

12th International Conference on Learning Representations

(ICLR 2024)

7-11 May 2024, Vienna, Austria

ISBN: 978-1-7138-9865-8

Title and Authors	Page#
Exposing Text-Image Inconsistency Using Diffusion Models Mingzhen Huang, Shan Jia, Zhou Zhou, Yan Ju, Jialing Cai, Siwei Lyu	1
I-PHYRE: Interactive Physical Reasoning Shiqian Li, Kewen Wu, Chi Zhang, Yixin Zhu	19
To the Cutoff... and Beyond? a Longitudinal Perspective on LLM Data Contamination Manley Roberts, Himanshu Thakur, Christine Herlihy, Colin White, Samuel Dooley	40
MuSc: Zero-Shot Industrial Anomaly Classification and Segmentation with Mutual Scoring of the Unlabeled Images Xurui Li, Ziming Huang, Feng Xue, Yu Zhou	108
Langevin Monte Carlo for Strongly Log-Concave Distributions: Randomized Midpoint Revisited Lu Yu, Avetik Karagulyan, Arnak S. Dalalyan	134
Rethinking Model Ensemble in Transfer-Based Adversarial Attacks Huanran Chen, Yichi Zhang, Yinpeng Dong, Xiao Yang, Hang Su, Jun Zhu	171
Recursive Generalization Transformer for Image Super-Resolution Zheng Chen, Yulun Zhang, Jinjin Gu, Linghe Kong, Xiaokang Yang	198
Early Stopping Against Label Noise Without Validation Data Suqin Yuan, Lei Feng, Tongliang Liu	210
Multi-Granularity Correspondence Learning from Long-Term Noisy Videos Yijie Lin, Jie Zhang, Zhenyu Huang, Jia Liu, Zujie Wen, Xi Peng	231
Towards Offline Opponent Modeling with In-Context Learning Yuheng Jing, Kai Li, Bingyun Liu, Yifan Zang, Haobo Fu, Qiang Fu, Junliang Xing, Jian Cheng	251
Divide and Not Forget: Ensemble of Selectively Trained Experts in Continual Learning Grzegorz Rypek, Sebastian Cygert, Valeriya Khan, Tomasz Trzcinski, Bartosz Michał, Zieliński, Bartłomiej Twardowski	286
Language Model Self-Improvement by Reinforcement Learning Contemplation Jing-Cheng Pang, Pengyuan Wang, Kaiyuan Li, Xiong-Hui Chen, Jiacheng Xu, Zongzhang Zhang, Yang Yu	304

3D-Aware Hypothesis & Verification for Generalizable Relative Object Pose Estimation Chen Zhao, Tong Zhang, Mathieu Salzmann	331
QLLM: Accurate and Efficient Low-Bitwidth Quantization for Large Language Models Jing Liu, Ruihao Gong, Xiuying Wei, Zhiwei Dong, Jianfei Cai, Bohan Zhuang	347
Threaten Spiking Neural Networks Through Combining Rate and Temporal Information Zecheng Hao, Tong Bu, Xinyu Shi, Zihan Huang, Zhaofei Yu, Tiejun Huang	370
A Symmetry-Aware Exploration of Bayesian Neural Network Posteriors Olivier Laurent, Emanuel Aldea, Gianni Franchi	387
Real3D-Portrait: One-Shot Realistic 3D Talking Portrait Synthesis Zhenhui Ye, Tianyun Zhong, Yi Ren, Jiaqi Yang, Weichuang Li, Jiawei Huang, Ziyue Jiang, Jinzheng He, Rongjie Huang, Jinglin Liu, Chen Zhang, Xiang Yin, Zejun Ma, Zhou Zhao	435
Mega-TTS 2: Boosting Prompting Mechanisms for Zero-Shot Speech Synthesis Ziyue Jiang, Jinglin Liu, Yi Ren, Jinzheng He, Zhenhui Ye, Shengpeng Ji, Qian Yang, Chen Zhang, Pengfei Wei, Chunfeng Wang, Xiang Yin, Zejun Ma, Zhou Zhao	458
A Progressive Training Framework for Spiking Neural Networks with Learnable Multi-Hierarchical Model Zecheng Hao, Xinyu Shi, Zihan Huang, Tong Bu, Zhaofei Yu, Tiejun Huang	479
SAS: Structured Activation Sparsification Yusuke Sekikawa, Shingo Yashima	496
Large-Vocabulary 3D Diffusion Model with Transformer Ziang Cao, Fangzhou Hong, Tong Wu, Liang Pan, Ziwei Liu	514
Neural Processing of Tri-Plane Hybrid Neural Fields Adriano Cardace, Pierluigi Zama Ramirez, Francesco Ballerini, Allan Zhou, Samuele Salti, Luigi Di Stefano	526
CriBo: Self-Supervised Learning Via Cross-Image Object-Level Bootstrapping Tim Lebailly, Thomas Stegmüller, Behzad Bozorgtabar, Jean-Philippe Thiran, Tinne Tuytelaars	553
TUVF: Learning Generalizable Texture UV Radiance Fields An-Chieh Cheng, Xueting Li, Sifei Liu, Xiaolong Wang	574
Compressed Context Memory for Online Language Model Interaction Jang-Hyun Kim, Junyoung Yeom, Sangdoo Yun, Hyun Oh Song	615
FedP3: Federated Personalized and Privacy-Friendly Network Pruning Under Model Heterogeneity Kai Yi, Nidham Gazagnadou, Peter Richtárik, Lingjuan Lyu	634
Multi-Scale Representations by Varying Window Attention for Semantic Segmentation Haotian Yan, Ming Wu, Chuang Zhang	667

Robust Angular Synchronization Via Directed Graph Neural Networks Yixuan He, Gesine Reinert, David Wipf, Mihai Cucuringu	684
Boosting Vanilla Lightweight Vision Transformers Via Re-Parameterization Zhentao Tan, Xiaodan Li, Yue Wu, Qi Chu, Le Lu, Nenghai Yu, Jieping Ye	729
MixSup: Mixed-Grained Supervision for Label-Efficient LiDAR-Based 3D Object Detection Yuxue Yang, Lue Fan, Zhaoxiang Zhang	749
Pareto Deep Long-Tailed Recognition: A Conflict-Averse Solution Zhipeng Zhou, Liu Liu, Peilin Zhao, Wei Gong	767
Constraint-Free Structure Learning with Smooth Acyclic Orientations Riccardo Massidda, Francesco Landolfi, Martina Cinquini, Davide Bacciu	794
Un-Mixing Test-Time Normalization Statistics: Combatting Label Temporal Correlation Devavrat Tomar, Guillaume Vray, Jean-Philippe Thiran, Behzad Bozorgtabar	824
GeoDiffusion: Text-Prompted Geometric Control for Object Detection Data Generation Kai Chen, Enze Xie, Zhe Chen, Yibo Wang, Lanqing Hong, Zhenguo Li, Dit-Yan Yeung	846
MogaNet: Multi-Order Gated Aggregation Network Siyuan Li, Zedong Wang, Zicheng Liu, Cheng Tan, Haitao Lin, Di Wu, Zhiyuan Chen, Jiangbin Zheng, Stan Z. Li	869
MagicDrive: Street View Generation with Diverse 3D Geometry Control Ruiyuan Gao, Kai Chen, Enze Xie, Lanqing Hong, Zhenguo Li, Dit-Yan Yeung, Qiang Xu	904
Gaining Wisdom from Setbacks: Aligning Large Language Models Via Mistake Analysis Kai Chen, Chunwei Wang, Kuo Yang, Jianhua Han, Lanqing Hong, Fei Mi, Hang Xu, Zhengying Liu, Wenying Huang, Zhenguo Li, Dit-Yan Yeung, Lifeng Shang	924
Online Stabilization of Spiking Neural Networks Yaoyu Zhu, Jianhao Ding, Tiejun Huang, Xiaodong Xie, Zhaoifei Yu	948
Towards Energy Efficient Spiking Neural Networks: An Unstructured Pruning Framework Xinyu Shi, Jianhao Ding, Zecheng Hao, Zhaoifei Yu	968
Koopman-Based Generalization Bound: New Aspect for Full-Rank Weights Yuka Hashimoto, Sho Sonoda, Isao Ishikawa, Atsushi Nitanda, Taiji Suzuki	989
Exploring Target Representations for Masked Autoencoders Xingbin Liu, Jinghao Zhou, Tao Kong, Xianming Lin, Rongrong Ji	1015
CLAP: Collaborative Adaptation for Patchwork Learning Sen Cui, Abudukelimu Wuerkaixi, Weishen Pan, Jian Liang, Lei Fang, Changshui Zhang, Fei Wang	1030
Path Choice Matters for Clear Attributions in Path Methods Borui Zhang, Wenzhao Zheng, Jie Zhou, Jiwen Lu	1051

REValueD: Regularised Ensemble Value-Decomposition for Factorisable Markov Decision Processes <i>David Ireland, Giovanni Montana</i>	1074
Geometry-Aware Projective Mapping for Unbounded Neural Radiance Fields <i>Junoh Lee, Hyunjun Jung, Jin-Hwi Park, Inhwan Bae, Hae-Gon Jeon</i>	1100
Guaranteed Approximation Bounds for Mixed-Precision Neural Operators <i>Renbo Tu, Colin White, Jean Kossaifi, Boris Bonev, Gennady Pekhimenko, Kamyar Azizzadenesheli, Anima Anandkumar</i>	1125
Fine-Tuning Multimodal LLMs to Follow Zero-Shot Demonstrative Instructions <i>Juncheng Li, Kaihang Pan, Zhiqi Ge, Minghe Gao, Wei Ji, Wenqiao Zhang, Tat-Seng Chua, Siliang Tang, Hanwang Zhang, Yueting Zhuang</i>	1155
OmniControl: Control Any Joint at Any Time for Human Motion Generation <i>Yiming Xie, Varun Jampani, Lei Zhong, Deqing Sun, Huaizu Jiang</i>	1181
Robust Adversarial Reinforcement Learning Via Bounded Rationality Curricula <i>Aryaman Reddi, Maximilian Tatveit, Jan Peters, Georgia Chalvatzaki, Carlo D'Eramo</i>	1200
SelfCheck: Using LLMs to Zero-Shot Check Their Own Step-By-Step Reasoning <i>Ning Miao, Yee Whye Teh, Tom Rainforth</i>	1225
Enhancing Contrastive Learning for Ordinal Regression Via Ordinal Content Preserved Data Augmentation <i>Jiyang Zheng, Yu Yao, Bo Han, Dadong Wang, Tongliang Liu</i>	1241
Large Language Models Are Not Robust Multiple Choice Selectors <i>Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, Minlie Huang</i>	1258
Mask-Based Modeling for Neural Radiance Fields <i>Ganlin Yang, Guoqiang Wei, Zhizheng Zhang, Yan Lu, Dong Liu</i>	1287
Variance-Enlarged Poisson Learning for Graph-Based Semi-Supervised Learning with Extremely Sparse Labeled Data <i>Xiong Zhou, Xianming Liu, Hao Yu, Jialiang Wang, Zeke Xie, Junjun Jiang, Xiangyang Ji</i>	1310
Zero-Mean Regularized Spectral Contrastive Learning: Implicitly Mitigating Wrong Connections in Positive-Pair Graphs <i>Xiong Zhou, Xianming Liu, Feilong Zhang, Gang Wu, Deming Zhai, Junjun Jiang, Xiangyang Ji</i>	1329
SEA: Sparse Linear Attention with Estimated Attention Mask <i>Heejun Lee, Jina Kim, Jeffrey Willette, Sung Ju Hwang</i>	1360
SYMBOL: Generating Flexible Black-Box Optimizers Through Symbolic Equation Learning <i>Jiacheng Chen, Zeyuan Ma, Hongshu Guo, Yining Ma, Jie Zhang, Yue-Jiao Gong</i>	1382

Towards Lossless Dataset Distillation Via Difficulty-Aligned Trajectory Matching Ziyao Guo, Kai Wang, George Cazenavette, Hui Li, Kaipeng Zhang, Yang You	1402
InfoBatch: Lossless Training Speed Up by Unbiased Dynamic Data Pruning Ziheng Qin, Kai Wang, Zangwei Zheng, Jianyang Gu, Xiangyu Peng, Xu Zhao Pan, Daquan Zhou, Lei Shang, Baigui Sun, Xuansong Xie, Yang You	1424
VDC: Versatile Data Cleanser Based on Visual-Linguistic Inconsistency by Multimodal Large Language Models Zihao Zhu, Mingda Zhang, Shaokui Wei, Bingzhe Wu, Baoyuan Wu	1446
GTA: A Geometry-Aware Attention Mechanism for Multi-View Transformers Takeru Miyato, Bernhard Jaeger, Max Welling, Andreas Geiger	1468
GRANDE: Gradient-Based Decision Tree Ensembles for Tabular Data Sascha Marton, Stefan Lã¼dtke, Christian Bartelt, Heiner Stuckenschmidt	1505
Bayesian Bi-Clustering of Neural Spiking Activity with Latent Structures Ganchao Wei	1532
ToRA: A Tool-Integrated Reasoning Agent for Mathematical Problem Solving Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Minlie Huang, Nan Duan, Weizhu Chen	1552
DFormer: Rethinking RGBD Representation Learning for Semantic Segmentation Bowen Yin, Xuying Zhang, Zhong-Yu Li, Li Liu, Ming-Ming Cheng, Qibin Hou	1586
Meta-Evolve: Continuous Robot Evolution for One-To-Many Policy Transfer Xingyu Liu, Deepak Pathak, Ding Zhao	1609
Unraveling the Enigma of Double Descent: An In-Depth Analysis Through the Lens of Learned Feature Space Yufei Gu, Xiaoqing Zheng, Tomaso Aste	1628
Less Or More from Teacher: Exploiting Trilateral Geometry for Knowledge Distillation Chengming Hu, Haolun Wu, Xuan Li, Chen Ma, Xi Chen, Boyu Wang, Jun Yan, Xue Liu	1640
Unprocessing Seven Years of Algorithmic Fairness Andrã© Cruz, Moritz Hardt	1656
Get What You Want, Not What You Don't: Image Content Suppression for Text-To-Image Diffusion Models Senmao Li, Joost Van De Weijer, Taihang Hu, Fahad Khan, Qibin Hou, Yaxing Wang, Jian Yang	1682
GeneOH Diffusion: Towards Generalizable Hand-Object Interaction Denoising Via Denoising Diffusion Xueyi Liu, Li Yi	1709

BrainSCUBA: Fine-Grained Natural Language Captions of Visual Cortex Selectivity Andrew Luo, Margaret Marie Henderson, Michael J. Tarr, Leila Wehbe	1744
Spike-Driven Transformer V2: Meta Spiking Neural Network Architecture Inspiring the Design of Next-Generation Neuromorphic Chips Man Yao, Jiakui Hu, Tianxiang Hu, Yifan Xu, Zhaokun Zhou, Yonghong Tian, Bo Xu, Guoqi Li	1781
LaneSegNet: Map Learning with Lane Segment Perception for Autonomous Driving Tianyu Li, Peijin Jia, Bangjun Wang, Li Chen, Kun Jiang, Junchi Yan, Hongyang Li	1804
Towards Domain-Invariant Self-Supervised Learning with Batch Styles Standardization Marin Scalbert, Maria Vakalopoulou, Florent Couzinie-Devy	1821
Parameter-Efficient Orthogonal Finetuning Via Butterfly Factorization Weiyang Liu, Zeju Qiu, Yao Feng, Yuliang Xiu, Yuxuan Xue, Longhui Yu, Haiwen Feng, Zhen Liu, Juyeon Heo, Songyou Peng, Yandong Wen, Michael J. Black, Adrian Weller, Bernhard Schölkopf	1840
Localizing and Editing Knowledge in Text-To-Image Generative Models Samyadeep Basu, Nanxuan Zhao, Vlad I Morariu, Soheil Feizi, Varun Manjunatha	1874
Online Continual Learning for Interactive Instruction Following Agents Byeonghwi Kim, Minhyuk Seo, Jonghyun Choi	1886
Language-Informed Visual Concept Learning Sharon Lee, Yunzhi Zhang, Shangzhe Wu, Jiajun Wu	1906
Real-Time Photorealistic Dynamic Scene Representation and Rendering with 4D Gaussian Splatting Zeyu Yang, Hongye Yang, Zijie Pan, Li Zhang	1934
Consistency-Guided Prompt Learning for Vision-Language Models Shuvendu Roy, Ali Etemad	1952
Strategic Preys Make Acute Predators: Enhancing Camouflaged Object Detectors by Generating Camouflaged Objects Chunming He, Kai Li, Yachao Zhang, Yulun Zhang, Chenyu You, Zhenhua Guo, Xiu Li, Martin Danelljan, Fisher Yu	1966
End-To-End (Instance)-Image Goal Navigation Through Correspondence as an Emergent Phenomenon Guillaume Bono, Leonid Antsfeld, Boris Chidlovskii, Philippe Weinzaepfel, Christian Wolf	1985
CausalLM is Not Optimal for In-Context Learning Nan Ding, Tomer Levinboim, Jialin Wu, Sebastian Goodman, Radu Soricut	2003
DreamClean: Restoring Clean Image Using Deep Diffusion Prior Jie Xiao, Ruili Feng, Han Zhang, Zhiheng Liu, Zhantao Yang, Yurui Zhu, Xueyang Fu, Kai Zhu, Yu Liu, Zheng-Jun Zha	2015

Weatherproofing Retrieval for Localization with Generative AI and Geometric Consistency Yannis Kalantidis, Mert Bărlent Sarăyă±Idă±z, Rafael S. Rezende, Philippe Weinzaepfel, Diane Larlus, Gabriela Csurka	2039
Transferring Labels to Solve Annotation Mismatches Across Object Detection Datasets Yuan-Hong Liao, David Acuna, Rafid Mahmood, James Lucas, Viraj Uday Prabhu, Sanja Fidler	2076
GROOT: Learning to Follow Instructions by Watching Gameplay Videos Shaofei Cai, Bowei Zhang, Zihao Wang, Xiaojian Ma, Anji Liu, Yitao Liang	2096
Enhancing High-Resolution 3D Generation Through Pixel-Wise Gradient Clipping Zijie Pan, Jiachen Lu, Xiatian Zhu, Li Zhang	2128
SpeechTokenizer: Unified Speech Tokenizer for Speech Language Models Xin Zhang, Dong Zhang, Shimin Li, Yaqian Zhou, Xipeng Qiu	2141
Energy-Based Concept Bottleneck Models: Unifying Prediction, Concept Intervention, and Probabilistic Interpretations Xinyue Xu, Yi Qin, Lu Mi, Hao Wang, Xiaomeng Li	2162
Space and Time Continuous Physics Simulation from Partial Observations Steeven Janny, Madiha Nadri, Julie Digne, Christian Wolf	2185
R&B: Region and Boundary Aware Zero-Shot Grounded Text-To-Image Generation Jiayu Xiao, Henglei Lv, Liang Li, Shuhui Wang, Qingming Huang	2209
ZeroFlow: Scalable Scene Flow Via Distillation Kyle Vedder, Neehar Peri, Nathaniel Eliot Chodosh, Ishan Khatri, Eric Eaton, Dinesh Jayaraman, Yang Liu, Deva Ramanan, James Hays	2238
HyperHuman: Hyper-Realistic Human Generation with Latent Structural Diffusion Xian Liu, Jian Ren, Aliaksandr Siarohin, Ivan Skorokhodov, Yanyu Li, Dahua Lin, Xihui Liu, Ziwei Liu, Sergey Tulyakov	2251
Tractable Probabilistic Graph Representation Learning with Graph-Induced Sum-Product Networks Federico Errica, Mathias Niepert	2285
ContextRef: Evaluating Referenceless Metrics for Image Description Generation Elisa Kreiss, Eric Zelikman, Christopher Potts, Nick Haber	2305
Large Language Models as Optimizers Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, Xinyun Chen	2326
BatchPrompt: Accomplish More with Less Jianzhe Lin, Maurice Diesendruck, Liang Du, Robin Abraham	2367
Reverse Forward Curriculum Learning for Extreme Sample and Demo Efficiency Stone Tao, Arth Shukla, Tse-Kai Chan, Hao Su	2390

Learning with a Mole: Transferable Latent Spatial Representations for Navigation Without Reconstruction Guillaume Bono, Leonid Antsfeld, Assem Sadek, Gianluca Monaci, Christian Wolf	2412
OpenNeRF: Open Set 3D Neural Scene Segmentation with Pixel-Wise Features and Rendered Novel Views Francis Engelmann, Fabian Manhardt, Michael Niemeyer, Keisuke Tateno, Federico Tombari	2431
Language Modeling is Compression Gregoire Deletang, Anian Ruoss, Paul-Ambroise Duquenne, Elliot Catt, Tim Genewein, Christopher Mattern, Jordi Grau-Moya, Li Kevin Wenliang, Matthew Aitchison, Laurent Orseau, Marcus Hutter, Joel Veness	2443
Learning the Greatest Common Divisor: Explaining Transformer Predictions Francois Charton	2460
LEAP: Liberate Sparse-View 3D Modeling from Camera Poses Hanwen Jiang, Zhenyu Jiang, Yue Zhao, Qixing Huang	2481
Statistically Optimal $\mathcal{K}$-Means Clustering Via Nonnegative Low-Rank Semidefinite Programming Yubo Zhuang, Xiaohui Chen, Yun Yang, Richard Y. Zhang	2501
Denoising Diffusion Via Image-Based Rendering Titas Anciukevičius, Fabian Manhardt, Federico Tombari, Paul Henderson	2534
Neural Polynomial Gabor Fields for Macro Motion Analysis Chen Geng, Hong-Xing Yu, Sida Peng, Xiaowei Zhou, Jiajun Wu	2557
Expected Flow Networks in Stochastic Environments and Two-Player Zero-Sum Games Marco Jiralerspong, Bilun Sun, Danilo Vucetic, Tianyu Zhang, Yoshua Bengio, Gauthier Gidel, Nikolay Malkin	2583
Weakly-Supervised Audio Separation Via Bi-Modal Semantic Similarity Tanvir Mahmud, Saeed Amizadeh, Kazuhito Koishida, Diana Marculescu	2604
DAM: Towards a Foundation Model for Forecasting Luke Nicholas Darlow, Qiwen Deng, Ahmed Hassan, Martin Asenov, Rajkarn Singh, Artjom Joosen, Adam Barker, Amos Storkey	2637
Leveraging Unpaired Data for Vision-Language Generative Models Via Cycle Consistency Tianhong Li, Sangnie Bhardwaj, Yonglong Tian, Han Zhang, Jarred Barber, Dina Katabi, Guillaume Lajoie, Huiwen Chang, Dilip Krishnan	2677
Approximately Piecewise $E(3)$ Equivariant Point Networks Matan Atzmon, Jiahui Huang, Francis Williams, Or Litany	2698

Self-Supervised Representation Learning from Random Data Projectors Yi Sui, Tongzi Wu, Jesse C. Cresswell, Ga Wu, George Stein, Xiao Shi Huang, Xiaochen Zhang, Maksims Volkovs	2716
Tensor Trust: Interpretable Prompt Injection Attacks from an Online Game Sam Toyer, Olivia Watkins, Ethan Adrian Mendes, Justin Svegliato, Luke Bailey, Tiffany Wang, Isaac Ong, Karim Elmaaroufi, Pieter Abbeel, Trevor Darrell, Alan Ritter, Stuart Russell	2739
ImageNet-OOD: Deciphering Modern Out-Of-Distribution Detection Algorithms William Yang, Byron Zhang, Olga Russakovsky	2772
Perceptual Group Tokenizer: Building Perception with Iterative Grouping Zhiwei Deng, Ting Chen, Yang Li	2802
Fine-Tuning Enhances Existing Mechanisms: A Case Study on Entity Tracking Nikhil Prakash, Tamar Rott Shaham, Tal Haklay, Yonatan Belinkov, David Bau	2821
Function Vectors in Large Language Models Eric Todd, Millicent Li, Arnab Sen Sharma, Aaron Mueller, Byron C Wallace, David Bau	2847
Learning No-Regret Sparse Generalized Linear Models with Varying Observation(s) Diyang Li, Charles Ling, Zhiqiang Xu, Huan Xiong, Bin Gu	2899
Adaptive Rational Activations to Boost Deep Reinforcement Learning Quentin Delfosse, Patrick Schramowski, Martin Mundt, Alejandro Molina, Kristian Kersting	2936
Feature Collapse Thomas Laurent, James Von Brecht, Xavier Bresson	2960
AutoCast++: Enhancing World Event Prediction with Zero-Shot Ranking-Based Context Retrieval Qi Yan, Raihan Seraj, Jiawei He, Lili Meng, Tristan Sylvain	3018
Universal Humanoid Motion Representations for Physics-Based Control Zhengyi Luo, Jinkun Cao, Josh Merel, Alexander Winkler, Jing Huang, Kris M. Kitani, Weipeng Xu	3035
FasterViT: Fast Vision Transformers with Hierarchical Attention Ali Hatamizadeh, Greg Heinrich, Hongxu Yin, Andrew Tao, Jose M. Alvarez, Jan Kautz, Pavlo Molchanov	3052
Tuning LayerNorm in Attention: Towards Efficient Multi-Modal LLM Finetuning Bingchen Zhao, Haoqin Tu, Chen Wei, Jieru Mei, Cihang Xie	3076
Dual-Encoders for Extreme Multi-Label Classification Nilesh Gupta, Fnu Devvrit, Ankit Singh Rawat, Srinadh Bhojanapalli, Prateek Jain, Inderjit S Dhillon	3091
VQGraph: Rethinking Graph Representation Space for Bridging GNNs and MLPs Ling Yang, Ye Tian, Minkai Xu, Zhongyi Liu, Shenda Hong, Wei Qu, Wentao Zhang, Bin Cui, Muhan Zhang, Jure Leskovec	3118

Order-Preserving GFlowNets Yihang Chen, Lukas Mauch	3138
Guiding Instruction-Based Image Editing Via Multimodal Large Language Models Tsu-Jui Fu, Wenze Hu, Xianzhi Du, William Yang Wang, Yinfei Yang, Zhe Gan	3176
WizardCoder: Empowering Code Large Language Models with Evol-Instruct Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, Daxin Jiang	3190
Gene Regulatory Network Inference in the Presence of Dropouts: A Causal View Haoyue Dai, Ignavier Ng, Gongxu Luo, Peter Spirtes, Petar Stojanov, Kun Zhang	3211
Structured Video-Language Modeling with Temporal Grouping and Spatial Grounding Yuanhao Xiong, Long Zhao, Boqing Gong, Ming-Hsuan Yang, Florian Schroff, Ting Liu, Cho-Jui Hsieh, Liangzhe Yuan	3239
AutoVP: An Automated Visual Prompting Framework and Benchmark Hsi-Ai Tsao, Lei Hsiung, Pin-Yu Chen, Sijia Liu, Tsung-Yi Ho	3257
Causal Structure Recovery with Latent Variables Under Milder Distributional and Graphical Assumptions Xiu-Chuan Li, Kun Zhang, Tongliang Liu	3286
3D Reconstruction with Generalizable Neural Fields Using Scene Priors Yang Fu, Shalini De Mello, Xueting Li, Amey Kulkarni, Jan Kautz, Xiaolong Wang, Sifei Liu	3318
Get More for Less: Principled Data Selection for Warming Up Fine-Tuning in LLMs Feiyang Kang, Hoang Anh Just, Yifan Sun, Himanshu Jahagirdar, Yuezhi Zhang, Rongxing Du, Anit Kumar Sahu, Ruoxi Jia	3338
Denoising Diffusion Step-Aware Models Shuai Yang, Yukang Chen, Luozhou Wang, Shu Liu, Ying-Cong Chen	3368
AnimateDiff: Animate Your Personalized Text-To-Image Diffusion Models Without Specific Tuning Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, Bo Dai	3384
AutoLoRa: An Automated Robust Fine-Tuning Framework Xilie Xu, Jingfeng Zhang, Mohan Kankanhalli	3403
An LLM Can Fool Itself: A Prompt-Based Adversarial Attack Xilie Xu, Keyi Kong, Ning Liu, Lizhen Cui, Di Wang, Jingfeng Zhang, Mohan Kankanhalli	3418
How Connectivity Structure Shapes Rich and Lazy Learning in Neural Circuits Yuhan Helena Liu, Aristide Baratin, Jonathan Cornford, Stefan Mihalas, Eric Todd Sheabrown, Guillaume Lajoie	3441

Stochastic Controlled Averaging for Federated Learning with Communication Compression Xinmeng Huang, Ping Li, Xiaoyun Li	3456
Diffusion Posterior Sampling for Linear Inverse Problem Solving: A Filtering Perspective Zehao Dou, Yang Song	3495
Adaptive Deep Spiking Neural Network with Global-Local Learning Via Balanced Excitatory and Inhibitory Mechanism Tingting Jiang, Qi Xu, Xuming Ran, Jiangrong Shen, Pan Lv, Qiang Zhang, Gang Pan	3522
TD-MPC2: Scalable, Robust World Models for Continuous Control Nicklas Hansen, Hao Su, Xiaolong Wang	3534
Progressive Fourier Neural Representation for Sequential Video Compilation Haeyong Kang, Jaehong Yoon, Dahyun Kim, Sung Ju Hwang, Chang D. Yoo	3564
Inducing High Energy-Latency of Large Vision-Language Models with Verbose Images Kuofeng Gao, Yang Bai, Jindong Gu, Shu-Tao Xia, Philip Torr, Zhifeng Li, Wei Liu	3589
ResFields: Residual Neural Fields for Spatiotemporal Signals Marko Mihajlovic, Sergey Prokudin, Marc Pollefeys, Siyu Tang	3616
Neuron Activation Coverage: Rethinking Out-Of-Distribution Detection and Generalization Yibing Liu, Chris Xing Tian, Haoliang Li, Lei Ma, Shiqi Wang	3635
STanHop: Sparse Tandem Hopfield Model for Memory-Enhanced Time Series Prediction Dennis Wu, Jerry Yao-Chieh Hu, Weijian Li, Bo-Yu Chen, Han Liu	3668
InterpGNN: Understand and Improve Generalization Ability of Transductive GNNs Through the Lens of Interplay Between Train and Test Nodes Jiawei Sun, Kailai Li, Ruoxin Chen, Jie Li, Chentao Wu, Yue Ding, Junchi Yan	3708
LMUFormer: Low Complexity Yet Powerful Spiking Model with Legendre Memory Units Zeyu Liu, Gourav Datta, Anni Li, Peter Anthony Beerel	3722
Mastering Symbolic Operations: Augmenting Language Models with Compiled Neural Networks Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Kang Liu, Jun Zhao	3737
R-EDL: Relaxing Nonessential Settings of Evidential Deep Learning Mengyuan Chen, Junyu Gao, Changsheng Xu	3768
ADDP: Learning General Representations for Image Recognition and Generation with Alternating Denoising Diffusion Process Changyao Tian, Chenxin Tao, Jifeng Dai, Hao Li, Ziheng Li, Lewei Lu, Xiaogang Wang, Hongsheng Li, Gao Huang, Xizhou Zhu	3794

Inner Classifier-Free Guidance and Its Taylor Expansion for Diffusion Models Shikun Sun, Longhui Wei, Zhicai Wang, Zixuan Wang, Junliang Xing, Jia Jia, Qi Tian	3824
FreeNoise: Tuning-Free Longer Video Diffusion Via Noise Rescheduling Haonan Qiu, Menghan Xia, Yong Zhang, Yingqing He, Xintao Wang, Ying Shan, Ziwei Liu	3841
Prompt Gradient Projection for Continual Learning Jingyang Qiao, Zhizhong Zhang, Xin Tan, Chengwei Chen, Yanyun Qu, Yong Peng, Yuan Xie	3856
Graph Lottery Ticket Automated Guibin Zhang, Kun Wang, Wei Huang, Yanwei Yue, Yang Wang, Roger Zimmermann, Aojun Zhou, Dawei Cheng, Jin Zeng, Yuxuan Liang	3878
Theoretical Understanding of Learning from Adversarial Perturbations Soichiro Kumano, Hiroshi Kera, Toshihiko Yamasaki	3903
De Novo Protein Design Using Geometric Vector Field Networks Weian Mao, Muzhi Zhu, Zheng Sun, Shuaike Shen, Lin Yuanbo Wu, Hao Chen, Chunhua Shen	3952
Spatio-Temporal Few-Shot Learning Via Diffusive Neural Network Generation Yuan Yuan, Chenyang Shao, Jingtao Ding, Depeng Jin, Yong Li	3976
ITransformer: Inverted Transformers Are Effective for Time Series Forecasting Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, Mingsheng Long	4004
Q-Bench: A Benchmark for General-Purpose Foundation Models on Low-Level Vision Haoning Wu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Annan Wang, Chunyi Li, Wenxiu Sun, Qiong Yan, Guangtao Zhai, Weisi Lin	4029
OctoPack: Instruction Tuning Code Large Language Models Niklas Muennighoff, Qian Liu, Armel Randy Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro Von Werra, Shayne Longpre	4056
MVSFormer++: Revealing the Devil in Transformer's Details for Multi-View Stereo Chenjie Cao, Xinlin Ren, Yanwei Fu	4076
Efficient Sharpness-Aware Minimization for Molecular Graph Transformer Models Yili Wang, Kaixiong Zhou, Ninghao Liu, Ying Wang, Xin Wang	4098
When Semantic Segmentation Meets Frequency Aliasing Linwei Chen, Lin Gu, Ying Fu	4121
Continuous-Multiple Image Outpainting in One-Step Via Positional Query and a Diffusion-Based Approach Shaofeng Zhang, Jinfa Huang, Qiang Zhou, Zhibin Wang, Fan Wang, Jiebo Luo, Junchi Yan	4144
TimeMixer: Decomposable Multiscale Mixing for Time Series Forecasting Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Y. Zhang, Jun Zhou	4166

P2Seg: Pointly-Supervised Segmentation Via Mutual Distillation Zipeng Wang, Xuehui Yu, Xumeng Han, Wenwen Yu, Zhixun Huang, Jianbin Jiao, Zhenjun Han	4193
Augmenting Transformers with Recursively Composed Multi-Grained Representations Xiang Hu, Qingyang Zhu, Kewei Tu, Wei Wu	4205
InsertNeRF: Instilling Generalizability into NeRF with HyperNet Modules Yanqi Bao, Tianyu Ding, Jing Huo, Wenbin Li, Yuxin Li, Yang Gao	4223
Vision-Language Foundation Models as Effective Robot Imitators Xinghang Li, Minghuan Liu, Hanbo Zhang, Cunjun Yu, Jie Xu, Hongtao Wu, Chilam Cheang, Ya Jing, Weinan Zhang, Huaping Liu, Hang Li, Tao Kong	4247
VDT: General-Purpose Video Diffusion Transformers Via Mask Modeling Haoyu Lu, Guoxing Yang, Nanyi Fei, Yuqi Huo, Zhiwu Lu, Ping Luo, Mingyu Ding	4266
Multi-View Representation is What You Need for Point-Cloud Pre-Training Siming Yan, Chen Song, Youkang Kong, Qixing Huang	4294
DePT: Decomposed Prompt Tuning for Parameter-Efficient Fine-Tuning Zhengxiang Shi, Aldo Lipani	4309
Raidar: geneRative AI Detection viA Rewriting Chengzhi Mao, Carl Vondrick, Hao Wang, Junfeng Yang	4330
Procedural Fairness Through Decoupling Objectionable Data Generating Components Zeyu Tang, Jialu Wang, Yang Liu, Peter Spirtes, Kun Zhang	4348
How Well Do Supervised 3D Models Transfer to Medical Imaging Tasks? Wenxuan Li, Alan Yuille, Zongwei Zhou	4378
LRM: Large Reconstruction Model for Single Image to 3D Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, Hao Tan	4394
PF-LRM: Pose-Free Large Reconstruction Model for Joint Pose and Shape Prediction Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexiang Xu, Kai Zhang	4419
Influencer Backdoor Attack on Semantic Segmentation Haoheng Lan, Jindong Gu, Philip Torr, Hengshuang Zhao	4441
DMV3D: Denoising Multi-View Diffusion Using 3D Large Reconstruction Model Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, Kai Zhang	4466
Image Translation as Diffusion Visual Programmers Cheng Han, James Chenhao Liang, Qifan Wang, Majid Rabbani, Sohail Dianat, Raghuveer Rao, Ying Nian Wu, Dongfang Liu	4483

Knowledge Distillation Based on Transformed Teacher Matching Kaixiang Zheng, En-Hui Yang	4508
Initializing Models with Larger Ones Zhiqiu Xu, Yanjie Chen, Kirill Vishniakov, Yida Yin, Zhiqiang Shen, Trevor Darrell, Lingjie Liu, Zhuang Liu	4525
Effective and Efficient Federated Tree Learning on Hybrid Data Qinbin Li, Chulin Xie, Xiaojun Xu, Xiaoyuan Liu, Ce Zhang, Bo Li, Bingsheng He, Dawn Song	4541
Instant3D: Fast Text-To-3D with Sparse-View Generation and Large Reconstruction Model Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, Sai Bi	4561
Noisy Interpolation Learning with Shallow Univariate ReLU Networks Nirmit Joshi, Gal Vardi, Nathan Srebro	4586
GeoLLM: Extracting Geospatial Knowledge from Large Language Models Rohin Manvi, Samar Khanna, Gengchen Mai, Marshall Burke, David B. Lobell, Stefano Ermon	4623
A Simple and Effective Pruning Approach for Large Language Models Mingjie Sun, Zhuang Liu, Anna Bair, J Zico Kolter	4640
Scaling for Training Time and Post-Hoc Out-Of-Distribution Detection Enhancement Kai Xu, Rongyu Chen, Gianni Franchi, Angela Yao	4663
DreamGaussian: Generative Gaussian Splatting for Efficient 3D Content Creation Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, Gang Zeng	4677
Multisize Dataset Condensation Yang He, Lingao Xiao, Joey Tianyi Zhou, Ivor Tsang	4695
Data-Independent Module-Aware Pruning for Hierarchical Vision Transformers Yang He, Joey Tianyi Zhou	4716
Unifying Feature and Cost Aggregation with Transformers for Semantic and Visual Correspondence Sunghwan Hong, Seokju Cho, Seungryong Kim, Stephen Lin	4732
Connect, Collapse, Corrupt: Learning Cross-Modal Tasks with Uni-Modal Data Yuhui Zhang, Elaine Sui, Serena Yeung	4746
IDEAL: Influence-Driven Selective Annotations Empower In-Context Learners in Large Language Models Shaokun Zhang, Xiaobo Xia, Zhaoqing Wang, Ling-Hao Chen, Jiale Liu, Qingyun Wu, Tongliang Liu	4772
ViDA: Homeostatic Visual Domain Adapter for Continual Test Time Adaptation Jiaming Liu, Senqiao Yang, Peidong Jia, Renrui Zhang, Ming Lu, Yandong Guo, Wei Xue, Shanghang Zhang	4798

PanoDiffusion: 360-Degree Panorama Outpainting Via Diffusion Tianhao Wu, Chuanxia Zheng, Tat-Jen Cham	4820
DoLa: Decoding by Contrasting Layers Improves Factuality in Large Language Models Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James R. Glass, Pengcheng He	4841
Forward χ^2 Divergence Based Variational Importance Sampling Chengrui Li, Yule Wang, Weihan Li, Anqi Wu	4867
Hebbian Learning Based Orthogonal Projection for Continual Learning of Spiking Neural Networks Mingqing Xiao, Qingyan Meng, Zongpeng Zhang, Di He, Zhouchen Lin	4885
Analyzing and Mitigating Object Hallucination in Large Vision-Language Models Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, Huaxiu Yao	4908
LayoutNUWA: Revealing the Hidden Layout Expertise of Large Language Models Zecheng Tang, Chenfei Wu, Juntao Li, Nan Duan	4938
Idempotence and Perceptual Image Compression Tongda Xu, Ziran Zhu, Dailan He, Yanghao Li, Lina Guo, Yuanyuan Wang, Zhe Wang, Hongwei Qin, Yan Wang, Jingjing Liu, Ya-Qin Zhang	4958
Matcher: Segment Anything with One Shot Using All-Purpose Feature Matching Yang Liu, Muzhi Zhu, Hengtao Li, Hao Chen, Xinlong Wang, Chunhua Shen	4986
Revisiting Deep Audio-Text Retrieval Through the Lens of Transportation Manh Luong, Khai Nguyen, Nhat Ho, Reza Haf, Dinh Phung, Lizhen Qu	5008
Defining and Extracting Generalizable Interaction Primitives from DNNs Lu Chen, Siyu Lou, Benhao Huang, Quanshi Zhang	5023
Fast-DetectGPT: Efficient Zero-Shot Detection of Machine-Generated Text Via Conditional Probability Curvature Guangsheng Bao, Yanbin Zhao, Zhiyang Teng, Linyi Yang, Yue Zhang	5046
Scale-Adaptive Diffusion Model for Complex Sketch Synthesis Jijin Hu, Ke Li, Yonggang Qi, Yi-Zhe Song	5069
Byzantine Robust Cooperative Multi-Agent Reinforcement Learning as a Bayesian Game Simin Li, Jun Guo, Jingqiao Xiu, Ruixiao Xu, Xin Yu, Jiakai Wang, Aishan Liu, Yaodong Yang, Xianglong Liu	5089
Consistent Video-To-Video Transfer Using Synthetic Dataset Jiaxin Cheng, Tianjun Xiao, Tong He	5116
RTFS-Net: Recurrent Time-Frequency Modelling for Efficient Audio-Visual Speech Separation Samuel Pegg, Kai Li, Xiaolin Hu	5129

Enabling Efficient Equivariant Operations in the Fourier Basis Via Gaunt Tensor Products Shengjie Luo, Tianlang Chen, Aditi S. Krishnapriyan	5148
PromptTTS 2: Describing and Generating Voices with Text Prompt Yichong Leng, Zhifang Guo, Kai Shen, Zeqian Ju, Xu Tan, Eric Liu, Yufei Liu, Dongchao Yang, Leying Zhang, Kaitao Song, Lei He, Xiangyang Li, Sheng Zhao, Tao Qin, Jiang Bian	5184
LLCP: Learning Latent Causal Processes for Reasoning-Based Video Question Answer Guangyi Chen, Yuke Li, Xiao Liu, Zijian Li, Eman Al Suradi, Donglai Wei, Kun Zhang	5201
Rethinking CNNs Generalization to Backdoor Attack from Frequency Domain Quanrui Rao, Lin Wang, Wuying Liu	5224
Quantifying the Sensitivity of Inverse Reinforcement Learning to Misspecification Joar Max Viktor Skalse, Alessandro Abate	5243
Local Search GFlowNets Minsu Kim, Taeyoung Yun, Emmanuel Bengio, Dinghui Zhang, Yoshua Bengio, Sungsoo Ahn, Jinkyoo Park	5269
BarLeRla: An Efficient Tuning Framework for Referring Image Segmentation Yaoming Wang, Jin Li, Xiaopeng Zhang, Bowen Shi, Chenglin Li, Wenrui Dai, Hongkai Xiong, Qi Tian	5287
Understanding Convergence and Generalization in Federated Learning Through Feature Learning Theory Wei Huang, Ye Shi, Zhongyi Cai, Taiji Suzuki	5303
Learning to Embed Time Series Patches Independently Seunghan Lee, Taeyoung Park, Kibok Lee	5335
Xformer: Hybrid X-Shaped Transformer for Image Denoising Jiale Zhang, Yulun Zhang, Jinjin Gu, Jiahua Dong, Linghe Kong, Xiaokang Yang	5357
Entropy-MCMC: Sampling from Flat Basins with Ease Bolian Li, Ruqi Zhang	5370
Mind Your Augmentation: The Key to Decoupling Dense Self-Supervised Learning Congpei Qiu, Tong Zhang, Yanhao Wu, Wei Ke, Mathieu Salzmann, Sabine Salsstrunk	5394
NetInfoF Framework: Measuring and Exploiting Network Usable Information Meng-Chieh Lee, Haiyang Yu, Jian Zhang, Vassilis N. Ioannidis, Xiang Song, Soji Adeshina, Da Zheng, Christos Faloutsos	5422
Implicit Regularization of Deep Residual Networks Towards Neural ODEs Pierre Marion, Yu-Han Wu, Michael Eli Sander, G�rard Biau	5441
GPAvatar: Generalizable and Precise Head Avatar from Image(s) Xuangeng Chu, Yu Li, Ailing Zeng, Tianyu Yang, Lijian Lin, Yunfei Liu, Tatsuya Harada	5481

Variational Inference for SDEs Driven by Fractional Noise Rembert Daems, Manfred Oppel, Guillaume Crevecoeur, Tolga Birdal	5501
ReTaSA: A Nonparametric Functional Estimation Approach for Addressing Continuous Target Shift Hwanwoo Kim, Xin Zhang, Jiwei Zhao, Qinglong Tian	5529
Towards Few-Shot Adaptation of Foundation Models Via Multitask Finetuning Zhuoyan Xu, Zhenmei Shi, Junyi Wei, Fangzhou Mu, Yin Li, Yingyu Liang	5540
Interpreting CLIP's Image Representation Via Text-Based Decomposition Yossi Gandelsman, Alexei A Efros, Jacob Steinhardt	5594
BioBridge: Bridging Biomedical Foundation Models Via Knowledge Graphs Zifeng Wang, Zichen Wang, Balasubramaniam Srinivasan, Vassilis N. Ioannidis, Huzefa Rangwala, Rishita Anubhai	5616
Time Travel in LLMs: Tracing Data Contamination in Large Language Models Shahriar Golchin, Mihai Surdeanu	5639
Query-Dependent Prompt Evaluation and Optimization with Offline Inverse RL Hao Sun, Alihan Hã¼yã¼k, Mihaela Van Der Schaar	5661
Self-Guided Masked Autoencoders for Domain-Agnostic Self-Supervised Learning Johnathan Wenjia Xie, Yoonho Lee, Annie S Chen, Chelsea Finn	5717
Adaptive Instrument Design for Indirect Experiments Yash Chandak, Shiv Shankar, Vasilis Syrgkanis, Emma Brunskill	5736
IMPUS: Image Morphing with Perceptually-Uniform Sampling Using Diffusion Models Zhaoyuan Yang, Zhengyang Yu, Zhiwei Xu, Jaskirat Singh, Jing Zhang, Dylan Campbell, Peter Tu, Richard Hartley	5786
Generalized Neural Sorting Networks with Error-Free Differentiable Swap Functions Jungtaek Kim, Jeongbeen Yoon, Minsu Cho	5814
Context is Environment Sharut Gupta, Stefanie Jegelka, David Lopez-Paz, Kartik Ahuja	5835
Neural Snowflakes: Universal Latent Graph Inference Via Trainable Latent Geometries Haitz Sã¡ez De Ocã¡riz Borde, Anastasis Kratsios	5877
Decodable and Sample Invariant Continuous Object Encoder Dehao Yuan, Furong Huang, Cornelia Fermuller, Yiannis Aloimonos	5892
Structuring Representation Geometry with Rotationally Equivariant Contrastive Learning Sharut Gupta, Joshua Robinson, Derek Lim, Soledad Villar, Stefanie Jegelka	5918

Attention Satisfies: A Constraint-Satisfaction Lens on Factual Errors of Language Models Mert Yuksekgonul, Varun Chandrasekaran, Erik Jones, Suriya Gunasekar, Ranjita Naik, Hamid Palangi, Ece Kamar, Besmira Nushi	5942
A Statistical Analysis of Wasserstein Autoencoders for Intrinsically Low-Dimensional Data Saptarshi Chakraborty, Peter Bartlett	5967
Attention-Guided Contrastive Role Representations for Multi-Agent Reinforcement Learning Zican Hu, Zongzhang Zhang, Huaxiong Li, Chunlin Chen, Hongyu Ding, Zhi Wang	5999
PINNsFormer: A Transformer-Based Framework for Physics-Informed Neural Networks Zhiyuan Zhao, Xueying Ding, B. Aditya Prakash	6022
Learning with Mixture of Prototypes for Out-Of-Distribution Detection Haodong Lu, Dong Gong, Shuo Wang, Jason Xue, Lina Yao, Kristen Moore	6039
Enhancing Human-AI Collaboration Through Logic-Guided Reasoning Chengzhi Cao, Yinghao Fu, Sheng Xu, Ruimao Zhang, Shuang Li	6062
A Benchmark Study on Calibration Linwei Tao, Younan Zhu, Haolan Guo, Minjing Dong, Chang Xu	6074
MAmmoTH: Building Math Generalist Models Through Hybrid Instruction Tuning Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, Wenhui Chen	6087
Lipschitz Singularities in Diffusion Models Zhantao Yang, Ruili Feng, Han Zhang, Yujun Shen, Kai Zhu, Lianghua Huang, Yifei Zhang, Yu Liu, Deli Zhao, Jingren Zhou, Fan Cheng	6109
PnP Inversion: Boosting Diffusion-Based Editing with 3 Lines of Code Xuan Ju, Ailing Zeng, Yuxuan Bian, Shaoteng Liu, Qiang Xu	6136
Learning Energy-Based Models by Cooperative Diffusion Recovery Likelihood Yaxuan Zhu, Jianwen Xie, Ying Nian Wu, Ruiqi Gao	6164
Only Pay for What is Uncertain: Variance-Adaptive Thompson Sampling Aadirupa Saha, Branislav Kveton	6195
SEAL: A Framework for Systematic Evaluation of Real-World Super-Resolution Wenlong Zhang, Xiaohui Li, Xiangyu Chen, Xiaoyun Zhang, Yu Qiao, Xiao-Ming Wu, Chao Dong	6207
Detecting Machine-Generated Texts by Multi-Population Aware Optimization for Maximum Mean Discrepancy Shuhai Zhang, Yiliao Song, Jiahao Yang, Yuanqing Li, Bo Han, Mingkui Tan	6229
Momentum Benefits Non-IID Federated Learning Simply and Provably Ziheng Cheng, Xinmeng Huang, Pengfei Wu, Kun Yuan	6265

Improved Regret Bounds for Non-Convex Online-Within-Online Meta Learning Jiechao Guan, Hui Xiong	6299
Decentralized Riemannian Conjugate Gradient Method on the Stiefel Manifold Jun Chen, Haishan Ye, Mengmeng Wang, Tianxin Huang, Guang Dai, Ivor Tsang, Yong Liu	6326
Entropy is Not Enough for Test-Time Adaptation: From the Perspective of Disentangled Factors Jonghyun Lee, Dahuin Jung, Saehyung Lee, Junsung Park, Juhyeon Shin, Uiwon Hwang, Sungroh Yoon	6345
Class Incremental Learning Via Likelihood Ratio Based Task Prediction Haowei Lin, Yijia Shao, Weinan Qian, Ningxin Pan, Yiduo Guo, Bing Liu	6371
Enhanced Face Recognition Using Intra-Class Incoherence Constraint Yuanqing Huang, Yinggui Wang, Le Yang, Lei Wang	6411
Tag2Text: Guiding Vision-Language Model Via Image Tagging Xinyu Huang, Youcai Zhang, Jinyu Ma, Weiwei Tian, Rui Feng, Yuejie Zhang, Yaqian Li, Yandong Guo, Lei Zhang	6425
Convolution Meets LoRA: Parameter Efficient Finetuning for Segment Anything Model Zihan Zhong, Zhiqiang Tang, Tong He, Haoyang Fang, Chun Yuan	6445
Conformal Prediction Via Regression-As-Classification Etash Kumar Guha, Shlok Natarajan, Thomas Mallenhoff, Mohammad Emtiyaz Khan, Eugene Ndiaye	6470
Efficient Planning with Latent Diffusion Wenhao Li	6494
IceFormer: Accelerated Inference with Long-Sequence Transformers on CPUs Yuzhen Mao, Martin Ester, Ke Li	6527
AgentVerse: Facilitating Multi-Agent Collaboration and Exploring Emergent Behaviors Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, Jie Zhou	6545
Pre-Training with Synthetic Data Helps Offline Reinforcement Learning Zecheng Wang, Che Wang, Zixuan Dong, Keith W. Ross	6588
Sample-Efficient Learning of POMDPs with Multiple Observations in Hindsight Jiacheng Guo, Minshuo Chen, Huan Wang, Caiming Xiong, Mengdi Wang, Yu Bai	6625
ToolLLM: Facilitating Large Language Models to Master 16000+ Real-World APIs Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, Maosong Sun	6644

Stylized Offline Reinforcement Learning: Extracting Diverse High-Quality Behaviors from Heterogeneous Datasets Yihuan Mao, Chengjie Wu, Xi Chen, Hao Hu, Ji Jiang, Tianze Zhou, Tangjie Lv, Changjie Fan, Zhipeng Hu, Yi Wu, Yujing Hu, Chongjie Zhang	6667
IRAD: Implicit Representation-Driven Image Resampling Against Adversarial Attacks Yue Cao, Tianlin Li, Xiaofeng Cao, Ivor Tsang, Yang Liu, Qing Guo	6688
VFLAIR: A Research Library and Benchmark for Vertical Federated Learning Tianyuan Zou, Zixuan Gu, Yu He, Hideaki Takahashi, Yang Liu, Ya-Qin Zhang	6712
Grounding Multimodal Large Language Models to the World Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, Qixiang Ye, Furu Wei	6751
Finite-State Autoregressive Entropy Coding for Efficient Learned Lossless Compression Yufeng Zhang, Hang Yu, Jianguo Li, Weiyao Lin	6775
Facing the Elephant in the Room: Visual Prompt Tuning Or Full Finetuning? Cheng Han, Qifan Wang, Yiming Cui, Wenguan Wang, Lifu Huang, Siyuan Qi, Dongfang Liu	6800
SyncDreamer: Generating Multiview-Consistent Images from a Single-View Image Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, Wenping Wang	6829
Hybrid Distillation: Connecting Masked Autoencoders with Contrastive Learners Bowen Shi, Xiaopeng Zhang, Yaoming Wang, Jin Li, Wenrui Dai, Junni Zou, Hongkai Xiong, Qi Tian	6851
ZipIt! Merging Models from Different Tasks Without Training George Stoica, Daniel Bolya, Jakob Brandt Bjorner, Pratik Ramesh, Taylor Hearn, Judy Hoffman	6869
Continuous Invariance Learning Lin Yong, Fan Zhou, Lu Tan, Lintao Ma, Jianmeng Liu, Yansu He, Yuan Yuan, Yu Liu, James Y. Zhang, Yujiu Yang, Hao Wang	6892
Be Careful What You Smooth For: Label Smoothing Can Be a Privacy Shield but Also a Catalyst for Model Inversion Attacks Lukas Struppek, Dominik Hintersdorf, Kristian Kersting	6917
Embodied Active Defense: Leveraging Recurrent Feedback to Counter Adversarial Patches Lingxuan Wu, Xiao Yang, Yinpeng Dong, Liuwei Xie, Hang Su, Jun Zhu	6948
Window Attention is Bugged: How Not to Interpolate Position Embeddings Daniel Bolya, Chaitanya Ryali, Judy Hoffman, Christoph Feichtenhofer	6977
LipVoicer: Generating Speech from Silent Videos Guided by Lip Reading Yochai Yemini, Aviv Shamsian, Lior Bracha, Sharon Gannot, Ethan Fetaya	6999

Fusion is Not Enough: Single Modal Attacks on Fusion Models for 3D Object Detection Zhiyuan Cheng, Hongjun Choi, Shiwei Feng, James Chenhao Liang, Guanhong Tao, Dongfang Liu, Michael Zuzak, Xiangyu Zhang	7019
Neural Field Classifiers Via Target Encoding and Classification Loss Xindi Yang, Zeke Xie, Xiong Zhou, Boyu Liu, Buhua Liu, Yi Liu, Haoran Wang, Yunfeng Cai, Mingming Sun	7044
Universal Jailbreak Backdoors from Poisoned Human Feedback Javier Rando, Florian Tramèr	7061
Plug-And-Play: An Efficient Post-Training Pruning Method for Large Language Models Yingtao Zhang, Haoli Bai, Haokun Lin, Jialin Zhao, Lu Hou, Carlo Vittorio Cannistraci	7089
Epitopological Learning and Cannistraci-Hebb Network Shape Intelligence Brain-Inspired Theory for Ultra-Sparse Advantage in Deep Learning Yingtao Zhang, Jialin Zhao, Wenjing Wu, Alessandro Muscoloni, Carlo Vittorio Cannistraci	7108
GIM: Learning Generalizable Image Matcher from Internet Videos Xuelun Shen, Zhipeng Cai, Wei Yin, Matthias Møller, Zijun Li, Kaixuan Wang, Xiaozhi Chen, Cheng Wang	7137
On the Generalization and Approximation Capacities of Neural Controlled Differential Equations Linus Bleistein, Agathe Guilloux	7153
Contrastive Learning is Spectral Clustering on Similarity Graph Zhiquan Tan, Yifan Zhang, Jingqin Yang, Yang Yuan	7190
Internal Cross-Layer Gradients for Extending Homogeneity to Heterogeneity in Federated Learning Yun-Hin Chan, Rui Zhou, Running Zhao, Zhihan Jiang, Edith C. H. Ngai	7207
Self-Supervised Speech Quality Estimation and Enhancement Using Only Clean Speech Szu-Wei Fu, Kuo-Hsuan Hung, Yu Tsao, Yu-Chiang Frank Wang	7236
Towards the Fundamental Limits of Knowledge Transfer Over Finite Domains Qingyue Zhao, Banghua Zhu	7258
GPT-4 is Too Smart to Be Safe: Stealthy Chat with LLMs Via Cipher Youliang Yuan, Wenxiang Jiao, Wenxuan Wang, Jen-Tse Huang, Pinjia He, Shuming Shi, Zhaopeng Tu	7291
Open-Ended VQA Benchmarking of Vision-Language Models by Exploiting Classification Datasets and Their Semantic Hierarchy Simon Ging, Maria Alejandra Bravo, Thomas Brox	7312
Learning Conditional Invariances Through Non-Commutativity Abhra Chaudhuri, Serban Georgescu, Anjan Dutta	7341

Real-Fake: Effective Training Data Synthesis Through Distribution Matching Jianhao Yuan, Jie Zhang, Shuyang Sun, Philip Torr, Bo Zhao	7362
A Real-World WebAgent with Planning, Long Context Understanding, and Program Synthesis Izzeddin Gur, Hiroki Furuta, Austin V Huang, Mustafa Safdari, Yutaka Matsuo, Douglas Eck, Aleksandra Faust	7383
Multimodal Web Navigation with Instruction-Finetuned Foundation Models Hiroki Furuta, Kuang-Huei Lee, Ofir Nachum, Yutaka Matsuo, Aleksandra Faust, Shixiang Shane Gu, Izzeddin Gur	7411
VertiBench: Advancing Feature Distribution Diversity in Vertical Federated Learning Benchmarks Zhaomin Wu, Junyi Hou, Bingsheng He	7442
G^S2N^S2^S : Weisfeiler and Lehman Go Grammatical Jason Piquenot, Aldo Moscatelli, Maxime Berar, Pierre Hăroux, Romain Raveaux, Jean-Yves Ramel, Săbastien Adam	7485
Symmetric Basis Convolutions for Learning Lagrangian Fluid Mechanics Rene Winchenbach, Nils Thuerey	7509
Efficient Backdoor Attacks for Deep Neural Networks in Real-World Scenarios Ziqiang Li, Hong Sun, Pengfei Xia, Heng Li, Beihao Xia, Yi Wu, Bin Li	7563
Scalable Language Model with Generalized Continual Learning Bohao Peng, Zhuotao Tian, Shu Liu, Ming-Chang Yang, Jiaya Jia	7602
CivRealm: A Learning and Reasoning Odyssey in Civilization for Decision-Making Agents Siyuan Qi, Shuo Chen, Yexin Li, Xiangyu Kong, Junqi Wang, Bangcheng Yang, Pring Wong, Yifan Zhong, Xiaoyuan Zhang, Zhaowei Zhang, Nian Liu, Yaodong Yang, Song-Chun Zhu	7625
Ground-A-Video: Zero-Shot Grounded Video Editing Using Text-To-Image Diffusion Models Hyeonho Jeong, Jong Chul Ye	7657
Spurious Feature Diversification Improves Out-Of-Distribution Generalization Lin Yong, Lu Tan, Yifan Hao, Ho Nam Wong, Hanze Dong, Weizhong Zhang, Yujiu Yang, Tong Zhang	7685
Modelling Complex Vector Drawings with Stroke-Clouds Alexander Ashcroft, Ayan Das, Yulia Gryaditskaya, Zhiyu Qu, Yi-Zhe Song	7742
PORF: POSE RESIDUAL FIELD FOR ACCURATE NEURAL SURFACE RECONSTRUCTION Jia-Wang Bian, Wenjing Bian, Victor Adrian Prisacariu, Philip Torr	7761
GNeRP: Gaussian-Guided Neural Reconstruction of Reflective Objects with Noisy Polarization Priors Li Yang, Ruizheng Wu, Jiyong Li, Ying-Cong Chen	7780

Accurate Forgetting for Heterogeneous Federated Continual Learning Abudukelimu Wuerkaixi, Sen Cui, Jingfeng Zhang, Kunda Yan, Bo Han, Gang Niu, Lei Fang, Changshui Zhang, Masashi Sugiyama	7799
Text-To-3D with Classifier Score Distillation Xin Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Song-Hai Zhang, Xiaojuan Qi	7818
Adversarial Causal Bayesian Optimization Scott Sussex, Pier Giuseppe Sessa, Anastasia Makarova, Andreas Krause	7833
Elastic Feature Consolidation for Cold Start Exemplar-Free Incremental Learning Simone Magistri, Tomaso Trinci, Albin Soutif, Joost Van De Weijer, Andrew D. Bagdanov	7855
Synapse: Trajectory-As-Exemplar Prompting with Memory for Computer Control Longtao Zheng, Rundong Wang, Xinrun Wang, Bo An	7878
GRAPH-CONSTRAINED DIFFUSION FOR END-TO-END PATH PLANNING Dingyuan Shi, Yongxin Tong, Zimu Zhou, Ke Xu, Zheng Wang, Jieping Ye	7909
Predicting Emergent Abilities with Infinite Resolution Evaluation Shengding Hu, Xin Liu, Xu Han, Xinrong Zhang, Chaoqun He, Weilin Zhao, Yankai Lin, Ning Ding, Zebin Ou, Guoyang Zeng, Zhiyuan Liu, Maosong Sun	7928
Batch Normalization is Sufficient for Universal Function Approximation in CNNs Rebekka Burkholz	7949
Communication-Efficient Gradient Descent-Accent Methods for Distributed Variational Inequalities: Unified Analysis and Local Updates Siqi Zhang, Sayantan Choudhury, Sebastian U Stich, Nicolas Loizou	7965
Deep Reinforcement Learning Guided Improvement Heuristic for Job Shop Scheduling Cong Zhang, Zhiguang Cao, Wen Song, Yaoxin Wu, Jie Zhang	7998
Uni3D: Exploring Unified 3D Representation at Scale Junsheng Zhou, Jinsheng Wang, Baorui Ma, Yu-Shen Liu, Tiejun Huang, Xinlong Wang	8019
Performance Gaps in Multi-View Clustering Under the Nested Matrix-Tensor Model Hugo Lebeau, Mohamed El Amine Seddik, Jos�� Henrique De Morais Goulart	8036
PBADet: A One-Stage Anchor-Free Approach for Part-Body Association Zhongpai Gao, Huayi Zhou, Abhishek Sharma, Meng Zheng, Benjamin Planche, Terrence Chen, Ziyang Wu	8065
RAIN: Your Language Models Can Align Themselves Without Finetuning Yuhui Li, Fangyun Wei, Jinjing Zhao, Chao Zhang, Hongyang Zhang	8082
DragonDiffusion: Enabling Drag-Style Manipulation on Diffusion Models Chong Mou, Xintao Wang, Jiechong Song, Ying Shan, Jian Zhang	8098

InstructDET: Diversifying Referring Object Detection with Generalized Instructions Ronghao Dang, Jiangyan Feng, Haodong Zhang, Chongjian Ge, Lin Song, Lijun Gong, Chengju Liu, Qijun Chen, Feng Zhu, Rui Zhao, Yibing Song	8110
TapMo: Shape-Aware Motion Generation of Skeleton-Free Characters Jiaxu Zhang, Shaoli Huang, Zhigang Tu, Xin Chen, Xiaohang Zhan, Gang Yu, Ying Shan	8139
Nearly $\mathcal{O}(n)$-Linear Convergence Bounds for Diffusion Models Via Stochastic Localization Joe Benton, Valentin De Bortoli, Arnaud Doucet, George Deligiannidis	8154
Bridging State and History Representations: Understanding Self-Predictive RL Tianwei Ni, Benjamin Eysenbach, Erfan Seyedsalehi, Michel Ma, Clement Gehring, Aditya Mahajan, Pierre-Luc Bacon	8175
ChatEval: Towards Better LLM-Based Evaluators Through Multi-Agent Debate Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, Zhiyuan Liu	8190
When Should We Prefer Decision Transformers for Offline Reinforcement Learning? Prajjwal Bhargava, Rohan Chitnis, Alborz Geramifard, Shagun Sodhani, Amy Zhang	8205
Scaling Laws for Sparsely-Connected Foundation Models Elias Frantar, Carlos Riquelme Ruiz, Neil Houlsby, Dan Alistarh, Utku Evci	8245
The Expressive Power of Transformers with Chain of Thought William Merrill, Ashish Sabharwal	8265
DENEUIL: TOWARDS DECIPHERING AND NAVIGATING THE ETHICAL VALUES OF LARGE LANGUAGE MODELS VIA INSTRUCTION LEARNING Shitong Duan, Xiaoyuan Yi, Peng Zhang, Tun Lu, Xing Xie, Ning Gu	8282
Function-Space Parameterization of Neural Networks for Sequential Learning Aidan Scannell, Riccardo Mereu, Paul Edmund Chang, Ella Tamir, Joni Pajarinen, Arno Solin	8322
On Penalty Methods for Nonconvex Bilevel Optimization and First-Order Stochastic Approximation Jeongyeol Kwon, Dohyun Kwon, Stephen Wright, Robert D Nowak	8351
Bayesian Low-Rank Adaptation for Large Language Models Adam X. Yang, Maxime Robeyns, Xi Wang, Laurence Aitchison	8397
LLMCarbon: Modeling the End-To-End Carbon Footprint of Large Language Models Ahmad Faiz, Sotaro Kaneda, Ruhan Wang, Rita Chukwunyere Osi, Prateek Sharma, Fan Chen, Lei Jiang	8428
Consciousness-Inspired Spatio-Temporal Abstractions for Better Generalization in Reinforcement Learning Harry Zhao, Safa Alver, Harm Van Seijen, Romain Laroche, Doina Precup, Yoshua Bengio	8443

Amortizing Intractable Inference in Large Language Models Edward J Hu, Moksh Jain, Eric Elmoznino, Younesse Kaddar, Guillaume Lajoie, Yoshua Bengio, Nikolay Malkin	8478
Cleanba: A Reproducible and Efficient Distributed Reinforcement Learning Platform Shengyi Huang, Jiayi Weng, Rujikorn Charakorn, Min Lin, Zhongwen Xu, Santiago Ontanon	8509
Learning to Act from Actionless Videos Through Dense Correspondences Po-Chen Ko, Jiayuan Mao, Yilun Du, Shao-Hua Sun, Joshua B. Tenenbaum	8521
Meta Inverse Constrained Reinforcement Learning: Convergence Guarantee and Generalization Analysis Shicheng Liu, Minghui Zhu	8542
DreamLLM: Synergistic Multimodal Comprehension and Creation Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, Xiangwen Kong, Xiangyu Zhang, Kaisheng Ma, Li Yi	8580
How to Capture Higher-Order Correlations? Generalizing Matrix Softmax Attention to Kronecker Computation Josh Alman, Zhao Song	8617
Structural Estimation of Partially Observed Linear Non-Gaussian Acyclic Model: A Practical Approach with Identifiability Songyao Jin, Feng Xie, Guangyi Chen, Biwei Huang, Zhengming Chen, Xinshuai Dong, Kun Zhang	8645
Rethinking Label Poisoning for GNNs: Pitfalls and Attacks Vijay Lingam, Mohammad Sadegh Akhondzadeh, Aleksandar Bojchevski	8672
Improving Protein Optimization with Smoothed Fitness Landscapes Andrew Kirjner, Jason Yim, Raman Samusevich, Shahar Bracha, Tommi S. Jaakkola, Regina Barzilay, Ila R Fiete	8701
Learning Implicit Representation for Reconstructing Articulated Objects Hao Zhang, Fang Li, Samyak Rawlekar, Narendra Ahuja	8720
Delta-AI: Local Objectives for Amortized Inference in Sparse Graphical Models Jean-Pierre Renã© Falet, Hae Beom Lee, Nikolay Malkin, Chen Sun, Dragos Secrieru, Dinghui Zhang, Guillaume Lajoie, Yoshua Bengio	8742
How Does Unlabeled Data Provably Help Out-Of-Distribution Detection? Xuefeng Du, Zhen Fang, Ilias Diakonikolas, Yixuan Li	8763
BECLR: Batch Enhanced Contrastive Few-Shot Learning Stylianios Poulakakis-Daktylidis, Hadi Jamali-Rad	8802
FLATTEN: Optical FLOW-Guided ATTENTION for Consistent Text-To-Video Editing Yuren Cong, Mengmeng Xu, Christian Simon, Shoufa Chen, Jiawei Ren, Yanping Xie, Juan-Manuel Perez-Rua, Bodo Rosenhahn, Tao Xiang, Sen He	8826

HIFA: High-Fidelity Text-To-3D Generation with Advanced Diffusion Guidance Junzhe Zhu, Peiye Zhuang, Sanmi Koyejo	8847
Unveiling Options with Neural Network Decomposition Mahdi Alikhasi, Levi Lelis	8872
NOLA: Compressing LoRA Using Linear Combination of Random Basis Soroush Abbasi Koohpayegani, Navaneet K L, Parsa Nooralinejad, Soheil Kolouri, Hamed Pirsiavash	8894
Patched Denoising Diffusion Models for High-Resolution Image Synthesis Zheng Ding, Mengqi Zhang, Jiajun Wu, Zhuowen Tu	8911
Complete and Efficient Graph Transformers for Crystal Material Property Prediction Keqiang Yan, Cong Fu, Xiaofeng Qian, Xiaoning Qian, Shuiwang Ji	8934
The Effect of Intrinsic Dataset Properties on Generalization: Unraveling Learning Differences Between Natural and Medical Images Nicholas Konz, Maciej A Mazurowski	8961
Pose Modulated Avatars from Video Chunjin Song, Bastian Wandt, Helge Rhodin	8973
Circumventing Concept Erasure Methods for Text-To-Image Generative Models Minh Pham, Kelly O. Marshall, Niv Cohen, Govind Mittal, Chinmay Hegde	8998
SWAP: Sparse Entropic Wasserstein Regression for Robust Network Pruning Lei You, Hei Victor Cheng	9024
Meaning Representations from Trajectories in Autoregressive Models Tian Yu Liu, Matthew Trager, Alessandro Achille, Pramuditha Perera, Luca Zancato, Stefano Soatto	9047
Trajenglish: Traffic Modeling as Next-Token Prediction Jonah Phillion, Xue Bin Peng, Sanja Fidler	9070
Retroformer: Retrospective Large Language Agents with Policy Gradient Optimization Weiran Yao, Shelby Heinecke, Juan Carlos Niebles, Zhiwei Liu, Yihao Feng, Le Xue, Rithesh R N, Zeyuan Chen, Jianguo Zhang, Devansh Arpit, Ran Xu, Phil L Mui, Huan Wang, Caiming Xiong, Silvio Savarese	9090
INViTE: Interpret and Control Vision-Language Models with Text Explanations Haozhe Chen, Junfeng Yang, Carl Vondrick, Chengzhi Mao	9111
Set Learning for Accurate and Calibrated Models Lukas Muttenthaler, Robert A. Vandermeulen, Qiuyi Zhang, Thomas Unterthiner, Klaus Robert Muller	9131
Symphony: Symmetry-Equivariant Point-Centered Spherical Harmonics for 3D Molecule Generation Ameya Daigavane, Song Eun Kim, Mario Geiger, Tess Smidt	9163

Accurate and Scalable Estimation of Epistemic Uncertainty for Graph Neural Networks <i>Puja Trivedi, Mark Heimann, Rushil Anirudh, Danai Koutra, Jayaraman J. Thiagarajan</i>	9191
Adaptive Regularization of Representation Rank as an Implicit Constraint of Bellman Equation <i>Qiang He, Tianyi Zhou, Meng Fang, Setareh Maghsudi</i>	9224
On-Policy Distillation of Language Models: Learning from Self-Generated Mistakes <i>Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos Garea, Matthieu Geist, Olivier Bachem</i>	9249
SalUn: Empowering Machine Unlearning Via Gradient-Based Weight Saliency in Both Image Classification and Generation <i>Chongyu Fan, Jiancheng Liu, Yihua Zhang, Eric Wong, Dennis Wei, Sijia Liu</i>	9267
Pooling Image Datasets with Multiple Covariate Shift and Imbalance <i>Sotirios Panagiotis Chytas, Vishnu Suresh Lokhande, Vikas Singh</i>	9298
Ferret: Refer and Ground Anything Anywhere at Any Granularity <i>Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, Yinfei Yang</i>	9321
Sin3DM: Learning a Diffusion Model from a Single 3D Textured Shape <i>Rundi Wu, Ruoshi Liu, Carl Vondrick, Changxi Zheng</i>	9349
TEST: Text Prototype Aligned Embedding to Activate LLM's Ability for Time Series <i>Chenxi Sun, Hongyan Li, Yaliang Li, Shenda Hong</i>	9372
DynaVol: Unsupervised Learning for Dynamic Scenes Through Object-Centric Voxelization <i>Yanpeng Zhao, Siyu Gao, Yunbo Wang, Xiaokang Yang</i>	9400
A Cognitive Model for Learning Abstract Relational Structures from Memory-Based Decision-Making Tasks <i>Haruo Hosoya</i>	9420
Making LLaMA SEE and Draw with SEED Tokenizer <i>Yuying Ge, Sijie Zhao, Ziyun Zeng, Yixiao Ge, Chen Li, Xintao Wang, Ying Shan</i>	9438
Domain-Inspired Sharpness-Aware Minimization Under Domain Shifts <i>Ruipeng Zhang, Ziqing Fan, Jiangchao Yao, Ya Zhang, Yanfeng Wang</i>	9464
Lewis's Signaling Game as beta-VAE for Natural Word Lengths and Segments <i>Ryo Ueda, Tadahiyo Taniguchi</i>	9491
LightHGNN: Distilling Hypergraph Neural Networks into MLPs for 100x Faster Inference <i>Yifan Feng, Yihe Luo, Shihui Ying, Yue Gao</i>	9518
NECO: NEural Collapse Based Out-Of-Distribution Detection <i>Mouā'n Ben Ammar, Nacim Belkhir, Sebastian Popescu, Antoine Manzanera, Gianni Franchi</i>	9539

Balancing Act: Constraining Disparate Impact in Sparse Models Meraj Hashemizadeh, Juan Ramirez, Rohan Sukumaran, Golnoosh Farnadi, Simon Lacoste-Julien, Jose Gallego-Posada	9569
Motion Guidance: Diffusion-Based Image Editing with Differentiable Motion Estimators Daniel Geng, Andrew Owens	9602
Octavius: Mitigating Task Interference in MLLMs Via LoRA-MoE Zeren Chen, Ziqin Wang, Zhen Wang, Huayang Liu, Zhenfei Yin, Si Liu, Lu Sheng, Wanli Ouyang, Jing Shao	9624
ReSimAD: Zero-Shot 3D Domain Transfer for Autonomous Driving with Source Reconstruction and Target Simulation Bo Zhang, Xinyu Cai, Jiakang Yuan, Donglin Yang, Jianfei Guo, Xiangchao Yan, Renqiu Xia, Botian Shi, Min Dou, Tao Chen, Si Liu, Junchi Yan, Yu Qiao	9646
VONet: Unsupervised Video Object Learning with Parallel U-Net Attention and Object-Wise Sequential VAE Haonan Yu, Wei Xu	9661
OmniQuant: Omnidirectionally Calibrated Quantization for Large Language Models Wenqi Shao, Mengzhao Chen, Zhaoyang Zhang, Peng Xu, Lirui Zhao, Zhiqian Li, Kaipeng Zhang, Peng Gao, Yu Qiao, Ping Luo	9680
DreamSmooth: Improving Model-Based Reinforcement Learning Via Reward Smoothing Vint Lee, Pieter Abbeel, Youngwoon Lee	9705
FreeDyG: Frequency Enhanced Continuous-Time Dynamic Graph Model for Link Prediction Yuxing Tian, Yiyang Qi, Fan Guo	9725
THOUGHT PROPAGATION: AN ANALOGICAL APPROACH TO COMPLEX REASONING WITH LARGE LANGUAGE MODELS Junchi Yu, Ran He, Zhitao Ying	9745
LongLoRA: Efficient Fine-Tuning of Long-Context Large Language Models Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, Jiaya Jia	9772
Zero-Shot Robustification of Zero-Shot Models Dyah Adila, Changho Shin, Linrong Cai, Frederic Sala	9791
HypeBoy: Generative Self-Supervised Representation Learning on Hypergraphs Sunwoo Kim, Shinhwan Kang, Fanchen Bu, Soo Yong Lee, Jaemin Yoo, Kijung Shin	9824
Symbol as Points: Panoptic Symbol Spotting Via Point-Based Representation Wenlong Liu, Tianyu Yang, Yuhang Wang, Qizhi Yu, Lei Zhang	9856
Enhancing One-Shot Federated Learning Through Data and Ensemble Co-Boosting Rong Dai, Yonggang Zhang, Ang Li, Tongliang Liu, Xun Yang, Bo Han	9877

EX-Graph: A Pioneering Dataset Bridging Ethereum and X Qian Wang, Zhen Zhang, Zemin Liu, Shengliang Lu, Bingqiao Luo, Bingsheng He	9898
Aux-NAS: Exploiting Auxiliary Labels with Negligibly Extra Inference Cost Yuan Gao, Weizhong Zhang, Wenhan Luo, Lin Ma, Jin-Gang Yu, Gui-Song Xia, Jiayi Ma	9927
What Does Automatic Differentiation Compute for Neural Networks? Sejun Park, Sanghyuk Chun, Wonyeol Lee	9943
Compositional Generative Inverse Design Tailin Wu, Takashi Maruyama, Long Wei, Tao Zhang, Yilun Du, Gianluca Iaccarino, Jure Leskovec	9970
Simplifying Transformer Blocks Bobby He, Thomas Hofmann	10000
Dynamic Sparse No Training: Training-Free Fine-Tuning for Sparse LLMs Yuxin Zhang, Lirui Zhao, Mingbao Lin, Sun Yunyun, Yiwu Yao, Xingjia Han, Jared Tanner, Shiwei Liu, Rongrong Ji	10029
SEABO: A Simple Search-Based Method for Offline Imitation Learning Jiafei Lyu, Xiaoteng Ma, Le Wan, Runze Liu, Xiu Li, Zongqing Lu	10045
Analyzing and Improving Optimal-Transport-Based Adversarial Networks Jaemoo Choi, Jaewoong Choi, Myungjoo Kang	10072
HYPPO: Hyperspherical Out-Of-Distribution Generalization Haoyue Bai, Yifei Ming, Julian Katz-Samuels, Yixuan Li	10099
Sentence-Level Prompts Benefit Composed Image Retrieval Yang Bai, Xinxing Xu, Yong Liu, Salman Khan, Fahad Khan, Wangmeng Zuo, Rick Siow Mong Goh, Chun-Mei Feng	10126
PixArt-\$\alpha\$: Fast Training of Diffusion Transformer for Photorealistic Text-To-Image Synthesis Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, Zhenguo Li	10144
SaProt: Protein Language Modeling with Structure-Aware Vocabulary Jin Su, Chenchen Han, Yuyang Zhou, Junjie Shan, Xibin Zhou, Fajie Yuan	10174
LLM-Assisted Code Cleaning for Training Accurate Code Generators Naman Jain, Tianjun Zhang, Wei-Lin Chiang, Joseph E. Gonzalez, Koushik Sen, Ion Stoica	10197
Are Human-Generated Demonstrations Necessary for In-Context Learning? Rui Li, Guoyin Wang, Jiwei Li	10228
EMO: EARTH MOVER DISTANCE OPTIMIZATION FOR AUTO-REGRESSIVE LANGUAGE MODELING Siyu Ren, Zhiyong Wu, Kenny Q. Zhu	10265

Matrix Manifold Neural Networks++ Xuan Son Nguyen, Shuo Yang, Aymeric Histace	10286
Seer: Language Instructed Video Prediction with Latent Diffusion Models Xianfan Gu, Chuan Wen, Weirui Ye, Jiaming Song, Yang Gao	10317
Quick-Tune: Quickly Learning Which Pretrained Model to Finetune and How Sebastian Pineda Arango, Fabio Ferreira, Arlind Kadra, Frank Hutter, Josif Grabocka	10349
A Foundation Model for Error Correction Codes Yoni Choukroun, Lior Wolf	10369
Debiasing Algorithm Through Model Adaptation Tomasz Limisiewicz, David Mareš, Tomáš Musil	10388
Fairret: A Framework for Differentiable Fairness Regularization Terms Maarten Buyt, Marybeth Defrance, Tijl De Bie	10408
Deep Reinforcement Learning for Modelling Protein Complexes Ziqi Gao, Tao Feng, Jiaxuan You, Chenyi Zi, Yan Zhou, Chen Zhang, Jia Li	10436
Continual Momentum Filtering on Parameter Space for Online Test-Time Adaptation Jae-Hong Lee, Joon-Hyuk Chang	10453
MINDE: Mutual Information Neural Diffusion Estimation Giulio Franzese, Mustapha Bounoua, Pietro Michiardi	10486
Boosting Graph Anomaly Detection with Adaptive Message Passing Jingyan Chen, Guanghui Zhu, Chunfeng Yuan, Yihua Huang	10518
KITAB: Evaluating LLMs on Constraint Satisfaction for Information Retrieval Marah I Abidin, Suriya Gunasekar, Varun Chandrasekaran, Jerry Li, Mert Yuksekgonul, Rahee Ghosh Peshawaria, Ranjita Naik, Besmira Nushi	10540
Towards Robust and Efficient Cloud-Edge Elastic Model Adaptation Via Selective Entropy Distillation Yaofu Chen, Shuaicheng Niu, Yaowei Wang, Shoukai Xu, Hengjie Song, Minghui Tan	10563
Rethinking Information-Theoretic Generalization: Loss Entropy Induced PAC Bounds Yuxin Dong, Tieliang Gong, Hong Chen, Shujian Yu, Chen Li	10594
Few-Shot Hybrid Domain Adaptation of Image Generator Hengjia Li, Yang Liu, Linxuan Xia, Yuqi Lin, Wenxiao Wang, Tu Zheng, Zheng Yang, Xiaohui Zhong, Xiaobo Ren, Xiaofei He	10641
Analysis of Learning a Flow-Based Generative Model from Limited Sample Complexity Hugo Cui, Florent Krzakala, Eric Vanden-Eijnden, Lenka Zdeborova	10661

Understanding In-Context Learning from Repetitions Jianhao Yan, Jin Xu, Chiyu Song, Chenming Wu, Yafu Li, Yue Zhang	10688
When Do Prompting and Prefix-Tuning Work? a Theory of Capabilities and Limitations Aleksandar Petrov, Philip Torr, Adel Bibi	10710
#InsTag: Instruction Tagging for Analyzing Supervised Fine-Tuning of Large Language Models Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, Jingren Zhou	10734
NfgTransformer: Equivariant Representation Learning for Normal-Form Games Siqi Liu, Luke Marris, Georgios Piliouras, Ian Gemp, Nicolas Heess	10753
BatteryML: An Open-Source Platform for Machine Learning on Battery Degradation Han Zhang, Xiaofan Gui, Shun Zheng, Ziheng Lu, Yuqi Li, Jiang Bian	10765
Physics-Regulated Deep Reinforcement Learning: Invariant Embeddings Hongpeng Cao, Yanbing Mao, Lui Sha, Marco Caccamo	10794
Spatio-Temporal Approximation: A Training-Free SNN Conversion for Transformers Yizhou Jiang, Kunlin Hu, Tianren Zhang, Haichuan Gao, Yuqian Liu, Ying Fang, Feng Chen	10839
Optimal Sample Complexity for Average Reward Markov Decision Processes Shengbo Wang, Jose Blanchet, Peter Glynn	10861
Single Motion Diffusion Sigal Raab, Inbal Leibovitch, Guy Tevet, Moab Arar, Amit Haim Bermano, Daniel Cohen-Or	10880
A Simple Romance Between Multi-Exit Vision Transformer and Token Reduction Dongyang Liu, Meina Kan, Shiguang Shan, Xilin Chen	10901
InstructScene: Instruction-Driven 3D Indoor Scene Synthesis with Semantic Graph Prior Chenguo Lin, Yadong Mu	10915
Long-Tailed Diffusion Models with Oriented Calibration Tianjiao Zhang, Huangjie Zheng, Jiangchao Yao, Xiangfeng Wang, Mingyuan Zhou, Ya Zhang, Yanfeng Wang	10947
Parallelizing Non-Linear Sequential Models Over the Sequence Length Yi Heng Lim, Qi Zhu, Joshua Selfridge, Muhammad Firmansyah Kasim	10969
Task Planning for Visual Room Rearrangement Under Partial Observability Karan Mirakhor, Sourav Ghosh, Dipanjan Das, Brojeshwar Bhowmick	10996
CIFAR-10-Warehouse: Broad and More Realistic Testbeds in Model Generalization Analysis Xiaoxiao Sun, Xingjian Leng, Zijian Wang, Yang Yang, Zi Huang, Liang Zheng	11021

The All-Seeing Project: Towards Panoptic Visual Recognition and Understanding of the Open World Weiyun Wang, Min Shi, Qingyun Li, Wenhai Wang, Zhenhang Huang, Linjie Xing, Zhe Chen, Hao Li, Xizhou Zhu, Zhiguo Cao, Yushi Chen, Tong Lu, Jifeng Dai, Yu Qiao	11049
Interventional Fairness on Partially Known Causal Graphs: A Constrained Optimization Approach Aoqi Zuo, Yiqing Li, Susan Wei, Mingming Gong	11082
FROSTER: Frozen CLIP is a Strong Teacher for Open-Vocabulary Action Recognition Xiaohu Huang, Hao Zhou, Kun Yao, Kai Han	11117
Towards Seamless Adaptation of Pre-Trained Models for Visual Place Recognition Feng Lu, Lijun Zhang, Xiangyuan Lan, Shuting Dong, Yaowei Wang, Chun Yuan	11133
π^2vec: Policy Representation with Successor Features Gianluca Scarpellini, Ksenia Konyushkova, Claudio Fantacci, Thomas Paine, Yutian Chen, Misha Denil	11155
Out-Of-Distribution Detection with Negative Prompts Jun Nie, Yonggang Zhang, Zhen Fang, Tongliang Liu, Bo Han, Xinmei Tian	11173
In-Context Exploration-Exploitation for Reinforcement Learning Zhenwen Dai, Federico Tomasi, Sina Ghiassian	11193
Mitigating the Curse of Dimensionality for Certified Robustness Via Dual Randomized Smoothing Song Xia, Yi Yu, Xudong Jiang, Henghui Ding	11209
RetroBridge: Modeling Retrosynthesis with Markov Bridges Ilia Igashov, Arne Schneuing, Marwin Segler, Michael M. Bronstein, Bruno Correia	11227
SF(DA)π^2: Source-Free Domain Adaptation Through the Lens of Data Augmentation Uiwon Hwang, Jonghyun Lee, Juhyeon Shin, Sungroh Yoon	11246
Conformal Risk Control Anastasios Nikolas Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, Tal Schuster	11265
Encoding Unitig-Level Assembly Graphs with Heterophilous Constraints for Metagenomic Contigs Binning Hansheng Xue, Vijini Mallawaarachchi, Lexing Xie, Vaibhav Rajan	11286
AffineQuant: Affine Transformation Quantization for Large Language Models Yuexiao Ma, Huixia Li, Xiawu Zheng, Feng Ling, Xuefeng Xiao, Rui Wang, Shilei Wen, Fei Chao, Rongrong Ji	11308
DAFA: Distance-Aware Fair Adversarial Training Hyungyu Lee, Saehyung Lee, Hyemi Jang, Junsung Park, Ho Bae, Sungroh Yoon	11328

Transport Meets Variational Inference: Controlled Monte Carlo Diffusions Francisco Vargas, Shreyas Padhy, Denis Blessing, Nikolas Nă¼sken	11358
Spoken Question Answering and Speech Continuation Using Spectrogram-Powered LLM Eliya Nachmani, Alon Levkovitch, Roy Hirsch, Julian Salazar, Chulayuth Asawaroengchai, Soroosh Mariooryad, Ehud Rivlin, Rj Skerry-Ryan, Michelle Tadmor Ramanovich	11401
Fast Equilibrium of SGD in Generic Situations Zhiyuan Li, Yi Wang, Zhiren Wang	11417
A Semantic Invariant Robust Watermark for Large Language Models Aiwei Liu, Leyi Pan, Xuming Hu, Shiao Meng, Lijie Wen	11446
Self-Supervised Contrastive Learning for Long-Term Forecasting Junwoo Park, Daehoon Gwak, Jaegul Choo, Edward Choi	11467
Separating Common from Salient Patterns with Contrastive Representation Learning Robin Louiset, Edouard Duchesnay, Antoine Grigis, Pietro Gori	11490
Faithful Explanations of Black-Box NLP Models Using LLM-Generated Counterfactuals Yair Ori Gat, Nitay Calderon, Amir Feder, Alexander Chapanin, Amit Sharma, Roi Reichart	11522
DecompOpt: Controllable and Decomposed Diffusion Models for Structure-Based Molecular Optimization Xiangxin Zhou, Xiwei Cheng, Yuwei Yang, Yu Bao, Liang Wang, Quanquan Gu	11556
Rethinking Channel Dimensions to Isolate Outliers for Low-Bit Weight Quantization of Large Language Models Jung Hwan Heo, Jeonghoon Kim, Beomseok Kwon, Byeongwook Kim, Se Jung Kwon, Dongsoo Lee	11579
Biased Temporal Convolution Graph Network for Time Series Forecasting with Missing Values Xiaodan Chen, Xiucheng Li, Bo Liu, Zhijun Li	11598
Most Discriminative Stimuli for Functional Cell Type Clustering Max F Burg, Thomas Zenkel, Michaela Vystră¼lovă¼, Jonathan Oesterle, Larissa Hă¼fling, Konstantin Friedrich Willeke, Jan Lause, Sarah Mă¼ller, Paul G. Fahey, Zhiwei Ding, Kelli Restivo, Shashwat Sridhar, Tim Gollisch, Philipp Berens, Andreas S. Tolias, Thomas Euler, Matthias Bethge, Alexander S Ecker	11617
A Data-Driven Measure of Relative Uncertainty for Misclassification Detection Eduardo Dadalto Că¼mara Gomes, Marco Romanelli, Georg Pichler, Pablo Piantanida	11643
Faithful and Efficient Explanations for Neural Networks Via Neural Tangent Kernel Surrogate Models Andrew William Engel, Zhichao Wang, Natalie Frank, Ioana Dumitriu, Sutanay Choudhury, Anand Sarwate, Tony Chiang	11666

Breaking Physical and Linguistic Borders: Multilingual Federated Prompt Tuning for Low-Resource Languages Wanru Zhao, Yihong Chen, Royson Lee, Xinchu Qiu, Yan Gao, Hongxiang Fan, Nicholas Donald Lane	11720
The Devil is in the Object Boundary: Towards Annotation-Free Instance Segmentation Using Foundation Models Cheng Shi, Sibe Yang	11736
On the Generalization Capacity of Neural Networks During Generic Multimodal Reasoning Takuya Ito, Soham Dan, Mattia Rigotti, James Kozloski, Murray Campbell	11762
Statistical Rejection Sampling Improves Preference Optimization Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J Liu, Jialu Liu	11781
DiffAR: Denoising Diffusion Autoregressive Model for Raw Speech Waveform Generation Roi Benita, Michael Elad, Joseph Keshet	11801
Discovering Failure Modes of Text-Guided Diffusion Models Via Adversarial Search Qihao Liu, Adam Kortylewski, Yutong Bai, Song Bai, Alan Yuille	11823
Grokking in Linear Estimators -- A Solvable Model that Groks Without Understanding Noam Itzhak Levi, Alon Beck, Yohai Bar-Sinai	11848
Detecting, Explaining, and Mitigating Memorization in Diffusion Models Yuxin Wen, Yuchen Liu, Chen Chen, Lingjuan Lyu	11865
Interpretable Meta-Learning of Physical Systems Matthieu Blanke, Marc Lelarge	11881
Finite Scalar Quantization: VQ-VAE Made Simple Fabian Mentzer, David Minnen, Eirikur Agustsson, Michael Tschannen	11899
Navigating Text-To-Image Customization: From LyCORIS Fine-Tuning to Model Evaluation Shih-Ying Yeh, Yu-Guan Hsieh, Zhidong Gao, Bernard B W Yang, Giyeong Oh, Yanmin Gong	11911
Learning Robust Generalizable Radiance Field with Visibility and Feature Augmented Point Representation Wang Jiaxu, Ziyi Zhang, Renjing Xu	11923
Plug-And-Play Posterior Sampling Under Mismatched Measurement and Prior Models Marien Renaud, Jiaming Liu, Valentin De Bortoli, Andres Almansa, Ulugbek Kamilov	11945
Can Transformers Capture Spatial Relations Between Objects? Chuan Wen, Dinesh Jayaraman, Yang Gao	11975
A Hierarchical Bayesian Model for Few-Shot Meta Learning Minyoung Kim, Timothy Hospedales	11996

COPlanner: Plan to Roll Out Conservatively but to Explore Optimistically for Model-Based RL <i>Xiyao Wang, Ruijie Zheng, Yanchao Sun, Ruonan Jia, Wichayaporn Wongkamjan, Huazhe Xu, Furong Huang</i>	12024
Provably Efficient CVaR RL in Low-Rank MDPs <i>Yulai Zhao, Wenhao Zhan, Xiaoyan Hu, Ho-Fung Leung, Farzan Farnia, Wen Sun, Jason D. Lee</i>	12046
ADOPD: A Large-Scale Document Page Decomposition Dataset <i>Jiuxiang Gu, Xiangxi Shi, Jason Kuen, Lu Qi, Ruiyi Zhang, Anqi Liu, Ani Nenkova, Tong Sun</i>	12076
Adaptive Chameleon Or Stubborn Sloth: Revealing the Behavior of Large Language Models in Knowledge Conflicts <i>Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, Yu Su</i>	12094
FedCompass: Efficient Cross-Silo Federated Learning on Heterogeneous Client Devices Using a Computing Power-Aware Scheduler <i>Zilinghan Li, Pranshu Chaturvedi, Shilan He, Han Chen, Gagandeep Singh, Volodymyr Kindratenko, Eliu A Huerta, Kibaek Kim, Ravi Madduri</i>	12118
Automatic Functional Differentiation in JAX <i>Min Lin</i>	12155
$\mathcal{B}$-Coder: Value-Based Deep Reinforcement Learning for Program Synthesis <i>Zishun Yu, Yunzhe Tao, Liyu Chen, Tao Sun, Hongxia Yang</i>	12176
High-Dimensional SGD Aligns with Emerging Outlier Eigenspaces <i>Gerard Ben Arous, Reza Gheissari, Jiaoyang Huang, Aukosh Jagannath</i>	12197
Accelerating Distributed Stochastic Optimization Via Self-Repellent Random Walks <i>Jie Hu, Vishwaraj Doshi, Do Young Eun</i>	12244
On Stationary Point Convergence of PPO-Clip <i>Ruinan Jin, Shuai Li, Baoxiang Wang</i>	12286
Butterfly Effects of SGD Noise: Error Amplification in Behavior Cloning and Autoregression <i>Adam Block, Dylan J Foster, Akshay Krishnamurthy, Max Simchowitz, Cyril Zhang</i>	12304
Beyond Task Performance: Evaluating and Reducing the Flaws of Large Multimodal Models with In-Context-Learning <i>Mustafa Shukor, Alexandre Rame, Corentin Dancette, Matthieu Cord</i>	12363
Video Decomposition Prior: Editing Videos Layer by Layer <i>Gaurav Shrivastava, Ser-Nam Lim, Abhinav Shrivastava</i>	12394
Generating Images with 3D Annotations Using Diffusion Models <i>Wufei Ma, Qihao Liu, Jiahao Wang, Angtian Wang, Xiaoding Yuan, Yi Zhang, Zihao Xiao, Guofeng Zhang, Beijia Lu, Ruxiao Duan, Yongrui Qi, Adam Kortylewski, Yaoyao Liu, Alan Yuille</i>	12414

Continual Learning on a Diet: Learning from Sparsely Labeled Streams Under Constrained Computation Wenxuan Zhang, Youssef Mohamed, Bernard Ghanem, Philip Torr, Adel Bibi, Mohamed Elhoseiny	12435
LpNTK: Better Generalisation with Less Data Via Sample Interaction During Learning Shangmin Guo, Yi Ren, Stefano V Albrecht, Kenny Smith	12452
Closing the Gap Between TD Learning and Supervised Learning - A Generalisation Point of View. Raj Ghugare, Matthieu Geist, Glen Berseth, Benjamin Eysenbach	12486
Improving Domain Generalization with Domain Relations Huaxiu Yao, Xinyu Yang, Xinyi Pan, Shengchao Liu, Pang Wei Koh, Chelsea Finn	12505
Learning Multi-Agent Communication with Contrastive Learning Yat Long Lo, Biswa Sengupta, Jakob Nicolaus Foerster, Michael Noukhovitch	12528
Domain Constraints Improve Risk Prediction When Outcome Data is Missing Sidhika Balachandar, Nikhil Garg, Emma Pierson	12546
Cameras as Rays: Pose Estimation Via Ray Diffusion Jason Y. Zhang, Amy Lin, Moneish Kumar, Tzu-Hsuan Yang, Deva Ramanan, Shubham Tulsiani	12563
Less is More: Fewer Interpretable Region Via Submodular Subset Selection Ruoyu Chen, Hua Zhang, Siyuan Liang, Jingzhi Li, Xiaochun Cao	12585
On Gauge Freedom, Conservativity and Intrinsic Dimensionality Estimation in Diffusion Models Christian Horvat, Jean-Pascal Pfister	12608
PRIME: Prioritizing Interpretability in Failure Mode Extraction Keivan Rezaei, Mehrdad Saberi, Mazda Moayeri, Soheil Feizi	12625
Point2SSM: Learning Morphological Variations of Anatomies from Point Clouds Jadie Adams, Shireen Elhabian	12649
Tell Your Model Where to Attend: Post-Hoc Attention Steering for LLMs Qingru Zhang, Chandan Singh, Liyuan Liu, Xiaodong Liu, Bin Yu, Jianfeng Gao, Tuo Zhao	12670
GraphPulse: Topological Representations for Temporal Graph Property Prediction Kiarash Shamsi, Farimah Poursafaei, Shenyang Huang, Bao Tran Gia Ngo, Baris Coskunuzer, Cuneyt Gurcan Akcora	12690
Principled Federated Domain Adaptation: Gradient Projection and Auto-Weighting Enyi Jiang, Yibo Jacky Zhang, Sanmi Koyejo	12714
Local Graph Clustering with Noisy Labels Artur Back De Luca, Kimon Fountoulakis, Shenghao Yang	12750

Incremental Randomized Smoothing Certification Shubham Ugare, Tarun Suresh, Debangshu Banerjee, Gagandeep Singh, Sasa Misailovic	12777
Denoising Diffusion Bridge Models Linqi Zhou, Aaron Lou, Samar Khanna, Stefano Ermon	12799
Chain-Of-Table: Evolving Tables in the Reasoning Chain for Table Understanding Zilong Wang, Hao Zhang, Chun-Liang Li, Julian Martin Eisenschlos, Vincent Perot, Zifeng Wang, Lesly Miculicich, Yasuhisa Fujii, Jingbo Shang, Chen-Yu Lee, Tomas Pfister	12811
What Does the Knowledge Neuron Thesis Have to Do with Knowledge? Jingcheng Niu, Andrew Liu, Zining Zhu, Gerald Penn	12835
Revisiting Link Prediction: A Data Perspective Haitao Mao, Juanhui Li, Harry Shomer, Bingheng Li, Wenqi Fan, Yao Ma, Tong Zhao, Neil Shah, Jiliang Tang	12861
Semantic Flow: Learning Semantic Fields of Dynamic Scenes from Monocular Videos Fengrui Tian, Yueqi Duan, Angtian Wang, Jianfei Guo, Shaoyi Du	12897
The Generative AI Paradox: “What it Can Create, it May Not Understand” Peter West, Ximing Lu, Nouha Dziri, Faeze Brahman, Linjie Li, Jena D. Hwang, Liwei Jiang, Jillian Fisher, Abhilasha Ravichander, Khyathi Chandu, Benjamin Newman, Pang Wei Koh, Allyson Ettinger, Yejin Choi	12915
On the Over-Memorization During Natural, Robust and Catastrophic Overfitting Runqi Lin, Chaojian Yu, Bo Han, Tongliang Liu	12944
Rethinking Adversarial Policies: A Generalized Attack Formulation and Provable Defense in RL Xiangyu Liu, Souradip Chakraborty, Yanchao Sun, Furong Huang	12962
Jointly Training Large Autoregressive Multimodal Models Emanuele Aiello, Lili Yu, Yixin Nie, Armen Aghajanyan, Barlas Oguz	12988
Threshold-Consistent Margin Loss for Open-World Deep Metric Learning Qin Zhang, Linghan Xu, Jun Fang, Qingming Tang, Ying Nian Wu, Joseph Tighe, Yifan Xing	13010
Stable Neural Stochastic Differential Equations in Analyzing Irregular Time Series Data Yongkyung Oh, Dongyoung Lim, Sungil Kim	13033
Stable Anisotropic Regularization William Rudman, Carsten Eickhoff	13065
Outliers with Opposing Signals Have an Outsized Effect on Neural Network Optimization Elan Rosenfeld, Andrej Risteski	13081

LUT-GEMM: Quantized Matrix Multiplication Based on LUTs for Efficient Inference in Large-Scale Generative Language Models Gunho Park, Baeseong Park, Minsub Kim, Sungjae Lee, Jeonghoon Kim, Beomseok Kwon, Se Jung Kwon, Byeongwook Kim, Youngjoo Lee, Dongsoo Lee	13128
Tree Cross Attention Leo Feng, Frederick Tung, Hossein Hajimirsadeghi, Yoshua Bengio, Mohamed Osama Ahmed	13146
Latent Trajectory Learning for Limited Timestamps Under Distribution Shift Over Time Qiu hao Zeng, Changjian Shui, Long-Kai Huang, Peng Liu, Xi Chen, Charles Ling, Boyu Wang	13167
Pathformer: Multi-Scale Transformers with Adaptive Pathways for Time Series Forecasting Peng Chen, Yingying Zhang, Yunyao Cheng, Yang Shu, Yihang Wang, Qingsong Wen, Bin Yang, Chenjuan Guo	13193
PTaRL: Prototype-Based Tabular Representation Learning Via Space Calibration Hangting Ye, Wei Fan, Xiaozhuang Song, Shun Zheng, He Zhao, Dan Dan Guo, Yi Chang	13212
Unified Human-Scene Interaction Via Prompted Chain-Of-Contacts Zeqi Xiao, Tai Wang, Jingbo Wang, Jinkun Cao, Wenwei Zhang, Bo Dai, Dahua Lin, Jiangmiao Pang	13231
The Importance of Feature Preprocessing for Differentially Private Linear Optimization Ziteng Sun, Ananda Theertha Suresh, Aditya Krishna Menon	13243
Towards Green AI in Fine-Tuning Large Language Models Via Adaptive Backpropagation Kai Huang, Hanyun Yin, Heng Huang, Wei Gao	13262
MiniLLM: Knowledge Distillation of Large Language Models Yuxian Gu, Li Dong, Furu Wei, Minlie Huang	13278
A Hard-To-Beat Baseline for Training-Free CLIP-Based Adaptation Zhengbo Wang, Jian Liang, Lijun Sheng, Ran He, Zilei Wang, Tieniu Tan	13302
LEMON: Lossless Model Expansion Yite Wang, Jiahao Su, Hanlin Lu, Cong Xie, Tianyi Liu, Jianbo Yuan, Haibin Lin, Ruoyu Sun, Hongxia Yang	13321
Poisoned Forgery Face: Towards Backdoor Attacks on Face Forgery Detection Jiawei Liang, Siyuan Liang, Aishan Liu, Xiaojun Jia, Junhao Kuang, Xiaochun Cao	13351
DIAGNOSIS: Detecting Unauthorized Data Usages in Text-To-Image Diffusion Models Zhenting Wang, Chen Chen, Lingjuan Lyu, Dimitris N. Metaxas, Shiqing Ma	13367
Adapting Large Language Models Via Reading Comprehension Daixuan Cheng, Shaohan Huang, Furu Wei	13388

UC-NERF: Neural Radiance Field for Under-Calibrated Multi-View Cameras in Autonomous Driving Kai Cheng, Xiaoxiao Long, Wei Yin, Jin Wang, Zhiqiang Wu, Yuexin Ma, Kaixuan Wang, Xiaozhi Chen, Xuejin Chen	13417
ImagenHub: Standardizing the Evaluation of Conditional Image Generation Models Max Ku, Tianle Li, Kai Zhang, Yujie Lu, Xingyu Fu, Wenwen Zhuang, Wenhui Chen	13439
Take a Step Back: Evoking Reasoning Via Abstraction in Large Language Models Huaixiu Steven Zheng, Swaroop Mishra, Xinyun Chen, Heng-Tze Cheng, Ed H. Chi, Quoc V Le, Denny Zhou	13473
FreeReg: Image-To-Point Cloud Registration Leveraging Pretrained Diffusion Models and Monocular Depth Estimators Haiping Wang, Yuan Liu, Bing Wang, Yujing Sun, Zhen Dong, Wenping Wang, Bisheng Yang	13511
FedHyper: A Universal and Robust Learning Rate Scheduler for Federated Learning with Hypergradient Descent Ziyao Wang, Jianyu Wang, Ang Li	13535
On the Analysis of GAN-Based Image-To-Image Translation with Gaussian Noise Injection Chaohua Shi, Kexin Huang, Lu Gan, Hongqing Liu, Mingrui Zhu, Nannan Wang, Xinbo Gao	13551
Towards Best Practices of Activation Patching in Language Models: Metrics and Methods Fred Zhang, Neel Nanda	13592
The Unlocking Spell on Base LLMs: Rethinking Alignment Via In-Context Learning Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, Yejin Choi	13620
Progressive3D: Progressively Local Editing for Text-To-3D Content Creation with Complex Semantic Prompts Xinhua Cheng, Tianyu Yang, Jianan Wang, Yu Li, Lei Zhang, Jian Zhang, Li Yuan	13647
LoTa-Bench: Benchmarking Language-Oriented Task Planners for Embodied Agents Jae-Woo Choi, Youngwoo Yoon, Hyobin Ong, Jaehong Kim, Minsu Jang	13668
Long-Term Typhoon Trajectory Prediction: A Physics-Conditioned Approach Without Reanalysis Data Young-Jae Park, Minseok Seo, Doyi Kim, Hyeri Kim, Sanghoon Choi, Beomkyu Choi, Jeongwon Ryu, Sohee Son, Hae-Gon Jeon, Yeji Choi	13695
PoSE: Efficient Context Window Extension of LLMs Via Positional Skip-Wise Training Dawei Zhu, Nan Yang, Liang Wang, Yifan Song, Wenhao Wu, Furu Wei, Sujian Li	13715
Integrating Planning and Deep Reinforcement Learning Via Automatic Induction of Task Substructures Jung-Chun Liu, Chi-Hsien Chang, Shao-Hua Sun, Tian-Li Yu	13730

Personalize Segment Anything Model with One Shot Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Yu Qiao, Peng Gao, Hongsheng Li	13754
Towards Reliable and Efficient Backdoor Trigger Inversion Via Decoupling Benign Features Xiong Xu, Kunzhe Huang, Yiming Li, Zhan Qin, Kui Ren	13784
TopoMLP: A Simple Yet Strong Pipeline for Driving Topology Reasoning Dongming Wu, Jiahao Chang, Fan Jia, Yingfei Liu, Tiancai Wang, Jianbing Shen	13809
FLD: Fourier Latent Dynamics for Structured Motion Representation and Learning Chenhao Li, Elijah Stanger-Jones, Steve Heim, Sang Bae Kim	13821
Simple Minimax Optimal Byzantine Robust Algorithm for Nonconvex Objectives with Uniform Gradient Heterogeneity Tomoya Murata, Kenta Niwa, Takumi Fukami, Ilfan Tyou	13848
Label-Focused Inductive Bias Over Latent Object Features in Visual Classification Ilmin Kang, Hyounyoung Bae, Kangil Kim	13885
Noise Map Guidance: Inversion with Spatial Context for Real Image Editing Hansam Cho, Jonghyun Lee, Seoung Bum Kim, Tae-Hyun Oh, Yonghyun Jeong	13903
Neural Fine-Tuning Search for Few-Shot Learning Panagiotis Eustratiadis, Łukasz Dudziak, Da Li, Timothy Hospedales	13923
IGraphMix: Input Graph Mixup Method for Node Classification Jongwon Jeong, Hyeop Lee, Hyui Geon Yoon, Beomyoung Lee, Junhee Heo, Geonsoo Kim, Kim Jin Seon	13940
Imitation Learning from Observation with Automatic Discount Scheduling Yuyang Liu, Weijun Dong, Yingdong Hu, Chuan Wen, Zhao-Heng Yin, Chongjie Zhang, Yang Gao	13972
Periodicity Decoupling Framework for Long-Term Series Forecasting Tao Dai, Beiliang Wu, Peiyuan Liu, Naiqi Li, Jigang Bao, Yong Jiang, Shu-Tao Xia	13990
LLM-Grounded Video Diffusion Models Long Lian, Baifeng Shi, Adam Yala, Trevor Darrell, Boyi Li	14002
MUSTARD: Mastering Uniform Synthesis of Theorem and Proof Data Yinya Huang, Xiaohan Lin, Zhengying Liu, Qingxing Cao, Huajian Xin, Haiming Wang, Zhenguo Li, Linqi Song, Xiaodan Liang	14023
Domain-Agnostic Molecular Generation with Chemical Feedback Yin Fang, Ningyu Zhang, Zhuo Chen, Lingbing Guo, Xiaohui Fan, Huajun Chen	14062
Is ImageNet Worth 1 Video? Learning Strong Image Encoders from 1 Long Unlabelled Video Shashanka Venkataramanan, Mamshad Nayeem Rizve, Joao Carreira, Yuki M Asano, Yannis Avrithis	14088

Protein Multimer Structure Prediction Via Prompt Learning Ziqi Gao, Xiangguo Sun, Zijiang Liu, Yu Li, Hong Cheng, Jia Li	14110
Prototypical Information Bottlenecking and Disentangling for Multimodal Cancer Survival Prediction Yilan Zhang, Yingxue Xu, Jianqi Chen, Fengying Xie, Hao Chen	14129
CoBIT: A Contrastive Bi-Directional Image-Text Generation Model Haoxuan You, Mandy Guo, Zhecan Wang, Kai-Wei Chang, Jason Michael Baldridge, Jiahui Yu	14152
Self-Supervised Heterogeneous Graph Learning: A Homophily and Heterogeneity View Yujie Mo, Feiping Nie, Ping Hu, Heng Tao Shen, Zheng Zhang, Xinchao Wang, Xiaofeng Zhu	14171
Think-On-Graph: Deep and Responsible Reasoning of Large Language Model on Knowledge Graph Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel Ni, Heung-Yeung Shum, Jian Guo	14199
Protein-Ligand Interaction Prior for Binding-Aware 3D Molecule Diffusion Models Zhilin Huang, Ling Yang, Xiangxin Zhou, Zhilong Zhang, Wentao Zhang, Xiawu Zheng, Jie Chen, Yu Wang, Bin Cui, Wenming Yang	14230
Demystifying Local & Global Fairness Trade-Offs in Federated Learning Using Partial Information Decomposition Faisal Hamman, Sanghamitra Dutta	14252
SparseFormer: Sparse Visual Recognition Via Limited Latent Tokens Ziteng Gao, Zhan Tong, Limin Wang, Mike Zheng Shou	14277
Let 2D Diffusion Model Know 3D-Consistency for Robust Text-To-3D Generation Junyoung Seo, Wooseok Jang, Min-Seop Kwak, Hyeonsu Kim, Jaehoon Ko, Junho Kim, Jin-Hwa Kim, Jiyoung Lee, Seungryong Kim	14296
LLaMA-Adapter: Efficient Fine-Tuning of Large Language Models with Zero-Initialized Attention Renrui Zhang, Jiaming Han, Chris Liu, Aojun Zhou, Pan Lu, Yu Qiao, Hongsheng Li, Peng Gao	14316
Input-Gradient Space Particle Inference for Neural Network Ensembles Trung Trinh, Markus Heinonen, Luigi Acerbi, Samuel Kaski	14346
AdjointDPM: Adjoint Sensitivity Method for Gradient Backpropagation of Diffusion Probabilistic Models Jiachun Pan, Jun Hao Liew, Vincent Tan, Jiashi Feng, Hanshu Yan	14365
M3C: A Framework Towards Convergent, Flexible, and Unsupervised Learning of Mixture Graph Matching and Clustering Jiaxin Lu, Zetian Jiang, Tianzhe Wang, Junchi Yan	14387

Going Beyond Neural Network Feature Similarity: The Network Feature Complexity and Its Interpretation Using Category Theory Yiting Chen, Zhanpeng Zhou, Junchi Yan	14415
Implicit Neural Representation Inference for Low-Dimensional Bayesian Deep Learning Panagiotis Dimitrakopoulos, Giorgos Sfikas, Christophoros Nikou	14433
Mean Field Theory in Deep Metric Learning Takuya Furusawa	14455
Composed Image Retrieval with Text Feedback Via Multi-Grained Uncertainty Regularization Yiyang Chen, Zhedong Zheng, Wei Ji, Leigang Qu, Tat-Seng Chua	14478
Diffusion Model for Dense Matching Jisu Nam, Gyuseong Lee, Sunwoo Kim, Hyeonsu Kim, Hyoungwon Cho, Seyeon Kim, Seungryong Kim	14491
MetaTool Benchmark for Large Language Models: Deciding Whether to Use Tools and Which to Use Yue Huang, Jiawen Shi, Yuan Li, Chenrui Fan, Siyuan Wu, Qihui Zhang, Yixin Liu, Pan Zhou, Yao Wan, Neil Zhenqiang Gong, Lichao Sun	14515
Self-Supervised High Dynamic Range Imaging with Multi-Exposure Images in Dynamic Scenes Zhilu Zhang, Haoyu Wang, Shuai Liu, Xiaotao Wang, Lei Lei, Wangmeng Zuo	14545
AdaMerging: Adaptive Model Merging for Multi-Task Learning Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, Dacheng Tao	14563
Multi-Resolution HuBERT: Multi-Resolution Speech Self-Supervised Learning with Masked Unit Prediction Jiatong Shi, Hirofumi Inaguma, Xutai Ma, Ilia Kulikov, Anna Sun	14584
New Insight of Variance Reduce in Zero-Order Hard-Thresholding: Mitigating Gradient Error and Expansivity Contradictions Xinzhe Yuan, William De Vazelhes, Bin Gu, Huan Xiong	14621
Human Motion Diffusion as a Generative Prior Yoni Shafir, Guy Tevet, Roy Kapon, Amit Haim Bermano	14632
In-Context Learning Learns Label Relationships but is Not Conventional Learning Jannik Kossen, Yarin Gal, Tom Rainforth	14649
OpenChat: Advancing Open-Source Language Models with Mixed-Quality Data Guan Wang, Sijie Cheng, Xianyuan Zhan, Xiangang Li, Sen Song, Yang Liu	14721
Transformer-Modulated Diffusion Models for Probabilistic Multivariate Time Series Forecasting Yuxin Li, Wenchao Chen, Xinyue Hu, Bo Chen, Baolin Sun, Mingyuan Zhou	14741

R-MAE: Regions Meet Masked Autoencoders Duy Kien Nguyen, Yanghao Li, Vaibhav Aggarwal, Martin R. Oswald, Alexander Kirillov, Cees G. M. Snoek, Xinlei Chen	14760
QA-LoRA: Quantization-Aware Low-Rank Adaptation of Large Language Models Yuhui Xu, Lingxi Xie, Xiaotao Gu, Xin Chen, Heng Chang, Hengheng Zhang, Zhengsu Chen, Xiaopeng Zhang, Qi Tian	14778
Image Inpainting Via Tractable Steering of Diffusion Models Anji Liu, Mathias Niepert, Guy Van Den Broeck	14796
Accelerated Sampling with Stacked Restricted Boltzmann Machines Jorge Fernandez-De-Cossio-Diaz, Clément Roussel, Simona Cocco, Remi Monasson	14818
STREAM: Spatio-TempoRal Evaluation and Analysis Metric for Video Generative Models Pum Jun Kim, Seojun Kim, Jaejun Yoo	14829
Duolando: Follower GPT with Off-Policy Reinforcement Learning for Dance Accompaniment Li Siyao, Tianpei Gu, Zhitao Yang, Zhengyu Lin, Ziwei Liu, Henghui Ding, Lei Yang, Chen Change Loy	14854
Optimal Transport Based Adversarial Patch to Leverage Large Scale Attack Transferability Pol Labarbarie, Adrien Chan-Hon-Tong, Stéphanie Herbin, Milad Leyli-Abadi	14874
Equivariant Matrix Function Neural Networks Ilyes Batatia, Lars Leon Schaaf, Gabor Csanyi, Christoph Ortner, Felix Andreas Faber	14902
LiDAR-PTQ: Post-Training Quantization for Point Cloud 3D Object Detection Sifan Zhou, Liang Li, Xinyu Zhang, Bo Zhang, Shipeng Bai, Miao Sun, Ziyu Zhao, Xiaobo Lu, Xiangxiang Chu	14922
ACRF: Compressing Explicit Neural Radiance Fields Via Attribute Compression Guangchi Fang, Qingyong Hu, Longguang Wang, Yulan Guo	14937
Retro-Fallback: Retrosynthetic Planning in an Uncertain World Austin Tripp, Krzysztof Maziarz, Sarah Lewis, Marwin Segler, Jos�� Miguel Hern��ndez-Lobato	14948
State Representation Learning Using an Unbalanced Atlas Li Meng, Morten Goodwin, Anis Yazidi, Paal E. Engelstad	15006
MovingParts: Motion-Based 3D Part Discovery in Dynamic Radiance Field Kaizhi Yang, Xiaoshuai Zhang, Zhiao Huang, Xuejin Chen, Zexiang Xu, Hao Su	15021
Convergence of Bayesian Bilevel Optimization Shi Fu, Fengxiang He, Xinmei Tian, Dacheng Tao	15038
EControl: Fast Distributed Optimization with Compression and Error Control Yuan Gao, Rustem Islamov, Sebastian U Stich	15064

Learning Nash Equilibria in Rank-1 Games Nikolas Patris, Ioannis Panageas	15096
Learning Invariant Representations of Time-Homogeneous Stochastic Dynamical Systems Vladimir R Kostic, Pietro Novelli, Riccardo Grazzi, Karim Lounici, Massimiliano Pontil	15118
BESA: Pruning Large Language Models with Blockwise Parameter-Efficient Sparsity Allocation Peng Xu, Wenqi Shao, Mengzhao Chen, Shitao Tang, Kaipeng Zhang, Peng Gao, Fengwei An, Yu Qiao, Ping Luo	15131
On the Humanity of Conversational AI: Evaluating the Psychological Portrayal of LLMs Jen-Tse Huang, Wenxuan Wang, Eric John Li, Man Ho Lam, Shujie Ren, Youliang Yuan, Wenxiang Jiao, Zhaopeng Tu, Michael Lyu	15146
Fast and Unified Path Gradient Estimators for Normalizing Flows Lorenz Vaitl, Ludwig Winkler, Lorenz Richter, Pan Kessel	15170
Amortized Network Intervention to Steer the Excitatory Point Processes Zitao Song, Wendi Ren, Shuang Li	15200
ODEFormer: Symbolic Regression of Dynamical Systems with Transformers Stéphane D'Ascoli, Sören Becker, Philippe Schwallier, Alexander Mathis, Niki Kilbertus	15224
Leveraging augmented-Lagrangian Techniques for Differentiating Over Infeasible Quadratic Programs in Machine Learning Antoine Bambade, Fabian Schramm, Adrien Taylor, Justin Carpentier	15258
DiffEnc: Variational Diffusion with a Learned Encoder Beatrix Miranda Ginn Nielsen, Anders Christensen, Andrea Dittadi, Ole Winther	15290
Decongestion by Representation: Learning to Improve Economic Welfare in Marketplaces Omer Nahum, Gali Noti, David C. Parkes, Nir Rosenfeld	15323
InstructPix2NeRF: Instructed 3D Portrait Editing from a Single Image Jianhui Li, Shilong Liu, Zidong Liu, Yikai Wang, Kaiwen Zheng, Jinghui Xu, Jianmin Li, Jun Zhu	15348
Unveiling the Pitfalls of Knowledge Editing for Large Language Models Zhoubo Li, Ningyu Zhang, Yunzhi Yao, Mengru Wang, Xi Chen, Huajun Chen	15376
LCOT: Linear Circular Optimal Transport Rocio P Diaz Martin, Ivan Vladimir Medri, Yikun Bai, Xinran Liu, Kangbai Yan, Gustavo Rohde, Soheil Kolouri	15393
Bespoke Solvers for Generative Flow Models Neta Shaul, Juan Perez, Ricky T. Q. Chen, Ali Thabet, Albert Pumarola, Yaron Lipman	15417
Image2Sentence Based Asymmetrical Zero-Shot Composed Image Retrieval Yongchao Du, Min Wang, Wengang Zhou, Shuping Hui, Houqiang Li	15454

Decoding Natural Images from EEG for Object Recognition Yonghao Song, Bingchuan Liu, Xiang Li, Nanlin Shi, Yijun Wang, Xiaorong Gao	15475
Image Inpainting Via Iteratively Decoupled Probabilistic Modeling Wenbo Li, Xin Yu, Kun Zhou, Yibing Song, Zhe Lin	15493
Align with Purpose: Optimize Desired Properties in CTC Models with a General Plug-And-Play Framework Eliya Segev, Maya Alroy, Ronen Katsir, Noam Wies, Ayana Shenhav, Yael Sapir Ben-Oren, David Zar, Oren Tadmor, Jacob Bitterman, Amnon Shashua, Tal Rosenwein	15519
Towards Enhancing Time Series Contrastive Learning: A Dynamic Bad Pair Mining Approach Xiang Lan, Hanshu Yan, Shenda Hong, Mengling Feng	15535
Robust Agents Learn Causal World Models Jonathan Richens, Tom Everitt	15556
CADS: Unleashing the Diversity of Diffusion Models Through Condition-Annealed Sampling Seyedmorteza Sadat, Jakob Buhmann, Derek Bradley, Otmar Hilliges, Romann M. Weber	15588
Fantastic Generalization Measures Are Nowhere to Be Found Michael Gastpar, Ido Nachum, Jonathan Shafer, Thomas Weinberger	15621
What's in a Prior? Learned Proximal Networks for Inverse Problems Zhenghan Fang, Sam Buchanan, Jeremias Sulam	15652
Theoretical Analysis of Robust Overfitting for Wide DNNs: An NTK Approach Shaopeng Fu, Di Wang	15693
Cross-Modal Contextualized Diffusion Models for Text-Guided Visual Generation and Editing Ling Yang, Zhilong Zhang, Zhaochen Yu, Jingwei Liu, Minkai Xu, Stefano Ermon, Bin Cui	15739
CO2: Efficient Distributed Training with Full Communication-Computation Overlap Weigao Sun, Zhen Qin, Weixuan Sun, Shidi Li, Dong Li, Xuyang Shen, Yu Qiao, Yiran Zhong	15765
Towards Poisoning Fair Representations Tianci Liu, Haoyu Wang, Feijie Wu, Hengtong Zhang, Pan Li, Lu Su, Jing Gao	15788
Local Composite Saddle Point Optimization Site Bai, Brian Bullins	15815
Multi-Task Reinforcement Learning with Mixture of Orthogonal Experts Ahmed Hendawy, Jan Peters, Carlo D'Eramo	15857
Neural Rate Control for Learned Video Compression Yiwei Zhang, Guo Lu, Yunuo Chen, Shen Wang, Yibo Shi, Jing Wang, Li Song	15879
Sliced Wasserstein Estimation with Control Variates Khai Nguyen, Nhat Ho	15894

Pre-Training and Fine-Tuning Generative Flow Networks <i>Ling Pan, Moksh Jain, Kanika Madan, Yoshua Bengio</i>	15914
Label-Free Node Classification on Graphs with Large Language Models (LLMs) <i>Zhikai Chen, Haitao Mao, Hongzhi Wen, Haoyu Han, Wei Jin, Haiyang Zhang, Hui Liu, Jiliang Tang</i>	15935
Compositional Preference Models for Aligning LMs <i>Dongyoung Go, Tomasz Korbak, Germăjn Kruszewski, Jos Rozen, Marc Dymetman</i>	15959
ARGS: Alignment as Reward-Guided Search <i>Maxim Khanov, Jirayu Burapachee, Yixuan Li</i>	15982
Generative Sliced MMD Flows with Riesz Kernels <i>Johannes Hertrich, Christian Wald, Fabian Altekraßer, Paul Hagemann</i>	16001
DMBP: Diffusion Model-Based Predictor for Robust Offline Reinforcement Learning Against State Observation Perturbations <i>Zhihe Yang, Yunjian Xu</i>	16028
OWL: A Large Language Model for it Operations <i>Hongcheng Guo, Jian Yang, Jiaheng Liu, Liqun Yang, Linzheng Chai, Jiaqi Bai, Junran Peng, Xiaorong Hu, Chao Chen, Dongfeng Zhang, Xu Shi, Tieqiao Zheng, Liangfan Zheng, Bo Zhang, Ke Xu, Zhoujun Li</i>	16058
Efficient Streaming Language Models with Attention Sinks <i>Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, Mike Lewis</i>	16082
Multi-Task Learning with 3D-Aware Regularization <i>Wei-Hong Li, Steven McDonagh, Ales Leonardis, Hakan Bilen</i>	16103
Data Filtering Networks <i>Alex Fang, Albin Madappally Jose, Amit Jain, Ludwig Schmidt, Alexander T Toshev, Vaishaal Shankar</i>	16119
A Branching Decoder for Set Generation <i>Zixian Huang, Gengyang Xiao, Yu Gu, Gong Cheng</i>	16136
V-DETR: DETR with Vertex Relative Position Encoding for 3D Object Detection <i>Yichao Shen, Zigang Geng, Yuhui Yuan, Yutong Lin, Ze Liu, Chunyu Wang, Han Hu, Nanning Zheng, Baining Guo</i>	16155
NuwaDynamics: Discovering and Updating in Causal Spatio-Temporal Modeling <i>Kun Wang, Hao Wu, Yifan Duan, Guibin Zhang, Kai Wang, Xiaojiang Peng, Yu Zheng, Yuxuan Liang, Yang Wang</i>	16174
Scaling Laws of RoPE-Based Extrapolation <i>Xiaoran Liu, Hang Yan, Chenxin An, Xipeng Qiu, Dahua Lin</i>	16199
Towards Image Compression with Perfect Realism at Ultra-Low Bitrates <i>Marlene Careil, Matthew J. Muckley, Jakob Verbeek, Stéphanie Lathuilière</i>	16225

Network Memory Footprint Compression Through Jointly Learnable Codebooks and Mappings Edouard Yvinec, Arnaud Dapogny, Kevin Bailly	16246
Towards Understanding Sycophancy in Language Models Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askeel, Samuel R. Bowman, Esin Durmus, Zac Hatfield-Dodds, Scott R Johnston, Shauna M Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, Ethan Perez	16258
SpQR: A Sparse-Quantized Representation for Near-Lossless LLM Weight Compression Tim Dettmers, Ruslan A. Svirschevski, Vage Egiazarian, Denis Kuznedelev, Elias Frantar, Saleh Ashkboos, Alexander Borzunov, Torsten Hoefer, Dan Alistarh	16293
Light-MILPopt: Solving Large-Scale Mixed Integer Linear Programs with Lightweight Optimizer and Small-Scale Training Dataset Huigen Ye, Hua Xu, Hongyan Wang	16322
Candidate Label Set Pruning: A Data-Centric Perspective for Deep Partial-Label Learning Shuo He, Chaojie Wang, Guowu Yang, Lei Feng	16346
RECOMBINER: Robust and Enhanced Compression with Bayesian Implicit Neural Representations Jiajun He, Gergely Flamich, Zongyu Guo, Jos�� Miguel Hern��ndez-Lobato	16364
Task Adaptation from Skills: Information Geometry, Disentanglement, and New Objectives for Unsupervised Reinforcement Learning Yucheng Yang, Tianyi Zhou, Qiang He, Lei Han, Mykola Pechenizkiy, Meng Fang	16391
Backdoor Federated Learning by Poisoning Backdoor-Critical Layers Haomin Zhuang, Mingxian Yu, Hao Wang, Yang Hua, Jian Li, Xu Yuan	16427
Horizon-Free Reinforcement Learning in Adversarial Linear Mixture MDPs Kaixuan Ji, Qingyue Zhao, Jiafan He, Weitong Zhang, Quanquan Gu	16453
Learning in Reverse Causal Strategic Environments with Ramifications on Two Sided Markets Seamus Somerstep, Yuekai Sun, Yaacov Ritov	16481
Learning Interactive Real-World Simulators Sherry Yang, Yilun Du, Seyed Kamyar Seyed Ghasemipour, Jonathan Tompson, Leslie Pack Kaelbling, Dale Schuurmans, Pieter Abbeel	16504
Stabilizing Backpropagation Through Time to Learn Complex Physics Patrick Schnell, Nils Thuerey	16529
Fast Imitation Via Behavior Foundation Models Matteo Pirota, Andrea Tirinzoni, Ahmed Touati, Alessandro Lazaric, Yann Ollivier	16554
Repelling Random Walks Isaac Reid, Eli Berger, Krzysztof Marcin Choromanski, Adrian Weller	16594

HyperAttention: Long-Context Attention in Near-Linear Time <i>Insu Han, Rajesh Jayaram, Amin Karbasi, Vahab Mirrokni, David Woodruff, Amir Zandieh</i>	16616
General Graph Random Features <i>Isaac Reid, Krzysztof Marcin Choromanski, Eli Berger, Adrian Weller</i>	16632
AlpaGasus: Training a Better Alpaca with Fewer Data <i>Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srinivasan, Tianyi Zhou, Heng Huang, Hongxia Jin</i>	16650
Intriguing Properties of Generative Classifiers <i>Priyank Jaini, Kevin Clark, Robert Geirhos</i>	16681
Democratizing Fine-Grained Visual Recognition with Large Language Models <i>Mingxuan Liu, Subhankar Roy, Wenjing Li, Zhun Zhong, Nicu Sebe, Elisa Ricci</i>	16707
Decomposed Diffusion Sampler for Accelerating Large-Scale Inverse Problems <i>Hyungjin Chung, Suhyeon Lee, Jong Chul Ye</i>	16739
Language Model Inversion <i>John Xavier Morris, Wenting Zhao, Justin T Chiu, Vitaly Shmatikov, Alexander M Rush</i>	16767
Perceptual Scales Predicted by Fisher Information Metrics <i>Jonathan Vacher, Pascal Mamassian</i>	16788
The LLM Surgeon <i>Tycho F. A. Van Der Ouderaa, Markus Nagel, Mart Van Baalen, Tijmen Blankevoort</i>	16809
MEND: Meta Demonstration Distillation for Efficient and Effective In-Context Learning <i>Yichuan Li, Xiyao Ma, Sixing Lu, Kyumin Lee, Xiaohu Liu, Chenlei Guo</i>	16820
Self-Alignment with Instruction Backtranslation <i>Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Omer Levy, Luke Zettlemoyer, Jason E Weston, Mike Lewis</i>	16835
Monte Carlo Guided Denoising Diffusion Models for Bayesian Linear Inverse Problems. <i>Gabriel Cardoso, Yazid Janati El Idrissi, Sylvain Le Corff, Eric Moulines</i>	16861
A Recipe for Improved Certifiable Robustness <i>Kai Hu, Klas Leino, Zifan Wang, Matt Fredrikson</i>	16898
Masked Completion Via Structured Diffusion with White-Box Transformers <i>Druv Pai, Sam Buchanan, Ziyang Wu, Yaodong Yu, Yi Ma</i>	16910
Let's Verify Step by Step <i>Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, Karl Cobbe</i>	16954

Habitat 3.0: A Co-Habitat for Humans, Avatars, and Robots <i>Xavier Puig, Eric Undersander, Andrew Szot, Mikael Dallaire Cote, Tsung-Yen Yang, Ruslan Partsey, Ruta Desai, Alexander Clegg, Michal Hlavac, So Yeon Min, Vladim�r Vondru�j, Theophile Gervet, Vincent-Pierre Berges, John M Turner, Oleksandr Maksymets, Zsolt Kira, Mrinal Kalakrishnan, Jitendra Malik, Devendra Singh Chaplot, Unnat Jain, Dhruv Batra, Akshara Rai, Roozbeh Mottaghi</i>	16978
Parameter-Efficient Multi-Task Model Fusion with Partial Linearization <i>Anke Tang, Li Shen, Yong Luo, Yibing Zhan, Han Hu, Bo Du, Yixin Chen, Dacheng Tao</i>	17009
Compressing Latent Space Via Least Volume <i>Qiuyi Chen, Mark Fuge</i>	17040
On the Stability of Iterative Retraining of Generative Models on Their Own Data <i>Quentin Bertrand, Joey Bose, Alexandre Duplessis, Marco Jiralerspong, Gauthier Gidel</i>	17064
Efficient Local Linearity Regularization to Overcome Catastrophic Overfitting <i>Elias Abad Rocamora, Fanghui Liu, Grigorios Chrysos, Pablo M. Olmos, Volkan Cevher</i>	17093
Neur2RO: Neural Two-Stage Robust Optimization <i>Justin Dumouchelle, Esther Julien, Jannis Kurtz, Elias Boutros Khalil</i>	17123
Mastering Memory Tasks with World Models <i>Mohammad Reza Samsami, Artem Zholus, Janarthanan Rajendran, Sarath Chandar</i>	17148
Manifold Diffusion Fields <i>Ahmed A. A. Elhag, Yuyang Wang, Joshua M. Susskind, Miguel ��ngel Bautista</i>	17185
Learning to Solve Bilevel Programs with Binary Tender <i>Bo Zhou, Ruiwei Jiang, Siqian Shen</i>	17210
Safety-Tuned LLaMAs: Lessons from Improving the Safety of Large Language Models that Follow Instructions <i>Federico Bianchi, Mirac Suzgun, Giuseppe Attanasio, Paul Rottger, Dan Jurafsky, Tatsunori Hashimoto, James Zou</i>	17233
Feature Emergence Via Margin Maximization: Case Studies in Algebraic Tasks <i>Depen Morwani, Benjamin L. Edelman, Costin-Andrei Oncescu, Rosie Zhao, Sham M. Kakade</i>	17254
Masked Audio Generation Using a Single Non-Autoregressive Transformer <i>Alon Ziv, Itai Gat, Gael Le Lan, Tal Remez, Felix Kreuk, Jade Copet, Alexandre D��fossez, Gabriel Synnaeve, Yossi Adi</i>	17292
MINT: Evaluating LLMs in Multi-Turn Interaction with Tools and Language Feedback <i>Xingyao Wang, Zihan Wang, Jiateng Liu, Yangyi Chen, Lifan Yuan, Hao Peng, Heng Ji</i>	17312
GenCorres: Consistent Shape Matching Via Coupled Implicit-Explicit Shape Generative Models <i>Haitao Yang, Xiangru Huang, Bo Sun, Chandrajit L. Bajaj, Qixing Huang</i>	17347

The Joint Effect of Task Similarity and Overparameterization on Catastrophic Forgetting – an Analytical Model Daniel Goldfarb, Itay Evron, Nir Weinberger, Daniel Soudry, Paul Hand	17367
Universal Backdoor Attacks Benjamin Schneider, Nils Lukas, Florian Kerschbaum	17453
Copilot4D: Learning Unsupervised World Models for Autonomous Driving Via Discrete Diffusion Lunjun Zhang, Yuwen Xiong, Ze Yang, Sergio Casas, Rui Hu, Raquel Urtasun	17467
Large Language Models Cannot Self-Correct Reasoning Yet Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, Denny Zhou	17498
ZeRO++: Extremely Efficient Collective Communication for Large Model Training Guanhua Wang, Heyang Qin, Sam Ade Jacobs, Xiaoxia Wu, Connor Holmes, Zhewei Yao, Samyam Rajbhandari, Olatunji Ruwase, Feng Yan, Lei Yang, Yuxiong He	17515
Leveraging Optimization for Adaptive Attacks on Image Watermarks Nils Lukas, Abdulrahman Diaa, Lucas Fenaux, Florian Kerschbaum	17534
Towards a Statistical Theory of Data Selection Under Weak Supervision Germain Kolossov, Andrea Montanari, Pulkit Tandon	17549
NAISR: A 3D Neural Additive Model for Interpretable Shape Representation Yining Jiao, Carlton Jude Zdanski, Julia S Kimbell, Andrew Prince, Cameron P Worden, Samuel Kirse, Christopher Rutter, Benjamin Shields, William Alexander Dunn, Jisan Mahmud, Marc Niethammer	17588
Headless Language Models: Learning Without Predicting with Contrastive Weight Tying Nathan Godey, Éric Villemonte De La Clergerie, Benoît Sagot	17621
Robust Model Based Reinforcement Learning Using $\mathcal{L}_1$ Adaptive Control Minjun Sung, Sambhu Harimanas Karumanchi, Aditya Gahlawat, Naira Hovakimyan	17642
MVDream: Multi-View Diffusion for 3D Generation Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, Xiao Yang	17664
Linear Attention is (maybe) All You Need (to Understand Transformer Optimization) Kwangjun Ahn, Xiang Cheng, Minhak Song, Chulhee Yun, Ali Jadbabaie, Suvrit Sra	17686
Classification with Conceptual Safeguards Hailey Joren, Charles Thomas Marx, Berk Ustun	17699
Machine Unlearning for Image-To-Image Generative Models Guihong Li, Hsiang Hsu, Chun-Fu Chen, Radu Marculescu	17710
OMNI: Open-Endedness Via Models of Human Notions of Interestingness Jenny Zhang, Joel Lehman, Kenneth Stanley, Jeff Clune	17745

Domain Randomization Via Entropy Maximization <i>Gabriele Tiboni, Pascal Klink, Jan Peters, Tatiana Tommasi, Carlo D'Eramo, Georgia Chalvatzaki</i>	17792
Effective Structural Encodings Via Local Curvature Profiles <i>Lukas Fesser, Melanie Weber</i>	17815
Branch-GAN: Improving Text Generation with (not So) Large Language Models <i>Fredrik Carlsson, Johan Broberg, Erik Hillbom, Magnus Sahlgren, Joakim Nivre</i>	17834
One for All: Towards Training One Graph Model for All Classification Tasks <i>Hao Liu, Jiarui Feng, Lecheng Kong, Ningyue Liang, Dacheng Tao, Yixin Chen, Muhan Zhang</i>	17862
An Image is Worth 1000 Lies: Transferability of Adversarial Images Across Prompts on Vision-Language Models <i>Haochen Luo, Jindong Gu, Fengyuan Liu, Philip Torr</i>	17885
Copula Conformal Prediction for Multi-Step Time Series Prediction <i>Sophia Huiwen Sun, Rose Yu</i>	17907
Dropout-Based Rashomon Set Exploration for Efficient Predictive Multiplicity Estimation <i>Hsiang Hsu, Guihong Li, Shaohan Hu, Chun-Fu Chen</i>	17929
Combining Axes Preconditioners Through Kronecker Approximation for Deep Learning <i>Sai Surya Duvvuri, Fnu Devvrit, Rohan Anil, Cho-Jui Hsieh, Inderjit S Dhillon</i>	17970
On the Hardness of Learning Under Symmetries <i>Bobak Kiani, Thien Le, Hannah Lawrence, Stefanie Jegelka, Melanie Weber</i>	17992
ARM: Refining Multivariate Forecasting with Adaptive Temporal-Contextual Learning <i>Jiecheng Lu, Xu Han, Shihao Yang</i>	18040
Diffusion-TS: Interpretable Diffusion for General Time Series Generation <i>Xinyu Yuan, Yan Qiao</i>	18062
Complex Priors and Flexible Inference in Recurrent Circuits with Dendritic Nonlinearities <i>Benjamin S. H. Lyo, Cristina Savin</i>	18091
Cascading Reinforcement Learning <i>Yihan Du, R. Srikant, Wei Chen</i>	18111
Plan-Seq-Learn: Language Model Guided RL for Solving Long Horizon Robotics Tasks <i>Murtaza Dalal, Tarun Chiruvolu, Devendra Singh Chaplot, Ruslan Salakhutdinov</i>	18153
Adversarial Training Should Be Cast as a Non-Zero-Sum Game <i>Alexander Robey, Fabian Latorre, George J. Pappas, Hamed Hassani, Volkan Cevher</i>	18182
A Poincaré Inequality and Consistency Results for Signal Sampling on Large Graphs <i>Thien Le, Luana Ruiz, Stefanie Jegelka</i>	18203

A Graph is Worth 1-Bit Spikes: When Graph Contrastive Learning Meets Spiking Neural Networks Jintang Li, Huizhe Zhang, Ruofan Wu, Zulun Zhu, Baokun Wang, Changhua Meng, Zibin Zheng, Liang Chen	18225
Quasi-Monte Carlo for 3D Sliced Wasserstein Khai Nguyen, Nicola Barileto, Nhat Ho	18249
ConR: Contrastive Regularizer for Deep Imbalanced Regression Mahsa Keramati, Lili Meng, R. David Evans	18275
Learning Performance-Improving Code Edits Alexander G Shypula, Aman Madaan, Yimeng Zeng, Uri Alon, Jacob R. Gardner, Yiming Yang, Milad Hashemi, Graham Neubig, Parthasarathy Ranganathan, Osbert Bastani, Amir Yazdanbakhsh	18299
Overcoming the Pitfalls of Vision-Language Model Finetuning for OOD Generalization Yuhang Zang, Hanlin Goh, Joshua M. Susskind, Chen Huang	18324
Robust Similarity Learning with Difference Alignment Regularization Shuo Chen, Gang Niu, Chen Gong, Okan Koc, Jian Yang, Masashi Sugiyama	18340
Decoupling Regularization from the Action Space Sobhan Mohammadpour, Emma Frejinger, Pierre-Luc Bacon	18362
Fast Ensembling with Diffusion Schrödinger Bridge Hyunsu Kim, Jongmin Yoon, Juho Lee	18378
Crystalformer: Infinitely Connected Attention for Periodic Structure Encoding Tatsunori Tanaii, Ryo Igarashi, Yuta Suzuki, Naoya Chiba, Kotaro Saito, Yoshitaka Ushiku, Kanta Ono	18395
The Unreasonable Effectiveness of Linear Prediction as a Perceptual Metric Daniel Severo, Lucas Theis, Johannes Ballé	18418
SPDER: Semiperiodic Damping-Enabled Object Representation Kathan Shah, Chawin Sitawarin	18429
NeurRev: Train Better Sparse Neural Network Practically Via Neuron Revitalization Gen Li, Lu Yin, Jie Ji, Wei Niu, Minghai Qin, Bin Ren, Linke Guo, Shiwei Liu, Xiaolong Ma	18460
Learning to Jointly Understand Visual and Tactile Signals Yichen Li, Yilun Du, Chao Liu, Chao Liu, Francis Williams, Michael Foshey, Benjamin Eckart, Jan Kautz, Joshua B. Tenenbaum, Antonio Torralba, Wojciech Matusik	18476
TEDDY: Trimming Edges with Degree-Based Discrimination Strategy Hyunjin Seo, Jihun Yun, Eunho Yang	18493
PRES: Toward Scalable Memory-Based Dynamic Graph Neural Networks Junwei Su, Difan Zou, Chuan Wu	18516

DittoGym: Learning to Control Soft Shape-Shifting Robots Suning Huang, Boyuan Chen, Huazhe Xu, Vincent Sitzmann	18545
SocioDojo: Building Lifelong Analytical Agents with Real-World Text and Time Series Junyan Cheng, Peter Chin	18564
AirPhyNet: Harnessing Physics-Guided Neural Networks for Air Quality Prediction Kethmi Hirushini Hettige, Jiahao Ji, Shili Xiang, Cheng Long, Gao Cong, Jingyuan Wang	18599
Bridging Neural and Symbolic Representations with Transitional Dictionary Learning Junyan Cheng, Peter Chin	18616
Mathematical Justification of Hard Negative Mining Via Isometric Approximation Theorem Albert Xu, Jhih-Yi Hsieh, Bhaskar Vundurthy, Nithya Kemp, Eliana Cohen, Lu Li, Howie Choset	18645
Frozen Transformers in Language Models Are Effective Visual Encoder Layers Ziqi Pang, Ziyang Xie, Yunze Man, Yu-Xiong Wang	18663
You Only Query Once: An Efficient Label-Only Membership Inference Attack Yutong Wu, Han Qiu, Shangwei Guo, Jiwei Li, Tianwei Zhang	18686
Towards Category Unification of 3D Single Object Tracking on Point Clouds Jiahao Nie, Zhiwei He, Xudong Lv, Xueyi Zhou, Dong-Kyu Chae, Fei Xie	18700
The Expressive Power of Low-Rank Adaptation Yuchen Zeng, Kangwook Lee	18718
Demystifying Embedding Spaces Using Large Language Models Guy Tennenholtz, Yinlam Chow, Chihwei Hsu, Jihwan Jeong, Lior Shani, Azamat Tulepbergenov, Deepak Ramachandran, Martin Mladenov, Craig Boutilier	18764
BTR: Binary Token Representations for Efficient Retrieval Augmented Language Models Qingqing Cao, Sewon Min, Yizhong Wang, Hannaneh Hajishirzi	18798
Weaker MVI Condition: Extragradient Methods with Multi-Step Exploration Yifeng Fan, Yongqiang Li, Bo Chen	18816
EfficientDM: Efficient Quantization-Aware Fine-Tuning of Low-Bit Diffusion Models Yefei He, Jing Liu, Weijia Wu, Hong Zhou, Bohan Zhuang	18842
Continuous Field Reconstruction from Sparse Observations with Implicit Neural Networks Xihaier Luo, Wei Xu, Balu Nadiga, Yihui Ren, Shinjae Yoo	18862
ASID: Active Exploration for System Identification in Robotic Manipulation Marius Memmel, Andrew Wagenmaker, Chuning Zhu, Dieter Fox, Abhishek Gupta	18887
DV-3DLane: End-To-End Multi-Modal 3D Lane Detection with Dual-View Representation Yueru Luo, Shuguang Cui, Zhen Li	18906

Where We Have Arrived in Proving the Emergence of Sparse Interaction Primitives in DNNs Qihan Ren, Jiayang Gao, Wen Shen, Quanshi Zhang	18920
Polynormer: Polynomial-Expressive Graph Transformer in Linear Time Chenhui Deng, Zichao Yue, Zhiru Zhang	18954
SetCSE: Set Operations Using Contrastive Learning of Sentence Embeddings Kang Liu	18980
Heterogeneous Personalized Federated Learning by Local-Global Updates Mixing Via Convergence Rate Meirui Jiang, Anjie Le, Xiaoxiao Li, Qi Dou	19007
MG-TSD: Multi-Granularity Time Series Diffusion Models with Guided Learning Process Xinyao Fan, Yueying Wu, Chang Xu, Yuhao Huang, Weiqing Liu, Jiang Bian	19036
A Unified and General Framework for Continual Learning Zhenyi Wang, Yan Li, Li Shen, Heng Huang	19055
LMSYS-Chat-1M: A Large-Scale Real-World LLM Conversation Dataset Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zhuohan Li, Zilong Jin, Eric Xing, Joseph E. Gonzalez, Ion Stoica, Hao Zhang	19073
SRL: Scaling Distributed Reinforcement Learning to Over Ten Thousand Cores Zhiyu Mei, Wei Fu, Jiaxuan Gao, Guangju Wang, Huanchen Zhang, Yi Wu	19106
PS²OT: Progressive Partial Optimal Transport for Deep Imbalanced Clustering Chuyu Zhang, Hui Ren, Xuming He	19130
Rayleigh Quotient Graph Neural Networks for Graph-Level Anomaly Detection Xiangyu Dong, Xingyi Zhang, Sibor Wang	19152
Sliced Denoising: A Physics-Informed Molecular Pre-Training Method Yuyan Ni, Shikun Feng, Wei-Ying Ma, Zhi-Ming Ma, Yanyan Lan	19171
Efficiently Computing Similarities to Private Datasets Arturs Backurs, Zinan Lin, Sepideh Mahabadi, Sandeep Silwal, Jakub Tarnawski	19196
Faithful Vision-Language Interpretation Via Concept Bottleneck Models Songning Lai, Lijie Hu, Junxiao Wang, Laure Berti-Equille, Di Wang	19225
Adversarial Attacks on Fairness of Graph Neural Networks Binchi Zhang, Yushun Dong, Chen Chen, Yada Zhu, Minnan Luo, Jundong Li	19249
SLiMe: Segment Like Me Aliasghar Khani, Saeid Asgari, Aditya Sanghi, Ali Mahdavi Amiri, Ghassan Hamarneh	19283
Protein-Ligand Binding Representation Learning from Fine-Grained Interactions Shikun Feng, Minghao Li, Yinjun Jia, Wei-Ying Ma, Yanyan Lan	19302

CausalTime: Realistically Generated Time-Series for Benchmarking of Causal Discovery Yuxiao Cheng, Ziqian Wang, Tingxiong Xiao, Qin Zhong, Jinli Suo, Kunlun He	19318
EBMDock: Neural Probabilistic Protein-Protein Docking Via a Differentiable Energy Model Huaijin Wu, Wei Liu, Yatao Bian, Jiaxiang Wu, Nianzu Yang, Junchi Yan	19341
Generative Pre-Training for Speech with Flow Matching Alexander H. Liu, Matthew Le, Apoorv Vyas, Bowen Shi, Andros Tjandra, Wei-Ning Hsu	19361
Functional Bayesian Tucker Decomposition for Continuous-Indexed Tensor Data Shikai Fang, Xin Yu, Zheng Wang, Shibo Li, Mike Kirby, Shandian Zhe	19381
FLASK: Fine-Grained Language Model Evaluation Based on Alignment Skill Sets Seonghyeon Ye, Doyoung Kim, Sungdong Kim, Hyeonbin Hwang, Seungone Kim, Yongrae Jo, James Thorne, Juho Kim, Minjoon Seo	19399
Interpretable Sparse System Identification: Beyond Recent Deep Learning Techniques on Time-Series Prediction Xiaoyi Liu, Duxin Chen, Wenjia Wei, Xia Zhu, Wenwu Yu	19453
Follow-Up Differential Descriptions: Language Models Resolve Ambiguities for Image Classification Reza Esfandiarpour, Stephen Bach	19467
Pseudo-Generalized Dynamic View Synthesis from a Video Xiaoming Zhao, R Alex Colburn, Fangchang Ma, Miguel Ángel Bautista, Joshua M. Susskind, Alex Schwing	19492
DreamCraft3D: Hierarchical 3D Generation with Bootstrapped Diffusion Prior Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, Yebin Liu	19515
DiffusionSat: A Generative Foundation Model for Satellite Imagery Samar Khanna, Patrick Liu, Linqi Zhou, Chenlin Meng, Robin Rombach, Marshall Burke, David B. Lobell, Stefano Ermon	19539
Dual Associated Encoder for Face Restoration Yu-Ju Tsai, Yu-Lun Liu, Lu Qi, Kelvin C. K. Chan, Ming-Hsuan Yang	19558
Solving Challenging Math Word Problems Using GPT-4 Code Interpreter with Code-Based Self-Verification Aojun Zhou, Ke Wang, Zimu Lu, Weikang Shi, Sichun Luo, Zipeng Qin, Shaoqing Lu, Anya Jia, Linqi Song, Mingjie Zhan, Hongsheng Li	19582
MathCoder: Seamless Code Integration in LLMs for Enhanced Mathematical Reasoning Ke Wang, Houxing Ren, Aojun Zhou, Zimu Lu, Sichun Luo, Weikang Shi, Renrui Zhang, Linqi Song, Mingjie Zhan, Hongsheng Li	19609

Beyond Reverse KL: Generalizing Direct Preference Optimization with Diverse Divergence Constraints Chaoqi Wang, Yibo Jiang, Chenghao Yang, Han Liu, Yuxin Chen	19643
Large Content and Behavior Models to Understand, Simulate, and Optimize Content and Behavior Ashmit Khandelwal, Aditya Agrawal, Aanisha Bhattacharyya, Yaman Kumar, Somesh Singh, Uttaran Bhattacharya, Ishita Dasgupta, Stefano Petrangeli, Rajiv Ratn Shah, Changyou Chen, Balaji Krishnamurthy	19674
Efficient ConvBN Blocks for Transfer Learning and Beyond Kaichao You, Guo Qin, Anchang Bao, Meng Cao, Ping Huang, Jiulong Shan, Mingsheng Long	19693
Self-Supervised Pocket Pretraining Via Protein Fragment-Surroundings Alignment Bowen Gao, Yinjun Jia, Yuanle Mo, Yuyan Ni, Wei-Ying Ma, Zhi-Ming Ma, Yanyan Lan	19714
Language Model Decoding as Direct Metrics Optimization Haozhe Ji, Pei Ke, Hongning Wang, Minlie Huang	19737
Hypothesis Search: Inductive Reasoning with Language Models Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, Noah Goodman	19770
Towards Codable Watermarking for Injecting Multi-Bits Information to LLMs Lean Wang, Wenkai Yang, Deli Chen, Hao Zhou, Yankai Lin, Fandong Meng, Jie Zhou, Xu Sun	19792
Rethinking the Symmetry-Preserving Circuits for Constrained Variational Quantum Algorithms Ge Yan, Hongxu Chen, Kaisen Pan, Junchi Yan	19819
Self-Supervised Dataset Distillation for Transfer Learning Dong Bok Lee, Seanie Lee, Joonho Ko, Kenji Kawaguchi, Juho Lee, Sung Ju Hwang	19839
Coordinate-Aware Modulation for Neural Fields Joo Chan Lee, Daniel Rho, Seungtae Nam, Jong Hwan Ko, Eunbyung Park	19856
At Which Training Stage Does Code Data Help LLMs Reasoning? Yingwei Ma, Yue Liu, Yue Yu, Yuanliang Zhang, Yu Jiang, Changjian Wang, Shanshan Li	19874
Minimum Width for Universal Approximation Using ReLU Networks on Compact Domain Namjun Kim, Chanho Min, Sejun Park	19894
Sample-Efficient Multi-Agent RL: An Optimization Perspective Nuoya Xiong, Zhihan Liu, Zhaoran Wang, Zhuoran Yang	19928
Peering Through Preferences: Unraveling Feedback Acquisition for Aligning Large Language Models Hritik Bansal, John Dang, Aditya Grover	19971

SAN: Inducing Metrizability of GAN with Discriminative Normalized Linear Layer Yuhta Takida, Masaaki Imaizumi, Takashi Shibuya, Chieh-Hsin Lai, Toshimitsu Uesaka, Naoki Murata, Yuki Mitsufuji	20003
Object-Aware Inversion and Reassembly for Image Editing Zhen Yang, Ganggui Ding, Wen Wang, Hao Chen, Bohan Zhuang, Chunhua Shen	20036
InstaFlow: One Step is Enough for High-Quality Diffusion-Based Text-To-Image Generation Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, Qiang Liu	20060
A Variational Framework for Estimating Continuous Treatment Effects with Measurement Error Erdun Gao, Howard Bondell, Wei Huang, Mingming Gong	20090
CRAFT: Customizing LLMs by Creating and Retrieving from Specialized Toolsets Lifan Yuan, Yangyi Chen, Xingyao Wang, Yi Fung, Hao Peng, Heng Ji	20116
PAC-FNO: Parallel-Structured All-Component Fourier Neural Operators for Recognizing Low-Quality Images Jinsung Jeon, Hyundong Jin, Jonghyun Choi, Sanghyun Hong, Dongeun Lee, Kookjin Lee, Noseong Park	20145
DiLu: A Knowledge-Driven Approach to Autonomous Driving with Large Language Models Licheng Wen, Daocheng Fu, Xin Li, Xinyu Cai, Tao Ma, Pinlong Cai, Min Dou, Botian Shi, Liang He, Yu Qiao	20165
Mixture-Of-Experts Meets Instruction Tuning: A Winning Combination for Large Language Models Sheng Shen, Le Hou, Yanqi Zhou, Nan Du, Shayne Longpre, Jason Wei, Hyung Won Chung, Barret Zoph, William Fedus, Xinyun Chen, Tu Vu, Yuxin Wu, Wuyang Chen, Albert Webson, Yunxuan Li, Vincent Y Zhao, Hongkun Yu, Kurt Keutzer, Trevor Darrell, Denny Zhou	20185
LEGO-Prover: Neural Theorem Proving with Growing Libraries Haiming Wang, Huajian Xin, Chuanyang Zheng, Zhengying Liu, Qingxing Cao, Yinya Huang, Jing Xiong, Han Shi, Enze Xie, Jian Yin, Zhenguo Li, Xiaodan Liang	20212
Dynamic Neighborhood Construction for Structured Large Discrete Action Spaces Fabian Akkerman, Julius Luy, Wouter Van Heeswijk, Maximilian Schiffer	20244
Efficient Backpropagation with Variance Controlled Adaptive Sampling Ziteng Wang, Jianfei Chen, Jun Zhu	20269
Unified Language-Vision Pretraining in LLM with Dynamic Discrete Visual Tokenization Yang Jin, Kun Xu, Kun Xu, Liwei Chen, Chao Liao, Jianchao Tan, Quzhe Huang, Bin Chen, Chengru Song, Dai Meng, Di Zhang, Wenwu Ou, Kun Gai, Yadong Mu	20286
AnyText: Multilingual Visual Text Generation and Editing Yuxiang Tuo, Wangmeng Xiang, Jun-Yan He, Yifeng Geng, Xuansong Xie	20307

NeRM: Learning Neural Representations for High-Framerate Human Motion Synthesis <i>Dong Wei, Huaijiang Sun, Bin Li, Xiaoning Sun, Shengxiang Hu, Weiqing Li, Jianfeng Lu</i>	20324
How Over-Parameterization Slows Down Gradient Descent in Matrix Sensing: The Curses of Symmetry and Initialization <i>Nuoya Xiong, Lijun Ding, Simon Shaolei Du</i>	20345
EventRPG: Event Data Augmentation with Relevance Propagation Guidance <i>Mingyuan Sun, Donghao Zhang, Zongyuan Ge, Wang Jiaxu, Jia Li, Zheng Fang, Renjing Xu</i>	20402
Efficient Multi-Agent Reinforcement Learning by Planning <i>Qihan Liu, Jianing Ye, Xiaoteng Ma, Jun Yang, Bin Liang, Chongjie Zhang</i>	20422
DrM: Mastering Visual Reinforcement Learning Through Dormant Ratio Minimization <i>Guowei Xu, Ruijie Zheng, Yongyuan Liang, Xiyao Wang, Zhecheng Yuan, Tianying Ji, Yu Luo, Xiaoyu Liu, Jiaxin Yuan, Pu Hua, Shuzhen Li, Yanjie Ze, Hal Daumã© III, Furong Huang, Huazhe Xu</i>	20444
Learning Multi-Agent Communication from Graph Modeling Perspective <i>Shengchao Hu, Li Shen, Ya Zhang, Dacheng Tao</i>	20469
Space Group Constrained Crystal Generation <i>Rui Jiao, Wenbing Huang, Yu Liu, Deli Zhao, Yang Liu</i>	20485
Towards Faithful XAI Evaluation Via Generalization-Limited Backdoor Watermark <i>Mengxi Ya, Yiming Li, Tao Dai, Bin Wang, Yong Jiang, Shu-Tao Xia</i>	20503
ScaleCrafter: Tuning-Free Higher-Resolution Visual Generation with Diffusion Models <i>Yingqing He, Shaoshu Yang, Haoxin Chen, Xiaodong Cun, Menghan Xia, Yong Zhang, Xintao Wang, Ran He, Qifeng Chen, Ying Shan</i>	20531
Reinforcement Symbolic Regression Machine <i>Yilong Xu, Yang Liu, Hao Sun</i>	20552
CARD: Channel Aligned Robust Blend Transformer for Time Series Forecasting <i>Xue Wang, Tian Zhou, Qingsong Wen, Jinyang Gao, Bolin Ding, Rong Jin</i>	20577
DreamTime: An Improved Optimization Strategy for Diffusion-Guided 3D Generation <i>Yukun Huang, Jianan Wang, Yukai Shi, Boshi Tang, Xianbiao Qi, Lei Zhang</i>	20616
TokenFlow: Consistent Diffusion Features for Consistent Video Editing <i>Michal Geyer, Omer Bar-Tal, Shai Bagon, Tali Dekel</i>	20637
Vision-Language Models Are Zero-Shot Reward Models for Reinforcement Learning <i>Juan Rocamonde, Victoriano Montesinos, Elvis Nava, Ethan Perez, David Lindner</i>	20650
Fast Updating Truncated SVD for Representation Learning with Sparse Matrices <i>Haoran Deng, Yang Yang, Jiahe Li, Cheng Chen, Weihao Jiang, Shiliang Pu</i>	20668

Improving the Convergence of Dynamic NeRFs Via Optimal Transport Sameera Ramasinghe, Violetta Shevchenko, Gil Avraham, Hisham Husain, Anton Van Den Hengel	20687
Multi-Source Diffusion Models for Simultaneous Music Generation and Separation Giorgio Mariani, Irene Tallini, Emilian Postolache, Michele Mancusi, Luca Cosmo, Emanuele Rodolà	20705
Relay Diffusion: Unifying Diffusion Process Across Resolutions for Image Synthesis Jiayan Teng, Wendi Zheng, Ming Ding, Wenyi Hong, Jianqiao Wangni, Zhuoyi Yang, Jie Tang	20726
Addressing Signal Delay in Deep Reinforcement Learning Wei Wang, Dongqi Han, Xufang Luo, Dongsheng Li	20749
Uni-RLHF: Universal Platform and Benchmark Suite for Reinforcement Learning with Diverse Human Feedback Yifu Yuan, Jianye Hao, Yi Ma, Zibin Dong, Hebin Liang, Jinyi Liu, Zhixin Feng, Kai Zhao, Yan Zheng	20774
HoloNets: Spectral Convolutions Do Extend to Directed Graphs Christian Koke, Daniel Cremers	20807
Dirichlet-Based Per-Sample Weighting by Transition Matrix for Noisy Label Learning Heesun Bae, Seungjae Shin, Byeonghu Na, Il-Chul Moon	20841
Emu: Generative Pretraining in Multimodality Quan Sun, Qiyang Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, Xinlong Wang	20876
Augmented Bayesian Policy Search Mahdi Kallel, Debabrota Basu, Riad Akrou, Carlo D'Eramo	20905
Hypergraph Dynamic System Jielong Yan, Yifan Feng, Shihui Ying, Yue Gao	20923
SweetDreamer: Aligning Geometric Priors in 2D Diffusion for Consistent Text-To-3D Weiyu Li, Rui Chen, Xuelin Chen, Ping Tan	20944
Exploring Diffusion Time-Steps for Unsupervised Representation Learning Zhongqi Yue, Jiankun Wang, Qianru Sun, Lei Ji, Eric I-Chao Chang, Hanwang Zhang	20967
Respect the Model: Fine-Grained and Robust Explanation with Sharing Ratio Decomposition Sangyu Han, Yearim Kim, Nojun Kwak	20978
Select to Perfect: Imitating Desired Behavior from Large Multi-Agent Data Tim Franzmeyer, Edith Elkind, Philip Torr, Jakob Nicolaus Foerster, Joao F. Henriques	21013
Illusory Attacks: Information-Theoretic Detectability Matters in Adversarial Attacks Tim Franzmeyer, Stephen Marcus McAleer, Joao F. Henriques, Jakob Nicolaus Foerster, Philip Torr, Adel Bibi, Christian Schroeder De Witt	21032

FedCDA: Federated Learning with Cross-Rounds Divergence-Aware Aggregation Haozhao Wang, Haoran Xu, Yichen Li, Yuan Xu, Ruixuan Li, Tianwei Zhang	21052
ReFusion: Improving Natural Language Understanding with Computation-Efficient Retrieval Representation Fusion Shangyu Wu, Ying Xiong, Yufei Cui, Xue Liu, Buzhou Tang, Tei-Wei Kuo, Chun Jason Xue	21064
SparseDFF: Sparse-View Feature Distillation for One-Shot Dexterous Manipulation Qianxu Wang, Haotong Zhang, Congyue Deng, Yang You, Hao Dong, Yixin Zhu, Leonidas Guibas	21080
DeepSPF: Spherical SO(3)-Equivariant Patches for Scan-To-CAD Estimation Driton Salihi, Adam Misik, Yuankai Wu, Constantin Patsch, Fabian Esteban Seguel, Eckehard Steinbach	21093
Idempotent Generative Network Assaf Shocher, Amil V Dravid, Yossi Gandelsman, Inbar Mosseri, Michael Rubinstein, Alexei A Efros	21114
Learning to Solve Class-Constrained Bin Packing Problems Via Encoder-Decoder Model Hanni Cheng, Ya Cong, Weihao Jiang, Shiliang Pu	21135
Unveiling and Manipulating Prompt Influence in Large Language Models Zijian Feng, Hanzhang Zhou, Zixiao Zhu, Junlang Qian, Kezhi Mao	21159
Skill Machines: Temporal Logic Skill Composition in Reinforcement Learning Geraud Nangue Tasse, Devon Jarvis, Steven James, Benjamin Rosman	21183
A Neural Framework for Generalized Causal Sensitivity Analysis Dennis Frauen, Fergus Imrie, Alicia Curth, Valentyn Melnychuk, Stefan Feuerriegel, Mihaela Van Der Schaar	21203
From Posterior Sampling to Meaningful Diversity in Image Restoration Noa Cohen, Hila Manor, Yuval Bahat, Tomer Michaeli	21235
DyVal: Dynamic Evaluation of Large Language Models for Reasoning Tasks Kaijie Zhu, Jiaao Chen, Jindong Wang, Neil Zhenqiang Gong, Diyi Yang, Xing Xie	21273
FedTrans: Client-Transparent Utility Estimation for Robust Federated Learning Mingkun Yang, Ran Zhu, Qing Wang, Jie Yang	21311
SGD Finds Then Tunes Features in Two-Layer Neural Networks with near-Optimal Sample Complexity: A Case Study in the XOR Problem Margalit Glasgow	21336
Knowledge Fusion of Large Language Models Fanqi Wan, Xinting Huang, Deng Cai, Xiaojun Quan, Wei Bi, Shuming Shi	21348
Approximating Nash Equilibria in Normal-Form Games Via Stochastic Optimization Ian Gemp, Luke Marris, Georgios Piliouras	21368

Information Bottleneck Analysis of Deep Neural Networks Via Lossy Compression Ivan Butakov, Alexander Tolmachev, Sofia Malanchuk, Anna Neopryatnaya, Alexey Frolov, Kirill Andreev	21416
RAPPER: Reinforced Rationale-Prompted Paradigm for Natural Language Explanation in Visual Question Answering Kai-Po Chang, Chi-Pin Huang, Wei-Yuan Cheng, Fu-En Yang, Chien-Yi Wang, Yung-Hsuan Lai, Yu-Chiang Frank Wang	21439
Learning Energy Decompositions for Partial Inference in GFlowNets Hyosoon Jang, Minsu Kim, Sungsoo Ahn	21457
Successor Heads: Recurring, Interpretable Attention Heads in the Wild Rhys Gould, Euan Ong, George Ogden, Arthur Conmy	21472
MMD Graph Kernel: Effective Metric Learning for Graphs Via Maximum Mean Discrepancy Yan Sun, Jicong Fan	21498
JointNet: Extending Text-To-Image Diffusion for Dense Distribution Modeling Jingyang Zhang, Shiwei Li, Yuanxun Lu, Tian Fang, David Neil McKinnon, Yanghai Tsin, Long Quan, Yao Yao	21532
An Analytical Solution to Gauss-Newton Loss for Direct Image Alignment Sergei Solonets, Daniil Sinitsyn, Lukas Von Stumberg, Nikita Araslanov, Daniel Cremers	21548
Learning Large DAGs is Harder than You Think: Many Losses Are Minimal for the Wrong DAG Jonas Seng, Matej Zečević, Devendra Singh Dhami, Kristian Kersting	21565
Fast, Expressive $\text{SE}(n)$ Equivariant Networks Through Weight-Sharing in Position-Orientation Space Erik J Bekkers, Sharvaree Vadgama, Rob Hesselink, Putri A Van Der Linden, David W. Romero	21576
Accelerated Convergence of Stochastic Heavy Ball Method Under Anisotropic Gradient Noise Rui Pan, Yuxing Liu, Xiaoyu Wang, Tong Zhang	21601
Elucidating the Exposure Bias in Diffusion Models Mang Ning, Mingxiao Li, Jianlin Su, Albert Ali Salah, Itir Onal Ertugrul	21641
Grokking as a First Order Phase Transition in Two Layer Networks Noa Rubin, Inbar Seroussi, Zohar Ringel	21664
Privacy-Preserving In-Context Learning with Differentially Private Few-Shot Generation Xinyu Tang, Richard Shin, Huseyin A Inan, Andre Manoel, Fatemehsadat Mireshghallah, Zinan Lin, Sivakanth Gopi, Janardhan Kulkarni, Robert Sim	21688
A Mutual Information Perspective on Federated Contrastive Learning Christos Louizos, Matthias Reisser, Denis Korzhenkov	21708

TOSS: High-Quality Text-Guided Novel View Synthesis from a Single Image Yukai Shi, Jianan Wang, He Cao, Boshi Tang, Xianbiao Qi, Tianyu Yang, Yukun Huang, Shilong Liu, Lei Zhang, Heung-Yeung Shum	21728
Inherently Interpretable Time Series Classification Via Multiple Instance Learning Joseph Early, Gavin Cheung, Kurt Cutajar, Hanting Xie, Jas Kandola, Niall Twomey	21750
Constrained Bi-Level Optimization: Proximal Lagrangian Value Function Approach and Hessian-Free Algorithm Wei Yao, Chengming Yu, Shangzhi Zeng, Jin Zhang	21779
FedInverse: Evaluating Privacy Leakage in Federated Learning Di Wu, Jun Bai, Yiliao Song, Junjun Chen, Wei Zhou, Yong Xiang, Atul Sajjanhar	21811
The Cost of Scaling Down Large Language Models: Reducing Model Size Affects Memory Before In-Context Learning Tian Jin, Nolan Clement, Xin Dong, Vaishnavh Nagarajan, Michael Carbin, Jonathan Ragan-Kelley, Gintare Karolina Dziugaite	21842
AuG-KD: Anchor-Based Mixup Generation for Out-Of-Domain Knowledge Distillation Zihao Tang, Zheqi Lv, Shengyu Zhang, Yifan Zhou, Xinyu Duan, Fei Wu, Kun Kuang	21870
True Knowledge Comes from Practice: Aligning Large Language Models with Embodied Environments Via Reinforcement Learning Weihao Tan, Wentao Zhang, Shanqi Liu, Longtao Zheng, Xinrun Wang, Bo An	21894
A Simple and Scalable Representation for Graph Generation Yunhui Jang, Seul Lee, Sungsoo Ahn	21942
Flat Minima in Linear Estimation and an Extended Gauss Markov Theorem Simon Segert	21969
Enhancing Neural Training Via a Correlated Dynamics Model Jonathan Brokman, Roy Betser, Rotem Turjeman, Tom Berkov, Ido Cohen, Guy Gilboa	21981
Can LLMs Express Their Uncertainty? an Empirical Evaluation of Confidence Elicitation in LLMs Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, Bryan Hooi	22016
Entity-Centric Reinforcement Learning for Object Manipulation from Pixels Dan Haramati, Tal Daniel, Aviv Tamar	22045
Unknown Domain Inconsistency Minimization for Domain Generalization Seungjae Shin, Heesun Bae, Byeonghu Na, Yoon-Yeong Kim, Il-Chul Moon	22080
On the Limitations of Temperature Scaling for Distributions with Overlaps Muthu Chidambaram, Rong Ge	22105

Separate and Diffuse: Using a Pretrained Diffusion Model for Better Source Separation Shahar Lutati, Eliya Nachmani, Lior Wolf	22132
Fixed-Budget Differentially Private Best Arm Identification Zhirui Chen, P. N. Karthik, Yeow Meng Chee, Vincent Tan	22146
Less is More: One-Shot Subgraph Reasoning on Large-Scale Knowledge Graphs Zhanke Zhou, Yongqi Zhang, Jiangchao Yao, Quanming Yao, Bo Han	22182
Pushing Mixture of Experts to the Limit: Extremely Parameter Efficient MoE for Instruction Tuning Ted Zadouri, Ahmet Ãœstüner, Arash Ahmadian, Beyza Ermis, Acyr Locatelli, Sara Hooker	22214
Adversarial Supervision Makes Layout-To-Image Diffusion Models Thrive Yumeng Li, Margret Keuper, Dan Zhang, Anna Khoreva	22234
Intelligent Switching for Reset-Free RL Darshan Patil, Janarthanan Rajendran, Glen Berseth, Sarath Chandar	22257
SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Möller, Joe Penna, Robin Rombach	22275
Out-Of-Distribution Detection by Leveraging Between-Layer Transformation Smoothness Fran Jelenić, Josip Jukić, Martin Tutek, Mate Puljiz, Jan Snajder	22288
Course Correcting Koopman Representations Mahan Fathi, Clement Gehring, Jonathan Pilault, David Kanaa, Pierre-Luc Bacon, Ross Goroshin	22320
Random Sparse Lifts: Construction, Analysis and Convergence of Finite Sparse Networks David A. R. Robin, Kevin Scaman, Marc Lelarge	22337
NoiseDiffusion: Correcting Noise for Image Interpolation with Diffusion Models Beyond Spherical Linear Interpolation Pengfei Zheng, Yonggang Zhang, Zhen Fang, Tongliang Liu, Defu Lian, Bo Han	22376
PB-LLM: Partially Binarized Large Language Models Zhihang Yuan, Yuzhang Shang, Zhen Dong	22401
Conformal Inductive Graph Neural Networks Soroush H. Zargarbashi, Aleksandar Bojchevski	22415
Does Writing with Language Models Reduce Content Diversity? Vishakh Padmakumar, He He	22438
Vision Transformers Need Registers Timothée Darcet, Maxime Oquab, Julien Mairal, Piotr Bojanowski	22466

Counterfactual Density Estimation Using Kernel Stein Discrepancies <i>Diego Martinez-Taboada, Edward Kennedy</i>	22487
A Framework for Inference Inspired by Human Memory Mechanisms <i>Xiangyu Zeng, Jie Lin, Piao Hu, Ruizheng Huang, Zhicheng Zhang</i>	22509
One Forward is Enough for Neural Network Training Via Likelihood Ratio Method <i>Jinyang Jiang, Zeliang Zhang, Chenliang Xu, Zhaoifei Yu, Yijie Peng</i>	22531
Causal Fairness Under Unobserved Confounding: A Neural Sensitivity Framework <i>Maresa Schröder, Dennis Frauen, Stefan Feuerriegel</i>	22555
ALAM: Averaged Low-Precision Activation for Memory-Efficient Training of Transformer Models <i>Sunghyeon Woo, Sunwoo Lee, Dongsuk Jeon</i>	22592
Selective Mixup Fine-Tuning for Optimizing Non-Decomposable Objectives <i>Shrinivas Ramasubramanian, Harsh Rangwani, Sho Takemori, Kunal Samanta, Yuhei Umeda, Venkatesh Babu Radhakrishnan</i>	22611
CAMBranch: Contrastive Learning with Augmented MILPs for Branching <i>Jiacheng Lin, Meng Xu, Zhihua Xiong, Huangang Wang</i>	22641
From Sparse to Soft Mixtures of Experts <i>Joan Puigcerver, Carlos Riquelme Ruiz, Basil Mustafa, Neil Houlsby</i>	22665
Learning 3D Particle-Based Simulators from RGB-D Videos <i>William F Whitney, Tatiana Lopez-Guevara, Tobias Pfaff, Yulia Rubanova, Thomas Kipf, Kim Stachenfeld, Kelsey R Allen</i>	22676
A 2-Dimensional State Space Layer for Spatial Inductive Bias <i>Ethan Baron, Itamar Zimerman, Lior Wolf</i>	22702
Large Multilingual Models Pivot Zero-Shot Multimodal Learning Across Languages <i>Jinyi Hu, Yuan Yao, Chongyi Wang, Shan Wang, Yinxu Pan, Qianyu Chen, Tianyu Yu, Hanghao Wu, Yue Zhao, Haoye Zhang, Xu Han, Yankai Lin, Jiao Xue, Dahai Li, Zhiyuan Liu, Maosong Sun</i>	22721
LOQA: Learning with Opponent Q-Learning Awareness <i>Milad Aghajohari, Juan Agustin Duque, Tim Cooijmans, Aaron Courville</i>	22747
It's Never Too Late: Fusing Acoustic Information into Large Language Models for Automatic Speech Recognition <i>Chen Chen, Ruizhe Li, Yuchen Hu, Sabato Marco Siniscalchi, Pin-Yu Chen, Engsiong Chng, Chao-Han Huck Yang</i>	22763
Tree-Planner: Efficient Close-Loop Task Planning with Large Language Models <i>Mengkang Hu, Yao Mu, Xinmiao Chelsey Yu, Mingyu Ding, Shiguang Wu, Wenqi Shao, Qiguang Chen, Bin Wang, Yu Qiao, Ping Luo</i>	22780

FairTune: Optimizing Parameter Efficient Fine Tuning for Fairness in Medical Image Analysis Raman Dutt, Ondrej Bohdal, Sotirios A. Tsaftaris, Timothy Hospedales	22808
Neural Atoms: Propagating Long-Range Interaction in Molecular Graphs Through Efficient Communication Channel Xuan Li, Zhanke Zhou, Jiangchao Yao, Yu Rong, Lu Zhang, Bo Han	22832
Discovering Modular Solutions that Generalize Compositionally Simon Schug, Seijin Kobayashi, Yassir Akram, Maciej Wolczyk, Alexandra Maria Proca, Johannes Von Oswald, Razvan Pascanu, Joao Sacramento, Angelika Steger	22865
Efficient and Scalable Graph Generation Through Iterative Local Expansion Andreas Bergmeister, Karolis Martinkus, Nathanaël Perraudin, Roger Wattenhofer	22918
SEPT: Towards Efficient Scene Representation Learning for Motion Prediction Zhiqian Lan, Yuxuan Jiang, Yao Mu, Chen Chen, Shengbo Eben Li	22957
An Investigation of Representation and Allocation Harms in Contrastive Learning Subha Maity, Mayank Agarwal, Mikhail Yurochkin, Yuekai Sun	22971
A Path-Norm Toolkit for Modern Networks: Consequences, Promises and Challenges Antoine Gonon, Nicolas Brisebarre, Elisa Riccietti, R�mi Gribonval	22996
Compressing LLMs: The Truth is Rarely Pure and Never Simple Ajay Kumar Jaiswal, Zhe Gan, Xianzhi Du, Bowen Zhang, Zhangyang Wang, Yinfei Yang	23041
Relaxing the Additivity Constraints in Decentralized No-Regret High-Dimensional Bayesian Optimization Anthony Bardou, Patrick Thiran, Thomas Begin	23062
CODE REPRESENTATION LEARNING AT SCALE Dejiao Zhang, Wasi Uddin Ahmad, Ming Tan, Hantian Ding, Ramesh Nallapati, Dan Roth, Xiaofei Ma, Bing Xiang	23089
Lemur: Harmonizing Natural Language and Code for Language Agents Yiheng Xu, Hongjin Su, Chen Xing, Boyu Mi, Qian Liu, Weijia Shi, Binyuan Hui, Fan Zhou, Yitao Liu, Tianbao Xie, Zhoujun Cheng, Siheng Zhao, Lingpeng Kong, Bailin Wang, Caiming Xiong, Tao Yu	23109
Efficient-3Dim: Learning a Generalizable Single-Image Novel-View Synthesizer in One Day Yifan Jiang, Hao Tang, Jen-Hao Rick Chang, Liangchen Song, Zhangyang Wang, Liangliang Cao	23130
How I Warped Your Noise: A Temporally-Correlated Noise Prior for Diffusion Models Pascal Chang, Jingwei Tang, Markus Gross, Vinicius C. Azevedo	23146
Depthwise Hyperparameter Transfer in Residual Networks: Dynamics and Scaling Limit Blake Bordelon, Lorenzo Noci, Mufan Bill Li, Boris Hanin, Cengiz Pehlevan	23178
MAMBA: An Effective World Model Approach for Meta-Reinforcement Learning Zohar Rimon, Tom Jurgenson, Orr Krupnik, Gilad Adler, Aviv Tamar	23218

A General Framework for User-Guided Bayesian Optimization Carl Hvarfner, Frank Hutter, Luigi Nardi	23238
Efficient Video Diffusion Models Via Content-Frame Motion-Latent Decomposition Sihyun Yu, Weili Nie, De-An Huang, Boyi Li, Jinwoo Shin, Anima Anandkumar	23256
Generalized Schrödinger Bridge Matching Guan-Hong Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou, Ricky T. Q. Chen	23285
FairerCLIP: Debiasing CLIP's Zero-Shot Predictions Using Functions in RKHSs Sepehr Dehdashtian, Lan Wang, Vishnu Boddeti	23311
Small-Scale Proxies for Large-Scale Transformer Training Instabilities Mitchell Wortsman, Peter J Liu, Lechao Xiao, Katie E Everett, Alexander A Alemi, Ben Adlam, John D Co-Reyes, Izzeddin Gur, Abhishek Kumar, Roman Novak, Jeffrey Pennington, Jascha Sohl-Dickstein, Kelvin Xu, Jaehoon Lee, Justin Gilmer, Simon Kornblith	23329
Zero-Shot Continuous Prompt Transfer: Generalizing Task Semantics Across Language Models Zijun Wu, Yongkang Wu, Lili Mou	23355
Oracle Efficient Algorithms for Groupwise Regret Krishna Acharya, Eshwar Ram Arunachaleswaran, Sampath Kannan, Aaron Roth, Juba Ziani	23370
Fast Value Tracking for Deep Reinforcement Learning Frank Shih, Faming Liang	23393
The Reasonableness Behind Unreasonable Translation Capability of Large Language Model Tingchen Fu, Lemao Liu, Deng Cai, Guoping Huang, Shuming Shi, Rui Yan	23418
Geographic Location Encoding with Spherical Harmonics and Sinusoidal Representation Networks Marc Rußwurm, Konstantin Klemmer, Esther Rolf, Robin Zbinden, Devis Tuia	23438
A Restoration Network as an Implicit Prior Yuyang Hu, Mauricio Delbracio, Peyman Milanfar, Ulugbek Kamilov	23452
αTC-VAE: On the Relationship Between Disentanglement and Diversity Cristian Meo, Louis Mahon, Anirudh Goyal, Justin Dauwels	23477
DP-OPT: Make Large Language Model Your Privacy-Preserving Prompt Engineer Junyuan Hong, Jiachen T. Wang, Chenhui Zhang, Zhangheng Li, Bo Li, Zhangyang Wang	23504
Model Merging by Uncertainty-Based Gradient Matching Nico Daheim, Thomas Mollenhoff, Edoardo Ponti, Iryna Gurevych, Mohammad Emtiyaz Khan	23528

Scalable Neural Network Kernels Arijit Sehanobish, Krzysztof Marcin Choromanski, Yunfan Zhao, Kumar Avinava Dubey, Valerii Likhoshesterov	23547
SE(3)-Stochastic Flow Matching for Protein Backbone Generation Joey Bose, Tara Akhound-Sadegh, Guillaume Huguet, Kilian Fatras, Jarrod Rector-Brooks, Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael M. Bronstein, Alexander Tong	23573
Matryoshka Diffusion Models Jiatao Gu, Shuangfei Zhai, Yizhe Zhang, Joshua M. Susskind, Navdeep Jaitly	23605
LILO: Learning Interpretable Libraries by Compressing and Documenting Code Gabriel Grand, Lionel Wong, Matthew Bowers, Theo X. Olausson, Muxin Liu, Joshua B. Tenenbaum, Jacob Andreas	23634
SAFLEX: Self-Adaptive Augmentation Via Feature Label Extrapolation Mucong Ding, Bang An, Yuancheng Xu, Anirudh Satheesh, Furong Huang	23682
Reward Design for Justifiable Sequential Decision-Making Aleksa Sukovic, Goran Radanovic	23695
Latent 3D Graph Diffusion Yuning You, Ruida Zhou, Jiwoong Park, Haotian Xu, Chao Tian, Zhangyang Wang, Yang Shen	23715
Learning Stackable and Skippable LEGO Bricks for Efficient, Reconfigurable, and Variable-Resolution Diffusion Modeling Huangjie Zheng, Zhendong Wang, Jianbo Yuan, Guanghan Ning, Pengcheng He, Quanzeng You, Hongxia Yang, Mingyuan Zhou	23742
DNABERT-2: Efficient Foundation Model and Benchmark for Multi-Species Genomes Zhihan Zhou, Yanrong Ji, Weijian Li, Pratik Dutta, Ramana V Davuluri, Han Liu	23768
A Linear Algebraic Framework for Counterfactual Generation Jong-Hoon Ahn, Akshay Vashist	23792
Differentially Private SGD Without Clipping Bias: An Error-Feedback Approach Xinwei Zhang, Zhiqi Bu, Steven Wu, Mingyi Hong	23814
Interpreting Robustness Proofs of Deep Neural Networks Debangshu Banerjee, Avaljot Singh, Gagandeep Singh	23841
MgNO: Efficient Parameterization of Linear Operators Via Multigrid Juncai He, Xinliang Liu, Jinchao Xu	23875
Is Self-Repair a Silver Bullet for Code Generation? Theo X. Olausson, Jeevana Priya Inala, Chenglong Wang, Jianfeng Gao, Armando Solar-Lezama	23895
Sudden Drops in the Loss: Syntax Acquisition, Phase Transitions, and Simplicity Bias in MLMs Angelica Chen, Ravid Shwartz-Ziv, Kyunghyun Cho, Matthew L Leavitt, Naomi Saphra	23944

Unified Generative Modeling of 3D Molecules with Bayesian Flow Networks Yuxuan Song, Jingjing Gong, Hao Zhou, Mingyue Zheng, Jingjing Liu, Wei-Ying Ma	23976
BaDExpert: Extracting Backdoor Functionality for Accurate Backdoor Input Detection Tinghao Xie, Xiangyu Qi, Ping He, Yiming Li, Jiachen T. Wang, Prateek Mittal	23995
SCHEMA: State CHangEs MATter for Procedure Planning in Instructional Videos Yulei Niu, Wenliang Guo, Long Chen, Xudong Lin, Shih-Fu Chang	24023
Improved Probabilistic Image-Text Representations Sanghyuk Chun	24043
Understanding Augmentation-Based Self-Supervised Representation Learning Via RKHS Approximation and Regression Runtian Zhai, Bingbin Liu, Andrej Risteski, J Zico Kolter, Pradeep Kumar Ravikumar	24072
Test-Time Training on Nearest Neighbors for Large Language Models Moritz Hardt, Yu Sun	24106
Prediction Error-Based Classification for Class-Incremental Learning Micha�, Zaji�c, Tinne Tuytelaars, Gido M Van De Ven	24122
Parametric Augmentation for Time Series Contrastive Learning Xu Zheng, Tianchun Wang, Wei Cheng, Aitian Ma, Haifeng Chen, Mo Sha, Dongsheng Luo	24151
Polynomial Width is Sufficient for Set Representation with High-Dimensional Features Peihao Wang, Shenghao Yang, Shu Li, Zhangyang Wang, Pan Li	24177
GenSim: Generating Robotic Simulation Tasks Via Large Language Models Lirui Wang, Yiyang Ling, Zhecheng Yuan, Mohit Shridhar, Chen Bao, Yuzhe Qin, Bailin Wang, Huazhe Xu, Xiaolong Wang	24207
Representation Deficiency in Masked Language Modeling Yu Meng, Jitin Krishnan, Sinong Wang, Qifan Wang, Yuning Mao, Han Fang, Marjan Ghazvininejad, Jiawei Han, Luke Zettlemoyer	24242
MOTOR: A Time-To-Event Foundation Model for Structured Medical Records Ethan Steinberg, Jason Alan Fries, Yizhe Xu, Nigam Shah	24263
Conformal Language Modeling Victor Quach, Adam Fisch, Tal Schuster, Adam Yala, Jae Ho Sohn, Tommi S. Jaakkola, Regina Barzilay	24303
Critical Learning Periods Emerge Even in Deep Linear Networks Michael Kleinman, Alessandro Achille, Stefano Soatto	24331
Robust NAS Under Adversarial Training: Benchmark, Theory, and Beyond Yongtao Wu, Fanghui Liu, Carl-Johann Simon-Gabriel, Grigorios Chrysos, Volkan Cevher	24348

Chameleon: Increasing Label-Only Membership Leakage with Adaptive Poisoning Harsh Chaudhari, Giorgio Severi, Alina Oprea, Jonathan Ullman	24385
EquiformerV2: Improved Equivariant Transformer for Scaling to Higher-Degree Representations Yi-Lun Liao, Brandon M Wood, Abhishek Das, Tess Smidt	24405
SOHES: Self-Supervised Open-World Hierarchical Entity Segmentation Shengcao Cao, Jiuxiang Gu, Jason Kuen, Hao Tan, Ruiyi Zhang, Handong Zhao, Ani Nenkova, Liangyan Gui, Tong Sun, Yu-Xiong Wang	24433
Look, Remember and Reason: Grounded Reasoning in Videos with Language Models Apratim Bhattacharyya, Sunny Panchal, Reza Pourreza, Mingu Lee, Pulkit Madan, Roland Memisevic	24455
Conditional Instrumental Variable Regression with Representation Learning for Causal Inference Debo Cheng, Ziqi Xu, Jiuyong Li, Lin Liu, Jixue Liu, Thuc Duy Le	24474
Evaluating the Zero-Shot Robustness of Instruction-Tuned Language Models Jiuding Sun, Chantal Shaib, Byron C Wallace	24491
SmartPlay : A Benchmark for LLMs as Intelligent Agents Yue Wu, Xuan Tang, Tom Mitchell, Yuanzhi Li	24530
Whole-Song Hierarchical Generation of Symbolic Music Using Cascaded Diffusion Models Ziyu Wang, Lejun Min, Gus Xia	24549
Belief-Enriched Pessimistic Q-Learning Against Adversarial State Perturbations Xiaolin Sun, Zizhan Zheng	24570
Skill Or Luck? Return Decomposition Via Advantage Functions Hsiao-Ru Pan, Bernhard Schölkopf	24594
Learning Dynamic Representations of the Functional Connectome in Neurobiological Networks Luciano Dyballa, Samuel Lang, Alexandra Haslund-Gourley, Eviatar Yemini, Steven W. Zucker	24617
Memory-Consistent Neural Networks for Imitation Learning Kaustubh Sridhar, Souradeep Dutta, Dinesh Jayaraman, James Weimer, Insup Lee	24633
CCIL: Continuity-Based Data Augmentation for Corrective Imitation Learning Liyiming Ke, Yunchu Zhang, Abhay Deshpande, Siddhartha Srinivasa, Abhishek Gupta	24659
Str2Str: A Score-Based Framework for Zero-Shot Protein Conformation Sampling Jiarui Lu, Bozitao Zhong, Zuobai Zhang, Jian Tang	24678
Provable Robust Watermarking for AI-Generated Text Xuandong Zhao, Prabhanjan Vijendra Ananth, Lei Li, Yu-Xiang Wang	24710

Principled Architecture-Aware Scaling of Hyperparameters Wuyang Chen, Junru Wu, Zhangyang Wang, Boris Hanin	24745
CALICO: Self-Supervised Camera-LiDAR Contrastive Pre-Training for BEV Perception Jiachun Sun, Haizhong Zheng, Qingzhao Zhang, Atul Prakash, Zhuoqing Mao, Chaowei Xiao	24765
Language Control Diffusion: Efficiently Scaling Through Space, Time, and Tasks Edwin Zhang, Yujie Lu, Shinda Huang, William Yang Wang, Amy Zhang	24786
CellPLM: Pre-Training of Cell Language Model Beyond Single Cells Hongzhi Wen, Wenzhuo Tang, Xinnan Dai, Jiayuan Ding, Wei Jin, Yuying Xie, Jiliang Tang	24817
Federated Causal Discovery from Heterogeneous Data Longkang Li, Ignavier Ng, Gongxu Luo, Biwei Huang, Guangyi Chen, Tongliang Liu, Bin Gu, Kun Zhang	24842
DiffusionNAG: Predictor-Guided Neural Architecture Generation with Diffusion Models Sohyun An, Hayeon Lee, Jaehyeong Jo, Seanie Lee, Sung Ju Hwang	24876
Low Rank Matrix Completion Via Robust Alternating Minimization in Nearly Linear Time Yuzhou Gu, Zhao Song, Junze Yin, Lichen Zhang	24901
Implicit Bias of SGD in $\mathcal{L}_2$-Regularized Linear DNNs: One-Way Jumps from High to Low Rank Zihan Wang, Arthur Jacot	24954
Project and Probe: Sample-Efficient Adaptation by Interpolating Orthogonal Features Annie S Chen, Yoonho Lee, Amrith Setlur, Sergey Levine, Chelsea Finn	24981
Reasoning with Latent Diffusion in Offline Reinforcement Learning Siddarth Venkatraman, Shivesh Khaitan, Ravi Tej Akella, John Dolan, Jeff Schneider, Glen Berseth	25005
The Hidden Language of Diffusion Models Hila Chefer, Oran Lang, Mor Geva, Volodymyr Polosukhin, Assaf Shocher, Michal Irani, Inbar Mosseri, Lior Wolf	25027
Self-Consuming Generative Models Go MAD Sina Alemohammad, Josue Casco-Rodriguez, Lorenzo Luzi, Ahmed Imtiaz Humayun, Hossein Babaei, Daniel Lejeune, Ali Siahkoobi, Richard Baraniuk	25060
Federated Orthogonal Training: Mitigating Global Catastrophic Forgetting in Continual Federated Learning Yavuz Faruk Bakman, Duygu Nur Yaldiz, Yahya H. Ezzeldin, Salman Avestimehr	25088
MBR and QE Finetuning: Training-Time Distillation of the Best and Most Expensive Decoding Methods Mara Finkelstein, Markus Freitag	25113

Adaptive Sharpness-Aware Pruning for Robust Sparse Networks Anna Bair, Hongxu Yin, Maying Shen, Pavlo Molchanov, Jose M. Alvarez	25132
OpenTab: Advancing Large Language Models as Open-Domain Table Reasoners Kezhi Kong, Jiani Zhang, Zhengyuan Shen, Balasubramaniam Srinivasan, Chuan Lei, Christos Faloutsos, Huzefa Rangwala, George Karypis	25154
Neurosymbolic Grounding for Compositional World Models Atharva Sehgal, Arya Grayeli, Jennifer J. Sun, Swarat Chaudhuri	25171
Entropy Coding of Unordered Data Structures Julius Kunze, Daniel Severo, Giulio Zani, Jan-Willem Van De Meent, James Townsend	25189
Zeroth-Order Optimization Meets Human Feedback: Provable Learning Via Ranking Oracles Zhiwei Tang, Dmitry Rybin, Tsung-Hui Chang	25206
RLCD: Reinforcement Learning from Contrastive Distillation for LM Alignment Kevin Yang, Dan Klein, Asli Celikyilmaz, Nanyun Peng, Yuandong Tian	25237
Deep Neural Networks Tend to Extrapolate Predictably Katie Kang, Amrith Setlur, Claire Tomlin, Sergey Levine	25277
Conditional Information Bottleneck Approach for Time Series Imputation Mingyu Choi, Changhee Lee	25306
Causal-StoNet: Causal Inference for High-Dimensional Complex Data Yaxin Fang, Faming Liang	25328
Det-CGD: Compressed Gradient Descent with Matrix Stepsizes for Non-Convex Optimization Hanmin Li, Avetik Karagulyan, Peter Richtárik	25352
On the Hardness of Constrained Cooperative Multi-Agent Reinforcement Learning Ziyi Chen, Yi Zhou, Heng Huang	25364
Views Can Be Deceiving: Improved SSL Through Feature Space Augmentation Kimia Hamidieh, Haoran Zhang, Swami Sankaranarayanan, Marzyeh Ghassemi	25401
Federated Text-Driven Prompt Generation for Vision-Language Models Chen Qiu, Xingyu Li, Chaithanya Kumar Mummadi, Madan Ravi Ganesh, Zhenzhen Li, Lu Peng, Wan-Yi Lin	25414
A Lie Group Approach to Riemannian Batch Normalization Ziheng Chen, Yue Song, Yunmei Liu, Nicu Sebe	25430
Time Fairness in Online Knapsack Problems Adam Lechowicz, Rik Sengupta, Bo Sun, Shahin Kamali, Mohammad Hajiesmaili	25455
GOAt: Explaining Graph Neural Networks Via Graph Output Attribution Shengyao Lu, Keith G Mills, Jiao He, Bang Liu, Di Niu	25481

<i>Listen, Think, and Understand</i> Yuan Gong, Hongyin Luo, Alexander H. Liu, Leonid Karlinsky, James R. Glass	25501
<i>Aligning Relational Learning with Lipschitz Fairness</i> Yaning Jia, Chunhui Zhang, Soroush Vosoughi	25531
<i>Fair Classifiers that Abstain Without Harm</i> Tongxin Yin, Jean-Francois Ton, Ruocheng Guo, Yuanshun Yao, Mingyan Liu, Yang Liu	25557
<i>DrS: Learning Reusable Dense Rewards for Multi-Stage Tasks</i> Tongzhou Mu, Minghua Liu, Hao Su	25584
<i>A Dynamical View of the Question of Why</i> Mehdi Fatemi, Sindhu C. M. Gowda	25597
<i>Variance Reduced Halpern Iteration for Finite-Sum Monotone Inclusions</i> Xufeng Cai, Ahmet Alacaoglu, Jelena Diakonikolas	25628
<i>ReMasker: Imputing Tabular Data with Masked Autoencoding</i> Tianyu Du, Luca Melis, Ting Wang	25661
<i>Teaching Arithmetic to Small Transformers</i> Nayoung Lee, Kartik Sreenivasan, Jason D. Lee, Kangwook Lee, Dimitris Papailiopoulos	25684
<i>MetaPhysiCa: Improving OOD Robustness in Physics-Informed Machine Learning</i> S Chandra Mouli, Muhammad Alam, Bruno Ribeiro	25738
<i>AUC-CL: A Batchsize-Robust Framework for Self-Supervised Contrastive Representation Learning</i> Rohan Sharma, Kaiyi Ji, Zhiqiang Xu, Changyou Chen	25764
<i>LabelDP-Pro: Learning with Label Differential Privacy Via Projections</i> Badih Ghazi, Yangsibo Huang, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, Chiyuan Zhang	25785
<i>Soft Mixture Denoising: Beyond the Expressive Bottleneck of Diffusion Models</i> Yangming Li, Boris Van Breugel, Mihaela Van Der Schaar	25810
<i>The Update-Equivalence Framework for Decision-Time Planning</i> Samuel Sokota, Gabriele Farina, David J Wu, Hengyuan Hu, Kevin A. Wang, J Zico Kolter, Noam Brown	25831
<i>On Error Propagation of Diffusion Models</i> Yangming Li, Mihaela Van Der Schaar	25848
<i>CLIP the Bias: How Useful is Balancing Data in Multimodal Learning?</i> Ibrahim Alabdulmohsin, Xiao Wang, Andreas Peter Steiner, Priya Goyal, Alexander D'Amour, Xiaohua Zhai	25865

Unleashing the Potential of Fractional Calculus in Graph Neural Networks with FROND Qiyu Kang, Kai Zhao, Qinxu Ding, Feng Ji, Xuhao Li, Wenfei Liang, Yang Song, Wee Peng Tay	25897
Multi-Resolution Diffusion Models for Time Series Forecasting Lifeng Shen, Weiyu Chen, James Kwok	25939
DeepZero: Scaling Up Zeroth-Order Optimization for Deep Model Training Aochuan Chen, Yimeng Zhang, Jinghan Jia, James Diffenderfer, Konstantinos Parasyris, Jiancheng Liu, Yihua Zhang, Zheng Zhang, Bhavya Kailkhura, Sijia Liu	25958
The Truth is in There: Improving Reasoning in Language Models with Layer-Selective Rank Reduction Pratyusha Sharma, Jordan T. Ash, Dipendra Misra	25980
Adaptive Retrieval and Scalable Indexing for k-NN Search with Cross-Encoders Nishant Yadav, Nicholas Monath, Manzil Zaheer, Rob Fergus, Andrew McCallum	26000
Provable Reward-Agnostic Preference-Based Reinforcement Learning Wenhao Zhan, Masatoshi Uehara, Wen Sun, Jason D. Lee	26028
Multimodal Learning Without Labeled Multimodal Data: Guarantees and Applications Paul Pu Liang, Chun Kai Ling, Yun Cheng, Alexander Obolenskiy, Yudong Liu, Rohan Pandey, Alex Wilf, Louis-Philippe Morency, Russ Salakhutdinov	26056
Can LLMs Keep a Secret? Testing Privacy Implications of Language Models Via Contextual Integrity Theory Niloofer Mireshghallah, Hyunwoo Kim, Xuhui Zhou, Yulia Tsvetkov, Maarten Sap, Reza Shokri, Yejin Choi	26088
Provable Offline Preference-Based Reinforcement Learning Wenhao Zhan, Masatoshi Uehara, Nathan Kallus, Jason D. Lee, Wen Sun	26112
Sample Efficient Myopic Exploration Through Multitask Reinforcement Learning with Diverse Tasks Ziping Xu, Zifan Xu, Runxuan Jiang, Peter Stone, Ambuj Tewari	26138
JoMA: Demystifying Multilayer Transformers Via Joint Dynamics of MLP and Attention Yuandong Tian, Yiping Wang, Zhenyu Zhang, Beidi Chen, Simon Shaolei Du	26151
Robustifying State-Space Models for Long Sequences Via Approximate Diagonalization Annan Yu, Arnur Nigmatov, Dmitriy Morozov, Michael W. Mahoney, N. Benjamin Erichson	26177
Understanding When Dynamics-Invariant Data Augmentations Benefit Model-Free Reinforcement Learning Updates Nicholas Corrado, Josiah P. Hanna	26210
Multimodal Molecular Pretraining Via Modality Blending Qiyang Yu, Yudi Zhang, Yuyan Ni, Shikun Feng, Yanyan Lan, Hao Zhou, Jingjing Liu	26235

Learning with Language-Guided State Abstractions Andi Peng, Ilia Sucholutsky, Belinda Z. Li, Theodore Summers, Thomas L. Griffiths, Jacob Andreas, Julie Shah	26254
ReLoRA: High-Rank Training Through Low-Rank Updates Vladislav Lialin, Sherin Muckatira, Namrata Shivagunde, Anna Rumshisky	26288
Fantastic Gains and Where to Find Them: On the Existence and Prospect of General Knowledge Transfer Between Any Pretrained Model Karsten Roth, Lukas Thede, A. Sophia Koepke, Oriol Vinyals, Olivier J Henaff, Zeynep Akata	26305
Enabling Lanuguage Models to Implicitly Learn Self-Improvement Ziqi Wang, Le Hou, Tianjian Lu, Yuexin Wu, Yunxuan Li, Hongkun Yu, Heng Ji	26329
Teaching Language Models to Hallucinate Less with Synthetic Tasks Erik Jones, Hamid Palangi, Clarisse Simões Ribeiro, Varun Chandrasekaran, Subhabrata Mukherjee, Arindam Mitra, Ahmed Hassan Awadallah, Ece Kamar	26360
Neural Neighborhood Search for Multi-Agent Path Finding Zhongxia Yan, Cathy Wu	26380
Generative Human Motion Stylization in Latent Space Chuan Guo, Yuxuan Mu, Xinxin Zuo, Peng Dai, Youliang Yan, Juwei Lu, Li Cheng	26402
On Differentially Private Federated Linear Contextual Bandits Xingyu Zhou, Sayak Ray Chowdhury	26422
Motif: Intrinsic Motivation from Artificial Intelligence Feedback Martin Klissarov, Pierluca D'Oro, Shagun Sodhani, Roberta Raileanu, Pierre-Luc Bacon, Pascal Vincent, Amy Zhang, Mikael Henaff	26453
Guess & Sketch: Language Model Guided Transpilation Celine Lee, Abdulrahman Mahmoud, Michal Kurek, Simone Campanoni, David Brooks, Stephen Chong, Gu-Yeon Wei, Alexander M Rush	26487
Towards Optimal Feature-Shaping Methods for Out-Of-Distribution Detection Qinyu Zhao, Ming Xu, Kartik Gupta, Akshay Asthana, Liang Zheng, Stephen Gould	26507
Scalable Diffusion for Materials Generation Sherry Yang, Kwanghwan Cho, Amil Merchant, Pieter Abbeel, Dale Schuurmans, Igor Mordatch, Ekin Dogus Cubuk	26525
SNIP: Bridging Mathematical Symbolic and Numeric Realms with Unified Pre-Training Kazem Meidani, Parshin Shojaei, Chandan K. Reddy, Amir Barati Farimani	26541
Score-Based Generative Models Break the Curse of Dimensionality in Learning a Family of sub-Gaussian Distributions Frank Cole, Yulong Lu	26555

Faithful Rule Extraction for Differentiable Rule Learning Models <i>Xiaxia Wang, David Jaime Tena Cucala, Bernardo Cuenca Grau, Ian Horrocks</i>	26567
Horizon-Free Regret for Linear Markov Decision Processes <i>Zhang Zihan, Jason D. Lee, Yuxin Chen, Simon Shaolei Du</i>	26590
Probabilistic Adaptation of Black-Box Text-To-Video Models <i>Sherry Yang, Yilun Du, Bo Dai, Dale Schuurmans, Joshua B. Tenenbaum, Pieter Abbeel</i>	26602
Smooth ECE: Principled Reliability Diagrams Via Kernel Smoothing <i>Jaroslav Blasiok, Preetum Nakkiran</i>	26619
GIO: Gradient Information Optimization for Training Dataset Selection <i>Dante Everaert, Christopher Potts</i>	26648
Empirical Likelihood for Fair Classification <i>Pangpang Liu, Yichuan Zhao</i>	26684
LDReg: Local Dimensionality Regularized Self-Supervised Learning <i>Hanxun Huang, Ricardo J. G. B. Campello, Sarah Monazam Erfani, Xingjun Ma, Michael E. Houle, James Bailey</i>	26699
WÄrstchen: An Efficient Architecture for Large-Scale Text-To-Image Diffusion Models <i>Pablo Pernias, Dominic Rampas, Mats Leon Richter, Christopher Pal, Marc Aubreville</i>	26725
FlashAttention-2: Faster Attention with Better Parallelism and Work Partitioning <i>Tri Dao</i>	26738
Improving Intrinsic Exploration by Creating Stationary Objectives <i>Roger Creus Castanyer, Joshua Romoff, Glen Berseth</i>	26752
Modulate Your Spectrum in Self-Supervised Learning <i>Xi Weng, Yunhao Ni, Tengwei Song, Jie Luo, Rao Muhammad Anwer, Salman Khan, Fahad Khan, Lei Huang</i>	26783
DATS: Difficulty-Aware Task Sampler for Meta-Learning Physics-Informed Neural Networks <i>Maryam Toloubidokhti, Yubo Ye, Ryan Missel, Xiajun Jiang, Nilesch Kumar, Ruby Shrestha, Linwei Wang</i>	26811
Ghost on the Shell: An Expressive Representation of General 3D Shapes <i>Zhen Liu, Yao Feng, Yuliang Xiu, Weiyang Liu, Liam Paull, Michael J. Black, Bernhard Schölkopf</i>	26832
Dynamic Discounted Counterfactual Regret Minimization <i>Hang Xu, Kai Li, Haobo Fu, Qiang Fu, Junliang Xing, Jian Cheng</i>	26862
SemiReward: A General Reward Model for Semi-Supervised Learning <i>Siyuan Li, Weiyang Jin, Zedong Wang, Fang Wu, Zicheng Liu, Cheng Tan, Stan Z. Li</i>	26880

Future Language Modeling from Temporal Document History Changmao Li, Jeffrey Flanigan	26905
Test-Time Adaptation with CLIP Reward for Zero-Shot Generalization in Vision-Language Models Shuai Zhao, Xiaohan Wang, Linchao Zhu, Yi Yang	26917
Identifiable Latent Polynomial Causal Models Through the Lens of Change Yuhang Liu, Zhen Zhang, Dong Gong, Mingming Gong, Biwei Huang, Anton Van Den Hengel, Kun Zhang, Javen Qinfeng Shi	26934
Federated Recommendation with Additive Personalization Zhiwei Li, Guodong Long, Tianyi Zhou	26956
Achieving Human Parity in Content-Grounded Datasets Generation Asaf Yehudai, Boaz Carmeli, Yosi Mass, Ofir Arviv, Nathaniel Mills, Eyal Shnarch, Leshem Choshen	26974
Teaching Large Language Models to Self-Debug Xinyun Chen, Maxwell Lin, Nathanael Schärli, Denny Zhou	26990
Flow Matching on General Geometries Ricky T. Q. Chen, Yaron Lipman	27070
A Versatile Causal Discovery Framework to Allow Causally-Related Hidden Variables Xinshuai Dong, Biwei Huang, Ignavier Ng, Xiangchen Song, Yujia Zheng, Songyao Jin, Roberto Legaspi, Peter Spirtes, Kun Zhang	27094
Magnitude Invariant Parametrizations Improve Hypernetwork Learning Jose Javier Gonzalez Ortiz, John Guttag, Adrian V Dalca	27129
Fusing Models with Complementary Expertise Hongyi Wang, Felipe Maia Polo, Yuekai Sun, Souvik Kundu, Eric Xing, Mikhail Yurochkin	27149
Improving Convergence and Generalization Using Parameter Symmetries Bo Zhao, Robert M. Gower, Robin Walters, Rose Yu	27172
SILO Language Models: Isolating Legal Risk in a Nonparametric Datastore Sewon Min, Suchin Gururangan, Eric Wallace, Weijia Shi, Hannaneh Hajishirzi, Noah A. Smith, Luke Zettlemoyer	27200
KW-Design: Pushing the Limit of Protein Design Via Knowledge Refinement Zhangyang Gao, Cheng Tan, Xingran Chen, Yijie Zhang, Jun Xia, Siyuan Li, Stan Z. Li	27228
RETSim: Resilient and Efficient Text Similarity Marina Zhang, Owen Skipper Vallis, Aysegul Bumin, Tanay Vakharia, Elie Bursztein	27247
ControlVideo: Training-Free Controllable Text-To-Video Generation Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, Qi Tian	27266

SALMON: Self-Alignment with Instructable Reward Models Zhiqing Sun, Yikang Shen, Hongxin Zhang, Qinzhong Zhou, Zhenfang Chen, David Daniel Cox, Yiming Yang, Chuang Gan	27287
More is Better: When Infinite Overparameterization is Optimal and Overfitting is Obligatory James B Simon, Dhruva Karkada, Nikhil Ghosh, Mikhail Belkin	27329
Towards Non-Asymptotic Convergence for Diffusion-Based Generative Models Gen Li, Yuting Wei, Yuxin Chen, Yuejie Chi	27369
COCO-Periph: Bridging the Gap Between Human and Machine Perception in the Periphery Anne Harrington, Vasha Dutell, Mark Hamilton, Ayush Tewari, Simon Stent, William T. Freeman, Ruth Rosenholtz	27381
Towards Foundation Models for Knowledge Graph Reasoning Mikhail Galkin, Xinyu Yuan, Hesham Mostafa, Jian Tang, Zhaocheng Zhu	27411
Exploring Effective Stimulus Encoding Via Vision System Modeling for Visual Prostheses Chuanqing Wang, Di Wu, Chaoming Fang, Jie Yang, Mohamad Sawan	27433
Adaptive Self-Training Framework for Fine-Grained Scene Graph Generation Kibum Kim, Kanghoon Yoon, Yeonjun In, Jinyoung Moon, Donghyun Kim, Chanyoung Park	27446
Masked Distillation Advances Self-Supervised Transformer Architecture Search Caixia Yan, Xiaojun Chang, Zhihui Li, Lina Yao, Minnan Luo, Qinghua Zheng	27471
Privacy-Preserving In-Context Learning for Large Language Models Tong Wu, Ashwinee Panda, Jiachen T. Wang, Prateek Mittal	27484
Latent Intuitive Physics: Learning to Transfer Hidden Physics from a 3D Video Xiangming Zhu, Huayu Deng, Haochen Yuan, Yunbo Wang, Xiaokang Yang	27520
Neural Common Neighbor with Completion for Link Prediction Xiyuan Wang, Haotong Yang, Muhan Zhang	27546
Manifold Preserving Guided Diffusion Yutong He, Naoki Murata, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Dongjun Kim, Wei-Hsiang Liao, Yuki Mitsufuji, J Zico Kolter, Ruslan Salakhutdinov, Stefano Ermon	27563
Graph Generation with K^2-Trees Yunhui Jang, Dongwoo Kim, Sungsoo Ahn	27595
Zero and Few-Shot Semantic Parsing with Ambiguous Inputs Elias Stengel-Eskin, Kyle Rawlins, Benjamin Van Durme	27619
Geometrically Aligned Transfer Encoder for Inductive Transfer in Regression Tasks Sung Moon Ko, Sumin Lee, Dae-Woong Jeong, Woohyung Lim, Sehui Han	27637

Sample-Efficient Quality-Diversity by Cooperative Coevolution Ke Xue, Ren-Jian Wang, Pengyi Li, Dong Li, Jianye Hao, Chao Qian	27660
CRITIC: Large Language Models Can Self-Correct with Tool-Interactive Critiquing Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, Weizhu Chen	27686
Large Language Models as Tool Makers Tianle Cai, Xuezhi Wang, Tengyu Ma, Xinyun Chen, Denny Zhou	27764
Memory-Assisted Sub-Prototype Mining for Universal Domain Adaptation Yuxiang Lai, Yi Zhou, Xinghong Liu, Tao Zhou	27787
RDesign: Hierarchical Data-Efficient Representation Learning for Tertiary Structure-Based RNA Design Cheng Tan, Yijie Zhang, Zhangyang Gao, Bozhen Hu, Siyuan Li, Zicheng Liu, Stan Z. Li	27802
H2O-SDF: Two-Phase Learning for 3D Indoor Reconstruction Using Object Surface Fields Minyoung Park, Mirae Do, Yeon Jae Shin, Jaeseok Yoo, Jongkwang Hong, Joongrock Kim, Chul Lee	27832
Embarrassingly Simple Dataset Distillation Yunzhen Feng, Shanmukha Ramakrishna Vedantam, Julia Kempe	27852
Offline Data Enhanced On-Policy Policy Gradient with Provable Guarantees Yifei Zhou, Ayush Sekhari, Yuda Song, Wen Sun	27876
BayesPrompt: Prompting Large-Scale Pre-Trained Language Models on Few-Shot Inference Via Debiased Domain Abstraction Jiangmeng Li, Fei Song, Yifan Jin, Wenwen Qiang, Changwen Zheng, Fuchun Sun, Hui Xiong	27912
Time-LLM: Time Series Forecasting by Reprogramming Large Language Models Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y. Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, Qingsong Wen	27928
Don't Judge by the Look: Towards Motion Coherent Video Representation Yitian Zhang, Yue Bai, Huan Wang, Yizhou Wang, Yun Fu	27952
Annealing Self-Distillation Rectification Improves Adversarial Training Yu-Yu Wu, Hung-Jui Wang, Shang-Tse Chen	27968
UniAdapter: Unified Parameter-Efficient Transfer Learning for Cross-Modal Modeling Haoyu Lu, Yuqi Huo, Guoxing Yang, Zhiwu Lu, Wei Zhan, Masayoshi Tomizuka, Mingyu Ding	27990
Large Language Models Are Efficient Learners of Noise-Robust Speech Recognition Yuchen Hu, Chen Chen, Chao-Han Huck Yang, Ruizhe Li, Chao Zhang, Pin-Yu Chen, Engsiong Chng	28010
Multi-Modal Gaussian Process Variational Autoencoders for Neural and Behavioral Data Rabia Gondur, Usama Bin Sikandar, Evan Schaffer, Mikio Christian Aoi, Stephen L Keeley	28035

Large Language Models as Analogical Reasoners Michihiro Yasunaga, Xinyun Chen, Yujia Li, Panupong Pasupat, Jure Leskovec, Percy Liang, Ed H. Chi, Denny Zhou	28060
Kernel Metric Learning for In-Sample Off-Policy Evaluation of Deterministic RL Policies Haanvid Lee, Tri Wahyu Guntara, Jongmin Lee, Yung-Kyun Noh, Kee-Eung Kim	28087
A Study of Bayesian Neural Network Surrogates for Bayesian Optimization Yucen Lily Li, Tim G. J. Rudner, Andrew Gordon Wilson	28110
Long-Short-Range Message-Passing: A Physics-Informed Framework to Capture Non-Local Interaction for Scalable Molecular Dynamics Simulation Yunyang Li, Yusong Wang, Lin Huang, Han Yang, Xinran Wei, Jia Zhang, Tong Wang, Zun Wang, Bin Shao, Tie-Yan Liu	28149
Exploring the Cloud of Feature Interaction Scores in a Rashomon Set Sichao Li, Rong Wang, Quanling Deng, Amanda S Barnard	28189
L2P-MIP: Learning to Presolve for Mixed Integer Programming Chang Liu, Zhichen Dong, Haobo Ma, Weilin Luo, Xijun Li, Bowen Pang, Jia Zeng, Junchi Yan	28212
TAIL: Task-Specific Adapters for Imitation Learning with Large Pretrained Models Zuxin Liu, Jesse Zhang, Kavosh Asadi, Yao Liu, Ding Zhao, Shoham Sabach, Rasool Fakoor	28234
Inverse Approximation Theory for Nonlinear Recurrent Neural Networks Shida Wang, Zhong Li, Qianxiao Li	28258
Meta Continual Learning Revisited: Implicitly Enhancing Online Hessian Approximation Via Variance Reduction Yichen Wu, Long-Kai Huang, Renzhen Wang, Deyu Meng, Ying Wei	28291
FedWon: Triumphant Multi-Domain Federated Learning Without Normalization Weiming Zhuang, Lingjuan Lyu	28321
MOFDiff: Coarse-Grained Diffusion for Metal-Organic Framework Design Xiang Fu, Tian Xie, Andrew Scott Rosen, Tommi S. Jaakkola, Jake Allen Smith	28346
Uncertainty Quantification Via Stable Distribution Propagation Felix Petersen, Aashwin Ananda Mishra, Hilde Kuehne, Christian Borgelt, Oliver Deussen, Mikhail Yurochkin	28369
Bidirectional Temporal Diffusion Model for Temporally Consistent Human Animation Tserendorj Adiya, Jae Shin Yoon, Jungeun Lee, Sanghun Kim, Hwasup Lim	28396
Meta-Learning Priors Using Unrolled Proximal Networks Yilang Zhang, Georgios B. Giannakis	28417

SpikePoint: An Efficient Point-Based Spiking Neural Network for Event Cameras Action Recognition Hongwei Ren, Yue Zhou, Xiaopeng Lin, Yulong Huang, Haotian Fu, Jie Song, Bojun Cheng	28442
Prompt Learning with Quaternion Networks Boya Shi, Zhengqin Xu, Shuai Jia, Chao Ma	28462
ImplicitSLIM and How it Improves Embedding-Based Collaborative Filtering Ilya Shenbin, Sergey Nikolenko	28474
Revisit and Outstrip Entity Alignment: A Perspective of Generative Models Lingbing Guo, Zhuo Chen, Jiaoyan Chen, Yin Fang, Wen Zhang, Huajun Chen	28496
Sample-Efficient Learning of Infinite-Horizon Average-Reward MDPs with General Function Approximation Jianliang He, Han Zhong, Zhuoran Yang	28514
InfoCon: Concept Discovery with Generative and Discriminative Informativeness Ruizhe Liu, Qian Luo, Yanchao Yang	28549
Robot Fleet Learning Via Policy Merging Lirui Wang, Kaiqing Zhang, Allan Zhou, Max Simchowitz, Russ Tedrake	28576
Teach LLMs to Phish: Stealing Private Information from Language Models Ashwinee Panda, Christopher A. Choquette-Choo, Zhengming Zhang, Yaoqing Yang, Prateek Mittal	28611
Elucidating the Design Space of Classifier-Guided Diffusion Generation Jiajun Ma, Tianyang Hu, Wenjia Wang, Jiacheng Sun	28636
Scaling Supervised Local Learning with Augmented Auxiliary Networks Chenxiang Ma, Jibin Wu, Chenyang Si, Kc Tan	28648
Ins-DetCLIP: Aligning Detection Model to Follow Human-Language Instruction Renjie Pi, Lewei Yao, Jianhua Han, Xiaodan Liang, Wei Zhang, Hang Xu	28666
A Paradigm Shift in Machine Translation: Boosting Translation Performance of Large Language Models Haoran Xu, Young Jin Kim, Amr Sharaf, Hany Hassan Awadalla	28684
Consistent Multi-Class Classification from Multiple Unlabeled Datasets Zixi Wei, Senlin Shu, Yuzhou Cao, Hongxin Wei, Bo An, Lei Feng	28705
Revisiting Plasticity in Visual Reinforcement Learning: Data, Modules and Training Stages Guozheng Ma, Lu Li, Sen Zhang, Zixuan Liu, Zhen Wang, Yixin Chen, Li Shen, Xueqian Wang, Dacheng Tao	28725
Unpaired Image-To-Image Translation Via Neural Schrödinger Bridge Beomsu Kim, Gihyun Kwon, Kwanyoung Kim, Jong Chul Ye	28750

Large Language Model Cascades with Mixture of Thought Representations for Cost-Efficient Reasoning Murong Yue, Jie Zhao, Min Zhang, Liang Du, Ziyu Yao	28770
Instructive Decoding: Instruction-Tuned Large Language Models Are Self-Refiner from Noisy Instructions Taehyeon Kim, Joonkee Kim, Gihun Lee, Se-Young Yun	28808
On the Effect of Batch Size in Byzantine-Robust Distributed Learning Yi-Rui Yang, Chang-Wei Shi, Wu-Jun Li	28836
Feature-Aligned N-BEATS with Sinkhorn Divergence Joonhun Lee, Myeongho Jeon, Myungjoo Kang, Kyunghyun Park	28871
Query-Policy Misalignment in Preference-Based Reinforcement Learning Xiao Hu, Jianxiong Li, Xianyuan Zhan, Qing-Shan Jia, Ya-Qin Zhang	28899
Synergistic Patch Pruning for Vision Transformer: Unifying Intra- & Inter-Layer Patch Importance Yuyao Zhang, Lan Wei, Nikolaos Freris	28924
MetaMath: Bootstrap Your Own Mathematical Questions for Large Language Models Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James Kwok, Zhenguo Li, Adrian Weller, Weiyang Liu	28943
From Graphs to Hypergraphs: Hypergraph Projection and Its Reconstruction Yanbang Wang, Jon Kleinberg	28965
Neural Active Learning Beyond Bandits Yikun Ban, Ishika Agarwal, Ziwei Wu, Yada Zhu, Kommy Weldemariam, Hanghang Tong, Jingrui He	28996
Shadow Cones: A Generalized Framework for Partial Order Embeddings Tao Yu, Toni J. B. Liu, Albert Tseng, Christopher De Sa	29036
Rethinking Complex Queries on Knowledge Graphs with Neural Link Predictors Hang Yin, Zihao Wang, Yangqiu Song	29057
Sheared LLaMA: Accelerating Language Model Pre-Training Via Structured Pruning Mengzhou Xia, Tianyu Gao, Zhiyuan Zeng, Danqi Chen	29082
SPTNet: An Efficient Alternative Framework for Generalized Category Discovery with Spatial Prompt Tuning Hongjun Wang, Sagar Vaze, Kai Han	29107
FITS: Modeling Time Series with \$10k\$ Parameters Zhijian Xu, Ailing Zeng, Qiang Xu	29143

LRR: Language-Driven Resamplable Continuous Representation Against Adversarial Tracking Attacks Jianlang Chen, Xuhong Ren, Qing Guo, Felix Juefei-Xu, Di Lin, Wei Feng, Lei Ma, Jianjun Zhao	29167
DQ-LoRe: Dual Queries with Low Rank Approximation Re-Ranking for In-Context Learning Jing Xiong, Zixuan Li, Chuanyang Zheng, Zhijiang Guo, Yichun Yin, Enze Xie, Zhicheng Yang, Qingxing Cao, Haiming Wang, Xiongwei Han, Jing Tang, Chengming Li, Xiaodan Liang	29180
Boosting the Adversarial Robustness of Graph Neural Networks: An OOD Perspective Kuan Li, Yiwen Chen, Yang Liu, Jin Wang, Qing He, Minhao Cheng, Xiang Ao	29205
DistillSpec: Improving Speculative Decoding Via Knowledge Distillation Yongchao Zhou, Kaifeng Lyu, Ankit Singh Rawat, Aditya Krishna Menon, Afshin Rostamizadeh, Sanjiv Kumar, Jean-François Kagy, Rishabh Agarwal	29230
Coeditor: Leveraging Repo-Level Diffs for Code Auto-Editing Jiayi Wei, Greg Durrett, Isil Dillig	29270
Towards Robust Multi-Modal Reasoning Via Model Selection Xiangyan Liu, Rongxue Li, Wei Ji, Tao Lin	29293
CodeChain: Towards Modular Code Generation Through Chain of Self-Revisions with Representative Sub-Modules Hung Le, Hailin Chen, Amrita Saha, Akash Gokul, Doyen Sahoo, Shafiq Joty	29315
The Blessing of Randomness: SDE Beats ODE in General Diffusion-Based Image Editing Shen Nie, Hanzhong Allan Guo, Cheng Lu, Yuhao Zhou, Chenyu Zheng, Chongxuan Li	29346
Identifying the Risks of LM Agents with an LM-Emulated Sandbox Yangjun Ruan, Honghua Dong, Andrew Wang, Silviu Pitis, Yongchao Zhou, Jimmy Ba, Yann Dubois, Chris J. Maddison, Tatsunori Hashimoto	29383
On the Parameterization of Second-Order Optimization Effective Towards the Infinite Width Satoki Ishikawa, Ryo Karakida	29451
Neural Auto-Designer for Enhanced Quantum Kernels Cong Lei, Yuxuan Du, Peng Mi, Jun Yu, Tongliang Liu	29485
One-Hot Generalized Linear Model for Switching Brain State Discovery Chengrui Li, Soon Ho Kim, Chris Rodgers, Hannah Choi, Anqi Wu	29510
Doubly Robust Proximal Causal Learning for Continuous Treatments Yong Wu, Yanwei Fu, Shouyan Wang, Xinwei Sun	29532
AlignDiff: Aligning Diverse Human Preferences Via Behavior-Customisable Diffusion Model Zibin Dong, Yifu Yuan, Jianye Hao, Fei Ni, Yao Mu, Yan Zheng, Yujing Hu, Tangjie Lv, Changjie Fan, Zhipeng Hu	29576

Active Test-Time Adaptation: Theoretical Analyses and an Algorithm Shurui Gui, Xiner Li, Shuiwang Ji	29604
A Lightweight Method for Tackling Unknown Participation Statistics in Federated Averaging Shiqiang Wang, Mingyue Ji	29653
Diffusion Sampling with Momentum for Mitigating Divergence Artifacts Suttisak Wizadwongsa, Worameth Chinchuthakun, Pramook Khungurn, Amit Raj, Supasorn Suwajanakorn	29693
Multiscale Positive-Unlabeled Detection of AI-Generated Texts Yuchuan Tian, Hanting Chen, Xutao Wang, Zheyuan Bai, Qinghua Zhang, Ruifeng Li, Chao Xu, Yunhe Wang	29739
Generative Learning for Financial Time Series with Irregular and Scale-Invariant Patterns Hongbin Huang, Minghua Chen, Xiao Qiao	29757
Enhancing Group Fairness in Online Settings Using Oblique Decision Forests Somnath Basu Roy Chowdhury, Nicholas Monath, Ahmad Beirami, Rahul Kidambi, Kumar Avinava Dubey, Amr Ahmed, Snigdha Chaturvedi	29778
Causal Inference with Conditional Front-Door Adjustment and Identifiable Variational Autoencoder Ziqi Xu, Debo Cheng, Jiuyong Li, Jixue Liu, Lin Liu, Kui Yu	29810
WildFusion: Learning 3D-Aware Latent Diffusion Models in View Space Katja Schwarz, Seung Wook Kim, Jun Gao, Sanja Fidler, Andreas Geiger, Karsten Kreis	29832
Sparse Weight Averaging with Multiple Particles for Iterative Magnitude Pruning Moonseok Choi, Hyungi Lee, Giung Nam, Juho Lee	29867
Connecting Large Language Models with Evolutionary Algorithms Yields Powerful Prompt Optimizers Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, Yujiu Yang	29890
WebArena: A Realistic Web Environment for Building Autonomous Agents Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, Graham Neubig	29914
Diffusion Models for Multi-Task Generative Modeling Changyou Chen, Han Ding, Bunyamin Sisman, Yi Xu, Ouye Xie, Benjamin Z. Yao, Son Dinh Tran, Belinda Zeng	29936
Free from Bellman Completeness: Trajectory Stitching Via Model-Based Return-Conditioned Supervised Learning Zhaoyi Zhou, Chuning Zhu, Runlong Zhou, Qiwen Cui, Abhishek Gupta, Simon Shaolei Du	29966

A Unified Framework for Bayesian Optimization Under Contextual Uncertainty Sebastian Shenghong Tay, Chuan-Sheng Foo, Daisuke Urano, Richalynn Leong, Bryan Kian Hsiang Low	29991
Compose and Conquer: Diffusion-Based 3D Depth Aware Composable Image Synthesis Jonghyun Lee, Hansam Cho, Youngjoon Yoo, Seoung Bum Kim, Yonghyun Jeong	30024
FedImpro: Measuring and Improving Client Update in Federated Learning Zhenheng Tang, Yonggang Zhang, Shaohuai Shi, Xinmei Tian, Tongliang Liu, Bo Han, Xiaowen Chu	30047
DNA-GPT: Divergent N-Gram Analysis for Training-Free Detection of GPT-Generated Text Xianjun Yang, Wei Cheng, Yue Wu, Linda Ruth Petzold, William Yang Wang, Haifeng Chen	30077
Provably Robust Conformal Prediction with Improved Efficiency Ge Yan, Yaniv Romano, Tsui-Wei Weng	30103
Learning Over Molecular Conformer Ensembles: Datasets and Benchmarks Yanqiao Zhu, Jeehyun Hwang, Keir Adams, Zhen Liu, Bozhao Nan, Brock Stenfors, Yuanqi Du, Jatin Chauhan, Olaf Wiest, Olexandr Isayev, Connor W. Coley, Yizhou Sun, Wei Wang	30140
Efficient Integrators for Diffusion Generative Models Kushagra Pandey, Maja Rudolph, Stephan Mandt	30159
Soft Contrastive Learning for Time Series Seunghan Lee, Taeyoung Park, Kibok Lee	30208
MOFI: Learning Image Representations from Noisy Entity Annotated Images Wentao Wu, Aleksei Timofeev, Chen Chen, Bowen Zhang, Kun Duan, Shuangning Liu, Yantao Zheng, Jonathon Shlens, Xianzhi Du, Yinfei Yang	30233
Graphical Multioutput Gaussian Process with Attention Yijue Dai, Wenzhong Yan, Feng Yin	30248
On the Vulnerability of Adversarially Trained Models Against Two-Faced Attacks Shengjie Zhou, Lue Tao, Yuzhou Cao, Tao Xiang, Bo An, Lei Feng	30271
METRA: Scalable Unsupervised RL with Metric-Aware Abstraction Seohong Park, Oleh Rybkin, Sergey Levine	30289
On the Hardness of Online Nonconvex Optimization with Single Oracle Feedback Ziwei Guan, Yi Zhou, Yingbin Liang	30314
Transformers as Decision Makers: Provable In-Context Reinforcement Learning Via Supervised Pretraining Licong Lin, Yu Bai, Song Mei	30338
Temporal Generalization Estimation in Evolving Graphs Bin Lu, Tingyan Ma, Xiaoying Gan, Xinbing Wang, Yunqiang Zhu, Chenghu Zhou, Shiyu Liang	30396

Knowledge Card: Filling LLMs' Knowledge Gaps with Plug-In Specialized Language Models Shangbin Feng, Weijia Shi, Yuyang Bai, Vidhisha Balachandran, Tianxing He, Yulia Tsvetkov	30422
Exploring the Common Appearance-Boundary Adaptation for Nighttime Optical Flow Hanyu Zhou, Yi Chang, Haoyue Liu, Yan Wending, Yuxing Duan, Zhiwei Shi, Luxin Yan	30447
Mayfly: A Neural Data Structure for Graph Stream Summarization Yuan Feng, Yukun Cao, Wang Hairu, Xike Xie, S Kevin Zhou	30458
HAZARD Challenge: Embodied Decision Making in Dynamically Changing Environments Qinhong Zhou, Sunli Chen, Yisong Wang, Haozhe Xu, Weihua Du, Hongxin Zhang, Yilun Du, Joshua B. Tenenbaum, Chuang Gan	30476
Learning from Aggregate Responses: Instance Level Versus Bag Level Loss Functions Adel Javanmard, Lin Chen, Vahab Mirrokni, Ashwinkumar Badanidiyuru, Gang Fu	30493
Don't Play Favorites: Minority Guidance for Diffusion Models Soobin Um, Suhyeon Lee, Jong Chul Ye	30514
Information Retention Via Learning Supplemental Features Zhipeng Xie, Yahe Li	30553
Solving Diffusion ODEs with Optimal Boundary Conditions for Better Image Super-Resolution Yiyang Ma, Huan Yang, Wenhan Yang, Jianlong Fu, Jiaying Liu	30570
Mol-Instructions: A Large-Scale Biomolecular Instruction Dataset for Large Language Models Yin Fang, Xiaozhuan Liang, Ningyu Zhang, Kangwei Liu, Rui Huang, Zhuo Chen, Xiaohui Fan, Huajun Chen	30589
An Efficient Membership Inference Attack for the Diffusion Model by Proximal Initialization Fei Kong, Jinhao Duan, Ruipeng Ma, Heng Tao Shen, Xiaoshuang Shi, Xiaofeng Zhu, Kaidi Xu	30620
LanguageBind: Extending Video-Language Pretraining to N-Modality by Language-Based Semantic Alignment Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiayi Cui, Wang Hongfa, Yatian Pang, Wenhao Jiang, Junwu Zhang, Zongwei Li, Cai Wan Zhang, Zhifeng Li, Wei Liu, Li Yuan	30641
Video Language Planning Yilun Du, Sherry Yang, Pete Florence, Fei Xia, Ayzaan Wahid, Brian Ichter, Pierre Sermanet, Tianhe Yu, Pieter Abbeel, Joshua B. Tenenbaum, Leslie Pack Kaelbling, Andy Zeng, Jonathan Tompson	30662
Large Language Models as Generalizable Policies for Embodied Tasks Andrew Szot, Max Schwarzer, Harsh Agrawal, Bogdan Mazouze, Rin Metcalf, Walter Talbott, Natalie Mackraz, R Devon Hjelm, Alexander T Toshev	30680
Alice Benchmarks: Connecting Real World Re-Identification with the Synthetic Xiaoxiao Sun, Yue Yao, Shengjin Wang, Hongdong Li, Liang Zheng	30713

Sample-Efficient Linear Representation Learning from Non-IID Non-Isotropic Data Thomas Tck Zhang, Leonardo Felipe Toso, James Anderson, Nikolai Matni	30732
Combinatorial Bandits for Maximum Value Reward Function Under Value-Index Feedback Yiliu Wang, Wei Chen, Milan Vojnovic	30744
EmerNeRF: Emergent Spatial-Temporal Scene Decomposition Via Self-Supervision Jiawei Yang, Boris Ivanovic, Or Litany, Xinshuo Weng, Seung Wook Kim, Boyi Li, Tong Che, Danfei Xu, Sanja Fidler, Marco Pavone, Yue Wang	30766
How Do Transformers Learn In-Context Beyond Simple Functions? a Case Study on Learning with Representations Tianyu Guo, Wei Hu, Song Mei, Huan Wang, Caiming Xiong, Silvio Savarese, Yu Bai	30794
AmortizedPeriod: Attention-Based Amortized Inference for Periodicity Identification Hang Yu, Cong Liao, Ruolan Liu, Jianguo Li, Hu Yun, Xinzhe Wang	30823
The False Promise of Imitating Proprietary Language Models Arnav Gudibande, Eric Wallace, Charlie Victor Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, Dawn Song	30853
MaGIC: Multi-Modality Guided Image Completion Hao Wang, Yongsheng Yu, Tiejian Luo, Heng Fan, Libo Zhang	30871
Solving Inverse Problems with Latent Diffusion Models Via Hard Data Consistency Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, Liyue Shen	30896
Fixed Non-Negative Orthogonal Classifier: Inducing Zero-Mean Neural Collapse with Feature Dimension Separation Hoyong Kim, Kangil Kim	30927
Making Pre-Trained Language Models Great on Tabular Prediction Jiahuan Yan, Bo Zheng, Hongxia Xu, Yiheng Zhu, Danny Chen, Jimeng Sun, Jian Wu, Jintai Chen	30956
Unsupervised Order Learning Seon-Ho Lee, Nyeong-Ho Shin, Chang-Su Kim	30980
Decoupled Marked Temporal Point Process Using Neural Ordinary Differential Equations Yujee Song, Donghyun Lee, Rui Meng, Won Hwa Kim	31000
VQ-TR: Vector Quantized Attention for Time Series Forecasting Kashif Rasul, Andrew Bennett, Pablo Vicente, Umang Gupta, Hena Ghonia, Anderson Schneider, Yuriy Nevmyvaka	31018
Transferring Learning Trajectories of Neural Networks Daiki Chijiwa	31040

Foundation Model-Oriented Robustness: Robust Image Model Evaluation with Pretrained Models <i>Peiyan Zhang, Haoyang Liu, Chaozhuo Li, Xing Xie, Sunghun Kim, Haohan Wang</i>	31066
Stochastic Modified Equations and Dynamics of Dropout Algorithm <i>Zhongwang Zhang, Yuqing Li, Tao Luo, Zhi-Qin John Xu</i>	31095
\$\infty\$-Diff: Infinite Resolution Diffusion with Subsampled Mollified States <i>Sam Bond-Taylor, Chris G. Willcocks</i>	31151
Neuron-Enhanced AutoEncoder Matrix Completion and Collaborative Filtering: Theory and Practice <i>Jicong Fan, Rui Chen, Zhao Zhang, Chris Ding</i>	31173
Submodular Reinforcement Learning <i>Manish Prajapat, Mojmir Mutny, Melanie Zeilinger, Andreas Krause</i>	31198
A Representation-Learning Game for Classes of Prediction Tasks <i>Neria Uzan, Nir Weinberger</i>	31221
A Unified Sampling Framework for Solver Searching of Diffusion Probabilistic Models <i>Enshu Liu, Xuefei Ning, Huazhong Yang, Yu Wang</i>	31264
Adaptive Window Pruning for Efficient Local Motion Deblurring <i>Haoying Li, Jixin Zhao, Shangchen Zhou, Huajun Feng, Chongyi Li, Chen Change Loy</i>	31303
ConjNorm: Tractable Density Estimation for Out-Of-Distribution Detection <i>Bo Peng, Yadan Luo, Yonggang Zhang, Yixuan Li, Zhen Fang</i>	31323
EasyTPP: Towards Open Benchmarking Temporal Point Processes <i>Siqiao Xue, Xiaoming Shi, Zhixuan Chu, Yan Wang, Hongyan Hao, Fan Zhou, Caigao Jiang, Chen Pan, James Y. Zhang, Qingsong Wen, Jun Zhou, Hongyuan Mei</i>	31345
Label-Noise Robust Diffusion Models <i>Byeonghu Na, Yeongmin Kim, Heesun Bae, Jung Hyun Lee, Se Jung Kwon, Wanmo Kang, Il-Chul Moon</i>	31367
Evoke: Evoking Critical Thinking Abilities in LLMs Via Reviewer-Author Prompt Editing <i>Xinyu Hu, Pengfei Tang, Simiao Zuo, Zihan Wang, Bowen Song, Qiang Lou, Jian Jiao, Denis X Charles</i>	31411
MMICL: Empowering Vision-Language Model with Multi-Modal In-Context Learning <i>Haozhe Zhao, Zefan Cai, Shuzheng Si, Xiaojian Ma, Kaikai An, Liang Chen, Zixuan Liu, Sheng Wang, Wenjuan Han, Baobao Chang</i>	31432
NaturalSpeech 2: Latent Diffusion Models Are Natural and Zero-Shot Speech and Singing Synthesizers <i>Kai Shen, Zeqian Ju, Xu Tan, Eric Liu, Yichong Leng, Lei He, Tao Qin, Sheng Zhao, Jiang Bian</i>	31471

Masked Structural Growth for 2x Faster Language Model Pre-Training Yiqun Yao, Zheng Zhang, Jing Li, Yequan Wang	31496
Learning Flexible Body Collision Dynamics with Hierarchical Contact Mesh Transformer Youn-Yeol Yu, Jeongwhan Choi, Woojin Cho, Kookjin Lee, Nayong Kim, Kiseok Chang, Changseung Woo, Ilho Kim, Seokwoo Lee, Joon Young Yang, Sooyoung Yoon, Noseong Park	31514
RT-Trajectory: Robotic Task Generalization Via Hindsight Trajectory Sketches Jiayuan Gu, Sean Kirmani, Paul Wohlhart, Yao Lu, Montserrat Gonzalez Arenas, Kanishka Rao, Wenhao Yu, Chuyuan Fu, Keerthana Gopalakrishnan, Zhuo Xu, Priya Sundaresan, Peng Xu, Hao Su, Karol Hausman, Chelsea Finn, Quan Vuong, Ted Xiao	31542
Invariance-Based Learning of Latent Dynamics Kai Lagemann, Christian Lagemann, Sach Mukherjee	31566
SWAP-NAS: Sample-Wise Activation Patterns for Ultra-Fast NAS Yameng Peng, Andy Song, Haytham M. Fayek, Vic Ciesielski, Xiaojun Chang	31611
FFB: A Fair Fairness Benchmark for In-Processing Group Fairness Methods Xiaotian Han, Jianfeng Chi, Yu Chen, Qifan Wang, Han Zhao, Na Zou, Xia Hu	31637
LLM Blueprint: Enabling Text-To-Image Generation with Complex and Detailed Prompts Hanan Gani, Shariq Farooq Bhat, Muzammal Naseer, Salman Khan, Peter Wonka	31663
Deep Geodesic Canonical Correlation Analysis for Covariance-Based Neuroimaging Data Ce Ju, Reinmar J Kobler, Liyao Tang, Cuntai Guan, Motoaki Kawanabe	31684
Finetuning Text-To-Image Diffusion Models for Fairness Xudong Shen, Chao Du, Tianyu Pang, Min Lin, Yongkang Wong, Mohan Kankanhalli	31720
Bounding Box Stability Against Feature Dropout Reflects Detector Generalization Across Environments Yang Yang, Wenhai Wang, Zhe Chen, Jifeng Dai, Liang Zheng	31763
Dictionary Contrastive Learning for Efficient Local Supervision Without Auxiliary Networks Suhwan Choi, Myeongho Jeon, Yeonjung Hwang, Jeonglyul Oh, Sungjun Lim, Joonseok Lee, Myungjoo Kang	31780
Uni-O4: Unifying Online and Offline Deep Reinforcement Learning with Multi-Step On-Policy Optimization Kun Lei, Zhengmao He, Chenhao Lu, Kaizhe Hu, Yang Gao, Huazhe Xu	31806
InternVid: A Large-Scale Video-Text Dataset for Multimodal Understanding and Generation Yi Wang, Yinan He, Yizhuo Li, Kunchang Li, Jiashuo Yu, Xin Ma, Xinhao Li, Guo Chen, Xinyuan Chen, Yaohui Wang, Ping Luo, Ziwei Liu, Yali Wang, Limin Wang, Yu Qiao	31840
Adversarial Training on Purification (AToP): Advancing Both Robustness and Generalization Guang Lin, Chao Li, Jianhai Zhang, Toshihisa Tanaka, Qibin Zhao	31865

Benchmarking and Improving Generator-Validator Consistency of Language Models <i>Xiang Lisa Li, Vaishnavi Shrivastava, Siyan Li, Tatsunori Hashimoto, Percy Liang</i>	31880
Skip-Attention: Improving Vision Transformers by Paying Less Attention <i>Shashanka Venkataramanan, Amir Ghodrati, Yuki M Asano, Fatih Porikli, Amir Habibian</i>	31898
Turning Large Language Models into Cognitive Models <i>Marcel Binz, Eric Schulz</i>	31917
LogicMP: A Neuro-Symbolic Approach for Encoding First-Order Logic Constraints <i>Weidi Xu, Jingwei Wang, Lele Xie, Jianshan He, Hongting Zhou, Taifeng Wang, Xiaopei Wan, Jingdong Chen, Chao Qu, Wei Chu</i>	31935
On the Posterior Distribution in Denoising: Application to Uncertainty Quantification <i>Hila Manor, Tomer Michaeli</i>	31964
Mixture of LoRA Experts <i>Xun Wu, Shaohan Huang, Furu Wei</i>	31995
Towards Meta-Pruning Via Optimal Transport <i>Alexander Theus, Olin Geimer, Friedrich Wicke, Thomas Hofmann, Sotiris Anagnostidis, Sidak Pal Singh</i>	32012
Transformer Fusion with Optimal Transport <i>Moritz Imfeld, Jacopo Galdi, Marco Giordano, Thomas Hofmann, Sotiris Anagnostidis, Sidak Pal Singh</i>	32053
Provable Memory Efficient Self-Play Algorithm for Model-Free Reinforcement Learning <i>Na Li, Yuchen Jiao, Hanguan Shan, Shufeng Yan</i>	32077
Towards Assessing and Benchmarking Risk-Return Tradeoff of Off-Policy Evaluation <i>Haruka Kiyohara, Ren Kishimoto, Kosuke Kawakami, Ken Kobayashi, Kazuhide Nakata, Yuta Saito</i>	32143
Modulated Phase Diffusor: Content-Oriented Feature Synthesis for Detecting Unknown Objects <i>Aming Wu, Cheng Deng</i>	32173
The Wasserstein Believer: Learning Belief Updates for Partially Observable Environments Through Reliable Latent Space Models <i>Raphaël Avalos, Florent Delgrange, Ann Nowe, Guillermo Perez, Diederik M Roijers</i>	32191
On Double Descent in Reinforcement Learning with LSTD and Random Features <i>David Brellmann, Eloise Berthier, David Filliat, Goran Frehse</i>	32231
Measuring Vision-Language STEM Skills of Neural Models <i>Jianhao Shen, Ye Yuan, Srubhi Mirzoyan, Ming Zhang, Chenguang Wang</i>	32307
MAP it to Visualize Representations <i>Robert Jenssen</i>	32348

FedLoGe: Joint Local and Generic Federated Learning Under Long-Tailed Data Zikai Xiao, Zihan Chen, Liyinglan Liu, Yang Feng, Joey Tianyi Zhou, Jian Wu, Wanlu Liu, Howard Hao Yang, Zuozhu Liu	32383
Decoupling Weighing and Selecting for Integrating Multiple Graph Pre-Training Tasks Tianyu Fan, Lirong Wu, Yufei Huang, Haitao Lin, Cheng Tan, Zhangyang Gao, Stan Z. Li	32403
Towards Imitation Learning to Branch for MIP: A Hybrid Reinforcement Learning Based Sample Augmentation Approach Changwen Zhang, Wenli Ouyang, Hao Yuan, Liming Gong, Yong Sun, Ziao Guo, Zhichen Dong, Junchi Yan	32422
Retrieval is Accurate Generation Bowen Cao, Deng Cai, Leyang Cui, Xuxin Cheng, Wei Bi, Yuexian Zou, Shuming Shi	32437
Scalable and Effective Implicit Graph Neural Networks on Large Graphs Juncheng Liu, Bryan Hooi, Kenji Kawaguchi, Yiwei Wang, Chaosheng Dong, Xiaokui Xiao	32455
AutoChunk: Automated Activation Chunk for Memory-Efficient Deep Learning Inference Xuanlei Zhao, Shenggan Cheng, Guangyang Lu, Haotian Zhou, Bin Jia, Yang You	32470
Uncertainty-Aware Constraint Inference in Inverse Constrained Reinforcement Learning Sheng Xu, Guiliang Liu	32481
Two-Timescale Extragradient for Finding Local Minimax Points Jiseok Chae, Kyuwon Kim, Donghwan Kim	32506
Algorithms for Caching and MTS with Reduced Number of Predictions Karim Ahmed Abdel Sadek, Marek Elias	32546
Interpretable Diffusion Via Information Decomposition Xianghao Kong, Ollie Liu, Han Li, Dani Yogatama, Greg Ver Steeg	32588
Generative Learning for Solving Non-Convex Problem with Multi-Valued Input-Solution Mapping Enming Liang, Minghua Chen	32602
Constructing Adversarial Examples for Vertical Federated Learning: Optimal Client Corruption Through Multi-Armed Bandit Duanyi Yao, Songze Li, Ye Xue, Jin Liu	32623
Text2Reward: Reward Shaping with Language Models for Reinforcement Learning Tianbao Xie, Siheng Zhao, Chen Henry Wu, Yitao Liu, Qian Luo, Victor Zhong, Yanchao Yang, Tao Yu	32646
An Unforgeable Publicly Verifiable Watermark for Large Language Models Aiwei Liu, Leyi Pan, Xuming Hu, Shuang Li, Lijie Wen, Irwin King, Philip S. Yu	32683

SOInter: A Novel Deep Energy-Based Interpretation Method for Explaining Structured Output Models <i>S. Fatemeh Seyyedsalehi, Mahdieh Soleymani Baghshah, Hamid R. Rabiee</i>	32700
COSA: Concatenated Sample Pretrained Vision-Language Foundation Model <i>Sihan Chen, Xingjian He, Handong Li, Xiaojie Jin, Jiashi Feng, Jing Liu</i>	32714
Win-Win: Training High-Resolution Vision Transformers from Two Windows <i>Vincent Leroy, Jerome Revaud, Thomas Lucas, Philippe Weinzaepfel</i>	32733
BRUSLEATTACK: A QUERY-EFFICIENT SCORE- BASED BLACK-BOX SPARSE ADVERSARIAL ATTACK <i>Quoc Viet Vo, Ehsan Abbasnejad, Damith Ranasinghe</i>	32752
AttEXplore: Attribution for Explanation with Model Parameters eXploration <i>Zhiyu Zhu, Huaming Chen, Jiayu Zhang, Xinyi Wang, Zhibo Jin, Jason Xue, Flora D. Salim</i>	32790
Effective Pruning of Web-Scale Datasets Based on Complexity of Concept Clusters <i>Amro Kamal Mohamed Abbas, Evgenia Rusak, Kushal Tirumala, Wieland Brendel, Kamalika Chaudhuri, Ari S. Morcos</i>	32817
Towards 3D Molecule-Text Interpretation in Language Models <i>Sihang Li, Zhiyuan Liu, Yanchen Luo, Xiang Wang, Xiangnan He, Kenji Kawaguchi, Tat-Seng Chua, Qi Tian</i>	32840
Bongard-OpenWorld: Few-Shot Reasoning for Free-Form Visual Concepts in the Real World <i>Rujie Wu, Xiaojian Ma, Zhenliang Zhang, Wei Wang, Qing Li, Song-Chun Zhu, Yizhou Wang</i>	32860
The Need for Speed: Pruning Transformers with One Recipe <i>Samir Khaki, Konstantinos N Plataniotis</i>	32891
RLIF: Interactive Imitation Learning as Reinforcement Learning <i>Jianlan Luo, Perry Dong, Yuexiang Zhai, Yi Ma, Sergey Levine</i>	32910
An Interpretable Error Correction Method for Enhancing Code-To-Code Translation <i>Min Xue, Artur Andrzejak, Marla Leuther</i>	32933
Enhancing Transferable Adversarial Attacks on Vision Transformers Through Gradient Normalization Scaling and High-Frequency Adaptation <i>Zhiyu Zhu, Xinyi Wang, Zhibo Jin, Jiayu Zhang, Huaming Chen</i>	32951
Scaling Laws for Associative Memories <i>Vivien Cabannes, Elvis Dohmatob, Alberto Bietti</i>	32964
Fully Hyperbolic Convolutional Neural Networks for Computer Vision <i>Ahmad Bdeir, Kristian Schwethelm, Niels Landwehr</i>	32994
Intriguing Properties of Data Attribution on Diffusion Models <i>Xiaosen Zheng, Tianyu Pang, Chao Du, Jing Jiang, Min Lin</i>	33019

MIntRec2.0: A Large-Scale Benchmark Dataset for Multimodal Intent Recognition and Out-Of-Scope Detection in Conversations Hanlei Zhang, Xin Wang, Hua Xu, Qianrui Zhou, Kai Gao, Jianhua Su, Jinyue Zhao, Wenrui Li, Yanting Chen	33055
Diverse Projection Ensembles for Distributional Reinforcement Learning Moritz Akiya Zanger, Wendelin Boehmer, Matthijs T. J. Spaan	33082
Consistent4D: Consistent 360° Dynamic Object Generation from Monocular Video Yanqin Jiang, Li Zhang, Jin Gao, Weiming Hu, Yao Yao	33108
Near-Optimal Quantum Algorithm for Minimizing the Maximal Loss Hao Wang, Chenyi Zhang, Tongyang Li	33126
Noise-Free Score Distillation Oren Katzir, Or Patashnik, Daniel Cohen-Or, Dani Lischinski	33147
Scalable Modular Network: A Framework for Adaptive Learning Via Agreement Routing Minyang Hu, Hong Chang, Bingpeng Ma, Shiguang Shan, Xilin Chen	33169
ValUES: A Framework for Systematic Validation of Uncertainty Estimation in Semantic Segmentation Kim-Celine Kahl, Carsten T. L��th, Maximilian Zenk, Klaus Maier-Hein, Paul F Jaeger	33188
Sparsistency for Inverse Optimal Transport Francisco Andrade, Gabriel Peyr��, Clarice Poon	33223
Neural Contractive Dynamical Systems Hadi Beik Mohammadi, S��ren Hauberg, Georgios Arvanitidis, Nadia Figueroa, Gerhard Neumann, Leonel Rozo	33255
OPTIMAL ROBUST MEMORIZATION WITH RELU NEURAL NETWORKS Lijia Yu, Xiao-Shan Gao, Lijun Zhang	33279
Emergent Mechanisms for Long Timescales Depend on Training Curriculum and Affect Performance in Memory Tasks Sina Khajehabdollahi, Roxana Zeraati, Emmanouil Giannakakis, Tim Jakob Sch��fer, Georg Martius, Anna Levina	33326
The Generalization Gap in Offline Reinforcement Learning Ishita Mediratta, Qingfei You, Mingqi Jiang, Roberta Raileanu	33354
Unconstrained Stochastic CCA: Unifying Multiview and Self-Supervised Learning James Chapman, Lennie Wells, Ana Lawry Aguila	33398
On Accelerating Diffusion-Based Sampling Processes Via Improved Integration Approximation Guoqiang Zhang, Kenta Niwa, W. Bastiaan Kleijn	33459

Robust Training of Federated Models with Extremely Label Deficiency Yonggang Zhang, Zhiqin Yang, Xinmei Tian, Nannan Wang, Tongliang Liu, Bo Han	33483
BadEdit: Backdooring Large Language Models by Model Editing Yanzhou Li, Tianlin Li, Kangjie Chen, Jian Zhang, Shangqing Liu, Wenhan Wang, Tianwei Zhang, Yang Liu	33505
Probabilistically Rewired Message-Passing Neural Networks Chendi Qian, Andrei Manolache, Kareem Ahmed, Zhe Zeng, Guy Van Den Broeck, Mathias Niepert, Christopher Morris	33523
Adversarial Adaptive Sampling: Unify PINN and Optimal Transport for the Approximation of PDEs Kejun Tang, Jiayu Zhai, Xiaoliang Wan, Chao Yang	33549
Magic123: One Image to High-Quality 3D Object Generation Using Both 2D and 3D Diffusion Priors Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, Bernard Ghanem	33565
Negative Label Guided OOD Detection with Pretrained Vision-Language Models Xue Jiang, Feng Liu, Zhen Fang, Hong Chen, Tongliang Liu, Feng Zheng, Bo Han	33583
MAPE-PPI: Towards Effective and Efficient Protein-Protein Interaction Prediction Via Microenvironment-Aware Protein Embedding Lirong Wu, Yijun Tian, Yufei Huang, Siyuan Li, Haitao Lin, Nitesh V Chawla, Stan Z. Li	33612
Rethinking Branching on Exact Combinatorial Optimization Solver: The First Deep Symbolic Discovery Framework Yufei Kuang, Jie Wang, Haoyang Liu, Fangzhou Zhu, Xijun Li, Jia Zeng, Jianye Hao, Bin Li, Feng Wu	33628
Continual Learning in the Presence of Spurious Correlations: Analyses and a Simple Baseline Donggyu Lee, Sangwon Jung, Taesup Moon	33649
Towards Cheaper Inference in Deep Networks with Lower Bit-Width Accumulators Yaniv Blumenfeld, Itay Hubara, Daniel Soudry	33673
On Harmonizing Implicit Subpopulations Feng Hong, Jiangchao Yao, Yueming Lyu, Zhihan Zhou, Ivor Tsang, Ya Zhang, Yanfeng Wang	33690
Neural Fourier Transform: A General Approach to Equivariant Representation Learning Masanori Koyama, Kenji Fukumizu, Kohei Hayashi, Takeru Miyato	33726
The Optimality of Kernel Classifiers in Sobolev Space Jianfa Lai, Zhifan Li, Dongming Huang, Qian Lin	33758
Towards Robust Out-Of-Distribution Generalization Bounds Via Sharpness Yingtian Zou, Kenji Kawaguchi, Yingnan Liu, Jiashuo Liu, Mong-Li Lee, Wynne Hsu	33770

Training Unbiased Diffusion Models from Biased Dataset Yeongmin Kim, Byeonghu Na, Minsang Park, Joonho Jang, Dongjun Kim, Wanmo Kang, Il-Chul Moon	33810
Out-Of-Variable Generalisation for Discriminative Models Siyuan Guo, Jonas Bernhard Wildberger, Bernhard Schölkopf	33850
Bandits with Replenishable Knapsacks: The Best of Both Worlds Martino Bernasconi, Matteo Castiglioni, Andrea Celli, Federico Fusco	33863
Game-Theoretic Robust Reinforcement Learning Handles Temporally-Coupled Perturbations Yongyuan Liang, Yanchao Sun, Ruijie Zheng, Xiangyu Liu, Benjamin Eysenbach, Tuomas Sandholm, Furong Huang, Stephen Marcus McAleer	33882
Learning Multi-Faceted Prototypical User Interests Nhu-Thuat Tran, Hady W. Lauw	33902
Online Information Acquisition: Hiring Multiple Agents Federico Cacciamani, Matteo Castiglioni, Nicola Gatti	33918
A Multi-Level Framework for Accelerating Training Transformer Models Longwei Zou, Han Zhang, Yangdong Deng	33954
AGILE3D: Attention Guided Interactive Multi-Object 3D Segmentation Yuanwen Yue, Sabarinath Mahadevan, Jonas Schult, Francis Engelmann, Bastian Leibe, Konrad Schindler, Theodora Kontogianni	33978
Learning Optimal Contracts: How to Exploit Small Action Spaces Francesco Bacchiocchi, Matteo Castiglioni, Alberto Marchesi, Nicola Gatti	33999
CPPO: Continual Learning for Reinforcement Learning with Human Feedback Han Zhang, Yu Lei, Lin Gui, Min Yang, Yulan He, Hui Wang, Ruifeng Xu	34026
How to Catch an AI Liar: Lie Detection in Black-Box LLMs by Asking Unrelated Questions Lorenzo Pacchiardi, Alex James Chan, Sören Mindermann, Ilan Moscovitz, Alexa Yue Pan, Yarin Gal, Owain Evans, Jan M. Brauner	34050
Reasoning on Graphs: Faithful and Interpretable Large Language Model Reasoning Linhao Luo, Yuan-Fang Li, Reza Haf, Shirui Pan	34088
First-Order ANIL Provably Learns Representations Despite Overparametrisation Oğuz Kaan Yıldız, Etienne Boursier, Nicolas Flammarion	34112
Posterior Sampling Based on Gradient Flows of the MMD with Negative Distance Kernel Paul Hagemann, Johannes Hertrich, Fabian Altekötter, Robert Beinert, Jannis Chemseddine, Gabriele Steidl	34154
Batch Calibration: Rethinking Calibration for In-Context Learning and Prompt Engineering Han Zhou, Xingchen Wan, Lev Proleev, Diana Mincu, Jilin Chen, Katherine A Heller, Subhrajit Roy	34187

ModernTCN: A Modern Pure Convolution Structure for General Time Series Analysis Luo Donghao, Wang Xue	34209
Human Feedback is Not Gold Standard Tom Hosking, Phil Blunsom, Max Bartolo	34252
Provable and Practical: Efficient Exploration in Reinforcement Learning Via Langevin Monte Carlo Haque Ishfaq, Qingfeng Lan, Pan Xu, A. Rupam Mahmood, Doina Precup, Anima Anandkumar, Kamyar Azizzadenesheli	34272
DP-SGD Without Clipping: The Lipschitz Neural Network Way Louis Băthune, Thomas Massena, Thibaut Boissin, Aurélien Bellet, Franck Mamalet, Yannick Prudent, Corentin Friedrich, Mathieu Serrurier, David Vigouroux	34318
Sparse Model Soups: A Recipe for Improved Pruning Via Model Averaging Max Zimmer, Christoph Spiegel, Sebastian Pokutta	34364
Improving Non-Transferable Representation Learning by Harnessing Content and Style Ziming Hong, Zhenyi Wang, Li Shen, Yu Yao, Zhuo Huang, Shiming Chen, Chuanwu Yang, Mingming Gong, Tongliang Liu	34395
PeFLL: Personalized Federated Learning by Learning to Learn Jonathan Scott, Hossein Zakerinia, Christoph H Lampert	34423
DEEP NEURAL NETWORK INITIALIZATION WITH SPARSITY INDUCING ACTIVATIONS Ilan Price, Nicholas Daultry Ball, Adam Christopher Jones, Samuel Chun Hei Lam, Jared Tanner	34454
Retrieval-Enhanced Contrastive Vision-Text Models Ahmet Iscen, Mathilde Caron, Alireza Fathi, Cordelia Schmid	34481
C-TPT: Calibrated Test-Time Prompt Tuning for Vision-Language Models Via Text Feature Dispersion Hee Suk Yoon, Eunseop Yoon, Joshua Tian Jin Tee, Mark A. Hasegawa-Johnson, Yingzhen Li, Chang D. Yoo	34495
Leveraging Uncertainty Estimates to Improve Classifier Performance Gundeep Arora, Srujana Merugu, Anoop Saladi, Rajeev Rastogi	34520
Beam Enumeration: Probabilistic Explainability for Sample Efficient Self-Conditioned Molecular Design Jeff Guo, Philippe Schwaller	34550
Reward-Free Curricula for Training Robust World Models Marc Rigter, Minqi Jiang, Ingmar Posner	34591
Light Schrödinger Bridge Alexander Korotin, Nikita Gushchin, Evgeny Burnaev	34619

Learning to Compose: Improving Object Centric Learning by Injecting Compositionality Whie Jung, Jaehoon Yoo, Sungjin Ahn, Seunghoon Hong	34649
ODICE: Revealing the Mystery of Distribution Correction Estimation Via Orthogonal-Gradient Update Liyuan Mao, Haoran Xu, Weinan Zhang, Xianyuan Zhan	34671
Magnushammer: A Transformer-Based Approach to Premise Selection Maciej Mikuć, Szymon Tworkowski, Szymon Antoniak, Bartosz Piotrowski, Albert Q. Jiang, Jin Peng Zhou, Christian Szegedy, Łukasz Kuciński, Piotr Miłoś, Yuhuai Wu	34696
Kalman Filter for Online Classification of Non-Stationary Data Michalis Titsias, Alexandre Galashov, Amal Rannen-Triki, Razvan Pascanu, Yee Whye Teh, Jorg Bornschein	34721
SALMONN: Towards Generic Hearing Abilities for Large Language Models Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, Chao Zhang	34740
Cauchy-Schwarz Divergence Information Bottleneck for Regression Shujian Yu, Xi Yu, Sigurd Lasse, Robert Jenssen, Jose C Principe	34763
Sign2GPT: Leveraging Large Language Models for Gloss-Free Sign Language Translation Ryan Wong, Necati Cihan Camgoz, Richard Bowden	34807
Toward Effective Protection Against Diffusion-Based Mimicry Through Score Distillation Haotian Xue, Chumeng Liang, Xiaoyu Wu, Yongxin Chen	34825
Improved Analysis of Sparse Linear Regression in Local Differential Privacy Model Liyang Zhu, Meng Ding, Vaneet Aggarwal, Jinhui Xu, Di Wang	34850
Simplicial Representation Learning with Neural $\mathbb{R}$-Forms Kelly Maggs, Celia Hacker, Bastian Rieck	34890
Prediction Without Preclusion: Recourse Verification with Reachable Sets Avni Kothari, Bogdan Kulynych, Tsui-Wei Weng, Berk Ustun	34911
Structural Inference with Dynamics Encoding and Partial Correlation Coefficients Aoran Wang, Jun Pang	34934
Revisiting Data Augmentation in Deep Reinforcement Learning Jianshu Hu, Yunpeng Jiang, Paul Weng	34961
Learning Hierarchical World Models with Adaptive Temporal Abstractions from Discrete Latent Dynamics Christian Gumbsch, Noor Sajid, Georg Martius, Martin V. Butz	34998
Improving Equilibrium Propagation Without Weight Symmetry Through Jacobian Homeostasis Axel Laborieux, Friedemann Zenke	35035

DisenBooth: Identity-Preserving Disentangled Tuning for Subject-Driven Text-To-Image Generation Hong Chen, Yipeng Zhang, Simin Wu, Xin Wang, Xuguang Duan, Yuwei Zhou, Wenwu Zhu	35055
INSIDE: LLMs' Internal States Retain the Power of Hallucination Detection Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, Jieping Ye	35078
ODE Discovery for Longitudinal Heterogeneous Treatment Effects Inference Krzysztof Kacprzyk, Samuel Holt, Jeroen Berrevoets, Zhaozhi Qian, Mihaela Van Der Schaar	35099
Generalization of Scaled Deep ResNets in the Mean-Field Regime Yihang Chen, Fanghui Liu, Yiping Lu, Grigorios Chrysos, Volkan Cevher	35136
From Zero to Turbulence: Generative Modeling for 3D Flow Simulation Marten Lienen, David Ldke, Jan Hansen-Palmus, Stephan Gnnemann	35168
GTMGC: Using Graph Transformer to Predict Molecule�s Ground-State Conformation Guikun Xu, Yongquan Jiang, Pengchuan Lei, Yan Yang, Jim Chen	35186
Towards Eliminating Hard Label Constraints in Gradient Inversion Attacks Yanbo Wang, Jian Liang, Ran He	35206
How Realistic is Your Synthetic Data? Constraining Deep Generative Models for Tabular Data Mihaela C Stoian, Salijona Dyrnishi, Maxime Cordy, Thomas Lukasiewicz, Eleonora Giunchiglia	35224
Guiding Masked Representation Learning to Capture Spatio-Temporal Relationship of Electrocardiogram Yeongyeon Na, Minje Park, Yunwon Tae, Sunghoon Joo	35262
Discovering Temporally-Aware Reinforcement Learning Algorithms Matthew Thomas Jackson, Chris Lu, Louis Kirsch, Robert Tjarko Lange, Shimon Whiteson, Jakob Nicolaus Foerster	35286
The Alignment Problem from a Deep Learning Perspective Richard Ngo, Lawrence Chan, Sren Mindermann	35305
Equivariant Scalar Fields for Molecular Docking with Fast Fourier Transforms Bowen Jing, Tommi S. Jaakkola, Bonnie Berger	35333
Towards LLM4QPE: Unsupervised Pretraining of Quantum Property Estimation and a Benchmark Yehui Tang, Hao Xiong, Nianzu Yang, Tailong Xiao, Junchi Yan	35361
Unleashing the Power of Pre-Trained Language Models for Offline Reinforcement Learning Ruizhe Shi, Yuyao Liu, Yanjie Ze, Simon Shaolei Du, Huazhe Xu	35380
Deep Generative Clustering with Multimodal Diffusion Variational Autoencoders Emanuele Palumbo, Laura Manduchi, Sonia Laguna, Daphn�� Chopard, Julia E Vogt	35404

Large Brain Model for Learning Generic Representations with Tremendous EEG Data in BCI Weibang Jiang, Liming Zhao, Bao-Liang Lu	35429
Unraveling the Key Components of OOD Generalization Via Diversification Harold Luc Benoit, Liangze Jiang, Andrei Atanov, Oguzhan Fatih Kar, Mattia Rigotti, Amir Zamir	35451
CLIPSelf: Vision Transformer Distills Itself for Open-Vocabulary Dense Prediction Size Wu, Wenwei Zhang, Lumin Xu, Sheng Jin, Xiangtai Li, Wentao Liu, Chen Change Loy	35483
Learning Mean Field Games on Sparse Graphs: A Hybrid Graphex Approach Christian Fabian, Kai Cui, Heinz Koepl	35503
FOSI: Hybrid First and Second Order Optimization Hadar Sivan, Moshe Gabel, Assaf Schuster	35542
A Differentiable Brain Simulator Bridging Brain Simulation and Brain-Inspired Computing Chaoming Wang, Tianqiu Zhang, Sichao He, Hongyaoxing Gu, Shangyang Li, Si Wu	35565
GAIA: A Benchmark for General AI Assistants Gr�goire Mialon, Cl�mentine Fourier, Thomas Wolf, Yann Lecun, Thomas Scialom	35596
STARC: A General Framework for Quantifying Differences Between Reward Functions Joar Max Viktor Skalse, Lucy Farnik, Sumeet Ramesh Motwani, Erik Jenner, Adam Gleave, Alessandro Abate	35621
Neural Structure Learning with Stochastic Differential Equations Benjie Wang, Joel Jennings, Wenbo Gong	35663
An Extensible Framework for Open Heterogeneous Collaborative Perception Yifan Lu, Yue Hu, Yiqi Zhong, Dequan Wang, Yanfeng Wang, Siheng Chen	35676
Off-Policy Primal-Dual Safe Reinforcement Learning Zifan Wu, Bo Tang, Qian Lin, Chao Yu, Shangqin Mao, Qianlong Xie, Xingxing Wang, Dong Wang	35694
Neuroformer: Multimodal and Multitask Generative Pretraining for Brain Data Antonis Antoniadis, Yiyi Yu, Joe S Canzano, William Yang Wang, Spencer Smith	35714
UniversalNER: Targeted Distillation from Large Language Models for Open Named Entity Recognition Wenxuan Zhou, Sheng Zhang, Yu Gu, Muhao Chen, Hoifung Poon	35739
Leveraging Generative Models for Unsupervised Alignment of Neural Time Series Data Ayesha Vermani, Il Memming Park, Josue Nassar	35758
Multilinear Operator Networks Yixin Cheng, Grigorios Chrysos, Markos Georgopoulos, Volkan Cevher	35774
The Curse of Diversity in Ensemble-Based Exploration Zhixuan Lin, Pierluca D'Oro, Evgenii Nikishin, Aaron Courville	35802

Memorization in Self-Supervised Learning Improves Downstream Generalization Wenhao Wang, Muhammad Ahmad Kaleem, Adam Dziedzic, Michael Backes, Nicolas Papernot, Franziska Boenisch	35838
PAC Prediction Sets Under Label Shift Wenwen Si, Sangdon Park, Insup Lee, Edgar Dobriban, Osbert Bastani	35863
Leveraging Low-Rank and Sparse Recurrent Connectivity for Robust Closed-Loop Control Neehal Tumma, Mathias Lechner, Noel Loo, Ramin Hasani, Daniela Rus	35884
Unbiased Watermark for Large Language Models Zhengmian Hu, Lichang Chen, Xidong Wu, Yihan Wu, Hongyang Zhang, Heng Huang	35926
On the Power of the Weisfeiler-Leman Test for Graph Motif Parameters Matthias Lanzinger, Pablo Barcelo	35955
MetaGPT: Meta Programming for a Multi-Agent Collaborative Framework Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, Jürgen Schmidhuber	35980
Sample-Efficiency in Multi-Batch Reinforcement Learning: The Need for Dimension-Dependent Adaptivity Emmeran Johnson, Ciara Pike-Burke, Patrick Rebeschini	36009
An Improved Analysis of Per-Sample and Per-Update Clipping in Federated Learning Bo Li, Xiaowen Jiang, Mikkel N. Schmidt, Tommy Sonne Alstrøm, Sebastian U Stich	36045
Symmetric Single Index Learning Aaron Zweig, Joan Bruna	36096
Hybrid Sharing for Multi-Label Image Classification Zihao Yin, Chen Gan, Kelei He, Yang Gao, Junfeng Zhang	36107
General Stability Analysis for Zeroth-Order Optimization Algorithms Xinyue Liu, Hualin Zhang, Bin Gu, Hong Chen	36121
Diving Segmentation Model into Pixels Chen Gan, Zihao Yin, Kelei He, Yang Gao, Junfeng Zhang	36167
On Adversarial Training Without Perturbing All Examples Max Losch, Mohamed Omran, David Stutz, Mario Fritz, Bernt Schiele	36184
Particle Guidance: non-I.I.D. Diverse Sampling with Diffusion Models Gabriele Corso, Yilun Xu, Valentin De Bortoli, Regina Barzilay, Tommi S. Jaakkola	36205
Hybrid Directional Graph Neural Network for Molecules Junyi An, Chao Qu, Zhipeng Zhou, Fenglei Cao, Xu Yinghui, Yuan Qi, Furao Shen	36233

StructComp: Substituting Propagation with Structural Compression in Training Graph Contrastive Learning Shengzhong Zhang, Wenjie Yang, Xinyuan Cao, Hongwei Zhang, Zengfeng Huang	36245
TAB: Temporal Accumulated Batch Normalization in Spiking Neural Networks Haiyan Jiang, Vincent Zoonekynd, Giulia De Masi, Bin Gu, Huan Xiong	36265
Learning Decentralized Partially Observable Mean Field Control for Artificial Collective Behavior Kai Cui, Sascha H. Hauck, Christian Fabian, Heinz Koeppl	36292
Maximum Entropy Heterogeneous-Agent Reinforcement Learning Jiarong Liu, Yifan Zhong, Siyi Hu, Haobo Fu, Qiang Fu, Xiaojun Chang, Yaodong Yang	36332
Rethinking the Uniformity Metric in Self-Supervised Learning Xianghong Fang, Jian Li, Qiang Sun, Benyou Wang	36374
Early Neuron Alignment in Two-Layer ReLU Networks with Small Initialization Hancheng Min, Enrique Mallada, Rene Vidal	36396
Neural Architecture Retrieval Xiaohuan Pei, Yanxi Li, Minjing Dong, Chang Xu	36439
Removing Biases from Molecular Representations Via Information Maximization Chenyu Wang, Sharut Gupta, Caroline Uhler, Tommi S. Jaakkola	36453
Incentive-Aware Federated Learning with Training-Time Model Rewards Zhaoxuan Wu, Mohammad Mohammadi Amiri, Ramesh Raskar, Bryan Kian Hsiang Low	36479
Gradual Domain Adaptation Via Gradient Flow Zhan Zhuang, Yu Zhang, Ying Wei	36515
Rigid Protein-Protein Docking Via Equivariant Elliptic-Paraboloid Interface Prediction Ziyang Yu, Wenbing Huang, Yang Liu	36541
Faster Approximation of Probabilistic and Distributional Values Via Least Squares Weida Li, Yaoliang Yu	36557
Distinguished in Uniform: Self-Attention Vs. Virtual Nodes Eran Rosenbluth, Jan Tã¶rnshoff, Martin Ritzert, Berke Kisin, Martin Grohe	36592
Advancing the Lower Bounds: An Accelerated, Stochastic, Second-Order Method with Optimal Adaptation to Inexactness Artem Agafonov, Dmitry Kamzolov, Alexander Gasnikov, Ali Kavis, Kimon Antonakopoulos, Volkan Cevher, Martin Takã¶jã¶	36613
Quantifying and Enhancing Multi-Modal Robustness with Modality Preference Zequn Yang, Yake Wei, Ce Liang, Di Hu	36645

Glucobench: Curated List of Continuous Glucose Monitoring Datasets with Prediction Benchmarks <i>Renat Sergazinov, Elizabeth Chun, Valeriya Rogovchenko, Nathaniel J Fernandes, Nicholas Kasman, Irina Gaynanova</i>	36658
Masks, Signs, and Learning Rate Rewinding <i>Advait Harshal Gadhikar, Rebekka Burkholz</i>	36686
TextField3D: Towards Enhancing Open-Vocabulary 3D Generation with Noisy Text Fields <i>Tianyu Huang, Yihan Zeng, Bowen Dong, Hang Xu, Songcen Xu, Rynson W. H. Lau, Wangmeng Zuo</i>	36708
Fine-Tuned Language Models Generate Stable Inorganic Materials as Text <i>Nate Gruver, Anuroop Sriram, Andrea Madotto, Andrew Gordon Wilson, C. Lawrence Zitnick, Zachary Ward Ulissi</i>	36731
Learning Semantic Proxies from Visual Prompts for Parameter-Efficient Fine-Tuning in Deep Metric Learning <i>Li Ren, Chen Chen, Liqiang Wang, Kien A. Hua</i>	36751
Zipformer: A Faster and Better Encoder for Automatic Speech Recognition <i>Zengwei Yao, Liyong Guo, Xiaoyu Yang, Wei Kang, Fangjun Kuang, Yifan Yang, Zengrui Jin, Long Lin, Daniel Povey</i>	36770
Energy-Guided Entropic Neural Optimal Transport <i>Petr Mokrov, Alexander Korotin, Alexander Kolesov, Nikita Gushchin, Evgeny Burnaev</i>	36786
Prioritized Soft Q-Decomposition for Lexicographic Reinforcement Learning <i>Finn Rietz, Erik Schaffernicht, Stefan Heinrich, Johannes A. Stork</i>	36813
Impact of Computation in Integral Reinforcement Learning for Continuous-Time Control <i>Wenhan Cao, Wei Pan</i>	36843
DSPy: Compiling Declarative Language Model Calls into State-Of-The-Art Pipelines <i>Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan A, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, Christopher Potts</i>	36865
Linear Log-Normal Attention with Unbiased Concentration <i>Yury Nahshan, Joseph Kampeas, Emir Haleva</i>	36896
Brain Decoding: Toward Real-Time Reconstruction of Visual Perception <i>Yohann Benchetrit, Hubert Banville, Jean-Remi King</i>	36921
Increasing Model Capacity for Free: A Simple Strategy for Parameter Efficient Fine-Tuning <i>Haobo Song, Hao Zhao, Soumajit Majumder, Tao Lin</i>	36934
Leave-One-Out Distinguishability in Machine Learning <i>Jiayuan Ye, Anastasia Borovykh, Soufiane Hayou, Reza Shokri</i>	36957

Delphic Offline Reinforcement Learning Under Nonidentifiable Hidden Confounding Aliz�e Pace, Hugo Ya�che, Bernhard Sch�lkopf, Gunnar Ratsch, Guy Tennenholtz	36994
Near-Optimal Solutions of Constrained Learning Problems Juan Elenter, Luiz F. O. Chamon, Alejandro Ribeiro	37023
On the Reliability of Watermarks for Large Language Models John Kirchenbauer, Jonas Geiping, Yuxin Wen, Manli Shu, Khalid Saifullah, Kezhi Kong, Kasun Fernando, Aniruddha Saha, Micah Goldblum, Tom Goldstein	37048
Chain-Of-Experts: When LLMs Meet Complex Operations Research Problems Ziyang Xiao, Dongxiang Zhang, Yangjun Wu, Lilin Xu, Yuan Jessica Wang, Xiongwei Han, Xiaojin Fu, Tao Zhong, Jia Zeng, Mingli Song, Gang Chen	37093
Weakly Supervised Virus Capsid Detection with Image-Level Annotations in Electron Microscopy Images Hannah Kniesel, Leon Sick, Tristan Payer, Tim Bergner, Kavitha Shaga Devan, Clarissa Read, Paul Walther, Timo Ropinski, Pedro Hermosilla	37112
Convolutional Deep Kernel Machines Edward Milsom, Ben Anson, Laurence Aitchison	37147
Bounds on Representation-Induced Confounding Bias for Treatment Effect Estimation Valentyn Melnychuk, Dennis Frauen, Stefan Feuerriegel	37176
Experimental Design for Multi-Channel Imaging Via Task-Driven Feature Selection Stefano B. Blumberg, Paddy J. Slator, Daniel C. Alexander	37200
SliceGPT: Compress Large Language Models by Deleting Rows and Columns Saleh Ashkboos, Maximilian L. Croci, Marcelo Gennari Do Nascimento, Torsten Hoefler, James Hensman	37227
Beyond Stationarity: Convergence Analysis of Stochastic Softmax Policy Gradient Methods Sara Klein, Simon Weissmann, Leif D�rling	37247
Bilevel Optimization Under Unbounded Smoothness: A New Algorithm and Convergence Analysis Jie Hao, Xiaochuan Gong, Mingrui Liu	37301
T-Rep: Representation Learning for Time Series Using Time-Embeddings Archibald Felix Fraikin, Adrien Bennetot, Stephanie Allasonniere	37350
DORSal: Diffusion for Object-Centric Representations of Scenes $\{et AI.\}$ Allan Jabri, Sjoerd Van Steenkiste, Emiel Hoogeboom, Mehdi S. M. Sajjadi, Thomas Kipf	37371
Probabilistic Self-Supervised Representation Learning Via Scoring Rules Minimization Amirhossein Vahidi, Simon Schosser, Lisa Wimmer, Yawei Li, Bernd Bischl, Eyke H�llermeier, Mina Rezaei	37394

DyST: Towards Dynamic Neural Scene Representations on Real-World Videos Maximilian Seitzer, Sjoerd Van Steenkiste, Thomas Kipf, Klaus Greff, Mehdi S. M. Sajjadi	37413
GAIA: Zero-Shot Talking Avatar Generation Tianyu He, Junliang Guo, Runyi Yu, Yuchi Wang, Jialiang Zhu, Kaikai An, Leyi Li, Xu Tan, Chunyu Wang, Han Hu, Hsiangtao Wu, Sheng Zhao, Jiang Bian	37427
CAMIL: Context-Aware Multiple Instance Learning for Cancer Detection and Subtyping in Whole Slide Images Olga Fourkioti, Matt De Vries, Chris Bakal	37445
Understanding the Effects of RLHF on LLM Generalisation and Diversity Robert Kirk, Ishita Mediratta, Christoforos Nalmpantis, Jelena Luketina, Eric Hambro, Edward Grefenstette, Roberta Raileanu	37461
RobustTSF: Towards Theory and Design of Robust Time Series Forecasting with Anomalies Hao Cheng, Qingsong Wen, Yang Liu, Liang Sun	37495
Benign Oscillation of Stochastic Gradient Descent with Large Learning Rate Miao Lu, Beining Wu, Xiaodong Yang, Difan Zou	37522
Vision-By-Language for Training-Free Compositional Image Retrieval Shyamgopal Karthik, Karsten Roth, Massimiliano Mancini, Zeynep Akata	37584
Unleashing Large-Scale Video Generative Pre-Training for Visual Robot Manipulation Hongtao Wu, Ya Jing, Chilam Cheang, Guangzeng Chen, Jiafeng Xu, Xinghang Li, Minghuan Liu, Hang Li, Tao Kong	37600
Clifford Group Equivariant Simplicial Message Passing Networks Cong Liu, David Ruhe, Floor Eijkelboom, Patrick Forrá©	37622
Fourier Transporter: Bi-Equivariant Robotic Manipulation in 3D Haojie Huang, Owen Lewis Howell, Dian Wang, Xupeng Zhu, Robert Platt, Robin Walters	37641
Federated Wasserstein Distance Alain Rakotomamonjy, Kimia Nadjahi, Liva Ralaivola	37663
On Trajectory Augmentations for Off-Policy Evaluation Ge Gao, Qitong Gao, Xi Yang, Song Ju, Miroslav Pajic, Min Chi	37687
Diffusion Generative Flow Samplers: Improving Learning Signals Through Partial Trajectory Optimization Dinghuai Zhang, Ricky T. Q. Chen, Cheng-Hao Liu, Aaron Courville, Yoshua Bengio	37731
Large-Scale Training of Foundation Models for Wearable Biosignals Salar Abbaspourazad, Oussama Elachqar, Andrew Miller, Saba Emrani, Udhyakumar Nallasamy, Ian Shapiro	37753

Incentivized Truthful Communication for Federated Bandits Zhepei Wei, Chuanhao Li, Tianze Ren, Haifeng Xu, Hongning Wang	37777
The Expressive Leaky Memory Neuron: An Efficient and Expressive Phenomenological Neuron Model Can Solve Long-Horizon Tasks. Aaron Spieler, Nasim Rahaman, Georg Martius, Bernhard Schölkopf, Anna Levina	37797
Training Socially Aligned Language Models on Simulated Social Interactions Ruibo Liu, Ruixin Yang, Chenyan Jia, Ge Zhang, Diyi Yang, Soroush Vosoughi	37822
Bayesian Neural Controlled Differential Equations for Treatment Effect Estimation Konstantin Hess, Valentyn Melnychuk, Dennis Frauen, Stefan Feuerriegel	37846
BENO: Boundary-Embedded Neural Operators for Elliptic PDEs Haixin Wang, Jiaxin Li, Anubhav Dwivedi, Kentaro Hara, Tailin Wu	37883
A Plug-And-Play Image Registration Network Junhao Hu, Weijie Gan, Zhixin Sun, Hongyu An, Ulugbek Kamilov	37905
AntGPT: Can Large Language Models Help Long-Term Action Anticipation from Videos? Qi Zhao, Shijie Wang, Ce Zhang, Changcheng Fu, Minh Quan Do, Nakul Agarwal, Kwonjoon Lee, Chen Sun	37926
VeRA: Vector-Based Random Matrix Adaptation Dawid Jan Kopiczko, Tijmen Blankevoort, Yuki M Asano	37947
PerceptionCLIP: Visual Classification by Inferring and Conditioning on Contexts Bang An, Sicheng Zhu, Michael-Andrei Panaitescu-Liess, Chaithanya Kumar Mummadi, Furong Huang	37968
Ring-A-Bell! How Reliable Are Concept Removal Methods for Diffusion Models? Yu-Lin Tsai, Chia-Yi Hsu, Chulin Xie, Chih-Hsun Lin, Jia You Chen, Bo Li, Pin-Yu Chen, Chia-Mu Yu, Chun-Ying Huang	37994
Object-Centric Learning with Slot Mixture Module Daniil Kirilenko, Vitaliy Vorobyov, Alexey Kovalev, Aleksandr Panov	38006
Orbit-Equivariant Graph Neural Networks Matthew Morris, Bernardo Cuenca Grau, Ian Horrocks	38025
Towards Unified Multi-Modal Personalization: Large Vision-Language Models for Generative Recommendation and Beyond Tianxin Wei, Bowen Jin, Ruirui Li, Hansi Zeng, Zhengyang Wang, Jianhui Sun, Qingyu Yin, Hanqing Lu, Suhang Wang, Jingrui He, Xianfeng Tang	38064
Image Clustering Via the Principle of Rate Reduction in the Age of Pretrained Models Tianzhe Chu, Shengbang Tong, Tianjiao Ding, Xili Dai, Benjamin David Haeffele, Rene Vidal, Yi Ma	38083

TRAM: Bridging Trust Regions and Sharpness Aware Minimization Tom Sherborne, Naomi Saphra, Pradeep Dasigi, Hao Peng	38106
Lifting Architectural Constraints of Injective Flows Peter Sorrenson, Felix Draxler, Armand Rousselot, Sander Hummerich, Lea Zimmermann, Ullrich Koethe	38127
Provably Efficient UCB-Type Algorithms for Learning Predictive State Representations Ruiquan Huang, Yingbin Liang, Jing Yang	38163
On the Markov Property of Neural Algorithmic Reasoning: Analyses and Methods Montgomery Bohde, Meng Liu, Alexandra Saxton, Shuiwang Ji	38206
Non-Exchangeable Conformal Risk Control Ant�nio Farinhas, Chrysoula Zerva, Dennis Thomas Ulmer, Andre Martins	38223
Latent Representation and Simulation of Markov Processes Via Time-Lagged Information Bottleneck Marco Federici, Patrick Forr�, Ryota Tomioka, Bastiaan S. Veeling	38238
Node2ket: Efficient High-Dimensional Network Embedding in Quantum Hilbert Space Hao Xiong, Yehui Tang, Yunlin He, Wei Tan, Junchi Yan	38275
Can We Get the Best of Both Binary Neural Networks and Spiking Neural Networks for Efficient Computer Vision? Gourav Datta, Zeyu Liu, Peter Anthony Beerel	38308
Emergent Communication with Conversational Repair Mitja Nikolaus	38329
Open the Black Box: Step-Based Policy Updates for Temporally-Correlated Episodic Reinforcement Learning Ge Li, Hongyi Zhou, Dominik Roth, Serge Thilges, Fabian Otto, Rudolf Lioutikov, Gerhard Neumann	38346
Retrieval-Based Disentangled Representation Learning with Natural Language Supervision Jiawei Zhou, Xiaoguang Li, Lifeng Shang, Xin Jiang, Qun Liu, Lei Chen	38379
Time-Efficient Reinforcement Learning with Stochastic Stateful Policies Firas Al-Hafez, Guoping Zhao, Jan Peters, Davide Tateo	38399
A Topological Perspective on Demystifying GNN-Based Link Prediction Performance Yu Wang, Tong Zhao, Yuying Zhao, Yunchao Liu, Xueqi Cheng, Neil Shah, Tyler Derr	38436
EQA-MX: Embodied Question Answering Using Multimodal Expression Md Mofijul Islam, Alexi Gladstone, Riashat Islam, Tariq Iqbal	38475
Neural Optimal Transport with General Cost Functionals Arip Asadulaev, Alexander Korotin, Vage Egiazarian, Petr Mokrov, Evgeny Burnaev	38499

Optimal Criterion for Feature Learning of Two-Layer Linear Neural Network in High Dimensional Interpolation Regime Keita Suzuki, Taiji Suzuki	38531
Piecewise Linear Parametrization of Policies: Towards Interpretable Deep Reinforcement Learning Maxime Wabartha, Joelle Pineau	38542
Variational Bayesian Last Layers James Harrison, John Willes, Jasper Snoek	38585
Towards Robust Fidelity for Evaluating Explainability of Graph Neural Networks Xu Zheng, Farhad Shirani, Tianchun Wang, Wei Cheng, Zhuomin Chen, Haifeng Chen, Hua Wei, Dongsheng Luo	38615
Towards Robust Offline Reinforcement Learning Under Diverse Data Corruption Rui Yang, Han Zhong, Jiawei Xu, Amy Zhang, Chongjie Zhang, Lei Han, Tong Zhang	38641
TorchRL: A Data-Driven Decision-Making Library for PyTorch Albert Bou, Matteo Bettini, Sebastian Dittert, Vikash Kumar, Shagun Sodhani, Xiaomeng Yang, Gianni De Fabritiis, Vincent Moens	38673
BEND: Benchmarking DNA Language Models on Biologically Meaningful Tasks Frederikke Isa Marin, Felix Teufel, Marc Horlacher, Dennis Madsen, Dennis Pultz, Ole Winther, Wouter Boomsma	38707
On the Provable Advantage of Unsupervised Pretraining Jiawei Ge, Shange Tang, Jianqing Fan, Chi Jin	38743
Subtractive Mixture Models Via Squaring: Representation and Learning Lorenzo Loconte, Aleksanteri Mikulus Sladek, Stefan Mengel, Martin Trapp, Arno Solin, Nicolas Gillis, Antonio Vergari	38755
From Latent Graph to Latent Topology Inference: Differentiable Cell Complex Module Claudio Battiloro, Indro Spinelli, Lev Telyatnikov, Michael M. Bronstein, Simone Scardapane, Paolo Di Lorenzo	38790
Better Neural PDE Solvers Through Data-Free Mesh Movers Peiyan Hu, Yue Wang, Zhi-Ming Ma	38812
Conditional Variational Diffusion Models Gabriel Della Maggiora, Luis Alberto Croquevielle, Nikita Deshpande, Harry Horsley, Thomas Heinis, Artur Yakimovich	38839
Understanding Transferable Representation Learning and Zero-Shot Transfer in CLIP Zixiang Chen, Yihe Deng, Yuanzhi Li, Quanquan Gu	38871

Replay Across Experiments: A Natural Extension of Off-Policy RL Dhruva Tirumala, Thomas Lampe, Jose Enrique Chen, Tuomas Haarnoja, Sandy Huang, Guy Lever, Ben Moran, Tim Hertweck, Leonard Hasenclever, Martin Riedmiller, Nicolas Heess, Markus Wulfmeier	38901
OpenWebMath: An Open Dataset of High-Quality Mathematical Web Text Keiran Paster, Marco Dos Santos, Zhangir Azerbayev, Jimmy Ba	38918
Solving High Frequency and Multi-Scale PDEs with Gaussian Processes Shikai Fang, Madison Cooley, Da Long, Shibo Li, Mike Kirby, Shandian Zhe	38941
Robust Model-Based Optimization for Challenging Fitness Landscapes Saba Ghaffari, Ehsan Saleh, Alex Schwing, Yu-Xiong Wang, Martin D. Burke, Saurabh Sinha	38958
A Newborn Embodied Turing Test for Comparing Object Segmentation Across Animals and Machines Manju Garimella, Denizhan Pak, Justin Newell Wood, Samantha Marie Waters Wood	38982
SS2\$AC: Energy-Based Reinforcement Learning with Stein Soft Actor Critic Safa Messaoud, Billel Mokeddem, Zhenghai Xue, Linsey Pang, Bo An, Haipeng Chen, Sanjay Chawla	39004
Hybrid Internal Model: Learning Agile Legged Locomotion with Simulated Robot Response Junfeng Long, Zirui Wang, Quanyi Li, Liu Cao, Jiawei Gao, Jiangmiao Pang	39018
Implicit Maximum a Posteriori Filtering Via Adaptive Optimization Gianluca Bencomo, Jake Snell, Thomas L. Griffiths	39035
On the Expressivity of Objective-Specification Formalisms in Reinforcement Learning Rohan Subramani, Marcus Williams, Max Heitmann, Halfdan Holm, Charlie Griffin, Joar Max Viktor Skalse	39047
Learning from Sparse Offline Datasets Via Conservative Density Estimation Zhepeng Cen, Zuxin Liu, Zitong Wang, Yihang Yao, Henry Lam, Ding Zhao	39110
Quality-Diversity Through AI Feedback Herbie Bradley, Andrew Dai, Hannah Benita Teufel, Jenny Zhang, Koen Oostermeijer, Marco Bellagente, Jeff Clune, Kenneth Stanley, Gregory Schott, Joel Lehman	39137
AUGCAL: Improving Sim2Real Adaptation by Uncertainty Calibration on Augmented Synthetic Images Prithvijit Chattopadhyay, Bharat Goyal, Boglarka Ecsedi, Viraj Uday Prabhu, Judy Hoffman	39249
VBH-GNN: Variational Bayesian Heterogeneous Graph Neural Networks for Cross-Subject Emotion Recognition Chenyu Liu, Xinliang Zhou, Zhengri Zhu, Liming Zhai, Ziyu Jia, Yang Liu	39270
Benign Overfitting and Grokking in ReLU Networks for XOR Cluster Data Zhiwei Xu, Yutong Wang, Spencer Frei, Gal Vardi, Wei Hu	39286

Chain of Log-Concave Markov Chains Saeed Saremi, Ji Won Park, Francis Bach	39331
OVOR: OnePrompt with Virtual Outlier Regularization for Rehearsal-Free Class-Incremental Learning Wei-Cheng Huang, Chun-Fu Chen, Hsiang Hsu	39355
Generative Modeling with Phase Stochastic Bridge Tianrong Chen, Jiatao Gu, Laurent Dinh, Evangelos Theodorou, Joshua M. Susskind, Shuangfei Zhai	39375
Linearity of Relation Decoding in Transformer Language Models Evan Hernandez, Arnab Sen Sharma, Tal Haklay, Kevin Meng, Martin Wattenberg, Jacob Andreas, Yonatan Belinkov, David Bau	39412
Sparse Autoencoders Find Highly Interpretable Features in Language Models Robert Huben, Hoagy Cunningham, Logan Riggs Smith, Aidan Ewart, Lee Sharkey	39435
Toward Optimal Policy Population Growth in Two-Player Zero-Sum Games Stephen Marcus McAleer, Jb Lanier, Kevin A. Wang, Pierre Baldi, Tuomas Sandholm, Roy Fox	39454
Lagrangian Flow Networks for Conservation Laws Fabricio Arend Torres, Marcello Massimo Negri, Marco Inversi, Jonathan Aellen, Volker Roth	39483
Masked Autoencoders with Multi-Window Local-Global Attention Are Better Audio Learners Sarthak Yadav, Sergios Theodoridis, Lars Kai Hansen, Zheng-Hua Tan	39508
FeatUp: A Model-Agnostic Framework for Features at Any Resolution Stephanie Fu, Mark Hamilton, Laura E. Brandt, Axel Feldmann, Zhoutong Zhang, William T. Freeman	39528
Pre-Training with Random Orthogonal Projection Image Modeling Maryam Haghighat, Peyman Moghadam, Shaheer Mohamed, Piotr Koniusz	39555
Functional Interpolation for Relative Positions Improves Long Context Transformers Shanda Li, Chong You, Guru Guruganesh, Joshua Ainslie, Santiago Ontanon, Manzil Zaheer, Sumit Sanghai, Yiming Yang, Sanjiv Kumar, Srinadh Bhojanapalli	39574
Accelerating Sinkhorn Algorithm with Sparse Newton Iterations Xun Tang, Michael Shavlovsky, Holakou Rahmanian, Elisa Tardini, Kiran Koshy Thekumparampil, Tesi Xiao, Lexing Ying	39600
YaRN: Efficient Context Window Extension of Large Language Models Bowen Peng, Jeffrey Quesnelle, Honglu Fan, Enrico Shippole	39620
Privately Aligning Language Models with Reinforcement Learning Fan Wu, Huseyin A Inan, Arturs Backurs, Varun Chandrasekaran, Janardhan Kulkarni, Robert Sim	39640
Sharpness-Aware Data Poisoning Attack Pengfei He, Han Xu, Jie Ren, Yingqian Cui, Shenglai Zeng, Hui Liu, Charu C. Aggarwal, Jiliang Tang	39659

Post-Hoc Bias Scoring is Optimal for Fair Classification Wenlong Chen, Yegor Klochkov, Yang Liu	39680
Talk Like a Graph: Encoding Graphs for Large Language Models Bahare Fatemi, Jonathan Halcrow, Bryan Perozzi	39708
Think Before You Speak: Training Language Models with Pause Tokens Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, Vaishnavh Nagarajan	39734
A Characterization Theorem for Equivariant Networks with Point-Wise Activations Marco Pacini, Xiaowen Dong, Bruno Lepri, Gabriele Santin	39762
Towards Generative Abstract Reasoning: Completing Raven’s Progressive Matrix Via Rule Abstraction and Selection Fan Shi, Bin Li, Xiangyang Xue	39783
When Can Transformers Reason with Abstract Symbols? Enric Boix-Adserà, Omid Saremi, Emmanuel Abbe, Samy Bengio, Etai Littwin, Joshua M. Susskind	39812
Nougat: Neural Optical Understanding for Academic Documents Lukas Blecher, Guillem Cucurull, Thomas Scialom, Robert Stojnic	39861
T-MARS: Improving Visual Representations by Circumventing Text Feature Learning Pratyush Maini, Sachin Goyal, Zachary Chase Lipton, J Zico Kolter, Aditi Raghunathan	39879
Evaluating Representation Learning on the Protein Structure Universe Arian Rokkum Jamasb, Alex Morehead, Chaitanya K. Joshi, Zuobai Zhang, Kieran Didi, Simon V Mathis, Charles Harris, Jian Tang, Jianlin Cheng, Pietro Lio, Tom Leon Blundell	39908
Quadratic Models for Understanding Catapult Dynamics of Neural Networks Libin Zhu, Chaoyue Liu, Adityanarayanan Radhakrishnan, Mikhail Belkin	39939
LLM Augmented LLMs: Expanding Capabilities Through Composition Rachit Bansal, Bidisha Samanta, Siddharth Dalmia, Nitish Gupta, Sriram Ganapathy, Abhishek Bapna, Prateek Jain, Partha Talukdar	39977
Estimating Conditional Mutual Information for Dynamic Feature Selection Soham Gadgil, Ian Connick Covert, Su-In Lee	39994
Reconciling Spatial and Temporal Abstractions for Goal Representation Mehdi Zadem, Sergio Mover, Sao Mai Nguyen	40026
Towards Diverse Behaviors: A Benchmark for Imitation Learning with Human Demonstrations Xiaogang Jia, Denis Blessing, Xinkai Jiang, Moritz Reuss, Atalay Donat, Rudolf Lioutikov, Gerhard Neumann	40045
Rethinking Backdoor Attacks on Dataset Distillation: A Kernel Method Perspective Ming-Yu Chung, Sheng-Yen Chou, Chia-Mu Yu, Pin-Yu Chen, Sy-Yen Kuo, Tsung-Yi Ho	40066

Transformers Can Optimally Learn Regression Mixture Models <i>Reese Pathak, Rajat Sen, Weihao Kong, Abhimanyu Das</i>	40078
Optimal Sample Complexity of Contrastive Learning <i>Noga Alon, Dmitrii Avdiukhin, Dor Elboim, Orr Fischer, Grigory Yaroslavtsev</i>	40101
Mixture of Weak and Strong Experts on Graphs <i>Hanqing Zeng, Hanjia Lyu, Diyi Hu, Yinglong Xia, Jiebo Luo</i>	40126
Learning Polynomial Problems with $SL(2, \mathbb{R})$-Equivariance <i>Hannah Lawrence, Mitchell Tong Harris</i>	40167
Towards Principled Representation Learning from Videos for Reinforcement Learning <i>Dipendra Misra, Akanksha Saran, Tengyang Xie, Alex Lamb, John Langford</i>	40191
The Marginal Value of Momentum for Small Learning Rate SGD <i>Runzhe Wang, Sadhika Malladi, Tianhao Wang, Kaifeng Lyu, Zhiyuan Li</i>	40231
Language Model Cascades: Token-Level Uncertainty and Beyond <i>Neha Gupta, Harikrishna Narasimhan, Wittawat Jitkittum, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar</i>	40275
Exploiting Causal Graph Priors with Posterior Sampling for Reinforcement Learning <i>Mirco Mutti, Riccardo De Santi, Marcello Restelli, Alexander Marx, Giorgia Ramponi</i>	40309
Grokking as the Transition from Lazy to Rich Training Dynamics <i>Tanishq Kumar, Blake Bordelon, Samuel J. Gershman, Cengiz Pehlevan</i>	40342
Towards Characterizing Domain Counterfactuals for Invertible Latent Causal Models <i>Zeyu Zhou, Ruqi Bai, Sean Kulinski, Murat Kocaoglu, David I. Inouye</i>	40368
Faster Sampling from Log-Concave Densities Over Polytopes Via Efficient Linear Solvers <i>Oren Mangoubi, Nisheeth K. Vishnoi</i>	40418
Causally Aligned Curriculum Learning <i>Mingxuan Li, Junzhe Zhang, Elias Bareinboim</i>	40441
Differentiable Euler Characteristic Transforms for Shape Classification <i>Ernst Rießell, Bastian Rieck</i>	40470
Searching for High-Value Molecules Using Reinforcement Learning and Transformers <i>Raj Ghugare, Santiago Miret, Adriana Hugessen, Mariano Phielipp, Glen Berseth</i>	40482
Jumanji: A Diverse Suite of Scalable Reinforcement Learning Environments in JAX <i>Clément Bonnet, Daniel Luo, Donal John Byrne, Shikha Surana, Sasha Abramowitz, Paul Duckworth, Vincent Coyette, Laurence Illing Midgley, Elshadai Tegegn, Tristan Kalloniatis, Omayma Mahjoub, Matthew Macfarlane, Andries Petrus Smit, Nathan Grinsztajn, Raphael Boige, Cemlyn Neil Waters, Mohamed Ali Ali Mimouni, Ulrich Arnel Mbou Sob, Ruan John De Kock, Siddarth Singh, Daniel Furelos-Blanco, Victor Le, Arnu Pretorius, Alexandre Laterre</i>	40512

Bayesian Optimization Through Gaussian Cox Process Models for Spatio-Temporal Data Yongsheng Mei, Mahdi Imani, Tian Lan	40542
Quantifying the Plausibility of Context Reliance in Neural Machine Translation Gabriele Sarti, Grzegorz Chrupała, Malvina Nissim, Arianna Bisazza	40563
Towards Establishing Guaranteed Error for Learned Database Operations Sepanta Zeighami, Cyrus Shahabi	40592
"What Data Benefits My Classifier?" Enhancing Model Performance and Interpretability Through Influence-Based Data Selection Anshuman Chhabra, Peizhao Li, Prasant Mohapatra, Hongfu Liu	40620
MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, Mohamed Elhoseiny	40652
Improved Sampling Via Learned Diffusions Lorenz Richter, Julius Berner	40669
Multimodal Patient Representation Learning with Missing Modalities and Labels Zhenbang Wu, Anant Dadu, Nicholas Tustison, Brian Avants, Mike Nalls, Jimeng Sun, Faraz Faghri	40697
Adapting to Distribution Shift by Visual Domain Prompt Generation Zhixiang Chi, Li Gu, Tao Zhong, Huan Liu, Yuanhao Yu, Konstantinos N Plataniotis, Yang Wang	40717
Locality-Aware Graph Rewiring in GNNs Federico Barbero, Ameya Velingker, Amin Saberi, Michael M. Bronstein, Francesco Di Giovanni	40737
Training Bayesian Neural Networks with Sparse Subspace Variational Inference Junbo Li, Zichen Miao, Qiang Qiu, Ruqi Zhang	40759
PhyloGFN: Phylogenetic Inference with Generative Flow Networks Ming Yang Zhou, Zichao Yan, Elliot Layne, Nikolay Malkin, Dinghuai Zhang, Moksh Jain, Mathieu Blanchette, Yoshua Bengio	40771
Improving LoRA in Privacy-Preserving Federated Learning Youbang Sun, Zitao Li, Yaliang Li, Bolin Ding	40805
Behaviour Distillation Andrei Lupu, Chris Lu, Jarek Luca Liesen, Robert Tjarko Lange, Jakob Nicolaus Foerster	40822
ECoFLaP: Efficient Coarse-To-Fine Layer-Wise Pruning for Vision-Language Models Yi-Lin Sung, Jaehong Yoon, Mohit Bansal	40839
Bandits Meet Mechanism Design to Combat Clickbait in Online Recommendation Thomas Kleine Buening, Aadirupa Saha, Christos Dimitrakakis, Haifeng Xu	40859
Privacy Amplification for Matrix Mechanisms Christopher A. Choquette-Choo, Arun Ganesh, Thomas Steinke, Abhradeep Guha Thakurta	40888

Understanding Reconstruction Attacks with the Neural Tangent Kernel and Dataset Distillation Noel Loo, Ramin Hasani, Mathias Lechner, Alexander Amini, Daniela Rus	40918
Dropout Enhanced Bilevel Training Peiran Yu, Junyi Li, Heng Huang	40964
Graph Metanetworks for Processing Diverse Neural Architectures Derek Lim, Haggai Maron, Marc T. Law, Jonathan Lorraine, James Lucas	40976
Channel Vision Transformers: An Image is Worth 1 X 16 X 16 Words Yujia Bao, Srinivasan Sivanandan, Theofanis Karaletsos	41005
Synaptic Weight Distributions Depend on the Geometry of Plasticity Roman Pogodin, Jonathan Cornford, Arna Ghosh, Gauthier Gidel, Guillaume Lajoie, Blake Aaron Richards	41032
Self-RAG: Learning to Retrieve, Generate, and Critique Through Self-Reflection Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, Hannaneh Hajishirzi	41056
Constrained Decoding for Cross-Lingual Label Projection Duong Minh Le, Yang Chen, Alan Ritter, Wei Xu	41086
Robustness of AI-Image Detectors: Fundamental Limits and Practical Attacks Mehrdad Saberi, Vinu Sankar Sadasivan, Keivan Rezaei, Aounon Kumar, Atoosa Chegini, Wenxiao Wang, Soheil Feizi	41108
TiC-CLIP: Continual Training of CLIP Models Saurabh Garg, Mehrdad Farajtabar, Hadi Pouransari, Raviteja Vemulapalli, Sachin Mehta, Oncel Tuzel, Vaishaal Shankar, Fartash Faghri	41135
A Primal-Dual Approach to Solving Variational Inequalities with General Constraints Tatjana Chavdarova, Tong Yang, Matteo Pagliardini, Michael Jordan	41171
A ROBUST DIFFERENTIAL NEURAL ODE OPTIMIZER Panagiotis Theodoropoulos, Guan-Horng Liu, Tianrong Chen, Augustinos D Saravanos, Evangelos Theodorou	41238
TACTIS-2: Better, Faster, Simpler Attentional Copulas for Multivariate Time Series Arjun Ashok, Étienne Marcotte, Valentina Zantedeschi, Nicolas Chapados, Alexandre Drouin	41267
Does Progress on Object Recognition Benchmarks Improve Generalization on Crowdsourced, Global Data? Megan Richards, Polina Kirichenko, Diane Bouchacourt, Mark Ibrahim	41295
LipSim: A Provably Robust Perceptual Similarity Metric Sara Ghazanfari, Alexandre Araujo, Prashanth Krishnamurthy, Farshad Khorrami, Siddharth Garg	41320
Unsupervised Pretraining for Fact Verification by Language Model Distillation Adrián Bazaga, Pietro Lio, Gos Micklem	41337

Adversarial Imitation Learning Via Boosting Jonathan Daniel Chang, Dhruv Sreenivas, Yingbing Huang, Kiantã© Brantley, Wen Sun	41352
Topological Data Analysis on Noisy Quantum Computers Ismail Yunus Akhalwaya, Shashanka Ubaru, Kenneth L. Clarkson, Mark S. Squillante, Vishnu Jejjala, Yang-Hui He, Kugendran Naidoo, Vasileios Kalantzis, Lior Horesh	41371
Understanding Prompt Engineering May Not Require Rethinking Generalization Victor Akinwande, Yiding Jiang, Dylan Sam, J Zico Kolter	41405
Curiosity-Driven Red-Teaming for Large Language Models Zhang-Wei Hong, Idan Shenfeld, Tsun-Hsuan Wang, Yung-Sung Chuang, Aldo Pareja, James R. Glass, Akash Srivastava, Pulkit Agrawal	41422
Sampling Multimodal Distributions with the Vanilla Score: Benefits of Data-Based Initialization Frederic Koehler, Thuy-Duong Vuong	41449
Tight Rates in Supervised Outlier Transfer Learning Mohammadreza Mousavi Kalan, Samory Kpotufe	41500
Learning Interpretable Control Inputs and Dynamics Underlying Animal Locomotion Thomas Soares Mullen, Marine Schimel, Guillaume Hennequin, Christian K. Machens, Michael Orger, Adrien Jouary	41517
Efficient Dynamics Modeling in Interactive Environments with Koopman Theory Arnab Kumar Mondal, Siba Smarak Panigrahi, Sai Rajeswar, Kaleem Siddiqi, Siamak Ravanbakhsh	41540
Dynamic Sparse Training with Structured Sparsity Mike Lasby, Anna Golubeva, Utku Evci, Mihai Nica, Yani Ioannou	41564
Plugin Estimators for Selective Classification with Out-Of-Distribution Detection Harikrishna Narasimhan, Aditya Krishna Menon, Wittawat Jitkrittum, Sanjiv Kumar	41593
On the Joint Interaction of Models, Data, and Features Yiding Jiang, Christina Baek, J Zico Kolter	41627
Towards Training Without Depth Limits: Batch Normalization Without Gradient Explosion Alexandru Meterez, Amir Joudaki, Francesco Orabona, Alexander Immer, Gunnar Ratsch, Hadi Daneshmand	41664
Let Models Speak Ciphers: Multiagent Debate Through Embeddings Chau Pham, Boyi Liu, Yingxiang Yang, Zhengyu Chen, Tianyi Liu, Jianbo Yuan, Bryan A. Plummer, Zhaoran Wang, Hongxia Yang	41704
On the Foundations of Shortcut Learning Katherine Hermann, Hossein Mobahi, Thomas Fel, Michael Curtis Mozer	41734
Learning to Reject Meets Long-Tail Learning Harikrishna Narasimhan, Aditya Krishna Menon, Wittawat Jitkrittum, Neha Gupta, Sanjiv Kumar	41771

Zoology: Measuring and Improving Recall in Efficient Language Models Simran Arora, Sabri Eyuboglu, Aman Timalsina, Isys Johnson, Michael Poli, James Zou, Atri Rudra, Christopher Re	41799
CrossLoco: Human Motion Driven Control of Legged Robots Via Guided Unsupervised Reinforcement Learning Tianyu Li, Hyunyoung Jung, Matthew Gombolay, Yong Cho, Sehoon Ha	41866
Can Sensitive Information Be Deleted from LLMs? Objectives for Defending Against Extraction Attacks Vaidehi Patil, Peter Hase, Mohit Bansal	41880
Adaptive Federated Learning with Auto-Tuned Clients Junhyung Lyle Kim, Taha Toghani, Cesar A Uribe, Anastasios Kyrillidis	41898
Energy-Conserving Equivariant GNN for Elasticity of Lattice Architected Metamaterials Ivan Grega, Ilyes Batatia, Gabor Csanyi, Sri Karlapati, Vikram Deshpande	41926
Prompt Risk Control: A Rigorous Framework for Responsible Deployment of Large Language Models Thomas P Zollo, Todd Morrill, Zhun Deng, Jake Snell, Toniann Pitassi, Richard Zemel	41951
A Black-Box Approach for Non-Stationary Multi-Agent Reinforcement Learning Haozhe Jiang, Qiwen Cui, Zhihan Xiong, Maryam Fazel, Simon Shaolei Du	41974
Rephrase, Augment, Reason: Visual Grounding of Questions for Vision-Language Models Archiki Prasad, Elias Stengel-Eskin, Mohit Bansal	42000
Thin-Shell Object Manipulations with Differentiable Physics Simulations Yian Wang, Juntian Zheng, Zhehuan Chen, Zhou Xian, Gu Zhang, Chao Liu, Chuang Gan	42023
ReLU Strikes Back: Exploiting Activation Sparsity in Large Language Models Seyed Iman Mirzadeh, Keivan Alizadeh-Vahid, Sachin Mehta, Carlo C Del Mundo, Oncel Tuzel, Golnoosh Samei, Mohammad Rastegari, Mehrdad Farajtabar	42044
Universal Guidance for Diffusion Models Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, Tom Goldstein	42067
Quantifying Language Models' Sensitivity to Spurious Features in Prompt Design Or: How I Learned to Start Worrying About Prompt Formatting Melanie Sclar, Yejin Choi, Yulia Tsvetkov, Alane Suhr	42087
Bayes Conditional Distribution Estimation for Knowledge Distillation Based on Conditional Mutual Information Linfeng Ye, Shayan Mohajer Hamidi, Renhao Tan, En-Hui Yang	42116

Tangent Transformers for Composition, Privacy and Removal Tian Yu Liu, Aditya Golatkar, Stefano Soatto	42149
DIFFTACTILE: A Physics-Based Differentiable Tactile Simulator for Contact-Rich Robotic Manipulation Zilin Si, Gu Zhang, Qingwei Ben, Branden Romero, Zhou Xian, Chao Liu, Chuang Gan	42167
Class Probability Matching with Calibrated Networks for Label Shift Adaption Hongwei Wen, Annika Betken, Hanyuan Hang	42187
Neural SDF Flow for 3D Reconstruction of Dynamic Scenes Wei Mao, Richard Hartley, Mathieu Salzmann, Miaomiao Liu	42218
Massively Scalable Inverse Reinforcement Learning in Google Maps Matt Barnes, Matthew Abueg, Oliver F. Lange, Matt Deeds, Jason Trader, Denali Molitor, Markus Wulfmeier, Shawn O'Banion	42239
Rethinking the Benefits of Steerable Features in 3D Equivariant Graph Neural Networks Shih-Hsin Wang, Yung-Chang Hsu, Justin Baker, Andrea L. Bertozzi, Jack Xin, Bao Wang	42261
Alleviating Exposure Bias in Diffusion Models Through Sampling with Shifted Time Steps Mingxiao Li, Tingyu Qu, Ruicong Yao, Wei Sun, Marie-Francine Moens	42293
Defining Expertise: Applications to Treatment Effect Estimation Alihan Hã¼yã¼k, Qiyao Wei, Alicia Curth, Mihaela Van Der Schaar	42316
Alt-Text with Context: Improving Accessibility for Images on Twitter Nikita Srivatsan, Sofia Samaniego, Omar Florez, Taylor Berg-Kirkpatrick	42332
Error Norm Truncation: Robust Training in the Presence of Data Noise for Text Generation Models Tianjian Li, Haoran Xu, Philipp Koehn, Daniel Khashabi, Kenton Murray	42346
Efficient Score Matching with Deep Equilibrium Layers Yuhao Huang, Qingsong Wang, Akwum Onwunta, Bao Wang	42366
Object Centric Architectures Enable Efficient Causal Representation Learning Amin Mansouri, Jason Hartford, Yan Zhang, Yoshua Bengio	42385
Data Distillation Can Be Like Vodka: Distilling More Times for Better Quality Xuxi Chen, Yu Yang, Zhangyang Wang, Baharan Mirzasoleiman	42409
SWE-Bench: Can Language Models Resolve Real-World Github Issues? Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, Karthik R Narasimhan	42422

BrainLM: A Foundation Model for Brain Activity Recordings Josue Ortega Caro, Antonio Henrique De Oliveira Fonseca, Syed A Rizvi, Matteo Rosati, Christopher Averill, James L Cross, Prateek Mittal, Emanuele Zappala, Rahul Madhav Dhodapkar, Chadi Abdallah, David Van Dijk	42473
Massive Editing for Large Language Models Via Meta Learning Chenmien Tan, Ge Zhang, Jie Fu	42485
Safe and Robust Watermark Injection with a Single OoD Image Shuyang Yu, Junyuan Hong, Haobo Zhang, Haotao Wang, Zhangyang Wang, Jiayu Zhou	42500
Few-Shot Detection of Machine-Generated Text Using Style Representations Rafael Alberto Rivera Soto, Kailin Koch, Aleem Khan, Barry Y. Chen, Marcus Bishop, Nicholas Andrews	42512
Boundary Denoising for Video Activity Localization Mengmeng Xu, Mattia Soldan, Jialin Gao, Shuming Liu, Juan-Manuel Perez-Rua, Bernard Ghanem	42533
A Probabilistic Framework for Modular Continual Learning Lazar Valkov, Akash Srivastava, Swarat Chaudhuri, Charles Sutton	42552
Poly-View Contrastive Learning Amitis Shidani, R Devon Hjelm, Jason Ramapuram, Russell Webb, Eeshan Gunesh Dhekane, Dan Busbridge	42578
Conversational Drug Editing Using Retrieval and Domain Feedback Shengchao Liu, Jiong Xiao Wang, Yijin Yang, Chengpeng Wang, Ling Liu, Hongyu Guo, Chaowei Xiao	42620
Sophia: A Scalable Stochastic Second-Order Optimizer for Language Model Pre-Training Hong Liu, Zhiyuan Li, David Leo Wright Hall, Percy Liang, Tengyu Ma	42657
Private Zeroth-Order Nonsmooth Nonconvex Optimization Qinzi Zhang, Hoang Tran, Ashok Cutkosky	42687
Fast Hyperboloid Decision Tree Algorithms Philippe Chlenski, Ethan Turok, Antonio Khalil Moretti, Itsik Pe'Er	42708
Building Cooperative Embodied Agents Modularly with Large Language Models Hongxin Zhang, Weihua Du, Jiaming Shan, Qinhong Zhou, Yilun Du, Joshua B. Tenenbaum, Tianmin Shu, Chuang Gan	42734
Tailoring Self-Rationalizers with Multi-Reward Distillation Sahana Ramnath, Brihi Joshi, Skyler Hallinan, Ximing Lu, Liunian Harold Li, Aaron Chan, Jack Hessel, Yejin Choi, Xiang Ren	42763
Learning to Act Without Actions Dominik Schmidt, Mingqi Jiang	42794

Learning Planning Abstractions from Language Weiyu Liu, Geng Chen, Joy Hsu, Jiayuan Mao, Jiajun Wu	42811
One-Shot Empirical Privacy Estimation for Federated Learning Galen Andrew, Peter Kairouz, Sewoong Oh, Alina Oprea, Hugh Brendan McMahan, Vinith Menon, Suriyakumar	42830
On the Fairness ROAD: Robust Optimization for Adversarial Debiasing Vincent Grari, Thibault Laugel, Tatsunori Hashimoto, Sylvain Lamprier, Marcin Detyniecki	42851
The Lipschitz-Variance-Margin Tradeoff for Enhanced Randomized Smoothing Blaise Delattre, Alexandre Araujo, Quentin Barthélemy, Alexandre Allauzen	42882
Remote Sensing Vision-Language Foundation Models Without Annotations Via Ground Remote Alignment Utkarsh Mall, Cheng Perng Phoo, Meilin Kelsey Liu, Carl Vondrick, Bharath Hariharan, Kavita Bala	42907
Privileged Sensing Scaffolds Reinforcement Learning Edward S. Hu, James Springer, Oleh Rybkin, Dinesh Jayaraman	42928
Reclaiming the Source of Programmatic Policies: Programmatic Versus Latent Spaces Tales Henrique Carvalho, Kenneth Tjhia, Levi Lelis	42952
An Efficient Tester-Learner for Halfspaces Aravind Gollakota, Adam Klivans, Konstantinos Stavropoulos, Arsen Vasilyan	42973
CoT3DRef: Chain-Of-Thoughts Data-Efficient 3D Visual Grounding Eslam Mohamed Bakr, Mohamed Ayman Mohamed, Mahmoud Ahmed, Habib Slim, Mohamed Elhoseiny	42997
Model Tells You What to Discard: Adaptive KV Cache Compression for LLMs Suyu Ge, Yunan Zhang, Liyuan Liu, Minjia Zhang, Jiawei Han, Jianfeng Gao	43023
Visual Data-Type Understanding Does Not Emerge from Scaling Vision-Language Models Vishaal Udandaraao, Max F Burg, Samuel Albanie, Matthias Bethge	43037
Improved Active Learning Via Dependent Leverage Score Sampling Atsushi Shimizu, Xiaou Cheng, Christopher Musco, Jonathan Weare	43076
Like Oil and Water: Group Robustness Methods and Poisoning Defenses May Be at Odds Michael-Andrei Panaitescu-Liess, Yigitcan Kaya, Sicheng Zhu, Furong Huang, Tudor Dumitras	43100
SOTOPIA: Interactive Evaluation for Social Intelligence in Language Agents Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, Maarten Sap	43122
RA-DIT: Retrieval-Augmented Dual Instruction Tuning Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Richard James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvasy, Mike Lewis, Luke Zettlemoyer, Wen-Tau Yih	43167

Debiasing Attention Mechanism in Transformer Without Demographics Shenyu Lu, Yipei Wang, Xiaoqian Wang	43192
Stack Attention: Improving the Ability of Transformers to Model Hierarchical Patterns Brian Dusell, David Chiang	43211
Time-Varying Propensity Score to Bridge the Gap Between the Past and Present Rasool Fakoor, Jonas Mueller, Zachary Chase Lipton, Pratik Chaudhari, Alex Smola	43231
Batched Low-Rank Adaptation of Foundation Models Yeming Wen, Swarat Chaudhuri	43253
Neural Network-Based Score Estimation in Diffusion Models: Optimization and Generalization Yinbin Han, Meisam Razaviyayn, Renyuan Xu	43269
Contrastive Preference Learning: Learning from Human Feedback Without Reinforcement Learning Joey Hejna, Rafael Rafailov, Harshit Sikchi, Chelsea Finn, Scott Niekum, W. Bradley Knox, Dorsa Sadigh	43308
Consistency Models as a Rich and Efficient Policy Class for Reinforcement Learning Zihan Ding, Chi Jin	43337
ExeDec: Execution Decomposition for Compositional Generalization in Neural Program Synthesis Kensen Shi, Joey Hong, Yinlin Deng, Pengcheng Yin, Manzil Zaheer, Charles Sutton	43357
DreamFlow: High-Quality text-To-3D Generation by Approximating Probability Flow Kyungmin Lee, Kihyuk Sohn, Jinwoo Shin	43387
Protein Discovery with Discrete Walk-Jump Sampling Nathan C. Frey, Dan Berenberg, Karina Zadorozhny, Joseph Kleinhenz, Julien Lafrance-Vanasse, Isidro Hotzel, Yan Wu, Stephen Ra, Richard Bonneau, Kyunghyun Cho, Andreas Loukas, Vladimir Gligorijevic, Saeed Saremi	43413
Jailbreak in Pieces: Compositional Adversarial Attacks on Multi-Modal Language Models Erfan Shayegani, Yue Dong, Nael Abu-Ghazaleh	43433
CLIP-MUSED: CLIP-Guided Multi-Subject Visual Neural Information Semantic Decoding Qiongyi Zhou, Changde Du, Shengpei Wang, Huiguang He	43466
LLMs Meet VLMs: Boost Open Vocabulary Object Detection with Fine-Grained Descriptors Sheng Jin, Xueying Jiang, Jiaxing Huang, Lewei Lu, Shijian Lu	43483
WizardLM: Empowering Large Pre-Trained Language Models to Follow Complex Instructions Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, Daxin Jiang	43495

Waxing-And-Waning: A Generic Similarity-Based Framework for Efficient Self-Supervised Learning Sheng Li, Chao Wu, Ao Li, Yanzhi Wang, Xulong Tang, Geng Yuan	43517
Learning Hierarchical Image Segmentation for Recognition and by Recognition Tsung-Wei Ke, Sangwoo Mo, Stella X. Yu	43535
WildChat: 1M ChatGPT Interaction Logs in the Wild Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, Yuntian Deng	43574
Understanding and Mitigating the Label Noise in Pre-Training on Downstream Tasks Hao Chen, Jindong Wang, Ankit Shah, Ran Tao, Hongxin Wei, Xing Xie, Masashi Sugiyama, Bhiksha Raj	43590
Tackling the Data Heterogeneity in Asynchronous Federated Learning with Cached Update Calibration Yujia Wang, Yuanpu Cao, Jingcheng Wu, Ruoyu Chen, Jinghui Chen	43621
Steve-Eye: Equipping LLM-Based Embodied Agents with Visual Perception in Open Worlds Sipeng Zheng, Jiazheng Liu, Yicheng Feng, Zongqing Lu	43654
Online GNN Evaluation Under Test-Time Graph Distribution Shifts Xin Zheng, Dongjin Song, Qingsong Wen, Bo Du, Shirui Pan	43675
PandaLM: An Automatic Evaluation Benchmark for LLM Instruction Tuning Optimization Yidong Wang, Zhuohao Yu, Wenjin Yao, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang, Rui Xie, Jindong Wang, Xing Xie, Wei Ye, Shikun Zhang, Yue Zhang	43697
Communication-Efficient Federated Non-Linear Bandit Optimization Chuanhao Li, Chong Liu, Yu-Xiang Wang	43718
DREAM: Dual Structured Exploration with Mixup for Open-Set Graph Domain Adaption Nan Yin, Mengzhu Wang, Zhenghan Chen, Li Shen, Huan Xiong, Bin Gu, Xiao Luo	43741
Towards Aligned Layout Generation Via Diffusion Model with Aesthetic Constraints Jian Chen, Ruiyi Zhang, Yufan Zhou, Changyou Chen	43763
Demystifying Linear MDPs and Novel Dynamics Aggregation Framework Joongkyu Lee, Min-Hwan Oh	43781
Treatment Effects Estimation by Uniform Transformer Ruoqi Yu, Shulei Wang	43811
CoVLM: Composing Visual Entities and Relationships in Large Language Models Via Communicative Decoding Junyan Li, Delin Chen, Yining Hong, Zhenfang Chen, Peihao Chen, Yikang Shen, Chuang Gan	43841
Deep Temporal Graph Clustering Meng Liu, Yue Liu, Ke Liang, Wenxuan Tu, Siwei Wang, Sihang Zhou, Xinwang Liu	43859

LQ-LoRA: Low-Rank Plus Quantized Matrix Decomposition for Efficient Language Model Finetuning Han Guo, Philip Greengard, Eric Xing, Yoon Kim	43877
Spectrally Transformed Kernel Regression Runtian Zhai, Rattana Pukdee, Roger Jin, Maria Florina Balcan, Pradeep Kumar Ravikumar	43895
Proper Laplacian Representation Learning Diego Gomez, Michael Bowling, Marlos C. Machado	43931
A Sublinear Adversarial Training Algorithm Yeqi Gao, Lianke Qin, Zhao Song, Yitan Wang	43955
SEGNO: Generalizing Equivariant Graph Neural Networks with Physical Inductive Biases Yang Liu, Jiashun Cheng, Haihong Zhao, Tingyang Xu, Peilin Zhao, Fugee Tsung, Jia Li, Yu Rong	43989
Language Model Beats Diffusion - Tokenizer is Key to Visual Generation Lijun Yu, Jose Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, Boqing Gong, Ming-Hsuan Yang, Irfan Essa, David A Ross, Lu Jiang	44013
Skeleton-Of-Thought: Prompting LLMs for Efficient Parallel Generation Xuefei Ning, Zinan Lin, Zixuan Zhou, Zifu Wang, Huazhong Yang, Yu Wang	44032
Deep Confident Steps to New Pockets: Strategies for Docking Generalization Gabriele Corso, Arthur Deng, Nicholas Polizzi, Regina Barzilay, Tommi S. Jaakkola	44083
The Human-AI Substitution Game: Active Learning from a Strategic Labeler Tom Yan, Chicheng Zhang	44104
GraphCare: Enhancing Healthcare Predictions with Personalized Knowledge Graphs Pengcheng Jiang, Cao Xiao, Adam Richard Cross, Jimeng Sun	44163
$\mathbb{D}^2$ Pruning: Message Passing for Balancing Diversity & Difficulty in Data Pruning Adyasha Maharana, Prateek Yadav, Mohit Bansal	44191
BayesDiff: Estimating Pixel-Wise Uncertainty in Diffusion Via Bayesian Inference Siqi Kou, Lei Gan, Dequan Wang, Chongxuan Li, Zhijie Deng	44211
MathVista: Evaluating Mathematical Reasoning of Foundation Models in Visual Contexts Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, Jianfeng Gao	44229
A Unique M-Pattern for Micro-Expression Spotting in Long Videos Jinxuan Wang, Shiting Xu, Tong Zhang	44345
Pre-Training Sequence, Structure, and Surface Features for Comprehensive Protein Representation Learning Youhan Lee, Hasun Yu, Jaemyung Lee, Jaehoon Kim	44362

Learning Grounded Action Abstractions from Language Lionel Wong, Jiayuan Mao, Pratyusha Sharma, Zachary S Siegel, Jiahai Feng, Noa Korneev, Joshua B. Tenenbaum, Jacob Andreas	44379
On Bias-Variance Alignment in Deep Models Lin Chen, Michal Lukasik, Wittawat Jitkrittum, Chong You, Sanjiv Kumar	44409
Learning Hierarchical Polynomials with Three-Layer Neural Networks Zihao Wang, Eshaan Nichani, Jason D. Lee	44438
Achieving Sample and Computational Efficient Reinforcement Learning by Action Space Reduction Via Grouping Yining Li, Peizhong Ju, Ness Shroff	44481
Negatively Correlated Ensemble Reinforcement Learning for Online Diverse Game Level Generation Ziqi Wang, Chengpeng Hu, Jialin Liu, Xin Yao	44515
Adaptive Regret for Bandits Made Possible: Two Queries Suffice Zhou Lu, Qiuyu Zhang, Xinyi Chen, Fred Zhang, David Woodruff, Elad Hazan	44547
Merge, Then Compress: Demystify Efficient SMOE with Hints from Its Routing Policy Pingzhi Li, Zhenyu Zhang, Prateek Yadav, Yi-Lin Sung, Yu Cheng, Mohit Bansal, Tianlong Chen	44558
Making RL with Preference-Based Feedback Efficient Via Randomization Runzhe Wu, Wen Sun	44581
Beyond Weisfeiler-Lehman: A Quantitative Framework for GNN Expressiveness Bohang Zhang, Jingchu Gai, Yiheng Du, Qiwei Ye, Di He, Liwei Wang	44638
Generating Pragmatic Examples to Train Neural Program Synthesizers Saujas Vaduguru, Daniel Fried, Yewen Pu	44711
Generative Modeling of Regular and Irregular Time Series Data Via Koopman VAEs Ilan Naiman, N. Benjamin Erichson, Pu Ren, Michael W. Mahoney, Omri Azencot	44729
Beyond Worst-Case Attacks: Robust RL with Adaptive Defense Via Non-Dominated Policies Xiangyu Liu, Chenghao Deng, Yanchao Sun, Yongyuan Liang, Furong Huang	44756
Lipsum-FT: Robust Fine-Tuning of Zero-Shot Models Using Random Text Guidance Giung Nam, Byeongho Heo, Juho Lee	44779
MUFFIN: Curating Multi-Faceted Instructions for Improving Instruction Following Renze Lou, Kai Zhang, Jian Xie, Yuxuan Sun, Janice Ahn, Hanzi Xu, Yu Su, Wenpeng Yin	44813
Concept Bottleneck Generative Models Aya Abdelsalam Ismail, Julius Adebayo, Hector Corrada Bravo, Stephen Ra, Kyunghyun Cho	44843

Robustifying and Boosting Training-Free Neural Architecture Search Zhenfeng He, Yao Shu, Zhongxiang Dai, Bryan Kian Hsiang Low	44862
Score Regularized Policy Optimization Through Diffusion Behavior Huayu Chen, Cheng Lu, Zhengyi Wang, Hang Su, Jun Zhu	44890
Goodhart's Law in Reinforcement Learning Jacek Karwowski, Oliver Hayman, Xingjian Bai, Klaus Kiendlhofer, Charlie Griffin, Joar Max Viktor Skalse	44910
Hyper Evidential Deep Learning to Quantify Composite Classification Uncertainty Changbin Li, Kangshuo Li, Yuzhe Ou, Lance M. Kaplan, Audun J,sang, Jin-Hee Cho, Dong Hyun Jeong, Feng Chen	44941
Grounding Language Plans in Demonstrations Through Counterfactual Perturbations Yanwei Wang, Tsun-Hsuan Wang, Jiayuan Mao, Michael Hagenow, Julie Shah	44976
Vanishing Gradients in Reinforcement Finetuning of Language Models Noam Razin, Hattie Zhou, Omid Saremi, Vimal Thilak, Arwen Bradley, Preetum Nakkiran, Joshua M. Susskind, Etai Littwin	44993
GENOME: Generative Neuro-Symbolic Visual Reasoning by Growing and Reusing Modules Zhenfang Chen, Rui Sun, Wenjun Liu, Yining Hong, Chuang Gan	45030
COLLIE: Systematic Construction of Constrained Text Generation Tasks Shunyu Yao, Howard Chen, Austin W. Hanjie, Runzhe Yang, Karthik R Narasimhan	45058
Understanding Expressivity of GNN in Rule Learning Haiquan Qiu, Yongqi Zhang, Yong Li, Quanming Yao	45070
Unmasking and Improving Data Credibility: A Study with Datasets for Training Harmless Language Models Zhaowei Zhu, Jialu Wang, Hao Cheng, Yang Liu	45094
Correlated Noise Provably Beats Independent Noise for Differentially Private Learning Christopher A. Choquette-Choo, Krishnamurthy Dj Dvijotham, Krishna Pillutla, Arun Ganesh, Thomas Steinke, Abhradeep Guha Thakurta	45121
Image Clustering Conditioned on Text Criteria Sehyun Kwon, Jaeseung Park, Minkyu Kim, Jaewoong Cho, Ernest K. Ryu, Kangwook Lee	45182
Non-Negative Contrastive Learning Yifei Wang, Qi Zhang, Yaoyu Guo, Yisen Wang	45222
Understanding Domain Generalization: A Noise Robustness Perspective Rui Qiao, Bryan Kian Hsiang Low	45244
DataInf: Efficiently Estimating Data Influence in LoRA-Tuned LLMs and Diffusion Models Yongchan Kwon, Eric Wu, Kevin Wu, James Zou	45274

Flag Aggregator: Scalable Distributed Training Under Failures and Augmented Losses Using Convex Optimization Hamidreza Almasi, Harsh Mishra, Balajee Vamanan, Sathya N. Ravi	45296
Neuro-Inspired Information-Theoretic Hierarchical Perception for Multimodal Learning Xiongye Xiao, Gengshuo Liu, Gaurav Gupta, Defu Cao, Shixuan Li, Yaxing Li, Tianqing Fang, Mingxi Cheng, Paul Bogdan	45322
Phenomenal Yet Puzzling: Testing Inductive Reasoning Capabilities of Language Models with Hypothesis Refinement Linlu Qiu, Liwei Jiang, Ximing Lu, Melanie Sclar, Valentina Pyatkin, Chandra Bhagavatula, Bailin Wang, Yoon Kim, Yejin Choi, Nouha Dziri, Xiang Ren	45351
Training-Free Multi-Objective Diffusion Model for 3D Molecule Generation Xu Han, Caihua Shan, Yifei Shen, Can Xu, Han Yang, Xiang Li, Dongsheng Li	45381
Beyond Vanilla Variational Autoencoders: Detecting Posterior Collapse in Conditional and Hierarchical Variational Autoencoders Hien Dang, Tho Tran Huu, Tan Minh Nguyen, Nhat Ho	45397
What Algorithms Can Transformers Learn? a Study in Length Generalization Hattie Zhou, Arwen Bradley, Etai Littwin, Noam Razin, Omid Saremi, Joshua M. Susskind, Samy Bengio, Preetum Nakkiran	45436
Ensemble Distillation for Unsupervised Constituency Parsing Behzad Shayegh, Yanshuai Cao, Xiaodan Zhu, Jackie Ck Cheung, Lili Mou	45465
ToolChain*: Efficient Action Space Navigation in Large Language Models with A* Search Yuchen Zhuang, Xiang Chen, Tong Yu, Saayan Mitra, Victor Bursztyn, Ryan A. Rossi, Somdeb Sarkhel, Chao Zhang	45483
Enhancing Transfer Learning with Flexible Nonparametric Posterior Sampling Hyungi Lee, Giung Nam, Edwin Fong, Juho Lee	45509
Provably Efficient Iterated CVaR Reinforcement Learning with Function Approximation and Human Feedback Yu Chen, Yihan Du, Pihe Hu, Siwei Wang, Desheng Wu, Longbo Huang	45533
The Devil is in the Neurons: Interpreting and Mitigating Social Biases in Language Models Yan Liu, Yu Liu, Xiaokang Chen, Pin-Yu Chen, Daoguang Zan, Min-Yen Kan, Tsung-Yi Ho	45571
Statistical Perspective of Top-K Sparse Softmax Gating Mixture of Experts Huy Nguyen, Pedram Akbarian, Fanqi Yan, Nhat Ho	45583
Confidence-Aware Reward Optimization for Fine-Tuning Text-To-Image Models Kyuyoung Kim, Jongheon Jeong, Minyong An, Mohammad Ghavamzadeh, Krishnamurthy Dj Dvijotham, Jinwoo Shin, Kimin Lee	45621

Dynamic Neural Response Tuning Tian Qiu, Wenxiang Xu, Lin Chen, Linyun Zhou, Zunlei Feng, Mingli Song	45641
On the Stability of Expressive Positional Encodings for Graphs Yinan Huang, William Lu, Joshua Robinson, Yu Yang, Muhan Zhang, Stefanie Jegelka, Pan Li	45657
Spatially-Aware Transformers for Embodied Agents Junmo Cho, Jaesik Yoon, Sungjin Ahn	45687
Grounded Object-Centric Learning Avinash Kori, Francesco Locatello, Fabio De Sousa Ribeiro, Francesca Toni, Ben Glocker	45714
Adversarial AutoMixup Huafeng Qin, Xin Jin, Yun Jiang, Mounâ@m El-Yacoubi, Xinbo Gao	45759
Mixed-Type Tabular Data Synthesis with Score-Based Diffusion in Latent Space Hengrui Zhang, Jiani Zhang, Zhengyuan Shen, Balasubramaniam Srinivasan, Xiao Qin, Christos Faloutsos, Huzefa Rangwala, George Karypis	45778
GNNX-BENCH: Unravelling the Utility of Perturbation-Based GNN Explainers Through In-Depth Benchmarking Mert Kosan, Samidha Verma, Burouj Armgaan, Khushbu Pahwa, Ambuj Singh, Sourav Medya, Sayan Ranu	45807
Explaining Time Series Via Contrastive and Locally Sparse Perturbations Zichuan Liu, Yingying Zhang, Tianchun Wang, Zefan Wang, Dongsheng Luo, Mengnan Du, Min Wu, Yi Wang, Chunlin Chen, Lunting Fan, Qingsong Wen	45837
SEINE: Short-To-Long Video Diffusion Model for Generative Transition and Prediction Xinyuan Chen, Yaohui Wang, Lingjun Zhang, Shaobin Zhuang, Xin Ma, Jiashuo Yu, Yali Wang, Dahua Lin, Yu Qiao, Ziwei Liu	45858
Fake it Till Make It: Federated Learning with Consensus-Oriented Generation Rui Ye, Yaxin Du, Zhenyang Ni, Yanfeng Wang, Siheng Chen	45873
Certified Adversarial Robustness for Rate Encoded Spiking Neural Networks Bhaskar Mukhoty, Hilal Alquabeh, Giulia De Masi, Huan Xiong, Bin Gu	45899
Language-Interfaced Tabular Oversampling Via Progressive Imputation and Self-Authentication June Yong Yang, Geondo Park, Joowon Kim, Hyeongwon Jang, Eunho Yang	45916
Are Transformers with One Layer Self-Attention Using Low-Rank Weight Matrices Universal Approximators? Tokio Kajitsuka, Issei Sato	45939
Is Attention Required for ICL? Exploring the Relationship Between Model Architecture and In-Context Learning Ability Ivan Lee, Nan Jiang, Taylor Berg-Kirkpatrick	45968

Learning Thresholds with Latent Values and Censored Feedback Jiahao Zhang, Tao Lin, Weiqiang Zheng, Zhe Feng, Yifeng Teng, Xiaotie Deng	45998
Sum-Product-Set Networks: Deep Tractable Models for Tree-Structured Graphs Milan Papez, Martin Rektoris, Vaclav Smidl, Tomáš Pevný	46015
Looped Transformers Are Better at Learning Learning Algorithms Liu Yang, Kangwook Lee, Robert D Nowak, Dimitris Papailiopoulos	46045
Whittle Index with Multiple Actions and State Constraint for Inventory Management Chuheng Zhang, Xiangsen Wang, Wei Jiang, Xianliang Yang, Siwei Wang, Lei Song, Jiang Bian	46065
MetaCoCo: A New Few-Shot Classification Benchmark with Spurious Correlation Min Zhang, Haoxuan Li, Fei Wu, Kun Kuang	46084
Demystifying CLIP Data Hu Xu, Saining Xie, Xiaoqing Tan, Po-Yao Huang, Russell Howes, Vasu Sharma, Shang-Wen Li, Gargi Ghosh, Luke Zettlemoyer, Christoph Feichtenhofer	46111
SafeDreamer: Safe Reinforcement Learning with World Models Weidong Huang, Jiaming Ji, Chunhe Xia, Borong Zhang, Yaodong Yang	46131
PARL: A Unified Framework for Policy Alignment in Reinforcement Learning from Human Feedback Souradip Chakraborty, Amrit Bedi, Alec Koppel, Huazheng Wang, Dinesh Manocha, Mengdi Wang, Furong Huang	46162
CAS: A Probability-Based Approach for Universal Condition Alignment Score Chunsan Hong, Byunghee Cha, Tae-Hyun Oh	46202
Backdoor Contrastive Learning Via Bi-Level Trigger Optimization Weiyu Sun, Xinyu Zhang, Hao Lu, Ying-Cong Chen, Ting Wang, Jinghui Chen, Lu Lin	46224
How to Fine-Tune Vision Models with SGD Ananya Kumar, Ruoqi Shen, Sebastien Bubeck, Suriya Gunasekar	46244
Evaluating Large Language Models at Evaluating Instruction Following Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, Danqi Chen	46272
Contextual Bandits with Online Neural Regression Rohan Deb, Yikun Ban, Shiliang Zuo, Jingrui He, Arindam Banerjee	46299
Bootstrapping Variational Information Pursuit with Large Language and Vision Models for Interpretable Image Classification Aditya Chattopadhyay, Kwan Ho Ryan Chan, Rene Vidal	46343
Plug-And-Play Policy Planner for Large Language Model Powered Dialogue Agents Yang Deng, Wenxuan Zhang, Wai Lam, See-Kiong Ng, Tat-Seng Chua	46370

Kill Two Birds with One Stone: Rethinking Data Augmentation for Deep Long-Tailed Learning <i>Binwu Wang, Pengkun Wang, Wei Xu, Xu Wang, Yudong Zhang, Kun Wang, Yang Wang</i>	46403
In Defense of Parameter Sharing for Model-Compression <i>Aditya Desai, Anshumali Shrivastava</i>	46427
Towards Understanding Factual Knowledge of Large Language Models <i>Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, Philip S. Yu, Zhijiang Guo</i>	46452
Beating Price of Anarchy and Gradient Descent Without Regret in Potential Games <i>Iosif Sakos, Stefanos Leonardos, Stelios Andrew Stavroulakis, William Overman, Ioannis Panageas, Georgios Piliouras</i>	46488
Understanding Addition in Transformers <i>Philip Quirke, Fazl Barez</i>	46511
Beyond IID Weights: Sparse and Low-Rank Deep Neural Networks Are Also Gaussian Processes <i>Thiziri Nait Saada, Alireza Naderi, Jared Tanner</i>	46532
Harnessing Joint Rain-/Detail-Aware Representations to Eliminate Intricate Rains <i>Wu Ran, Peirong Ma, Zhiquan He, Hao Ren, Hong Lu</i>	46561
Differentially Private Synthetic Data Via Foundation Model APIs 1: Images <i>Zinan Lin, Sivakanth Gopi, Janardhan Kulkarni, Harsha Nori, Sergey Yekhanin</i>	46582
Image Background Serves as Good Proxy for Out-Of-Distribution Data <i>Sen Pei</i>	46633
AgentBench: Evaluating LLMs as Agents <i>Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, Jie Tang</i>	46649
Adaptive Stochastic Gradient Algorithm for Black-Box Multi-Objective Learning <i>Feiyang Ye, Yueming Lyu, Xuehao Wang, Yu Zhang, Ivor Tsang</i>	46707
Training Graph Transformers Via Curriculum-Enhanced Attention Distillation <i>Yisong Huang, Jin Li, Xinlong Chen, Yang-Geng Fu</i>	46734
Nemesis: Normalizing the Soft-Prompt Vectors of Vision-Language Models <i>Shuai Fu, Xiequn Wang, Qiusi Huang, Yu Zhang</i>	46757
Identifying Policy Gradient Subspaces <i>Jan Schneider, Pierre Schumacher, Simon Guist, Le Chen, Daniel Haeufle, Bernhard Schölkopf, Dieter Borchert</i>	46781
Asymptotically Free Sketched Ridge Ensembles: Risks, Cross-Validation, and Tuning <i>Pratik Patil, Daniel Lejeune</i>	46803

Yet Another ICU Benchmark: A Flexible Multi-Center Framework for Clinical ML Robin Van De Water, Hendrik Nils Aurel Schmidt, Paul Elbers, Patrick Thorat, Bert Arnrich, Patrick Rockenschaub	46839
Controlling Vision-Language Models for Multi-Task Image Restoration Ziwei Luo, Fredrik K. Gustafsson, Zheng Zhao, Jens Sjöstrand, Thomas B. Schön	46895
Label-Agnostic Forgetting: A Supervision-Free Unlearning in Deep Models Shaofei Shen, Chenhao Zhang, Yawen Zhao, Alina Bialkowski, Weitong Tony Chen, Miao Xu	46916
VersVideo: Leveraging Enhanced Temporal Diffusion Models for Versatile Video Generation Jinxi Xiang, Ricong Huang, Jun Zhang, Guanbin Li, Xiao Han, Yang Wei	46935
Analyzing Feed-Forward Blocks in Transformers Through the Lens of Attention Maps Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, Kentaro Inui	46954
Making Retrieval-Augmented Language Models Robust to Irrelevant Context Ori Yoran, Tomer Wolfson, Ori Ram, Jonathan Berant	46991
On the Role of General Function Approximation in Offline Reinforcement Learning Chenjie Mao, Qiaosheng Zhang, Zhen Wang, Xuelong Li	47013
Evaluating Language Model Agency Through Negotiations Tim Ruben Davidson, Veniamin Veselovsky, Michal Kosinski, Robert West	47048
Some Fundamental Aspects About Lipschitz Continuity of Neural Networks Grigory Khromov, Sidak Pal Singh	47074
BroGNet: Momentum-Conserving Graph Neural Stochastic Differential Equation for Learning Brownian Dynamics Suresh Bishnoi, Jayadeva Jayadeva, Sayan Ranu, N M Anoop Krishnan	47119
Efficient Continual Finite-Sum Minimization Ioannis Mavrothalassitis, Stratis Skoulakis, Leello Tadesse Dadi, Volkan Cevher	47142
Attention-Based Iterative Decomposition for Tensor Product Representation Taewon Park, Inchul Choi, Minho Lee	47177
Querying Easily Flip-Flopped Samples for Deep Active Learning Seong Jin Cho, Gwangsu Kim, Junghyun Lee, Jinwoo Shin, Chang D. Yoo	47201
ClimODE: Climate and Weather Forecasting with Physics-Informed Neural ODEs Yogesh Verma, Markus Heinonen, Vikas Garg	47235
Prometheus: Inducing Fine-Grained Evaluation Capability in Language Models Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoo Yun, Seongjin Shin, Sungdong Kim, James Thorne, Minjoon Seo	47258

One Step of Gradient Descent is Provably the Optimal In-Context Learner with One Layer of Linear Self-Attention Arvind V. Mahankali, Tatsunori Hashimoto, Tengyu Ma	47294
Advancing Pose-Guided Image Synthesis with Progressive Conditional Diffusion Models Fei Shen, Hu Ye, Jun Zhang, Cong Wang, Xiao Han, Yang Wei	47314
PAE: Reinforcement Learning from External Knowledge for Efficient Exploration Zhe Wu, Hao-fei Lu, Junliang Xing, You Wu, Renye Yan, Yaozhong Gan, Yuanchun Shi	47337
Deep Orthogonal Hypersphere Compression for Anomaly Detection Yunhe Zhang, Yan Sun, Jinyu Cai, Jicong Fan	47362
Learning Personalized Causally Invariant Representations for Heterogeneous Federated Clients Xueyang Tang, Song Guo, Jie Zhang, Jingcai Guo	47387
ED-NeRF: Efficient Text-Guided Editing of 3D Scene with Latent Space NeRF Jangho Park, Gihyun Kwon, Jong Chul Ye	47409
Topic Modeling as Multi-Objective Contrastive Optimization Thong Thanh Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy T Nguyen, See-Kiong Ng, Anh Tuan Luu	47430
Kernelised Normalising Flows Eshant English, Matthias Kirchler, Christoph Lippert	47450
Accelerating Data Generation for Neural Operators Via Krylov Subspace Recycling Hong Wang, Zhongkai Hao, Jie Wang, Zijie Geng, Zhen Wang, Bin Li, Feng Wu	47471
Flow to Better: Offline Preference-Based Reinforcement Learning Via Preferred Trajectory Generation Zhilong Zhang, Yihao Sun, Junyin Ye, Tian-Shuo Liu, Jiaji Zhang, Yang Yu	47511
Unlocking the Power of Representations in Long-Term Novelty-Based Exploration Alaa Saade, Steven Kapturowski, Daniele Calandriello, Charles Blundell, Pablo Sprechmann, Leopoldo Sarra, Oliver Groth, Michal Valko, Bilal Piot	47530
Compositional Conservatism: A Transductive Approach in Offline Reinforcement Learning Yeda Song, Dongwook Lee, Gunhee Kim	47559
H-GAP: Humanoid Control with a Generalist Planner Zhengyao Jiang, Yingchen Xu, Nolan Wagener, Yicheng Luo, Michael Janner, Edward Grefenstette, Tim Rocktäschel, Yuandong Tian	47578
CoRe-GD: A Hierarchical Framework for Scalable Graph Visualization with GNNs Florian Gräßl, Joël Mathys, Robert Veres, Roger Wattenhofer	47596

Zero Bubble (Almost) Pipeline Parallelism Penghui Qi, Xinyi Wan, Guangxing Huang, Min Lin	47619
Exploring Weight Balancing on Long-Tailed Recognition Problem Naoya Hasegawa, Issei Sato	47635
Blending Imitation and Reinforcement Learning for Robust Policy Improvement Xuefeng Liu, Takuma Yoneda, Rick Stevens, Matthew Walter, Yuxin Chen	47663
ASMR: Activation-Sharing Multi-Resolution Coordinate Networks for Efficient Inference Jason Chun Lok Li, Steven Tin Sui Luo, Le Xu, Ngai Wong	47685
Do Generated Data Always Help Contrastive Learning? Yifei Wang, Jizhe Zhang, Yisen Wang	47706
Safe Offline Reinforcement Learning with Feasibility-Guided Diffusion Model Yinan Zheng, Jianxiong Li, Dongjie Yu, Yujie Yang, Shengbo Eben Li, Xianyuan Zhan, Jingjing Liu	47725
Identifying Representations for Intervention Extrapolation Sorawit Saengkyongam, Elan Rosenfeld, Pradeep Kumar Ravikumar, Niklas Pfister, Jonas Peters	47755
Hybrid LLM: Cost-Efficient and Quality-Aware Query Routing Dujian Ding, Ankur Mallick, Chi Wang, Robert Sim, Subhabrata Mukherjee, Victor Răzile, Laks V. S. Lakshmanan, Ahmed Hassan Awadallah	47781
Understanding the Robustness of Randomized Feature Defense Against Query-Based Adversarial Attacks Nguyen Hung-Quang, Yingjie Lao, Tung Pham, Kok-Seng Wong, Khoa D Doan	47800
Efficient Episodic Memory Utilization of Cooperative Multi-Agent Reinforcement Learning Hyungho Na, Yunkyeong Seo, Il-Chul Moon	47821
CrossQ: Batch Normalization in Deep Reinforcement Learning for Greater Sample Efficiency and Simplicity Aditya Bhatt, Daniel Palenicek, Boris Belousov, Max Argus, Artemij Amiranashvili, Thomas Brox, Jan Peters	47854
GoLLIE: Annotation Guidelines Improve Zero-Shot Information-Extraction Oscar Sainz, Iker García-Ferrero, Rodrigo Agerri, Oier Lopez De Lacalle, German Rigau, Eneko Agirre	47873
Leveraging Hyperbolic Embeddings for Coarse-To-Fine Robot Design Heng Dong, Junyu Zhang, Chongjie Zhang	47898
Benchmarking Algorithms for Federated Domain Generalization Ruqi Bai, Saurabh Bagchi, David I. Inouye	47926

CircuitNet 2.0: An Advanced Dataset for Promoting Machine Learning Innovations in Realistic Chip Design Environment Xun Jiang, Zhuomin Chai, Yuxiang Zhao, Yibo Lin, Runsheng Wang, Ru Huang	47952
Frequency-Aware Transformer for Learned Image Compression Han Li, Shaohui Li, Wenrui Dai, Chenglin Li, Junni Zou, Hongkai Xiong	47972
Safe RLHF: Safe Reinforcement Learning from Human Feedback Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, Yaodong Yang	47991
Reward Model Ensembles Help Mitigate Overoptimization Thomas Coste, Usman Anwar, Robert Kirk, David Krueger	48019
CABINET: Content Relevance-Based Noise Reduction for Table Question Answering Sohan Patnaik, Heril Changwal, Milan Aggarwal, Sumit Bhatia, Yaman Kumar, Balaji Krishnamurthy	48046
Denoising Task Routing for Diffusion Models Byeongjun Park, Sangmin Woo, Hyojun Go, Jin-Young Kim, Changick Kim	48070
DOS: Diverse Outlier Sampling for Out-Of-Distribution Detection Wenyu Jiang, Hao Cheng, Mingcai Chen, Chongjun Wang, Hongxin Wei	48095
Multi-View Causal Representation Learning with Partial Observability Dingling Yao, Danru Xu, Sebastien Lachapelle, Sara Magliacane, Perouz Taslakian, Georg Martius, Julius Von K��gelgen, Francesco Locatello	48111
Global Optimality for Non-Linear Constrained Restoration Problems Via Invexity Samuel Pinilla, Jeyan Thiyagalingam	48143
Can We Evaluate Domain Adaptation Models Without Target-Domain Labels? Jianfei Yang, Hanjie Qian, Yuecong Xu, Kai Wang, Lihua Xie	48181
LLM-CXR: Instruction-Finetuned LLM for CXR Image Understanding and Generation Suhyeon Lee, Won Jun Kim, Jinho Chang, Jong Chul Ye	48202
Stabilizing Contrastive RL: Techniques for Robotic Goal Reaching from Offline Data Chongyi Zheng, Benjamin Eysenbach, Homer Rich Walke, Patrick Yin, Kuan Fang, Ruslan Salakhutdinov, Sergey Levine	48223
Navigating the Design Space of Equivariant Diffusion-Based Generative Models for De Novo 3D Molecule Generation Tuan Le, Julian Cremer, Frank Noe, Djork-Arn�� Clevert, Kristof T Sch��tt	48252
Hindsight PRIORs for Reward Learning from Human Preferences Mudit Verma, Katherine Metcalf	48275
Solving Homogeneous and Heterogeneous Cooperative Tasks with Greedy Sequential Execution Shanqi Liu, Dong Xing, Pengjie Gu, Xinrun Wang, Bo An, Yong Liu	48294

Conserve-Update-Revise to Cure Generalization and Robustness Trade-Off in Adversarial Training Shruthi Gowda, Bahram Zonooz, Elahe Arani	48311
Traveling Waves Encode the Recent Past and Enhance Sequence Learning T. Anderson Keller, Lyle Muller, Terrence Sejnowski, Max Welling	48332
Mitigating Emergent Robustness Degradation While Scaling Graph Learning Xiangchi Yuan, Chunhui Zhang, Yijun Tian, Yanfang Ye, Chuxu Zhang	48359
InstructCV: Instruction-Tuned Text-To-Image Diffusion Models as Vision Generalists Yulu Gan, Sungwoo Park, Alexander Marcel Schubert, Anthony Philippakis, Ahmed Alaa	48383
Customizable Combination of Parameter-Efficient Modules for Multi-Task Learning Haowen Wang, Tao Sun, Congyun Jin, Yingbo Wang, Yibo Fan, Yunqi Xu, Yuliang Du, Cong Fan	48400
FairSeg: A Large-Scale Medical Image Segmentation Dataset for Fairness Learning Using Segment Anything Model with Fair Error-Bound Scaling Yu Tian, Min Shi, Yan Luo, Ava Kouhana, Tobias Elze, Mengyu Wang	48422
Logical Languages Accepted by Transformer Encoders with Hard Attention Pablo Barcelo, Alexander Kozachinskiy, Anthony Widjaja Lin, Vladimir Podolskii	48437
GNNCert: Deterministic Certification of Graph Neural Networks Against Adversarial Perturbations Zaishuo Xia, Han Yang, Binghui Wang, Jinyuan Jia	48448
Implicit Gaussian Process Representation of Vector Fields Over Arbitrary Latent Manifolds Robert Peach, Matteo Vaino-Carl, Nir Grossman, Michael David, Emma Mallas, David J. Sharp, Paresh A. Malhotra, Pierre Vandergheynst, Adam Gosztolai	48466
CLEX: Continuous Length Extrapolation for Large Language Models Guanzheng Chen, Xin Li, Zaiqiao Meng, Shangsong Liang, Lidong Bing	48482
EmerDiff: Emerging Pixel-Level Semantic Knowledge in Diffusion Models Koichi Namekata, Amirmojtaba Sabour, Sanja Fidler, Seung Wook Kim	48500
Graph Neural Networks for Learning Equivariant Representations of Neural Networks Miltiadis Kofinas, Boris Knyazev, Yan Zhang, Yunlu Chen, Gertjan J. Burghouts, Efstratios Gavves, Cees G. M. Snoek, David W. Zhang	48523
Expressivity of ReLU-Networks Under Convex Relaxations Maximilian Baader, Mark Niklas Mueller, Yuhao Mao, Martin Vechev	48542
Adversarial Feature Map Pruning for Backdoor Dong Huang, Qingwen Bu	48567

PolyVoice: Language Models for Speech to Speech Translation Qian Qian Dong, Zhiying Huang, Qiao Tian, Chen Xu, Tom Ko, Yunlong Zhao, Siyuan Feng, Tang Li, Kexin Wang, Xuxin Cheng, Fengpeng Yue, Ye Bai, Xi Chen, Lu Lu, Zejun Ma, Yuping Wang, Mingxuan Wang, Yuxuan Wang	48583
UniTabE: A Universal Pretraining Protocol for Tabular Foundation Model in Data Science Yazheng Yang, Yuqi Wang, Guang Liu, Ledell Wu, Qi Liu	48597
PolyGCL: GRAPH CONTRASTIVE LEARNING Via Learnable Spectral Polynomial Filters Jingyu Chen, Runlin Lei, Zhewei Wei	48613
Catastrophic Jailbreak of Open-Source LLMs Via Exploiting Generation Yangsibo Huang, Samyak Gupta, Mengzhou Xia, Kai Li, Danqi Chen	48643
Scaling Convex Neural Networks with Burer-Monteiro Factorization Arda Sahiner, Tolga Ergen, Batu Ozturkler, John M. Pauly, Morteza Mardani, Mert Pilanci	48664
LoftQ: LoRA-Fine-Tuning-Aware Quantization for Large Language Models Yixiao Li, Yifan Yu, Chen Liang, Nikos Karampatziakis, Pengcheng He, Weizhu Chen, Tuo Zhao	48695
CoLiDE: Concomitant Linear DAG Estimation Seyed Saman Saboksayr, Gonzalo Mateos, Mariano Tepper	48711
REBAR: Retrieval-Based Reconstruction for Time-Series Contrastive Learning Maxwell Xu, Alexander Moreno, Hui Wei, Benjamin Marlin, James Matthew Rehg	48740
Escape Sky-High Cost: Early-Stopping Self-Consistency for Multi-Step Reasoning Yiwei Li, Peiwen Yuan, Shaoxiong Feng, Boyuan Pan, Xinglin Wang, Bin Sun, Heda Wang, Kan Li	48765
Achieving Fairness in Multi-Agent MDP Using Reinforcement Learning Peizhong Ju, Arnob Ghosh, Ness Shroff	48783
A Policy Gradient Method for Confounded POMDPs Mao Hong, Zhengling Qi, Yanxun Xu	48817
HiGen: Hierarchical Graph Generative Networks Mahdi Karami	48892
MCM: Masked Cell Modeling for Anomaly Detection in Tabular Data Jiaxin Yin, Yuanyuan Qiao, Zitang Zhou, Xiangchao Wang, Jie Yang	48920
A Fast and Provable Algorithm for Sparse Phase Retrieval Jian-Feng Cai, Yu Long, Ruixue Wen, Jiaxi Ying	48943
Towards Transparent Time Series Forecasting Krzysztof Kacprzyk, Tennison Liu, Mihaela Van Der Schaar	48969

Out-Of-Domain Unlabeled Data Improves Generalization Seyed Amir Hossein Saberi, Amir Najafi, Alireza Heidari, Mohammad Hosein Movasaghinia, Abolfazl Motahari, Babak Khalaj	49005
Effective Data Augmentation with Diffusion Models Brandon Trabucco, Kyle Doherty, Max A Gurinas, Ruslan Salakhutdinov	49034
Beyond Imitation: Leveraging Fine-Grained Quality Signals for Alignment Geyang Guo, Ranchi Zhao, Tianyi Tang, Xin Zhao, Ji-Rong Wen	49057
Contrastive Difference Predictive Coding Chongyi Zheng, Ruslan Salakhutdinov, Benjamin Eysenbach	49073
COLEP: Certifiably Robust Learning-Reasoning Conformal Prediction Via Probabilistic Circuits Mintong Kang, Nezihe Merve G��rel, Linyi Li, Bo Li	49098
Supervised Knowledge Makes Large Language Models Better In-Context Learners Linyi Yang, Shuibai Zhang, Zhuohao Yu, Guangsheng Bao, Yidong Wang, Jindong Wang, Ruochen Xu, Wei Ye, Xing Xie, Weizhu Chen, Yue Zhang	49139
USB-NeRF: Unrolling Shutter Bundle Adjusted Neural Radiance Fields Moyang Li, Peng Wang, Lingzhe Zhao, Bangyan Liao, Peidong Liu	49156
Hiding in Plain Sight: Disguising Data Stealing Attacks in Federated Learning Kostadin Garov, Dimitar Iliev Dimitrov, Nikola Jovanovi��, Martin Vechev	49177
Score Models for Offline Goal-Conditioned Reinforcement Learning Harshit Sikchi, Rohan Chitnis, Ahmed Touati, Alborz Geramifard, Amy Zhang, Scott Niekum	49204
Bounding the Expected Robustness of Graph Neural Networks Subject to Node Feature Attacks Yassine Abbahaddou, Sofiane Ennadir, Johannes F. Lutzeyer, Michalis Vazirgiannis, Henrik Bostr��m	49230
Diffusion in Diffusion: Cyclic One-Way Diffusion for Text-Vision-Conditioned Generation Ruoyu Wang, Yongqi Yang, Zhihao Qian, Ye Zhu, Yu Wu	49262
Hierarchical Context Merging: Better Long Context Understanding for Pre-Trained LLMs Woomin Song, Seunghyuk Oh, Sangwoo Mo, Jaehyung Kim, Sukmin Yun, Jung-Woo Ha, Jinwoo Shin	49283
CompA: Addressing the Gap in Compositional Reasoning in Audio-Language Models Sreyan Ghosh, Ashish Seth, Sonal Kumar, Utkarsh Tyagi, Chandra Kiran Reddy Evuru, Ramaneswaran S, S Sakshi, Oriol Nieto, Ramani Duraiswami, Dinesh Manocha	49302
Dual RL: Unification and New Methods for Reinforcement and Imitation Learning Harshit Sikchi, Qinqing Zheng, Amy Zhang, Scott Niekum	49328
A Precise Characterization of SGD Stability Using Loss Surface Geometry Gregory Dexter, Borja Ocejo, Sathiya Keerthi, Aman Gupta, Ayan Acharya, Rajiv Khanna	49376

<i>Towards Optimal Regret in Adversarial Linear MDPs with Bandit Feedback</i> <i>Haolin Liu, Chen-Yu Wei, Julian Zimmert</i>	49394
<i>ViLMA: A Zero-Shot Benchmark for Linguistic and Temporal Grounding in Video-Language Models</i> <i>Ilker Kesen, Andrea Pedrotti, Mustafa Dogan, Michele Cafagna, Emre Can Acikgoz, Letitia Parcalabescu, Iacer Calixto, Anette Frank, Albert Gatt, Aykut Erdem, Erkut Erdem</i>	49435
<i>Lemur: Integrating Large Language Models in Automated Program Verification</i> <i>Haoze Wu, Clark Barrett, Nina Narodytska</i>	49486
<i>Deep SE(3)-Equivariant Geometric Reasoning for Precise Placement Tasks</i> <i>Ben Eisner, Yi Yang, Todor Davchev, Mel Vecerik, Jonathan Scholz, David Held</i>	49497
<i>Error Feedback Reloaded: From Quadratic to Arithmetic Mean of Smoothness Constants</i> <i>Peter Richtárik, Elnur Gasanov, Konstantin Pavlovich Burlachenko</i>	49518
<i>Backdoor Secrets Unveiled: Identifying Backdoor Data with Optimized Scaled Prediction Consistency</i> <i>Soumyadeep Pal, Yuguang Yao, Ren Wang, Bingquan Shen, Sijia Liu</i>	49585
<i>Dissecting Sample Hardness: A Fine-Grained Analysis of Hardness Characterization Methods for Data-Centric AI</i> <i>Nabeel Seedat, Fergus Imrie, Mihaela Van Der Schaar</i>	49608
<i>Consistency Trajectory Models: Learning Probability Flow ODE Trajectory of Diffusion</i> <i>Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka, Yutong He, Yuki Mitsufuji, Stefano Ermon</i>	49678
<i>Reward-Consistent Dynamics Models Are Strongly Generalizable for Offline Reinforcement Learning</i> <i>Fan-Ming Luo, Tian Xu, Xingchen Cao, Yang Yu</i>	49711
<i>Dynamic Layer Tying for Parameter-Efficient Transformers</i> <i>Tamir David Hay, Lior Wolf</i>	49744
<i>Consistent Algorithms for Multi-Label Classification with Macro-At-\$k\$ Metrics</i> <i>Erik Schultheis, Wojciech Kotłowski, Marek Wydmuch, Rohit Babbar, Strom Borman, Krzysztof Dembczynski</i>	49759
<i>Minimax Optimality of Convolutional Neural Networks for Infinite Dimensional Input-Output Problems and Separation from Kernel Methods</i> <i>Yuto Nishimura, Taiji Suzuki</i>	49789
<i>Tool-Augmented Reward Modeling</i> <i>Lei Li, Yekun Chai, Shuohuan Wang, Yu Sun, Hao Tian, Ningyu Zhang, Hua Wu</i>	49820

A Differentially Private Clustering Algorithm for Well-Clustered Graphs Weiqiang He, Hendrik Fichtenberger, Pan Peng	49848
Navigating Dataset Documentations in AI: A Large-Scale Analysis of Dataset Cards on HuggingFace Xinyu Yang, Weixin Liang, James Zou	49863
PROGRAM: PROtotype GRaph Model Based Pseudo-Label Learning for Test-Time Adaptation Haopeng Sun, Lumin Xu, Sheng Jin, Ping Luo, Chen Qian, Wentao Liu	49881
Accurate Retraining-Free Pruning for Pretrained Encoder-Based Language Models Seungcheol Park, Hojun Choi, U Kang	49912
Scalable Monotonic Neural Networks Hyunho Kim, Jong-Seok Lee	49932
A Stable, Fast, and Fully Automatic Learning Algorithm for Predictive Coding Networks Tommaso Salvatori, Yuhang Song, Yordan Yordanov, Beren Millidge, Lei Sha, Cornelius Emde, Zhenghua Xu, Rafal Bogacz, Thomas Lukasiewicz	49949
How Many Pretraining Tasks Are Needed for In-Context Learning of Linear Regression? Jingfeng Wu, Difan Zou, Zixiang Chen, Vladimir Braverman, Quanquan Gu, Peter Bartlett	49974
MixSATGEN: Learning Graph Mixing for SAT Instance Generation Xinyan Chen, Yang Li, Runzhong Wang, Junchi Yan	50029
Decision ConvFormer: Local Filtering in MetaFormer is Sufficient for Decision Making Jeonghye Kim, Suyoung Lee, Woojun Kim, Youngchul Sung	50046
AutoDAN: Generating Stealthy Jailbreak Prompts on Aligned Large Language Models Xiaogeng Liu, Nan Xu, Muhao Chen, Chaowei Xiao	50066
The Reversal Curse: LLMs Trained on “A is B” Fail to Learn “B is A” Lukas Berglund, Meg Tong, Maximilian Kaufmann, Mikita Balesni, Asa Cooper Stickland, Tomasz Korbak, Owain Evans	50087
Learning Delays in Spiking Neural Networks Using Dilated Convolutions with Learnable Spacings Ilyass Hammouamri, Ismail Khalfaoui-Hassani, Timoth�e Masquelier	50107
Provable Benefits of Multi-Task RL Under Non-Markovian Decision Making Processes Ruiquan Huang, Yuan Cheng, Jing Yang, Vincent Tan, Yingbin Liang	50121
Improving Offline RL by Blending Heuristics Sinong Geng, Aldo Pacchiano, Andrey Kolobov, Ching-An Cheng	50150
Explaining Kernel Clustering Via Decision Trees Maximilian Fleissner, Leena Chennuru Vankadara, Debarghya Ghoshdastidar	50180

Never Train from Scratch: Fair Comparison of Long-Sequence Models Requires Data-Driven Priors Ido Amos, Jonathan Berant, Ankit Gupta	50203
An Agnostic View on the Cost of Overfitting in (Kernel) Ridge Regression Lijia Zhou, James B Simon, Gal Vardi, Nathan Srebro	50222
Vocos: Closing the Gap Between Time-Domain and Fourier-Based Neural Vocoders for High-Quality Audio Synthesis Hubert Siuzdak	50241
Enhancing Human Experience in Human-Agent Collaboration: A Human-Centered Modeling Approach Based on Positive Human Gain Yiming Gao, Feiyu Liu, Liang Wang, Dehua Zheng, Zhenjie Lian, Weixuan Wang, Wenjin Yang, Siqin Li, Xianliang Wang, Wenhui Chen, Jing Dai, Qiang Fu, Yang Wei, Lanxiao Huang, Wei Liu	50256
What Makes a Good Prune? Maximal Unstructured Pruning for Maximal Cosine Similarity Gabryel Mason-Williams, Fredrik Dahlqvist	50285
Enhancing Small Medical Learners with Privacy-Preserving Contextual Prompting Xinlu Zhang, Shiyang Li, Xianjun Yang, Chenxin Tian, Yao Qin, Linda Ruth Petzold	50300
Retrieval-Guided Reinforcement Learning for Boolean Circuit Minimization Animesh Basak Chowdhury, Marco Romanelli, Benjamin Tan, Ramesh Karri, Siddharth Garg	50329
Neural Language of Thought Models Yi-Fu Wu, Minseung Lee, Sungjin Ahn	50351
Modeling State-Dependent Communication Between Brain Regions with Switching Nonlinear Dynamical Systems Orren Karniol-Tambour, David M. Zoltowski, E. Mika Diamanti, Lucas Pinto, Carlos D Brody, David W. Tank, Jonathan W. Pillow	50371
Learning from Label Proportions: Bootstrapping Supervised Learners Via Belief Propagation Shreyas Havaladar, Navodita Sharma, Shubhi Sareen, Karthikeyan Shanmugam, Aravindan Raghuv eer	50392
Fiber Monte Carlo Nick Richardson, Deniz Oktay, Yaniv Ovadia, James C Bowden, Ryan P Adams	50416
Implicit Neural Representations and the Algebra of Complex Wavelets T Mitchell Roddenberry, Vishwanath Saragadam, Maarten V. De Hoop, Richard Baraniuk	50432
Unbalancedness in Neural Monge Maps Improves Unpaired Domain Translation Luca Eyring, Dominik Klein, Th��o Uscidda, Giovanni Palla, Niki Kilbertus, Zeynep Akata, Fabian J Theis	50447
Novel Quadratic Constraints for Extending LipSDP Beyond Slope-Restricted Activations Patricia Pauli, Aaron J Havens, Alexandre Araujo, Siddharth Garg, Farshad Khorrami, Frank Allg��wer, Bin Hu	50475

A Benchmark for Learning to Translate a New Language from One Grammar Book Garrett Tanzer, Mirac Suzgun, Eline Visser, Dan Jurafsky, Luke Melas-Kyriazi	50501
How Do Language Models Bind Entities in Context? Jiahai Feng, Jacob Steinhardt	50532
Active Retrosynthetic Planning Aware of Route Quality Luotian Yuan, Yemin Yu, Ying Wei, Yongwei Wang, Zhihua Wang, Fei Wu	50555
Parsing Neural Dynamics with Infinite Recurrent Switching Linear Dynamical Systems Victor Geadah, International Brain Laboratory, Jonathan W. Pillow	50575
Efficient Subgraph GNNs by Learning Effective Selection Policies Beatrice Bevilacqua, Moshe Eliasof, Eli Meirom, Bruno Ribeiro, Haggai Maron	50594
Distributionally Robust Optimization with Bias and Variance Reduction Ronak Mehta, Vincent Roulet, Krishna Pillutla, Zaid Harchaoui	50621
Lion Secretly Solves a Constrained Optimization: As Lyapunov Predicts Lizhang Chen, Bo Liu, Kaizhao Liang, Qiang Liu	50687
Understanding Catastrophic Forgetting in Language Models Via Implicit Inference Suhas Kotha, Jacob Mitchell Springer, Aditi Raghunathan	50723
3D Feature Prediction for Masked-AutoEncoder-Based Point Cloud Pretraining Siming Yan, Yuqi Yang, Yu-Xiao Guo, Hao Pan, Peng-Shuai Wang, Xin Tong, Yang Liu, Qixing Huang	50753
Mediator Interpretation and Faster Learning Algorithms for Linear Correlated Equilibria in General Sequential Games Brian Hu Zhang, Gabriele Farina, Tuomas Sandholm	50765
Closing the Curious Case of Neural Text Degeneration Matthew Finlayson, John Hewitt, Alexander Koller, Swabha Swayamdipta, Ashish Sabharwal	50784
Source-Free and Image-Only Unsupervised Domain Adaptation for Category Level Object Pose Estimation Prakhar Kaushik, Aayush Mishra, Adam Kortylewski, Alan Yuille	50803
F-FERM: A Scalable Framework for Robust Fair Empirical Risk Minimization Sina Baharlouei, Shivam Patel, Meisam Razaviyayn	50839
Eureka: Human-Level Reward Design Via Coding Large Language Models Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, Anima Anandkumar	50853
Learning Semilinear Neural Operators: A Unified Recursive Framework for Prediction and Data Assimilation. Ashutosh Singh, Ricardo Augusto Borsoi, Deniz Erdogmus, Tales Imbiriba	50898

Estimating Shape Distances on Neural Representations with Limited Samples Dean A Pospisil, Brett W. Larsen, Sarah E Harvey, Alex H Williams	50916
Distributional Preference Learning: Understanding and Accounting for Hidden Context in RLHF Anand Siththaranjan, Cassidy Laidlaw, Dylan Hadfield-Menell	50940
Role of Locality and Weight Sharing in Image-Based Tasks: A Sample Complexity Separation Between CNNs, LCNs, and FCNs Aakash Lahoti, Stefani Karp, Ezra Winston, Aarti Singh, Yuanzhi Li	50964
Improved Algorithm and Bounds for Successive Projection Jiashun Jin, Tracy Ke, Gabriel Moryoussef, Jiajun Tang, Jingming Wang	51004
Disentangling Time Series Representations Via Contrastive Independence-Of-Support on I-Variational Inference Khalid Oublal, Said Ladjal, David Benhaïem, Emmanuel Le Borgne, François Roueff	51015
One-Shot Active Learning Based on Lewis Weight Sampling for Multiple Deep Models Sheng-Jun Huang, Yi Li, Yiming Sun, Ying-Peng Tang	51029
Fine-Tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To! Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, Peter Henderson	51053
A Simple Interpretable Transformer for Fine-Grained Image Classification and Analysis Dipanjoy Paul, Arpita Chowdhury, Xinqi Xiong, Feng-Ju Chang, David Edward Carlyn, Samuel Stevens, Kaiya L Provost, Anuj Karpatne, Bryan Carstens, Daniel Rubenstein, Charles Stewart, Tanya Berger-Wolf, Yu Su, Wei-Lun Chao	51109
Can LLM-Generated Misinformation Be Detected? Canyu Chen, Kai Shu	51133
Sharpness-Aware Minimization Enhances Feature Quality Via Balanced Learning Jacob Mitchell Springer, Vaishnavh Nagarajan, Aditi Raghunathan	51173
AutomaTikZ: Text-Guided Synthesis of Scientific Vector Graphics with TikZ Jonas Belouadi, Anne Lauscher, Steffen Eger	51198
Social Reward: Evaluating and Enhancing Generative AI Through Million-User Feedback from an Online Creative Community Arman Isajanyan, Artur Shatveryan, David Kocharian, Zhangyang Wang, Humphrey Shi	51225
PubDef: Defending Against Transfer Attacks from Public Models Chawin Sitawarin, Jaewon Chang, David Huang, Wesson Altoyian, David Wagner	51241
On Diffusion Modeling for Anomaly Detection Victor Liversnoche, Vineet Jain, Yashar Hezaveh, Siamak Ravanbakhsh	51267

Don't Trust: Verify -- Grounding LLM Quantitative Reasoning with Autoformalization Jin Peng Zhou, Charles E Staats, Wenda Li, Christian Szegedy, Kilian Q Weinberger, Yuhuai Wu	51298
Detecting Pretraining Data from Large Language Models Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, Luke Zettlemoyer	51318
Sparse MoE with Language Guided Routing for Multilingual Machine Translation Xinyu Zhao, Xuxi Chen, Yu Cheng, Tianlong Chen	51336
Let's Do the Time-Warp-Attend: Learning Topological Invariants of Dynamical Systems Noa Moriel, Matt Ricci, Mor Nitzan	51355
REFACTOR: Learning to Extract Theorems from Proofs Jin Peng Zhou, Yuhuai Wu, Qiyang Li, Roger Baker Grosse	51382
Learning Adaptive Multiresolution Transforms Via Meta-Framelet-Based Graph Convolutional Network Tianze Luo, Zhanfeng Mo, Sinno Jialin Pan	51403
What's in My Big Data? Yanai Elazar, Akshita Bhagia, Ian Helgi Magnusson, Abhilasha Ravichander, Dustin Schwenk, Alane Suhr, Evan Pete Walsh, Dirk Groeneveld, Luca Soldaini, Sameer Singh, Hannaneh Hajishirzi, Noah A. Smith, Jesse Dodge	51429
In-Context Pretraining: Language Modeling Beyond Document Boundaries Weijia Shi, Sewon Min, Maria Lomeli, Chunting Zhou, Margaret Li, Xi Victoria Lin, Noah A. Smith, Luke Zettlemoyer, Wen-Tau Yih, Mike Lewis	51485
Expressive Losses for Verified Robustness Via Convex Combinations Alessandro De Palma, Rudy R Bunel, Krishnamurthy Dj Dvijotham, M. Pawan Kumar, Robert Stanforth, Alessio Lomuscio	51499
Two-Stage LLM Fine-Tuning with Less Specialization and More Generalization Yihan Wang, Si Si, Daliang Li, Michal Lukasik, Felix Yu, Cho-Jui Hsieh, Inderjit S Dhillon, Sanjiv Kumar	51527
An Operator Preconditioning Perspective on Training in Physics-Informed Machine Learning Tim De Ryck, Florent Bonnet, Siddhartha Mishra, Emmanuel De Bezenac	51546
NEFTune: Noisy Embeddings Improve Instruction Finetuning Neel Jain, Ping-Yeh Chiang, Yuxin Wen, John Kirchenbauer, Hong-Min Chu, Gowthami Somepalli, Brian R. Bartoldson, Bhavya Kailkhura, Avi Schwarzschild, Aniruddha Saha, Micah Goldblum, Jonas Geiping, Tom Goldstein	51575
Graph-Based Virtual Sensing from Sparse and Partial Multivariate Observations Giovanni De Felice, Andrea Cini, Daniele Zambon, Vladimir Gusev, Cesare Alippi	51601
Offline RL with Observation Histories: Analyzing and Improving Sample Complexity Joey Hong, Anca Dragan, Sergey Levine	51623

Understanding Certified Training with Interval Bound Propagation Yuhao Mao, Mark Niklas Mueller, Marc Fischer, Martin Vechev	51643
Layer-Wise Linear Mode Connectivity Linara Adilova, Maksym Andriushchenko, Michael Kamp, Asja Fischer, Martin Jaggi	51666
SequenceMatch: Imitation Learning for Autoregressive Sequence Modelling with Backtracking Chris Cundy, Stefano Ermon	51694
Understanding the Robustness of Multi-Modal Contrastive Learning to Distribution Shift Yihao Xue, Siddharth Joshi, Dang Nguyen, Baharan Mirzasoleiman	51715
Confidential-DPproof: Confidential Proof of Differentially Private Training Ali Shahin Shamsabadi, Gefei Tan, Tudor Ioan Cebere, Aurélien Bellet, Hamed Haddadi, Nicolas Papernot, Xiao Wang, Adrian Weller	51748
Jointly-Learned Exit and Inference for a Dynamic Neural Network Florence Regol, Joud Chataoui, Mark Coates	51764
Improved Statistical and Computational Complexity of the Mean-Field Langevin Dynamics Under Structured Data Atsushi Nitanda, Kazusato Oko, Taiji Suzuki, Denny Wu	51793
Can Large Language Models Infer Causation from Correlation? Zhijing Jin, Jiarui Liu, Zhiheng Lyu, Spencer Poff, Mrinmaya Sachan, Rada Mihalcea, Mona T. Diab, Bernhard Schölkopf	51819
A Variational Perspective on Solving Inverse Problems with Diffusion Models Morteza Mardani, Jiaming Song, Jan Kautz, Arash Vahdat	51836
Investigating the Benefits of Projection Head for Representation Learning Yihao Xue, Eric Gan, Jiayi Ni, Siddharth Joshi, Baharan Mirzasoleiman	51863
Sparse Spiking Neural Network: Exploiting Heterogeneity in Timescales for Pruning Recurrent SNN Biswadeep Chakraborty, Beomseok Kang, Harshit Kumar, Saibal Mukhopadhyay	51886
GNNBoundary: Towards Explaining Graph Neural Networks Through the Lens of Decision Boundaries Xiaoqi Wang, Han Wei Shen	51918
SineNet: Learning Temporal Dynamics in Time-Dependent Partial Differential Equations Xuan Zhang, Jacob Helwig, Yuchao Lin, Yaochen Xie, Cong Fu, Stephan Wojtowytsch, Shuiwang Ji	51933
Empirical Analysis of Model Selection for Heterogeneous Causal Effect Estimation Divyat Mahajan, Ioannis Mitliagkas, Brady Neal, Vasilis Syrgkanis	51975

<i>Towards Identifiable Unsupervised Domain Translation: A Diversified Distribution Matching Approach</i> Sagar Shrestha, Xiao Fu	52005
<i>Likelihood Training of Cascaded Diffusion Models Via Hierarchical Volume-Preserving Maps</i> Henry Li, Ronen Basri, Yuval Kluger	52039
<i>Large Language Models to Enhance Bayesian Optimization</i> Tennison Liu, Nicol�js Astorga, Nabeel Seedat, Mihaela Van Der Schaar	52060
<i>Circuit Component Reuse Across Tasks in Transformer Language Models</i> Jack Merullo, Carsten Eickhoff, Ellie Pavlick	52093
<i>On the Variance of Neural Network Training with Respect to Test Sets and Distributions</i> Keller Jordan	52122
<i>DRSM: De-Randomized Smoothing on Malware Classifier Providing Certified Robustness</i> Shoumik Saha, Wenxiao Wang, Yigitcan Kaya, Soheil Feizi, Tudor Dumitras	52145
<i>Optimistic Bayesian Optimization with Unknown Constraints</i> Quoc Phong Nguyen, Wan Theng Ruth Chew, Le Song, Bryan Kian Hsiang Low, Patrick Jaillet	52166
<i>VCR-Graphormer: A Mini-Batch Graph Transformer Via Virtual Connections</i> Dongqi Fu, Zhigang Hua, Yan Xie, Jin Fang, Si Zhang, Kaan Sancak, Hao Wu, Andrey Malevich, Jingrui He, Bo Long	52178
<i>To Grok Or Not to Grok: Disentangling Generalization and Memorization on Corrupted Algorithmic Datasets</i> Darshil Doshi, Aritra Das, Tianyu He, Andrey Gromov	52198
<i>Efficient Heterogeneous Meta-Learning Via Channel Shuffling Modulation</i> Minh Hoang, Carl Kingsford	52227
<i>Lie Group Decompositions for Equivariant Neural Networks</i> Mircea Mironenco, Patrick Forr��	52238
<i>Diffeomorphic Mesh Deformation Via Efficient Optimal Transport for Cortical Surface Reconstruction</i> Thanh Tung Le, Khai Nguyen, Shanlin Sun, Kun Han, Nhat Ho, Xiaohui Xie	52273
<i>The Effectiveness of Random Forgetting for Robust Generalization</i> Vijaya Raghavan T Ramkumar, Bahram Zonooz, Elahe Arani	52294
<i>Modeling Boundedly Rational Agents with Latent Inference Budgets</i> Athul Paul Jacob, Abhishek Gupta, Jacob Andreas	52311
<i>Memorization Capacity of Multi-Head Attention in Transformers</i> Sadegh Mahdavi, Renjie Liao, Christos Thrampoulidis	52326

Pre-Training Goal-Based Models for Sample-Efficient Reinforcement Learning Haoqi Yuan, Zhancun Mu, Feiyang Xie, Zongqing Lu	52352
Proximal Policy Gradient Arborescence for Quality Diversity Reinforcement Learning Sumeet Batra, Bryon Tjanaka, Matthew Christopher Fontaine, Aleksei Petrenko, Stefanos Nikolaidis, Gaurav S. Sukhatme	52374
From Bricks to Bridges: Product of Invariances to Enhance Latent Space Communication Irene Cannistraci, Luca Moschella, Marco Fumero, Valentino Maiorca, Emanuele Rodolà	52395
Ito Diffusion Approximation of Universal Ito Chains for Sampling, Optimization and Boosting Aleksei Ustimenko, Aleksandr Beznosikov	52436
Beyond Accuracy: Evaluating Self-Consistency of Code Large Language Models with IdentityChain Marcus J. Min, Yangruibo Ding, Luca Buratti, Saurabh Pujar, Gail Kaiser, Suman Jana, Baishakhi Ray	52463
Context-Aware Meta-Learning Christopher Fifty, Dennis Duan, Ronald Guenther Junkins, Ehsan Amid, Jure Leskovec, Christopher Re, Sebastian Thrun	52479
Tree Search-Based Policy Optimization Under Stochastic Execution Delay David Valensi, Esther Derman, Shie Mannor, Gal Dalal	52501
Designing Skill-Compatible AI: Methodologies and Frameworks in Chess Karim Hamade, Reid McIlroy-Young, Siddhartha Sen, Jon Kleinberg, Ashton Anderson	52522
A Flexible Generative Model for Heterogeneous Tabular EHR with Missing Modality Huan He, William Hao, Yuanzhe Xi, Yong Chen, Bradley Malin, Joyce Ho	52540
Differentiable Learning of Generalized Structured Matrices for Efficient Deep Neural Networks Changwoo Lee, Hun-Seok Kim	52557
Maximum Likelihood Estimation is All You Need for Well-Specified Covariate Shift Jiawei Ge, Shange Tang, Jianqing Fan, Cong Ma, Chi Jin	52578
Predictive, Scalable and Interpretable Knowledge Tracing on Structured Domains Hanqi Zhou, Robert Bamler, Charley M Wu, Álvaro Tejero-Cantero	52590
Directly Fine-Tuning Diffusion Models on Differentiable Rewards Kevin Clark, Paul Vicol, Kevin Swersky, David J. Fleet	52604
Llemma: An Open Language Model for Mathematics Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen Marcus McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, Sean Welleck	52634
Task Structure and Nonlinearity Jointly Determine Learned Representational Geometry Matteo Alleman, Jack Lindsey, Stefano Fusi	52662

Risk Bounds of Accelerated SGD for Overparameterized Linear Regression <i>Xuheng Li, Yihe Deng, Jingfeng Wu, Dongruo Zhou, Quanquan Gu</i>	52677
Predictive Auxiliary Objectives in Deep RL Mimic Learning in the Brain <i>Ching Fang, Kim Stachenfeld</i>	52746
A Discretization Framework for Robust Contextual Stochastic Optimization <i>Rares C Cristian, Georgia Perakis</i>	52760
Variance-Aware Regret Bounds for Stochastic Contextual Dueling Bandits <i>Qiwei Di, Tao Jin, Yue Wu, Heyang Zhao, Farzad Farnoud, Quanquan Gu</i>	52781
MERT: Acoustic Music Understanding Model with Large-Scale Self-Supervised Training <i>Yizhi Li, Ruibin Yuan, Ge Zhang, Yinghao Ma, Xingran Chen, Hanzhi Yin, Chenghao Xiao, Chenghua Lin, Anton Ragni, Emmanouil Benetos, Norbert Gyenge, Roger Dannenberg, Ruiho Liu, Wenhui Chen, Gus Xia, Yemin Shi, Wenhao Huang, Zili Wang, Yike Guo, Jie Fu</i>	52805
Harnessing Density Ratios for Online Reinforcement Learning <i>Philip Amortila, Dylan J Foster, Nan Jiang, Ayush Sekhari, Tengyang Xie</i>	52829
MuSR: Testing the Limits of Chain-Of-Thought with Multistep Soft Reasoning <i>Zayne Rea Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, Greg Durrett</i>	52875
Group Preference Optimization: Few-Shot Alignment of Large Language Models <i>Siyan Zhao, John Dang, Aditya Grover</i>	52934
NeuroBack: Improving CDCL SAT Solving Using Graph Neural Networks <i>Wenxi Wang, Yang Hu, Mohit Tiwari, Sarfraz Khurshid, Kenneth McMillan, Risto Miikkulainen</i>	52957
BadChain: Backdoor Chain-Of-Thought Prompting for Large Language Models <i>Zhen Xiang, Fengqing Jiang, Zidi Xiong, Bhaskar Ramasubramanian, Radha Poovendran, Bo Li</i>	52972
L2MAC: Large Language Model Automatic Computer for Extensive Code Generation <i>Samuel Holt, Max Ruiz Luyten, Mihaela Van Der Schaar</i>	53000
Learning to Design Protein-Protein Interactions with Enhanced Generalization <i>Anton Bushuiev, Roman Bushuiev, Petr Kouba, Anatolii Filkin, Marketa Gabrielova, Michal Gabriel, Jiri Sedlar, Tomas Pluskal, Jiri Damborsky, Stanislav Mazurenko, Josef Sivic</i>	53061
When Scaling Meets LLM Finetuning: The Effect of Data, Model and Finetuning Method <i>Biao Zhang, Zhongtao Liu, Colin Cherry, Orhan Firat</i>	53087
Unified Projection-Free Algorithms for Adversarial DR-Submodular Optimization <i>Mohammad Pedramfar, Yididiya Y. Nadew, Christopher John Quinn, Vaneet Aggarwal</i>	53107
Does CLIP’s Generalization Performance Mainly Stem from High Train-Test Similarity? <i>Prasanna Mayilvahanan, Thaddäus Wiedemer, Evgenia Rusak, Matthias Bethge, Wieland Brendel</i>	53138

Curriculum Reinforcement Learning for Quantum Architecture Search Under Hardware Errors Yash J. Patel, Akash Kundu, Mateusz Ostaszewski, Xavier Bonet-Monroig, Vedran Dunjko, Onur Danaci	53173
MT-Ranker: Reference-Free Machine Translation Evaluation by Inter-System Ranking Ibraheem Muhammad Moosa, Rui Zhang, Wenpeng Yin	53205
Forward Learning of Graph Neural Networks Namyong Park, Xing Wang, Antoine Simoulin, Shuai Yang, Grey Yang, Ryan A. Rossi, Puja Trivedi, Nesreen K. Ahmed	53223
Sufficient Conditions for Offline Reactivation in Recurrent Neural Networks Nanda H Krishna, Colin Bredenberg, Daniel Levenstein, Blake Aaron Richards, Guillaume Lajoie	53259
Cycle Consistency Driven Object Discovery Aniket Rajiv Didolkar, Anirudh Goyal, Yoshua Bengio	53279
Beyond Spatio-Temporal Representations: Evolving Fourier Transform for Temporal Graphs Anson Bastos, Kuldeep Singh, Abhishek Nadgeri, Manish Singh, Toyotaro Suzumura	53298
Provable Compositional Generalization for Object-Centric Learning Thaddäus Wiedemer, Jack Brady, Alexander Panfilov, Attila Juhos, Matthias Bethge, Wieland Brendel	53329
Learning to Make Adherence-Aware Advice Guanting Chen, Xiaocheng Li, Chunlin Sun, Hanzhao Wang	53357
Demystifying Poisoning Backdoor Attacks from a Statistical Perspective Ganghua Wang, Xun Xian, Ashish Kundu, Jayanth Srinivasa, Xuan Bi, Mingyi Hong, Jie Ding	53369
Improved Techniques for Training Consistency Models Yang Song, Prafulla Dhariwal	53387
Tensor Programs VI: Feature Learning in Infinite Depth Neural Networks Greg Yang, Dingli Yu, Chen Zhu, Soufiane Hayou	53407
Soft Robust MDPs and Risk-Sensitive MDPs: Equivalence, Policy Gradient, and Sample Complexity Runyu Zhang, Yang Hu, Na Li	53459
Overthinking the Truth: Understanding How Language Models Process False Demonstrations Danny Halawi, Jean-Stanislas Denain, Jacob Steinhardt	53474
Improved Efficiency Based on Learned Saccade and Continuous Scene Reconstruction from Foveated Visual Sampling Jiayang Liu, Yiming Bu, Daniel Tso, Qinru Qiu	53513

LiDAR: Sensing Linear Probing Performance in Joint Embedding SSL Architectures Vimal Thilak, Chen Huang, Omid Saremi, Laurent Dinh, Hanlin Goh, Preetum Nakkiran, Joshua M. Susskind, Etai Littwin	53531
Fine-Tuning Language Models for Factuality Katherine Tian, Eric Mitchell, Huaxiu Yao, Christopher D Manning, Chelsea Finn	53571
CNN Kernels Can Be the Best Shapelets Eric Qu, Yansen Wang, Xufang Luo, Wenqiang He, Kan Ren, Dongsheng Li	53587
Revisiting the Last-Iterate Convergence of Stochastic Gradient Methods Zijian Liu, Zhengyuan Zhou	53613
Why is SAM Robust to Label Noise? Christina Baek, J Zico Kolter, Aditi Raghunathan	53648
GAFormer: Enhancing Timeseries Transformers Through Group-Aware Embeddings Jingyun Xiao, Ran Liu, Eva L Dyer	53667
Confronting Reward Model Overoptimization with Constrained RLHF Ted Moskovitz, Aaditya K Singh, Dj Strouse, Tuomas Sandholm, Ruslan Salakhutdinov, Anca Dragan, Stephen Marcus McAleer	53683
Stochastic Gradient Descent for Gaussian Processes Done Right Jihao Andreas Lin, Shreyas Padhy, Javier Antoran, Austin Tripp, Alexander Terenin, Csaba Szepesvari, Jos�� Miguel Hern��ndez-Lobato, David Janz	53711
Simple Hierarchical Planning with Diffusion Chang Chen, Fei Deng, Kenji Kawaguchi, Caglar Gulcehre, Sungjin Ahn	53729
Zero-Shot Robotic Manipulation with Pre-Trained Image-Editing Diffusion Models Kevin Black, Mitsuhiko Nakamoto, Pranav Atreya, Homer Rich Walke, Chelsea Finn, Aviral Kumar, Sergey Levine	53753
Language Model Detectors Are Easily Optimized Against Charlotte Nicks, Eric Mitchell, Rafael Rafailov, Archit Sharma, Christopher D Manning, Chelsea Finn, Stefano Ermon	53775
PILOT: An $\mathcal{O}(1/K)$-Convergent Approach for Policy Evaluation with Nonlinear Function Approximation Zhuqing Liu, Xin Zhang, Jia Liu, Zhengyuan Zhu, Songtao Lu	53795
SpaCE: The Spatial Confounding Environment Mauricio Tec, Ana Trisovic, Michelle Audirac, Sophie Mirabai Woodward, Jie Kate Hu, Naeem Khoshnevis, Francesca Dominici	53808
The Mechanistic Basis of Data Dependence and Abrupt Learning in an In-Context Classification Task Gautam Reddy	53830

Maximum Entropy Model Correction in Reinforcement Learning Amin Rakhsha, Mete Kemertas, Mohammad Ghavamzadeh, Amir-Massoud Farahmand	53844
Fast-ELECTRA for Efficient Pre-Training Chengyu Dong, Liyuan Liu, Hao Cheng, Jingbo Shang, Jianfeng Gao, Xiaodong Liu	53881
Causal Modelling Agents: Causal Graph Discovery Through Synergising Metadata- And Data-Driven Reasoning Ahmed Abdulaal, Adamos Hadjivasiliou, Nina Montana-Brown, Tiantian He, Ayodeji Ijishakin, Ivana Drobnjak, Daniel C. Castro, Daniel C. Alexander	53897
Language Models Represent Space and Time Wes Gurnee, Max Tegmark	53949
Toward Student-Oriented Teacher Network Training for Knowledge Distillation Chengyu Dong, Liyuan Liu, Jingbo Shang	53970
An Emulator for Fine-Tuning Large Language Models Using Small Language Models Eric Mitchell, Rafael Rafailov, Archit Sharma, Chelsea Finn, Christopher D Manning	53993
Pre-Training LiDAR-Based 3D Object Detectors Through Colorization Tai-Yu Pan, Chenyang Ma, Tianle Chen, Cheng Perng Phoo, Katie Z Luo, Yurong You, Mark Campbell, Kilian Q Weinberger, Bharath Hariharan, Wei-Lun Chao	54009
Efficient Modulation for Vision Networks Xu Ma, Xiyang Dai, Jianwei Yang, Bin Xiao, Yinpeng Chen, Yun Fu, Lu Yuan	54023
The Trickle-Down Impact of Reward Inconsistency on RLHF Lingfeng Shen, Sihao Chen, Linfeng Song, Lifeng Jin, Baolin Peng, Haitao Mi, Daniel Khashabi, Dong Yu	54042
Federated Q-Learning: Linear Regret Speedup with Low Communication Cost Zhong Zheng, Fengyu Gao, Lingzhou Xue, Jing Yang	54071
Finite-Time Analysis of On-Policy Heterogeneous Federated Reinforcement Learning Chenyu Zhang, Han Wang, Aritra Mitra, James Anderson	54116
Training Diffusion Models with Reinforcement Learning Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, Sergey Levine	54161
Doubly Robust Instance-Reweighted Adversarial Training Daouda Sow, Sen Lin, Zhangyang Wang, Yingbin Liang	54184
Abstractors and Relational Cross-Attention: An Inductive Bias for Explicit Relational Reasoning in Transformers Awni Altabaa, Taylor Whittington Webb, Jonathan D. Cohen, John Lafferty	54211
The Hedgehog & the Porcupine: Expressive Linear Attentions with Softmax Mimicry Michael Zhang, Kush Bhatia, Hermann Kumbong, Christopher Re	54228

Understanding In-Context Learning in Transformers and LLMs by Learning to Learn Discrete Functions Satwik Bhattamishra, Arkil Patel, Phil Blunsom, Varun Kanade	54258
Transformer-VQ: Linear-Time Transformers Via Vector Quantization Lucas Dax Lingle	54291
FlashFFTConv: Efficient Convolutions for Long Sequences with Tensor Cores Daniel Y Fu, Hermann Kumbong, Eric Nguyen, Christopher Re	54321
Independent-Set Design of Experiments for Estimating Treatment and Spillover Effects Under Network Interference Chencheng Cai, Xu Zhang, Edoardo Airoidi	54350
Towards Foundational Models for Molecular Learning on Large-Scale Multi-Task Datasets Dominique Beaini, Shenyang Huang, Joao Alex Cunha, Zhiyi Li, Gabriela Moisesescu-Pareja, Oleksandr Dymov, Samuel Maddrell-Mander, Callum McLean, Frederik Wenkel, Luis Mller, Jama Hussein Mohamud, Ali Parviz, Michael Craig, Micha, Koziarski, Jiarui Lu, Zhaocheng Zhu, Cristian Gabellini, Kerstin Klaser, Josef Dean, Cas Wognum, Maciej Sypetkowski, Guillaume Rabusseau, Reihaneh Rabbany, Jian Tang, Christopher Morris, Mirco Ravanelli, Guy Wolf, Prudencio Tossou, Hadrien Mary, Therence Bois, Andrew William Fitzgibbon, Blazej Banaszewski, Chad Martin, Dominic Masters	54366
AMAGO: Scalable In-Context Reinforcement Learning for Adaptive Agents Jake Grigsby, Linxi Fan, Yuke Zhu	54399
From Molecules to Materials: Pre-Training Large Generalizable Models for Atomic Property Prediction Nima Shoghi, Adeesh Kolluru, John R. Kitchin, Zachary Ward Ulissi, C. Lawrence Zitnick, Brandon M Wood	54433
PlaSma: Procedural Knowledge Models for Language-Based Planning and Re-Planning Faeze Brahman, Chandra Bhagavatula, Valentina Pyatkin, Jena D. Hwang, Xiang Lorraine Li, Hirona Jacqueline Arai, Soumya Sanyal, Keisuke Sakaguchi, Xiang Ren, Yejin Choi	54463
Are Models Biased on Text Without Gender-Related Language? Catarina G Belm, Preethi Seshadri, Yasaman Razeghi, Sameer Singh	54494
Counting Graph Substructures with Graph Neural Networks Charilaos Kanatsoulis, Alejandro Ribeiro	54534
Reverse Diffusion Monte Carlo Xunpeng Huang, Hanze Dong, Yifan Hao, Yian Ma, Tong Zhang	54566
Optimal Sketching for Residual Error Estimation for Matrix and Vector Norms Yi Li, Honghao Lin, David Woodruff	54610
Tractable MCMC for Private Learning with Pure and Gaussian Differential Privacy Yingyu Lin, Yian Ma, Yu-Xiang Wang, Rachel Emily Redberg, Zhiqi Bu	54621

Chain of Thought Empowers Transformers to Solve Inherently Serial Problems Zhiyuan Li, Hong Liu, Denny Zhou, Tengyu Ma	54660
The Consensus Game: Language Model Generation Via Equilibrium Search Athul Paul Jacob, Yikang Shen, Gabriele Farina, Jacob Andreas	54693
Davidsonian Scene Graph: Improving Reliability in Fine-Grained Evaluation for Text-To-Image Generation Jaemin Cho, Yushi Hu, Jason Michael Baldrige, Roopal Garg, Peter Anderson, Ranjay Krishna, Mohit Bansal, Jordi Pont-Tuset, Su Wang	54710
Neural-Symbolic Recursive Machine for Systematic Generalization Qing Li, Yixin Zhu, Yitao Liang, Ying Nian Wu, Song-Chun Zhu, Siyuan Huang	54731
Structural Fairness-Aware Active Learning for Graph Neural Networks Haoyu Han, Xiaorui Liu, Li Ma, Mohamadali Torkamani, Hui Liu, Jiliang Tang, Makoto Yamada	54752
Generalization in Diffusion Models Arises from Geometry-Adaptive Harmonic Representations Zahra Kadkhodaie, Florentin Guth, Eero P Simoncelli, Stéphanie Mallat	54769
Scaling Physics-Informed Hard Constraints with Mixture-Of-Experts Nithin Chalapathi, Yiheng Du, Aditi S. Krishnapriyan	54794
TEMPO: Prompt-Based Generative Pre-Trained Transformer for Time Series Forecasting Defu Cao, Furong Jia, Serkan O Arik, Tomas Pfister, Yixiang Zheng, Wen Ye, Yan Liu	54813
Exploring the Promise and Limits of Real-Time Recurrent Learning Kazuki Irie, Anand Gopalakrishnan, Jürgen Schmidhuber	54846
Are Bert Family Good Instruction Followers? a Study on Their Potential and Limitations Yisheng Xiao, Juntao Li, Zechen Sun, Zechang Li, Qingrong Xia, Xinyu Duan, Zhefeng Wang, Min Zhang	54870
UNR-Explainer: Counterfactual Explanations for Unsupervised Node Representation Learning Models Hyunju Kang, Geonhee Han, Hogun Park	54890
Unveiling the Unseen: Identifiable Clusters in Trained Depthwise Convolutional Kernels Zahra Babaiee, Peyman Kiasari, Daniela Rus, Radu Grosu	54910
CLaM-TTS: Improving Neural Codec Language Model for Zero-Shot Text-To-Speech Jaehyeon Kim, Keon Lee, Seungjun Chung, Jaewoong Cho	54942
Large Language Models as Automated Aligners for Benchmarking Vision-Language Models Yuanfeng Ji, Chongjian Ge, Weikai Kong, Enze Xie, Zhengying Liu, Zhenguo Li, Ping Luo	54966
On the Role of Discrete Tokenization in Visual Representation Learning Tianqi Du, Yifei Wang, Yisen Wang	54993

On the Scalability and Memory Efficiency of Semidefinite Programs for Lipschitz Constant Estimation of Neural Networks <i>Zi Wang, Bin Hu, Aaron J Havens, Alexandre Araujo, Yang Zheng, Yudong Chen, Somesh Jha</i>	55010
Graph Transformers on EHRs: Better Representation Improves Downstream Performance <i>Raphael Poulain, Rahmatollah Beheshti</i>	55037
Efficient Inverse Multiagent Learning <i>Denizalp Goktas, Amy Greenwald, Sadie Zhao, Alec Koppel, Sumitra Ganesh</i>	55054
Generative Adversarial Equilibrium Solvers <i>Denizalp Goktas, David C. Parkes, Ian Gemp, Luke Marris, Georgios Piliouras, Romuald Elie, Guy Lever, Andrea Tacchetti</i>	55068
Uncertainty-Aware Graph-Based Hyperspectral Image Classification <i>Linlin Yu, Yifei Lou, Feng Chen</i>	55084
Realistic Evaluation of Semi-Supervised Learning Algorithms in Open Environments <i>Lin-Han Jia, Lan-Zhe Guo, Zhi Zhou, Yu-Feng Li</i>	55116
SRGB Real Noise Modeling Via Noise-Aware Sampling with Normalizing Flows <i>Dongjin Kim, Donggoo Jung, Sungyong Baik, Tae Hyun Kim</i>	55146
Pushing Boundaries: Mixup's Influence on Neural Collapse <i>Quinn Leblanc Fisher, Haoming Meng, Vardan Papayan</i>	55161
Rotation Has Two Sides: Evaluating Data Augmentation for Deep One-Class Classification <i>Guodong Wang, Yunhong Wang, Xiuguo Bao, Di Huang</i>	55181
Enhancing Instance-Level Image Classification with Set-Level Labels <i>Renyu Zhang, Aly A Khan, Yuxin Chen, Robert L. Grossman</i>	55195
Bias Runs Deep: Implicit Reasoning Biases in Persona-Assigned LLMs <i>Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, Tushar Khot</i>	55230
Repeated Random Sampling for Minimizing the Time-To-Accuracy of Learning <i>Patrik Okanovic, Roger Waleffe, Vasilis Mageirakos, Konstantinos Nikolakakis, Amin Karbasi, Dionysios Kalogerias, Nezihe Merve G��rel, Theodoros Rekatsinas</i>	55256
PromptAgent: Strategic Planning with Language Models Enables Expert-Level Prompt Optimization <i>Xinyuan Wang, Chenxi Li, Zhen Wang, Fan Bai, Haotian Luo, Jiayou Zhang, Nebojsa Jojic, Eric Xing, Zhiting Hu</i>	55297
Be Aware of the Neighborhood Effect: Modeling Selection Bias Under Interference <i>Haoxuan Li, Chunyuan Zheng, Sihao Ding, Peng Wu, Zhi Geng, Fuli Feng, Xiangnan He</i>	55332

Dichotomy of Early and Late Phase Implicit Biases Can Provably Induce Grokking Kaifeng Lyu, Jikai Jin, Zhiyuan Li, Simon Shaolei Du, Jason D. Lee, Wei Hu	55363
In-Context Learning Dynamics with Random Binary Sequences Eric J Bigelow, Ekdeep Singh Lubana, Robert P. Dick, Hidenori Tanaka, Tomer Ullman	55403
Gen-Z: Generative Zero-Shot Text Classification with Contextualized Label Descriptions Sachin Kumar, Chan Young Park, Yulia Tsvetkov	55447
RECOMP: Improving Retrieval-Augmented LMs with Context Compression and Selective Augmentation Fangyuan Xu, Weijia Shi, Eunsol Choi	55471
Leftover Lunch: Advantage-Based Offline Reinforcement Learning for Language Models Ashutosh Baheti, Ximing Lu, Faeze Brahman, Ronan Le Bras, Maarten Sap, Mark Riedl	55496
ArchLock: Locking DNN Transferability at the Architecture Level with a Zero-Cost Binary Predictor Tong Zhou, Shaolei Ren, Xiaolin Xu	55527
A Quadratic Synchronization Rule for Distributed Deep Learning Xinran Gu, Kaifeng Lyu, Sanjeev Arora, Jingzhao Zhang, Longbo Huang	55543
SKILL-MIX: A Flexible and Expandable Family of Evaluations for AI Models Dingli Yu, Simran Kaur, Arushi Gupta, Jonah Brown-Cohen, Anirudh Goyal, Sanjeev Arora	55578
Pessimistic Nonlinear Least-Squares Value Iteration for Offline Reinforcement Learning Qiwei Di, Heyang Zhao, Jiafan He, Quanquan Gu	55612
Learning to Relax: Setting Solver Parameters Across a Sequence of Linear System Instances Mikhail Khodak, Edmond Chow, Maria Florina Balcan, Ameet Talwalkar	55646
Data Debugging with Shapley Importance Over Machine Learning Pipelines Bojan Karlađi, David Dao, Matteo Interlandi, Sebastian Schelter, Wentao Wu, Ce Zhang	55679
Meta-VBO: Utilizing Prior Tasks in Optimizing Risk Measures with Gaussian Processes Quoc Phong Nguyen, Bryan Kian Hsiang Low, Patrick Jaillet	55714
Extending Power of Nature from Binary to Real-Valued Graph Learning in Real World Chunshu Wu, Ruibing Song, Chuan Liu, Yunan Yang, Ang Li, Michael Huang, Tong Geng	55725
BooookScore: A Systematic Exploration of Book-Length Summarization in the Era of LLMs Yapei Chang, Kyle Lo, Tanya Goyal, Mohit Iyyer	55738
FedDA: Faster Adaptive Gradient Methods for Federated Constrained Optimization Junyi Li, Feihu Huang, Heng Huang	55773
PINNACLE: PINN Adaptive CoLocation and Experimental Points Selection Gregory Kang Ruey Lau, Apivich Hemachandra, See-Kiong Ng, Bryan Kian Hsiang Low	55786

What Matters to You? Towards Visual Representation Alignment for Robot Learning Thomas Tian, Chenfeng Xu, Masayoshi Tomizuka, Jitendra Malik, Andrea Bajcsy	55822
Deceptive Fairness Attacks on Graphs Via Meta Learning Jian Kang, Yinglong Xia, Ross Maciejewski, Jiebo Luo, Hanghang Tong	55842
Energy-Based Automated Model Evaluation Ru Peng, Heming Zou, Haobo Wang, Yawen Zeng, Zenan Huang, Junbo Zhao	55866
Policy Rehearsing: Training Generalizable Policies for Reinforcement Learning Chengxing Jia, Chenxiao Gao, Hao Yin, Fuxiang Zhang, Xiong-Hui Chen, Tian Xu, Lei Yuan, Zongzhang Zhang, Zhi-Hua Zhou, Yang Yu	55905
Chain-Of-Knowledge: Grounding Large Language Models Via Dynamic Knowledge Adapting Over Heterogeneous Sources Xingxuan Li, Ruochen Zhao, Yew Ken Chia, Bosheng Ding, Shafiq Joty, Soujanya Poria, Lidong Bing	55930
Improving Generalization of Alignment with Human Preferences Through Group Invariant Learning Rui Zheng, Wei Shen, Yuan Hua, Wenbin Lai, Shihan Dou, Yuhao Zhou, Zhiheng Xi, Xiao Wang, Haoran Huang, Tao Gui, Qi Zhang, Xuanjing Huang	55953
LitCab: Lightweight Language Model Calibration Over Short- And Long-Form Responses Xin Liu, Muhammad Khalifa, Lu Wang	55976
Assessing Uncertainty in Similarity Scoring: Performance & Fairness in Face Recognition Jean-R��my Conti, Stephan Cl��men��son	55991
Selective Visual Representations Improve Convergence and Generalization for Embodied AI Ainaz Eftekhari, Kuo-Hao Zeng, Jiafei Duan, Ali Farhadi, Aniruddha Kembhavi, Ranjay Krishna	56003
Kosmos-G: Generating Images in Context with Multimodal Large Language Models Xichen Pan, Li Dong, Shaohan Huang, Zhiliang Peng, Wenhui Chen, Furu Wei	56025
Neural Spectral Methods: Self-Supervised Learning in the Spectral Domain Yiheng Du, Nithin Chalapathi, Aditi S. Krishnapriyan	56041
Retrieval Meets Long Context Large Language Models Peng Xu, Wei Ping, Xianchao Wu, Lawrence McAfee, Chen Zhu, Zihan Liu, Sandeep Subramanian, Evelina Bakhturina, Mohammad Shoeybi, Bryan Catanzaro	56059
SuRe: Summarizing Retrievals Using Answer Candidates for Open-Domain QA of LLMs Jaehyung Kim, Jaehyun Nam, Sangwoo Mo, Jongjin Park, Sang-Woo Lee, Minjoon Seo, Jung-Woo Ha, Jinwoo Shin	56075
The Effective Horizon Explains Deep RL Performance in Stochastic Environments Cassidy Laidlaw, Banghua Zhu, Stuart Russell, Anca Dragan	56108

Harnessing Explanations: LLM-To-LM Interpreter for Enhanced Text-Attributed Graph Representation Learning Xiaoxin He, Xavier Bresson, Thomas Laurent, Adam Perold, Yann Lecun, Bryan Hooi	56138
Seeking Neural Nuggets: Knowledge Transfer in Large Language Models from a Parametric Perspective Ming Zhong, Chenxin An, Weizhu Chen, Jiawei Han, Pengcheng He	56160
RepoBench: Benchmarking Repository-Level Code Auto-Completion Systems Tianyang Liu, Canwen Xu, Julian McAuley	56181
Mechanistically Analyzing the Effects of Fine-Tuning on Procedurally Defined Tasks Samyak Jain, Robert Kirk, Ekdeep Singh Lubana, Robert P. Dick, Hidenori Tanaka, Tim Rocktäschel, Edward Grefenstette, David Krueger	56200
Diagnosing Transformers: Illuminating Feature Spaces for Clinical Decision-Making Aliyah R. Hsu, Yeshwanth Cherapanamjeri, Briton Park, Tristan Naumann, Anobel Odisho, Bin Yu	56272
A Good Learner Can Teach Better: Teacher-Student Collaborative Knowledge Distillation Ayan Sengupta, Shantanu Dixit, Md Shad Akhtar, Tanmoy Chakraborty	56302
Addressing Loss of Plasticity and Catastrophic Forgetting in Continual Learning Mohamed Elsayed, A. Rupam Mahmood	56321
On Representation Complexity of Model-Based and Model-Free Reinforcement Learning Hanlin Zhu, Baihe Huang, Stuart Russell	56353
Safe Collaborative Filtering Riku Togashi, Tatsushi Oka, Naoto Ohsaka, Tetsuro Morimura	56376
GraphChef: Decision-Tree Recipes to Explain Graph Neural Networks Peter Mäüller, Lukas Faber, Karolis Martinkus, Roger Wattenhofer	56407
Chain of Hindsight Aligns Language Models with Feedback Hao Liu, Carmelo Sferrazza, Pieter Abbeel	56436
RingAttention with Blockwise Transformers for Near-Infinite Context Hao Liu, Matei Zaharia, Pieter Abbeel	56456
In-Context Learning Through the Bayesian Prism Madhur Panwar, Kabir Ahuja, Navin Goyal	56473
Fair and Efficient Contribution Valuation for Vertical Federated Learning Zhenan Fan, Huang Fang, Xinglu Wang, Zirui Zhou, Jian Pei, Michael Friedlander, Yong Zhang	56528
RAPTOR: Recursive Abstractive Processing for Tree-Organized Retrieval Parth Sarthi, Salman Abdullah, Aditi Tuli, Shubh Khanna, Anna Goldie, Christopher D Manning	56548

Rethinking Channel Dependence for Multivariate Time Series Forecasting: Learning from Leading Indicators Lifan Zhao, Yanyan Shen	56570
Proving Test Set Contamination in Black-Box Language Models Yonatan Oren, Nicole Meister, Niladri S. Chatterji, Faisal Ladhak, Tatsunori Hashimoto	56589
Multimarginal Generative Modeling with Stochastic Interpolants Michael Samuel Albergo, Nicholas Matthew Boffi, Michael Lindsey, Eric Vanden-Eijnden	56608
Mitigating Hallucination in Large Multi-Modal Models Via Robust Instruction Tuning Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, Lijuan Wang	56626
In-Context Autoencoder for Context Compression in a Large Language Model Tao Ge, Hu Jing, Lei Wang, Xun Wang, Si-Qing Chen, Furu Wei	56671
Bayesian Coreset Optimization for Personalized Federated Learning Prateek Chanda, Shrey Modi, Ganesh Ramakrishnan	56688
DDMI: Domain-Agnostic Latent Diffusion Models for Synthesizing High-Quality Implicit Neural Representations Dogyun Park, Sihyeon Kim, Sojin Lee, Hyunwoo J. Kim	56716
Manipulating Dropout Reveals an Optimal Balance of Efficiency and Robustness in Biological and Machine Visual Systems Jacob S. Prince, Gabriel Fajardo, George A. Alvarez, Talia Konkle	56742
Symmetric Neural-Collapse Representations with Supervised Contrastive Loss: The Impact of ReLU and Batching Ganesh Ramachandra Kini, Vala Vakilian, Tina Behnia, Jaidev Gill, Christos Thrampoulidis	56758
Debiased Collaborative Filtering with Kernel-Based Causal Balancing Haoxuan Li, Chunyuan Zheng, Yanghao Xiao, Peng Wu, Zhi Geng, Xu Chen, Peng Cui	56791
Self-Contradictory Hallucinations of Large Language Models: Evaluation, Detection and Mitigation Niels Mndler, Jingxuan He, Slobodan Jenko, Martin Vechev	56810
Partitioning Message Passing for Graph Fraud Detection Wei Zhuo, Zemin Liu, Bryan Hooi, Bingsheng He, Guang Tan, Rizal Fathony, Jia Chen	56840
Dynamics-Informed Protein Design with Structure Conditioning Urszula Julia Komorowska, Simon V Mathis, Kieran Didi, Francisco Vargas, Pietro Lio, Mateja Jamnik	56859
Learning to Reject with a Fixed Predictor: Application to Decontextualization Christopher Mohri, Daniel Andor, Eunsol Choi, Michael Collins, Anqi Mao, Yutao Zhong	56883

Generalization Error of Spectral Algorithms Maksim Velikanov, Maxim Panov, Dmitry Yarotsky	56906
Robust Classification Via Regression for Learning with Noisy Labels Erik Englesson, Hossein Azizpour	56956
Social-Transmotion: Promptable Human Trajectory Prediction Saeed Saadatnejad, Yang Gao, Kaouther Messaoud, Alexandre Alahi	56977
LUM-ViT: Learnable Under-Sampling Mask Vision Transformer for Bandwidth Limited Optical Signal Acquisition Lingfeng Liu, Dong Ni, Hangjie Yuan	56995
Learning from Simplicial Data Based on Random Walks and 1D Convolutions Florian Frantzen, Michael T Schaub	57015
TESTAM: A Time-Enhanced Spatio-Temporal Attention Model with Mixture of Experts Hyunwook Lee, Sungahn Ko	57030
Graph Parsing Networks Yunchong Song, Siyuan Huang, Xinbing Wang, Chenghu Zhou, Zhouhan Lin	57049
KoLA: Carefully Benchmarking World Knowledge of Large Language Models Jifan Yu, Xiaozhi Wang, Shangqing Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun Yao, Xiaohan Zhang, Hanming Li, Chunyang Li, Zheyuan Zhang, Yushi Bai, Yantao Liu, Amy Xin, Kaifeng Yun, Linlu Gong, Nianyi Lin, Jianhui Chen, Zhili Wu, Yunjia Qi, Weikai Li, Yong Guan, Kaisheng Zeng, Ji Qi, Hailong Jin, Jinxin Liu, Yu Gu, Yuan Yao, Ning Ding, Lei Hou, Zhiyuan Liu, Xu Bin, Jie Tang, Juanzi Li	57070
An Intuitive Multi-Frequency Feature Representation for $SO(3)$-Equivariant Networks Dongwon Son, Jaehyung Kim, Sanghyeon Son, Beomjoon Kim	57114
CORN: Contact-Based Object Representation for Nonprehensile Manipulation of General Unseen Objects Yoonyoung Cho, Junhyek Han, Yoontae Cho, Beomjoon Kim	57141
Towards Faithful Explanations: Boosting Rationalization with Shortcuts Discovery Linan Yue, Qi Liu, Yichao Du, Li Wang, Weibo Gao, Yanqing An	57158
AnomalyCLIP: Object-Agnostic Prompt Learning for Zero-Shot Anomaly Detection Qihang Zhou, Guansong Pang, Yu Tian, Shibo He, Jiming Chen	57178
Gradual Optimization Learning for Conformational Energy Minimization Artem Tsybin, Leonid Anatolievich Ugadiarov, Kuzma Khrabrov, Alexander Telepov, Egor Rumiantsev, Alexey Skrynnik, Aleksandr Panov, Dmitry P. Vetrov, Elena Tutubalina, Artur Kadurin	57211
Rethinking the Power of Graph Canonization in Graph Representation Learning with Stability Zehao Dong, Muhan Zhang, Philip Payne, Michael A Province, Carlos Cruchaga, Tianyu Zhao, Fuhai Li, Yixin Chen	57229

St³-Variational Autoencoder: Learning Heavy-Tailed Data with Student's T and Power Divergence Juno Kim, Jaehyuk Kwon, Mincheol Cho, Hyunjong Lee, Joong-Ho Won	57250
Generalized Policy Iteration Using Tensor Approximation for Hybrid Control Suhan Shetty, Teng Xue, Sylvain Calinon	57277
Multilingual Jailbreak Challenges in Large Language Models Yue Deng, Wenxuan Zhang, Sinno Jialin Pan, Lidong Bing	57296
Demonstration-Regularized RL Daniil Tiapkin, Denis Belomestny, Daniele Calandriello, Eric Moulines, Alexey Naumov, Pierre Perrault, Michal Valko, Pierre Menard	57314
Symmetric Mean-Field Langevin Dynamics for Distributional Minimax Problems Juno Kim, Kakei Yamamoto, Kazusato Oko, Zhuoran Yang, Taiji Suzuki	57379
Is This the Subspace You Are Looking For? an Interpretability Illusion for Subspace Activation Patching Aleksandar Makelov, Georg Lange, Atticus Geiger, Neel Nanda	57422
Enhancing Tail Performance in Extreme Classifiers by Label Variance Reduction Anirudh Buvanesh, Rahul Chand, Jatin Prakash, Bhawna Paliwal, Mudit Dhawan, Neelabh Madan, Deepesh Hada, Vidit Jain, Sonu Mehta, Yashoteja Prabhu, Manish Gupta, Ramachandran Ramjee, Manik Varma	57452
Bellman Optimal Stepsize Straightening of Flow-Matching Models Bao Nguyen, Binh Nguyen, Viet Anh Nguyen	57485
Consistency Training with Learnable Data Augmentation for Graph Anomaly Detection with Limited Supervision Nan Chen, Zemin Liu, Bryan Hooi, Bingsheng He, Rizal Fathony, Jun Hu, Jia Chen	57506
On the Learnability of Watermarks for Language Models Chenchen Gu, Xiang Lisa Li, Percy Liang, Tatsunori Hashimoto	57534
Mirage: Model-Agnostic Graph Distillation for Graph Classification Mridul Gupta, Sahil Manchanda, Hariprasad Kodamana, Sayan Ranu	57558
Test-Time Adaptation Against Multi-Modal Reliability Bias Mouxing Yang, Yunfan Li, Changqing Zhang, Peng Hu, Xi Peng	57577
Dissecting Learning and Forgetting in Language Model Finetuning Xiao Zhang, Ji Wu	57597
Causality-Inspired Spatial-Temporal Explanations for Dynamic Graph Neural Networks Kesen Zhao, Liang Zhang	57618

What Makes Good Data for Alignment? a Comprehensive Study of Automatic Data Selection in Instruction Tuning Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, Junxian He	57631
Towards Cross Domain Generalization of Hamiltonian Representation Via Meta Learning Yeongwoo Song, Hawoong Jeong	57652
Forward Learning with Top-Down Feedback: Empirical and Analytical Characterization Ravi Francesco Srinivasan, Francesca Mignacco, Martino Sorbaro, Maria Refinetti, Avi Cooper, Gabriel Kreiman, Giorgia Dellafererra	57673
Controlled Text Generation Via Language Model Arithmetic Jasper Dekoninck, Marc Fischer, Luca Beurer-Kellner, Martin Vechev	57697
Learning Model Uncertainty as Variance-Minimizing Instance Weights Nishant Jain, Karthikeyan Shanmugam, Pradeep Shenoy	57725
Rethinking and Extending the Probabilistic Inference Capacity of GNNs Tuo Xu, Lei Zou	57752
Generative Judge for Evaluating Alignment Junlong Li, Shichao Sun, Weizhe Yuan, Run-Ze Fan, Hai Zhao, Pengfei Liu	57796
Bridging Vision and Language Spaces with Assignment Prediction Jungin Park, Jiyoung Lee, Kwanghoon Sohn	57824
Enhancing Neural Subset Selection: Integrating Background Information into Set Representations Binghui Xie, Yatao Bian, Kaiwen Zhou, Yongqiang Chen, Peilin Zhao, Bo Han, Wei Meng, James Cheng	57843
Locality Sensitive Sparse Encoding for Learning World Models Online Zichen Liu, Chao Du, Wee Sun Lee, Min Lin	57865
Beyond Memorization: Violating Privacy Via Inference with Large Language Models Robin Staab, Mark Vero, Mislav Balunovic, Martin Vechev	57884
Boosting of Thoughts: Trial-And-Error Problem Solving with Large Language Models Sijia Chen, Baochun Li, Di Niu	57931
SaNN: Simple Yet Powerful Simplicial-Aware Neural Networks Sravanthi Gurugubelli, Sundeep Prabhakar Chepuri	57959
TabR: Tabular Deep Learning Meets Nearest Neighbors Yury Gorishniy, Ivan Rubachev, Nikolay Kartashev, Daniil Shlenskii, Akim Kotelnikov, Artem Babenko	57985