Reinforcement Learning with General Value Function Approximation <i>Taehyun Cho, Seungyub Han, Seokhun Ju, Dohyeong Kim, Kyungjae Lee, Jungwoo Lee</i>	17095
Sample Complexity of Branch-length Estimation by Maximum Likelihood <i>David Clancy Jr., Hanbaek Lyu, Sebastien Roch</i>	17120

ReVISE: Learning to Refine at Test-Time via Intrinsic Self-Verification <i>Hyunseok Lee, Seunghyuk Oh, Jaehyung Kim, Jinwoo Shin, Jihoon Tack</i>	17142
Learning Policy Committees for Effective Personalization in MDPs with Diverse Tasks <i>Luise Ge, Michael Lanier, Anindya Sarkar, Bengisu Guresti, Chongjie Zhang, Yevgeniy Vorobeychik</i>	17161
VIP: Vision Instructed Pre-training for Robotic Manipulation <i>Zhuoling Li, LiangLiang Ren, Jinrong Yang, Yong Zhao, Xiaoyang Wu, Zhenhua Xu, Xiang Bai, Hengshuang Zhao</i>	17184
OR-Bench: An Over-Refusal Benchmark for Large Language Models <i>Justin Cui, Wei-Lin Chiang, Ion Stoica, Cho-Jui Hsieh</i>	17194
From Black Boxes to Transparent Minds: Evaluating and Enhancing the Theory of Mind in Multimodal Large Language Models <i>Xinyang Li, Siqi Liu, Bochao Zou, Jiansheng Chen, Huimin Ma</i>	17222
Improving Soft Unification with Knowledge Graph Embedding Methods <i>Xuanming Cui, Chionh Wei Peng, Adriel Kuek, Ser-Nam Lim</i>	17246
Leveraging Predictive Equivalence in Decision Trees <i>Hayden McTavish, Zachery Boner, Jon Donnelly, Margo Seltzer, Cynthia Rudin</i>	17266
Emergent Response Planning in LLMs <i>Zhichen Dong, Zhanhui Zhou, Zhixuan Liu, Chao Yang, Chaochao Lu</i>	17302
Bounded Rationality for LLMs: Satisficing Alignment at Inference-Time <i>Mohamad Fares El Hajj Chehade, Soumya Suvra Ghosal, Souradip Chakraborty, Avinash Reddy, Dinesh Manocha, Hao Zhu, Amrit Singh Bedi</i>	17322
Controlling Large Language Model with Latent Action <i>Chengxing Jia, Ziniu Li, Pengyuan Wang, Yi-Chen Li, Zhenyu Hou, Yuxiao Dong, Yang Yu</i>	17339
Balancing Efficiency and Expressiveness: Subgraph GNNs with Walk-Based Centrality <i>Joshua Southern, Yam Eitan, Guy Bar-Shalom, Michael M. Bronstein, Haggai Maron, Fabrizio Frasca</i>	17381
Improving LLMs for Recommendation with Out-Of-Vocabulary Tokens <i>Ting-Ji Huang, Jia-Qi Yang, Chunxu Shen, Kai-Qi Liu, De-Chuan Zhan, Han-Jia Ye</i>	17416
SAN: Hypothesizing Long-Term Synaptic Development and Neural Engram Mechanism in Scalable Model's Parameter-Efficient Fine-Tuning <i>Gaole Dai, Chun-Kai Fan, Yiming Tang, Zhi Zhang, Yuan Zhang, Yulu Gan, Qizhe Zhang, Cheng-Ching Tseng, Shanghang Zhang, Tiejun Huang</i>	17433
Theoretical Limitations of Ensembles in the Age of Overparameterization <i>Niclas Dorn, John Patrick Cunningham, Geoff Pleiss</i>	17458
Comparing Comparisons: Informative and Easy Human Feedback with Distinguishability Queries <i>Xuening Feng, Zhaohui JIANG, Timo Kaufmann, Eyke Hüllermeier, Paul Weng, Yifei Zhu</i>	17491

A Computationally Efficient Algorithm for Infinite-Horizon Average-Reward Linear MDPs <i>Kihyuk Hong, Ambuj Tewari</i>	17508
Playmate: Flexible Control of Portrait Animation via 3D-Implicit Space Guided Diffusion <i>Xingpei Ma, Jiaran Cai, Yuansheng Guan, Shenneng Huang, Qiang Zhang, Shunsi Zhang</i>	17531
Momentum-Driven Adaptivity: Towards Tuning-Free Asynchronous Federated Learning <i>Wenjing Yan, Xiangyu Zhong, Xiaolu Wang, Ying-Jun Angela Zhang</i>	17547
Evaluating Judges as Evaluators: The JETTS Benchmark of LLM-as-Judges as Test-Time Scaling Evaluators <i>Yilun Zhou, Austin Xu, PeiFeng Wang, Caiming Xiong, Shafiq Joty</i>	17579
On the Similarities of Embeddings in Contrastive Learning <i>Chungpa Lee, Sehee Lim, Kibok Lee, Jy-yong Sohn</i>	17615
Peri-LN: Revisiting Normalization Layer in the Transformer Architecture <i>Jeonghoon Kim, Byeongchan Lee, Cheonbok Park, Yeontaek Oh, Beomjun Kim, Taehwan Yoo, Seongjin Shin, Dongyoon Han, Jinwoo Shin, Kang Min Yoo</i>	17646
Tensor Product Neural Networks for Functional ANOVA Model <i>Seokhun Park, Insung Kong, yongchan Choi, Chanmoo Park, Yongdai Kim</i>	17683
A Comprehensive Framework for Analyzing the Convergence of Adam: Bridging the Gap with SGD <i>Ruinan Jin, Xiao Li, Yaoliang Yu, Baoxiang Wang</i>	17728
Empower Structure-Based Molecule Optimization with Gradient Guided Bayesian Flow Networks <i>Keyue Qiu, Yuxuan Song, Jie Yu, Hongbo Ma, Ziyao Cao, Zhilong Zhang, Yushuai Wu, Mingyue Zheng, Hao Zhou, Wei-Ying Ma</i>	17780
Incentivize without Bonus: Provably Efficient Model-based Online Multi-agent RL for Markov Games <i>Tong Yang, Bo Dai, Lin Xiao, Yuejie Chi</i>	17807
EvFocus: Learning to Reconstruct Sharp Images from Out-of-Focus Event Streams <i>Lin Zhu, Xiantao Ma, Xiao Wang, Lizhi Wang, Hua Huang</i>	17848
LLMs can see and hear without any training <i>Kumar Ashutosh, Yossi Gandelsman, Xinlei Chen, Ishan Misra, Rohit Girdhar</i>	17870
EVOLvE: Evaluating and Optimizing LLMs For In-Context Exploration <i>Allen Nie, Yi Su, Bo Chang, Jonathan Lee, Ed H. Chi, Quoc V Le, Minmin Chen</i>	17885
NeuroTree: Hierarchical Functional Brain Pathway Decoding for Mental Health Disorders <i>Jun-En Ding, Dongsheng Luo, Chenwei Wu, Feng Liu</i>	17916

On the Alignment between Fairness and Accuracy: from the Perspective of Adversarial Robustness <i>Junyi Chai, Taeuk Jang, Jing Gao, Xiaoqian Wang</i>	17941
Learning Fused State Representations for Control from Multi-View Observations <i>Zeyu Wang, Yao-Hui Li, Xin Li, Hongyu Zang, Romain Laroche, Riashat Islam</i>	17965
Learning from True-False Labels via Multi-modal Prompt Retrieving <i>Zhongnian Li, Jinghao Xu, Peng Ying, Meng Wei, Xinzheng Xu</i>	17964
An Interpretable N-gram Perplexity Threat Model for Large Language Model Jailbreaks <i>Valentyn Boreiko, Alexander Panfilov, Vaclav Voracek, Matthias Hein, Jonas Geiping</i>	17979
LaRA: Benchmarking Retrieval-Augmented Generation and Long-Context LLMs – No Silver Bullet for LC or RAG Routing <i>Kuan Li, Liwen Zhang, Yong Jiang, Pengjun Xie, Fei Huang, Shuai Wang, Minhao Cheng</i>	18007
Bootstrapping Self-Improvement of Language Model Programs for Zero-Shot Schema Matching <i>Nabeel Seedat, Mihaela van der Schaar</i>	18029
Cape: Context-Aware Prompt Perturbation Mechanism with Differential Privacy <i>Haoqi Wu, Wei Dai, Wang Li, Qiang Yan</i>	18065
VCT: Training Consistency Models with Variational Noise Coupling <i>Gianluigi Silvestri, Luca Ambrogioni, Chieh-Hsin Lai, Yuhta Takida, Yuki Mitsufuji</i>	18083
Spurious Correlations in High Dimensional Regression: The Roles of Regularization, Simplicity Bias and Over-Parameterization <i>Simone Bombari, Marco Mondelli</i>	18110
An Augmentation-Aware Theory for Self-Supervised Contrastive Learning <i>Jingyi Cui, Hongwei Wen, Yisen Wang</i>	18145
EARTH: Epidemiology-Aware Neural ODE with Continuous Disease Transmission Graph <i>Guancheng Wan, Zewen Liu, Xiaojun Shan, Max SY Lau, B. Aditya Prakash, Wei Jin</i>	18164
Optimizing Robustness and Accuracy in Mixture of Experts: A Dual-Model Approach <i>Xu Zhang, Kaidi Xu, Ziqing Hu, Ren Wang</i>	18177
Unsupervised Learning for Class Distribution Mismatch <i>Pan Du, Wangbo Zhao, Xinai Lu, Nian Liu, Zhikai Li, Chaoyu Gong, Suyun Zhao, Hong Chen, Cuiping Li, Kai Wang, Yang You</i>	18192
The Impact of On-Policy Parallelized Data Collection on Deep Reinforcement Learning Networks <i>Walter Mayor, Johan Obando-Ceron, Aaron Courville, Pablo Samuel Castro</i>	18229
Gaussian Mixture Flow Matching Models <i>Hansheng Chen, Kai Zhang, Hao Tan, Zexiang Xu, Fujun Luan, Leonidas Guibas, Gordon Wetzstein, Sai Bi</i>	18251

Cache Me If You Must: Adaptive Key-Value Quantization for Large Language Models <i>Alina Shutova, Vladimir Malinovskii, Vage Egiazarian, Denis Kuznedelev, Denis Mazur, Surkov Nikita, Ivan Ermakov, Dan Alistarh</i>	18271
Tracking Most Significant Shifts in Infinite-Armed Bandits <i>Joe Suk, Jung-hun Kim</i>	18294
Maximum Entropy Reinforcement Learning with Diffusion Policy <i>Xiaoyi Dong, Jian Cheng, Xi Sheryl Zhang</i>	18319
Riemann Tensor Neural Networks: Learning Conservative Systems with Physics-Constrained Networks <i>Anas Jnini, Lorenzo Breschi, Flavio Vella</i>	18340
Training Software Engineering Agents and Verifiers with SWE-Gym <i>Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, Yizhe Zhang</i>	18363
Conformal Anomaly Detection in Event Sequences <i>Shuai Zhang, Chuan Zhou, Yang Liu, Peng Zhang, Xixun Lin, Shirui Pan</i>	18384
Compressed Image Generation with Denoising Diffusion Codebook Models <i>Guy Ohayon, Hila Manor, Tomer Michaeli, Michael Elad</i>	18401
Reward-Guided Prompt Evolving in Reinforcement Learning for LLMs <i>Ziyu Ye, Rishabh Agarwal, Tianqi Liu, Rishabh Joshi, Sarmishta Velury, Quoc V Le, Qijun Tan, Yuan Liu</i>	18447
SWAN: SGD with Normalization and Whitening Enables Stateless LLM Training <i>Chao Ma, Wenbo Gong, Meyer Scetbon, Edward Meeds</i>	18475
Not All Tokens Matter All The Time: Dynamic Token Aggregation Towards Efficient Detection Transformers <i>Jiacheng Cheng, Xiwen Yao, Xiang Yuan, Junwei Han</i>	18511
Chip Placement with Diffusion Models <i>Vint Lee, Minh Nguyen, Leena Elzeiny, Chun Deng, Pieter Abbeel, John Wawrzynek</i>	18526
Reinforcement Learning with Random Time Horizons <i>Enric Ribera Borrell, Lorenz Richter, Christof Schuette</i>	18542
Guided Zeroth-Order Methods for Stochastic Non-convex Problems with Decision-Dependent Distributions <i>Yuya Hikima, Hiroshi Sawada, Akinori Fujino</i>	18565
CaDA: Cross-Problem Routing Solver with Constraint-Aware Dual-Attention <i>Han Li, Fei Liu, Zhi Zheng, Yu Zhang, Zhenkun Wang</i>	18591
On the Provable Separation of Scales in Maximal Update Parameterization <i>Letong Hong, Zhangyang Wang</i>	18610

Faster Stochastic Optimization with Arbitrary Delays via Adaptive Asynchronous Mini-Batching <i>Amit Attia, Ofir Gaash, Tomer Koren</i>	18623
The Limits of Tractable Marginalization <i>Oliver Broadrick, Sanyam Agarwal, Guy Van den Broeck, Markus Bläser</i>	18642
Geometric Feature Embedding for Effective 3D Few-Shot Class Incremental Learning <i>Xiangqi Li, Libo Huang, Zhulin An, Weilun Feng, Chuanguang Yang, Boyu Diao, Fei Wang, Yongjun Xu</i>	18658
Gradient Descent Converges Arbitrarily Fast for Logistic Regression via Large and Adaptive Stepsizes <i>Ruiqi Zhang, Jingfeng Wu, Peter Bartlett</i>	18672
InfoCons: Identifying Interpretable Critical Concepts in Point Clouds via Information Theory <i>Feifei Li, Mi Zhang, Zhaoxiang Wang, Min Yang</i>	18696
Comparing Few to Rank Many: Active Human Preference Learning Using Randomized Frank-Wolfe Method <i>Kiran Koshy Thekumparampil, Gaurush Hiranandani, Kousha Kalantari, Shoham Sabach, Branislav Kveton</i>	18714
Sample Efficient Demonstration Selection for In-Context Learning <i>Kiran Purohit, Venkatesh V, Sourangshu Bhattacharya, Avishek Anand</i>	18736
Tuning LLM Judge Design Decisions for 1/1000 of the Cost <i>David Salinas, Omar Swelam, Frank Hutter</i>	18760
More Than Meets the Eye: Enhancing Multi-Object Tracking Even with Prolonged Occlusions <i>Bishoy Galoaa, Somaieh Amraee, Sarah Ostadabbas</i>	18777
Nonparametric Identification of Latent Concepts <i>Yujia Zheng, Shaoan Xie, Kun Zhang</i>	18788
Faster Approximation Algorithms for k-Center via Data Reduction <i>Arnold Filtser, Shaofeng H.-C. Jiang, Yi Li, Anurag Murty Naredla, Ioannis Psarros, Qiaoyuan Yang, Qin Zhang</i>	18819
Learning Utilities from Demonstrations in Markov Decision Processes <i>Filippo Lazzati, Alberto Maria Metelli</i>	18833
Learning Input Encodings for Kernel-Optimal Implicit Neural Representations <i>Zhemini Li, Liyuan Ma, Hongxia Wang, Yaoyun Zeng, Xiaolong Han</i>	18900
CERTAIN: Context Uncertainty-aware One-Shot Adaptation for Context-based Offline Meta Reinforcement Learning <i>Hongtu Zhou, Ruiling Yang, Yakun Zhu, Haoqi Zhao, Hai Zhang, Di Zhang, Junqiao Zhao, Chen Ye, Changjun Jiang</i>	18917

Generalization Performance of Ensemble Clustering: From Theory to Algorithm <i>Xu Zhang, Haoye Qiu, Weixuan Liang, Hui LIU, Junhui Hou, Yuheng Jia</i>	18936
LapSum - One Method to Differentiate Them All: Ranking, Sorting and Top-k Selection <i>Łukasz Struski, Michal B. Bednarczyk, Igor T. Podolak, Jacek Tabor</i>	18963
SketchDNN: Joint Continuous-Discrete Diffusion for CAD Sketch Generation <i>Sathvik Reddy Chereddy, John Femiani</i>	18981
Near-Optimal Sample Complexity for MDPs via Anchoring <i>Jongmin Lee, Mario Bravo, Roberto Cominetti</i>	18998
LOB-Bench: Benchmarking Generative AI for Finance - an Application to Limit Order Book Data <i>Peer Nagy, Sascha Yves Frey, Kang Li, Bidipta Sarkar, Svitlana Vyetenko, Stefan Zohren, Ani Calinescu, Jakob Nicolaus Foerster</i>	19021
EditLord: Learning Code Transformation Rules for Code Editing <i>Weichen Li, Albert Jan, Baishakhi Ray, Junfeng Yang, Chengzhi Mao, Kexin Pei</i>	19045
Rethinking Aleatoric and Epistemic Uncertainty <i>Freddie Bickford Smith, Jannik Kossen, Eleanor Trollope, Mark van der Wilk, Adam Foster, Tom Rainforth</i>	19073
Curse of High Dimensionality Issue in Transformer for Long Context Modeling <i>Shuhai Zhang, Zeng You, Yaofu Chen, Zhiquan Wen, Qianyu Wang, Zhijie Qiu, Yuanqing Li, Minghui Tan</i>	19088
Portable Reward Tuning: Towards Reusable Fine-Tuning across Different Pretrained Models <i>Daiki Chijiwa, Taku Hasegawa, Kyosuke Nishida, Kuniko Saito, Susumu Takeuchi</i>	19113
Leveraging Online Olympiad-Level Math Problems for LLMs Training and Contamination-Resistant Evaluation <i>Sadegh Mahdavi, Muchen Li, Kaiwen Liu, Christos Thrampoulidis, Leonid Sigal, Renjie Liao</i>	19134
ResearchTown: Simulator of Human Research Community <i>Haofei Yu, Zhaochen Hong, Zirui Cheng, Kunlun Zhu, Keyang Xuan, Jinwei Yao, Tao Feng, Jiaxuan You</i>	19159
Revisiting Non-Acyclic GFlowNets in Discrete Environments <i>Nikita Morozov, Ian Maksimov, Daniil Tiapkin, Sergey Samsonov</i>	19205
Deep Sturm–Liouville: From Sample-Based to 1D Regularization with Learnable Orthogonal Basis Functions <i>David Vigouroux, Joseba Dalmau, Louis Béthune, Victor Boutin</i>	19229
Improved Discretization Complexity Analysis of Consistency Models: Variance Exploding Forward Process and Decay Discretization Scheme <i>Ruofeng Yang, Bo Jiang, Cheng Chen, Shuai Li</i>	19247

MultiPDENet: PDE-embedded Learning with Multi-time-stepping for Accelerated Flow Simulation <i>Qi Wang, Yuan Mi, Wang Haoyun, Yi Zhang, Ruizhi Chengze, Hongsheng Liu, Ji-Rong Wen, Hao Sun</i>	19267
Linear Q-Learning Does Not Diverge in L^2: Convergence Rates to a Bounded Set <i>Xinyu Liu, Zixuan Xie, Shangtong Zhang</i>	19290
Sanity Checking Causal Representation Learning on a Simple Real-World System <i>Juan L. Gamella, Simon Bing, Jakob Runge</i>	19317
Adjustment for Confounding using Pre-Trained Representations <i>Rickmer Schulte, David Rügamer, Thomas Nagler</i>	19344
Grokking at the Edge of Linear Separability <i>Alon Beck, Noam Itzhak Levi, Yohai Bar-Sinai</i>	19368
Ranking with Multiple Oracles: From Weak to Strong Stochastic Transitivity <i>Tao Jin, Yue Wu, Quanquan Gu, Farzad Farnoud</i>	19396
CHATS: Combining Human-Aligned Optimization and Test-Time Sampling for Text-to-Image Generation <i>Minghao Fu, Guo-Hua Wang, Liangfu Cao, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang</i>	19415
DINO-WM: World Models on Pre-trained Visual Features enable Zero-shot Planning <i>Gaoyue Zhou, Hengkai Pan, Yann LeCun, Lerrel Pinto</i>	19438
ELMO : Efficiency via Low-precision and Peak Memory Optimization in Large Output Spaces <i>Jinbin Zhang, Nasib Ullah, Erik Schultheis, Rohit Babbar</i>	19459
Provably Cost-Sensitive Adversarial Defense via Randomized Smoothing <i>Yuan Xin, Dingfan Chen, Michael Backes, Xiao Zhang</i>	19475
Pareto Merging: Multi-Objective Optimization for Preference-Aware Model Merging <i>Weiyu Chen, James Kwok</i>	19497
SkipGPT: Each Token is One of a Kind <i>Anhao Zhao, Fanghua Ye, Yingqi Fan, Junlong Tong, Jing Xiong, Zhiwei Fei, Hui Su, Xiaoyu Shen</i>	19516
Multi-Timescale Dynamics Model Bayesian Optimization for Plasma Stabilization in Tokamaks <i>Rohit Sonker, Alexandre Capone, Andrew Rothstein, Hiro Josep Farre Kaga, Egemen Kolemen, Jeff Schneider</i>	19536
BAnG: Bidirectional Anchored Generation for Conditional RNA Design <i>Roman Klypa, Alberto Bietti, Sergei Grudinin</i>	19551
Online Episodic Convex Reinforcement Learning <i>Bianca Marin Moreno, Khaled Eldowa, Pierre Gaillard, Margaux Brégère, Nadia Oudjane</i>	19575

Distilling the Knowledge in Data Pruning <i>Emanuel Ben Baruch, Adam Botach, Igor Kviatkovsky, Manoj Aggarwal, Gerard Medioni</i>	19625
RobustZero: Enhancing MuZero Reinforcement Learning Robustness to State Perturbations <i>Yushuai Li, Hengyu Liu, Torben Bach Pedersen, Yuqiang He, Kim Guldstrand Larsen, Lu Chen, Christian S. Jensen, Jiachen Xu, TIANYI LI</i>	19643
From Feature Interaction to Feature Generation: A Generative Paradigm of CTR Prediction Models <i>Mingjia Yin, Junwei Pan, Hao Wang, Ximei Wang, Shangyu Zhang, Jie Jiang, Defu Lian, Enhong Chen</i>	19667
Provable In-Context Vector Arithmetic via Retrieving Task Concepts <i>Dake Bu, Wei Huang, Andi Han, Atsushi Nitanda, Qingfu Zhang, Hau-San Wong, Taiji Suzuki</i>	19688
TRUST-VLM: Thorough Red-Teaming for Uncovering Safety Threats in Vision-Language Models <i>Kangjie Chen, Li Muyang, Guanlin Li, Shudong Zhang, Shangwei Guo, Tianwei Zhang</i>	19744
Multilayer Matrix Factorization via Dimension-Reducing Diffusion Variational Inference <i>Junbin Liu, Farzan Farnia, Wing-Kin Ma</i>	19778
Metadata Conditioning Accelerates Language Model Pre-training <i>Tianyu Gao, Alexander Wettig, Luxi He, Yihe Dong, Sathika Malladi, Danqi Chen</i>	19797
Visual Abstraction: A Plug-and-Play Approach for Text-Visual Retrieval <i>Guofeng Ding, Yiding Lu, Peng Hu, Mouxing Yang, Yijie Lin, Xi Peng</i>	19815
Scaling Laws for Differentially Private Language Models <i>Ryan McKenna, Yangsibo Huang, Amer Sinha, Borja Balle, Zachary Charles, Christopher A. Choquette-Choo, Badih Ghazi, Georgios Kaissis, Ravi Kumar, Ruibo Liu, Da Yu, Chiyuan Zhang</i>	19835
Streamline Without Sacrifice - Squeeze out Computation Redundancy in LMM <i>Penghao Wu, Lewei Lu, Ziwei Liu</i>	19859
FlexiReID: Adaptive Mixture of Expert for Multi-Modal Person Re-Identification <i>Zhen Sun, Lei Tan, Yunhang Shen, Chengmao Cai, Xing Sun, Pingyang Dai, Liujuan Cao, Rongrong Ji</i>	19874
Bayesian Basis Function Approximation for Scalable Gaussian Process Priors in Deep Generative Models <i>Mehmet Yiğit Balık, Maksim Sinelnikov, Priscilla Ong, Harri Lähdesmäki</i>	19888
Fixing the Loose Brake: Exponential-Tailed Stopping Time in Best Arm Identification <i>Kapilan Balagopalan, Ngo Tuan Nguyen, Yao Zhao, Kwang-Sung Jun</i>	19912
Generalization Principles for Inference over Text-Attributed Graphs with Large Language Models <i>Haoyu Peter Wang, Shikun Liu, Rongzhe Wei, Pan Li</i>	19955
Auto-reconfiguration for Latency Minimization in CPU-based DNN Serving <i>Ankit Bhardwaj, Amar Phanishayee, Deepak Narayanan, Ryan Stutsman</i>	19977

FlexTok: Resampling Images into 1D Token Sequences of Flexible Length <i>Roman Bachmann, Jesse Allardice, David Mizrahi, Enrico Fini, Oğuzhan Fatih Kar, Elmira Amirloo, Alaaeldin El-Nouby, Amir Zamir, Afshin Dehghan</i>	19992
EmbodiedBench: Comprehensive Benchmarking Multi-modal Large Language Models for Vision-Driven Embodied Agents <i>Rui Yang, Hanyang Chen, Junyu Zhang, Mark Zhao, Cheng Qian, Kangrui Wang, Qineng Wang, Teja Venkat Koripella, Marziyeh Movahedi, Manling Li, Heng Ji, Huan Zhang, Tong Zhang</i>	20044
Angle Domain Guidance: Latent Diffusion Requires Rotation Rather Than Extrapolation <i>Cheng Jin, Zhenyu Xiao, Chutao Liu, Yuantao Gu</i>	20100
Testing Conditional Mean Independence Using Generative Neural Networks <i>Yi Zhang, Linjun Huang, Yun Yang, Xiaofeng Shao</i>	20126
Towards World Simulator: Crafting Physical Commonsense-Based Benchmark for Video Generation <i>Fanqing Meng, Jiaqi Liao, Xinyu Tan, Quanfeng Lu, Wenqi Shao, Kaipeng Zhang, Yu Cheng, Dianqi Li, Ping Luo</i>	20156
Identifiable Object Representations under Spatial Ambiguities <i>Avinash Kori, Francesca Toni, Ben Glocker</i>	20182
Knowledge Retention in Continual Model-Based Reinforcement Learning <i>Haotian Fu, Yixiang Sun, Michael Littman, George Konidaris</i>	20215
PhantomWiki: On-Demand Datasets for Reasoning and Retrieval Evaluation <i>Albert Gong, Kamilė Stankevičiūtė, Chao Wan, Anmol Kabra, Raphael Thesmar, Johann Lee, Julius Klenke, Carla P Gomes, Kilian Q Weinberger</i>	20235
Learning Bayesian Nash Equilibrium in Auction Games via Approximate Best Response <i>Kexin Huang, Ziqian Chen, Xue Wang, Chongming Gao, Jinyang Gao, Bolin Ding, Xiang Wang</i>	20267
Efficient Robotic Policy Learning via Latent Space Backward Planning <i>Dongxiu Liu, Haoyi Niu, Zhihao Wang, Jinliang Zheng, Yinan Zheng, Zhonghong Ou, Jianming HU, Jianxiong Li, Xianyu Zhan</i>	20293
Train for the Worst, Plan for the Best: Understanding Token Ordering in Masked Diffusions <i>Jaeyeon Kim, Kulin Shah, Vasilis Kontonis, Sham M. Kakade, Sitan Chen</i>	20308
Neural Event-Triggered Control with Optimal Scheduling <i>Luan Yang, Jingdong Zhang, Qunxi Zhu, Wei Lin</i>	20328
Exploiting Curvature in Online Convex Optimization with Delayed Feedback <i>Hao Qiu, Emmanuel Esposito, Mengxiao Zhang</i>	20352
Certifiably Robust Model Evaluation in Federated Learning under Meta-Distributional Shifts <i>Amir Najafi, Samin Mahdizadeh Sani, Farzan Farnia</i>	20384

Learning to Trust Bellman Updates: Selective State-Adaptive Regularization for Offline RL <i>Qin-Wen Luo, Ming-Kun Xie, Ye-Wen Wang, Sheng-Jun Huang</i>	20420
ShieldAgent: Shielding Agents via Verifiable Safety Policy Reasoning <i>Zhaorun Chen, Mintong Kang, Bo Li</i>	20438
StealthInk: A Multi-bit and Stealthy Watermark for Large Language Models <i>Ya Jiang, Chuxiong Wu, Massieh Kordi Boroujeny, Brian Mark, Kai Zeng</i>	20470
General agents need world models <i>Jonathan Richens, Tom Everitt, David Abel</i>	20495
How Far Is Video Generation from World Model: A Physical Law Perspective <i>Bingyi Kang, Yang Yue, Rui Lu, Zhijie Lin, Yang Zhao, Kaixin Wang, Gao Huang, Jiashi Feng</i>	20524
Physics-Informed Weakly Supervised Learning For Interatomic Potentials <i>Makoto Takamoto, Viktor Zaverkin, Mathias Niepert</i>	20551
CollabLLM: From Passive Responders to Active Collaborators <i>Shirley Wu, Michel Galley, Baolin Peng, Hao Cheng, Gavin Li, Yao Dou, Weixin Cai, James Zou, Jure Leskovec, Jianfeng Gao</i>	20580
Measuring Diversity in Synthetic Datasets <i>Yuchang Zhu, Huizhe Zhang, Bingzhe Wu, Jintang Li, Zibin Zheng, Peilin Zhao, Liang Chen, Yatao Bian</i>	20604
Graph-Based Algorithms for Diverse Similarity Search <i>Piyush Anand, Piotr Indyk, Ravishankar Krishnaswamy, Sepideh Mahabadi, Vikas C. Raykar, Kirankumar Shiragur, Haike Xu</i>	20629
Spectral-Aware Reservoir Computing for Fast and Accurate Time Series Classification <i>Shikang Liu, Chuyang Wei, Xiren Zhou, Huanhuan Chen</i>	20651
Towards flexible perception with visual memory <i>Robert Geirhos, Priyank Jaini, Austin Stone, Sourabh Medapati, Xi Yi, George Toderici, Abhijit Ogale, Jonathon Shlens</i>	20666
Online Differentially Private Conformal Prediction for Uncertainty Quantification <i>Qiangqiang Zhang, Ting Li, Xinwei Feng, Xiaodong Yan, Jinhan Xie</i>	20692
Understanding Synthetic Context Extension via Retrieval Heads <i>Xinyu Zhao, Fangcong Yin, Greg Durrett</i>	20732
Stochastic Control for Fine-tuning Diffusion Models: Optimality, Regularity, and Convergence <i>Yinbin Han, Meisam Razaviyayn, Renyuan Xu</i>	20758
System-Aware Unlearning Algorithms: Use Lesser, Forget Faster <i>Linda Lu, Ayush Sekhari, Karthik Sridharan</i>	20785

Monte Carlo Tree Search for Comprehensive Exploration in LLM-Based Automatic Heuristic Design <i>Zhi Zheng, Zhuoliang Xie, Zhenkun Wang, Bryan Hooi</i>	20818
Towards Global-level Mechanistic Interpretability: A Perspective of Modular Circuits of Large Language Models <i>Yinhan He, Wendy Zheng, Yushun Dong, Yaochen Zhu, Chen Chen, Jundong Li</i>	20854
Provable Maximum Entropy Manifold Exploration via Diffusion Models <i>Riccardo De Santi, Marin Vlastelica, Ya-Ping Hsieh, Zebang Shen, Niao He, Andreas Krause</i>	20870
Towards an Explainable Comparison and Alignment of Feature Embeddings <i>Mohammad Jalali, Bahar Dibaei Nia, Farzan Farnia</i>	20893
RealRAG: Retrieval-augmented Realistic Image Generation via Self-reflective Contrastive Learning <i>Yuanhuiyi Lyu, Xu Zheng, Lutao Jiang, Yibo Yan, Xin Zou, Huiyu Zhou, Linfeng Zhang, Xuming Hu</i>	20933
Inverse Flow and Consistency Models <i>Yuchen Zhang, Jian Zhou</i>	20952
Orthus: Autoregressive Interleaved Image-Text Generation with Modality-Specific Heads <i>Siqi Kou, Jiachun Jin, Zhihong Liu, Chang Liu, Ye Ma, Jian Jia, Quan Chen, Peng Jiang, Zhijie Deng</i>	20977
AdvPrompter: Fast Adaptive Adversarial Prompting for LLMs <i>Anselm Paulus, Arman Zharmagambetov, Chuan Guo, Brandon Amos, Yuandong Tian</i>	20992
Distributionally Robust Multi-Agent Reinforcement Learning for Dynamic Chute Mapping <i>Guangyi Liu, Suzan Iloglu, Michael Caldara, Joseph W Durham, Michael M. Zavlanos</i>	21023
When Will It Fail?: Anomaly to Prompt for Forecasting Future Anomalies in Time Series <i>Min-Yeong Park, Won-Jeong Lee, Seong Tae Kim, Gyeong-Moon Park</i>	21045
Circumventing Backdoor Space via Weight Symmetry <i>Jie Peng, Hongwei Yang, Jing Zhao, Hengji Dong, Hui He, Weizhe Zhang, Haoyu He</i>	21063
FIC-TSC: Learning Time Series Classification with Fisher Information Constraint <i>Xiwen Chen, Wenhui Zhu, Peijie Qiu, Hao Wang, Huayu Li, ZIHAN LI, Yalin Wang, Aristeidis Sotiras, Abolfazl Razi</i>	21100
Temporal Distance-aware Transition Augmentation for Offline Model-based Reinforcement Learning <i>Dongsu Lee, Minhae Kwon</i>	21122
Preserving AUC Fairness in Learning with Noisy Protected Groups <i>Mingyang Wu, Li Lin, Wenbin Zhang, Xin Wang, Zhenhuan Yang, Shu Hu</i>	21138

Zero-Inflated Bandits <i>Haoyu Wei, Runzhe Wan, Lei Shi, Rui Song</i>	21160
Product of Experts with LLMs: Boosting Performance on ARC Is a Matter of Perspective <i>Daniel Franzen, Jan Disselhoff, David Hartmann</i>	21225
Solving Zero-Sum Convex Markov Games <i>Fivos Kalogiannis, Emmanouil-Vasileios Vlatakis-Gkaragkounis, Ian Gemp, Georgios Piliouras</i>	21240
An Empirical Study on Configuring In-Context Learning Demonstrations for Unleashing MLLMs' Sentimental Perception Capability <i>Daqing Wu, Dongbao Yang, Sicheng Zhao, Can Ma, Yu Zhou</i>	21300
Hypo3D: Exploring Hypothetical Reasoning in 3D <i>Ye Mao, Weixun Luo, Junpeng Jing, Anlan Qiu, Krystian Mikolajczyk</i>	21314
Modularized Self-Reflected Video Reasoner for Multimodal LLM with Application to Video Question Answering <i>Zihan Song, Xin Wang, Zi Qian, Hong Chen, Longtao Huang, Hui Xue, Wenwu Zhu</i>	21338
Beyond Sensor Data: Foundation Models of Behavioral Data from Wearables Improve Health Predictions <i>Eray Erturk, Fahad Kamran, Salar Abbaspourazad, Sean Jewell, Harsh Sharma, Yujie Li, Sinead Williamson, Nicholas J Foti, Joseph Futoma</i>	21363
Revisiting Diffusion Models: From Generative Pre-training to One-Step Generation <i>Bowen Zheng, Tianming Yang</i>	21389
Robust Multi-bit Text Watermark with LLM-based Paraphrasers <i>Xiaojun Xu, Jinghan Jia, Yuanshun Yao, Yang Liu, Hang Li</i>	21409
The Complexity of Learning Sparse Superposed Features with Feedback <i>Akash Kumar</i>	21423
Diffusion Instruction Tuning <i>Chen Jin, Ryutaro Tanno, Amrutha Saseendran, Tom Diethe, Philip Alexander Teare</i>	21457
Text-to-CAD Generation Through Infusing Visual Feedback in Large Language Models <i>Ruiyu Wang, Yu Yuan, Shizhao Sun, Jiang Bian</i>	21498
A Closer Look at Backdoor Attacks on CLIP <i>Shuo He, Zhifang Zhang, Feng Liu, Roy Ka-Wei Lee, Bo An, Lei Feng</i>	21518
Addressing Imbalanced Domain-Incremental Learning through Dual-Balance Collaborative Experts <i>Lan Li, Da-Wei Zhou, Han-Jia Ye, De-Chuan Zhan</i>	21535

Spatial Reasoning with Denoising Models <i>Christopher Wewer, Bartłomiej Pogodzinski, Bernt Schiele, Jan Eric Lenssen</i>	21554
BARK: A Fully Bayesian Tree Kernel for Black-box Optimization <i>Toby Boyne, Jose Pablo Folch, Robert Matthew Lee, Behrang Shafei, Ruth Misener</i>	21574
Logits are All We Need to Adapt Closed Models <i>Gaurush Hiranandani, Haolun Wu, Subhojyoti Mukherjee, Sanmi Koyejo</i>	21602
SFT Memorizes, RL Generalizes: A Comparative Study of Foundation Model Post-training <i>Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V Le, Sergey Levine, Yi Ma</i>	21631
Transfer Q-Learning with Composite MDP Structures <i>Jinhang Chai, Elynn Chen, Lin Yang</i>	21652
Contract Design Under Approximate Best Responses <i>Francesco Bacchiocchi, Jiarui Gan, Matteo Castiglioni, Alberto Marchesi, Nicola Gatti</i>	21670
The Sharpness Disparity Principle in Transformers for Accelerating Language Model Pre-Training <i>Jinbo Wang, Mingze Wang, Zhanpeng Zhou, Junchi Yan, Weinan E, Lei Wu</i>	21688
Interaction-Aware Gaussian Weighting for Clustered Federated Learning <i>Alessandro Licciardi, Davide Leo, Eros Fani, Barbara Caputo, Marco Ciccone</i>	21709
How Distributed Collaboration Influences the Diffusion Model Training? A Theoretical Perspective <i>Jing Qiao, Yu Liu, YUAN YUAN, Xiao Zhang, Zhipeng Cai, Dongxiao Yu</i>	21734
Perceptually Constrained Precipitation Nowcasting Model <i>Wenzhi Feng, Xutao Li, Zhe Wu, Kenghong Lin, Demin Yu, Yunming Ye, Yaowei Wang</i>	21752
FisherSFT: Data-Efficient Supervised Fine-Tuning of Language Models Using Information Gain <i>Rohan Deb, Kiran Koshy Thekumparampil, Kousha Kalantari, Gaurush Hiranandani, Shoham Sabach, Branislav Kveton</i>	21771
Neural Solver Selection for Combinatorial Optimization <i>Chengrui Gao, Haopu Shang, Ke Xue, Chao Qian</i>	21791
Generative Social Choice: The Next Generation <i>Niclas Boehmer, Sara Fish, Ariel D. Procaccia</i>	21813
Temporal Query Network for Efficient Multivariate Time Series Forecasting <i>Shengsheng Lin, Haojun Chen, Haijie Wu, Chunyun Qiu, Weiwei Lin</i>	21846
One-Step Diffusion Policy: Fast Visuomotor Policies via Diffusion Distillation <i>Zhendong Wang, Max Li, Ajay Mandlekar, Zhenjia Xu, Jiaojiao Fan, Yashraj Narang, Linxi Fan, Yuke Zhu, Yogesh Balaji, Mingyuan Zhou, Ming-Yu Liu, Yu Zeng</i>	21864

Contextual Optimization Under Model Misspecification: A Tractable and Generalizable Approach <i>Omar Bennouna, Jiawei Zhang, Saurabh Amin, Asuman E. Ozdaglar</i>	21882
The Sample Complexity of Online Strategic Decision Making with Information Asymmetry and Knowledge Transportability <i>Jiachen Hu, Rui Ai, Han Zhong, Xiaoyu Chen, Liwei Wang, Zhaoran Wang, Zhuoran Yang</i>	21909
MMedPO: Aligning Medical Vision-Language Models with Clinical-Aware Multimodal Preference Optimization <i>Kangyu Zhu, Peng Xia, Yun Li, Hongtu Zhu, Sheng Wang, Huaxiu Yao</i>	21938
The Logical Implication Steering Method for Conditional Interventions on Transformer Generation <i>Damjan Kalajdzievski</i>	21954
Understanding Mode Connectivity via Parameter Space Symmetry <i>Bo Zhao, Nima Dehmamy, Robin Walters, Rose Yu</i>	21986
Safety Alignment Can Be Not Superficial With Explicit Safety Signals <i>Jianwei Li, Jung-Eun Kim</i>	22006
Constrained Online Convex Optimization with Polyak Feasibility Steps <i>Spencer Hutchinson, Mahnoosh Alizadeh</i>	22041
GRU: Mitigating the Trade-off between Unlearning and Retention for LLMs <i>Yue Wang, Qizhou Wang, Feng Liu, Wei Huang, Yali Du, Xiaojiang Du, Bo Han</i>	22056
Algorithms and Hardness for Active Learning on Graphs <i>Vincent Cohen-Addad, Silvio Lattanzi, Simon Meierhans</i>	22077
Return Capping: Sample Efficient CVaR Policy Gradient Optimisation <i>Harry Mead, Clarissa Costen, Bruno Lacerda, Nick Hawes</i>	22092
Hierarchical Refinement: Optimal Transport to Infinity and Beyond <i>Peter Halmos, Julian Gold, Xinhao Liu, Benjamin Raphael</i>	22108
Statistical Hypothesis Testing for Auditing Robustness in Language Models <i>Paulius Rauba, Qiyao Wei, Mihaela van der Schaar</i>	22141
Unified Analysis of Continuous Weak Features Learning with Applications to Learning from Missing Data <i>Kosuke Sugiyama, Masato Uchida</i>	22158
Language Models over Canonical Byte-Pair Encodings <i>Tim Vieira, Tianyu Liu, Clemente Pasti, Yahya Emara, Brian DuSell, Benjamin LeBrun, Mario Giulianelli, Juan Luis Gastaldi, Timothy J. O'Donnell, Ryan Cotterell</i>	22199

The Importance of Being Lazy: Scaling Limits of Continual Learning <i>Jacopo Galdi, Alessandro Breccia, Giulia Lanzillotta, Thomas Hofmann, Lorenzo Noci</i>	22230
FACTER: Fairness-Aware Conformal Thresholding and Prompt Engineering for Enabling Fair LLM-Based Recommender Systems <i>Arya Fayyazi, Mehdi Kamal, Massoud Pedram</i>	22272
Sidechain conditioning and modeling for full-atom protein sequence design with FAMPNN <i>Talal Widatalla, Richard W. Shuai, Brian Hie, Possu Huang</i>	22288
Automatically Interpreting Millions of Features in Large Language Models <i>Gonçalo Santos Paulo, Alex Troy Mallen, Caden Juang, Nora Belrose</i>	22314
In-Context Reinforcement Learning From Suboptimal Historical Data <i>Juncheng Dong, Moyang Guo, Ethan X Fang, Zhuoran Yang, Vahid Tarokh</i>	22343
Active Learning with Selective Time-Step Acquisition for PDEs <i>Yegon Kim, Hyunsu Kim, Gyeonghoon Ko, Juho Lee</i>	22362
Reflection-Bench: Evaluating Epistemic Agency in Large Language Models <i>Lingyu Li, Yixu Wang, Haiquan Zhao, Shuqi Kong, Yan Teng, Chunbo Li, Yingchun Wang</i>	22387
TS-SNN: Temporal Shift Module for Spiking Neural Networks <i>Kairong Yu, Tianqing Zhang, Qi Xu, Gang Pan, Hongwei Wang</i>	22416
Wyckoff Transformer: Generation of Symmetric Crystals <i>Nikita Kazeev, Wei Nong, Ignat Romanov, Ruiming Zhu, Andrey E Ustyuzhanin, Shuya Yamazaki, Kedar Hippalgaonkar</i>	22431
ExpProof : Operationalizing Explanations for Confidential Models with ZKPs <i>Chhavi Yadav, Evan Laufer, Dan Boneh, Kamalika Chaudhuri</i>	22463
Differentially Private Federated k-Means Clustering with Server-Side Data <i>Jonathan Scott, Christoph H. Lampert, David Saulpic</i>	22485
Bifurcate then Alienate: Incomplete Multi-view Clustering via Coupled Distribution Learning with Linear Overhead <i>Shengju Yu, Yiu-ming Cheung, Siwei Wang, Xinwang Liu, En Zhu</i>	22519
Global Convergence and Rich Feature Learning in L-Layer Infinite-Width Neural Networks under μ Parametrization <i>Zixiang Chen, Greg Yang, Qingyue Zhao, Quanquan Gu</i>	22561
Universal Approximation Theorem of Deep Q-Networks <i>Qian Qi</i>	22589
Causality-Aware Contrastive Learning for Robust Multivariate Time-Series Anomaly Detection <i>HyunGi Kim, Jisoo Mok, Dongjun Lee, Jaihyun Lew, Sungjae Kim, Sungroh Yoon</i>	22612

Instance-Optimal Pure Exploration for Linear Bandits on Continuous Arms <i>Sho Takemori, Yuhei Umeda, Aditya Gopalan</i>	22630
An Architecture Search Framework for Inference-Time Techniques <i>Jon Saad-Falcon, Adrian Gamarra Lafuente, Shlok Natarajan, Nahum Maru, Hristo Todorov, Etash Kumar Guha, E. Kelly Buchanan, Mayee F Chen, Neel Guha, Christopher Re, Azalia Mirhoseini</i>	22658
GTR: A General, Multi-View, and Dynamic Framework for Trajectory Representation Learning <i>Xiangheng Wang, Ziquan Fang, Chenglong Huang, Danlei Hu, Lu Chen, Yunjun Gao</i>	22691
Local Manifold Approximation and Projection for Manifold-Aware Diffusion Planning <i>Kyowoon Lee, Jaesik Choi</i>	22711
Active Evaluation Acquisition for Efficient LLM Benchmarking <i>Yang Li, Jie Ma, Miguel Ballesteros, Yassine Benajiba, Graham Horwood</i>	22729
From Theory to Practice: Rethinking Green and Martin Kernels for Unleashing Graph Transformers <i>Yoon Hyeok Lee, Jaemin Park, Taejin Paik, Doyun Kim, Bosun Hwang</i>	22751
Wrapped Gaussian on the manifold of Symmetric Positive Definite Matrices <i>Thibault de Surrel, Fabien Lotte, Sylvain Chevallier, Florian Yger</i>	22778
MindLLM: A Subject-Agnostic and Versatile Model for fMRI-to-text Decoding <i>Weikang Qiu, Zheng Huang, Haoyu Hu, Aosong Feng, Yujun Yan, Rex Ying</i>	22802
QUTE: Quantifying Uncertainty in TinyML models with Early-exit-assisted ensembles for model-monitoring <i>Nikhil Pratap Ghanathe, Steven J E Wilton</i>	22824
Elucidating the design space of language models for image generation <i>Xuantong LIU, Shaozhe Hao, Xianbiao Qi, Tianyang Hu, Jun Wang, Rong Xiao, Yuan Yao</i>	22845
FG-CLIP: Fine-Grained Visual and Textual Alignment <i>Chunyu Xie, Bin Wang, Fanjing Kong, Jincheng Li, Dawei Liang, Gengshen Zhang, Dawei Leng, Yuhui Yin</i>	22870
Core Knowledge Deficits in Multi-Modal Language Models <i>Yijiang Li, Qingying Gao, Tianwei Zhao, Bingyang Wang, Haoran Sun, Haiyun Lyu, Robert D. Hawkins, Nuno Vasconcelos, Tal Golan, Dezhi Luo, Hokin Deng</i>	22887
QMamba: On First Exploration of Vision Mamba for Image Quality Assessment <i>Fengbin Guan, Xin Li, Zihao Yu, Yiting Lu, Zhibo Chen</i>	22918
CFP-Gen: Combinatorial Functional Protein Generation via Diffusion Language Models <i>Junbo Yin, Chao Zha, Wenjia He, Chencheng Xu, Xin Gao</i>	22932
Flow-of-Options: Diversified and Improved LLM Reasoning by Thinking Through Options <i>Lakshmi Nair, Ian Trase, J. Mark Kim</i>	22952

Connecting Thompson Sampling and UCB: Towards More Efficient Trade-offs Between Privacy and Regret <i>Bingshan Hu, ZHIMING HUANG, Tianyue H. Zhang, Mathias Lécuyer, Nidhi Hegde</i>	22984
PAC-Bayes Analysis for Recalibration in Classification <i>Masahiro Fujisawa, Futoshi Futami</i>	23005
Online Curvature-Aware Replay: Leveraging 2nd Order Information for Online Continual Learning <i>Edoardo Urettini, Antonio Carta</i>	23043
Integrating Intermediate Layer Optimization and Projected Gradient Descent for Solving Inverse Problems with Diffusion Models <i>Yang Zheng, Wen Li, Zhaoqiang Liu</i>	23063
Look Twice Before You Answer: Memory-Space Visual Retracing for Hallucination Mitigation in Multimodal Large Language Models <i>Xin Zou, Yizhou Wang, Yibo Yan, Yuanhuiyi Lyu, Kening Zheng, Sirui Huang, Junkai Chen, Peijie Jiang, Jia Liu, Chang Tang, Xuming Hu</i>	23081
Mitigating Plasticity Loss in Continual Reinforcement Learning by Reducing Churn <i>Hongyao Tang, Johan Obando-Ceron, Pablo Samuel Castro, Aaron Courville, Glen Berseth</i>	23108
Binary Hypothesis Testing for Softmax Models and Leverage Score Models <i>Yuzhou Gu, Zhao Song, Junze Yin</i>	23130
Explaining the role of Intrinsic Dimensionality in Adversarial Training <i>Enes Altinisik, Safa Messaoud, Husrev Taha Sencar, Hassan Sajjad, Sanjay Chawla</i>	23152
An Effective and Secure Federated Multi-View Clustering Method with Information-Theoretic Perspective <i>Xinyue Chen, Jinfeng Peng, YuHao Li, Xiaorong Pu, Yang Yang, Yazhou Ren</i>	23168
Provably Efficient Exploration in Inverse Constrained Reinforcement Learning <i>Bo Yue, Jian Li, Guiliang Liu</i>	23187
HashAttention: Semantic Sparsity for Faster Inference <i>Aditya Desai, Shuo Yang, Alejandro Cuadron, Matei Zaharia, Joseph E. Gonzalez, Ion Stoica</i>	23230
Vulnerability-Aware Alignment: Mitigating Uneven Forgetting in Harmful Fine-Tuning <i>Liang CHEN, Xueting Han, Li Shen, Jing Bai, Kam-Fai Wong</i>	23247
Shifting Time: Time-series Forecasting with Khatri-Rao Neural Operators <i>Srinath Dama, Kevin Course, Prasanth B. Nair</i>	23259
UGPhysics: A Comprehensive Benchmark for Undergraduate Physics Reasoning with Large Language Models <i>Xin Xu, Qiyun Xu, Tong Xiao, Tianhao Chen, Yuchen Yan, Jiaxin ZHANG, Shizhe Diao, Can Yang, Yang Wang</i>	23293

Predicting mutational effects on protein binding from folding energy <i>Arthur Deng, Karsten D. Householder, Fang Wu, K. Christopher Garcia, Brian L. Trippe</i>	23322
Function-Space Learning Rates <i>Edward Milsom, Ben Anson, Laurence Aitchison</i>	23345
MuLan: Adapting Multilingual Diffusion Models for Hundreds of Languages with Negligible Cost <i>Sen Xing, Muyan Zhong, Zeqiang Lai, Liangchen Li, Jiawen Liu, Yaohui Wang, Jifeng Dai, Wenhai Wang</i>	23372
Deterministic Sparse Fourier Transform for Continuous Signals with Frequency Gap <i>Xiaoyu Li, Zhao Song, Shenghao Xie</i>	23389
On-Device Collaborative Language Modeling via a Mixture of Generalists and Specialists <i>Dongyang Fan, Bettina Messmer, Nikita Doikov, Martin Jaggi</i>	23427
Learnable Spatial-Temporal Positional Encoding for Link Prediction <i>Katherine Tieu, Dongqi Fu, Zihao Li, Ross Maciejewski, Jingrui He</i>	23456
A Cognac Shot To Forget Bad Memories: Corrective Unlearning for Graph Neural Networks <i>Varshita Kolipaka, Akshit Sinha, Debangan Mishra, Sumit Kumar, Arvinth Arun, Shashwat Goel, Ponnurangam Kumaraguru</i>	23484
Modeling Multi-Task Model Merging as Adaptive Projective Gradient Descent <i>Yongxian Wei, Anke Tang, Li Shen, Zixuan Hu, Chun Yuan, Xiaochun Cao</i>	23510
Offline Model-based Optimization for Real-World Molecular Discovery <i>Dong-Hee Shin, Young-Han Son, Hyun Jung Lee, Deok-Joong Lee, Tae-Eui Kam</i>	23526
Differentially Private Boxplots <i>Kelly Ramsay, Jairo Diaz-Rodriguez</i>	23576
Designing Cyclic Peptides via Harmonic SDE with Atom-Bond Modeling <i>Xiangxin Zhou, Mingyu Li, Yi Xiao, Jiahao Li, Dongyu Xue, Zaixiang Zheng, Jianzhu Ma, Quanquan Gu</i>	23599
Can Diffusion Models Learn Hidden Inter-Feature Rules Behind Images? <i>Yujin Han, Andi Han, Wei Huang, Chaochao Lu, Difan Zou</i>	23631
Strategy Coopetition Explains the Emergence and Transience of In-Context Learning <i>Aaditya K Singh, Ted Moskowitz, Sara Dragutinović, Felix Hill, Stephanie C.Y. Chan, Andrew M Saxe</i>	23660
Online Clustering of Dueling Bandits <i>Zhiyong Wang, Jiahang Sun, Mingze Kong, Jize Xie, Qinghua Hu, John C.S. Lui, Zhongxiang Dai</i>	23680
Reinforcement Learning for Quantum Control under Physical Constraints <i>Jan Ole Ernst, Aniket Chatterjee, Tim Franzmeyer, Axel Kuhn</i>	23709
Equivariant Polynomial Functional Networks <i>Thieu Vo, Hoang V. Tran, Tho Tran Huu, An Nguyen The, Thanh Tran, Minh-Khoi Nguyen-Nhat, Duy-Tung Pham, Tan Minh Nguyen</i>	23736

Unbiased Evaluation of Large Language Models from a Causal Perspective <i>Meilin Chen, Jian Tian, Liang Ma, Di Xie, Weijie Chen, Jiang Zhu</i>	23792
AlphaDPO: Adaptive Reward Margin for Direct Preference Optimization <i>Junkang Wu, Xue Wang, Zhengyi Yang, Jiancan Wu, Jinyang Gao, Bolin Ding, Xiang Wang, Xiangnan He</i>	23808
General framework for online-to-nonconvex conversion: Schedule-free SGD is also effective for nonconvex optimization <i>Kwangjun Ahn, Gagik Magakyan, Ashok Cutkosky</i>	23825
Uncertainty-Based Extensible Codebook for Discrete Federated Learning in Heterogeneous Data Silos <i>Tianyi Zhang, Yu Cao, Dianbo Liu</i>	23849
QPRL : Learning Optimal Policies with Quasi-Potential Functions for Asymmetric Traversal <i>Jumman Hossain, Nirmalya Roy</i>	23867
Emotional Face-to-Speech <i>Jiaxin Ye, Boyuan Cao, Hongming Shan</i>	23885
I²MoE: Interpretable Multimodal Interaction-aware Mixture-of-Experts <i>Jiayi Xin, Sukwon Yun, Jie Peng, Inyoung Choi, Jenna L. Ballard, Tianlong Chen, Qi Long</i>	23906
Aligning LLMs by Predicting Preferences from User Writing Samples <i>Stéphane Aroca-Ouellette, Natalie Mackraz, Barry-John Theobald, Katherine Metcalf</i>	23925
COGNATE: Acceleration of Sparse Tensor Programs on Emerging Hardware using Transfer Learning <i>Chamika Sudusinghe, Gerasimos Gerogiannis, Damitha Lenadora, Charles Block, Josep Torrellas, Charith Mendis</i>	23957
DiffMS: Diffusion Generation of Molecules Conditioned on Mass Spectra <i>Montgomery Bohde, Mrunali Manjrekar, Runzhong Wang, Shuiwang Ji, Connor W. Coley</i>	23975
The Price of Freedom: Exploring Expressivity and Runtime Tradeoffs in Equivariant Tensor Products <i>YuQing Xie, Ameya Daigavane, Mit Kotak, Tess Smidt</i>	23995
SK-VQA: Synthetic Knowledge Generation at Scale for Training Context-Augmented Multimodal LLMs <i>Xin Su, Man Luo, Kris W Pan, Tien Pei Chou, Vasudev Lal, Phillip Howard</i>	24022
DTZO: Distributed Trilevel Zeroth Order Learning with Provable Non-Asymptotic Convergence <i>Yang Jiao, Kai Yang, Chengtao Jian</i>	24044
SquarexPO: Differentially Private and Robust χ^2-Preference Optimization in Offline Direct Alignment <i>Xingyu Zhou, Yulian Wu, Wenqian Weng, Francesco Orabona</i>	24088

Posterior Inference with Diffusion Models for High-dimensional Black-box Optimization <i>Taeyoung Yun, Kiyoung Om, Jaewoo Lee, Sujin Yun, Jinkyoo Park</i>	24113
Learning-Order Autoregressive Models with Application to Molecular Graph Generation <i>Zhe Wang, Jiaxin Shi, Nicolas Heess, Arthur Gretton, Michalis Titsias</i>	24134
CAD-Editor: A Locate-then-Infill Framework with Automated Training Data Synthesis for Text-Based CAD Editing <i>Yu Yuan, Shizhao Sun, Qi Liu, Jiang Bian</i>	24154
CMoS: Rethinking Time Series Prediction Through the Lens of Chunk-wise Spatial Correlations <i>Haotian Si, Changhua Pei, Jianhui Li, Dan Pei, Gaogang Xie</i>	24170
MoH: Multi-Head Attention as Mixture-of-Head Attention <i>Peng Jin, Bo Zhu, Li Yuan, Shuicheng YAN</i>	24188
Policy-Regret Minimization in Markov Games with Function Approximation <i>Thanh Nguyen-Tang, Raman Arora</i>	24211
Discrete and Continuous Difference of Submodular Minimization <i>George Orfanides, Tim Hoheisel, Marwa El Halabi</i>	24234
CodeSteer: Symbolic-Augmented Language Models via Code/Text Guidance <i>Yongchao Chen, Yilun Hao, Yueying Liu, Yang Zhang, Chuchu Fan</i>	24271
TAROT: Targeted Data Selection via Optimal Transport <i>Lan Feng, Fan Nie, Yuejiang Liu, Alexandre Alahi</i>	24299
Benchmarking Quantum Reinforcement Learning <i>Nico Meyer, Christian Ufrecht, George Yammine, Georgios Kontes, Christopher Mutschler, Daniel Scherer</i>	24315
Statistical Query Hardness of Multiclass Linear Classification with Random Classification Noise <i>Ilias Diakonikolas, Mingchen Ma, Lisheng Ren, Christos Tzamos</i>	24346
Demeaned Sparse: Efficient Anomaly Detection by Residual Estimate <i>Yifan Fang, Yifei Fang, Ruizhe Chen, Haote Xu, Xinghao Ding, Yue Huang</i>	24366
In-Context Denoising with One-Layer Transformers: Connections between Attention and Associative Memory Retrieval <i>Matthew Smart, Alberto Bietti, Anirvan M. Sengupta</i>	24390
Ranked from Within: Ranking Large Multimodal Models Without Labels <i>Weijie Tu, Weijian Deng, Dylan Campbell, Yu Yao, Jiyang Zheng, Tom Gedeon, Tongliang Liu</i>	24412
A Theoretical Justification for Asymmetric Actor-Critic Algorithms <i>Gaspard Lambrechts, Damien Ernst, Aditya Mahajan</i>	24435
Global curvature for second-order optimization of neural networks <i>Alberto Bernacchia</i>	24466

On Volume Minimization in Conformal Regression <i>Batiste Le bars, Pierre Humbert</i>	24500
ReQFlow: Rectified Quaternion Flow for Efficient and High-Quality Protein Backbone Generation <i>Angxiao Yue, Zichong Wang, Hongteng Xu</i>	24524
Causal Invariance-aware Augmentation for Brain Graph Contrastive Learning <i>Minqi Yu, Jinduo Liu, Junzhong Ji</i>	24547
Proactive Agents for Multi-Turn Text-to-Image Generation Under Uncertainty <i>Meera Hahn, Wenjun Zeng, Nithish Kannen, Rich Galt, Kartikeya Badola, Been Kim, Zi Wang</i>	24567
The Canary's Echo: Auditing Privacy Risks of LLM-Generated Synthetic Text <i>Matthieu Meeus, Lukas Wutschitz, Santiago Zanella-Beguelin, Shruti Tople, Reza Shokri</i>	24605
Provable Policy Gradient for Robust Average-Reward MDPs Beyond Rectangularity <i>Qiu hao Wang, Yuqi Zha, Chin Pang Ho, Marek Petrik</i>	24629
Eigen Analysis of Conjugate Kernel and Neural Tangent Kernel <i>Xiangchao Li, Xiao Han, Qing Yang</i>	24661
Categorical Distributional Reinforcement Learning with Kullback-Leibler Divergence: Convergence and Asymptotics <i>Tyler Kastner, Mark Rowland, Yunhao Tang, Murat A Erdogdu, Amir-massoud Farahmand</i>	24680
Elucidating Flow Matching ODE Dynamics via Data Geometry and Denoisers <i>Zhengchao Wan, Qingsong Wang, Gal Mishne, Yusu Wang</i>	24707
Constrained Belief Updates Explain Geometric Structures in Transformer Representations <i>Mateusz Piotrowski, Paul M. Riechers, Daniel Filan, Adam Shai</i>	24771
TreeLoRA: Efficient Continual Learning via Layer-Wise LoRAs Guided by a Hierarchical Gradient-Similarity Tree <i>Yu-Yang Qian, Yuan-Ze Xu, Zhen-Yu Zhang, Peng Zhao, Zhi-Hua Zhou</i>	24792
QLASS: Boosting Language Agent Inference via Q-Guided Stepwise Search <i>Zongyu Lin, Yao Tang, Xingcheng Yao, Da Yin, Ziniu Hu, Yizhou Sun, Kai-Wei Chang</i>	24812
What can large language models do for sustainable food? <i>Anna Thomas, Adam Yee, Andrew Mayne, Maya B. Mathur, Dan Jurafsky, Kristina Gligorić</i>	24829
Exact Recovery of Sparse Binary Vectors from Generalized Linear Measurements <i>Arya Mazumdar, Neha Sangwan</i>	24886
How Do Images Align and Complement LiDAR? Towards a Harmonized Multi-modal 3D Panoptic Segmentation <i>Yining Pan, Qiongjie Cui, Xulei Yang, Na Zhao</i>	24908

Efficient Generative Modeling with Residual Vector Quantization-Based Tokens <i>Jaehyeon Kim, Taehong Moon, Keon Lee, Jaewoong Cho</i>	24925
DragSolver: A Multi-Scale Transformer for Real-World Automotive Drag Coefficient Estimation <i>Ye Liu, Yuntian Chen</i>	24947
Rethinking the Temperature for Federated Heterogeneous Distillation <i>Fan Qi, Daxu Shi, Chuokun Xu, Shuai Li, Changsheng Xu</i>	24964
From Pixels to Perception: Interpretable Predictions via Instance-wise Grouped Feature Selection <i>Moritz Vandenhardt, Julia E Vogt</i>	24982
Graph4MM: Weaving Multimodal Learning with Structural Information <i>Xuying Ning, Dongqi Fu, Tianxin Wei, Wujiang Xu, Jingrui He</i>	25006
Modeling All-Atom Glycan Structures via Hierarchical Message Passing and Multi-Scale Pre-training <i>Minghao Xu, Jiaze Song, Keming Wu, Xiangxin Zhou, Bin CUI, Wentao Zhang</i>	25031
Towards scientific discovery with dictionary learning: Extracting biological concepts from microscopy foundation models <i>Konstantin Donhauser, Kristina Ulicna, Gemma Elyse Moran, Aditya Ravuri, Kian Kenyon-Dean, Cian Eastwood, Jason Hartford</i>	25045
Linear Contextual Bandits With Interference <i>Yang Xu, Wenbin Lu, Rui Song</i>	25068
CASE-Bench: Context-Aware SafeTy Benchmark for Large Language Models <i>Guangzhi Sun, Xiao Zhan, Shutong Feng, Phil Woodland, Jose Such</i>	25106
Synthesizing Privacy-Preserving Text Data via Finetuning *without* Finetuning Billion-Scale LLMs <i>Bowen Tan, Zheng Xu, Eric P. Xing, Zhiting Hu, Shanshan Wu</i>	25129
AutoGFM: Automated Graph Foundation Model with Adaptive Architecture Customization <i>Haibo Chen, Xin Wang, Zeyang Zhang, Haoyang Li, Ling Feng, Wenwu Zhu</i>	25150
Unifying 2D and 3D Vision-Language Understanding <i>Ayush Jain, Alexander Swerdlow, Yuzhou Wang, Sergio Arnaud, Ada Martin, Alexander Sax, Franziska Meier, Katerina Fragkiadaki</i>	25171
Idiosyncrasies in Large Language Models <i>Mingjie Sun, Yida Yin, Zhiqiu Xu, J Zico Kolter, Zhuang Liu</i>	25194
Reducing Confounding Bias without Data Splitting for Causal Inference via Optimal Transport <i>Yuguang Yan, Zongyu Li, Haolin Yang, Zeqin Yang, Hao Zhou, Ruichu Cai, Zhifeng Hao</i>	25226
MATS: An Audio Language Model under Text-only Supervision <i>Wen Wang, RuiBing Hou, Hong Chang, Shiguang Shan, Xilin Chen</i>	25245

Revealing Weaknesses in Text Watermarking Through Self-Information Rewrite Attacks <i>Yixin Cheng, Hongcheng Guo, Yangming Li, Leonid Sigal</i>	25265
CoDy: Counterfactual Explainers for Dynamic Graphs <i>Zhan Qu, Daniel Gomm, Michael Färber</i>	25293
CodeIO: Condensing Reasoning Patterns via Code Input-Output Prediction <i>Junlong Li, Daya Guo, Dejian Yang, Runxin Xu, Yu Wu, Junxian He</i>	25317
A Bayesian Model Selection Criterion for Selecting Pretraining Checkpoints <i>Michael Munn, Susan Wei</i>	25336
Improving Flow Matching by Aligning Flow Divergence <i>Yuhao Huang, Taos Transue, Shih-Hsin Wang, William M Feldman, Hong Zhang, Bao Wang</i>	25352
Faster Global Minimum Cut with Predictions <i>Helia Niarparast, Benjamin Moseley, Karan Singh</i>	25374
Breaking the Curse of Multiagency in Robust Multi-Agent Reinforcement Learning <i>Laixi Shi, Jingchu Gai, Eric Mazumdar, Yuejie Chi, Adam Wierman</i>	25389
An Optimistic Algorithm for online CMDPS with Anytime Adversarial Constraints <i>Jiahui Zhu, Kihyun Yu, Dabeen Lee, Xin Liu, Honghao Wei</i>	25404
Physics-Informed Generative Modeling of Wireless Channels <i>Benedikt Böck, Andreas Oeldemann, Timo Mayer, Francesco Rossetto, Wolfgang Utschick</i>	25430
Active Learning for Efficient Discovery of Optimal Combinatorial Perturbations <i>Jason Qin, Hans-Hermann Wessels, Carlos Fernandez-Granda, Yuhan Hao</i>	25455
High-Fidelity Simultaneous Speech-To-Speech Translation <i>Tom Labiausse, Laurent Mazaré, Edouard Grave, Alexandre Défossez, Neil Zeghidour</i>	25476
Generative Intervention Models for Causal Perturbation Modeling <i>Nora Schneider, Lars Lorch, Niki Kilbertus, Bernhard Schölkopf, Andreas Krause</i>	25490
Calibrated Value-Aware Model Learning with Probabilistic Environment Models <i>Claas A Voelcker, Anastasiia Pedan, Arash Ahmadian, Romina Abachi, Igor Gilitschenski, Amir-massoud Farahmand</i>	25515
Cover learning for large-scale topology representation <i>Luis Scoccola, Uzu Lim, Heather A. Harrington</i>	25539
Variance as a Catalyst: Efficient and Transferable Semantic Erasure Adversarial Attack for Customized Diffusion Models <i>Jiachen Yang, Yusong Wang, Yanmei Fang, Yunshu Dai, Fangjun Huang</i>	25568
Censor Dependent Variational Inference <i>Chuanhui Liu, Xiao Wang</i>	25595

Maximal Update Parametrization and Zero-Shot Hyperparameter Transfer for Fourier Neural Operators <i>Shanda Li, Shinjae Yoo, Yiming Yang</i>	25620
Identifying Metric Structures of Deep Latent Variable Models <i>Stas Syrota, Yevgen Zainchkovskyy, Johnny Xi, Benjamin Bloem-Reddy, Søren Hauberg</i>	25635
Autoencoder-Based Hybrid Replay for Class-Incremental Learning <i>Milad Khademi Nori, IL MIN KIM, Guanghai Wang</i>	25658
Machines and Mathematical Mutations: Using GNNs to Characterize Quiver Mutation Classes <i>Jesse He, Helen Jenne, Herman Chau, Davis Brown, Mark Raugas, Sara C. Billey, Henry Kvinge</i>	25673
Whitened CLIP as a Likelihood Surrogate of Images and Captions <i>Roy Betser, Meir Yossef Levi, Guy Gilboa</i>	25695
A Variational Perspective on Generative Protein Fitness Optimization <i>Lea Bogensperger, Dominik Narnhofer, Ahmed Allam, Konrad Schindler, Michael Krauthammer</i>	25722
Falcon: Fast Visuomotor Policies via Partial Denoising <i>Haojun Chen, Minghao Liu, Chengdong Ma, Xiaojian Ma, Zailin Ma, Huimin Wu, Yuanpei Chen, Yifan Zhong, Mingzhi Wang, Qing Li, Yaodong Yang</i>	25735
AutoAdvExBench: Benchmarking Autonomous Exploitation of Adversarial Example Defenses <i>Nicholas Carlini, Edoardo Debenedetti, Javier Rando, Milad Nasr, Florian Tramèr</i>	25757
BDC-CLIP: Brownian Distance Covariance for Adapting CLIP to Action Recognition <i>Fei Long, Xiaou Li, Jiaming Lv, Haoyuan Yang, Xianjun Cheng, Peihua Li</i>	25773
Exploring Large Action Sets with Hyperspherical Embeddings using von Mises-Fisher Sampling <i>Walid Bendada, Guillaume Salha-Galvan, Romain Hennequin, Théo Bontempelli, Thomas Bouabça, Tristan Cazenave</i>	25790
Learn from Downstream and Be Yourself in Multimodal Large Language Models Fine-Tuning <i>Wenke Huang, Jian Liang, Zekun Shi, Didi Zhu, Guancheng Wan, He Li, Bo Du, Dacheng Tao, Mang Ye</i>	25825
Double Machine Learning for Causal Inference under Shared-State Interference <i>Chris Hays, Manish Raghavan</i>	25839
Are Large Language Models Ready for Multi-Turn Tabular Data Analysis? <i>Jinyang Li, Nan Huo, Yan Gao, Jiayi Shi, Yingxiu Zhao, Ge Qu, Bowen Qin, Yurong Wu, Xiaodong Li, Chenhao Ma, Jian-Guang Lou, Reynold Cheng</i>	25872
Value-Based Deep RL Scales Predictably <i>Oleh Rybkin, Michal Nauman, Preston Fu, Charlie Victor Snell, Pieter Abbeel, Sergey Levine, Aviral Kumar</i>	25913
Quadratic Upper Bound for Boosting Robustness <i>Euijin You, Hyang-Won Lee</i>	25935

BOPO: Neural Combinatorial Optimization via Best-anchored and Objective-guided Preference Optimization <i>Zijun Liao, Jinbiao Chen, Debing Wang, Zizhen Zhang, Jiahai Wang</i>	25956
LowRA: Accurate and Efficient LoRA Fine-Tuning of LLMs under 2 Bits <i>Zikai Zhou, Qizheng Zhang, Hermann Kumbong, Kunle Olukotun</i>	25976
Revisiting Unbiased Implicit Variational Inference <i>Tobias Pielok, Bernd Bischl, David Rügamer</i>	26001
Primphormer: Efficient Graph Transformers with Primal Representations <i>Mingzhen He, Ruikai Yang, Hanling Tian, Youmei Qiu, Xiaolin Huang</i>	26014
SBGD: Improving Graph Diffusion Generative Model via Stochastic Block Diffusion <i>Junwei Su, shan Wu</i>	26038
Graph Generative Pre-trained Transformer <i>Xiaohui Chen, Yinkai Wang, Jiaying He, Yuanqi Du, Soha Hassoun, Xiaolin Xu, Liping Liu</i>	26057
BSemiFL: Semi-supervised Federated Learning via a Bayesian Approach <i>Haozhao Wang, Shengyu Wang, Jiaming Li, Hao Ren, XINGSHUO HAN, Wenchao Xu, Shangwei Guo, Tianwei Zhang, Ruixuan Li</i>	26079
Revisiting Neural Networks for Few-Shot Learning: A Zero-Cost NAS Perspective <i>Haidong Kang</i>	26094
Critical Tokens Matter: Token-Level Contrastive Estimation Enhances LLM's Reasoning Capability <i>Zicheng Lin, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xing Wang, Ruilin Luo, Chufan Shi, Siheng Li, Yujiu Yang, Zhaopeng Tu</i>	26112
Generative Modeling Reinvents Supervised Learning: Label Repurposing with Predictive Consistency Learning <i>Yang Li, Jiale Ma, Yebin Yang, Qitian Wu, Hongyuan Zha, Junchi Yan</i>	26125
MathConstruct: Challenging LLM Reasoning with Constructive Proofs <i>Mislav Balunovic, Jasper Dekoninck, Nikola Jovanović, Ivo Petrov, Martin Vechev</i>	26144
Near-optimal Regret Using Policy Optimization in Online MDPs with Aggregate Bandit Feedback <i>Tal Lincewicz, Yishay Mansour</i>	26172
PepTune: De Novo Generation of Therapeutic Peptides with Multi-Objective-Guided Discrete Diffusion <i>Sophia Tang, Yinyao Zhang, Pranam Chatterjee</i>	26197
Graph-constrained Reasoning: Faithful Reasoning on Knowledge Graphs with Large Language Models <i>Linhao Luo, Zicheng Zhao, Gholamreza Haffari, Yuan-Fang Li, Chen Gong, Shirui Pan</i>	26246

When Maximum Entropy Misleads Policy Optimization <i>Ruipeng Zhang, Ya-Chien Chang, Sicun Gao</i>	26272
Grokking Beyond the Euclidean Norm of Model Parameters <i>Tikeng Notsawo Pascal Junior, Guillaume Dumas, Guillaume Rabusseau</i>	26294
LEMoN: Label Error Detection using Multimodal Neighbors <i>Haoran Zhang, Aparna Balagopalan, Nassim Oufattole, Hyewon Jeong, Yan Wu, Jiacheng Zhu, Marzyeh Ghassemi</i>	26361
From Low Rank Gradient Subspace Stabilization to Low-Rank Weights: Observations, Theories, and Applications <i>AJAY KUMAR JAISWAL, Yifan Wang, Lu Yin, Shiwei Liu, Runjin Chen, Jiawei Zhao, Ananth Grama, Yuandong Tian, Zhangyang Wang</i>	26404
CombimOTS: Combinatorial Multi-Objective Tree Search for Dual-Target Molecule Generation <i>Thibaud Southiratr, Bonil Koo, Yijingxiu Lu, Sun Kim</i>	26421
Pfeife: Automatic Pipeline Parallelism for PyTorch <i>Ho Young Jhoo, Chung-Kil Hur, Nuno P. Lopes</i>	26463
Contour Integration Underlies Human-Like Vision <i>Ben Lonnqvist, Elsa Scialom, Abdulkadir Gokce, Zehra Merchant, Michael Herzog, Martin Schrimpf</i>	26479
Learning-Augmented Hierarchical Clustering <i>Vladimir Braverman, Jon C. Ergun, Chen Wang, Samson Zhou</i>	26501
To Steer or Not to Steer? Mechanistic Error Reduction with Abstention for Language Models <i>Anna Hedström, Salim I. Amoukou, Tom Bewley, Saumitra Mishra, Manuela Veloso</i>	26555
HEAP: Hyper Extended A-PDHG Operator for Constrained High-dim PDEs <i>Mingquan Feng, Weixin Liao, Yixin Huang, Yifan Fu, Qifu Zheng, Junchi Yan</i>	26577
From Token to Rhythm: A Multi-Scale Approach for ECG-Language Pretraining <i>Fuying Wang, Jiacheng Xu, Lequan Yu</i>	26589
Textural or Textual: How Vision-Language Models Read Text in Images <i>Hanzhang Wang, Qingyuan Ma</i>	26605
Identifying and Understanding Cross-Class Features in Adversarial Training <i>Zeming Wei, Steven Y. Guo, Yisen Wang</i>	26620
Learning Cascade Ranking as One Network <i>Yunli Wang, ZhenZhang, Zhiqiang Wang, Zixuan Yang, Yu Li, Jian Yang, Shiyang Wen, Peng Jiang, Kun Gai</i>	26638
Scaling Collapse Reveals Universal Dynamics in Compute-Optimally Trained Neural Networks <i>Shikai Qiu, Lechao Xiao, Andrew Gordon Wilson, Jeffrey Pennington, Atish Agarwala</i>	26655

MIPT: Multilevel Informed Prompt Tuning for Robust Molecular Property Prediction <i>yeyunchen, Jiangming Shi</i>	26679
Predicting High-precision Depth on Low-Precision Devices Using 2D Hilbert Curves <i>Mykhail Uss, Ruslan Yermolenko, Oleksii Shashko, Olena Kolodiazhna, Ivan Safonov, Volodymyr Savin, Yoonjae Yeo, Seowon Ji, Jaeyun Jeong</i>	26695
Hierarchical Reinforcement Learning with Targeted Causal Interventions <i>Mohammadsadegh Khorasani, Saber Salehkaleybar, Negar Kiyavash, Matthias Grossglauser</i>	26717
Towards a Formal Theory of Representational Compositionality <i>Eric Elmoznino, Thomas Jiralerspong, Yoshua Bengio, Guillaume Lajoie</i>	26761
Learning Gaussian DAG Models without Condition Number Bounds <i>Constantinos Costis Daskalakis, ANTHIMOS VARDIS KANDIROU, Rui Yao</i>	26791
Improved Online Confidence Bounds for Multinomial Logistic Bandits <i>Joongkyu Lee, Min-hwan Oh</i>	26827
Multi-band Frequency Reconstruction for Neural Psychoacoustic Coding <i>Dianwen Ng, Kun Zhou, Yi-Wen Chao, Zhiwei Xiong, Bin Ma, EngSiong Chng</i>	26867
Importance Sampling for Nonlinear Models <i>Prakash Palanivelu Rajmohan, Fred Roosta</i>	26887
MONA: Myopic Optimization with Non-myopic Approval Can Mitigate Multi-step Reward Hacking <i>Sebastian Farquhar, Vikrant Varma, David Lindner, David Elson, Caleb Biddulph, Ian Goodfellow, Rohin Shah</i>	26908
Stabilizing Sample Similarity in Representation via Mitigating Random Consistency <i>Jieting Wang, ZhangZelong, Feijiang Li, Yuhua Qian, Xinyan Liang</i>	26944
Differentiable Structure Learning with Ancestral Constraints <i>Taiyu Ban, Changxin Rong, Xiangyu Wang, Lyuzhou Chen, Xin Wang, Derui Lyu, Qinrui Zhu, Huanhuan Chen</i>	26976
O-MAPL: Offline Multi-agent Preference Learning <i>The Viet Bui, Tien Anh Mai, Thanh Hong Nguyen</i>	27011
WildChat-50M: A Deep Dive Into the Role of Synthetic Data in Post-Training <i>Benjamin Feuer, Chinmay Hegde</i>	27042
Robust Offline Reinforcement Learning with Linearly Structured f-Divergence Regularization <i>Cheng Tang, Zhishuai Liu, Pan Xu</i>	27073
Cross-regularization: Adaptive Model Complexity through Validation Gradients <i>Carlos Stein Brito</i>	27114

PEAKS: Selecting Key Training Examples Incrementally via Prediction Error Anchored by Kernel Similarity <i>Mustafa Burak Gurbuz, Xingyu Zheng, Constantine Dovrolis</i>	27134
Enhancing Graph Contrastive Learning for Protein Graphs from Perspective of Invariance <i>Yusong Wang, Shiyin Tan, Jialun Shen, Yicheng Xu, Haobo Song, Qi Xu, Prayag Tiwari, Mingkun Xu</i>	27157
On Explaining Equivariant Graph Networks via Improved Relevance Propagation <i>Hongyi Ling, Haiyang Yu, Zhimeng Jiang, Na Zou, Shuiwang Ji</i>	27185
Unveiling Markov heads in Pretrained Language Models for Offline Reinforcement Learning <i>Wenhao Zhao, Qiushui Xu, Linjie Xu, Lei Song, Jinyu Wang, Chunlai Zhou, Jiang Bian</i>	27201
Multivariate Conformal Selection <i>Tian Bai, Yue Zhao, Xiang Yu, Archer Y. Yang</i>	27216
Smoothed Preference Optimization via ReNoise Inversion for Aligning Diffusion Models with Varied Human Preferences <i>Yunhong Lu, Qichao Wang, Hengyuan Cao, Xiaoyin Xu, Min Zhang</i>	27241
UncertainSAM: Fast and Efficient Uncertainty Quantification of the Segment Anything Model <i>Timo Kaiser, Thomas Norrenbrock, Bodo Rosenhahn</i>	27258
QoS-Efficient Serving of Multiple Mixture-of-Expert LLMs Using Partial Runtime Reconfiguration <i>HamidReza Imani, Jiaxin Peng, Peiman Mohseni, Abdolah Amirany, Tarek El-Ghazawi</i>	27277
Stability and Generalization Analysis of Decentralized SGD: Sharper Bounds Beyond Lipschitzness and Smoothness <i>Shuang Zeng, Yunwen Lei</i>	27290
Cavia: Camera-controllable Multi-view Video Diffusion with View-Integrated Attention <i>Dejia Xu, Yifan Jiang, Chen Huang, Liangchen Song, Thorsten Gernoth, Liangliang Cao, Zhangyang Wang, Hao Tang</i>	27325
Learning Imperfect Information Extensive-form Games with Last-iterate Convergence under Bandit Feedback <i>Canzhe Zhao, Yutian Cheng, Jing Dong, Baoxiang Wang, Shuai Li</i>	27350
LRA-QViT: Integrating Low-Rank Approximation and Quantization for Robust and Efficient Vision Transformers <i>Beom Jin Kang, NamJoon Kim, Hyun Kim</i>	27386
Guided Structural Inference: Leveraging Priors with Soft Gating Mechanisms <i>Aoran Wang, Xinnan Dai, Jun Pang</i>	27402
Towards Rationale-Answer Alignment of LVLMs via Self-Rationale Calibration <i>Yuanchen Wu, Ke Yan, Shouhong Ding, Ziyin Zhou, Xiaoqiang Li</i>	27425

Beyond Cropped Regions: New Benchmark and Corresponding Baseline for Chinese Scene Text Retrieval in Diverse Layouts <i>Gengluo Li, Huawen Shen, Yu Zhou</i>	27446
A Unified Comparative Study with Generalized Conformity Scores for Multi-Output Conformal Regression <i>Victor Dheur, Matteo Fontana, Yorick Estievenart, Naomi Desobry, Souhaib Ben Taieb</i>	27466
Learning Classifiers That Induce Markets <i>Yonatan Sommer, Ivri Hikri, Lotan Amit, Nir Rosenfeld</i>	27508
A Parametric Contextual Online Learning Theory of Brokerage <i>François Bachoc, Tommaso Cesari, Roberto Colomboni</i>	27533
Guarantees of a Preconditioned Subgradient Algorithm for Overparameterized Asymmetric Low-rank Matrix Recovery <i>Paris Giampouras, HanQin Cai, Rene Vidal</i>	27550
Diss-I-ECT: Dissecting Graph Data with Local Euler Characteristic Transforms <i>Julius von Rohrscheidt, Bastian Rieck</i>	27572
Lightweight Online Adaption for Time Series Foundation Model Forecasts <i>Thomas L Lee, William Toner, Rajkarn Singh, Artjom Joosen, Martin Asenov</i>	27592
Robust Secure Swap: Responsible Face Swap With Persons of Interest Redaction and Provenance Traceability <i>Yunshu Dai, Jianwei Fei, Fangjun Huang, Chip Hong Chang</i>	27621
Learning Invariant Causal Mechanism from Vision-Language Models <i>Zeen Song, Siyu Zhao, Xingyu Zhang, Jiangmeng Li, Changwen Zheng, Wenwen Qiang</i>	27637
Over-Tokenized Transformer: Vocabulary is Generally Worth Scaling <i>Hongzhi Huang, Defa Zhu, Banggu Wu, Yutao Zeng, Ya Wang, Qiyang Min, Zhou Xun</i>	27673
ITFormer: Bridging Time Series and Natural Language for Multi-Modal QA with Large-Scale Multitask Dataset <i>Yilin wang, Peixuan Lei, Jie Song, Yuzhe Hao, Tao Chen, Yuxuan Zhang, LEI JIA, Yuanxiang Li, zhongyu wei</i>	27695
ProofAug: Efficient Neural Theorem Proving via Fine-grained Proof Structure Analysis <i>Haoxiong Liu, Jiacheng Sun, Zhenguo Li, Andrew C Yao</i>	27716
Adversaries Can Misuse Combinations of Safe Models <i>Erik Jones, Anca Dragan, Jacob Steinhardt</i>	27735
Morse: Dual-Sampling for Lossless Acceleration of Diffusion Models <i>Chao Li, Jiawei Fan, Anbang Yao</i>	27758

GS-Bias: Global-Spatial Bias Learner for Single-Image Test-Time Adaptation of Vision-Language Models <i>Zhaohong Huang, Yuxin Zhang, JingJing Xie, Fei Chao, Rongrong Ji</i>	27779
Model-Based Exploration in Monitored Markov Decision Processes <i>Alireza Kazemipour, Matthew E. Taylor, Michael Bowling</i>	27792
Adaptive Self-improvement LLM Agentic System for ML Library Development <i>Genghan Zhang, Weixin Liang, Olivia Hsu, Kunle Olukotun</i>	27822
Uncertainty Estimation for Heterophilic Graphs Through the Lens of Information Theory <i>Dominik Fuchsluger, Tom Wollschläger, Johannes Bordne, Stephan Günnemann</i>	27848
LSCD: Lomb--Scargle Conditioned Diffusion for Time series Imputation <i>Elizabeth Fons, Alejandro Sztrajman, Yousef El-Laham, Luciana Ferrer, Svitlana Vyetrenko, Manuela Veloso</i>	27880
QuanONet: Quantum Neural Operator with Application to Differential Equation <i>Ruocheng Wang, Zhuo Xia, Ge Yan, Junchi Yan</i>	27906
AutoElicit: Using Large Language Models for Expert Prior Elicitation in Predictive Modelling <i>Alexander Capstick, Rahul Krishnan, Payam Barnaghi</i>	27927
Orthogonal Subspace Decomposition for Generalizable AI-Generated Image Detection <i>Zhiyuan Yan, Jiangming Wang, Peng Jin, Ke-Yue Zhang, Chengchun Liu, Shen Chen, Taiping Yao, Shouhong Ding, Baoyuan Wu, Li Yuan</i>	27959
Upcycling Text-to-Image Diffusion Models for Multi-Task Capabilities <i>Ruchika Chavhan, Abhinav Mehrotra, Malcolm Chadwick, Alberto Gil Couto Pimentel Ramos, Luca Morreale, Mehdi Noroozi, Sourav Bhattacharya</i>	27980
Matrix Completion with Incomplete Side Information via Orthogonal Complement Projection <i>Gengshuo Chang, Wei Zhang, Lehan Zhang</i>	27997
Ergodic Generative Flows <i>Leo Maxime Brunswic, Mateo Cl��mente, Rui Heng Yang, Adam Sigal, Amir Rasouli, Yinchuan Li</i>	28021
A Variational Information Theoretic Approach to Out-of-Distribution Detection <i>Sudeepa Mondal, Zhuolin Jiang, Ganesh Sundaramoorthi</i>	28041
Does Generation Require Memorization? Creative Diffusion Models using Ambient Diffusion <i>Kulin Shah, Alkis Kalavasis, Adam Klivans, Giannis Daras</i>	28066
Efficient Quantification of Multimodal Interaction at Sample Level <i>Zequn Yang, Hongfa Wang, Di Hu</i>	28090
Contextual Bandits for Unbounded Context Distributions <i>Puning Zhao, Rongfei Fan, Shaowei Wang, Li Shen, Qixin Zhang, ZongKe, Tianhang Zheng</i>	28106

Variational Counterfactual Intervention Planning to Achieve Target Outcomes <i>Xin Wang, Shengfei Lyu, Chi Luo, Xiren Zhou, Huanhuan Chen</i>	28140
Open Materials Generation with Stochastic Interpolants <i>Philipp Höllmer, Thomas Egg, Maya Martirosyan, Eric Fuemmeler, Zeren Shui, Amit Gupta, Pawan Prakash, Adrian Roitberg, Mingjie Liu, George Karypis, Mark Transtrum, Richard Hennig, Ellad B. Tadmor, Stefano Martiniani</i>	28162
BanditSpec: Adaptive Speculative Decoding via Bandit Algorithms <i>Yunlong Hou, Fengzhuo Zhang, Cunxiao Du, Xuan Zhang, Jiachun Pan, Tianyu Pang, Chao Du, Vincent Y. F. Tan, Zhuoran Yang</i>	28196
Trajectory Inference with Smooth Schrödinger Bridges <i>Wanli Hong, Yuliang Shi, Jonathan Niles-Weed</i>	28231
Pre-Training Graph Contrastive Masked Autoencoders are Strong Distillers for EEG <i>Xinxu Wei, kanhao zhao, Yong Jiao, Hua Xie, Lifang He, Yu Zhang</i>	28260
Vintix: Action Model via In-Context Reinforcement Learning <i>Andrei Polubarov, Lyubaykin Nikita, Alexander Derevyagin, Ilya Zisman, Denis Tarasov, Alexander Nikulin, Vladislav Kurenkov</i>	28280
Implicit degree bias in the link prediction task <i>Rachith Aiyappa, Xin Wang, Munjung Kim, Ozgur Can Seckin, Yong-Yeol Ahn, Sadamori Kojaku</i>	28314
Best of Both Worlds: Advantages of Hybrid Graph Sequence Models <i>Ali Behrouz, Ali Parviz, Mahdi Karami, Clayton Sanford, Bryan Perozzi, Vahab Mirrokni</i>	28349
The Noisy Laplacian: a Threshold Phenomenon for Non-Linear Dimension Reduction <i>Alex Kokot, Octavian-Vlad Murad, Marina Meila</i>	28378
Shortcut-connected Expert Parallelism for Accelerating Mixture of Experts <i>Weilin Cai, Juyong Jiang, Le Qin, junweicui, Sunghun Kim, Jiayi Huang</i>	28402
Deep Streaming View Clustering <i>Honglin Yuan, Xingfeng Li, Jian Dai, Xiaojian You, Yuan Sun, Zhenwen Ren</i>	28420
KGMark: A Diffusion Watermark for Knowledge Graphs <i>Hongrui Peng, Haolang Lu, Yuanlong Yu, WeiYe Fu, Kun Wang, Guoshun Nan</i>	28435
Permutation-Free High-Order Interaction Tests <i>Zhaolu Liu, Robert Peach, Mauricio Barahona</i>	28455
Visual Generation Without Guidance <i>Huayu Chen, Kai Jiang, Kaiwen Zheng, Jianfei Chen, Hang Su, Jun Zhu</i>	28472
Neutral residues: revisiting adapters for model extension <i>Franck SIGNE TALLA, Edouard Grave, Herve Jegou</i>	28493

Speak Easy: Eliciting Harmful Jailbreaks from LLMs with Simple Interactions <i>Yik Siu Chan, Narutatsu Ri, Yuxin Xiao, Marzyeh Ghassemi</i>	28510
Enhancing Ligand Validity and Affinity in Structure-Based Drug Design with Multi-Reward Optimization <i>Seungbeom Lee, Munsun Jo, Jungseul Ok, Dongwoo Kim</i>	28536
WATCH: Adaptive Monitoring for AI Deployments via Weighted-Conformal Martingales <i>Drew Prinster, Xing Han, Anqi Liu, Suchi Saria</i>	28550
Algorithmic Recourse for Long-Term Improvement <i>Kentaro Kanamori, Ken Kobayashi, Satoshi Hara, Takuya Takagi</i>	28580
AuPair: Golden Example Pairs for Code Repair <i>Aditi Mavalankar, Hassan Mansoor, Zita Marinho, Mariia Samsikova, Tom Schaul</i>	28609
Risk and cross validation in ridge regression with correlated samples <i>Alexander Atanasov, Jacob A Zavatore-Veth, Cengiz Pehlevan</i>	28635
Density Ratio Estimation with Conditional Probability Paths <i>Hanlin Yu, Arto Klami, Aapo Hyvarinen, Anna Korba, Omar Chehab</i>	28679
EasyRef: Omni-Generalized Group Image Reference for Diffusion Models via Multimodal LLM <i>Zhuofan Zong, Dongzhi Jiang, Bingqi Ma, Guanglu Song, Hao Shao, Dazhong Shen, Yu Liu, Hongsheng Li</i>	28708
Exploring Representations and Interventions in Time Series Foundation Models <i>Michał Wiliński, Mononito Goswami, Willa Potosnak, Nina Żukowska, Artur Dubrawski</i>	28725
The underlying structures of self-attention: symmetry, directionality, and emergent dynamics in Transformer training <i>Matteo Saponati, Pascal Josef Sager, Pau Vilimelis Aceituno, Thilo Stadelmann, Benjamin F Grewe</i>	28751
Lego Sketch: A Scalable Memory-augmented Neural Network for Sketching Data Streams <i>Yuan Feng, Yukun Cao, Hairu Wang, Xike Xie, S Kevin Zhou</i>	28788
Prediction-Aware Learning in Multi-Agent Systems <i>Aymeric Capitaine, Etienne Boursier, Eric Moulines, Michael I. Jordan, Alain Oliviero Durmus</i>	28803
Flow Matching for Denoised Social Recommendation <i>Yinxuan Huang, KE LIANG, Zhuofan Dong, Xiaodong Qu, Wang Tianxiang, Yue Han, Jingao Xu, Bin Zhou, Ye Wang</i>	28830
Galileo: Learning Global & Local Features of Many Remote Sensing Modalities <i>Gabriel Tseng, Anthony Fuller, Marlena Reil, Henry Herzog, Patrick Beukema, Favyen Bastani, James R Green, Evan Shelhamer, Hannah Kerner, David Rolnick</i>	28842
Feasible Action Search for Bandit Linear Programs via Thompson Sampling <i>Aditya Gangrade, Aldo Pacchiano, Clayton Scott, Venkatesh Saligrama</i>	28863

ELEMENTAL: Interactive Learning from Demonstrations and Vision-Language Models for Reward Design in Robotics <i>Letian Chen, Nina Marie Moorman, Matthew Craig Gombolay</i>	28892
Balancing the Scales: A Theoretical and Algorithmic Framework for Learning from Imbalanced Data <i>Corinna Cortes, Anqi Mao, Mehryar Mohri, Yutao Zhong</i>	28918
Active Reward Modeling: Adaptive Preference Labeling for Large Language Model Alignment <i>Yunyi Shen, Hao Sun, Jean-Francois Ton</i>	28951
Diagonal Symmetrization of Neural Network Solvers for the Many-Electron Schrödinger Equation <i>Kevin Han Huang, Ni Zhan, Elif Ertekin, Peter Orbanz, Ryan P Adams</i>	28972
MELON: Provable Defense Against Indirect Prompt Injection Attacks in AI Agents <i>Kaijie Zhu, Xianjun Yang, Jindong Wang, Wenbo Guo, William Yang Wang</i>	29001
Structure-Guided Large Language Models for Text-to-SQL Generation <i>Qinggang Zhang, Hao Chen, Junnan Dong, Shengyuan Chen, Feiran Huang, Xiao Huang</i>	29021
Learning to Steer Learners in Games <i>Yizhou Zhang, Yian Ma, Eric Mazumdar</i>	29042
Learning from Sample Stability for Deep Clustering <i>Zhixin Li, Yuheng Jia, Hui LIU, Junhui Hou</i>	29081
PPDiff: Diffusing in Hybrid Sequence-Structure Space for Protein-Protein Complex Design <i>Zhenqiao Song, Tianxiao Li, Lei Li, Martin Renqiang Min</i>	29097
BECAME: Bayesian Continual Learning with Adaptive Model Merging <i>Mei Li, Yuxiang Lu, Qinyan Dai, Suizhi Huang, Yue Ding, Hongtao Lu</i>	29115
Retrieval Augmented Time Series Forecasting <i>Sungwon Han, Seungeon Lee, Meeyoung Cha, Sercan O Arik, Jinsung Yoon</i>	29136
Auditing Prompt Caching in Language Model APIs <i>Chenchen Gu, Xiang Lisa Li, Rohith Kuditipudi, Percy Liang, Tatsunori Hashimoto</i>	29160
FedECADO: A Dynamical System Model of Federated Learning <i>Aayushya Agarwal, Gauri Joshi, Lawrence Pileggi</i>	29180
Distinguishing Cause from Effect with Causal Velocity Models <i>Johnny Xi, Hugh Dance, Peter Orbanz, Benjamin Bloem-Reddy</i>	29199
Beyond Log-Concavity and Score Regularity: Improved Convergence Bounds for Score-Based Generative Models in W2-distance <i>Marta Gentiloni Silveri, Antonio Ocello</i>	29223

How does Labeling Error Impact Contrastive Learning? A Perspective from Data Dimensionality Reduction <i>Jun Chen, Hong Chen, Yonghua Yu, Yiming Ying</i>	29249
Structure-informed Risk Minimization for Robust Ensemble Learning <i>Fengchun Qiao, Yanlin Chen, Xi Peng</i>	29271
Slimming the Fat-Tail: Morphing-Flow for Adaptive Time Series Modeling <i>Tianyu Liu, Kai Sun, Fuchun Sun, Yu Luo, Yuanlong Zhang</i>	29283
Unisoma: A Unified Transformer-based Solver for Multi-Solid Systems <i>Shilong Tao, Zhe Feng, Haonan Sun, Zhanxing Zhu, Yunhuai Liu</i>	29320
PiD: Generalized AI-Generated Images Detection with Pixelwise Decomposition Residuals <i>Xinghe Fu, Zhiyuan Yan, Zheng Yang, Taiping Yao, Yandan Zhao, Shouhong Ding, Xi Li</i>	29342
Time-Aware World Model for Adaptive Prediction and Control <i>Anh N Nhu, Sanghyun Son, Ming Lin</i>	29357
Prototriever: End-to-End Differentiable Protein Homology Search for Fitness Prediction <i>Ruben Weitzman, Peter Mørch Groth, Lood Van Niekerk, Aoi Otani, Yarin Gal, Debora Susan Marks, Pascal Notin</i>	29380
KoNODE: Koopman-Driven Neural Ordinary Differential Equations with Evolving Parameters for Time Series Analysis <i>Hanru Bai, Weiyang Ding</i>	29402
Structured Preconditioners in Adaptive Optimization: A Unified Analysis <i>Shuo Xie, Tianhao Wang, Sashank J. Reddi, Sanjiv Kumar, Zhiyuan Li</i>	29438
GLGENN: A Novel Parameter-Light Equivariant Neural Networks Architecture Based on Clifford Geometric Algebras <i>Ekaterina Filimoshina, Dmitry Shirokov</i>	29469
A Square Peg in a Square Hole: Meta-Expert for Long-Tailed Semi-Supervised Learning <i>Yaxin Hou, Yuheng Jia</i>	29505
Fragments to Facts: Partial-Information Fragment Inference from LLMs <i>Lucas Rosenblatt, Bin HAN, Robert Wolfe, Bill Howe</i>	29521
On the Learnability of Distribution Classes with Adaptive Adversaries <i>Tosca Lechner, Alex Bie, Gautam Kamath</i>	29545
RoSTE: An Efficient Quantization-Aware Supervised Fine-Tuning Approach for Large Language Models <i>Quan Wei, Chung-Yiu Yau, Hoi To Wai, Yang Zhao, Dongyeop Kang, Youngsuk Park, Mingyi Hong</i>	29570

CoastalBench: A Decade-Long High-Resolution Dataset to Emulate Complex Coastal Processes <i>Zelin Xu, Yupu Zhang, Tingsong Xiao, Maitane Olabarrieta Lizaso, Jose M. Gonzalez-Ondina, Zibo Liu, Shigang Chen, Zhe Jiang</i>	29590
Adversarial Robustness in Two-Stage Learning-to-Defer: Algorithms and Guarantees <i>Yannis Montreuil, Axel Carlier, Lai Xing Ng, Wei Tsang Ooi</i>	29603
On Expressive Power of Looped Transformers: Theoretical Analysis and Enhancement via Timestep Encoding <i>Kevin Xu, Issei Sato</i>	29630
Probabilistic Group Mask Guided Discrete Optimization for Incremental Learning <i>Fengqiang Wan, Yang Yang</i>	29664
Masked Generative Nested Transformers with Decode Time Scaling <i>Sahil Goyal, Debapriya Tula, Gagan Jain, Pradeep Shenoy, Prateek Jain, Sujoy Paul</i>	29678
Benign Overfitting in Token Selection of Attention Mechanism <i>Keitaro Sakamoto, Issei Sato</i>	29697
Towards Attributions of Input Variables in a Coalition <i>Xinhao Zheng, Huiqi Deng, Quanshi Zhang</i>	29781
Towards Efficient Online Tuning of VLM Agents via Counterfactual Soft Reinforcement Learning <i>Lang Feng, Weihao Tan, Zhiyi Lyu, Longtao Zheng, Haiyang Xu, Ming Yan, Fei Huang, Bo An</i>	29805
MetaOptimize: A Framework for Optimizing Step Sizes and Other Meta-parameters <i>Arsalan Sharifnassab, Saber Salehkaleybar, Richard S. Sutton</i>	29825
Conformal Tail Risk Control for Large Language Model Alignment <i>Catherine Chen, Jingyan Shen, Zhun Deng, Lihua Lei</i>	29851
Hessian Geometry of Latent Space in Generative Models <i>Alexander Lobashev, Dmitry Guskov, Maria Larchenko, Mikhail Tamm</i>	29875
Mixture of Hidden-Dimensions: Not All Hidden-States' Dimensions are Needed in Transformer <i>Yilong Chen, Junyuan Shang, Zhenyu Zhang, Jiawei Sheng, Tingwen Liu, Shuohuan Wang, Yu Sun, Hua Wu, Haifeng Wang</i>	29901
False Coverage Proportion Control for Conformal Prediction <i>Alexandre Blain, Bertrand Thirion, Pierre Neuvial</i>	29925
Dendritic Localized Learning: Toward Biologically Plausible Algorithm <i>Changze Lv, Jingwen Xu, Yiyang Lu, Xiaohua Wang, Zhenghua Wang, Zhibo Xu, Di Yu, Xin Du, Xiaoqing Zheng, Xuanjing Huang</i>	29936
NTK-DfL: Enhancing Decentralized Federated Learning in Heterogeneous Settings via Neural Tangent Kernel <i>Gabriel Thompson, Kai Yue, Chau-Wai Wong, Huaiyu Dai</i>	29955

Energy-Based Flow Matching for Generating 3D Molecular Structure <i>Wenyin Zhou, Christopher Iliffe Sprague, Vsevolod Viliuga, Matteo Tadiello, Arne Elofsson, Hossein Azizpour</i>	29977
iDPA: Instance Decoupled Prompt Attention for Incremental Medical Object Detection <i>Huahui Yi, Wei Xu, Ziyuan Qin, Xi Chen, Xiaohu Wu, Kang Li, Qicheng Lao</i>	30001
Sample-Optimal Agnostic Boosting with Unlabeled Data <i>Udaya Ghai, Karan Singh</i>	30020
ELoRA: Low-Rank Adaptation for Equivariant GNNs <i>Chen Wang, Siyu Hu, Guangming Tan, Weile Jia</i>	30040
Reliable and Efficient Amortized Model-based Evaluation <i>Sang T. Truong, Yuheng Tu, Percy Liang, Bo Li, Sanmi Koyejo</i>	30063
Improving Multimodal Learning Balance and Sufficiency through Data Remixing <i>Xiaoyu Ma, Hao Chen, Yongjian Deng</i>	30091
Exploring Invariance in Images through One-way Wave Equations <i>Yinpeng Chen, Dongdong Chen, Xiyang Dai, Mengchen Liu, Yinan Feng, Youzuo Lin, Lu Yuan, Zicheng Liu</i>	30101
Janus: Dual-Server Multi-Round Secure Aggregation with Verifiability for Federated Learning <i>Lang Pu, Jingjing Gu, Chao Lin, Xinyi Huang</i>	30136
R.I.P.: Better Models by Survival of the Fittest Prompts <i>Ping Yu, Weizhe Yuan, Olga Golovneva, Tianhao Wu, Sainbayar Sukhbaatar, Jason E Weston, Jing Xu</i>	30155
Data-Driven Selection of Instrumental Variables for Additive Nonlinear, Constant Effects Models <i>Xichen Guo, Feng Xie, Yan Zeng, Hao Zhang, Zhi Geng</i>	30180
UniMC: Taming Diffusion Transformer for Unified Keypoint-Guided Multi-Class Image Generation <i>Qin Guo, Ailing Zeng, Dongxu Yue, Ceyuan Yang, Yang Cao, Hanzhong Guo, Fei Shen, Wei Liu, Xihui Liu, Dan Xu</i>	30201
Modular Duality in Deep Learning <i>Jeremy Bernstein, Laker Newhouse</i>	30219
From Kernels to Features: A Multi-Scale Adaptive Theory of Feature Learning <i>Noa Rubin, Kirsten Fischer, Javed Lindner, Inbar Seroussi, Zohar Ringel, Michael Krämer, Moritz Helias</i>	30230
LAION-C: An Out-of-Distribution Benchmark for Web-Scale Vision Models <i>Fanfei Li, Thomas Klein, Wieland Brendel, Robert Geirhos, Roland S. Zimmermann</i>	30263
Overcoming Non-monotonicity in Transducer-based Streaming Generation <i>Zhengrui Ma, Yang Feng, Min zhang</i>	30295

Contextual Online Decision Making with Infinite-Dimensional Functional Regression <i>Haichen Hu, Rui Ai, Stephen Bates, David Simchi-Levi</i>	30312
DipLLM: Fine-Tuning LLM for Strategic Decision-making in Diplomacy <i>Kaixuan Xu, Jiajun Chai, Sicheng Li, Yuqian Fu, Yuanheng Zhu, Dongbin Zhao</i>	30339
Multiobjective distribution matching <i>Xiaoyuan Zhang, Peijie Li, Yingying Yu, Yichi Zhang, Han Zhao, Qingfu Zhang</i>	30357
Redundancy Undermines the Trustworthiness of Self-Interpretable GNNs <i>Wenxin Tai, Ting Zhong, Goce Trajcevski, Fan Zhou</i>	30371
Secant Line Search for Frank-Wolfe Algorithms <i>Deborah Hendrych, Sebastian Pokutta, Mathieu Besançon, David Martínez-Rubio</i>	30391
Leveraging Model Guidance to Extract Training Data from Personalized Diffusion Models <i>Xiaoyu Wu, Jiaru Zhang, Steven Wu</i>	30416
Olica: Efficient Structured Pruning of Large Language Models without Retraining <i>Jiujun He, Huazhen Lin</i>	30434
Scaling Laws for Pre-training Agents and World Models <i>Tim Pearce, Tabish Rashid, David Bignell, Raluca Georgescu, Sam Devlin, Katja Hofmann</i>	30449
CFPT: Empowering Time Series Forecasting through Cross-Frequency Interaction and Periodic-Aware Timestamp Modeling <i>Feifei Kou, Jiahao Wang, Lei Shi, Yuhan Yao, Yawen Li, Suguo Zhu, Zhongbao Zhang, Junping Du</i>	30470
Roll the dice & look before you leap: Going beyond the creative limits of next-token prediction <i>Vaishnavh Nagarajan, Chen Henry Wu, Charles Ding, Aditi Raghunathan</i>	30484
InfAlign: Inference-aware language model alignment <i>Ananth Balashankar, Ziteng Sun, Jonathan Berant, Jacob Eisenstein, Michael Collins, Adrian Hutter, Jong Lee, Chirag Nagpal, Flavien Prost, Aradhana Sinha, Ananda Theertha Suresh, Ahmad Beirami</i>	30526
A Mixture-Based Framework for Guiding Diffusion Models <i>Yazid Janati, Badr MOUFAD, Mehdi Abou El Qassime, Alain Oliviero Durmus, Eric Moulines, Jimmy Olsson</i>	30553
De-AntiFake: Rethinking the Protective Perturbations Against Voice Cloning Attacks <i>Wei Fan, Kejiang Chen, Chang Liu, Weiming Zhang, Nenghai Yu</i>	30600
WMarkGPT: Watermarked Image Understanding via Multimodal Large Language Models <i>Songbai Tan, Xuerui Qiu, Yao Shu, Gang Xu, Linrui Xu, Xiangyu Xu, Huiping Zhuang, Ming Li, Fei Yu</i>	30619
Bipartite Ranking From Multiple Labels: On Loss Versus Label Aggregation <i>Michal Lukasik, Lin Chen, Harikrishna Narasimhan, Aditya Krishna Menon, Wittawat Jitkrittum, Felix X. Yu, Sashank J. Reddi, Gang Fu, Mohammadhossein Bateni, Sanjiv Kumar</i>	30635

A Variational Framework for Improving Naturalness in Generative Spoken Language Models <i>Li-Wei Chen, Takuya Higuchi, Zakaria Aldeneh, Ahmed Hussen Abdelaziz, Alexander Rudnicky</i>	30664
On the Generalization Ability of Next-Token-Prediction Pretraining <i>Zhihao Li, Xue Jiang, Liyuan Liu, Xuelin Zhang, Hong Chen, Feng Zheng</i>	30683
C-3PO: Compact Plug-and-Play Proxy Optimization to Achieve Human-like Retrieval-Augmented Generation <i>Guoxin Chen, Minpeng Liao, Peiying Yu, Dingmin Wang, Zile Qiao, Chao Yang, Xin Zhao, Kai Fan</i>	30716
Measuring Variable Importance in Heterogeneous Treatment Effects with Confidence <i>Joseph Paillard, Angel David REYERO LOBO, Vitaliy Kolodyazhniy, Bertrand Thirion, Denis-Alexander Engemann</i>	30738
Laplace Transform Based Low-Complexity Learning of Continuous Markov Semigroups <i>Vladimir R Kostic, Karim Lounici, Hélène Halconrui, Timothée Devergne, Pietro Novelli, Massimiliano Pontil</i>	30760
HPS: Hard Preference Sampling for Human Preference Alignment <i>Xiandong Zou, Wanyu Lin, Yuchen Li, Pan Zhou</i>	30790
Active Fine-Tuning of Multi-Task Policies <i>Marco Bagatella, Jonas Hübötter, Georg Martius, Andreas Krause</i>	30814
ExtPose: Robust and Coherent Pose Estimation by Extending ViTs <i>Rongyu Chen, Li'an Zhuo, Linlin Yang, Qi WANG, Liefeng Bo, Bang Zhang, Angela Yao</i>	30847
Fixed-Confidence Multiple Change Point Identification under Bandit Feedback <i>Joseph Lazzaro, Ciara Pike-Burke</i>	30861
LMRL Gym: Benchmarks for Multi-Turn Reinforcement Learning with Language Models <i>Marwa Abdulhai, Isadora White, Charlie Victor Snell, Charles Sun, Joey Hong, Yuexiang Zhai, Kelvin Xu, Sergey Levine</i>	30890
What If We Recaption Billions of Web Images with LLaMA-3? <i>Xianhang Li, Haoqin Tu, Mude Hui, Zeyu Wang, Bingchen Zhao, Junfei Xiao, Sucheng Ren, Jieru Mei, Qing Liu, Huangjie Zheng, Yuyin Zhou, Cihang Xie</i>	30918
FSL-SAGE: Accelerating Federated Split Learning via Smashed Activation Gradient Estimation <i>Srijith Nair, Michael Lin, Peizhong Ju, Amirreza Talebi, Elizabeth Serena Bentley, Jia Liu</i>	30938
CEGA: A Cost-Effective Approach for Graph-Based Model Extraction and Acquisition <i>Zebin Wang, Menghan Lin, Bolin Shen, Ken Anderson, Molei Liu, Tianxi Cai, Yushun Dong</i>	30960
Optimal Algorithm for Max-Min Fair Bandit <i>Zilong Wang, Zhiyao Zhang, Shuai Li</i>	30977

Fully Dynamic Embedding into ℓ_p Spaces <i>Kiarash Banihashem, Xiang Chen, MohammadTaghi Hajiaghayi, Sungchul Kim, Kanak Mahadik, Ryan A. Rossi, Tong Yu</i>	30993
LEAPS: A discrete neural sampler via locally equivariant networks <i>Peter Holderrieth, Michael Samuel Albergo, Tommi Jaakkola</i>	31009
Positive-unlabeled AUC Maximization under Covariate Shift <i>Atsutoshi Kumagai, Tomoharu Iwata, Hiroshi Takahashi, Taishi Nishiyama, Kazuki Adachi, Yasuhiro Fujiwara</i>	31029
Improving Generalization in Federated Learning with Highly Heterogeneous Data via Momentum-Based Stochastic Controlled Weight Averaging <i>Junkang Liu, Yuanyuan Liu, Fanhua Shang, Hongying Liu, Jin Liu, Wei Feng</i>	31045
Voronoi-grid-based Pareto Front Learning and Its Application to Collaborative Federated Learning <i>Mengmeng Chen, Xiaohu Wu, QIQI LIU, Tiantian He, Yew-Soon Ong, Yaochu Jin, Qicheng Lao, Han Yu</i>	31091
MAGELLAN: Metacognitive predictions of learning progress guide autotelic LLM agents in large goal spaces <i>Loris Gaven, Thomas Carta, Clément ROMAC, Cédric Colas, sylvain lamprier, Olivier Sigaud, Pierre-Yves Oudeyer</i>	31111
Joker: Joint Optimization Framework for Lightweight Kernel Machines <i>Junhong Zhang, Zhihui Lai</i>	31145
Prompt-based Depth Pruning of Large Language Models <i>Juyun Wee, Minjae Park, Jaeho Lee</i>	31162
Discovering Spoofing Attempts on Language Model Watermarks <i>Thibaud Gloaguen, Nikola Jovanović, Robin Staab, Martin Vechev</i>	31175
Retraining-free Merging of Sparse MoE via Hierarchical Clustering <i>I-Chun Chen, Hsu-Shen Liu, Wei-Fang Sun, Chen-Hao Chao, Yen-Chang Hsu, Chun-Yi Lee</i>	31198
Q-Supervised Contrastive Representation: A State Decoupling Framework for Safe Offline Reinforcement Learning <i>Zhihe Yang, Yunjian Xu, Yang Zhang</i>	31225
CoreMatching: A Co-adaptive Sparse Inference Framework with Token and Neuron Pruning for Comprehensive Acceleration of Vision-Language Models <i>Qinsi Wang, Hancheng Ye, Ming-Yu Chung, Yudong Liu, Yueqian Lin, Martin Kuo, Mingyuan Ma, Jianyi Zhang, Yiran Chen</i>	31253
Craftium: Bridging Flexibility and Efficiency for Rich 3D Single- and Multi-Agent Environments <i>Mikel Malagón, Josu Ceberio, Jose A. Lozano</i>	31270

Stay Hungry, Keep Learning: Sustainable Plasticity for Deep Reinforcement Learning <i>Huaicheng Zhou, Zifeng Zhuang, Donglin Wang</i>	31300
Universal Biological Sequence Reranking for Improved De Novo Peptide Sequencing <i>Zijie Qiu, Jiaqi Wei, Xiang Zhang, Sheng Xu, Kai Zou, Zhi Jin, ZhiQiang Gao, Nanqing Dong, Siqi Sun</i>	31329
Automatic Reward Shaping from Confounded Offline Data <i>Mingxuan Li, Junzhe Zhang, Elias Bareinboim</i>	31354
Learning Vision and Language Concepts for Controllable Image Generation <i>Shaoan Xie, Lingjing Kong, Yujia Zheng, Zeyu Tang, Eric P. Xing, Guangyi Chen, Kun Zhang</i>	31383
Low-Rank Tensor Transitions (LoRT) for Transferable Tensor Regression <i>Andong Wang, Yuning Qiu, Zhong Jin, Guoxu Zhou, Qibin Zhao</i>	31406
PyTDC: A multimodal machine learning training, evaluation, and inference platform for biomedical foundation models <i>Alejandro Velez-Arce, Marinka Zitnik</i>	31456
PRIME: Deep Imbalanced Regression with Proxies <i>Jongin Lim, Sucheol Lee, Daeho Um, Sung-Un Park, Jinwoo Shin</i>	31496
Private Model Personalization Revisited <i>Conor Snedeker, Xinyu Zhou, Raef Bassily</i>	31521
CoMemo: LVLMs Need Image Context with Image Memory <i>Shi Liu, Weijie Su, Xizhou Zhu, Wenhai Wang, Jifeng Dai</i>	31571
Efficient Bisection Projection to Ensure Neural-Network Solution Feasibility for Optimization over General Set <i>Enming Liang, Minghua Chen</i>	31588
Understanding the Skill Gap in Recurrent Language Models: The Role of the Gather-and-Aggregate Mechanism <i>Aviv Bick, Eric P. Xing, Albert Gu</i>	31617
Activation Space Interventions Can Be Transferred Between Large Language Models <i>Narmeen Fatimah Oozeer, Dhruv Nathawani, Nirmalendu Prakash, Michael Lan, Abir HARRASSE, Amir Abdullah</i>	31638
Weakly Supervised Anomaly Detection via Dual-Tailed Kernel <i>Walid Durani, Tobias Nitzl, Claudia Plant, Christian Böhm</i>	31713
Token Assorted: Mixing Latent and Text Tokens for Improved Language Model Reasoning <i>DiJia Su, Hanlin Zhu, Yingchen Xu, Jiantao Jiao, Yuandong Tian, Qinqing Zheng</i>	31747
Graph Minimum Factorization Distance and Its Application to Large-Scale Graph Data Clustering <i>Jicong Fan</i>	31767

Integration-free Kernels for Equivariant Gaussian Process Modelling <i>Tim Steinert, David Ginsbourger, August Lykke-Møller, Ove Christiansen, Henry Moss</i>	31793
The impact of uncertainty on regularized learning in games <i>Pierre-Louis Cauvin, Davide Legacci, Panayotis Mertikopoulos</i>	31818
ReFrame: Layer Caching for Accelerated Inference in Real-Time Rendering <i>Lufei Liu, Tor M. Aamodt</i>	31856
Learning Robust Neural Processes with Risk-Averse Stochastic Optimization <i>Huafeng Liu, Yiran Fu, Liping Jing, Hui Li, Shuyang Lin, Jingyue Shi, Deqiang Ouyang, Jian Yu</i>	31873
Efficient Parallel Training Methods for Spiking Neural Networks with Constant Time Complexity <i>Wanjin Feng, Xingyu Gao, Wenqian Du, Hailong Shi, Peilin Zhao, Pengcheng Wu, Chunyan Miao</i>	31897
Discovering Physics Laws of Dynamical Systems via Invariant Function Learning <i>Shurui Gui, Xiner Li, Shuiwang Ji</i>	31913
Learning Joint Interventional Effects from Single-Variable Interventions in Additive Models <i>Armin Kekić, Sergio Hernan Garrido Mejia, Bernhard Schölkopf</i>	31945
QuEST: Stable Training of LLMs with 1-Bit Weights and Activations <i>Andrei Panferov, Jiale Chen, Soroush Tabesh, Mahdi Nikdan, Dan Alistarh</i>	31964
Counterfactual Voting Adjustment for Quality Assessment and Fairer Voting in Online Platforms with Helpfulness Evaluation <i>Chang Liu, Yixin Wang, Moontae Lee</i>	31981
Why Has Predicting Downstream Capabilities of Frontier AI Models with Scale Remained Elusive? <i>Rylan Schaeffer, Hailey Schoelkopf, Brando Miranda, Gabriel Mukobi, Varun Madan, Adam Ibrahim, Herbie Bradley, Stella Biderman, Sanmi Koyejo</i>	32000
Flowing Datasets with Wasserstein over Wasserstein Gradient Flows <i>Clément Bonet, Christophe Vauthier, Anna Korba</i>	32090
Divide and Conquer: Learning Label Distribution with Subtasks <i>Haitao Wu, Weiwei Li, Xiuyi Jia</i>	32147
FloE: On-the-Fly MoE Inference on Memory-constrained GPU <i>Yuxin Zhou, Zheng Li, Jun Zhang, Jue Wang, Yiping Wang, Zhongle Xie, Ke Chen, Lidan Shou</i>	32166
Contrastive Learning with Simplicial Convolutional Networks for Short-Text Classification <i>Huang Liang, Benedict Lee, Daniel Hui Loong Ng, Kelin Xia</i>	32190
Human-Aligned Image Models Improve Visual Decoding from the Brain <i>Nona Rajabi, Antonio H. Ribeiro, Miguel Vasco, Farzaneh Taleb, Mårten Björkman, Danica Kragic</i>	32204

On Mitigating Affinity Bias through Bandits with Evolving Biased Feedback <i>Matthew Faw, Constantine Caramanis, Jessica Hoffmann</i>	32234
Adaptive Elicitation of Latent Information Using Natural Language <i>Jimmy Wang, Thomas P Zollo, Richard Zemel, Hongseok Namkoong</i>	32282
FairPFN: A Tabular Foundation Model for Causal Fairness <i>Jake Robertson, Noah Hollmann, Samuel Müller, Noor Awad, Frank Hutter</i>	32303
Quantifying Treatment Effects: Estimating Risk Ratios via Observational Studies <i>Ahmed BOUGHDIRI, Julie Josse, Erwan Scornet</i>	32325
DiMa: Understanding the Hardness of Online Matching Problems via Diffusion Models <i>Boyu Zhang, Aocheng Shen, Bing Liu, Qiankun Zhang, Bin Yuan, Jing Wang, Shenghao Liu, Xianjun Deng</i>	32387
Deep Electromagnetic Structure Design Under Limited Evaluation Budgets <i>Shijian Zheng, Fangxiao Jin, Shuhai Zhang, Quan Xue, Mingkui Tan</i>	32411
What Has a Foundation Model Found? Using Inductive Bias to Probe for World Models <i>Keyon Vafa, Peter G. Chang, Ashesh Rambachan, Sendhil Mullainathan</i>	32431
Understanding Multimodal LLMs Under Distribution Shifts: An Information-Theoretic Approach <i>Changdae Oh, Zhen Fang, Shawn Im, Xuefeng Du, Yixuan Li</i>	32452
Strengthen Out-of-Distribution Detection Capability with Progressive Self-Knowledge Distillation <i>Yang Yang, Haonan Xu</i>	32480
Improving Transformer World Models for Data-Efficient RL <i>Antoine Dedieu, Joseph Ortiz, Xinghua Lou, Carter Wendelken, J Swaroop Guntupalli, Wolfgang Lehrach, Miguel Lazaro-Gredilla, Kevin Patrick Murphy</i>	32498
Low-Rank Thinning <i>Annabelle Michael Carrell, Albert Gong, Abhishek Shetty, Raaz Dwivedi, Lester Mackey</i>	32523
Synonymous Variational Inference for Perceptual Image Compression <i>Zijian Liang, Kai Niu, Changshuo Wang, Jin Xu, Ping Zhang</i>	32561
Deliberation in Latent Space via Differentiable Cache Augmentation <i>Luyang Liu, Jonas Pfeiffer, Jiaying Wu, Jun Xie, Arthur Szlam</i>	32592
Minerva: A Programmable Memory Test Benchmark for Language Models <i>Menglin Xia, Victor Rühle, Saravan Rajmohan, Reza Shokri</i>	32606
Byzantine-Resilient Federated Alternating Gradient Descent and Minimization for Partly-Decoupled Low Rank Matrix Learning <i>Ankit Pratap Singh, Ahmed Ali Abbasi, Namrata Vaswani</i>	32625

Fraud-Proof Revenue Division on Subscription Platforms <i>Abheek Ghosh, Tzeh Yuan Neoh, Nicholas Teh, Giannis Tyrovolas</i>	32643
Tractable Transformers for Flexible Conditional Generation <i>Anji Liu, Xuejie Liu, Dayuan Zhao, Mathias Niepert, Yitao Liang, Guy Van den Broeck</i>	32654
Test-Time Learning for Large Language Models <i>Jinwu Hu, Zitian Zhang, Guohao Chen, Xutao Wen, Chao Shuai, Wei Luo, Bin Xiao, Yuanqing Li, Minghui Tan</i>	32678
On the Local Complexity of Linear Regions in Deep ReLU Networks <i>Niket Nikul Patel, Guido Montufar</i>	32705
LASER: Attention with Exponential Transformation <i>Sai Surya Duvvuri, Inderjit S Dhillon</i>	32741
Stochastic Online Conformal Prediction with Semi-Bandit Feedback <i>Haosen Ge, Hamsa Bastani, Osbert Bastani</i>	32757
AffinityFlow: Guided Flows for Antibody Affinity Maturation <i>Can Chen, Karla-Luise Herpoldt, Chenchao Zhao, Zichen Wang, Marcus D. Collins, Shang Shang, Ron Benson</i>	32769
Beyond Confidence: Exploiting Homogeneous Pattern for Semi-Supervised Semantic Segmentation <i>Rui Sun, Huayu Mai, Wangkai Li, Yujia Chen, Naisong Luo, Yuan Wang, Tianzhu Zhang</i>	32783
Radio: Rate-Distortion Optimization for Large Language Model Compression <i>Sean I. Young</i>	32803
Context is Key: A Benchmark for Forecasting with Essential Textual Information <i>Andrew Robert Williams, Arjun Ashok, Étienne Marcotte, Valentina Zantedeschi, Jithendaraa Subramanian, Roland Riachi, James Requeima, Alexandre Lacoste, Irina Rish, Nicolas Chapados, Alexandre Drouin</i>	32821
Boost-and-Skip: A Simple Guidance-Free Diffusion for Minority Generation <i>Soobin Um, Beomsu Kim, Jong Chul Ye</i>	32879
Kandinsky Conformal Prediction: Beyond Class- and Covariate-Conditional Coverage <i>Konstantina Bairaktari, Jiayun Wu, Steven Wu</i>	32908
Safely Learning Optimal Auctions: A Testable Learning Framework for Mechanism Design <i>Vikram Kher, Manolis Zampetakis</i>	32930
How to Synthesize Text Data without Model Collapse? <i>Xuekai Zhu, Daixuan Cheng, Hengli Li, Kaiyan Zhang, Ermo Hua, Xingtai Lv, Ning Ding, Zhouhan Lin, Zilong Zheng, Bowen Zhou</i>	32947

EmoGrowth: Incremental Multi-label Emotion Decoding with Augmented Emotional Relation Graph <i>Kaicheng Fu, Changde Du, Jie Peng, Kunpeng Wang, Shuangchen Zhao, Xiaoyu Chen, Huiguang He</i>	32973
Learn Beneficial Noise as Graph Augmentation <i>Siqi Huang, Yanchen Xu, Hongyuan Zhang, Xuelong Li</i>	32991
Extracting Rare Dependence Patterns via Adaptive Sample Reweighting <i>Yiqing Li, Yewei Xia, Xiaofei Wang, Zhengming Chen, Lihua Peng, Mingming Gong, Kun Zhang</i>	33006
Feature Importance Metrics in the Presence of Missing Data <i>Henrik von Kleist, Joshua Wendland, Ilya Shpitser, Carsten Marr</i>	33041
Perceptual-GS: Scene-adaptive Perceptual Densification for Gaussian Splatting <i>Hongbi Zhou, Zhangkai Ni</i>	33062
Rethinking Causal Ranking: A Balanced Perspective on Uplift Model Evaluation <i>Minqin Zhu, Zexu Sun, Ruoxuan Xiong, Anpeng Wu, Baohong Li, Caizhi Tang, JUN ZHOU, Fei Wu, Kun Kuang</i>	33088
Understanding Chain-of-Thought in LLMs through Information Theory <i>Jean-Francois Ton, Muhammad Faaiz Taufiq, Yang Liu</i>	33109
3D Question Answering via only 2D Vision-Language Models <i>FENGYUN WANG, Sicheng Yu, Jiawei Wu, Jinhui Tang, Hanwang Zhang, Qianru Sun</i>	33137
Ensemble Learned Bloom Filters: Two Oracles are Better than One <i>Ming Lin, Lin CHEN</i>	33153
ENSUR: Equitable and Statistically Unbiased Recommendation <i>Nitin Bisht, Xiuwen Gong, Guandong Xu</i>	33171
An analytic theory of creativity in convolutional diffusion models <i>Mason Kamb, Surya Ganguli</i>	33191
Multi-agent Architecture Search via Agentic Supernet <i>Guibin Zhang, Luyang Niu, Junfeng Fang, Kun Wang, LEI BAI, Xiang Wang</i>	33228
Subobject-level Image Tokenization <i>Delong Chen, Samuel Cahyawijaya, Jianfeng Liu, Baoyuan Wang, Pascale Fung</i>	33247
High Dynamic Range Novel View Synthesis with Single Exposure <i>Kaixuan Zhang, HuWang, Minxian Li, Mingwu Ren, Mao Ye, Xiatian Zhu</i>	33267
Unifews: You Need Fewer Operations for Efficient Graph Neural Networks <i>Ningyi Liao, Zihao Yu, Ruixiao Zeng, Siqiang Luo</i>	33279
Discrete Neural Algorithmic Reasoning <i>Gleb Rodionov, Liudmila Prokhorenkova</i>	33302

RePaViT: Scalable Vision Transformer Acceleration via Structural Reparameterization on Feedforward Network Layers <i>Xuwei Xu, Yang Li, Yudong Chen, Jiajun Liu, Sen Wang</i>	33317
Pessimism Principle Can Be Effective: Towards a Framework for Zero-Shot Transfer Reinforcement Learning <i>Chi Zhang, Ziyang Jia, George K. Atia, Sihong He, Yue Wang</i>	33331
Foundation Molecular Grammar: Multi-Modal Foundation Models Induce Interpretable Molecular Graph Languages <i>Michael Sun, Weize Yuan, Gang Liu, Wojciech Matusik, Jie Chen</i>	33377
ZeroFlow: Overcoming Catastrophic Forgetting is Easier than You Think <i>Tao Feng, Wei Li, Didi Zhu, Hangjie Yuan, Wendi Zheng, Dan Zhang, Jie Tang</i>	33429
FlipAttack: Jailbreak LLMs via Flipping <i>Yue Liu, Xiaoxin He, Miao Xiong, Jinlan Fu, Shumin Deng, YINGWEI MA, Jiaheng Zhang, Bryan Hooi</i>	33446
Self-Supervised Transformers as Iterative Solution Improvers for Constraint Satisfaction <i>Yudong Xu, Wenhao Li, Scott Sanner, Elias Boutros Khalil</i>	33487
Understanding Generalization in Quantum Machine Learning with Margins <i>Tak Hur, Daniel K. Park</i>	33506
Learngene Tells You How to Customize: Task-Aware Parameter Initialization at Flexible Scales <i>Jiaze Xu, Shiyu Xia, Xu Yang, Jiaqi Lv, Xin Geng</i>	33529
ML²-GCL: Manifold Learning Inspired Lightweight Graph Contrastive Learning <i>Jianqing Liang, Zhiqiang Li, Xinkai Wei, Yuan Liu, Zhiqiang Wang</i>	33542
On the Training Convergence of Transformers for In-Context Classification of Gaussian Mixtures <i>Wei Shen, Ruida Zhou, Jing Yang, Cong Shen</i>	33557
How to set AdamW's weight decay as you scale model and dataset size <i>Xi Wang, Laurence Aitchison</i>	33597
AdvI2I: Adversarial Image Attack on Image-to-Image Diffusion Models <i>Yaopei Zeng, Yuanpu Cao, Bochuan Cao, Yurui Chang, Jinghui Chen, Lu Lin</i>	33626
Annealing Flow Generative Models Towards Sampling High-Dimensional and Multi-Modal Distributions <i>Dongze Wu, Yao Xie</i>	33640
Automated Hypothesis Validation with Agentic Sequential Falsifications <i>Kexin Huang, Ying Jin, Ryan Li, Michael Y. Li, Emmanuel Candes, Jure Leskovec</i>	33669
Attention-Only Transformers via Unrolled Subspace Denoising <i>Peng Wang, Yifu Lu, Yaodong Yu, Druv Pai, Qing Qu, Yi Ma</i>	33735

Preference Learning for AI Alignment: a Causal Perspective <i>Kasia Kobalczyk, Mihaela van der Schaar</i>	33755
Meta Optimality for Demographic Parity Constrained Regression via Post-Processing <i>Kazuto Fukuchi</i>	33776
Random Feature Representation Boosting <i>Nikita Zozoulenko, Thomas Cass, Lukas Gonon</i>	33799
The Batch Complexity of Bandit Pure Exploration <i>Adrienne Tuynman, Rémy Degenne</i>	33832
Optimal Decision Tree Pruning Revisited: Algorithms and Complexity <i>Juha Harviainen, Frank Sommer, Manuel Sorge, Stefan Szeider</i>	33859
LLMs on the Line: Data Determines Loss-to-Loss Scaling Laws <i>Prasanna Mayilvahanan, Thaddäus Wiedemer, Sayak Mallick, Matthias Bethge, Wieland Brendel</i>	33886
Complex Wavelet Mutual Information Loss: A Multi-Scale Loss Function for Semantic Segmentation <i>Renhao Lu</i>	33915
Explicit Preference Optimization: No Need for an Implicit Reward Model <i>Xiangkun Hu, Lemin Kong, Tong He, David Wipf</i>	33927
BSLoRA: Enhancing the Parameter Efficiency of LoRA with Intra-Layer and Inter-Layer Sharing <i>Yuhua Zhou, Ruifeng Li, Changhai Zhou, Fei Yang, Aimin PAN</i>	33950
FrameBridge: Improving Image-to-Video Generation with Bridge Models <i>Yuji Wang, Zehua Chen, Chen Xiaoyu, Yixiang Wei, Jun Zhu, Jianfei Chen</i>	33970
ExLM: Rethinking the Impact of [MASK] Tokens in Masked Language Models <i>Kangjie Zheng, Junwei Yang, Siyue Liang, Bin Feng, Zequn Liu, Wei Ju, Zhiping Xiao, Ming Zhang</i>	34001
MedXpertQA: Benchmarking Expert-Level Medical Reasoning and Understanding <i>Yuxin Zuo, Shang Qu, Yifei Li, Zhang-Ren Chen, Xuekai Zhu, Ermo Hua, Kaiyan Zhang, Ning Ding, Bowen Zhou</i>	34030
Task Generalization with Autoregressive Compositional Structure: Can Learning from $\mathcal{D}$ Tasks Generalize to $\mathcal{D}^T$ Tasks? <i>Amirhesam Abedsoltan, Huaqing Zhang, Kaiyue Wen, Hongzhou Lin, Jingzhao Zhang, Mikhail Belkin</i>	34060
Unpaired Point Cloud Completion via Unbalanced Optimal Transport <i>Taekyung Lee, Jaemoo Choi, Jaewoong Choi, Myungjoo Kang</i>	34080
Causal Abstraction Learning based on the Semantic Embedding Principle <i>Gabriele D'Acunto, Fabio Massimo Zennaro, Yorgos Felekis, Paolo Di Lorenzo</i>	34102

On Fine-Grained Distinct Element Estimation <i>Ilias Diakonikolas, Daniel Kane, Jasper C.H. Lee, Thanasis Pittas, David Woodruff, Samson Zhou</i>	34139
BARNN: A Bayesian Autoregressive and Recurrent Neural Network <i>Dario Coscia, Max Welling, Nicola Demo, Gianluigi Rozza</i>	34175
A Large Recurrent Action Model: xLSTM enables Fast Inference for Robotics Tasks <i>Thomas Schmied, Thomas Adler, Vihang Prakash Patil, Maximilian Beck, Korbinian Pöppel, Johannes Brandstetter, Günter Klambauer, Razvan Pascanu, Sepp Hochreiter</i>	34201
One-Pass Feature Evolvable Learning with Theoretical Guarantees <i>Cun-Yuan Xing, Meng-Zhang Qian, Wu-Yang Chen, Wei Gao, Zhi-Hua Zhou</i>	34246
Mind Your Step (by Step): Chain-of-Thought can Reduce Performance on Tasks where Thinking Makes Humans Worse <i>Ryan Liu, Jiayi Geng, Addison J. Wu, Ilia Sucholutsky, Tania Lombrozo, Thomas L. Griffiths</i>	34271
DeepCrossAttention: Supercharging Transformer Residual Connections <i>Mike Heddes, Adel Javanmard, Kyriakos Axiotis, Gang Fu, Mohammadhossein Bateni, Vahab Mirrokni</i>	34300
Satori: Reinforcement Learning with Chain-of-Action-Thought Enhances LLM Reasoning via Autoregressive Search <i>Maohao Shen, Guangtao Zeng, Zhenting Qi, Zhang-Wei Hong, Zhenfang Chen, Wei Lu, Gregory W. Wornell, Subhro Das, David Daniel Cox, Chuang Gan</i>	34323
Lean and Mean Adaptive Optimization via Subset-Norm and Subspace-Momentum with Convergence Guarantees <i>Thien Hang Nguyen, Huy Nguyen</i>	34361
Gradient Inversion of Multimodal Models <i>Omri Ben Hemo, Alon Zolfi, Oryan Yehezkel, Omer Hofman, Roman Vainshtein, Hisashi Kojima, Yuval Elovici, Asaf Shabtai</i>	34407
Are Large Brainwave Foundation Models Capable Yet ? Insights from Fine-Tuning <i>Na Lee, Konstantinos Barmpas, Yannis Panagakis, Dimitrios Adamos, Nikolaos Laskaris, Stefanos Zafeiriou</i>	34424
Enhancing Decision-Making of Large Language Models via Actor-Critic <i>Heng Dong, Kefei Duan, Chongjie Zhang</i>	34435
Online Sparsification of Bipartite-Like Clusters in Graphs <i>Joyentanuj Das, Suranjan De, He Sun</i>	34472
Temporal Difference Flows <i>Jesse Farebrother, Matteo Pirota, Andrea Tirinzoni, Remi Munos, Alessandro Lazaric, Ahmed Touati</i>	34491
Layer-wise Quantization for Quantized Optimistic Dual Averaging <i>Anh Duc Nguyen, Ilia Markov, Frank Zhengqing Wu, Ali Ramezani-Kebrya, Kimon Antonakopoulos, Dan Alistarh, Volkan Cevher</i>	34533

SecEmb: Sparsity-Aware Secure Federated Learning of On-Device Recommender System with Large Embedding <i>Peihua Mai, Youlong Ding, Ziyang Lyu, Minxin Du, Yan Pang</i>	34580
Clustering Properties of Self-Supervised Learning <i>Xi Weng, Jianing An, Xudong Ma, Binhang Qi, Jie Luo, Xi Yang, Jin Song Dong, Lei Huang</i>	34606
Causal Discovery from Conditionally Stationary Time Series <i>Carles Ballells-Rodas, Xavier Sumba, Tanmayee Narendra, Ruibo Tu, Gabriele Schweikert, Hedvig Kjellstrom, Yingzhen Li</i>	34626
History-Guided Video Diffusion <i>Kiwhan Song, Boyuan Chen, Max Simchowitz, Yilun Du, Russ Tedrake, Vincent Sitzmann</i>	34653
Active feature acquisition via explainability-driven ranking <i>Osman Berke Guney, Ketan Suhaas Saichandran, Karim Elzokm, Ziming Zhang, Vijaya B Kolachalama</i>	34692
WorldSimBench: Towards Video Generation Models as World Simulators <i>Yiran Qin, Zhelun Shi, Jiwen Yu, Xijun Wang, Enshen Zhou, Lijun Li, Zhenfei Yin, Xihui Liu, Lu Sheng, Jing Shao, LEI BAI, Ruimao Zhang</i>	34710
EquivaMap: Leveraging LLMs for Automatic Equivalence Checking of Optimization Formulations <i>Haotian Zhai, Connor Lawless, Ellen Vitercik, Liu Leqi</i>	34735
The Relationship Between No-Regret Learning and Online Conformal Prediction <i>Ramya Ramalingam, Shayan Kiyani, Aaron Roth</i>	34753
Learning Event Completeness for Weakly Supervised Video Anomaly Detection <i>Yu Wang, Shiwei Chen</i>	34772
EvoControl: Multi-Frequency Bi-Level Control for High-Frequency Continuous Control <i>Samuel Holt, Todor Davchev, Dhruva Tirumala, Ben Moran, Atıl İscen, Antoine Laurens, Yixin Lin, Erik Frey, Markus Wulfmeier, Francesco Romano, Nicolas Heess</i>	34785
Fast, Accurate Manifold Denoising by Tunneling Riemannian Optimization <i>Shiyu Wang, Mariam Avagyan, Yihan Shen, Arnaud Lamy, Tingran Wang, Szabolcs Marka, Zsuzsanna Marka, John Wright</i>	34847
Efficient Graph Continual Learning via Lightweight Graph Neural Tangent Kernels-based Dataset Distillation <i>Rihong Qiu, Xinke Jiang, Yuchen Fang, Hongbin Lai, Hao Miao, Xu Chu, Junfeng Zhao, Yasha Wang</i>	34895
Distillation of Discrete Diffusion through Dimensional Correlations <i>Satoshi Hayakawa, Yuhta Takida, Masaaki Imaizumi, Hiromi Wakaki, Yuki Mitsufuji</i>	34920
Graph Neural Network Generalization With Gaussian Mixture Model Based Augmentation <i>Yassine ABBAHADDOU, Fragkiskos D. Malliaros, Johannes F. Lutzeyer, Amine M. Aboussalah, Michalis Vazirgiannis</i>	34959

Diff-MoE: Diffusion Transformer with Time-Aware and Space-Adaptive Experts <i>Kun Cheng, Xiao He, Lei Yu, Zhijun Tu, Mingrui Zhu, Nannan Wang, Xinbo Gao, Jie Hu</i>	34983
Enhancing Target-unspecific Tasks through a Features Matrix <i>Fangming Cui, Yonggang Zhang, Xuan Wang, Xinmei Tian, Jun Yu</i>	34998
Symmetry-Aware GFlowNets <i>Hohyun Kim, Seunggeun Lee, Min-hwan Oh</i>	35011
Provably Efficient Algorithm for Best Scoring Rule Identification in Online Principal-Agent Information Acquisition <i>Zichen Wang, Chuanhao Li, Huazheng Wang</i>	35040
Direct Density Ratio Optimization: A Statistically Consistent Approach to Aligning Large Language Models <i>Rei Higuchi, Taiji Suzuki</i>	35068
Global Context-aware Representation Learning for Spatially Resolved Transcriptomics <i>Yunhak Oh, Junseok Lee, Yeongmin Kim, Sangwoo Seo, Namkyeong Lee, Chanyoung Park</i>	35084
Disparate Conditional Prediction in Multiclass Classifiers <i>Sivan Sabato, Eran Treister, Elad Yom-Tov</i>	35111
FireFlow: Fast Inversion of Rectified Flow for Image Semantic Editing <i>Yingying Deng, Xiangyu He, Changwang Mei, Peisong Wang, Fan Tang</i>	35129
FlowAR: Scale-wise Autoregressive Image Generation Meets Flow Matching <i>Sucheng Ren, Qihang Yu, Ju He, Xiaohui Shen, Alan Yuille, Liang-Chieh Chen</i>	35148
Adaptive Estimation and Learning under Temporal Distribution Shift <i>Dheeraj Baby, Yifei Tang, Hieu Duy Nguyen, Yu-Xiang Wang, Rohit Pyati</i>	35162
Accelerating Spectral Clustering under Fairness Constraints <i>Francesco Tonin, Alex Lambert, Johan Suykens, Volkan Cevher</i>	35189
Enabling Optimal Decisions in Rehearsal Learning under CARE Condition <i>Wen-Bo Du, Hao-Yi Lei, Lue Tao, Tian-Zuo Wang, Zhi-Hua Zhou</i>	35210
Counterfactual Effect Decomposition in Multi-Agent Sequential Decision Making <i>Stelios Triantafyllou, Aleksa Sukovic, Yasaman Zolfimoselo, Goran Radanovic</i>	35236
Continual Generalized Category Discovery: Learning and Forgetting from a Bayesian Perspective <i>Hao Dai, Jagmohan Chauhan</i>	35263
BoxLM: Unifying Structures and Semantics of Medical Concepts for Diagnosis Prediction in Healthcare <i>Yanchao Tan, Hang Lv, Yunfei Zhan, Guofang Ma, Bo Xiong, Carl Yang</i>	35283

MedRAX: Medical Reasoning Agent for Chest X-ray <i>Adibvafa Fallahpour, Jun Ma, Alif Munim, Hongwei Lyu, BO WANG</i>	35301
Fair Clustering via Alignment <i>Kunwoong Kim, Jihu Lee, Sangchul Park, Yongdai Kim</i>	35317
HGOT: Self-supervised Heterogeneous Graph Neural Network with Optimal Transport <i>Yanbei Liu, Chongxu Wang, Zhitao Xiao, Lei Geng, Yanwei Pang, Xiao Wang</i>	35354
Reconstructing Cell Lineage Trees from Phenotypic Features with Metric Learning <i>Da Kuang, GuanWen Qiu, Junhyong Kim</i>	35367
Continuously Updating Digital Twins using Large Language Models <i>Harry Amad, Nicolás Astorga, Mihaela van der Schaar</i>	35411
Enforcing Idempotency in Neural Networks <i>Nikolaj Banke Jensen, Jamie Vicary</i>	35435
Deep Ridgelet Transform and Unified Universality Theorem for Deep and Shallow Joint-Group-Equivariant Machines <i>Sho Sonoda, Yuka Hashimoto, Isao Ishikawa, Masahiro Ikeda</i>	35456
Three-Dimensional Trajectory Prediction with 3DMoTraj Dataset <i>Hao Zhou, Xu Yang, Mingyu Fan, Lu Qi, Xiangtai Li, Ming-Hsuan Yang, Fei Luo</i>	35475
Domain-Adapted Diffusion Model for PROTAC Linker Design Through the Lens of Density Ratio in Chemical Space <i>Zixing Song, Ziqiao Meng, José Miguel Hernández-Lobato</i>	35493
Bi-perspective Splitting Defense: Achieving Clean-Seed-Free Backdoor Security <i>Yangyang Shen, Xiao Tan, Dian Shen, Meng Wang, Beilun Wang</i>	35519
Direct Prediction Set Minimization via Bilevel Conformal Classifier Training <i>Yuanjie Shi, Hooman Shahrokhi, Xuesong Jia, Xiongzhi Chen, Jana Doppa, Yan Yan</i>	35545
Behavior-agnostic Task Inference for Robust Offline In-context Reinforcement Learning <i>Long Ma, Fangwei Zhong, Yizhou Wang</i>	35575
Leveraging Offline Data in Linear Latent Contextual Bandits <i>Chinmaya Kausik, Kevin Tan, Ambuj Tewari</i>	35591
Compositional Scene Understanding through Inverse Generative Modeling <i>Yanbo Wang, Justin Dauwels, Yilun Du</i>	35636
Test-Time Canonicalization by Foundation Models for Robust Perception <i>Utkarsh Singhal, Ryan Feng, Stella X. Yu, Atul Prakash</i>	35656
When, Where and Why to Average Weights? <i>Niccolò Ajroldi, Antonio Orvieto, Jonas Geiping</i>	35678

Two Tickets are Better than One: Fair and Accurate Hiring Under Strategic LLM Manipulations <i>Lee Cohen, Connie Hong, Jack Hsieh, Judy Hanwen Shen</i>	35695
SEMU: Singular Value Decomposition for Efficient Machine Unlearning <i>Marcin Sendera, Łukasz Struski, Kamil Książek, Kryspin Musiol, Jacek Tabor, Dawid Damian Rymarczyk</i>	35726
Parametric Scaling Law of Tuning Bias in Conformal Prediction <i>Hao Zeng, Kangdao Liu, Bingyi Jing, Hongxin Wei</i>	35750
Optimization Proxies using Limited Labeled Data and Training Time -- A Semi-Supervised Bayesian Neural Network Approach <i>Parikshit Pareek, Abhijith Jayakumar, Kaarthik Sundar, Sidhant Misra, Deepjyoti Deka</i>	35774
Sort Before You Prune: Improved Worst-Case Guarantees of the DiskANN Family of Graphs <i>Siddharth Gollapudi, Ravishankar Krishnaswamy, Kirankumar Shiragur, Harsh Wardhan</i>	35792
ITBench: Evaluating AI Agents across Diverse Real-World IT Automation Tasks <i>Saurabh Jha, Rohan R. Arora, Yuji Watanabe, Takumi Yanagawa, Yinfang Chen, Jackson Clark, Bhavya Bhavya, Mudit Verma, Harshit Kumar, Hirokuni Kitahara, Noah Zheutlin, Saki Takano, Divya Pathak, Felix George, Xinbo Wu, Bekir O Turkkkan, Gerard Vanloo, Michael Nidd, Ting Dai, Oishik Chatterjee, Pranjal Gupta, Suranjana Samanta, Pooja Aggarwal, Rong Lee, Jae-wook Ahn, Debanjana Kar, Amit Paradkar, Yu Deng, Pratibha Moogi, Prateeti Mohapatra, Naoki Abe, Chandrasekhar Narayanaswami, Tianyin Xu, Lav R. Varshney, Ruchi Mahindru, Anca Sailer, Laura Shwartz, Daby Sow, Nicholas C. M. Fuller, Ruchir Puri</i>	35814
A General Framework for Inference-time Scaling and Steering of Diffusion Models <i>Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen McKeown, Rajesh Ranganath</i>	35878
Inference-Time Alignment of Diffusion Models with Direct Noise Optimization <i>Zhiwei Tang, Jiangweizhi Peng, Jiasheng Tang, Mingyi Hong, Fan Wang, Tsung-Hui Chang</i>	35896
Revisiting Cooperative Off-Policy Multi-Agent Reinforcement Learning <i>Yueheng Li, Guangming Xie, Zongqing Lu</i>	35922
SKIM: Any-bit Quantization Pushing The Limits of Post-Training Quantization <i>Runsheng Bai, Bo Liu, qiang liu</i>	35938
Visual Autoregressive Modeling for Image Super-Resolution <i>Yunpeng Qu, Kun Yuan, Jinhua Hao, Kai Zhao, Qizhi Xie, Ming Sun, Chao Zhou</i>	35953
SHE: Streaming-media Hashing Retrieval <i>Ruitao Pu, Yang Qin, Xiaomin Song, Dezhong Peng, Zhenwen Ren, Yuan Sun</i>	35976
Variational Learning of Fractional Posteriors <i>Kian Ming A. Chai, Edwin V. Bonilla</i>	35989
Step-DAD: Semi-Amortized Policy-Based Bayesian Experimental Design <i>Marcel Hedman, Desi R. Ivanova, Cong Guan, Tom Rainforth</i>	36020

Trustworthy Machine Learning through Data-Specific Indistinguishability <i>Hanshen Xiao, Zhen Yang, G. Edward Suh</i>	36040
polybasic Speculative Decoding Through a Theoretical Perspective <i>Ruilin Wang, Huixia Li, Yuexiao Ma, Xiawu Zheng, Fei Chao, Xuefeng Xiao, Rongrong Ji</i>	36062
Understanding and Mitigating Memorization in Diffusion Models for Tabular Data <i>Zhengyu Fang, Zhimeng Jiang, Huiyuan Chen, Xiao Li, Jing Li</i>	36074
Rethinking the Bias of Foundation Model under Long-tailed Distribution <i>Jiahao Chen, Bin Qin, Jiangmeng Li, Hao Chen, Bing Su</i>	36104
A Mixed-Curvature based Pre-training Paradigm for Multi-Task Vehicle Routing Solver <i>Suyu Liu, Zhiguang Cao, Shanshan Feng, Yew-Soon Ong</i>	36125
Testing the Limits of Fine-Tuning for Improving Visual Cognition in Vision Language Models <i>Luca M. Schulze Buschoff, Konstantinos Voudouris, Elif Akata, Matthias Bethge, Joshua B. Tenenbaum, Eric Schulz</i>	36161
Black-Box Adversarial Attacks on LLM-Based Code Completion <i>Slobodan Jenko, Niels Mündler, Jingxuan He, Mark Vero, Martin Vechev</i>	36179
Fast Estimation of Partial Dependence Functions using Trees <i>Jinyang Liu, Tessa Steensgaard, Marvin N. Wright, Niklas Pfister, Munir Hiabu</i>	36199
Instruction-Following Pruning for Large Language Models <i>Bairu Hou, Qibin Chen, Jianyu Wang, Guoli Yin, Chong Wang, Nan Du, Ruoming Pang, Shiyu Chang, Tao Lei</i>	36238
Unnatural Languages Are Not Bugs but Features for LLMs <i>Keyu Duan, Yiran Zhao, Zhili Feng, Jinjie Ni, Tianyu Pang, Qian Liu, Tianle Cai, Longxu Dou, Kenji Kawaguchi, Anirudh Goyal, J Zico Kolter, Michael Qizhe Shieh</i>	36254
Spherical-Nested Diffusion Model for Panoramic Image Outpainting <i>Xiancheng Sun, Senmao Ma, Shengxi Li, Mai Xu, Jingyuan Xia, Lai Jiang, Xin Deng, Jiali Wang</i>	36269
A Parameter-Free and Near-Optimal Zeroth-Order Algorithm for Stochastic Convex Optimization <i>Kunjie Ren, Luo Luo</i>	36286
Neural Interpretable PDEs: Harmonizing Fourier Insights with Attention for Scalable and Interpretable Physics Discovery <i>Ning Liu, Yue Yu</i>	36307
Preference Optimization for Combinatorial Optimization Problems <i>Mingjun Pan, Guanquan Lin, You-Wei Luo, Bin Zhu, Zhien Dai, Lijun Sun, Chun Yuan</i>	36322
Federated Oriented Learning: A Practical One-Shot Personalized Federated Learning Framework <i>Guan Huang, Tao Shu</i>	36347

LongRoPE2: Near-Lossless LLM Context Window Scaling <i>Ning Shang, Li Lyna Zhang, Siyuan Wang, Gaokai Zhang, Gilsinia Lopez, Fan Yang, Weizhu Chen, Mao Yang</i>	36376
BalancEdit: Dynamically Balancing the Generality-Locality Trade-off in Multi-modal Model Editing <i>Dongliang Guo, Mengxuan Hu, Zihan Guan, Thomas Hartvigsen, Sheng Li</i>	36392
QuEst: Enhancing Estimates of Quantile-Based Distributional Measures Using Model Predictions <i>Zhun Deng, Thomas P Zollo, Benjamin Eyre, Amogh Inamdar, David Madras, Richard Zemel</i>	36407
DUNIA: Pixel-Sized Embeddings via Cross-Modal Alignment for Earth Observation Applications <i>Ibrahim Fayad, Max Zimmer, Martin Schwartz, Fabian Gieseke, Philippe CIAIS, Gabriel Belouze, Sarah Brood, Aurélien de Truchis, Alexandre d'Aspremont</i>	36429
Evolving Prompts In-Context: An Open-ended, Self-replicating Perspective <i>Jianyu Wang, Zhiqiang Hu, Lidong Bing</i>	36461
Learning Changes in Graphon Attachment Network Models <i>Xinyuan Fan, Bufan Li, Chenlei Leng, Weichi Wu</i>	36513
Aggregation of Dependent Expert Distributions in Multimodal Variational Autoencoders <i>Rogelio A. Mancisidor, Robert Jenssen, Shujian Yu, Michael Kampffmeyer</i>	36538
Assessing Safety Risks and Quantization-aware Safety Patching for Quantized Large Language Models <i>Kejia Chen, Jiawen Zhang, Jiacong Hu, Yu Wang, Jian Lou, Zunlei Feng, Mingli Song</i>	36564
Textual Unlearning Gives a False Sense of Unlearning <i>Jiacheng Du, Zhibo Wang, Jie Zhang, Xiaoyi Pang, Jiahui Hu, Kui Ren</i>	36583
CSTrack: Enhancing RGB-X Tracking via Compact Spatiotemporal Features <i>Xiaokun Feng, Dailing Zhang, Shiyu Hu, Xuchen Li, Meiqi Wu, Jing Zhang, Xiaotang Chen, Kaiqi Huang</i>	36602
Tackling Dimensional Collapse toward Comprehensive Universal Domain Adaptation <i>Hung-Chieh Fang, Po-Yi Lu, Hsuan-Tien Lin</i>	36621
Revisiting Instance-Optimal Cluster Recovery in the Labeled Stochastic Block Model <i>Kaito Ariu, Alexandre Proutiere, Se-Young Yun</i>	36639
WeGeFT: Weight-Generative Fine-Tuning for Multi-Faceted Efficient Adaptation of Large Models <i>Chinmay Savadikar, Xi Song, Tianfu Wu</i>	36679
PDUDT: Provable Decentralized Unlearning under Dynamic Topologies <i>Jing Qiao, Yu Liu, Zengzhe Chen, Mingyi Li, YUAN YUAN, Xiao Zhang, Dongxiao Yu</i>	36702
LLM Alignment as Retriever Optimization: An Information Retrieval Perspective <i>Bowen Jin, Jinsung Yoon, Zhen Qin, Ziqi Wang, Wei Xiong, Yu Meng, Jiawei Han, Serkan O Arik</i>	36731

Efficient and Scalable Density Functional Theory Hamiltonian Prediction through Adaptive Sparsity <i>Erpai Luo, Xinran Wei, Lin Huang, Yunyang Li, Han Yang, Zaishuo Xia, Zun Wang, Chang Liu, Bin Shao, Jia Zhang</i>	36751
IntLoRA: Integral Low-rank Adaptation of Quantized Diffusion Models <i>Hang Guo, Yawei Li, Tao Dai, Shu-Tao Xia, Luca Benini</i>	36774
Test-Time Adaptation for Online Vision-Language Navigation with Feedback-based Reinforcement Learning <i>Sungjune Kim, Gyeongrok Oh, Heeju Ko, Daehyun Ji, Dongwook Lee, Byung-Jun Lee, Sujin Jang, Sangpil Kim</i>	36796
Dynamical Modeling of Behaviorally Relevant Spatiotemporal Patterns in Neural Imaging Data <i>Sayed Mohammad Hosseini, Maryam Shanechi</i>	36814
Improving Model Alignment Through Collective Intelligence of Open-Source Models <i>Junlin Wang, Roy Xie, Shang Zhu, Jue WANG, Ben Athiwaratkun, Bhuwan Dhingra, Shuaiwen Leon Song, Ce Zhang, James Zou</i>	36841
Nemotron-CORTEXA: Enhancing LLM Agents for Software Engineering Tasks via Improved Localization and Solution Diversity <i>Atefeh Sohrabizadeh, Jialin Song, Mingjie Liu, Rajarshi Roy, Chankyu Lee, Jonathan Raiman, Bryan Catanzaro</i>	36863
Enhancing Rating-Based Reinforcement Learning to Effectively Leverage Feedback from Large Vision-Language Models <i>Tung Minh Luu, Younghwan Lee, Donghoon Lee, Sunho Kim, Min Jun Kim, Chang D. Yoo</i>	36879
Why Is Spatial Reasoning Hard for VLMs? An Attention Mechanism Perspective on Focus Areas <i>Shiqi Chen, Tongyao Zhu, Ruochen Zhou, Jinghan Zhang, Siyang Gao, Juan Carlos Niebles, Mor Geva, Junxian He, Jiajun Wu, Manling Li</i>	36907
On Efficient Estimation of Distributional Treatment Effects under Covariate-Adaptive Randomization <i>Undral Byambadalai, Tomu Hirata, Tatsushi Oka, Shota Yasui</i>	36930
VerbalTS: Generating Time Series from Texts <i>Shuqi Gu, Chuyue Li, Baoyu Jing, Kan Ren</i>	36954
Locate-then-edit for Multi-hop Factual Recall under Knowledge Editing <i>Zhuoran Zhang, Yongxiang Li, Zijian Kan, Keyuan Cheng, Lijie Hu, Di Wang</i>	36983
Learning Single Index Models with Diffusion Priors <i>Anqi Tang, Youming Chen, Shuchen Xue, Zhaoqiang Liu</i>	37006
A Reasoning-Based Approach to Cryptic Crossword Clue Solving <i>Martin Andrews, Sam Witteveen</i>	37030

Provable Benefit of Random Permutations over Uniform Sampling in Stochastic Coordinate Descent <i>Donghwa Kim, Jaewook Lee, Chulhee Yun</i>	37050
How to Evaluate and Mitigate IP Infringement in Visual Generative AI? <i>Zhenting Wang, Chen Chen, Vikash Sehwal, Minzhou Pan, Lingjuan Lyu</i>	37118
Graph-Supported Dynamic Algorithm Configuration for Multi-Objective Combinatorial Optimization <i>Robbert Reijnen, Yaoxin Wu, Zaharah Bukhsh, Yingqian Zhang</i>	37138
A Unified Framework for Generalization Error Analysis of Learning with Arbitrary Discrete Weak Features <i>Kosuke Sugiyama, Masato Uchida</i>	37154
Conditional Diffusion Model with Nonlinear Data Transformation for Time Series Forecasting <i>J Rishi, GVS Mothish, Deepak Subramani</i>	37175
Hardware and Software Platform Inference <i>Cheng Zhang, Hanna Foerster, Robert D. Mullins, Yiren Zhao, Ilia Shumailov</i>	37196
Optimal transport-based conformal prediction <i>Gauthier Thurin, Kimia Nadjahi, Claire Boyer</i>	37210
On Differential Privacy for Adaptively Solving Search Problems via Sketching <i>Shiyuan Feng, Ying Feng, George Zhaoqi Li, Zhao Song, David Woodruff, Lichen Zhang</i>	37229
Learning With Multi-Group Guarantees For Clusterable Subpopulations <i>Jessica Dai, Nika Haghtalab, Eric Zhao</i>	37265
GradPS: Resolving Futile Neurons in Parameter Sharing Network for Multi-Agent Reinforcement Learning <i>Haoyuan Qin, Zhengzhu Liu, Chenxing Lin, Chennan Ma, Songzhu Mei, Siqi Shen, Cheng Wang</i>	37287
From Crowdsourced Data to High-quality Benchmarks: Arena-Hard and Benchbuilder Pipeline <i>Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E. Gonzalez, Ion Stoica</i>	37310
XAttention: Block Sparse Attention with Antidiagonal Scoring <i>Ruyi Xu, Guangxuan Xiao, Haofeng Huang, Junxian Guo, Song Han</i>	37333
Counting in Small Transformers: The Delicate Interplay between Attention and Feed-Forward Layers <i>Freya Behrens, Luca Biggio, Lenka Zdeborova</i>	37346
Alpha-SQL: Zero-Shot Text-to-SQL using Monte Carlo Tree Search <i>Boyan Li, Jiayi Zhang, Ju Fan, Yanwei Xu, Chong Chen, Nan Tang, Yuyu Luo</i>	37379

Learning dynamics in linear recurrent neural networks <i>Alexandra Maria Proca, Cl��mentine Carla Juliette Domin��, Murray Shanahan, Pedro A. M. Mediano</i>	37400
Learning Time-Aware Causal Representation for Model Generalization in Evolving Domains <i>Zhuo He, Shuang Li, Wenze Song, Longhui Yuan, Jian Liang, Han Li, Kun Gai</i>	37443
A Closer Look at Transformers for Time Series Forecasting: Understanding Why They Work and Where They Struggle <i>Yu Chen, Nathalia C��spedes, Payam Barnaghi</i>	37468
Online Detection of LLM-Generated Texts via Sequential Hypothesis Testing by Betting <i>Can Chen, Jun-Kun Wang</i>	37486
Temporal Misalignment in ANN-SNN Conversion and its Mitigation via Probabilistic Spiking Neurons <i>Velibor Bojkovic, Xiaofeng Wu, Bin Gu</i>	37532
Domain2Vec: Vectorizing Datasets to Find the Optimal Data Mixture without Training <i>Mozhi Zhang, Howe Tissue, Lu Wang, Xipeng Qiu</i>	37564
Improving LLM Safety Alignment with Dual-Objective Optimization <i>Xuandong Zhao, Will Cai, Tianneng Shi, David Huang, Licong Lin, Song Mei, Dawn Song</i>	37584
A Generalization Theory for Zero-Shot Prediction <i>Ronak Mehta, Zaid Harchaoui</i>	37605
Exponential Family Variational Flow Matching for Tabular Data Generation <i>Andr��s Guzm��n-Cordero, Floor Eijkelboom, Jan-Willem van de Meent</i>	37663
SPMC: Self-Purifying Federated Backdoor Defense via Margin Contribution <i>Wenwen He, Wenke Huang, Bin Yang, ShuKan Liu, Mang Ye</i>	37677
LIMEFLDL: A Local Interpretable Model-Agnostic Explanations Approach for Label Distribution Learning <i>Xiuyi Jia, Jinchi Li, Yunan Lu, Weiwei Li</i>	37689
Generalized Venn and Venn-Abers Calibration with Applications in Conformal Prediction <i>Lars van der Laan, Ahmed Alaa</i>	37713
B-score: Detecting biases in large language models using response history <i>An Vo, Mohammad Reza Taesiri, Daeyoung Kim, Anh Totti Nguyen</i>	37729
SPRI: Aligning Large Language Models with Context-Situated Principles <i>Hongli Zhan, Muneeza Azmat, Raya Horesh, Junyi Jessy Li, Mikhail Yurochkin</i>	37766
Stochastic Deep Restoration Priors for Imaging Inverse Problems <i>Yuyang Hu, Albert Peng, Weijie Gan, Peyman Milanfar, Mauricio Delbracio, Ulugbek S. Kamilov</i>	37802

GenMol: A Drug Discovery Generalist with Discrete Diffusion <i>Seul Lee, Karsten Kreis, Srimukh Prasad Veccham, Meng Liu, Danny Reidenbach, Yuxing Peng, Saeed Gopal Paliwal, Weili Nie, Arash Vahdat</i>	37834
COExpander: Adaptive Solution Expansion for Combinatorial Optimization <i>Jiale Ma, Wenzheng Pan, Yang Li, Junchi Yan</i>	37856
Towards Understanding Parametric Generalized Category Discovery on Graphs <i>Bowen Deng, Lele Fu, Jialong Chen, Sheng Huang, Tianchi Liao, Zhang Tao, Chuan Chen</i>	37891
Explainable Concept Generation through Vision-Language Preference Learning for Understanding Neural Networks' Internal Representations <i>Aditya Taparia, Som Sagar, Ransalu Senanayake</i>	37932
Falsification of Unconfoundedness by Testing Independence of Causal Mechanisms <i>Rickard Karlsson, JH Krijthe</i>	37960
Subgoal-Guided Policy Heuristic Search with Learned Subgoals <i>Jake Tuero, Michael Buro, Levi Lelis</i>	37980
Generalization in Federated Learning: A Conditional Mutual Information Framework <i>Ziqiao Wang, Cheng Long, Yongyi Mao</i>	37999
Controllable Data Generation with Hierarchical Neural Representations <i>Sheyang Tang, Xiaoyu Xu, Jiayan Qiu, Zhou Wang</i>	38031
What Limits Bidirectional Model's Generative Capabilities? A Uni-Bi-Directional Mixture-of-Expert Method For Bidirectional Fine-tuning <i>Zuchao Li, Yonghua Hei, Qiwei Li, Lefei Zhang, Ping Wang, hai zhao, Baoyuan Qi, Liu Guoming</i>	38051
DeFoG: Discrete Flow Matching for Graph Generation <i>Yiming QIN, Manuel Madeira, Dorina Thanou, Pascal Frossard</i>	38065
Simple and Critical Iterative Denoising: A Recasting of Discrete Diffusion in Graph Generation <i>Yoann Boget</i>	38123
Optimal Transfer Learning for Missing Not-at-Random Matrix Completion <i>Akhil Jalan, Yassir Jedra, Arya Mazumdar, Soumendu Sundar Mukherjee, Purnamrita Sarkar</i>	38147
Zero Shot Generalization of Vision-Based RL Without Data Augmentation <i>Sumeet Batra, Gaurav S. Sukhatme</i>	38180
Putnam-AXIOM: A Functional & Static Benchmark for Measuring Higher Level Mathematical Reasoning in LLMs <i>Aryan Gulati, Brando Miranda, Eric Chen, Emily Xia, Kai Fronsdal, Bruno de Moraes Dumont, Sanmi Koyejo</i>	38203
What makes an Ensemble (Un) Interpretable? <i>Shahaf Bassan, Guy Amir, Meirav Zehavi, Guy Katz</i>	38228

AnalogGenie-Lite: Enhancing Scalability and Precision in Circuit Topology Discovery through Lightweight Graph Modeling <i>Jian Gao, Weidong Cao, Xuan Zhang</i>	38288
Regression for the Mean: Auto-Evaluation and Inference with Few Labels through Post-hoc Regression <i>Benjamin Eyre, David Madras</i>	38301
FastCAV: Efficient Computation of Concept Activation Vectors for Explaining Deep Neural Networks <i>Laines Schmalwasser, Niklas Penzel, Joachim Denzler, Julia Niebling</i>	38318
Rethinking Chain-of-Thought from the Perspective of Self-Training <i>Zongqian Wu, Baoduo Xu, Ruochen Cui, Mengmeng Zhan, Xiaofeng Zhu, Lei Feng</i>	38345
The Elicitation Game: Evaluating Capability Elicitation Techniques <i>Felix Hofstätter, Teun van der Weij, Jayden Teoh, Rada Djoneva, Henning Bartsch, Francis Rhys Ward</i>	38366
SMART-PC: Skeletal Model Adaptation for Robust Test-Time Training in Point Clouds <i>Ali Bahri, Moslem Yazdanpanah, Sahar Dastani, Mehrdad Noori, Gustavo Adolfo Vargas Hakim, David OSOWIECHI, Farzad Bezaee, Ismail Ben Ayed, Christian Desrosiers</i>	38393
Self-Supervised Learning of Intertwined Content and Positional Features for Object Detection <i>Kang-Jun Liu, Masanori Suganuma, Takayuki Okatani</i>	38409
An Expressive and Self-Adaptive Dynamical System for Efficient Function Learning <i>Chuan Liu, Chunshu Wu, Ruibing Song, Ang Li, Ying Nian Wu, Tong Geng</i>	38425
Stray Intrusive Outliers-Based Feature Selection on Intra-Class Asymmetric Instance Distribution or Multiple High-Density Clusters <i>Lixin Yuan, Yirui Wu, Wenxiao Zhang, Minglei Yuan, Jun Liu</i>	38437
A Hitchhiker's Guide to Scaling Law Estimation <i>Leshem Choshen, Yang Zhang, Jacob Andreas</i>	38455
LV-XAttn: Distributed Cross-Attention for Long Visual Inputs in Multimodal Large Language Models <i>Tzu-Tao Chang, Shivaram Venkataraman</i>	38472
UniMoMo: Unified Generative Modeling of 3D Molecules for De Novo Binder Design <i>Xiangzhe Kong, Zishen Zhang, Ziting Zhang, Rui Jiao, Jianzhu Ma, Wenbing Huang, Kai Liu, Yang Liu</i>	38485
Learning Along the Arrow of Time: Hyperbolic Geometry for Backward-Compatible Representation Learning <i>Ngoc Bui, Menglin Yang, Runjin Chen, Leonardo Neves, Mingxuan Ju, Rex Ying, Neil Shah, Tong Zhao</i>	38507
Nonlinearly Preconditioned Gradient Methods under Generalized Smoothness <i>Konstantinos Oikonomidis, Jan Quan, Emanuel Laude, Panagiotis Patrinos</i>	38521

Flow Q-Learning <i>Seohong Park, Qiyang Li, Sergey Levine</i>	38544
ADIOS: Antibody Development via Opponent Shaping <i>Sebastian Rene Towers, Aleksandra Kalisz, Philippe A. Robert, Alicia Higuieruelo, Francesca Vianello, Chloe Ming-Han Tsai, Harrison Steel, Jakob Nicolaus Foerster</i>	38568
How Do Transformers Learn Variable Binding in Symbolic Programs? <i>Yiwei Wu, Atticus Geiger, Raphaël Millière</i>	38588
Online Conformal Prediction via Online Optimization <i>Felipe Areces, Christopher Mohri, Tatsunori Hashimoto, John Duchi</i>	38604
On the Role of Label Noise in the Feature Learning Process <i>Andi Han, Wei Huang, Zhanpeng Zhou, Gang Niu, Wuyang Chen, Junchi Yan, Akiko Takeda, Taiji Suzuki</i>	38650
LoRA-One: One-Step Full Gradient Could Suffice for Fine-Tuning Large Language Models, Provably and Efficiently <i>Yuanhe Zhang, Fanghui Liu, Yudong Chen</i>	38691
A Recipe for Causal Graph Regression: Confounding Effects Revisited <i>Yujia Yin, Tianyi Qu, Zihao Wang, Yifan Chen</i>	38753
Generalization Analysis for Supervised Contrastive Representation Learning under Non-IID Settings <i>Nong Minh Hieu, Antoine Ledent</i>	38767
Poly2Vec: Polymorphic Fourier-Based Encoding of Geospatial Objects for GeoAI Applications <i>Maria Despoina Siampou, Jialiang Li, John Krumm, Cyrus Shahabi, Hua Lu</i>	38807
Hybrid Quantum-Classical Multi-Agent Pathfinding <i>Thore Gerlach, Loong Kuan Lee, Frederic BARBARESCO, Nico Piatkowski</i>	38829
Aligning Spoken Dialogue Models from User Interactions <i>Anne Wu, Laurent Mazaré, Neil Zeghidour, Alexandre Défossez</i>	38840
FuseUNet: A Multi-Scale Feature Fusion Method for U-like Networks <i>Quansong He, Xiangde Min, Kaishen Wang, Tao He</i>	38863
A Generic Family of Graphical Models: Diversity, Efficiency, and Heterogeneity <i>Yufei Huang, Changhu Wang, Junjie Tang, Weichi Wu, Ruibin Xi</i>	38880
Measuring Representational Shifts in Continual Learning: A Linear Transformation Perspective <i>Joonkyu Kim, Yejin Kim, Jy-yong Sohn</i>	38912
Towards Memorization Estimation: Fast, Formal and Free <i>Deepak Ravikumar, Efstathia Soufleri, Abolfazl Hashemi, Kaushik Roy</i>	38935

A Reductions Approach to Risk-Sensitive Reinforcement Learning with Optimized Certainty Equivalents <i>Kaiwen Wang, Dawen Liang, Nathan Kallus, Wen Sun</i>	38962
Explicit Discovery of Nonlinear Symmetries from Dynamic Data <i>Lexiang Hu, Yikang Li, Zhouchen Lin</i>	38988
Correlated Errors in Large Language Models <i>Elliot Myunghoon Kim, Avi Garg, Kenny Peng, Nikhil Garg</i>	39014
VersaPRM: Multi-Domain Process Reward Model via Synthetic Reasoning Data <i>Thomas Zeng, Shuibai Zhang, Shutong Wu, Christian Classen, Daewon Chae, Ethan Ewer, Minjae Lee, Heeju Kim, Wonjun Kang, Jackson Kunde, Ying Fan, Jungtaek Kim, Hyung Il Koo, Kannan Ramchandran, Dimitris Papailiopoulos, Kangwook Lee</i>	39043
The Butterfly Effect: Neural Network Training Trajectories Are Highly Sensitive to Initial Conditions <i>Gül Sena Altıntaş, Devin Kwok, Colin Raffel, David Rolnick</i>	39086
Curvature Enhanced Data Augmentation for Regression <i>Ilya Kaufman, Omri Azencot</i>	39115
Decoding Rewards in Competitive Games: Inverse Game Theory with Entropy Regularization <i>Junyi Liao, Zihan Zhu, Ethan X Fang, Zhuoran Yang, Vahid Tarokh</i>	39139
Scalable Non-Equivariant 3D Molecule Generation via Rotational Alignment <i>Yuhui Ding, Thomas Hofmann</i>	39152
Flexible and Efficient Grammar-Constrained Decoding <i>Kanghee Park, Timothy Zhou, Loris D'Antoni</i>	39163
LDMol: A Text-to-Molecule Diffusion Model with Structurally Informative Latent Space Surpasses AR Models <i>Jinho Chang, Jong Chul Ye</i>	39177
Provably Improving Generalization of Few-shot models with Synthetic Data <i>Lan-Cuong Nguyen, Quan Nguyen-Tri, Bang Tran Khanh, Dung D. Le, Long Tran-Thanh, Khoat Than</i>	39196
Does Low Rank Adaptation Lead to Lower Robustness against Training-Time Attacks? <i>Zi Liang, Haibo Hu, Qingqing Ye, Yaxin Xiao, RongHua Li</i>	39214
Limitations of measure-first protocols in quantum machine learning <i>Casper Gyurik, Riccardo Molteni, Vedran Dunjko</i>	39241
EvoPress: Accurate Dynamic Model Compression via Evolutionary Search <i>Oliver Sieberling, Denis Kuznedelev, Eldar Kurtic, Dan Alistarh</i>	39257
Training High Performance Spiking Neural Network by Temporal Model Calibration <i>Jiaqi Yan, Changping Wang, De Ma, Huajin Tang, Qian Zheng, Gang Pan</i>	39292

Should Decision-Makers Reveal Classifiers in Online Strategic Classification? <i>Han Shao, Shuo Xie, Kunhe Yang</i>	39312
Policy Filtration for RLHF to Mitigate Noise in Reward Models <i>Chuheng Zhang, Wei Shen, Li Zhao, Xuyun Zhang, Xiaolong Xu, Wanchun Dou, Jiang Bian</i>	39333
On the Importance of Embedding Norms in Self-Supervised Learning <i>Andrew Draganov, Sharvaree Vadgama, Sebastian Damrich, Jan Niklas Böhm, Lucas Maes, Dmitry Kobak, Erik J Bekkers</i>	39357
Self-Discriminative Modeling for Anomalous Graph Detection <i>Jinyu Cai, Yunhe Zhang, Jicong Fan</i>	39379
Parameters vs FLOPs: Scaling Laws for Optimal Sparsity for Mixture-of-Experts Language Models <i>Samira Abnar, Harshay Shah, Dan Busbridge, Alaaeldin El-Nouby, Joshua M. Susskind, Vimal Thilak</i>	39401
On the Diversity of Adversarial Ensemble Learning <i>Jun-Qi Guo, Meng-Zhang Qian, Wei Gao, Zhi-Hua Zhou</i>	39428
Rethinking External Slow-Thinking: From Snowball Errors to Probability of Correct Reasoning <i>Zeyu Gan, Yun Liao, Yong Liu</i>	39451
Directed Graph Grammars for Sequence-based Learning <i>Michael Sun, Orion Foo, Gang Liu, Wojciech Matusik, Jie Chen</i>	39470
LangDAug: Langevin Data Augmentation for Multi-Source Domain Generalization in Medical Image Segmentation <i>Piyush Tiwary, Kinjawl Bhattacharyya, Prathosh AP</i>	39498
A Novel Characterization of the Population Area Under the Risk Coverage Curve (AURC) and Rates of Finite Sample Estimators <i>Han Zhou, Jordy Van Landeghem, Teodora Popordanoska, Matthew B. Blaschko</i>	39534
FreeMesh: Boosting Mesh Generation with Coordinates Merging <i>Jian Liu, Haohan Weng, Biwen Lei, Xianghui Yang, Zibo Zhao, Zhuo Chen, Song Guo, Tao Han, Chunchao Guo</i>	39562
A Unified Framework for Entropy Search and Expected Improvement in Bayesian Optimization <i>Nuojin Cheng, Leonard Papenmeier, Stephen Becker, Luigi Nardi</i>	39576
Disentangling and Integrating Relational and Sensory Information in Transformer Architectures <i>Awni Altabaa, John Lafferty</i>	39591
Test-time Adaptation on Graphs via Adaptive Subgraph-based Selection and Regularized Prototypes <i>Yusheng Zhao, Qixin Zhang, Xiao Luo, Junyu Luo, Wei Ju, Zhiping Xiao, Ming Zhang</i>	39618

From Passive to Active Reasoning: Can Large Language Models Ask the Right Questions under Incomplete Information? <i>Zhanke Zhou, Xiao Feng, Zhaocheng Zhu, Jiangchao Yao, Sanmi Koyejo, Bo Han</i>	39638
Expected Variational Inequalities <i>Brian Hu Zhang, Ioannis Anagnostides, Emanuel Tewolde, Ratip Emin Berker, Gabriele Farina, Vincent Conitzer, Tuomas Sandholm</i>	39683
A Near Linear Query Lower Bound for Submodular Maximization <i>Binghui Peng, Aviad Rubinstein</i>	39708
Focal-SAM: Focal Sharpness-Aware Minimization for Long-Tailed Classification <i>Sicong Li, Qianqian Xu, Zhiyong Yang, Zitai Wang, Linchao Zhang, Xiaochun Cao, Qingming Huang</i>	39728
Preference Adaptive and Sequential Text-to-Image Generation <i>Ofir Nabati, Guy Tennenholtz, ChihWei Hsu, Moonkyung Ryu, Deepak Ramachandran, Yinlam Chow, Xiang Li, Craig Boutilier</i>	39756
Accelerating Quantum Reinforcement Learning with a Quantum Natural Policy Gradient Based Approach <i>Yang Xu, Vaneet Aggarwal</i>	39789
Low-Rank Adapting Models for Sparse Autoencoders <i>Matthew Chen, Joshua Engels, Max Tegmark</i>	39808
Making Hard Problems Easier with Custom Data Distributions and Loss Regularization: A Case Study in Modular Arithmetic <i>Eshika Saxena, Alberto Alfarano, Emily Wenger, Kristin E. Lauter</i>	39824
Private Lossless Multiple Release <i>Joel Daniel Andersson, Lukas Retschmeier, Boel Nelson, Rasmus Pagh</i>	39839
Inverse Optimization via Learning Feasible Regions <i>Ke Ren, Peyman Mohajerin Esfahani, Angelos Georghiou</i>	39859
Accelerating PDE-Constrained Optimization by the Derivative of Neural Operators <i>Ze Cheng, Zhuoyu Li, Wang Xiaoqiang, Jianing Huang, Zhizhou Zhang, Zhongkai Hao, Hang Su</i>	39877
WILTING Trees: Interpreting the Distance Between MPNN Embeddings <i>Masahiro Negishi, Thomas Gärtner, Pascal Welke</i>	39893
Deep Bayesian Filter for Bayes-Faithful Data Assimilation <i>Yuta Tarumi, Keisuke Fukuda, Shin-ichi Maeda</i>	39918
Ranked Entropy Minimization for Continual Test-Time Adaptation <i>Jisu Han, Jaemin Na, Wonjun Hwang</i>	39946

TRACE Back from the Future: A Probabilistic Reasoning Approach to Controllable Language Generation <i>Gwen Yidou Weng, Benjie Wang, Guy Van den Broeck</i>	39962
Best of Both Worlds: Regret Minimization versus Minimax Play <i>Adrian Müller, Jon Schneider, Stratis Skoulakis, Luca Viano, Volkan Cevher</i>	39978
Self-Bootstrapping for Versatile Test-Time Adaptation <i>Shuaicheng Niu, Guohao Chen, Peilin Zhao, Tianyi Wang, Pengcheng Wu, Zhiqi Shen</i>	40005
LBI-FL: Low-Bit Integerized Federated Learning with Temporally Dynamic Bit-Width Allocation <i>Li Ding, Hao Zhang, Wenrui Dai, Chenglin Li, Weijia Lu, ZHIFEI YANG, xiaodong Zhang, Xiaofeng Ma, Junni Zou, Hongkai Xiong</i>	40023
LGDM: Latent Guidance in Diffusion Models for Perceptual Evaluations <i>Shreshth Saini, Ru-Ling Liao, Yan Ye, Alan Bovik</i>	40038
FeatSharp: Your Vision Model Features, Sharper <i>Mike Ranzinger, Greg Heinrich, Pavlo Molchanov, Bryan Catanzaro, Andrew Tao</i>	40062
Discovering Global False Negatives On the Fly for Self-supervised Contrastive Learning <i>Vicente Balmaseda, Bokun Wang, Ching-Long Lin, Tianbao Yang</i>	40089
Incremental Gradient Descent with Small Epoch Counts is Surprisingly Slow on III-Conditioned Problems <i>Yujun Kim, Jaeyoung Cha, Chulhee Yun</i>	40107
Learning to Stop: Deep Learning for Mean Field Optimal Stopping <i>Lorenzo Magnino, Yuchen Zhu, Mathieu Lauriere</i>	40163
Self-Disentanglement and Re-Composition for Cross-Domain Few-Shot Segmentation <i>Jintao Tong, Yixiong Zou, Guangyao Chen, Yuhua Li, Ruixuan Li</i>	40197
Armijo Line-search Can Make (Stochastic) Gradient Descent Provably Faster <i>Sharan Vaswani, Reza Babanezhad Harikandeh</i>	40213
Power Mean Estimation in Stochastic Continuous Monte-Carlo Tree Search <i>Tuan Quang Dam</i>	40249
An Adaptive Orthogonal Convolution Scheme for Efficient and Flexible CNN Architectures <i>Thibaut Boissin, Franck Mamalet, Thomas Fel, Agustin Martin Picard, Thomas Massena, Mathieu Serrurier</i>	40282
LETS Forecast: Learning Embedology for Time Series Forecasting <i>Abrar Majeedi, Viswanatha Reddy Gajjala, Satya Sai Srinath Namburi GNVV, Nada Magdi Elkordi, Yin Li</i>	40316
Joint Localization and Activation Editing for Low-Resource Fine-Tuning <i>Wen Lai, Alexander Fraser, Ivan Titov</i>	40339

Differential Privacy Guarantees of Markov Chain Monte Carlo Algorithms <i>Andrea Bertazzi, Tim Johnston, Gareth O. Roberts, Alain Oliviero Durmus</i>	40361
Hyperband-based Bayesian Optimization for Black-box Prompt Selection <i>Lennart Schneider, Martin Wistuba, Aaron Klein, Jacek Golebiowski, Giovanni Zappella, Felice Antonio Merra</i>	40382
AlphaPO: Reward Shape Matters for LLM Alignment <i>Aman Gupta, Shao Tang, Qingquan Song, Sirou Zhu, Jiwoo Hong, Ankan Saha, Viral Gupta, Noah Lee, Eunki Kim, Siyu Zhu, Parag Agrawal, Natesh S. Pillai, Sathiya Keerthi</i>	40408
HALoS: Hierarchical Asynchronous Local SGD over Slow Networks for Geo-Distributed Large Language Model Training <i>Geon-Woo Kim, Junbo Li, Shashidhar Gandham, Omar Baldonado, Adithya Gangidi, Pavan Balaji, Zhangyang Wang, Aditya Akella</i>	40434
Efficient Multivariate Robust Mean Estimation Under Mean-Shift Contamination <i>Ilias Diakonikolas, Giannis Iakovidis, Daniel Kane, Thanasis Pittas</i>	40453
Offline Opponent Modeling with Truncated Q-driven Instant Policy Refinement <i>Yuheng Jing, Kai Li, Bingyun Liu, Ziwen Zhang, Haobo Fu, QIANG FU, Junliang Xing, Jian Cheng</i>	40484
Learning to Generate Projections for Reducing Dimensionality of Heterogeneous Linear Programming Problems <i>Tomoharu Iwata, Shinsaku Sakaue</i>	40516
Clipping Improves Adam-Norm and AdaGrad-Norm when the Noise Is Heavy-Tailed <i>Savelii Chezhegov, Klyukin Yaroslav, Andrei Semenov, Aleksandr Beznosikov, Alexander Gasnikov, Samuel Horváth, Martin Takáč, Eduard Gorbunov</i>	40531
Hgformer: Hyperbolic Graph Transformer for Collaborative Filtering <i>Xin Yang, Xingrun Li, Heng Chang, Yang jinze, Xihong Yang, Shengyu Tao, Maiko Shigeno, Ningkan Chang, Junfeng Wang, Dawei Yin, Erxue Min</i>	40596
Hi Robot: Open-Ended Instruction Following with Hierarchical Vision-Language-Action Models <i>Lucy Xiaoyang Shi, brian ichter, Michael Robert Equi, Liyiming Ke, Karl Pertsch, Quan Vuong, James Tanner, Anna Walling, Haohuan Wang, Niccolo Fusai, Adrian Li-Bell, Danny Driess, Lachy Groom, Sergey Levine, Chelsea Finn</i>	40616
SDMG: Smoothing Your Diffusion Models for Powerful Graph Representation Learning <i>Junyou Zhu, Langzhou He, Chao Gao, Dongpeng Hou, Zhen Su, Philip S. Yu, Juergen Kurths, Frank Hellmann</i>	40631
Sundial: A Family of Highly Capable Time Series Foundation Models <i>Yong Liu, Guo Qin, Zhiyuan Shi, Zhi Chen, Caiyin Yang, Xiangdong Huang, Jianmin Wang, Mingsheng Long</i>	40652
Branches: Efficiently Seeking Optimal Sparse Decision Trees via AO* <i>Ayman Chaouki, Jesse Read, Albert Bifet</i>	40675

Context-Informed Neural ODEs Unexpectedly Identify Broken Symmetries: Insights from the Poincaré–Hopf Theorem <i>In Huh, changwook jeong, Muhammad Alam</i>	40730
Enhancing Treatment Effect Estimation via Active Learning: A Counterfactual Covering Perspective <i>Hechuan Wen, Tong Chen, Mingming Gong, Li Kheng Chai, Shazia Sadiq, Hongzhi Yin</i>	40757
WOMD-Reasoning: A Large-Scale Dataset for Interaction Reasoning in Driving <i>Yiheng Li, Cunxin Fan, Chongjian GE, Seth Z. Zhao, Chenran Li, Chenfeng Xu, Huaxiu Yao, Masayoshi Tomizuka, Bolei Zhou, Chen Tang, Mingyu Ding, Wei Zhan</i>	40787
Rhomboid Tiling for Geometric Graph Deep Learning <i>Yipeng Zhang, Longlong Li, Kelin Xia</i>	40811
Oscillation-Reduced MXFP4 Training for Vision Transformers <i>Yuxiang Chen, Haocheng Xi, Jun Zhu, Jianfei Chen</i>	40832
xLSTM 7B: A Recurrent LLM for Fast and Efficient Inference <i>Maximilian Beck, Korbinian Pöppel, Phillip Lippe, Richard Kurl, Patrick M Blies, Günter Klambauer, Sebastian Böck, Sepp Hochreiter</i>	40847
UltraTWD: Optimizing Ultrametric Trees for Tree-Wasserstein Distance <i>Fangchen Yu, Yanzhen Chen, Jiaying Wei, Jianfeng Mao, Wenye Li, Qiang Sun</i>	40870
CAN: Leveraging Clients As Navigators for Generative Replay in Federated Continual Learning <i>Xuankun Rong, Jianshu Zhang, Kun He, Mang Ye</i>	40886
Generalized Smooth Bilevel Optimization with Nonconvex Lower-Level <i>Siqi Zhang, Xing Huang, Feihu Huang</i>	40902
On the Clean Generalization and Robust Overfitting in Adversarial Training from Two Theoretical Views: Representation Complexity and Training Dynamics <i>Binghui Li, Yuanzhi Li</i>	40935
Faster Rates for Private Adversarial Bandits <i>Hilal Asi, Vinod Raman, Kunal Talwar</i>	40963
Code-Generated Graph Representations Using Multiple LLM Agents for Material Properties Prediction <i>Jiao Huang, Qianli Xing, Jinglong Ji, Bo Yang</i>	41000
Learning curves theory for hierarchically compositional data with power-law distributed features <i>Francesco Cagnetta, Hyunmo Kang, Matthieu Wyart</i>	41015
Learning to Match Unpaired Data with Minimum Entropy Coupling <i>Mustapha BOUNOUA, Giulio Franzese, Pietro Michiardi</i>	41031

From RAG to Memory: Non-Parametric Continual Learning for Large Language Models <i>Bernal Jiménez Gutiérrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, Yu Su</i>	41054
Avoiding Catastrophe in Online Learning by Asking for Help <i>Benjamin Plaut, Hanlin Zhu, Stuart Russell</i>	41073
KAN-AD: Time Series Anomaly Detection with Kolmogorov–Arnold Networks <i>Quan Zhou, Changhua Pei, Fei Sun, HanJing, Zhengwei Gao, Haiming Zhang, Gaogang Xie, Dan Pei, Jianhui Li</i>	41105
Fundamental Bias in Inverting Random Sampling Matrices with Application to Sub-sampled Newton <i>Chengmei Niu, Zhenyu Liao, Zenan Ling, Michael W. Mahoney</i>	41119
Tensor Decomposition Based Memory-Efficient Incremental Learning <i>Yuhang Li, Guoxu Zhou, Zhenhao Huang, Xinqi Chen, Yuning Qiu, Qibin Zhao</i>	41163
An Asymptotically Optimal Approximation Algorithm for Multiobjective Submodular Maximization at Scale <i>Fabian Christian Spaeh, Atsushi Miyauchi</i>	41178
The Underlying Universal Statistical Structure of Natural Datasets <i>Noam Itzhak Levi, Yaron Oz</i>	41205
Function Encoders: A Principled Approach to Transfer Learning in Hilbert Spaces <i>Tyler Ingebrand, Adam Thorpe, ufuk topcu</i>	41229
KV Shifting Attention Enhances Language Modeling <i>Mingyu Xu, Bingning Wang, Weipeng Chen</i>	41255
OrcaLoca: An LLM Agent Framework for Software Issue Localization <i>Zhongming Yu, Hejia Zhang, Yujie Zhao, Hanxian Huang, Matrix Yao, Ke Ding, Jishen Zhao</i>	41279
Grokking in the Wild: Data Augmentation for Real-World Multi-Hop Reasoning with Transformers <i>Roman Abramov, Felix Steinbauer, Gjergji Kasneci</i>	41300
Cooperation of Experts: Fusing Heterogeneous Information with Large Margin <i>Shuo Wang, Shunyang Huang, Jinghui Yuan, Zhixiang Shen, zhao kang</i>	41315
SING: Spatial Context in Large Language Model for Next-Gen Wearables <i>Ayushi Mishra, Yang Bai, Priyadarshan Narayanasamy, Nakul Garg, Nirupam Roy</i>	41332
Benefits of Early Stopping in Gradient Descent for Overparameterized Logistic Regression <i>Jingfeng Wu, Peter Bartlett, Matus Telgarsky, Bin Yu</i>	41351
A Bregman Proximal Viewpoint on Neural Operators <i>Abdel-Rahim Mezidi, Jordan Patracone, Saverio Salzo, Amaury Habrard, Massimiliano Pontil, Rémi Emonet, Marc Sebban</i>	41381

Local Pan-privacy for Federated Analytics <i>Vitaly Feldman, Audra McMillan, Guy N. Rothblum, Kunal Talwar</i>	41406
Online Robust Reinforcement Learning Through Monte-Carlo Planning <i>Tuan Quang Dam, Kishan Panaganti, Brahim Driss, Adam Wierman</i>	41422
Learning Multi-Level Features with Matryoshka Sparse Autoencoders <i>Bart Bussmann, Noa Nabeshima, Adam Karvonen, Neel Nanda</i>	41452
DiLQR: Differentiable Iterative Linear Quadratic Regulator via Implicit Differentiation <i>Shuyuan Wang, Philip D Loewen, Michael Forbes, Bhushan Gopaluni, Wei Pan</i>	41477
Does Data Scaling Lead to Visual Compositional Generalization? <i>Arnas Uselis, Andrea Dittadi, Seong Joon Oh</i>	41492
An Instrumental Value for Data Production and its Application to Data Pricing <i>Rui Ai, Boxiang Lyu, Zhaoran Wang, Zhuoran Yang, Haifeng Xu</i>	41517
TokenSwift: Lossless Acceleration of Ultra Long Sequence Generation <i>Tong Wu, Junzhe Shen, Zixia Jia, Yuxuan Wang, Zilong Zheng</i>	41552
Navigating Semantic Drift in Task-Agnostic Class-Incremental Learning <i>Fangwen Wu, Lechao Cheng, Shengeng Tang, Xiaofeng Zhu, Chaowei Fang, Dingwen Zhang, Meng Wang</i>	41572
FicGCN: Unveiling the Homomorphic Encryption Efficiency from Irregular Graph Convolutional Networks <i>Zhaoxuan Kan, Husheng Han, shangyi shi, Tenghui Hua, Hang Lu, Xiaowei Li, Jianan Mu, Xing Hu</i>	41584
SAND: One-Shot Feature Selection with Additive Noise Distortion <i>Pedram Pad, Hadi Hammoud, Mohamad Dia, Nadim Maamari, Liza Andrea Dunbar</i>	41601
GaussMarker: Robust Dual-Domain Watermark for Diffusion Models <i>Kecen Li, Zhicong Huang, Xinwen Hou, Cheng Hong</i>	41614
LLMScan: Causal Scan for LLM Misbehavior Detection <i>Mengdi Zhang, Goh Kai Kiat, Peixin Zhang, Jun Sun, Lin Xin Rose, Hongyu Zhang</i>	41628
Fundamental Limits of Visual Autoregressive Transformers: Universal Approximation Abilities <i>Yifang Chen, Xiaoyu Li, Yingyu Liang, Zhenmei Shi, Zhao Song</i>	41658
Core Context Aware Transformers for Long Context Language Modeling <i>Yaofo Chen, Zeng You, Shuhai Zhang, Haokun Li, Yirui Li, Yaowei Wang, Mingkui Tan</i>	41685
Permutation Equivariant Neural Networks for Symmetric Tensors <i>Edward Pearce-Crump</i>	41705
Supercharging Graph Transformers with Advective Diffusion <i>Qitian Wu, Chenxiao Yang, Kaipeng Zeng, Michael M. Bronstein</i>	41745

Score Matching with Missing Data <i>Josh Givens, Song Liu, Henry Reeve</i>	41772
Improved Theoretically-Grounded Evolutionary Algorithms for Subset Selection with a Linear Cost Constraint <i>Dan-Xuan Liu, Chao Qian</i>	41811
Closed-Loop Long-Horizon Robotic Planning via Equilibrium Sequence Modeling <i>Jinghan Li, Zhicheng Sun, Yadong MU</i>	41831
Chameleon: A Flexible Data-mixing Framework for Language Model Pretraining and Finetuning <i>Wanyun Xie, Francesco Tonin, Volkan Cevher</i>	41856
The Lock-in Hypothesis: Stagnation by Algorithm <i>Tianyi Qiu, Zhonghao He, Tejasveer Chugh, Max Kleiman-Weiner</i>	41877
MMInference: Accelerating Pre-filling for Long-Context Visual Language Models via Modality-Aware Permutation Sparse Attention <i>Yucheng Li, Huiqiang Jiang, Chengruidong Zhang, Qianhui Wu, Xufang Luo, Surin Ahn, Amir H. Abdi, Dongsheng Li, Jianfeng Gao, Yuqing Yang, Lili Qiu</i>	41923
BOOD: Boundary-based Out-Of-Distribution Data Generation <i>Qilin Liao, Shuo Yang, Bo Zhao, Ping Luo, Hengshuang Zhao</i>	41946
Reliable Algorithm Selection for Machine Learning-Guided Design <i>Clara Fannjiang, Ji Won Park</i>	41960
Soup-of-Experts: Pretraining Specialist Models via Parameters Averaging <i>Pierre Ablin, Angelos Katharopoulos, Skyler Seto, David Grangier</i>	41984
Primitive Vision: Improving Diagram Understanding in MLLMs <i>Shan Zhang, Aotian Chen, Yanpeng Sun, Jindong Gu, Yi-Yu Zheng, Piotr Koniusz, Kai Zou, Anton van den Hengel, Yuan Xue</i>	42003
HYGMA: Hypergraph Coordination Networks with Dynamic Grouping for Multi-Agent Reinforcement Learning <i>Chiqiang Liu, Dazi Li</i>	42027
Projection Pursuit Density Ratio Estimation <i>Meilin Wang, Wei Huang, Mingming Gong, Zheng Zhang</i>	42049
Exploring Criteria of Loss Reweighting to Enhance LLM Unlearning <i>Puning Yang, Qizhou Wang, Zhuo Huang, Tongliang Liu, Chengqi Zhang, Bo Han</i>	42082
IT³: Idempotent Test-Time Training <i>Nikita Durasov, Assaf Shocher, Doruk Oner, Gal Chechik, Alexei A Efros, Pascal Fua</i>	42122

SepLLM: Accelerate Large Language Models by Compressing One Segment into One Separator <i>Guoxuan Chen, Han Shi, Jiawei Li, Yihang Gao, Xiaozhe Ren, Yimeng Chen, Xin Jiang, Zhenguo Li, Weiyang Liu, Chao Huang</i>	42139
Concurrent Reinforcement Learning with Aggregated States via Randomized Least Squares Value Iteration <i>Yan Chen, Qinxun Bai, Yiteng Zhang, Maria Dimakopoulou, Shi Dong, Qi Sun, Zhengyuan Zhou</i>	42160
Synthetic Text Generation for Training Large Language Models via Gradient Matching <i>Dang Nguyen, Zeman Li, Mohammadhossein Bateni, Vahab Mirrokni, Meisam Razaviyayn, Baharan Mirzsoleiman</i>	42193
Learning from Loss Landscape: Generalizable Mixed-Precision Quantization via Adaptive Sharpness-Aware Gradient Aligning <i>Lianbo Ma, Jianlun Ma, Yuee Zhou, Guoyang Xie, Qiang He, Zhichao Lu</i>	42211
Network Sparsity Unlocks the Scaling Potential of Deep Reinforcement Learning <i>Guozheng Ma, Lu Li, Zilin Wang, Li Shen, Pierre-Luc Bacon, Dacheng Tao</i>	42226
TTFSFormer: A TTFS-based Lossless Conversion of Spiking Transformer <i>Lusen Zhao, Zihan Huang, Jianhao Ding, Zhaofei Yu</i>	42247
Stochastic Encodings for Active Feature Acquisition <i>Alexander Luke Ian Norcliffe, Changhee Lee, Fergus Imrie, Mihaela van der Schaar, Pietro Lio</i>	42261
Beyond Topological Self-Explainable GNNs: A Formal Explainability Perspective <i>Steve Azzolin, SAGAR MALHOTRA, Andrea Passerini, Stefano Teso</i>	42292
Mirror, Mirror of the Flow: How Does Regularization Shape Implicit Bias? <i>Tom Jacobs, Chao Zhou, Rebekka Burkholz</i>	42328
Optimization for Neural Operators can Benefit from Width <i>Pedro Cisneros-Velarde, Bhavesh Shrimali, Arindam Banerjee</i>	42355
Boosting Masked ECG-Text Auto-Encoders as Discriminative Learners <i>Manh Pham Hung, Aaqib Saeed, Dong Ma</i>	42403
Understanding the Statistical Accuracy-Communication Trade-off in Personalized Federated Learning with Minimax Guarantees <i>Xin Yu, Zelin He, Ying Sun, Lingzhou Xue, Runze Li</i>	42418
DEALing with Image Reconstruction: Deep Attentive Least Squares <i>Mehrsa Pourya, Erich Kobler, Michael Unser, Sebastian Neumayer</i>	42460
Reasoning-as-Logic-Units: Scaling Test-Time Reasoning in Large Language Models Through Logic Unit Alignment <i>Cheryl Li, Tianyuan Xu, Steven Y. Guo</i>	42480

Boosting Adversarial Robustness with CLAT: Criticality Leveraged Adversarial Training <i>Bhavna Gopal, Huanrui Yang, Jingyang Zhang, Mark Horton, Yiran Chen</i>	42501
Impossible Videos <i>Zechen Bai, Hai Ci, Mike Zheng Shou</i>	42521
Competing Bandits in Matching Markets via Super Stability <i>Soumya Basu</i>	42547
MAPLE: Many-Shot Adaptive Pseudo-Labeling for In-Context Learning <i>Zihan Chen, Song Wang, Zhen Tan, Jundong Li, Cong Shen</i>	42572
Scaling Sparse Feature Circuits For Studying In-Context Learning <i>Dmitrii Kharlapenko, Stepan Shabalin, Arthur Conmy, Neel Nanda</i>	42590
Retrieval Augmented Zero-Shot Enzyme Generation for Specified Substrate <i>Jiahe Du, Kaixiong Zhou, Xinyu Hong, Zhaozhuo Xu, Jinbo Xu, Xiao Huang</i>	42623
Overestimation in LLM Evaluation: A Controlled Large-Scale Study on Data Contamination's Impact on Machine Translation <i>Muhammed Yusuf Kocyigit, Eleftheria Briakou, Daniel Deutsch, Jiaming Luo, Colin Cherry, Markus Freitag</i>	42639
Pointwise Information Measures as Confidence Estimators in Deep Neural Networks: A Comparative Study <i>Shelvya Wongso, Rohan Ghosh, Mehul Motani</i>	42667
PAK-UCB Contextual Bandit: An Online Learning Approach to Prompt-Aware Selection of Generative Models and LLMs <i>Xiaoyan Hu, Ho-fung Leung, Farzan Farnia</i>	42710
AEQA-NAT : Adaptive End-to-end Quantization Alignment Training Framework for Non-autoregressive Machine Translation <i>Xiangyu Qu, guojing liu, Liang Li</i>	42745
Fusing Reward and Dueling Feedback in Stochastic Bandits <i>Xuchuang Wang, Qirun Zeng, Jinhang Zuo, Xutong Liu, Mohammad Hajiesmaili, John C.S. Lui, Adam Wierman</i>	42762
Exact risk curves of signSGD in High-Dimensions: quantifying preconditioning and noise-compression effects <i>Ke Liang Xiao, Noah Marshall, Atish Agarwala, Elliot Paquette</i>	42784
CogMath: Assessing LLMs' Authentic Mathematical Ability from a Human Cognitive Perspective <i>Jiayu Liu, Zhenya Huang, Wei Dai, Cheng Cheng, Jinze Wu, Jing Sha, Song Li, Qi Liu, Shijin Wang, Enhong Chen</i>	42833

An Improved Clique-Picking Algorithm for Counting Markov Equivalent DAGs via Super Cliques Transfer <i>Lifu Liu, Shiyuan He, Jianhua Guo</i>	42849
BackSlash: Rate Constrained Optimized Training of Large Language Models <i>Jun Wu, Jiangtao Wen, Yuxing Han</i>	42865
Risk-Sensitive Theory of Mind: Coordinating with Agents of Unknown Bias using Cumulative Prospect Theory <i>Mason O. Smith, Wenlong Zhang</i>	42877
Wasserstein Flow Matching: Generative Modeling Over Families of Distributions <i>Doron Haviv, Aram-Alexandre Pooladian, Dana Pe'er, Brandon Amos</i>	42896
Erwin: A Tree-based Hierarchical Transformer for Large-scale Physical Systems <i>Maksim Zhdanov, Max Welling, Jan-Willem van de Meent</i>	42917
Do NOT Think That Much for 2+3=? On the Overthinking of Long Reasoning Models <i>Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, Dong Yu</i>	42935
RISE: Radius of Influence based Subgraph Extraction for 3D Molecular Graph Explanation <i>Jingxiang Qu, Wenhan Gao, Jiaxing Zhang, Xufeng Liu, Hua Wei, Haibin Ling, Yi Liu</i>	42948
Enhancing Logits Distillation with Plug&Play Kendall's τ Ranking Loss <i>Yuchen Guan, Runxi Cheng, Kang Liu, Chun Yuan</i>	42966
PIGDreamer: Privileged Information Guided World Models for Safe Partially Observable Reinforcement Learning <i>Dongchi Huang, Jiaqi WANG, Yang Li, Chunhe Xia, Tianle Zhang, Kaige Zhang</i>	42982
Proto Successor Measure: Representing the Behavior Space of an RL Agent <i>Siddhant Agarwal, Harshit Sikchi, Peter Stone, Amy Zhang</i>	42999
Learnings from Scaling Visual Tokenizers for Reconstruction and Generation <i>Philippe Hansen-Estruch, David Yan, Ching-Yao Chuang, Orr Zohar, Jialiang Wang, Tingbo Hou, Tao Xu, Sriram Vishwanath, Peter Vajda, Xinlei Chen</i>	43020
Tensorized Multi-View Multi-Label Classification via Laplace Tensor Rank <i>Qiyu Zhong, Yi Shan, Haobo Wang, Zhen Yang, Gengyu Lyu</i>	43041
Pareto-Optimal Fronts for Benchmarking Symbolic Regression Algorithms <i>Kei Sen Fong, Mehul Motani</i>	43056
Confidential Guardian: Cryptographically Prohibiting the Abuse of Model Abstention <i>Stephan Rabanser, Ali Shahin Shamsabadi, Olive Franzese, Xiao Wang, Adrian Weller, Nicolas Papernot</i>	43075
Inverse Problem Sampling in Latent Space Using Sequential Monte Carlo <i>Idan Achituve, Hai Victor Habi, Amir Rosenfeld, Arnon Netzer, Idit Diamant, Ethan Fetaya</i>	43098

Towards Cost-Effective Reward Guided Text Generation <i>Ahmad Rashid, Ruotian Wu, Rongqi Fan, Hongliang Li, Agustinus Kristiadi, Pascal Poupart</i>	43122
How Expressive are Knowledge Graph Foundation Models? <i>Xingyue Huang, Pablo Barcelo, Michael M. Bronstein, Ismail Ilkan Ceylan, Mikhail Galkin, Juan L Reutter, Miguel Romero Orth</i>	43140
Improving the Statistical Efficiency of Cross-Conformal Prediction <i>Matteo Gasparin, Aaditya Ramdas</i>	43178
VinePPO: Refining Credit Assignment in RL Training of LLMs <i>Amirhossein Kazemnejad, Milad Aghajohari, Eva Portelance, Alessandro Sordoni, Siva Reddy, Aaron Courville, Nicolas Le Roux</i>	43198
Implicit Riemannian Optimism with Applications to Min-Max Problems <i>Christophe Roux, David Martínez-Rubio, Sebastian Pokutta</i>	43232
Toward a Unified Theory of Gradient Descent under Generalized Smoothness <i>Alexander Tyurin</i>	43266
Ferret: Federated Full-Parameter Tuning at Scale for Large Language Models <i>Yao Shu, Wenyang Hu, See-Kiong Ng, Bryan Kian Hsiang Low, Fei Yu</i>	43288
Modalities Contribute Unequally: Enhancing Medical Multi-modal Learning through Adaptive Modality Token Re-balancing <i>Jie Peng, Jenna L. Ballard, Mohan Zhang, Sukwon Yun, Jiayi Xin, Qi Long, Yanyong Zhang, Tianlong Chen</i>	43312
ReinboT: Amplifying Robot Visual-Language Manipulation with Reinforcement Learning <i>Hongyin Zhang, Zifeng Zhuang, Han Zhao, Pengxiang Ding, Hongchao Lu, Donglin Wang</i>	43331
Diffusion Models are Secretly Exchangeable: Parallelizing DDPMs via Auto Speculation <i>Hengyuan Hu, Aniket Das, Dorsa Sadigh, Nima Anari</i>	43349
Trust-Region Twisted Policy Improvement <i>Joery A. de Vries, Jinke He, Yaniv Oren, Matthijs T. J. Spaan</i>	43369
Hyper: Hyperparameter Robust Efficient Exploration in Reinforcement Learning <i>Yiran Wang, Chenshu Liu, Yunfan Li, Sanae Amani, Bolei Zhou, Lin Yang</i>	43392
Rethinking Latent Redundancy in Behavior Cloning: An Information Bottleneck Approach for Robot Manipulation <i>Shuanghao Bai, Wanqi Zhou, Pengxiang Ding, Wei Zhao, Donglin Wang, Badong Chen</i>	43419
Learning Strategic Language Agents in the Werewolf Game with Iterative Latent Space Policy Optimization <i>Zelai Xu, Wanjun Gu, Chao Yu, Yi Wu, Yu Wang</i>	43440
EasyInv: Toward Fast and Better DDIM Inversion <i>Ziyue Zhang, Mingbao Lin, Shuicheng YAN, Rongrong Ji</i>	43463

GSM-∞: How Do your LLMs Behave over Infinitely Increasing Reasoning Complexity and Context Length? <i>Yang Zhou, Hongyi Liu, Zhuoming Chen, Yuandong Tian, Beidi Chen</i>	43473
Identifying Neural Dynamics Using Interventional State Space Models <i>Amin Nejatbakhsh, Yixin Wang</i>	43524
Efficiently Serving Large Multimodal Models Using EPD Disaggregation <i>Gursimran Singh, Xinglu Wang, Yifan Hu, Timothy Tin Long Yu, Linzi Xing, Wei Jiang, Zhefeng Wang, Bai Xiaolong, Yi Li, Ying Xiong, Yong Zhang, Zhenan Fan</i>	43542
Distributional Diffusion Models with Scoring Rules <i>Valentin De Bortoli, Alexandre Galashov, J Swaroop Guntupalli, Guangyao Zhou, Kevin Patrick Murphy, Arthur Gretton, Arnaud Doucet</i>	43559
Adaptive Data Collection for Robust Learning Across Multiple Distributions <i>Chengbo Zang, Mehmet Kerem Turkcan, Gil Zussman, Zoran Kostic, Javad Ghaderi</i>	43604
GHOST: Generalizable One-Shot Federated Graph Learning with Proxy-Based Topology Knowledge Retention <i>Jiaru Qian, Guancheng Wan, Wenke Huang, Guibin Zhang, Yuxin Wu, Bo Du, Mang Ye</i>	43625
Hierarchical Equivariant Policy via Frame Transfer <i>Haibo Zhao, Dian Wang, Yizhe Zhu, Xupeng Zhu, Owen Lewis Howell, Linfeng Zhao, Yaoyao Qian, Robin Walters, Robert Platt</i>	43644
Concept Reachability in Diffusion Models: Beyond Dataset Constraints <i>Marta Aparicio Rodriguez, Xenia Miscouridou, Anastasia Borovykh</i>	43664
Hi-Patch: Hierarchical Patch GNN for Irregular Multivariate Time Series <i>Yicheng Luo, Bowen Zhang, Zhen Liu, Qianli Ma</i>	43688
Identifying biological perturbation targets through causal differential networks <i>Menghua Wu, Umesh Padia, Sean H. Murphy, Regina Barzilay, Tommi Jaakkola</i>	43714
RestoreGrad: Signal Restoration Using Conditional Denoising Diffusion Models with Jointly Learned Prior <i>Ching-Hua Lee, Chouchang Yang, Jaejin Cho, Yashas Malur Saidutta, Rakshith Sharma Srinivasa, Yilin Shen, Hongxia Jin</i>	43739
Floating-Point Neural Networks Can Represent Almost All Floating-Point Functions <i>Geonho Hwang, Yeachan Park, Wonyeol Lee, Sejun Park</i>	43770
SageAttention2: Efficient Attention with Thorough Outlier Smoothing and Per-thread INT4 Quantization <i>Jintao Zhang, Haofeng Huang, Pengl Zhang, Jia wei, Jun Zhu, Jianfei Chen</i>	43791
The Best of Both Worlds: Bridging Quality and Diversity in Data Selection with Bipartite Graph <i>Minghao Wu, Thuy-Trang Vu, Lizhen Qu, Gholamreza Haffari</i>	43814

Larger or Smaller Reward Margins to Select Preferences for LLM Alignment? <i>Kexin Huang, Junkang Wu, Ziqian Chen, Xue Wang, Jinyang Gao, Bolin Ding, Jiancan Wu, Xiangnan He, Xiang Wang</i>	43829
DA-KD: Difficulty-Aware Knowledge Distillation for Efficient Large Language Models <i>Changyi He, Yifu Ding, Jinyang Guo, Ruihao Gong, Haotong Qin, Xianglong Liu</i>	43854
Reasoning Limitations of Multimodal Large Language Models. A case study of Bongard Problems <i>Mikołaj Mańkiński, Szymon Pawlonka, Jacek Mańdziuk</i>	43867
Training Diffusion-based Generative Models with Limited Data <i>Zhaoyu Zhang, Yang Hua, Guanxiong Sun, Hui Wang, Seán McLoone</i>	43906
Maximum Coverage in Turnstile Streams with Applications to Fingerprinting Measures <i>Alina Ene, Alessandro Epasto, Vahab Mirrokni, Hoai-An Nguyen, Huy Nguyen, David Woodruff, Peilin Zhong</i>	43922
Understanding Model Reprogramming for CLIP via Decoupling Visual Prompts <i>Chengyi Cai, Zesheng Ye, Lei Feng, Jianzhong Qi, Feng Liu</i>	43946
Stability and Generalization Capability of Subgraph Reasoning Models for Inductive Knowledge Graph Completion <i>Minsung Hwang, Jaejun Lee, Joyce Jiyoung Whang</i>	43971
Finite-Sample Convergence Bounds for Trust Region Policy Optimization in Mean Field Games <i>Antonio Ocello, Daniil Tiapkin, Lorenzo Mancini, Mathieu Lauriere, Eric Moulines</i>	44007
DS-VLM: Diffusion Supervision Vision Language Model <i>Zhen Sun, Yunhang Shen, Jie Li, Xing Sun, Pingyang Dai, Liujuan Cao, Rongrong Ji</i>	44053
Scalable Private Partition Selection via Adaptive Weighting <i>Justin Y. Chen, Vincent Cohen-Addad, Alessandro Epasto, Morteza Zadimoghaddam</i>	44066
GeoPixel: Pixel Grounding Large Multimodal Model in Remote Sensing <i>Akashah Shabbir, Mohammed Zumri, Mohammed Bennamoun, Fahad Shahbaz Khan, Salman Khan</i>	44091
Self-Organizing Visual Prototypes for Non-Parametric Representation Learning <i>Thalles Silva, Helio Pedrini, Adín Ramírez Rivera</i>	44108
TextCenGen: Attention-Guided Text-Centric Background Adaptation for Text-to-Image Generation <i>Tianyi Liang, Jiangqi Liu, Yifei Huang, Shiqi Jiang, Jianshen Shi, Changbo Wang, Chenhui Li</i>	44125
An Efficient Pruner for Large Language Model with Theoretical Guarantee <i>Canhong Wen, Yihong Zuo, Wenliang Pan</i>	44145
Triple-Optimistic Learning for Stochastic Contextual Bandits with General Constraints <i>Hengquan Guo, Ling kai Zu, Xin Liu</i>	44167

LLM-Augmented Chemical Synthesis and Design Decision Programs <i>Haorui Wang, Jeff Guo, Lingkai Kong, Rampi Ramprasad, Philippe Schwaller, Yuanqi Du, Chao Zhang</i>	44192
MissScore: High-Order Score Estimation in the Presence of Missing Data <i>Wenqin Liu, Haoze Hou, Erdun Gao, Biwei Huang, Qihong Ke, Howard Bondell, Mingming Gong</i>	44214
MGD³ : Mode-Guided Dataset Distillation using Diffusion Models <i>Jeffrey A Chan Santiago, praveen tirupattur, Gaurav Kumar Nayak, Gaowen Liu, Mubarak Shah</i>	44242
Mitigating Over-Squashing in Graph Neural Networks by Spectrum-Preserving Sparsification <i>Langzhang Liang, Fanchen Bu, Zixing Song, Zenglin Xu, Shirui Pan, Kijung Shin</i>	44258
One Wave To Explain Them All: A Unifying Perspective On Feature Attribution <i>Gabriel Kasmi, Amandine Brunetto, Thomas Fel, Jayneel Parekh</i>	44278
Flow Matching for Few-Trial Neural Adaptation with Stable Latent Dynamics <i>Puli Wang, Yu Qi, Yueming Wang, Gang Pan</i>	44307
SafeAuto: Knowledge-Enhanced Safe Autonomous Driving with Multimodal Foundation Models <i>Jiawei Zhang, Xuan Yang, Taiqi Wang, Yu Yao, Aleksandr Petiushko, Bo Li</i>	44332
Determinant Estimation under Memory Constraints and Neural Scaling Laws <i>Siavash Ameli, Chris van der Heide, Liam Hodgkinson, Fred Roosta, Michael W. Mahoney</i>	44353
Provable and Practical Online Learning Rate Adaptation with Hypergradient Descent <i>Ya-Chi Chu, Wenzhi Gao, Yinyu Ye, Madeleine Udell</i>	44387
Adaptive kernel predictors from feature-learning infinite limits of neural networks <i>Clarissa Lauditi, Blake Bordelon, Cengiz Pehlevan</i>	44420
You Always Recognize Me (YARM): Robust Texture Synthesis Against Multi-View Corruption <i>Weihang Ran, Wei Yuan, Yinqiang Zheng</i>	44452
Guided Search Strategies in Non-Serializable Environments with Applications to Software Engineering Agents <i>Karina Zainullina, Alexander Golubev, Maria Trofimova, Sergei Polezhaev, Ibragim Badertdinov, Daria Litvintseva, Simon Karasik, Filipp Fisin, Sergei Skvortsov, Maksim Nekrashevich, Anton Shevtsov, Boris Yangel</i>	44468
BRiTE: Bootstrapping Reinforced Thinking Process to Enhance Language Model Reasoning <i>Han Zhong, Yutong Yin, Shenao Zhang, Xiaojun Xu, Yuanxin Liu, Yifei Zuo, Zhihan Liu, Boyi Liu, Sirui Zheng, Hongyi Guo, Liwei Wang, Mingyi Hong, Zhaoran Wang</i>	44485
Agent-as-a-Judge: Evaluate Agents with Agents <i>Mingchen Zhuge, Changsheng Zhao, Dylan R. Ashley, Wenyi Wang, Dmitrii Khizbullin, Yunyang Xiong, Zechun Liu, Ernie Chang, Raghuraman Krishnamoorthi, Yuandong Tian, Yangyang Shi, Vikas Chandra, Jürgen Schmidhuber</i>	44505

Pixel-level Certified Explanations via Randomized Smoothing <i>Alaa Anani, Tobias Lorenz, Mario Fritz, Bernt Schiele</i>	44548
MDDM: Practical Message-Driven Generative Image Steganography Based on Diffusion Models <i>ZiHao Xu, Dawei xu, Zihan Li, Chuan Zhang</i>	44577
Latent Score-Based Reweighting for Robust Classification on Imbalanced Tabular Data <i>Yunze Tong, Fengda Zhang, Zihao Tang, Kaifeng Gao, Kai Huang, Pengfei Lyu, Jun Xiao, Kun Kuang</i>	44594
IL-SOAR : Imitation Learning with Soft Optimistic Actor cRitic <i>Stefano Viel, Luca Viano, Volkan Cevher</i>	44615
Policy Guided Tree Search for Enhanced LLM Reasoning <i>Yang Li</i>	44651
Learning State-Based Node Representations from a Class Hierarchy for Fine-Grained Open-Set Detection <i>Spandan Pyakurel, Qi Yu</i>	44672
Sparse Spectral Training and Inference on Euclidean and Hyperbolic Neural Networks <i>Jialin Zhao, Yingtao Zhang, Xinghang Li, Huaping Liu, Carlo Vittorio Cannistraci</i>	44695
Explaining, Fast and Slow: Abstraction and Refinement of Provable Explanations <i>Shahaf Bassan, Yizhak Yisrael Elboher, Tobias Ladner, Matthias Althoff, Guy Katz</i>	44722
Towards Universal Offline Black-Box Optimization via Learning Language Model Embeddings <i>Rong-Xi Tan, Ming Chen, Ke Xue, Yao Wang, Yaoyuan Wang, Fu Sheng, Chao Qian</i>	44746
Refining Adaptive Zeroth-Order Optimization at Ease <i>Yao Shu, Qixin Zhang, Kun He, Zhongxiang Dai</i>	44792
Improving Value Estimation Critically Enhances Vanilla Policy Gradient <i>Tao Wang, Ruipeng Zhang, Sicun Gao</i>	44816
Controlling Underestimation Bias in Constrained Reinforcement Learning for Safe Exploration <i>Shiqing Gao, Jiaxin Ding, Luoyi Fu, Xinbing Wang</i>	44831
Online Learning with Unknown Constraints <i>Karthik Sridharan, Seung Won Wilson Yoo</i>	44858
HarmoniCa: Harmonizing Training and Inference for Better Feature Caching in Diffusion Transformer Acceleration <i>Yushi Huang, Zining Wang, Ruihao Gong, Jing Liu, Xinjie Zhang, Jinyang Guo, Xianglong Liu, Jun Zhang</i>	44888
MARS: Unleashing the Power of Variance Reduction for Training Large Models <i>Huizhuo Yuan, Yifeng Liu, Shuang Wu, Zhou Xun, Quanquan Gu</i>	44912

Balancing Preservation and Modification: A Region and Semantic Aware Metric for Instruction-Based Image Editing <i>Zhuoying Li, Zhu Xu, Yuxin Peng, Yang Liu</i>	44947
A New Concentration Inequality for Sampling Without Replacement and Its Application for Transductive Learning <i>Yingzhen Yang</i>	44968
TINED: GNNs-to-MLPs by Teacher Injection and Dirichlet Energy Distillation <i>Ziang Zhou, Zhihao Ding, Jieming Shi, Li Qing, Shiqi Shen</i>	44994
EFDR: Learnable Elliptical Fourier Descriptor Transformer for Instance Segmentation <i>Jiawei Cao, Chaochen Gu, Hao Cheng, Xiaofeng Zhang, Kaijie Wu, Changsheng Lu</i>	45011
Underestimated Privacy Risks for Minority Populations in Large Language Model Unlearning <i>Rongzhe Wei, Mufei Li, Mohsen Ghassemi, Eleonora Kreacic, Yifan Li, Xiang Yue, Bo Li, Vamsi K. Potluru, Pan Li, Eli Chien</i>	45022
MERIT: Maximum-normalized Element-wise Ratio for Language Model Large-batch Training <i>Yang Luo, Zangwei Zheng, Ziheng Qin, Zirui Zhu, Yong Liu, Yang You</i>	45051
Agent Workflow Memory <i>Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, Graham Neubig</i>	45071
Bring Reason to Vision: Understanding Perception and Reasoning through Model Merging <i>Shiqi Chen, Jinghan Zhang, Tongyao Zhu, Wei Liu, Siyang Gao, Miao Xiong, Manling Li, Junxian He</i>	45086
DMM: Distributed Matrix Mechanism for Differentially-Private Federated Learning Based on Constant-Overhead Linear Secret Resharing <i>Alexander Bienstock, Ujjwal Kumar, Antigoni Polychroniadou</i>	45101
Scalable Gaussian Processes with Latent Kronecker Structure <i>Jihao Andreas Lin, Sebastian Ament, Maximilian Balandat, David Eriksson, José Miguel Hernández-Lobato, Eytan Bakshy</i>	45125
An Efficient Matrix Multiplication Algorithm for Accelerating Inference in Binary and Ternary Neural Networks <i>Mohsen Dehghankar, Mahdi Erfanian, Abolfazl Asudeh</i>	45140
AUTOCIRCUIT-RL: Reinforcement Learning-Driven LLM for Automated Circuit Topology Generation <i>Prashanth Vijayaraghavan, Luyao Shi, Ehsan Degan, Vandana Mukherjee, Xin Zhang</i>	45158
Scalable Meta-Learning via Mixed-Mode Differentiation <i>Iurii Kemaev, Dan A. Calian, Luisa M Zintgraf, Gregory Farquhar, Hado van Hasselt</i>	45173

Neural Graph Matching Improves Retrieval Augmented Generation in Molecular Machine Learning <i>Runzhong Wang, Rui-Xi Wang, Mrunali Manjrekar, Connor W. Coley</i>	45192
Leveraging Sparsity for Sample-Efficient Preference Learning: A Theoretical Perspective <i>Yunzhen Yao, Lie He, Michael Gastpar</i>	45208
CogReact: A Reinforced Framework to Model Human Cognitive Reaction Modulated by Dynamic Intervention <i>Songlin Xu, Xinyu Zhang</i>	45234
Online Laplacian-Based Representation Learning in Reinforcement Learning <i>Maheed H. Ahmed, Jayanth Bhargav, Mahsa Ghasemi</i>	45264
Bridging Fairness and Efficiency in Conformal Inference: A Surrogate-Assisted Group-Clustered Approach <i>Chenyin Gao, Peter B. Gilbert, Larry Han</i>	45286
TinyMIG: Transferring Generalization from Vision Foundation Models to Single-Domain Medical Imaging <i>Chuang LIU, Hongyan Xu, Yichao Cao, Xiu Su, Zhe Qu, Tianfa Li, Shan An, Haogang Zhu</i>	45306
When Model Knowledge meets Diffusion Model: Diffusion-assisted Data-free Image Synthesis with Alignment of Domain and Class <i>Yujin Kim, Hyunsoo Kim, Hyunwoo J. Kim, Suhyun Kim</i>	45323
Splitting with Importance-aware Updating for Heterogeneous Federated Learning with Large Language Models <i>Yangxu Liao, Wenke Huang, Guancheng Wan, Jian Liang, Bin Yang, Mang Ye</i>	45346
Uncertainty Quantification for LLM-Based Survey Simulations <i>Chengpiao Huang, Yuhang Wu, Kaizheng Wang</i>	45362
Sample Complexity of Correlation Detection in the Gaussian Wigner Model <i>Dong Huang, Pengkun Yang</i>	45387
Action-Constrained Imitation Learning <i>Chia-Han Yeh, Tse-Sheng Nan, Risto Vuorio, Wei Hung, Hung Yen Wu, Shao-Hua Sun, Ping-Chun Hsieh</i>	45408
Near Optimal Non-asymptotic Sample Complexity of 1-Identification <i>ZITIAN LI, Wang Chi Cheung</i>	45423
Diversity By Design: Leveraging Distribution Matching for Offline Model-Based Optimization <i>Michael S Yao, James Gee, Osbert Bastani</i>	45464
GRAIL: Graph Edit Distance and Node Alignment using LLM-Generated Code <i>Samidha Verma, Arushi Goyal, Ananya Mathur, Ankit Anand, Sayan Ranu</i>	45516

PipeOffload: Improving Scalability of Pipeline Parallelism with Memory Optimization <i>Xinyi Wan, Penghui Qi, Guangxing Huang, Min Lin, Jialin Li</i>	45538
Accurate Identification of Communication Between Multiple Interacting Neural Populations <i>Belle Liu, Jacob Sacks, Matthew D. Golub</i>	45552
SafeMap: Robust HD Map Construction from Incomplete Observations <i>Xiaoshuai Hao, Lingdong Kong, Rong Yin, Pengwei Wang, Jing Zhang, Yunfeng Diao, Shu Zhao</i>	45576
Optimal Transport Barycenter via Nonconvex-Concave Minimax Optimization <i>Kaheon Kim, Rentian Yao, Changbo Zhu, Xiaohui Chen</i>	45588
ViSTA: Video Hierarchical Spatial-Temporal Direct Preference Optimization for Large Video Models <i>Haojian Huang, Haodong Chen, Shengqiong Wu, Meng Luo, Jinlan Fu, Xinya Du, Hanwang Zhang, Hao Fei</i>	45609
How to Move Your Dragon: Text-to-Motion Synthesis for Large-Vocabulary Objects <i>Wonkwang Lee, Jongwon Jeong, Taehong Moon, Hyeon-Jong Kim, Jaehyeon Kim, Gunhee Kim, Byeong-Uk Lee</i>	45628
Heterogeneous Data Game: Characterizing the Model Competition Across Multiple Data Sources <i>Renzhe Xu, Kang Wang, Bo Li</i>	45647
Differentially Private Space-Efficient Algorithms for Counting Distinct Elements in the Turnstile Model <i>Rachel Cummings, Alessandro Epasto, Jieming Mao, Tamalika Mukherjee, Tingting Ou, Peilin Zhong</i>	45670
CurvGAD: Leveraging Curvature for Enhanced Graph Anomaly Detection <i>Karish Grover, Geoffrey J. Gordon, Christos Faloutsos</i>	45700
When Do LLMs Help With Node Classification? A Comprehensive Analysis <i>Xixi Wu, Yifei Shen, Fangzhou Ge, Caihua Shan, Yizhu Jiao, Xiangguo Sun, Hong Cheng</i>	45719
Protein Structure Tokenization: Benchmarking and New Recipe <i>Xinyu Yuan, Zichen Wang, Marcus D. Collins, Huzefa Rangwala</i>	45746
Geometry-Informed Neural Networks <i>Arturs Berzins, Andreas Radler, Eric Volkmann, Sebastian Sanokowski, Sepp Hochreiter, Johannes Brandstetter</i>	45772
SPACE: Your Genomic Profile Predictor is a Powerful DNA Foundation Model <i>Zhao Yang, Jiwei Zhu, Bing Su</i>	45801
Dataflow-Guided Neuro-Symbolic Language Models for Type Inference <i>Gen Li, Yao Wan, Hongyu Zhang, Zhou Zhao, Wenbin Jiang, Xuanhua Shi, Hai Jin, Zheng Wang</i>	45821
Investigating the Overlooked Hessian Structure: From CNNs to LLMs <i>Qian-Yuan Tang, Yufei Gu, Yunfeng Cai, Mingming Sun, Ping Li, Zhou Xun, Zeke Xie</i>	45842

Dimension-Independent Rates for Structured Neural Density Estimation <i>Robert A. Vandermeulen, Wai Ming Tai, Bryon Aragam</i>	45869
Can Biologically Plausible Temporal Credit Assignment Rules Match BPTT for Neural Similarity? E-prop as an Example <i>Yuhan Helena Liu, Guangyu Robert Yang, Christopher J Cueva</i>	45892
HyperIV: Real-time Implied Volatility Smoothing <i>Yongxin Yang, Wenqi Chen, Chao Shu, Timothy Hospedales</i>	45917
LoRA Training Provably Converges to a Low-Rank Global Minimum Or It Fails Loudly (But it Probably Won't Fail) <i>Junsu Kim, Jaeyeon Kim, Ernest K. Ryu</i>	45932
One Image is Worth a Thousand Words: A Usability Preservable Text-Image Collaborative Erasing Framework <i>Feiran Li, Qianqian Xu, Shilong Bao, Zhiyong Yang, Xiaochun Cao, Qingming Huang</i>	45956
Efficient Curvature-Aware Hypergradient Approximation for Bilevel Optimization <i>Youran Dong, Junfeng Yang, Wei Yao, Jin Zhang</i>	46000
Wasserstein Policy Optimization <i>David Pfau, Ian Davies, Diana L Borsa, João Guilherme Madeira Araújo, Brendan Daniel Tracey, Hado van Hasselt</i>	46029
Tackling View-Dependent Semantics in 3D Language Gaussian Splatting <i>Jiazhong Cen, Xudong Zhou, Jiemin Fang, Changsong Wen, Lingxi Xie, XIAOPENG ZHANG, Wei Shen, Qi Tian</i>	46051
Holistic Physics Solver: Learning PDEs in a Unified Spectral-Physical Space <i>Xihang Yue, Yi Yang, Linchao Zhu</i>	46073
Minimalist Concept Erasure in Generative Models <i>Yang Zhang, Er Jin, Yanfei Dong, Yixuan Wu, Philip Torr, Ashkan Khakzar, Johannes Stegmaier, Kenji Kawaguchi</i>	46093
Task-Gated Multi-Expert Collaboration Network for Degraded Multi-Modal Image Fusion <i>Yiming Sun, Xin Li, Pengfei Zhu, Qinghua Hu, Dongwei Ren, Huiying Xu, Xinzhong Zhu</i>	46117
M+: Extending MemoryLLM with Scalable Long-Term Memory <i>Yu Wang, Dmitry Krotov, Yuanzhe Hu, Yifan Gao, Wangchunshu Zhou, Julian McAuley, Dan Gutfreund, Rogerio Feris, Zexue He</i>	46133
Rethinking Point Cloud Data Augmentation: Topologically Consistent Deformation <i>Jian Bi, Qianliang Wu, Xiang Li, Shuo Chen, Jianjun Qian, lei luo, Jian Yang</i>	46149
Reinforced Learning Explicit Circuit Representations for Quantum State Characterization from Local Measurements <i>Manwen Liao, Yan Zhu, Weitian Zhang, Yuxiang Yang</i>	46170

The dark side of the forces: assessing non-conservative force models for atomistic machine learning <i>Filippo Bigi, Marcel F. Langer, Michele Ceriotti</i>	46189
Q-VDiT: Towards Accurate Quantization and Distillation of Video-Generation Diffusion Transformers <i>Weilun Feng, Chuanguang Yang, Haotong Qin, Xiangqi Li, Yu Wang, Zhulin An, Libo Huang, Boyu Diao, Zixiang Zhao, Yongjun Xu, Michele Magno</i>	46220
Potemkin Understanding in Large Language Models <i>Marina Mancoridis, Bec Weeks, Keyon Vafa, Sendhil Mullainathan</i>	46241
Empirical Privacy Variance <i>Yuzheng Hu, Fan Wu, Ruicheng Xian, Yuhang Liu, Lydia Zakyntinou, Pritish Kamath, Chiyuan Zhang, David Forsyth</i>	46266
Physics-Informed DeepONets for drift-diffusion on metric graphs: simulation and parameter identification <i>Jan Blechschmidt, Tom-Christian Riemer, Max Winkler, Martin STOLL, Jan-Frederik Pietschmann</i>	46310
EnIGMA: Interactive Tools Substantially Assist LM Agents in Finding Security Vulnerabilities <i>Talor Abramovich, Meet Udeshi, Minghao Shao, Kilian Lieret, Haoran Xi, Kimberly Milner, Sofija Jancheska, John Yang, Carlos E Jimenez, Farshad Khorrami, Prashanth Krishnamurthy, Brendan Dolan-Gavitt, Muhammad Shafique, Karthik R Narasimhan, Ramesh Karri, Ofir Press</i>	46335
Internal Causal Mechanisms Robustly Predict Language Model Out-of-Distribution Behaviors <i>Jing Huang, Junyi Tao, Thomas Icard, Diyi Yang, Christopher Potts</i>	46445
CALM: Consensus-Aware Localized Merging for Multi-Task Learning <i>Kunda Yan, Min Zhang, Sen Cui, Qu Zikun, Bo Jiang, Feng Liu, Changshui Zhang</i>	46467
Lower Bounds for Chain-of-Thought Reasoning in Hard-Attention Transformers <i>Alireza Amiri Bavandpour, Xinting Huang, Mark Rofin, Michael Hahn</i>	46488
Token Coordinated Prompt Attention is Needed for Visual Prompting <i>Zichen Liu, Xu Zou, Gang Hua, Jiahuan Zhou</i>	46521
WyckoffDiff -- A Generative Diffusion Model for Crystal Symmetry <i>Filip Ekström Kelvinius, Oskar B. Andersson, Abhijith S Parackal, Dong Qian, Rickard Armiento, Fredrik Lindsten</i>	46534
A Versatile Influence Function for Data Attribution with Non-Decomposable Loss <i>Junwei Deng, Weijing Tang, Jiaqi W. Ma</i>	46552
Compositional Causal Reasoning Evaluation in Language Models <i>Jacqueline R. M. A. Maasch, Alihan Hüyük, Xinnuo Xu, Aditya V. Nori, Javier Gonzalez</i>	46573
On Temperature Scaling and Conformal Prediction of Deep Classifiers <i>Lahav Dabah, Tom Tirer</i>	46614

DMOSpeech: Direct Metric Optimization via Distilled Diffusion Model in Zero-Shot Speech Synthesis <i>Yinghao Aaron Li, Rithesh Kumar, Zeyu Jin</i>	46647
From Local Details to Global Context: Advancing Vision-Language Models with Attention-Based Selection <i>Lincan Cai, Jingxuan Kang, Shuang Li, Wenxuan Ma, Binhui Xie, Zhida Qin, Jian Liang</i>	46670
Enhancing Spectral GNNs: From Topology and Perturbation Perspectives <i>Taoyang Qin, Ke-Jia CHEN, Zheng Liu</i>	46684
Fundamental limits of learning in sequence multi-index models and deep attention networks: high-dimensional asymptotics and sharp thresholds <i>Emanuele Troiani, Hugo Cui, Yatin Dandi, Florent Krzakala, Lenka Zdeborova</i>	46705
Stable Fair Graph Representation Learning with Lipschitz Constraint <i>Qiang Chen, Zhongze Wu, Xiu Su, Xi Lin, Zhe Qu, Shan You, Shuo Yang, Chang Xu</i>	46741
Linear Bandits with Partially Observable Features <i>Wonyoung Kim, Sungwoo PARK, Garud Iyengar, assaf zeevi, Min-hwan Oh</i>	46764
Hypothesis Testing for Generalized Thurstone Models <i>Anuran Makur, Japneet Singh</i>	46796
Epsilon-VAE: Denoising as Visual Decoding <i>Long Zhao, Sanghyun Woo, Ziyu Wan, YANDONG LI, Han Zhang, Boqing Gong, Hartwig Adam, Xuhui Jia, Ting Liu</i>	46831
Nested Expectations with Kernel Quadrature <i>Zonghao Chen, Masha Naslidnyk, Francois-Xavier Briol</i>	46851
Physics-informed Temporal Alignment for Auto-regressive PDE Foundation Models <i>Congcong Zhu, Xiaoyan Xu, Jiayue Han, Jingrun Chen</i>	46885
POQD: Performance-Oriented Query Decomposer for Multi-vector retrieval <i>Yaoyang Liu, Junlin Li, Yinjun Wu, zhen chen</i>	46921
Boosting Virtual Agent Learning and Reasoning: A Step-Wise, Multi-Dimensional, and Generalist Reward Model with Benchmark <i>Bingchen Miao, Yang Wu, Minghe Gao, Qifan Yu, Wendong Bu, Wenqiao Zhang, Yunfei Li, Siliang Tang, Tat-Seng Chua, Juncheng Li</i>	46941
Demystifying Long Chain-of-Thought Reasoning <i>Shiming Yang, Yuxuan Tong, Xinyao Niu, Graham Neubig, Xiang Yue</i>	46963
Privacy-Shielded Image Compression: Defending Against Exploitation from Vision-Language Pretrained Models <i>Xuelin Shen, Jiayin Xu, Kangsheng Yin, Wenhan Yang</i>	46996

LensLLM: Unveiling Fine-Tuning Dynamics for LLM Selection <i>Xinyue Zeng, Haohui Wang, Junhong Lin, Jun Wu, Tyler Cody, Dawei Zhou</i>	47007
Learning Safe Strategies for Value Maximizing Buyers in Uniform Price Auctions <i>Negin Golrezaei, Sourav Sahoo</i>	47029
The Harder Path: Last Iterate Convergence for Uncoupled Learning in Zero-Sum Games with Bandit Feedback <i>Côme Fiegel, Pierre Menard, Tadashi Kozuno, Michal Valko, Vianney Perchet</i>	47066
Learning In-context n-grams with Transformers: Sub-n-grams Are Near-Stationary Points <i>Aditya Varre, Gizem Yüce, Nicolas Flammarion</i>	47088
An Efficient Private GPT Never Autoregressively Decodes <i>Zhengyi Li, Yue Guan, Kang Yang, Yu Feng, Ning Liu, Yu Yu, Jingwen Leng, Minyi Guo</i>	47128
Volume Optimality in Conformal Prediction with Structured Prediction Sets <i>Chao Gao, Liren Shan, Vaidehi Srinivas, Aravindan Vijayaraghavan</i>	47147
Empowering World Models with Reflection for Embodied Video Prediction <i>Xiaowei Chi, Chun-Kai Fan, Hengyuan Zhang, Xingqun Qi, Rongyu Zhang, Anthony Chen, Chi-Min Chan, Wei Xue, Qifeng Liu, Shanghang Zhang, Yike Guo</i>	47180
Efficient Fine-Grained Guidance for Diffusion Model Based Symbolic Music Generation <i>Tingyu Zhu, Haoyu Liu, Ziyu Wang, Zhimin Jiang, Zeyu Zheng</i>	47208
Minimax Optimal Regret Bound for Reinforcement Learning with Trajectory Feedback <i>Zihan Zhang, Yuxin Chen, Jason D. Lee, Simon Shaolei Du, Ruosong Wang</i>	47233
Concentration Distribution Learning from Label Distributions <i>Jiawei Tang, Yuheng Jia</i>	47255
Training Flexible Models of Genetic Variant Effects from Functional Annotations using Accelerated Linear Algebra <i>Alan Nawzad Amin, Andres Potapczynski, Andrew Gordon Wilson</i>	47265
Equivariant Neural Tangent Kernels <i>Philipp Misof, Pan Kessel, Jan E Gerken</i>	47283
On the Interplay between Graph Structure and Learning Algorithms in Graph Neural Networks <i>Junwei Su, Chuan Wu</i>	47317
Unified Breakdown Analysis for Byzantine Robust Gossip <i>Renaud Gaucher, Aymeric Dieuleveut, Hadrien Hendrikx</i>	47345
Towards Understanding Gradient Dynamics of the Sliced-Wasserstein Distance via Critical Point Analysis <i>Christophe Vauthier, Anna Korba, Quentin Mérigot</i>	47374

Accelerating Large Language Model Reasoning via Speculative Search <i>Zhihai Wang, Jie Wang, Jilai Pan, Xilin Xia, Huiling Zhen, Mingxuan Yuan, Jianye HAO, Feng Wu</i>	47411
A Theory for Conditional Generative Modeling on Multiple Data Sources <i>Rongzhen Wang, Yan Zhang, Chenyu Zheng, Chongxuan Li, Guoqiang Wu</i>	47444
R2-T2: Re-Router in Test-Time for Multimodal Mixture-of-Experts <i>Zhongyang Li, Ziyue Li, Tianyi Zhou</i>	47481
Learn Singularly Perturbed Solutions via Homotopy Dynamics <i>Chuqi CHEN, Yahong Yang, Yang Xiang, Wenrui Hao</i>	47506
Aligned Multi Objective Optimization <i>Yonathan Efroni, Ben Kretzu, Daniel Jiang, Jalaj Bhandari, Zheqing Zhu, Karen Ullrich</i>	47530
Contrastive Private Data Synthesis via Weighted Multi-PLM Fusion <i>Tianyuan Zou, Yang Liu, Peng Li, Yufei Xiong, Jianqing Zhang, Jingjing Liu, Xiaozhou Ye, Ye Ouyang, Ya-Qin Zhang</i>	47559
Heterogeneous Treatment Effect in Time-to-Event Outcomes: Harnessing Censored Data with Recursively Imputed Trees <i>Tomer Meir, Uri Shalit, Malka Gorfine</i>	47586
Data-driven Design of Randomized Control Trials with Guaranteed Treatment Effects <i>Santiago Cortes-Gomez, Naveen Janaki Raman, Aarti Singh, Bryan Wilder</i>	47619
A Classification View on Meta Learning Bandits <i>Mirco Mutti, Jeongyeol Kwon, Shie Mannor, Aviv Tamar</i>	47634
QT-DoG: Quantization-Aware Training for Domain Generalization <i>Saqib Javed, Hieu Le, Mathieu Salzmann</i>	47654
Zero-Shot Adaptation of Parameter-Efficient Fine-Tuning in Diffusion Models <i>Farzad Farhadzadeh, Debasmit Das, Shubhankar Borse, Fatih Porikli</i>	47678
FedBEs: One-Shot Federated Learning based on Bayesian Ensemble <i>Jacopo Talpini, Marco Savi, Giovanni Neglia</i>	47695
AMPO: Active Multi Preference Optimization for Self-play Preference Selection <i>Taneesh Gupta, Rahul Madhavan, Xuchao Zhang, Chetan Bansal, Saravan Rajmohan</i>	47714
ExPLoRA: Parameter-Efficient Extended Pre-Training to Adapt Vision Transformers under Domain Shifts <i>Samar Khanna, Medhanie Irgau, David B. Lobell, Stefano Ermon</i>	47741
Enhancing Adversarial Robustness with Conformal Prediction: A Framework for Guaranteed Model Reliability <i>Jie Bao, Chuangyin Dang, Rui Luo, Hanwei Zhang, Zhixin Zhou</i>	47761

Generation from Noisy Examples <i>Ananth Raman, Vinod Raman</i>	47781
Differentiable Solver Search for Fast Diffusion Sampling <i>Shuai Wang, Zexian Li, Qipeng zhang, Tianhui Song, Xubin Li, Tiezheng Ge, Bo Zheng, Limin Wang</i>	47796
Collaborative Mean Estimation Among Heterogeneous Strategic Agents: Individual Rationality, Fairness, and Truthful Contribution <i>Alex Clinton, Yiding Chen, Jerry Zhu, Kirthevasan Kandasamy</i>	47816
Weak-to-Strong Generalization Even in Random Feature Networks, Provably <i>Marko Medvedev, Kaifeng Lyu, Dingli Yu, Sanjeev Arora, Zhiyuan Li, Nathan Srebro</i>	47866
When do neural networks learn world models? <i>Tianren Zhang, Guanyu Chen, Feng Chen</i>	47904
Reward-free World Models for Online Imitation Learning <i>Shangzhe Li, Zhiao Huang, Hao Su</i>	47936
AdaSplash: Adaptive Sparse Flash Attention <i>Nuno Gonçalves, Marcos V Treviso, Andre Martins</i>	47959
On the Private Estimation of Smooth Transport Maps <i>Clément Lalanne, Franck Iutzeler, Jean-Michel Loubes, Julien Chhor</i>	47978
Adapting While Learning: Grounding LLMs for Scientific Problems with Tool Usage Adaptation <i>Bohan Lyu, Yadi Cao, Duncan Watson-Parris, Leon Bergen, Taylor Berg-Kirkpatrick, Rose Yu</i>	48011
Staged and Physics-Grounded Learning Framework with Hyperintensity Prior for Pre-Contrast MRI Synthesis <i>Dayang Wang, Srivathsa Pasumarthi Venkata, Ajit Shankaranarayanan, Greg Zaharchuk</i>	48048
Differential Coding for Training-Free ANN-to-SNN Conversion <i>Zihan Huang, Wei Fang, Tong Bu, Peng Xue, Zecheng Hao, Wenxuan Liu, Yuanhong Tang, Zhaoifei Yu, Tiejun Huang</i>	48066
OW-VAP: Visual Attribute Parsing for Open World Object Detection <i>Xing Xi, Xing Fu, Weiqiang Wang, Ronghua Luo</i>	48096
BiAssemble: Learning Collaborative Affordance for Bimanual Geometric Assembly <i>Yan Shen, Ruihai Wu, Yubin Ke, Xinyuan Song, Zeyi Li, Xiaoqi Li, Hongwei Fan, Haoran Lu, Hao Dong</i>	48114
SAH-Drive: A Scenario-Aware Hybrid Planner for Closed-Loop Vehicle Trajectory Generation <i>Yuqi Fan, Zhiyong Cui, Zhenning Li, Yilong Ren, Haiyang Yu</i>	48137
Separating Knowledge and Perception with Procedural Data <i>Adrian Rodriguez-Munoz, Manel Baradad, Phillip Isola, Antonio Torralba</i>	48154

FSTLLM: Spatio-Temporal LLM for Few Shot Time Series Forecasting <i>YUE JIANG, Yile Chen, Xiucheng Li, Qin Chao, SHUAI LIU, Gao Cong</i>	48169
Smooth Interpolation for Improved Discrete Graph Generative Models <i>Yuxuan Song, Juntong Shi, Jingjing Gong, Minkai Xu, Stefano Ermon, Hao Zhou, Wei-Ying Ma</i>	48187
Fast and Low-Cost Genomic Foundation Models via Outlier Removal <i>Haozheng Luo, Chenghao Qiu, Maojiang Su, Zhihan Zhou, Zoe Mehta, Guo Ye, Jerry Yao-Chieh Hu, Han Liu</i>	48213
Unifying Specialized Visual Encoders for Video Language Models <i>Jihoon Chung, Tyler Zhu, Max Gonzalez Saez-Diez, Juan Carlos Niebles, Honglu Zhou, Olga Russakovsky</i>	48249
LoRA-Gen: Specializing Large Language Model via Online LoRA Generation <i>Yicheng Xiao, Lin Song, Rui Yang, Cheng Cheng, Yixiao Ge, Xiu Li, Ying Shan</i>	48271
Correlation Clustering Beyond the Pivot Algorithm <i>Soheil Behnezhad, Moses Charikar, Vincent Cohen-Addad, Alma Ghafari, Weiyun ma</i>	48284
Auditing f-differential privacy in one run <i>Saeed Mahloujifar, Luca Melis, Kamalika Chaudhuri</i>	48311
MATH-Perturb: Benchmarking LLMs' Math Reasoning Abilities against Hard Perturbations <i>Kaixuan Huang, Jiacheng Guo, Zihao Li, Xiang Ji, Jiawei Ge, Wenzhe Li, Yingqing Guo, Tianle Cai, Hui Yuan, Runzhe Wang, Yue Wu, Ming Yin, Shange Tang, Yangsibo Huang, Chi Jin, Xinyun Chen, Chiyuan Zhang, Mengdi Wang</i>	48338
Subgroups Matter for Robust Bias Mitigation <i>Anissa Alloula, Charles Jones, Ben Glocker, Bartłomiej Papiez</i>	48356
Gradient-based Explanations for Deep Learning Survival Models <i>Sophie Hanna Langbein, Niklas Koenen, Marvin N. Wright</i>	48378
Diversified Flow Matching with Translation Identifiability <i>Sagar Shrestha, Xiao Fu</i>	48409
Patch-wise Structural Loss for Time Series Forecasting <i>Dilfira Kudrat, Zongxia Xie, Yanru Sun, Tianyu Jia, Qinghua Hu</i>	48428
Taming Diffusion for Dataset Distillation with High Representativeness <i>Lin Zhao, Yushu Wu, Xinru Jiang, Jianyang Gu, Yanzhi Wang, Xiaolin Xu, Pu Zhao, Xue Lin</i>	48447
AutoML-Agent: A Multi-Agent LLM Framework for Full-Pipeline AutoML <i>Patara Trirat, Wonyong Jeong, Sung Ju Hwang</i>	48468
A Tale of Two Structures: Do LLMs Capture the Fractal Complexity of Language? <i>Ibrahim Alabdulmohsin, Andreas Peter Steiner</i>	48516

Grammar-Forced Translation of Natural Language to Temporal Logic using LLMs <i>William H English, Dominic Simon, Sumit Kumar Jha, Rickard Ewetz</i>	48578
Unlocking the Power of Rehearsal in Continual Learning: A Theoretical Perspective <i>Junze Deng, Qinhang Wu, Peizhong Ju, Sen Lin, Yingbin Liang, Ness Shroff</i>	48592
Stealix: Model Stealing via Prompt Evolution <i>Zhixiong Zhuang, Hui-Po Wang, Maria-Irina Nicolae, Mario Fritz</i>	48642
Reflection-Window Decoding: Text Generation with Selective Refinement <i>Zeyu Tang, Zhenhao Chen, Xiangchen Song, Loka Li, Yunlong Deng, Yifan Shen, Guangyi Chen, Peter Spirtes, Kun Zhang</i>	48660
Fast and Provable Algorithms for Sparse PCA with Improved Sample Complexity <i>Jian-Feng Cai, Zhuozhi XIAN, Jiayi Ying</i>	48686
A Closer Look at Generalized BH Algorithm for Out-of-Distribution Detection <i>Xinsong Ma, Jie Wu, Weiwei Liu</i>	48708
Synthesizing Software Engineering Data in a Test-Driven Manner <i>Lei Zhang, Jiayi Yang, Min Yang, Jian Yang, Mouxiang Chen, Jiajun Zhang, Zeyu Cui, Binyuan Hui, Junyang Lin</i>	48724
PlaySlot: Learning Inverse Latent Dynamics for Controllable Object-Centric Video Prediction and Planning <i>Angel Villar-Corrales, Sven Behnke</i>	48747
DataDecide: How to Predict Best Pretraining Data with Small Experiments <i>Ian Magnusson, Nguyen Tai, Ben Bogin, David Heineman, Jena D. Hwang, Luca Soldaini, Akshita Bhagia, Jiacheng Liu, Dirk Groeneveld, Oyvind Tafford, Noah A. Smith, Pang Wei Koh, Jesse Dodge</i>	48777
From Complex to Atomic: Enhancing Augmented Generation via Knowledge-Aware Dual Rewriting and Reasoning <i>Jinyu Wang, Jingjing Fu, Rui Wang, Lei Song, Jiang Bian</i>	48793
One Stone, Two Birds: Enhancing Adversarial Defense Through the Lens of Distributional Discrepancy <i>Jiacheng Zhang, Benjamin I. P. Rubinstein, Jingfeng Zhang, Feng Liu</i>	48819
Random Policy Evaluation Uncovers Policies of Generative Flow Networks <i>Haoran He, Emmanuel Bengio, Qingpeng Cai, Ling Pan</i>	48839
Improved Expressivity of Hypergraph Neural Networks through High-Dimensional Generalized Weisfeiler-Leman Algorithms <i>Detian Zhang, Chengqiang Zhang, Yanghui Rao, Li Qing, Chunjiang Zhu</i>	48856

Adaptive Median Smoothing: Adversarial Defense for Unlearned Text-to-Image Diffusion Models at Inference Time <i>Xiaoxuan Han, Songlin Yang, Wei Wang, Yang Li, Jing Dong</i>	48885
Divide and Conquer: Grounding LLMs as Efficient Decision-Making Agents via Offline Hierarchical Reinforcement Learning <i>Zican Hu, Wei Liu, Xiaoye Qu, Xiangyu Yue, Chunlin Chen, Zhi Wang, Yu Cheng</i>	48901
e-GAI: e-value-based Generalized α-Investing for Online False Discovery Rate Control <i>Yifan Zhang, Zijian Wei, Haojie Ren, Changliang Zou</i>	48922
When Data-Free Knowledge Distillation Meets Non-Transferable Teacher: Escaping Out-of-Distribution Trap is All You Need <i>Ziming Hong, Runnan Chen, Zengmao Wang, Bo Han, Bo Du, Tongliang Liu</i>	48946
Leveraging Skills from Unlabeled Prior Data for Efficient Online Exploration <i>Max Wilcoxson, Qiyang Li, Kevin Frans, Sergey Levine</i>	48983
How Much Can We Forget about Data Contamination? <i>Sebastian Bordt, Suraj Srinivas, Valentyn Boreiko, Ulrike von Luxburg</i>	49011
Demystifying Catastrophic Forgetting in Two-Stage Incremental Object Detector <i>Qirui Wu, Shizhou Zhang, De Cheng, Yinghui Xing, Di Xu, PENG WANG, Yanning Zhang</i>	49030
Meta-Reinforcement Learning with Adaptation from Human Feedback via Preference-Order-Preserving Task Embedding <i>Siyuan Xu, Minghui Zhu</i>	49045
S²FGL: Spatial Spectral Federated Graph Learning <i>Zihan Tan, Suyuan Huang, Guancheng Wan, Wenke Huang, He Li, Mang Ye</i>	49070
The Brain's Bitter Lesson: Scaling Speech Decoding With Self-Supervised Learning <i>Dulhan Jayalath, Gilad Landau, Brendan Shillingford, Mark Woolrich, Oiwi Parker Jones</i>	49082
De-coupled NeuroGF for Shortest Path Distance Approximations on Large Terrain Graphs <i>Samantha Chen, Pankaj K Agarwal, Yusu Wang</i>	49098
Explanatory Instructions: Towards Unified Vision Tasks Understanding and Zero-shot Generalization <i>Yang Shen, Xiu-Shen Wei, Yifan Sun, YuXin Song, Tao Yuan, Jian Jin, He-Yang Xu, Yazhou Yao, Errui Ding</i>	49114
EncryptedLLM: Privacy-Preserving Large Language Model Inference via GPU-Accelerated Fully Homomorphic Encryption <i>Leo de Castro, Daniel Escudero, Adya Agrawal, Antigoni Polychroniadou, Manuela Veloso</i>	49158
Learning Representations of Instruments for Partial Identification of Treatment Effects <i>Jonas Schweisthal, Dennis Frauen, Maresa Schröder, Konstantin Hess, Niki Kilbertus, Stefan Feuerriegel</i>	49170

Model Uncertainty Quantification by Conformal Prediction in Continual Learning <i>Rui Gao, Weiwei Liu</i>	49195
Matryoshka Quantization <i>Pranav Ajit Nair, Puranjay Datta, Jeff Dean, Prateek Jain, Aditya Kusupati</i>	49212
Federated Learning for Feature Generalization with Convex Constraints <i>Dongwon Kim, Donghee Kim, Sung Kuk Shyn, Kwangsu Kim</i>	49235
Integer Programming for Generalized Causal Bootstrap Designs <i>Jennifer Rogers Brennan, Sebastien Lahaie, Adel Javanmard, Nick Doudchenko, Jean Pouget-Abadie</i>	49254
Representative Ranking for Deliberation in the Public Sphere <i>Manon Revel, Smitha Milli, Tyler Lu, Jamelle Watson-Daniels, Maximilian Nickel</i>	49277
Tuning Sequential Monte Carlo Samplers via Greedy Incremental Divergence Minimization <i>Kyurae Kim, Zuheng Xu, Jacob R. Gardner, Trevor Campbell</i>	49308
Towards Trustworthy Federated Learning with Untrusted Participants <i>Youssef Allouah, Rachid Guerraoui, John Stephan</i>	49361
Sample Complexity of Distributionally Robust Off-Dynamics Reinforcement Learning with Online Interaction <i>Yiting He, Zhishuai Liu, Weixin Wang, Pan Xu</i>	49405
GANQ: GPU-Adaptive Non-Uniform Quantization for Large Language Models <i>Pengxiang Zhao, Xiaoming Yuan</i>	49457
Robust Reward Alignment via Hypothesis Space Batch Cutting <i>Zhixian Xie, Haode Zhang, Yizhe Feng, Wanxin Jin</i>	49472
Aguvis: Unified Pure Vision Agents for Autonomous GUI Interaction <i>Yiheng Xu, Zekun Wang, Junli Wang, Dunjie Lu, Tianbao Xie, Amrita Saha, Doyen Sahoo, Tao Yu, Caiming Xiong</i>	49491
Reasoning Through Execution: Unifying Process and Outcome Rewards for Code Generation <i>Zhuohao Yu, Weizheng Gu, Yidong Wang, Xingru Jiang, Zhengran Zeng, Jindong Wang, Wei Ye, Shikun Zhang</i>	49525
PROTOCOL: Partial Optimal Transport-enhanced Contrastive Learning for Imbalanced Multi-view Clustering <i>Xuqian Xue, Yiming Lei, Qi Cai, Hongming Shan, Junping Zhang</i>	49548
SSHR: More Secure Generative Steganography with High-Quality Revealed Secret Images <i>Jiannian Wang, Yao Lu, Guangming Lu</i>	49563
Conformal Prediction as Bayesian Quadrature <i>Jake C. Snell, Thomas L. Griffiths</i>	49579

Learning to Plan & Reason for Evaluation with Thinking-LLM-as-a-Judge <i>Swarnadeep Saha, Xian Li, Marjan Ghazvininejad, Jason E Weston, Tianlu Wang</i>	49596
Inducing, Detecting and Characterising Neural Modules: A Pipeline for Functional Interpretability in Reinforcement Learning <i>Anna Soligo, Pietro Ferraro, David Boyle</i>	49615
Preconditioned Riemannian Gradient Descent Algorithm for Low-Multilinear-Rank Tensor Completion <i>Yuanwei Zhang, Fengmiao Bian, Xiaoqun Zhang, Jian-Feng Cai</i>	49641
Semantics-aware Test-time Adaptation for 3D Human Pose Estimation <i>Qiuxia Lin, Rongyu Chen, Kerui Gu, Angela Yao</i>	49666
Rethinking Time Encoding via Learnable Transformation Functions <i>Xi Chen, Yateng Tang, Jiarong Xu, Jiawei Zhang, Siwei Zhang, Sijia Peng, Xuehao Zheng, Yun Xiong</i>	49683
SLiM-LLM: Salience-Driven Mixed-Precision Quantization for Large Language Models <i>Wei Huang, Haotong Qin, Yangdong Liu, Yawei Li, Qinshuo Liu, Xianglong Liu, Luca Benini, Michele Magno, Shiming Zhang, XIAOJUAN QI</i>	49715
Learning Time-Varying Multi-Region Brain Communications via Scalable Markovian Gaussian Processes <i>Weihan Li, Yule Wang, Chengrui Li, Anqi Wu</i>	49736
Improving Rationality in the Reasoning Process of Language Models through Self-playing Game <i>Pinzheng Wang, Juntao Li, Zecheng Tang, Haijia Gui, Min zhang</i>	49757
A Generalization Result for Convergence in Learning-to-Optimize <i>Michael Sucker, Peter Ochs</i>	49779
On-the-Fly Adaptive Distillation of Transformer to Dual-State Linear Attention for Long-Context LLM Serving <i>Yeonju Ro, Zhenyu Zhang, Souvik Kundu, Zhangyang Wang, Aditya Akella</i>	49798
SpeCache: Speculative Key-Value Caching for Efficient Generation of LLMs <i>Shibo Jie, Yehui Tang, Kai Han, Zhi-Hong Deng, Jing Han</i>	49815
One-Shot Heterogeneous Federated Learning with Local Model-Guided Diffusion Models <i>Mingzhao Yang, Shangchao Su, Bin Li, Xiangyang Xue</i>	49827
Channel Normalization for Time Series Channel Identification <i>Seunghan Lee, Taeyoung Park, Kibok Lee</i>	49847
Information Bottleneck-guided MLPs for Robust Spatial-temporal Forecasting <i>Min Chen, Guansong Pang, Wenjun Wang, Cheng Yan</i>	49880

SToFM: a Multi-scale Foundation Model for Spatial Transcriptomics <i>Suyuan Zhao, YIZHEN LUO, Ganbo Yang, Yan Zhong, Hao Zhou, Zaiqing Nie</i>	49915
Reflect-then-Plan: Offline Model-Based Planning through a Doubly Bayesian Lens <i>Jihwan Jeong, Xiaoyu Wang, Jingmin Wang, Scott Sanner, Pascal Poupart</i>	49934
MiraGe: Editable 2D Images using Gaussian Splatting <i>Joanna Waczynska, Tomasz Szczepanik, Piotr Borycki, Slawomir Tadeja, Thomas Bohné, Przemysław Spurek</i>	49955
FlexControl: Computation-Aware Conditional Control with Differentiable Router for Text-to-Image Generation <i>Zheng Fang, Lichuan Xiang, Xu Cai, Kaicheng Zhou, Hongkai Wen</i>	49972
SPEX: Scaling Feature Interaction Explanations for LLMs <i>Justin Singh Kang, Landon Butler, Abhineet Agarwal, Yigit Efe Erginbas, Ramtin Pedarsani, Bin Yu, Kannan Ramchandran</i>	49992
UDora: A Unified Red Teaming Framework against LLM Agents by Dynamically Hijacking Their Own Reasoning <i>Jiawei Zhang, Shuang Yang, Bo Li</i>	50018
Efficient Heterogeneity-Aware Federated Active Data Selection <i>Ying-Peng Tang, Chao Ren, Xiaoli Tang, Sheng-Jun Huang, Lizhen Cui, Han Yu</i>	50038
Effective and Efficient Masked Image Generation Models <i>Zebin You, Jingyang Ou, Xiaolu Zhang, Jun Hu, JUN ZHOU, Chongxuan Li</i>	50051
QEM-Bench: Benchmarking Learning-based Quantum Error Mitigation and QEMFormer as a Multi-ranged Context Learning Baseline <i>Tianyi Bao, Ruizhe Zhong, Xinyu Ye, Yehui Tang, Junchi Yan</i>	50068
Human Cognition-Inspired Hierarchical Fuzzy Learning Machine <i>Junbiao Cui, Qin Yue, Jianqing Liang, Jiye Liang</i>	50083
Learning Efficient Robotic Garment Manipulation with Standardization <i>Changshi Zhou, Feng Luan, huijarui, Shaoqiang Meng, Zhipeng Wang, Yanchao Dong, Yanmin Zhou, Bin He</i>	50112
Is Noise Conditioning Necessary for Denoising Generative Models? <i>Qiao Sun, Zhicheng Jiang, Hanhong Zhao, Kaiming He</i>	50128
Steering Protein Language Models <i>Long-Kai Huang, Rongyi Zhu, Bing He, Jianhua Yao</i>	50162
Behavior-Regularized Diffusion Policy Optimization for Offline Reinforcement Learning <i>Chen-Xiao Gao, Chenyang Wu, Mingjun Cao, Chenjun Xiao, Yang Yu, Zongzhang Zhang</i>	50176

Improving Multi-Class Calibration through Normalization-Aware Isotonic Techniques <i>Alon Arad, Saharon Rosset</i>	50204
POROver: Improving Safety and Reducing Overrefusal in Large Language Models with Overgeneration and Preference Optimization <i>Batuhan K. Karaman, ishman zabir, Alon Benhaim, Vishrav Chaudhary, Mert R. Sabuncu, Xia Song</i>	50234
On the Convergence of Continuous Single-timescale Actor-critic <i>Xuyang Chen, Lin Zhao</i>	50262
Conformity Score Averaging for Classification <i>Rui Luo, Zhixin Zhou</i>	50293
G-Sim: Generative Simulations with Large Language Models and Gradient-Free Calibration <i>Samuel Holt, Max Ruiz Luyten, Antonin Berthon, Mihaela van der Schaar</i>	50311
Deep Reinforcement Learning from Hierarchical Preference Design <i>Alexander Bukharin, Yixiao Li, Pengcheng He, Tuo Zhao</i>	50360
Inductive Moment Matching <i>Linqi Zhou, Stefano Ermon, Jiaming Song</i>	50379
M2PDE: Compositional Generative Multiphysics and Multi-component PDE Simulation <i>Tao Zhang, Zhenhai Liu, Feipeng Qi, Yongjun Jiao, Tailin Wu</i>	50415
On the Duality between Gradient Transformations and Adapters <i>Lucas Torroba Hennigen, Hunter Lang, Han Guo, Yoon Kim</i>	50444
Compact Matrix Quantum Group Equivariant Neural Networks <i>Edward Pearce-Crump</i>	50461
Causality Inspired Federated Learning for OOD Generalization <i>Jiayuan Zhang, Xuefeng Liu, Jianwei Niu, Shaojie Tang, Haotian Yang, Xinghao Wu</i>	50493
AI for Global Climate Cooperation: Modeling Global Climate Negotiations, Agreements, and Long-Term Cooperation in RICE-N <i>Tianyu Zhang, Andrew Robert Williams, Phillip Wozny, Kai-Hendrik Cohrs, Koen Ponse, Marco Jiralerspong, Soham Rajesh Phade, Sunil Srinivasa, Lu Li, Yang Zhang, Prateek Gupta, Erman Acar, Irina Rish, Yoshua Bengio, Stephan Zheng</i>	50516
Geometric Median (GM) Matching for Robust k-Subset Selection from Noisy Data <i>Anish Acharya, Sujay Sanghavi, Alex Dimakis, Inderjit S Dhillon</i>	50545
Revisiting Convergence: Shuffling Complexity Beyond Lipschitz Smoothness <i>Qi He, Peiran Yu, Ziyi Chen, Heng Huang</i>	50593
MxMoE: Mixed-precision Quantization for MoE with Accuracy and Performance Co-Design <i>Haojie Duanmu, Xiuhong Li, Zhihang Yuan, Size Zheng, Jiangfei Duan, Xingcheng Zhang, Dahua Lin</i>	50613

Trajectory World Models for Heterogeneous Environments <i>Shaofeng Yin, Jialong Wu, Siqiao Huang, Xingjian Su, Xu He, Jianye HAO, Mingsheng Long</i>	50627
Learning with Selectively Labeled Data from Multiple Decision-makers <i>Jian Chen, Zhehao Li, Xiaojie Mao</i>	50650
Generalized Category Discovery via Reciprocal Learning and Class-Wise Distribution Regularization <i>Duo Liu, Zhiquan Tan, Linglan Zhao, Zhongqiang Zhang, Xiangzhong Fang, Weiran Huang</i>	50690
TimeBridge: Non-Stationarity Matters for Long-term Time Series Forecasting <i>Peiyuan Liu, Beiliang Wu, Yifan Hu, Naiqi Li, Tao Dai, Jigang Bao, Shu-Tao Xia</i>	50708
Inductive Gradient Adjustment for Spectral Bias in Implicit Neural Representations <i>Kexuan Shi, Hai Chen, Leheng Zhang, Shuhang Gu</i>	50734
Improving Generalization with Flat Hilbert Bayesian Inference <i>Tuan Truong, Quyen Tran, Ngoc-Quan Pham, Nhat Ho, Dinh Phung, Trung Le</i>	50762
ICLShield: Exploring and Mitigating In-Context Learning Backdoor Attacks <i>Zhiyao Ren, Siyuan Liang, Aishan Liu, Dacheng Tao</i>	50782
Beyond Self-Interest: How Group Strategies Reshape Content Creation in Recommendation Platforms? <i>Yaolong Yu, Fan Yao, Sinno Jialin Pan</i>	50796
DIS-CO: Discovering Copyrighted Content in VLMs Training Data <i>André V. Duarte, Xuandong Zhao, Arlindo L. Oliveira, Lei Li</i>	50820
Zero-Shot Cyclic Peptide Design via Composable Geometric Constraints <i>Dapeng Jiang, Xiangzhe Kong, Jiaqi Han, Mingyu Li, Rui Jiao, Wenbing Huang, Stefano Ermon, Jianzhu Ma, Yang Liu</i>	50846
Emoji Attack: Enhancing Jailbreak Attacks Against Judge LLM Detection <i>Zhipeng Wei, Yuqi Liu, N. Benjamin Erichson</i>	50862
GraphCL: Graph-based Clustering for Semi-Supervised Medical Image Segmentation <i>Mengzhu Wang, houcheng su, Jiao Li, Chuan Li, Nan Yin, Li Shen, Jingcai Guo</i>	50877
Of Mice and Machines: A Comparison of Learning Between Real World Mice and RL Agents <i>Shuo Han, German Espinosa, Junda Huang, Daniel Dombek, Malcolm A. MacIver, Bradly C. Stadie</i>	50887
DiffAdvMAP: Flexible Diffusion-Based Framework for Generating Natural Unrestricted Adversarial Examples <i>Zhengzhao Pan, Hua Chen, Xiaogang Zhang</i>	50906
Simple Randomized Rounding for Max-Min Eigenvalue Augmentation <i>Jourdain Lamperski, Haeseong Yang, Oleg Prokopyev</i>	50931

Optimizing Social Network Interventions via Hypergradient-Based Recommender System Design <i>Marino Kühne, Panagiotis D. Grontas, Giulia De Pasquale, Giuseppe Belgioioso, Florian Dorfler, John Lygeros</i>	50945
Revisiting the Predictability of Performative, Social Events <i>Juan Carlos Perdomo</i>	50961
Beyond Task-Specific Reasoning: A Unified Conditional Generative Framework for Abstract Visual Reasoning <i>Fan Shi, Bin Li, Xiangyang Xue</i>	50975
Finite-Time Convergence Rates in Stochastic Stackelberg Games with Smooth Algorithmic Agents <i>Eric Frankel, Kshitij Kulkarni, Dmitriy Drusvyatskiy, Sewoong Oh, Lillian J. Ratliff</i>	50996
Robust Spatio-Temporal Centralized Interaction for OOD Learning <i>Jiaming Ma, Binwu Wang, Pengkun Wang, Zhengyang Zhou, Xu Wang, Yang Wang</i>	51078
OmniBal: Towards Fast Instruction-Tuning for Vision-Language Models via Omniverse Computation Balance <i>Yongqiang Yao, Jingru Tan, Feizhao Zhang, Jiahao Hu, Yazhe Niu, JinXin, Bo Li, Pengfei Liu, Ruihao Gong, Dahua Lin, Ningyi Xu</i>	51106
Self-Play Q-Learners Can Provably Collude in the Iterated Prisoner's Dilemma <i>Quentin Bertrand, Juan Agustin Duque, Emilio Calvano, Gauthier Gidel</i>	51121
Discrepancies are Virtue: Weak-to-Strong Generalization through Lens of Intrinsic Dimension <i>Yijun Dong, Yicheng Li, Yunai Li, Jason D. Lee, Qi Lei</i>	51145
NExtLong: Toward Effective Long-Context Training without Long Documents <i>Chaochen Gao, Xing W, Zijia Lin, Debing Zhang, Songlin Hu</i>	51180
Zero-shot Meta-learning for Tabular Prediction Tasks with Adversarially Pre-trained Transformer <i>Yulun Wu, Doron L Bergman</i>	51199
A Causal World Model Underlying Next Token Prediction: Exploring GPT in a Controlled Environment <i>Raanan Yehezkel Rohekar, Yaniv Gurwicz, Sungduk Yu, Estelle Aflalo, Vasudev Lal</i>	51216
Simplicity Bias and Optimization Threshold in Two-Layer ReLU Networks <i>Etienne Boursier, Nicolas Flammarion</i>	51230
Persistent Topological Features in Large Language Models <i>Yuri Gardinazzi, Karthik Viswanathan, Giada Panerai, Alessio ansuini, Alberto Cazzaniga, Matteo Biagetti</i>	51265
Linear Mode Connectivity between Multiple Models modulo Permutation Symmetries <i>Akira Ito, Masanori Yamada, Atsutoshi Kumagai</i>	51285

Edge-Colored Clustering in Hypergraphs: Beyond Minimizing Unsatisfied Edges <i>Alex Crane, Thomas Stanley, Blair D. Sullivan, Nate Veldt</i>	51301
PILAF: Optimal Human Preference Sampling for Reward Modeling <i>Yunzhen Feng, Ariel Kwiatkowski, Kunhao Zheng, Julia Kempe, Yaqi Duan</i>	51331
Retraining with Predicted Hard Labels Provably Increases Model Accuracy <i>Rudrajit Das, Inderjit S Dhillon, Alessandro Epasto, Adel Javanmard, Jieming Mao, Vahab Mirrokni, Sujay Sanghavi, Peilin Zhong</i>	51364
AutoCATE: End-to-End, Automated Treatment Effect Estimation <i>Toon Vanderschueren, Tim Verdonck, Mihaela van der Schaar, Wouter Verbeke</i>	51394
Boosting Protein Graph Representations through Static-Dynamic Fusion <i>Pengkang Guo, Bruno Correia, Pierre Vanderghenst, Daniel Probst</i>	51419
Validating Mechanistic Interpretations: An Axiomatic Approach <i>Nils Palumbo, Ravi Mangal, Zifan Wang, Saranya Vijayakumar, Corina S. Pasareanu, Somesh Jha</i>	51435
MemFreezing: A Novel Adversarial Attack on Temporal Graph Neural Networks under Limited Future Knowledge <i>Yue Dai, Liang Liu, Xulong Tang, Youtao Zhang, Jun Yang</i>	51471
NeuronTune: Towards Self-Guided Spurious Bias Mitigation <i>Guangtao Zheng, Wenqian Ye, Aidong Zhang</i>	51511
GPTAQ: Efficient Finetuning-Free Quantization for Asymmetric Calibration <i>Yuhang Li, Ruokai Yin, Donghyun Lee, Shiting Xiao, Priyadarshini Panda</i>	51535
Telling Peer Direct Effects from Indirect Effects in Observational Network Data <i>Xiaojing Du, Jiuyong Li, Debo Cheng, Lin Liu, Wentao Gao, XIONGREN CHEN, Ziqi Xu</i>	51552
Active Treatment Effect Estimation via Limited Samples <i>Zhiheng Zhang, Haoxiang Wang, Haoxuan Li, Zhouchen Lin</i>	51569
Self-cross Feature based Spiking Neural Networks for Efficient Few-shot Learning <i>Qi Xu, Junyang Zhu, Dongdong Zhou, Hao Chen, Yang Liu, Jiangrong Shen, Qiang Zhang</i>	51590
OmniAudio: Generating Spatial Audio from 360-Degree Video <i>Huadai Liu, Tianyi Luo, Kaicheng Luo, Qikai Jiang, Peiwen Sun, Jialei Wang, Rongjie Huang, Qian Chen, Wen Wang, Xiangtai Li, ShiLiang Zhang, Zhijie Yan, Zhou Zhao, Wei Xue</i>	51603
STAMP Your Content: Proving Dataset Membership via Watermarked Rephrasings <i>Saksham Rastogi, Pratyush Maini, Danish Pruthi</i>	51628
Hidden No More: Attacking and Defending Private Third-Party LLM Inference <i>Rahul Krishna Thomas, Louai Zahran, Erica Choi, Akilesh Potti, Micah Goldblum, Arka Pal</i>	51654

Exactly Tight Information-theoretic Generalization Bounds via Binary Jensen-Shannon Divergence <i>Yuxin Dong, Haoran Guo, Tieliang Gong, Wen Wen, Chen Li</i>	51690
Propagate and Inject: Revisiting Propagation-Based Feature Imputation for Graphs with Partially Observed Features <i>Daeho Um, Sunoh Kim, Jiwoong Park, Jongin Lim, Seong Jin Ahn, Seulki Park</i>	51711
MERGE³: Efficient Evolutionary Merging on Consumer-grade GPUs <i>Tommaso Mencattini, Robert Adrian Minut, Donato Crisostomi, Andrea Santilli, Emanuele Rodolà</i>	51742
Stream-level Flow Matching with Gaussian Processes <i>Ganchao Wei, Li Ma</i>	51764
Cost-efficient Collaboration between On-device and Cloud Language Models <i>Avanika Narayan, Dan Biderman, Sabri Eyuboglu, Avner May, Scott Linderman, James Zou, Christopher Re</i>	51782
Causal-PIK: Causality-based Physical Reasoning with a Physics-Informed Kernel <i>Carlota Parés Morlans, Michelle Yi, Claire Chen, Sarah A Wu, Rika Antonova, Tobias Gerstenberg, Jeannette Bohg</i>	51820
The Double-Ellipsoid Geometry of CLIP <i>Meir Yossef Levi, Guy Gilboa</i>	51833
Implicit Subgraph Neural Network <i>Yongjian Zhong, Liao Zhu, Hieu Vu, Bijaya Adhikari</i>	51854
Large Language-Geometry Model: When LLM meets Equivariance <i>Zongzhao Li, Jiacheng Cen, Bing Su, Tingyang Xu, Yu Rong, Deli Zhao, Wenbing Huang</i>	51868
Combinatorial Reinforcement Learning with Preference Feedback <i>Joongkyu Lee, Min-hwan Oh</i>	51888
Handling Imbalanced Pseudolabels for Vision-Language Models with Concept Alignment and Confusion-Aware Calibrated Margin <i>Yuchen Wang, Xuefeng Bai, Xiucheng Li, Weili Guan, Liqiang Nie, Xinyang Chen</i>	51949
INRFlow: Flow Matching for INRs in Ambient Space <i>Yuyang Wang, Anurag Ranjan, Joshua M. Susskind, Miguel Ángel Bautista</i>	51965
EPIC: Efficient Position-Independent Caching for Serving Large Language Models <i>Junhao Hu, Wenrui Huang, Weidong Wang, Haoyi Wang, tiancheng hu, zhang qin, Hao Feng, Xusheng Chen, Yizhou Shan, Tao Xie</i>	51987
Nonconvex Theory of M-estimators with Decomposable Regularizers <i>Weiwei Liu</i>	51999

Improving the Effective Receptive Field of Message-Passing Neural Networks <i>Shahaf E. Finder, Ron Shapira Weber, Moshe Eliasof, Oren Freifeld, Eran Treister</i>	52008
RE-IMAGINE: Symbolic Benchmark Synthesis for Reasoning Evaluation <i>Xinnuo Xu, Rachel Lawrence, Kshitij Dubey, Atharva Pandey, Risa Ueno, Fabian Falck, Aditya V. Nori, Rahul Sharma, Amit Sharma, Javier Gonzalez</i>	52026
Neurosymbolic World Models for Sequential Decision Making <i>Leonardo Hernandez Cano, Maxine Perroni-Scharf, Neil Dhir, Arun Ramamurthy, Armando Solar-Lezama</i>	52058
MUDDFormer: Breaking Residual Bottlenecks in Transformers via Multiway Dynamic Dense Connections <i>Da Xiao, Qingye Meng, Shengping Li, Xingyuan Yuan</i>	52074
Robust Sparsification via Sensitivity <i>Chansophea Wathanak In, Yi Li, David Woodruff, Xuan Wu</i>	52093
Federated Causal Structure Learning with Non-identical Variable Sets <i>Yunxia Wang, Fuyuan CAO, Kui Yu, Jiye Liang</i>	52111
Distributed Event-Based Learning via ADMM <i>Guner Dilsad ER, Sebastian Trimpe, Michael Muehlebach</i>	52133
A Two-Stage Learning-to-Defer Approach for Multi-Task Learning <i>Yannis Montreuil, Yeo Shu Heng, Axel Carlier, Lai Xing Ng, Wei Tsang Ooi</i>	52168
Toward Data-centric Directed Graph Learning: An Entropy-driven Approach <i>Xunkai Li, Zhengyu Wu, Kaichi Yu, Hongchao Qin, Guang Zeng, Rong-Hua Li, Guoren Wang</i>	52192
An Online Adaptive Sampling Algorithm for Stochastic Difference-of-convex Optimization with Time-varying Distributions <i>Yuhan Ye, Ying Cui, Jingyi Wang</i>	52222
Rényi Neural Processes <i>Xuesong Wang, He Zhao, Edwin V. Bonilla</i>	52242
Generalization Bounds via Meta-Learned Model Representations: PAC-Bayes and Sample Compression Hypernetworks <i>Benjamin Leblanc, Mathieu Bazinet, Nathaniel D'Amours, Alexandre Drouin, Pascal Germain</i>	52263
Automated Benchmark Generation for Repository-Level Coding Tasks <i>Konstantinos Vergopoulos, Mark Niklas Mueller, Martin Vechev</i>	52286
Is Best-of-N the Best of Them? Coverage, Scaling, and Optimality in Inference-Time Alignment <i>Audrey Huang, Adam Block, Qinghua Liu, Nan Jiang, Akshay Krishnamurthy, Dylan J Foster</i>	52312
µnit Scaling: Simple and Scalable FP8 LLM Training <i>Saaketh Narayan, Abhay Gupta, Mansheej Paul, Davis Blalock</i>	52364

Latent Imputation before Prediction: A New Computational Paradigm for De Novo Peptide Sequencing <i>Ye Du, Chen Yang, Nanxi Yu, Wanyu Lin, Qian Zhao, Shujun Wang</i>	52381
Direct Motion Models for Assessing Generated Videos <i>Kelsey R Allen, Carl Doersch, Guangyao Zhou, Mohammed Suhail, Danny Driess, Ignacio Rocco, Yulia Rubanova, Thomas Kipf, Mehdi S. M. Sajjadi, Kevin Patrick Murphy, Joao Carreira, Sjoerd van Steenkiste</i>	52394
Adaptive Sample Sharing for Multi Agent Linear Bandits <i>Hamza Cherkaoui, Merwan Barlier, Igor Colin</i>	52419
AdaWorld: Learning Adaptable World Models with Latent Actions <i>Shenyuan Gao, Siyuan Zhou, Yilun Du, Jun Zhang, Chuang Gan</i>	52452
NETS: A Non-equilibrium Transport Sampler <i>Michael Samuel Albergo, Eric Vanden-Eijnden</i>	52480
WAVE: Weighted Autoregressive Varying Gate for Time Series Forecasting <i>Jiecheng Lu, Xu Han, Yan Sun, Shihao Yang</i>	52510
How Do Large Language Monkeys Get Their Power (Laws)? <i>Rylan Schaeffer, Joshua Kazdan, John Hughes, Jordan Juravsky, Sara Price, Aengus Lynch, Erik Jones, Robert Kirk, Azalia Mirhoseini, Sanmi Koyejo</i>	52537
Policy Design for Two-sided Platforms with Participation Dynamics <i>Haruka Kiyohara, Fan Yao, Sarah Dean</i>	52582
COMRECGC: Global Graph Counterfactual Explainer through Common Recourse <i>Gregoire Fournier, Sourav Medya</i>	52602
All-Purpose Mean Estimation over R: Optimal Sub-Gaussianity with Outlier Robustness and Low Moments Performance <i>Jasper C.H. Lee, Walter McKelvie, Maoyuan Song, Paul Valiant</i>	52622
Evolving Minds: Logic-Informed Inference from Temporal Action Patterns <i>Chao Yang, Shuting Cui, Yang Yang, Shuang Li</i>	52663
Solving Linear-Gaussian Bayesian Inverse Problems with Decoupled Diffusion Sequential Monte Carlo <i>Filip Ekström Kelvinius, Zheng Zhao, Fredrik Lindsten</i>	52687
SAEBench: A Comprehensive Benchmark for Sparse Autoencoders in Language Model Interpretability <i>Adam Karvonen, Can Rager, Johnny Lin, Curt Tigges, Joseph Isaac Bloom, David Chanin, Yeu-Tong Lau, Eoin Farrell, Callum Stuart McDougall, Kola Ayonrinde, Demian Till, Matthew Wearden, Arthur Conmy, Samuel Marks, Neel Nanda</i>	52721
Safe Delta: Consistently Preserving Safety when Fine-Tuning LLMs on Diverse Datasets <i>Ning Lu, Shengcai Liu, Jiahao Wu, Weiyu Chen, Zhirui Zhang, Yew-Soon Ong, Qi Wang, Ke Tang</i>	52763

How Effective Can Dropout Be in Multiple Instance Learning ? <i>Wenhui Zhu, Peijie Qiu, Xiwen Chen, Zhangsihao Yang, Aristeidis Sotiras, Abolfazl Razi, Yalin Wang</i>	52786
"Who experiences large model decay and why?" A Hierarchical Framework for Diagnosing Heterogeneous Performance Drift <i>Harvineet Singh, Fan Xia, Alexej Gossmann, Andrew Chuang, Julian C. Hong, Jean Feng</i>	52803
ALMTOKENIZER: A Low-bitrate and Semantic-rich Audio Codec Tokenizer for Audio Language Modeling <i>Dongchao Yang, Songxiang Liu, Haohan Guo, Jiankun Zhao, Yuanyuan Wang, Helin Wang, Zeqian Ju, Xubo Liu, Xueyuan Chen, Xu Tan, Xixin Wu, Helen M. Meng</i>	52834
Improved and Oracle-Efficient Online ℓ_1-Multicalibration <i>Rohan Ghuge, Vidya Muthukumar, Sahil Singla</i>	52857
Mastering Massive Multi-Task Reinforcement Learning via Mixture-of-Expert Decision Transformer <i>Yilun Kong, Guozheng Ma, Qi Zhao, Haoyu Wang, Li Shen, Xueqian Wang, Dacheng Tao</i>	52878
KBQA-o1: Agentic Knowledge Base Question Answering with Monte Carlo Tree Search <i>Haoran Luo, Haihong E, Yikai Guo, Qika Lin, Xiaobao Wu, Xinyu Mu, Wenhao Liu, Meina Song, Yifan Zhu, Anh Tuan Luu</i>	52896
AutoStep: Locally adaptive involutive MCMC <i>Tiange Liu, Nikola Surjanovic, Miguel Biron-Lattes, Alexandre Bouchard-Cote, Trevor Campbell</i>	52919
Principled Data Selection for Alignment: The Hidden Risks of Difficult Examples <i>Chengqian Gao, Haonan Li, Liu Liu, Zeke Xie, Peilin Zhao, zhiqiang xu</i>	52946
Diverse Prototypical Ensembles Improve Robustness to Subpopulation Shift <i>Nguyen Nhat Minh To, Paul F R Wilson, Viet Nguyen, Mohamed Harmanani, Michael Cooper, Fahimeh Fooladgar, Purang Abolmaesumi, Parvin Mousavi, Rahul Krishnan</i>	52970
Mechanisms of Projective Composition of Diffusion Models <i>Arwen Bradley, Preetum Nakkiran, David Berthelot, James Thornton, Joshua M. Susskind</i>	52993
Blink of an eye: a simple theory for feature localization in generative models <i>Marvin Li, Aayush Karan, Sitan Chen</i>	53029
Generative Human Trajectory Recovery via Embedding-Space Conditional Diffusion <i>Kaijun Liu, Sijie Ruan, Liang Zhang, Cheng Long, Shuliang Wang, liang yu</i>	53063
Componential Prompt-Knowledge Alignment for Domain Incremental Learning <i>Kunlun Xu, Xu Zou, Gang Hua, Jiahuan Zhou</i>	53078
Diverging Preferences: When do Annotators Disagree and do Models Know? <i>Michael JQ Zhang, Zhilin Wang, Jena D. Hwang, Yi Dong, Olivier Delalleau, Yejin Choi, Eunsol Choi, Xiang Ren, Valentina Pyatkin</i>	53093

Overcoming Vocabulary Mismatch: Vocabulary-agnostic Teacher Guided Language Modeling <i>Haebin Shin, Lei Ji, Xiao Liu, Yeyun Gong</i>	53113
Action-Minimization Meets Generative Modeling: Efficient Transition Path Sampling with the Onsager-Machlup Functional <i>Sanjeev Raja, Martin Sipka, Michael Psenka, Tobias Kreiman, Michal Pavelka, Aditi S. Krishnapriyan</i>	53127
REINFORCE Adversarial Attacks on Large Language Models: An Adaptive, Distributional, and Semantic Objective <i>Simon Geisler, Tom Wollschläger, M. H. I. Abdalla, Vincent Cohen-Addad, Johannes Gasteiger, Stephan Günnemann</i>	53164
FedOne: Query-Efficient Federated Learning for Black-box Discrete Prompt Learning <i>Ganyu Wang, Jinjie Fang, Maxwell Juncheng Yin, Bin Gu, Xi Chen, Boyu Wang, Yi Chang, Charles Ling</i>	53194
Beyond Entropy: Region Confidence Proxy for Wild Test-Time Adaptation <i>Zixuan Hu, Yichun Hu, Xiaotong Li, SHIXIANG TANG, LINGYU DUAN</i>	53227
Scalable Attribute-Missing Graph Clustering via Neighborhood Differentiation <i>Yaowenhu, Wenxuan Tu, Yue Liu, Xinhang Wan, Junyi Yan, Taichun Zhou, Xinwang Liu</i>	53247
Synthetic Face Datasets Generation via Latent Space Exploration from Brownian Identity Diffusion <i>David Geissbühler, Hatef Otroschi Shahreza, Sébastien Marcel</i>	53266
On Teacher Hacking in Language Model Distillation <i>Daniil Tiapkin, Daniele Calandriello, Johan Ferret, Sarah Perrin, Nino Vieillard, Alexandre Rame, Mathieu Blondel</i>	53291
Set Valued Predictions For Robust Domain Generalization <i>Ron Tsibulsky, Daniel Nevo, Uri Shalit</i>	53309
Star Attention: Efficient LLM Inference over Long Sequences <i>Shantanu Acharya, Fei Jia, Boris Ginsburg</i>	53344
GEFA: A General Feature Attribution Framework Using Proxy Gradient Estimation <i>Yi Cai, Thibaud Ardoin, Gerhard Wunder</i>	53360
Flow of Reasoning: Training LLMs for Divergent Reasoning with Minimal Examples <i>Fangxu Yu, Lai Jiang, Haoqiang Kang, Shibo Hao, Lianhui Qin</i>	53388
Inverse problems with experiment-guided AlphaFold <i>Sai Advaith Maddipatla, Nadav Bojan, Meital Bojan, Sanketh Vedula, Paul Schanda, Ailie Marx, Alexander Bronstein</i>	53419
R*: Efficient Reward Design via Reward Structure Evolution and Parameter Alignment Optimization with Large Language Models <i>Pengyi Li, Jianye HAO, Hongyao Tang, Yifu Yuan, Jinbin Qiao, Zibin Dong, YAN ZHENG</i>	53447

RollingQ: Reviving the Cooperation Dynamics in Multimodal Transformer <i>HaoTian Ni, Yake Wei, Hang Liu, Gong Chen, Chong Peng, Hao Lin, Di Hu</i>	53466
Learning Smooth and Expressive Interatomic Potentials for Physical Property Prediction <i>Xiang Fu, Brandon M Wood, Luis Barroso-Luque, Daniel S. Levine, Meng Gao, Misko Dzamba, C. Lawrence Zitnick</i>	53483
Editable Concept Bottleneck Models <i>Lijie Hu, Chenyang Ren, Zhengyu Hu, Hongbin Lin, Cheng-Long Wang, Zhen Tan, Weimin Lyu, Jingfeng Zhang, Hui Xiong, Di Wang</i>	53502
Adversarial Inputs for Linear Algebra Backends <i>Jonas Möller, Lukas Pirch, Felix Weissberg, Sebastian Baunsgaard, Thorsten Eisenhofer, Konrad Rieck</i>	53551
Unisolver: PDE-Conditional Transformers Towards Universal Neural PDE Solvers <i>Hang Zhou, Yuezhou Ma, Haixu Wu, Haowen Wang, Mingsheng Long</i>	53563
"Why Is There a Tumor?": Tell Me the Reason, Show Me the Evidence <i>Mengmeng Ma, Tang Li, Yunxiang Peng, Lu Lin, Volkan Beylergil, Binsheng Zhao, Oguz Akin, Xi Peng</i>	53591
Bridging Protein Sequences and Microscopy Images with Unified Diffusion Models <i>Dihan Zheng, Bo Huang</i>	53608
Models of Heavy-Tailed Mechanistic Universality <i>Liam Hodgkinson, Zhichao Wang, Michael W. Mahoney</i>	53625
Bongard in Wonderland: Visual Puzzles that Still Make AI Go Mad? <i>Antonia Wüst, Tim Tobiasch, Lukas Helff, Inga Ibs, Wolfgang Stammer, Devendra Singh Dhami, Constantin A. Rothkopf, Kristian Kersting</i>	53665
Sharp Optimality of Simple, Plug-in Estimation of the Fisher Information of a Smoothed Density <i>Subhodh Kotikal</i>	53690
Generalizable Multi-Camera 3D Object Detection from a Single Source via Fourier Cross-View Learning <i>Xue Zhao, Qinying Gu, Xinbing Wang, Chenghu Zhou, Nanyang Ye</i>	53719
Sum-of-Parts: Self-Attributing Neural Networks with End-to-End Learning of Feature Groups <i>Weiqliu You, Helen Qu, Marco Gatti, Bhuvnesh Jain, Eric Wong</i>	53737
Theoretically Unmasking Inference Attacks Against LDP-Protected Clients in Federated Vision Models <i>Quan Minh Nguyen, Minh N. Vu, Truc Nguyen, My T. Thai</i>	53776
Distributed Conformal Prediction via Message Passing <i>Haifeng Wen, Hong Xing, Osvaldo Simeone</i>	53807

Zero-Shot Generalization of GNNs over Distinct Attribute Domains <i>Yangyi Shen, Jincheng Zhou, Beatrice Bevilacqua, Joshua Robinson, Charilaos Kanatsoulis, Jure Leskovec, Bruno Ribeiro</i>	53826
Approximate Forest Completion and Learning-Augmented Algorithms for Metric Minimum Spanning Trees <i>Nate Veldt, Thomas Stanley, Benjamin W Priest, Trevor Steil, Keita Iwabuchi, T.S. Jayram, Geoffrey Sanders</i>	53857
Adaptive Partitioning Schemes for Optimistic Optimization <i>Raja Sunkara, Ardhendu Tripathy</i>	53881
Advancing Constrained Monotonic Neural Networks: Achieving Universal Approximation Beyond Bounded Activations <i>Davide Sartor, Alberto Sinigaglia, Gian Antonio Susto</i>	53917
Sparse Causal Discovery with Generative Intervention for Unsupervised Graph Domain Adaptation <i>Junyu Luo, Yuhao Tang, yiwei Fu, Xiao Luo, Zhizhuo KOU, Zhiping Xiao, Wei Ju, Wentao Zhang, Ming Zhang</i>	53941
Breaking the Quadratic Barrier: Robust Cardinality Sketches for Adaptive Queries <i>Edith Cohen, Mihir Singhal, Uri Stemmer</i>	53956
EpiCoder: Encompassing Diversity and Complexity in Code Generation <i>Yaoxiang Wang, Haoling Li, Xin Zhang, Jie Wu, Xiao Liu, Wenxiang Hu, Zhongxin Guo, Yangyu Huang, Ying Xin, Yujiu Yang, Jinsong Su, Qi Chen, Scarlett Li</i>	53971
Focus On This, Not That! Steering LLMs with Adaptive Feature Specification <i>Tom A. Lamb, Adam Davies, Alasdair Paren, Philip Torr, Francesco Pinto</i>	54009
Categorical Schrödinger Bridge Matching <i>Grigoriy Ksenofontov, Alexander Korotin</i>	54045
Learning Compact Semantic Information for Incomplete Multi-View Missing Multi-Label Classification <i>Jie Wen, Yadong Liu, Zhanyan Tang, Yuting He, Yulong Chen, Mu Li, Chengliang Liu</i>	54070
No Task Left Behind: Isotropic Model Merging with Common and Task-Specific Subspaces <i>Daniel Marczak, Simone Magistri, Sebastian Cygert, Bartłomiej Twardowski, Andrew D. Bagdanov, Joost van de Weijer</i>	54084
DistiLLM-2: A Contrastive Approach Boosts the Distillation of LLMs <i>Jongwoo Ko, Tianyi Chen, Sungnyun Kim, Tianyu Ding, Luming Liang, Ilya Zharkov, Se-Young Yun</i>	54107
ConceptAttention: Diffusion Transformers Learn Highly Interpretable Features <i>Alec Helbling, Tuna Han Salih Meral, Benjamin Hoover, Pinar Yanardag, Duen Horng Chau</i>	54126

Symmetry-Robust 3D Orientation Estimation <i>Christopher Scarvelis, David Benhaim, Paul Zhang</i>	54144
Randomized Dimensionality Reduction for Euclidean Maximization and Diversity Measures <i>Jie Gao, Rajesh Jayaram, Benedikt Kolbe, Shay Sapir, Chris Schwiegelshohn, Sandeep Silwal, Erik Waingarten</i>	54161
Haste Makes Waste: A Simple Approach for Scaling Graph Neural Networks <i>Rui Xue, Tong Zhao, Neil Shah, Xiaorui Liu</i>	54184
Highly Compressed Tokenizer Can Generate Without Training <i>Lukas Lao Beyer, Tianhong Li, Xinlei Chen, Sertac Karaman, Kaiming He</i>	54207
Improved Coresets for Vertical Federated Learning: Regularized Linear and Logistic Regressions <i>Supratim Shit, Gurmehak kaur chadha, Surendra kumar, Bapi Chatterjee</i>	54226
Expressive Power of Graph Neural Networks for (Mixed-Integer) Quadratic Programs <i>Ziang Chen, Xiaohan Chen, Jialin Liu, Xinshang Wang, Wotao Yin</i>	54248
Tight and Fast Bounds for Multi-Label Learning <i>Yi-Fan Zhang, Min-Ling Zhang</i>	54280
Securing Equal Share: A Principled Approach for Learning Multiplayer Symmetric Games <i>Jiawei Ge, Yuanhao Wang, Wenzhe Li, Chi Jin</i>	54310
MoRAgent: Parameter Efficient Agent Tuning with Mixture-of-Roles <i>Jing Han, Binwei Yan, Tianyu Guo, Zheyuan Bai, Mengyu Zheng, Hanting Chen, Ying Nie</i>	54332
Bivariate Causal Discovery with Proxy Variables: Integral Solving and Beyond <i>Yong Wu, Yanwei Fu, Shouyan Wang, Xinwei Sun</i>	54348
Toward Robust Hyper-Detailed Image Captioning: A Multiagent Approach and Dual Evaluation Metrics for Factuality and Coverage <i>Saehyung Lee, Seunghyun Yoon, Trung Bui, Jing Shi, Sungroh Yoon</i>	54406
Generalizing from SIMPLE to HARD Visual Reasoning: Can We Mitigate Modality Imbalance in VLMs? <i>Simon Park, Abhishek Panigrahi, Yun Cheng, Dingli Yu, Anirudh Goyal, Sanjeev Arora</i>	54424
ReferSplat: Referring Segmentation in 3D Gaussian Splatting <i>Shuting He, Guangquan Jie, Changshuo Wang, Yun Zhou, Shuming Hu, Guanbin Li, Henghui Ding</i>	54477
PISA Experiments: Exploring Physics Post-Training for Video Diffusion Models by Watching Stuff Drop <i>Chenyu Li, Oscar Michel, Xichen Pan, Sainan Liu, Mike Roberts, Saining Xie</i>	54489
Efficient Network Automatic Relevance Determination <i>Hongwei Zhang, Ziqi Ye, Xinyuan Wang, Xin Guo, Zenglin Xu, Yuan Cheng, Zixin Hu, Yuan Qi</i>	54514

An Efficient Search-and-Score Algorithm for Ancestral Graphs using Multivariate Information Scores for Complex Non-linear and Categorical Data <i>Nikita Lagrange, Herve Isambert</i>	54536
SCISSOR: Mitigating Semantic Bias through Cluster-Aware Siamese Networks for Robust Classification <i>Shuo Yang, Bardh Prenkaj, Gjergji Kasneci</i>	54560
Lightspeed Geometric Dataset Distance via Sliced Optimal Transport <i>Khai Nguyen, Hai Nguyen, Tuan Pham, Nhat Ho</i>	54580
AAAR-1.0: Assessing AI's Potential to Assist Research <i>Renze Lou, Hanzi Xu, Sijia Wang, Jiangshu Du, Ryo Kamoi, Xiaoxin Lu, Jian Xie, Yuxuan Sun, Yusen Zhang, Jihyun Janice Ahn, Hongchao Fang, Zhuoyang Zou, Wenchao Ma, Xi Li, Kai Zhang, Congying Xia, Lifu Huang, Wenpeng Yin</i>	54596
Capturing Temporal Dynamics in Large-Scale Canopy Tree Height Estimation <i>Jan Pauls, Max Zimmer, Berkant Turan, Sassan Saatchi, Philippe CIAIS, Sebastian Pokutta, Fabian Gieseke</i>	54619
One Diffusion Step to Real-World Super-Resolution via Flow Trajectory Distillation <i>Jianze Li, Jiezhong Cao, Yong Guo, Wenbo Li, Yulun Zhang</i>	54636
Observation Interference in Partially Observable Assistance Games <i>Scott Emmons, Caspar Oesterheld, Vincent Conitzer, Stuart Russell</i>	54646
Toward Efficient Kernel-Based Solvers for Nonlinear PDEs <i>Zhitong Xu, Da Long, Yiming Xu, Guang Yang, Shandian Zhe, Houman Owhadi</i>	54672
Variational Rectified Flow Matching <i>Pengsheng Guo, Alex Schwing</i>	54691
Schwarz–Schur Involution: Lightspeed Differentiable Sparse Linear Solvers <i>Yu Wang, Mazdak Abulnaga, Yaël Balbastre, Bruce Fischl</i>	54711
Prediction-Powered E-Values <i>Daniel Csillag, Claudio Jose Struchiner, Guilherme Tegoni Goedert</i>	54761
Ringmaster ASGD: The First Asynchronous SGD with Optimal Time Complexity <i>Arto Maranjyan, Alexander Tyurin, Peter Richtárik</i>	54783
Compositional Generalization via Forced Rendering of Disentangled Latents <i>Qiyao Liang, Daoyuan Qian, Liu Ziyin, Ila R Fiete</i>	54803
SelfCite: Self-Supervised Alignment for Context Attribution in Large Language Models <i>Yung-Sung Chuang, Benjamin Cohen-Wang, Zejiang Shen, Zhaofeng Wu, Hu Xu, Xi Victoria Lin, James R. Glass, Shang-Wen Li, Wen-tau Yih</i>	54829

Can Compressed LLMs Truly Act? An Empirical Evaluation of Agentic Capabilities in LLM Compression <i>Peijie Dong, Zhenheng Tang, Xiang Liu, Lujun Li, Xiaowen Chu, Bo Li</i>	54849
RATE: Causal Explainability of Reward Models with Imperfect Counterfactuals <i>David Reber, Sean M Richardson, Todd Nief, Cristina Garbacea, Victor Veitch</i>	54883
Features are fate: a theory of transfer learning in high-dimensional regression <i>Javan Tahir, Surya Ganguli, Grant M. Rotskoff</i>	54911
H-Tuning: Toward Low-Cost and Efficient ECG-based Cardiovascular Disease Detection with Pre-Trained Models <i>Rushuang Zhou, Yuanting Zhang, Yining Dong</i>	54938
Improved Algorithm for Deep Active Learning under Imbalance via Optimal Separation <i>Shyam Nuggehalli, Jifan Zhang, Lalit K Jain, Robert D Nowak</i>	54960
Large Displacement Motion Transfer with Unsupervised Anytime Interpolation <i>Guixiang Wang, Jianjun Li</i>	54982
Instruct2See: Learning to Remove Any Obstructions Across Distributions <i>Junhang Li, Yu Guo, Chuhua XIAN, Shengfeng He</i>	54992
Optimizing Temperature for Language Models with Multi-Sample Inference <i>Weihua Du, Yiming Yang, Sean Welleck</i>	55010
Almost Optimal Fully Dynamic k-Center Clustering with Recourse <i>Sayan Bhattacharya, Martin Costa, Ermiya Farokhnejad, Silvio Lattanzi, Nikos Parotsidis</i>	55031
Learning Attribute-Aware Hash Codes for Fine-Grained Image Retrieval via Query Optimization <i>Peng Wang, Yong Li, Lin Zhao, Xiu-Shen Wei</i>	55045
Are Sparse Autoencoders Useful? A Case Study in Sparse Probing <i>Subhash Kantamneni, Joshua Engels, Senthoran Rajamanoharan, Max Tegmark, Neel Nanda</i>	55059
Understanding Input Selectivity in Mamba: Impact on Approximation Power, Memorization, and Associative Recall Capacity <i>Ningyuan Teresa Huang, Miguel Sarabia, Abhinav Moudgil, Pau Rodriguez, Luca Zappella, Federico Danieli</i>	55091
Universal Length Generalization with Turing Programs <i>Kaiying Hou, David Brandfonbrener, Sham M. Kakade, Samy Jelassi, Eran Malach</i>	55126
HiRemate: Hierarchical Approach for Efficient Re-materialization of Neural Networks <i>Julia Gusak, Xunyi Zhao, Théotime Le Hellard, Zhe Li, Lionel Eyraud-Dubois, Olivier Beaumont</i>	55147
PINNsAgent: Automated PDE Surrogation with Large Language Models <i>Qingpo Wuwu, Chonghan Gao, Tianyu Chen, Yihang Huang, Yuekai Zhang, Jianing Wang, Jianxin Li, Haoyi Zhou, Shanghang Zhang</i>	55173

Customizing the Inductive Biases of Softmax Attention using Structured Matrices <i>Yilun Kuang, Noah Amsel, Sanae Lotfi, Shikai Qiu, Andres Potapczynski, Andrew Gordon Wilson</i>	55196
Neural Representational Consistency Emerges from Probabilistic Neural-Behavioral Representation Alignment <i>Yu Zhu, Chunfeng Song, Wanli Ouyang, Shan Yu, Tiejun Huang</i>	55221
A Market for Accuracy: Classification Under Competition <i>Ohad Einav, Nir Rosenfeld</i>	55248
A Sample Efficient Conditional Independence Test in the Presence of Discretization <i>Boyang Sun, Yu Yao, Xinshuai Dong, Zongfang Liu, Tongliang Liu, Yumou Qiu, Kun Zhang</i>	55274
Analytical Lyapunov Function Discovery: An RL-based Generative Approach <i>Haohan Zou, Jie Feng, Hao Zhao, Yuanyuan Shi</i>	55300
Temperature-Annealed Boltzmann Generators <i>Henrik Schopmans, Pascal Friederich</i>	55329
From Debate to Equilibrium: Belief-Driven Multi-Agent LLM Reasoning via Bayesian Nash Equilibrium <i>Xie Yi, Zhanke Zhou, Chentao Cao, Qiyu Niu, Tongliang Liu, Bo Han</i>	55363
Residual Matrix Transformers: Scaling the Size of the Residual Stream <i>Brian Mak, Jeffrey Flanigan</i>	55403
Bayesian Weight Enhancement with Steady-State Adaptation for Test-time Adaptation in Dynamic Environments <i>Jae-Hong Lee</i>	55421
Constant Stepsize Local GD for Logistic Regression: Acceleration by Instability <i>Michael Crawshaw, Blake Woodworth, Mingrui Liu</i>	55439
ViTally Consistent: Scaling Biological Representation Learning for Cell Microscopy <i>Kian Kenyon-Dean, Zitong Jerry Wang, John Urbanik, Konstantin Donhauser, Jason Hartford, Saber Saberian, Nil Sahin, Ihab Bendif, Safiye Celik, Juan Sebastián Rodríguez Vera, Marta Fay, Imran S Haque, Oren Kraus</i>	55467
Vision Graph Prompting via Semantic Low-Rank Decomposition <i>Zixiang Ai, Zichen Liu, Jiahuan Zhou</i>	55485
Improving Diversity in Language Models: When Temperature Fails, Change the Loss <i>Alexandre Verine, Florian Le Bronnec, Kunhao Zheng, Alexandre Allauzen, Yann Chevalere, benjamin negrevergne</i>	55498
N2GON: Neural Networks for Graph-of-Net with Position Awareness <i>Yejiang Wang, Yuhai Zhao, Zhengkui Wang, Wen Shan, Ling Li, Qian Li, Miaomiao Huang, Meixia Wang, Shirui Pan, Xingwei Wang</i>	55533

Multi-objective Linear Reinforcement Learning with Lexicographic Rewards <i>Bo Xue, Dake Bu, Ji Cheng, Yuanyu Wan, Qingfu Zhang</i>	55545
High Probability Bound for Cross-Learning Contextual Bandits with Unknown Context Distributions <i>Ruiyuan Huang, Zengfeng Huang</i>	55566
LAuReL: Learned Augmented Residual Layer <i>Gaurav Menghani, Ravi Kumar, Sanjiv Kumar</i>	55593
Adapting to Evolving Adversaries with Regularized Continual Robust Training <i>Sihui Dai, Christian Cianfarani, Vikash Sehwag, Prateek Mittal, Arjun Bhagoji</i>	55604
Enhancing Visual Localization with Cross-Domain Image Generation <i>Yuanze Wang, Yichao Yan, Shiming Song, Songchang Jin, Yilan Huang, Xingdong Sheng, Dianxi Shi</i>	55651
Fast Inference with Kronecker-Sparse Matrices <i>Antoine Gonon, Léon Zheng, Pascal Carrivain, TUNG QUOC LE</i>	55667
Implicit Bias of Gradient Descent for Non-Homogeneous Deep Networks <i>Yuhang Cai, Kangjie Zhou, Jingfeng Wu, Song Mei, Michael Lindsey, Peter Bartlett</i>	55695
Reward-Augmented Data Enhances Direct Preference Alignment of LLMs <i>Shenao Zhang, Zhihan Liu, Boyi Liu, Yufeng Zhang, Yingxiang Yang, Yongfei Liu, Liyu Chen, Tao Sun, Zhaoran Wang</i>	55774
MASS: Mathematical Data Selection via Skill Graphs for Pretraining Large Language Models <i>Jiazheng Li, Lu Yu, Qing Cui, Zhiqiang Zhang, JUN ZHOU, Yanfang Ye, Chuxu Zhang</i>	55798
Generalized Interpolating Discrete Diffusion <i>Dimitri von Rütte, Janis Fluri, Yuhui Ding, Antonio Orvieto, Bernhard Schölkopf, Thomas Hofmann</i>	55819
Adapting to Linear Separable Subsets with Large-Margin in Differentially Private Learning <i>Erchi Wang, Yuqing Zhu, Yu-Xiang Wang</i>	55853
TeDS: Joint Learning of Diachronic and Synchronic Perspectives in Quaternion Space for Temporal Knowledge Graph Completion <i>Jiujiang Guo, Mankun Zhao, Wenbin Zhang, Tianyi Xu, Linying Xu, Yu Jian, Yu Mei, Yu Ruiguo</i>	55883
DSBRouter: End-to-end Global Routing via Diffusion Schrödinger Bridge <i>Liangliang Shi, Shenhui Zhang, Xingbo Du, Nianzu Yang, Junchi Yan</i>	55903
Chaos Meets Attention: Transformers for Large-Scale Dynamical Prediction <i>Yi He, Yiming Yang, Xiaoyuan Cheng, Hai Wang, Xiao Xue, Boli Chen, Yukun Hu</i>	55920
GRAM: A Generative Foundation Reward Model for Reward Generalization <i>Chenglong Wang, Yang Gan, Yifu Huo, Yongyu Mu, Qiaozhi He, MuRun Yang, Bei Li, Tong Xiao, Chunliang Zhang, Tongran Liu, JingBo Zhu</i>	55945

Puzzle: Distillation-Based NAS for Inference-Optimized LLMs <i>Akhiad Bercovich, Tomer Ronen, Talor Abramovich, Nir Ailon, Nave Assaf, Mohammed Dabbah, Ido Galil, Amnon Geifman, Yonatan Geifman, Izhak Golan, Netanel Haber, Ehud Dov Karpas, Roi Koren, Itay Levy, Pavlo Molchanov, Shahar Mor, Zach Moshe, Najeeb Nabwani, Omri Puny, Ran Rubin, Itamar Schen, Ido Shahaf, Oren Tropp, Omer Ullman Argov, Ran Zilberstein, Ran El-Yaniv</i>	55966
Towards the Efficient Inference by Incorporating Automated Computational Phenotypes under Covariate Shift <i>chao ying, Jun Jin, Yi Guo, Xiudi Li, Muxuan Liang, Jiwei Zhao</i>	55991
RocketKV: Accelerating Long-Context LLM Inference via Two-Stage KV Cache Compression <i>Payman Behnam, Yaosheng Fu, Ritchie Zhao, Po-An Tsai, Zhiding Yu, Alexey Tumanov</i>	56021
Efficient Long Context Fine-tuning with Chunk Flow <i>Xiulong Yuan, Hongtao Xu, Wenting Shen, Ang Wang, Xiafei Qiu, Jie Zhang, Yuqiong Liu, Bowen Yu, Junyang Lin, Mingzhen Li, Weile Jia, Yong Li, Wei Lin</i>	56056
Dual Feature Reduction for the Sparse-group Lasso and its Adaptive Variant <i>Fabio Feser, Marina Evangelou</i>	56067
Efficient Diffusion Models for Symmetric Manifolds <i>Oren Mangoubi, Neil He, Nisheeth K. Vishnoi</i>	56099
ZipAR: Parallel Autoregressive Image Generation through Spatial Locality <i>Yefei He, Feng Chen, Yuanyu He, Shaoxuan He, Hong Zhou, Kaipeng Zhang, Bohan Zhuang</i>	56136
Learning Adversarial MDPs with Stochastic Hard Constraints <i>Francesco Emanuele Stradi, Matteo Castiglioni, Alberto Marchesi, Nicola Gatti</i>	56147
Targeted Low-rank Refinement: Enhancing Sparse Language Models with Precision <i>Li Shen, Anke Tang, Yong Luo, Tao Sun, Han Hu, Xiaochun Cao</i>	56179
Stable Offline Value Function Learning with Bisimulation-based Representations <i>Brahma S Pavse, Yudong Chen, Qiaomin Xie, Josiah P. Hanna</i>	56198
Freeze-Omni: A Smart and Low Latency Speech-to-speech Dialogue Model with Frozen LLM <i>Xiong Wang, Yangze Li, Chaoyou Fu, Yike Zhang, Yunhang Shen, Lei Xie, Ke Li, Xing Sun, Long MA</i>	56228
Variational Phylogenetic Inference with Products over Bipartitions <i>Evan Sidrow, Alexandre Bouchard-Cote, Lloyd T Elliott</i>	56238
Selective Preference Aggregation <i>Shreyas Kadekodi, Hayden McTavish, Berk Ustun</i>	56261
SAE-V: Interpreting Multimodal Models for Enhanced Alignment <i>Hantao Lou, Changye Li, Jiaming Ji, Yaodong Yang</i>	56287
Softmax is not Enough (for Sharp Size Generalisation) <i>Petar Veličković, Christos Perivolaropoulos, Federico Barbero, Razvan Pascanu</i>	56304

Robust Conformal Outlier Detection under Contaminated Reference Data <i>Meshi Bashari, Matteo Sesia, Yaniv Romano</i>	56326
TopoTune: A Framework for Generalized Combinatorial Complex Neural Networks <i>Mathilde Papillon, Guillermo Bernardez, Claudio Battiloro, Nina Miolane</i>	56377
TimePro: Efficient Multivariate Long-term Time Series Forecasting with Variable- and Time-Aware Hyper-state <i>Xiaowen Ma, Zhen-Liang Ni, Shuai Xiao, Xinghao Chen</i>	56406
Agent Reviewers: Domain-specific Multimodal Agents with Shared Memory for Paper Review <i>Kai Lu, Shixiong Xu, Jinqiu Li, Kun Ding, Gaofeng Meng</i>	56422
AutoEval Done Right: Using Synthetic Data for Model Evaluation <i>Pierre Boyeau, Anastasios Nikolas Angelopoulos, Tianle Li, Nir Yosef, Jitendra Malik, Michael I. Jordan</i>	56450
The Role of Sparsity for Length Generalization in LLMs <i>Noah Golowich, Samy Jelassi, David Brandfonbrener, Sham M. Kakade, Eran Malach</i>	56465
Directly Forecasting Belief for Reinforcement Learning with Delays <i>Qingyuan Wu, Yuhui Wang, Simon Sinong Zhan, Yixuan Wang, Chung-Wei Lin, Chen Lv, Qi Zhu, Jürgen Schmidhuber, Chao Huang</i>	56497
Finding Wasserstein Ball Center: Efficient Algorithm and The Applications in Fairness <i>Yuntao Wang, Yuxuan Li, Qingyuan Yang, Hu Ding</i>	56520
Fairness Overfitting in Machine Learning: An Information-Theoretic Perspective <i>Firas Laakom, Haobo Chen, Jürgen Schmidhuber, Yuheng Bu</i>	56546
Provable Efficiency of Guidance in Diffusion Models for General Data Distribution <i>Gen Li, Yuchen Jiao</i>	56584
Time Series Representations with Hard-Coded Invariances <i>Thibaut Germain, Chrysoula Kosma, Laurent Oudre</i>	56597
Conformal Prediction with Cellwise Outliers: A Detect-then-Impute Approach <i>Qian Peng, Yajie Bao, Haojie Ren, Zhaojun Wang, Changliang Zou</i>	56621
CommVQ: Commutative Vector Quantization for KV Cache Compression <i>Junyan Li, Yang Zhang, Muhammad Yusuf Hassan, Talha Chafekar, Tianle Cai, Zhile Ren, Pengsheng Guo, Foroozan Karimzadeh, Colorado Reed, Chong Wang, Chuang Gan</i>	56645
Sparsing Law: Towards Large Language Models with Greater Activation Sparsity <i>Yuqi Luo, Chenyang Song, Xu Han, Yingfa Chen, Chaojun Xiao, Xiaojun Meng, Liqun Deng, Jiansheng Wei, Zhiyuan Liu, Maosong Sun</i>	56660

DriveGPT: Scaling Autoregressive Behavior Models for Driving <i>Xin Huang, Eric M Wolff, Paul Vernaza, Tung Phan-Minh, Hongge Chen, David S Hayden, Mark Edmonds, Brian Pierce, Xinxin Chen, Pratik Elias Jacob, Xiaobai Chen, Chingiz Tairbekov, Pratik Agarwal, Tianshi Gao, Yuning Chai, Siddhartha Srinivasa</i>	56680
SECOND: Mitigating Perceptual Hallucination in Vision-Language Models via Selective and Contrastive Decoding <i>Woohyeon Park, Woojin Kim, Jaeik Kim, Jaeyoung Do</i>	56694
A Theoretical Framework For Overfitting In Energy-based Modeling <i>Giovanni Catania, Aurélien Decelle, Cyril Furtlehner, Beatriz Seoane</i>	56708
Breaking Silos: Adaptive Model Fusion Unlocks Better Time Series Forecasting <i>Zhining Liu, Ze Yang, Xiao Lin, Ruizhong Qiu, Tianxin Wei, Yada Zhu, Hendrik Hamann, Jingrui He, Hanghang Tong</i>	56737
Multi-Objective Causal Bayesian Optimization <i>Shriya Bhatija, Paul-David Zuercher, Jakob Thumm, Thomas Bohné</i>	56758
LaCache: Ladder-Shaped KV Caching for Efficient Long-Context Modeling of Large Language Models <i>Dachuan Shi, Yonggan Fu, Xiangchi Yuan, Zhongzhi Yu, Haoran You, Sixu Li, Xin Dong, Jan Kautz, Pavlo Molchanov, Yingyan Celine Lin</i>	56782
Deep Linear Network Training Dynamics from Random Initialization: Data, Width, Depth, and Hyperparameter Transfer <i>Blake Bordelon, Cengiz Pehlevan</i>	56794
Provably Efficient RL for Linear MDPs under Instantaneous Safety Constraints in Non-Convex Feature Spaces <i>Amirhossein Roknilamouki, Arnob Ghosh, Ming Shi, Fatemeh Nourzad, Eylem Ekici, Ness Shroff</i>	56824
Analyze Feature Flow to Enhance Interpretation and Steering in Language Models <i>Daniil Laptev, Nikita Balagansky, Yaroslav Aksenov, Daniil Gavrilov</i>	56863
TraceGrad: a Framework Learning Expressive SO(3)-equivariant Non-linear Representations for Electronic-Structure Hamiltonian Prediction <i>Shi Yin, Xinyang Pan, Fengyan Wang, Lixin He</i>	56887
Selective Response Strategies for GenAI <i>Boaz Taitler, Omer Ben-Porat</i>	56916
Diversifying Policy Behaviors with Extrinsic Behavioral Curiosity <i>Zhenglin Wan, Xingrui Yu, David Mark Bossens, Yueming Lyu, Qing Guo, Flint Xiaofeng Fan, Yew-Soon Ong, Ivor Tsang</i>	56954
SPHINX: Structural Prediction using Hypergraph Inference Network <i>Iulia Duta, Pietro Lio</i>	56974

Simple Policy Optimization <i>Zhengpeng Xie, Qiang Zhang, Fan Yang, Marco Hutter, Renjing Xu</i>	56992
RepLoRA: Reparameterizing Low-rank Adaptation via the Perspective of Mixture of Experts <i>Tuan Truong, Chau Nguyen, Huy Nguyen, Minh Le, Trung Le, Nhat Ho</i>	57004
Strong and Weak Identifiability of Optimization-based Causal Discovery in Non-linear Additive Noise Models <i>Mingjia Li, Hong Qian, Tian-Zuo Wang, ShujunLi, Min Zhang, Aimin Zhou</i>	57039
Differential Privacy Under Class Imbalance: Methods and Empirical Insights <i>Lucas Rosenblatt, Yuliia Lut, Ethan Turok, Marco Avella Medina, Rachel Cummings</i>	57055
STD-FD: Spatio-Temporal Distribution Fitting Deviation for AIGC Forgery Identification <i>Hengrui Lou, Zunlei Feng, Jinsong Geng, Erteng Liu, Jie Lei, Lechao Cheng, Jie Song, Mingli Song, Yijun Bei</i>	57100
Multiaccuracy and Multicalibration via Proxy Groups <i>Beepul Bharti, Mary Versa Clemens-Sewall, Paul Yi, Jeremias Sulam</i>	57116
Contrastive Localized Language-Image Pre-Training <i>Hong-You Chen, Zhengfeng Lai, Haotian Zhang, Xinze Wang, Marcin Eichner, Keen You, Meng Cao, Bowen Zhang, Yinfei Yang, Zhe Gan</i>	57139
Fast Large Language Model Collaborative Decoding via Speculation <i>Jiale Fu, Yuchu Jiang, Junkai Chen, Jiaming Fan, Xin Geng, Xu Yang</i>	57156
Evaluating LLMs Across Multi-Cognitive Levels: From Medical Knowledge Mastery to Scenario-Based Problem Solving <i>Yuxuan Zhou, Xien Liu, Chenwei Yan, Chen Ning, Xiao Zhang, Boxun Li, Xiangling Fu, Shijin Wang, Guoping Hu, Yu Wang, Ji Wu</i>	57175
Task-Agnostic Pre-training and Task-Guided Fine-tuning for Versatile Diffusion Planner <i>Chenyou Fan, Chenjia Bai, Zhao Shan, Haoran He, Yang Zhang, Zhen Wang</i>	57195
KIND: Knowledge Integration and Diversion for Training Decomposable Models <i>Yucheng Xie, Fu Feng, Ruixiao Shi, Jing Wang, Yong Rui, Xin Geng</i>	57218
Simplifying DINO via Coding Rate Regularization <i>Ziyang Wu, Jingyuan Zhang, Druv Pai, XuDong Wang, Chandan Singh, Jianwei Yang, Jianfeng Gao, Yi Ma</i>	57238
Update Your Transformer to the Latest Release: Re-Basin of Task Vectors <i>Filippo Rinaldi, Giacomo Capitani, Lorenzo Bonicelli, Donato Crisostomi, Federico Bolelli, ELISA FICARRA, Emanuele Rodolà, Simone Calderara, Angelo Porrello</i>	57262
Boosting Multi-Domain Fine-Tuning of Large Language Models through Evolving Interactions between Samples <i>Xize Liang, Lin Yang, Jie Wang, Yiyang Lu, Runyu Wu, Hanzhu Chen, Jianye HAO</i>	57277

Gamma Distribution PCA-Enhanced Feature Learning for Angle-Robust SAR Target Recognition <i>Chong Zhang, Peng Zhang, Mengke Li</i>	57292
Synthesizing Images on Perceptual Boundaries of ANNs for Uncovering and Manipulating Human Perceptual Variability <i>Chen Wei, Chi Zhang, Jiachen Zou, Haotian Deng, Dietmar Heinke, Quanying Liu</i>	57309
Breaking the Barrier of Hard Samples: A Data-Centric Approach to Synthetic Data for Medical Tasks <i>MAYNARA DONATO DE SOUZA, Cleber Zanchettin</i>	57345
Propagation of Chaos for Mean-Field Langevin Dynamics and its Application to Model Ensemble <i>Atsushi Nitanda, Anzelle Lee, Damian Tan Xing Kai, Mizuki Sakaguchi, Taiji Suzuki</i>	57399
Does Graph Prompt Work? A Data Operation Perspective with Theoretical Analysis <i>Qunzhong Wang, Xiangguo Sun, Hong Cheng</i>	57424
Plausible Token Amplification for Improving Accuracy of Differentially Private In-Context Learning Based on Implicit Bayesian Inference <i>Yusuke Yamasaki, Kenta Niwa, Daiki Chijiwa, Takumi Fukami, Takayuki Miura</i>	57450
Nearly Optimal Sample Complexity for Learning with Label Proportions <i>Robert Istvan Busa-Fekete, Travis Dick, Claudio Gentile, Haim Kaplan, Tomer Koren, Uri Stemmer</i>	57494
Delay-DSGN: A Dynamic Spiking Graph Neural Network with Delay Mechanisms for Evolving Graph <i>Zhiqiang Wang, Jianghao Wen, Jianqing Liang</i>	57528
The Perils of Optimizing Learned Reward Functions: Low Training Error Does Not Guarantee Low Regret <i>Lukas Fluri, Leon Lang, Alessandro Abate, Patrick Forré, David Krueger, Joar Max Viktor Skalse</i>	57543
Are High-Quality AI-Generated Images More Difficult for Models to Detect? <i>Yao Xiao, Binbin Yang, Weiyan Chen, Jiahao Chen, Zijie Cao, ZiYi Dong, Xiangyang Ji, Liang Lin, Wei Ke, Pengxu Wei</i>	57615
DANCE: Dual Unbiased Expansion with Group-acquired Alignment for Out-of-distribution Graph Fairness Learning <i>Yifan Wang, Hourun Li, Yue Ling, Zhiping Xiao, Jia Yang, Changling Zhou, Wei Ju, Ming Zhang, Xiao Luo</i>	57637
Policy Optimization for CMDPs with Bandit Feedback: Learning Stochastic and Adversarial Constraints <i>Francesco Emanuele Stradi, Anna Lunghi, Matteo Castiglioni, Alberto Marchesi, Nicola Gatti</i>	57651
Off-Policy Evaluation under Nonignorable Missing Data <i>Han Wang, Yang Xu, Wenbin Lu, Rui Song</i>	57689

Calibrated Language Models and How to Find Them with Label Smoothing <i>Jerry Huang, Peng Lu, QIUHAO Zeng</i>	57728
On the Vulnerability of Applying Retrieval-Augmented Generation within Knowledge-Intensive Application Domains <i>Xun Xian, Ganghua Wang, Xuan Bi, Rui Zhang, Jayanth Srinivasa, Ashish Kundu, Charles Fleming, Mingyi Hong, Jie Ding</i>	57748
SCENIR: Visual Semantic Clarity through Unsupervised Scene Graph Retrieval <i>Nikolaos Chaidos, Angeliki Dimitriou, Maria Lymperaiou, Giorgos Stamou</i>	57772
Learning Distances from Data with Normalizing Flows and Score Matching <i>Peter Sorrenson, Daniel Behrend-Urriarte, Christoph Schnoerr, Ullrich Koethe</i>	57790
Robust Multimodal Large Language Models Against Modality Conflict <i>Zongmeng Zhang, Wengang Zhou, Jie Zhao, Houqiang Li</i>	57808
Simultaneous Multi-Robot Motion Planning with Projected Diffusion Models <i>Jinhao Liang, Jacob K Christopher, Sven Koenig, Ferdinando Fioretto</i>	57829
LIVS: A Pluralistic Alignment Dataset for Inclusive Public Spaces <i>Rashid Mushkani, Shravan Nayak, Hugo Berard, Allison Cohen, Shin Koseki, Hadrien Bertrand</i>	57848
Improving Continual Learning Performance and Efficiency with Auxiliary Classifiers <i>Filip Szatkowski, Yaoyue Zheng, Fei Yang, Tomasz Trzcinski, Bartłomiej Twardowski, Joost van de Weijer</i>	57879
Addressing Concept Mislabeling in Concept Bottleneck Models Through Preference Optimization <i>Emiliano Penalzoa, Tianyue H. Zhang, Laurent Charlin, Mateo Espinosa Zarlenga</i>	57915
Regret-Free Reinforcement Learning for Temporal Logic Specifications <i>R Majumdar, Mahmoud Salamati, Sadegh Soudjani</i>	57942
Learning Condensed Graph via Differentiable Atom Mapping for Reaction Yield Prediction <i>Ankit Ghosh, Gargee Kashyap, Sarthak Mittal, Nupur Jain, Raghavan B Sunoj, Abir De</i>	57963
Towards Theoretical Understanding of Sequential Decision Making with Preference Feedback <i>Simone Drago, Marco Mussi, Alberto Maria Metelli</i>	57994
Moirai-MoE: Empowering Time Series Foundation Models with Sparse Mixture of Experts <i>Xu Liu, Juncheng Liu, Gerald Woo, Taha Aksu, Yuxuan Liang, Roger Zimmermann, Chenghao Liu, Junnan Li, Silvio Savarese, Caiming Xiong, Doyen Sahoo</i>	58010
A Physics-Informed Machine Learning Framework for Safe and Optimal Control of Autonomous Systems <i>Manan Tayal, Aditya Singh, Shishir Kolathaya, Somil Bansal</i>	58033
Memorization Sinks: Isolating Memorization during LLM Training <i>Gaurav Rohit Ghosal, Pratyush Maini, Aditi Raghunathan</i>	58055

Mind the Gap: a Spectral Analysis of Rank Collapse and Signal Propagation in Attention Layers <i>Thiziri Nait Saada, Alireza Naderi, Jared Tanner</i>	58075
Distributed Parallel Gradient Stacking(DPGS): Solving Whole Slide Image Stacking Challenge in Multi-Instance Learning <i>Boyuan Wu, ZEFENG WANG, Xianwei Lin, Jiachun Xu, Jikai Yu, Zhou Shicheng, Hongda Chen, Lianxin Hu</i>	58102
When Bad Data Leads to Good Models <i>Kenneth Li, Yida Chen, Fernanda Viégas, Martin Wattenberg</i>	58113
MIB: A Mechanistic Interpretability Benchmark <i>Aaron Mueller, Atticus Geiger, Sarah Wiegreffe, Dana Arad, Iván Arcuschin, Adam Belfki, Yik Siu Chan, Jaden Fried Fiotto-Kaufman, Tal Haklay, Michael Hanna, Jing Huang, Rohan Gupta, Yaniv Nikankin, Hadas Orgad, Nikhil Prakash, Anja Reusch, Aruna Sankaranarayanan, Shun Shao, Alessandro Stolfo, Martin Tutek, Amir Zur, David Bau, Yonatan Belinkov</i>	58127
ZebraLogic: On the Scaling Limits of LLMs for Logical Reasoning <i>Bill Yuchen Lin, Ronan Le Bras, Kyle Richardson, Ashish Sabharwal, Radha Poovendran, Peter Clark, Yejin Choi</i>	58167
Tree-Sliced Wasserstein Distance: A Geometric Perspective <i>Hoang V. Tran, Huyen Trang Pham, Tho Tran Huu, Minh-Khoi Nguyen-Nhat, Thanh Chu, Tam Le, Tan Minh Nguyen</i>	58184
A Multi-Region Brain Model to Elucidate the Role of Hippocampus in Spatially Embedded Decision-Making <i>Yi Xie, Jaedong Hwang, Carlos D Brody, David W. Tank, Ila R Fiete</i>	58216
Rectifying Conformity Scores for Better Conditional Coverage <i>Vincent Plassier, Alexander Fishkov, Victor Dheur, Mohsen Guizani, Souhaib Ben Taieb, Maxim Panov, Eric Moulines</i>	58241
A Sharper Global Convergence Analysis for Average Reward Reinforcement Learning via an Actor-Critic Approach <i>Swetha Ganesh, Washim Uddin Mondal, Vaneet Aggarwal</i>	58275
Dynamic Similarity Graph Construction with Kernel Density Estimation <i>Steinar Laenen, Peter Macgregor, He Sun</i>	58297
Predicting the Susceptibility of Examples to Catastrophic Forgetting <i>Guy Hacohen, Tinne Tuytelaars</i>	58331
Demonstration Selection for In-Context Learning via Reinforcement Learning <i>Xubin Wang, Jianfei Wu, Yuan Yichen, Deyu Cai, Mingzhe Li, Weijia Jia</i>	58355
Aggregation Buffer: Revisiting DropEdge with a New Parameter Block <i>Dooho Lee, Myeong Kong, Sagad Hamid, Cheonwoo Lee, Jaemin Yoo</i>	58373

Stochastic Poisson Surface Reconstruction with One Solve using Geometric Gaussian Processes <i>Sidhanth Holalkere, David Bindel, Silvia Sellán, Alexander Terenin</i>	58397
Make LoRA Great Again: Boosting LoRA with Adaptive Singular Values and Mixture-of-Experts Optimization Alignment <i>Chenghao Fan, Zhenyi Lu, Sichen Liu, Chengfeng Gu, Xiaoye Qu, Wei Wei, Yu Cheng</i>	58412
Dialogue Without Limits: Constant-Sized KV Caches for Extended Response in LLMs <i>Ravi Ghadia, Avinash Kumar, Gaurav Jain, Prashant J. Nair, Poulami Das</i>	58441
Solving Probabilistic Verification Problems of Neural Networks using Branch and Bound <i>David Boetius, Stefan Leue, Tobias Sutter</i>	58456
On the Power of Learning-Augmented Search Trees <i>Jingbang Chen, Xinyuan Cao, Alicia Stepin, Li Chen</i>	58496
ResKoopNet: Learning Koopman Representations for Complex Dynamics with Spectral Residuals <i>Yuanchao Xu, Kaidi Shao, Nikos K. Logothetis, Zhongwei Shen</i>	58520
DockS-RAG: Optimizing Document-Level Relation Extraction through LLM-Enhanced Hybrid Prompt Tuning <i>Xiaolong Xu, Yibo Zhou, Haolong Xiang, Xiaoyong Li, Xuyun Zhang, Lianyong Qi, Wanchun Dou</i>	58548
SNS-Bench: Defining, Building, and Assessing Capabilities of Large Language Models in Social Networking Services <i>Hongcheng Guo, wangyue, Shaosheng Cao, Fei zhao, Boyang Wang, Lei Li, Liang Chen, Xinze Lyu, Zhe Xu, Yao Hu, Zhoujun Li</i>	58562
Generative Audio Language Modeling with Continuous-valued Tokens and Masked Next-Token Prediction <i>Shu-wen Yang, Byeonggeun Kim, Kuan-Po Huang, Qingming Tang, HUY PHAN, Bo-Ru Lu, Harshavardhan Sundar, Shalini Ghosh, Hung-yi Lee, Chieh-Chi Kao, Chao Wang</i>	58599
The Ripple Effect: On Unforeseen Complications of Backdoor Attacks <i>Rui Zhang, Yun Shen, Hongwei Li, Wenbo Jiang, Hanxiao Chen, Yuan Zhang, Guowen Xu, Yang Zhang</i>	58619
Bayesian Inference for Correlated Human Experts and Classifiers <i>Markelle Kelly, Alex James Boyd, Sam Showalter, Mark Steyvers, Padhraic Smyth</i>	58647
Concept-Centric Token Interpretation for Vector-Quantized Generative Models <i>Tianze Yang, Yucheng Shi, Mengnan Du, Xuansheng Wu, Qiaoyu Tan, Jin Sun, Ninghao Liu</i>	58664
Hierarchical Masked Autoregressive Models with Low-Resolution Token Pivots <i>Guangting Zheng, Yehao Li, Yingwei Pan, Jiajun Deng, Ting Yao, Yanyong Zhang, Tao Mei</i>	58681
Tightening Causal Bounds via Covariate-Aware Optimal Transport <i>Sirui Lin, Zijun Gao, Jose Blanchet, Peter Glynn</i>	58694

Tree-Sliced Wasserstein Distance with Nonlinear Projection <i>Thanh Tran, Hoang V. Tran, Thanh Chu, Huyen Trang Pham, Laurent Ghaoui, Tam Le, Tan Minh Nguyen</i>	58725
Linear Transformers as VAR Models: Aligning Autoregressive Attention Mechanisms with Autoregressive Forecasting <i>Jiecheng Lu, Shihao Yang</i>	58758
Long-Short Alignment for Effective Long-Context Modeling in LLMs <i>Tianqi Du, Haotian Huang, Yifei Wang, Yisen Wang</i>	58778
The Case for Learned Provenance-based System Behavior Baseline <i>Yao Zhu, Zhenyuan Li, Yangyang Wei, Shouling Ji</i>	58799
Efficient Federated Incomplete Multi-View Clustering <i>Suyuan Liu, Hao Yu, Hao Tan, KE LIANG, Siwei Wang, Shengju Yu, En Zhu, Xinwang Liu</i>	58814
LLM-SRBench: A New Benchmark for Scientific Equation Discovery with Large Language Models <i>Parshin Shojaee, Ngoc-Hieu Nguyen, Kazem Meidani, Amir Barati Farimani, Khoa D Doan, Chandan K. Reddy</i>	58829
Average Certified Radius is a Poor Metric for Randomized Smoothing <i>Chenhao Sun, Yuhao Mao, Mark Niklas Mueller, Martin Vechev</i>	58864
PTTA: Purifying Malicious Samples for Test-Time Model Adaptation <i>Jing Ma, Hanlin Li, Xiang Xiang</i>	58880
PertEval-scFM: Benchmarking Single-Cell Foundation Models for Perturbation Effect Prediction <i>Aaron Wenteler, Martina Occhetta, Nikhil Branson, Victor Curean, Magdalena Huebner, William Dee, William Connell, Siu Pui Chung, Alex Hawkins-Hooker, Yasha Ektefaie, César Miguel Valdez Córdova, Amaya Gallagher-Syed</i>	58912
From Logits to Hierarchies: Hierarchical Clustering made Simple <i>Emanuele Palumbo, Moritz Vandenhirtz, Alain Ryser, Imant Daunhawer, Julia E Vogt</i>	58957
In-Context Learning and Occam's Razor <i>Eric Elmoznino, Tom Marty, Tejas Kasetty, Leo Gagnon, Sarthak Mittal, Mahan Fathi, Dhanya Sridhar, Guillaume Lajoie</i>	58977
Universal Neural Optimal Transport <i>Jonathan Geuter, Gregor Kornhardt, Ingimar Tomasson, Vaios Laschos</i>	59001
Deep Principal Support Vector Machines for Nonlinear Sufficient Dimension Reduction <i>Yinfeng Chen, Jin Liu, Rui Qiu</i>	59038
Simple Path Structural Encoding for Graph Transformers <i>Louis Airale, Antonio Longa, Mattia Rigon, Andrea Passerini, Roberto Passerone</i>	59058

MENTOR: Mixture-of-Experts Network with Task-Oriented Perturbation for Visual Reinforcement Learning <i>Suning Huang, Zheyu Aqa Zhang, Tianhai Liang, Yihan Xu, Zhehao Kou, Chenhao Lu, Guowei Xu, Zhengrong Xue, Huazhe Xu</i>	59075
Energy-Based Preference Model Offers Better Offline Alignment than the Bradley-Terry Preference Model <i>Yuzhong Hong, Hanshan Zhang, Junwei Bao, Hongfei Jiang, yang song</i>	59094
A Rescaling-Invariant Lipschitz Bound Based on Path-Metrics for Modern ReLU Network Parameterizations <i>Antoine Gonon, Nicolas Brisebarre, Elisa Riccietti, Rémi Gribonval</i>	59112
Heterogeneous Sufficient Dimension Reduction and Subspace Clustering <i>Lei Yan, Xin Zhang, Qing Mai</i>	59140
A Physics-Augmented Deep Learning Framework for Classifying Single Molecule Force Spectroscopy Data <i>Cailong Hua, Sivaraman Rajaganapathy, Rebecca A Slick, Joseph Vavra, Joseph M. Muretta, James M. Ervasti, Murti Salapaka</i>	59220
Stacey: Promoting Stochastic Steepest Descent via Accelerated ℓ_p-Smooth Nonconvex Optimization <i>Xinyu Luo, Site Bai, Bolian Li, Petros Drineas, Ruqi Zhang, Brian Bullins</i>	59245
Deep Unsupervised Hashing via External Guidance <i>Qihong Song, XitingLiu, Hongyuan Zhu, Joey Tianyi Zhou, Xi Peng, Peng Hu</i>	59262
Skrr: Skip and Re-use Text Encoder Layers for Memory Efficient Text-to-Image Generation <i>Hoigi Seo, Wongi Jeong, Jae-sun Seo, Se Young Chun</i>	59277
The Power of Random Features and the Limits of Distribution-Free Gradient Descent <i>Ari Karchmer, Eran Malach</i>	59316
Robot-Gated Interactive Imitation Learning with Adaptive Intervention Mechanism <i>Haoyuan Cai, Zhenghao Peng, Bolei Zhou</i>	59336
Adaptive Localization of Knowledge Negation for Continual LLM Unlearning <i>Abudukelimu Wuerkaixi, Qizhou Wang, Sen Cui, Wutong Xu, Bo Han, Gang Niu, Masashi Sugiyama, Changshui Zhang</i>	59350
RuleAdapter: Dynamic Rules for training Safety Reward Models in RLHF <i>Xiaomin Li, Mingye Gao, Zhiwei Zhang, Jingxuan Fan, Weiyu Li</i>	59374
Enhancing Cooperative Multi-Agent Reinforcement Learning with State Modelling and Adversarial Exploration <i>Andreas Kontogiannis, Konstantinos Papathanasiou, Yi Shen, Giorgos Stamou, Michael M. Zavlanos, George Vouras</i>	59398

CLIMB: Data Foundations for Large Scale Multimodal Clinical Foundation Models <i>Wei Dai, Peilin Chen, Malinda Lu, Daniel A Li, Haowen Wei, Hejie Cui, Paul Pu Liang</i>	59428
Exploiting Similarity for Computation and Communication-Efficient Decentralized Optimization <i>Yuki Takezawa, Xiaowen Jiang, Anton Rodomanov, Sebastian U Stich</i>	59478
Action-Dependent Optimality-Preserving Reward Shaping <i>Grant Collier Forbes, Jianxun Wang, Leonardo Villalobos-Arias, Arnav Jhala, David Roberts</i>	59525
Non-stationary Online Learning for Curved Losses: Improved Dynamic Regret via Mixability <i>Yu-Jie Zhang, Peng Zhao, Masashi Sugiyama</i>	59540
SEFE: Superficial and Essential Forgetting Eliminator for Multimodal Continual Instruction Tuning <i>Jinpeng Chen, Runmin Cong, Yuzhi Zhao, Hongzheng Yang, Guangneng Hu, Horace Ip, Sam Kwong</i>	59565
On the Robustness of Reward Models for Language Model Alignment <i>Jiwoo Hong, Noah Lee, Euniki Kim, Guijin Son, Woojin Chung, Aman Gupta, Shao Tang, James Thorne</i>	59585
Proxy-FDA: Proxy-based Feature Distribution Alignment for Fine-tuning Vision Foundation Models without Forgetting <i>Chen Huang, Skyler Seto, Hadi Pouransari, Mehrdad Farajtabar, Raviteja Vemulapalli, Fartash Faghri, Oncel Tuzel, Barry-John Theobald, Joshua M. Susskind</i>	59603
BEST-Route: Adaptive LLM Routing with Test-Time Optimal Compute <i>Dujian Ding, Ankur Mallick, Shaokun Zhang, Chi Wang, Daniel Madrigal, Mirian Del Carmen Hipolito Garcia, Menglin Xia, Laks V. S. Lakshmanan, Qingyun Wu, Victor Rühle</i>	59625
Tokenized Bandit for LLM Decoding and Alignment <i>Suho Shin, Chenghao Yang, Haifeng Xu, MohammadTaghi Hajiaghayi</i>	59640
Optimal Fair Learning Robust to Adversarial Distribution Shift <i>Sushant Agarwal, Amit Deshpande, Rajmohan Rajaraman, Ravi Sundaram</i>	59668
Semantic Shift Estimation via Dual-Projection and Classifier Reconstruction for Exemplar-Free Class-Incremental Learning <i>Run He, Di Fang, Yicheng Xu, Yawen Cui, Ming Li, Cen Chen, Ziqian Zeng, Huiping Zhuang</i>	59686
Noisy SIGNSGD Is More Differentially Private Than You (Might) Think <i>Richeng Jin, Huaiyu Dai</i>	59701
Adapting Precomputed Features for Efficient Graph Condensation <i>Yuan Li, Jun Hu, Zemin Liu, Bryan Hooi, Jia Chen, Bingsheng He</i>	59723
SADA: Stability-guided Adaptive Diffusion Acceleration <i>Ting Jiang, Yixiao Wang, Hancheng Ye, Zishan Shao, Jingwei Sun, Jingyang Zhang, Zekai Chen, Jianyi Zhang, Yiran Chen, Hai Li</i>	59737

Reinforcement Learning with Adaptive Reward Modeling for Expensive-to-Evaluate Systems <i>Hongyuan Su, Yu Zheng, Yuan Yuan, Yuming Lin, Depeng Jin, Yong Li</i>	59758
TSP: A Two-Sided Smoothed Primal-Dual Method for Nonconvex Bilevel Optimization <i>Songtao Lu</i>	59771
RUN: Reversible Unfolding Network for Concealed Object Segmentation <i>Chunming He, Rihan Zhang, Fengyang Xiao, Chengyu Fang, Longxiang Tang, Yulun Zhang, Linghe Kong, Deng-Ping Fan, Kai Li, Sina Farsiu</i>	59815
Just Enough Shifts: Mitigating Over-Refusal in Aligned Language Models with Targeted Representation Fine-Tuning <i>Mahavir Dabas, Si Chen, Charles Fleming, Ming Jin, Ruoxi Jia</i>	59827
ROS: A GNN-based Relax-Optimize-and-Sample Framework for Max-k-Cut Problems <i>Yeqing Qiu, Ye XUE, Akang Wang, Yiheng Wang, Qingjiang Shi, Zhi-Quan Luo</i>	59843
Retrieval-Augmented Language Model for Knowledge-aware Protein Encoding <i>Jiasheng Zhang, Delvin Ce Zhang, Shuang Liang, Zhengpin Li, Rex Ying, Jie Shao</i>	59866
any4: Learned 4-bit Numeric Representation for LLMs <i>Mostafa Elhoushi, Jeff Johnson</i>	59882
The Disparate Benefits of Deep Ensembles <i>Kajetan Schweighofer, Adrian Arnaiz-Rodriguez, Sepp Hochreiter, Nuria M Oliver</i>	59899
Noise-Guided Predicate Representation Extraction and Diffusion-Enhanced Discretization for Scene Graph Generation <i>Guoqing Zhang, Shichao Kan, Fanghui Zhang, Wanru Xu, Yue Zhang, Yigang Cen</i>	59939
TGDPO: Harnessing Token-Level Reward Guidance for Enhancing Direct Preference Optimization <i>Mingkang Zhu, Xi Chen, Zhongdao Wang, Bei Yu, Hengshuang Zhao, Jiaya Jia</i>	59958
AutoAL: Automated Active Learning with Differentiable Query Strategy Search <i>Yifeng Wang, Xueying Zhan, Siyu Huang</i>	59977
Ad Hoc Teamwork via Offline Goal-Based Decision Transformers <i>Xinzhi Zhang, Hohei Chan, Deheng Ye, Yi Cai, Mengchen Zhao</i>	59989
Behavioral Exploration: Learning to Explore via In-Context Adaptation <i>Andrew Wagenmaker, Zhiyuan Zhou, Sergey Levine</i>	60004
Machine Learning meets Algebraic Combinatorics: A Suite of Datasets Capturing Research-level Conjecturing Ability in Pure Mathematics <i>Herman Chau, Helen Jenne, Davis Brown, Jesse He, Mark Raugas, Sara C. Billey, Henry Kvinge</i>	60032
Homophily Enhanced Graph Domain Adaptation <i>Ruiyi Fang, Bingheng Li, Jingyu Zhao, Ruizhi Pu, QIUHAO Zeng, Gezheng Xu, Charles Ling, Boyu Wang</i>	60072

Learning Latent Graph Structures and their Uncertainty <i>Alessandro Manenti, Daniele Zambon, Cesare Alippi</i>	60095
T1: Advancing Language Model Reasoning through Reinforcement Learning and Inference Scaling <i>Zhenyu Hou, Xin Lv, Rui Lu, Jiajie Zhang, Yujiang Li, Zijun Yao, Juanzi Li, Jie Tang, Yuxiao Dong</i>	60115
Locality Preserving Markovian Transition for Instance Retrieval <i>Jifei Luo, Wenzheng Wu, Hantao Yao, Lu Yu, Changsheng Xu</i>	60143
Scaling Large Motion Models with Million-Level Human Motions <i>Ye Wang, Sipeng Zheng, Bin Cao, Qianshan Wei, Weishuai Zeng, Qin Jin, Zongqing Lu</i>	60168
VideoRoPE: What Makes for Good Video Rotary Position Embedding? <i>Xilin Wei, Xiaoran Liu, Yuhang Zang, Xiaoyi Dong, Pan Zhang, Yuhang Cao, Jian Tong, Haodong Duan, Qipeng Guo, Jiaqi Wang, Xipeng Qiu, Dahua Lin</i>	60194
Privacy Amplification Through Synthetic Data: Insights from Linear Regression <i>Clément Pierquin, Aurélien Bellet, Marc Tommasi, Matthieu Boussard</i>	60213
Quantum Algorithms for Finite-horizon Markov Decision Processes <i>Bin Luo, Yuwen Huang, Jonathan Allcock, Xiaojun Lin, Shengyu Zhang, John C.S. Lui</i>	60239
HyperNear: Unnoticeable Node Injection Attacks on Hypergraph Neural Networks <i>Tingyi Cai, Yunliang Jiang, Ming Li, Lu Bai, Changqin Huang, Yi Wang</i>	60274
One Leaf Reveals the Season: Occlusion-Based Contrastive Learning with Semantic-Aware Views for Efficient Visual Representation <i>Xiaoyu Yang, Lijian Xu, Hongsheng Li, Shaoting Zhang</i>	60292
World Model Implanting for Test-time Adaptation of Embodied Agents <i>Minjong Yoo, Jinwoo Jang, Sihyung Yoon, Honguk Woo</i>	60308
LAST SToP for Modeling Asynchronous Time Series <i>Shubham Gupta, Thibaut Durand, Graham W. Taylor, Lilian Bialozozowicz</i>	60326
Jacobian Sparse Autoencoders: Sparsify Computations, Not Just Activations <i>Lucy Farnik, Tim Lawson, Conor Houghton, Laurence Aitchison</i>	60351
Optimizing Test-Time Compute via Meta Reinforcement Finetuning <i>Yuxiao Qu, Matthew Y. R. Yang, Amrith Setlur, Lewis Tunstall, Edward Emanuel Beeching, Ruslan Salakhutdinov, Aviral Kumar</i>	60410
Unlocking the Power of SAM 2 for Few-Shot Segmentation <i>Qianxiong Xu, Lanyun Zhu, Xuanyi Liu, Guosheng Lin, Cheng Long, Ziyue Li, Rui Zhao</i>	60443
Double-Filter: Efficient Fine-tuning of Pre-trained Vision-Language Models via Patch&Layer Filtering <i>Yaoqin He, Junchen Fu, Kaiwen Zheng, Songpei Xu, Fuhai Chen, Jie Li, Joemon M. Jose, Xuri Ge</i>	60460

Predictive Data Selection: The Data That Predicts Is the Data That Teaches <i>KaShun SHUM, Yuzhen Huang, Hongjian Zou, dingqi, YiXuan Liao, Xiaoxin Chen, Qian Liu, Junxian He</i>	60475
AlphaVerus: Bootstrapping Formally Verified Code Generation through Self-Improving Translation and Treefinement <i>Pranjal Aggarwal, Bryan Parno, Sean Welleck</i>	60499
Federated In-Context Learning: Iterative Refinement for Improved Answer Quality <i>Ruhan Wang, Zhiyong Wang, Chengkai Huang, Rui Wang, Tong Yu, Lina Yao, John C.S. Lui, Dongruo Zhou</i>	60528
FactTest: Factuality Testing in Large Language Models with Finite-Sample and Distribution-Free Guarantees <i>Fan Nie, Xiaotian Hou, Shuhang Lin, James Zou, Huaxiu Yao, Linjun Zhang</i>	60555
The Emperor's New Clothes in Benchmarking? A Rigorous Examination of Mitigation Strategies for LLM Benchmark Data Contamination <i>Yifan Sun, Han Wang, Dongbai Li, Gang Wang, Huan Zhang</i>	60581
Does learning the right latent variables necessarily improve in-context learning? <i>Sarthak Mittal, Eric Elmoznino, Leo Gagnon, Sangnie Bhardwaj, Guillaume Lajoie, Dhanya Sridhar</i>	60607
Mind the Gap: A Practical Attack on GGUF Quantization <i>Kazuki Egashira, Robin Staab, Mark Vero, Jingxuan He, Martin Vechev</i>	60634
One-Step Generalization Ratio Guided Optimization for Domain Generalization <i>Sumin Cho, Dongwon Kim, Kwangsu Kim</i>	60662
Distributionally Robust Active Learning for Gaussian Process Regression <i>Shion Takeno, Yoshito Okura, Yu Inatsu, Aoyama Tatsuya, Tomonari Tanaka, Akahane Satoshi, Hiroyuki Hanada, Noriaki Hashimoto, Taro Murayama, Hanju Lee, Shinya Kojima, Ichiro Takeuchi</i>	60691
Understanding and Mitigating Miscalibration in Prompt Tuning for Vision-Language Models <i>Shuoyuan Wang, Yixuan Li, Hongxin Wei</i>	60711
Aequa: Fair Model Rewards in Collaborative Learning via Slimmable Networks <i>Nurbek Tastan, Samuel Horváth, Karthik Nandakumar</i>	60734
RepoAudit: An Autonomous LLM-Agent for Repository-Level Code Auditing <i>Jinyao Guo, Chengpeng Wang, Xiangzhe Xu, Zian Su, Xiangyu Zhang</i>	60761
Griffin: Towards a Graph-Centric Relational Database Foundation Model <i>Yanbo Wang, Xiyuan Wang, Quan Gan, Minjie Wang, Qibin Yang, David Wipf, Muhan Zhang</i>	60779
Distributionally Robust Policy Learning under Concept Drifts <i>Jingyuan Wang, Zhimei Ren, Ruohan Zhan, Zhengyuan Zhou</i>	60803
Token Cleaning: Fine-Grained Data Selection for LLM Supervised Fine-Tuning <i>Jinlong Pang, Na Di, Zhaowei Zhu, Jiaheng Wei, Hao Cheng, Chen Qian, Yang Liu</i>	60841

PieClam: A Universal Graph Autoencoder Based on Overlapping Inclusive and Exclusive Communities <i>Daniel Zilberg, Ron Levie</i>	60863
M3-JEPA: Multimodal Alignment via Multi-gate MoE based on the Joint-Embedding Predictive Architecture <i>Hongyang Lei, Xiaolong Cheng, Qi Qin, Dan Wang, Huazhen Huang, Qingqing Gu, Yetao Wu, Luo Ji</i>	60896
Understanding the difficulties of posterior predictive estimation <i>Abhinav Agrawal, Justin Domke</i>	60912
Fast Exact Unlearning for In-Context Learning Data for LLMs <i>Andrei Ioan Muresanu, Anvith Thudi, Michael R. Zhang, Nicolas Papernot</i>	60953
FairICP: Encouraging Equalized Odds via Inverse Conditional Permutation <i>Yuheng Lai, Leying Guan</i>	60970
Feedforward Few-shot Species Range Estimation <i>Christian Lange, Max Hamilton, Elijah Cole, Alexander Shepard, Samuel Heinrich, Angela Zhu, Subhansu Maji, Grant Van Horn, Oisin Mac Aodha</i>	60988
Learning to Route LLMs with Confidence Tokens <i>Yu-Neng Chuang, Prathusha Kameswara Sarma, Parikshit Gopalan, John Boccio, Sara Bolouki, Xia Hu, Helen Zhou</i>	61027
ARS: Adaptive Reward Scaling for Multi-Task Reinforcement Learning <i>Myungsik Cho, Jongeui Park, Jeonghye Kim, Youngchul Sung</i>	61047
In-Context Adaptation to Concept Drift for Learned Database Operations <i>Jiaqi Zhu, Shaofeng Cai, Yanyan Shen, Gang Chen, Fang Deng, Beng Chin Ooi</i>	61066
LADA: Scalable Label-Specific CLIP Adapter for Continual Learning <i>Mao-Lin Luo, Zi-Hao Zhou, Tong Wei, Min-Ling Zhang</i>	61094
Learning Curves of Stochastic Gradient Descent in Kernel Regression <i>Haihan Zhang, Weicheng Lin, Yuanshi Liu, Cong Fang</i>	61110
Theoretical guarantees on the best-of-n alignment policy <i>Ahmad Beirami, Alekh Agarwal, Jonathan Berant, Alexander Nicholas D'Amour, Jacob Eisenstein, Chirag Nagpal, Ananda Theertha Suresh</i>	61145
TopInG: Topologically Interpretable Graph Learning via Persistent Rationale Filtration <i>Cheng Xin, Fan Xu, Xin Ding, Jie Gao, Jiaxin Ding</i>	61168
SE(3)-Equivariant Diffusion Policy in Spherical Fourier Space <i>Xupeng Zhu, Fan Wang, Robin Walters, Jane Shi</i>	61191
Cross-City Latent Space Alignment for Consistency Region Embedding <i>Meng Chen, Hongwei Jia, Zechen Li, Wenzhen Jia, Kai Zhao, Hongjun Dai, Weiming Huang</i>	61211

Optimal Information Retention for Time-Series Explanations <i>Jinghang Yue, Jing Wang, Lu Zhang, Shuo Zhang, Da Li, Zhaoyang Ma, Youfang Lin</i>	61225
Fast Video Generation with Sliding Tile Attention <i>Peiyuan Zhang, Yongqi Chen, Runlong Su, Hangliang Ding, Ion Stoica, Zhengzhong Liu, Hao Zhang</i>	61251
Sleeping Reinforcement Learning <i>Simone Drago, Marco Mussi, Alberto Maria Metelli</i>	61269
CateKV: On Sequential Consistency for Long-Context LLM Inference Acceleration <i>Haoyun Jiang, Haolin li, jianwei zhang, Fei Huang, Qiang Hu, Minmin Sun, Shuai Xiao, Yong Li, Junyang Lin, Jiangchao Yao</i>	61329
Scalable Equilibrium Sampling with Sequential Boltzmann Generators <i>Charlie B. Tan, Joey Bose, Chen Lin, Leon Klein, Michael M. Bronstein, Alexander Tong</i>	61346
Sparse Video-Gen: Accelerating Video Diffusion Transformers with Spatial-Temporal Sparsity <i>Haocheng Xi, Shuo Yang, Yilong Zhao, Chenfeng Xu, Muyang Li, Xiuyu Li, Yujun Lin, Han Cai, Jintao Zhang, Dacheng Li, Jianfei Chen, Ion Stoica, Kurt Keutzer, Song Han</i>	61378
UnHiPPO: Uncertainty-aware Initialization for State Space Models <i>Marten Lienen, Abdullah Saydemir, Stephan Günnemann</i>	61395
Neural Genetic Search in Discrete Spaces <i>Hyeonah Kim, Sanghyeok Choi, Jiwoo Son, Jinkyoo Park, Changhyun Kwon</i>	61410
HyperIMTS: Hypergraph Neural Network for Irregular Multivariate Time Series Forecasting <i>Boyuan Li, Yicheng Luo, Zhen Liu, Junhao Zheng, Jianming Lv, Qianli Ma</i>	61429
Preference learning made easy: Everything should be understood through win rate <i>Lily H Zhang, Rajesh Ranganath</i>	61446
Multimodal Medical Code Tokenizer <i>Xiaorui Su, Shvat Messica, Yepeng Huang, Ruth Johnson, Lukas Fesser, Shanghua Gao, Faryad Sahneh, Marinka Zitnik</i>	61469
Progressive Tempering Sampler with Diffusion <i>Severi Rissanen, RuiKang OuYang, Jiajun He, Wenlin Chen, Markus Heinonen, Arno Solin, José Miguel Hernández-Lobato</i>	61485
Rethinking Benign Overfitting in Two-Layer Neural Networks <i>Ruichen Xu, Kexin Chen</i>	61508
Distributed Retraction-Free and Communication-Efficient Optimization on the Stiefel Manifold <i>Yilong Song, Peijin Li, Bin Gao, Kun Yuan</i>	61539
Inverse Bridge Matching Distillation <i>Nikita Gushchin, David Li, Daniil Selikhanovych, Evgeny Burnaev, Dmitry Baranchuk, Alexander Korotin</i>	61577

Proposer-Agent-Evaluator (PAE): Autonomous Skill Discovery For Foundation Model Internet Agents <i>Yifei Zhou, Qianlan Yang, Kaixiang Lin, Min Bai, Xiong Zhou, Yu-Xiong Wang, Sergey Levine, Li Erran Li</i>	61603
Super Deep Contrastive Information Bottleneck for Multi-modal Clustering <i>Zhengzheng Lou, Ke Zhang, Yucong Wu, Shizhe Hu</i>	61642
A Sub-Problem Quantum Alternating Operator Ansatz for Correlation Clustering <i>Lucas Fabian Naumann, Jannik Irmair, Bjoern Andres</i>	61653
SIMPLEMIX: Frustratingly Simple Mixing of Off- and On-policy Data in Language Model Preference Learning <i>Tianjian Li, Daniel Khashabi</i>	61669
Taming Rectified Flow for Inversion and Editing <i>Jiangshan Wang, Junfu Pu, Zhongang Qi, Jiayi Guo, Yue Ma, Nisha Huang, Yuxin Chen, Xiu Li, Ying Shan</i>	61689
Video-Enhanced Offline Reinforcement Learning: A Model-Based Approach <i>Minting Pan, Yitao Zheng, Jiajian Li, Yunbo Wang, Xiaokang Yang</i>	61704
High-Dimensional Tensor Regression With Oracle Properties <i>Wenbin Wang, Yu Shi, Ziping Zhao</i>	61721
Training a Generally Curious Agent <i>Fahim Tajwar, Yiding Jiang, Abitha Thankaraj, Sumaita Sadia Rahman, J Zico Kolter, Jeff Schneider, Russ Salakhutdinov</i>	61762
Efficient Core-set Selection for Deep Learning Through Squared Loss Minimization <i>Jianting Chen</i>	61817
Scaling Laws in Patchification: An Image Is Worth 50,176 Tokens And More <i>Feng Wang, Yaodong Yu, Wei Shao, Yuyin Zhou, Alan Yuille, Cihang Xie</i>	61829
When to Forget? Complexity Trade-offs in Machine Unlearning <i>Martin Van Waerebeke, Marco Lorenzi, Giovanni Neglia, Kevin Scaman</i>	61842
Token Signature: Predicting Chain-of-Thought Gains with Token Decoding Feature in Large Language Models <i>Peijie Liu, Fengli Xu, Yong Li</i>	61859
Beyond Bradley-Terry Models: A General Preference Model for Language Model Alignment <i>Yifan Zhang, Ge Zhang, Yue Wu, Kangping Xu, Quanquan Gu</i>	61880
OTTER: A Vision-Language-Action Model with Text-Aware Visual Feature Extraction <i>Huang Huang, Fangchen Liu, Letian Fu, Tingfan Wu, Mustafa Mukadam, Jitendra Malik, Ken Goldberg, Pieter Abbeel</i>	61907

Introducing 3D Representation for Dense Volume-to-Volume Translation via Score Fusion <i>Xiyue Zhu, Dou Hoon Kwark, Ruike Zhu, Kaiwen Hong, Yiqi Tao, Shirui Luo, Yudu Li, Zhi-Pei Liang, Volodymyr Kindratenko</i>	61922
Model Immunization from a Condition Number Perspective <i>Amber Yijia Zheng, Site Bai, Brian Bullins, Raymond A. Yeh</i>	61957
Navigating Conflicting Views: Harnessing Trust for Learning <i>Jueqing Lu, Wray Buntine, Yuanyuan Qi, Joanna Dipnall, Belinda Gabbe, Lan Du</i>	61983
Improving Out-of-Distribution Detection via Dynamic Covariance Calibration <i>Kaiyu Guo, Zijian Wang, Tan Pan, Brian C. Lovell, Mahsa Baktashmotlagh</i>	62008
Hyperbolic-PDE GNN: Spectral Graph Neural Networks in the Perspective of A System of Hyperbolic Partial Differential Equations <i>Juwei Yue, Haikuo Li, Jiawei Sheng, Xiaodong Li, Taoyu Su, Tingwen Liu, Li Guo</i>	62024
Optimizing Large Language Model Training Using FP4 Quantization <i>Ruizhe Wang, Yeyun Gong, Xiao Liu, Guoshuai Zhao, Ziyue Yang, Baining Guo, Zheng-Jun Zha, Peng CHENG</i>	62042
Probably Approximately Global Robustness Certification <i>Peter Blohm, Patrick Indri, Thomas Gärtner, SAGAR MALHOTRA</i>	62063
Can DBNNs Robust to Environmental Noise for Resource-constrained Scenarios? <i>Wendong Zheng, Junyang Chen, Husheng Guo, Wenjian Wang</i>	62081
Natural Perturbations for Black-box Training of Neural Networks by Zeroth-Order Optimization <i>Hiroshi Sawada, Kazuo Aoyama, Yuya Hikima</i>	62101
Towards Black-Box Membership Inference Attack for Diffusion Models <i>Jingwei Li, Jing Dong, Tianxing He, Jingzhao Zhang</i>	62118
MM-RLHF: The Next Step Forward in Multimodal LLM Alignment <i>YiFan Zhang, Tao Yu, Haochen Tian, Chaoyou Fu, Peiyan Li, Jianshu Zeng, Wulin Xie, Yang Shi, Huanyu Zhang, Junkang Wu, Xue Wang, Yibo Hu, Bin Wen, Tingting Gao, Zhang Zhang, Fan Yang, Di ZHANG, Liang Wang, Rong Jin</i>	62135
OmniArch: Building Foundation Model for Scientific Computing <i>Tianyu Chen, Haoyi Zhou, Ying Li, Hao Wang, Chonghan Gao, Rongye Shi, Shanghang Zhang, Jianxin Li</i>	62165
Privacy-Preserving Federated Convex Optimization: Balancing Partial-Participation and Efficiency via Noise Cancellation <i>Roie Reshef, Kfir Yehuda Levy</i>	62193
Towards Robust Influence Functions with Flat Validation Minima <i>Xichen Ye, Yifan Wu, WEIZHONG ZHANG, Cheng Jin, Yifan Chen</i>	62221

Steer LLM Latents for Hallucination Detection <i>Seongheon Park, Xuefeng Du, Min-Hsuan Yeh, Haobo Wang, Yixuan Li</i>	62242
DEFAME: Dynamic Evidence-based Fact-checking with Multimodal Experts <i>Tobias Braun, Mark Rothermel, Marcus Rohrbach, Anna Rohrbach</i>	62262
Non-asymptotic Error Bounds in ℓ_2-Distance with Sqrt(d) Dimension Dependence and First Order Convergence for Langevin Monte Carlo beyond Log-Concavity <i>Bin Yang, Xiaojie Wang</i>	62297
REG: Rectified Gradient Guidance for Conditional Diffusion Models <i>Zhengqi Gao, Kaiwen Zha, Tianyuan Zhang, Zihui Xue, Duane S Boning</i>	62322
Accelerating Linear Recurrent Neural Networks for the Edge with Unstructured Sparsity <i>Alessandro Pierro, Steven Abreu, Jonathan Timcheck, Philipp Stratmann, Andreas Wild, Sumit Bam Shrestha</i>	62342
Constrained Exploitability Descent: An Offline Reinforcement Learning Method for Finding Mixed-Strategy Nash Equilibrium <i>Runyu Lu, Yuanheng Zhu, Dongbin Zhao</i>	62359
Universal Sparse Autoencoders: Interpretable Cross-Model Concept Alignment <i>Harrish Thasarthan, Julian Forsyth, Thomas Fel, Matthew Kowal, Konstantinos G. Derpanis</i>	62377
COSDA: Counterfactual-based Susceptibility Risk Framework for Open-Set Domain Adaptation <i>Wenxu Wang, Rui Zhou, Jing Wang, Yun Zhou, Cheng Zhu, Ruichun Tang, Bo Han, Nevin L. Zhang</i>	62399
RWKVQuant: Quantizing the RWKV Family with Proxy Guided Hybrid of Scalar and Vector Quantization <i>XUCHEN, Yuxuan Yue, Zukang Xu, Xing Hu, JiangyongYu, Zhixuan Chen, Sifan Zhou, Zhihang Yuan, Dawei Yang</i>	62421
Fully Dynamic Euclidean Bi-Chromatic Matching in Sublinear Update Time <i>Gramoz Goranci, Peter Kiss, Neel Patel, Martin P. Seybold, Eva Szilagyi, Da Wei Zheng</i>	62439
Joint Learning of Energy-based Models and their Partition Function <i>Michael Eli Sander, Vincent Roulet, Tianlin Liu, Mathieu Blondel</i>	62464
What Makes In-context Learning Effective for Mathematical Reasoning <i>Jiayu Liu, Zhenya Huang, Chaokun Wang, Xunpeng Huang, ChengXiang Zhai, Enhong Chen</i>	62484
Unifying Knowledge from Diverse Datasets to Enhance Spatial-Temporal Modeling: A Granularity-Adaptive Geographical Embedding Approach <i>Zhigaoyuan Wang, Ying Sun, Hengshu Zhu</i>	62498
ESPFormer: Doubly-Stochastic Attention with Expected Sliced Transport Plans <i>Ashkan Shahbazi, Elaheh Akbari, Darian Salehi, Xinran Liu, Navid NaderiAlizadeh, Soheil Kolouri</i>	62513

RULEBREAKERS: Challenging LLMs at the Crossroads between Formal Logic and Human-like Reasoning <i>Jason Chan, Robert J. Gaizauskas, Zhixue Zhao</i>	62530
High-Dimensional Prediction for Sequential Decision Making <i>Georgy Noarov, Ramya Ramalingam, Aaron Roth, Stephan Xie</i>	62560
Global-Local Dirichlet Processes for Clustering Grouped Data in the Presence of Group-Specific Idiosyncratic Variables <i>Arhit Chakrabarti, Yang Ni, Debdeep Pati, Bani Mallick</i>	62582
BRIDGE: Bootstrapping Text to Control Time-Series Generation via Multi-Agent Iterative Optimization and Diffusion Modeling <i>Hao Li, Yu-Hao Huang, Chang Xu, Viktor Schlegel, Renhe Jiang, Riza Batista-Navarro, Goran Nenadic, Jiang Bian</i>	62618
The Generalized Skew Spectrum of Graphs <i>Armando Bellante, Martin Plávala, Alessandro Luongo</i>	62650
Gradient Boosting Reinforcement Learning <i>Benjamin Fuhrer, Chen Tessler, Gal Dalal</i>	62670
Point Cloud Dataset Distillation <i>Deyu Bo, Xinchao Wang</i>	62696
TabSDS: a Lightweight, Fully Non-Parametric, and Model Free Approach for Generating Synthetic Tabular Data <i>Elias Chaibub Neto</i>	62710
FOUNDER: Grounding Foundation Models in World Models for Open-Ended Embodied Decision Making <i>Yucen Wang, Rui Yu, Shenghua Wan, Le Gan, De-Chuan Zhan</i>	62775
LLMs Can Reason Faster Only If We Let Them <i>Bilgehan Sel, Lifu Huang, Naren Ramakrishnan, Ruoxi Jia, Ming Jin</i>	62794
A General Graph Spectral Wavelet Convolution via Chebyshev Order Decomposition <i>Nian Liu, Xiaoxin He, Thomas Laurent, Francesco Di Giovanni, Michael M. Bronstein, Xavier Bresson</i>	62810
FlatQuant: Flatness Matters for LLM Quantization <i>Yuxuan Sun, Ruikang Liu, Haoli Bai, Han Bao, Kang Zhao, Yuening Li, JiaxinHu, Xianzhi Yu, Lu Hou, Chun Yuan, Xin Jiang, Wulong Liu, Jun Yao</i>	62835
ConText: Driving In-context Learning for Text Removal and Segmentation <i>Fei Zhang, Pei Zhang, Baosong Yang, Fei Huang, Yanfeng Wang, Ya Zhang</i>	62862
Multi-Modal Object Re-identification via Sparse Mixture-of-Experts <i>Yingying Feng, Jie Li, Chi Xie, Lei Tan, Jiayi Ji</i>	62881

BoA: Attention-aware Post-training Quantization without Backpropagation <i>Junhan Kim, Ho-young Kim, Eulrang Cho, Chungman Lee, Joonyoung Kim, Yongkweon Jeon</i>	62893
Teaching Language Models to Critique via Reinforcement Learning <i>Zhihui Xie, Jie chen, Liyu Chen, Weichao Mao, Jingjing Xu, Lingpeng Kong</i>	62914
Few-Shot Learner Generalizes Across AI-Generated Image Detection <i>Shiyu Wu, Jing Liu, Jing Li, Yequan Wang</i>	62933
Maximizing Intermediate Checkpoint Value in LLM Pretraining with Bayesian Optimization <i>Deyuan Liu, Zecheng Wang, Bingning Wang, Weipeng Chen, Chunshan Li, Zhiying Tu, Dianhui Chu, Dianbo Sui</i>	62945
Balancing Interference and Correlation in Spatial Experimental Designs: A Causal Graph Cut Approach <i>Jin Zhu, Jingyi Li, Hongyi Zhou, Yinan Lin, Zhenhua Lin, Chengchun Shi</i>	62974
EQ-VAE: Equivariance Regularized Latent Space for Improved Generative Image Modeling <i>Theodoros Kouzelis, Ioannis Kakogeorgiou, Spyros Gidaris, Nikos Komodakis</i>	62994
LLM Enhancers for GNNs: An Analysis from the Perspective of Causal Mechanism Identification <i>Hang Gao, Huang Wenxuan, Fengge Wu, Zhao Junsuo, Changwen Zheng, Huaping Liu</i>	63013
Self-Consuming Generative Models with Adversarially Curated Data <i>Xiukun Wei, Xueru Zhang</i>	63040
MOGIC: Metadata-infused Oracle Guidance for Improved Extreme Classification <i>Suchith Chidananda Prabhu, Bhavyajeet Singh, Anshul Mittal, Siddarth Asokan, Shikhar Mohan, Deepak Saini, Yashoteja Prabhu, Lakshya Kumar, Jian Jiao, Amit S, Niket Tandon, Manish Gupta, Sumeet Agarwal, Manik Varma</i>	63062
Counting atoms faster: policy-based nuclear magnetic resonance pulse sequencing for atomic abundance measurement <i>Rohan Shenoy, Evan Austen Coleman, Hans Gaensbauer, Elsa Olivetti</i>	63090
In-Context Fine-Tuning for Time-Series Foundation Models <i>Matthew Faw, Rajat Sen, Yichen Zhou, Abhimanyu Das</i>	63105
Noise Conditional Variational Score Distillation <i>Xinyu Peng, Ziyang Zheng, Yaoming Wang, Han Li, Nuowen Kan, Wenrui Dai, Chenglin Li, Junni Zou, Hongkai Xiong</i>	63125
FourierMamba: Fourier Learning Integration with State Space Models for Image Deraining <i>Dong Li, Yidi Liu, Xueyang Fu, Jie Huang, Senyan Xu, Qi Zhu, Zheng-Jun Zha</i>	63150
Scaling Probabilistic Circuits via Monarch Matrices <i>Honghua Zhang, Meihua Dang, Benjie Wang, Stefano Ermon, Nanyun Peng, Guy Van den Broeck</i>	63174

A-PSRO: A Unified Strategy Learning Method with Advantage Metric for Normal-form Games <i>Yudong Hu, Haoran Li, Congying Han, Tiande Guo, Bonan Li, Mingqiang Li</i>	63189
AGAV-Rater: Adapting Large Multimodal Model for AI-Generated Audio-Visual Quality Assessment <i>Yuqin Cao, Xiongkuo Min, Yixuan Gao, Wei Sun, Guangtao Zhai</i>	63216
Suitability Filter: A Statistical Framework for Classifier Evaluation in Real-World Deployment Settings <i>Angéline Pouget, Mohammad Yaghini, Stephan Rabanser, Nicolas Papernot</i>	63235
Quantum Optimization via Gradient-Based Hamiltonian Descent <i>Jiaqi Leng, Bin Shi</i>	63260
FedSMU: Communication-Efficient and Generalization-Enhanced Federated Learning through Symbolic Model Updates <i>Xinyi Lu, Hao Zhang, Chenglin Li, Weijia Lu, ZHIFEI YANG, Wenrui Dai, xiaodong Zhang, Xiaofeng Ma, Can Zhang, Junni Zou, Hongkai Xiong</i>	63279
Loss Functions and Operators Generated by f-Divergences <i>Vincent Roulet, Tianlin Liu, Nino Vieillard, Michael Eli Sander, Mathieu Blondel</i>	63306
Can MLLMs Reason in Multimodality? EMMA: An Enhanced MultiModal ReAsoning Benchmark <i>Yunzhuo Hao, Jiawei Gu, Huichen Will Wang, Linjie Li, Zhengyuan Yang, Lijuan Wang, Yu Cheng</i>	63335
BSO: Binary Spiking Online Optimization Algorithm <i>Yu Liang, Yu Yang, Wenjie Wei, Ammar Belatreche, Shuai Wang, Malu Zhang, Yang Yang</i>	63382
TimeDART: A Diffusion Autoregressive Transformer for Self-Supervised Time Series Representation <i>Daoyu Wang, Mingyue Cheng, Zhiding Liu, Qi Liu</i>	63396
Geometric Contact Flows: Contactomorphisms for Dynamics and Control <i>Andrea Testa, Søren Hauberg, Tamim Asfour, Leonel Roza</i>	63421
Are LLMs Prescient? A Continuous Evaluation using Daily News as the Oracle <i>Hui Dai, Ryan Teehan, Mengye Ren</i>	63447
MODULI: Unlocking Preference Generalization via Diffusion Models for Offline Multi-Objective Reinforcement Learning <i>Yifu Yuan, Zhenrui Zheng, Zibin Dong, Jianye HAO</i>	63482
RIFLEx: A Free Lunch for Length Extrapolation in Video Diffusion Transformers <i>Min Zhao, Guande He, Yixiao Chen, Hongzhou Zhu, Chongxuan Li, Jun Zhu</i>	63506
Beyond The Rainbow: High Performance Deep Reinforcement Learning on a Desktop PC <i>Tyler Clark, Mark Towers, Christine Evers, Jonathon Hare</i>	63525

Knowledge-Guided Wasserstein Distributionally Robust Optimization <i>Zitao Wang, Ziyuan Wang, Molei Liu, Nian Si</i>	63553
Beyond Communication Overhead: A Multilevel Monte Carlo Approach for Mitigating Compression Bias in Distributed Learning <i>Ze'ev Zukerman, Bassel Hamoud, Kfir Yehuda Levy</i>	63576
Score as Action: Fine Tuning Diffusion Generative Models by Continuous-time Reinforcement Learning <i>Hanyang Zhao, Haoxian Chen, Ji Zhang, David Yao, Wenpin Tang</i>	63604
Regress, Don't Guess: A Regression-like Loss on Number Tokens for Language Models <i>Jonas Zausinger, Lars Pennig, Anamarija Kozina, Sean Sdahl, Julian Sikora, Adrian Dendorfer, Timofey Kuznetsov, Mohamad Hagog, Nina Wiedemann, Kacper Chlodny, Vincent Limbach, Anna Ketteler, Thorben Prein, Vishwa Mohan Singh, Michael Danziger, Jannis Born</i>	63623
Modified K-means Algorithm with Local Optimality Guarantees <i>Mingyi Li, Michael R. Metel, Akiko Takeda</i>	63646
Reaction Graph: Towards Reaction-Level Modeling for Chemical Reactions with 3D Structures <i>Yingzhao Jian, Yue Zhang, Ying Wei, Hehe Fan, Yi Yang</i>	63675
Outlier Gradient Analysis: Efficiently Identifying Detrimental Training Samples for Deep Learning Models <i>Anshuman Chhabra, Bo Li, Jian Chen, Prasant Mohapatra, Hongfu Liu</i>	63741
CUPS: Improving Human Pose-Shape Estimators with Conformalized Deep Uncertainty <i>Harry Zhang, Luca Carlone</i>	63761
UP-VLA: A Unified Understanding and Prediction Model for Embodied Agent <i>Jianke Zhang, Yanjiang Guo, Yucheng Hu, Xiaoyu Chen, Xiang Zhu, Jianyu Chen</i>	63780
Sub-Sequential Physics-Informed Learning with State Space Model <i>Chenhui Xu, Dancheng Liu, Yuting Hu, Jiajie Li, Ruiyang Qin, Qingxiao Zheng, Jinjun Xiong</i>	63792
Latent Variable Estimation in Bayesian Black-Litterman Models <i>Thomas Yuan-Lung Lin, Jerry Yao-Chieh Hu, Paul W. Chiou, Peter Lin</i>	63811
Finite-Time Analysis of Discrete-Time Stochastic Interpolants <i>Yuhao Liu, Yu Chen, Rui Hu, Longbo Huang</i>	63839
Objective drives the consistency of representational similarity across datasets <i>Laure Ciernik, Lorenz Linhardt, Marco Morik, Jonas Dippel, Simon Kornblith, Lukas Muttenthaler</i>	63864
The Hidden Joules: Evaluating the Energy Consumption of Vision Backbones for Progress Towards More Efficient Model Inference <i>Zeyu Yang, Wesley Armour</i>	63893

Permutation-based Rank Test in the Presence of Discretization and Application in Causal Discovery with Mixed Data <i>Xinshuai Dong, Ignavier Ng, Boyang Sun, Haoyue Dai, Guang-Yuan Hao, Shunxing Fan, Peter Spirtes, Yumou Qiu, Kun Zhang</i>	63910
Clone-Robust AI Alignment <i>Ariel D. Procaccia, Benjamin Schiffer, Shirley Zhang</i>	63925
Improved Lower Bounds for First-order Stochastic Non-convex Optimization under Markov Sampling <i>Zhenyu Sun, Ermin Wei</i>	63949
Understanding the Unfairness in Network Quantization <i>Bing Liu, Wenjun Miao, Boyu Zhang, Qiankun Zhang, Bin Yuan, Jing Wang, Shenghao Liu, Xianjun Deng</i>	63968
Improved Learning via k-DTW: A Novel Dissimilarity Measure for Curves <i>Amer Krivošija, Alexander Munteanu, André Nusser, Chris Schwiegelshohn</i>	63988
Recommendations with Sparse Comparison Data: Provably Fast Convergence for Nonconvex Matrix Factorization <i>Suryanarayana Sankagiri, Jalal Etesami, Matthias Grossglauser</i>	64019
A Chaotic Dynamics Framework Inspired by Dorsal Stream for Event Signal Processing <i>Yu Chen, Jing Lian, Zhaofei Yu, Jizhao Liu, Jisheng Dang, Gang Wang</i>	64054
Multidimensional Adaptive Coefficient for Inference Trajectory Optimization in Flow and Diffusion <i>Dohoon Lee, Jaehyun Park, Hyunwoo J. Kim, Kyogu Lee</i>	64064
Unlocking the Capabilities of Large Vision-Language Models for Generalizable and Explainable Deepfake Detection <i>Peipeng Yu, Jianwei Fei, Hui Gao, Xuan Feng, Zhihua Xia, Chip Hong Chang</i>	64084
Towards Lifelong Model Editing via Simulating Ideal Editor <i>Yaming Guo, Siyang Guo, Hengshu Zhu, Ying Sun</i>	64103
DCTdiff: Intriguing Properties of Image Generative Modeling in the DCT Space <i>Mang Ning, Mingxiao Li, Jianlin Su, Jia Haozhe, Lanmiao Liu, Martin Benes, Wenshuo Chen, Albert Ali Salah, Itir Onal Ertugrul</i>	64135
DiffusionVLA: Scaling Robot Foundation Models via Unified Diffusion and Autoregression <i>Junjie Wen, Yichen Zhu, Minjie Zhu, Zhibin Tang, Jinming Li, Zhongyi Zhou, Xiaoyu Liu, Chaomin Shen, Yaxin Peng, Feifei Feng</i>	64162
Canonical Rank Adaptation: An Efficient Fine-Tuning Strategy for Vision Transformers <i>Lokesh Veeramacheni, Moritz Wolter, Hilde Kuehne, Juergen Gall</i>	64179

unMORE: Unsupervised Multi-Object Segmentation via Center-Boundary Reasoning <i>Yafei YANG, Zihui Zhang, Bo Yang</i>	64197
Computing Voting Rules with Improvement Feedback <i>Evi Micha, Vasilis Varsamis</i>	64223
Online Linear Classification with Massart Noise <i>Ilias Diakonikolas, Vasilis Kontonis, Christos Tzamos, Nikos Zarifis</i>	64251
LangTime: A Language-Guided Unified Model for Time Series Forecasting with Proximal Policy Optimization <i>Wenzhe Niu, Zongxia Xie, Yanru Sun, Wei He, Man Xu, Chao Hao</i>	64265
Hierarchical Planning for Complex Tasks with Knowledge Graph-RAG and Symbolic Verification <i>Flavio Petruzzellis, Cristina Cornelio, Pietro Lio</i>	64288
AKORN: Adaptive Knots generated Online for RegressionN splines <i>Sunil Madhow, Dheeraj Baby, Yu-Xiang Wang</i>	64311
Expressive Score-Based Priors for Distribution Matching with Geometry-Preserving Regularization <i>Ziyu Gong, Jim Lim, David I. Inouye</i>	64338
Occult: Optimizing Collaborative Communications across Experts for Accelerated Parallel MoE Training and Inference <i>Shuqing Luo, Pingzhi Li, Jie Peng, Yang Zhao, Yu Cao, Yu Cheng, Tianlong Chen</i>	64370
Scalable Sobolev IPM for Probability Measures on a Graph <i>Tam Le, Truyen Nguyen, Hideitsu Hino, Kenji Fukumizu</i>	64389
SCENT: Robust Spatiotemporal Learning for Continuous Scientific Data via Scalable Conditioned Neural Fields <i>David Keetae Park, Xihaier Luo, Guang Zhao, Seungjun Lee, Miruna Oprescu, Shinjae Yoo</i>	64421
LLM-Assisted Semantically Diverse Teammate Generation for Efficient Multi-agent Coordination <i>Lihe Li, Lei Yuan, Pengsen Liu, Tao Jiang, Yang Yu</i>	64446
A Forget-and-Grow Strategy for Deep Reinforcement Learning Scaling in Continuous Control <i>Zilin Kang, Chenyuan Hu, Yu Luo, Zhecheng Yuan, Ruijie Zheng, Huazhe Xu</i>	64468
S2-Track: A Simple yet Strong Approach for End-to-End 3D Multi-Object Tracking <i>Tao Tang, Lijun Zhou, Pengkun Hao, Zihang He, Kalok Ho, Shuo Gu, Zhihui Hao, Haiyang Sun, Kun Zhan, Peng Jia, XianPeng Lang, Xiaodan Liang</i>	64490
How to Train Your Multi-Exit Model? Analyzing the Impact of Training Strategies <i>Piotr Kubaty, Bartosz Wójcik, Bartłomiej Tomasz Krzepkowski, Monika Michaluk, Tomasz Trzcinski, Jary Pomponi, Kamil Adamczewski</i>	64507

AdaptiveStep: Automatically Dividing Reasoning Step through Model Confidence <i>Yuliang Liu, Junjie Lu, Chaofeng Qu, Zhaoling Chen, Zefan Cai, Jason Klein Liu, Chonghan Liu, Yunhui Xia, Li Zhao, Jiang Bian, Chuheng Zhang, Wei Shen, Zhouhan Lin</i>	64527
Adversarial Robustness via Deformable Convolution with Stochasticity <i>Yanxiang Ma, Zixuan Huang, Minjing Dong, Shan You, Chang Xu</i>	64543
Certification for Differentially Private Prediction in Gradient-Based Training <i>Matthew Robert Wicker, Philip Sosnin, Igor Shilov, Adrianna Janik, Mark Niklas Mueller, Yves-Alexandre de Montjoye, Adrian Weller, Calvin Tsay</i>	64559
PF3plat: Pose-Free Feed-Forward 3D Gaussian Splatting for Novel View Synthesis <i>Sunghwan Hong, Jaewoo Jung, Heeseong Shin, Jisang Han, Jiaolong Yang, Chong Luo, Seungryong Kim</i>	64579
Learning Dynamics in Continual Pre-Training for Large Language Models <i>Xingjin Wang, Howe Tissue, Lu Wang, Linjing Li, Daniel Dajun Zeng</i>	64599
MuseControlLite: Multifunctional Music Generation with Lightweight Conditioners <i>Fang-Duo Tsai, Shih-Lun Wu, Weijaw Lee, Sheng-Ping Yang, Bo-Rui Chen, Hao-Chung Cheng, Yi-Hsuan Yang</i>	64623
The Four Color Theorem for Cell Instance Segmentation <i>Ye Zhang, Yu Zhou, Yifeng Wang, Jun Xiao, Ziyue Wang, Yongbing Zhang, Jianxu Chen</i>	64637
VTGaussian-SLAM: RGBD SLAM for Large Scale Scenes with Splatting View-Tied 3D Gaussians <i>Pengchong Hu, Zhizhong Han</i>	64659
Sliding Puzzles Gym: A Scalable Benchmark for State Representation in Visual Reinforcement Learning <i>Bryan Lincoln Marques de Oliveira, Luana Guedes Barros Martins, Bruno Brandão, Murilo Lopes da Luz, telma woerle de lima soares, Luckeciano Carvalho Melo</i>	64683
Info-Coevolution: An Efficient Framework for Data Model Coevolution <i>Ziheng Qin, Hailun Xu, Wei Chee Yew, Qi Jia, Yang Luo, Kanchan Sarkar, Danhui Guan, Kai Wang, Yang You</i>	64712
Counterfactual Contrastive Learning with Normalizing Flows for Robust Treatment Effect Estimation <i>Jiaxuan Zhang, Emadeldeen Eldele, Fuyuan CAO, Yang Wang, Xiaoli Li, Jiye Liang</i>	64725
Stay-Positive: A Case for Ignoring Real Image Features in Fake Image Detection <i>Anirudh Sundara Rajan, Yong Jae Lee</i>	64746
AdaDecode: Accelerating LLM Decoding with Adaptive Layer Parallelism <i>Zhepei Wei, Wei-Lin Chen, Xinyu Zhu, Yu Meng</i>	64767
Doubly Protected Estimation for Survival Outcomes Utilizing External Controls for Randomized Clinical Trials <i>Chenyin Gao, Shu Yang, Mingyang Shan, Wenyu Ye, Ilya Lipkovich, Douglas Faries</i>	64783

CLARIFY: Contrastive Preference Reinforcement Learning for Untangling Ambiguous Queries <i>Ni Mu, Hao Hu, Xiao Hu, Yiqin Yang, Bo XU, Qing-Shan Jia</i>	64809
Mitigating Heterogeneous Token Overfitting in LLM Knowledge Editing <i>Tianci Liu, Ruirui Li, Zihan Dong, Hui Liu, Xianfeng Tang, Qingyu Yin, Linjun Zhang, Haoyu Wang, Jing Gao</i>	64828
MetaAgent: Automatically Constructing Multi-Agent Systems Based on Finite State Machines <i>Yaolun Zhang, Xiaogeng Liu, Chaowei Xiao</i>	64851
Symmetry-Driven Discovery of Dynamical Variables in Molecular Simulations <i>Jeet Mohapatra, Nima Dehmamy, Csaba Both, Subhro Das, Tommi Jaakkola</i>	64879
Sampling from Binary Quadratic Distributions via Stochastic Localization <i>Chenguang Wang, Kaiyuan Cui, Weichen Zhao, Tianshu Yu</i>	64900
Accelerating LLM Inference with Lossless Speculative Decoding Algorithms for Heterogeneous Vocabularies <i>Nadav Timor, Jonathan Mamou, Daniel Korat, Moshe Berchansky, Gaurav Jain, Oren Pereg, Moshe Wasserblat, David Harel</i>	64932
Enhancing Performance of Explainable AI Models with Constrained Concept Refinement <i>Geyu Liang, Senne Michielssen, Salar Fattahi</i>	64955
Universal Approximation of Mean-Field Models via Transformers <i>Shiba Biswal, Karthik Elamvazhuthi, Rishi Sonthalia</i>	64995
R3DM: Enabling Role Discovery and Diversity Through Dynamics Models in Multi-agent Reinforcement Learning <i>Harsh Goel, Mohammad Omama, Behdad Chalaki, Vaishnav Tadiparthi, Ehsan Moradi Pari, Sandeep P. Chinchali</i>	65011
On Path to Multimodal Generalist: General-Level and General-Bench <i>Hao Fei, Yuan Zhou, Juncheng Li, Xiangtai Li, Qingshan Xu, Bobo Li, Shengqiong Wu, Yaoting Wang, Junbao Zhou, Jiahao Meng, Qingyu Shi, Zhiyuan Zhou, Liangtao Shi, Minghe Gao, Daoan Zhang, Zhiqi Ge, Siliang Tang, Kaihang Pan, Yaobo Ye, Haobo Yuan, Tao Zhang, Weiming Wu, Tianjie Ju, Zixiang Meng, Shilin Xu, Liyu Jia, Wentao Hu, Meng Luo, Jiebo Luo, Tat-Seng Chua, Shuicheng YAN, Hanwang Zhang</i>	65032
Convergence of Consistency Model with Multistep Sampling under General Data Assumptions <i>Yiding Chen, Yiyi Zhang, Owen Oertell, Wen Sun</i>	65152
Watch Out Your Album! On the Inadvertent Privacy Memorization in Multi-Modal Large Language Models <i>Tianjie Ju, Yi Hua, Hao Fei, Zhenyu Shao, Yubin Zheng, Haodong Zhao, Mong-Li Lee, Wynne Hsu, Zhuosheng Zhang, Gongshen Liu</i>	65171

ABKD: Pursuing a Proper Allocation of the Probability Mass in Knowledge Distillation via α-β-Divergence <i>Guanghui Wang, Zhiyong Yang, Zitai Wang, Shi Wang, Qianqian Xu, Qingming Huang</i>	65188
Reidentify: Context-Aware Identity Generation for Contextual Multi-Agent Reinforcement Learning <i>Zhiwei Xu, Kun Hu, Xin Xin, Weiliang Meng, Yiwei Shi, Hangyu Mao, Bin Zhang, Dapeng Li, Jiangjin Yin</i>	65234
Mahalanobis+ +: Improving OOD Detection via Feature Normalization <i>Maximilian Müller, Matthias Hein</i>	65251
EraseAnything: Enabling Concept Erasure in Rectified Flow Transformers <i>Daiheng Gao, Shilin Lu, Wenbo Zhou, Jiaming Chu, Jie Zhang, Mengxi Jia, Bang Zhang, Zhaoxin Fan, Weiming Zhang</i>	65285
LOGO --- Long cOntext aliGnment via efficient preference Optimization <i>Zecheng Tang, Zechen Sun, Juntao Li, Qiaoming Zhu, Min Zhang</i>	65310
On Linear Convergence in Smooth Convex-Concave Bilinearly-Coupled Saddle-Point Optimization: Lower Bounds and Optimal Algorithms <i>Ekaterina Borodich, Alexander Gasnikov, Dmitry Kovalev</i>	65335
Gravity-Bench-v1: A Benchmark on Gravitational Physics Discovery for Agents <i>Nolan Koblichke, Hyunseok Jang, Kristen Menou, Mohamad Ali-Dib</i>	65391
Feature out! Let Raw Image as Your Condition for Blind Face Restoration <i>Xinmin Qiu, Chen Gege, Bonan Li, Congying Han, Tiande Guo, Zicheng Zhang</i>	65412
Learning Monotonic Probabilities with a Generative Cost Model <i>Yongxiang Tang, Yanhua cheng, Xiaocheng Liu, Jiaochenchen, Yanxiang Zeng, Ning Luo, Pengjia Yuan, Xialong Liu, Peng Jiang</i>	65436
Scaling Laws for Forgetting during Finetuning with Pretraining Data Injection <i>Louis Béthune, David Grangier, Dan Busbridge, Eleonora Gualdoni, marco cuturi, Pierre Ablin</i>	65452
Towards Learning to Complete Anything in Lidar <i>Ayça Takmaz, Cristiano Saltori, Neehar Peri, Tim Meinhardt, Riccardo de Lutio, Laura Leal-Taixé, Aljosa Osep</i>	65475
Curvature-aware Graph Attention for PDEs on Manifolds <i>Yunfeng Liao, Jiawen Guan, Xiucheng Li</i>	65500
Batch List-Decodable Linear Regression via Higher Moments <i>Ilias Diakonikolas, Daniel Kane, Sushrut Karmalkar, Sihan Liu, Thanasis Pittas</i>	65519
No-Regret is not enough! Bandits with General Constraints through Adaptive Regret Minimization <i>Martino Bernasconi, Matteo Castiglioni, Andrea Celli</i>	65540

Attributes Shape the Embedding Space of Face Recognition Models <i>Pierrick Leroy, Antonio Mastropietro, Marco Nurisso, Francesco Vaccarino</i>	65562
Avoiding spurious sharpness minimization broadens applicability of SAM <i>Sidak Pal Singh, Hossein Mobahi, Atish Agarwala, Yann Dauphin</i>	65586
The Courage to Stop: Overcoming Sunk Cost Fallacy in Deep Reinforcement Learning <i>Jiashun Liu, Johan Obando-Ceron, Pablo Samuel Castro, Aaron Courville, Ling Pan</i>	65604
Contradiction Retrieval via Contrastive Learning with Sparsity <i>Haikexu Xu, Zongyu Lin, Kai-Wei Chang, Yizhou Sun, Piotr Indyk</i>	65623
Active Learning of Deep Neural Networks via Gradient-Free Cutting Planes <i>Erica Zhang, Fangzhao Zhang, Mert Pilanci</i>	65652
Extractive Structures Learned in Pretraining Enable Generalization on Finetuned Facts <i>Jiahai Feng, Stuart Russell, Jacob Steinhardt</i>	65701
Adaptive Sensitivity Analysis for Robust Augmentation against Natural Corruptions in Image Segmentation <i>Laura Yu Zheng, Wenjie Wei, Tony Wu, Jacob Clements, Shreelekha Revankar, Andre Harrison, Yu Shen, Ming Lin</i>	65732
MTL-UE: Learning to Learn Nothing for Multi-Task Learning <i>Yi Yu, Song Xia, SIYUAN YANG, Chenqi Kong, Wenhan Yang, Shijian Lu, Yap-Peng Tan, Alex Kot</i>	65755
Fast Min-ϵ Segmented Regression using Constant-Time Segment Merging <i>Ansgar L��ber, Max Schlecht, Florian Schintke, Joel Witzke, Matthias Weidlich, Bj��rn Scheuermann</i>	65773
Rethinking Confidence Scores and Thresholds in Pseudolabeling-based SSL <i>Harit Vishwakarma, Yi Chen, Satya Sai Srinath Namburi GNVV, Sui Jiet Tay, Ramya Korlakai Vinayak, Frederic Sala</i>	65789
UniSim: A Unified Simulator for Time-Coarsened Dynamics of Biomolecules <i>Ziyang Yu, Wenbing Huang, Yang Liu</i>	65808
Outlier-Aware Post-Training Quantization for Discrete Graph Diffusion Models <i>Zheng Gong, Ying Sun</i>	65826
ReverB-SNN: Reversing Bit of the Weight and Activation for Spiking Neural Networks <i>Yufei Guo, Yuhan Zhang, Zhou Jie, Xiaode Liu, Xin Tong, Yuanpei Chen, Weihang Peng, Zhe Ma</i>	65846
Efficient Robust Conformal Prediction via Lipschitz-Bounded Networks <i>Thomas Massena, L��o And��ol, Thibaut Boissin, Franck Mamalet, Corentin Friedrich, Mathieu Serrurier, S��bastien Gerchinovitz</i>	65858
Communicating Activations Between Language Model Agents <i>Vignav Ramesh, Kenneth Li</i>	65881

Consensus Based Stochastic Optimal Control <i>Liyao Lyu, Jingrun Chen</i>	65904
A Cross Modal Knowledge Distillation & Data Augmentation Recipe for Improving Transcriptomics Representations through Morphological Features <i>Ihab Bendidi, Yassir El Mesbahi, Alisandra Kaye Denton, Karush Suri, Kian Kenyon-Dean, Auguste Genovesio, Emmanuel Noutahi</i>	65938
Advancing Personalized Learning with Neural Collapse for Long-Tail Challenge <i>Hanglei Hu, Yingying Guo, Zhikang Chen, Sen Cui, Fei Wu, Kun Kuang, Min Zhang, Bo Jiang</i>	65957
From Thousands to Billions: 3D Visual Language Grounding via Render-Supervised Distillation from 2D VLMs <i>Ang Cao, Sergio Arnaud, Oleksandr Maksymets, Jianing Yang, Ayush Jain, Ada Martin, Vincent-Pierre Berges, Paul McVay, Ruslan Partsey, Aravind Rajeswaran, Franziska Meier, Justin Johnson, Jeong Joon Park, Alexander Sax</i>	65971
What Makes a Good Feedforward Computational Graph? <i>Alex Vitvitskyi, João Guilherme Madeira Araújo, Marc Lackenby, Petar Veličković</i>	65988
Windows Agent Arena: Evaluating Multi-Modal OS Agents at Scale <i>Rogério Bonatti, Dan Zhao, Francesco Bonacci, Dillon Dupont, Sara Abdali, Yinheng Li, Yadong Lu, Justin Wagle, Kazuhito Koishida, Arthur Buckner, Lawrence Keunho Jang, Zheng Hui</i>	66005
Latent Variable Causal Discovery under Selection Bias <i>Haoyue Dai, Yiwen Qiu, Ignavier Ng, Xinshuai Dong, Peter Spirtes, Kun Zhang</i>	66042
Ensemble Distribution Distillation via Flow Matching <i>Jonggeon Park, Giung Nam, Hyunsu Kim, Jongmin Yoon, Juho Lee</i>	66060
Generating Hypotheses of Dynamic Causal Graphs in Neuroscience: Leveraging Generative Factor Models of Observed Time Series <i>Zachary C. Brown, David Carlson</i>	66082
Bayesian Active Learning for Bivariate Causal Discovery <i>Yuxuan Wang, Mingzhou Liu, Xinwei Sun, Wei Wang, Yizhou Wang</i>	66116
TabNAT: A Continuous-Discrete Joint Generative Framework for Tabular Data <i>Hengrui Zhang, Liancheng Fang, Qitian Wu, Philip S. Yu</i>	66137
Curriculum Learning for Biological Sequence Prediction: The Case of De Novo Peptide Sequencing <i>Xiang Zhang, Jiaqi Wei, Zijie Qiu, Sheng Xu, Nanqing Dong, ZhiQiang Gao, Siqi Sun</i>	66161
KinDEL: DNA-Encoded Library Dataset for Kinase Inhibitors <i>Benson Chen, Tomasz Danel, Gabriel H. S. Dreiman, Patrick J. McEnaney, Nikhil Jain, Kirill Novikov, Spurti Umesh Akki, Joshua L. Turnbull, Virja Atul Pandya, Boris P. Belotserkovskii, Jared Bryce Weaver, Ankita Biswas, Dat Nguyen, Kent Gorday, Mohammad Sultan, Nathaniel Stanley, Daniel M Whalen, Divya Kanichar, Christoph Klein, Emily Fox, R. Edward Watts</i>	66180

Optimizing Noise Distributions for Differential Privacy <i>Atefeh Gilani, Juan Felipe Gomez, Shahab Asoodeh, Flavio Calmon, Oliver Kosut, Lalitha Sankar</i>	66200
Optimal and Practical Batched Linear Bandit Algorithm <i>Sanghoon Yu, Min-hwan Oh</i>	66218
The Sparse-Plus-Low-Rank Quasi-Newton Method for Entropic-Regularized Optimal Transport <i>Chenrui Wang, Yixuan Qiu</i>	66242
When Diffusion Models Memorize: Inductive Biases in Probability Flow of Minimum-Norm Shallow Neural Nets <i>Chen Zeno, Hila Manor, Greg Ongie, Nir Weinberger, Tomer Michaeli, Daniel Soudry</i>	66262
Beyond Atoms: Enhancing Molecular Pretrained Representations with 3D Space Modeling <i>Shuqi Lu, Xiaohong Ji, Bohang Zhang, Lin Yao, Siyuan Liu, Zhifeng Gao, Linfeng Zhang, Guolin Ke</i>	66293
Dissecting Submission Limit in Desk-Rejections: A Mathematical Analysis of Fairness in AI Conference Policies <i>Yuefan Cao, Xiaoyu Li, Yingyu Liang, Zhizhou Sha, Zhenmei Shi, Zhao Song, Jiahao Zhang</i>	66307
Broadband Ground Motion Synthesis by Diffusion Model with Minimal Condition <i>Jaeheun Jung, Jaehyuk Lee, Chang-Hae Jung, Hanyoung Kim, Bosung Jung, Donghun Lee</i>	66330
Causal Abstraction Inference under Lossy Representations <i>Kevin Muiyuan Xia, Elias Bareinboim</i>	66354
Beyond CVaR: Leveraging Static Spectral Risk Measures for Enhanced Decision-Making in Distributional Reinforcement Learning <i>Mehrdad Moghimi, Hyejin Ku</i>	66365
Reducing Tool Hallucination via Reliability Alignment <i>Hongshen Xu, Zichen Zhu, Lei Pan, Zihan Wang, Su Zhu, Da Ma, Ruisheng Cao, Lu Chen, Kai Yu</i>	66388
TtBA: Two-third Bridge Approach for Decision-Based Adversarial Attack <i>Feiyang Wang, Xingquan Zuo, Hai Huang, Gang Chen</i>	66403
DocVXQA: Context-Aware Visual Explanations for Document Question Answering <i>Mohamed Ali Souibgui, Changkyu Choi, Andrey Barsky, Kangsoo Jung, Ernest Valveny, Dimosthenis Karatzas</i>	66421
GrokFormer: Graph Fourier Kolmogorov-Arnold Transformers <i>GuoguoAi, Guansong Pang, Hezhe Qiao, YuanGao, Hui Yan</i>	66442
LightningDrag: Lightning Fast and Accurate Drag-based Image Editing Emerging from Videos <i>Yujun Shi, Jun Hao Liew, Hanshu Yan, Vincent Y. F. Tan, Jiashi Feng</i>	66462
LEVIS: Large Exact Verifiable Input Spaces for Neural Networks <i>Mohamad Fares El Hajj Chehade, WENTING LI, Brian Wesley Bell, Russell Bent, Saif R. Kazi, Hao Zhu</i>	66475

Open-Det: An Efficient Learning Framework for Open-Ended Detection <i>Guiping Cao, Tao Wang, Wenjian Huang, Xiangyuan Lan, Jianguo Zhang, Dongmei Jiang</i>	66489
Learning the Electronic Hamiltonian of Large Atomic Structures <i>Chen Hao Xia, Manasa Kaniselvan, Alexandros Nikolaos Ziogas, Marko Mladenović, Rayen Mahjoub, Alexander Maeder, Mathieu Luisier</i>	66510
The Hidden Dimensions of LLM Alignment: A Multi-Dimensional Analysis of Orthogonal Safety Directions <i>Wenbo Pan, Zhichao Liu, Qiguang Chen, Xiangyang Zhou, Yu Haining, Xiaohua Jia</i>	66532
Rethinking Addressing in Language Models via Contextualized Equivariant Positional Encoding <i>Jiajun Zhu, Peihao Wang, Ruisi Cai, Jason D. Lee, Pan Li, Zhangyang Wang</i>	66552
Clipped SGD Algorithms for Performative Prediction: Tight Bounds for Stochastic Bias and Remedies <i>Qiang Li, Michal Yemini, Hoi To Wai</i>	66584
Bayesian Optimization from Human Feedback: Near-Optimal Regret Bounds <i>Aya Kayal, Sattar Vakili, Laura Toni, Da-shan Shiu, Alberto Bernacchia</i>	66607
Layer by Layer: Uncovering Hidden Representations in Language Models <i>Oscar Skean, Md Rifat Arefin, Dan Zhao, Niket Nikul Patel, Jalal Naghiyev, Yann LeCun, Ravid Shwartz-Ziv</i>	66627
Identification of Latent Confounders via Investigating the Tensor Ranks of the Nonlinear Observations <i>Zhengming Chen, Yewei Xia, Feng Xie, Jie Qiao, Zhifeng Hao, Ruichu Cai, Kun Zhang</i>	66649
On Exact Bit-level Reversible Transformers Without Changing Architecture <i>Guoqiang Zhang, JP Lewis, W. Bastiaan Kleijn</i>	66677
Tracking The Best Expert Privately <i>Hilal Asi, Vinod Raman, Aadirupa Saha</i>	66692
Online Learning in the Random-Order Model <i>Martino Bernasconi, Andrea Celli, Riccardo Colini Baldeschi, Federico Fusco, Stefano Leonardi, Matteo Russo</i>	66711
Collapse-Proof Non-Contrastive Self-Supervised Learning <i>Emanuele Sansone, Tim Lebailly, Tinne Tuytelaars</i>	66732
DreamDPO: Aligning Text-to-3D Generation with Human Preferences via Direct Preference Optimization <i>Zhenglin Zhou, Xiaobo Xia, Fan Ma, Hehe Fan, Yi Yang, Tat-Seng Chua</i>	66760
Relating Misfit to Gain in Weak-to-Strong Generalization Beyond the Squared Loss <i>Abhijeet Mulgund, Chirag Pabbaraju</i>	66782

Task-Aware Virtual Training: Enhancing Generalization in Meta-Reinforcement Learning for Out-of-Distribution Tasks <i>Jeongmo Kim, Yisak Park, Minung Kim, Seungyul Han</i>	66804
Learning Mixtures of Experts with EM: A Mirror Descent Perspective <i>Quentin Fruytier, Aryan Mokhtari, Sujay Sanghavi</i>	66833
Towards Understanding Catastrophic Forgetting in Two-layer Convolutional Neural Networks <i>Boqi Li, Youjun Wang, Weiwei Liu</i>	66870
CRANE: Reasoning with constrained LLM generation <i>Debangshu Banerjee, Tarun Suresh, Shubham Ugare, Sasa Misailovic, Gagandeep Singh</i>	66909
Sample, Scrutinize and Scale: Effective Inference-Time Search by Scaling Verification <i>Eric Zhao, Pranjal Awasthi, Sreenivas Gollapudi</i>	66931
Less is More: Federated Graph Learning with Alleviating Topology Heterogeneity from A Causal Perspective <i>Lele Fu, Bowen Deng, Sheng Huang, Tianchi Liao, Shirui Pan, Chuan Chen</i>	66969
Dueling Convex Optimization with General Preferences <i>Aadirupa Saha, Tomer Koren, Yishay Mansour</i>	66988
Computing Optimal Transport Maps and Wasserstein Barycenters Using Conditional Normalizing Flows <i>Gabriele Visentin, Patrick Cheridito</i>	67001
Oracle-MoE: Locality-preserving Routing in the Oracle Space for Memory-constrained Large Language Model Inference <i>Jixian Zhou, Fang Dong, Ruijun Huang, Hengjie Cao, Mengyi Chen, Yifeng Yang, Anrui Chen, Mingzhi Dong, Yujia Wang, Dongsheng Li, David A. Clifton, Qin Lv, Rui Zhu, Chun Zhang, Fan Yang, Tun Lu, Ning Gu, Li Shang</i>	67019
Efficient and Privacy-Preserving Soft Prompt Transfer for LLMs <i>Xun Wang, Jing Xu, Franziska Boenisch, Michael Backes, Christopher A. Choquette-Choo, Adam Dziedzic</i>	67037
Compute or Load KV Cache? Why Not Both? <i>Shuwei Jin, Xueshen Liu, Qingzhao Zhang, Zhuoqing Mao</i>	67059
EnsLoss: Stochastic Calibrated Loss Ensembles for Preventing Overfitting in Classification <i>Ben Dai</i>	67072
MixMin: Finding Data Mixtures via Convex Minimization <i>Anvith Thudi, Evianne Rovers, Yangjun Ruan, Tristan Thrush, Chris J. Maddison</i>	67093
Hierarchical Graph Tokenization for Molecule-Language Alignment <i>Yongqiang Chen, Quanming Yao, Juzheng Zhang, James Cheng, Yatao Bian</i>	67110

Discovering Latent Causal Graphs from Spatiotemporal Data <i>Kun Wang, Sumanth Varambally, Duncan Watson-Parris, Yian Ma, Rose Yu</i>	67137
Mixed-curvature decision trees and random forests <i>Philippe Chlenski, Quentin Chu, Raiyan R. Khan, Kaizhu Du, Antonio Khalil Moretti, Itsik Pe'er</i>	67174
Diffusion Counterfactual Generation with Semantic Abduction <i>Rajat R Rasal, Avinash Kori, Fabio De Sousa Ribeiro, Tian Xia, Ben Glocker</i>	67204
A Theoretical Study of (Hyper) Self-Attention through the Lens of Interactions: Representation, Training, Generalization <i>Muhammed Ustaomeroglu, Guannan Qu</i>	67232
Stochastic Forward-Backward Deconvolution: Training Diffusion Models with Finite Noisy Datasets <i>Haoye Lu, Qifan Wu, Yaoliang Yu</i>	67286
Towards Escaping from Class Dependency Modeling for Multi-Dimensional Classification <i>Teng Huang, Bin-Bin Jia, Min-Ling Zhang</i>	67314
Fast Incomplete Multi-view Clustering by Flexible Anchor Learning <i>Yalan Qin, Guorui Feng, Xinpeng Zhang</i>	67330
Scaling Inference-Efficient Language Models <i>Song Bian, Minghao Yan, Shivaram Venkataraman</i>	67341
MVA: Linear Attention with High-order Query-Keys Integration and Multi-level Vocabulary Decomposition <i>wang ning, Zekun Li, Tongxin Bai, Man Yao, Zhen Qin, Guoqi Li</i>	67362
Revisiting Chain-of-Thought in Code Generation: Do Language Models Need to Learn Reasoning before Coding? <i>Ren-Biao Liu, Anqi Li, Chaoding Yang, Hui Sun, Ming Li</i>	67387
Learning Initial Basis Selection for Linear Programming via Duality-Inspired Tripartite Graph Representation and Comprehensive Supervision <i>Anqi Lu, Junchi Yan</i>	67405
Mixture of Lookup Experts <i>Shibo Jie, Yehui Tang, Kai Han, Yitong Li, Duyu Tang, Zhi-Hong Deng, Yunhe Wang</i>	67422
SafetyAnalyst: Interpretable, Transparent, and Steerable Safety Moderation for AI Behavior <i>Jing-Jing Li, Valentina Pyatkin, Max Kleiman-Weiner, Liwei Jiang, Nouha Dziri, Anne Collins, Jana Schaich Borg, Maarten Sap, Yejin Choi, Sydney Levine</i>	67434
Right Time to Learn: Promoting Generalization via Bio-inspired Spacing Effect in Knowledge Distillation <i>Guanglong Sun, Hongwei Yan, Liyuan Wang, Qian Li, Bo Lei, Yi Zhong</i>	67456

Enhancing Statistical Validity and Power in Hybrid Controlled Trials: A Randomization Inference Approach with Conformal Selective Borrowing <i>Ke Zhu, Shu Yang, Xiaofei Wang</i>	67474
Generalization and Robustness of the Tilted Empirical Risk <i>Gholamali Aminian, Amir R. Asadi, Tian Li, Ahmad Beirami, Gesine Reinert, Samuel N. Cohen</i>	67502
How Contaminated Is Your Benchmark? Measuring Dataset Leakage in Large Language Models with Kernel Divergence <i>Hyeong Kyu Choi, Maxim Khanov, Hongxin Wei, Yixuan Li</i>	67545
D-Fusion: Direct Preference Optimization for Aligning Diffusion Models with Visually Consistent Samples <i>Zijing Hu, Fengda Zhang, Kun Kuang</i>	67562
Interpreting the Repeated Token Phenomenon in Large Language Models <i>Itay Yona, Ilia Shumailov, Jamie Hayes, Yossi Gandelsman</i>	67586
Automatic Differentiation of Optimization Algorithms with Time-Varying Updates <i>Sheheryar Mehmood, Peter Ochs</i>	67607
A Unified View on Learning Unnormalized Distributions via Noise-Contrastive Estimation <i>Jongha Jon Ryu, Abhin Shah, Gregory W. Wornell</i>	67629
Relational Conformal Prediction for Correlated Time Series <i>Andrea Cini, Alexander Jenkins, Danilo Mandic, Cesare Alippi, Filippo Maria Bianchi</i>	67660
Sparse Autoencoders, Again? <i>Yin Lu, Xuening Zhu, Tong He, David Wipf</i>	67677
Approximating Latent Manifolds in Neural Networks via Vanishing Ideals <i>Nico Pelleriti, Max Zimmer, Elias Samuel Wirth, Sebastian Pokutta</i>	67702
Faster and Stronger: When ANN-SNN Conversion Meets Parallel Spiking Calculation <i>Zecheng Hao, Qichao Ma, Kang Chen, Yi Zhang, Zhaofei Yu, Tiejun Huang</i>	67730
Hybrid Spiking Vision Transformer for Object Detection with Event Cameras <i>Qi Xu, Jie Deng, Jiangrong Shen, Biwu Chen, Huajin Tang, Gang Pan</i>	67743
Non-Asymptotic Length Generalization <i>Thomas Chen, Tengyu Ma, Zhiyuan Li</i>	67756
Safety-Polarized and Prioritized Reinforcement Learning <i>Ke Fan, Jinpeng Zhang, Xuefeng Zhang, Yunze Wu, Jingyu Cao, Yuan Zhou, Jianzhu Ma</i>	67818
Measuring In-Context Computation Complexity via Hidden State Prediction <i>Vincent Herrmann, Róbert Csordás, Jürgen Schmidhuber</i>	67843

Heads up! Large Language Models Can Perform Tasks Without Your Instruction via Selective Attention Head Masking <i>Senyu Han, Hongchuan Zeng, Kai Yu, Lu Chen</i>	67863
Invariance Makes LLM Unlearning Resilient Even to Unanticipated Downstream Fine-Tuning <i>Changsheng Wang, Yihua Zhang, Jinghan Jia, Parikshit Ram, Dennis Wei, Yuguang Yao, Soumyadeep Pal, Nathalie Baracaldo, Sijia Liu</i>	67889
Competitively Consistent Clustering <i>Niv Buchbinder, Roie Levin, Yue Yang</i>	67905
Diving into Self-Evolving Training for Multimodal Reasoning <i>Wei Liu, Junlong Li, Xiwen Zhang, Fan Zhou, Yu Cheng, Junxian He</i>	67923
Feature Shift Localization Network <i>Míriam Barrabés, Daniel Mas Montserrat, Kapal Dev, Alexander G. Ioannidis</i>	67938
From Weight-Based to State-Based Fine-Tuning: Further Memory Reduction on LoRA with Parallel Control <i>Chi Zhang, REN Lianhai, Jingpu Cheng, Qianxiao Li</i>	67965
Orient Anything: Learning Robust Object Orientation Estimation from Rendering 3D Models <i>Zehan Wang, Ziang Zhang, Tianyu Pang, Chao Du, Hengshuang Zhao, Zhou Zhao</i>	67980
Do Bayesian Neural Networks Actually Behave Like Bayesian Models? <i>Gábor Pituk, Vik Shirvaikar, Tom Rainforth</i>	67999
Learning Safe Control via On-the-Fly Bandit Exploration <i>Alexandre Capone, Ryan Kazuo Cosner, Aaron Ames, Sandra Hirche</i>	68038
UniMate: A Unified Model for Mechanical Metamaterial Generation, Property Prediction, and Condition Confirmation <i>Wangzhi Zhan, Jianpeng Chen, Dongqi Fu, Dawei Zhou</i>	68058
When to retrain a machine learning model <i>Florence Regol, Leo Schwinn, Kyle Sprague, Mark Coates, Thomas Markovich</i>	68074
Phase transitions for the existence of unregularized M-estimators in single index models <i>Takuya Koriyama, Pierre C Bellec</i>	68110
GoIRL: Graph-Oriented Inverse Reinforcement Learning for Multimodal Trajectory Prediction <i>Muleilan Pei, Shaoshuai Shi, Lu Zhang, Peiliang Li, Shaojie Shen</i>	68132
Retrieval-Augmented Perception: High-resolution Image Perception Meets Visual RAG <i>Wenbin Wang, Yongcheng Jing, Liang Ding, Yingjie Wang, Li Shen, Yong Luo, Bo Du, Dacheng Tao</i>	68147
Shielded Diffusion: Generating Novel and Diverse Images using Sparse Repellency <i>Michael Kirchhof, James Thornton, Louis Béthune, Pierre Ablin, Eugene Ndiaye, marco cuturi</i>	68165

Non-Asymptotic and Non-Lipschitzian Bounds on Optimal Values in Stochastic Optimization Under Heavy Tails <i>Jindong Tong, Hongcheng Liu, Johannes O. Royset</i>	68197
On the Adversarial Robustness of Multi-Kernel Clustering <i>Hao Yu, Weixuan Liang, KE LIANG, Suyuan Liu, Meng Liu, Xinwang Liu</i>	68214
Ex-VAD: Explainable Fine-grained Video Anomaly Detection Based on Visual-Language Models <i>Chao Huang, Yushu Shi, Jie Wen, Wei Wang, Yong Xu, Xiaochun Cao</i>	68236
DeepLayout: Learning Neural Representations of Circuit Placement Layout <i>Yuxiang Zhao, zhuomin chai, Xun Jiang, Qiang Xu, Runsheng Wang, Yibo Lin</i>	68248
BAME: Block-Aware Mask Evolution for Efficient N:M Sparse Training <i>Chenyi yang, Wenjie Nie, Yuxin Zhang, Yuhang Wu, Xiawu Zheng, GUANNAN JIANG, Rongrong Ji</i>	68259
Diffusion on Language Model Encodings for Protein Sequence Generation <i>Viacheslav Meshchaninov, Pavel Strashnov, Andrey Shevtsov, Fedor Nikolaev, Nikita Ivanisenko, Olga Kardymon, Dmitry Vetrov</i>	68269
No Metric to Rule Them All: Toward Principled Evaluations of Graph-Learning Datasets <i>Corinna Coupette, Jeremy Wayland, Emily Simons, Bastian Rieck</i>	68311
Speculate, then Collaborate: Fusing Knowledge of Language Models during Decoding <i>Ziyao Wang, Muneeza Azmat, Ang Li, Raya Horesh, Mikhail Yurochkin</i>	68341
FOCoOp: Enhancing Out-of-Distribution Robustness in Federated Prompt Learning for Vision-Language Models <i>Xinting Liao, Weiming Liu, Jiaming Qian, Pengyang Zhou, Jiahe Xu, Wenjie Wang, Chaochao Chen, Xiaolin Zheng, Tat-Seng Chua</i>	68353
Non-Stationary Predictions May Be More Informative: Exploring Pseudo-Labels with a Two-Phase Pattern of Training Dynamics <i>Hongbin Pei, Jingxin Hai, Yu Li, Huiqi Deng, Denghao Ma, Jie Ma, Pinghui Wang, Jing Tao, Xiaohong Guan</i>	68380
Latent Thought Models with Variational Bayes Inference-Time Computation <i>Deqian Kong, Minglu Zhao, Dehong Xu, Bo Pang, Shu Wang, Edouardo Honig, Zhangzhang Si, Chuan Li, Jianwen Xie, Sirui Xie, Ying Nian Wu</i>	68397
Speeding up Policy Simulation in Supply Chain RL <i>Vivek Farias, Joren Gijsbrechts, Aryan I. Khojandi, Tianyi Peng, Andrew Zheng</i>	68415
Convergence Analysis of Policy Gradient Methods with Dynamic Stochasticity <i>Alessandro Montenegro, Marco Mussi, Matteo Papini, Alberto Maria Metelli</i>	68432
Continuous Semi-Implicit Models <i>Longlin Yu, Jiajun Zha, Tong Yang, Tianyu Xie, Xiangyu Zhang, S.-H. Chan, Cheng Zhang</i>	68479

PaperBench: Evaluating AI's Ability to Replicate AI Research <i>Giulio Starace, Oliver Jaffe, Dane Sherburn, James Aung, Jun Shern Chan, Leon Maksin, Rachel Dias, Evan Mays, Benjamin Kinsella, Wyatt Thompson, Johannes Heidecke, Amelia Glaese, Tejal Patwardhan</i>	68505
Understanding Sharpness Dynamics in NN Training with a Minimalist Example: The Effects of Dataset Difficulty, Depth, Stochasticity, and More <i>Geonhui Yoo, Minhak Song, Chulhee Yun</i>	68536
Quadruple Attention in Many-body Systems for Accurate Molecular Property Predictions <i>Jiahua Rao, Dahao Xu, Wentao Wei, Yicong Chen, Mingjun Yang, Yuedong Yang</i>	68580
Interpolating Neural Network-Tensor Decomposition (INN-TD): a scalable and interpretable approach for large-scale physics-based problems <i>Jiachen Guo, Xiaoyu Xie, Chanwook Park, Hantao Zhang, Matthew J. Politis, Gino Domel, Wing Kam Liu</i>	68598
PIPA: Preference Alignment as Prior-Informed Statistical Estimation <i>Junbo Li, Zhangyang Wang, qiang liu</i>	68623
SKOLR: Structured Koopman Operator Linear RNN for Time-Series Forecasting <i>Yitian Zhang, Liheng Ma, Antonios Valkanias, Boris N. Oreshkin, Mark Coates</i>	68638
From Uncertain to Safe: Conformal Adaptation of Diffusion Models for Safe PDE Control <i>Peiyan Hu, Xiaowei Qian, Wenhao Deng, Rui Wang, Haodong Feng, Ruiqi Feng, Tao Zhang, Long Wei, Yue Wang, Zhi-Ming Ma, Tailin Wu</i>	68661
A Non-Asymptotic Convergent Analysis for Score-Based Graph Generative Model via a System of Stochastic Differential Equations <i>Junwei Su, Chuan Wu</i>	68686
LIFT the Veil for the Truth: Principal Weights Emerge after Rank Reduction for Reasoning-Focused Supervised Fine-Tuning <i>Zihang Liu, Tianyu Pang, Oleg Balabanov, Chaoqun Yang, Tianjin Huang, Lu Yin, Yaoqing Yang, Shiwei Liu</i>	68724
Projection Optimization: A General Framework for Multi-Objective and Multi-Group RLHF <i>Nuoya Xiong, Aarti Singh</i>	68749
Ultra Lowrate Image Compression with Semantic Residual Coding and Compression-aware Diffusion <i>Anle Ke, Xu Zhang, Tong Chen, Ming Lu, Chao Zhou, Jiawen Gu, Zhan Ma</i>	68791
Habitizing Diffusion Planning for Efficient and Effective Decision Making <i>Haofei Lu, Yifei Shen, Dongsheng Li, Junliang Xing, Dongqi Han</i>	68816
KEA: Keeping Exploration Alive by Proactively Coordinating Exploration Strategies <i>Shih-Min Yang, Martin Magnusson, Johannes A. Stork, Todor Stoyanov</i>	68837

FunBO: Discovering Acquisition Functions for Bayesian Optimization with FunSearch <i>Virginia Aglietti, Ira Ktena, Jessica Schrouff, Eleni Sgouritsa, Francisco Ruiz, Alan Malek, Alexis Bellot, Silvia Chiappa</i>	68853
Graph World Model <i>Tao Feng, Yexin Wu, Guanyu Lin, Jiaxuan You</i>	68876
Fast Tensor Completion via Approximate Richardson Iteration <i>Mehrdad Ghadiri, Matthew Fahrback, Yunbum Kook, Ali Jadbabaie</i>	68901
Adversarial Cooperative Rationalization: The Risk of Spurious Correlations in Even Clean Datasets <i>Wei Liu, Zhongyu Niu, Lang Gao, Zhiying Deng, Jun Wang, Haozhao Wang, Ruixuan Li</i>	68919
Strategic Planning: A Top-Down Approach to Option Generation <i>Max Ruiz Luyten, Antonin Berthon, Mihaela van der Schaar</i>	68940
Dimensionality Reduction on Complex Vector Spaces for Euclidean Distance with Dynamic Weights <i>Simone Moretti, Paolo Pellizzoni, Francesco Silvestri</i>	68985
K²IE: Kernel Method-based Kernel Intensity Estimators for Inhomogeneous Poisson Processes <i>Hideaki Kim, Tomoharu Iwata, Akinori Fujino</i>	69005
Nonparametric Modern Hopfield Models <i>Jerry Yao-Chieh Hu, Bo-Yu Chen, Dennis Wu, Feng Ruan, Han Liu</i>	69019
TypyBench: Evaluating LLM Type Inference for Untyped Python Repositories <i>Honghua Dong, Jiacheng Yang, Xun Deng, Yuhe Jiang, Gennady Pekhimenko, Fan Long, Xujie Si</i>	69057
Neighbour-Driven Gaussian Process Variational Autoencoders for Scalable Structured Latent Modelling <i>Xinxing Shi, Xiaoyu Jiang, Mauricio A Álvarez</i>	69085
AffectGPT: A New Dataset, Model, and Benchmark for Emotion Understanding with Multimodal Large Language Models <i>Zheng Lian, Haoyu Chen, Lan Chen, Haiyang Sun, Licai Sun, Yong Ren, Zebang Cheng, Bin Liu, Rui Liu, Xiaojiang Peng, Jiangyan Yi, Jianhua Tao</i>	69111
Demystifying Cost-Efficiency in LLM Serving over Heterogeneous GPUs <i>YOUHE JIANG, Fangcheng Fu, Xiaozhe Yao, Guoliang HE, Xupeng Miao, Ana Klimovic, Bin CUI, Binhang Yuan, Eiko Yoneki</i>	69133
OWLS: Scaling Laws for Multilingual Speech Recognition and Translation Models <i>William Chen, Jinchuan Tian, Yifan Peng, Brian Yan, Chao-Han Huck Yang, Shinji Watanabe</i>	69152

SOLD: Slot Object-Centric Latent Dynamics Models for Relational Manipulation Learning from Pixels <i>Malte Mosbach, Jan Niklas Ewertz, Angel Villar-Corrales, Sven Behnke</i>	69177
Arbitrarily-Conditioned Multi-Functional Diffusion for Multi-Physics Emulation <i>Da Long, Zhitong Xu, Guang Yang, Akil Narayan, Shandian Zhe</i>	69202
Improving the Continuity of Goal-Achievement Ability via Policy Self-Regularization for Goal-Conditioned Reinforcement Learning <i>Xudong Gong, Sen Yang, Feng Dawei, Kele Xu, Bo Ding, Huaimin Wang, Yong Dou</i>	69222
Unconstrained Robust Online Convex Optimization <i>Jiujia Zhang, Ashok Cutkosky</i>	69242
Nonlinear transformers can perform inference-time feature learning <i>Naoki Nishikawa, Yujin Song, Kazusato Oko, Denny Wu, Taiji Suzuki</i>	69273
Whoever Started the interference Should End It: Guiding Data-Free Model Merging via Task Vectors <i>Runxi Cheng, Feng Xiong, Yongxian Wei, Wanyun Zhu, Chun Yuan</i>	69305
Fluctuations of the largest eigenvalues of transformed spiked Wigner matrices <i>Aro Lee, Ji Oon Lee</i>	69328
Consensus Is All You Get: The Role of Attention in Transformers <i>Álvaro Rodríguez Abella, João Pedro Silvestre, Paulo Tabuada</i>	69358
Large Continual Instruction Assistant <i>Jingyang Qiao, zhizhong zhang, Xin Tan, Yanyun Qu, Shouhong Ding, Yuan Xie</i>	69369
When Dynamic Data Selection Meets Data Augmentation: Achieving Enhanced Training Acceleration <i>Suorong Yang, Peng Ye, Furao Shen, Dongzhan Zhou</i>	69394
Bayesian Neural Scaling Law Extrapolation with Prior-Data Fitted Networks <i>Dongwoo Lee, Dong Bok Lee, Steven Adriaensen, Juho Lee, Sung Ju Hwang, Frank Hutter, Seon Joo Kim, Hae Beom Lee</i>	69407
Variance-Reduced Forward-Reflected-Backward Splitting Methods for Nonmonotone Generalized Equations <i>Quoc Tran-Dinh</i>	69440
FlexiClip: Locality-Preserving Free-Form Character Animation <i>Anant Khandelwal</i>	69478
A Memory Efficient Randomized Subspace Optimization Method for Training Large Language Models <i>Yiming Chen, Yuan Zhang, Yin Liu, Kun Yuan, Zaiwen Wen</i>	69493

EARL-BO: Reinforcement Learning for Multi-Step Lookahead, High-Dimensional Bayesian Optimization <i>Mujin Cheon, Jay H Lee, Dong-Yeun Koh, Calvin Tsay</i>	69510
Test-time Adapted Reinforcement Learning with Action Entropy Regularization <i>Shoukai Xu, ZihaoLian, Mingkui Tan, Liu Liu, Zhong Zhang, Peilin Zhao</i>	69527
Beyond Induction Heads: In-Context Meta Learning Induces Multi-Phase Circuit Emergence <i>Gouki Minegishi, Hiroki Furuta, Shohei Taniguchi, Yusuke Iwasawa, Yutaka Matsuo</i>	69539
SUICA: Learning Super-high Dimensional Sparse Implicit Neural Representations for Spatial Transcriptomics <i>Qingtian Zhu, Yumin Zheng, Yuling Sang, Yifan Zhan, Ziyang Zhu, Jun Ding, Yinqiang Zheng</i>	69563
Audio Flamingo 2: An Audio-Language Model with Long-Audio Understanding and Expert Reasoning Abilities <i>Sreyan Ghosh, Zhifeng Kong, Sonal Kumar, S Sakshi, Jaehyeon Kim, Wei Ping, Rafael Valle, Dinesh Manocha, Bryan Catanzaro</i>	69578
Scalable Model Merging with Progressive Layer-wise Distillation <i>Jing Xu, Jiazheng Li, Jingzhao Zhang</i>	69626
RelGNN: Composite Message Passing for Relational Deep Learning <i>Tianlang Chen, Charilaos Kanatsoulis, Jure Leskovec</i>	69653
WMAdapter: Adding WaterMark Control to Latent Diffusion Models <i>Hai Ci, Yiren Song, Pei Yang, Jinheng Xie, Mike Zheng Shou</i>	69670
Vector Grimoire: Codebook-based Shape Generation under Raster Image Supervision <i>Marco Cipriano, Moritz Feuerpfeil, Gerard de Melo</i>	69689
ADHMR: Aligning Diffusion-based Human Mesh Recovery via Direct Preference Optimization <i>Wenhao Shen, Wanqi Yin, Xiaofeng Yang, Cheng Chen, Chaoyue Song, Zhongang Cai, Lei Yang, Hao Wang, Guosheng Lin</i>	69717
SWE-Lancer: Can Frontier LLMs Earn \$1 Million from Real-World Freelance Software Engineering? <i>Samuel Miserendino, Michele Wang, Tejal Patwardhan, Johannes Heidecke</i>	69729
LongVU: Spatiotemporal Adaptive Compression for Long Video-Language Understanding <i>Xiaoqian Shen, Yuniyang Xiong, Changsheng Zhao, Lemeng Wu, Jun Chen, Chenchen Zhu, Zechun Liu, Fanyi Xiao, Balakrishnan Varadarajan, Florian Bordes, Zhuang Liu, Hu Xu, Hyunwoo J. Kim, Bilge Soran, Raghuraman Krishnamoorthi, Mohamed Elhoseiny, Vikas Chandra</i>	69768
Activation by Interval-wise Dropout: A Simple Way to Prevent Neural Networks from Plasticity Loss <i>Sangyeon Park, Isaac Han, Seungwon Oh, KyungJoong Kim</i>	69786

From Jack of All Trades to Master of One: Specializing LLM-based Autoraters to a Test Set <i>Mara Finkelstein, Daniel Deutsch, Parker Riley, Juraj Juraska, Geza Kovacs, Markus Freitag</i>	69822
BlockDialect: Block-wise Fine-grained Mixed Format Quantization for Energy-Efficient LLM Inference <i>Wonsuk Jang, Thierry Tambe</i>	69840
LaMAGIC2: Advanced Circuit Formulations for Language Model-Based Analog Topology Generation <i>Chen-Chia Chang, Wan-Hsuan Lin, Yikang Shen, Yiran Chen, Xin Zhang</i>	69861
Emergent Symbolic Mechanisms Support Abstract Reasoning in Large Language Models <i>Yukang Yang, Declan Iain Campbell, Kaixuan Huang, Mengdi Wang, Jonathan D. Cohen, Taylor Whittington Webb</i>	69871
Efficient Time Series Processing for Transformers and State-Space Models through Token Merging <i>Leon Götze, Marcel Kollovieh, Stephan Günnemann, Leo Schwinn</i>	69906
Learning Mean Field Control on Sparse Graphs <i>Christian Fabian, Kai Cui, Heinz Koepl</i>	69927
Addressing Misspecification in Simulation-based Inference through Data-driven Calibration <i>Antoine Wehenkel, Juan L. Gamella, Ozan Sener, Jens Behrmann, Guillermo Sapiro, Joern-Henrik Jacobsen, marco cuturi</i>	69951
Guardians of Image Quality: Benchmarking Defenses Against Adversarial Attacks on Image Quality Metrics <i>Aleksandr Gushchin, Khaled Abud, Georgii Bychkov, Ekaterina Shumitskaya, Anna Chistyakova, Sergey Lavrushkin, Bader Rasheed, Kirill Malyshev, Dmitriy S. Vatolin, Anastasia Antsiferova</i>	69983
Latent Mamba Operator for Partial Differential Equations <i>Karn Tiwari, Niladri Dutta, N M Anoop Krishnan, Prathosh AP</i>	70010
video-SALMONN-o1: Reasoning-enhanced Audio-visual Large Language Model <i>Guangzhi Sun, Yudong Yang, Jimin Zhuang, Changli Tang, Yixuan Li, Wei Li, Zejun MA, Chao Zhang</i>	70039
Gradient Flow Provably Learns Robust Classifiers for Orthonormal GMMs <i>Hancheng Min, Rene Vidal</i>	70059
Maintaining Proportional Committees with Dynamic Candidate Sets <i>Chris Dong, Jannik Peters</i>	70118
FedPHA: Federated Prompt Learning for Heterogeneous Client Adaptation <i>Chengying Fang, Wenke Huang, Guancheng Wan, Yihao Yang, Mang Ye</i>	70135

Breaking Barriers: Combinatorial Algorithms for Non-Monotone Submodular Maximization with Sublinear Adaptivity and 1/e Approximation <i>Yixin Chen, Wenjing Chen, Alan Kuhnle</i>	70151
Gandalf the Red: Adaptive Security for LLMs <i>Niklas Pfister, Václav Volhejn, Manuel Knott, Santiago Arias, Julia Bazinska, Mykhailo Bichurin, Alan Y. Commike, Janet Darling, Peter Dienes, Matthew Fiedler, David Haber, Matthias Kraft, Marco Lancini, Max Mathys, Damian Pascual-Ortiz, Jakub Podolak, Adrià Romero-López, Kyriacos Shiarlis, Andreas Signer, Zolt Terek, Athanasios Theocharis, Daniel Timbrell, Samuel Trautwein, Samuel Watts, Yun-Han Wu, Mateo Rojas-Carulla</i>	70191
Modulated Diffusion: Accelerating Generative Modeling with Modulated Quantization <i>Weizhi Gao, Zhichao Hou, Junqi Yin, Feiyi Wang, Linyu Peng, Xiaorui Liu</i>	70236
OV-MER: Towards Open-Vocabulary Multimodal Emotion Recognition <i>Zheng Lian, Haiyang Sun, Licai Sun, Haoyu Chen, Lan Chen, Hao Gu, Zhuofan Wen, Shun Chen, Zhang Siyuan, Hailiang Yao, Bin Liu, Rui Liu, Shan Liang, Ya Li, Jiangyan Yi, Jianhua Tao</i>	70262
Contrastive Visual Data Augmentation <i>Yu Zhou, Bingxuan Li, Tang Mohan, Xiaomeng Jin, Te-Lin Wu, Kuan-Hao Huang, Heng Ji, Kai-Wei Chang, Nanyun Peng</i>	70298
Thickness-aware E(3)-Equivariant 3D Mesh Neural Networks <i>Sungwon Kim, Namkyeong Lee, Yunyoung Doh, Seungmin Shin, Guimok Cho, Seung-Won Jeon, Sangkook Kim, Chanyoung Park</i>	70313
Reinforcement Learning Control of a Physical Robot Device for Assisted Human Walking without a Simulator <i>Junmin Zhong, Emiliano Quinones Yumbla, Seyed Yousef Soltanian, Ruofan Wu, Wenlong Zhang, Jennie Si</i>	70331
EEG-Language Pretraining for Highly Label-Efficient Clinical Phenotyping <i>Sam Gijsen, Kerstin Ritter</i>	70358
Plan-and-Act: Improving Planning of Agents for Long-Horizon Tasks <i>Lutfi Eren Erdogan, Nicholas Lee, Sehoon Kim, Suhong Moon, Hiroki Furuta, Gopala Anumanchipalli, Kurt Keutzer, Amir Gholami</i>	70383
Dynamical phases of short-term memory mechanisms in RNNs <i>Bariscan Kurtkaya, Fatih Dinc, Mert Yuksekgonul, Marta Blanco-Pozo, Ege Cirakman, Mark Schnitzer, Yucel Yemez, Hidenori Tanaka, Peng Yuan, Nina Miolane</i>	70427
Improved Regret Analysis in Gaussian Process Bandits: Optimality for Noiseless Reward, RKHS norm, and Non-Stationary Variance <i>Shogo Iwazaki, Shion Takeno</i>	70458
PatchPilot: A Cost-Efficient Software Engineering Agent with Early Attempts on Formal Verification <i>Hongwei Li, Yuheng Tang, Shiqi Wang, Wenbo Guo</i>	70489

Ca2-VDM: Efficient Autoregressive Video Diffusion Model with Causal Generation and Cache Sharing <i>Kaifeng Gao, Jiaxin Shi, Hanwang Zhang, Chunping Wang, Jun Xiao, Long Chen</i>	70509
Improving Compositional Generation with Diffusion Models Using Lift Scores <i>Chenning Yu, Sicun Gao</i>	70525
Unified K-Means Clustering with Label-Guided Manifold Learning <i>Qianqian Wang, Mengping Jiang, Zhengming Ding, Quanyue Gao</i>	70553
Solving Satisfiability Modulo Counting Exactly with Probabilistic Circuits <i>Jinzhao Li, Nan Jiang, Yexiang Xue</i>	70567
Learning with Expected Signatures: Theory and Applications <i>Lorenzo Lucchese, Mikko S. Pakkanen, Almut E. D. Veraart</i>	70589
Rethink the Role of Deep Learning towards Large-scale Quantum Systems <i>Yusheng Zhao, Chi Zhang, Yuxuan Du</i>	70650
AdaPTS: Adapting Univariate Foundation Models to Probabilistic Multivariate Time Series Forecasting <i>Abdelhakim Benezhab, Vasilii Feofanov, Giuseppe Paolo, Albert Thomas, Maurizio Filippone, Balázs Kégl</i>	70678
KernelBench: Can LLMs Write Efficient GPU Kernels? <i>Anne Ouyang, Simon Guo, Simran Arora, Alex L Zhang, William Hu, Christopher Re, Azalia Mirhoseini</i>	70696
Disentangling Invariant Subgraph via Variance Contrastive Estimation under Distribution Shifts <i>Haoyang Li, Xin Wang, Xueling Zhu, Weigao Wen, Wenwu Zhu</i>	70756
Theoretical Performance Guarantees for Partial Domain Adaptation via Partial Optimal Transport <i>Jayadev Naram, Fredrik Hellström, Ziming Wang, Rebecca Jörnsten, Giuseppe Durisi</i>	70772
A Non-isotropic Time Series Diffusion Model with Moving Average Transitions <i>Chenxi Wang, Linxiao Yang, Zhixian Wang, Liang Sun, Yi Wang</i>	70791
TUMTraf VideoQA: Dataset and Benchmark for Unified Spatio-Temporal Video Understanding in Traffic Scenes <i>Xingcheng Zhou, Konstantinos Larintzakis, Hao Guo, Walter Zimmer, Mingyu Liu, Hu Cao, Jiajie Zhang, Venkatnarayanan Lakshminarasimhan, Leah Strand, Alois Knoll</i>	70814
GaussMark: A Practical Approach for Structural Watermarking of Language Models <i>Adam Block, Alexander Rakhlin, Ayush Sekhari</i>	70840
Robust and Conjugate Spatio-Temporal Gaussian Processes <i>William Laplante, Matias Altamirano, Andrew B. Duncan, Jeremias Knoblauch, Francois-Xavier Briol</i>	70903

Large Language Models are Demonstration Pre-Selectors for Themselves <i>Jiarui Jin, Yuwei Wu, Haoxuan Li, Xiaoting He, Weinan Zhang, Yiming Yang, Yong Yu, Jun Wang, Mengyue Yang</i>	70934
RobustLight: Improving Robustness via Diffusion Reinforcement Learning for Traffic Signal Control <i>Mingyuan Li, Jiahao Wang, Guangsheng Yu, Xu Wang, Qianrun Chen, Wei Ni, Lixiang Li, Haipeng Peng</i>	70964
Lexico: Extreme KV Cache Compression via Sparse Coding over Universal Dictionaries <i>Junhyuck Kim, Jongho Park, Jaewoong Cho, Dimitris Papailiopoulos</i>	70987
Hyper-Transforming Latent Diffusion Models <i>Ignacio Peis, Batuhan Koyuncu, Isabel Valera, Jes Frellsen</i>	71003
Geometric Algebra Planes: Convex Implicit Neural Volumes <i>Irmak Sivgin, Sara Fridovich-Keil, Gordon Wetzstein, Mert Pilanci</i>	71023
Principal-Agent Bandit Games with Self-Interested and Exploratory Learning Agents <i>Junyan Liu, Lillian J. Ratliff</i>	71049
Convex Markov Games: A New Frontier for Multi-Agent Reinforcement Learning <i>Ian Gemp, Andreas Alexander Haupt, Luke Marris, Siqi Liu, Georgios Piliouras</i>	71097
On the Impact of Performative Risk Minimization for Binary Random Variables <i>Nikita Tsoy, Ivan Kirev, Negin Rahimiyazdi, Nikola Konstantinov</i>	71121
LlavaGuard: An Open VLM-based Framework for Safeguarding Vision Datasets and Models <i>Lukas Helff, Felix Friedrich, Manuel Brack, Kristian Kersting, Patrick Schramowski</i>	71147
Policy Gradient with Tree Expansion <i>Gal Dalal, Assaf Hallak, Gagan Thoppe, Shie Mannor, Gal Chechik</i>	71171
Adaptive Flow Matching for Resolving Small-Scale Physics <i>Stathi Fotiadis, Noah D Brenowitz, Tomas Geffner, Yair Cohen, Michael Pritchard, Arash Vahdat, Morteza Mardani</i>	71198
Symmetric Reinforcement Learning Loss for Robust Learning on Diverse Tasks and Model Scales <i>Ju-Seung Byun, Andrew Perrault</i>	71231
Exact Upper and Lower Bounds for the Output Distribution of Neural Networks with Random Inputs <i>Andrey Kofnov, Daniel Kapla, Ezio Bartocci, Efstathia Bura</i>	71254
Quantifying Memory Utilization with Effective State-Size <i>Rom Parnichkun, Neehal Tumma, Armin W Thomas, Alessandro Moro, Qi An, Taiji Suzuki, Atsushi Yamashita, Michael Poli, Stefano Massaroli</i>	71279
When can in-context learning generalize out of task distribution? <i>Chase Goddard, Lindsay M. Smith, Vudtiwat Ngampruetikorn, David J. Schwab</i>	71338

Reward Translation via Reward Machine in Semi-Alignable MDPs <i>Yun Hua, Haosheng Chen, Wenhao Li, Bo Jin, Baoxiang Wang, Hongyuan Zha, Xiangfeng Wang</i>	71353
ProSec: Fortifying Code LLMs with Proactive Security Alignment <i>Xiangzhe Xu, Zian Su, Jinyao Guo, Kaiyuan Zhang, Zhenting Wang, Xiangyu Zhang</i>	71373
Self-supervised Adversarial Purification for Graph Neural Networks <i>Woohyun Lee, Hogun Park</i>	71389
LieRE: Lie Rotational Positional Encodings <i>Sophie Ostmeier, Brian Axelrod, Maya Varma, Michael Moseley, Akshay S Chaudhari, Curtis Langlotz</i>	71410
VideoJAM: Joint Appearance-Motion Representations for Enhanced Motion Generation in Video Models <i>Hila Chefer, Uriel Singer, Amit Zohar, Yuval Kirstain, Adam Polyak, Yaniv Taigman, Lior Wolf, Shelly Sheynin</i>	71427
ABNet: Adaptive explicit-Barrier Net for Safe and Scalable Robot Learning <i>Wei Xiao, Tsun-Hsuan Wang, Chuang Gan, Daniela Rus</i>	71449
Principled Algorithms for Optimizing Generalized Metrics in Binary Classification <i>Anqi Mao, Mehryar Mohri, Yutao Zhong</i>	71468
Fixing the Double Penalty in Data-Driven Weather Forecasting Through a Modified Spherical Harmonic Loss Function <i>Christopher Subich, Syed Zahid Husain, Leo Separovic, Jing Yang</i>	71490
Fleet of Agents: Coordinated Problem Solving with Large Language Models <i>Lars Henning Klein, Nearchos Potamitis, Roland Aydin, Robert West, Caglar Gulcehre, Akhil Arora</i>	71511
Rethinking the Stability-Plasticity Trade-off in Continual Learning from an Architectural Perspective <i>Aojun Lu, Hangjie Yuan, Tao Feng, Yanan Sun</i>	71545
Generative Point Cloud Registration <i>Haobo Jiang, Jin Xie, jian Yang, liang yu, Jianmin Zheng</i>	71560
Robust Autonomy Emerges from Self-Play <i>Marco Francis Cusumano-Towner, David Hafner, Alexander Hertzberg, Brody Huval, Aleksei Petrenko, Eugene Vinitzky, Erik Wijmans, Taylor W. Killian, Stuart Bowers, Ozan Sener, Philipp Kraehenbuehl, Vladlen Koltun</i>	71574
Improving the Variance of Differentially Private Randomized Experiments through Clustering <i>Adel Javanmard, Vahab Mirrokni, Jean Pouget-Abadie</i>	71602
Approximate Differential Privacy of the ℓ_2 Mechanism <i>Matthew Joseph, Alex Kulesza, Alexander Yu</i>	71637

Representations Shape Weak-to-Strong Generalization: Theoretical Insights and Empirical Predictions <i>Yihao Xue, Jiping Li, Baharan Mirzasoleiman</i>	71653
Diffusion Sampling Correction via Approximately 10 Parameters <i>Guangyi Wang, Wei Peng, Iijiang Li, Wenyu Chen, Yuren Cai, Song-Zhi Su</i>	71693
Importance Corrected Neural JKO Sampling <i>Johannes Hertrich, Robert Gruhlke</i>	71714
Learning to Quantize for Training Vector-Quantized Networks <i>Peijia Qin, Jianguo Zhang</i>	71751
Perception in Reflection <i>Yana Wei, Liang Zhao, Kangheng Lin, En Yu, Yuang Peng, Runpei Dong, Jianjian Sun, Haoran Wei, Zheng Ge, Xiangyu Zhang, Vishal M. Patel</i>	71764
Catching Two Birds with One Stone: Reward Shaping with Dual Random Networks for Balancing Exploration and Exploitation <i>Haozhe Ma, Fangling Li, Jing Yu Lim, Zhengding Luo, Thanh Vinh Vo, Tze-Yun Leong</i>	71783
A Stronger Mixture of Low-Rank Experts for Fine-Tuning Foundation Models <i>Mengyang Sun, Yihao Wang, Tao Feng, Dan Zhang, Yifan Zhu, Jie Tang</i>	71800
BaWA: Automatic Optimizing Pruning Metric for Large Language Models with Balanced Weight and Activation <i>Lian Liu, Xiangdong Zhao, Guanchen Li, Dong Li, Mengdi Wang, Yinhe Han, Xiaowei Li, ying wang</i>	71816
TANGO: Clustering with Typicality-Aware Nonlocal Mode-Seeking and Graph-Cut Optimization <i>Haowen Ma, Zhiguo Long, Hua Meng</i>	71834
Understanding Model Ensemble in Transferable Adversarial Attack <i>Wei Yao, Zeliang Zhang, Huayi Tang, Yong Liu</i>	71853
Controlled Generation with Equivariant Variational Flow Matching <i>Floor Eijkelboom, Heiko Zimmermann, Sharvaree Vadgama, Erik J Bekkers, Max Welling, Christian A. Naesseth, Jan-Willem van de Meent</i>	71889
Think Twice, Act Once: A Co-Evolution Framework of LLM and RL for Large-Scale Decision Making <i>Xu Wan, Wenyue Xu, Chao Yang, Mingyang Sun</i>	71902
How Transformers Learn Regular Language Recognition: A Theoretical Study on Training Dynamics and Implicit Bias <i>Ruiquan Huang, Yingbin Liang, Jing Yang</i>	71921
Representation Preserving Multiclass Agnostic to Realizable Reduction <i>Steve Hanneke, Qinglin Meng, Amirreza Shaeiri</i>	71955

Divide and Conquer: Exploring Language-centric Tree Reasoning for Video Question-Answering <i>Zhaohe Liao, Jiangtong Li, Siyu Sun, Qingyang Liu, Fengshun Xiao, Tianjiao Li, Qiang Zhang, Guang Chen, Li Niu, Changjun Jiang, Liqing Zhang</i>	71969
Incorporating Arbitrary Matrix Group Equivariance into KANs <i>Lexiang Hu, Yisen Wang, Zhouchen Lin</i>	71986
Adversarial Inception Backdoor Attacks against Reinforcement Learning <i>Ethan Rathbun, Alina Oprea, Christopher Amato</i>	72003
Promoting Ensemble Diversity with Interactive Bayesian Distributional Robustness for Fine-tuning Foundation Models <i>Ngoc-Quan Pham, Tuan Truong, Quyen Tran, Tan Minh Nguyen, Dinh Phung, Trung Le</i>	72027
On the Emergence of Position Bias in Transformers <i>Xinyi Wu, Yifei Wang, Stefanie Jegelka, Ali Jadbabaie</i>	72045
EcoMapper: Generative Modeling for Climate-Aware Satellite Imagery <i>Muhammed Goktepe, Amir hossein Shamseddin, Erencan Uysal, Javier Muinelo Monteagudo, Lukas Drees, Aysim Toker, Senthold Asseng, Malte von Bloh</i>	72071
Inverse Reinforcement Learning with Switching Rewards and History Dependency for Characterizing Animal Behaviors <i>Jingyang Ke, Feiyang Wu, Jiayi Wang, Jeffrey Markowitz, Anqi Wu</i>	72092
Imitation Learning from a Single Temporally Misaligned Video <i>William Huey, Huaxiaoyue Wang, Anne Wu, Yoav Artzi, Sanjiban Choudhury</i>	72110
BiMaCoSR: Binary One-Step Diffusion Model Leveraging Flexible Matrix Compression for Real Super-Resolution <i>Kai Liu, Kaicheng Yang, Zheng Chen, Zhiteng Li, Yong Guo, Wenbo Li, Linghe Kong, Yulun Zhang</i>	72138
Transformer-Based Spatial-Temporal Counterfactual Outcomes Estimation <i>He Li, Haoang Chi, Mingyu Liu, Wanrong Huang, Liyang Xu, Wenjing Yang</i>	72148
Overtrained Language Models Are Harder to Fine-Tune <i>Jacob Mitchell Springer, Sachin Goyal, Kaiyue Wen, Tanishq Kumar, Xiang Yue, Sadhika Malladi, Graham Neubig, Aditi Raghunathan</i>	72172
MTSTRec: Multimodal Time-Aligned Shared Token Recommender <i>Ming-Yi Hong, Yen-Jung Hsu, Miao-Chen Chiang, Che Lin</i>	72243
Scaling Value Iteration Networks to 5000 Layers for Extreme Long-Term Planning <i>Yuhui Wang, Qingyuan Wu, Dylan R. Ashley, Francesco Faccio, Weida Li, Chao Huang, Jürgen Schmidhuber</i>	72265

Adaptive Message Passing: A General Framework to Mitigate Oversmoothing, Oversquashing, and Underreaching <i>Federico Errica, Henrik Christiansen, Viktor Zaverkin, Takashi Maruyama, Mathias Niepert, Francesco Alesiani</i>	72288
From Individual Experience to Collective Evidence: A Reporting-Based Framework for Identifying Systemic Harms <i>Jessica Dai, Paula Gradu, Inioluwa Deborah Raji, Benjamin Recht</i>	72314
Beyond Low-rank Decomposition: A Shortcut Approach for Efficient On-Device Learning <i>Le-Trung Nguyen, Aël Quéllennec, Van-Tam Nguyen, Enzo Tartaglione</i>	72335
Parrot: Multilingual Visual Instruction Tuning <i>Hai-Long Sun, Da-Wei Zhou, Yang Li, Shiyin Lu, Chao Yi, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, De-Chuan Zhan, Han-Jia Ye</i>	72350
NEAR: Neural Electromagnetic Array Response <i>Yinyan Bu, Jiajie Yu, Kai Zheng, Xinyu Zhang, Piya Pal</i>	72374
TuCo: Measuring the Contribution of Fine-Tuning to Individual Responses of LLMs <i>Felipe Pinto Coelho Nuti, Tim Franzmeyer, Joao F. Henriques</i>	72400
SGD Jittering: A Training Strategy for Robust and Accurate Model-Based Architectures <i>Peimeng Guan, Mark A. Davenport</i>	72440
MME-CoT: Benchmarking Chain-of-Thought in Large Multimodal Models for Reasoning Quality, Robustness, and Efficiency <i>Dongzhi Jiang, Renrui Zhang, Ziyu Guo, Yanwei Li, Yu Qi, Xinyan Chen, Lihui Wang, Jianhan Jin, Claire Guo, Shen Yan, Bo Zhang, Chaoyou Fu, Peng Gao, Hongsheng Li</i>	72459
Variational Control for Guidance in Diffusion Models <i>Kushagra Pandey, Farrin Marouf Sofian, Felix Draxler, Theofanis Karaletsos, Stephan Mandt</i>	72497
Isolated Causal Effects of Natural Language <i>Victoria Lin, Louis-Philippe Morency, Eli Ben-Michael</i>	72523
Beyond Matryoshka: Revisiting Sparse Coding for Adaptive Representation <i>Tiansheng Wen, Yifei Wang, Zequn Zeng, Zhong Peng, Yudi Su, Xinyang Liu, Bo Chen, Hongwei Liu, Stefanie Jegelka, Chenyu You</i>	72546
CursorCore: Assist Programming through Aligning Anything <i>Hao Jiang, Qi Liu, Rui Li, Shengyu Ye, Shijin Wang</i>	72565
On the Importance of Gaussianizing Representations <i>Daniel Eftekhari, Vardan Papyan</i>	72603
Calibrated Physics-Informed Uncertainty Quantification <i>Vignesh Gopakumar, Ander Gray, Lorenzo Zanisi, Timothy Nunn, Daniel Giles, Matt Kusner, Stanislas Pamela, Marc Peter Deisenroth</i>	72623

Learnware Specification via Dual Alignment <i>Wei Chen, Jun-Xiang Mao, Xiaozheng Wang, Min-Ling Zhang</i>	72662
Settling the Maximin Share Fairness for Scheduling among Groups of Machines <i>Bo Li, Fangxiao Wang, Xing Shiji</i>	72679
Self-Improving Language Models for Evolutionary Program Synthesis: A Case Study on ARC-AGI <i>Julien Pourcel, Cédric Colas, Pierre-Yves Oudeyer</i>	72690
Unified Screening for Multiple Diseases <i>Yiğit Narter, Alihan Hüyük, Mihaela van der Schaar, Cem Tekin</i>	72720
LightGTS: A Lightweight General Time Series Forecasting Model <i>Yihang Wang, Yuying Qiu, Peng Chen, Yang Shu, Zhongwen Rao, Lujia Pan, Bin Yang, Chenjuan Guo</i>	72753
Editable Noise Map Inversion: Encoding Target-image into Noise For High-Fidelity Image Manipulation <i>Mingyu Kang, Yong Suk Choi</i>	72771
DRAG: Data Reconstruction Attack using Guided Diffusion <i>Wa-Kin Lei, Jun-Cheng Chen, Shang-Tse Chen</i>	72788
Flexible Tails for Normalizing Flows <i>Tennessee Hickling, Dennis Prangle</i>	72807
Thinking LLMs: General Instruction Following with Thought Generation <i>Tianhao Wu, Janice Lan, Weizhe Yuan, Jiantao Jiao, Jason E Weston, Sainbayar Sukhbaatar</i>	72831
CtrlSynth: Controllable Image Text Synthesis for Data-Efficient Multimodal Learning <i>Qingqing Cao, Mahyar Najibi, Sachin Mehta</i>	72857
CLOVER: Cross-Layer Orthogonal Vectors Pruning <i>Fanxu Meng, Pingzhi Tang, Fan Jiang, Muhan Zhang</i>	72877
SparseLoRA: Accelerating LLM Fine-Tuning with Contextual Sparsity <i>Samir Khaki, Xiuyu Li, Junxian Guo, Ligeng Zhu, Konstantinos N. Plataniotis, Amir Yazdanbakhsh, Kurt Keutzer, Song Han, Zhijian Liu</i>	72896
ETTA: Elucidating the Design Space of Text-to-Audio Models <i>Sang-gil Lee, Zhifeng Kong, Arushi Goel, Sungwon Kim, Rafael Valle, Bryan Catanzaro</i>	72912
No Free Lunch from Random Feature Ensembles: Scaling Laws and Near-Optimality Conditions <i>Benjamin Samuel Ruben, William Lingxiao Tong, Hamza Tahir Chaudhry, Cengiz Pehlevan</i>	72936
Survival Analysis via Density Estimation <i>Hiroki Yanagisawa, Shunta Akiyama</i>	72963

BounDr.E: Predicting Drug-likeness via Biomedical Knowledge Alignment and EM-like One-Class Boundary Optimization <i>Dongmin Bang, Inyoung Sung, Yinhua Piao, Sangseon Lee, Sun Kim</i>	72993
Equivalence is All: A Unified View for Self-supervised Graph Learning <i>Yejiang Wang, Yuhai Zhao, Zhengkui Wang, Ling Li, Jiapu Wang, Fangting Li, Miaomiao Huang, Shirui Pan, Xingwei Wang</i>	73029
GuidedQuant: Large Language Model Quantization via Exploiting End Loss Guidance <i>Jinuk Kim, Marwa El Halabi, Wonpyo Park, Clemens JS Schaefer, Deokjae Lee, Yeonhong Park, Jae W. Lee, Hyun Oh Song</i>	73043
Scaling Laws for Upcycling Mixture-of-Experts Language Models <i>Seng Pei Liew, Takuya Kato, Sho Takase</i>	73070
Cross-environment Cooperation Enables Zero-shot Multi-agent Coordination <i>Kunal Jha, Wilka Carvalho, Yancheng Liang, Simon Shaolei Du, Max Kleiman-Weiner, Natasha Jaques</i>	73093
Transfer Learning for Nonparametric Contextual Dynamic Pricing <i>Fan Wang, Feiyu Jiang, Zifeng Zhao, Yi Yu</i>	73116
NegMerge: Sign-Consensual Weight Merging for Machine Unlearning <i>Hyo Seo Kim, Dongyoon Han, Junsuk Choe</i>	73150
A Near-Optimal Single-Loop Stochastic Algorithm for Convex Finite-Sum Coupled Compositional Optimization <i>Bokun Wang, Tianbao Yang</i>	73169
Adversarial Combinatorial Semi-bandits with Graph Feedback <i>Yuxiao Wen</i>	73200
Finite-Time Global Optimality Convergence in Deep Neural Actor-Critic Methods for Decentralized Multi-Agent Reinforcement Learning <i>Zhiyao Zhang, Myeung Suk Oh, FNU Hairi, Ziyue Luo, Alvaro Velasquez, Jia Liu</i>	73218
LARM: Large Auto-Regressive Model for Long-Horizon Embodied Intelligence <i>Zhuoling Li, Xiaogang Xu, Zhenhua Xu, Ser-Nam Lim, Hengshuang Zhao</i>	73243
Prediction via Shapley Value Regression <i>Amr Alkhatib, Roman Bresson, Henrik Boström, Michalis Vazirgiannis</i>	73253
Can Large Language Models Understand Intermediate Representations in Compilers? <i>Hailong Jiang, Jianfeng Zhu, Yao Wan, Bo Fang, Hongyu Zhang, Ruoming Jin, Qiang Guan</i>	73299
You Get What You Give: Reciprocally Fair Federated Learning <i>Aniket Murhekar, Jiaxin Song, Parnian Shahkar, Bhaskar Ray Chaudhury, Ruta Mehta</i>	73321
SENSEI: Semantic Exploration Guided by Foundation Models to Learn Versatile World Models <i>Cansu Sancaktar, Christian Gumbsch, Andrii Zadaianchuk, Pavel Kolev, Georg Martius</i>	73343

Foundation Model Insights and a Multi-Model Approach for Superior Fine-Grained One-shot Subset Selection <i>Zhijing Wan, Zhixiang Wang, Zheng Wang, Xin Xu, Shin'ichi Satoh</i>	73376
ERICT: Enhancing Robustness by Identifying Concept Tokens in Zero-Shot Vision Language Models <i>Xinpeng Dong, Min Zhang, Didi Zhu, Ye Jun Jian, Zhang Keli, Aimin Zhou, Fei Wu, Kun Kuang</i>	73394
KVTuner: Sensitivity-Aware Layer-Wise Mixed-Precision KV Cache Quantization for Efficient and Nearly Lossless LLM Inference <i>Xing Li, Zeyu XING, Yiming Li, Linping Qu, Hui-Ling Zhen, Yiwu Yao, Wulong Liu, Sinno Jialin Pan, Mingxuan Yuan</i>	73410
Beyond One-Hot Labels: Semantic Mixing for Model Calibration <i>Haoyang Luo, Linwei Tao, Mingjing Dong, Chang Xu</i>	73445
Goal-Oriented Skill Abstraction for Offline Multi-Task Reinforcement Learning <i>Jinmin He, Kai Li, Yifan Zang, Haobo Fu, QIANG FU, Junliang Xing, Jian Cheng</i>	73467
Backdoor Attacks in Token Selection of Attention Mechanism <i>Yunjuan Wang, Raman Arora</i>	73481
Exploring and Mitigating Adversarial Manipulation of Voting-Based Leaderboards <i>Yangsibo Huang, Milad Nasr, Anastasios Nikolas Angelopoulos, Nicholas Carlini, Wei-Lin Chiang, Christopher A. Choquette-Choo, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Ken Liu, Ion Stoica, Florian Tramèr, Chiyuan Zhang</i>	73527
Generalists vs. Specialists: Evaluating LLMs on Highly-Constrained Biophysical Sequence Optimization Tasks <i>Angelica Chen, Samuel Don Stanton, Frances Ding, Robert G Alberstein, Andrew Martin Watkins, Richard Bonneau, Vladimir Gligorijevic, Kyunghyun Cho, Nathan C. Frey</i>	73545
Fourier Position Embedding: Enhancing Attention's Periodic Extension for Length Generalization <i>Ermo Hua, Che Jiang, Xingtai Lv, Kaiyan Zhang, Youbang Sun, Yuchen Fan, Xuekai Zhu, Biqing Qi, Ning Ding, Bowen Zhou</i>	73589
Let LLM Tell What to Prune and How Much to Prune <i>Mingzhe Yang, Sihao Lin, Changlin Li, Xiaojun Chang</i>	73607
Learning to Keep a Promise: Scaling Language Model Decoding Parallelism with Learned Asynchronous Decoding <i>Tian Jin, Ellie Y Cheng, Zachary Ankner, Nikunj Saunshi, Blake M Elias, Amir Yazdanbakhsh, Jonathan Ragan-Kelley, Suvinay Subramanian, Michael Carbin</i>	73624
ConfPO: Exploiting Policy Model Confidence for Critical Token Selection in Preference Optimization <i>Hee Suk Yoon, Eunseop Yoon, Mark A. Hasegawa-Johnson, Sungwoong Kim, Chang D. Yoo</i>	73640

Enhancing Parallelism in Decentralized Stochastic Convex Optimization <i>Ofri Eisen, Ron Dorfman, Kfir Yehuda Levy</i>	73655
CROW: Eliminating Backdoors from Large Language Models via Internal Consistency Regularization <i>Nay Myat Min, Long H. Pham, Yige Li, Jun Sun</i>	73680
Can Classic GNNs Be Strong Baselines for Graph-level Tasks? Simple Architectures Meet Excellence <i>Yuankai Luo, Lei Shi, Xiao-Ming Wu</i>	73700
Come Together, But Not Right Now: A Progressive Strategy to Boost Low-Rank Adaptation <i>Zhan Zhuang, Xiequn Wang, Wei Li, Yulong Zhang, Qiushi Huang, Shuhao Chen, Xuehao Wang, Yanbin Wei, Yuhe Nie, Kede Ma, Yu Zhang, Ying Wei</i>	73721
DynaMind: Reasoning over Abstract Video Dynamics for Embodied Decision-Making <i>Ziru Wang, Mengmeng Wang, Jade Dai, Teli Ma, Guo-Jun Qi, Yong Liu, Guang Dai, Jingdong Wang</i>	73739
Delta Decompression for MoE-based LLMs Compression <i>Hao Gu, Wei Li, Lujun Li, Zhu Qiyuan, Mark G. Lee, Shengjie Sun, Wei Xue, Yike Guo</i>	73757
On the Out-of-Distribution Generalization of Self-Supervised Learning <i>Wenwen Qiang, Jingyao Wang, Zeen Song, Jiangmeng Li, Changwen Zheng</i>	73775
Surrogate Prompt Learning: Towards Efficient and Diverse Prompt Learning for Vision-Language Models <i>Liangchen Liu, Nannan Wang, Xi Yang, Xinbo Gao, Tongliang Liu</i>	73802
Improved Sample Complexity for Private Nonsmooth Nonconvex Optimization <i>Guy Kornowski, Daogao Liu, Kunal Talwar</i>	73821
Point-Level Topological Representation Learning on Point Clouds <i>Vincent Peter Grande, Michael T Schaub</i>	73840
Refined generalization analysis of the Deep Ritz Method and Physics-Informed Neural Networks <i>Xianliang Xu, Ye Li, Zhongyi Huang</i>	73871
PROXSPARSE: REGULARIZED LEARNING OF SEMI-STRUCTURED SPARSITY MASKS FOR PRETRAINED LLMS <i>Hongyi Liu, Rajarshi Saha, Zhen Jia, Youngsuk Park, Jiaji Huang, Shoham Sabach, Yu-Xiang Wang, George Karypis</i>	73926
MoHAVE: Mixture of Hierarchical Audio-Visual Experts for Robust Speech Recognition <i>Sungnyun Kim, Kangwook Jang, Sangmin Bae, Sungwoo Cho, Se-Young Yun</i>	73941
Contextual Linear Bandits with Delay as Payoff <i>Mengxiao Zhang, Yingfei Wang, Haipeng Luo</i>	73960

Flow-field inference from neural data using deep recurrent networks <i>Timothy Doyeon Kim, Thomas Zhihao Luo, Tankut Can, Kamesh Krishnamurthy, Jonathan W. Pillow, Carlos D Brody</i>	73987
Alberta Wells Dataset: Pinpointing Oil and Gas Wells from Satellite Imagery <i>Pratinav Seth, Michelle Lin, BREFO DWAMENA YAW, Jade Boutot, Mary Kang, David Rolnick</i>	74011
Aligning Multimodal Representations through an Information Bottleneck <i>Antonio Almudévar, José Miguel Hernández-Lobato, Sameer Khurana, Ricard Marxer, Alfonso Ortega</i>	74039
Multi-Marginal Stochastic Flow Matching for High-Dimensional Snapshot Data at Irregular Time Points <i>Justin Lee, Behnaz Moradjamei, Heman Shakeri</i>	74060
Quamba2: A Robust and Scalable Post-training Quantization Framework for Selective State Space Models <i>Hung-Yueh Chiang, Chi-Chih Chang, Natalia Frumkin, Kai-Chiang Wu, Mohamed S. Abdelfattah, Diana Marculescu</i>	74083
PARM: Multi-Objective Test-Time Alignment via Preference-Aware Autoregressive Reward Model <i>Baijiong Lin, Weisen Jiang, Yuancheng Xu, Hao Chen, Ying-Cong Chen</i>	74100
MARGE: Improving Math Reasoning with Guided Exploration <i>Jingyue Gao, Runji Lin, Keming Lu, Bowen Yu, Junyang Lin, Jianyu Chen</i>	74115
Target Concrete Score Matching: A Holistic Framework for Discrete Diffusion <i>Ruixiang ZHANG, Shuangfei Zhai, Yizhe Zhang, James Thornton, Zijing Ou, Joshua M. Susskind, Navdeep Jaitly</i>	74138
Harnessing Heterogeneous Statistical Strength for Personalized Federated Learning via Hierarchical Bayesian Inference <i>Mahendra Singh Thapa, Rui Li</i>	74176
Quantum Speedup for Hypergraph Sparsification <i>Chenghua Liu, Minbo Gao, Zhengfeng Ji, Mingsheng Ying</i>	74195
Accurate and Efficient World Modeling with Masked Latent Transformers <i>Maxime Burchi, Radu Timofte</i>	74214
Dynamic Mixture of Curriculum LoRA Experts for Continual Multimodal Instruction Tuning <i>Chendi Ge, Xin Wang, Zeyang Zhang, Hong Chen, Jiapei Fan, Longtao Huang, Hui Xue, Wenwu Zhu</i>	74233
Physics Aware Neural Networks for Unsupervised Binding Energy Prediction <i>Ke Liu, Hao Chen, Chunhua Shen</i>	74256
Do Vision-Language Models Really Understand Visual Language? <i>Yifan Hou, Buse Giledereli, Yilei Tu, Mrinmaya Sachan</i>	74275

Dimension-Free Adaptive Subgradient Methods with Frequent Directions <i>Sifan Yang, Yuanyu Wan, Peijia Li, Yibo Wang, Xiao Zhang, Zhewei Wei, Lijun Zhang</i>	74325
MIRROR: Make Your Object-Level Multi-View Generation More Consistent with Training-Free Rectification <i>Tianchi Xing, Bonan Li, Congying Han, Xinmin Qiu, Zicheng Zhang, Tiande Guo</i>	74351
Don't Restart, Just Reuse: Reoptimizing MILPs with Dynamic Parameters <i>Sijia Zhang, Shuli Zeng, Shaoang Li, Feng Wu, Shaojie Tang, Xiangyang Li</i>	74379
Drug-TTA: Test-Time Adaptation for Drug Virtual Screening via Multi-task Meta-Auxiliary Learning <i>Ao Shen, Mingzhi Yuan, Yingfan MA, Jie Du, Qiao Huang, Manning Wang</i>	74402
Low-Dimension-to-High-Dimension Generalization and Its Implications for Length Generalization <i>Yang Chen, Long Yang, Yitao Liang, Zhouchen Lin</i>	74421
Stealing That Free Lunch: Exposing the Limits of Dyna-Style Reinforcement Learning <i>Brett Barkley, David Fridovich-Keil</i>	74445
Pareto-Optimality, Smoothness, and Stochasticity in Learning-Augmented One-Max-Search <i>Ziyad Benomar, Lorenzo Croissant, Vianney Perchet, Spyros Angelopoulos</i>	74470
Self-Improving Transformers Overcome Easy-to-Hard and Length Generalization Challenges <i>Nayoung Lee, Ziyang Cai, Avi Schwarzschild, Kangwook Lee, Dimitris Papailiopoulos</i>	74500
Volume-Aware Distance for Robust Similarity Learning <i>Shuo Chen, Chen Gong, Jun Li, Jian Yang</i>	74535
Private Federated Learning using Preference-Optimized Synthetic Data <i>Charlie Hou, Mei-Yu Wang, Yige Zhu, Daniel Lazar, Giulia Fanti</i>	74557
QuRe: Query-Relevant Retrieval through Hard Negative Sampling in Composed Image Retrieval <i>Jaehyun Kwak, Ramahdani Muhammad Izaaz Inhar, Se-Young Yun, Sung-Ju Lee</i>	74577
XAttnMark: Learning Robust Audio Watermarking with Cross-Attention <i>Yixin Liu, Lie Lu, Jihui Jin, Lichao Sun, Andrea Fanelli</i>	74592
In-Context Learning as Conditioned Associative Memory Retrieval <i>Weimin Wu, Teng-Yun Hsiao, Jerry Yao-Chieh Hu, Wenxin Zhang, Han Liu</i>	74621
A Mathematical Framework for AI-Human Integration in Work <i>L. Elisa Celis, Lingxiao Huang, Nisheeth K. Vishnoi</i>	74647
Visual Attention Never Fades: Selective Progressive Attention ReCalibration for Detailed Image Captioning in Multimodal Large Language Models <i>Mingi Jung, Saehyung Lee, Eunji Kim, Sungroh Yoon</i>	74682

Hybrid Batch Normalisation: Resolving the Dilemma of Batch Normalisation in Federated Learning <i>Hongyao Chen, Tianyang Xu, Xiaojun Wu, Josef Kittler</i>	74707
GCAL: Adapting Graph Models to Evolving Domain Shifts <i>Ziyue Qiao, Qianyi Cai, Hao Dong, Jiawei Gu, Pengyang Wang, Meng Xiao, Xiao Luo, Hui Xiong</i>	74723
ENAHPool: The Edge-Node Attention-based Hierarchical Pooling for Graph Neural Networks <i>Zhehan Zhao, Lu Bai, Lixin Cui, Ming Li, Ziyu Lyu, Lixiang Xu, Yue Wang, Edwin Hancock</i>	74740
Efficient Logit-based Knowledge Distillation of Deep Spiking Neural Networks for Full-Range Timestep Deployment <i>Chengting Yu, Xiaochen Zhao, Lei Liu, Shu Yang, Gaoang Wang, Erping Li, Aili Wang</i>	74751
Optimizing Language Models for Inference Time Objectives using Reinforcement Learning <i>Yunhao Tang, Kunhao Zheng, Gabriel Synnaeve, Remi Munos</i>	74769
BiMark: Unbiased Multilayer Watermarking for Large Language Models <i>Xiaoyan Feng, He Zhang, Yanjun Zhang, Leo Yu Zhang, Shirui Pan</i>	74789
Dequantified Diffusion-Schrödinger Bridge for Density Ratio Estimation <i>Wei Chen, Shigui Li, Jiacheng Li, Junmei Yang, John Paisley, Delu Zeng</i>	74808
SpikeVideoFormer: An Efficient Spike-Driven Video Transformer with Hamming Attention and $\mathcal{O}(T)$ Complexity <i>Shihao Zou, Qingfeng Li, Wei Ji, Jingjing Li, Yongkui Yang, Guoqi Li, Chao Dong</i>	74834
DAMA: Data- and Model-aware Alignment of Multi-modal LLMs <i>Jinda Lu, Junkang Wu, Jinghan Li, Xiaojun Jia, Shuo Wang, YiFan Zhang, Junfeng Fang, Xiang Wang, Xiangnan He</i>	74851
STP: Self-play LLM Theorem Provers with Iterative Conjecturing and Proving <i>Kefan Dong, Tengyu Ma</i>	74866
Unraveling the Interplay between Carryover Effects and Reward Autocorrelations in Switchback Experiments <i>Qianglin Wen, Chengchun Shi, YingYang, Niansheng Tang, Hongtu Zhu</i>	74889
Concept-Based Unsupervised Domain Adaptation <i>Xinyue Xu, Yueying Hu, Hui Tang, Yi Qin, Lu Mi, Hao Wang, Xiaomeng Li</i>	74928
Text-to-LoRA: Instant Transformer Adaption <i>Rujikorn Charakorn, Edoardo Cetin, Yujin Tang, Robert Tjarko Lange</i>	74950
Approximation to Smooth Functions by Low-Rank Swish Networks <i>Zimeng Li, Hongjun Li, Jingyuan Wang, Ke Tang</i>	74980

Latent Preference Coding: Aligning Large Language Models via Discrete Latent Codes <i>Zhuocheng Gong, Jian Guan, Wei Wu, Huishuai Zhang, Dongyan Zhao</i>	75013
Pretraining Generative Flow Networks with Inexpensive Rewards for Molecular Graph Generation <i>Mohit Pandey, Gopeshh Subbaraj, Artem Cherkasov, Martin Ester, Emmanuel Bengio</i>	75028
CAT Merging: A Training-Free Approach for Resolving Conflicts in Model Merging <i>Wenju Sun, Qingyong Li, Yangliao Geng, Boyang Li</i>	75067
MP-Nav: Enhancing Data Poisoning Attacks against Multimodal Learning <i>Jingfeng Zhang, Prashanth Krishnamurthy, Naman Patel, Anthony Tzes, Farshad Khorrami</i>	75088
Continuous Bayesian Model Selection for Multivariate Causal Discovery <i>Anish Dhir, Ruby Sedgwick, Avinash Kori, Ben Glocker, Mark van der Wilk</i>	75102
Decoupled SGDA for Games with Intermittent Strategy Communication <i>Ali Zindari, Parham Yazdkhasti, Anton Rodomanov, Tatjana Chavdarova, Sebastian U Stich</i>	75129
Towards LLM Unlearning Resilient to Relearning Attacks: A Sharpness-Aware Minimization Perspective and Beyond <i>Chongyu Fan, Jinghan Jia, Yihua Zhang, Anil Ramakrishna, Mingyi Hong, Sijia Liu</i>	75175
Deep Fuzzy Multi-view Learning for Reliable Classification <i>Siyuan Duan, Yuan Sun, Dezhong Peng, Guiduo Duan, Xi Peng, Peng Hu</i>	75192
EvoMesh: Adaptive Physical Simulation with Hierarchical Graph Evolutions <i>Huayu Deng, Xiangming Zhu, Yunbo Wang, Xiaokang Yang</i>	75216
Sketch to Adapt: Fine-Tunable Sketches for Efficient LLM Adaptation <i>Tianyi Zhang, Junda Su, Aditya Desai, Oscar Wu, Zhaozhuo Xu, Anshumali Shrivastava</i>	75236
CABS: Conflict-Aware and Balanced Sparsification for Enhancing Model Merging <i>zongzhen yang, Binhang Qi, Hailong Sun, Wenrui Long, Ruobing Zhao, Xiang Gao</i>	75259