

International Conference on Machine Learning (ICML 2024)

Proceedings of Machine Learning Research Volume 235

Vienna, Austria
21-27 July 2024

Part 1 of 75

ISBN: 979-8-3313-0223-8

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2024) by the Authors
All rights reserved.

Printed with permission by Curran Associates, Inc. (2024)

For permission requests, please contact International Machine Learning Society (IMLS)
at the address below.

International Machine Learning Society (IMLS)
c/o Priscilla Rasmussen
209 North Eighth Street
Stroudsburg PA 18360
USA

Phone: (570) 476-8006

Fax: (570) 476-0860

rasmussen@ptd.net

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

TABLE OF CONTENTS

PART 1

Revisiting Character-Level Adversarial Attacks for Language Models	1
<i>Elias Abad Rocamora, Yongtao Wu, Fanghui Liu, Grigorios G. Chrysos, Volkan Cevher</i>	
Adaptively Perturbed Mirror Descent for Learning in Games	31
<i>Kenshi Abe, Kaito Ariu, Mitsuki Sakamoto, Atsushi Iwasaki</i>	
InferCept: Efficient Intercept Support for Augmented Large Language Model Inference	81
<i>Reyna Abhyankar, Zijian He, Vikranth Srivatsa, Hao Zhang, Yiyang Zhang</i>	
Balancing Feature Similarity and Label Variability for Optimal Size-Aware One-Shot Subset Selection	96
<i>Abhinab Acharya, Dayou Yu, Qi Yu, Xumin Liu</i>	
Bayesian Uncertainty for Gradient Aggregation in Multi-Task Learning	117
<i>Idan Achituve, Idit Diamant, Arnon Netzer, Gal Chechik, Ethan Fetaya</i>	
AttnLRP: Attention-Aware Layer-Wise Relevance Propagation for Transformers	135
<i>Reduan Achitbat, Sayed Mohammad Vakilzadeh Hatefi, Maximilian Dreyer, Aakriti Jain, Thomas Wiegand, Sebastian Lapuschkin, Wojciech Samek</i>	
A Unified Framework for Learning with Nonlinear Model Classes from Arbitrary Linear Samples.....	169
<i>Ben Adcock, Juan M. Cardenas, Nick Dexter</i>	
FrameQuant: Flexible Low-Bit Quantization for Transformers	203
<i>Harshavardhan Adep, Zhanpeng Zeng, Li Zhang, Vikas Singh</i>	
BeigeMaps: Behavioral Eigenmaps for Reinforcement Learning from Images.....	228
<i>Sandesh Adhikary, Anqi Li, Byron Boots</i>	
Discovering Bias in Latent Space: An Unsupervised Debiasing Approach.....	246
<i>Dyah Adila, Shuai Zhang, Boran Han, Yuyang Wang</i>	
Optimal Coresets for Low-Dimensional Geometric Median.....	262
<i>Peyman Afshani, Chris Schwiegelshohn</i>	
REST: Efficient and Accelerated EEG Seizure Analysis Through Residual State Updates	271
<i>Arshia Afzal, Grigorios Chrysos, Volkan Cevher, Mahsa Shoaran</i>	
CHAI: Clustered Head Attention for Efficient LLM Inference	291
<i>Saurabh Agarwal, Bilge Acun, Basil Hosmer, Mostafa Elhoushi, Yejin Lee, Shivaram Venkataraman, Dimitris Papailiopoulos, Carole-Jean Wu</i>	
Learning to Play Atari in a World of Tokens	313
<i>Pranav Agarwal, Sheldon Andrews, Samira Ebrahimi Kahou</i>	
Probabilistic Generating Circuits—Demystified	329
<i>Sanyam Agarwal, Markus Bläser</i>	
Improved Differentially Private and Lazy Online Convex Optimization: Lower Regret Without Smoothness Requirements.....	343
<i>Naman Agarwal, Satyen Kale, Karan Singh, Abhradeep Thakurta</i>	
The Non-Linear F-Design and Applications to Interactive Learning	362
<i>Alekh Agarwal, Jian Qian, Alexander Rakhlin, Tong Zhang</i>	
ACPO: A Policy Optimization Algorithm for Average MDPs with Constraints	397
<i>Akhil Agnihotri, Rahul Jain, Haipeng Luo</i>	

CosPGD: An Efficient White-Box Adversarial Attack for Pixel-Wise Prediction Tasks.....	416
<i>Shashank Agnihotri, Steffen Jung, Margret Keuper</i>	
LeaPformer: Enabling Linear Transformers for Autoregressive and Simultaneous Tasks Via Learned Proportions	452
<i>Victor Agostinelli, Sanghyun Hong, Lizhong Chen</i>	
Policy Evaluation for Variance in Average Reward Reinforcement Learning.....	471
<i>Shubhada Agrawal, Prashanth L. A., Siva Theja Maguluri</i>	
Distinguishing the Knowable from the Unknowable with Language Models.....	503
<i>Gustaf Ahdritz, Tian Qin, Nikhil Vyas, Boaz Barak, Benjamin L. Edelman</i>	
Unmasking Vulnerabilities: Cardinality Sketches Under Adaptive Inputs	550
<i>Sara Ahmadian, Edith Cohen</i>	
OptiMUS: Scalable Optimization Modeling with (MI)LP Solvers and Large Language Models.....	577
<i>Ali Ahmaditeshnizi, Wenzhi Gao, Madeleine Udell</i>	
How to Escape Sharp Minima with Random Perturbations	597
<i>Kwangjun Ahn, Ali Jadbabaie, Suvrit Sra</i>	
Understanding Adam Optimizer Via Online Learning of Updates: Adam is FTRL in Disguise	619
<i>Kwangjun Ahn, Zhiyu Zhang, Yunbum Kook, Yan Dai</i>	

PART 2

Not All Distributional Shifts Are Equal: Fine-Grained Robust Conformal Inference	641
<i>Jiahao Ai, Zhimei Ren</i>	
Triple Changes Estimator for Targeted Policies	666
<i>Sina Akbari, Negar Kiyavash</i>	
Improving Computational Complexity in Statistical Models with Local Curvature Information	696
<i>Pedram Akbarian, Tongzheng Ren, Jiacheng Zhuo, Sujay Sanghavi, Nhat Ho</i>	
Learning Mixtures of Gaussian Processes Through Random Projection.....	720
<i>Emmanuel Akeweje, Mimi Zhang</i>	
Encodings for Prediction-Based Neural Architecture Search.....	740
<i>Yash Akhauri, Mohamed S. Abdelfattah</i>	
Iterated Denoising Energy Matching for Sampling from Boltzmann Densities	760
<i>Tara Akhound-Sadegh, Jarrod Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, Alexander Tong</i>	
In-Context Language Learning: Architectures and Algorithms.....	787
<i>Ekin Akyurek, Bailin Wang, Yoon Kim, Jacob Andreas</i>	
Nonlinear Filtering with Brenier Optimal Transport Maps	813
<i>Mohammad Al-Jarrah, Niyizhen Jin, Bamdad Hosseini, Amirhossein Taghvaei</i>	
Revisiting Inexact Fixed-Point Iterations for Min-Max Problems: Stochasticity and Structured Nonconvexity	840
<i>Ahmet Alacaoglu, Donghwan Kim, Stephen J. Wright</i>	
Gaussian Processes on Cellular Complexes	879
<i>Mathieu Alain, So Takao, Brooks Paige, Marc Peter Deisenroth</i>	
Remembering to Be Fair: Non-Markovian Fairness in Sequential Decision Making	906
<i>Parand A. Alamdari, Toryn Q. Klassen, Elliot Creager, Sheila A. McIlraith</i>	

Stochastic Interpolants with Data-Dependent Couplings	921
<i>Michael S. Albergo, Mark Goldstein, Nicholas M. Boffi, Rajesh Ranganath, Eric Vanden-Eijnden</i>	
Evaluating Model Bias Requires Characterizing Its Mistakes	938
<i>Isabela Albuquerque, Jessica Schrouff, David Warde-Farley, Taylan Cemgil, Sven Gowal, Olivia Wiles</i>	
Energy-Efficient Gaussian Processes Using Low-Precision Arithmetic.....	955
<i>Nicolas Alder, Ralf Herbrich</i>	
Evaluation of Test-Time Adaptation Under Computational Time Constraints	976
<i>Motaseem Alfarra, Hani Itani, Alejandro Pardo, Shyma Alhuwaider, Merey Ramazanova, Juan C. Pérez, Zhipeng Cai, Matthias Müller, Bernard Ghanem</i>	
On the Weight Dynamics of Deep Normalized Networks.....	992
<i>Christian H. X. Ali Mehmeti-Göpel, Michael Wand</i>	
No Dimensional Sampling Coresets for Classification	1008
<i>Meysam Alishahi, Jeff M. Phillips</i>	
Unsupervised Evaluation of Code LLMs with Round-Trip Correctness.....	1050
<i>Miltiadis Allamanis, Sheena Panthaplackel, Pengcheng Yin</i>	
Physics of Language Models: Part 3.1, Knowledge Storage and Extraction.....	1067
<i>Zeyuan Allen-Zhu, Yuanzhi Li</i>	
Byzantine-Robust Federated Learning: Impact of Client Subsampling and Local Updates.....	1078
<i>Youssef Allouah, Sadegh Farhadkhani, Rachid Guerraoui, Nirupam Gupta, Rafael Pinot, Geovani Rizk, Sasha Voitovych</i>	
The Privacy Power of Correlated Noise in Decentralized Learning.....	1115
<i>Youssef Allouah, Anastasia Koloskova, Aymane El Firdoussi, Martin Jaggi, Rachid Guerraoui</i>	
Predicting Dose-Response Curves with Deep Neural Networks.....	1144
<i>Pedro A. Campana, Paul Prasse, Tobias Scheffer</i>	
Robust and Conjugate Gaussian Process Regression	1155
<i>Matias Altamirano, François-Xavier Briol, Jeremias Knoblauch</i>	
Beyond the Norms: Detecting Prediction Errors in Regression Models	1186
<i>Andres Altieri, Marco Romanelli, Georg Pichler, Florence Alberge, Pablo Piantanida</i>	
Position: Stop Making Unscientific AGI Performance Claims	1222
<i>Patrick Altmeyer, Andrew M. Demetriou, Antony Bartlett, Cynthia C. S. Liem</i>	
Hyperbolic Optimizer as a Dynamical System.....	1243
<i>Nico Alvarado, Hans Lobel</i>	
Stationarity Without Mean Reversion in Improper Gaussian Processes.....	1261
<i>Luca Ambrogioni</i>	
Robust Graph Matching When Nodes Are Corrupt.....	1276
<i>Taha Ameen, Bruce Hajek</i>	
Fast Algorithms for Hypergraph PageRank with Applications to Semi-Supervised Learning	1306
<i>Konstantinos Ameranis, Adela Depavia, Lorenzo Orecchia, Erasmo Tani</i>	
Scalable and Flexible Causal Discovery with an Efficient Test for Adjacency	1331
<i>Alan Nawzad Amin, Andrew Gordon Wilson</i>	
Generalization Error of Graph Neural Networks in the Mean-Field Regime.....	1359
<i>Gholamali Aminian, Yixuan He, Gesine Reinert, Lukasz Szpruch, Samuel N. Cohen</i>	
Scalable Online Exploration Via Coverability.....	1392
<i>Philip Amortila, Dylan J. Foster, Akshay Krishnamurthy</i>	

WAVES: Benchmarking the Robustness of Image Watermarks	1456
<i>Bang An, Mucong Ding, Tahseen Rabbani, Aakriti Agrawal, Yuancheng Xu, Chenghao Deng, Sicheng Zhu, Abdirisak Mohamed, Yuxin Wen, Tom Goldstein, Furong Huang</i>	
Training-Free Long-Context Scaling of Large Language Models.....	1493
<i>Chenxin An, Fei Huang, Jun Zhang, Shansan Gong, Xipeng Qiu, Chang Zhou, Lingpeng Kong</i>	

PART 3

Navigating Scaling Laws: Compute Optimality in Adaptive Model Training.....	1511
<i>Sotiris Anagnostidis, Gregor Bachmann, Imanol Schlag, Thomas Hofmann</i>	
Adaptive Hierarchical Certification for Segmentation Using Randomized Smoothing	1531
<i>Alaa Anani, Tobias Lorenz, Bernt Schiele, Mario Fritz</i>	
Adaptive Observation Cost Control for Variational Quantum Eigensolvers	1557
<i>Christopher J. Anders, Kim A. Nicoli, Bingting Wu, Naima Elosegui, Samuele Pedrielli, Lena Funcke, Karl Jansen, Stefan Kühn, Shinichi Nakajima</i>	
Fast, Scalable, Warm-Start Semidefinite Programming with Spectral Bundling and Sketching.....	1579
<i>Rico Angell, Andrew McCallum</i>	
Online Conformal Prediction with Decaying Step Sizes.....	1616
<i>Anastasios N. Angelopoulos, Rina Foygel Barber, Stephen Bates</i>	
A Rate-Distortion View of Uncertainty Quantification	1631
<i>Ifigeneia Apostolopoulou, Benjamin Eysenbach, Frank Nielsen, Artur Dubrawski</i>	
Practical Performance Guarantees for Pipelined DNN Inference.....	1655
<i>Aaron Archer, Matthew Fahrback, Kuikui Liu, Prakash Prabhu</i>	
Unsupervised Concept Discovery Mitigates Spurious Correlations.....	1672
<i>Md Rifat Arefin, Yan Zhang, Aristide Baratin, Francesco Locatello, Irina Rish, Dianbo Liu, Kenji Kawaguchi</i>	
Accelerating Legacy Numerical Solvers by Non-Intrusive Gradient-Based Meta-Solving	1689
<i>Sohei Arisaka, Qianxiao Li</i>	
Causal Action Influence Aware Counterfactual Data Augmentation.....	1709
<i>Núria Armengol Urpí, Marco Bagatella, Marin Vlastelica, Georg Martius</i>	
Online Learning and Information Exponents: The Importance of Batch Size & Time / Complexity Tradeoffs.....	1730
<i>Luca Arnaboldi, Yatin Dandi, Florent Krzakala, Bruno Loureiro, Luca Pesce, Ludovic Stephan</i>	
Simple Linear Attention Language Models Balance the Recall-Throughput Tradeoff.....	1763
<i>Simran Arora, Sabri Eyuboglu, Michael Zhang, Aman Timalsina, Silas Alberti, James Zou, Atri Rudra, Christopher Ré</i>	
Inferring Change Points in High-Dimensional Linear Regression Via Approximate Message Passing.....	1841
<i>Gabriel Arpino, Xiaoqi Liu, Ramji Venkataramanan</i>	
An Amortized Approach to Non-Linear Mixed-Effects Modelingbased on Neural Posterior Estimation.....	1865
<i>Jonas Arruda, Yannik Schälte, Clemens Peiter, Olga Teylytska, Ulrich Jaehde, Jan Hasenauer</i>	
Learning the Target Network in Function Space	1902
<i>Kavosh Asadi, Yao Liu, Shoham Sabach, Ming Yin, Rasool Fakoor</i>	

Translation Equivariant Transformer Neural Processes.....	1924
<i>Matthew Ashman, Cristiana Diaconu, Junhyuck Kim, Lakee Sivaraya, Stratis Markou, James Requeima, Wessel P. Bruinsma, Richard E. Turner</i>	
Private Vector Mean Estimation in the Shuffle Model: Optimal Rates Require Many Messages	1945
<i>Hilal Asi, Vitaly Feldman, Jelani Nelson, Huy L. Nguyen, Kunal Talwar, Samson Zhou</i>	
Bifurcated Attention for Single-Context Large-Batch Sampling.....	1971
<i>Ben Athiwaratkun, Sujan Kumar Gonugondla, Sanjay Krishna Gouda, Haifeng Qian, Hantian Ding, Qing Sun, Jun Wang, Jiacheng Guo, Liangfu Chen, Parminder Bhatia, Ramesh Nallapati, Sudipta Sengupta, Bing Xiang</i>	
Delaunay Graph: Addressing Over-Squashing and Over-Smoothing Using Delaunay Triangulation.....	1992
<i>Hugo Attali, Davide Buscaldi, Nathalie Pernelle</i>	
How Free is Parameter-Free Stochastic Optimization?	2009
<i>Amit Attia, Tomer Koren</i>	
Information Complexity of Stochastic Convex Optimization: Applications to Generalization, Memorization, and Tracing	2035
<i>Idan Attias, Gintare Karolina Dziugaite, Mahdi Haghifam, Roi Livni, Daniel M. Roy</i>	
Agnostic Sample Compression Schemes for Regression	2069
<i>Idan Attias, Steve Hanneke, Aryeh Kontorovich, Menachem Sadigurschi</i>	
Data-Efficient Learning Via Clustering-Based Sensitivity Sampling: Foundation Models and Beyond	2086
<i>Kyriakos Axiotis, Vincent Cohen-Addad, Monika Henzinger, Sammy Jerome, Vahab Mirrokni, David Saulpic, David P. Woodruff, Michael Wunder</i>	
Random Features Models: A Way to Study the Success of Naive Imputation	2108
<i>Alexis Ayme, Claire Boyer, Aymeric Dieuleveut, Erwan Scornet</i>	
Switching the Loss Reduces the Cost in Batch Reinforcement Learning.....	2135
<i>Alex Ayoub, Kaiwen Wang, Vincent Liu, Samuel Robertson, James McInerney, Dawen Liang, Nathan Kallus, Csaba Szepesvári</i>	
Bipartite Matching in Massive Graphs: A Tight Analysis of EDCS	2159
<i>Amir Azarmehr, Soheil Behnezhad, Mohammad Roghani</i>	
What is the Long-Run Distribution of Stochastic Gradient Descent? a Large Deviations Analysis	2168
<i>Waïss Azizian, Franck Iutzeler, Jérôme Malick, Panayotis Mertikopoulos</i>	
HyperFields: Towards Zero-Shot Generation of NeRFs from Text.....	2230
<i>Sudarshan Babu, Richard Liu, Avery Zhou, Michael Maire, Greg Shakhnarovich, Rana Hanocka</i>	
Online Matrix Completion: A Collaborative Approach with Hott Items	2248
<i>Dheeraj Baby, Soumyabrata Pal</i>	
Differentiable Weightless Neural Networks	2277
<i>Alan T. L. Bacellar, Zachary Susskind, Mauricio Breternitz Jr., Eugene John, Lizy K. John, Priscila M. V. Lima, Felipe M. G. França</i>	
The Pitfalls of Next-Token Prediction	2296
<i>Gregor Bachmann, Vaishnavh Nagarajan</i>	
Simulation of Graph Algorithms with Looped Transformers	2319
<i>Artur Back De Luca, Kimon Fountoulakis</i>	

PART 4

QBKM: Quantum-Based Matching Kernels for Un-Attributed Graphs.....	2364
<i>Lu Bai, Lixin Cui, Ming Li, Yuewang, Edwin R. Hancock</i>	

Diffusion Models Demand Contrastive Guidance for Adversarial Purification to Advance	2375
<i>Mingyuan Bai, Wei Huang, Tenghui Li, Andong Wang, Junbin Gao, Cesar F Caiafa, Qibin Zhao</i>	
On the Complexity of Finite-Sum Smooth Optimization Under the Polyak–Łojasiewicz Condition.....	2392
<i>Yunyan Bai, Yuxing Liu, Luo Luo</i>	
Constrained Ensemble Exploration for Unsupervised Skill Discovery	2418
<i>Chenjia Bai, Rushuai Yang, Qiaosheng Zhang, Kang Xu, Yi Chen, Ting Xiao, Xuelong Li</i>	
Image Hijacks: Adversarial Images Can Control Generative Models at Runtime.....	2443
<i>Luke Bailey, Euan Ong, Stuart Russell, Scott Emmons</i>	
An Explicit Frame Construction for Normalizing 3D Point Clouds.....	2456
<i>Justin Baker, Shih-Hsin Wang, Tommaso De Fernex, Bao Wang</i>	
Disentanglement Learning Via Topology	2474
<i>Nikita Balabin, Daria Voronkova, Ilya Trofimov, Evgeny Burnaev, Serguei Barannikov</i>	
Adversarial Attacks on Combinatorial Multi-Armed Bandits	2505
<i>Rishab Balasubramanian, Jiawei Li, Prasad Tadepalli, Huazheng Wang, Qingyun Wu, Haoyu Zhao</i>	
Memory Consolidation Enables Long-Context Video Understanding	2527
<i>Ivana Balažević, Yuge Shi, Pinelopi Papalampidi, Rahma Chaabouni, Skanda Koppula, Olivier J. Hénaff</i>	
Characterizing Large Language Model Geometry Helps Solve Toxicity Detection and Generation	2543
<i>Randall Balestriero, Romain Cosentino, Sarath Shekkizhar</i>	
How Learning by Reconstruction Produces Uninformative Features for Perception.....	2566
<i>Randall Balestriero, Yann Lecun</i>	
Combinatorial Approximations for Cluster Deletion: Simpler, Faster, and Better	2586
<i>Vicente Balmaseda, Ying Xu, Yixin Cao, Nate Veldt</i>	
A Field Guide for Pacing Budget and ROS Constraints	2607
<i>Santiago R. Balseiro, Kshipra Bhawalkar, Zhe Feng, Haihao Lu, Vahab Mirrokni, Balasubramanian Sivan, Di Wang</i>	
On the Identifiability of Switching Dynamical Systems	2639
<i>Carles Balsells-Rodas, Yixin Wang, Yingzhen Li</i>	
Analyzing D^{α} Seeding for K-Means.....	2673
<i>Etienne Bamas, Sai Ganesh Nagarajan, Ola Svensson</i>	
Parsimonious Learning-Augmented Approximations for Dense Instances of NP-Hard Problems.....	2700
<i>Evrpidis Bampis, Bruno Escoffier, Michalis Xeferis</i>	
Fair Resource Allocation in Multi-Task Learning	2715
<i>Hao Ban, Kaiyi Ji</i>	
Linguistic Calibration of Long-Form Generations	2732
<i>Neil Band, Xuechen Li, Tengyu Ma, Tatsunori Hashimoto</i>	
Relational DNN Verification with Cross Executional Bound Refinement	2779
<i>Debangshu Banerjee, Gagandeep Singh</i>	
A Dynamic Algorithm for Weighted Submodular Cover Problem	2808
<i>Kiarash Banihashem, Samira Goudarzi, Mohammadtaghi Hajiaghayi, Peyman Jabbarzade, Morteza Monemizadeh</i>	
Dynamic Metric Embedding into ℓ_p Space	2831
<i>Kiarash Banihashem, Mohammadtaghi Hajiaghayi, Dariusz R. Kowalski, Jan Olkowski, Max Springer</i>	

VNN: Verification-Friendly Neural Networks with Hard Robustness Guarantees.....	2846
<i>Anahita Baninajjar, Ahmed Rezine, Amir Aminifar</i>	
Provable Benefits of Local Steps in Heterogeneous Federated Learning for Neural Networks: A Feature Learning Perspective	2857
<i>Yajie Bao, Michael Crawshaw, Mingrui Liu</i>	
Self-Attention Networks Localize When QK-Eigenspectrum Concentrates	2903
<i>Han Bao, Ryuichiro Hataya, Ryo Karakida</i>	
Graph Out-Of-Distribution Detection Goes Neighborhood Shaping	2923
<i>Tianyi Bao, Qitian Wu, Zetian Jiang, Yiting Chen, Jiawei Sun, Junchi Yan</i>	
Stochastic Positional Embeddings Improve Masked Image Modeling	2944
<i>Amir Bar, Florian Bordes, Assaf Shocher, Mahmoud Assran, Pascal Vincent, Nicolas Ballas, Trevor Darrell, Amir Globerson, Yann Lecun</i>	
Subgraphormer: Unifying Subgraph GNNs and Graph Transformers Via Graph Products	2959
<i>Guy Bar-Shalom, Beatrice Bevilacqua, Haggai Maron</i>	
Scale-Free Image Keypoints Using Differentiable Persistent Homology.....	2990
<i>Giovanni Barbarani, Francesco Vaccarino, Gabriele Trivigno, Marco Guerra, Gabriele Berton, Carlo Masone</i>	
To Each (Textual Sequence) Its Own: Improving Memorized-Data Unlearning in Large Language Models.....	3003
<i>George-Octavian Bărbulescu, Peter Triantafillou</i>	
Sliding Down the Stairs: How Correlated Latent Variables Accelerate Learning with Neural Networks	3024
<i>Lorenzo Bardone, Sebastian Goldt</i>	
Adversarial Robustness Limits Via Scaling-Law and Human-Alignment Studies.....	3046
<i>Brian R. Bartoldson, James Diffenderfer, Konstantinos Parasyris, Bhavya Kailkhura</i>	
Neural Diffusion Models	3073
<i>Grigory Bartosh, Dmitry Vetrov, Christian A. Naeseth</i>	
Generalization in Kernel Regression Under Realistic Assumptions.....	3096
<i>Daniel Barzilai, Ohad Shamir</i>	
Local Vs. Global Interpretability: A Computational Complexity Perspective.....	3133
<i>Shahaf Bassan, Guy Amir, Guy Katz</i>	
Differentially Private Domain Adaptation with Theoretical Guarantees	3168
<i>Raef Bassily, Corinna Cortes, Anqi Mao, Mehryar Mohri</i>	

PART 5

A Statistical Framework for Data-Dependent Retrieval-Augmented Models	3197
<i>Soumya Basu, Ankit Singh Rawat, Manzil Zaheer</i>	
On Mechanistic Knowledge Localization in Text-To-Image Generative Models	3224
<i>Samyadeep Basu, Keivan Rezaei, Priyatham Kattakinda, Vlad I Morariu, Nanxuan Zhao, Ryan A Rossi, Varun Manjunatha, Soheil Feizi</i>	
Monotone Individual Fairness	3266
<i>Yahav Bechavod</i>	
Graph Neural Networks Use Graphs When They Shouldn't.....	3284
<i>Maya Bechler-Speicher, Ido Amos, Ran Gilad-Bachrach, Amir Globerson</i>	

Diffusion Tempering Improves Parameter Estimation with Probabilistic Integrators for Ordinary Differential Equations.....	3305
<i>Jonas Beck, Nathanael Bosch, Michael Deistler, Kyra L. Kadhim, Jakob H. Macke, Philipp Hennig, Philipp Berens</i>	
Standardized Interpretable Fairness Measures for Continuous Risk Scores.....	3327
<i>Ann-Kristin Becker, Oana Dumitrasc, Klaus Broelemann</i>	
Unsupervised Representation Learning of Brain Activity Via Bridging Voxel Activity and Functional Connectivity	3347
<i>Ali Behrouz, Parsa Delavari, Farnoosh Hashemi</i>	
Neural Networks Learn Statistics of Increasing Complexity	3382
<i>Nora Belrose, Quintin Pope, Lucia Quirke, Alex Mallen, Xiaoli Fern</i>	
Accelerating Federated Learning with Quick Distributed Mean Estimation.....	3410
<i>Ran Ben-Basat, Shay Vargaftik, Amit Portnoy, Gil Einziger, Yaniv Ben-Itzhak, Michael Mitzenmacher</i>	
The Role of Learning Algorithms in Collective Action.....	3443
<i>Omri Ben-Dov, Jake Fawkes, Samira Samadi, Amartya Sanyal</i>	
D-Flow: Differentiating Through Flows for Controlled Generation	3462
<i>Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, Yaron Lipman</i>	
Transitional Uncertainty with Layered Intermediate Predictions	3484
<i>Ryan Benkert, Mohit Prabhushankar, Ghassan Alregib</i>	
Non-Clairvoyant Scheduling with Partial Predictions.....	3506
<i>Ziyad Benomar, Vianney Perchet</i>	
Sequential Disentanglement by Extracting Static Information from a Single Sequence Element.....	3539
<i>Nimrod Berman, Ilan Naiman, Idan Arbiv, Gal Fadlon, Omri Azencot</i>	
CoLoRA: Continuous Low-Rank Adaptation for Reduced Implicit Neural Modeling of Parameterized Partial Differential Equations	3565
<i>Jules Berman, Benjamin Peherstorfer</i>	
By Tying Embeddings You Are Assuming the Distributional Hypothesis.....	3584
<i>Francesco Bertolotti, Walter Cazzola</i>	
Controlling Behavioral Diversity in Multi-Agent Reinforcement Learning.....	3611
<i>Matteo Bettini, Ryan Kortvelesy, Amanda Prorok</i>	
Refining Minimax Regret for Unsupervised Environment Design	3637
<i>Michael Beukman, Samuel Coward, Michael Matthews, Mattie Fellows, Minqi Jiang, Michael Dennis, Jakob Foerster</i>	
Guiding LLMs the Right Way: Fast, Non-Invasive Constrained Generation.....	3658
<i>Luca Beurer-Kellner, Marc Fischer, Martin Vechev</i>	
Prompt Sketching for Large Language Models.....	3674
<i>Luca Beurer-Kellner, Mark Niklas Mueller, Marc Fischer, Martin Vechev</i>	
Counterfactual Metarules for Local and Global Recourse	3707
<i>Tom Bewley, Salim I. Amoukou, Saumitra Mishra, Daniele Magazzeni, Manuela Veloso</i>	
Sarah Frank-Wolfe: Methods for Constrained Optimization with Best Rates and Practical Features.....	3725
<i>Aleksandr Beznosikov, David Dobre, Gauthier Gidel</i>	
Position: Scaling Simulation is Neither Necessary nor Sufficient for In-The-Wild Robot Manipulation	3751
<i>Homanga Bharadhwaj</i>	
Dynamic Facility Location in High Dimensional Euclidean Spaces.....	3763
<i>Sayan Bhattacharya, Gramoz Goranci, Shaofeng H.-C. Jiang, Yi Qian, Yubo Zhang</i>	

Total Variation Distance Meets Probabilistic Inference.....	3776
<i>Arnab Bhattacharyya, Sutanu Gayen, Kuldeep S. Meel, Dimitrios Myrisiotis, A. Pavan, N. V. Vinodchandran</i>	
Hierarchical State Space Models for Continuous Sequence-To-Sequence Modeling	3795
<i>Raunaq Bhirangi, Chenyu Wang, Venkatesh Pattabiraman, Carmel Majidi, Abhinav Gupta, Tess Hellebrekers, Lerrel Pinto</i>	
Why Do Variational Autoencoders Really Promote Disentanglement?.....	3817
<i>Pratik Bhowal, Achint Soni, Sirisha Rambhatla</i>	
Best of Both Worlds Guarantees for Smoothed Online Quadratic Optimization.....	3850
<i>Neelkamal Bhuyan, Debankur Mukherjee, Adam Wierman</i>	
Multi-Patch Prediction: Adapting Language Models for Time Series Representation Learning.....	3889
<i>Yuxuan Bian, Xuan Ju, Jiangtong Li, Zhijian Xu, Dawei Cheng, Qiang Xu</i>	
Naive Bayes Classifiers Over Missing Data: Decision and Poisoning.....	3913
<i>Song Bian, Xiating Ouyang, Zhiwei Fan, Paraschos Koutris</i>	
How Well Can LLMs Negotiate? mediumpurpleNegotiationArena Platform and Analysis	3935
<i>Federico Bianchi, Patrick John Chia, Mert Yuksekgonul, Jacopo Tagliabue, Dan Jurafsky, James Zou</i>	
Scalable Safe Policy Improvement for Factored Multi-Agent MDPs	3952
<i>Federico Bianchi, Edoardo Zorzi, Alberto Castellini, Thiago D. Simão, Matthijs T. J. Spaan, Alessandro Farinelli</i>	
Improving Fine-Grained Understanding in Image-Text Pre-Training	3974
<i>Ioana Bica, Anastasija Ilic, Matthias Bauer, Goker Erdogan, Matko Bošnjak, Christos Kaplanis, Alexey A. Gritsenko, Matthias Minderer, Charles Blundell, Razvan Pa. Canu, Jovana Mitrovic</i>	
Position: Explain to Question Not to Justify	3996
<i>Przemyslaw Biecek, Wojciech Samek</i>	
ETHER: Efficient Finetuning of Large-Scale Models with Hyperplane Reflections	4007
<i>Massimo Bini, Karsten Roth, Zeynep Akata, Anna Khoreva</i>	

PART 6

Incorporating Information into Shapley Values: Reweighting Via a Maximum Entropy Approach.....	4027
<i>Darya Biparva, Donatello Materassi</i>	
Graph2Tac: Online Representation Learning of Formal Math Concepts	4046
<i>Lasse Blaauwbroek, Miroslav Olšák, Jason Rute, Fidel Ivan Schaposnik Massolo, Jelle Piepenbrock, Vasily Pestun</i>	
Biharmonic Distance of Graphs and Its Higher-Order Variants: Theoretical Properties with Applications to Centrality and Clustering	4077
<i>Mitchell Black, Lucy Lin, Weng-Keen Wong, Amir Nayyeri</i>	
Comparing Graph Transformers Via Positional Encodings.....	4103
<i>Mitchell Black, Zhengchao Wan, Gal Mishne, Amir Nayyeri, Yusu Wang</i>	
Stability Evaluation Through Distributional Perturbation Analysis	4140
<i>Jose Blanchet, Peng Cui, Jiajin Li, Jiashuo Liu</i>	
Dynamic Survival Analysis with Controlled Latent States.....	4160
<i>Linus Bleistein, Van-Tuan Nguyen, Adeline Fermanian, Agathe Guilloux</i>	
Beyond ELBOs: A Large-Scale Evaluation of Variational Methods for Sampling	4205
<i>Denis Blessing, Xiaogang Jia, Johannes Esslinger, Francisco Vargas, Gerhard Neumann</i>	

Shifted Interpolation for Differential Privacy.....	4230
<i>Jinho Bok, Weijie J. Su, Jason Altschuler</i>	
How Spurious Features Are Memorized: Precise Analysis for Random and NTK Features.....	4267
<i>Simone Bombari, Marco Mondelli</i>	
Towards Understanding the Word Sensitivity of Attention Layers: A Study Via Random Features	4300
<i>Simone Bombari, Marco Mondelli</i>	
Position: Machine Learning-Powered Assessments of the EU Digital Services Act Aid Quantify Policy Impacts on Online Harms.....	4329
<i>Eleonora Bonel, Luca Nannini, Davide Bassi, Michele Joshua Maggini</i>	
A Dynamical Model of Neural Scaling Laws.....	4345
<i>Blake Bordelon, Alexander Atanasov, Cengiz Pehlevan</i>	
A New Computationally Efficient Algorithm to Solve Feature Selection for Functional Data Classification in High-Dimensional Spaces	4383
<i>Tobia Boschi, Francesca Bonin, Rodrigo Ordonez-Hurtado, Alessandra Pascale, Jonathan Epperlein</i>	
Random Matrix Theory Improved Fréchet Mean of Symmetric Positive Definite Matrices	4403
<i>Florent Bouchard, Ammar Mian, Malik Tiomoko, Guillaume Ginolhac, Frédéric Pascal</i>	
Improving Neural Additive Models with Bayesian Principles	4416
<i>Kouroche Bouchiat, Alexander Immer, Hugo Yèche, Gunnar Rätsch, Vincent Fortuin</i>	
S Ω I: Score-Based O-INFORMATION Estimation	4444
<i>Mustapha Bounoua, Giulio Franzese, Pietro Michiardi</i>	
On Dimensionality of Feature Vectors in MPNNs	4472
<i>César Bravo, Alexander Kozachinskiy, Cristóbal Rojas</i>	
Integrating Multimodal Data for Joint Generative Modeling of Complex Dynamics	4482
<i>Manuel Brenner, Florian Hess, Georgia Koppe, Daniel Durstewitz</i>	
Fully-Dynamic Approximate Decision Trees with Worst-Case Update Time Guarantees	4517
<i>Marco Bressan, Mauro Sozio</i>	
Applying Language Models to Algebraic Topology: Generating Simplicial Cycles Using Multi- Labeling in Wu's Formula	4542
<i>Kirill Brilliantov, Fedor Pavutnitskiy, Dmitry Pasechnyuk, German Magai</i>	
Private Gradient Descent for Linear Regression: Tighter Error Bounds and Instance-Specific Uncertainty Estimation.....	4561
<i>Gavin Brown, Krishnamurthy (Dj) Dvijotham, Georgina Evans, Daogao Liu, Adam Smith, Abhradeep Thakurta</i>	
Scalable AI Safety Via Doubly-Efficient Debate.....	4585
<i>Jonah Brown-Cohen, Geoffrey Irving, Georgios Piliouras</i>	
Genie: Generative Interactive Environments.....	4603
<i>Jake Bruce, Michael Dennis, Ashley Edwards, Jack Parker-Holder, Yuge (Jimmy) Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, Yusuf Aytar, Sarah Bechtle, Feryal Behbahani, Stephanie Chan, Nicolas Heess, Lucy Gonzalez, Simon Osindero, Sherjil Ozair, Scott Reed, Jingwei Zhang, Konrad Zolna, Jeff Clune, Nando De Freitas, Satinder Singh, Tim Rocktäschel</i>	
HAMLET: Graph Transformer Neural Operator for Partial Differential Equations.....	4624
<i>Andrey Bryutkin, Jiahao Huang, Zhongying Deng, Guang Yang, Carola-Bibiane Schönlieb, Angelica Aviles-Rivero</i>	
Provably Neural Active Learning Succeeds Via Prioritizing Perplexing Samples	4642
<i>Dake Bu, Wei Huang, Taiji Suzuki, Ji Cheng, Qingfu Zhang, Zhiqiang Xu, Hau-San Wong</i>	

Tackling Prevalent Conditions in Unsupervised Combinatorial Optimization: Cardinality, Minimum, Covering, and More.....	4696
<i>Fanchen Bu, Hyeonsoo Jo, Soo Yong Lee, Sungsoo Ahn, Kijung Shin</i>	
Differentially Private Bias-Term Fine-Tuning of Foundation Models	4730
<i>Zhiqi Bu, Yu-Xiang Wang, Sheng Zha, George Karypis</i>	
Bayesian Optimization of Function Networks with Partial Evaluations	4752
<i>Poompol Buathong, Jiayue Wan, Raul Astudillo, Samuel Daulton, Maximilian Balandat, Peter I. Frazier</i>	
Robustness of Nonlinear Representation Learning	4785
<i>Simon Buchholz, Bernhard Schölkopf</i>	
Density-Softmax: Efficient Test-Time Model for Uncertainty Estimation and Robustness Under Distribution Shifts	4822
<i>Ha Manh Bui, Anqi Liu</i>	
Explaining Graph Neural Networks Via Structure-Aware Interaction Index.....	4854
<i>Ngoc Bui, Hieu Trung Nguyen, Viet Anh Nguyen, Rex Ying</i>	

PART 7

Assessing Large Language Models on Climate Information.....	4884
<i>Jannis Bulian, Mike S. Schäfer, Afra Amini, Heidi Lam, Massimiliano Ciaramita, Ben Gaiarin, Michelle Chen Hübscher, Christian Buck, Niels G. Mede, Markus Leippold, Nadine Strauß</i>	
Semantically-Correlated Memories in a Dense Associative Model.....	4936
<i>Thomas F Burns</i>	
Weak-To-Strong Generalization: Eliciting Strong Capabilities with Weak Supervision	4971
<i>Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, Jeff Wu</i>	
CodeIt: Self-Improving Language Models with Prioritized Hindsight Replay	5013
<i>Natasha Butt, Blazej Manczak, Auke Wiggers, Corrado Rainone, David W. Zhang, Michaël Defferrard, Taco Cohen</i>	
How Uniform Random Weights Induce Non-Uniform Bias: Typical Interpolating Neural Networks Generalize with Narrow Teachers.....	5035
<i>Gon Buzaglo, Itamar Harel, Mor Shpigel Nacson, Alon Brutzkus, Nathan Srebro, Daniel Soudry</i>	
Estimating Distributional Treatment Effects in Randomized Experiments: Machine Learning for Variance Reduction.....	5082
<i>Undral Byambadalai, Tatsushi Oka, Shota Yasui</i>	
Learning Associative Memories with Gradient Descent	5114
<i>Vivien Cabannes, Berfin Simsek, Alberto Bietti</i>	
Bridging Environments and Language with Rendering Functions and Vision-Language Models	5135
<i>Théo Cachet, Christopher R. Dance, Olivier Sigaud</i>	
Vocabulary for Universal Approximation: A Linguistic Perspective of Mapping Compositions	5189
<i>Yongqiang Cai</i>	
Medusa: Simple LLM Inference Acceleration Framework with Multiple Decoding Heads	5209
<i>Tianle Cai, Yuhong Li, Zhengyang Geng, Hongwu Peng, Jason D. Lee, Deming Chen, Tri Dao</i>	
Enhancing Cross-Modal Fine-Tuning with Gradually Intermediate Modality Generation	5236
<i>Lincan Cai, Shuang Li, Wenxuan Ma, Jingxuan Kang, Binhui Xie, Zixun Sun, Chengwei Zhu</i>	

Batch and Match: Black-Box Variational Inference with a Score-Based Divergence	5258
<i>Diana Cai, Chirag Modi, Loucas Pillaud-Vivien, Charles C. Margossian, Robert M. Gower, David M. Blei, Lawrence K. Saul</i>	
Flextron: Many-In-One Flexible Large Language Model	5298
<i>Ruisi Cai, Saurav Muralidharan, Greg Heinrich, Hongxu Yin, Zhangyang Wang, Jan Kautz, Pavlo Molchanov</i>	
Accelerated Algorithms for Constrained Nonconvex-Nonconcave Min-Max Optimization and Comonotone Inclusion.....	5312
<i>Yang Cai, Argyris Oikonomou, Weiqiang Zheng</i>	
LoCoCo: Dropping in Convolutions for Long Context Compression.....	5348
<i>Ruisi Cai, Yuandong Tian, Zhangyang Wang, Beidi Chen</i>	
On Gradient-Like Explanation Under a Black-Box Setting: When Black-Box Explanations Become as Good as White-Box.....	5360
<i>Yi Cai, Gerhard Wunder</i>	
Sample-Specific Masks for Visual Reprogramming-Based Prompting.....	5383
<i>Chengyi Cai, Zesheng Ye, Lei Feng, Jianzhong Qi, Feng Liu</i>	
Human Alignment of Large Language Models Through Online Preference Optimisation	5409
<i>Daniele Calandriello, Zhaohan Daniel Guo, Remi Munos, Mark Rowland, Yunhao Tang, Bernardo Avila Pires, Pierre Harvey Richemond, Charline Le Lan, Michal Valko, Tianqi Liu, Rishabh Joshi, Zeyu Zheng, Bilal Piot</i>	
Partially Stochastic Infinitely Deep Bayesian Neural Networks	5436
<i>Sergio Calvo-Ordoñez, Matthieu Meunier, Francesco Piatti, Yuantao Shi</i>	
Generative Flows on Discrete State-Spaces: Enabling Multimodal Flows with Applications to Protein Co-Design	5453
<i>Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, Tommi Jaakkola</i>	
Mimicking Better by Matching the Approximate Action Distribution.....	5513
<i>João A. Cândido Ramos, Lionel Blondé, Naoya Takeishi, Alexandros Kalousis</i>	
Graph Positional and Structural Encoder	5533
<i>Semih Cantürk, Renming Liu, Olivier Lapointe-Gagné, Vincent Létourneau, Guy Wolf, Dominique Beaini, Ladislav Rampásek</i>	
Successor Features for Efficient Multi-Subject Controlled Text Generation.....	5567
<i>Meng Cao, Mehdi Fatemi, Jackie Chi Kit Cheung, Samira Shabianian</i>	
Limited Preference Aided Imitation Learning from Imperfect Demonstrations.....	5584
<i>Xingchen Cao, Fan-Ming Luo, Junyin Ye, Tian Xu, Zhilong Zhang, Yang Yu</i>	
Predictive Dynamic Fusion	5608
<i>Bing Cao, Yinan Xia, Yi Ding, Changqing Zhang, Qinghua Hu</i>	
Envisioning Outlier Exposure by Large Language Models for Out-Of-Distribution Detection.....	5629
<i>Chentao Cao, Zhun Zhong, Zhanke Zhou, Yang Liu, Tongliang Liu, Bo Han</i>	
Can a Few Decide for Many? the Metric Distortion of Sortition	5660
<i>Ioannis Caragiannis, Evi Micha, Jannik Peters</i>	
Stealing Part of a Production Language Model.....	5680
<i>Nicholas Carlini, Daniel Paleka, Krishnamurthy (Dj) Dvijotham, Thomas Steinke, Jonathan Hayase, A. Feder Cooper, Katherine Lee, Matthew Jagielski, Milad Nasr, Arthur Conmy, Eric Wallace, David Rolnick, Florian Tramèr</i>	

PART 8

AI Alignment with Changing and Influenceable Reward Functions	5706
<i>Micah Carroll, Davis Foote, Anand Siththaranjan, Stuart Russell, Anca Dragan</i>	
Near-Optimal Regret in Linear MDPs with Aggregate Bandit Feedback.....	5757
<i>Asaf Cassel, Haipeng Luo, Aviv Rosenberg, Dmitry Sotnikov</i>	
Online Learning Under Budget and ROI Constraints Via Weak Adaptivity.....	5792
<i>Matteo Castiglioni, Andrea Celli, Christian Kroer</i>	
How Smooth is Attention?.....	5817
<i>Valérie Castin, Pierre Ablin, Gabriel Peyré</i>	
Hierarchical Integral Probability Metrics: A Distance on Random Probability Measures with Low Sample Complexity	5841
<i>Marta Catalano, Hugo Lavenant</i>	
On the Implicit Bias of Adam.....	5862
<i>Matias D. Cattaneo, Jason M. Klusowski, Boris Shigida</i>	
Acquiring Diverse Skills Using Curriculum Reinforcement Learning with Mixture of Experts	5907
<i>Onur Celik, Aleksandar Taranovic, Gerhard Neumann</i>	
Centralized Selection with Preferences in the Presence of Biases.....	5934
<i>L. Elisa Celis, Amit Kumar, Nisheeth K. Vishnoi, Andrew Xu</i>	
Using Left and Right Brains Together: Towards Vision and Language Planning	5982
<i>Jun Cen, Chenfei Wu, Xiao Liu, Shengming Yin, Yixuan Pei, Jinglong Yang, Qifeng Chen, Nan Duan, Jianguo Zhang</i>	
Feasibility Consistent Representation Learning for Safe Reinforcement Learning.....	6002
<i>Zhepeng Cen, Yihang Yao, Zuxin Liu, Ding Zhao</i>	
Simple Ingredients for Offline Reinforcement Learning.....	6020
<i>Edoardo Ceting, Andrea Tirinzoni, Matteo Pirota, Alessandro Lazaric, Yann Ollivier, Ahmed Touati</i>	
Regularizing with Pseudo-Negatives for Continual Self-Supervised Learning.....	6048
<i>Sungmin Cha, Kyunghyun Cho, Taesup Moon</i>	
Auditing Private Prediction	6066
<i>Karan Chadha, Matthew Jagielski, Nicolas Papernot, Christopher A. Choquette-Choo, Milad Nasr</i>	
Position: On the Possibilities of AI-Generated Text Detection.....	6093
<i>Souradip Chakraborty, Amrit Singh Bedi, Sicheng Zhu, Bang An, Dinesh Manocha, Furong Huang</i>	
MaxMin-RLHF: Alignment with Diverse Human Preferences	6116
<i>Souradip Chakraborty, Jiahao Qiu, Hui Yuan, Alec Koppel, Dinesh Manocha, Furong Huang, Amrit Singh Bedi, Mengdi Wang</i>	
Dense Reward for Free in Reinforcement Learning from Human Feedback	6136
<i>Alex J. Chan, Hao Sun, Samuel Holt, Mihaela Van Der Schaar</i>	
Scribble-Supervised Semantic Segmentation with Prototype-Based Feature Augmentation	6155
<i>Guiyang Chan, Pengcheng Zhang, Hai Dong, Shunhui Ji, Bainian Chen</i>	
Feature Importance Disparities for Data Bias Investigations	6170
<i>Peter W Chang, Leor Fishman, Seth Neal</i>	
Inferring Dynamic Networks from Marginals with Iterative Proportional Fitting	6202
<i>Serina Chang, Frederic Koehler, Zhaonan Qu, Jure Leskovec, Johan Ugander</i>	

LaMAGIC: Language-Model-Based Topology Generation for Analog Integrated Circuits.....	6253
<i>Chen-Chia Chang, Yikang Shen, Shaoze Fan, Jing Li, Shun Zhang, Ningyuan Cao, Yiran Chen, Xin Zhang</i>	
MagicPose: Realistic Human Poses and Facial Expressions Retargeting with Identity-Aware Diffusion.....	6263
<i>Di Chang, Yichun Shi, Quankai Gao, Hongyi Xu, Jessica Fu, Guoxian Song, Qing Yan, Yizhe Zhu, Xiao Yang, Mohammad Soleymani</i>	
From Biased Selective Labels to Pseudo-Labels: An Expectation-Maximization Framework for Learning from Biased Decisions	6286
<i>Trenton Chang, Jenna Wiens</i>	
On the Role of Edge Dependency in Graph Generative Models	6325
<i>Sudhanshu Chanpuriya, Cameron Musco, Konstantinos Sotiropoulos, Charalampos Tsourakakis</i>	
Performance Bounds for Active Binary Testing with Information Maximization.....	6346
<i>Aditya Chattopadhyay, Benjamin D. Haeffele, René Vidal, Donald Geman</i>	
Target Networks and Over-Parameterization Stabilize Off-Policy Bootstrapping with Function Approximation.....	6372
<i>Fengdi Che, Chenjun Xiao, Jincheng Mei, Bo Dai, Ramki Gummadi, Oscar A Ramirez, Christopher K Harris, A. Rupam Mahmood, Dale Schuurmans</i>	
PlanDQ: Hierarchical Plan Orchestration Via D-Conductor and Q-Performer	6397
<i>Chang Chen, Junyeob Baek, Fei Deng, Kenji Kawaguchi, Caglar Gulcehre, Sungjin Ahn</i>	
How Interpretable Are Interpretable Graph Neural Networks?.....	6413
<i>Yongqiang Chen, Yatao Bian, Bo Han, James Cheng</i>	
Doubly Robust Causal Effect Estimation Under Networked Interference Via Targeted Learning	6457
<i>Weilin Chen, Ruichu Cai, Zeqin Yang, Jie Qiao, Yuguang Yan, Zijian Li, Zhifeng Hao</i>	
Feature Attribution with Necessity and Sufficiency via Dual-Stage Perturbation Test for Causal Explanation.....	6486
<i>Xuexin Chen, Ruichu Cai, Zhengting Huang, Yuxuan Zhu, Julien Horwood, Zhifeng Hao, Zijian Li, José Miguel Hernández-Lobato</i>	
InstructZero: Efficient Instruction Optimization for Black-Box Large Language Models.....	6503
<i>Lichang Chen, Jiu hai Chen, Tom Goldstein, Heng Huang, Tianyi Zhou</i>	
MaSS: Multi-Attribute Selective Suppression for Utility-Preserving Data Transformation from an Information-Theoretic Perspective	6519
<i>Yizhuo Chen, Chun-Fu (Richard) Chen, Hsiang Hsu, Shaohan Hu, Marco Pistoia, Tarek Abdelzaher</i>	
Policy-Conditioned Environment Models Are More Generalizable.....	6539
<i>Ruifeng Chen, Xiong-Hui Chen, Yihao Sun, Siyuan Xiao, Minhui Li, Yang Yu</i>	

PART 9

MLLM-As-A-Judge: Assessing Multimodal LLM-As-A-Judge with Vision-Language Benchmark	6562
<i>Dongping Chen, Ruoxi Chen, Shilin Zhang, Yaochen Wang, YINUO Liu, Huichi Zhou, Qihui Zhang, Yao Wan, Pan Zhou, Lichao Sun</i>	
Premise Order Matters in Reasoning with Large Language Models	6596
<i>Xinyun Chen, Ryan A. Chi, Xuezhi Wang, Denny Zhou</i>	
Self-Play Fine-Tuning Converts Weak Language Models to Strong Language Models	6621
<i>Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, Quanquan Gu</i>	
Robust Classification Via a Single Diffusion Model.....	6643
<i>Huanran Chen, Yinpeng Dong, Zhengyi Wang, Xiao Yang, Chengqi Duan, Hang Su, Jun Zhu</i>	

Relational Learning in Pre-Trained Models: A Theory from Hypergraph Recovery Perspective	6666
<i>Yang Chen, Cong Fang, Zhouchen Lin, Bing Liu</i>	
Towards AutoAI: Optimizing a Machine Learning System with Black-Box and Differentiable Components.....	6699
<i>Zhiliang Chen, Chuan Sheng Foo, Bryan Kian Hsiang Low</i>	
Probabilistic Forecasting with Stochastic Interpolants and Föllmer Processes	6728
<i>Yifan Chen, Mark Goldstein, Mengjian Hua, Michael S. Albergó, Nicholas M. Boffi, Eric Vanden-Eijnden</i>	
CogDPM: Diffusion Probabilistic Models Via Cognitive Predictive Coding.....	6757
<i>Kaiyuan Chen, Xingzhuo Guo, Yu Zhang, Jianmin Wang, Mingsheng Long</i>	
On Interpolating Experts and Multi-Armed Bandits	6776
<i>Houshuang Chen, Yuchen He, Chihao Zhang</i>	
Bagged Deep Image Prior for Recovering Images in the Presence of Speckle Noise	6803
<i>Xi Chen, Zhewen Hou, Christopher A. Metzler, Arian Maleki, Shirin Jalali</i>	
Exact Conversion of In-Context Learning to Model Weights in Linearized-Attention Transformers.....	6833
<i>Brian K Chen, Tianyang Hu, Hui Jin, Hwee Kuan Lee, Kenji Kawaguchi</i>	
Accelerated Policy Gradient for s-Rectangular Robust MDPs with Large State Spaces.....	6847
<i>Ziyi Chen, Heng Huang</i>	
Accelerated Policy Gradient: On the Convergence Rates of the Nesterov Momentum for Reinforcement Learning.....	6881
<i>Yen-Ju Chen, Nai-Chieh Huang, Ching-Pei Lee, Ping-Chun Hsieh</i>	
From Yes-Men to Truth-Tellers: Addressing Sycophancy in Large Language Models with Pinpoint Tuning	6950
<i>Wei Chen, Zhen Huang, Liang Xie, Binbin Lin, Houqiang Li, Le Lu, Xinmei Tian, Deng Cai, Yonggang Zhang, Wenxiao Wan, Xu Shen, Jieping Ye</i>	
Improved Communication-Privacy Trade-Offs in L_2 Mean Estimation Under Streaming Differential Privacy.....	6973
<i>Wei-Ning Chen, Berivan Isik, Peter Kairouz, Albert No, Sewoong Oh, Zheng Xu</i>	
Offline Transition Modeling Via Contrastive Energy Learning	6992
<i>Ruifeng Chen, Chengxing Jia, Zefang Huang, Tian-Shuo Liu, Xu-Hui Liu, Yang Yu</i>	
Efficient Pareto Manifold Learning with Low-Rank Structure	7015
<i>Weiyu Chen, James T. Kwok</i>	
Toward Adaptive Reasoning in Large Language Models with Thought Rollback	7033
<i>Sijia Chen, Baochun Li</i>	
Identifiability Matters: Revealing the Hidden Recoverable Condition in Unbiased Learning to Rank.....	7057
<i>Mouxian Chen, Chenghao Liu, Zemin Liu, Zhuo Li, Jianling Sun</i>	
High-Dimensional Kernel Methods Under Covariate Shift: Data-Dependent Implicit Regularization.....	7081
<i>Yihang Chen, Fanghui Liu, Taiji Suzuki, Volkan Cevher</i>	
DiJiang: Efficient Large Language Models Through Compact Kernelization	7103
<i>Hanting Chen, Zhicheng Liu, Xutao Wang, Yuchuan Tian, Yunhe Wang</i>	
GeoMFormer: A General Architecture for Geometric Molecular Representation Learning.....	7118
<i>Tianlang Chen, Shengjie Luo, Di He, Shuxin Zheng, Tie-Yan Liu, Liwei Wang</i>	
TENG: Time-Evolving Natural Gradient for Solving PDEs with Deep Neural Nets Toward Machine Precision	7143
<i>Zhuo Chen, Jacob McCarran, Esteban Vizcaino, Marin Soljaèia, Di Luo</i>	

EE-LLM: Large-Scale Training and Inference of Early-Exit Large Language Models with 3D Parallelism	7163
<i>Yanxi Chen, Xuchen Pan, Yaliang Li, Bolin Ding, Jingren Zhou</i>	
TimeMIL: Advancing Multivariate Time Series Classification Via a Time-Aware Multiple Instance Learning	7190
<i>Xiwen Chen, Peijie Qiu, Wenhui Zhu, Huayu Li, Hao Wang, Aristeidis Sotiras, Yalin Wang, Abolfazl Razi</i>	
AegisFL: Efficient and Flexible Privacy-Preserving Byzantine-Robust Cross-Silo Federated Learning	7207
<i>Dong Chen, Hongyuan Qu, Guangwu Xu</i>	
MAGDi: Structured Distillation of Multi-Agent Interaction Graphs Improves Reasoning in Smaller Language Models	7220
<i>Justin Chih-Yao Chen, Swarnadeep Saha, Elias Stengel-Eskin, Mohit Bansal</i>	
CaRiNG: Learning Temporal Causal Representation Under Non-Invertible Generation Process.....	7236
<i>Guangyi Chen, Yifan Shen, Zhenhao Chen, Xiangchen Song, Yuewen Sun, Weiran Yao, Xiao Liu, Kun Zhang</i>	
GRATH: Gradual Self-Truthifying for Large Language Models.....	7260
<i>Weixin Chen, Dawn Song, Bo Li</i>	
Unleashing the Power of Meta-Tuning for Few-Shot Generalization Through Sparse Interpolated Experts.....	7280
<i>Shengzhuang Chen, Jihoon Tack, Yunqiao Yang, Yee Whye Teh, Jonathan Richard Schwarz, Ying Wei</i>	
Performative Prediction with Bandit Feedback: Learning Through Reparameterization.....	7298
<i>Yatong Chen, Wei Tang, Chien-Ju Ho, Yang Liu</i>	
Self-Attention Through Kernel-Eigen Pair Sparse Variational Gaussian Processes.....	7325
<i>Yingyi Chen, Qinghua Tao, Francesco Tonin, Johan A. K. Suykens</i>	
Recovering Labels from Local Updates in Federated Learning	7346
<i>Huancheng Chen, Haris Vikalo</i>	
SelfIE: Self-Interpretation of Large Language Model Embeddings.....	7373
<i>Haozhe Chen, Carl Vondrick, Chengzhi Mao</i>	

PART 10

Locally Differentially Private Decentralized Stochastic Bilevel Optimization with Guaranteed Convergence Accuracy	7389
<i>Ziqin Chen, Yongqiang Wang</i>	
Subequivariant Reinforcement Learning in 3D Multi-Entity Physical Environments	7440
<i>Runfa Chen, Ling Wang, Yu Du, Tianrui Xue, Fuchun Sun, Jianwei Zhang, Wenbing Huang</i>	
A General Framework for Learning from Weak Supervision.....	7462
<i>Hao Chen, Jindong Wang, Lei Feng, Xiang Li, Yidong Wang, Xing Xie, Masashi Sugiyama, Rita Singh, Bhiksha Raj</i>	
Diffusion Model-Augmented Behavioral Cloning	7486
<i>Shang-Fu Chen, Hsiang-Chun Wang, Ming-Hao Hsu, Chun-Mao Lai, Shao-Hua Sun</i>	
AutoOS: Make Your OS More Powerful by Exploiting Large Language Models	7511
<i>Huilai Chen, Yuanbo Wen, Limin Cheng, Shouxu Kuang, Yumeng Liu, Weijia Li, Ling Li, Rui Zhang, Xinkai Song, Wei Li, Qi Guo, Yunji Chen</i>	
Positional Knowledge is All You Need: Position-Induced Transformer (PiT) for Operator Learning	7526
<i>Junfeng Chen, Kailiang Wu</i>	

In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation.....	7553
<i>Shiqi Chen, Miao Xiong, Junteng Liu, Zhengxuan Wu, Teng Xiao, Siyang Gao, Junxian He</i>	
Split-Ensemble: Efficient OOD-Aware Ensemble Via Task and Model Splitting	7568
<i>Anthony Chen, Huanrui Yang, Yulu Gan, Denis A Gudovskiy, Zhen Dong, Haofan Wang, Tomoyuki Okuno, Yohei Nakata, Kurt Keutzer, Shanghang Zhang</i>	
Deep Demonstration Tracing: Learning Generalizable Imitator Policy for Runtime Imitation from a Single Demonstration	7586
<i>Xiong-Hui Chen, Junyin Ye, Hang Zhao, Yi-Chen Li, Xu-Hui Liu, Haoran Shi, Yu-Yan Xu, Zhihao Ye, Si-Hang Yang, Yang Yu, Anqi Huang, Kai Xu, Zongzhang Zhang</i>	
DRCT: Diffusion Reconstruction Contrastive Training Towards Universal Detection of Diffusion Generated Images	7621
<i>Baoying Chen, Jishen Zeng, Jianquan Yang, Rui Yang</i>	
E(3)-Equivariant Actor-Critic Methods for Cooperative Multi-Agent Reinforcement Learning	7640
<i>Dingyang Chen, Qi Zhang</i>	
FedMBridge: Bridgeable Multimodal Federated Learning	7667
<i>Jiayi Chen, Aidong Zhang</i>	
Revealing the Dark Secrets of Extremely Large Kernel ConvNets on Robustness.....	7687
<i>Honghao Chen, Yurong Zhang, Xiaokun Feng, Xiangxiang Chu, Kaiqi Huang</i>	
One for All: A Universal Generator for Concept Unlearnability Via Multi-Modal Alignment.....	7700
<i>Chaochao Chen, Jiaming Zhang, Yuyuan Li, Zhongxuan Han</i>	
Generating In-Distribution Proxy Graphs for Explaining Graph Neural Networks	7712
<i>Zhuomin Chen, Jiaying Zhang, Jingchao Ni, Xiaoting Li, Yuchen Bian, Md Mezbahul Islam, Ananda Mohan Mondal, Hua Wei, Dongsheng Luo</i>	
Diffusive Gibbs Sampling	7731
<i>Wenlin Chen, Mingtian Zhang, Brooks Paige, José Miguel Hernández-Lobato, David Barber</i>	
Provable Risk-Sensitive Distributional Reinforcement Learning with General Function Approximation.....	7748
<i>Yu Chen, Xiangcheng Zhang, Siwei Wang, Longbo Huang</i>	
MoE-RBench: Towards Building Reliable Language Models with Sparse Mixture-Of-Experts	7792
<i>Guanjie Chen, Xinyu Zhao, Tianlong Chen, Yu Cheng</i>	
LLaGA: Large Language and Graph Assistant.....	7809
<i>Runjin Chen, Tong Zhao, Ajay Jaiswal, Neil Shah, Zhangyang Wang</i>	
HALC: Object Hallucination Reduction Via Adaptive Focal-Contrast Decoding.....	7824
<i>Zhaorun Chen, Zhuokai Zhao, Hongyin Luo, Huaxiu Yao, Bo Li, Jiawei Zhou</i>	
Compact Optimality Verification for Optimization Proxies	7847
<i>Wenbo Chen, Haoruo Zhao, Mathieu Tanneau, Pascal Van Hentenryck</i>	
Enhancing Implicit Shape Generators Using Topological Regularizations.....	7864
<i>Liyan Chen, Yan Zhen, Yang Li, Lohit Anirudh Jagarapu, Haoxiang Li, Hao Kang, Gang Hua, Qixing Huang</i>	
Do Models Explain Themselves? Counterfactual Simulatability of Natural Language Explanations	7880
<i>Yanda Chen, Ruiqi Zhong, Narutatsu Ri, Chen Zhao, He He, Jacob Steinhardt, Zhou Yu, Kathleen McKeown</i>	
On the Trajectory Regularity of ODE-Based Diffusion Sampling	7905
<i>Defang Chen, Zhenyu Zhou, Can Wang, Chunhua Shen, Siwei Lyu</i>	
Odin: Disentangled Reward Mitigates Hacking in RLHF	7935
<i>Lichang Chen, Chen Zhu, Jiu-hai Chen, Davit Soselia, Tianyi Zhou, Tom Goldstein, Heng Huang, Mohammad Shoeybi, Bryan Catanzaro</i>	

Stacking Deep Set Networks and Pooling by Quantiles.....	7953
<i>Zhuojun Chen, Xinghua Zhu, Dongzhe Su, Justin C. I. Chuang</i>	
What Can Transformer Learn with Varying Depth? Case Studies on Sequence Learning Tasks	7972
<i>Xingwu Chen, Difan Zou</i>	
Transformers Implement Functional Gradient Descent to Learn Non-Linear Functions in Context.....	8002
<i>Xiang Cheng, Yuxin Chen, Suvrit Sra</i>	
Layerwise Change of Knowledge in Neural Networks	8038
<i>Xu Cheng, Lei Cheng, Zhaoran Peng, Yang Xu, Tian Han, Quanshi Zhang</i>	
Reference Neural Operators: Learning the Smooth Dependence of Solutions of PDEs on Geometric Deformations	8060
<i>Ze Cheng, Zhongkai Hao, Xiaoqiang Wang, Jianing Huang, Youjia Wu, Xudan Liu, Yiru Zhao, Songming Liu, Hang Su</i>	
Causal Inference Out of Control: Estimating Performativity Without Treatment Randomization	8077
<i>Gary Cheng, Moritz Hardt, Celestine Mendler-Dünnler</i>	
BadPart: Unified Black-Box Adversarial Patch Attacks Against Pixel-Wise Regression Tasks	8104
<i>Zhiyuan Cheng, Zhaoyi Liu, Tengda Guo, Shiwei Feng, Dongfang Liu, Mingjie Tang, Xiangyu Zhang</i>	
GaussianPro: 3D Gaussian Splatting with Progressive Propagation	8123
<i>Kai Cheng, Xiaoxiao Long, Kaizhi Yang, Yao Yao, Wei Yin, Yuexin Ma, Wenping Wang, Xuejin Chen</i>	
Characterizing Overfitting in Kernel Ridgeless Regression Through the Eigenspectrum.....	8141
<i>Tin Sum Cheng, Aurelien Lucchi, Anastasis Kratsios, David Belius</i>	
Efficient Black-Box Adversarial Attacks Via Bayesian Optimization Guided by a Function Prior	8163
<i>Shuyu Cheng, Yibo Miao, Yinpeng Dong, Xiao Yang, Xiao-Shan Gao, Jun Zhu</i>	
Can AI Assistants Know What They Don't Know?	8184
<i>Qinyuan Cheng, Tianxiang Sun, Xiangyang Liu, Wenwei Zhang, Zhangyue Yin, Shimin Li, Linyang Li, Zhengfu He, Kai Chen, Xipeng Qiu</i>	
RICE: Breaking Through the Training Bottlenecks of Reinforcement Learning with Explanation	8203
<i>Zelei Cheng, Xian Wu, Jiahao Yu, Sabrina Yang, Gang Wang, Xinyu Xing</i>	
RIME: Robust Preference-Based Reinforcement Learning with Noisy Preferences.....	8229
<i>Jie Cheng, Gang Xiong, Xingyuan Dai, Qinghai Miao, Yisheng Lv, Fei-Yue Wang</i>	

PART 11

Kernel Semi-Implicit Variational Inference.....	8248
<i>Ziheng Cheng, Longlin Yu, Tianyu Xie, Shiyue Zhang, Cheng Zhang</i>	
Creative Text-To-Audio Generation Via Synthesizer Programming	8270
<i>Manuel Cherep, Nikhil Singh, Jessica Shand</i>	
Leveraging (Biased) Information: Multi-Armed Bandits with Offline Data	8286
<i>Wang Chi Cheung, Lixing Lyu</i>	
Language Models as Science Tutors	8310
<i>Alexis Chevalier, Jiayi Geng, Alexander Wettig, Howard Chen, Sebastian Mizera, Toni Annala, Max Jameson Aragon, Arturo Rodriguez Fanlo, Simon Frieder, Simon Machado, Akshara Prabhakar, Ellie Thieu, Jiachen T. Wang, Zirui Wang, Xindi Wu, Mengzhou Xia, Wenhan Xia, Jiatong Yu, Jun-Jie Zhu, Zhiyong Jason Ren, Sanjeev Arora, Danqi Chen</i>	
Expert Proximity as Surrogate Rewards for Single Demonstration Imitation Learning.....	8336
<i>Chia-Cheng Chiang, Li-Cheng Lan, Wei-Fang Sun, Chien Feng, Cho-Jui Hsieh, Chun-Yi Lee</i>	

Chatbot Arena: An Open Platform for Evaluating LLMs by Human Preference	8359
<i>Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios N. Angelopoulos, Tianle Li, Dacheng Li, Banghua Zhu, Hao Zhang, Michael I. Jordan, Joseph E. Gonzalez, Ion Stoica</i>	
MS-TIP: Imputation Aware Pedestrian Trajectory Prediction	8389
<i>Pranav Singh Chib, Achintya Nath, Paritosh Kabra, Ishu Gupta, Pravendra Singh</i>	
Enhancing Trajectory Prediction Through Self-Supervised Waypoint Distortion Prediction.....	8403
<i>Pranav Singh Chib, Pravendra Singh</i>	
How Flawed is ECE? an Analysis Via Logit Smoothing.....	8417
<i>Muthu Chidambaram, Holden Lee, Colin McSwiggen, Semon Rezchikov</i>	
Safe Exploration in Dose Finding Clinical Trials with Heterogeneous Participants	8436
<i>Isabel Chien, Wessel Bruinsma, Javier Gonzalez Hernandez, Richard E. Turner</i>	
Prompting4Debugging: Red-Teaming Text-To-Image Diffusion Models by Finding Problematic Prompts.....	8468
<i>Zhi-Yi Chin, Chieh-Ming Jiang, Ching-Chun Huang, Pin-Yu Chen, Wei-Chen Chiu</i>	
Peeking with PEAK: Sequential, Nonparametric Composite Hypothesis Tests for Means of Multiple Data Streams.....	8487
<i>Brian Cho, Kyra Gan, Nathan Kallus</i>	
Parameterized Physics-Informed Neural Networks for Parameterized PDEs	8510
<i>Woojin Cho, Minju Jo, Haksoo Lim, Kookjin Lee, Dongeun Lee, Sanghyun Hong, Noseong Park</i>	
Kernel Debiased Plug-In Estimation: Simultaneous, Automated Debiasing Without Influence Functions for Many Target Parameters.....	8534
<i>Brian Cho, Yaroslav Mukhin, Kyra Gan, Ivana Malenica</i>	
Hard Tasks First: Multi-Task Reinforcement Learning Through Task Scheduling	8556
<i>Myungsik Cho, Jongeui Park, Suyoung Lee, Youngchul Sung</i>	
KV-Runahead: Scalable Causal LLM Inference by Parallel Key-Value Cache Generation	8578
<i>Minsik Cho, Mohammad Rastegari, Devang Naik</i>	
Neurodegenerative Brain Network Classification Via Adaptive Diffusion with Temporal Regularization	8593
<i>Hyuna Cho, Jaeyoon Sim, Guorong Wu, Won Hwa Kim</i>	
Tilt and Average : Geometric Adjustment of the Last Layer for Recalibration	8609
<i>Gyusang Cho, Chan-Hyun Youn</i>	
Scalable Wasserstein Gradient Flow for Generative Modeling Through Unbalanced Optimal Transport.....	8629
<i>Jaemoo Choi, Jaewoong Choi, Myungjoo Kang</i>	
Listwise Reward Estimation for Offline Preference-Based Reinforcement Learning	8651
<i>Heewoong Choi, Sangwon Jung, Hongjoon Ahn, Taesup Moon</i>	
BWS: Best Window Selection Based on Sample Scores for Data Pruning Across Broad Ranges.....	8672
<i>Hoyong Choi, Nohyun Ki, Hye Won Chung</i>	
Embodied CoT Distillation from LLM to Off-The-Shelf Agents	8702
<i>Wonje Choi, Woo Kyung Kim, Minjong Yoo, Honguk Woo</i>	
PICLe: Eliciting Diverse Behaviors from Large Language Models with Persona In-Context Learning	8722
<i>Hyeong Kyu Choi, Yixuan Li</i>	
PANDA: Expanded Width-Aware Message Passing Beyond Rewiring	8740
<i>Jeongwhan Choi, Sumin Park, Hyowon Wi, Sung-Bae Cho, Noseong Park</i>	

Online Bipartite Matching with Imperfect Advice	8762
<i>Davin Choo, Themis Gouleakis, Chun Kai Ling, Arnab Bhattacharyya</i>	
A Connection Between Tempering and Entropic Mirror Descent	8782
<i>Nicolas Chopin, Francesca Crucinio, Anna Korba</i>	
Learning Linear Block Error Correction Codes	8801
<i>Yoni Choukroun, Lior Wolf</i>	
A Provably Effective Method for Pruning Experts in Fine-Tuned Sparse Mixture-Of-Experts	8815
<i>Mohammed Nowaz Rabbani Chowdhury, Meng Wang, Kaoutar El Maghraoui, Naigang Wang, Pin-Yu Chen, Christopher Carothers</i>	
SPABA: A Single-Loop and Probabilistic Stochastic Bilevel Algorithm Achieving Optimal Sample Complexity	8848
<i>Tianshu Chu, Dachuan Xu, Wei Yao, Jin Zhang</i>	
How Private Are DP-SGD Implementations?	8904
<i>Lynn Chua, Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, Amer Sinha, Chiyuan Zhang</i>	
Sampling-Based Multi-Dimensional Recalibration.....	8919
<i>Youngseog Chung, Ian Char, Jeff Schneider</i>	
Prompt-Tuning Latent Diffusion Models for Inverse Problems	8941
<i>Hyungjin Chung, Jong Chul Ye, Peyman Milanfar, Mauricio Delbracio</i>	
MusicRL: Aligning Music Generation to Human Preferences	8968
<i>Geoffrey Cideron, Sertan Girgin, Mauro Verzetti, Damien Vincent, Matej Kastelic, Zalán Borsos, Brian McWilliams, Victor Ungureanu, Olivier Bachem, Olivier Pietquin, Matthieu Geist, Léonard Hussenot, Neil Zeghidour, Andrea Agostinelli</i>	
Graph-Based Time Series Clustering for End-To-End Hierarchical Forecasting	8985
<i>Andrea Cini, Danilo Mandic, Cesare Alippi</i>	
Studying K-FAC Heuristics by Viewing Adam Through a Second-Order Lens.....	9000
<i>Ross M. Clarke, José Miguel Hernández-Lobato</i>	
VITS : Variational Inference Thompson Sampling for Contextual Bandits	9033
<i>Pierre Clavier, Tom Huix, Alain Durmus</i>	
CogBench: A Large Language Model Walks into a Psychology Lab	9076
<i>Julian Coda-Forno, Marcel Binz, Jane X. Wang, Eric Schulz</i>	

PART 12

Slicedit: Zero-Shot Video Editing with Text-To-Image Diffusion Models Using Spatio-Temporal Slices	9109
<i>Nathaniel Cohen, Vladimir Kulikov, Matan Kleiner, Inbar Huberman-Spiegelglas, Tomer Michaeli</i>	
Improving Token-Based World Models with Parallel Observation Prediction.....	9138
<i>Lior Cohen, Kaixin Wang, Bingyi Kang, Shie Mannor</i>	
Perturb-And-Project: Differentially Private Similarities and Marginals.....	9161
<i>Vincent Cohen-Addad, Tommaso D'Orsi, Alessandro Epasto, Vahab Mirrokni, Peilin Zhong</i>	
Multi-View Stochastic Block Models.....	9180
<i>Vincent Cohen-Addad, Tommaso D'Orsi, Silvio Lattanzi, Rajai Nasser</i>	
A Near-Linear Time Approximation Algorithm for Beyond-Worst-Case Graph Clustering	9208
<i>Vincent Cohen-Addad, Tommaso D'Orsi, Aida Mousavifar</i>	
Dynamic Correlation Clustering in Sublinear Update Time.....	9230
<i>Vincent Cohen-Addad, Silvio Lattanzi, Andreas Maggiori, Nikos Parotsidis</i>	

A2Q+: Improving Accumulator-Aware Weight Quantization	9275
<i>Ian Colbert, Alessandro Pappalardo, Jakoba Petri-Koenig, Yaman Umuroglu</i>	
Provable Multi-Task Representation Learning by Two-Layer ReLU Neural Networks.....	9292
<i>Liam Collins, Hamed Hassani, Mahdi Soltanolkotabi, Aryan Mokhtari, Sanjay Shakkottai</i>	
Position: Social Choice Should Guide AI Alignment in Dealing with Diverse Human Feedback	9346
<i>Vincent Conitzer, Rachel Freedman, Jobst Heitzig, Wesley H. Holliday, Bob M. Jacobs, Nathan Lambert, Milan Mossé, Eric Pacuit, Stuart Russell, Hailey Schoelkopf, Emanuel Tewolde, William S. Zwicker</i>	
Statistical Inference Under Constrained Selection Bias	9361
<i>Santiago Cortes-Gomez, Mateo Dulce, Carlos Patino, Bryan Wilder</i>	
Scaling Laws for the Value of Individual Data Points in Machine Learning.....	9380
<i>Ian Covert, Wenlong Ji, Tatsunori Hashimoto, James Zou</i>	
Time Series Diffusion in the Frequency Domain	9407
<i>Jonathan Crabbé, Nicolas Huynh, Jan Stanczuk, Mihaela Van Der Schaar</i>	
Conformal Prediction Sets Improve Human Decision Making	9439
<i>Jesse C. Cresswell, Yi Sui, Bhargava Kumar, Noël Vouitsis</i>	
Agent Instructs Large Language Models to Be General Zero-Shot Reasoners	9458
<i>Nicholas Crispino, Kyle Montgomery, Fankun Zeng, Dawn Song, Chenguang Wang</i>	
Scalable High-Resolution Pixel-Space Image Synthesis with Hourglass Diffusion Transformers	9550
<i>Katherine Crowson, Stefan Andreas Baumann, Alex Birch, Tanishq Mathew Abraham, Daniel Z. Kaplan, Enrico Shippole</i>	
Generalization Bounds for Causal Regression: Insights, Guarantees and Sensitivity Analysis	9576
<i>Daniel Csillag, Claudio José Struchiner, Guilherme Tegoni Goedert</i>	
Major-Minor Mean Field Multi-Agent Reinforcement Learning	9603
<i>Kai Cui, Christian Fabian, Anam Tahir, Heinz Koepl</i>	
Learning Latent Space Hierarchical EBM Diffusion Models.....	9633
<i>Jiali Cui, Tian Han</i>	
Harmonizing Generalization and Personalization in Federated Prompt Learning.....	9646
<i>Tianyu Cui, Hongxia Li, Jingya Wang, Ye Shi</i>	
Asymptotics of Feature Learning in Two-Layer Networks After One Gradient-Step	9662
<i>Hugo Cui, Luca Pesce, Yatin Dandi, Florent Krzakala, Yue M. Lu, Lenka Zdeborová, Bruno Loureiro</i>	
Ameliorate Spurious Correlations in Dataset Condensation	9696
<i>Justin Cui, Ruochen Wang, Yuanhao Xiong, Cho-Jui Hsieh</i>	
UltraFeedback: Boosting Language Models with Scaled AI Feedback.....	9722
<i>Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Bingxiang He, Wei Zhu, Yuan Ni, Guotong Xie, Ruobing Xie, Yankai Lin, Zhiyuan Liu, Maosong Sun</i>	
Et Tu Certifications: Robustness Certificates Yield Better Adversarial Examples	9745
<i>Andrew C. Cullen, Shijie Liu, Paul Montague, Sarah M. Erfani, Benjamin I. P. Rubinstein</i>	
Differentially Private Decentralized Learning with Random Walks.....	9762
<i>Edwige Cyffers, Aurélien Bellet, Jalaj Upadhyay</i>	
Getting the Most Out of Your Tokenizer for Pre-Training and Domain Adaptation.....	9784
<i>Gautier Dagan, Gabriel Synnaeve, Baptiste Rozière</i>	
Fault Tolerant ML: Efficient Meta-Aggregation and Synchronous Training.....	9806
<i>Tehila Dahan, Kfir Y. Levy</i>	

POSITION: Beyond Personhood: Agency, Accountability, and the Limits of Anthropomorphic Ethical Analysis	9834
<i>Jessica Dai</i>	
Multi-View Clustering by Inter-Cluster Connectivity Guided Reward	9846
<i>Hao Dai, Yang Liu, Peng Su, Hecheng Cai, Shudong Huang, Jiancheng Lv</i>	
High-Order Contrastive Learning with Fine-Grained Comparative Levels for Sparse Ordinal Tensor Completion	9856
<i>Yu Dai, Junchen Shen, Zijie Zhai, Danlin Liu, Jingyang Chen, Yu Sun, Ping Li, Jie Zhang, Kai Zhang</i>	
Safe Reinforcement Learning Using Finite-Horizon Gradient-Based Estimation	9872
<i>Juntao Dai, Yaodong Yang, Qian Zheng, Gang Pan</i>	
Averaging N-Step Returns Reduces Variance in Reinforcement Learning	9904
<i>Brett Daley, Martha White, Marlos C. Machado</i>	

PART 13

Pruned Pivot: Correlation Clustering Algorithm for Dynamic, Parallel, and Local Computation Models	9931
<i>Mina Dalirrooyfard, Konstantin Makarychev, Slobodan Mitrovic</i>	
Physics and Lie Symmetry Informed Gaussian Processes	9953
<i>David Dalton, Dirk Husmeier, Hao Gao</i>	
Exploring the Enigma of Neural Dynamics Through a Scattering-Transform Mixer Landscape for Riemannian Manifold	9976
<i>Tingting Dan, Ziquan Wei, Won Hwa Kim, Guorong Wu</i>	
The Benefits of Reusing Batches for Gradient Descent in Two-Layer Networks: Breaking the Curse of Information and Leap Exponents	9991
<i>Yatin Dandi, Emanuele Troiani, Luca Arnaboldi, Luca Pesce, Lenka Zdeborova, Florent Krzakala</i>	
Neural Collapse for Cross-Entropy Class-Imbalanced Learning with Unconstrained ReLU Features Model	10017
<i>Hien Dang, Tho Tran, Tan Nguyen, Nhat Ho</i>	
Transformers Are SSMS: Generalized Models and Efficient Algorithms Through Structured State Space Duality	10041
<i>Tri Dao, Albert Gu</i>	
Boosting Offline Optimizers with Surrogate Sensitivity	10072
<i>Manh Cuong Dao, Phi Le Nguyen, Thao Nguyen Truong, Trong Nghia Hoang</i>	
Consistent Diffusion Meets Tweedie: Training Exact Ambient Diffusion Models with Noisy Data	10091
<i>Giannis Daras, Alexandros G. Dimakis, Constantinos Daskalakis</i>	
Larimar: Large Language Models with Episodic Memory Control	10109
<i>Payel Das, Subhajit Chaudhury, Elliot Nelson, Igor Melnyk, Sarathkrishna Swaminathan, Sihui Dai, Aurelie Lozano, Georgios Kollias, Vijil Chenthamarakshan, Jiri Navratil, Soham Dan, Pin-Yu Chen</i>	
Understanding the Training Speedup from Sampling with Approximate Losses	10127
<i>Rudrajit Das, Xi Chen, Bertram Ieong, Parikshit Bansal, Sujay Sanghavi</i>	
A Decoder-Only Foundation Model for Time-Series Forecasting	10148
<i>Abhimanyu Das, Weihao Kong, Rajat Sen, Yichen Zhou</i>	
Disparate Impact on Group Accuracy of Linearization for Private Inference	10168
<i>Saswat Das, Marco Romanelli, Ferdinando Fioretto</i>	

New Bounds on the Cohesion of Complete-Link and Other Linkage Methods for Agglomerative Clustering	10185
<i>Sanjoy Dasgupta, Eduardo Laber</i>	
Geometric Active Exploration in Markov Decision Processes: The Benefit of Abstraction	10206
<i>Riccardo De Santi, Federico Arangath Joseph, Noah Liniger, Mirco Mutti, Andreas Krause</i>	
Global Reinforcement Learning: Beyond Linear and Convex Rewards Via Submodular Semi-Gradient Methods	10235
<i>Riccardo De Santi, Manish Prajapat, Andreas Krause</i>	
Provably Better Explanations with Optimized Aggregation of Feature Attributions.....	10267
<i>Thomas Decker, Ananta R. Bhattarai, Jindong Gu, Volker Tresp, Florian Buettner</i>	
Learning Cognitive Maps from Transformer Representations for Efficient Planning in Partially Observed Environments	10287
<i>Antoine Dedieu, Wolfgang Lehrach, Guangyao Zhou, Dileep George, Miguel Lázaro-Gredilla</i>	
Asymptotically Optimal and Computationally Efficient Average Treatment Effect Estimation in A/B Testing.....	10317
<i>Vikas Deep, Achal Bassamboo, Sandeep Juneja</i>	
Predicting Lagrangian Multipliers for Mixed Integer Linear Programs	10368
<i>Francesco Demelas, Joseph Le Roux, Mathieu Lacroix, Axel Parmentier</i>	
Prediction-Powered Generalization of Causal Inferences	10385
<i>Ilker Demirel, Ahmed Alaa, Anthony Philippakis, David Sontag</i>	
An Unsupervised Approach for Periodic Source Detection in Time Series.....	10409
<i>Berken Utku Demirel, Christian Holz</i>	
Multi-Group Learning for Hierarchical Groups	10440
<i>Samuel Deng, Daniel Hsu</i>	
A3S: A General Active Clustering Method with Pairwise Constraints.....	10488
<i>Xun Deng, Junlong Liu, Han Zhong, Fuli Feng, Chen Shen, Xiangnan He, Jieping Ye, Zheng Wang</i>	
Variational Schrödinger Diffusion Models	10506
<i>Wei Deng, Weijian Luo, Yixin Tan, Marin Biloš, Yu Chen, Yuriy Nevmyvaka, Ricky T. Q. Chen</i>	
Collaborative Learning with Different Labeling Functions.....	10530
<i>Yuyang Deng, Mingda Qiao</i>	
Exploring the Low-Pass Filtering Behavior in Image Super-Resolution.....	10553
<i>Haoyu Deng, Zijing Xu, Yule Duan, Xiao Wu, Wenjie Shu, Liang-Jian Deng</i>	
Network Tight Community Detection	10574
<i>Jiayi Deng, Xiaodong Yang, Jun Yu, Jun S. Liu, Zhaiming Shen, Danyang Huang, Huimin Cheng</i>	
Going Beyond Compositions, DDPMs Can Produce Zero-Shot Interpolations.....	10597
<i>Justin Deschenaux, Igor Krawczuk, Grigorios G. Chrysos, Volkan Cevher</i>	
Multicalibration for Confidence Scoring in LLMs.....	10624
<i>Gianluca Detommaso, Martin Bertran, Riccardo Fogliato, Aaron Roth</i>	
Contextualized Policy Recovery: Modeling and Interpreting Medical Decisions with Adaptive Imitation Learning.....	10642
<i>Jannik Deuschel, Caleb N. Ellington, Yingtao Luo, Benjamin J. Lengerich, Pascal Friederich, Eric P. Xing</i>	
Stability and Multigroup Fairness in Ranking with Uncertain Predictions	10661
<i>Siddhartha Devic, Aleksandra Korolova, David Kempe, Vatsal Sharan</i>	

Locally Interdependent Multi-Agent MDP: Theoretical Framework for Decentralized Agents with Dynamic Dependencies	10687
<i>Alex Deweese, Guannan Qu</i>	
Bivariate Causal Discovery Using Bayesian Model Selection	10710
<i>Anish Dhir, Samuel Power, Mark Van Der Wilk</i>	
Trust Regions for Explanations Via Black-Box Probabilistic Certification.....	10736
<i>Amit Dhurandhar, Swagatam Halder, Dennis Wei, Karthikeyan Natesan Ramamurthy</i>	
Double Stochasticity Gazes Faster: Snap-Shot Decentralized Stochastic Gradient Tracking Methods	10765
<i>Hao Di, Haishan Ye, Xiangyu Chang, Guang Dai, Ivor W. Tsang</i>	

PART 14

Double Variance Reduction: A Smoothing Trick for Composite Optimization Problems Without First-Order Gradient	10792
<i>Hao Di, Haishan Ye, Yueling Zhang, Xiangyu Chang, Guang Dai, Ivor W. Tsang</i>	
Robust Sparse Estimation for Gaussians with Optimal Error Under Huber Contamination	10811
<i>Ilias Diakonikolas, Daniel Kane, Sushrut Karmalkar, Ankit Pensia, Thanasis Pittas</i>	
Fast Co-Training Under Weak Dependence Via Stream-Based Active Learning	10841
<i>Ilias Diakonikolas, Mingchen Ma, Lisheng Ren, Christos Tzamos</i>	
Convex and Bilevel Optimization for Neural-Symbolic Inference and Learning	10865
<i>Charles Dickens, Changyu Gao, Connor Pryor, Stephen Wright, Lise Getoor</i>	
Structure Your Data: Towards Semantic Graph Counterfactuals.....	10897
<i>Angeliki Dimitriou, Maria Lymperaioi, Giorgos Filandrianos, Konstantinos Thomas, Giorgos Stamou</i>	
Efficient Algorithms for Sum-Of-Minimum Optimization.....	10927
<i>Lisang Ding, Ziang Chen, Xinshang Wang, Wotao Yin</i>	
AMPA: Adaptive Mixed Precision Allocation for Low-Bit Integer Training.....	10960
<i>Li Ding, Wen Fei, Yuyang Huang, Shuangrui Ding, Wenrui Dai, Chenglin Li, Junni Zou, Hongkai Xiong</i>	
Understanding Forgetting in Continual Learning with Linear Regression	10978
<i>Meng Ding, Kaiyi Ji, Di Wang, Jinhui Xu</i>	
Recurrent Distance Filtering for Graph Representation Learning	11002
<i>Yuhui Ding, Antonio Orvieto, Bobby He, Thomas Hofmann</i>	
Robust Stable Spiking Neural Networks	11016
<i>Jianhao Ding, Zhiyu Pan, Yujia Liu, Zhaoifei Yu, Tiejun Huang</i>	
Fewer Truncations Improve Language Modeling.....	11030
<i>Hantian Ding, Zijian Wang, Giovanni Paolini, Varun Kumar, Anoop Deoras, Dan Roth, Stefano Soatto</i>	
Delving into Differentially Private Transformer	11049
<i>Youlong Ding, Xueyang Wu, Yining Meng, Yonggang Luo, Hao Wang, Weike Pan</i>	
Quality Diversity Through Human Feedback: Towards Open-Ended Diversity-Driven Optimization	11072
<i>Li Ding, Jenny Zhang, Jeff Clune, Lee Spector, Joel Lehman</i>	
LongRoPE: Extending LLM Context Window Beyond 2 Million Tokens	11091
<i>Yiran Ding, Li Lyna Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, Mao Yang</i>	
Learning-Rate-Free Stochastic Optimization Over Riemannian Manifolds.....	11105
<i>Daniel Dodd, Louis Sharrock, Christopher Nemeth</i>	

Consistent Adversarially Robust Linear Classification: Non-Parametric Setting.....	11149
<i>Elvis Dohmatob</i>	
A Tale of Tails: Model Collapse as a Change of Scaling Laws	11165
<i>Elvis Dohmatob, Yunzhen Feng, Pu Yang, Francois Charton, Julia Kempe</i>	
Precise Accuracy / Robustness Tradeoffs in Regression: Case of General Norms.....	11198
<i>Elvis Dohmatob, Meyer Scetbon</i>	
Spectral Preconditioning for Gradient Methodson Graded Non-Convex Functions	11227
<i>Nikita Doikov, Sebastian U. Stich, Martin Jaggi</i>	
Impact of Decentralized Learning on Player Utilities in Stackelberg Games.....	11253
<i>Kate Donahue, Nicole Immorlica, Meena Jagadeesan, Brendan Lucier, Aleksandrs Slivkins</i>	
Towards Generalization Beyond Pointwise Learning: A Unified Information-Theoretic Perspective	11311
<i>Yuxin Dong, Tieliang Gong, Hong Chen, Zhongjiang He, Mengxiang Li, Shuangyong Song, Chen Li</i>	
Pruner-Zero: Evolving Symbolic Pruning Metric from Scratch for Large Language Models.....	11346
<i>Peijie Dong, Lujun Li, Zhenheng Tang, Xiang Liu, Xinglin Pan, Qiang Wang, Xiaowen Chu</i>	
Position: Building Guardrails for Large Language Models Requires Systematic Design	11375
<i>Yi Dong, Ronghui Mu, Gaojie Jin, Yi Qi, Jinwei Hu, Xingyu Zhao, Jie Meng, Wenjie Ruan, Xiaowei Huang</i>	
Accelerating PDE Data Generation Via Differential Operator Action in Solution Space	11395
<i>Huanshuo Dong, Hong Wang, Haoyang Liu, Jian Luo, Jie Wang</i>	
TimeSiam: A Pre-Training Framework for Siamese Time-Series Modeling.....	11412
<i>Jiaxiang Dong, Haixu Wu, Yuxuan Wang, Yunzhong Qiu, Li Zhang, Jianmin Wang, Mingsheng Long</i>	
Get More with LESS: Synthesizing Recurrence with KV Cache Compression for Efficient LLM Inference.....	11437
<i>Harry Dong, Xinyu Yang, Zhenyu Zhang, Zhangyang Wang, Yuejie Chi, Beidi Chen</i>	
Privacy-Preserving Data Release Leveraging Optimal Transport and Particle Gradient Descent.....	11453
<i>Konstantin Donhauser, Javier Abad, Neha Hulkund, Fanny Yang</i>	
Spike Distance Function as a Learning Objective for Spike Prediction	11474
<i>Kevin Doran, Marvin Seifert, Carola A. M. Yovanovich, Tom Baden</i>	
Dynamic Byzantine-Robust Learning: Adapting to Switching Byzantine Workers	11501
<i>Ron Dorfman, Naseem Yehya, Kfir Y. Levy</i>	
Don't Label Twice: Quantity Beats Quality When Comparing Binary Classifiers on a Budget.....	11544
<i>Florian E. Dorner, Moritz Hardt</i>	
Learning from Students: Applying t-Distributions to Explore Accurate and Efficient Formats for LLMs.....	11573
<i>Jordan Dotzel, Yuzong Chen, Bahaa Kotb, Sushma Prasad, Gang Wu, Sheng Li, Mohamed S. Abdelfattah, Zhiru Zhang</i>	
Theory of Consistency Diffusion Models: Distribution Estimation Meets Fast Sampling.....	11592
<i>Zehao Dou, Minshuo Chen, Mengdi Wang, Zhuoran Yang</i>	

PART 15

On the Universality of Volume-Preserving and Coupling-Based Normalizing Flows	11613
<i>Felix Draxler, Stefan Wahl, Christoph Schnörr, Ullrich Köthe</i>	
WorkArena: How Capable Are Web Agents at Solving Common Knowledge Work Tasks?	11642
<i>Alexandre Drouin, Maxime Gasse, Massimo Caccia, Issam H. Laradji, Manuel Del Verme, Tom Marty, David Vazquez, Nicolas Chapados, Alexandre Lacoste</i>	

Principled Gradient-Based MCMC for Conditional Sampling of Text.....	11663
<i>Li Du, Afra Amini, Lucas Torroba Hennigen, Xinyan Velocity Yu, Holden Lee, Jason Eisner, Ryan Cotterell</i>	
SimPro: A Simple Probabilistic Framework Towards Realistic Long-Tailed Semi-Supervised Learning	11686
<i>Chaoqun Du, Yizeng Han, Gao Huang</i>	
GliDe with a CaPE: A Low-Hassle Method to Accelerate Speculative Decoding	11704
<i>Cunxiao Du, Jing Jiang, Xu Yuanchen, Jiawei Wu, Sicheng Yu, Yongqi Li, Shenggui Li, Kai Xu, Liqiang Nie, Zhaopeng Tu, Yang You</i>	
Position: Compositional Generative Modeling a Single Model is Not All You Need	11721
<i>Yilun Du, Leslie Kaelbling</i>	
Improving Factuality and Reasoning in Language Models Through Multiagent Debate.....	11733
<i>Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, Igor Mordatch</i>	
Learning Iterative Reasoning Through Energy Diffusion	11764
<i>Yilun Du, Jiayuan Mao, Joshua Tenenbaum</i>	
When and How Does In-Distribution Label Help Out-Of-Distribution Detection?.....	11777
<i>Xuefeng Du, Yiyu Sun, Yixuan Li</i>	
AnyTool: Self-Reflective, Hierarchical Agents for Large-Scale API Calls.....	11812
<i>Yu Du, Fangyun Wei, Hongyang Zhang</i>	
Exploration-Driven Policy Optimization in RLHF: Theoretical Insights on Efficient Data Utilization.....	11830
<i>Yihan Du, Anna Winnicki, Gal Dalal, Shie Mannor, R. Srikant</i>	
Bottleneck-Minimal Indexing for Generative Document Retrieval	11888
<i>Xin Du, Lixin Xiu, Kumiko Tanaka-Ishii</i>	
MuxServe: Flexible Spatial-Temporal Multiplexing for Multiple LLM Serving	11905
<i>Jiangfei Duan, Runyu Lu, Haojie Duanmu, Xiuhong Li, Xingcheng Zhang, Dahua Lin, Ion Stoica, Hao Zhang</i>	
MF-CLR: Multi-Frequency Contrastive Learning Representation for Time Series	11918
<i>Jufang Duan, Wei Zheng, Yangzhou Du, Wenfa Wu, Haipeng Jiang, Hongsheng Qi</i>	
DE-COP: Detecting Copyrighted Content in Language Models Training Data.....	11940
<i>André V. Duarte, Xuandong Zhao, Arlindo L. Oliveira, Lei Li</i>	
Unveiling the Potential of AI for Nanomaterial Morphology Prediction	11957
<i>Ivan Dubrovsky, Andrei Dmitrenko, Aleksei Dmitrenko, Nikita Serov, Vladimir Vinogradov</i>	
Consistent Submodular Maximization	11979
<i>Paul Dütting, Federico Fusco, Silvio Lattanzi, Ashkan Norouzi-Fard, Morteza Zadimoghaddam</i>	
Observable Propagation: Uncovering Feature Vectors in Transformers.....	11992
<i>Jacob Dunefsky, Arman Cohan</i>	
Sharpness-Aware Data Generation for Zero-Shot Quantization.....	12034
<i>Hoang Anh Dung, Cuong Pham, Trung Le, Jianfei Cai, Thanh-Toan Do</i>	
Making Old Things New: A Unified Algorithm for Differentially Private Clustering	12046
<i>Max Dupré La Tour, Monika Henzinger, David Saulpic</i>	
Generalization Bounds for Heavy-Tailed SDEs Through the Fractional Fokker-Planck Equation	12087
<i>Benjamin Dupuis, Umut Simsekli</i>	
Outlier-Robust Kalman Filtering Through Generalised Bayes.....	12138
<i>Gerardo Duran-Martin, Matias Altamirano, Alexander Y. Shestopaloff, Leandro Sánchez-Betancourt, Jeremias Knoblauch, Matt Jones, François-Xavier Briol, Kevin Murphy</i>	

Enabling Uncertainty Estimation in Iterative Neural Networks	12172
<i>Nikita Durasov, Doruk Oner, Jonathan Donier, Hieu Le, Pascal Fua</i>	
Barrier Algorithms for Constrained Non-Convex Optimization.....	12190
<i>Pavel Dvurechensky, Mathias Staudigl</i>	
Efficient Exploration for LLMs.....	12215
<i>Vikranth Dwaracherla, Seyed Mohammad Asghari, Botao Hao, Benjamin Van Roy</i>	
Equivariant Frames and the Impossibility of Continuous Canonicalization.....	12228
<i>Nadav Dym, Hannah Lawrence, Jonathan W Siegel</i>	
Position: Insights from Survey Methodology Can Improve Training Data	12268
<i>Stephanie Eckman, Barbara Plank, Frauke Kreuter</i>	
Extreme Compression of Large Language Models Via Additive Quantization.....	12284
<i>Vage Egiazarian, Andrei Panferov, Denis Kuznedelev, Elias Frantar, Artem Babenko, Dan Alistarh</i>	
Ai-Sampler: Adversarial Learning of Markov Kernels with Involutive Maps	12304
<i>Evgenii Egorov, Riccardo Valperga, Efstratios Gavves</i>	
Efficient Error Certification for Physics-Informed Neural Networks.....	12318
<i>Francisco Eiras, Adel Bibi, Rudy Bunel, Krishnamurthy Dj Dvijotham, Philip H. S. Torr, M. Pawan Kumar</i>	
Position: Near to Mid-Term Risks and Opportunities of Open-Source Generative AI.....	12348
<i>Francisco Eiras, Aleksandar Petrov, Bertie Vidgen, Christian Schroeder De Witt, Fabio Pizzati, Katherine Elkins, Supratik Mukhopadhyay, Adel Bibi, Botos Csaba, Fabro Steibel, Fazl Barez, Genevieve Smith, Gianluca Guadagni, Jon Chun, Jordi Cabot, Joseph Marvin Imperial, Juan A. Nolasco-Flores, Lori Landay, Matthew Jackson, Paul Röttger, Philip H. S. Torr, Trevor Darrell, Yong Suk Lee, Jakob Foerster</i>	
Scalable Pre-Training of Large Autoregressive Image Models	12371
<i>Alaaeldin El-Nouby, Michal Klein, Shuangfei Zhai, Miguel Angel Bautista, Vaishaal Shankar, Alexander Toshev, Joshua M Suskind, Armand Joulin</i>	
Adaptive Online Experimental Design for Causal Discovery	12385
<i>Muhammad Qasim Elahi, Lai Wei, Murat Kocaoglu, Mahsa Ghasemi</i>	
TSLANet: Rethinking Transformers for Time Series Representation Learning.....	12409
<i>Emadeldeen Eldele, Mohamed Ragab, Zhenghua Chen, Min Wu, Xiaoli Li</i>	
Neuro-Visualizer: A Novel Auto-Encoder-Based Loss Landscape Visualization Method with an Application in Knowledge-Guided Machine Learning	12429
<i>Mohannad Elhamod, Anuj Karpatne</i>	
Revisiting Scalable Hessian Diagonal Approximations for Applications in Reinforcement Learning.....	12448
<i>Mohamed Elsayed, Homayoon Farrahi, Felix Dangel, A. Rupam Mahmood</i>	

PART 16

Approximate Nearest Neighbor Search with Window Filters	12469
<i>Joshua Engels, Benjamin Landrum, Shangdi Yu, Laxman Dhulipala, Julian Shun</i>	
DsDm: Model-Aware Dataset Selection with Datamodels.....	12491
<i>Logan Engstrom, Axel Feldmann, Aleksander Madry</i>	
Compositional Curvature Bounds for Deep Neural Networks	12527
<i>Taha Entesari, Sina Sharifi, Mahyar Fazlyab</i>	
Disentangled 3D Scene Generation with Layout Learning	12547
<i>Dave Epstein, Ben Poole, Ben Mildenhall, Alexei A. Efros, Aleksander Holynski</i>	

PAC-Bayesian Error Bound, Via Rényi Divergence, for a Class of Linear Time-Invariant State-Space Models	12560
<i>Deividas Eringis, John Leth, Zheng-Hua Tan, Rafal Wisniewski, Mihály Petreczky</i>	
High-Dimensional Geometric Streaming for Nearly Low Rank Data.....	12588
<i>Hossein Esfandiari, Praneeth Kacham, Vahab Mirrokni, David P. Woodruff, Peilin Zhong</i>	
Scaling Rectified Flow Transformers for High-Resolution Image Synthesis.....	12606
<i>Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, Robin Rombach</i>	
Model Alignment as Prospect Theoretic Optimization.....	12634
<i>Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, Douwe Kiela</i>	
Fast Timing-Conditioned Latent Audio Diffusion.....	12652
<i>Zach Evans, Cj Carr, Josiah Taylor, Scott H. Hawley, Jordi Pons</i>	
Scaling Exponents Across Parameterizations and Optimizers.....	12666
<i>Katie Everett, Lechao Xiao, Mitchell Wortsman, Alexander A. Alemi, Roman Novak, Peter J. Liu, Izzeddin Gur, Jascha Sohl-Dickstein, Leslie Pack Kaelbling, Jaehoon Lee, Jeffrey Pennington</i>	
Out of the Ordinary: Spectrally Adapting Regression for Covariate Shift	12701
<i>Benjamin Eyre, Elliot Creager, David Madras, Vardan Papayan, Richard Zemel</i>	
Adapt and Diffuse: Sample-Adaptive Reconstruction Via Latent Diffusion Models	12723
<i>Zalan Fabian, Berk Tinaz, Mahdi Soltanolkotabi</i>	
DiracDiffusion: Denoising and Incremental Reconstruction with Assured Data-Consistency.....	12754
<i>Zalan Fabian, Berk Tinaz, Mahdi Soltanolkotabi</i>	
Is In-Context Learning in Large Language Models Bayesian? a Martingale Perspective.....	12784
<i>Fabian Falck, Ziyu Wang, Chris Holmes</i>	
On the Recoverability of Causal Relations from Temporally Aggregated I.I.D. Data.....	12806
<i>Shunxing Fan, Mingming Gong, Kun Zhang</i>	
On the Convergence of Projected Bures-Wasserstein Gradient Descent Under Euclidean Strong Convexity	12832
<i>Junyi Fan, Yuxuan Han, Zijian Liu, Jian-Feng Cai, Yang Wang, Zhengyuan Zhou</i>	
Locally Estimated Global Perturbations Are Better than Local Perturbations for Federated Sharpness-Aware Minimization.....	12858
<i>Ziqing Fan, Shengchao Hu, Jiangchao Yao, Gang Niu, Ya Zhang, Masashi Sugiyama, Yanfeng Wang</i>	
Revisit the Essence of Distilling Knowledge Through Calibration.....	12882
<i>Wen-Shu Fan, Su Lu, Xin-Chun Li, De-Chuan Zhan, Le Gan</i>	
DoGE: Domain Reweighting with Generalization Estimation.....	12895
<i>Simin Fan, Matteo Pagliardini, Martin Jaggi</i>	
Path-Guided Particle-Based Sampling	12916
<i>Mingzhou Fan, Ruida Zhou, Chao Tian, Xiaoning Qian</i>	
Bayesian Knowledge Distillation: A Bayesian Perspective of Distillation with Uncertainty Quantification.....	12935
<i>Luyang Fang, Yongkai Chen, Wenxuan Zhong, Ping Ma</i>	
Exploring Correlations of Self-Supervised Tasks for Graphs.....	12957
<i>Taoran Fang, Wei Zhou, Yifei Sun, Kaiqiao Han, Lvbin Ma, Yang Yang</i>	
INViT: A Generalizable Routing Problem Solver with Invariant Nested View Transformer	12973
<i>Han Fang, Zhihao Song, Paul Weng, Yutong Ban</i>	

BayOTIDE: Bayesian Online Multivariate Time Series Imputation with Functional Decomposition	12993
<i>Shikai Fang, Qingsong Wen, Yingtao Luo, Shandian Zhe, Liang Sun</i>	
StackSight: Unveiling WebAssembly Through Large Language Models and Neurosymbolic Chain-Of-Thought Decompilation	13010
<i>Weike Fang, Zhejian Zhou, Junzhou He, Weihang Wang</i>	
Accelerating Heterogeneous Federated Learning with Closed-Form Classifiers	13029
<i>Eros Fani, Raffaello Camoriano, Barbara Caputo, Marco Ciccone</i>	
Stop Regressing: Training Value Functions Via Classification for Scalable Deep RL	13049
<i>Jesse Farebrother, Jordi Orbay, Quan Vuong, Adrien Ali Taïga, Yevgen Chebotar, Ted Xiao, Alex Irpan, Sergey Levine, Pablo Samuel Castro, Aleksandra Faust, Aviral Kumar, Rishabh Agarwal</i>	
Position: Cracking the Code of Cascading Disparity Towards Marginalized Communities	13072
<i>Golnoosh Farnadi, Mohammad Havaei, Negar Rostamzadeh</i>	
From Geometry to Causality- Ricci Curvature and the Reliability of Causal Inference on Networks	13086
<i>Amirhossein Farzam, Allen Tannenbaum, Guillermo Sapiro</i>	
Video-Of-Thought: Step-By-Step Video Reasoning from Perception to Cognition.....	13109
<i>Hao Fei, Shengqiong Wu, Wei Ji, Hanwang Zhang, Meishan Zhang, Mong Li Lee, Wynne Hsu</i>	
Bayesian Exploration Networks	13126
<i>Mattie Fellows, Brandon Kaplowitz, Christian Schroeder De Witt, Shimon Whiteson</i>	
Auto-Linear Phenomenon in Subsurface Imaging.....	13153
<i>Yinan Feng, Yinpeng Chen, Peng Jin, Shihang Feng, Youzuo Lin</i>	
Resisting Stochastic Risks in Diffusion Planners with the Trajectory Aggregation Tree	13175
<i>Lang Feng, Pengjie Gu, Bo An, Gang Pan</i>	
Fast White-Box Adversarial Streaming Without a Random Oracle.....	13199
<i>Ying Feng, Aayush Jain, David P. Woodruff</i>	
DSD-DA: Distillation-Based Source Debiasing for Domain Adaptive Object Detection	13225
<i>Yongchao Feng, Shiwei Li, Yingjie Gao, Ziyue Huang, Yanan Zhang, Qingjie Liu, Yunhong Wang</i>	
Keypoint-Based Progressive Chain-Of-Thought Distillation for LLMs	13241
<i>Kaituo Feng, Changsheng Li, Xiaolu Zhang, Jun Zhou, Ye Yuan, Guoren Wang</i>	
UniCorn: A Unified Contrastive Learning Approach for Multi-View Molecular Representation Learning	13256
<i>Shikun Feng, Yuyan Ni, Minghao Li, Yanwen Huang, Zhi-Ming Ma, Wei-Ying Ma, Yanyan Lan</i>	
Prediction Accuracy of Learning in Games : Follow-The-Regularized-Leader Meets Heisenberg	13278
<i>Yi Feng, Georgios Piliouras, Xiao Wang</i>	

PART 17

Privacy Backdoors: Stealing Data with Corrupted Pretrained Models.....	13326
<i>Shanglun Feng, Florian Tramèr</i>	
Memory Efficient Neural Processes Via Constant Memory Attention Block.....	13365
<i>Leo Feng, Frederick Tung, Hossein Hajimirsadeghi, Yoshua Bengio, Mohamed Osama Ahmed</i>	
Improving Sample Efficiency of Model-Free Algorithms for Zero-Sum Markov Games.....	13387
<i>Songtao Feng, Ming Yin, Yu-Xiang Wang, Jing Yang, Yingbin Liang</i>	
AquaLoRA: Toward White-Box Protection for Customized Stable Diffusion Models Via Watermark LoRA	13423
<i>Weitao Feng, Wenbo Zhou, Jiyan He, Jie Zhang, Tianyi Wei, Guanlin Li, Tianwei Zhang, Weiming Zhang, Nenghai Yu</i>	

GenCO: Generating Diverse Designs with Combinatorial Constraints.....	13445
<i>Aaron M. Ferber, Arman Zharmagambetov, Taoan Huang, Bistra Dilkina, Yuandong Tian</i>	
Active Ranking and Matchmaking, with Perfect Matchings	13460
<i>Hafedh El Ferchichi, Matthieu Lerasle, Vianney Perchet</i>	
Promptbreeder: Self-Referential Self-Improvement Via Prompt Evolution.....	13481
<i>Chrisantha Fernando, Dylan Banarse, Henryk Michalewski, Simon Osindero, Tim Rocktäschel</i>	
Trained Random Forests Completely Reveal Your Dataset.....	13545
<i>Julien Ferry, Ricardo Fukasawa, Timothée Pascal, Thibaut Vidal</i>	
Reservoir Computing for Short High-Dimensional Time Series: An Application to SARS-CoV-2 Hospitalization Forecast	13570
<i>Thomas Féré, Dan Dutartre, Boris Hejblum, Romain Griffier, Vianney Jouhet, Rodolphe Thiébaud, Pierrick Legrand, Xavier Hinaut</i>	
Position: Relational Deep Learning - Graph Representation Learning on Relational Databases	13592
<i>Matthias Fey, Weihua Hu, Kexin Huang, Jan Eric Lenssen, Rishabh Ranjan, Joshua Robinson, Rex Ying, Jiaxuan You, Jure Leskovec</i>	
On Statistical Learning Theory for Distributional Inputs	13608
<i>Christian Fiedler, Pierre-François Massiani, Friedrich Solowjow, Sebastian Trimpe</i>	
Cooperative Graph Neural Networks	13633
<i>Ben Finkelshtein, Xingyue Huang, Michael Bronstein, Ismail Ilkan Ceylan</i>	
Critical Feature Learning in Deep Neural Networks	13660
<i>Kirsten Fischer, Javed Lindner, David Dahmen, Zohar Ringel, Michael Krämer, Moritz Helias</i>	
Neuroexplicit Diffusion Models for Inpainting of Optical Flow Fields	13691
<i>Tom Fischer, Pascal Peter, Joachim Weickert, Eddy Ilg</i>	
Inverse-Variance Weighting for Estimation of Heterogeneous Treatment Effects	13706
<i>Aaron Fisher</i>	
Stochastic Q-Learning for Large Discrete Action Spaces	13734
<i>Fares Fourati, Vaneet Aggarwal, Mohamed-Slim Alouini</i>	
Federated Combinatorial Multi-Agent Multi-Armed Bandits	13760
<i>Fares Fourati, Mohamed-Slim Alouini, Vaneet Aggarwal</i>	
Initial Guessing Bias: How Untrained Networks Favor Some Classes	13783
<i>Emanuele Francazi, Aurelien Lucchi, Marco Baity-Jesi</i>	
Explaining Probabilistic Models with Distributional Values.....	13840
<i>Luca Franceschi, Michele Donini, Cédric Archambeau, Matthias Seeger</i>	
Hyperbolic Active Learning for Semantic Segmentation Under Domain Shift.....	13864
<i>Luca Franco, Paolo Mandica, Konstantinos Kallidromitis, Devin Guillory, Yu-Teng Li, Trevor Darrell, Fabio Galasso</i>	
Weisfeiler–Leman at the Margin: When More Expressivity Matters	13885
<i>Billy J. Franks, Christopher Morris, Ameya Velingker, Floris Geerts</i>	
Unsupervised Zero-Shot Reinforcement Learning Via Functional Reward Encodings.....	13927
<i>Kevin Frans, Seohong Park, Pieter Abbeel, Sergey Levine</i>	
Fair Off-Policy Learning from Observational Data.....	13943
<i>Dennis Frauen, Valentyn Melnychuk, Stefan Feuerriegel</i>	
Trust the Model Where it Trusts Itself Model-Based Actor-Critic with Uncertainty-Aware Rollout Adaption	13973
<i>Bernd Frauenknecht, Artur Eisele, Devdutt Subhasish, Friedrich Solowjow, Sebastian Trimpe</i>	

Trustworthy Actionable Perturbations	14006
<i>Jesse Friedbaum, Sudarshan Adiga, Ravi Tandon</i>	
Interpretability Illusions in the Generalization of Simplified Models	14035
<i>Dan Friedman, Andrew Lampinen, Lucas Dixon, Danqi Chen, Asma Ghandeharioun</i>	
Break the Sequential Dependency of LLM Inference Using Lookahead Decoding.....	14060
<i>Yichao Fu, Peter Bailis, Ion Stoica, Hao Zhang</i>	
A Touch, Vision, and Language Dataset for Multimodal Alignment.....	14080
<i>Letian Fu, Gaurav Datta, Huang Huang, William Chung-Ho Panitch, Jaimyn Drake, Joseph Ortiz, Mustafa Mukadam, Mike Lambeta, Roberto Calandra, Ken Goldberg</i>	
Hyperbolic Geometric Latent Diffusion Model for Graph Generation	14102
<i>Xingcheng Fu, Yisen Gao, Yuecen Wei, Qingyun Sun, Hao Peng, Jianxin Li, Xianxian Li</i>	
Data Engineering for Scaling Language Models to 128K Context	14125
<i>Yao Fu, Rameswar Panda, Xinyao Niu, Xiang Yue, Hannaneh Hajishirzi, Yoon Kim, Hao Peng</i>	
Language-Guided Skill Learning with Temporal Variational Inference.....	14135
<i>Haotian Fu, Pratyusha Sharma, Elias Stengel-Eskin, George Konidaris, Nicolas Le Roux, Marc-Alexandre Côté , Xingdi Yuan</i>	

PART 18

PinNet: Pinpoint Instructive Information for Retrieval Augmented Code-To-Text Generation	14157
<i>Han Fu, Jian Tan, Pinhan Zhang, Feifei Li, Jianling Sun</i>	
Mean-Field Underdamped Langevin Dynamics and Its Spacetime Discretization	14175
<i>Qiang Fu, Ashia Wilson</i>	
Breaking Through the Learning Plateaus of In-Context Learning in Transformer.....	14207
<i>Jingwen Fu, Tao Yang, Yuwang Wang, Yan Lu, Nanning Zheng</i>	
Towards Theoretical Understandings of Self-Consuming Generative Models.....	14228
<i>Shi Fu, Sen Zhang, Yingjie Wang, Xinmei Tian, Dacheng Tao</i>	
FuRL: Visual-Language Models as Fuzzy Rewards for Reinforcement Learning	14256
<i>Yuwei Fu, Haichao Zhang, Di Wu, Wei Xu, Benoit Boulet</i>	
Uncertainty for Active Learning on Graphs	14275
<i>Dominik Fuchsgruber, Tom Wollschläger, Bertrand Charpentier, Antonio Oroz, Stephan Günnemann</i>	
KernelSHAP-IQ: Weighted Least Square Optimization for Shapley Interactions.....	14308
<i>Fabian Fumagalli, Maximilian Muschalik, Patrick Kolpaczki, Eyke Hüllermeier, Barbara Hammer</i>	
Beyond the ROC Curve: Classification Trees Using Cost-Optimal Curves, with Application to Imbalanced Datasets.....	14343
<i>Magzhan Gabidolla, Arman Zharmagambetov, Miguel Á. Carreira-Perpiñán</i>	
Positive Concave Deep Equilibrium Models.....	14365
<i>Mateusz Gabor, Tomasz Piotrowski, Renato L. G. Cavalcante</i>	
Let Go of Your Labels with Unsupervised Transfer	14382
<i>Artyom Gadetsky, Yulun Jiang, Maria Brbic</i>	
Bring Your Own (Non-Robust) Algorithm to Solve Robust MDPs by Estimating the Worst Kernel.....	14408
<i>Uri Gadot, Kaixin Wang, Navdeep Kumar, Kfir Y. Levy, Shie Mannor</i>	
Leverage Class-Specific Accuracy to Guide Data Generation for Improving Image Classification	14433
<i>Jay Gala, Pengtao Xie</i>	

Erasing the Bias: Fine-Tuning Foundation Models for Semi-Supervised Learning	14453
<i>Kai Gan, Tong Wei</i>	
Reflective Policy Optimization	14471
<i>Yaozhong Gan, Renye Yan, Zhe Wu, Junliang Xing</i>	
Confidence Aware Inverse Constrained Reinforcement Learning	14491
<i>Sriram Ganapathi Subramanian, Guiliang Liu, Mohammed Elmahgiubi, Kasra Rezaee, Pascal Poupart</i>	
Testing the Feasibility of Linear Programs with Bandit Feedback	14513
<i>Aditya Gangrade, Aditya Gopalan, Venkatesh Saligrama, Clayton Scott</i>	
A Doubly Recursive Stochastic Compositional Gradient Descent Method for Federated Multi-Level Compositional Optimization.....	14540
<i>Hongchang Gao</i>	
Energy-Based Backdoor Defense Without Task-Specific Samples and Model Retraining	14611
<i>Yudong Gao, Honglong Chen, Peng Sun, Zhe Li, Junjian Li, Huajie Shao</i>	
An Intrinsic Vector Heat Network	14638
<i>Alexander Gao, Maurice Chu, Mubbasir Kapadia, Ming C. Lin, Hsueh-Ti Derek Liu</i>	
Stochastic Weakly Convex Optimization Beyond Lipschitz Continuity	14651
<i>Wenzhi Gao, Qi Deng</i>	
A Graph is Worth K Words: Euclideanizing Graph Using Pure Transformer.....	14681
<i>Zhangyang Gao, Daize Dong, Cheng Tan, Jun Xia, Bozhen Hu, Stan Z. Li</i>	
Linear Alignment: A Closed-Form Solution for Aligning Human Preferences Without Tuning and Feedback.....	14702
<i>Songyang Gao, Qiming Ge, Wei Shen, Shihan Dou, Junjie Ye, Xiao Wang, Rui Zheng, Yicheng Zou, Zhi Chen, Hang Yan, Qi Zhang, Dahua Lin</i>	
Multi-Agent Reinforcement Learning Meets Leaf Sequencing in Radiotherapy	14723
<i>Riqiang Gao, Florin C. Ghesu, Simon Arberet, Shahab Basiri, Esa Kuusela, Martin Kraus, Dorin Comaniciu, Ali Kamen</i>	
DMTG: One-Shot Differentiable Multi-Task Grouping.....	14747
<i>Yuan Gao, Shuguo Jiang, Moran Li, Jin-Gang Yu, Gui-Song Xia</i>	
Private Heterogeneous Federated Learning Without a Trusted Server Revisited: Error-Optimal and Communication-Efficient Algorithms for Convex Losses.....	14763
<i>Changyu Gao, Andrew Lowy, Xingyu Zhou, Stephen J. Wright</i>	
Speech Self-Supervised Learning Using Diffusion Model Synthetic Data	14790
<i>Heting Gao, Kaizhi Qian, Junrui Ni, Chuang Gan, Mark Hasegawa-Johnson, Shiyu Chang, Yang Zhang</i>	
Rethinking Specificity in SBDD: Leveraging Delta Score and Energy-Guided Diffusion	14811
<i>Bowen Gao, Minsi Ren, Yuyan Ni, Yanwen Huang, Bo Qiang, Zhi-Ming Ma, Wei-Ying Ma, Yanyan Lan</i>	
Non-Convex Stochastic Composite Optimization with Polyak Momentum	14826
<i>Yuan Gao, Anton Rodomanov, Sebastian U. Stich</i>	
Adaptive-Gradient Policy Optimization: Enhancing Policy Learning in Non-Smooth Differentiable Simulations.....	14844
<i>Feng Gao, Liangzhi Shi, Shenao Zhang, Zhaoran Wang, Yi Wu</i>	
Decoupling Learning and Decision-Making: Breaking the $O(\sqrt{T})$ Barrier in Online Resource Allocation with First-Order Methods	14859
<i>Wenzhi Gao, Chunlin Sun, Chenyu Xue, Yinyu Ye</i>	
Parameter-Efficient Fine-Tuning with Discrete Fourier Transform.....	14884
<i>Ziqi Gao, Qichao Wang, Aochuan Chen, Zijing Liu, Bingzhe Wu, Liang Chen, Jia Li</i>	

Fast-Slow Test-Time Adaptation for Online Vision-And-Language Navigation.....	14902
<i>Junyu Gao, Xuan Yao, Changsheng Xu</i>	
Causal Customer Churn Analysis with Low-Rank Tensor Block Hazard Model.....	14920
<i>Chenyin Gao, Zhiming Zhang, Shu Yang</i>	
Prompt-Based Visual Alignment for Zero-Shot Policy Transfer	14954
<i>Haihan Gao, Rui Zhang, Qi Yi, Hantao Yao, Haochen Li, Jiaming Guo, Shaohui Peng, Yunkai Gao, Qicheng Wang, Xing Hu, Yuanbo Wen, Zihao Zhang, Zidong Du, Ling Li, Qi Guo, Yunji Chen</i>	
Distribution Alignment Optimization Through Neural Collapse for Long-Tailed Classification.....	14969
<i>Jintong Gao, He Zhao, Dandan Guo, Hongyuan Zha</i>	

PART 19

Projection-Free Online Convex Optimization with Time-Varying Constraints.....	14988
<i>Dan Garber, Ben Kretzu</i>	
DRED: Zero-Shot Transfer in Reinforcement Learning Via Data-Regularised Environment Design.....	15006
<i>Samuel Garcin, James Doran, Shangmin Guo, Christopher G. Lucas, Stefano V. Albrecht</i>	
LLark : A Multimodal Instruction-Following Language Model for Music	15037
<i>Josh Gardner, Simon Durand, Daniel Stoller, Rachel Bittner</i>	
Memorization Through the Lens of Curvature of Loss Function Around Samples.....	15083
<i>Isha Garg, Deepak Ravikumar, Kaushik Roy</i>	
Simplicity Bias Via Global Convergence of Sharpness Minimization	15102
<i>Khashayar Gatmiry, Zhiyuan Li, Sashank Reddi, Stefanie Jegelka</i>	
Can Looped Transformers Learn to Implement Multi-Step Gradient Descent for In-Context Learning?.....	15130
<i>Khashayar Gatmiry, Nikunj Saunshi, Sashank Reddi, Stefanie Jegelka, Sanjiv Kumar</i>	
Closing the Gap: Achieving Global Convergence (Last Iterate) of Actor-Critic Under Markovian Sampling with Neural Network Parametrization.....	15153
<i>Mudit Gaur, Amrit Singh Bedi, Di Wang, Vaneet Aggarwal</i>	
Variance-Reduced Zeroth-Order Methods for Fine-Tuning Language Models.....	15180
<i>Tanmay Gautam, Youngsuk Park, Hao Zhou, Parameswaran Raman, Wooseok Ha</i>	
Position: Categorical Deep Learning is an Algebraic Theory of All Architectures	15209
<i>Bruno Gavranovi., Paul Lessard, Andrew Dudzik, Tamara Von Glehn, Joao G. M. Araujo, Petar Veli. Kovi.</i>	
Masked Face Recognition with Generative-To-Discriminative Representations	15242
<i>Shiming Ge, Weijia Guo, Chenyu Li, Junzheng Zhang, Yong Li, Dan Zeng</i>	
Safe and Robust Subgame Exploitation in Imperfect Information Games	15255
<i>Zhenxing Ge, Zheng Xu, Tianyu Ding, Linjian Meng, Bo An, Wenbin Li, Yang Gao</i>	
No Double Descent in Principal Component Regression: A High-Dimensional Analysis	15271
<i>Daniel Gedon, Antônio H. Ribeiro, Thomas B. Schön</i>	
Don't Trust Your Eyes: On the (un)reliability of Feature Visualizations	15294
<i>Robert Geirhos, Roland S. Zimmermann, Blair Bilodeau, Wieland Brendel, Been Kim</i>	
Learning with 3D Rotations: A Hitchhiker's Guide to SO(3).....	15331
<i>A. René Geist, Jonas Frey, Mikel Zhubro, Anna Levina, Georg Martius</i>	
Graph-Triggered Rising Bandits.....	15351
<i>Gianmarco Genalti, Marco Mussi, Nicola Gatti, Marcello Restelli, Matteo Castiglioni, Alberto Maria Metelli</i>	

Improving Adversarial Energy-Based Model Via Diffusion Process	15381
<i>Cong Geng, Tian Han, Peng-Tao Jiang, Hao Zhang, Jinwei Chen, Søren Hauberg, Bo Li</i>	
Reinforcement Learning Within Tree Search for Fast Macro Placement	15402
<i>Zijie Geng, Jie Wang, Ziyang Liu, Siyuan Xu, Zhentao Tang, Mingxuan Yuan, Jianye Hao, Yongdong Zhang, Feng Wu</i>	
Adaptive Horizon Actor-Critic for Policy Learning in Contact-Rich Differentiable Simulation	15418
<i>Ignat Georgiev, Krishnan Srinivasan, Jie Xu, Eric Heiden, Animesh Garg</i>	
Emergent Equivariance in Deep Ensembles	15438
<i>Jan E. Gerken, Pan Kessel</i>	
Patchscopes: A Unifying Framework for Inspecting Hidden Representations of Language Models	15466
<i>Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, Mor Geva</i>	
Individualized Privacy Accounting Via Subsampling with Applications in Combinatorial Optimization	15491
<i>Badi Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, Adam Sealfon</i>	
State-Constrained Zero-Sum Differential Games with One-Sided Information	15512
<i>Mukesh Ghimire, Lei Zhang, Zhe Xu, Yi Ren</i>	
Understanding Finetuning for Factual Knowledge Extraction	15540
<i>Gaurav Ghosal, Tatsunori Hashimoto, Aditi Raghunathan</i>	
A Closer Look at the Limitations of Instruction Tuning	15559
<i>Sreyan Ghosh, Chandra Kiran Reddy Evuru, Sonal Kumar, Ramaneswaran S, Deepali Aneja, Zeyu Jin, Ramani Duraiswami, Dinesh Manocha</i>	
Agnostic Learning of Mixed Linear Regressions with EM and AM Algorithms	15590
<i>Avishek Ghosh, Arya Mazumdar</i>	
Optimal Eye Surgeon: Finding Image Priors Through Sparse Generators at Initialization	15610
<i>Avrajit Ghosh, Xitong Zhang, Kenneth Sun, Qing Qu, Saiprasad Ravishankar, Rongrong Wang</i>	
Self-Correcting Self-Consuming Loops for Generative Model Training	15646
<i>Nate Gillman, Michael Freeman, Daksh Aggarwal, Chia-Hong Hsu, Calvin Luo, Yonglong Tian, Chen Sun</i>	
Kernel-Based Evaluation of Conditional Biological Sequence Models	15678
<i>Pierre Glaser, Steffanie Paul, Alissa M. Hummer, Charlotte M. Deane, Debora S. Marks, Alan N. Amin</i>	
Better & Faster Large Language Models Via Multi-Token Prediction	15706
<i>Fabian Gloeckle, Badr Youbi Idrissi, Baptiste Rozière, David Lopez-Paz, Gabriel Synnaeve</i>	
All-In-One Simulation-Based Inference	15735
<i>Manuel Gloeckler, Michael Deistler, Christian Weilbach, Frank Wood, Jakob H. Macke</i>	
Position: Fundamental Limitations of LLM Censorship Necessitate New Approaches	15767
<i>David Glukhov, Iliia Shumailov, Yarin Gal, Nicolas Papernot, Vardan Papayan</i>	
Position: The No Free Lunch Theorem, Kolmogorov Complexity, and the Role of Inductive Biases in Machine Learning	15788
<i>Micah Goldblum, Marc Finzi, Keefer Rowan, Andrew Gordon Wilson</i>	
CasCast: Skillful High-Resolution Precipitation Nowcasting Via Cascaded Modelling	15809
<i>Junchao Gong, Lei Bai, Peng Ye, Wanghan Xu, Na Liu, Jianhua Dai, Xiaokang Yang, Wanli Ouyang</i>	
Does Label Smoothing Help Deep Partial Label Learning?	15823
<i>Xiuwen Gong, Nitin Bisht, Guandong Xu</i>	

PART 20

AST-T5: Structure-Aware Pretraining for Code Generation and Understanding	15839
<i>Linyuan Gong, Mostafa Elhoushi, Alvin Cheung</i>	
A Nearly Optimal Single Loop Algorithm for Stochastic Bilevel Optimization Under Unbounded Smoothness.....	15854
<i>Xiaochuan Gong, Jie Hao, Mingrui Liu</i>	
Evolution-Inspired Loss Functions for Protein Representation Learning	15893
<i>Chengyue Gong, Adam Klivans, James Madigan Loy, Tianlong Chen, Qiang Liu, Daniel Diaz</i>	
Evaluation of LLMs on Syntax-Aware Code Fill-In-The-Middle Tasks	15907
<i>Linyuan Gong, Sida Wang, Mostafa Elhoushi, Alvin Cheung</i>	
E ² GAN: Efficient Training of Efficient GANs for Image-To-Image Translation.....	15929
<i>Yifan Gong, Zheng Zhan, Qing Jin, Yanyu Li, Yerlan Idelbayev, Xian Liu, Andrey Zharkov, Kfir Aberman, Sergey Tulyakov, Yanzhi Wang, Jian Ren</i>	
High-Probability Convergence for Composite and Distributed Stochastic Minimization and Variational Inequalities with Heavy-Tailed Noise	15951
<i>Eduard Gorbunov, Abdurakhmon Sadiev, Marina Danilova, Samuel Horváth, Gauthier Gidel, Pavel Dvurechensky, Alexander Gasnikov, Peter Richtárik</i>	
Out-Of-Domain Generalization in Dynamical Systems Reconstruction	16071
<i>Niclas Göring, Florian Hess, Manuel Brenner, Zahra Monfared, Daniel Durstewitz</i>	
MOMENT: A Family of Open Time-Series Foundation Models.....	16115
<i>Mononito Goswami, Konrad Szafer, Arjun Choudhry, Yifu Cai, Shuo Li, Artur Dubrawski</i>	
Weighted Distance Nearest Neighbor Condensing.....	16153
<i>Lee-Ad Gottlieb, Timor Sharabi, Roi Weiss</i>	
Test-Time Degradation Adaptation for Open-Set Image Restoration.....	16167
<i>Yuanbiao Gou, Haiyu Zhao, Boyun Li, Xinyan Xiao, Xi Peng</i>	
Learning Universal Predictors	16178
<i>Jordi Grau-Moya, Tim Genewein, Marcus Hutter, Laurent Orseau, Grégoire Déletang, Elliot Catt, Anian Ruoss, Li Kevin Wenliang, Christopher Mattern, Matthew Aitchison, Joel Veness</i>	
Long Range Propagation on Continuous-Time Dynamic Graphs	16206
<i>Alessio Gravina, Giulio Lovisotto, Claudio Gallicchio, Davide Bacciu, Claas Grohnfeldt</i>	
The Expressive Power of Path-Based Graph Neural Networks.....	16226
<i>Caterina Graziani, Tamara Drucks, Fabian Jögl, Monica Bianchini, Franco Scarselli, Thomas Gärtner</i>	
Nonsmooth Implicit Differentiation: Deterministic and Stochastic Convergence Rates.....	16250
<i>Riccardo Grazi, Massimiliano Pontil, Saverio Salzo</i>	
Fine-Grained Classes and How to Find Them.....	16275
<i>Matej Grčić, Artyom Gadetsky, Maria Brbic</i>	
AI Control: Improving Safety Despite Intentional Subversion	16295
<i>Ryan Greenblatt, Buck Shlegeris, Kshitij Sachan, Fabien Roger</i>	
Stochastic Localization Via Iterative Posterior Sampling.....	16337
<i>Louis Grenioux, Maxence Noble, Marylou Gabrié, Alain Durmus</i>	
A Bayesian Approach to Online Planning	16377
<i>Nir Greshler, David Ben Eli, Carmel Rabinovitz, Gabi Guetta, Liran Gispan, Guy Zohar, Aviv Tamar</i>	
Scaling Down Deep Learning with MNIST-1D	16404
<i>Sam Greydanus, Dmitry Kobak</i>	

Quality-Diversity Actor-Critic: Learning High-Performing and Diverse Behaviors via Value and Successor Features Critics	16416
<i>Luca Grillotti, Maxence Faldor, Borja González León, Antoine Cully</i>	
A Bias-Variance-Covariance Decomposition of Kernel Scores for Generative Models.....	16460
<i>Sebastian G. Gruber, Florian Buettner</i>	
Overcoming Saturation in Density Ratio Estimation by Iterated Regularization	16502
<i>Lukas Gruber, Markus Holzleitner, Johannes Lehner, Sepp Hochreiter, Werner Zellinger</i>	
On the Calibration of Human Pose Estimation.....	16530
<i>Kerui Gu, Rongyu Chen, Xuanlong Yu, Angela Yao</i>	
EDISON: Enhanced Dictionary-Induced Tensorized Incomplete Multi-View Clustering with Gaussian Error Rank Minimization.....	16548
<i>Zhibin Gu, Zhendong Li, Songhe Feng</i>	
CRUXEval: A Benchmark for Code Reasoning, Understanding and Execution.....	16568
<i>Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, Sida I. Wang</i>	
Data-Free Distillation of Diffusion Models with Bootstrapping	16622
<i>Jiatao Gu, Chen Wang, Shuangfei Zhai, Yizhe Zhang, Lingjie Liu, Joshua M. Susskind</i>	
Agent Smith: A Single Image Can Jailbreak One Million Multimodal LLM Agents Exponentially Fast	16647
<i>Xiangming Gu, Xiaosen Zheng, Tianyu Pang, Chao Du, Qian Liu, Ye Wang, Jing Jiang, Min Lin</i>	

PART 21

Predictive Performance Comparison of Decision Policies Under Confounding	16673
<i>Luke Guerdan, Amanda Coston, Kenneth Holstein, Zhiwei Steven Wu</i>	
On the Diminishing Returns of Width for Continual Learning	16706
<i>Etash Guha, Vihan Lakshman</i>	
Vector Quantization Pretraining for EEG Time Series with Random Projection and Phase Alignment.....	16731
<i>Haokun Gui, Xiucheng Li, Xinyang Chen</i>	
No Wrong Turns: The Simple Geometry of Neural Networks Optimization Paths.....	16751
<i>Charles Guille-Escuret, Hiroki Naganuma, Kilian Fatras, Ioannis Mitliagkas</i>	
Automated Evaluation of Retrieval-Augmented Language Models with Task-Specific Exam Generation	16773
<i>Gauthier Guinet, Behrooz Omidvar-Tehrani, Anoop Deoras, Laurent Callot</i>	
SLAB: Efficient Transformers with Simplified Linear Attention and Progressive Re-Parameterized Batch Normalization.....	16802
<i>Jialong Guo, Xinghao Chen, Yehui Tang, Yunhe Wang</i>	
DS-Agent: Automated Data Science by Empowering Large Language Models with Case-Based Reasoning	16813
<i>Siyuan Guo, Cheng Deng, Ying Wen, Hechang Chen, Yi Chang, Jun Wang</i>	
Collaborative Heterogeneous Causal Inference Beyond Meta-Analysis.....	16849
<i>Tianyu Guo, Sai Praneeth Karimireddy, Michael I. Jordan</i>	
ACM-MILP: Adaptive Constraint Modification Via Grouping and Selection for Hardness-Preserving MILP Instance Generation.....	16869
<i>Ziao Guo, Yang Li, Chang Liu, Wenli Ouyang, Junchi Yan</i>	

On the Embedding Collapse When Scaling Up Recommendation Models.....	16891
<i>Xingzhuo Guo, Junwei Pan, Ximei Wang, Baixu Chen, Jie Jiang, Mingsheng Long</i>	
FedRC: Tackling Diverse Distribution Shifts Challenge in Federated Learning by Robust Clustering	16910
<i>Yongxin Guo, Xiaoying Tang, Tao Lin</i>	
Compressing Large Language Models by Joint Sparsification and Quantization	16945
<i>Jinyang Guo, Jianyu Wu, Zining Wang, Jiaheng Liu, Ge Yang, Yifu Ding, Ruihao Gong, Haotong Qin, Xianglong Liu</i>	
Automated Loss Function Search for Class-Imbalanced Node Classification	16958
<i>Xinyu Guo, Kai Wu, Xiaoyu Zhang, Jing Liu</i>	
COLD-Attack: Jailbreaking LLMs with Stealthiness and Controllability.....	16974
<i>Xingang Guo, Fangxu Yu, Huan Zhang, Lianhui Qin, Bin Hu</i>	
Temporal Logic Specification-Conditioned Decision Transformer for Offline Safe Reinforcement Learning	17003
<i>Zijian Guo, Weichao Zhou, Wenchao Li</i>	
Diffusion Posterior Sampling is Computationally Intractable.....	17020
<i>Shivam Gupta, Ajil Jalal, Aditya Parulekar, Eric Price, Zhiyang Xun</i>	
xT: Nested Tokenization for Larger Context in Large Images.....	17060
<i>Ritwik Gupta, Shufan Li, Tyler Zhu, Jitendra Malik, Trevor Darrell, Karttikeya Mangalam</i>	
GistScore: Learning Better Representations for In-Context Example Selection with Gist Bottlenecks	17072
<i>Shivanshu Gupta, Clemens Rosenbaum, Ethan R. Elenberg</i>	
Light and Optimal Schrödinger Bridge Matching	17100
<i>Nikita Gushchin, Sergei Kholkin, Evgeny Burnaev, Alexander Korotin</i>	
Robust Yet Efficient Conformal Prediction Sets	17123
<i>Soroush H. Zargarbashi, Mohammad Sadegh Akhondzadeh, Aleksandar Bojchevski</i>	
Optimally Improving Cooperative Learning in a Social Setting	17148
<i>Shahrzad Haddadan, Cheng Xin, Jie Gao</i>	
Vectorized Conditional Neural Fields: A Framework for Solving Time-Dependent Parametric Partial Differential Equations	17189
<i>Jan Hagnberger, Marimuthu Kalimuthu, Daniel Musekamp, Mathias Niepert</i>	
Isometric Representation Learning for Disentangled Latent Space of Diffusion Models	17224
<i>Jaehoon Hahn, Junho Lee, Sunghyun Kim, Joonseok Lee</i>	
Pursuing Overall Welfare in Federated Learning Through Sequential Decision Making.....	17246
<i>Seok-Ju Hahn, Gi-Soo Kim, Junghye Lee</i>	
Incorporating Probabilistic Domain Knowledge into Deep Multiple Instance Learning	17279
<i>Ghadi S. Al Hajj, Aliaksandr Hubin, Chakravarthi Kanduri, Milena Pavlovic, Knut Rand, Michael Widrich, Anne Solberg, Victor Greiff, Johan Pensar, Günter Klambauer, Geir Kjetil Sandve</i>	
Covert Malicious Finetuning: Challenges in Safeguarding LLM Adaptation	17298
<i>Danny Halawi, Alexander Wei, Eric Wallace, Tony Wang, Nika Haghtalab, Jacob Steinhardt</i>	
A Study of First-Order Methods with a Deterministic Relative-Error Gradient Oracle	17313
<i>Nadav Hallak, Kfir Y. Levy</i>	
Dr. Strategy: Model-Based Generalist Agents with Strategic Dreaming.....	17333
<i>Hany Hamed, Subin Kim, Dongyeong Kim, Jaesik Yoon, Sungjin Ahn</i>	
UGrid: An Efficient-And-Rigorous Neural Multigrid Solver for Linear PDEs.....	17354
<i>Xi Han, Fei Hou, Hong Qin</i>	

Model Assessment and Selection Under Temporal Distribution Shift.....	17374
<i>Elise Han, Chengpiao Huang, Kaizheng Wang</i>	
Riemannian Coordinate Descent Algorithms on Matrix Manifolds	17393
<i>Andi Han, Pratik Jawanpuria, Bamdev Mishra</i>	
Prototypical Transformer as Unified Motion Learners.....	17416
<i>Cheng Han, Yawen Lu, Guohao Sun, James C. Liang, Zhiwen Cao, Qifan Wang, Qiang Guan, Sohail A. Dianat, Raghuvver M. Rao, Tong Geng, Zhiqiang Tao, Dongfang Liu</i>	
SIN: Selective and Interpretable Normalization for Long-Term Time Series Forecasting	17437
<i>Lu Han, Han-Jia Ye, De-Chuan Zhan</i>	
Large Language Models Can Automatically Engineer Features for Few-Shot Tabular Learning	17454
<i>Sungwon Han, Jinsung Yoon, Sercan Ö. Arik, Tomas Pfister</i>	
Improving Group Robustness on Spurious Correlation Requires Preciser Group Inference.....	17480
<i>Yujin Han, Difan Zou</i>	
Binary Decomposition: A Problem Transformation Perspective for Open-Set Semi-Supervised Learning	17505
<i>Jun-Yi Hang, Min-Ling Zhang</i>	

PART 22

Spotting LLMs with Binoculars: Zero-Shot Detection of Machine-Generated Text.....	17519
<i>Abhimanyu Hans, Avi Schwarzschild, Valeriia Cherepanova, Hamid Kazemi, Aniruddha Saha, Micah Goldblum, Jonas Geiping, Tom Goldstein</i>	
Interpreting Equivariant Representations	17538
<i>Andreas Abildtrup Hansen, Anna Calissano, Aasa Feragen</i>	
FLORA: Low-Rank Adapters Are Secretly Gradient Compressors	17554
<i>Yongchang Hao, Yanshuai Cao, Lili Mou</i>	
Data-Efficient Large Vision Models Through Sequential Autoregression.....	17572
<i>Zhiwei Hao, Jianyuan Guo, Chengcheng Wang, Yehui Tang, Han Wu, Han Hu, Kai Han, Chang Xu</i>	
MGit: A Model Versioning and Management System	17597
<i>Wei Hao, Daniel Mendoza, Rafael Da Silva, Deepak Narayanan, Amar Phanishayee, Asaf Cidon, Junfeng Yang</i>	
DPOT: Auto-Regressive Denoising Operator Transformer for Large-Scale PDE Pre-Training.....	17616
<i>Zhongkai Hao, Chang Su, Songming Liu, Julius Berner, Chengyang Ying, Hang Su, Anima Anandkumar, Jian Song, Jun Zhu</i>	
Convergence Guarantees for the DeepWalk Embedding on Block Models	17636
<i>Christopher Harker, Aditya Bhaskara</i>	
Estimating the Permanent by Nesting Importance Sampling	17657
<i>Juha Harviainen, Mikko Koivisto</i>	
Position: C*-Algebraic Machine Learning — Moving in a New Direction.....	17667
<i>Yuka Hashimoto, Masahiro Ikeda, Hachem Kadri</i>	
Fine-Grained Local Sensitivity Analysis of Standard Dot-Product Self-Attention	17680
<i>Aaron Havens, Alexandre Araujo, Huan Zhang, Bin Hu</i>	
Wasserstein Wormhole: Scalable Optimal Transport Distance with Transformer	17697
<i>Doron Haviv, Russell Zhang Kunes, Thomas Dougherty, Cassandra Burdziak, Tal Nawy, Anna Gilbert, Dana Pe'Er</i>	

GLoRe: When, Where, and How to Improve LLM Reasoning Via Global and Local Refinements	17719
<i>Alex Havrilla, Sharath Rapparthi, Christoforos Nalmpantis, Jane Dwivedi-Yu, Maksym Zhuravynski, Eric Hambro, Roberta Raileanu</i>	
Understanding MLP-Mixer as a Wide and Sparse MLP	17734
<i>Tomohiro Hayase, Ryo Karakida</i>	
MAGNOLIA: Matching Algorithms Via GNNs for Online Value-To-Go Approximation.....	17759
<i>Alexandre Hayderi, Amin Saberi, Ellen Vitercik, Anders Wikum</i>	
LoRA+: Efficient Low Rank Adaptation of Large Models	17783
<i>Soufiane Hayou, Nikhil Ghosh, Bin Yu</i>	
From Words to Actions: Unveiling the Theoretical Underpinnings of LLM-Driven Autonomous Systems.....	17807
<i>Jianliang He, Siyu Chen, Fengzhuo Zhang, Zhuoran Yang</i>	
Deep Neural Room Acoustics Primitive.....	17842
<i>Yuhang He, Anoop Cherian, Gordon Wichern, Andrew Markham</i>	
Two Stones Hit One Bird: Bilevel Positional Encoding for Better Length Extrapolation.....	17858
<i>Zhenyu He, Guhao Feng, Shengjie Luo, Kai Yang, Liwei Wang, Jingjing Xu, Zhi Zhang, Hongxia Yang, Di He</i>	
Distributed Bilevel Optimization with Communication Compression	17877
<i>Yutong He, Jie Hu, Xinmeng Huang, Songtao Lu, Bin Wang, Kun Yuan</i>	
ReDiffuser: Reliable Decision-Making Using a Diffuserwith Confidence Estimation	17921
<i>Nantian He, Shaohui Li, Zhi Li, Yu Liu, You He</i>	
Domain-Wise Data Acquisition to Improve Performance Under Distribution Shift	17934
<i>Yue He, Dongbai Li, Pengfei Tian, Han Yu, Jiashuo Liu, Hao Zou, Peng Cui</i>	
Quantum Algorithm for Online Exp-Concave Optimization.....	17946
<i>Jianhao He, Chengchang Liu, Xutong Liu, Lvzhou Li, John C. S. Lui</i>	
Riemannian Accelerated Zeroth-Order Algorithm: Improved Robustness and Lower Query Complexity	17972
<i>Chang He, Zhaoye Pan, Xiao Wang, Bo Jiang</i>	
The Effect of Weight Precision on the Neuron Count in Deep ReLU Networks.....	18010
<i>Songhua He, Periklis A. Papakonstantinou</i>	
Ambiguity-Aware Abductive Learning	18019
<i>Hao-Yuan He, Hui Sun, Zheng Xie, Ming Li</i>	
Instruction Tuning for Secure Code Generation.....	18043
<i>Jingxuan He, Mark Vero, Gabriela Krasnopolska, Martin Vechev</i>	
Be Your Own Neighborhood: Detecting Adversarial Examples by the Neighborhood Relations Built on Self-Supervised Learning	18063
<i>Zhiyuan He, Yijun Yang, Pin-Yu Chen, Qiang Xu, Tsung-Yi Ho</i>	
SFC: Achieve Accurate Fast Convolution Under Low-Precision Arithmetic.....	18081
<i>Liulu He, Yufei Zhao, Rui Gao, Yuan Du, Li Du</i>	
Robust Multi-Task Learning with Excess Risks.....	18094
<i>Yifei He, Shiji Zhou, Guojun Zhang, Hyokun Yun, Yi Xu, Belinda Zeng, Trishul Chilimbi, Han Zhao</i>	
DynSyn: Dynamical Synergistic Representation for Efficient Learning and Control in Overactuated Embodied Systems	18115
<i>Kaibo He, Chenhui Zuo, Chengtian Ma, Yanan Sui</i>	
DeepPolar: Inventing Nonlinear Large-Kernel Polar Codes Via Deep Learning	18133
<i>S Ashwin Hebbar, Sravan Kumar Ankireddy, Hyeji Kim, Sewoong Oh, Pramod Viswanath</i>	

Differentially Private Sum-Product Networks.....	18155
<i>Xenia Heilmann, Mattia Cerrato, Ernst Althaus</i>	
Optimal Recurrent Network Topologies for Dynamical Systems Reconstruction.....	18174
<i>Christoph Jürgen Hemmer, Manuel Brenner, Florian Hess, Daniel Durstewitz</i>	
Learning Useful Representations of Recurrent Neural Network Weight Matrices.....	18205
<i>Vincent Herrmann, Francesco Faccio, Jürgen Schmidhuber</i>	
Position: Why We Must Rethink Empirical Research in Machine Learning.....	18228
<i>Moritz Herrmann, F. Julian D. Lange, Katharina Eggenberger, Giuseppe Casalicchio, Marcel Wever, Matthias Feurer, David Rügamer, Eyke Hüllermeier, Anne-Laure Boulesteix, Bernd Bischl</i>	
Randomized Confidence Bounds for Stochastic Partial Monitoring.....	18248
<i>Maxime Heuillet, Ola Ahmad, Audrey Durand</i>	
Listening to the Noise: Blind Denoising with Gibbs Diffusion.....	18284
<i>David Heurtel-Depeiges, Charles C. Margossian, Ruben Ohana, Bruno Régildo-Saint Blancard</i>	
Balanced Resonate-And-Fire Neurons	18305
<i>Saya Higuchi, Sebastian Kairat, Sander M. Bohté, Sebastian Otte</i>	
Understanding Diffusion Models by Feynman's Path Integral	18324
<i>Yuji Hirono, Akinori Tanaka, Kenji Fukushima</i>	
RVI-SAC: Average Reward Off-Policy Deep Reinforcement Learning.....	18352
<i>Yukinari Hisaki, Isao Ono</i>	

PART 23

Learning Surrogates for Offline Black-Box Optimization Via Gradient Matching.....	18374
<i>Minh Hoang, Azza Fadhel, Aryan Deshwal, Janardhan Rao Doppa, Trong Nghia Hoang</i>	
Estimating Unknown Population Sizes Using the Hypergeometric Distribution	18394
<i>Liam Hodgson, Danilo Bzdok</i>	
Eureka-Moments in Transformers:Multi-Step Tasks Reveal Softmax Induced Optimization Problems.....	18409
<i>David T. Hoffmann, Simon Schrod, Jelena Bratulic, Nadine Behrmann, Volker Fischer, Thomas Brox</i>	
Position: Is Machine Learning Good Or Bad for the Natural Sciences?	18439
<i>David W. Hogg, Soledad Villar</i>	
Sparse Dimensionality Reduction Revisited	18454
<i>Mikael Møller Høgsgaard, Lior Kamma, Kasper Green Larsen, Jelani Nelson, Chris Schwiengelshohn</i>	
Two Tales of Single-Phase Contrastive Hebbian Learning.....	18470
<i>Rasmus Høier, Christopher Zach</i>	
Verifying Message-Passing Neural Networks Via Topology-Based Bounds Tightening	18489
<i>Christopher Hojny, Shiqiang Zhang, Juan S. Campos, Ruth Misener</i>	
Φ_{Flow} : Differentiable Simulations for PyTorch, TensorFlow and Jax.....	18515
<i>Philipp Holl, Nils Thuerey</i>	
Criterion Collapse and Loss Distribution Control	18547
<i>Matthew J. Holland</i>	
Removing Spurious Concepts from Neural Network Representations Via Joint Subspace Estimation	18568
<i>Floris Holstege, Bram Wouters, Noud Van Giersbergen, Cees Diks</i>	

Decoding Compressed Trust: Scrutinizing the Trustworthiness of Efficient LLMs Under Compression.....	18611
<i>Junyuan Hong, Jinhao Duan, Chenhui Zhang, Zhangheng Li, Chulin Xie, Kelsey Lieberman, James Diffenderfer, Brian Bartoldson, Ajay Jaiswal, Kaidi Xu, Bhavya Kailkhura, Dan Hendrycks, Dawn Song, Zhangyang Wang, Bo Li</i>	
Enhancing Sufficient Dimension Reduction Via Hellinger Correlation	18634
<i>Seungbeom Hong, Ilmun Kim, Jun Song</i>	
Diversified Batch Selection for Training Acceleration.....	18648
<i>Feng Hong, Yueming Lyu, Jiangchao Yao, Ya Zhang, Ivor W. Tsang, Yanfeng Wang</i>	
Model-Based Reinforcement Learning for Confounded POMDPs	18668
<i>Mao Hong, Zhengling Qi, Yanxun Xu</i>	
A Primal-Dual Algorithm for Offline Constrained Reinforcement Learning with Linear MDPs.....	18711
<i>Kihyuk Hong, Ambuj Tewari</i>	
Do Large Code Models Understand Programming Concepts? Counterfactual Analysis for Code Predicates.....	18738
<i>Ashish Hooda, Mihai Christodorescu, Miltiadis Allamanis, Aaron Wilson, Kassem Fawaz, Somesh Jha</i>	
Weisfeiler Lemma for Euclidean Equivariant Machine Learning	18749
<i>Snir Hordan, Tal Amir, Nadav Dym</i>	
Graph Neural PDE Solvers with Conservation and Similarity-Equivariance.....	18785
<i>Masanobu Horie, Naoto Mitsume</i>	
Harmony in Diversity: Merging Neural Networks with Canonical Correlation Analysis	18815
<i>Stefan Horoi, Albert Manuel Orozco Camacho, Eugene Belilovsky, Guy Wolf</i>	
Classification Under Strategic Self-Selection.....	18833
<i>Guy Horowitz, Yonatan Sommer, Moran Koren, Nir Rosenfeld</i>	
Maestro: Uncovering Low-Rank Structures Via Trainable Decomposition	18859
<i>Samuel Horváth, Stefanos Laskaridis, Shashank Rajput, Hongyi Wang</i>	
Recovering the Pre-Fine-Tuning Weights of Generative Models	18882
<i>Eliahu Horwitz, Jonathan Kahana, Yedid Hoshen</i>	
Equilibrium of Data Markets with Externality	18905
<i>Safwan Hossain, Yiling Chen</i>	
A Persuasive Approach to Combating Misinformation	18926
<i>Safwan Hossain, Andjela Mladenovic, Yiling Chen, Gauthier Gidel</i>	
Multi-Sender Persuasion: A Computational Perspective	18944
<i>Safwan Hossain, Tonghan Wang, Tao Lin, Yiling Chen, David C. Parkes, Haifeng Xu</i>	
Efficient Mixture Learning in Black-Box Variational Inference	18972
<i>Alexandra Hotti, Oskar Kviman, Ricky Molén, Victor Elvira, Jens Lagergren</i>	
IBD-PSC: Input-Level Backdoor Detection Via Parameter-Oriented Scaling Consistency	18992
<i>Linshan Hou, Ruili Feng, Zhongyun Hua, Wei Luo, Leo Yu Zhang, Yiming Li</i>	
Decomposing Uncertainty for Large Language Models Through Input Clarification Ensembling.....	19023
<i>Bairu Hou, Yujian Liu, Kaizhi Qian, Jacob Andreas, Shiyu Chang, Yang Zhang</i>	
PrE-Text: Training Language Models on Private Federated Data in the Age of LLMs.....	19043
<i>Charlie Hou, Akshat Shrivastava, Hongyuan Zhan, Rylan Conway, Trang Le, Adithya Sagar, Giulia Fanti, Daniel Lazar</i>	
Loss Shaping Constraints for Long-Term Time Series Forecasting	19062
<i>Ignacio Hounie, Javier Porras-Valenzuela, Alejandro Ribeiro</i>	

Careful with that Scalpel: Improving Gradient Surgery with an EMA	19085
<i>Yu-Guan Hsieh, James Thornton, Eugene Ndiaye, Michal Klein, Marco Cuturi, Pierre Ablin</i>	
Tripod: Three Complementary Inductive Biases for Disentangled Representation Learning	19101
<i>Kyle Hsu, Jubayer Ibn Hamid, Kaylee Burns, Chelsea Finn, Jiajun Wu</i>	
Outlier-Efficient Hopfield Layers for Large Transformer-Based Models	19123
<i>Jerry Yao-Chieh Hu, Pei-Hsuan Chang, Haozheng Luo, Hong-Yu Chen, Weijian Li, Wei-Po Wang, Han Liu</i>	
Task-Aware Orthogonal Sparse Network for Exploring Shared Knowledge in Continual Learning	19153
<i>Yusong Hu, De Cheng, Dingwen Zhang, Nannan Wang, Tongliang Liu, Xinbo Gao</i>	
Q-Value Regularized Transformer for Offline Reinforcement Learning	19165
<i>Shengchao Hu, Ziqing Fan, Chaoqin Huang, Li Shen, Ya Zhang, Yanfeng Wang, Dacheng Tao</i>	
HarmoDT: Harmony Multi-Task Decision Transformer for Offline Reinforcement Learning	19182
<i>Shengchao Hu, Ziqing Fan, Li Shen, Ya Zhang, Yanfeng Wang, Dacheng Tao</i>	

PART 24

An Information Theoretic Approach to Interaction-Grounded Learning.....	19198
<i>Xiaoyan Hu, Farzan Farnia, Ho-Fung Leung</i>	
Accelerated Speculative Sampling Based on Tree Monte Carlo	19216
<i>Zhengmian Hu, Heng Huang</i>	
SceneCraft: An LLM Agent for Synthesizing 3D Scenes as Blender Code	19252
<i>Ziniu Hu, Ahmet Iscen, Aashi Jain, Thomas Kipf, Yisong Yue, David A Ross, Cordelia Schmid, Alireza Fathi</i>	
InfoNet: Neural Estimation of Mutual Information Without Test-Time Optimization.....	19283
<i>Zhengyang Hu, Song Kang, Qunsong Zeng, Kaibin Huang, Yanchao Yang</i>	
Pseudo-Calibration: Improving Predictive Uncertainty Estimation in Unsupervised Domain Adaptation	19304
<i>Dapeng Hu, Jian Liang, Xinchao Wang, Chuan-Sheng Foo</i>	
On Computational Limits of Modern Hopfield Models: A Fine-Grained Complexity Analysis	19327
<i>Jerry Yao-Chieh Hu, Thomas Lin, Zhao Song, Han Liu</i>	
Improving Interpretation Faithfulness for Vision Transformers	19344
<i>Lijie Hu, Yixin Liu, Ninghao Liu, Mengdi Huai, Lichao Sun, Di Wang</i>	
Multigroup Robustness.....	19371
<i>Lunjia Hu, Charlotte Peale, Judy Hanwen Shen</i>	
Provable Privacy with Non-Private Pre-Processing	19402
<i>Yaxi Hu, Amartya Sanyal, Bernhard Schölkopf</i>	
Case-Based Or Rule-Based: How Do Transformers Do the Math?.....	19438
<i>Yi Hu, Xiaojuan Tang, Haotong Yang, Muhan Zhang</i>	
Sparse Model Inversion: Efficient Inversion of Vision Transformers for Data-Free Applications.....	19475
<i>Zixuan Hu, Yongxian Wei, Li Shen, Zhenyi Wang, Lei Li, Chun Yuan, Dacheng Tao</i>	
Bayesian Design Principles for Offline-To-Online Reinforcement Learning.....	19491
<i>Hao Hu, Yiqin Yang, Jianing Ye, Chengjie Wu, Ziqing Mai, Yujing Hu, Tangjie Lv, Changjie Fan, Qianchuan Zhao, Chongjie Zhang</i>	
High-Performance Temporal Reversible Spiking Neural Networks with O(L) Training Memory and O(1) Inference Cost	19516
<i>Jiakui Hu, Yao Man, Xuerui Qiu, Yuhong Chou, Yuxuan Cai, Ning Qiao, Yonghong Tian, Bo Xu, Guoqi Li</i>	

Accelerating Transformer Pre-Training with 2:4 Sparsity.....	19531
<i>Yuezhou Hu, Kang Zhao, Weiyu Huang, Jianfei Chen, Jun Zhu</i>	
InfiAgent-DABench: Evaluating Agents on Data Analysis Tasks.....	19544
<i>Xueyu Hu, Ziyu Zhao, Shuang Wei, Ziwei Chai, Qianli Ma, Guoyin Wang, Xuwu Wang, Jing Su, Jingjing Xu, Ming Zhu, Yao Cheng, Jianbo Yuan, Jiwei Li, Kun Kuang, Yang Yang, Hongxia Yang, Fei Wu</i>	
ReconBoost: Boosting Can Achieve Modality Reconciliation.....	19573
<i>Cong Hua, Qianqian Xu, Shilong Bao, Zhiyong Yang, Qingming Huang</i>	
Optimal Hessian/Jacobian-Free Nonconvex-PL Bilevel Optimization.....	19598
<i>Feihu Huang</i>	
NeuralIndicator: Implicit Surface Reconstruction from Neural Indicator Priors.....	19622
<i>Shi-Sheng Huang, Guo Chen, Li Heng Chen, Hua Huang</i>	
Auctionformer: A Unified Deep Learning Algorithm for Solving Equilibrium Strategies in Auction Games.....	19635
<i>Kexin Huang, Ziqian Chen, Xue Wang, Chongming Gao, Jinyang Gao, Bolin Ding, Xiang Wang</i>	
In-Context Convergence of Transformers.....	19660
<i>Yu Huang, Yuan Cheng, Yingbin Liang</i>	
Near-Linear Time Approximation Algorithms for K-Means with Outliers.....	19723
<i>Junyu Huang, Qilong Feng, Ziyun Huang, Jinhui Xu, Jianxin Wang</i>	
Contrastive Predict-And-Search for Mixed Integer Linear Programs.....	19757
<i>Taoan Huang, Aaron Ferber, Arman Zharmagambetov, Yuandong Tian, Bistra Dilikina</i>	
Symbolic Music Generation with Non-Differentiable Rule Guided Diffusion.....	19772
<i>Yujia Huang, Adishree Ghatare, Yanzhe Liu, Ziniu Hu, Qinsheng Zhang, Chandramouli S Sastry, Siddharth Gururani, Sageev Oore, Yisong Yue</i>	
DiffDA: A Diffusion Model for Weather-Scale Data Assimilation.....	19798
<i>Langwen Huang, Lukas Gianinazzi, Yuejiang Yu, Peter D. Dueben, Torsten Hoefler</i>	
Model-Based RL for Mean-Field Games is Not Statistically Harder than Single-Agent RL.....	19816
<i>Jiawei Huang, Niao He, Andreas Krause</i>	
In-Context Decision Transformer: Reinforcement Learning Via Hierarchical Chain-Of-Thought.....	19871
<i>Sili Huang, Jifeng Hu, Hechang Chen, Lichao Sun, Bo Yang</i>	
InstructSpeech: Following Speech Editing Instructions Via Large Language Models.....	19886
<i>Rongjie Huang, Ruofan Hu, Yongqi Wang, Zehan Wang, Xize Cheng, Ziyue Jiang, Zhenhui Ye, Dongchao Yang, Luping Liu, Peng Gao, Zhou Zhao</i>	
Bayesian Power Steering: An Effective Approach for Domain Adaptation of Diffusion Models.....	19904
<i>Ding Huang, Ting Li, Jian Huang</i>	
AttNS: Attention-Inspired Numerical Solving for Limited Data Scenarios.....	19929
<i>Zhongzhan Huang, Mingfu Liang, Shanshan Zhong, Liang Lin</i>	
CLIF: Complementary Leaky Integrate-And-Fire Neuron for Spiking Neural Networks.....	19949
<i>Yulong Huang, Xiaopeng Lin, Hongwei Ren, Haotian Fu, Yue Zhou, Zunchang Liu, Biao Pan, Bojun Cheng</i>	
Machine Vision Therapy: Multimodal Large Language Models Can Enhance Visual Robustness Via Denoising In-Context Learning.....	19973
<i>Zhuo Huang, Chang Liu, Yinpeng Dong, Hang Su, Shibao Zheng, Tongliang Liu</i>	
Efficient Adaptation in Mixed-Motive Environments Via Hierarchical Opponent Modeling and Planning.....	20004
<i>Yizhe Huang, Anji Liu, Fanqi Kong, Yaodong Yang, Song-Chun Zhu, Xue Feng</i>	

BiLLM: Pushing the Limit of Post-Training Quantization for LLMs	20023
<i>Wei Huang, Yangdong Liu, Haotong Qin, Ying Li, Shiming Zhang, Xianglong Liu, Michele Magno, Xiaojuan Qi</i>	

PART 25

An Empirical Examination of Balancing Strategy for Counterfactual Estimation on Time Series	20043
<i>Qiang Huang, Chuizheng Meng, Defu Cao, Biwei Huang, Yi Chang, Yan Liu</i>	
MFTN: A Multi-Scale Feature Transfer Network Based on IMatchFormer for Hyperspectral Image Super-Resolution	20063
<i>Shuying Huang, Mingyang Ren, Yong Yang, Xiaozheng Wang, Yingzhi Wei</i>	
On Which Nodes Does GCN Fail? Enhancing GCN from the Node Perspective	20073
<i>Jincheng Huang, Jialie Shen, Xiaoshuang Shi, Xiaofeng Zhu</i>	
Self-Driven Entropy Aggregation for Byzantine-Robust Heterogeneous Federated Learning	20096
<i>Wenke Huang, Zekun Shi, Mang Ye, He Li, Bo Du</i>	
Overcoming Data and Model Heterogeneities in Decentralized Federated Learning Via Synthetic Anchors	20111
<i>Chun-Yin Huang, Kartik Srinivas, Xin Zhang, Xiaoxiao Li</i>	
Quasi-Monte Carlo Features for Kernel Approximation	20134
<i>Zhen Huang, Jiajin Sun, Yian Huang</i>	
Position: TrustLLM: Trustworthiness in Large Language Models	20166
<i>Yue Huang, Lichao Sun, Haoran Wang, Siyuan Wu, Qihui Zhang, Yuan Li, Chujie Gao, Yixin Huang, Wenhan Lyu, Yixuan Zhang, Xiner Li, Hanchi Sun, Zhengliang Liu, Yixin Liu, Yijue Wang, Zhikun Zhang, Bertie Vidgen, Bhavya Kailkhura, Caiming Xiong, Chaowei Xiao, Chunyuan Li, Eric Xing, Furong Huang, Hao Liu, Heng Ji, Hongyi Wang, Huan Zhang, Huaxiu Yao, Manolis Kellis, Marinka Zitnik, Meng Jiang, Mohit Bansal, James Zou, Jian Pei, Jian Liu, Jianfeng Gao, Jiawei Han, Jieyu Zhao, Jiliang Tang, Jindong Wang, Joaquin Vanschoren, John C Mitchell, Kai Shu, Kaidi Xu, Kai-Wei Chang, Lifang He, Lifu Huang, Michael Backes, Neil Zhenqiang Gong, Philip S. Yu, Pin-Yu Chen, Quanquan Gu, Ran Xu, Rex Ying, Shuiwang Ji, Suman Jana, Tianlong Chen, Tianming Liu, Tianyi Zhou, William Wang, Xiang Li, Xiangliang Zhang, Xiao Wang, Xing Xie, Xun Chen, Xuyu Wang, Yan Liu, Yanfang Ye, Yinzhi Cao, Yong Chen, Yue Zhao</i>	
MLAgentBench: Evaluating Language Agents on Machine Learning Experimentation	20271
<i>Qian Huang, Jian Vora, Percy Liang, Jure Leskovec</i>	
How Universal Polynomial Bases Enhance Spectral Graph Neural Networks: Heterophily, Over-Smoothing, and Over-Squashing	20310
<i>Keke Huang, Yu Guang Wang, Ming Li, Pietro Liò</i>	
Conformal Prediction for Deep Classifier Via Label Ranking	20331
<i>Jianguo Huang, Huajun Xi, Linjun Zhang, Huaxiu Yao, Yue Qiu, Hongxin Wei</i>	
Interaction-Based Retrieval-Augmented Diffusion Models for Protein-Specific 3D Molecule Generation	20348
<i>Zhilin Huang, Ling Yang, Xiangxin Zhou, Chujun Qin, Yijie Yu, Xiawu Zheng, Zikun Zhou, Wentao Zhang, Yu Wang, Wenming Yang</i>	
Enhancing Size Generalization in Graph Neural Networks Through Disentangled Representation Learning	20365
<i>Zheng Huang, Qihui Yang, Dawei Zhou, Yujun Yan</i>	
Triadic-OCD: Asynchronous Online Change Detection with Provable Robustness, Optimality, and Convergence	20382
<i>Yancheng Huang, Kai Yang, Zelin Zhu, Leian Chen</i>	

An Embodied Generalist Agent in 3D World	20413
<i>Jiangyong Huang, Silong Yong, Xiaojian Ma, Xiongkun Linghu, Puhao Li, Yan Wang, Qing Li, Song-Chun Zhu, Baoxiong Jia, Siyuan Huang</i>	
InterLUDE: Interactions Between Labeled and Unlabeled Data to Enhance Semi-Supervised Learning	20452
<i>Zhe Huang, Xiaowei Yu, Dajiang Zhu, Michael C. Hughes</i>	
Re-Dock: Towards Flexible and Realistic Molecular Docking with Diffusion Bridge	20474
<i>Yufei Huang, Odin Zhang, Lirong Wu, Cheng Tan, Haitao Lin, Zhangyang Gao, Siyuan Li, Stan. Z. Li</i>	
Faster Adaptive Decentralized Learning Algorithms.....	20490
<i>Feihu Huang, Jianyu Zhao</i>	
Adversarially Robust Deep Multi-View Clustering: A Novel Attack and Defense Framework	20526
<i>Haonan Huang, Guoxu Zhou, Yangang Zheng, Yuning Qiu, Andong Wang, Qibin Zhao</i>	
Faster Sampling Via Stochastic Gradient Proximal Sampler.....	20559
<i>Xunpeng Huang, Difan Zou, Hanze Dong, Yi-An Ma, Tong Zhang</i>	
Position: Open-Endedness is Essential for Artificial Superhuman Intelligence	20597
<i>Edward Hughes, Michael Dennis, Jack Parker-Holder, Feryal Behbahani, Aditi Mavalankar, Yuge Shi, Tom Schaul, Tim Rocktäschel</i>	
Position: The Platonic Representation Hypothesis	20617
<i>Minyoung Huh, Brian Cheung, Tongzhou Wang, Phillip Isola</i>	
Nash Incentive-Compatible Online Mechanism Learning Via Weakly Differentially Private Online Learning	20643
<i>Joon Suk Huh, Kirthevasan Kandasamy</i>	
Make-A-Shape: A Ten-Million-Scale 3D Shape Model	20660
<i>Ka-Hei Hui, Aditya Sanghi, Arianna Rampini, Kamal Rahimi Malekshan, Zhengzhe Liu, Hooman Shayani, Chi-Wing Fu</i>	
Residual Quantization with Implicit Neural Codebooks	20682
<i>Iris A. M. Huijben, Matthijs Douze, Matthew Muckley, Ruud J. G. Van Sloun, Jakob Verbeek</i>	
Theoretical Guarantees for Variational Inference with Fixed-Variance Mixture of Gaussians.....	20700
<i>Tom Huix, Anna Korba, Alain Durmus, Eric Moulines</i>	
Deep Networks Always Grok and Here is Why	20722
<i>Ahmed Imtiaz Humayun, Randall Balestriero, Richard Baraniuk</i>	
Token-Specific Watermarking with Enhanced Detectability and Semantic Coherence for Large Language Models	20746
<i>Mingjia Huo, Sai Ashish Somayajula, Youwei Liang, Ruisi Zhang, Farinaz Koushanfar, Pengtao Xie</i>	
Triplet Interaction Improves Graph Transformers: Accurate Molecular Graph Learning with Triplet Graph Transformers.....	20768
<i>Md Shamim Hussain, Mohammed J. Zaki, Dharmashankar Subramanian</i>	
Vanilla Bayesian Optimization Performs Great in High Dimensions.....	20793
<i>Carl Hvarfner, Erik O. Hellsten, Luigi Nardi</i>	
On Positivity Condition for Causal Inference	20818
<i>Inwoo Hwang, Yesong Choe, Yeahoon Kwon, Sanghack Lee</i>	
Fine-Grained Causal Dynamics Learning with Quantization for Improving Robustness in Reinforcement Learning	20842
<i>Inwoo Hwang, Yunhyeok Kwak, Suhyung Choi, Byoung-Tak Zhang, Sanghack Lee</i>	
Adapting Pretrained ViTs with Convolution Injector for Visuo-Motor Control.....	20871
<i>Dongyoon Hwang, Byungkun Lee, Hojoon Lee, Hyunseung Kim, Jaegul Choo</i>	

PART 26

EVEREST: Efficient Masked Video Autoencoder by Removing Redundant Spatiotemporal Tokens	20889
<i>Sunil Hwang, Jaehong Yoon, Youngwan Lee, Sung Ju Hwang</i>	
Smooth Min-Max Monotonic Networks	20908
<i>Christian Igel</i>	
SAMformer: Unlocking the Potential of Transformers in Time Series Forecasting with Sharpness-Aware Minimization and Channel-Wise Attention	20924
<i>Romain Ilbert, Ambroise Odonnat, Vasilii Feofanov, Aladin Virmaux, Giuseppe Paolo, Themis Palpanas, Ievgen Redko</i>	
From Self-Attention to Markov Models: Unveiling the Dynamics of Generative Transformers	20955
<i>M. Emrullah Ildiz, Yixiao Huang, Yingcong Li, Ankit Singh Rawat, Samet Oymak</i>	
Understanding the Learning Dynamics of Alignment with Human Feedback	20983
<i>Shawn Im, Yixuan Li</i>	
Zero-Shot Reinforcement Learning Via Function Encoders.....	21007
<i>Tyler Ingebrand, Amy Zhang, Ufuk Topcu</i>	
PASOA- PArticle baSed Bayesian Optimal Adaptive Design	21020
<i>Jacopo Iollo, Christophe Heinkelé, Pierre Alliez, Florence Forbes</i>	
Nesting Particle Filters for Experimental Design in Dynamical Systems	21047
<i>Sahel Iqbal, Adrien Corenflos, Simo Särkkä, Hany Abdulsamad</i>	
PAGER: Accurate Failure Characterization in Deep Regression Models.....	21069
<i>Jayaraman J. Thiagarajan, Vivek Narayanaswamy, Puja Trivedi, Rushil Anirudh</i>	
Online Linear Regression in Dynamic Environments Via Discounting	21083
<i>Andrew Jacobsen, Ashok Cutkosky</i>	
Human-Like Category Learning by Injecting Ecological Priors from Large Language Models into Neural Networks.....	21121
<i>Akshay K. Jagadish, Julian Coda-Forno, Mirko Thalmann, Eric Schulz, Marcel Binz</i>	
Position: Scarce Resource Allocations that Rely on Machine Learning Should Be Randomized	21148
<i>Shomik Jain, Kathleen Creel, Ashia Wilson</i>	
Learning to Reach Goals Via Diffusion.....	21170
<i>Vineet Jain, Siamak Ravanbakhsh</i>	
R2E: Turning Any GitHub Repository into a Programming Agent Environment	21196
<i>Naman Jain, Manish Shetty, Tianjun Zhang, King Han, Koushik Sen, Ion Stoica</i>	
On the Sample Complexity of Conditional Independence Testing with Von Mises Estimator with Application to Causal Discovery.....	21225
<i>Fateme Jamshidi, Luca Ganassali, Negar Kiyavash</i>	
Rethinking DP-SGD in Discrete Domain: Exploring Logistic Distribution in the Realm of signSGD.....	21241
<i>Jonggyu Jang, Seongjin Hwang, Hyun Jong Yang</i>	
Degeneration-Free Policy Optimization: RL Fine-Tuning for Language Models Without Degeneration	21266
<i>Youngsoo Jang, Geon-Hyeong Kim, Byoungjip Kim, Yu Jin Kim, Honglak Lee, Moontae Lee</i>	
Visual Representation Learning with Stochastic Frame Prediction.....	21289
<i>Huiwon Jang, Dongyoung Kim, Junsu Kim, Jinwoo Shin, Pieter Abbeel, Younggyo Seo</i>	
LoRA Training in the NTK Regime Has No Spurious Local Minima.....	21306
<i>Uijeong Jang, Jason D. Lee, Ernest K. Ryu</i>	

Efficient Low-Rank Matrix Estimation, Experimental Design, and Arm-Set-Dependent Low-Rank Bandits.....	21329
<i>Kyoungseok Jang, Chicheng Zhang, Kwang-Sung Jun</i>	
Bringing Motion Taxonomies to Continuous Domains Via GPLVM on Hyperbolic Manifolds	21373
<i>Noémie Jaquier, Leonel Rozo, Miguel González-Duque, Viacheslav Borovitskiy, Tamim Asfour</i>	
PriorBoost: An Adaptive Algorithm for Learning from Aggregate Responses.....	21410
<i>Adel Javanmard, Matthew Fahrbach, Vahab Mirrokni</i>	
Low-Rank Bandits Via Tight Two-To-Infinity Singular Subspace Recovery.....	21430
<i>Yassir Jedra, William Réveillard, Stefan Stojanovic, Alexandre Proutiere</i>	
Gradual Divergence for Seamless Adaptation: A Novel Domain Incremental Learning Method	21486
<i>Kishaan Jeeveswaran, Elahe Arani, Bahram Zonooz</i>	
Repeat After Me: Transformers Are Better than State Space Models at Copying Transformers Are Better than State Space Models at Copying	21502
<i>Samy Jelassi, David Brandfonbrener, Sham M. Kakade, Eran Malach</i>	
An Information-Theoretic Analysis of In-Context Learning.....	21522
<i>Hong Jun Jeon, Jason D. Lee, Qi Lei, Benjamin Van Roy</i>	
Finite Volume Features, Global Geometry Representations, and Residual Training for Deep Learning-Based CFD Simulation	21555
<i>Loh Sher En Jessica, Naheed Anjum Arafat, Wei Xian Lim, Wai Lee Chan, Adams Wai Kin Kong</i>	
ReLU to the Rescue: Improve Your On-Policy Actor-Critic with Positive Advantages.....	21577
<i>Andrew Jesson, Chris Lu, Gunshi Gupta, Nicolas Beltran-Velez, Angelos Filos, Jakob N. Foerster, Yarin Gal</i>	
Advancing Dynamic Sparse Training by Exploring Optimization Opportunities	21606
<i>Jie Ji, Gen Li, Lu Yin, Minghai Qin, Geng Yuan, Linke Guo, Shiwei Liu, Xiaolong Ma</i>	
ACE : Off-Policy Actor-Critic with Causality-Aware Entropy Regularization	21620
<i>Tianying Ji, Yongyuan Liang, Yan Zeng, Yu Luo, Guowei Xu, Jiawei Guo, Ruijie Zheng, Furong Huang, Fuchun Sun, Huazhe Xu</i>	
Towards Efficient Exact Optimization of Language Model Alignment	21648
<i>Haozhe Ji, Cheng Lu, Yilin Niu, Pei Ke, Hongning Wang, Jun Zhu, Jie Tang, Minlie Huang</i>	
Seizing Serendipity: Exploiting the Value of Past Success in Off-Policy Actor-Critic	21672
<i>Tianying Ji, Yu Luo, Fuchun Sun, Xianyuan Zhan, Jianwei Zhang, Huazhe Xu</i>	
Discrete Latent Perspective Learning for Segmentation and Detection	21719
<i>Deyi Ji, Feng Zhao, Lanyun Zhu, Wenwei Jin, Hongtao Lu, Jieping Ye</i>	

PART 27

Simulation-Based Inference with Quantile Regression.....	21731
<i>He Jia</i>	
GeminiFusion: Efficient Pixel-Wise Multimodal Fusion for Vision Transformer.....	21753
<i>Ding Jia, Jianyuan Guo, Kai Han, Han Wu, Chao Zhang, Chang Xu, Xinghao Chen</i>	
Chain-Of-Thought Predictive Control.....	21768
<i>Zhiwei Jia, Vineet Thumulari, Fangchen Liu, Linghao Chen, Zhiao Huang, Hao Su</i>	
Pre-Training Protein Bi-Level Representation Through Span Mask Strategy on 3D Protein Chains	21791
<i>Jiale Zhao, Wanru Zhuang, Jia Song, Yaqi Li, Shuqi Lu</i>	
NDOT: Neuronal Dynamics-Based Online Training for Spiking Neural Networks.....	21806
<i>Haiyan Jiang, Giulia De Masi, Huan Xiong, Bin Gu</i>	

Conditional Common Entropy for Instrumental Variable Testing and Partial Identification.....	21824
<i>Ziwei Jiang, Murat Kocaoglu</i>	
ProtoGate: Prototype-Based Neural Networks with Global-To-Local Feature Selection for Tabular Biomedical Data	21844
<i>Xiangjian Jiang, Andrei Margeloiu, Nikola Simidjievski, Mateja Jamnik</i>	
On the Origins of Linear Representations in Large Language Models	21879
<i>Yibo Jiang, Goutham Rajendran, Pradeep Ravikumar, Bryon Aragam, Victor Veitch</i>	
Federated Optimization with Doubly Regularized Drift Correction	21912
<i>Xiaowen Jiang, Anton Rodomanov, Sebastian U. Stich</i>	
HexGen: Generative Inference of Large Language Model Over Heterogeneous Environment	21946
<i>Youhe Jiang, Ran Yan, Xiaozhe Yao, Yang Zhou, Beidi Chen, Binhang Yuan</i>	
Projection-Free Variance Reduction Methods for Stochastic Constrained Multi-Level Compositional Optimization.....	21962
<i>Wei Jiang, Sifan Yang, Wenhao Yang, Yibo Wang, Yuanyu Wan, Lijun Zhang</i>	
Tabular Insights, Visual Impacts: Transferring Expertise from Tables to Images	21988
<i>Jun-Peng Jiang, Han-Jia Ye, Leye Wang, Yang Yang, Yuan Jiang, De-Chuan Zhan</i>	
Generalized Neural Collapse for a Large Number of Classes	22010
<i>Jiachen Jiang, Jinxin Zhou, Peng Wang, Qing Qu, Dustin G. Mixon, Chong You, Zhihui Zhu</i>	
SuDA: Support-Based Domain Adaptation for Sim2Real blackHinge Joint Tracking with Flexible Sensors	22042
<i>Jiawei Fang, Haishan Song, Chengxu Zuo, Xiaoxia Gao, Xiaowei Chen, Shihui Guo, Yipeng Qin</i>	
Memory-Space Visual Prompting for Efficient Vision-Language Fine-Tuning	22062
<i>Shibo Jie, Yehui Tang, Ning Ding, Zhi-Hong Deng, Kai Han, Yunhe Wang</i>	
Homomorphism Counts for Graph Neural Networks: All About that Basis.....	22075
<i>Emily Jin, Michael Bronstein, Ismail Ilkan Ceylan, Matthias Lanzinger</i>	
LLM Maybe LongLM: SelfExtend LLM Context Window Without Tuning.....	22099
<i>Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia-Yuan Chang, Huiyuan Chen, Xia Hu</i>	
On the Maximal Local Disparity of Fairness-Aware Classifiers	22115
<i>Jinqiu Jin, Haoxuan Li, Fuli Feng</i>	
What Will My Model Forget? Forecasting Forgotten Examples in Language Model Refinement	22145
<i>Xisen Jin, Xiang Ren</i>	
Emergent Representations of Program Semantics in Language Models Trained on Programs	22160
<i>Charles Jin, Martin Rinard</i>	
Video-LaVIT: Unified Video-Language Pre-Training with Decoupled Visual-Motional Tokenization	22185
<i>Yang Jin, Zhicheng Sun, Kun Xu, Kun Xu, Liwei Chen, Hao Jiang, Quzhe Huang, Chengru Song, Yuliang Liu, Di Zhang, Yang Song, Kun Gai, Yadong Mu</i>	
An Image is Worth Multiple Words: Discovering Object Level Concepts Using Multi-Concept Prompt Learning.....	22210
<i>Chen Jin, Ryutaro Tanno, Amrutha Saseendran, Tom Diethe, Philip Teare</i>	
Language Models as Semantic Indexers.....	22244
<i>Bowen Jin, Hansi Zeng, Guoyin Wang, Xiusi Chen, Tianxin Wei, Ruirui Li, Zhengyang Wang, Zheng Li, Yang Li, Hanqing Lu, Suhang Wang, Jiawei Han, Xianfeng Tang</i>	
Position: What Can Large Language Models Tell Us About Time Series Analysis	22260
<i>Ming Jin, Yifan Zhang, Wei Chen, Kexin Zhang, Yuxuan Liang, Bin Yang, Jindong Wang, Shirui Pan, Qingsong Wen</i>	

AlphaFold Meets Flow Matching for Generating Protein Ensembles	22277
<i>Bowen Jing, Bonnie Berger, Tommi Jaakkola</i>	
FedSC: Provable Federated Self-Supervised Learning with Spectral Contrastive Objective Over Non-I.i.d. Data.....	22304
<i>Shusen Jing, Anlan Yu, Shuai Zhang, Songyang Zhang</i>	
Model-Based Minimum Bayes Risk Decoding for Text Generation	22326
<i>Yuu Jinnai, Tetsuro Morimura, Ukyo Honda, Kaito Ariu, Kenshi Abe</i>	
Generative Modeling on Manifolds Through Mixture of Riemannian Diffusion Processes	22348
<i>Jaehyeong Jo, Sung Ju Hwang</i>	
Graph Generation with Diffusion Mixture	22371
<i>Jaehyeong Jo, Dongki Kim, Sung Ju Hwang</i>	
Experts Don't Cheat: Learning What You Don't Know by Predicting Pairs	22406
<i>Daniel D. Johnson, Daniel Tarlow, David Duvenaud, Chris J. Maddison</i>	
Learning to Infer Generative Template Programs for Visual Concepts	22465
<i>R. Kenny Jones, Siddhartha Chaudhuri, Daniel Ritchie</i>	
Learning Coverage Paths in Unknown Environments with Deep Reinforcement Learning	22491
<i>Arvi Jonnarth, Jie Zhao, Michael Felsberg</i>	
IW-GAE: Importance Weighted Group Accuracy Estimation for Improved Calibration and Model Selection in Unsupervised Domain Adaptation.....	22509
<i>Taejong Joo, Diego Klabjan</i>	
Decoupling Feature Extraction and Classification Layers for Calibrated Neural Networks	22530
<i>Mikkel Jordahn, Pablo M. Olmos</i>	
Position: Benchmarking is Limited in Reinforcement Learning Research.....	22551
<i>Scott M. Jordan, Adam White, Bruno Castro Da Silva, Martha White, Philip S. Thomas</i>	

PART 28

Watermark Stealing in Large Language Models.....	22570
<i>Nikola Jovanovic, Robin Staab, Martin Vechev</i>	
Hypergraph-Enhanced Dual Semi-Supervised Graph Classification	22594
<i>Wei Ju, Zhengyang Mao, Siyu Yi, Yifang Qin, Yiyang Gu, Zhiping Xiao, Yifan Wang, Xiao Luo, Ming Zhang</i>	
NaturalSpeech 3: Zero-Shot Speech Synthesiswith Factorized Codec and Diffusion Models	22605
<i>Zeqian Ju, Yuancheng Wang, Kai Shen, Xu Tan, Detai Xin, Dongchao Yang, Yanqing Liu, Yichong Leng, Kaitao Song, Siliang Tang, Zhizheng Wu, Tao Qin, Xiang-Yang Li, Wei Ye, Shikun Zhang, Jiang Bian, Lei He, Jinyu Li, Sheng Zhao</i>	
Is Epistemic Uncertainty Faithfully Represented by Evidential Deep Learning Methods?	22624
<i>Mira Jürgens, Nis Meinert, Viktor Bengs, Eyke Hüllermeier, Willem Waegeman</i>	
Noise-Adaptive Confidence Sets for Linear Bandits and Application to Bayesian Optimization.....	22643
<i>Kwang-Sung Jun, Jungtaek Kim</i>	
Unsupervised Episode Generation for Graph Meta-Learning	22672
<i>Jihyeong Jung, Sangwoo Seo, Sungwon Kim, Chanyoung Park</i>	
PruNeRF: Segment-Centric Dataset Pruning Via 3D Spatial Consistency.....	22696
<i>Yeonsung Jung, Heecheol Yun, Joonhyung Park, Jin-Hwa Kim, Eunho Yang</i>	
Leveraging VLM-Based Pipelines to Annotate 3D Objects.....	22710
<i>Rishabh Kabra, Loic Matthey, Alexander Lerchner, Niloy J. Mitra</i>	

PolySketchFormer: Fast Transformers Via Sketching Polynomial Kernels	22748
<i>Praneeth Kacham, Vahab Mirrokni, Peilin Zhong</i>	
Surprisingly Strong Performance Prediction with Neural Graph Features	22771
<i>Gabriela Kadlecová, Jovita Lukasik, Martin Pilát, Petra Vidnerová, Mahmoud Safari, Roman Neruda, Frank Hutter</i>	
EvTexture: Event-Driven Texture Enhancement for Video Super-Resolution	22817
<i>Dachun Kai, Jiayao Lu, Yueyi Zhang, Xiaoyan Sun</i>	
Beyond the Calibration Point: Mechanism Comparison in Differential Privacy	22840
<i>Georgios Kaissis, Stefan Kolek, Borja Balle, Jamie Hayes, Daniel Rueckert</i>	
Replicable Learning of Large-Margin Halfspaces	22861
<i>Alkis Kalavasis, Amin Karbasi, Kasper Green Larsen, Grigoris Velegkas, Felix Zhou</i>	
Tell, Don't Show!: Language Guidance Eases Transfer Across Domains in Images and Videos	22879
<i>Tarun Kalluri, Bodhisattwa Prasad Majumder, Manmohan Chandraker</i>	
Position: LLMs Can't Plan, but Can Help Planning in LLM-Modulo Frameworks	22895
<i>Subbarao Kambhampati, Karthik Valmееkam, Lin Guan, Mudit Verma, Kaya Stechly, Siddhant Bhambri, Lucas Saldyt, Anil Murthy</i>	
A Geometric Explanation of the Likelihood OOD Detection Paradox	22908
<i>Hamidreza Kamkari, Brendan Leigh Ross, Jesse C. Cresswell, Anthony L. Caterini, Rahul G. Krishnan, Gabriel Loaiza-Ganem</i>	
Learning Decision Trees and Forests with Algorithmic Recourse	22936
<i>Kentaro Kanamori, Takuya Takagi, Ken Kobayashi, Yuichi Ike</i>	
C-RAG: Certified Generation Risks for Retrieval-Augmented Language Models	22963
<i>Mintong Kang, Nezihe Merve Gürel, Ning Yu, Dawn Song, Bo Li</i>	
Think Before You Act: Decision Transformers with Working Memory	23001
<i>Jikun Kang, Romain Laroche, Xingdi Yuan, Adam Trischler, Xue Liu, Jie Fu</i>	
Certiably Byzantine-Robust Federated Conformal Prediction	23022
<i>Mintong Kang, Zhen Lin, Jimeng Sun, Cao Xiao, Bo Li</i>	
Neural Tangent Kernels for Axis-Aligned Tree Ensembles	23058
<i>Ryuichi Kanoh, Mahito Sugiyama</i>	
Position: On the Societal Impact of Open Foundation Models	23082
<i>Sayash Kapoor, Rishi Bommasani, Kevin Klyman, Shayne Longpre, Ashwin Ramaswami, Peter Cihon, Aspen Hopkins, Kevin Bankston, Stella Biderman, Miranda Bogen, Rumman Chowdhury, Alex Engler, Peter Henderson, Yacine Jernite, Seth Lazar, Stefano Maffulli, Alondra Nelson, Joelle Pineau, Aviya Skowron, Dawn Song, Victor Storchan, Daniel Zhang, Daniel E. Ho, Percy Liang, Arvind Narayanan</i>	
Adaptive Robust Learning Using Latent Bernoulli Variables	23105
<i>Aleksandr Karakulev, Dave Zachariah, Prashant Singh</i>	
Prismatic VLMs: Investigating the Design Space of Visually-Conditioned Language Models	23123
<i>Siddharth Karamcheti, Suraj Nair, Ashwin Balakrishna, Percy Liang, Thomas Kollar, Dorsa Sadigh</i>	
On Stronger Computational Separations Between Multimodal and Unimodal Machine Learning	23145
<i>Ari Karchmer</i>	
On the Generalization of Equivariant Graph Neural Networks	23159
<i>Rafa Karczewski, Amauri H. Souza, Vikas Garg</i>	
Infinite-Horizon Distributionally Robust Regret-Optimal Control	23187
<i>Taylan Kargin, Joudi Hajar, Vikrant Malik, Babak Hassibi</i>	

Challenges and Considerations in the Evaluation of Bayesian Causal Discovery.....	23215
<i>Amir Mohammad Karimi Mamaghan, Panagiotis Tigas, Karl Henrik Johansson, Yarin Gal, Yashas Annadani, Stefan Bauer</i>	
Progressive Inference: Explaining Decoder-Only Sequence Classification Models Using Intermediate Predictions.....	23238
<i>Sanjay Kariyappa, Freddy Lécué, Saumitra Mishra, Christopher Pond, Daniele Magazzeni, Manuela Veloso</i>	
Position: Embracing Negative Results in Machine Learning.....	23256
<i>Florian Karl, Lukas Malte Kemeter, Gabriel Dax, Paulina Sierak</i>	
Hidden Traveling Waves Bind Working Memory Variables in Recurrent Neural Networks.....	23266
<i>Arjun Karuvally, Terrence J. Sejnowski, Hava T. Siegelmann</i>	
Active Adaptive Experimental Design for Treatment Effect Estimation with Covariate Choice.....	23291
<i>Masahiro Kato, Akihiro Oga, Wataru Komatsubara, Ryo Inokuchi</i>	
First-Order Manifold Data Augmentation for Regression Learning.....	23324
<i>Ilya Kaufman, Omri Azencot</i>	
Balanced Data, Imbalanced Spectra: Unveiling Class Disparities with Spectral Imbalance.....	23343
<i>Chiraag Kaushik, Ran Liu, Chi-Heng Lin, Amrit Khera, Matthew Y Jin, Wenrui Ma, Vidya Muthukumar, Eva L Dyer</i>	
Pluvial Flood Emulation with Hydraulics-Informed Message Passing.....	23367
<i>Arnold Kazadi, James Doss-Gollin, Arlei Silva</i>	
Accelerating Convergence in Bayesian Few-Shot Classification.....	23391
<i>Tianjun Ke, Haoqun Cao, Feng Zhou</i>	
An Improved Finite-Time Analysis of Temporal Difference Learning with Deep Neural Networks.....	23407
<i>Zhifa Ke, Zaiwen Wen, Junyu Zhang</i>	

PART 29

DUPLEX: Dual GAT for Complex Embedding of Directed Graphs.....	23430
<i>Zhaoru Ke, Hang Yu, Jianguo Li, Haipeng Zhang</i>	
Transformers Get Stable: An End-To-End Signal Propagation Theory for Language Models.....	23449
<i>Akhil Kedia, Mohd Abbas Zaidi, Sushil Khyalia, Jungho Jung, Harshith Goka, Haejun Lee</i>	
A Universal Transfer Theorem for Convex Optimization Algorithms Using Inexact First-Order Oracles.....	23532
<i>Phillip Kerger, Marco Molinaro, Hongyi Jiang, Amitabh Basu</i>	
Fair Classification with Partial Feedback: An Exploration-Based Data Collection Approach.....	23547
<i>Vijay Keswani, Anay Mehrotra, L. Elisa Celis</i>	
Neural Tangent Kernels Motivate Cross-Covariance Graphs in Neural Networks.....	23577
<i>Shervin Khalafi, Saurabh Sihag, Alejandro Ribeiro</i>	
Tuning-Free Stochastic Optimization.....	23622
<i>Ahmed Khaled, Chi Jin</i>	
Debating with More Persuasive LLMs Leads to More Truthful Answers.....	23662
<i>Akbir Khan, John Hughes, Dan Valentine, Laura Ruis, Kshitij Sachan, Ansh Radhakrishnan, Edward Grefenstette, Samuel R. Bowman, Tim Rocktäschel, Ethan Perez</i>	
Off-Policy Evaluation Beyond Overlap: Sharp Partial Identification Under Smoothness.....	23734
<i>Samir Khan, Martin Saveski, Johan Ugander</i>	

Towards an Understanding of Stepwise Inference in Transformers: A Synthetic Graph Navigation Model	23758
<i>Mikail Khona, Maya Okawa, Jan Hula, Rahul Ramesh, Kento Nishi, Robert Dick, Ekdeep Singh Lubana, Hidenori Tanaka</i>	
Can Machines Learn the True Probabilities?	23781
<i>Jinsook Kim</i>	
Gaussian Plane-Wave Neural Operator for Electron Density Estimation.....	23805
<i>Seongsu Kim, Sungsoo Ahn</i>	
LayerMerge: Neural Network Depth Compression Through Layer Pruning and Merging	23825
<i>Jinuk Kim, Marwa El Halabi, Mingi Ji, Hyun Oh Song</i>	
CARTE: Pretraining and Transfer for Tabular Learning	23843
<i>Myung Jun Kim, Léo Grinsztajn, Gaël Varoquaux</i>	
Achieving Lossless Gradient Sparsification Via Mapping to Alternative Space in Federated Learning	23867
<i>Do-Yeon Kim, Dong-Jun Han, Jun Seo, Jaekyun Moon</i>	
SqueezeLLM: Dense-And-Sparse Quantization.....	23901
<i>Sehoon Kim, Coleman Hooper, Amir Gholami, Zhen Dong, Xiuyu Li, Sheng Shen, Michael W. Mahoney, Kurt Keutzer</i>	
Active Label Correction for Semantic Segmentation with Foundation Models.....	23924
<i>Hoyoung Kim, Sehyun Hwang, Suha Kwak, Jungseul Ok</i>	
ODIM: Outlier Detection Via Likelihood of Under-Fitted Generative Models.....	23941
<i>Dongha Kim, Jaesung Hwang, Jongjin Lee, Kunwoong Kim, Yongdai Kim</i>	
Hybrid Neural Representations for Spherical Data	23972
<i>Hyomin Kim, Yunhui Jang, Jaeho Lee, Sungsoo Ahn</i>	
Synergistic Integration of Coordinate Network and Tensorial Feature for Improving Neural Radiance Fields from Sparse Inputs	23986
<i>Mingyu Kim, Jun-Seong Kim, Se-Young Yun, Jin-Hwa Kim</i>	
Learning to Explore for Stochastic Gradient MCMC.....	24015
<i>Seunghyun Kim, Seohyeon Jung, Seonghyeon Kim, Juho Lee</i>	
Improving Robustness to Multiple Spurious Correlations by Multi-Objective Optimization	24040
<i>Nayeong Kim, Juwon Kang, Sungsoo Ahn, Jungseul Ok, Suha Kwak</i>	
Double-Step Alternating Extragradient with Increasing Timescale Separation for Finding Local Minimax Points: Provable Improvements	24059
<i>Kyuwon Kim, Donghwan Kim</i>	
Scene Graph Generation Strategy with Co-Occurrence Knowledge and Learnable Term Frequency	24094
<i>Hyeongjin Kim, Sangwon Kim, Dasom Ahn, Jong Taek Lee, Byoung Chul Ko</i>	
Symmetric Replay Training: Enhancing Sample Efficiency in Deep Reinforcement Learning for Combinatorial Optimization	24110
<i>Hyeonah Kim, Minsu Kim, Sungsoo Ahn, Jinkyoo Park</i>	
Clustered Federated Learning Via Gradient-Based Partitioning.....	24137
<i>Heasung Kim, Hyeji Kim, Gustavo De Veciana</i>	
Variational Partial Group Convolutions.....	24194
<i>Hyunsu Kim, Yegon Kim, Hongseok Yang, Juho Lee</i>	
Demystifying SGD with Doubly Stochastic Gradients	24210
<i>Kyurae Kim, Joohwan Ko, Yi-An Ma, Jacob R. Gardner</i>	

Learning to Scale Logits for Temperature-Conditional GFlowNets.....	24248
<i>Minsu Kim, Joohwan Ko, Taeyoung Yun, Dinghuai Zhang, Ling Pan, Woo Chang Kim, Jinkyoo Park, Emmanuel Bengio, Yoshua Bengio</i>	

PART 30

Attribute Based Interpretable Evaluation Metrics for Generative Models	24271
<i>Dongkyun Kim, Mingi Kwon, Youngjung Uh</i>	
Investigating Pre-Training Objectives for Generalization in Vision-Based Reinforcement Learning.....	24294
<i>Donghu Kim, Hojoon Lee, Kyungmin Lee, Dongyoon Hwang, Jaegul Choo</i>	
EquiAV: Leveraging Equivariance for Audio-Visual Contrastive Learning.....	24327
<i>Jongsuk Kim, Hyeongkeun Lee, Kyeongha Rho, Junmo Kim, Joon Son Chung</i>	
Scaling Beyond the GPU Memory Limit for Large Mixture-Of-Experts Model Training	24342
<i>Yechan Kim, Hwijoon Lim, Dongsu Han</i>	
Risk-Sensitive Policy Optimization Via Predictive CVaR Policy Gradient.....	24354
<i>Ju-Hyun Kim, Seungki Min</i>	
An LLM Compiler for Parallel Function Calling	24370
<i>Sehoon Kim, Suhong Moon, Ryan Tabrizi, Nicholas Lee, Michael W. Mahoney, Kurt Keutzer, Amir Gholami</i>	
Data-Efficient Molecular Generation with Hierarchical Textual Inversion.....	24392
<i>Seojin Kim, Jaehyun Nam, Sihyun Yu, Younghoon Shin, Jinwoo Shin</i>	
Translating Subgraphs to Nodes Makes Simple GNNs Strong and Efficient for Subgraph Representation Learning.....	24415
<i>Dongkwan Kim, Alice Oh</i>	
Privacy-Preserving Embedding Via Look-Up Table Evaluation with Fully Homomorphic Encryption	24437
<i>Jae-Yun Kim, Saerom Park, Joohee Lee, Jung Hee Cheon</i>	
Convex Relaxations of ReLU Neural Networks Approximate Global Optima in Polynomial Time.....	24458
<i>Sungyoon Kim, Mert Pilanci</i>	
USTAD: Unified Single-Model Training Achieving Diverse Scores for Information Retrieval	24486
<i>Seungyeon Kim, Ankit Singh Rawat, Manzil Zaheer, Wittawat Jitkittum, Veeranjanyulu Sadhanala, Sadeep Jayasumana, Aditya Krishna Menon, Rob Fergus, Sanjiv Kumar</i>	
Polynomial-Based Self-Attention for Table Representation Learning	24509
<i>Jayoung Kim, Yehjin Shin, Jeongwhan Choi, Hyowon Wi, Noseong Park</i>	
Transformers Learn Nonlinear Features in Context:Nonconvex Mean-Field Dynamics on the Attention Landscape.....	24527
<i>Juno Kim, Taiji Suzuki</i>	
Discovering Features with Synergistic Interactions in Multiple Views	24562
<i>Chohee Kim, Mihaela Van Der Schaar, Changhee Lee</i>	
An Infinite-Width Analysis on the Jacobian-Regularised Training of a Neural Network	24584
<i>Taeyoung Kim, Hongseok Yang</i>	
One Size Fits All for Semantic Shifts: Adaptive Prompt Tuning for Continual Learning	24658
<i>Doyoung Kim, Susik Yoon, Dongmin Park, Youngjun Lee, Hwanjun Song, Jihwan Bang, Jae-Gil Lee</i>	
Do Topological Characteristics Help in Knowledge Distillation?	24674
<i>Jungeun Kim, Junwon You, Dongjin Lee, Ha Young Kim, Jae-Hun Jung</i>	

A Unified Linear Programming Framework for Offline Reward Learning from Human Demonstrations and Feedback.....	24694
<i>Kihyun Kim, Jiawei Zhang, Asuman Ozdaglar, Pablo Parrilo</i>	
The Merit of River Network Topology for Neural Flood Forecasting.....	24713
<i>Nikolas Kirschstein, Yixuan Sun</i>	
From Neurons to Neutrons: A Case Study in Interpretability.....	24726
<i>Ouail Kitouni, Niklas Nolte, Víctor Samuel Pérez-Díaz, Sokratis Trifinopoulos, Mike Williams</i>	
Conformal Prediction with Learned Features.....	24749
<i>Shayan Kiyani, George Pappas, Hamed Hassani</i>	
Context-Guided Diffusion for Out-Of-Distribution Molecular and Protein Design.....	24770
<i>Leo Klärner, Tim G. J. Rudner, Garrett M. Morris, Charlotte M. Deane, Yee Whye Teh</i>	
Environment Design for Inverse Reinforcement Learning.....	24808
<i>Thomas Kleine Buening, Victor Villin, Christos Dimitrakakis</i>	
What Would Gauss Say About Representations? Probing Pretrained Image Models Using Synthetic Gaussian Benchmarks	24829
<i>Ching-Yun Ko, Pin-Yu Chen, Payel Das, Jeet Mohapatra, Luca Daniel</i>	
Universal Consistency of Wide and Deep ReLU Neural Networks and Minimax Optimal Convergence Rates for Kolmogorov-Donoho Optimal Function Classes	24859
<i>Hyunouk Ko, Xiaoming Huo</i>	
DistiLLM: Towards Streamlined Distillation for Large Language Models.....	24872
<i>Jongwoo Ko, Sungnyun Kim, Tianyi Chen, Se-Young Yun</i>	
Provably Scalable Black-Box Variational Inference with Structured Variational Families	24896
<i>Joohwan Ko, Kyurae Kim, Woo Chang Kim, Jacob R. Gardner</i>	
Stochastic Conditional Diffusion Models for Robust Semantic Image Synthesis	24932
<i>Juyeon Ko, Inho Kong, Dogyun Park, Hyunwoo J. Kim</i>	
Compression of Structured Data with Autoencoders: Provable Benefit of Nonlinearities and Depth.....	24964
<i>Kevin Kögler, Alexander Shevchenko, Hamed Hassani, Marco Mondelli</i>	
Estimating Barycenters of Distributions with Neural Optimal Transport.....	25016
<i>Alexander Kolesov, Petr Mokrov, Igor Udovichenko, Milena Gazdieva, Gudmund Pammer, Evgeny Burnaev, Alexander Korotin</i>	
AdsorbDiff: Adsorbate Placement Via Conditional Denoising Diffusion	25042
<i>Adeesh Kolluru, John R. Kitchin</i>	
On Convergence of Incremental Gradient for Non-Convex Smooth Functions.....	25058
<i>Anastasia Koloskova, Nikita Doikov, Sebastian U. Stich, Martin Jaggi</i>	
Logistic Variational Bayes Revisited.....	25087
<i>Michael Komodromos, Marina Evangelou, Sarah Filippi</i>	

PART 31

VideoPoet: A Large Language Model for Zero-Shot Video Generation.....	25105
<i>Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Grant Schindler, Rachel Hornung, Vighnesh Birodkar, Jimmy Yan, Ming-Chang Chiu, Krishna Somandepalli, Hassan Akbari, Yair Alon, Yong Cheng, Josh Dillon, Agrim Gupta, Meera Hahn, Anja Hauth, David Hendon, Alonso Martinez, David Minnen, Mikhail Sirotenko, Kihyuk Sohn, Xuan Yang, Hartwig Adam, Ming-Hsuan Yang, Irfan Essa, Huisheng Wang, David A. Ross, Bryan Seybold, Lu Jiang</i>	
Audio Flamingo: A Novel Audio Language Model with Few-Shot Learning and Dialogue Abilities	25125
<i>Zhifeng Kong, Arushi Goel, Rohan Badlani, Wei Ping, Rafael Valle, Bryan Catanzaro</i>	

Generalist Equivariant Transformer Towards 3D Molecular Interaction Learning	25149
<i>Xiangzhe Kong, Wenbing Huang, Yang Liu</i>	
ELF: Encoding Speaker-Specific Latent Speech Feature for Speech Synthesis.....	25176
<i>Jungil Kong, Junmo Lee, Jeongmin Kim, Beomjeong Kim, Jihoon Park, Dohee Kong, Changheon Lee, Sangjin Kim</i>	
The Computational Complexity of Finding Second-Order Stationary Points	25191
<i>Andreas Kontogiannis, Vasilis Pollatos, Sotiris Kanellopoulos, Panayotis Mertikopoulos, Aris Pagourtzis, Ioannis Panageas</i>	
Information-Directed Pessimism for Offline Reinforcement Learning	25226
<i>Alec Koppel, Sujay Bhatt, Jiacheng Guo, Joe Eappen, Mengdi Wang, Sumitra Ganesh</i>	
Understanding and Diagnosing Deep Reinforcement Learning	25265
<i>Ezgi Korkmaz</i>	
Bridging Mini-Batch and Asymptotic Analysis in Contrastive Learning: From InfoNCE to Kernel-Based Losses	25276
<i>Panagiotis Koromilas, Giorgos Bouritsas, Theodoros Giannakopoulos, Mihalis A. Nicolaou, Yannis Panagakis</i>	
convSeq: Fast and Scalable Method for Detecting Patterns in Spike Data.....	25302
<i>Roman Koshkin, Tomoki Fukai</i>	
Privacy Profiles for Private Selection.....	25313
<i>Antti Koskela, Rachel Redberg, Yu-Xiang Wang</i>	
Rotational Equilibrium: How Weight Decay Balances Learning Across Neural Networks	25333
<i>Atli Kossou, Bettina Messmer, Martin Jaggi</i>	
Consistent Long-Term Forecasting of Ergodic Dynamical Systems	25370
<i>Vladimir Kostic, Karim Lounici, Prune Inzerili, Pietro Novelli, Massimiliano Pontil</i>	
A General Online Algorithm for Optimizing Complex Performance Metrics	25396
<i>Wojciech Kotlowski, Marek Wydmuch, Erik Schultheis, Rohit Babbar, Krzysztof Dembczyński</i>	
CLLMs: Consistency Large Language Models	25426
<i>Siqi Kou, Lanxiang Hu, Zhezhi He, Zhijie Deng, Hao Zhang</i>	
KISA: A Unified Keyframe Identifier and Skill Annotator for Long-Horizon Robotics Demonstrations.....	25441
<i>Longxin Kou, Fei Ni, Yan Zheng, Jinyi Liu, Yifu Yuan, Zibin Dong, Jianye Hao</i>	
PcLast: Discovering Plannable Continuous Latent States	25475
<i>Anurag Koul, Shivakanth Sujit, Shaoru Chen, Ben Evans, Lili Wu, Byron Xu, Rajan Chari, Riashat Islam, Raihan Seraj, Yonathan Efroni, Lekan Molu, Miro Dudik, John Langford, Alex Lamb</i>	
Sobolev Space Regularised Pre Density Models.....	25494
<i>Mark Kozdoba, Binyamin Perets, Shie Mannor</i>	
Implicit Meta-Learning May Lead Language Models to Trust More Reliable Sources	25534
<i>Dmitrii Krasheninnikov, Egor Krasheninnikov, Bruno Mlodozeniec, Tegan Maharaj, David Krueger</i>	
Geometry-Aware Instrumental Variable Regression.....	25560
<i>Heiner Kremer, Bernhard Schölkopf</i>	
Understanding the Effects of Iterative Prompting on Truthfulness	25583
<i>Satyapriya Krishna, Chirag Agarwal, Himabindu Lakkaraju</i>	
A Sober Look at LLMs for Material Discovery: Are They Actually Good for Bayesian Optimization Over Molecules?.....	25603
<i>Agustinus Kristiadi, Felix Strieth-Kalthoff, Marta Skreta, Pascal Poupart, Alán Aspuru-Guzik, Geoff Pleiss</i>	

Towards General Algorithm Discovery for Combinatorial Optimization: Learning Symbolic Branching Policy from Bipartite Graph.....	25623
<i>Yufei Kuang, Jie Wang, Yuyan Zhou, Xijun Li, Fangzhou Zhu, Jianye Hao, Feng Wu</i>	
Mean Estimation in the Add-Remove Model of Differential Privacy	25642
<i>Alex Kulesza, Ananda Theertha Suresh, Yuyan Wang</i>	
No Free Prune: Information-Theoretic Barriers to Pruning at Initialization.....	25662
<i>Tanishq Kumar, Kevin Luo, Mark Sellke</i>	
Efficient Value Iteration for S-Rectangular Robust Markov Decision Processes	25682
<i>Navdeep Kumar, Kaixinwang, Kfir Levy, Shie Mannor</i>	
Minimum Norm Interpolation Meets the Local Theory of Banach Spaces.....	25726
<i>Gil Kur, Pedro Abdalla, Pierre Bizeul, Fanny Yang</i>	
Prospective Side Information for Latent MDPs	25755
<i>Jeongyeol Kwon, Yonathan Efroni, Shie Mannor, Constantine Caramanis</i>	
On the Complexity of First-Order Methods in Stochastic Bilevel Optimization.....	25784
<i>Jeongyeol Kwon, Dohyun Kwon, Hanbaek Lyu</i>	
TinyTrain: Resource-Aware Task-Adaptive Sparse Training of DNNs at the Data-Scarce Edge.....	25812
<i>Young D. Kwon, Rui Li, Stylianos I. Venieris, Jagmohan Chauhan, Nicholas D. Lane, Cecilia Mascolo</i>	
Dynamic Spectral Clustering with Provable Approximation Guarantee	25844
<i>Steinar Laenen, He Sun</i>	
Collective Certified Robustness Against Graph Injection Attacks	25871
<i>Yuni Lai, Bailin Pan, Kaihuang Chen, Yancheng Yuan, Kai Zhou</i>	
Position: Evolving AI Collectives Enhance Human Diversity and Enable Self-Regulation	25892
<i>Shiyang Lai, Yujin Potter, Junsol Kim, Richard Zhuang, Dawn Song, James Evans</i>	
Invariant Risk Minimization is a Total Variation Model	25913
<i>Zhao-Rong Lai, Weiwen Wang</i>	

PART 32

Privately Learning Smooth Distributions on the Hypercube by Projections	25936
<i>Clément Lalanne, Sébastien Gadat</i>	
A Neural-Preconditioned Poisson Solver for Mixed Dirichlet and Neumann Boundary Conditions	25976
<i>Kai Weixian Lan, Elias Gueidon, Ayano Kaneda, Julian Panetta, Joseph Teran</i>	
Sub-Token ViT Embedding Via Stochastic Resonance Transformers	25995
<i>Dong Lao, Yangchao Wu, Tian Yu Liu, Alex Wong, Stefano Soatto</i>	
Single-Model Attribution of Generative Models Through Final-Layer Inversion.....	26007
<i>Mike Laszkiewicz, Jonas Ricker, Johannes Lederer, Asja Fischer</i>	
Towards Understanding Inductive Bias in Transformers: A View from Infinity	26043
<i>Itay Lavie, Guy Gur-Ari, Zohar Ringel</i>	
Modeling Caption Diversity in Contrastive Vision-Language Pretraining.....	26070
<i>Samuel Lavoie, Polina Kirichenko, Mark Ibrahim, Mahmoud Assran, Andrew Gordon Wilson, Aaron Courville, Nicolas Ballas</i>	
Offline Inverse RL: New Solution Concepts and Provably Efficient Algorithms	26085
<i>Filippo Lazzati, Mirco Mutti, Alberto Maria Metelli</i>	
Generalized Sobolev Transport for Probability Measures on a Graph	26152
<i>Tam Le, Truyen Nguyen, Kenji Fukumizu</i>	

Robust Inverse Graphics Via Probabilistic Inference	26175
<i>Tuan Anh Le, Pavel Sountsov, Matthew D. Hoffman, Ben Lee, Brian Patton, Rif A. Saurous</i>	
Knowledge Graphs Can Be Learned with Just Intersection Features	26199
<i>Duy Le, Shaochen (Henry) Zhong, Zirui Liu, Shuai Xu, Vipin Chaudhary, Kaixiong Zhou, Zhaozhuo Xu</i>	
Improved Stability and Generalization Guarantees of the Decentralized SGD Algorithm.....	26215
<i>Batiste Le Bars, Aurélien Bellet, Marc Tommasi, Kevin Scaman, Giovanni Neglia</i>	
Run-Time Task Composition with Safety Semantics	26241
<i>Kevin Leahy, Makai Mann, Zachary Serlin</i>	
Chasing Convex Functions with Long-Term Constraints.....	26259
<i>Adam Lechowicz, Nicolas Christianson, Bo Sun, Noman Bashir, Mohammad Hajiesmaili, Adam Wierman, Prashant Shenoy</i>	
Generalization Analysis of Deep Non-Linear Matrix Completion	26290
<i>Antoine Ledent, Rodrigo Alves</i>	
A Mechanistic Understanding of Alignment Algorithms: A Case Study on DPO and Toxicity	26361
<i>Andrew Lee, Xiaoyan Bai, Itamar Pres, Martin Wattenberg, Jonathan K. Kummerfeld, Rada Mihalcea</i>	
Stationary Latent Weight Inference for Unreliable Observations from Online Test-Time Adaptation.....	26379
<i>Jae-Hong Lee, Joon-Hyuk Chang</i>	
A Human-Inspired Reading Agent with Gist Memory of Very Long Contexts	26396
<i>Kuang-Huei Lee, Xinyun Chen, Hiroki Furuta, John Canny, Ian Fischer</i>	
Slow and Steady Wins the Race Maintaining Plasticity with Hare and Tortoise Networks	26416
<i>Hojoon Lee, Hyeonseong Cho, Hyunseung Kim, Donghu Kim, Dugki Min, Jaegul Choo, Clare Lyle</i>	
Fundamental Benefit of Alternating Updates in Minimax Optimization.....	26439
<i>Jaewook Lee, Hanseul Cho, Chulhee Yun</i>	
DataFreeShield: Defending Adversarial Attacks Without Training Data	26515
<i>Hyeyoon Lee, Kanghyun Choi, Dain Kwon, Sunjong Park, Mayoore Selvarasa Jaiswal, Noseong Park, Jonghyun Choi, Jinho Lee</i>	
SelMatch: Effectively Scaling Up Dataset Distillation Via Selection-Based Initialization and Partial Updates by Trajectory Matching	26546
<i>Yongmin Lee, Hye Won Chung</i>	
Recurrent Early Exits for Federated Learning with Heterogeneous Clients.....	26568
<i>Royson Lee, Javier Fernandez-Marques, Shell Xu Hu, Da Li, Stefanos Laskaridis, Łukasz Dudziak, Timothy Hospedales, Ferenc Huszár, Nicholas D. Lane</i>	
PAC-Bayesian Generalization Bounds for Knowledge Graph Representation Learning	26589
<i>Jaejun Lee, Minsung Hwang, Joyce Jiyoung Whang</i>	
Learning to Continually Learn with the Bayesian Principle.....	26621
<i>Soochan Lee, Hyeonseong Jeon, Jaehyeon Son, Gunhee Kim</i>	
Graph Neural Networks with a Distribution of Parametrized Graphs	26640
<i>See Hian Lee, Feng Ji, Kelin Xia, Wee Peng Tay</i>	
Pausing Policy Learning in Non-Stationary Reinforcement Learning	26661
<i>Hyunin Lee, Ming Jin, Javad Lavaei, Somayeh Sojoudi</i>	
Feature Distribution on Graph Topology Mediates the Effect of Graph Convolution: Homophily Perspective.....	26686
<i>Soo Yong Lee, Sunwoo Kim, Fanchen Bu, Jaemin Yoo, Jiliang Tang, Kijung Shin</i>	

Neural Image Compression with Text-Guided Encoding for Both Pixel-Level and Perceptual Fidelity	26715
<i>Hagyeeong Lee, Minkyu Kim, Jun-Hyuk Kim, Seungeon Kim, Dokwan Oh, Jaeho Lee</i>	
Drug Discovery with Dynamic Goal-Aware Fragments.....	26731
<i>Seul Lee, Seanie Lee, Kenji Kawaguchi, Sung Ju Hwang</i>	
Supervised Matrix Factorization: Local Landscape Analysis and Applications.....	26752
<i>Joowon Lee, Hanbaek Lyu, Weixin Yao</i>	

PART 33

Defining Neural Network Architecture Through Polytope Structures of Datasets	26789
<i>Sangmin Lee, Abbas Mammadov, Jong Chul Ye</i>	
Why Do Animals Need Shaping? a Theory of Task Composition and Curriculum Learning	26837
<i>Jin Hwa Lee, Stefano Sarao Mannelli, Andrew Saxe</i>	
3D Geometric Shape Assembly Via Efficient Point Cloud Matching	26856
<i>Nahyuk Lee, Juhong Min, Junha Lee, Seungwook Kim, Kanghee Lee, Jaesik Park, Minsu Cho</i>	
RLAIF Vs. RLHF: Scaling Reinforcement Learning from Human Feedback with AI Feedback	26874
<i>Harrison Lee, Samrat Phatale, Hassan Mansoor, Thomas Mesnard, Johan Ferret, Kellie Lu, Colton Bishop, Ethan Hall, Victor Carbune, Abhinav Rastogi, Sushant Prakash</i>	
StrWAEs to Invariant Representations	26902
<i>Hyunjong Lee, Yedarm Seong, Sungdong Lee, Joong-Ho Won</i>	
Binning as a Pretext Task: Improving Self-Supervised Learning in Tabular Domains	26929
<i>Kyungeun Lee, Ye Seul Sim, Hye-Seung Cho, Moonjung Eo, Suhee Yoon, Sanghyu Yoon, Woohyung Lim</i>	
Training Greedy Policy for Proposal Batch Selection in Expensive Multi-Objective Combinatorial Optimization.....	26948
<i>Deokjae Lee, Hyun Oh Song, Kyunghyun Cho</i>	
Robust Optimization in Protein Fitness Landscapes Using Reinforcement Learning in Latent Space	26976
<i>Minji Lee, Luiz Felipe Vecchiatti, Hyunkyu Jung, Hyunjoon Ro, Meeyoung Cha, Ho Min Kim</i>	
Behavior Generation with Latent Actions	26991
<i>Seungjae Lee, Yibin Wang, Haritheja Etukuru, H. Jin Kim, Nur Muhammad Mahi Shaftullah, Lerrel Pinto</i>	
Improving Instruction Following in Language Models Through Proxy-Based Uncertainty Estimation.....	27009
<i>Joonho Lee, Jae Oh Woo, Juree Seok, Parisa Hassanzadeh, Wooseok Jang, Juyoun Son, Sima Didari, Baruch Gutow, Heng Hao, Hankyu Moon, Wenjun Hu, Yeong-Dae Kwon, Taehee Lee, Seungjai Min</i>	
Rethinking the Flat Minima Searching in Federated Learning.....	27037
<i>Taehwan Lee, Sung Whan Yoon</i>	
BECoTTA: Input-Dependent Online Blending of Experts for Continual Test-Time Adaptation	27072
<i>Daeun Lee, Jaehong Yoon, Sung Ju Hwang</i>	
STELLA: Continual Audio-Video Pre-Training with Spatio-Temporal Localized Alignment.....	27094
<i>Jaewoo Lee, Jaehong Yoon, Wonjae Kim, Yunji Kim, Sung Ju Hwang</i>	
Sign Rank Limitations for Inner Product Graph Decoders.....	27118
<i>Su Hyeong Lee, Qingqi Zhang, Risi Kondor</i>	
A Geometric Decomposition of Finite Games: Convergence Vs. Recurrence Under Exponential Weights.....	27137
<i>Davide Legacci, Panayotis Mertikopoulos, Bary Pradelski</i>	

Langevin Policy for Safe Reinforcement Learning	27174
<i>Fenghao Lei, Long Yang, Shiting Wen, Zhixiong Huang, Zhiwang Zhang, Chaoyi Pang</i>	
Sliced-Wasserstein Estimation with Spherical Harmonics as Control Variates	27191
<i>Rémi Leluc, Aymeric Dieuleveut, François Portier, Johan Segers, Aigerim Zhuman</i>	
An Independence-Promoting Loss for Music Generation with Language Models.....	27215
<i>Jean-Marie Lemerrier, Simon Rouard, Jade Copet, Yossi Adi, Alexandre Defossez</i>	
Improving Gradient-Guided Nested Sampling for Posterior Inference	27230
<i>Pablo Lemos, Nikolay Malkin, Will Handley, Yoshua Bengio, Yashar Hezaveh, Laurence Perreault-Levasseur</i>	
Winner-Takes-All Learners Are Geometry-Aware Conditional Density Estimators	27254
<i>Victor Letzelter, David Perera, Cédric Rommel, Mathieu Fontaine, Slim ESSID, Gaël Richard, Patrick Pérez</i>	
Online Isolation Forest	27288
<i>Filippo Leveni, Guilherme Weigert Cassales, Bernhard Pfahringer, Albert Bifet, Giacomo Boracchi</i>	
Cell2Sentence: Teaching Large Language Models the Language of Biology	27299
<i>Daniel Levine, Syed Asad Rizvi, Sacha Lévy, Nazreen Pallikkavaliyaveetil, David Zhang, Xingyu Chen, Sina Ghadermarzi, Ruiming Wu, Zihe Zheng, Ivan Vrkcic, Anna Zhong, Daphne Raskin, Insu Han, Antonio Henrique De Oliveira Fonseca, Josue Ortega Caro, Amin Karbasi, Rahul M. Dhodapkar, David Van Dijk</i>	
Eluder-Based Regret for Stochastic Contextual MDPs	27326
<i>Orin Levy, Asaf Cassel, Alon Cohen, Yishay Mansour</i>	
Feature Reuse and Scaling: Understanding Transfer Learning with Protein Language Models.....	27351
<i>Francesca-Zhoufan Li, Ava P. Amini, Yisong Yue, Kevin K. Yang, Alex X. Lu</i>	
Convergence and Complexity Guarantee for Inexact First-Order Riemannian Optimization Algorithms.....	27376
<i>Yuchen Li, Laura Balzano, Deanna Needell, Hanbaek Lyu</i>	
DetKDS: Knowledge Distillation Search for Object Detectors.....	27399
<i>Lujun Li, Yufan Bao, Peijie Dong, Chuanguang Yang, Anggeng Li, Wenhan Luo, Qifeng Liu, Wei Xue, Yike Guo</i>	
Denoising Autoregressive Representation Learning.....	27417
<i>Yazhe Li, Jorg Bornschein, Ting Chen</i>	
Purifying Quantization-Conditioned Backdoors Via Layer-Wise Activation Correction with Distribution Approximation.....	27439
<i>Boheng Li, Yishuo Cai, Jisong Cai, Yiming Li, Han Qiu, Run Wang, Tianwei Zhang</i>	
Improving Neural Logic Machines Via Failure Reflection	27457
<i>Zhiming Li, Yushi Cao, Yan Zheng, Xu Liu, Bozhi Wu, Tianlin Li, Xiufeng Xu, Junzhe Jiang, Yon Shin Teo, Shang-Wei Lin, Yang Liu</i>	
Critical Windows: Non-Asymptotic Theory for Feature Emergence in Diffusion Models	27474
<i>Marvin Li, Sitan Chen</i>	
Learning Causal Domain-Invariant Temporal Dynamics for Few-Shot Action Recognition	27499
<i>Yuke Li, Guangyi Chen, Ben Abramowitz, Stefano Anzellotti, Donglai Wei</i>	
GiLOT: Interpreting Generative Language Models Via Optimal Transport.....	27515
<i>Xuhong Li, Jiamin Chen, Yekun Chai, Haoyi Xiong</i>	
Completing Visual Objects Via Bridging Generation and Segmentation	27531
<i>Xiang Li, Yinpeng Chen, Chung-Ching Lin, Hao Chen, Kai Hu, Rita Singh, Bhiksha Raj, Lijuan Wang, Zicheng Liu</i>	

Evolving Subnetwork Training for Large Language Models	27547
<i>Hanqi Li, Lu Chen, Da Ma, Zijian Wu, Su Zhu, Kai Yu</i>	
Data Poisoning Attacks Against Conformal Prediction	27563
<i>Yangyi Li, Aobo Chen, Wei Qian, Chenxu Zhao, Divya Lidder, Mengdi Huai</i>	
ExCP: Extreme LLM Checkpoint Compression Via Weight-Momentum Joint Shrinking.....	27575
<i>Wenshuo Li, Xinghao Chen, Han Shu, Yehui Tang, Yunhe Wang</i>	
Two-Sided Competing Matching Recommendation Markets with Quota and Complementary Preferences Constraints	27589
<i>Yuantong Li, Guang Cheng, Xiaowu Dai</i>	
Full-Atom Peptide Design Based on Multi-Modal Flow Matching	27615
<i>Jiahao Li, Chaoran Cheng, Zuofan Wu, Ruihan Guo, Shitong Luo, Zhizhou Ren, Jian Peng, Jianzhu Ma</i>	

PART 34

Positive and Unlabeled Learning with Controlled Probability Boundary Fence.....	27641
<i>Changchun Li, Yuanchao Dai, Lei Feng, Ximing Li, Bing Wang, Jihong Ouyang</i>	
FightLadder: A Benchmark for Competitive Multi-Agent Reinforcement Learning	27653
<i>Wenzhe Li, Zihan Ding, Seth Karten, Chi Jin</i>	
Debiased Distribution Compression	27675
<i>Lingxiao Li, Raaz Dwivedi, Lester Mackey</i>	
Improving Context Understanding in Multimodal Large Language Models Via Multimodal Composition Learning	27732
<i>Wei Li, Hehe Fan, Yongkang Wong, Yi Yang, Mohan Kankanhalli</i>	
RL-CFR: Improving Action Abstraction for Imperfect Information Extensive-Form Games with Reinforcement Learning	27752
<i>Boning Li, Zhixuan Fang, Longbo Huang</i>	
Vague Prototype-Oriented Diffusion Model for Multi-Class Anomaly Detection	27771
<i>Yuxin Li, Yaoxuan Feng, Bo Chen, Wenchao Chen, Yubiao Wang, Xinyue Hu, Baoling Sun, Chunhui Qu, Mingyuan Zhou</i>	
Automated Statistical Model Discovery with Language Models	27791
<i>Michael Y. Li, Emily B. Fox, Noah D. Goodman</i>	
Adaptive Feature Selection for No-Reference Image Quality Assessment by Mitigating Semantic Noise Sensitivity.....	27808
<i>Xudong Li, Timin Gao, Runze Hu, Yan Zhang, Shengchuan Zhang, Xiawu Zheng, Jingyuan Zheng, Yunhang Shen, Ke Li, Yutao Liu, Pingyang Dai, Rongrong Ji</i>	
Mastering Robot Manipulation with Multimodal Prompts Through Pretraining and Multi-Task Fine-Tuning	27822
<i>Jiachen Li, Qiaozi Gao, Michael Johnston, Xiaofeng Gao, Xuehai He, Hangjie Shi, Suhaila Shakiah, Reza Ghanadan, William Yang Wang</i>	
Graph Structure Extrapolation for Out-Of-Distribution Generalization.....	27846
<i>Xiner Li, Shurui Gui, Youzhi Luo, Shuiwang Ji</i>	
Value-Evolutionary-Based Reinforcement Learning.....	27875
<i>Pengyi Li, Jianye Hao, Hongyao Tang, Yan Zheng, Fazl Barez</i>	
Image Clustering with External Guidance.....	27890
<i>Yunfan Li, Peng Hu, Dezhong Peng, Jiancheng Lv, Jianping Fan, Xi Peng</i>	
VisionGraph: Leveraging Large Multimodal Models for Graph Theory Problems in Visual Context	27903
<i>Yunxin Li, Baotian Hu, Haoyuan Shi, Wei Wang, Longyue Wang, Min Zhang</i>	

Integrating Global Context Contrast and Local Sensitivity for Blind Image Quality Assessment	27920
<i>Xudong Li, Runze Hu, Jingyuan Zheng, Yan Zhang, Shengchuan Zhang, Xiawu Zheng, Ke Li, Yunhang Shen, Yutao Liu, Pingyang Dai, Rongrong Ji</i>	
Accelerating Convergence of Score-Based Diffusion Models, Provably	27942
<i>Gen Li, Yu Huang, Timofey Efimov, Yuting Wei, Yuejie Chi, Yuxin Chen</i>	
Q-Probe: A Lightweight Approach to Reward Maximization for Language Models	27955
<i>Kenneth Li, Samy Jelassi, Hugh Zhang, Sham Kakade, Martin Wattenberg, David Brandfonbrener</i>	
Promises and Pitfalls of Generative Masked Language Modeling: Theoretical Framework and Practical Guidelines.....	27969
<i>Yuchen Li, Alexandre Kirchmeyer, Aashay Mehta, Yilong Qin, Boris Dadachev, Kishore Papineni, Sanjiv Kumar, Andrej Risteski</i>	
Visual-Text Cross Alignment: Refining the Similarity Score in Vision-Language Models.....	28018
<i>Jinhao Li, Haopeng Li, Sarah M. Erfani, Lei Feng, James Bailey, Feng Liu</i>	
Generalizing Knowledge Graph Embedding with Universal Orthogonal Parameterization	28040
<i>Rui Li, Chaozhuo Li, Yanming Shen, Zeyu Zhang, Xu Chen</i>	
Neural Collapse in Multi-Label Learning with Pick-All-Label Loss	28060
<i>Pengyu Li, Xiao Li, Yutong Wang, Qing Qu</i>	
A Differentiable Partially Observable Generalized Linear Model with Forward-Backward Message Passing.....	28095
<i>Chengrui Li, Weihan Li, Yule Wang, Anqi Wu</i>	
Multi-Region Markovian Gaussian Process: An Efficient Method to Discover Directional Communications Across Multiple Brain Regions	28112
<i>Weihan Li, Chengrui Li, Yule Wang, Anqi Wu</i>	
A Generative Approach for Treatment Effect Estimation Under Collider Bias: From an Out-Of-Distribution Perspective	28132
<i>Baohong Li, Haoxuan Li, Anpeng Wu, Minqin Zhu, Shiyuan Peng, Qingyu Cao, Kun Kuang</i>	
Learning Shadow Variable Representation for Treatment Effect Estimation Under Collider Bias	28146
<i>Baohong Li, Haoxuan Li, Ruoxuan Xiong, Anpeng Wu, Fei Wu, Kun Kuang</i>	
Configurable Mirror Descent: Towards a Unification of Decision Making	28164
<i>Pengdeng Li, Shuxin Li, Chang Yang, Xinrun Wang, Shuyue Hu, Xiao Huang, Hau Chan, Bo An</i>	
Enhancing Class-Imbalanced Learning with Pre-Trained Guidance Through Class-Conditional Knowledge Distillation.....	28204
<i>Lan Li, Xin-Chun Li, Han-Jia Ye, De-Chuan Zhan</i>	
A Neural-Guided Dynamic Symbolic Network for Exploring Mathematical Expressions from Data	28222
<i>Wenqiang Li, Weijun Li, Lina Yu, Min Wu, Linjun Sun, Jingyi Liu, Yanjie Li, Shu Wei, Yusong Deng, Meilan Hao</i>	
Cascade-CLIP: Cascaded Vision-Language Embeddings Alignment for Zero-Shot Semantic Segmentation	28243
<i>Yunheng Li, Zhong-Yu Li, Quansheng Zeng, Qibin Hou, Ming-Ming Cheng</i>	
Chain of Code: Reasoning with a Language Model-Augmented Code Emulator	28259
<i>Chengshu Li, Jacky Liang, Andy Zeng, Xinyun Chen, Karol Hausman, Dorsa Sadigh, Sergey Levine, Li Fei-Fei, Fei Xia, Brian Ichter</i>	
Preventing Model Collapse in Gaussian Process Latent Variable Models	28278
<i>Ying Li, Zhidi Lin, Feng Yin, Michael Minyi Zhang</i>	
A Theoretical Analysis of Backdoor Poisoning Attacks in Convolutional Neural Networks.....	28309
<i>Boqi Li, Weiwei Liu</i>	

Concentration Inequalities for General Functions of Heavy-Tailed Random Variables.....	28343
<i>Shaojie Li, Yong Liu</i>	
Sparse Cocktail: Every Sparse Pattern Every Sparse Ratio All at Once.....	28368
<i>Zhangheng Li, Shiwei Liu, Tianlong Chen, Ajay Kumar Jaiswal, Zhenyu Zhang, Dilin Wang, Raghuraman Krishnamoorthi, Shiyu Chang, Zhangyang Wang</i>	
Individual Contributions as Intrinsic Exploration Scaffolds for Multi-Agent Reinforcement Learning	28387
<i>Xinran Li, Zifan Liu, Shibo Chen, Jun Zhang</i>	
Learning Adaptive and View-Invariant Vision Transformer for Real-Time UAV Tracking	28403
<i>Yongxin Li, Mengyuan Liu, You Wu, Xucheng Wang, Xiangyang Yang, Shuiwang Li</i>	
PID: Prompt-Independent Data Protection Against Latent Diffusion Models.....	28421
<i>Ang Li, Yichuan Mo, Mingjie Li, Yisen Wang</i>	
A Contextual Combinatorial Bandit Approach to Negotiation.....	28448
<i>Yexin Li, Zhancun Mu, Siyuan Qi</i>	
Improving Prototypical Visual Explanations with Reward Reweighting, Reselection, and Retraining	28466
<i>Aaron J. Li, Robin Netzorg, Zhihan Cheng, Zhuoqin Zhang, Bin Yu</i>	

PART 35

Evaluating Quantized Large Language Models.....	28480
<i>Shiyao Li, Xuefei Ning, Luning Wang, Tengxuan Liu, Xiangsheng Shi, Shengen Yan, Guohao Dai, Huazhong Yang, Yu Wang</i>	
The WMDP Benchmark: Measuring and Reducing Malicious Use with Unlearning	28525
<i>Nathaniel Li, Alexander Pan, Anjali Gopal, Summer Yue, Daniel Berrios, Alice Gatti, Justin D. Li, Ann-Kathrin Dombrowski, Shashwat Goel, Gabriel Mukobi, Nathan Helm-Burger, Rassin Lababidi, Lennart Justen, Andrew B. Liu, Michael Chen, Isabelle Barrass, Oliver Zhang, Xiaoyuan Zhu, Rishub Tamirisa, Bhruhu Bharathi, Ariel Herbert-Voss, Cort B. Breuer, Andy Zou, Mantas Mazeika, Zifan Wang, Palash Oswal, Weiran Lin, Adam A. Hunt, Justin Tienken-Harder, Kevin Y. Shih, Kemper Talley, John Guan, Ian Steneker, David Campbell, Brad Jokubaitis, Steven Basart, Stephen Fitz, Ponnurangam Kumaraguru, Kallol Krishna Karmakar, Uday Tupakula, Vijay Varadharajan, Yan Shoshitaishvili, Jimmy Ba, Kevin M. Esvelt, Alexandr Wang, Dan Hendrycks</i>	
Graph Neural Network Explanations Are Fragile.....	28551
<i>Jiate Li, Meng Pang, Yun Dong, Jinyuan Jia, Binghui Wang</i>	
When Do Skills Help Reinforcement Learning? A Theoretical Analysis of Temporal Abstractions.....	28568
<i>Zhening Li, Gabriel Poesia, Armando Solar-Lezama</i>	
DiffStitch: Boosting Offline Reinforcement Learning with Diffusion-Based Trajectory Stitching.....	28597
<i>Guanghe Li, Yixiang Shan, Zhengbang Zhu, Ting Long, Weinan Zhang</i>	
Privacy Preserving Adaptive Experiment Design.....	28610
<i>Jiachun Li, Kaining Shi, David Simchi-Levi</i>	
Combining Experimental and Historical Data for Policy Evaluation.....	28630
<i>Ting Li, Chengchun Shi, Qianglin Wen, Yang Sui, Yongli Qin, Chunbo Lai, Hongtu Zhu</i>	
LoRAP: Transformer Sub-Layers Deserve Differentiated Structured Compression for Large Language Models	28657
<i>Guangyan Li, Yongqiang Tang, Wensheng Zhang</i>	
DiffFPR: Diffusion Prior for Oversampled Fourier Phase Retrieval.....	28673
<i>Ji Li, Chao Wang</i>	
Harnessing Neural Unit Dynamics for Effective and Scalable Class-Incremental Learning.....	28688
<i>Depeng Li, Tianqi Wang, Junwei Chen, Wei Dai, Zhigang Zeng</i>	

Compress Clean Signal from Noisy Raw Image: A Self-Supervised Approach.....	28706
<i>Zhihao Li, Yufei Wang, Alex Kot, Bihan Wen</i>	
VQDNA: Unleashing the Power of Vector Quantization for Multi-Species Genomic Sequence Modeling	28717
<i>Siyuan Li, Zedong Wang, Zicheng Liu, Di Wu, Cheng Tan, Jiangbin Zheng, Yufei Huang, Stan Z. Li</i>	
How Do Nonlinear Transformers Learn and Generalize in In-Context Learning?.....	28734
<i>Hongkang Li, Meng Wang, Songtao Lu, Xiaodong Cui, Pin-Yu Chen</i>	
What Improves the Generalization of Graph Transformers? A Theoretical Dive into the Self-Attention and Positional Encoding.....	28784
<i>Hongkang Li, Meng Wang, Tengfei Ma, Sijia Liu, Zaixi Zhang, Pin-Yu Chen</i>	
OODRobustBench: A Benchmark and Large-Scale Analysis of Adversarial Robustness Under Distribution Shift.....	28830
<i>Lin Li, Yifei Wang, Chawin Sitawarin, Michael Spratling</i>	
Improved Bounds for Pure Private Agnostic Learning: Item-Level and User-Level Privacy	28870
<i>Bo Li, Wei Wang, Peng Ye</i>	
Disentangled Graph Self-Supervised Learning for Out-Of-Distribution Generalization	28890
<i>Haoyang Li, Xin Wang, Zeyang Zhang, Haibo Chen, Ziwei Zhang, Wenwu Zhu</i>	
The Good, the Bad, and Why: Unveiling Emotions in Generative AI.....	28905
<i>Cheng Li, Jindong Wang, Yixuan Zhang, Kaijie Zhu, Xinyi Wang, Wenxin Hou, Jianxun Lian, Fang Luo, Qiang Yang, Xing Xie</i>	
EAGLE: Speculative Sampling Requires Rethinking Feature Uncertainty.....	28935
<i>Yuhui Li, Fangyun Wei, Chao Zhang, Hongyang Zhang</i>	
Two-Stage Shadow Inclusion Estimation: An IV Approach for Causal Inference Under Latent Confounding and Collider Bias	28949
<i>Baohong Li, Anpeng Wu, Ruoxuan Xiong, Kun Kuang</i>	
Towards Realistic Model Selection for Semi-Supervised Learning	28965
<i>Muyang Li, Xiaobo Xia, Runze Wu, Fengming Huang, Jun Yu, Bo Han, Tongliang Liu</i>	
FlashST: A Simple and Universal Prompt-Tuning Framework for Traffic Prediction.....	28978
<i>Zhonghang Li, Lianghao Xia, Yong Xu, Chao Huang</i>	
Size-Invariance Matters: Rethinking Metrics and Losses for Imbalanced Multi-Object Salient Object Detection.....	28989
<i>Feiran Li, Qianqian Xu, Shilong Bao, Zhiyong Yang, Runmin Cong, Xiaochun Cao, Qingming Huang</i>	
Q-Star Meets Scalable Posterior Sampling: Bridging Theory and Practice Via HyperAgent	29022
<i>Yingru Li, Jiawei Xu, Lei Han, Zhi-Quan Luo</i>	
Towards Efficient Deep Spiking Neural Networks Construction with Spiking Activity Based Pruning	29063
<i>Yaxin Li, Qi Xu, Jiangrong Shen, Hongming Xu, Long Chen, Gang Pan</i>	
FedBAT: Communication-Efficient Federated Learning Via Learnable Binarization	29074
<i>Shiwei Li, Wenchao Xu, Haozhao Wang, Xing Tang, Yining Qi, Shijie Xu, Weihong Luo, Yuhua Li, Xiuqiang He, Ruixuan Li</i>	
Beyond Point Prediction: Score Matching-Based Pseudolikelihood Estimation of Neural Marked Spatio-Temporal Point Process.....	29096
<i>Zichong Li, Qunzhi Xu, Zhenghao Xu, Yajun Mei, Tuo Zhao, Hongyuan Zha</i>	
Statistical Properties of Robust Satisficing.....	29112
<i>Zhiyi Li, Yunbei Xu, Ruohan Zhan</i>	

ReMax: A Simple, Effective, and Efficient Reinforcement Learning Method for Aligning Large Language Models	29128
<i>Ziniu Li, Tian Xu, Yushun Zhang, Zhihang Lin, Yang Yu, Ruoyu Sun, Zhi-Quan Luo</i>	
PDHG-Unrolled Learning-To-Optimize Method for Large-Scale Linear Programming	29164
<i>Bingheng Li, Linxin Yang, Yupeng Chen, Senmiao Wang, Haitao Mao, Qian Chen, Yao Ma, Akang Wang, Tian Ding, Jiliang Tang, Ruoyu Sun</i>	
IIANet: An Intra- And Inter-Modality Attention Network for Audio-Visual Speech Separation	29181
<i>Kai Li, Runxuan Yang, Fuchun Sun, Xiaolin Hu</i>	
KernelWarehouse: Rethinking the Design of Dynamic Convolution	29201
<i>Chao Li, Anbang Yao</i>	
GeoReasoner: Geo-Localization with Reasoning in Street Views Using a Large Vision-Language Model	29222
<i>Ling Li, Yu Ye, Bingchuan Jiang, Wei Zeng</i>	
Learning the Uncertainty Sets of Linear Control Systems Via Set Membership: A Non-Asymptotic Analysis	29234
<i>Yingying Li, Jing Yu, Lauren Conger, Taylan Kargin, Adam Wierman</i>	
Seesaw: Compensating for Nonlinear Reduction with Linear Computations for Private Inference.....	29266
<i>Fabing Li, Yuanhao Zhai, Shuangyu Cai, Mingyu Gao</i>	
Agnostic Interactive Imitation Learning: New Theory and Practical Algorithms	29278
<i>Yichen Li, Chicheng Zhang</i>	

PART 36

Towards Optimal Adversarial Robust Q-Learning with Bellman Infinity-Error	29324
<i>Haoran Li, Zicheng Zhang, Wang Luo, Congying Han, Yudong Hu, Tiande Guo, Shichen Liao</i>	
Predicting and Interpreting Energy Barriers of Metallic Glasses with Graph Neural Networks	29373
<i>Haoyu Li, Shichang Zhang, Longwen Tang, Mathieu Bauchy, Yizhou Sun</i>	
From Fourier to Neural ODEs: Flow Matching for Modeling Complex Systems.....	29390
<i>Xin Li, Jingdong Zhang, Qunxi Zhu, Chengli Zhao, Xue Zhang, Xiaojun Duan, Wei Lin</i>	
Feel-Good Thompson Sampling for Contextual Dueling Bandits	29406
<i>Xuheng Li, Heyang Zhao, Quanquan Gu</i>	
EvoRainbow: Combining Improvements in Evolutionary Reinforcement Learning for Policy Search	29427
<i>Pengyi Li, Yan Zheng, Hongyao Tang, Xian Fu, Jianye Hao</i>	
Relaxing the Accurate Imputation Assumption in Doubly Robust Learningfor Debiased Collaborative Filtering.....	29448
<i>Haoxuan Li, Chunyuan Zheng, Shuyi Wang, Kunhan Wu, Hao Wang, Peng Wu, Zhi Geng, Xu Chen, Xiao-Hua Zhou</i>	
DecisionNCE: Embodied Multimodal Representations Via Implicit Preference Learning	29461
<i>Jianxiong Li, Jinliang Zheng, Yinan Zheng, Liyuan Mao, Xiao Hu, Sijie Cheng, Haoyi Niu, Jihao Liu, Yu Liu, Jingjing Liu, Ya-Qin Zhang, Xianyuan Zhan</i>	
Algorithmic Stability Unleashed: Generalization Bounds with Unbounded Losses	29489
<i>Shaojie Li, Bowei Zhu, Yong Liu</i>	
Kepler Codebook.....	29511
<i>Junrong Lian, Ziyue Dong, Pengxu Wei, Wei Ke, Chang Liu, Qixiang Ye, Xiangyang Ji, Liang Lin</i>	
Receptive Fields as Experts in Convolutional Neural Architectures	29531
<i>Dongze Lian, Weihao Yu, Xinchao Wang</i>	

Diving into Underwater: Segment Anything Model Guided Underwater Salient Instance Segmentation and a Large-Scale Dataset.....	29545
<i>Shijie Lian, Ziyi Zhang, Hua Li, Wenjie Li, Laurence Tianruo Yang, Sam Kwong, Runmin Cong</i>	
Graph External Attention Enhanced Transformer	29560
<i>Jianqing Liang, Min Chen, Jiye Liang</i>	
Monitoring AI-Modified Content at Scale: A Case Study on the Impact of ChatGPT on AI Conference Peer Reviews.....	29575
<i>Weixin Liang, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, Zhi Huang, Daniel A. McFarland, James Y. Zou</i>	
Sign is Not a Remedy: Multiset-To-Multiset Message Passing for Learning on Heterophilic Graphs	29621
<i>Langzhang Liang, Sunwoo Kim, Kijung Shin, Zenglin Xu, Shirui Pan, Yuan Qi</i>	
Single-Trajectory Distributionally Robust Reinforcement Learning.....	29644
<i>Zhipeng Liang, Xiaoteng Ma, Jose Blanchet, Jun Yang, Jiheng Zhang, Zhengyuan Zhou</i>	
Realistic Unsupervised CLIP Fine-Tuning with Universal Entropy Optimization.....	29667
<i>Jian Liang, Lijun Sheng, Zhengbo Wang, Ran He, Tieniu Tan</i>	
Efficient Precision and Recall Metrics for Assessing Generative Models Using Hubness-Aware Sampling.....	29682
<i>Yuanbang Liang, Jing Wu, Yu-Kun Lai, Yipeng Qin</i>	
Scalable Multiple Kernel Clustering: Learning Clustering Structure from Expectation	29700
<i>Weixuan Liang, En Zhu, Shengju Yu, Huiying Xu, Xinzhong Zhu, Xinwang Liu</i>	
On the Error-Propagation of Inexact Hotelling's Deflation for Principal Component Analysis	29720
<i>Fangshuo Liao, Junhyung Lyle Kim, Cruz Barnum, Anastasios Kyrillidis</i>	
Bootstrapping Fisher Market Equilibrium and First-Price Pacing Equilibrium	29748
<i>Luofeng Liao, Christian Kroer</i>	
Enhancing Value Function Estimation Through First-Order State-Action Dynamics in Offline Reinforcement Learning.....	29782
<i>Yun-Hsuan Lien, Ping-Chun Hsieh, Tzu-Mao Li, Yu-Shuen Wang</i>	
Graph Geometry-Preserving Autoencoders.....	29795
<i>Junbin Lim, Jihwan Kim, Yonghyeon Lee, Cheongjae Jang, Frank Chongwoo Park</i>	
Momentum Particle Maximum Likelihood	29816
<i>Jen Ning Lim, Juan Kuntz, Samuel Power, Adam M. Johansen</i>	
An Effective Dynamic Gradient Calibration Method for Continual Learning	29872
<i>Weichen Lin, Jiayang Chen, Ruomin Huang, Hu Ding</i>	
Equivariant Diffusion for Crystal Structure Prediction	29890
<i>Peijia Lin, Pin Chen, Rui Jiao, Qing Mo, Jianhuan Cen, Wenbing Huang, Yang Liu, Dan Huang, Yutong Lu</i>	
Revisiting the Role of Language Priors in Vision-Language Models.....	29914
<i>Zhiqiu Lin, Xinyue Chen, Deepak Pathak, Pengchuan Zhang, Deva Ramanan</i>	
Non-Confusing Generation of Customized Concepts in Diffusion Models	29935
<i>Wang Lin, Jingyuan Chen, Jiaxin Shi, Yichen Zhu, Chen Liang, Junzhong Miao, Tao Jin, Zhou Zhao, Fei Wu, Shuicheng Yan, Hanwang Zhang</i>	
Can We Remove the Square-Root in Adaptive Gradient Methods?	29949
<i>Wu Lin, Felix Dangel, Runa Eschenhagen, Juhan Bae, Richard E. Turner, Alireza Makhzani</i>	
Structured Inverse-Free Natural Gradient Descent: Memory-Efficient & Numerically-Stable KFAC	29974
<i>Wu Lin, Felix Dangel, Runa Eschenhagen, Kirill Neklyudov, Agustinus Kristiadi, Richard E. Turner, Alireza Makhzani</i>	

Learning to Model the World with Language.....	29992
<i>Jessy Lin, Yuqing Du, Olivia Watkins, Danijar Hafner, Pieter Abbeel, Dan Klein, Anca Dragan</i>	
Robustness of Deep Learning for Accelerated MRI: Benefits of Diverse Training Data.....	30018
<i>Kang Lin, Reinhard Heckel</i>	
Equivariance Via Minimal Frame Averaging for More Symmetries and Efficiency.....	30042
<i>Yuchao Lin, Jacob Helwig, Shurui Gui, Shuiwang Ji</i>	
Selecting Large Language Model to Fine-Tune Via Rectified Scaling Law.....	30080
<i>Haowei Lin, Baizhou Huang, Haotian Ye, Qinyu Chen, Zihao Wang, Sujian Li, Jianzhu Ma, Xiaojun Wan, James Zou, Yitao Liang</i>	
Graph-Enhanced Large Language Models in Asynchronous Plan Reasoning.....	30108
<i>Fangru Lin, Emanuele La Malfa, Valentin Hofmann, Elle Michelle Yang, Anthony G. Cohn, Janet B. Pierrehumbert</i>	
Dual Operating Modes of In-Context Learning.....	30135
<i>Ziqian Lin, Kangwook Lee</i>	

PART 37

HGAP: Boosting Permutation Invariant and Permutation Equivariant in Multi-Agent Reinforcement Learning Via Graph Attention Network.....	30189
<i>Bor Jiun Lin, Chun-Yi Lee</i>	
SparseTSF: Modeling Long-Term Time Series Forecasting with 1k Parameters.....	30211
<i>Shengsheng Lin, Weiwei Lin, Wentai Wu, Haojun Chen, Junjie Yang</i>	
Fast and Sample Efficient Multi-Task Representation Learning in Stochastic Contextual Bandits.....	30227
<i>Jiabin Lin, Shana Moothedath, Namrata Vaswani</i>	
On Hypothesis Transfer Learning of Functional Linear Models.....	30252
<i>Haotian Lin, Matthew Reimherr</i>	
Smoothness Adaptive Hypothesis Transfer Learning.....	30286
<i>Haotian Lin, Matthew Reimherr</i>	
Use Your INSTINCT: INSTRUCTION Optimization for LLMs using Neural Bandits Coupled with Transformers.....	30317
<i>Xiaoqiang Lin, Zhaoxuan Wu, Zhongxiang Dai, Wenyang Hu, Yao Shu, See-Kiong Ng, Patrick Jaillet, Bryan Kian Hsiang Low</i>	
GeoAB: Towards Realistic Antibody Design and Reliable Affinity Maturation.....	30346
<i>Haitao Lin, Lirong Wu, Yufei Huang, Yunfan Liu, Odin Zhang, Yuanqing Zhou, Rui Sun, Stan Z. Li</i>	
Distributionally Robust Data Valuation.....	30362
<i>Xiaoqiang Lin, Xinyi Xu, Zhaoxuan Wu, See-Kiong Ng, Bryan Kian Hsiang Low</i>	
A Single-Loop Robust Policy Gradient Method for Robust Markov Decision Processes.....	30392
<i>Zhenwei Lin, Chenyu Xue, Qi Deng, Yinyu Ye</i>	
Layer-Aware Analysis of Catastrophic Overfitting: Revealing the Pseudo-Robust Shortcut Dependency.....	30427
<i>Runqi Lin, Chaojian Yu, Bo Han, Hang Su, Tongliang Liu</i>	
Autonomous Sparse Mean-CVaR Portfolio Optimization.....	30440
<i>Yizun Lin, Yangyu Zhang, Zhao-Rong Lai, Cheng Li</i>	
Graph Neural Stochastic Diffusion for Estimating Uncertainty in Node Classification.....	30457
<i>Xixun Lin, Wenxiao Zhang, Fengzhao Shi, Chuan Zhou, Lixin Zou, Xiangyu Zhao, Dawei Yin, Shirui Pan, Yanan Cao</i>	

Smooth Tchebycheff Scalarization for Multi-Objective Optimization	30479
<i>Xi Lin, Xiaoyuan Zhang, Zhiyuan Yang, Fei Liu, Zhenkun Wang, Qingfu Zhang</i>	
PPFlow: Target-Aware Peptide Design with Torsional Flow Matching	30510
<i>Haitao Lin, Odin Zhang, Hui Feng Zhao, Dejun Jiang, Lirong Wu, Zicheng Liu, Yufei Huang, Stan Z. Li</i>	
Lie Neurons: Adjoint-Equivariant Neural Networks for Semisimple Lie Algebras	30529
<i>Tzu-Yuan Lin, Minghan Zhu, Maani Ghaffari</i>	
Plug-In Performative Optimization	30546
<i>Licong Lin, Tijana Zrnic</i>	
Position: A Call to Action for a Human-Centered AutoML Paradigm.....	30566
<i>Marius Lindauer, Florian Karl, Anne Klier, Julia Moosbauer, Alexander Tornede, Andreas Mueller, Frank Hutter, Matthias Feurer, Bernd Bischl</i>	
Deep Equilibrium Models Are Almost Equivalent to Not-So-Deep Explicit Models for High-Dimensional Gaussian Mixtures.....	30585
<i>Zenan Ling, Longbo Li, Zhanbo Feng, Yixuan Zhang, Feng Zhou, Robert C. Qiu, Zhenyu Liao</i>	
Beyond Regular Grids: Fourier-Based Neural Operators on Arbitrary Domains	30610
<i>Levi Lingsch, Mike Y. Michelis, Emmanuel De Bézenac, Sirani M. Perera, Robert K. Katzschmann, Siddhartha Mishra</i>	
Scaling Tractable Probabilistic Circuits: A Systems Perspective	30630
<i>Anji Liu, Kareem Ahmed, Guy Van Den Broeck</i>	
Causal Bandits: The Pareto Optimal Frontier of Adaptivity, a Reduction to Linear Bandits, and Limitations Around Unknown Marginals	30647
<i>Ziyi Liu, Idan Attias, Daniel M. Roy</i>	
Orthogonal Bootstrap: Efficient Simulation of Input Uncertainty.....	30669
<i>Kaizhao Liu, Jose Blanchet, Lexing Ying, Yiping Lu</i>	
Graph Distillation with Eigenbasis Matching.....	30702
<i>Yang Liu, Deyu Bo, Chuan Shi</i>	
Adaptive Text Watermark for Large Language Models.....	30718
<i>Yepeng Liu, Yuheng Bu</i>	
Enhancing Adversarial Robustness in SNNs with Sparse Gradients.....	30738
<i>Yujia Liu, Tong Bu, Jianhao Ding, Zecheng Hao, Tiejun Huang, Zhaofei Yu</i>	
ReLU Network with Width $d+O(1)$ Can Achieve Optimal Approximation Rate	30755
<i>Chenghao Liu, Minghua Chen</i>	
Graph Adversarial Diffusion Convolution.....	30789
<i>Songtao Liu, Jinghui Chen, Tianfan Fu, Lu Lin, Marinka Zitnik, Dinghao Wu</i>	
Decentralized Convex Finite-Sum Optimization with Better Dependence on Condition Numbers	30807
<i>Yuxing Liu, Lesi Chen, Luo Luo</i>	
Zeroth-Order Methods for Constrained Nonconvex Nonsmooth Stochastic Optimization.....	30842
<i>Zhuanghua Liu, Cheng Chen, Luo Luo, Bryan Kian Hsiang Low</i>	
Unifying Image Processing as Visual Prompting Question Answering.....	30873
<i>Yihao Liu, Xiangyu Chen, Xianzheng Ma, Xintao Wang, Jiantao Zhou, Yu Qiao, Chao Dong</i>	
ESNet: Evolution and Succession Network for High-Resolution Salient Object Detection	30892
<i>Hongyu Liu, Runmin Cong, Hua Li, Qianqian Xu, Qingming Huang, Wei Zhang</i>	
The Pitfalls and Promise of Conformal Inference Under Adversarial Attacks	30908
<i>Ziquan Liu, Yufei Cui, Yan Yan, Yi Xu, Xiangyang Ji, Xue Liu, Antoni B. Chan</i>	
Preference Optimization for Molecule Synthesis with Conditional Residual Energy-Based Models	30929
<i>Songtao Liu, Hanjun Dai, Yue Zhao, Peng Liu</i>	

PEARL: Zero-Shot Cross-Task Preference Alignment and Robust Reward Learning for Robotic Manipulation	30946
<i>Runze Liu, Yali Du, Fengshuo Bai, Jiafei Lyu, Xiu Li</i>	
Harnessing the Power of Neural Operators with Automatically Encoded Conservation Laws	30965
<i>Ning Liu, Yiming Fan, Xianyi Zeng, Milan Klöwer, Lu Zhang, Yue Yu</i>	
Mitigating Privacy Risk in Membership Inference by Convex-Concave Loss.....	30998
<i>Zhenlong Liu, Lei Feng, Huiping Zhuang, Xiaofeng Cao, Hongxin Wei</i>	

PART 38

Decoding-Time Realignment of Language Models.....	31015
<i>Tianlin Liu, Shangmin Guo, Leonardo Bianco, Daniele Calandriello, Quentin Berthet, Felipe Llinares, Jessica Hoffmann, Lucas Dixon, Michal Valko, Mathieu Blondel</i>	
DIDI: Diffusion-Guided Diversity for Offline Behavioral Generation	31032
<i>Jinxin Liu, Xinghong Guo, Zifeng Zhuang, Donglin Wang</i>	
Bidirectional Reciprocal Information Communication for Few-Shot Semantic Segmentation.....	31048
<i>Yuanwei Liu, Junwei Han, Xiwen Yao, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, Nian Liu, Fahad Shahbaz Khan</i>	
Unlock the Cognitive Generalization of Deep Reinforcement Learning Via Granular Ball Representation.....	31062
<i>Jiashun Liu, Jianye Hao, Yi Ma, Shuyin Xia</i>	
PAPM: A Physics-Aware Proxy Model for Process Systems	31080
<i>Pengwei Liu, Zhongkai Hao, Xingyu Ren, Hangjie Yuan, Jiayang Ren, Dong Ni</i>	
ELTA: An Enhancer Against Long-Tail for Aesthetics-Oriented Models.....	31106
<i>Limin Liu, Shuai He, Anlong Ming, Rui Xie, Huadong Ma</i>	
On the Feasibility of Single-Pass Full-Capacity Learning in Linear Threshold Neurons with Binary Input Vectors.....	31119
<i>Ruipeng Liu, Borui He, Naveed Tahir, Garrett Ethan Katz</i>	
Online Speculative Decoding	31131
<i>Xiaoxuan Liu, Lanxiang Hu, Peter Bailis, Alvin Cheung, Zhijie Deng, Ion Stoica, Hao Zhang</i>	
Tuning-Free Estimation and Inference of Cumulative Distribution Function Under Local Differential Privacy	31147
<i>Yi Liu, Qirui Hu, Linglong Kong</i>	
Referee Can Play: An Alternative Approach to Conditional Generation Via Model Inversion.....	31165
<i>Xuantong Liu, Tianyang Hu, Wenjia Wang, Kenji Kawaguchi, Yuan Yao</i>	
Reason for Future, Act for Now: A Principled Architecture for Autonomous LLM Agents.....	31186
<i>Zhihan Liu, Hao Hu, Shenao Zhang, Hongyi Guo, Shuqi Ke, Boyi Liu, Zhaoran Wang</i>	
Weakly-Supervised Residual Evidential Learning for Multi-Instance Uncertainty Estimation	31262
<i>Pei Liu, Luping Ji</i>	
StyDeSty: Min-Max Stylization and Destylization for Single Domain Generalization	31293
<i>Songhua Liu, Xin Jin, Xingyi Yang, Jingwen Ye, Xinchao Wang</i>	
Time-Series Forecasting for Out-Of-Distribution Generalization Using Invariant Learning.....	31312
<i>Haoxin Liu, Harshavardhan Kamarthi, Lingkai Kong, Zhiyuan Zhao, Chao Zhang, B. Aditya Prakash</i>	
Stereo Risk: A Continuous Modeling Approach to Stereo Matching	31326
<i>Ce Liu, Suryansh Kumar, Shuhang Gu, Radu Timofte, Yao Yao, Luc Van Gool</i>	
Multi-Source Conformal Inference Under Distribution Shift.....	31344
<i>Yi Liu, Alexander W. Levis, Sharon-Lise Normand, Larry Han</i>	

From Generalization Analysis to Optimization Designs for State Space Models	31383
<i>Fusheng Liu, Qianxiao Li</i>	
Adaptive Advantage-Guided Policy Regularization for Offline Reinforcement Learning.....	31406
<i>Tenglong Liu, Yang Li, Yixing Lan, Hao Gao, Wei Pan, Xin Xu</i>	
Two-Timescale Derivative Free Optimization for Performative Prediction with Markovian Data	31425
<i>Haitong Liu, Qiang Li, Hoi-To Wai</i>	
Short-Long Convolutions Help Hardware-Efficient Linear Attention to Focus on Long Sequences.....	31451
<i>Zicheng Liu, Siyuan Li, Li Wang, Zedong Wang, Yunfan Liu, Stan Z. Li</i>	
Differentiable Combinatorial Scheduling at Scale	31464
<i>Mingju Liu, Yingjie Li, Jiaqi Yin, Zhiru Zhang, Cunxi Yu</i>	
Characterizing ResNet's Universal Approximation Capability.....	31477
<i>Chenghao Liu, Enming Liang, Minghua Chen</i>	
Convergence of Online Learning Algorithm for a Mixture of Multiple Linear Regressions.....	31516
<i>Yujing Liu, Zhixin Liu, Lei Guo</i>	
Energy-Guided Diffusion Sampling for Offline-To-Online Reinforcement Learning.....	31541
<i>Xu-Hui Liu, Tian-Shuo Liu, Shengyi Jiang, Ruifeng Chen, Zhilong Zhang, Xinwei Chen, Yang Yu</i>	
Moreau Envelope for Nonconvex Bi-Level Optimization: A Single-Loop and Hessian-Free Solution Strategy	31566
<i>Risheng Liu, Zhu Liu, Wei Yao, Shangzhi Zeng, Jin Zhang</i>	
Position: Foundation Agents as the Paradigm Shift for Decision Making	31597
<i>Xiaoqian Liu, Xingzhou Lou, Jianbin Jiao, Junge Zhang</i>	
Learning with Partial-Label and Unlabeled Data: A Uniform Treatment for Supervision Redundancy and Insufficiency	31614
<i>Yangfan Liu, Jiaqi Lv, Xin Geng, Ning Xu</i>	
Revisiting Context Aggregation for Image Matting	31629
<i>Qinglin Liu, Xiaoqian Lv, Quanling Meng, Zonglin Li, Xiangyuan Lan, Shuo Yang, Shengping Zhang, Liqiang Nie</i>	
Amortized Equation Discovery in Hybrid Dynamical Systems	31645
<i>Yongtuo Liu, Sara Magliacane, Miltiadis Kofinas, Efstratios Gavves</i>	
KnowFormer: Revisiting Transformers for Knowledge Graph Reasoning	31669
<i>Junnan Liu, Qianren Mao, Weifeng Jiang, Jianxin Li</i>	
Floating Anchor Diffusion Model for Multi-Motif Scaffolding	31691
<i>Ke Liu, Weian Mao, Shuaike Shen, Xiaoran Jiao, Zheng Sun, Hao Chen, Chunhua Shen</i>	
Deep Functional Factor Models: Forecasting High-Dimensional Functional Time Series Via Bayesian Nonparametric Factorization.....	31709
<i>Yirui Liu, Xinghao Qiao, Yulong Pei, Liying Wang</i>	
Fast Decision Boundary Based Out-Of-Distribution Detector	31728
<i>Litian Liu, Yao Qin</i>	
Class-Imbalanced Graph Learning Without Class Rebalancing.....	31747
<i>Zhining Liu, Ruizhong Qiu, Zhichen Zeng, Hyunsik Yoo, David Zhou, Zhe Xu, Yada Zhu, Kommy Weldemariam, Jingrui He, Hanghang Tong</i>	
Generative Marginalization Models	31773
<i>Sulin Liu, Peter J. Ramadge, Ryan P. Adams</i>	
Federated Representation Learning in the Under-Parameterized Regime.....	31808
<i>Renpu Liu, Cong Shen, Jing Yang</i>	

PART 39

How Do Large Language Models Navigate Conflicts Between Honesty and Helpfulness?	31844
<i>Ryan Liu, Theodore R. Sumers, Ishita Dasgupta, Thomas L. Griffiths</i>	
Causal Discovery Via Conditional Independence Testing with Proxy Variables.....	31866
<i>Mingzhou Liu, Xinwei Sun, Yu Qiao, Yizhou Wang</i>	
Perfect Alignment May Be Poisonous to Graph Contrastive Learning	31890
<i>Jingyu Liu, Huayi Tang, Yong Liu</i>	
Entropy-Reinforced Planning with Large Language Models for Drug Discovery	31917
<i>Xuefeng Liu, Chih-Chan Tien, Peng Ding, Songhao Jiang, Rick L. Stevens</i>	
New Sample Complexity Bounds for Sample Average Approximation in Heavy-Tailed Stochastic Programming.....	31933
<i>Hongcheng Liu, Jindong Tong</i>	
Zero-Shot ECG Classification with Multimodal Learning and Test-Time Clinical Knowledge Enhancement	31949
<i>Che Liu, Zhongwei Wan, Cheng Ouyang, Anand Shah, Wenjia Bai, Rossella Arcucci</i>	
Unified Generation, Reconstruction, and Representation: Generalized Diffusion with Adaptive Latent Encoding-Decoding.....	31964
<i>Guangyi Liu, Yu Wang, Zeyu Feng, Qiyu Wu, Liping Tang, Yuan Gao, Zhen Li, Shuguang Cui, Julian McAuley, Zichao Yang, Eric P. Xing, Zhiting Hu</i>	
Differentiable Model Scaling Using Differentiable Topk.....	31994
<i>Kai Liu, Ruohui Wang, Jianfei Gao, Kai Chen</i>	
Symmetric Matrix Completion with ReLU Sampling	32015
<i>Huikang Liu, Peng Wang, Longxiu Huang, Qing Qu, Laura Balzano</i>	
DNA-SE: Towards Deep Neural-Nets Assisted Semiparametric Estimation	32041
<i>Qinshuo Liu, Zixin Wang, Xi-An Li, Xinyao Ji, Lei Zhang, Lin Liu, Zhonghua Liu</i>	
TimeX++: Learning Time-Series Explanations with Information Bottleneck.....	32062
<i>Zichuan Liu, Tianchun Wang, Jimeng Shi, Xu Zheng, Zhuomin Chen, Lei Song, Wenqian Dong, Jayantha Obeysekera, Farhad Shirani, Dongsheng Luo</i>	
LIDAO: Towards Limited Interventions for Debiasing (Large) Language Models	32083
<i>Tianci Liu, Haoyu Wang, Shiyang Wang, Yu Cheng, Jing Gao</i>	
DoRA: Weight-Decomposed Low-Rank Adaptation.....	32100
<i>Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, Min-Hung Chen</i>	
High-Probability Bound for Non-Smooth Non-Convex Stochastic Optimization with Heavy Tails	32122
<i>Langqi Liu, Yibo Wang, Lijun Zhang</i>	
Combinatorial Multivariate Multi-Armed Bandits with Applications to Episodic Reinforcement Learning and Beyond	32139
<i>Xutong Liu, Siwei Wang, Jinhang Zuo, Han Zhong, Xuchuang Wang, Zhiyong Wang, Shuai Li, Mohammad Hajiesmaili, John C. S. Lui, Wei Chen</i>	
Language-Driven Cross-Modal Classifier for Zero-Shot Multi-Label Image Recognition.....	32173
<i>Yicheng Liu, Jie Wen, Chengliang Liu, Xiaozhao Fang, Zuoyong Li, Yong Xu, Zheng Zhang</i>	
Geometry-Calibrated DRO: Combating Over-Pessimism with Free Energy Implications.....	32184
<i>Jiashuo Liu, Jiayun Wu, Tianyu Wang, Hao Zou, Bo Li, Peng Cui</i>	

Evolution of Heuristics: Towards Efficient Automatic Algorithm Design Using Large Language Model	32201
<i>Fei Liu, Xialiang Tong, Mingxuan Yuan, Xi Lin, Fu Luo, Zhenkun Wang, Zhichao Lu, Qingfu Zhang</i>	
Correlation-Induced Label Prior for Semi-Supervised Multi-Label Learning	32224
<i>Biao Liu, Ning Xu, Xiangyu Fang, Xin Geng</i>	
Causality Based Front-Door Defense Against Backdoor Attack on Language Models.....	32239
<i>Yiran Liu, Xiaoang Xu, Zhiyi Hou, Yang Yu</i>	
Partial Multi-View Multi-Label Classification Via Semantic Invariance Learning and Prototype Modeling	32253
<i>Chengliang Liu, Gehui Xu, Jie Wen, Yabo Liu, Chao Huang, Yong Xu</i>	
Building Socially-Equitable Public Models	32268
<i>Yejia Liu, Jianyi Yang, Pengfei Li, Tongxin Li, Shaolei Ren</i>	
In-Context Vectors: Making in Context Learning More Effective and Controllable Through Latent Space Steering.....	32287
<i>Sheng Liu, Haotian Ye, Lei Xing, James Zou</i>	
Minimizing f-Divergences by Interpolating Velocity Fields	32308
<i>Song Liu, Jiahao Yu, Jack Simons, Mingxuan Yi, Mark Beaumont</i>	
KIVI: A Tuning-Free Asymmetric 2bit Quantization for KV Cache	32332
<i>Zirui Liu, Jiayi Yuan, Hongye Jin, Shaochen (Henry) Zhong, Zhaozhuo Xu, Vladimir Braverman, Beidi Chen, Xia Hu</i>	
Efficient Policy Evaluation with Offline Data Informed Behavior Policy Design	32345
<i>Shuze Liu, Shangtong Zhang</i>	
Timer: Generative Pre-Trained Transformers Are Large Time Series Models	32369
<i>Yong Liu, Haoran Zhang, Chenyu Li, Xiangdong Huang, Jianmin Wang, Mingsheng Long</i>	
SPHINX-X: Scaling Data and Parameters for a Family of Multi-Modal Large Language Models	32400
<i>Dongyang Liu, Renrui Zhang, Longtian Qiu, Siyuan Huang, Weifeng Lin, Shitian Zhao, Shijie Geng, Ziyi Lin, Peng Jin, Kaipeng Zhang, Wenqi Shao, Chao Xu, Conghui He, Junjun He, Hao Shao, Pan Lu, Yu Qiao, Hongsheng Li, Peng Gao</i>	
DFD: Distilling the Feature Disparity Differently for Detectors.....	32421
<i>Kang Liu, Yingyi Zhang, Jingyun Zhang, Jinmin Li, Jun Wang, Shaoming Wang, Chun Yuan, Rizen Guo</i>	
MobileLLM: Optimizing Sub-Billion Parameter Language Models for On-Device Use Cases	32431
<i>Zechun Liu, Changsheng Zhao, Forrest Iandola, Chen Lai, Yuandong Tian, Igor Fedorov, Yunyang Xiong, Ernie Chang, Yangyang Shi, Raghuraman Krishnamoorthi, Liangzhen Lai, Vikas Chandra</i>	
Reducing Item Discrepancy Via Differentially Private Robust Embedding Alignment for Privacy-Preserving Cross Domain Recommendation	32455
<i>Weiming Liu, Xiaolin Zheng, Chaochao Chen, Jiahe Xu, Xinting Liao, Fan Wang, Yanchao Tan, Yew-Soon Ong</i>	
On the Last-Iterate Convergence of Shuffling Gradient Methods	32471
<i>Zijian Liu, Zhengyuan Zhou</i>	
FedLMT: Tackling System Heterogeneity of Federated Learning Via Low-Rank Model Training with Theoretical Guarantees.....	32509
<i>Jiahao Liu, Yipeng Zhou, Di Wu, Miao Hu, Mohsen Guizani, Quan Z. Sheng</i>	
Pairwise Alignment Improves Graph Domain Adaptation	32552
<i>Shikun Liu, Deyu Zou, Han Zhao, Pan Li</i>	

Neural Operators with Localized Integral and Differential Kernels.....	32576
<i>Miguel Liu-Schiaffini, Julius Berner, Boris Bonev, Thorsten Kurth, Kamyar Azizzadenesheli, Anima Anandkumar</i>	
A Tensor Decomposition Perspective on Second-Order RNNs.....	32595
<i>Maude Lizaire, Michael Rizvi-Martel, Marawan Gamal Abdel Hameed, Guillaume Rabusseau</i>	
Coarse-To-Fine Tensor Trains for Compact Visual Representations	32612
<i>Sebastian Loeschcke, Dan Wang, Christian Leth-Espensen, Serge Belongie, Michael J. Kastoryano, Sagie Benaim</i>	
Restoring Balance: Principled Under/Oversampling of Data for Optimal Classification	32643
<i>Emanuele Loffredo, Mauro Pastore, Simona Cocco, Rémi Monasson</i>	
Position: The Causal Revolution Needs Scientific Pragmatism	32671
<i>Joshua R. Loftus</i>	
Reparameterized Importance Sampling for Robust Variational Bayesian Neural Networks.....	32680
<i>Yunfei Long, Zilin Tian, Liguozhang, Huosheng Xu</i>	

PART 40

Position: A Safe Harbor for AI Evaluation and Red Teaming	32691
<i>Shayne Longpre, Sayash Kapoor, Kevin Klyman, Ashwin Ramaswami, Rishi Bommasani, Borhane Blili-Hamelin, Yangsibo Huang, Aviya Skowron, Zheng-Xin Yong, Suhas Kotha, Yi Zeng, Weiyang Shi, Xianjun Yang, Reid Southen, Alexander Robey, Patrick Chao, Diyi Yang, Ruoxi Jia, Daniel Kang, Sandy Pentland, Arvind Narayanan, Percy Liang, Peter Henderson</i>	
Position: Data Authenticity, Consent, & Provenance for AI Are All Broken: What Will it Take to Fix Them?.....	32711
<i>Shayne Longpre, Robert Mahari, Naana Obeng-Marnu, William Brannon, Tobin South, Katy Gero, Sandy Pentland, Jad Kabbara</i>	
Latent Noise Segmentation: How Neural Noise Leads to the Emergence of Segmentation and Grouping.....	32726
<i>Ben Lonnqvist, Zhengqing Wu, Michael H. Herzog</i>	
Large Scale Dataset Distillation with Domain Shift.....	32759
<i>Noel Loo, Alaa Maalouf, Ramin Hasani, Mathias Lechner, Alexander Amini, Daniel Rus</i>	
Attention Meets Post-Hoc Interpretability: A Mathematical Perspective	32781
<i>Gianluigi Lopardo, Frederic Precioso, Damien Garreau</i>	
Non-Vacuous Generalization Bounds for Large Language Models	32801
<i>Sanae Lotfi, Marc Finzi, Yilun Kuang, Tim G. J. Rudner, Micah Goldblum, Andrew Gordon Wilson</i>	
Discrete Diffusion Modeling by Estimating the Ratios of the Data Distribution	32819
<i>Aaron Lou, Chenlin Meng, Stefano Ermon</i>	
Optimal Differentially Private Model Training with Public Data.....	32849
<i>Andrew Lowy, Zeman Li, Tianjian Huang, Meisam Razaviyayn</i>	
How to Make the Gradients Small Privately: Improved Rates for Differentially Private Non-Convex Optimization.....	32904
<i>Andrew Lowy, Jonathan Ullman, Stephen J. Wright</i>	
Beyond Sole Strength: Customized Ensembles for Generalized Vision-Language Models.....	32924
<i>Zhihe Lu, Jiawang Bai, Xin Li, Zeyu Xiao, Xinchao Wang</i>	
HumanTOMATO: Text-Aligned Whole-Body Motion Generation.....	32939
<i>Shunlin Lu, Ling-Hao Chen, Ailing Zeng, Jing Lin, Ruimao Zhang, Lei Zhang, Heung-Yeung Shum</i>	

Position: Exploring the Robustness of Pipeline-Parallelism-Based Decentralized Training.....	32978
<i>Lin Lu, Chenxi Dai, Wangcheng Tao, Binhang Yuan, Yanan Sun, Pan Zhou</i>	
CATS: Enhancing Multivariate Time Series Forecasting by Constructing Auxiliary Time Series as Exogenous Variables	32990
<i>Jiecheng Lu, Xu Han, Yan Sun, Shihao Yang</i>	
WebLINX: Real-World Website Navigation with Multi-Turn Dialogue	33007
<i>Xing Han Lu, Zdenek Kasner, Siva Reddy</i>	
Open-Domain Text Evaluation Via Contrastive Distribution Methods	33057
<i>Sidi Lu, Hongyi Liu, Asli Celikyilmaz, Tianlu Wang, Nanyun Peng</i>	
EiG-Search: Generating Edge-Induced Subgraphs for GNN Explanation in Linear Time.....	33069
<i>Shengyao Lu, Bang Liu, Keith G. Mills, Jiao He, Di Niu</i>	
Rethinking Transformers in Solving POMDPs	33089
<i>Chenhao Lu, Ruizhe Shi, Yuyao Liu, Kaizhe Hu, Simon S. Du, Huazhe Xu</i>	
CauDiTS: Causal Disentangled Domain Adaptation of Multivariate Time Series	33113
<i>Junxin Lu, Shiliang Sun</i>	
NeWRF: A Deep Learning Framework for Wireless Radiation Field Reconstruction and Channel Prediction	33147
<i>Haofan Lu, Christopher Vatheuer, Baharan Mirzsoleiman, Omid Abari</i>	
FiT: Flexible Vision Transformer for Diffusion Model	33160
<i>Zeyu Lu, Zidong Wang, Di Huang, Chengyue Wu, Xihui Liu, Wanli Ouyang, Lei Bai</i>	
Probabilistic Routing for Graph-Based Approximate Nearest Neighbor Search.....	33177
<i>Kejing Lu, Chuan Xiao, Yoshiharu Ishikawa</i>	
Disguised Copyright Infringement of Latent Diffusion Models.....	33196
<i>Yiwei Lu, Matthew Y. R. Yang, Zuoqiu Liu, Gautam Kamath, Yaoliang Yu</i>	
OXYGENERATOR: Reconstructing Global Ocean Deoxygenation Over a Century with Deep Learning	33219
<i>Bin Lu, Ze Zhao, Luyu Han, Xiaoying Gan, Yuntao Zhou, Lei Zhou, Luoyi Fu, Xinbing Wang, Chenghu Zhou, Jing Zhang</i>	
DiNADO: Norm-Disentangled Neurally-Decomposed Oracles for Controlling Language Models	33243
<i>Sidi Lu, Wenbo Zhao, Chenyang Tao, Arpit Gupta, Shanchan Wu, Tagyoung Chung, Nanyun Peng</i>	
SPP: Sparsity-Preserved Parameter-Efficient Fine-Tuning for Large Language Models.....	33254
<i>Xudong Lu, Aojun Zhou, Yuhui Xu, Renrui Zhang, Peng Gao, Hongsheng Li</i>	
Scaling Laws for Fine-Grained Mixture of Experts	33270
<i>Jan Ludziejewski, Jakub Krajewski, Kamil Adamczewski, Maciej Pióro, Michal Krutul, Szymon Antoniak, Kamil Ciebiera, Krystian Król, Tomasz Odrzygózd, Piotr Sankowski, Marek Cygan, Sebastian Jaszczur</i>	
Projecting Molecules into Synthesizable Chemical Spaces.....	33289
<i>Shitong Luo, Wenhao Gao, Zuofan Wu, Jian Peng, Connor W. Coley, Jianzhu Ma</i>	
PGODE: Towards High-Quality System Dynamics Modeling.....	33305
<i>Xiao Luo, Yiyang Gu, Huiyu Jiang, Hang Zhou, Jinsheng Huang, Wei Ju, Zhiping Xiao, Ming Zhang, Yizhou Sun</i>	
Unveiling the Cycloid Trajectory of EM Iterations in Mixed Linear Regression.....	33329
<i>Zhankun Luo, Abolfazl Hashemi</i>	
OMPO: A Unified Framework for RL Under Policy and Dynamics Shifts.....	33383
<i>Yu Luo, Tianying Ji, Fuchun Sun, Jianwei Zhang, Huazhe Xu, Xianyuan Zhan</i>	

Offline-Boosted Actor-Critic: Adaptively Blending Optimal Historical Behaviors in Deep Off-Policy RL.....	33411
<i>Yu Luo, Tianying Ji, Fuchun Sun, Jianwei Zhang, Huazhe Xu, Xianyuan Zhan</i>	
Position: Reinforcement Learning in Dynamic Treatment Regimes Needs Critical Reexamination	33432
<i>Zhiyao Luo, Yangchen Pan, Peter Watkinson, Tingting Zhu</i>	
Hierarchical Neural Operator Transformer with Learnable Frequency-Aware Loss Prior for Arbitrary-Scale Super-Resolution	33466
<i>Xihaier Luo, Xiaoning Qian, Byung-Jun Yoon</i>	
Potential Based Diffusion Motion Planning	33486
<i>Yunhao Luo, Chen Sun, Joshua B. Tenenbaum, Yilun Du</i>	
Cluster-Aware Similarity Diffusion for Instance Retrieval.....	33511
<i>Jifei Luo, Hantao Yao, Changsheng Xu</i>	

PART 41

End-To-End Neuro-Symbolic Reinforcement Learning with Textual Explanations	33533
<i>Lirui Luo, Guoxi Zhang, Hongming Xu, Yaodong Yang, Cong Fang, Qing Li</i>	
RoboMP ² : A Robotic Multimodal Perception-Planning Framework with Multimodal Large Language Models	33558
<i>Qi Lv, Hao Li, Xiang Deng, Rui Shao, Michael Yu Wang, Liqiang Nie</i>	
Contamination-Resilient Anomaly Detection Via Adversarial Learning on Partially-Observed Normal and Anomalous Data	33575
<i>Wenxi Lv, Qinliang Su, Hai Wan, Hongteng Xu, Wenchao Xu</i>	
Coarse-To-Fine Highlighting: Reducing Knowledge Hallucination in Large Language Models.....	33594
<i>Qitan Lv, Jie Wang, Hanzhu Chen, Bin Li, Yongdong Zhang, Feng Wu</i>	
Efficient and Effective Time-Series Forecasting with Spiking Neural Networks.....	33624
<i>Changze Lv, Yansen Wang, Dongqi Han, Xiaoqing Zheng, Xuanjing Huang, Dongsheng Li</i>	
Cross-Domain Policy Adaptation by Capturing Representation Mismatch.....	33638
<i>Jiafei Lyu, Chenjia Bai, Jingwen Yang, Zongqing Lu, Xiu Li</i>	
Sampling is as Easy as Keeping the Consistency: Convergence Guarantee for Consistency Models.....	33664
<i>Junlong Lyu, Zhitang Chen, Shoubo Feng</i>	
Parameter Efficient Quasi-Orthogonal Fine-Tuning Via Givens Rotation	33686
<i>Xinyu Ma, Xu Chu, Zhibang Yang, Yang Lin, Xin Gao, Junfeng Zhao</i>	
Rethinking Decision Transformer Via Hierarchical Reinforcement Learning.....	33730
<i>Yi Ma, Jianye Hao, Hebin Liang, Chenjun Xiao</i>	
Better Locally Private Sparse Estimation Given Multiple Samples Per User	33746
<i>Yuheng Ma, Ke Jia, Hanfang Yang</i>	
Learning Modality Knowledge Alignment for Cross-Modality Transfer	33777
<i>Wenxuan Ma, Shuang Li, Lincan Cai, Jingxuan Kang</i>	
Beyond the Federation: Topology-Aware Federated Learning for Generalization to Unseen Clients.....	33794
<i>Mengmeng Ma, Tang Li, Xi Peng</i>	
Outlier-Aware Slicing for Post-Training Quantization in Vision Transformer	33811
<i>Yuexiao Ma, Huixia Li, Xiawu Zheng, Feng Ling, Xuefeng Xiao, Rui Wang, Shilei Wen, Fei Chao, Rongrong Ji</i>	
X-Oscar: A Progressive Framework for High-Quality Text-Guided 3D Animatable Avatar Generation	33826
<i>Yiwei Ma, Zhekai Lin, Jiayi Ji, Yijun Fan, Xiaoshuai Sun, Rongrong Ji</i>	

Neighboring Perturbations of Knowledge Editing on Large Language Models.....	33839
<i>Jun-Yu Ma, Zhen-Hua Ling, Ningyu Zhang, Jia-Chen Gu</i>	
Do Transformer World Models Give Better Policy Gradients?	33855
<i>Michel Ma, Tianwei Ni, Clement Gehring, Pierluca D'Oro, Pierre-Luc Bacon</i>	
Differentiable Distributionally Robust Optimization Layers.....	33880
<i>Xutao Ma, Chao Ning, Wenli Du</i>	
CKGConv: General Graph Convolution with Continuous Kernels.....	33902
<i>Liheng Ma, Soumyasundar Pal, Yitian Zhang, Jiaming Zhou, Yingxue Zhang, Mark Coates</i>	
Reward Shaping for Reinforcement Learning with an Assistant Reward Agent.....	33925
<i>Haozhe Ma, Kuankuan Sima, Thanh Vinh Vo, Di Fu, Tze-Yun Leong</i>	
LLM and Simulation as Bilevel Optimizers: A New Paradigm to Advance Physical Scientific Discovery	33940
<i>Pingchuan Ma, Tsun-Hsuan Wang, Minghao Guo, Zhiqing Sun, Joshua B. Tenenbaum, Daniela Rus, Chuang Gan, Wojciech Matusik</i>	
Fast Peer Adaptation with Context-Aware Exploration.....	33963
<i>Long Ma, Yuanfei Wang, Fangwei Zhong, Song-Chun Zhu, Yizhou Wang</i>	
HarmonyDream: Task Harmonization Inside World Models	33983
<i>Haoyu Ma, Jialong Wu, Ningya Feng, Chenjun Xiao, Dong Li, Jianye Hao, Jianmin Wang, Mingsheng Long</i>	
High-Dimensional Linear Bandits with Knapsacks.....	34008
<i>Wanteng Ma, Dong Xia, Jiashuo Jiang</i>	
SyCoCa: Symmetrizing Contrastive Captioners with Attentive Masking for Multimodal Alignment.....	34038
<i>Ziping Ma, Furong Xu, Jian Liu, Ming Yang, Qingpei Guo</i>	
The Surprising Effectiveness of Skip-Tuning in Diffusion Sampling	34053
<i>Jiajun Ma, Shuchen Xue, Tianyang Hu, Wenjia Wang, Zhaoqiang Liu, Zhenguo Li, Zhi-Ming Ma, Kenji Kawaguchi</i>	
Correcting Diffusion-Based Perceptual Image Compression with Privileged End-To-End Decoder.....	34075
<i>Yiyang Ma, Wenhan Yang, Jiaying Liu</i>	
A Provable Decision Rule for Out-Of-Distribution Detection	34094
<i>Xinsong Ma, Xin Zou, Weiwei Liu</i>	
Prospector Heads: Generalized Feature Attribution for Large Models & Data	34115
<i>Gautam Machiraju, Alexander Derry, Arjun Desai, Neel Guha, Amir-Hossein Karimi, James Zou, Russ B. Altman, Christopher Ré, Parag Mallick</i>	
Unsupervised Parameter-Free Simplicial Representation Learning with Scattering Transforms.....	34145
<i>Hiren Madhu, Sravanthi Gurugubelli, Sundeep Prabhakar Chepuri</i>	
Faithfulness Measurable Masked Language Models.....	34161
<i>Andreas Madsen, Siva Reddy, Sarath Chandar</i>	
On the Hardness of Probabilistic Neurosymbolic Learning	34203
<i>Jaron Maene, Vincent Derkinderen, Luc De Raedt</i>	
Random Latent Exploration for Deep Reinforcement Learning.....	34219
<i>Srinath Mahankali, Zhang-Wei Hong, Ayush Sekhari, Alexander Rakhlin, Pulkit Agrawal</i>	
Mastering Zero-Shot Interactions in Cooperative and Competitive Simultaneous Games.....	34253
<i>Yannik Mahlau, Frederik Schubert, Bodo Rosenhahn</i>	
Split-And-Denoise: Protect Large Language Model Inference with Local Differential Privacy.....	34281
<i>Peihua Mai, Ran Yan, Zhe Huang, Youjia Yang, Yan Pang</i>	

TinyBenchmarks: Evaluating LLMs with Fewer Examples.....	34303
<i>Felipe Maia Polo, Lucas Weber, Leshem Choshen, Yuekai Sun, Gongjun Xu, Mikhail Yurochkin</i>	
SCoRe: Submodular Combinatorial Representation Learning	34327
<i>Anay Majeed, Suraj Kothawade, Krishnateja Killamsetty, Rishabh Iyer</i>	
Position: Data-Driven Discovery with Large Generative Models	34350
<i>Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv Agarwal, Sanchaita Hazra, Ashish Sabharwal, Peter Clark</i>	

PART 42

LASER: Linear Compression in Wireless Distributed Optimization	34383
<i>Ashok Vardhan Makkuva, Marco Bondaschi, Thijs Vogels, Martin Jaggi, Hyeji Kim, Michael Gastpar</i>	
Auto-Regressive Next-Token Predictors Are Universal Learners	34417
<i>Eran Malach</i>	
Self-Composing Policies for Scalable Continual Reinforcement Learning.....	34432
<i>Mikel Malagón, Josu Ceberio, Jose A. Lozano</i>	
Noise-Aware Algorithm for Heterogeneous Differentially Private Federated Learning.....	34461
<i>Saber Malekmohammadi, Yaoliang Yu, Yang Cao</i>	
Measures of Diversity and Space-Filling Designs for Categorical Data	34499
<i>Cedric Malherbe, Emilio Domínguez-Sánchez, Merwan Barlier, Igor Colin, Haitham Bou-Ammar, Tom Diethe</i>	
COPAL: Continual Pruning in Large Language Generative Models.....	34529
<i>Srikanth Malla, Joon Hee Choi, Chiho Choi</i>	
Minimum-Norm Interpolation Under Covariate Shift.....	34543
<i>Neil Mallinar, Austin Zane, Spencer Frei, Bin Yu</i>	
Tilting the Odds at the Lottery: the Interplay of Overparameterisation and Curricula in Neural Networks	34586
<i>Stefano Sarao Mannelli, Yaroslau Ivashynka, Andrew Saxe, Luca Saglietti</i>	
Zero-Shot Unsupervised and Text-Based Audio Editing Using DDPM Inversion	34603
<i>Hila Manor, Tomer Michaeli</i>	
Submodular Framework for Structured-Sparse Optimal Transport.....	34630
<i>Piyushi Manupriya, Pratik Jawanpuria, Karthik S. Gurumoorthy, Sakethanath Jagarlapudi, Bamdev Mishra</i>	
Large Language Models Are Geographically Biased	34654
<i>Rohin Manvi, Samar Khanna, Marshall Burke, David Lobell, Stefano Ermon</i>	
Position: Graph Foundation Models Are Already Here	34670
<i>Haitao Mao, Zhikai Chen, Wenzhuo Tang, Jianan Zhao, Yao Ma, Tong Zhao, Neil Shah, Mikhail Galkin, Jiliang Tang</i>	
Towards General Neural Surrogate Solvers with Specialized Neural Accelerators.....	34693
<i>Chenkai Mao, Robert Lupoiu, Tianxiang Dai, Mingkun Chen, Jonathan A. Fan</i>	
H-Consistency Guarantees for Regression	34712
<i>Anqi Mao, Mehryar Mohri, Yutao Zhong</i>	
Regression with Multi-Expert Deferral	34738
<i>Anqi Mao, Mehryar Mohri, Yutao Zhong</i>	
No-Regret Reinforcement Learning in Smooth MDPs.....	34760
<i>Davide Maran, Alberto Maria Metelli, Matteo Papini, Marcello Restelli</i>	

Keep the Momentum: Conservation Laws Beyond Euclidean Gradient Flows	34790
<i>Sibylle Marcotte, Rémi Gribonval, Gabriel Peyré</i>	
Quantum Theory and Application of Contextual Optimal Transport.....	34822
<i>Nicola Mariella, Albert Akhriev, Francesco Tacchino, Christa Zoufal, Juan Carlos Gonzalez-Espitia, Benedek Harsanyi, Eugene Koskin, Ivano Tavernelli, Stefan Woerner, Marianna Rapsomaniki, Sergiy Zhuk, Jannis Born</i>	
Graph-Based Forecasting with Missing Data Through Spatiotemporal Downsampling.....	34846
<i>Ivan Marisca, Cesare Alippi, Filippo Maria Bianchi</i>	
A Unified View of FANOVA: A Comprehensive Bayesian Framework for Component Selection and Estimation.....	34866
<i>Yosra Marnissi, Maxime Leiber</i>	
Expand-And-Cluster: Parameter Recovery of Neural Networks.....	34895
<i>Flavio Martinelli, Berfin Simsek, Wulfram Gerstner, Johanni Brea</i>	
Convergence and Trade-Offs in Riemannian Gradient Descent and Riemannian Proximal Point	34920
<i>David Martínez-Rubio, Christophe Roux, Sebastian Pokutta</i>	
Using AI Uncertainty Quantification to Improve Human Decision-Making.....	34949
<i>Laura R. Marusich, Jonathan Z. Bakdash, Yan Zhou, Murat Kantarcioglu</i>	
On the Tractability of SHAP Explanations Under Markovian Distributions.....	34961
<i>Reda Marzouk, Colin De La Higuera</i>	
Classification Under Nuisance Parameters and Generalized Label Shift in Likelihood-Free Inference.....	34987
<i>Luca Masserano, Alex Shen, Michele Doro, Tommaso Dorigo, Rafael Izbicki, Ann B. Lee</i>	
On the Consistency of Kernel Methods with Dependent Observations.....	35013
<i>Pierre-François Massiani, Sebastian Trimpe, Friedrich Solowjow</i>	
Online Variational Sequential Monte Carlo.....	35039
<i>Alessandro Mastrototaro, Jimmy Olsson</i>	
Amortized Variational Deep Kernel Learning	35063
<i>Alan L. S. Matias, César Lincoln C. Mattos, João P. P. Gomes, Diego P. P. Mesquita</i>	
Hieros: Hierarchical Imagination on Structured State Space Sequence World Models.....	35079
<i>Paul Mattes, Rainer Schlosser, Ralf Herbrich</i>	
Craftax: A Lightning-Fast Benchmark for Open-Ended Reinforcement Learning	35104
<i>Michael Matthews, Michael Beukman, Benjamin Ellis, Mikayel Samvelyan, Matthew Jackson, Samuel Coward, Jakob Foerster</i>	
Sampling in Unit Time with Kernel Fisher–Rao Flow	35138
<i>Aimee Maurais, Youssef Marzouk</i>	
Joint Composite Latent Space Bayesian Optimization.....	35163
<i>Natalie Maus, Zhiyuan Jerry Lin, Maximilian Balandat, Eytan Bakshy</i>	
HarmBench: A Standardized Evaluation Framework for Automated Red Teaming and Robust Refusal.....	35181
<i>Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, David Forsyth, Dan Hendrycks</i>	

PART 43

Deep Fusion: Efficient Network Training Via Pre-Trained Initializations	35225
<i>Hanna Mazzawi, Xavi Gonzalvo, Michael Wunder, Sammy Jerome, Benoit Dherin</i>	
Incremental Topological Ordering and Cycle Detection with Predictions	35240
<i>Samuel McCauley, Benjamin Moseley, Aidin Niaparast, Shikha Singh</i>	

Position: Standardization of Behavioral Use Clauses is Necessary for the Adoption of Responsible Licensing of AI.....	35255
<i>Daniel McDuff, Tim Korjakow, Scott Cambo, Jesse Josua Benjamin, Jenny Lee, Yacine Jernite, Carlos Muñoz Ferrandis, Aaron Gokaslan, Alek Tarkowski, Joseph Lindley, A. Feder Cooper, Danish Contractor</i>	
Roping in Uncertainty: Robustness and Regularization in Markov Games	35267
<i>Jeremy McMahan, Giovanni Artiglio, Qiaomin Xie</i>	
Copyright Traps for Large Language Models.....	35296
<i>Matthieu Meeus, Igor Shilov, Manuel Faysse, Yves-Alexandre De Montjoye</i>	
IM-3D: Iterative Multiview Diffusion and Reconstruction for High-Quality 3D Generation.....	35310
<i>Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, Natalia Neverova, Andrea Vedaldi, Oran Gafni, Filippos Kokkinos</i>	
On Learning Deep O(n)-Equivariant Hyperspheres	35324
<i>Pavlo Melnyk, Michael Felsberg, Mårten Wadenbäck, Andreas Robinson, Cuong Le</i>	
Position: Tensor Networks Are a Valuable Asset for Green AI	35340
<i>Eva Memmel, Clara Menzen, Jetze Schuurmans, Frederiek Wesel, Kim Batselier</i>	
OSSCAR: One-Shot Structured Pruning in Vision and Language Models with Combinatorial Optimization.....	35354
<i>Xiang Meng, Shibal Ibrahim, Kayhan Behdin, Hussein Hazimeh, Natalia Ponomareva, Rahul Mazumder</i>	
Physics-Informed Neural Network Policy Iteration: Algorithms, Convergence, and Verification	35378
<i>Yiming Meng, Ruikun Zhou, Amartya Mukherjee, Maxwell Fitzsimmons, Christopher Song, Jun Liu</i>	
Benign Overfitting in Two-Layer ReLU Convolutional Neural Networks for XOR Data	35404
<i>Xuran Meng, Difan Zou, Yuan Cao</i>	
Spectral Phase Transition and Optimal PCA in Block-Structured Spiked Models	35470
<i>Pierre Mergny, Justin Ko, Florent Krzakala</i>	
The Illusion of State in State-Space Models.....	35492
<i>William Merrill, Jackson Petty, Ashish Sabharwal</i>	
Superposition Prompting: Improving and Accelerating Retrieval-Augmented Generation	35507
<i>Thomas Merth, Qichen Fu, Mohammad Rastegari, Mahyar Najibi</i>	
How Deep Do We Need: Accelerating Training and Inference of Neural ODEs Via Control Perspective.....	35528
<i>Keyan Miao, Konstantinos Gatsis</i>	
Locality-Sensitive Hashing-Based Efficient Point Transformer with Applications in High-Energy Physics.....	35546
<i>Siqi Miao, Zhiyuan Lu, Mia Liu, Javier Duarte, Pan Li</i>	
Rethinking Independent Cross-Entropy Loss for Graph-Structured Data	35570
<i>Rui Miao, Kaixiong Zhou, Yili Wang, Ninghao Liu, Ying Wang, Xin Wang</i>	
DFlow: A Generative Model Combining Denoising AutoEncoder and Normalizing Flow for High Fidelity Waveform Generation	35590
<i>Chenfeng Miao, Qingying Zhu, Minchuan Chen, Wei Hu, Zijian Li, Shaojun Wang, Jing Xiao</i>	
Rethinking Momentum Knowledge Distillation in Online Continual Learning.....	35607
<i>Nicolas Michel, Maorong Wang, Ling Xiao, Toshihiko Yamasaki</i>	
Efficient World Models with Context-Aware Tokenization.....	35623
<i>Vincent Micheli, Eloi Alonso, François Fleuret</i>	

Interplay of ROC and Precision-Recall AUCs: Theoretical Limits and Practical Implications in Binary Classification	35639
<i>Martin Mihelich, François Castagnos, Charles Dognin</i>	
CLIPZyme: Reaction-Conditioned Virtual Screening of Enzymes	35647
<i>Peter G. Mikhael, Itamar Chinn, Regina Barzilay</i>	
FlowMM: Generating Materials with Riemannian Flow Matching	35664
<i>Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, Brandon M. Wood</i>	
Can Implicit Bias Imply Adversarial Robustness?	35687
<i>Hancheng Min, René Vidal</i>	
Understanding Retrieval-Augmented Task Adaptation for Vision-Language Models.....	35719
<i>Yifei Ming, Yixuan Li</i>	
RODEO: Robust Outlier Detection Via Exposing Adaptive Out-Of-Distribution Samples	35744
<i>Hossein Mirzaei, Mohammad Jafari, Hamid Reza Dehbashi, Ali Ansari, Sepehr Ghobadi, Masoud Hadi, Arshia Soltani Moakhar, Mohammad Azizmalayeri, Mahdieh Soleymani Baghshah, Mohammad Hossein Rohban</i>	
Prodigy: An Exponentially Adaptive Parameter-Free Learner	35779
<i>Konstantin Mishchenko, Aaron Defazio</i>	
From Inverse Optimization to Feasibility to ERM	35805
<i>Saurabh Mishra, Anant Raj, Sharan Vaswani</i>	
Provable Interactive Learning with Hindsight Instruction Feedback	35829
<i>Dipendra Misra, Aldo Pacchiano, Robert Schapire</i>	
Position: Intent-Aligned AI Systems Must Optimize for Agency Preservation.....	35851
<i>Catalin Mitelut, Ben Smith, Peter Vamplew</i>	
Faster Streaming and Scalable Algorithms for Finding Directed Dense Subgraphs in Large Graphs.....	35876
<i>Slobodan Mitrovic, Theodore Pan</i>	
TERD: A Unified Framework for Safeguarding Diffusion Models Against Backdoors.....	35892
<i>Yichuan Mo, Hui Huang, Mingjie Li, Ang Li, Yisen Wang</i>	
Error Feedback Can Accurately Compress Preconditioners	35910
<i>Ionut-Vlad Modoranu, Aleksei Kalinov, Eldar Kurtic, Elias Frantar, Dan Alistarh</i>	
Why Do You Grok? A Theoretical Analysis on Grokking Modular Addition	35934
<i>Mohamad Amin Mohamadi, Zhiyuan Li, Lei Wu, Danica J. Sutherland</i>	
Straight-Through Meets Sparse Recovery: The Support Exploration Algorithm	35968
<i>Mimoun Mohamed, François Malgouyres, Valentin Emiya, Caroline Chaux</i>	
OAK: Enriching Document Representations Using Auxiliary Knowledge for Extreme Classification	36012
<i>Shikhar Mohan, Deepak Saini, Anshul Mittal, Sayak Ray Chowdhury, Bhawna Paliwal, Jian Jiao, Manish Gupta, Manik Varma</i>	
Language Models with Conformal Factuality Guarantees	36029
<i>Christopher Mohri, Tatsunori Hashimoto</i>	
Finding NEM-U: Explaining Unsupervised Representation Learning Through Neural Network Generated Explanation Masks	36048
<i>Bjørn Leth Møller, Christian Igel, Kristoffer Knutsen Wickstrøm, Jon Sparring, Robert Jenssen, Bulat Ibragimov</i>	

PART 44

A Fresh Take on Stale Embeddings: Improving Dense Retriever Training with Corrector Networks	36072
<i>Nicholas Monath, Will Grathwohl, Michael Boratko, Rob Fergus, Andrew McCallum, Manzil Zaheer</i>	
Slot Abstractors: Toward Scalable Abstract Visual Reasoning	36088
<i>Shanka Subhra Mondal, Jonathan D. Cohen, Taylor W. Webb</i>	
A Theory of Non-Linear Feature Learning with One Gradient Step in Two-Layer Neural Networks	36106
<i>Behrad Moniri, Donghwan Lee, Hamed Hassani, Edgar Dobriban</i>	
Learning Optimal Deterministic Policies with Stochastic Policy Gradients	36160
<i>Alessandro Montenegro, Marco Mussi, Alberto Maria Metelli, Matteo Papini</i>	
A Simple Early Exiting Framework for Accelerated Sampling in Diffusion Models.....	36212
<i>Taehong Moon, Moonseok Choi, Eunggu Yun, Jongmin Yoon, Gayoung Lee, Jaewoong Cho, Juho Lee</i>	
Lookbehind-SAM: K Steps Back, 1 Step Forward.....	36229
<i>Gonçalo Mordido, Pranshu Malviya, Aristide Baratin, Sarath Chandar</i>	
Causal Representation Learning Made Identifiable by Grouping of Observational Variables	36249
<i>Hiroshi Morioka, Aapo Hyvärinen</i>	
Position: Future Directions in the Theory of Graph Machine Learning	36294
<i>Christopher Morris, Fabrizio Frasca, Nadav Dym, Haggai Maron, Ismail Ilkan Ceylan, Ron Levie, Derek Lim, Michael Bronstein, Martin Grohe, Stefanie Jegelka</i>	
Position: Levels of AGI for Operationalizing Progress on the Path to AGI	36308
<i>Meredith Ringel Morris, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, Shane Legg</i>	
Gibbs Sampling of Continuous Potentials on a Quantum Computer.....	36322
<i>Arsalan Motamedi, Pooya Ronagh</i>	
Using Uncertainty Quantification to Characterize and Improve Out-Of-Domain Learning for PDEs.....	36372
<i>S. Chandra Mouli, Danielle C. Maddix, Shima Alizadeh, Gaurav Gupta, Andrew Stuart, Michael W. Mahoney, Yuyang Wang</i>	
Privacy Attacks in Decentralized Learning	36419
<i>Abdellah El Mrini, Edwige Cyffers, Aurélien Bellet</i>	
RoboCodeX: Multimodal Code Generation for Robotic Behavior Synthesis	36434
<i>Yao Mu, Junting Chen, Qinglong Zhang, Shoufa Chen, Qiaojun Yu, Chongjian Ge, Runjian Chen, Zhixuan Liang, Mengkang Hu, Chaofan Tao, Peize Sun, Haibao Yu, Chao Yang, Wenqi Shao, Wenhai Wang, Jifeng Dai, Yu Qiao, Mingyu Ding, Ping Luo</i>	
On the Second-Order Convergence of Biased Policy Gradient Algorithms	36455
<i>Siqiao Mu, Diego Klabjan</i>	
Controlled Decoding from Language Models	36486
<i>Sidharth Mudgal, Jong Lee, Harish Ganapathy, Yaguang Li, Tao Wang, Yanping Huang, Zhifeng Chen, Heng-Tze Cheng, Michael Collins, Trevor Strohman, Jilin Chen, Alex Beutel, Ahmad Beirami</i>	
SiBBIInGS: Similarity-Driven Building-Block Inference Using Graphs Across States.....	36504
<i>Noga Mudrik, Gal Mishne, Adam S. Charles</i>	
SaVeR: Optimal Data Collection Strategy for Safe Policy Evaluation in Tabular MDP	36531
<i>Subhojyoti Mukherjee, Josiah P. Hanna, Robert Nowak</i>	
Active Preference Learning for Large Language Models.....	36577
<i>William Muldrew, Peter Hayes, Mingtian Zhang, David Barber</i>	

Is Kernel Prediction More Powerful than Gating in Convolutional Neural Networks?	36591
<i>Lorenz K. Müller</i>	
Truly No-Regret Learning in Constrained MDPs.....	36605
<i>Adrian Müller, Pragnya Alatur, Volkan Cevher, Giorgia Ramponi, Niao He</i>	
Aligning Transformers with Weisfeiler–Leman	36654
<i>Luis Müller, Christopher Morris</i>	
Position: Optimization in SciML Should Employ the Function Space Geometry.....	36705
<i>Johannes Müller, Marius Zeinhofer</i>	
Individual Fairness in Graph Decomposition.....	36723
<i>Kamesh Munagala, Govind S. Sankar</i>	
Nash Learning from Human Feedback.....	36743
<i>Rémi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Côme Fiegel, Andrea Michi, Marco Selvi, Sertan Girgin, Nikola Momchev, Olivier Bachem, Daniel J. Mankowitz, Doina Precup, Bilal Piot</i>	
Optimal Bounds for ℓ_p Sensitivity Sampling Via ℓ_2 Augmentation.....	36769
<i>Alexander Munteanu, Simon Omlor</i>	
Turnstile ℓ_p Leverage Score Sampling with Applications	36797
<i>Alexander Munteanu, Simon Omlor</i>	
Learning to Route Among Specialized Experts for Zero-Shot Generalization.....	36829
<i>Mohammed Muqeeth, Haokun Liu, Yufan Liu, Colin Raffel</i>	
Autoformalizing Euclidean Geometry.....	36847
<i>Logan Murphy, Kaiyu Yang, Jialiang Sun, Zhaoyu Li, Anima Anandkumar, Xujie Si</i>	
BAGEL: Bootstrapping Agents by Guiding Exploration with Language.....	36894
<i>Shikhar Murty, Christopher D. Manning, Peter Shaw, Mandar Joshi, Kenton Lee</i>	

PART 45

Factored-Reward Bandits with Intermediate Observations	36911
<i>Marco Mussi, Simone Drago, Marcello Restelli, Alberto Maria Metelli</i>	
Best Arm Identification for Stochastic Rising Bandits.....	36953
<i>Marco Mussi, Alessandro Montenegro, Francesco Trovò, Marcello Restelli, Alberto Maria Metelli</i>	
GATE: How to Keep Out Intrusive Neighbors.....	36990
<i>Nimrah Mustafa, Rebekka Burkholz</i>	
Test-Time Regret Minimization in Meta Reinforcement Learning	37016
<i>Mirco Mutti, Aviv Tamar</i>	
Saliency Strikes Back: How Filtering Out High Frequencies Improves White-Box Explanations	37041
<i>Sabine Muzellec, Thomas Fel, Victor Boutin, Léo Andéol, Rufin Vanrullen, Thomas Serre</i>	
Learning Temporal Distances: Contrastive Successor Features Can Provide a Metric Structure for Decision-Making	37076
<i>Vivek Myers, Chongyi Zheng, Anca Dragan, Sergey Levine, Benjamin Eysenbach</i>	
Diffusion Rejection Sampling	37097
<i>Byeonghu Na, Yeongmin Kim, Minsang Park, Donghyeok Shin, Wanmo Kang, Il-Chul Moon</i>	
LAGMA: LAtent Goal-Guided Multi-Agent Reinforcement Learning	37122
<i>Hyungho Na, Il-Chul Moon</i>	

Learning in Deep Factor Graphs with Gaussian Belief Propagation	37141
<i>Seth Nabarro, Mark Van Der Wilk, Andrew J. Davison</i>	
Don't Be So Negative! Score-Based Generative Modeling with Oracle-Assisted Guidance	37164
<i>Saeid Naderiparizi, Xiaoxuan Liang, Setareh Cohan, Berend Zwartsenberg, Frank Wood</i>	
Conditionally-Conjugate Gaussian Process Factor Analysis for Spike Count Data Via Data Augmentation	37188
<i>Yididiya Y. Nadew, Xuhui Fan, Christopher J. Quinn</i>	
Slicing Mutual Information Generalization Bounds for Neural Networks	37213
<i>Kimia Nadjahi, Kristjan Greenewald, Rickard Brüel Gabrielsson, Justin Solomon</i>	
PairNet: Training with Observed Pairs to Estimate Individual Treatment Effect.....	37237
<i>Lokesh Nagalapatti, Pranava Singhal, Avishek Ghosh, Sunita Sarawagi</i>	
Density Ratio Estimation with Doubly Strong Robustness	37260
<i>Ryosuke Nagumo, Hironori Fujisawa</i>	
Solving Poisson Equations Using Neural Walk-On-Spheres.....	37277
<i>Hong Chul Nam, Julius Berner, Anima Anandkumar</i>	
Time Weaver: A Conditional Time Series Generation Model.....	37293
<i>Sai Shankar Narasimhan, Shubhankar Agarwal, Oguzhan Akcin, Sujay Sanghavi, Sandeep Chinchali</i>	
PIVOT: Iterative Visual Prompting Elicits Actionable Knowledge for VLMs.....	37321
<i>Soroush Nasiriany, Fei Xia, Wenhao Yu, Ted Xiao, Jacky Liang, Ishita Dasgupta, Annie Xie, Danny Driess, Ayzaan Wahid, Zhuo Xu, Quan Vuong, Tingnan Zhang, Tsang-Wei Edward Lee, Kuang-Huei Lee, Peng Xu, Sean Kirmani, Yuke Zhu, Andy Zeng, Karol Hausman, Nicolas Heess, Chelsea Finn, Sergey Levine, Brian Ichter</i>	
Overestimation, Overfitting, and Plasticity in Actor-Critic: The Bitter Lesson of Reinforcement Learning	37342
<i>Michal Nauman, Michal Borkiewicz, Piotr Milos, Tomasz Trzcinski, Mateusz Ostaszewski, Marek Cygan</i>	
Box Facets and Cut Facets of Lifted Multicut Polytopes	37365
<i>Lucas Fabian Naumann, Jannik Irmay, Shengxian Zhao, Bjoern Andres</i>	
Equivariant Deep Weight Space Alignment	37376
<i>Aviv Navon, Aviv Shamsian, Ethan Fetaya, Gal Chechik, Nadav Dym, Haggai Maron</i>	
Dynamic Memory Compression: Retrofitting LLMs for Accelerated Inference.....	37396
<i>Piotr Nawrot, Adrian Łańcucki, Marcin Chochowski, David Tarjan, Edoardo M. Ponti</i>	
Stable Differentiable Causal Discovery.....	37413
<i>Achille Nazaret, Justin Hong, Elham Azizi, David Blei</i>	
SelfVC: Voice Conversion with Iterative Refinement Using Self Transformations.....	37446
<i>Paarth Neekhara, Shehzeen Hussain, Rafael Valle, Boris Ginsburg, Rishabh Ranjan, Shlomo Dubnov, Farinaz Koushanfar, Julian McAuley</i>	
A Computational Framework for Solving Wasserstein Lagrangian Flows.....	37461
<i>Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, Alireza Makhzani</i>	
On the Fairness Impacts of Hardware Selection in Machine Learning	37486
<i>Sree Harsha Nelaturu, Nishaanth K Ravichandran, Cuong Tran, Sara Hooker, Ferdinando Fioretto</i>	
Dealing with Unbounded Gradients in Stochastic Saddle-Point Optimization	37508
<i>Gergely Neu, Nneka Okolo</i>	
Score-Based Causal Discovery of Latent Variable Causal Models	37531
<i>Ignavier Ng, Xinshuai Dong, Haoyue Dai, Biwei Huang, Peter Spirtes, Kun Zhang</i>	

Measuring Stochastic Data Complexity with Boltzmann Influence Functions	37553
<i>Nathan Ng, Roger Grosse, Marzyeh Ghassemi</i>	
Is Temperature Sample Efficient for Softmax Gaussian Mixture of Experts?	37570
<i>Huy Nguyen, Pedram Akbarian, Nhat Ho</i>	
A General Theory for Softmax Gating Multinomial Logistic Mixture of Experts	37617
<i>Huy Nguyen, Pedram Akbarian, Trungtin Nguyen, Nhat Ho</i>	
PARCv2: Physics-Aware Recurrent Convolutional Neural Networks for Spatiotemporal Dynamics Modeling	37649
<i>Phong C. H Nguyen, Xinlun Cheng, Shahab Azarfar, Pradeep Seshadri, Yen T. Nguyen, Munho Kim, Sanghun Choi, H. S. Udaykumar, Stephen Baek</i>	
Quality-Weighted Vendi Scores And Their Application to Diverse Experimental Design	37667
<i>Quan Nguyen, Adji Bousso Dieng</i>	
Multiplicative Weights Update, Area Convexity and Random Coordinate Descent for Densest Subgraph Problems.....	37683
<i>Ta Duy Nguyen, Alina Ene</i>	
On Least Square Estimation in Softmax Gating Mixture of Experts.....	37707
<i>Huy Nguyen, Nhat Ho, Alessandro Rinaldo</i>	
Structure-Aware E(3)-Invariant Molecular Conformer Aggregation Networks	37736
<i>Duy M. H. Nguyen, Nina Lukashina, Tai Nguyen, An T. Le, Trungtin Nguyen, Nhat Ho, Jan Peters, Daniel Sonntag, Viktor Zaverkin, Mathias Niepert</i>	

PART 46

Generative Conditional Distributions by Neural (Entropic) Optimal Transport.....	37761
<i>Bao Nguyen, Binh Nguyen, Hieu Trung Nguyen, Viet Anh Nguyen</i>	
PIDformer: Transformer Meets Control Theory.....	37776
<i>Tam Nguyen, César A. Uribe, Tan M. Nguyen, Richard G. Baraniuk</i>	
Differentially Private Exact Recovery for Stochastic Block Models.....	37798
<i>Dung Nguyen, Anil Vullikanti</i>	
Novel Spectral Algorithms for the Partial Credit Model	37840
<i>Duc Nguyen, Anderson Ye Zhang</i>	
Sliced Wasserstein with Random-Path Projecting Directions	37879
<i>Khai Nguyen, Shujian Zhang, Tam Le, Nhat Ho</i>	
On the Statistical Complexity of Offline Decision-Making	37900
<i>Thanh Nguyen-Tang, Raman Arora</i>	
NExT: Teaching Large Language Models to Reason About Code Execution	37929
<i>Ansong Ni, Miltiadis Allamanis, Arman Cohan, Yinlin Deng, Kensen Shi, Charles Sutton, Pengcheng Yin</i>	
On the Nonlinearity of Layer Normalization.....	37957
<i>Yunhao Ni, Yuxin Guo, Junlong Jia, Lei Huang</i>	
Risk-Sensitive Reward-Free Reinforcement Learning with CVaR.....	37999
<i>Xinyi Ni, Guanlin Liu, Lifeng Lai</i>	
How Transformers Learn Causal Structure with Gradient Descent.....	38018
<i>Eshaan Nichani, Alex Damian, Jason D. Lee</i>	
Online Cascade Learning for Efficient Inference Over Streams	38071
<i>Lunyu Nie, Zhimin Ding, Erdong Hu, Christopher Jermaine, Swarat Chaudhuri</i>	

Compositional Text-To-Image Generation with Dense Blob Representations	38091
<i>Weili Nie, Sifei Liu, Morteza Mardani, Chao Liu, Benjamin Eckart, Arash Vahdat</i>	
Nearest Neighbour Score Estimators for Diffusion Generative Models.....	38117
<i>Matthew Niedoba, Dylan Green, Saeid Naderiparizi, Vasileios Lioutas, Jonathan Wilder Lavington, Xiaoxuan Liang, Yunpeng Liu, Ke Zhang, Setareh Dabiri, Adam Scibior, Berend Zwartsenberg, Frank Wood</i>	
Reward Model Learning Vs. Direct Policy Optimization: A Comparative Analysis of Learning from Human Preferences.....	38145
<i>Andi Nika, Debmalya Mandal, Parameswaran Kamalaruban, Georgios Tzannetos, Goran Radanovic, Adish Singla</i>	
RoSA: Accurate Parameter-Efficient Fine-Tuning Via Robust Adaptation.....	38187
<i>Mahdi Nikdan, Soroush Tabesh, Elvir Crncevic, Dan Alistarh</i>	
REMEDI: Corrective Transformations for Improved Neural Entropy Estimation	38207
<i>Viktor Nilsson, Anirban Samaddar, Sandeep Madireddy, Pierre Nyquist</i>	
Indirectly Parameterized Concrete Autoencoders.....	38237
<i>Alfred Nilsson, Klas Wijk, Sai Bharath Chandra Gutha, Erik Englesson, Alexandra Hotti, Carlo Saccardi, Oskar Kviman, Jens Lagergren, Ricardo Vinuesa, Hossein Azizpour</i>	
Understanding the Impact of Introducing Constraints at Inference Time on Generalization Error	38253
<i>Masaaki Nishino, Kengo Nakamura, Norihito Yasuda</i>	
Risk Aware Benchmarking of Large Language Models	38264
<i>Apoorva Nitsure, Youssef Mroueh, Mattia Rigotti, Kristjan Greenewald, Brian Belgodere, Mikhail Yurochkin, Jiri Navratil, Igor Melnyk, Jarret Ross</i>	
Test-Time Model Adaptation with Only Forward Passes	38298
<i>Shuaicheng Niu, Chunyan Miao, Guohao Chen, Pengcheng Wu, Peilin Zhao</i>	
Latent Optimal Paths by Gumbel Propagation for Variational Bayesian Dynamic Programming	38316
<i>Xinlei Niu, Christian Walder, Jing Zhang, Charles Patrick Martin</i>	
GFlowNet Training by Policy Gradients	38344
<i>Puhua Niu, Shili Wu, Mingzhou Fan, Xiaoning Qian</i>	
Multi-Fidelity Residual Neural Processes for Scalable Surrogate Modeling	38381
<i>Ruijia Niu, Dongxia Wu, Kai Kim, Yi-An Ma, Duncan Watson-Parris, Rose Yu</i>	
RNAFlow: RNA Structure & Sequence Design Via Inverse Folding-Based Flow Matching	38395
<i>Divya Nori, Wengong Jin</i>	
Skill Set Optimization: Reinforcing Language Model Behavior Via Transferable Skills	38409
<i>Kolby Nottingham, Bodhisattwa Prasad Majumder, Bhavana Dalvi Mishra, Sameer Singh, Peter Clark, Roy Fox</i>	
DITTO: Diffusion Inference-Time T-Optimization for Music Generation.....	38426
<i>Zachary Novack, Julian McAuley, Taylor Berg-Kirkpatrick, Nicholas J. Bryan</i>	
f-Divergence Based Classification: Beyond the Use of Cross-Entropy.....	38448
<i>Nicola Novello, Andrea M. Tonello</i>	
Sparser, Better, Deeper, Stronger: Improving Static Sparse Training with Exact Orthogonal Initialization.....	38474
<i>Aleksandra Irena Nowak, Łukasz Gniecki, Filip Szatkowski, Jacek Tabor</i>	
In Value-Based Deep Reinforcement Learning, a Pruned Network is a Good Network	38495
<i>Johan Obando-Ceron, Aaron Courville, Pablo Samuel Castro</i>	
Mixtures of Experts Unlock Parameter Scaling for Deep RL	38520
<i>Johan Obando-Ceron, Ghada Sokar, Timon Willi, Clare Lyle, Jesse Farebrother, Jakob Foerster, Karolina Dziugaite, Doina Precup, Pablo Samuel Castro</i>	

On the Effectiveness of Supervision in Asymmetric Non-Contrastive Learning	38541
<i>Jeongheon Oh, Kibok Lee</i>	
Sign Gradient Descent-Based Neuronal Dynamics: ANN-To-SNN Conversion Beyond ReLU Network.....	38562
<i>Hyunseok Oh, Youngki Lee</i>	

PART 47

The Perception-Robustness Tradeoff in Deterministic Image Restoration.....	38599
<i>Guy Ohayon, Tomer Michaeli, Michael Elad</i>	
Linear Explanations for Individual Neurons	38639
<i>Tuomas Oikarinen, Tsui-Wei Weng</i>	
Adaptive Proximal Gradient Methods Are Universal Without Approximation.....	38663
<i>Konstantinos Oikonomidis, Emanuel Laude, Puya Latafat, Andreas Themelis, Panagiotis Patrinos</i>	
SILVER: Single-Loop Variance Reduction and Application to Federated Learning.....	38683
<i>Kazusato Oko, Shunta Akiyama, Denny Wu, Tomoya Murata, Taiji Suzuki</i>	
Local Feature Selection Without Label Or Feature Leakage for Interpretable Machine Learning Predictions.....	38740
<i>Harrie Oosterhuis, Lijun Lyu, Avishek Anand</i>	
Do Language Models Exhibit the Same Cognitive Biases in Problem Solving as Human Learners?=-1.....	38762
<i>Andreas Opedal, Alessandro Stolfo, Haruki Shirakami, Ying Jiao, Ryan Cotterell, Bernhard Schölkopf, Abulhair Saparov, Mrinmaya Sachan</i>	
Deep Stochastic Mechanics	38779
<i>Elena Orlova, Aleksei Ustimenko, Ruoxi Jiang, Peter Y. Lu, Rebecca Willett</i>	
Variational Linearized Laplace Approximation for Bayesian Deep Learning.....	38815
<i>Luis A. Ortega, Simón Rodríguez Santana, Daniel Hernández-Lobato</i>	
Universality of Linear Recurrences Followed by Non-Linear Projections: Finite-Width Guarantees and Benefits of Complex Eigenvalues.....	38837
<i>Antonio Orvieto, Soham De, Caglar Gulcehre, Razvan Pascanu, Samuel L. Smith</i>	
Discovering Multiple Solutions from a Single Task in Offline Reinforcement Learning	38864
<i>Takayuki Osa, Tatsuya Harada</i>	
Towards Modular LLMs by Building and Reusing a Library of LoRAs.....	38885
<i>Oleksiy Ostapenko, Zhan Su, Edoardo Maria Ponti, Laurent Charlin, Nicolas Le Roux, Lucas Caccia, Alessandro Sordani</i>	
Few-Shot Unsupervised Implicit Neural Shape Representation Learning with Spatial Adversaries.....	38905
<i>Amine Ouasfi, Adnane Boukhayma</i>	
Differentiable Mapper for Topological Optimization of Data Representation	38919
<i>Ziyad Oulhaj, Mathieu Carrière, Bertrand Michel</i>	
Structured Chemistry Reasoning with Large Language Models	38937
<i>Siru Ouyang, Zhuosheng Zhang, Bing Yan, Xuan Liu, Yejin Choi, Jiawei Han, Lianhui Qin</i>	
Stochastic Quantum Sampling for Non-Logconcave Distributions and Estimating Partition Functions.....	38953
<i>Guneykan Ozgul, Xiantao Li, Mehrdad Mahdavi, Chunhao Wang</i>	
MADA: Meta-Adaptive Optimizers Through Hyper-Gradient Descent.....	38983
<i>Kaan Ozkara, Can Karakus, Parameswaran Raman, Mingyi Hong, Shoham Sabach, Branislav Kveton, Volkan Cevher</i>	

Listenable Maps for Audio Classifiers	39009
<i>Francesco Paissan, Mirco Ravanelli, Cem Subakan</i>	
Implicit Representations Via Operator Learning	39022
<i>Sourav Pal, Harshvardhan Adepur, Clinton Wang, Polina Golland, Vikas Singh</i>	
Bayesian Program Learning by Decompiling Amortized Knowledge.....	39042
<i>Alessandro B. Palmarini, Christopher G. Lucas, N. Siddharth</i>	
Two Heads Are Actually Better than One: Towards Better Adversarial Robustness Via Transduction and Rejection.....	39056
<i>Nils Palumbo, Yang Guo, Xi Wu, Jiefeng Chen, Yingyu Liang, Somesh Jha</i>	
Counterfactual Image Editing.....	39087
<i>Yushu Pan, Elias Bareinboim</i>	
MALIBO: Meta-Learning for Likelihood-Free Bayesian Optimization	39102
<i>Jiarong Pan, Stefan Falkner, Felix Berkenkamp, Joaquin Vanschoren</i>	
S ² IP-LLM: Semantic Space Informed Prompt Learning with LLM for Time Series Forecasting.....	39135
<i>Zijie Pan, Yushan Jiang, Sahil Garg, Anderson Schneider, Yuriy Nevmyvaka, Dongjin Song</i>	
Feedback Loops with Language Models Drive In-Context Reward Hacking	39154
<i>Alexander Pan, Erik Jones, Meena Jagadeesan, Jacob Steinhardt</i>	
Stability and Generalization for Stochastic Recursive Momentum-Based Algorithms for (Strongly-)Convex One to K-Level Stochastic Optimizations	39201
<i>Xiaokang Pan, Xingyu Li, Jin Liu, Tao Sun, Kai Sun, Lixing Chen, Zhe Qu</i>	
RMIB: Representation Matching Information Bottleneck for Matching Text Representations	39276
<i>Haihui Pan, Zhifang Liao, Wenrui Xie, Kun Han</i>	
Coprocessor Actor Critic: A Model-Based Reinforcement Learning Approach for Adaptive Brain Stimulation	39292
<i>Michelle Pan, Mariah Schrum, Vivek Myers, Erdem Biyik, Anca Dragan</i>	
Auto-Encoding Morph-Tokens for Multimodal LLM	39308
<i>Kaihang Pan, Siliang Tang, Juncheng Li, Zhaoyu Fan, Wei Chow, Shuicheng Yan, Tat-Seng Chua, Yueting Zhuang, Hanwang Zhang</i>	
Model-Free Robust ϕ -Divergence Reinforcement Learning Using Both Offline and Online Data.....	39324
<i>Kishan Panaganti, Adam Wierman, Eric Mazumdar</i>	
A New Linear Scaling Rule for Private Adaptive Hyperparameter Optimization	39364
<i>Ashwinee Panda, Xinyu Tang, Saeed Mahloujifar, Vikash Sehwal, Prateek Mittal</i>	
BRAIn: Bayesian Reward-Conditioned Amortized Inference for Natural Language Generation from Feedback.....	39400
<i>Gaurav Pandey, Yatin Nandwani, Tahira Naseem, Mayank Mishra, Guangxuan Xu, Dinesh Raghu, Sachindra Joshi, Asim Munawar, Ramón Fernandez Astudillo</i>	
Self-Alignment of Large Language Models Via Monopolylogue-Based Social Scene Simulation.....	39416
<i>Xianghe Pang, Shuo Tang, Rui Ye, Yuxin Xiong, Bolun Zhang, Yanfeng Wang, Siheng Chen</i>	

PART 48

Trainable Transformer in Transformer	39448
<i>Abhishek Panigrahi, Sadhika Malladi, Mengzhou Xia, Sanjeev Arora</i>	
Position: A Call for Embodied AI.....	39493
<i>Giuseppe Paolo, Jonas Gonzalez-Billandon, Balázs Kégl</i>	
Arrows of Time for Large Language Models	39509
<i>Vassilis Papadopoulos, Jérémie Wenger, Clément Hongler</i>	

Position: Topological Deep Learning is the New Frontier for Relational Learning	39529
<i>Theodore Papamarkou, Tolga Birdal, Michael Bronstein, Gunnar Carlsson, Justin Curry, Yue Gao, Mustafa Hajij, Roland Kwitt, Pietro Lio, Paolo Di Lorenzo, Vasileios Maroulas, Nina Miolane, Farzana Nasrin, Karthikeyan Natesan Ramamurthy, Bastian Rieck, Simone Scardapane, Michael T. Schaub, Petar Veli. Kovi., Bei Wang, Yusu Wang, Guo-Wei Wei, Ghada Zamzmi</i>	
Position: Bayesian Deep Learning is Needed in the Age of Large-Scale AI	39556
<i>Theodore Papamarkou, Maria Skoularidou, Konstantina Palla, Laurence Aitchison, Julyan Arbel, David Dunson, Maurizio Filippone, Vincent Fortuin, Philipp Hennig, José Miguel Hernández-Lobato, Aliaksandr Hubin, Alexander Immer, Theofanis Karaletsos, Mohammad Emtiyaz Khan, Agustinus Kristiadi, Yingzhen Li, Stephan Mandt, Christopher Nemeth, Michael A. Osborne, Tim G. J. Rudner, David Rügamer, Yee Whye Teh, Max Welling, Andrew Gordon Wilson, Ruqi Zhang</i>	
Memoria: Resolving Fateful Forgetting Problem Through Human-Inspired Memory Architecture	39587
<i>Sangjun Park, Jinyeong Bak</i>	
The Max-Min Formulation of Multi-Objective Reinforcement Learning: From Theory to a Model-Free Algorithm	39616
<i>Giseung Park, Woohyeon Byeon, Seongmin Kim, Elad Havakuk, Amir Leshem, Youngchul Sung</i>	
The Linear Representation Hypothesis and the Geometry of Large Language Models	39643
<i>Kiho Park, Yo Joong Choe, Victor Veitch</i>	
Mitigating Oversmoothing Through Reverse Process of GNNs for Heterophilic Graphs.....	39667
<i>Moonjeong Park, Jaeseung Heo, Dongwoo Kim</i>	
Any-Precision LLM: Low-Cost Deployment of Multiple, Different-Sized LLMs	39682
<i>Yeonhong Park, Jake Hyun, Sanglyul Cho, Bonggeun Sim, Jae W. Lee</i>	
Mean-Field Chaos Diffusion Models	39702
<i>Sungwoo Park, Dongjun Kim, Ahmed M. Alaa</i>	
Foundation Policies with Hilbert Representations.....	39737
<i>Seohong Park, Tobias Kreiman, Sergey Levine</i>	
SignSGD with Federated Defense: Harnessing Adversarial Attacks Through Gradient Sign Decoding	39762
<i>Chanho Park, Namyoon Lee</i>	
Position: Automatic Environment Shaping is the Next Frontier in RL	39781
<i>Younghyo Park, Gabriel B. Margolis, Pulkit Agrawal</i>	
Can Mamba Learn How to Learn? a Comparative Study on In-Context Learning Tasks	39793
<i>Jongho Park, Jaeseung Park, Zheyang Xiong, Nayoung Lee, Jaewoong Cho, Samet Oymak, Kangwook Lee, Dimitris Papailiopoulos</i>	
BOtied: Multi-Objective Bayesian Optimization with Tied Multivariate Ranks.....	39813
<i>Ji Won Park, Nataša Tagasovska, Michael Maser, Stephen Ra, Kyunghyun Cho</i>	
State-Free Inference of State-Space Models: The Transfer Function Approach.....	39834
<i>Rom N. Parnichkun, Stefano Massaroli, Alessandro Moro, Jimmy T. H. Smith, Ramin Hasani, Mathias Lechner, Qi An, Christopher Ré, Hajime Asama, Stefano Ermon, Taiji Suzuki, Michael Poli, Atsushi Yamashita</i>	
Variational Inference with Coverage Guarantees in Simulation-Based Inference	39861
<i>Yash Patel, Declan McNamara, Jackson Loper, Jeffrey Regier, Ambuj Tewari</i>	
Towards Global Optimality for Practical Average Reward Reinforcement Learning Without Mixing Time Oracles	39889
<i>Bhrij Patel, Wesley A. Suttle, Alec Koppel, Vaneet Aggarwal, Brian M. Sadler, Dinesh Manocha, Amrit Singh Bedi</i>	

Optimal Ridge Regularization for Out-Of-Distribution Prediction	39908
<i>Pratik Patil, Jin-Hong Du, Ryan J. Tibshirani</i>	
PerceptAnon: Exploring the Human Perception of Image Anonymization Beyond Pseudonymization for GDPR	39955
<i>Kartik Patwari, Chen-Nee Chuah, Lingjuan Lyu, Vivek Sharma</i>	
Estimating Canopy Height at Scale	39972
<i>Jan Pauls, Max Zimmer, Una M. Kelly, Martin Schwartz, Sassan Saatchi, Philippe Ciais, Sebastian Pokutta, Martin Brandt, Fabian Gieseke</i>	
LPGD: A General Framework for Backpropagation Through Embedded Optimization Layers	39989
<i>Anselm Paulus, Georg Martius, Vit Musil</i>	
Learning to Stabilize Online Reinforcement Learning in Unbounded State Spaces	40015
<i>Brahma S. Pavse, Matthew Zurek, Yudong Chen, Qiaomin Xie, Josiah P. Hanna</i>	
In-Context Unlearning: Language Models as Few-Shot Unlearners	40034
<i>Martin Pawelczyk, Seth Neel, Himabindu Lakkaraju</i>	
Graph Automorphism Group Equivariant Neural Networks	40051
<i>Edward Pearce-Crump, William J. Knottenbelt</i>	
Multi-Track Message Passing: Tackling Oversmoothing and Oversquashing in Graph Learning Via Preventing Heterophily Mixing	40078
<i>Hongbin Pei, Yu Li, Huiqi Deng, Jingxin Hai, Pinghui Wang, Jie Ma, Jing Tao, Yuheng Xiong, Xiaohong Guan</i>	
Exploiting Code Symmetries for Learning Program Semantics	40092
<i>Kexin Pei, Weichen Li, Qirui Jin, Shuyang Liu, Scott Geng, Lorenzo Cavallaro, Junfeng Yang, Suman Jana</i>	
Modeling Language Tokens as Functionals of Semantic Fields	40114
<i>Zhengqi Pei, Anran Zhang, Shuhui Wang, Qingming Huang</i>	
Data-Free Neural Representation Compression with Riemannian Neural Dynamics	40129
<i>Zhengqi Pei, Anran Zhang, Shuhui Wang, Xiangyang Ji, Qingming Huang</i>	
BetterV: Controlled Verilog Generation with Discriminative Guidance	40145
<i>Zehua Pei, Hui-Ling Zhen, Mingxuan Yuan, Yu Huang, Bei Yu</i>	
Bias of Stochastic Gradient Descent Or the Architecture: Disentangling the Effects of Overparameterization of Neural Networks	40154
<i>Amit Peleg, Matthias Hein</i>	
Knowledge Distillation with Auxiliary Variable	40185
<i>Bo Peng, Zhen Fang, Guangquan Zhang, Jie Lu</i>	
UPAM: Unified Prompt Attack in Text-To-Image Generation Models Against Both Textual Filters and Visual Checkers	40200
<i>Duo Peng, Qihong Ke, Jun Liu</i>	
eCeLLM: Generalizing Large Language Models for E-Commerce from Large-Scale, High-Quality Instruction Data	40215
<i>Bo Peng, Xinyi Ling, Ziru Chen, Huan Sun, Xia Ning</i>	
Pragmatic Feature Preferences: Learning Reward-Relevant Preferences from Human Input	40258
<i>Andi Peng, Yuying Sun, Tianmin Shu, David Abel</i>	
UPOCR: Towards Unified Pixel-Level OCR Interface	40271
<i>Dezhi Peng, Zhenhua Yang, Jiaxin Zhang, Chongyu Liu, Yongxin Shi, Kai Ding, Fengjun Guo, Lianwen Jin</i>	

PART 49

Block Acceleration Without Momentum: On Optimal Stepsizes of Block Gradient Descent for Least-Squares	40295
<i>Liangzu Peng, Wotao Yin</i>	
FedCal: Achieving Local and Global Calibration in Federated Learning Via Aggregated Parameterized Scaler	40331
<i>Hongyi Peng, Han Yu, Xiaoli Tang, Xiaoxiao Li</i>	
Improving Diffusion Models for Inverse Problems Using Optimal Posterior Covariance	40347
<i>Xinyu Peng, Ziyang Zheng, Wenrui Dai, Nuoqian Xiao, Chenglin Li, Junni Zou, Hongkai Xiong</i>	
A Subquadratic Time Algorithm for Robust Sparse Mean Estimation	40371
<i>Ankit Pensia</i>	
CaPS: Collaborative and Private Synthetic Data Generation from Distributed Sources	40397
<i>Sikha Pentylala, Mayana Pereira, Martine De Cock</i>	
Solving Hierarchical Information-Sharing Dec-POMDPs: An Extensive-Form Game Approach	40414
<i>Johan Peralez, Aurélien Delage, Olivier Buffet, Jilles S. Dibangoye</i>	
The Relative Value of Prediction in Algorithmic Decision Making	40439
<i>Juan Carlos Perdomo</i>	
Interpreting and Improving Diffusion Models from an Optimization Perspective	40461
<i>Frank Permenter, Chenyang Yuan</i>	
Mechanistic Neural Networks for Scientific Machine Learning	40484
<i>Adeel Pervez, Francesco Locatello, Efstratios Gavves</i>	
Bayesian Regret Minimization in Offline Bandits.....	40502
<i>Marek Petrik, Guy Tennenholtz, Mohammad Ghavamzadeh</i>	
Prompting a Pretrained Transformer Can Be a Universal Approximator	40523
<i>Aleksandar Petrov, Philip H. S. Torr, Adel Bibi</i>	
Discovering Environments with XRM	40551
<i>Mohammad Pezeshki, Diane Bouchacourt, Mark Ibrahim, Nicolas Ballas, Pascal Vincent, David Lopez-Paz</i>	
Transport of Algebraic Structure to Latent Embeddings	40570
<i>Samuel Pfommer, Brendon G. Anderson, Somayeh Sojoudi</i>	
Neural NeRF Compression.....	40592
<i>Tuan Pham, Stephan Mandt</i>	
Cross-View Masked Diffusion Transformers for Person Image Synthesis.....	40611
<i>Trung X. Pham, Zhang Kang, Chang D. Yoo</i>	
When is Transfer Learning Possible?	40642
<i>My Phan, Kianté Brantley, Stephanie Milani, Soroush Mehri, Gokul Swamy, Geoffrey J. Gordon</i>	
Controllable Prompt Tuning for Balancing Group Distributional Robustness	40667
<i>Hoang Phan, Andrew Gordon Wilson, Qi Lei</i>	
Particle Denoising Diffusion Sampler	40688
<i>Angus Phillips, Hai-Dang Dau, Michael John Hutchinson, Valentin De Bortoli, George Deligiannidis, Arnaud Doucet</i>	
Federated Continual Learning Via Prompt-Based Dual Knowledge Transfer	40725
<i>Hongming Piao, Yichen Wu, Dapeng Wu, Ying Wei</i>	

Detecting Influence Structures in Multi-Agent Reinforcement Learning.....	40740
<i>Fabian R. Pieroth, Katherine Fitch, Lenz Belzner</i>	
Rényi Pufferfish Privacy: General Additive Noise Mechanisms and Privacy Amplification by Iteration Via Shift Reduction Lemmas	40762
<i>Clément Pierquin, Aurélien Bellet, Marc Tommasi, Matthieu Boussard</i>	
Structure-Based Drug Design by Denoising Voxel Grids.....	40795
<i>Pedro O. Pinheiro, Arian Jamasb, Omar Mahmood, Vishnu Sresht, Saeed Saremi</i>	
Extracting Training Data from Document-Based VQA Models.....	40813
<i>Francesco Pinto, Nathalie Rauschmayr, Florian Tramèr, Philip Torr, Federico Tombari</i>	
Contrasting Multiple Representations with the Multi-Marginal Matching Gap.....	40827
<i>Zoe Piran, Michal Klein, James Thornton, Marco Cuturi</i>	
diff History for Neural Language Agents	40843
<i>Ulyana Piterbarg, Lerrel Pinto, Rob Fergus</i>	
Zero-Sum Positional Differential Games as a Framework for Robust Reinforcement Learning: Deep Q-Learning Approach.....	40869
<i>Anton Plaksin, Vitaly Kalev</i>	
Adaptive Conformal Inference by Betting	40886
<i>Aleksandr Podkopaev, Darren Xu, Kuang-Chih Lee</i>	
Mechanistic Design and Scaling of Hybrid Architectures.....	40908
<i>Michael Poli, Armin W Thomas, Eric Nguyen, Pragaash Ponnusamy, Björn Deiseroth, Kristian Kersting, Taiji Suzuki, Brian Hie, Stefano Ermon, Christopher Ré, Ce Zhang, Stefano Massaroli</i>	
Relaxed Quantile Regression: Prediction Intervals for Asymmetric Noise	40951
<i>Thomas Pouplin, Alan Jeffares, Nabeel Seedat, Mihaela Van Der Schaar</i>	
Robust Data-Driven Prescriptiveness Optimization.....	40982
<i>Mehran Poursoltani, Erick Delage, Angelos Georghiou</i>	
Stochastic Optimization with Arbitrary Recurrent Data Sampling.....	41000
<i>William Powell, Hanbaek Lyu</i>	
Learning Multiple Secrets in Mastermind	41039
<i>Milind Prabhu, David Woodruff</i>	
MusicFlow: Cascaded Flow Matching for Text Guided Music Generation	41052
<i>K R Prajwal, Bowen Shi, Matthew Lee, Apoorv Vyas, Andros Tjandra, Mahi Luthra, Baishan Guo, Huiyu Wang, Triantafyllos Afouras, David Kant, Wei-Ning Hsu</i>	
The Entropy Enigma: Success and Failure of Entropy Minimization	41064
<i>Ori Press, Ravid Shwartz-Ziv, Yann LeCun, Matthias Bethge</i>	
Conformal Validity Guarantees Exist for Any Data Distribution (and How to Find Them).....	41086
<i>Drew Prinster, Samuel Stanton, Anqi Liu, Suchi Saria</i>	

PART 50

Autoencoding Conditional Neural Processes for Representation Learning.....	41119
<i>Victor Prokhorov, Ivan Titov, N. Siddharth</i>	
Efficient Exploration in Average-Reward Constrained Reinforcement Learning: Achieving Near- Optimal Regret with Posterior Sampling.....	41144
<i>Danil Provodin, Maurits Kaptein, Mykola Pechenizkiy</i>	
Learning a Diffusion Model Policy from Rewards Via Q-Score Matching.....	41163
<i>Michael Psenka, Alejandro Escontrela, Pieter Abbeel, Yi Ma</i>	

Learning-Efficient Yet Generalizable Collaborative Filtering for Item Recommendation	41183
<i>Yuanhao Pu, Xiaolong Chen, Xu Huang, Jin Chen, Defu Lian, Enhong Chen</i>	
Unsupervised Domain Adaptation for Anatomical Structure Detection in Ultrasound Images.....	41204
<i>Bin Pu, Xingguo Lv, Jiewen Yang, Guannan He, Xingbo Dong, Yiqun Lin, Shengli Li, Ying Tan, Fei Liu, Ming Chen, Zhe Jin, Kenli Li, Xiaomeng Li</i>	
Amortizing Pragmatic Program Synthesis with Rankings.....	41221
<i>Yewen Pu, Saujas Vaduguru, Priyan Vaithilingam, Elena Glassman, Daniel Fried</i>	
Learning to Remove Cuts in Integer Linear Programming	41235
<i>Pol Puigdemont, Stratis Skoulakis, Grigorios G Chrysos, Volkan Cevher</i>	
On the Minimal Degree Bias in Generalization on the Unseen for non-Boolean Functions	41256
<i>Denys Pushkin, Raphaël Berthier, Emmanuel Abbe</i>	
Hierarchical Novelty Detection Via Fine-Grained Evidence Allocation	41280
<i>Spandan Pyakurel, Qi Yu</i>	
Conformalized Survival Distributions: A Generic Post-Process to Increase Calibration	41303
<i>Shi-Ang Qi, Yakun Yu, Russell Greiner</i>	
Momentor: Advancing Video Large Language Model with Fine-Grained Temporal Reasoning	41340
<i>Long Qian, Juncheng Li, Yu Wu, Yaobo Ye, Hao Fei, Tat-Seng Chua, Yueting Zhuang, Siliang Tang</i>	
ByMI: Byzantine Machine Identification with False Discovery Rate Control.....	41357
<i>Chengde Qian, Mengyuan Wang, Haojie Ren, Changliang Zou</i>	
Efficient Non-Stationary Online Learning by Wavelets with Applications to Online Distribution Shift Adaptation.....	41383
<i>Yu-Yang Qian, Peng Zhao, Yu-Jie Zhang, Masashi Sugiyama, Zhi-Hua Zhou</i>	
Ensemble Pruning for Out-Of-Distribution Generalization.....	41416
<i>Fengchun Qiao, Xi Peng</i>	
Near-Optimal Reinforcement Learning with Self-Play Under Adaptivity Constraints	41430
<i>Dan Qiao, Yu-Xiang Wang</i>	
ULAREF: A Unified Label Refinement Framework for Learning with Inaccurate Supervision	41456
<i>Congyu Qiao, Ning Xu, Yihao Hu, Xin Geng</i>	
Federated Full-Parameter Tuning of Billion-Sized Language Models with Communication Cost Under 18 Kilobytes.....	41473
<i>Zhen Qin, Daoyuan Chen, Bingchen Qian, Bolin Ding, Yaliang Li, Shuiguang Deng</i>	
Accurate LoRA-Finetuning Quantization of LLMs Via Information Retention.....	41498
<i>Haotong Qin, Xudong Ma, Xingyu Zheng, Xiaoyang Li, Yang Zhang, Shouda Liu, Jie Luo, Xianglong Liu, Michele Magno</i>	
Various Lengths, Constant Speed: Efficient Language Modeling with Lightning Attention.....	41517
<i>Zhen Qin, Weigao Sun, Dong Li, Xuyang Shen, Weixuan Sun, Yiran Zhong</i>	
Feasible Reachable Policy Iteration	41536
<i>Shentao Qin, Yujie Yang, Yao Mu, Jie Li, Wenjun Zou, Jingliang Duan, Shengbo Eben Li</i>	
Efficient PAC Learnability of Dynamical Systems Over Multilayer Networks	41557
<i>Zirou Qiu, Abhijin Adiga, Madhav V. Marathe, S. S. Ravi, Daniel J. Rosenkrantz, Richard E. Stearns, Anil Vullikanti</i>	
Learning High-Order Relationships of Brain Regions	41582
<i>Weikang Qiu, Huangrui Chu, Selena Wang, Haolan Zuo, Xiaoxiao Li, Yize Zhao, Rex Ying</i>	
To Cool Or Not to Cool? Temperature Network Meets Large Foundation Models Via DRO.....	41604
<i>Zi-Hao Qiu, Siqi Guo, Mao Xu, Tuo Zhao, Lijun Zhang, Tianbao Yang</i>	

Transferring Knowledge from Large Foundation Models to Small Downstream Models	41644
<i>Shikai Qiu, Boran Han, Danielle C. Maddix, Shuai Zhang, Yuyang Wang, Andrew Gordon Wilson</i>	
Complexity Matters: Feature Learning in the Presence of Spurious Correlations.....	41658
<i>Guanwen Qiu, Da Kuang, Surbhi Goel</i>	
Compute Better Spent: Replacing Dense Layers with Structured Matrices	41698
<i>Shikai Qiu, Andres Potapczynski, Marc Finzi, Micah Goldblum, Andrew Gordon Wilson</i>	
Gradient Compressed Sensing: A Query-Efficient Gradient Estimator for High-Dimensional Zeroth-Order Optimization.....	41717
<i>Ruizhong Qiu, Hanghang Tong</i>	
MolCRAFT: Structure-Based Drug Design in Continuous Parameter Space.....	41749
<i>Yanru Qu, Keyue Qiu, Yuxuan Song, Jingjing Gong, Jiawei Han, Mingyue Zheng, Hao Zhou, Wei-Ying Ma</i>	
Connect Later: Improving Fine-Tuning for Robustness with Targeted Augmentations	41769
<i>Helen Qu, Sang Michael Xie</i>	
Learning Constraints from Offline Demonstrations Via Superior Distribution Correction Estimation.....	41787
<i>Guorui Quan, Zhiqiang Xu, Guiliang Liu</i>	
Augmenting Decision with Hypothesis in Reinforcement Learning.....	41804
<i>Nguyen Minh Quang, Hady W. Lauw</i>	
Multiply-Robust Causal Change Attribution	41821
<i>Victor Quintas-Martinez, Mohammad Taha Bahadori, Eduardo Santiago, Jeff Mu, David Heckerman</i>	
Decomposable Submodular Maximization in Federated Setting.....	41841
<i>Akbar Rafiey</i>	
A New Robust Partial p-Wasserstein-Based Metric for Comparing Distributions	41867
<i>Sharath Raghvendra, Pouyan Shirzadian, Kaiyi Zhang</i>	
Modular Learning of Deep Causal Generative Models for High-Dimensional Causal Inference	41886
<i>Md Musfiqur Rahman, Murat Kocaoglu</i>	
Fundamental Limits of Distributed Covariance Matrix Estimation Under Communication Constraints.....	41927
<i>Mohammad Reza Rahmani, Mohammad Hossein Yassaee, Mohammad Ali Maddah-Ali, Mohammad Reza Aref</i>	

PART 51

Subsampling is Not Magic: Why Large Batch Sizes Work for Differentially Private Stochastic Optimisation	41959
<i>Ossi Räisä, Joonas Jälkö, Antti Honkela</i>	
In-Context Freeze-Thaw Bayesian Optimization for Hyperparameter Optimization.....	41982
<i>Herilalaina Rakotoarison, Steven Adriaensen, Neeratyoy Mallik, Samir Gariboy, Edward Bergman, Frank Hutter</i>	
Understanding Inter-Concept Relationships in Concept-Based Models.....	42009
<i>Naveen Raman, Mateo Espinosa Zarlenga, Mateja Jamnik</i>	
STEER: Assessing the Economic Rationality of Large Language Models	42026
<i>Narun Raman, Taylor Lundy, Samuel Joseph Amouyal, Yoav Levine, Kevin Leyton-Brown, Moshe Tennenholtz</i>	

WARM: On the Benefits of Weight Averaged Reward Models.....	42048
<i>Alexandre Ramé, Nino Vieillard, Léonard Hussenot, Robert Dadashi, Geoffrey Cideron, Olivier Bachem, Johan Ferret</i>	
Compositional Capabilities of Autoregressive Transformers: A Study on Synthetic, Interpretable Tasks.....	42074
<i>Rahul Ramesh, Ekdeep Singh Lubana, Mikail Khona, Robert P. Dick, Hidenori Tanaka</i>	
Sequence Compression Speeds Up Credit Assignment in Reinforcement Learning.....	42104
<i>Aditya A. Ramesh, Kenny Young, Louis Kirsch, Jürgen Schmidhuber</i>	
Position: The Reasonable Person Standard for AI.....	42127
<i>Sunayana Rane</i>	
Generalization to New Sequential Decision Making Tasks with In-Context Learning	42138
<i>Sharath Chandra Raparthy, Eric Hambro, Robert Kirk, Mikael Henaff, Roberta Raileanu</i>	
Challenges in Training PINNs: A Loss Landscape Perspective.....	42159
<i>Pratik Rathore, Weimu Lei, Zachary Frangella, Lu Lu, Madeleine Udell</i>	
Unveiling Privacy, Memorization, and Input Curvature Links.....	42192
<i>Deepak Ravikumar, Efstathia Soufleri, Abolfazl Hashemi, Kaushik Roy</i>	
Dissecting Multimodality in VideoQA Transformer Models by Impairing Modality Fusion.....	42213
<i>Ishaan Singh Rawal, Alexander Matyasko, Shantanu Jaiswal, Basura Fernando, Cheston Tan</i>	
Fair Federated Learning Via the Proportional Veto Core	42245
<i>Bhaskar Ray Chaudhury, Aniket Murhekar, Zhuowen Yuan, Bo Li, Ruta Mehta, Ariel D. Procaccia</i>	
Provably Robust DPO: Aligning Language Models with Noisy Feedback	42258
<i>Sayak Ray Chowdhury, Anush Kini, Nagarajan Natarajan</i>	
Implicit Bias of Policy Gradient in Linear Quadratic Control: Extrapolation to Unseen Initial States	42275
<i>Noam Razin, Yotam Alexander, Edo Cohen-Karlik, Raja Giryes, Amir Globerson, Nadav Cohen</i>	
Dynamic Anisotropic Smoothing for Noisy Derivative-Free Optimization	42332
<i>Sam Reifenstein, Timothee Leleu, Yoshihisa Yamamoto</i>	
Position: Understanding LLMs Requires More than Statistical Generalization.....	42365
<i>Patrik Reizinger, Szilvia Ujváry, Anna Mészáros, Anna Kerekes, Wieland Brendel, Ferenc Huszár</i>	
Optimal Batched Linear Bandits	42391
<i>Xuanfei Ren, Tianyuan Jin, Pan Xu</i>	
TabLog: Test-Time Adaptation for Tabular Data Using Logic Rules	42417
<i>Weijieying Ren, Xiaoting Li, Huiyuan Chen, Vineeth Rakesh, Zhuoyi Wang, Mahashweta Das, Vasant Honavar</i>	
Hybrid Inverse Reinforcement Learning	42428
<i>Juntao Ren, Gokul Swamy, Zhiwei Steven Wu, J. Andrew Bagnell, Sanjiban Choudhury</i>	
Rejuvenating image-GPT as Strong Visual Representation Learners.....	42449
<i>Sucheng Ren, Zeyu Wang, Hongru Zhu, Junfei Xiao, Alan Yuille, Cihang Xie</i>	
CarbonNovo: Joint Design of Protein Structure and Sequence Using a Unified Energy-Based Model	42462
<i>Milong Ren, Tian Zhu, Haicang Zhang</i>	
Plug-And-Play Image Restoration with Stochastic deNOising REgularization	42484
<i>Marien Renaud, Jean Prost, Arthur Leclaire, Nicolas Papadakis</i>	

Private and Federated Stochastic Convex Optimization: Efficient Strategies for Centralized Systems.....	42521
<i>Roie Reshef, Kfir Yehuda Levy</i>	
Position: Technical Research and Talent is Needed for Effective AI Governance	42543
<i>Anka Reuel, Lisa Soder, Ben Bucknall, Trond Arne Undheim</i>	
SparQ Attention: Bandwidth-Efficient LLM Inference	42558
<i>Luka Ribar, Ivan Chelombiev, Luke Hudlass-Galley, Charlie Blake, Carlo Luschi, Douglas Orr</i>	
Early Time Classification with Accumulated Accuracy Gap Control.....	42584
<i>Liran Ringel, Regev Cohen, Daniel Freedman, Michael Elad, Yaniv Romano</i>	
Implicit Regularization in Feedback Alignment Learning Mechanisms for Neural Networks.....	42601
<i>Zachary Robertson, Oluwasanmi Koyejo</i>	
Universal Gradient Methods for Stochastic Convex Optimization	42620
<i>Anton Rodomanov, Ali Kavis, Yongtao Wu, Kimon Antonakopoulos, Volkan Cevher</i>	
Position: Key Claims in LLM Research Have a Long Tail of Footnotes	42647
<i>Anna Rogers, Alexandra Sasha Luccioni</i>	
LEVI: Generalizable Fine-Tuning Via Layer-Wise Ensemble of Different Views.....	42666
<i>Yuji Roh, Qingyun Liu, Huan Gui, Zhe Yuan, Yujin Tang, Steven E. Whang, Liang Liu, Shuchao Bi, Lichan Hong, Ed H. Chi, Zhe Zhao</i>	
Position: Mission Critical – Satellite Data is a Distinct Modality in Machine Learning.....	42691
<i>Esther Rolf, Konstantin Klemmer, Caleb Robinson, Hannah Kerner</i>	
Position: Application-Driven Innovation in Machine Learning	42707
<i>David Rolnick, Alan Aspuru-Guzik, Sara Beery, Bistra Dilkina, Priya L. Donti, Marzyeh Ghassemi, Hannah Kerner, Claire Monteleoni, Esther Rolf, Milind Tambe, Adam White</i>	
One-Shot Strategic Classification Under Unknown Costs	42719
<i>Elan Rosenfeld, Nir Rosenfeld</i>	
Modelling Microbial Communities with Graph Neural Networks	42742
<i>Albane Ruaud, Cansu Sancaktar, Marco Bagatella, Christoph Ratzke, Georg Martius</i>	
Neural Operators Meet Conjugate Gradients: The FCG-NO Method for Efficient PDE Solving.....	42766
<i>Alexander Rudikov, Vladimir Fanaskov, Ekaterina Muravleva, Yuri M. Laevsky, Ivan Oseledets</i>	
Position: Amazing Things Come from Having Many Good Models.....	42783
<i>Cynthia Rudin, Chudi Zhong, Lesia Semenova, Margo Seltzer, Ronald Parr, Jiachang Liu, Srikar Katta, Jon Donnelly, Harry Chen, Zachery Boner</i>	
Generalizing Orthogonalization for Models with Non-Linearities.....	42796
<i>David Rügamer, Chris Kolb, Tobias Weber, Lucas Kook, Thomas Nagler</i>	

PART 52

Rolling Diffusion Models.....	42818
<i>David Ruhe, Jonathan Heek, Tim Salimans, Emiel Hoogeboom</i>	
Explorations of Self-Repair in Language Models	42836
<i>Cody Rushing, Neel Nanda</i>	
Gambling-Based Confidence Sequences for Bounded Random Vectors.....	42856
<i>J. Jon Ryu, Gregory W. Wornell</i>	
Operator SVD with Neural Networks Via Nested Low-Rank Approximation	42870
<i>J. Jon Ryu, Xiangxiang Xu, H. S. Melihcan Erol, Yuheng Bu, Lizhong Zheng, Gregory W. Wornell</i>	

Tandem Transformers for Inference Efficient LLMs.....	42906
<i>Aishwarya P S, Pranav Ajit Nair, Yashas Samaga, Toby Boyd, Sanjiv Kumar, Prateek Jain, Praneeth Netrapalli</i>	
Benchmarking and Building Long-Context Retrieval Models with LoCo and M2-BERT.....	42918
<i>Jon Saad-Falcon, Daniel Y. Fu, Simran Arora, Neel Guha, Christopher Ré</i>	
Align Your Steps: Optimizing Sampling Schedules in Diffusion Models.....	42947
<i>Amirmojtaba Sabour, Sanja Fidler, Karsten Kreis</i>	
Fast Adversarial Attacks on Language Models in One GPU Minute.....	42976
<i>Vinu Sankar Sadasivan, Shoumik Saha, Gaurang Sriramanan, Priyatham Kattakinda, Atoosa Chegini, Soheil Feizi</i>	
Failures Are Fated, but Can Be Faded: Characterizing and Mitigating Unwanted Behaviors in Large-Scale Vision and Language Models.....	42999
<i>Som Sagar, Aditya Taparia, Ransalu Senanayake</i>	
I/O Complexity of Attention, Or How Optimal is FlashAttention?.....	43024
<i>Barna Saha, Christopher Ye</i>	
An Online Optimization Perspective on First-Order and Zero-Order Decentralized Nonsmooth Nonconvex Stochastic Optimization.....	43043
<i>Emre Sahinoglu, Shahin Shahrampour</i>	
Second-Order Uncertainty Quantification: A Distance-Based Approach.....	43060
<i>Yusuf Sale, Viktor Bengs, Michele Caprio, Eyke Hüllermeier</i>	
Learning from Integral Losses in Physics Informed Neural Networks.....	43077
<i>Ehsan Saleh, Saba Ghaffari, Timothy Bretl, Luke Olson, Matthew West</i>	
Random Exploration in Bayesian Optimization: Order-Optimal Regret and Computational Efficiency.....	43112
<i>Sudeep Salgia, Sattar Vakili, Qing Zhao</i>	
Predictive Coding Beyond Correlations.....	43142
<i>Tommaso Salvatori, Luca Pinchetti, Amine M'Charrak, Beren Millidge, Thomas Lukasiewicz</i>	
Proactive Detection of Voice Cloning with Localized Watermarking.....	43180
<i>Robin San Roman, Pierre Fernandez, Hady Elsahar, Alexandre Défossez, Teddy Furon, Tuan Tran</i>	
Stay on Topic with Classifier-Free Guidance.....	43197
<i>Guillaume V. Sanchez, Alexander Spangher, Honglu Fan, Elad Levi, Stella Biderman</i>	
How Do Transformers Perform In-Context Autoregressive Learning?.....	43235
<i>Michaël E. Sander, Raja Giryes, Taiji Suzuki, Mathieu Blondel, Gabriel Peyré</i>	
Differentially Private Representation Learning Via Image Captioning.....	43255
<i>Tom Sander, Yaodong Yu, Maziar Sanjabi, Alain Durmus, Yi Ma, Kamalika Chaudhuri, Chuan Guo</i>	
Transformers, Parallel Computation, and Logarithmic Depth.....	43276
<i>Clayton Sanford, Daniel Hsu, Matus Telgarsky</i>	
Promoting External and Internal Equities Under Ex-Ante/Ex-Post Metrics in Online Resource Allocation.....	43328
<i>Karthik Abinav Sankararaman, Aravind Srinivasan, Pan Xu</i>	
A Diffusion Model Framework for Unsupervised Neural Combinatorial Optimization.....	43346
<i>Sebastian Sanokowski, Sepp Hochreiter, Sebastian Lehner</i>	
Sparse and Structured Hopfield Networks.....	43368
<i>Saul Santos, Vlad Niculae, Daniel McNamee, André F. T. Martins</i>	

Meta Evidential Transformer for Few-Shot Open-Set Recognition	43389
<i>Hitesh Sapkota, Krishna Prasad Neupane, Qi Yu</i>	
EvIL: Evolution Strategies for Generalisable Imitation Learning	43407
<i>Silvia Sapora, Gokul Swamy, Chris Lu, Yee Whye Teh, Jakob Nicolaus Foerster</i>	
A Sampling Theory Perspective on Activations for Implicit Neural Representations	43422
<i>Hemanth Saratchandran, Sameera Ramasinghe, Violetta Shevchenko, Alexander Long, Simon Lucey</i>	
Beyond Chinchilla-Optimal: Accounting for Inference in Language Model Scaling Laws	43445
<i>Nikhil Sardana, Jacob Portes, Sasha Doubov, Jonathan Frankle</i>	
Position: Quo Vadis, Unsupervised Time Series Anomaly Detection?	43461
<i>M. Saquib Sarfraz, Mei-Yen Chen, Lukas Layer, Kunyu Peng, Marios Koulakis</i>	
A Fast Algorithm to Simulate Nonlinear Resistive Networks	43477
<i>Benjamin Scellier</i>	
A Fixed-Point Approach for Causal Generative Modeling.....	43504
<i>Meyer Scetbon, Joel Jennings, Agrin Hilmkil, Cheng Zhang, Chao Ma</i>	
MoMo: Momentum Models for Adaptive Learning Rates	43542
<i>Fabian Schaipp, Ruben Ohana, Michael Eickenberg, Aaron Defazio, Robert M. Gower</i>	
Parallel Affine Transformation Tuning of Markov Chain Monte Carlo	43571
<i>Philip Schär, Michael Habeck, Daniel Rudolf</i>	
Incentivized Learning in Principal-Agent Bandit Games	43608
<i>Antoine Scheid, Daniil Tiapkin, Etienne Boursier, Aymeric Capitaine, Eric Moulines, Michael I. Jordan, El-Mahdi El-Mhamdi, Alain Durmus</i>	
Caduceus: Bi-Directional Equivariant Long-Range DNA Sequence Modeling	43632
<i>Yair Schiff, Chia-Hsiang Kao, Aaron Gokaslan, Tri Dao, Albert Gu, Volodymyr Kuleshov</i>	

PART 53

DySLIM: Dynamics Stable Learning by Invariant Measure for Chaotic Systems	43649
<i>Yair Schiff, Zhong Yi Wan, Jeffrey B. Parker, Stephan Hoyer, Volodymyr Kuleshov, Fei Sha, Leonardo Zepeda-Núñez</i>	
Robust CLIP: Unsupervised Adversarial Fine-Tuning of Vision Embeddings for Robust Large Vision-Language Models.....	43685
<i>Christian Schlarman, Naman Deep Singh, Francesco Croce, Matthias Hein</i>	
Tilt Your Head: Activating the Hidden Spatial-Invariance of Classifiers	43705
<i>Johann Schmidt, Sebastian Stober</i>	
Leveraging Self-Consistency for Data-Efficient Amortized Bayesian Inference	43723
<i>Marvin Schmitt, Desi R. Ivanova, Daniel Habermann, Ullrich Köthe, Paul-Christian Bürkner, Stefan T. Radev</i>	
Regularized Q-Learning Through Robust Averaging	43742
<i>Peter Schmitt-Förster, Tobias Sutter</i>	
Implicit Representations for Constrained Image Segmentation	43765
<i>Jan Philipp Schneider, Mishal Fatima, Jovita Lukasik, Andreas Kolb, Margret Keuper, Michael Moeller</i>	
Online Learning with Bounded Recall	43791
<i>Jon Schneider, Kiran Vodrahalli</i>	
Conditional Normalizing Flows for Active Learning of Coarse-Grained Molecular Representations	43804
<i>Henrik Schopmans, Pascal Friederich</i>	

Provably Efficient Long-Horizon Exploration in Monte Carlo Tree Search Through State Occupancy Regularization.....	43828
<i>Liam Schramm, Abdeslam Boularias</i>	
Asymptotics of Learning with Deep Structured (Random) Features.....	43862
<i>Dominik Schröder, Daniil Dmitriev, Hugo Cui, Bruno Loureiro</i>	
Simultaneous Identification of Models and Parameters of Scientific Simulators.....	43895
<i>Cornelius Schröder, Jakob H. Macke</i>	
In-Context Learning Agents Are Asymmetric Belief Updaters	43928
<i>Johannes A. Schubert, Akshay K. Jagadish, Marcel Binz, Eric Schulz</i>	
Towards Scalable and Versatile Weight Space Learning	43947
<i>Konstantin Schürholt, Michael W. Mahoney, Damian Borth</i>	
Meta-Learners for Partially-Identified Treatment Effects Across Multiple Environments.....	43967
<i>Jonas Schweisthal, Dennis Frauen, Mihaela Van Der Schaar, Stefan Feuerriegel</i>	
Differentiability and Optimization of Multiparameter Persistent Homology	43986
<i>Luis Scoccola, Siddharth Setlur, David Loiseaux, Mathieu Carrière, Steve Oudot</i>	
Improved Modelling of Federated Datasets Using Mixtures-Of-Dirichlet-Multinomials	44012
<i>Jonathan Scott, Áine Cahill</i>	
MindEye2: Shared-Subject Models Enable fMRI-To-Image with 1 Hour of Data	44038
<i>Paul S. Scotti, Mihir Tripathy, Cesar Kadir Torrico Villanueva, Reese Kneeland, Tong Chen, Ashutosh Narang, Charan Santhirasegaran, Jonathan Xu, Thomas Naselaris, Kenneth A. Norman, Tanishq Mathew Abraham</i>	
Curated LLM: Synergy of LLMs and Data Curation for Tabular Augmentation in Low-Data Regimes	44060
<i>Nabeel Seedat, Nicolas Huynh, Boris Van Breugel, Mihaela Van Der Schaar</i>	
Lessons from Generalization Error Analysis of Federated Learning : You May Communicate Less Often!	44093
<i>Milad Sefidgaran, Romain Chor, Abdellatif Zaidi, Yijun Wan</i>	
Algorithm of Thoughts: Enhancing Exploration of Ideas in Large Language Models.....	44136
<i>Bilgehan Sel, Ahmad Al-Tawaha, Vanshaj Khattar, Ruoxi Jia, Ming Jin</i>	
Retrieval-Augmented Score Distillation for Text-To-3D Generation.....	44190
<i>Junyoung Seo, Susung Hong, Wooseok Jang, Inès Hyeonsu Kim, Min-Seop Kwak, Doyup Lee, Seungryoung Kim</i>	
Self-Supervised Interpretable End-To-End Learning Via Latent Functional Modularity	44212
<i>Hyunki Seong, David Hyunchul Shim</i>	
Prompting is a Double-Edged Sword: Improving Worst-Group Robustness of Foundation Models	44224
<i>Amrith Setlur, Saurabh Garg, Virginia Smith, Sergey Levine</i>	
Decomposing and Editing Predictions by Modeling Model Computation	44244
<i>Harshay Shah, Andrew Ilyas, Aleksander Madry</i>	
A Multimodal Automated Interpretability Agent.....	44293
<i>Tamar Rott Shaham, Sarah Schwetmann, Franklin Wang, Achyuta Rajaram, Evan Hernandez, Jacob Andreas, Antonio Torralba</i>	
Evaluation of Trajectory Distribution Predictions with Energy Score.....	44322
<i>Novin Shahroudi, Mihkel Lepson, Meelis Kull</i>	
The Balanced-Pairwise-Affinities Feature Transform	44342
<i>Daniel Shalam, Simon Korman</i>	
DNCs Require More Planning Steps	44358
<i>Yara Shamshoum, Nitzan Hodos, Yuval Sieradzki, Assaf Schuster</i>	

Improved Generalization of Weight Space Networks Via Augmentations	44378
<i>Aviv Shamsian, Aviv Navon, David W. Zhang, Yan Zhang, Ethan Fetaya, Gal Chechik, Haggai Maron</i>	
On Online Experimentation Without Device Identifiers	44394
<i>Shiv Shankar, Ritwik Sinha, Madalina Fiterau</i>	
Improved Dimensionality Dependence for Zeroth-Order Optimisation Over Cross-Polytopes	44413
<i>Weijia Shao</i>	
On Multi-Armed Bandit with Impatient Arms	44429
<i>Yuming Shao, Zhixuan Fang</i>	
Language Generation with Strictly Proper Scoring Rules	44474
<i>Chenze Shao, Fandong Meng, Yijin Liu, Jie Zhou</i>	

PART 54

Learning Decision Policies with Instrumental Variables Through Double Machine Learning	44489
<i>Daqian Shao, Ashkan Soleymani, Francesco Quinzan, Marta Kwiatkowska</i>	
How Far Can Fairness Constraints Help Recover from Biased Data?	44515
<i>Mohit Sharma, Amit Jayant Deshpande</i>	
Diffuse, Sample, Project: Plug-And-Play Controllable Graph Generation	44545
<i>Kartik Sharma, Srijan Kumar, Rakshit S Trivedi</i>	
Sequential Neural Score Estimation: Likelihood-Free Inference with Conditional Score Based Diffusion Models	44565
<i>Louis Sharrock, Jack Simons, Song Liu, Mark Beaumont</i>	
Bespoke Non-Stationary Solvers for Fast Sampling of Diffusion and Flow Models	44603
<i>Neta Shaul, Uriel Singer, Ricky T. Q. Chen, Matthew Le, Ali Thabet, Albert Pumarola, Yaron Lipman</i>	
Reducing Sequential Change Detection to Sequential Estimation	44628
<i>Shubhanshu Shekhar, Aaditya Ramdas</i>	
Exploring the Complexity of Deep Neural Networks Through Functional Equivalence	44643
<i>Guohao Shen</i>	
Variational Learning is Effective for Large Deep Networks	44665
<i>Yuesong Shen, Nico Daheim, Bai Cong, Peter Nickl, Gian Maria Marconi, Clement Bazan, Rio Yokota, Iryna Gurevych, Daniel Cremers, Mohammad Emtiyaz Khan, Thomas Möllenhoff</i>	
Thermometer: Towards Universal Calibration for Large Language Models	44687
<i>Maohao Shen, Subhro Das, Kristjan Greenewald, Prasanna Sattigeri, Gregory Wornell, Soumya Ghosh</i>	
Position: Do Pretrained Transformers Learn In-Context by Gradient Descent?	44712
<i>Lingfeng Shen, Aayush Mishra, Daniel Khashabi</i>	
Adaptive Stabilization Based on Machine Learning for Column Generation	44741
<i>Yunzhuang Shen, Yuan Sun, Xiaodong Li, Zhiguang Cao, Andrew Eberhard, Guangquan Zhang</i>	
Tag-LLM: Repurposing General-Purpose LLMs for Specialized Domains	44759
<i>Junhong Shen, Neil Tenenholtz, James Brian Hall, David Alvarez-Melis, Nicolò Fusi</i>	
Principled Penalty-Based Methods for Bilevel Reinforcement Learning and RLHF	44774
<i>Han Shen, Zhuoran Yang, Tianyi Chen</i>	
ReLU's Are Sufficient for Learning Implicit Neural Representations	44800
<i>Joseph Shenouda, Yamin Zhou, Robert D. Nowak</i>	

Rate-Optimal Policy Optimization for Linear Markov Decision Processes	44815
<i>Uri Sherman, Alon Cohen, Tomer Koren, Yishay Mansour</i>	
Double Momentum Method for Lower-Level Constrained Bilevel Optimization	44838
<i>Wanli Shi, Yi Chang, Bin Gu</i>	
OT-CLIP: Understanding and Generalizing CLIP Via Optimal Transport	44865
<i>Liangliang Shi, Jack Fan, Junchi Yan</i>	
LCA-On-The-Line: Benchmarking Out-Of-Distribution Generalization with Class Taxonomies	44887
<i>Jia Shi, Gautam Gare, Jinjin Tian, Siqi Chai, Zhiqiu Lin, Arun Vasudevan, Di Feng, Francesco Ferroni, Shu Kong</i>	
Sample-Efficient Robust Multi-Agent Reinforcement Learning in the Face of Environmental Uncertainty	44909
<i>Laixi Shi, Eric Mazumdar, Yuejie Chi, Adam Wierman</i>	
CrossGET: Cross-Guided Ensemble of Tokens for Accelerating Vision-Language Transformers	44960
<i>Dachuan Shi, Chaofan Tao, Anyi Rao, Zhendong Yang, Chun Yuan, Jiaqi Wang</i>	
Why Larger Language Models Do In-Context Learning Differently?	44991
<i>Zhenmei Shi, Junyi Wei, Zhuoyan Xu, Yingyu Liang</i>	
Long-Tail Learning with Foundation Model: Heavy Fine-Tuning Hurts	45014
<i>Jiang-Xin Shi, Tong Wei, Zhi Zhou, Jie-Jing Shao, Xin-Yan Han, Yu-Feng Li</i>	
Neural-Kernel Conditional Mean Embeddings	45040
<i>Eiki Shimizu, Kenji Fukumizu, Dino Sejdinovic</i>	
Causal Discovery with Fewer Conditional Independence Tests	45060
<i>Kirankumar Shiragur, Jiaqi Zhang, Caroline Uhler</i>	
Statistical Test for Attention Maps in Vision Transformers	45079
<i>Tomohiro Shiraishi, Daiki Miwa, Teruyuki Katsuoka, Vo Nguyen Le Duy, Kouichi Taji, Ichiro Takeuchi</i>	
Longitudinal Targeted Minimum Loss-Based Estimation with Temporal-Difference Heterogeneous Transformer	45097
<i>Toru Shirakawa, Yi Li, Yulun Wu, Sky Qiu, Yuxuan Li, Mingduo Zhao, Hiroyasu Iso, Mark Van Der Laan</i>	
Allocation Requires Prediction Only if Inequality is Low	45114
<i>Ali Shirali, Rediet Abebe, Moritz Hardt</i>	
Prior Mismatch and Adaptation in PnP-ADMM with a Nonconvex Convergence Analysis	45154
<i>Shirin Shoushtari, Jiaming Liu, Edward P. Chandler, M. Salman Asif, Ulugbek S. Kamilov</i>	
Effects of Exponential Gaussian Distribution on (Double Sampling) Randomized Smoothing	45183
<i>Youwei Shu, Xi Xiao, Derui Wang, Yuxin Cao, Siji Chen, Minhui Xue, Linyi Li, Bo Li</i>	
WISER: Weak Supervision and Supervised Representation Learning to Improve Drug Response Prediction in Cancer	45228
<i>Kumar Shubham, Aishwarya Jayagopal, Syed Mohammed Danish, Prathosh Ap, Vaibhav Rajan</i>	
TIC-TAC: A Framework for Improved Covariance Estimation in Deep Heteroscedastic Regression	45244
<i>Megh Shukla, Mathieu Salzmann, Alexandre Alahi</i>	
Towards a Better Theoretical Understanding of Independent Subnetwork Training	45258
<i>Egor Shulgin, Peter Richtárik</i>	
Weakly Convex Regularisers for Inverse Problems: Convergence of Critical Points and Primal-Dual Optimisation	45286
<i>Zakhar Shumaylov, Jeremy Budd, Subhadip Mukherjee, Carola-Bibiane Schönlieb</i>	

Self-Supervised Coarsening of Unstructured Grid with Automatic Differentiation	45315
<i>Sergei Shumilin, Alexander Ryabov, Nikolay Yavich, Evgeny Burnaev, Vladimir Vanovski</i>	
IOI: Invisible One-Iteration Adversarial Attack on No-Reference Image- And Video-Quality Metrics.....	45329
<i>Ekaterina Shumitskaya, Anastasia Antsiferova, Dmitriy Vatolin</i>	

PART 55

InterpreTabNet: Distilling Predictive Signals from Tabular Data by Salient Feature Interpretation	45353
<i>Jacob Si, Wendy Yusi Cheng, Michael Cooper, Rahul G. Krishnan</i>	
Embarrassingly Parallel GFlowNets	45406
<i>Tiago Da Silva, Luiz Max Carvalho, Amauri Souza, Samuel Kaski, Diego Mesquita</i>	
On the Unexpected Effectiveness of Reinforcement Learning for Sequential Recommendation	45432
<i>Álvaro Labarca, Denis Parra, Rodrigo Toro Icarte</i>	
Learning from Memory: Non-Parametric Memory Augmented Self-Supervised Learning of Visual Features	45451
<i>Thalles Silva, Helio Pedrini, Adin Ramirez Rivera</i>	
Deletion-Anticipative Data Selection with a Limited Budget	45468
<i>Rachael Hwee Ling Sim, Jue Fan, Xiao Tian, Patrick Jaillet, Bryan Kian Hsiang Low</i>	
Position: AI-Powered Autonomous Weapons Risk Geopolitical Instability and Threaten AI Research	45508
<i>Riley Simmons-Edler, Ryan P. Badman, Shayne Longpre, Kanaka Rajan</i>	
Latent Variable Model for High-Dimensional Point Process with Structured Missingness	45525
<i>Maksim Sinelnikov, Manuel Hausmann, Harri Lähdesmäki</i>	
Domain Generalisation Via Imprecise Learning.....	45544
<i>Anurag Singh, Siu Lun Chau, Shahine Bouabid, Krikamol Muandet</i>	
Finite Time Logarithmic Regret Bounds for Self-Tuning Regulation	45571
<i>Rahul Singh, Akshay Mete, Avik Kar, P. R. Kumar</i>	
What Needs to Go Right for an Induction Head? a Mechanistic Study of In-Context Learning Circuits and Their Formation	45637
<i>Aaditya K. Singh, Ted Moskovitz, Felix Hill, Stephanie C. Y. Chan, Andrew M. Saxe</i>	
Representation Surgery: Theory and Practice of Affine Steering	45663
<i>Shashwat Singh, Shauli Ravfogel, Jonathan Herzig, Roei Aharoni, Ryan Cotterell, Ponnurangam Kumaraguru</i>	
PIPER: Primitive-Informed Preference-Based Hierarchical Reinforcement Learning Via Hindsight Relabeling.....	45681
<i>Utsav Singh, Wesley A. Suttle, Brian M. Sadler, Vinay P. Namboodiri, Amrit Singh Bedi</i>	
Byzantine Resilient and Fast Federated Few-Shot Learning	45696
<i>Ankit Pratap Singh, Namrata Vaswani</i>	
Parallelized Spatiotemporal Slot Binding for Videos	45707
<i>Gautam Singh, Yue Wang, Jiawei Yang, Boris Ivanovic, Sungjin Ahn, Marco Pavone, Tong Che</i>	
What's the Score? Automated Denoising Score Matching for Nonlinear Diffusions.....	45734
<i>Raghav Singhal, Mark Goldstein, Rajesh Ranganath</i>	
SAPG: Split and Aggregate Policy Gradients.....	45759
<i>Jayesh Singla, Ananye Agarwal, Deepak Pathak</i>	
In-Context Reinforcement Learning for Variable Action Spaces	45773
<i>Viacheslav Sinii, Alexander Nikulin, Vladislav Kurenkov, Ilya Zisman, Sergey Kolesnikov</i>	

Subhomogeneous Deep Equilibrium Models	45794
<i>Pietro Sittoni, Francesco Tudisco</i>	
Multi-Agent Reinforcement Learning with Hierarchical Coordination for Emergency Responder Stationing	45813
<i>Amutheezan Sivagnanam, Ava Pettet, Hunter Lee, Ayan Mukhopadhyay, Abhishek Dubey, Aron Laszka</i>	
Inexact Newton-Type Methods for Optimisation with Nonnegativity Constraints	45835
<i>Oscar Smeets, Fred Roosta</i>	
Should We Be Going MAD? a Look at Multi-Agent Debate Strategies for LLMs.....	45883
<i>Andries Smit, Nathan Grinsztajn, Paul Duckworth, Thomas D. Barrett, Arnu Pretorius</i>	
Probabilistic Modeling of Interpersonal Coordination Processes.....	45906
<i>Paulo Soares, Adarsh Pyarelal, Meghavarshini Krishnaswamy, Emily Butler, Kobus Barnard</i>	
On PI Controllers for Updating Lagrange Multipliers in Constrained Optimization	45922
<i>Motahareh Sohrabi, Juan Ramirez, Tianyue H. Zhang, Simon Lacoste-Julien, Jose Gallego-Posada</i>	
SPADE: Sparsity-Guided Debugging for Deep Neural Networks.....	45955
<i>Arshia Soltani Moakhar, Eugenia Iofinova, Elias Frantar, Dan Alistarh</i>	
Connecting the Dots: Is Mode-Connectedness the Key to Feasible Sample-Based Inference in Bayesian Neural Networks?	45988
<i>Emanuel Sommer, Lisa Wimmer, Theodore Papamarkou, Ludwig Bothmann, Bernd Bischl, David Rügamer</i>	
Hybrid Reinforcement Learning from Offline Observation Alone.....	46019
<i>Yuda Song, J. Andrew Bagnell, Aarti Singh</i>	
Multimodal Prototyping for Cancer Survival Prediction.....	46050
<i>Andrew H. Song, Richard J. Chen, Guillaume Jaume, Anurag Vaidya, Alexander S. Baras, Faisal Mahmood</i>	
SurfPro: Functional Protein Design Based on Continuous Surface.....	46074
<i>Zhenqiao Song, Tinglin Huang, Lei Li, Wengong Jin</i>	
OSN: Infinite Representations of Dynamic 3D Scenes from Monocular Videos	46089
<i>Ziyang Song, Jinxi Li, Bo Yang</i>	
Sparse is Enough in Fine-Tuning Pre-Trained Large Language Models.....	46121
<i>Weixi Song, Zuchao Li, Lefei Zhang, Hai Zhao, Bo Du</i>	
SLEB: Streamlining LLMs Through Redundancy Verification and Elimination of Transformer Blocks.....	46136
<i>Jiwon Song, Kyungseok Oh, Taesu Kim, Hyungjun Kim, Yulhwa Kim, Jae-Joon Kim</i>	
Unveiling the Dynamics of Information Interplay in Supervised Learning	46156
<i>Kun Song, Zhiquan Tan, Bochao Zou, Huimin Ma, Weiran Huang</i>	
Position: Leverage Foundational Models for Black-Box Optimization	46168
<i>Xingyou Song, Yingtao Tian, Robert Tjarko Lange, Chansoo Lee, Yujin Tang, Yutian Chen</i>	

PART 56

Rich-Observation Reinforcement Learning with Continuous Latent Dynamics	46181
<i>Yuda Song, Lili Wu, Dylan J. Foster, Akshay Krishnamurthy</i>	
Latent Logic Tree Extraction for Event Sequence Explanation from LLMs.....	46238
<i>Zitao Song, Chao Yang, Chaojie Wang, Bo An, Shuang Li</i>	
Generative Enzyme Design Guided by Functionally Important Sites and Small-Molecule Substrates.....	46259
<i>Zhenqiao Song, Yunlong Zhao, Wenxian Shi, Wengong Jin, Yang Yang, Lei Li</i>	

Position: A Roadmap to Pluralistic Alignment	46280
<i>Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Miresghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, Tim Althoff, Yejin Choi</i>	
Position: Opportunities Exist for Machine Learning in Magnetic Fusion Energy.....	46303
<i>Lucas Spangher, Allen Wang, Andrew Maris, Myles Stabelberg, Viraj Mehta, Alex Sparstein, Stephen Lane-Walsh, Akshata Kishore Moharir, Alessandro Pau, Cristina Rea</i>	
Offline Actor-Critic Reinforcement Learning Scales to Large Models	46323
<i>Jost Tobias Springenberg, Abbas Abdolmaleki, Jingwei Zhang, Oliver Groth, Michael Bloesch, Thomas Lampe, Philemon Brakel, Sarah Bechtle, Steven Kapturowski, Roland Hafner, Nicolas Heess, Martin Riedmiller</i>	
ChemReasoner: Heuristic Search Over a Large Language Model’s Knowledge Space Using Quantum-Chemical Feedback	46351
<i>Henry W. Sprueill, Carl Edwards, Khushbu Agarwal, Mariefel V. Olarte, Udishnu Sanyal, Conrad Johnston, Hongbin Liu, Heng Ji, Sutanay Choudhury</i>	
Contextual Feature Selection with Conditional Stochastic Gates.....	46375
<i>Ram Dyuthi Sristi, Ofir Lindenbaum, Shira Lifshitz, Maria Lavzin, Jackie Schiller, Gal Mishne, Hadas Benisty</i>	
The Fundamental Limits of Least-Privilege Learning.....	46393
<i>Theresa Stadler, Bogdan Kulynych, Michael C. Gastpar, Nicolas Papernot, Carmela Troncoso</i>	
Diffusion Models Encode the Intrinsic Dimension of Data Manifolds	46412
<i>Jan Pawel Stanczuk, Georgios Batzolis, Teo Deveney, Carola-Bibiane Schönlieb</i>	
Grokking Group Multiplication with Cosets	46441
<i>Dashiell Stander, Qinan Yu, Honglu Fan, Stella Biderman</i>	
Harmonic Self-Conditioned Flow Matching for Joint Multi-Ligand Docking and Binding Site Design.....	46468
<i>Hannes Stark, Bowen Jing, Regina Barzilay, Tommi Jaakkola</i>	
Dirichlet Flow Matching with Applications to DNA Sequence Design	46495
<i>Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, Tommi Jaakkola</i>	
Partial Optimality in the Linear Ordering Problem	46514
<i>David Stein, Bjoern Andres</i>	
Towards Compositionality in Concept Learning	46530
<i>Adam Stein, Aaditya Naik, Yinjun Wu, Mayur Naik, Eric Wong</i>	
Learning to Intervene on Concept Bottlenecks.....	46556
<i>David Steinmann, Wolfgang Stammer, Felix Friedrich, Kristian Kersting</i>	
QORA: Zero-Shot Transfer Via Interpretable Object-Relational Model Learning.....	46572
<i>Gabriel Stella, Dmitri Loguinov</i>	
Private Truly-Everlasting Robust-Prediction.....	46591
<i>Uri Stemmer</i>	
ReGAL: Refactoring Programs to Discover Generalizable Abstractions	46605
<i>Elias Stengel-Eskin, Archiki Prasad, Mohit Bansal</i>	
RLVF: Learning from Verbal Feedback Without Overgeneralization	46625
<i>Moritz Stephan, Alexander Khazatsky, Eric Mitchell, Annie S Chen, Sheryl Hsu, Archit Sharma, Chelsea Finn</i>	
Causal Inference from Competing Treatments.....	46657
<i>Ana-Andreea Stoica, Vivian Yvonne Nastl, Moritz Hardt</i>	

Online Learning in CMDPs: Handling Stochastic and Adversarial Constraints.....	46692
<i>Francesco Emanuele Stradi, Jacopo Germano, Gianmarco Genalti, Matteo Castiglioni, Alberto Marchesi, Nicola Gatti</i>	
Designing Decision Support Systems Using Counterfactual Prediction Sets.....	46722
<i>Eleni Straitouri, Manuel Gomez Rodriguez</i>	
DéjàVu: KV-Cache Streaming for Fast, Fault-Tolerant Generative LLM Serving.....	46745
<i>Foteini Strati, Sara McAllister, Amar Phanishayee, Jakub Tarnawski, Ana Klimovic</i>	
Learning from Streaming Data When Users Choose.....	46772
<i>Jinyan Su, Sarah Dean</i>	
From Vision to Audio and Beyond: A Unified Model for Audio-Visual Representation and Generation.....	46804
<i>Kun Su, Xiulong Liu, Eli Shlizerman</i>	
Compositional Image Decomposition with Diffusion Models.....	46823
<i>Jocelin Su, Nan Liu, Yanbo Wang, Joshua B. Tenenbaum, Yilun Du</i>	
Whispering Experts: Neural Interventions for Toxicity Mitigation in Language Models.....	46843
<i>Xavier Suau, Pieter Delobelle, Katherine Metcalf, Armand Joulin, Nicholas Apostoloff, Luca Zappella, Pau Rodriguez</i>	
Revealing Vision-Language Integration in the Brain with Multimodal Networks.....	46868
<i>Vighnesh Subramaniam, Colin Conwell, Christopher Wang, Gabriel Kreiman, Boris Katz, Ignacio Cases, Andrei Barbu</i>	
Networked Inequality: Preferential Attachment Bias in Graph Neural Network Link Prediction.....	46891
<i>Arjun Subramonian, Levent Sagun, Yizhou Sun</i>	
Enforcing Constraints in RNA Secondary Structure Predictions: A Post-Processing Framework Based on the Assignment Problem.....	46926
<i>Geewon Suh, Gyeongjo Hwang, Seokjun Kang, Doojin Baek, Mingeun Kang</i>	
ED-Copilot: Reduce Emergency Department Wait Time with Language Model Diagnostic Assistance.....	46942
<i>Liwen Sun, Abhineet Agarwal, Aaron Kornblith, Bin Yu, Chenyan Xiong</i>	
Breaking the Barrier: Enhanced Utility and Robustness in Smoothed DRL Agents.....	46957
<i>Chung-En Sun, Sicun Gao, Tsui-Wei Weng</i>	
Representing Molecules as Random Walks Over Interpretable Grammars.....	46988
<i>Michael Sun, Minghao Guo, Weize Yuan, Veronika Thost, Crystal Elaine Owens, Aristotle Franklin Grosz, Sharvaa Selvan, Katelyn Zhou, Hassan Mohiuddin, Benjamin J Pedretti, Zachary P Smith, Jie Chen, Wojciech Matusik</i>	

PART 57

Constrained Reinforcement Learning Under Model Mismatch.....	47017
<i>Zhongchang Sun, Sihong He, Fei Miao, Shaofeng Zou</i>	
DFA-RAG: Conversational Semantic Router for Large Language Model with Definite Finite Automaton.....	47033
<i>Yiyu Sun, Junjie Hu, Wei Cheng, Haifeng Chen</i>	
Online Algorithms with Uncertainty-Quantified Predictions.....	47056
<i>Bo Sun, Jerry Huang, Nicolas Christianson, Mohammad Hajiesmaili, Adam Wierman, Raouf Boutaba</i>	
LSEnet: Lorentz Structural Entropy Neural Network for Deep Graph Clustering.....	47078
<i>Li Sun, Zhenhao Huang, Hao Peng, Yujie Wang, Chunyang Liu, Philip S. Yu</i>	

Online Adaptive Anomaly Thresholding with Confidence Sequences	47105
<i>Sophia Sun, Abishek Sankararaman, Balakrishnan (Murali) Narayanaswamy</i>	
Learning Graph Representation Via Graph Entropy Maximization	47133
<i>Ziheng Sun, Xudong Wang, Chris Ding, Jicong Fan</i>	
FedBPT: Efficient Federated Black-Box Prompt Tuning for Large Language Models	47159
<i>Jingwei Sun, Ziyue Xu, Hongxu Yin, Dong Yang, Daguang Xu, Yudong Liu, Zhixu Du, Yiran Chen, Holger R. Roth</i>	
Regression Learning with Limited Observations of Multivariate Outcomes and Features	47174
<i>Yifan Sun, Grace Y. Yi</i>	
Video-SALMONN: Speech-Enhanced Audio-Visual Large Language Models	47198
<i>Guangzhi Sun, Wenyi Yu, Changli Tang, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, Yuxuan Wang, Chao Zhang</i>	
Learning High-Frequency Functions Made Easy with Sinusoidal Positional Encoding	47218
<i>Chuanhao Sun, Zhihang Yuan, Kai Xu, Luo Mai, Siddharth N, Shuo Chen, Mahesh K. Marina</i>	
Learning Latent Dynamic Robust Representations for World Models	47234
<i>Ruixiang Sun, Hongyu Zang, Xin Li, Riashat Islam</i>	
Self-Cognitive Denoising in the Presence of Multiple Noisy Label Sources	47261
<i>Yi-Xuan Sun, Ya-Lin Zhang, Bin Han, Longfei Li, Jun Zhou</i>	
BBox-Adapter: Lightweight Adapting for Black-Box Large Language Models	47280
<i>Haotian Sun, Yuchen Zhuang, Wei Wei, Chao Zhang, Bo Dai</i>	
On a Combinatorial Problem Arising in Machine Teaching	47305
<i>Joakim Sunde, Brigit Håvardstun, Jan Kratochvíl, Jan Arne Telle</i>	
Interpretable Deep Clustering for Tabular Data	47314
<i>Jonathan Svirsky, Ofir Lindenbaum</i>	
Reinforcement Learning from Reachability Specifications: PAC Guarantees with Expected Conditional Distance	47331
<i>Jakub Svoboda, Suguman Bansal, Krishnendu Chatterjee</i>	
A Minimaximalist Approach to Reinforcement Learning from Human Feedback	47345
<i>Gokul Swamy, Christoph Dann, Rahul Kidambi, Zhiwei Steven Wu, Alekh Agarwal</i>	
Fast Sampling-Based Sketches for Tensors	47378
<i>William Swartworth, David Woodruff</i>	
Sample Complexity Bounds for Estimating Probability Divergences Under Invariances	47396
<i>Behrooz Tahmasebi, Stefanie Jegelka</i>	
A Universal Class of Sharpness-Aware Minimization Algorithms	47418
<i>Behrooz Tahmasebi, Ashkan Soleymani, Dara Bahri, Stefanie Jegelka, Patrick Jaillet</i>	
Preference Fine-Tuning of LLMs Should Leverage Suboptimal, On-Policy Data	47441
<i>Fahim Tajwar, Anikait Singh, Archit Sharma, Rafael Rafailov, Jeff Schneider, Tengyang Xie, Stefano Ermon, Chelsea Finn, Aviral Kumar</i>	
Mean-Field Analysis on Two-Layer Neural Networks from a Kernel Perspective	47475
<i>Shokichi Takakura, Taiji Suzuki</i>	
Posterior Sampling-Based Bayesian Optimization with Tighter Bayesian Regret Bounds	47510
<i>Shion Takeno, Yu Inatsu, Masayuki Karasuyama, Ichiro Takeuchi</i>	
Codebook Features: Sparse and Discrete Interpretability for Neural Networks	47535
<i>Alex Tamkin, Mohammad Tafteeque, Noah D. Goodman</i>	
Deciphering RNA Secondary Structure Prediction: A Probabilistic K-Rook Matching Perspective	47564
<i>Cheng Tan, Zhangyang Gao, Hanqun Cao, Xingran Chen, Ge Wang, Lirong Wu, Jun Xia, Jiangbin Zheng, Stan Z. Li</i>	

Community-Invariant Graph Contrastive Learning.....	47579
<i>Shiyin Tan, Dongyuan Li, Renhe Jiang, Ying Zhang, Manabu Okumura</i>	
Fourier Controller Networks for Real-Time Decision-Making in Embodied Learning.....	47607
<i>Hengkai Tan, Songming Liu, Kai Ma, Chengyang Ying, Xingxing Zhang, Hang Su, Jun Zhu</i>	
Learning Solution-Aware Transformers for Efficiently Solving Quadratic Assignment Problem.....	47627
<i>Zhentao Tan, Yadong Mu</i>	
Information Flow in Self-Supervised Learning.....	47649
<i>Zhiquan Tan, Jingqin Yang, Weiran Huang, Yang Yuan, Yifan Zhang</i>	
OTMatch: Improving Semi-Supervised Learning with Optimal Transport.....	47667
<i>Zhiquan Tan, Kaipeng Zheng, Weiran Huang</i>	
Post-Hoc Part-Prototype Networks	47681
<i>Andong Tan, Fengtao Zhou, Hao Chen</i>	
Causally Motivated Personalized Federated Invariant Learning with Shortcut-Averse Information-Theoretic Regularization	47694
<i>Xueyang Tang, Song Guo, Jingcai Guo, Jie Zhang, Yue Yu</i>	
Generalized Preference Optimization: A Unified Approach to Offline Alignment.....	47725
<i>Yunhao Tang, Daniel Zhaohan Guo, Zeyu Zheng, Daniele Calandriello, Rémi Munos, Mark Rowland, Pierre Harvey Richemond, Michal Valko, Bernardo Ávila Pires, Bilal Piot</i>	
Rethinking Optimization and Architecture for Tiny Language Models	47743
<i>Yehui Tang, Kai Han, Fangcheng Liu, Yunsheng Ni, Yuchuan Tian, Zheyuan Bai, Yi-Qi Hu, Sichao Liu, Shangling Jui, Yunhe Wang</i>	
Residual-Conditioned Optimal Transport: Towards Structure-Preserving Unpaired and Paired Image Restoration.....	47757
<i>Xiaole Tang, Xin Hu, Xiang Gu, Jian Sun</i>	
Merging Multi-Task Models Via Weight-Ensembling Mixture of Experts.....	47778
<i>Anke Tang, Li Shen, Yong Luo, Nan Yin, Lefei Zhang, Dacheng Tao</i>	
Accelerating Parallel Sampling of Diffusion Models.....	47800
<i>Zhiwei Tang, Jiasheng Tang, Hao Luo, Fan Wang, Tsung-Hui Chang</i>	
Membership Inference Attacks on Diffusion Models Via Quantile Regression	47819
<i>Shuai Tang, Zhiwei Steven Wu, Sergul Aydore, Michael Kearns, Aaron Roth</i>	
StrokeNUWA—Tokenizing Strokes for Vector Graphic Synthesis.....	47830
<i>Zecheng Tang, Chenfei Wu, Zekai Zhang, Minheng Ni, Shengming Yin, Yu Liu, Zhengyuan Yang, Lijuan Wang, Zicheng Liu, Juntao Li, Nan Duan</i>	
SSL4Q: Semi-Supervised Learning of Quantum Data with Application to Quantum State Classification.....	47846
<i>Yehui Tang, Nianzu Yang, Mabiao Long, Junchi Yan</i>	

PART 58

Finite Smoothing Algorithm for High-Dimensional Support Vector Machines and Quantile Regression	47865
<i>Qian Tang, Yikai Zhang, Boxiang Wang</i>	
MathScale: Scaling Instruction Tuning for Mathematical Reasoning	47885
<i>Zhengyang Tang, Xingxing Zhang, Benyou Wang, Furu Wei</i>	
QUEST: Query-Aware Sparsity for Efficient Long-Context LLM Inference.....	47901
<i>Jiaming Tang, Yilong Zhao, Kan Zhu, Guangxuan Xiao, Baris Kasikci, Song Han</i>	
MLI Formula: A Nearly Scale-Invariant Solution with Noise Perturbation.....	47912
<i>Bowen Tao, Xin-Chun Li, De-Chuan Zhan</i>	

Learning in Feature Spaces Via Coupled Covariances: Asymmetric Kernel SVD and Nyström Method	47929
<i>Qinghua Tao, Francesco Tonin, Alex Lambert, Yingyi Chen, Panagiotis Patrinos, Johan A. K. Suykens</i>	
Selective Mixup Helps with Distribution Shifts,But Not (Only) Because of Mixup.....	47948
<i>Damien Teney, Jindong Wang, Ehsan Abbasnejad</i>	
Quantum Positional Encodings for Graph Neural Networks.....	47965
<i>Slimane Thabet, Mehdi Djellabi, Igor O. Sokolov, Sachin Kasture, Louis-Paul Henry, Loïc Henriot</i>	
Sparse-IFT: Sparse Iso-FLOP Transformations for Maximizing Training Efficiency	47997
<i>Vithursan Thangarasa, Shreyas Saxena, Abhay Gupta, Sean Lie</i>	
SleepFM: Multi-Modal Representation Learning for Sleep Across Brain Activity, ECG and Respiratory Signals.....	48019
<i>Rahul Thapa, Bryan He, Magnus Ruud Kjær, Hyatt Moore, Gauri Ganjoo, Emmanuel Mignot, James Zou</i>	
Bayesian Adaptation of Network Depth and Width for Continual Learning.....	48038
<i>Jeevan Thapa, Rui Li</i>	
Position: What Makes an Image Realistic?	48062
<i>Lucas Theis</i>	
A New Branch-And-Bound Pruning Framework for ℓ_0 -Regularized Problems	48077
<i>Théo Guyard, Cédric Herzet, Clément Elvira, Ayse-Nur Arslan</i>	
Beyond Individual Input for Deep Anomaly Detection on Tabular Data	48097
<i>Hugo Thimonier, Fabrice Popineau, Arpad Rimmel, Bich-Liên Doan</i>	
Risk Estimation in a Markov Cost Process: Lower and Upper Bounds	48124
<i>Gugan Thoppe, Prashanth L. A., Sanjay P. Bhat</i>	
Collapse-Aware Triplet Decoupling for Adversarially Robust Image Retrieval.....	48139
<i>Qiwei Tian, Chenhao Lin, Zhengyu Zhao, Qian Li, Chao Shen</i>	
MOKD: Cross-Domain Finetuning for Few-Shot Classification Via Maximizing Optimized Kernel Dependence	48154
<i>Hongduan Tian, Feng Liu, Tongliang Liu, Bo Du, Yiu-Ming Cheung, Bo Han</i>	
Liouville Flow Importance Sampler.....	48186
<i>Yifeng Tian, Nishant Panda, Yen Ting Lin</i>	
Ranking-Based Client Imitation Selection for Efficient Federated Learning	48211
<i>Chunlin Tian, Zhan Shi, Xinpeng Qin, Li Li, Chengzhong Xu</i>	
Towards the Theory of Unsupervised Federated Learning: Non-Asymptotic Analysis of Federated EM Algorithms	48226
<i>Ye Tian, Haolei Weng, Yang Feng</i>	
Copula-Nested Spectral Kernel Network	48280
<i>Jinyue Tian, Hui Xue, Yanfang Xue, Pengfei Fang</i>	
Boundary Exploration for Bayesian Optimization with Unknown Physical Constraints	48295
<i>Yunsheng Tian, Ane Zuniga, Xinwei Zhang, Johannes P. Dürholt, Payel Das, Jie Chen, Wojciech Matusik, Mina Konakovic Lukovic</i>	
FRAPPÉ: A Group Fairness Framework for Post-Processing Everything.....	48321
<i>Alexandru Tifrea, Preethi Lahoti, Ben Packer, Yoni Halpern, Ahmad Beirami, Flavien Prost</i>	
Faster Maximum Inner Product Search in High Dimensions	48344
<i>Mo Tiwari, Ryan Kang, Je-Yong Lee, Donghyun Lee, Chris Piech, Sebastian Thrun, Ilan Shomorony, Martin Jinye Zhang</i>	

Position: Enforced Amnesia as a Way to Mitigate the Potential Risk of Silent Suffering in the Conscious AI	48362
<i>Yegor Tkachenko</i>	
How Deep Networks Learn Sparse and Hierarchical Data: The Sparse Random Hierarchy Model	48369
<i>Umberto Tomasini, Matthieu Wyard</i>	
Position: Do Not Explain Vision Models Without Context	48390
<i>Paulina Tomaszewska, Przemyslaw Biecek</i>	
An Analysis of Linear Time Series Forecasting Models	48404
<i>William Toner, Luke Darlow</i>	
Neural SPH: Improved Neural Modeling of Lagrangian Fluid Dynamics	48428
<i>Artur P. Toshev, Jonas A. Erbesdobler, Nikolaus A. Adams, Johannes Brandstetter</i>	
Position: Considerations for Differentially Private Learning with Large-Scale Public Pretraining	48453
<i>Florian Tramèr, Gautam Kamath, Nicholas Carlini</i>	
Causal Effect Identification in LiNGAM Models with Latent Confounders.....	48468
<i>Daniele Tramontano, Yaroslav Kivva, Saber Salehkaleybar, Mathias Drton, Negar Kiyavash</i>	
Stereographic Spherical Sliced Wasserstein Distances.....	48494
<i>Huy Tran, Yikun Bai, Abihith Kothapalli, Ashkan Shahbazi, Xinran Liu, Rocio Diaz Martin, Soheil Kolouri</i>	
Inferring the Long-Term Causal Effects of Long-Term Treatments from Short-Term Experiments	48565
<i>Allen Tran, Aurélien Bibaut, Nathan Kallus</i>	
Agent-Specific Effects: A Causal Effect Propagation Analysis in Multi-Agent MDPs.....	48578
<i>Stelios Triantafyllou, Aleksa Sukovic, Debmalya Mandal, Goran Radanovic</i>	
Editing Partially Observable Networks Via Graph Diffusion Models.....	48608
<i>Puja Trivedi, Ryan A. Rossi, David Arbour, Tong Yu, Franck Deroncourt, Sungchul Kim, Nedin Lipka, Namyong Park, Nesreen K. Ahmed, Danai Koutra</i>	
QUIP#: Even Better LLM Quantization with Hadamard Incoherence and Lattice Codebooks	48630
<i>Albert Tseng, Jerry Chee, Qingyao Sun, Volodymyr Kuleshov, Christopher De Sa</i>	
Predictive Linear Online Tracking for Unknown Targets.....	48657
<i>Anastasios Tsiamis, Aren Karapetyan, Yueshan Li, Efe C. Balta, John Lygeros</i>	
Overcoming the Optimizer's Curse: Obtaining Realistic Prescriptions from Neural Networks	48695
<i>Asterios Tsiourvas, Georgia Perakis</i>	

PART 59

Learning Optimal Projection for Forecast Reconciliation of Hierarchical Time Series	48713
<i>Asterios Tsiourvas, Wei Sun, Georgia Perakis, Pin-Yu Chen, Yada Zhu</i>	
Simplicity Bias of Two-Layer Networks Beyond Linearly Separable Data	48728
<i>Nikita Tsoy, Nikola Konstantinov</i>	
Exploration by Optimization with Hybrid Regularizers: Logarithmic Regret with Adversarial Robustness in Partial Monitoring	48768
<i>Taira Tsuchiya, Shinji Ito, Junya Honda</i>	
An Empirical Study into What Matters for Calibrating Vision–Language Models	48791
<i>Weijie Tu, Weijian Deng, Dylan Campbell, Stephen Gould, Tom Gedeon</i>	
Coactive Learning for Large Language Models Using Implicit User Feedback	48809
<i>Aaron D. Tucker, Kianté Brantley, Adam Cahall, Thorsten Joachims</i>	

An Efficient Self-Learning Framework for Interactive Spoken Dialog Systems	48823
<i>Hitesh Tulsiani, David M. Chan, Shalini Ghosh, Garima Lalwani, Prabhat Pandey, Ankish Bansal, Sri Garimella, Ariya Rastrow, Björn Hoffmeister</i>	
ALERT-Transformer: Bridging Asynchronous and Synchronous Machine Learning for Real-Time Event-Based Spatio-Temporal Data	48837
<i>Carmen Martin-Turrero, Maxence Bouvier, Manuel Breitenstein, Pietro Zanuttigh, Vincent Parret</i>	
Matroid Semi-Bandits in Sublinear Time	48855
<i>Ruo-Chun Tzeng, Naoto Ohsaka, Kaito Ariu</i>	
Improving Antibody Humanness Prediction Using Patent Data.....	48878
<i>Talip Uçar, Aubin Ramon, Dino Oglic, Rebecca Croasdale-Wood, Tom Diethe, Pietro Sormanni</i>	
Feedback Efficient Online Fine-Tuning of Diffusion Models	48892
<i>Masatoshi Uehara, Yulai Zhao, Kevin Black, Ehsan Hajiramezani, Gabriele Scalia, Nathaniel Lee Diamant, Alex M Tseng, Sergey Levine, Tommaso Biancalani</i>	
Do Large Language Models Perform the Way People Expect? Measuring the Human Generalization Function	48919
<i>Keyon Vafa, Ashesh Rambachan, Sendhil Mullainathan</i>	
Reward-Free Kernel-Based Reinforcement Learning	48938
<i>Sattar Vakili, Farhang Nabiei, Da-Shan Shiu, Alberto Bernacchia</i>	
Acquisition Conditioned Oracle for Nongreedy Active Feature Acquisition	48957
<i>Michael Valancius, Max Lennon, Junier B. Oliva</i>	
Position: Why Tabular Foundation Models Should Be a Research Priority	48976
<i>Boris Van Breugel, Mihaela Van Der Schaar</i>	
Piecewise Constant and Linear Regression Trees:An Optimal Dynamic Programming Approach.....	48994
<i>Mim Van Den Bos, Jacobus G. M. Van Der Linden, Emir Demirovic</i>	
Algorithm and Hardness for Dynamic Attention Maintenance in Large Language Models.....	49008
<i>Jan Van Den Brand, Zhao Song, Tianyi Zhou</i>	
Proactive DP: A Multiple Target Optimization Framework for DP-SGD	49029
<i>Marten Van Dijk, Nhung V. Nguyen, Toan N. Nguyen, Lam M. Nguyen, Phuong Ha Nguyen</i>	
On the Independence Assumption in Neurosymbolic Learning	49078
<i>Emile Van Krieken, Pasquale Minervini, Edoardo M. Ponti, Antonio Vergari</i>	
When Representations Align: Universality in Representation Learning Dynamics	49098
<i>Loek Van Rossem, Andrew M. Saxe</i>	
Forget Sharpness: Perturbed Forgetting of Model Biases Within SAM Dynamics.....	49122
<i>Ankit Vani, Frederick Tung, Gabriel L. Oliveira, Hossein Sharifi-Noghabi</i>	
Generalized Smooth Variational Inequalities: Methods with Adaptive Stepsizes	49137
<i>Daniil Vankov, Angelia Nedic, Lalitha Sankar</i>	
Discovering Mixtures of Structural Causal Models from Time Series Data	49171
<i>Sumanth Varambally, Yi-An Ma, Rose Yu</i>	
Statistically Optimal Generative Modeling with Maximum Deviation from the Empirical Distribution.....	49203
<i>Elen Vardanyan, Sona Hunanyan, Tigran Galstyan, Arshak Minasyan, Arnak Dalalyan</i>	
Optimization Without Retraction on the Random Generalized Stiefel Manifold.....	49226
<i>Simon Vary, Pierre Ablin, Bin Gao, P.-A. Absil</i>	
A Probabilistic Approach to Learning the Degree of Equivariance in Steerable CNNs.....	49249
<i>Lars Veefkind, Gabriele Cesa</i>	

Stochastic Gradient Flow Dynamics of Test Risk and Its Exact Solution for Weak Features	49310
<i>Rodrigo Veiga, Anastasia Remizova, Nicolas Macris</i>	
Knowledge Transfer from Vision Foundation Models for Efficient Training of Small Task-Specific Models.....	49345
<i>Raviteja Vemulapalli, Hadi Pouransari, Fartash Faghri, Sachin Mehta, Mehrdad Farajtabar, Mohammad Rastegari, Oncel Tuzel</i>	
Code as Reward: Empowering Reinforcement Learning with VLMs	49368
<i>David Venuto, Sami Nur Islam, Martin Klissarov, Doina Precup, Sherry Yang, Ankit Anand</i>	
Topological Neural Networks Go Persistent, Equivariant, and Continuous.....	49388
<i>Yogesh Verma, Amauri H. Souza, Vikas Garg</i>	
CuTS: Customizable Tabular Synthetic Data Generation.....	49408
<i>Mark Vero, Mislav Balunovic, Martin Vechev</i>	
On a Neural Implementation of Brenier's Polar Factorization.....	49434
<i>Nina Vesseron, Marco Cuturi</i>	
To the Max: Reinventing Reward in Reinforcement Learning.....	49455
<i>Grigoriĭ Vevirko, Wendelin Böhmer, Mathijs De Weerd</i>	
Imitation Learning in Discounted Linear MDPs Without Exploration Assumptions	49471
<i>Luca Viano, Stratis Skoulakis, Volkan Cevher</i>	
Position: An Inner Interpretability Framework for AI Inspired by Lessons from Cognitive Neuroscience	49506
<i>Martina G. Vilas, Federico Adolfini, David Poeppel, Gemma Roig</i>	
Position: Will We Run Out of Data? Limits of LLM Scaling Based on Human-Generated Data	49523
<i>Pablo Villalobos, Anson Ho, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, Marius Hobbhahn</i>	

PART 60

ConvNet Vs Transformer, Supervised Vs CLIP: Beyond ImageNet Accuracy	49545
<i>Kirill Vishniakov, Zhiqiang Shen, Zhuang Liu</i>	
CF-OPT: Counterfactual Explanations for Structured Prediction.....	49558
<i>Germain Vivier-Ardisson, Alexandre Forel, Axel Parmentier, Thibaut Vidal</i>	
Parameter Estimation in DAGs from Incomplete Data Via Optimal Transport.....	49580
<i>Vy Vo, Trung Le, Tung-Long Vuong, He Zhao, Edwin V. Bonilla, Dinh Phung</i>	
Optimal Transport for Structure Learning Under Missing Data	49605
<i>Vy Vo, He Zhao, Trung Le, Edwin V. Bonilla, Dinh Phung</i>	
ArtWhisperer: A Dataset for Characterizing Human-AI Interactions in Artistic Creations.....	49627
<i>Kailas Vodrahalli, James Zou</i>	
A Language Model's Guide Through Latent Space.....	49655
<i>Dimitri Von Rütte, Sotiris Anagnostidis, Gregor Bachmann, Thomas Hofmann</i>	
Convergence of Some Convex Message Passing Algorithms to a Fixed Point	49688
<i>Václav Voráček, Tomáš Werner</i>	
Beyond Implicit Bias: The Insignificance of SGD Noise in Online Learning.....	49698
<i>Nikhil Vyas, Depen Morwani, Rosie Zhao, Gal Kaplun, Sham Kakade, Boaz Barak</i>	
Deconstructing the Goldilocks Zone of Neural Network Initialization.....	49717
<i>Artem Vysogorets, Anna Dawid, Julia Kempe</i>	

ConTextual: Evaluating Context-Sensitive Text-Rich Visual Reasoning in Large Multimodal Models.....	49733
<i>Rohan Wadhawan, Hritik Bansal, Kai-Wei Chang, Nanyun Peng</i>	
Just Cluster It: An Approach for Exploration in High-Dimensions Using Clustering and Pre-Trained Representations	49788
<i>Stefan Sylvius Wagner, Stefan Harmeling</i>	
Trustless Audits Without Revealing Data Or Models.....	49808
<i>Suppakit Waiwitlikhit, Ion Stoica, Yi Sun, Tatsunori Hashimoto, Daniel Kang</i>	
Log Neural Controlled Differential Equations: The Lie Brackets Make a Difference	49822
<i>Benjamin Walker, Andrew D. McLeod, Tiexin Qin, Yichuan Cheng, Haoliang Li, Terry Lyons</i>	
Implicit Compressibility of Overparametrized Neural Networks Trained with Heavy-Tailed SGD	49845
<i>Yijun Wan, Melih Barsbey, Abdellatif Zaidi, Umut Simsekli</i>	
SeMOPO: Learning High-Quality Model and Policy from Low-Quality Offline Visual Datasets.....	49867
<i>Shenghua Wan, Ziyuan Chen, Le Gan, Shuai Feng, De-Chuan Zhan</i>	
AlphaZero-Like Tree-Search Can Guide Large Language Model Decoding and Training	49890
<i>Ziyu Wan, Xidong Feng, Muning Wen, Stephen Marcus McAleer, Ying Wen, Weinan Zhang, Jun Wang</i>	
VinT-6D: A Large-Scale Object-In-Hand Dataset from Vision, Touch and Proprioception	49921
<i>Zhaoliang Wan, Yonggen Ling, Senlin Yi, Lu Qi, Wangwei Lee, Minglei Lu, Sicheng Yang, Xiao Teng, Peng Lu, Xu Yang, Ming-Hsuan Yang, Hui Cheng</i>	
Decouple Then Classify: A Dynamic Multi-View Labeling Strategy with Shared and Specific Information.....	49941
<i>Xinhang Wan, Jiyuan Liu, Xinwang Liu, Yi Wen, Hao Yu, Siwei Wang, Shengju Yu, Tianjiao Wan, Jun Wang, En Zhu</i>	
Superpoint Gaussian Splatting for Real-Time High-Fidelity Dynamic Scene Reconstruction	49957
<i>Diwen Wan, Ruijie Lu, Gang Zeng</i>	
S3GCL: Spectral, Swift, Spatial Graph Contrastive Learning	49973
<i>Guancheng Wan, Yijun Tian, Wenke Huang, Nitesh V Chawla, Mang Ye</i>	
Non-Stationary Online Convex Optimization with Arbitrary Delays.....	49991
<i>Yuanyu Wan, Chang Yao, Mingli Song, Lijun Zhang</i>	
Towards Unified Multi-Granularity Text Detection with Interactive Attention.....	50012
<i>Xingyu Wan, Chengquan Zhang, Pengyuan Lyu, Sen Fan, Zihan Ni, Kun Yao, Errui Ding, Jingdong Wang</i>	
Understanding Reasoning Ability of Language Models from the Perspective of Reasoning Paths Aggregation.....	50026
<i>Xinyi Wang, Alfonso Amayuelas, Kexun Zhang, Liangming Pan, Wenhui Chen, William Yang Wang</i>	
One Prompt is Not Enough: Automated Construction of a Mixture-Of-Expert Prompts	50043
<i>Ruochen Wang, Sohyun An, Minhao Cheng, Tianyi Zhou, Sung Ju Hwang, Cho-Jui Hsieh</i>	
On Universally Optimal Algorithms for A/B Testing	50065
<i>Po-An Wang, Kaito Ariu, Alexandre Proutiere</i>	
Adversarially Robust Hypothesis Transfer Learning.....	50092
<i>Yunjuan Wang, Raman Arora</i>	
Towards Theoretical Understanding of Learning Large-Scale Dependent Data Via Random Features	50118
<i>Chao Wang, Xin Bing, Xin He, Caixing Wang</i>	
Non-Parametric Online Change Point Detection on Riemannian Manifolds	50143
<i>Xiuheng Wang, Ricardo Augusto Borsoi, Cédric Richard</i>	

A Circuit Domain Generalization Framework for Efficient Logic Synthesis in Chip Design.....	50163
<i>Zhihai Wang, Lei Chen, Jie Wang, Yinqi Bai, Xing Li, Xijun Li, Mingxuan Yuan, Jianye Hao, Yongdong Zhang, Feng Wu</i>	
Executable Code Actions Elicit Better LLM Agents	50208
<i>Xingyao Wang, Yangyi Chen, Lifan Yuan, Yizhe Zhang, Yunzhu Li, Hao Peng, Heng Ji</i>	
Revisiting the Power of Prompt for Visual Tuning.....	50233
<i>Yuzhu Wang, Lechao Cheng, Chaowei Fang, Dingwen Zhang, Manni Duan, Meng Wang</i>	
TVE: Learning Meta-Attribution for Transferable Vision Explainer	50248
<i>Guanchu Wang, Yu-Neng Chuang, Fan Yang, Mengnan Du, Chia-Yuan Chang, Shaochen Zhong, Zirui Liu, Zhaozhuo Xu, Kaixiong Zhou, Xuanting Cai, Xia Hu</i>	
Localizing Task Information for Improved Model Merging and Compression.....	50268
<i>Ke Wang, Nikolaos Dimitriadis, Guillermo Ortiz-Jiménez, François Fleuret, Pascal Frossard</i>	
Adaptively Learning to Select-Rank in Online Platforms	50288
<i>Jingyuan Wang, Perry Dong, Ying Jin, Ruohan Zhan, Zhengyuan Zhou</i>	
Imitation Learning from Purified Demonstrations	50313
<i>Yunke Wang, Minjing Dong, Yukun Zhao, Bo Du, Chang Xu</i>	
Diagnosing the Compositional Knowledge of Vision Language Models from a Game-Theoretic View	50332
<i>Jin Wang, Shichao Dong, Yapeng Zhu, Kelu Yao, Weidong Zhao, Chao Li, Ping Luo</i>	
An Efficient Maximal Ancestral Graph Listing Algorithm	50353
<i>Tian-Zuo Wang, Wen-Bo Du, Zhi-Hua Zhou</i>	
Monotone, Bi-Lipschitz, and Polyak-Łojasiewicz Networks	50379
<i>Ruigang Wang, Krishnamurthy (Dj) Dvijotham, Ian R. Manchester</i>	

PART 61

Swallowing the Bitter Pill: Simplified Scalable Conformer Generation	50400
<i>Yuyang Wang, Ahmed A. Elhag, Navdeep Jaitly, Joshua M. Susskind, Miguel Ángel Bautista</i>	
Optimal Kernel Quantile Learning with Random Features	50419
<i>Caixing Wang, Xingdong Feng</i>	
MemoryLLM: Towards Self-Updatable Large Language Models	50453
<i>Yu Wang, Yifan Gao, Xiusi Chen, Haoming Jiang, Shiyang Li, Jingfeng Yang, Qingyu Yin, Zheng Li, Xian Li, Bing Yin, Jingbo Shang, Julian McAuley</i>	
Identification and Estimation for Nonignorable Missing Data: A Data Fusion Approach.....	50467
<i>Zixiao Wang, Amiremad Ghassami, Ilya Shpitser</i>	
Understanding Heterophily for Graph Neural Networks.....	50489
<i>Junfu Wang, Yuanfang Guo, Liang Yang, Yunhong Wang</i>	
Momentum for the Win: Collaborative Federated Reinforcement Learning Across Heterogeneous Environments.....	50530
<i>Han Wang, Sihong He, Zhili Zhang, Fei Miao, James Anderson</i>	
EvGGS: A Collaborative Learning Framework for Event-Based Generalizable Gaussian Splatting.....	50561
<i>Jiaxu Wang, Junhao He, Ziyi Zhang, Mingyuan Sun, Jingkai Sun, Renjing Xu</i>	
Mollification Effects of Policy Gradient Methods.....	50580
<i>Tao Wang, Sylvia Herbert, Sicun Gao</i>	
Discovering Symmetry Breaking in Physical Systems with Relaxed Group Convolution.....	50599
<i>Rui Wang, Elyssa Hofgard, Han Gao, Robin Walters, Tess E. Smidt</i>	

SciBench: Evaluating College-Level Scientific Problem-Solving Abilities of Large Language Models.....	50622
<i>Xiaoxuan Wang, Ziniu Hu, Pan Lu, Yanqiao Zhu, Jieyu Zhang, Satyen Subramaniam, Arjun R. Loomba, Shichang Zhang, Yizhou Sun, Wei Wang</i>	
Optimal Kernel Choice for Score Function-Based Causal Discovery.....	50650
<i>Wenjie Wang, Biwei Huang, Feng Liu, Xinge You, Tongliang Liu, Kun Zhang, Mingming Gong</i>	
Rapid Learning Without Catastrophic Forgetting in the Morris Water Maze.....	50669
<i>Raymond L Wang, Jaedong Hwang, Akhilan Boopathy, Ila R Fiete</i>	
Learning with Complementary Labels Revisited: The Selected-Completely-At-Random Setting is More Practical.....	50683
<i>Wei Wang, Takashi Ishida, Yu-Jie Zhang, Gang Niu, Masashi Sugiyama</i>	
Total Variation Floodgate for Variable Importance Inference in Classification.....	50711
<i>Wenshuo Wang, Lucas Janson, Lihua Lei, Aaditya Ramdas</i>	
In-Context Learning on Function Classes Unveiled for Transformers.....	50726
<i>Zhijie Wang, Bo Jiang, Shuai Li</i>	
MS ³ D: A RG Flow-Based Regularization for GAN Training with Limited Data.....	50746
<i>Jian Wang, Xin Lan, Yuxin Tian, Jiancheng Lv</i>	
StableSSM: Alleviating the Curse of Memory in State-Space Models Through Stable Reparameterization.....	50766
<i>Shida Wang, Qianxiao Li</i>	
Bootstrap AutoEncoders with Contrastive Paradigm for Self-Supervised Gaze Estimation.....	50794
<i>Yaoming Wang, Jin Li, Wenrui Dai, Bowen Shi, Xiaopeng Zhang, Chenglin Li, Hongkai Xiong</i>	
Highway Value Iteration Networks.....	50807
<i>Yuhui Wang, Weida Li, Francesco Faccio, Qingyuan Wu, Jürgen Schmidhuber</i>	
Improving Generalization in Offline Reinforcement Learning Via Adversarial Data Splitting.....	50822
<i>Da Wang, Lin Li, Wei Wei, Qixian Yu, Jianye Hao, Jiye Liang</i>	
Visual Transformer with Differentiable Channel Selection: An Information Bottleneck Inspired Approach.....	50840
<i>Yancheng Wang, Ping Li, Yingzhen Yang</i>	
An Iterative Min-Min Optimization Method for Sparse Bayesian Learning.....	50859
<i>Yasen Wang, Junlin Li, Zuogong Yue, Ye Yuan</i>	
Graph as Point Set.....	50874
<i>Xiyuan Wang, Pan Li, Muhan Zhang</i>	
Open Ad Hoc Teamwork with Cooperative Game Theory.....	50902
<i>Jianhong Wang, Yang Li, Yuan Zhang, Wei Pan, Samuel Kaski</i>	
Connecting the Dots: Collaborative Fine-Tuning for Black-Box Vision-Language Models.....	50931
<i>Zhengbo Wang, Jian Liang, Ran He, Zilei Wang, Tieniu Tan</i>	
Bridging Data Gaps in Diffusion Models with Adversarial Noise-Based Transfer Learning.....	50944
<i>Xiyu Wang, Baijiong Lin, Daochang Liu, Ying-Cong Chen, Chang Xu</i>	
Helpful Or Harmful Data? Fine-Tuning-Free Shapley Attribution for Explaining Language Model Predictions.....	50960
<i>Jingtang Wang, Xiaoqiang Lin, Rui Qiao, Chuan-Sheng Foo, Bryan Kian Hsiang Low</i>	
On Discrete Prompt Optimization for Diffusion Models.....	50992
<i>Ruo Chen Wang, Ting Liu, Cho-Jui Hsieh, Boqing Gong</i>	
A Global Geometric Analysis of Maximal Coding Rate Reduction.....	51012
<i>Peng Wang, Huikang Liu, Druv Pai, Yaodong Yu, Zhihui Zhu, Qing Qu, Yi Ma</i>	

EfficientZero V2: Mastering Discrete and Continuous Control with Limited Data	51041
<i>Shengjie Wang, Shaohuai Liu, Weirui Ye, Jiacheng You, Yang Gao</i>	
A Dual-Module Framework for Counterfactual Estimation Over Time	51063
<i>Xin Wang, Shengfei Lyu, Lishan Yang, Yibing Zhan, Huanhuan Chen</i>	
MC-GTA: Metric-Constrained Model-Based Clustering Using Goodness-Of-Fit Tests with Autocorrelations	51086
<i>Zhangyu Wang, Gengchen Mai, Krzysztof Janowicz, Ni Lao</i>	
EvoluNet: Advancing Dynamic Non-IID Transfer Learning on Graphs	51105
<i>Haohui Wang, Yuzhen Mao, Yujun Yan, Yaoqing Yang, Jianhui Sun, Kevin Choi, Balaji Veeramani, Alison Hu, Edward Bowen, Tyler Cody, Dawei Zhou</i>	
Achieving Margin Maximization Exponentially Fast Via Progressive Norm Rescaling.....	51124
<i>Mingze Wang, Zeping Min, Lei Wu</i>	
Transforming and Combining Rewards for Aligning Large Language Models	51161
<i>Zihao Wang, Chirag Nagpal, Jonathan Berant, Jacob Eisenstein, Alex D'Amour, Sanmi Koyejo, Victor Veitch</i>	
TroVE: Inducing Verifiable and Efficient Toolboxes for Solving Programmatic Tasks.....	51177
<i>Zora Zhiruo Wang, Graham Neubig, Daniel Fried</i>	
More Benefits of Being Distributional: Second-Order Bounds for Reinforcement Learning	51192
<i>Kaiwen Wang, Owen Oertell, Alekh Agarwal, Nathan Kallus, Wen Sun</i>	
Pi-DUAL: Using Privileged Information to Distinguish Clean from Noisy Labels.....	51214
<i>Ke Wang, Guillermo Ortiz-Jimenez, Rodolphe Jenatton, Mark Collier, Efi Kokiopoulou, Pascal Frossard</i>	

PART 62

Sample Average Approximation for Conditional Stochastic Optimization with Dependent Data.....	51237
<i>Yafei Wang, Bo Pan, Mei Li, Jianya Lu, Lingcheng Kong, Bei Jiang, Linglong Kong</i>	
InstructRetro: Instruction Tuning Post Retrieval-Augmented Pretraining.....	51255
<i>Boxin Wang, Wei Ping, Lawrence McAfee, Peng Xu, Bo Li, Mohammad Shoeybi, Bryan Catanzaro</i>	
A Fine-Grained Analysis of Fitted Q-Evaluation: Beyond Parametric Models.....	51273
<i>Jiayi Wang, Zhengling Qi, Raymond K. W. Wong</i>	
Probabilistic Constrained Reinforcement Learning with Formal Interpretability	51303
<i>Yanran Wang, Qiuchen Qian, David Boyle</i>	
Efficient Online Set-Valued Classification with Bandit Feedback	51328
<i>Zhou Wang, Xingye Qiao</i>	
LLM-Empowered State Representation for Reinforcement Learning.....	51348
<i>Boyuan Wang, Yun Qu, Yuhang Jiang, Jianzhun Shao, Chang Liu, Wenming Yang, Xiangyang Ji</i>	
Proteus: Exploring Protein Structure Generation for Enhanced Designability and Efficiency.....	51376
<i>Chentong Wang, Yannan Qu, Zhangzhi Peng, Yukai Wang, Hongli Zhu, Dachuan Chen, Longxing Cao</i>	
How to Trace Latent Generative Model Generated Images Without Artificial Watermark?	51396
<i>Zhenting Wang, Vikash Sehwal, Chen Chen, Lingjuan Lyu, Dimitris N. Metaxas, Shiqing Ma</i>	
Distributed High-Dimensional Quantile Regression: Estimation Efficiency and Support Recovery.....	51415
<i>Caixing Wang, Ziliang Shen</i>	
Generalization Analysis of Stochastic Weight Averaging with General Sampling.....	51442
<i>Peng Wang, Li Shen, Zerui Tao, Shuaida He, Dacheng Tao</i>	

The Stronger the Diffusion Model, the Easier the Backdoor: Data Poisoning to Induce Copyright Breaches Without Adjusting Finetuning Pipeline.....	51465
<i>Haonan Wang, Qianli Shen, Yao Tong, Yang Zhang, Kenji Kawaguchi</i>	
RL-VLM-F: Reinforcement Learning from Vision Language Foundation Model Feedback.....	51484
<i>Yufei Wang, Zhanyi Sun, Jesse Zhang, Zhou Xian, Erdem Biyik, David Held, Zackory Erickson</i>	
Probabilistic Conceptual Explainers: Trustworthy Conceptual Explanations for Vision Foundation Models.....	51502
<i>Hengyi Wang, Shiwei Tan, Hao Wang</i>	
Integrated Hardware Architecture and Device Placement Search.....	51523
<i>Irene Wang, Jakub Tarnawski, Amar Phanishayee, Divya Mahajan</i>	
AD3: Implicit Action is the Key for World Models to Distinguish the Diverse Visual Distractors.....	51546
<i>Yucen Wang, Shenghua Wan, Le Gan, Shuai Feng, De-Chuan Zhan</i>	
Benchmarking Deletion Metrics with the Principled Explanations.....	51569
<i>Yipei Wang, Xiaoqian Wang</i>	
CW Complex Hypothesis for Image Data.....	51596
<i>Yi Wang, Zhiren Wang</i>	
Optimal Exact Recovery in Semi-Supervised Learning: A Study of Spectral Methods and Graph Convolutional Networks.....	51614
<i>Haixiao Wang, Zhichao Wang</i>	
A New Theoretical Perspective on Data Heterogeneity in Federated Optimization.....	51650
<i>Jiayi Wang, Shiqiang Wang, Rong-Rong Chen, Mingyue Ji</i>	
FADAS: Towards Federated Adaptive Asynchronous Optimization.....	51701
<i>Yujia Wang, Shiqiang Wang, Songtao Lu, Jinghui Chen</i>	
Open-Vocabulary Calibration for Fine-Tuned CLIP.....	51734
<i>Shuoyuan Wang, Jindong Wang, Guoqing Wang, Bob Zhang, Kaiyang Zhou, Hongxin Wei</i>	
Probabilistic Subgoal Representations for Hierarchical Reinforcement Learning.....	51755
<i>Vivienne Huiling Wang, Tinghuai Wang, Wenyan Yang, Joni-Kristian Kämäräinen, Joni Pajarinen</i>	
Non-Asymptotic Analysis for Single-Loop (Natural) Actor-Critic with Compatible Function Approximation.....	51771
<i>Yudan Wang, Yue Wang, Yi Zhou, Shaofeng Zou</i>	
A Hierarchical Adaptive Multi-Task Reinforcement Learning Framework for Multiplier Circuit Design.....	51825
<i>Zhihai Wang, Jie Wang, Dongsheng Zuo, Yunjie Ji, Xilin Xia, Yuzhe Ma, Jianye Hao, Mingxuan Yuan, Yongdong Zhang, Feng Wu</i>	
Transformers Provably Learn Sparse Token Selection While Fully-Connected Nets Cannot.....	51854
<i>Zixuan Wang, Stanley Wei, Daniel Hsu, Jason D. Lee</i>	
Defense Against Model Extraction Attack by Bayesian Active Watermarking.....	51913
<i>Zhenyi Wang, Yihan Wu, Heng Huang</i>	
RoboGen: Towards Unleashing Infinite Data for Automated Robot Learning Via Generative Simulation.....	51936
<i>Yufei Wang, Zhou Xian, Feng Chen, Tsun-Hsuan Wang, Yian Wang, Katerina Fragkiadaki, Zackory Erickson, David Held, Chuang Gan</i>	
Quality-Diversity with Limited Resources.....	51984
<i>Ren-Jian Wang, Ke Xue, Cong Guan, Chao Qian</i>	
Probability Distribution of Hypervolume Improvement in Bi-Objective Bayesian Optimization.....	52002
<i>Hao Wang, Kaifeng Yang, Michael Affenzeller</i>	

Vision Transformers as Probabilistic Expansion from LearnGene.....	52019
<i>Qiufeng Wang, Xu Yang, Haokun Chen, Xin Geng</i>	
Rethinking Data Shapley for Data Selection Tasks: Misleads and Merits.....	52033
<i>Jiachen T. Wang, Tianji Yang, James Zou, Yongchan Kwon, Ruoxi Jia</i>	
Can Gaussian Sketching Converge Faster on a Preconditioned Landscape?.....	52064
<i>Yilong Wang, Haishan Ye, Guang Dai, Ivor W. Tsang</i>	

PART 63

Autaptic Synaptic Circuit Enhances Spatio-Temporal Predictive Learning of Spiking Neural Networks	52083
<i>Lihao Wang, Zhaofei Yu</i>	
Learning with Adaptive Resource Allocation.....	52099
<i>Jing Wang, Miao Yu, Peng Zhao, Zhi-Hua Zhou</i>	
Taylor Videos for Action Recognition.....	52117
<i>Lei Wang, Xiuyuan Yuan, Tom Gedeon, Liang Zheng</i>	
An Empirical Study of Realized GNN Expressiveness	52134
<i>Yanbo Wang, Muhan Zhang</i>	
Calibration Bottleneck: Over-Compressed Representations Are Less Calibratable	52156
<i>Deng-Bao Wang, Min-Ling Zhang</i>	
Benign Overfitting in Adversarial Training of Neural Networks	52171
<i>Yunjuan Wang, Kaibo Zhang, Raman Arora</i>	
FreeBind: Free Lunch in Unified Multimodal Space Via Knowledge Fusion.....	52233
<i>Zehan Wang, Ziang Zhang, Xize Cheng, Rongjie Huang, Luping Liu, Zhenhui Ye, Haifeng Huang, Yang Zhao, Tao Jin, Peng Gao, Zhou Zhao</i>	
Exploring Intrinsic Dimension for Vision-Language Model Pruning.....	52247
<i>Hanzhang Wang, Jiawen Zhang, Qingyuan Ma</i>	
Knowledge-Aware Reinforced Language Models for Protein Directed Evolution	52260
<i>Yuhao Wang, Qiang Zhang, Ming Qin, Xiang Zhuang, Xiaotong Li, Zhichen Gong, Zeyuan Wang, Yu Zhao, Jianhua Yao, Keyan Ding, Huajun Chen</i>	
Boximator: Generating Rich and Controllable Motions for Video Synthesis.....	52274
<i>Jiawei Wang, Yuchen Zhang, Jiaxin Zou, Yan Zeng, Guoqiang Wei, Liping Yuan, Hang Li</i>	
Bridging Model Heterogeneity in Federated Learning Via Uncertainty-Based Asymmetrical Reciprocity Learning	52290
<i>Jiaqi Wang, Chenxu Zhao, Lingjuan Lyu, Quanzeng You, Mengdi Huai, Fenglong Ma</i>	
Diffusion Language Models Are Versatile Protein Learners	52309
<i>Xinyou Wang, Zaixiang Zheng, Fei Ye, Dongyu Xue, Shujian Huang, Quanquan Gu</i>	
Neural Collapse Meets Differential Privacy: Curious Behaviors of NoisyGD with Near-Perfect Representation Learning.....	52334
<i>Chendi Wang, Yuqing Zhu, Weijie J. Su, Yu-Xiang Wang</i>	
Batch Singular Value Polarization and Weighted Semantic Augmentation for Universal Domain Adaptation	52361
<i>Ziqi Wang, Wei Wang, Chao Huang, Jie Wen, Cong Wang</i>	
Mapping the Multiverse of Latent Representations.....	52372
<i>Jeremy Wayland, Corinna Coupette, Bastian Rieck</i>	
Reinforcement Learning and Regret Bounds for Admission Control.....	52403
<i>Lucas Weber, Ana Bušić, Jiamin Zhu</i>	

Learning to Compile Programs to Neural Networks	52428
<i>Logan Weber, Jesse Michel, Alex Renda, Michael Carbin</i>	
Exact Soft Analytical Side-Channel Attacks Using Tractable Circuits.....	52472
<i>Thomas Wedenig, Rishub Nagpal, Gaëtan Cassiers, Stefan Mangard, Robert Peharz</i>	
Exploiting Human-AI Dependence for Learning to Defer	52484
<i>Zixi Wei, Yuzhou Cao, Lei Feng</i>	
Learning Pseudo-Contractive Denoisers for Inverse Problems	52500
<i>Deliang Wei, Peng Chen, Fang Li</i>	
Rethinking Generative Large Language Model Evaluation for Semantic Comprehension.....	52525
<i>Fangyun Wei, Xi Chen, Lin Luo</i>	
MMPareto: Boosting Multimodal Learning with Innocent Unimodal Assistance.....	52559
<i>Yake Wei, Di Hu</i>	
Task Groupings Regularization:Data-Free Meta-Learning with Heterogeneous Pre-Trained Models	52573
<i>Yongxian Wei, Zixuan Hu, Li Shen, Zhenyi Wang, Yu Li, Chun Yuan, Dacheng Tao</i>	
Assessing the Brittleness of Safety Alignment Via Pruning and Low-Rank Modifications	52588
<i>Boyi Wei, Kaixuan Huang, Yangsibo Huang, Tinghao Xie, Xiangyu Qi, Mengzhou Xia, Prateek Mittal, Mengdi Wang, Peter Henderson</i>	
Learning Label Shift Correction for Test-Agnostic Long-Tailed Recognition	52611
<i>Tong Wei, Zhen Mao, Zi-Hao Zhou, Yuanyu Wan, Min-Ling Zhang</i>	
Magocoder: Empowering Code Generation with OSS-Instruct	52632
<i>Yuxiang Wei, Zhe Wang, Jiawei Liu, Yifeng Ding, Lingming Zhang</i>	
Extending Test-Time Augmentation with Metamorphic Relations for Combinatorial Problems	52658
<i>Siwei Wei, Xudong Zhang, Zhiyang Zhou, Yan Cai</i>	
Position: AI/ML Influencers Have a Place in the Academic Process	52680
<i>Iain Xie Weissburg, Mehir Arora, Xinyi Wang, Liangming Pan, William Yang Wang</i>	
SiT: Symmetry-Invariant Transformers for Generalisation in Reinforcement Learning	52695
<i>Matthias Weissenbacher, Rishabh Agarwal, Yoshinobu Kawahara</i>	
Contrastive Representation for Data Filtering in Cross-Domain Offline Reinforcement Learning	52720
<i>Xiaoyu Wen, Chenjia Bai, Kang Xu, Xudong Yu, Yang Zhang, Xuelong Li, Zhen Wang</i>	
Cross-Domain Open-World Discovery.....	52744
<i>Shuo Wen, Maria Brbic</i>	
Diffusion-Based Missing-View Generation with the Application on Incomplete Multi-View Clustering	52762
<i>Jie Wen, Shijie Deng, Waikeng Wong, Guoqing Chao, Chao Huang, Lunke Fei, Yong Xu</i>	
Which Frequencies Do CNNs Need? Emergent Bottleneck Structure in Feature Learning.....	52779
<i>Yuxiao Wen, Arthur Jacot</i>	
From Coarse to Fine: Enable Comprehensive Graph Self-Supervised Learning with Multi-Granular Semantic Ensemble	52801
<i>Qianlong Wen, Mingxuan Ju, Zhongyu Ouyang, Chuxu Zhang, Yanfang Ye</i>	
Provable Contrastive Continual Learning	52819
<i>Yichen Wen, Zhiquan Tan, Kaipeng Zheng, Chuanlong Xie, Weiran Huang</i>	
Distributional Bellman Operators Over Mean Embeddings.....	52839
<i>Li Kevin Wenliang, Grégoire Delétang, Matthew Aitchison, Marcus Hutter, Anian Ruoss, Arthur Gretton, Mark Rowland</i>	
On the Asymptotic Distribution of the Minimum Empirical Risk.....	52869
<i>Jacob Westerhout, Trungtin Nguyen, Xin Guo, Hien Duy Nguyen</i>	

Stability-Informed Initialization of Neural Ordinary Differential Equations 52903
Theodor Westny, Arman Mohammadi, Daniel Jung, Erik Frisk

PART 64

QuRating: Selecting High-Quality Data for Training Language Models 52915
Alexander Wettig, Aatmik Gupta, Saumya Malik, Danqi Chen

Multiply Robust Estimation for Local Distribution Shifts with Multiple Domains 52972
Steven Wilkins-Reeves, Xu Chen, Qi Ma, Christine Agarwal, Aude Hofleitner

A Distributional Analogue to the Successor Representation 52994
Harley Wiltzer, Jesse Farebrother, Arthur Gretton, Yunhao Tang, André Barreto, Will Dabney, Marc G. Bellemare, Mark Rowland

Bridging Discrete and Continuous State spaces: Exploring the Ehrenfest Process in Time-Continuous Diffusion Models 53017
Ludwig Winkler, Lorenz Richter, Manfred Oppen

Fine-Tuning Reinforcement Learning Models is Secretly a Forgetting Mitigation Problem 53039
Maciej Wolczyk, Bartłomiej Cupiał, Mateusz Ostaszewski, Michał Bortkiewicz, Michał Zajac, Razvan Pascanu, Łukasz Kuciński, Piotr Miłoś

Fundamental Limitations of Alignment in Large Language Models 53079
Yotam Wolf, Noam Wies, Oshri Avnery, Yoav Levine, Amnon Shashua

Expressivity and Generalization: Fragment-Biases for Molecular GNNs 53113
Tom Wollschläger, Niklas Kemper, Leon Hetzel, Johanna Sommer, Stephan Günnemann

Unified Training of Universal Time Series Forecasting Transformers 53140
Gerald Woo, Chenghao Liu, Akshat Kumar, Caiming Xiong, Silvio Savarese, Doyen Sahoo

Federated Offline Reinforcement Learning: Collaborative Single-Policy Coverage Suffices 53165
Jiin Woo, Laixi Shi, Gauri Joshi, Yuejie Chi

Coresets for Multiple ℓ_p Regression 53202
David P. Woodruff, Taisuke Yasuda

Reweighted Solutions for Weighted Low Rank Approximation 53234
David P. Woodruff, Taisuke Yasuda

Optimizing Watermarks for Large Language Models 53251
Bram Wouters

Repoformer: Selective Retrieval for Repository-Level Code Completion 53270
Di Wu, Wasi Uddin Ahmad, Dejiao Zhang, Murali Krishna Ramanathan, Xiaofei Ma

Theoretical Insights for Diffusion Guidance: A Case Study for Gaussian Mixture Models 53291
Yuchen Wu, Minshuo Chen, Zihao Li, Mengdi Wang, Yuting Wei

Adaptive Accompaniment with RealChords 53328
Yusong Wu, Tim Cooijmans, Kyle Kastner, Adam Roberts, Ian Simon, Alexander Scarlatos, Chris Donahue, Cassie Tarakajian, Shayegan Omidshafiei, Aaron Courville, Pablo Samuel Castro, Natasha Jaques, Cheng-Zhi Anna Huang

Ditto: Quantization-Aware Secure Inference of Transformers Upon MPC 53346
Haoqi Wu, Wenjing Fang, Yancheng Zheng, Junming Ma, Jin Tan, Lei Wang

NExT-GPT: Any-To-Any Multimodal LLM 53366
Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, Tat-Seng Chua

Understanding Stochastic Natural Gradient Variational Inference 53398
Kaiwen Wu, Jacob R. Gardner

FAFE: Immune Complex Modeling with Geodesic Distance Loss on Noisy Group Frames.....	53422
<i>Ruidong Wu, Ruihan Guo, Rui Wang, Shitong Luo, Yue Xu, Jiahao Li, Jianzhu Ma, Qiang Liu, Yunan Luo, Jian Peng</i>	
A Resilient and Accessible Distribution-Preserving Watermark for Large Language Models	53443
<i>Yihan Wu, Zhengmian Hu, Junfeng Guo, Hongyang Zhang, Heng Huang</i>	
Uniform Memory Retrieval with Larger Capacity for Modern Hopfield Models	53471
<i>Dennis Wu, Jerry Yao-Chieh Hu, Teng-Yun Hsiao, Han Liu</i>	
Planning, Fast and Slow: Online Reinforcement Learning with Action-Free Offline Data Via Multiscale Planners	53515
<i>Chengjie Wu, Hao Hu, Yiqin Yang, Ning Zhang, Chongjie Zhang</i>	
PointMC: Multi-Instance Point Cloud Registration Based on Maximal Cliques	53542
<i>Yue Wu, Xidao Hu, Yongzhe Yuan, Xiaolong Fan, Maoguo Gong, Hao Li, Mingyang Zhang, Qiguang Miao, Wenping Ma</i>	
Evaluating and Analyzing Relationship Hallucinations in Large Vision-Language Models	53553
<i>Mingrui Wu, Jiayi Ji, Oucheng Huang, Jiale Li, Yuhang Wu, Xiaoshuai Sun, Rongrong Ji</i>	
Borda Regret Minimization for Generalized Linear Dueling Bandits	53571
<i>Yue Wu, Tao Jin, Qiwei Di, Hao Lou, Farzad Farnoud, Quanquan Gu</i>	
DISCRET: Synthesizing Faithful Explanations for Treatment Effect Estimation	53597
<i>Yinjun Wu, Mayank Keoliya, Kan Chen, Neelay Velingker, Ziyang Li, Emily J Getzen, Qi Long, Mayur Naik, Ravi B Parikh, Eric Wong</i>	
Surface-VQMAE: Vector-Quantized Masked Auto-Encoders on Molecular Surfaces.....	53619
<i>Fang Wu, Stan Z. Li</i>	
Learning Causal Relations from Subsampled Time Series with Two Time-Slices	53635
<i>Anpeng Wu, Haoxuan Li, Kun Kuang, Keli Zhang, Fei Wu</i>	
AND: Audio Network Dissection for Interpreting Deep Acoustic Models	53656
<i>Tung-Yu Wu, Yu-Xiang Lin, Tsui-Wei Weng</i>	
Transolver: A Fast Transformer Solver for PDEs on General Geometries	53681
<i>Haixu Wu, Huakun Luo, Haowen Wang, Jianmin Wang, Mingsheng Long</i>	
Confidence-Aware Contrastive Learning for Selective Classification	53706
<i>Yu-Chang Wu, Shen-Huan Lyu, Haopu Shang, Xiangyu Wang, Chao Qian</i>	
Minimally Modifying a Markov Game to Achieve Any Nash Equilibrium and Value.....	53730
<i>Young Wu, Jeremy McMahan, Yiding Chen, Yudong Chen, Xiaojin Zhu, Qiaomin Xie</i>	
VoroNav: Voronoi-Based Zero-Shot Object Navigation with Large Language Model	53757
<i>Pengying Wu, Yao Mu, Bingxian Wu, Yi Hou, Ji Ma, Shanghang Zhang, Chang Liu</i>	

PART 65

Learning Divergence Fields for Shift-Robust Graph Representations	53776
<i>Qitian Wu, Fan Nie, Chenxiao Yang, Junchi Yan</i>	
Profile Reconstruction from Private Sketches	53793
<i>Hao Wu, Rasmus Pagh</i>	
Policy Learning for Balancing Short-Term and Long-Term Rewards.....	53817
<i>Peng Wu, Ziyu Shen, Feng Xie, Zhongyao Wang, Chunchen Liu, Yan Zeng</i>	
Learning to Predict Mutational Effects of Protein-Protein Interactions by Microenvironment-Aware Hierarchical Prompt Learning	53847
<i>Lirong Wu, Yijun Tian, Haitao Lin, Yufei Huang, Siyuan Li, Nitesh V Chawla, Stan Z. Li</i>	

A Theory of Fault-Tolerant Learning.....	53860
<i>Changlong Wu, Yifan Wang, Ananth Grama</i>	
Prometheus: Out-Of-Distribution Fluid Dynamics Modeling with Disentangled Graph ODE	53870
<i>Hao Wu, Huiyuan Wang, Kun Wang, Weiyan Wang, Changan Ye, Yangyu Tao, Chong Chen, Xian-Sheng Hua, Xiao Luo</i>	
Mitigating Catastrophic Forgetting in Online Continual Learning by Modeling Previous Task Interrelations Via Pareto Optimization	53892
<i>Yichen Wu, Hong Wang, Peilin Zhao, Yefeng Zheng, Ying Wei, Long-Kai Huang</i>	
Detecting Any Instruction-To-Answer Interaction Relationship: Universal Instruction-To-Answer Navigator for Med-VQA	53909
<i>Zhongze Wu, Hongyan Xu, Yitian Long, Shan You, Xiu Su, Jun Long, Yueyi Luo, Chang Xu</i>	
Unraveling the Impact of Heterophilic Structures on Graph Positive-Unlabeled Learning	53928
<i>Yuhao Wu, Jiangchao Yao, Bo Han, Lina Yao, Tongliang Liu</i>	
Mitigating Label Noise on Graphs Via Topological Sample Selection	53944
<i>Yuhao Wu, Jiangchao Yao, Xiaobo Xia, Jun Yu, Ruxin Wang, Bo Han, Tongliang Liu</i>	
Boosting Reinforcement Learning with Strongly Delayed Feedback Through Auxiliary Short Delays.....	53973
<i>Qingyuan Wu, Simon Sinong Zhan, Yixuan Wang, Yuhui Wang, Chung-Wei Lin, Chen Lv, Qi Zhu, Jürgen Schmidhuber, Chao Huang</i>	
HGCN2SP: Hierarchical Graph Convolutional Network for Two-Stage Stochastic Programming	53999
<i>Yang Wu, Yifan Zhang, Zhenxing Liang, Jian Cheng</i>	
Q-Align: Teaching LMMs for Visual Scoring Via Discrete Text-Defined Levels.....	54015
<i>Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, Qiong Yan, Xiongkuo Min, Guangtao Zhai, Weisi Lin</i>	
Training Large Language Models for Reasoning Through Reverse Curriculum Reinforcement Learning	54030
<i>Zhiheng Xi, Wenxiang Chen, Boyang Hong, Senjie Jin, Rui Zheng, Wei He, Yiwen Ding, Shichun Liu, Xin Guo, Junzhe Wang, Honglin Guo, Wei Shen, Xiaoran Fan, Yuhao Zhou, Shihan Dou, Xiao Wang, Xinbo Zhang, Peng Sun, Tao Gui, Qi Zhang, Xuanjing Huang</i>	
Jetfire: Efficient and Accurate Transformer Pretraining with INT8 Data Flow and Per-Block Quantization	54049
<i>Haocheng Xi, Yuxiang Chen, Kang Zhao, Kai Jun Teh, Jianfei Chen, Jun Zhu</i>	
Unbiased Multi-Label Learning from Crowdsourced Annotations	54064
<i>Mingxuan Xia, Zenan Huang, Runze Wu, Gengyu Lyu, Junbo Zhao, Gang Chen, Haobo Wang</i>	
Refined Coreset Selection: Towards Minimal Coreset Size Under Model Performance Constraints.....	54082
<i>Xiaobo Xia, Jiale Liu, Shaokun Zhang, Qingyun Wu, Hongxin Wei, Tongliang Liu</i>	
LESS: Selecting Influential Data for Targeted Instruction Tuning.....	54104
<i>Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, Danqi Chen</i>	
SMArt: Improving GANs with Score Matching Regularity.....	54133
<i>Mengfei Xia, Yujun Shen, Ceyuan Yang, Ran Yi, Wenping Wang, Yong-Jin Liu</i>	
Contrastive Learning for Clinical Outcome Prediction with Partial Data Sources.....	54156
<i>Meng Xia, Jonathan Wilson, Benjamin Goldstein, Ricardo Henao</i>	
Position: Rethinking Post-Hoc Search-Based Neural Approaches for Solving Large-Scale Traveling Salesman Problems.....	54178
<i>Yifan Xia, Xianliang Yang, Zichuan Liu, Zhihao Liu, Lei Song, Jiang Bian</i>	
Delving into the Convergence of Generalized Smooth Minimax Optimization.....	54191
<i>Wenhan Xian, Ziyi Chen, Heng Huang</i>	

Differentially Private Post-Processing for Fair Regression	54212
<i>Ruicheng Xian, Qiaobo Li, Gautam Kamath, Han Zhao</i>	
Towards Neural Architecture Search Through Hierarchical Generative Modeling	54236
<i>Lichuan Xiang, Łukasz Dudziak, Mohamed S. Abdelfattah, Abhinav Mehrotra, Nicholas D. Lane, Hongkai Wen</i>	
Towards a Self-Contained Data-Driven Global Weather Forecasting Framework	54255
<i>Yi Xiao, Lei Bai, Wei Xue, Hao Chen, Kun Chen, Kang Chen, Tao Han, Wanli Ouyang</i>	
Category-Aware Active Domain Adaptation	54276
<i>Wenxiao Xiao, Jiuxiang Gu, Hongfu Liu</i>	
Improved Operator Learning by Orthogonal Attention	54288
<i>Zipeng Xiao, Zhongkai Hao, Bokai Lin, Zhijie Deng, Hang Su</i>	
Improving Transformers with Dynamically Composable Multi-Head Attention	54300
<i>Da Xiao, Qingye Meng, Shengping Li, Xingyuan Yuan</i>	
Uniformly Stable Algorithms for Adversarial Training and Beyond.....	54319
<i>Jiancong Xiao, Jiawei Zhang, Zhi-Quan Luo, Asuman Ozdaglar</i>	
Temporal Spiking Neural Networks with Synaptic Delay for Graph Reasoning	54341
<i>Mingqing Xiao, Yixin Zhu, Di He, Zhouchen Lin</i>	
Efficient Contrastive Learning for Fast and Accurate Inference on Graphs	54363
<i>Teng Xiao, Huaisheng Zhu, Zhiwei Zhang, Zhimeng Guo, Charu C. Aggarwal, Suhang Wang, Vasant G Honavar</i>	
CCM: Real-Time Controllable Visual Content Creation Using Text-To-Image Consistency Models.....	54382
<i>Jie Xiao, Kai Zhu, Han Zhang, Zhiheng Liu, Yujun Shen, Zhantao Yang, Ruili Feng, Yu Liu, Xueyang Fu, Zheng-Jun Zha</i>	
Intersecting-Boundary-Sensitive Fingerprinting for Tampering Detection of DNN Models	54402
<i>Xiaofan Bai, Chaoxiang He, Xiaojing Ma, Bin Benjamin Zhu, Hai Jin</i>	
Learning to Explore in POMDPs with Informational Rewards.....	54414
<i>Annie Xie, Logan Mondal Bhamidipaty, Evan Zheran Liu, Joey Hong, Sergey Levine, Chelsea Finn</i>	
Automating the Selection of Proxy Variables of Unmeasured Confounders.....	54430
<i>Feng Xie, Zhengming Chen, Shanshan Luo, Wang Miao, Ruichu Cai, Zhi Geng</i>	
FedREDefense: Defending Against Model Poisoning Attacks for Federated Learning Using Model Update Reconstruction Error	54460
<i>Yueqi Xie, Minghong Fang, Neil Zhenqiang Gong</i>	
Improving SAM Requires Rethinking Its Optimization Formulation	54475
<i>Wanyun Xie, Fabian Latorre, Kimon Antonakopoulos, Thomas Pethick, Volkan Cevher</i>	
Implicit Bias of AdamW: ℓ_∞ -Norm Constrained Optimization	54488
<i>Shuo Xie, Zhiyuan Li</i>	
Local Causal Structure Learning in the Presence of Latent Variables	54511
<i>Feng Xie, Zheng Li, Peng Wu, Yan Zeng, Chunchen Liu, Zhi Geng</i>	
Differentially Private Synthetic Data Via Foundation Model APIs 2: Text	54531
<i>Chulin Xie, Zinan Lin, Arturs Backurs, Sivakanth Gopi, Da Yu, Huseyin Inan, Harsha Nori, Haotian Jiang, Huishuai Zhang, Yin Tat Lee, Bo Li, Sergey Yekhanin</i>	
MH-pFLID: Model Heterogeneous Personalized Federated Learning Via Injection and Distillation for Medical Data Analysis	54561
<i>Luyuan Xie, Manqing Lin, Tianyu Luan, Cong Li, Yuejian Fang, Qingni Shen, Zhonghai Wu</i>	
Counterfactual Reasoning for Multi-Label Image Classification Via Patching-Based Training	54576
<i>Ming-Kun Xie, Jia-Hao Xiao, Pei Peng, Gang Niu, Masashi Sugiyama, Sheng-Jun Huang</i>	

TravelPlanner: A Benchmark for Real-World Planning with Language Agents.....	54590
<i>Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu, Renze Lou, Yuandong Tian, Yanghua Xiao, Yu Su</i>	

PART 66

Reflected Flow Matching	54614
<i>Tianyu Xie, Yu Zhu, Longlin Yu, Tong Yang, Ziheng Cheng, Shiyue Zhang, Xiangyu Zhang, Cheng Zhang</i>	
Federated Neuro-Symbolic Learning	54635
<i>Pengwei Xing, Songtao Lu, Han Yu</i>	
Less is More: On the Over-Globalizing Problem in Graph Transformers	54656
<i>Yujie Xing, Xiao Wang, Yibo Li, Hai Huang, Chuan Shi</i>	
HelmFluid: Learning Helmholtz Dynamics for Interpretable Fluid Prediction.....	54673
<i>Lanxiang Xing, Haixu Wu, Yuezhou Ma, Jianmin Wang, Mingsheng Long</i>	
SpikeLM: Towards General Spike-Driven Language Modeling Via Elastic Bi-Spiking Mechanisms.....	54698
<i>Xingrun Xing, Zheng Zhang, Ziyi Ni, Shitao Xiao, Yiming Ju, Siqu Fan, Yequan Wang, Jiajun Zhang, Guoqi Li</i>	
Iterative Preference Learning from Human Feedback: Bridging Theory and Practice for RLHF Under KL-Constraint.....	54715
<i>Wei Xiong, Hanze Dong, Chenlu Ye, Ziqi Wang, Han Zhong, Heng Ji, Nan Jiang, Tong Zhang</i>	
Provably Efficient Reinforcement Learning for Adversarial Restless Multi-Armed Bandits with Unknown Transitions and Bandit Feedback	54755
<i>Guojun Xiong, Jian Li</i>	
Distilling Morphology-Conditioned Hypernetworks for Efficient Universal Morphology Control.....	54777
<i>Zheng Xiong, Risto Vuorio, Jacob Beck, Matthieu Zimmer, Kun Shao, Shimon Whiteson</i>	
A General Framework for Sequential Decision-Making Under Adaptivity Constraints	54792
<i>Nuoya Xiong, Zhaoran Wang, Zhuoran Yang</i>	
Parameter-Dependent Competitive Analysis for Online Capacitated Coverage Maximization Through Boostings and Attenuations	54831
<i>Pan Xu</i>	
Reprompting: Automated Chain-Of-Thought Prompt Inference Through Gibbs Sampling.....	54852
<i>Weijia Xu, Andrzej Banburski-Fahey, Nebojsa Jojic</i>	
Stochastic Bandits with ReLU Neural Networks	54866
<i>Kan Xu, Hamsa Bastani, Surbhi Goel, Osbert Bastani</i>	
Intersectional Unfairness Discovery.....	54888
<i>Gezheng Xu, Qi Chen, Charles X. Ling, Boyu Wang, Changjian Shui</i>	
Semantic-Aware Human Object Interaction Image Generation.....	54918
<i>Zhu Xu, Qingchao Chen, Yuxin Peng, Yang Liu</i>	
DisCo-Diff: Enhancing Continuous Diffusion Models with Discrete Latents	54933
<i>Yilun Xu, Gabriele Corso, Tommi Jaakkola, Arash Vahdat, Karsten Kreis</i>	
Adapting Static Fairness to Sequential Decision-Making: Bias Mitigation Strategies Towards Equal Long-Term Benefit Rate.....	54962
<i>Yuancheng Xu, Chenghao Deng, Yanchao Sun, Ruijie Zheng, Xiyao Wang, Jieyu Zhao, Furong Huang</i>	
Is DPO Superior to PPO for LLM Alignment? A Comprehensive Study	54983
<i>Shusheng Xu, Wei Fu, Jiaxuan Gao, Wenjie Ye, Weilin Liu, Zhiyu Mei, Guangju Wang, Chao Yu, Yi Wu</i>	

Practical Hamiltonian Monte Carlo on Riemannian Manifolds via Relativity Theory.....	54999
<i>Kai Xu, Hong Ge</i>	
Equivariant Graph Neural Operator for Modeling 3D Dynamics.....	55015
<i>Minkai Xu, Jiaqi Han, Aaron Lou, Jean Kossaiji, Arvind Ramanathan, Kamyar Azizzadenesheli, Jure Leskovec, Stefano Ermon, Anima Anandkumar</i>	
Aligned Objective for Soft-Pseudo-Label Generation in Supervised Learning.....	55033
<i>Ning Xu, Yihao Hu, Congyu Qiao, Xin Geng</i>	
BiSHop: Bi-Directional Cellular Learning for Tabular Data with Generalized Sparse Modern Hopfield Model	55048
<i>Chenwei Xu, Yu-Chao Huang, Jerry Yao-Chieh Hu, Weijian Li, Ammar Gilani, Hsi-Sheng Goan, Han Liu</i>	
Conformal Prediction for Multi-Dimensional Time Series by Ellipsoidal Sets.....	55076
<i>Chen Xu, Hanyang Jiang, Yao Xie</i>	
Enhancing Vision Transformer: Amplifying Non-Linearity in Feedforward Network Module	55100
<i>Yixing Xu, Chao Li, Dong Li, Xiao Sheng, Fan Jiang, Lu Tian, Ashish Sirasao, Emad Barsoum</i>	
Meta-Reinforcement Learning Robust to Distributional Shift Via Performing Lifelong In-Context Learning	55112
<i>Tengye Xu, Zihao Li, Qinyuan Ren</i>	
Prompt-Guided Precise Audio Editing with Diffusion Models	55126
<i>Manjie Xu, Chenxing Li, Duzhen Zhang, Dan Su, Wei Liang, Dong Yu</i>	
Low-Rank Similarity Mining for Multimodal Dataset Distillation	55144
<i>Yue Xu, Zhilin Lin, Yusong Qiu, Cewu Lu, Yong-Lu Li</i>	
Robust Inverse Constrained Reinforcement Learning Under Model Misspecification	55162
<i>Sheng Xu, Guiliang Liu</i>	
Soft Prompt Recovers Compressed LLMs, Transferably	55186
<i>Zhaozhuo Xu, Zirui Liu, Beidi Chen, Shaochen (Henry) Zhong, Yuxin Tang, Jue Wang, Kaixiong Zhou, Xia Hu, Anshumali Shrivastava</i>	
Contrastive Preference Optimization: Pushing the Boundaries of LLM Performance in Machine Translation.....	55204
<i>Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, Young Jin Kim</i>	
Adaptive Group Personalization for Federated Mutual Transfer Learning	55225
<i>Haoqing Xu, Dian Shen, Meng Wang, Beilun Wang</i>	
Robust Universal Adversarial Perturbations.....	55241
<i>Changming Xu, Gagandeep Singh</i>	
Learning Exceptional Subgroups by End-To-End Maximizing KL-Divergence.....	55267
<i>Sascha Xu, Nils Philipp Walter, Janis Kalofolias, Jilles Vreeken</i>	
Pricing with Contextual Elasticity and Heteroscedastic Valuation.....	55286
<i>Jianyu Xu, Yu-Xiang Wang</i>	
Principled Preferential Bayesian Optimization.....	55305
<i>Wenjie Xu, Wenbin Wang, Yuning Jiang, Bratislav Svetozarevic, Colin N. Jones</i>	
Learning 1-Bit Tiny Object Detector with Discriminative Feature Refinement.....	55337
<i>Sheng Xu, Mingze Wang, Yanjing Li, Mingbao Lin, Baochang Zhang, David Doermann, Xiao Sun</i>	
SloG: An Inductive Spectral Graph Neural Network Beyond Polynomial Filter	55348
<i>Haobo Xu, Yuchen Yan, Dingsu Wang, Zhe Xu, Zhichen Zeng, Tarek F. Abdelzaher, Jiawei Han, Hanghang Tong</i>	

Libra: Building Decoupled Vision System on Large Language Models	55371
<i>Yifan Xu, Xiaoshan Yang, Yaguang Song, Changsheng Xu</i>	
A Sparsity Principle for Partially Observable Causal Representation Learning.....	55389
<i>Danru Xu, Dingling Yao, Sébastien Lachapelle, Perouz Taslakian, Julius Von Kügelgen, Francesco Locatello, Sara Magliacane</i>	
Language Agents with Reinforcement Learning for Strategic Play in the Werewolf Game.....	55434
<i>Zelai Xu, Chao Yu, Fei Fang, Yu Wang, Yi Wu</i>	

PART 67

Out-Of-Distribution Detection Via Deep Multi-Comprehension Ensemble.....	55465
<i>Chenhui Xu, Fuxun Yu, Zirui Xu, Nathan Inkawhich, Xiang Chen</i>	
Sparse Inducing Points in Deep Gaussian Processes: Enhancing Modeling with Denoising Diffusion Variational Inference	55490
<i>Jian Xu, Delu Zeng, John Paisley</i>	
Random Masking Finds Winning Tickets for Parameter Efficient Fine-Tuning	55501
<i>Jing Xu, Jingzhao Zhang</i>	
Exponential Spectral Pursuit: An Effective Initialization Method for Sparse Phase Retrieval.....	55525
<i>Mengchu Xu, Yuxuan Zhang, Jian Wang</i>	
Iterative Regularized Policy Optimization with Imperfect Demonstrations.....	55547
<i>Xudong Gong, Dawei Feng, Kele Xu, Yuanzhao Zhai, Chengkang Yao, Weijia Wang, Bo Ding, Huaimin Wang</i>	
Few-Shot Adaptation to Distribution Shifts by Mixing Source and Target Embeddings	55569
<i>Yihao Xue, Ali Payani, Yu Yang, Baharan Mirzsoleiman</i>	
Offline Multi-Objective Optimization.....	55595
<i>Ke Xue, Rong-Xi Tan, Xiaobin Huang, Chao Qian</i>	
OpenMoE: An Early Effort on Open Mixture-Of-Experts Language Models.....	55625
<i>Fuzhao Xue, Zian Zheng, Yao Fu, Jinjie Ni, Zangwei Zheng, Wangchunshu Zhou, Yang You</i>	
Unifying Bayesian Flow Networks and Diffusion Models Through Stochastic Differential Equations.....	55656
<i>Kaiwen Xue, Yuhao Zhou, Shen Nie, Xu Min, Xiaolu Zhang, Jun Zhou, Chongxuan Li</i>	
FairProof: Confidential and Certifiable Fairness for Neural Networks.....	55682
<i>Chhavi Yadav, Amrita Roy Chowdhury, Dan Boneh, Kamalika Chaudhuri</i>	
Mean Field Langevin Actor-Critic: Faster Convergence and Global Optimality Beyond Lazy Learning	55706
<i>Kakei Yamamoto, Kazusato Oko, Zhuoran Yang, Taiji Suzuki</i>	
Balancing Similarity and Complementarity for Federated Learning.....	55739
<i>Kunda Yan, Sen Cui, Abudukelimu Wuerkaixi, Jingfeng Zhang, Bo Han, Gang Niu, Masashi Sugiyama, Changshui Zhang</i>	
Probabilistic Time Series Modeling with Decomposable Denoising Diffusion Model	55759
<i>Tijin Yan, Hengheng Gong, Yongping He, Yufeng Zhan, Yuanqing Xia</i>	
Exploring the LLM Journey from Cognition to Expression with Linear Representations	55778
<i>Yuzi Yan, Jialian Li, Yipin Zhang, Dong Yan</i>	
A Space Group Symmetry Informed Network for O(3) Equivariant Crystal Tensor Prediction	55797
<i>Keqiang Yan, Alexandra Saxton, Xiaofeng Qian, Xiaoning Qian, Shuiwang Ji</i>	
Offline Imitation from Observation Via Primal Wasserstein State Occupancy Matching.....	55814
<i>Kai Yan, Alexander G. Schwing, Yu-Xiong Wang</i>	

Handling Heterogeneous Curvatures in Bandit LQR Control	55839
<i>Yu-Hu Yan, Jing Wang, Peng Zhao</i>	
Foundations of Testing for Finite-Sample Causal Discovery	55859
<i>Tom Yan, Ziyu Xu, Zachary Lipton</i>	
Retrieval Across Any Domains Via Large-Scale Pre-Trained Model	55901
<i>Jiexi Yan, Zhihui Yin, Chenghao Xu, Cheng Deng, Heng Huang</i>	
Reducing Balancing Error for Causal Inference Via Optimal Transport	55913
<i>Yuguang Yan, Hao Zhou, Zeqin Yang, Weilin Chen, Ruichu Cai, Zhifeng Hao</i>	
Do Efficient Transformers Really Save Computation?	55928
<i>Kai Yang, Jan Ackermann, Zhenyu He, Guhao Feng, Bohang Zhang, Yunzhen Feng, Qiwei Ye, Di He, Liwei Wang</i>	
Mind the Boundary: Coreset Selection Via Reconstructing the Decision Boundary	55948
<i>Shuo Yang, Zhe Cao, Sheng Guo, Ruiheng Zhang, Ping Luo, Shengping Zhang, Liqiang Nie</i>	
Sample-Efficient Multiagent Reinforcement Learning with Reset Replay	55961
<i>Yaodong Yang, Guangyong Chen, Jianye Hao, Pheng Ann Heng</i>	
DoraemonGPT: Toward Understanding Dynamic Scenes with Large Language Models (Exemplified as a Video Agent).....	55976
<i>Zongxin Yang, Guikun Chen, Xiaodi Li, Wenguan Wang, Yi Yang</i>	
A Dense Reward View on Aligning Text-To-Image Diffusion with Preference	55998
<i>Shentao Yang, Tianqi Chen, Mingyuan Zhou</i>	
Lyapunov-Stable Neural Control for State and Output Feedback: A Novel Formulation	56033
<i>Lujie Yang, Hongkai Dai, Zhouxing Shi, Cho-Jui Hsieh, Russ Tedrake, Huan Zhang</i>	
Latent Space Symmetry Discovery	56047
<i>Jianke Yang, Nima Dehmamy, Robin Walters, Rose Yu</i>	
Guidance with Spherical Gaussian Constraint for Conditional Diffusion	56071
<i>Lingxiao Yang, Shutong Ding, Yifan Cai, Jingyi Yu, Jingya Wang, Ye Shi</i>	
Better Safe than Sorry: Pre-Training CLIP Against Targeted Data Poisoning and Backdoor Attacks.....	56096
<i>Wenhan Yang, Jingdong Gao, Baharan Mirzasoleiman</i>	
Deeper Or Wider: A Perspective from Optimal Generalization Error with Sobolev Loss.....	56109
<i>Yahong Yang, Juncai He</i>	
SAM as the Guide: Mastering Pseudo-Label Refinement in Semi-Supervised Referring Expression Segmentation	56139
<i>Danni Yang, Jiayi Ji, Yiwei Ma, Tianyu Guo, Haowei Wang, Xiaoshuai Sun, Rongrong Ji</i>	
Small-Loss Adaptive Regret for Online Convex Optimization	56156
<i>Wenhao Yang, Wei Jiang, Yibo Wang, Ping Yang, Yao Hu, Lijun Zhang</i>	
Towards Interpretable Deep Local Learning with Successive Gradient Reconciliation.....	56196
<i>Yibo Yang, Xiaojie Li, Motasem Alfarra, Hasan Hammoud, Adel Bibi, Philip Torr, Bernard Ghanem</i>	
Bounded and Uniform Energy-Based Out-Of-Distribution Detection for Graphs	56216
<i>Shenzhi Yang, Bin Liang, An Liu, Lin Gui, Xingkai Yao, Xiaofang Zhang</i>	
Position: Towards Implicit Prompt for Text-To-Image Models.....	56235
<i>Yue Yang, Yuqi Lin, Hong Liu, Wenqi Shao, Runjian Chen, Hailong Shang, Yu Wang, Yu Qiao, Kaipeng Zhang, Ping Luo</i>	
Position: Towards Unified Alignment Between Agents, Humans, and Environment.....	56251
<i>Zonghan Yang, An Liu, Zijun Liu, Kaiming Liu, Fangzhou Xiong, Yile Wang, Zeyuan Yang, Qingyuan Hu, Xinrui Chen, Zhenhe Zhang, Fuwen Luo, Zhicheng Guo, Peng Li, Yang Liu</i>	

Rewards-In-Context: Multi-Objective Alignment of Foundation Models with Dynamic Preference Adjustment	56276
<i>Rui Yang, Xiaoman Pan, Feng Luo, Shuang Qiu, Han Zhong, Dong Yu, Jianshu Chen</i>	
Understanding Server-Assisted Federated Learning in the Presence of Incomplete Client Participation.....	56298
<i>Haibo Yang, Peiwen Qiu, Prashant Khanduri, Minghong Fang, Jia Liu</i>	
ILILT: Implicit Learning of Inverse Lithography Technologies.....	56319
<i>Haoyu Yang, Haoxing Ren</i>	

PART 68

Representation Surgery for Multi-Task Model Merging	56332
<i>Enneng Yang, Li Shen, Zhenyi Wang, Guibing Guo, Xiaojun Chen, Xingwei Wang, Dacheng Tao</i>	
Reducing Fine-Tuning Memory Overhead by Approximate and Memory-Sharing Backpropagation	56357
<i>Yuchen Yang, Yingdong Shi, Cheems Wang, Xiantong Zhen, Yuxuan Shi, Jun Xu</i>	
Unlocking the Power of Spatial and Temporal Information in Medical Multimodal Pre-Training	56382
<i>Jinxia Yang, Bing Su, Wayne Xin Zhao, Ji-Rong Wen</i>	
Exploration and Anti-Exploration with Distributional Random Network Distillation	56397
<i>Kai Yang, Jian Tao, Jiafei Lyu, Xiu Li</i>	
UniAudio: Towards Universal Audio Generation with Large Language Models.....	56422
<i>Dongchao Yang, Jinchuan Tian, Xu Tan, Rongjie Huang, Songxiang Liu, Haohan Guo, Xuankai Chang, Jiatong Shi, Sheng Zhao, Jiang Bian, Zhou Zhao, Xixin Wu, Helen Meng</i>	
Explain Temporal Black-Box Models Via Functional Decomposition.....	56448
<i>Linxiao Yang, Yunze Tong, Xinyue Gu, Liang Sun</i>	
Position: Video as the New Language for Real-World Decision Making.....	56465
<i>Sherry Yang, Jacob Walker, Jack Parker-Holder, Yilun Du, Jake Bruce, Andre Barreto, Pieter Abbeel, Dale Schuurmans</i>	
ATraDiff: Accelerating Online Reinforcement Learning with Imaginary Trajectories.....	56485
<i>Qianlan Yang, Yu-Xiong Wang</i>	
Gated Linear Attention Transformers with Hardware-Efficient Training.....	56501
<i>Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, Yoon Kim</i>	
Analysis for Abductive Learning and Neural-Symbolic Reasoning Shortcuts	56524
<i>Xiao-Wen Yang, Wen-Da Wei, Jie-Jing Shao, Yu-Feng Li, Zhi-Hua Zhou</i>	
Stability and Generalization of Stochastic Compositional Gradient Descent Algorithms.....	56542
<i>Ming Yang, Xiyuan Wei, Tianbao Yang, Yiming Ying</i>	
How Graph Neural Networks Learn: Lessons from Training Dynamics.....	56594
<i>Chenxiao Yang, Qitian Wu, David Wipf, Ruoyu Sun, Junchi Yan</i>	
Harnessing Hierarchical Label Distribution Variations in Test Agnostic Long-Tail Recognition	56624
<i>Zhiyong Yang, Qianqian Xu, Zitai Wang, Sicong Li, Boyu Han, Shilong Bao, Xiaochun Cao, Qingming Huang</i>	
Neuro-Symbolic Temporal Point Processes	56665
<i>Yang Yang, Chao Yang, Boyang Li, Yinghao Fu, Shuang Li</i>	
One Meta-Tuned Transformer is What You Need for Few-Shot Learning	56681
<i>Xu Yang, Huaxiu Yao, Ying Wei</i>	
Mastering Text-To-Image Diffusion: Recaptioning, Planning, and Generating with Multimodal LLMs.....	56704
<i>Ling Yang, Zhaochen Yu, Chenlin Meng, Minkai Xu, Stefano Ermon, Bin Cui</i>	

Accelerating Look-Ahead in Bayesian Optimization: Multilevel Monte Carlo is All You Need	56722
<i>Shangda Yang, Vitaly Zankin, Maximilian Balandat, Stefan Scherer, Kevin Carlberg, Neil Walton, Kody J. H. Law</i>	
MorphGrower: A Synchronized Layer-By-Layer Growing Approach for Plausible Neuronal Morphology Generation	56749
<i>Nianzu Yang, Kaipeng Zeng, Haotian Lu, Yexin Wu, Zexin Yuan, Danni Chen, Shengdian Jiang, Jiaxiang Wu, Yimin Wang, Junchi Yan</i>	
Mol-AE: Auto-Encoder Based Molecular Representation Learning with 3D Cloze Test Objective	56793
<i>Junwei Yang, Kangjie Zheng, Siyu Long, Zaiqing Nie, Ming Zhang, Xinyu Dai, Wei-Ying Ma, Hao Zhou</i>	
What is Dataset Distillation Learning?.....	56812
<i>William Yang, Ye Zhu, Zhiwei Deng, Olga Russakovsky</i>	
Protein Conformation Generation Via Force-Guided SE(3) Diffusion Models.....	56835
<i>Yan Wang, Lihao Wang, Yuning Shen, Yiqun Wang, Huizhuo Yuan, Yue Wu, Quanquan Gu</i>	
Empowering Graph Invariance Learning with Deep Spurious Infomax	56860
<i>Tianjun Yao, Yongqiang Chen, Zhenhao Chen, Kai Hu, Zhiqiang Shen, Kun Zhang</i>	
Human Vs. Generative AI in Content Creation Competition: Symbiosis Or Conflict?.....	56885
<i>Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, Haifeng Xu</i>	
Mobile Attention: Mobile-Friendly Linear-Attention for Vision Transformers.....	56914
<i>Zhiyu Yao, Jian Wang, Haixu Wu, Jingdong Wang, Mingsheng Long</i>	
Socialized Learning: Making Each Other Better Through Multi-Agent Collaboration.....	56927
<i>Xinjie Yao, Yu Wang, Pengfei Zhu, Wanyu Lin, Jialu Li, Weihao Li, Qinghua Hu</i>	
Compressible Dynamics in Deep Overparameterized Low-Rank Learning & Adaptation	56946
<i>Can Yaras, Peng Wang, Laura Balzano, Qing Qu</i>	
EMC ² : Efficient MCMC Negative Sampling for Contrastive Learning with Global Convergence	56966
<i>Chung-Yiu Yau, Hoi-To Wai, Parameswaran Raman, Soumajyoti Sarkar, Mingyi Hong</i>	
Towards Robust Model-Based Reinforcement Learning Against Adversarial Corruption.....	56982
<i>Chenlu Ye, Jiafan He, Quanquan Gu, Tong Zhang</i>	
Adaptive Sampling of K-Space in Magnetic Resonance for Rapid Pathology Prediction	57018
<i>Chen-Yu Yen, Raghav Singhal, Umang Sharma, Rajesh Ranganath, Sumit Chopra, Lerrel Pinto</i>	
StableMask: Refining Causal Masking in Decoder-Only Transformer	57033
<i>Qingyu Yin, Xuzheng He, Xiang Zhuang, Yu Zhao, Jianhua Yao, Xiaoyu Shen, Qiang Zhang</i>	
Junk DNA Hypothesis: Pruning Small Pre-Trained Weights Irreversibly and Monotonically Impairs ``Difficult'' Downstream Tasks in LLMs.....	57053
<i>Lu Yin, Ajay Jaiswal, Shiwei Liu, Souvik Kundu, Zhangyang Wang</i>	
Characterizing Truthfulness in Large Language Model Generations with Local Intrinsic Dimension.....	57069
<i>Fan Yin, Jayanth Srinivasa, Kai-Wei Chang</i>	
High-Dimensional Bayesian Optimization Via Semi-Supervised Learning with Optimized Unlabeled Data Sampling.....	57085
<i>Yuxuan Yin, Yu Wang, Peng Li</i>	
Outlier Weighed Layerwise Sparsity (OWL): A Missing Secret Sauce for Pruning LLMs to High Sparsity	57101
<i>Lu Yin, You Wu, Zhenyu Zhang, Cheng-Yu Hsieh, Yaqing Wang, Yiling Jia, Gen Li, Ajay Jaiswal, Mykola Pechenizkiy, Yi Liang, Michael Bendersky, Zhangyang Wang, Shiwei Liu</i>	

MMT-Bench: A Comprehensive Multimodal Benchmark for Evaluating Large Vision-Language Models Towards Multitask AGI	57116
<i>Kaining Ying, Fanqing Meng, Jin Wang, Zhiqian Li, Han Lin, Yue Yang, Hao Zhang, Wenbo Zhang, Yuqi Lin, Shuo Liu, Jiayi Lei, Quanfeng Lu, Runjian Chen, Peng Xu, Renrui Zhang, Haozhe Zhang, Peng Gao, Yali Wang, Yu Qiao, Ping Luo, Kaipeng Zhang, Wenqi Shao</i>	

PART 69

Layerwise Proximal Replay: A Proximal Point Method for Online Continual Learning	57199
<i>Jason Yoo, Yunpeng Liu, Frank Wood, Geoff Pleiss</i>	
Uncertainty Estimation by Density Aware Evidential Deep Learning.....	57217
<i>Taeseong Yoon, Heeyoung Kim</i>	
Optimal Acceleration for Minimax and Fixed-Point Problems is Not Unique.....	57244
<i>Taeho Yoon, Jaeyeon Kim, Jaewook J. Suh, Ernest K. Ryu</i>	
FRAG: Frequency Adapting Group for Diffusion Video Editing.....	57315
<i>Sunjae Yoon, Gwanhyeong Koo, Geonwoo Kim, Chang D. Yoo</i>	
Breadth-First Exploration on Adaptive Grid for Reinforcement Learning	57331
<i>Youngsik Yoon, Gangbok Lee, Sungsoo Ahn, Jungseul Ok</i>	
When Linear Attention Meets Autoregressive Decoding: Towards More Effective and Efficient Linearized Large Language Models	57350
<i>Haoran You, Yichao Fu, Zheng Wang, Amir Yazdanbakhsh, Yingyan (Celine) Lin</i>	
SpikeZIP-TF: Conversion is All You Need for Transformer-Based SNN	57367
<i>Kang You, Zekai Xu, Chen Nie, Zhijie Deng, Qinghai Guo, Xiang Wang, Zhezhi He</i>	
Efficient Algorithms for Empirical Group Distributionally Robust Optimization and Beyond	57384
<i>Dingzhi Yu, Yunuo Cai, Wei Jiang, Lijun Zhang</i>	
Towards Resource-Friendly, Extensible and Stable Incomplete Multi-View Clustering.....	57415
<i>Shengju Yu, Zhibing Dong, Siwei Wang, Xinhang Wan, Yue Liu, Weixuan Liang, Pei Zhang, Wenxuan Tu, Xinwang Liu</i>	
Activation-Descent Regularization for Input Optimization of ReLU Networks	57441
<i>Hongzhan Yu, Sicun Gao</i>	
Collage: Light-Weight Low-Precision Strategy for LLM Training	57459
<i>Tao Yu, Gaurav Gupta, Karthick Gopalswamy, Amith Mamidala, Hao Zhou, Jeffrey Huynh, Youngsuk Park, Ron Diamant, Anoop Deoras, Luke Huan</i>	
Privacy-Preserving Instructions for Aligning Large Language Models	57480
<i>Da Yu, Peter Kairouz, Sewoong Oh, Zheng Xu</i>	
Learning Latent Structures in Network Games Via Data-Dependent Gated-Prior Graph Variational Autoencoders.....	57507
<i>Xue Yu, Muchen Li, Yan Leng, Renjie Liao</i>	
Learning Scale-Aware Spatio-Temporal Implicit Representation for Event-Based Motion Deblurring	57527
<i>Wei Yu, Jianing Li, Shengping Zhang, Xiangyang Ji</i>	
Enabling Few-Shot Learning with PID Control: A Layer Adaptive Optimizer	57544
<i>Le Yu, Xinde Li, Pengfei Zhang, Zhentong Zhang, Fir Dunkin</i>	
Generalization Bound and New Algorithm for Clean-Label Backdoor Attack	57559
<i>Lijia Yu, Shuang Liu, Yibo Miao, Xiao-Shan Gao, Lijun Zhang</i>	
Learning Causal Dynamics Models in Object-Oriented Environments.....	57597
<i>Zhongwei Yu, Jingqing Ruan, Dengpeng Xing</i>	

ViP: A Differentially Private Foundation Model for Computer Vision	57639
<i>Yaodong Yu, Maziar Sanjabi, Yi Ma, Kamalika Chaudhuri, Chuan Guo</i>	
Unveiling and Harnessing Hidden Attention Sinks: Enhancing Large Language Models Without Training Through Attention Calibration	57659
<i>Zhongzhi Yu, Zheng Wang, Yonggan Fu, Huihong Shi, Khalid Shaikh, Yingyan (Celine) Lin</i>	
Purify Unlearnable Examples Via Rate-Constrained Variational Autoencoders.....	57678
<i>Yi Yu, Yufei Wang, Song Xia, Wenhan Yang, Shijian Lu, Yap-Peng Tan, Alex C. Kot</i>	
Few-Shot Character Understanding in Movies as an Assessment to Meta-Learning of Theory-Of-Mind	57703
<i>Mo Yu, Qiuqing Wang, Shunchi Zhang, Yisi Sang, Kangsheng Pu, Zekai Wei, Han Wang, Liyan Xu, Jing Li, Yue Yu, Jie Zhou</i>	
MM-Vet: Evaluating Large Multimodal Models for Integrated Capabilities	57730
<i>Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, Lijuan Wang</i>	
Language Models Are Super Mario: Absorbing Abilities from Homologous Models as a Free Lunch.....	57755
<i>Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, Yongbin Li</i>	
Improving Sharpness-Aware Minimization by Lookahead	57776
<i>Runsheng Yu, Youzhi Zhang, James T. Kwok</i>	
A Unified Adaptive Testing System Enabled by Hierarchical Structure Search.....	57803
<i>Junhao Yu, Yan Zhuang, Zhenya Huang, Qi Liu, Xin Li, Rui Li, Enhong Chen</i>	
Revitalizing Multivariate Time Series Forecasting: Learnable Decomposition with Inter-Series Dependencies and Intra-Series Variations Modeling.....	57818
<i>Guoqi Yu, Jing Zou, Xiaowei Hu, Angelica I. Aviles-Rivero, Jing Qin, Shujun Wang</i>	
Smoothing Proximal Gradient Methods for Nonsmooth Sparsity Constrained Optimization: Optimality Conditions and Global Convergence.....	57842
<i>Ganzhao Yuan</i>	
A Linear Time and Space Local Point Cloud Geometry Encodervia Vectorized Kernel Mixture (VecKM).....	57871
<i>Dehao Yuan, Cornelia Fermüller, Tahseen Rabbani, Furong Huang, Yiannis Aloimonos</i>	
SHINE: Shielding Backdoors in Deep Reinforcement Learning.....	57887
<i>Zhuowen Yuan, Wenbo Guo, Jinyuan Jia, Bo Li, Dawn Song</i>	
Self-Rewarding Language Models	57905
<i>Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, Jason Weston</i>	
Not Just Pretty Pictures: Toward Interventional Data Augmentation Using Text-To-Image Generators	57924
<i>Jianhao Yuan, Francesco Pinto, Adam Davies, Philip Torr</i>	
RigorLLM: Resilient Guardrails for Large Language Models Against Undesired Content	57953
<i>Zhuowen Yuan, Zidi Xiong, Yi Zeng, Ning Yu, Ruoxi Jia, Dawn Song, Bo Li</i>	
OLLIE: Imitation Learning from Offline Pretraining to Online Finetuning.....	57966
<i>Sheng Yue, Xingyuan Hua, Ju Ren, Sen Lin, Junshan Zhang, Yaoxue Zhang</i>	

PART 70

Federated Self-Explaining GNNs with Anti-Shortcut Augmentations	58019
<i>Linan Yue, Qi Liu, Weibo Gao, Ye Liu, Kai Zhang, Yichao Du, Li Wang, Fangzhou Yao</i>	

How to Leverage Diverse Demonstrations in Offline Imitation Learning	58037
<i>Sheng Yue, Jiani Liu, Xingyuan Hua, Ju Ren, Sen Lin, Junshan Zhang, Yaoxue Zhang</i>	
Image Restoration Through Generalized Ornstein-Uhlenbeck Bridge	58068
<i>Conghan Yue, Zhengwei Peng, Junlong Ma, Shiyao Du, Pengxu Wei, Dongyu Zhang</i>	
Manifold Integrated Gradients: Riemannian Geometry for Feature Attribution	58090
<i>Eslam Zaher, Maciej Trzaskowski, Quan Nguyen, Fred Roosta</i>	
Sample as You Infer: Predictive Coding with Langevin Dynamics.....	58105
<i>Umair Zahid, Qinghai Guo, Zafeirios Fountas</i>	
More Flexible PAC-Bayesian Meta-Learning by Learning Learning Algorithms.....	58122
<i>Hossein Zakerinia, Amin Behjati, Christoph H. Lampert</i>	
How to Explore with Belief: State Entropy Maximization in POMDPs	58140
<i>Riccardo Zamboni, Duilio Cirino, Marcello Restelli, Mirco Mutti</i>	
See More Details: Efficient Image Super-Resolution by Experts Mining	58158
<i>Eduard Zamfir, Zongwei Wu, Nancy Mehta, Yulun Zhang, Radu Timofte</i>	
DiffAug: Enhance Unsupervised Contrastive Learning with Domain-Knowledge-Free Diffusion- Based Data Augmentation	58174
<i>Zelin Zang, Hao Luo, Kai Wang, Panpan Zhang, Fan Wang, Stan. Z Li, Yang You</i>	
Robustly Learning Single-Index Models Via Alignment Sharpness.....	58197
<i>Nikos Zarifis, Puqian Wang, Ilias Diakonikolas, Jelena Diakonikolas</i>	
Low-Cost High-Power Membership Inference Attacks.....	58244
<i>Sajjad Zarifzadeh, Philippe Liu, Reza Shokri</i>	
Theoretical Analysis of Learned Database Operations Under Distribution Shift Through Distribution Learnability	58283
<i>Sepanta Zeighami, Cyrus Shahabi</i>	
Continuous Treatment Effects with Surrogate Outcomes.....	58306
<i>Zhenghao Zeng, David Arbour, Avi Feller, Raghavendra Addanki, Ryan Rossi, Ritwik Sinha, Edward H. Kennedy</i>	
tnGPS: Discovering Unknown Tensor Network Structure Search Algorithms Via Large Language Models (LLMs)	58329
<i>Junhua Zeng, Chao Li, Zhun Sun, Qibin Zhao, Guoxu Zhou</i>	
Token-Level Direct Preference Optimization.....	58348
<i>Yongcheng Zeng, Guoqing Liu, Weiyu Ma, Ning Yang, Haifeng Zhang, Jun Wang</i>	
Learning Reward for Robot Skills Using Large Language Models Via Self-Alignment.....	58366
<i>Yuwei Zeng, Yao Mu, Lin Shao</i>	
Graph Mixup on Approximate Gromov–Wasserstein Geodesics	58387
<i>Zhichen Zeng, Ruizhong Qiu, Zhe Xu, Zhining Liu, Yuchen Yan, Tianxin Wei, Lei Ying, Jingrui He, Hanghang Tong</i>	
Interacting Diffusion Processes for Event Sequence Forecasting.....	58407
<i>Mai Zeng, Florence Regol, Mark Coates</i>	
IM-Unpack: Training and Inference with Arbitrarily Low Precision Integers	58431
<i>Zhanpeng Zeng, Karthikeyan Sankaralingam, Vikas Singh</i>	
Differentiable Annealed Importance Sampling Minimizes the Jensen-Shannon Divergence Between Initial and Target Distribution.....	58448
<i>Johannes Zenn, Robert Bamler</i>	
Robust Learning-Augmented Dictionaries	58470
<i>Ali Zeynali, Shahin Kamali, Mohammad Hajiesmaili</i>	

Actions Speak Louder than Words: Trillion-Parameter Sequential Transducers for Generative Recommendations	58484
<i>Jiaqi Zhai, Lucy Liao, Xing Liu, Yueming Wang, Rui Li, Xuan Cao, Leon Gao, Zhaojie Gong, Fangda Gu, Jiayuan He, Yinghai Lu, Yu Shi</i>	
Tight Partial Identification of Causal Effects with Marginal Distribution of Unmeasured Confounders	58510
<i>Zhiheng Zhang</i>	
DAG-Based Column Generation for Adversarial Team Games	58563
<i>Youzhi Zhang, Bo An, Daniel Dajun Zeng</i>	
SAM-E: Leveraging Visual Foundation Model with Sequence Imitation for Embodied Manipulation	58579
<i>Junjie Zhang, Chenjia Bai, Haoran He, Zhigang Wang, Bin Zhao, Xiu Li, Xuelong Li</i>	
Efficient Stochastic Approximation of Minimax Excess Risk Optimization.....	58599
<i>Lijun Zhang, Haomin Bai, Wei-Wei Tu, Ping Yang, Yao Hu</i>	
Discounted Adaptive Online Learning: Towards Better Regularization.....	58631
<i>Zhiyu Zhang, David Bombara, Heng Yang</i>	
Improving Equivariant Graph Neural Networks on Large Geometric Graphs Via Virtual Nodes Learning	58662
<i>Yuelin Zhang, Jiacheng Cen, Jiaqi Han, Zhiqiang Zhang, Jun Zhou, Wenbing Huang</i>	
Provably Efficient Partially Observable Risk-Sensitive Reinforcement Learning with Hindsight Observation	58680
<i>Tonghe Zhang, Yu Chen, Longbo Huang</i>	
Verification of Machine Unlearning is Fragile	58717
<i>Binchi Zhang, Zihan Chen, Cong Shen, Jundong Li</i>	
Reshape and Adapt for Output Quantization (RAOQ): Quantization-Aware Training for In-Memory Computing Systems.....	58739
<i>Bonan Zhang, Chia-Yu Chen, Naveen Verma</i>	
LQER: Low-Rank Quantization Error Reconstruction for LLMs	58763
<i>Cheng Zhang, Jianyi Cheng, George A. Constantinides, Yiren Zhao</i>	
Random Scaling and Momentum for Non-Smooth Non-Convex Optimization.....	58780
<i>Qinzi Zhang, Ashok Cutkosky</i>	
Towards Certified Unlearning for Deep Neural Networks	58800
<i>Binchi Zhang, Yushun Dong, Tianhao Wang, Jundong Li</i>	

PART 71

Look Ahead Or Look Around? A Theoretical Comparison Between Autoregressive and Masked Pretraining	58819
<i>Qi Zhang, Tianqi Du, Haotian Huang, Yifei Wang, Yisen Wang</i>	
CaM: Cache Merging for Memory-Efficient LLMs Inference.....	58840
<i>Yuxin Zhang, Yuxuan Du, Gen Luo, Yunshan Zhong, Zhenyu Zhang, Shiwei Liu, Rongrong Ji</i>	
Watermarks in the Sand: Impossibility of Strong Watermarking for Language Models	58851
<i>Hanlin Zhang, Benjamin L. Edelman, Danilo Francati, Daniele Venturi, Giuseppe Ateniese, Boaz Barak</i>	
MILP-FBGen: LP/MILP Instance Generation with Feasibility/Boundedness	58881
<i>Yahong Zhang, Chenchen Fan, Donghui Chen, Congrui Li, Wenli Ouyang, Mingda Zhu, Junchi Yan</i>	

SF-DQN: Provable Knowledge Transfer Using Successor Feature for Deep Reinforcement Learning	58897
<i>Shuai Zhang, Heshan Devaka Fernando, Miao Liu, Keerthiram Murugesan, Songtao Lu, Pin-Yu Chen, Tianyi Chen, Meng Wang</i>	
Model-Based Reinforcement Learning for Parameterized Action Spaces.....	58935
<i>Renhao Zhang, Haotian Fu, Yilin Miao, George Konidakis</i>	
Easing Concept Bleeding in Diffusion Via Entity Localization and Anchoring.....	58955
<i>Jiewei Zhang, Song Guo, Peiran Dong, Jie Zhang, Ziming Liu, Yue Yu, Xiao-Ming Wu</i>	
Generating Chain-Of-Thoughts with a Pairwise-Comparison Approach to Searching for the Most Promising Intermediate Thought	58967
<i>Zhen-Yu Zhang, Siwei Han, Huaxiu Yao, Gang Niu, Masashi Sugiyama</i>	
Inherent Trade-Offs Between Diversity and Stability in Multi-Task Benchmarks	58984
<i>Guanhua Zhang, Moritz Hardt</i>	
Advancing DRL Agents in Commercial Fighting Games: Training, Integration, and Agent-Human Alignment.....	59003
<i>Chen Zhang, Qiang He, Yuan Zhou, Elvis S. Liu, Hong Wang, Jian Zhao, Yang Wang</i>	
On the Duality Between Sharpness-Aware Minimization and Adversarial Training.....	59024
<i>Yihao Zhang, Hangzhou He, Jingyu Zhu, Huanran Chen, Yifei Wang, Zeming Wei</i>	
Towards Causal Foundation Model: on Duality Between Optimal Balancing and Attention.....	59042
<i>Jiaqi Zhang, Joel Jennings, Agrin Hilmkil, Nick Pawlowski, Cheng Zhang, Chao Ma</i>	
Parameter-Efficient Fine-Tuning with Controls	59066
<i>Chi Zhang, Jingpu Cheng, Yanyu Xu, Qianxiao Li</i>	
Deep Regression Representation Learning with Topology.....	59080
<i>Shihao Zhang, Kenji Kawaguchi, Angela Yao</i>	
Understanding Unimodal Bias in Multimodal Deep Linear Networks.....	59100
<i>Yedi Zhang, Peter E. Latham, Andrew Saxe</i>	
Guarantees for Nonlinear Representation Learning: Non-Identical Covariates, Dependent Data, Fewer Samples	59126
<i>Thomas T. Zhang, Bruce D. Lee, Ingvar Ziemann, George J. Pappas, Nikolai Matni</i>	
An Interpretable Evaluation of Entropy-Based Novelty of Generative Models.....	59148
<i>Jingwei Zhang, Cheuk Ting Li, Farzan Farnia</i>	
Revisiting Zeroth-Order Optimization for Memory-Efficient LLM Fine-Tuning: A Benchmark	59173
<i>Yihua Zhang, Pingzhi Li, Junyuan Hong, Jiayang Li, Yimeng Zhang, Wenqing Zheng, Pin-Yu Chen, Jason D. Lee, Wotao Yin, Mingyi Hong, Zhangyang Wang, Sijia Liu, Tianlong Chen</i>	
S3O: A Dual-Phase Approach for Reconstructing Dynamic Shape and Skeleton of Articulated Objects from Single Monocular Video	59191
<i>Hao Zhang, Fang Li, Samyak Rawlekar, Narendra Ahuja</i>	
DPZero: Private Fine-Tuning of Language Models Without Backpropagation.....	59210
<i>Liang Zhang, Bingcong Li, Kiran Koshy Thekumparampil, Sewoong Oh, Niao He</i>	
Conditional Language Learning with Context.....	59247
<i>Xiao Zhang, Miao Li, Ji Wu</i>	
Tackling Non-Stationarity in Reinforcement Learning Via Causal-Origin Representation.....	59264
<i>Wanpeng Zhang, Yilin Li, Boyu Yang, Zongqing Lu</i>	
BLO-SAM: Bi-Level Optimization Based Finetuning of the Segment Anything Model for Overfitting-Preventing Semantic Segmentation	59289
<i>Li Zhang, Youwei Liang, Ruiyi Zhang, Amirhosein Javadi, Pengtao Xie</i>	

Multi-Factor Adaptive Vision Selection for Egocentric Video Question Answering.....	59310
<i>Haoyu Zhang, Meng Liu, Zixin Liu, Xuemeng Song, Yaowei Wang, Liqiang Nie</i>	
Self-Consistency Training for Density-Functional-Theory Hamiltonian Prediction	59329
<i>He Zhang, Chang Liu, Zun Wang, Xinran Wei, Siyuan Liu, Nanning Zheng, Bin Shao, Tie-Yan Liu</i>	
UP2ME: Univariate Pre-Training to Multivariate Fine-Tuning as a General-Purpose Framework for Multivariate Time Series Analysis.....	59358
<i>Yunhao Zhang, Minghao Liu, Shengyang Zhou, Junchi Yan</i>	
Improving Accuracy-Robustness Trade-Off Via Pixel Reweighted Adversarial Training.....	59382
<i>Jiacheng Zhang, Feng Liu, Dawei Zhou, Jingfeng Zhang, Tongliang Liu</i>	
MagicLens: Self-Supervised Image Retrieval with Open-Ended Instructions	59403
<i>Kai Zhang, Yi Luan, Hexiang Hu, Kenton Lee, Siyuan Qiao, Wenhui Chen, Yu Su, Ming-Wei Chang</i>	
Wukong: Towards a Scaling Law for Large-Scale Recommendation	59421
<i>Buyun Zhang, Liang Luo, Yuxin Chen, Jade Nie, Xi Liu, Shen Li, Yanli Zhao, Yuchen Hao, Yantao Yao, Ellie Dingqiao Wen, Jongsoo Park, Maxim Naumov, Wenlin Chen</i>	
Nonparametric Teaching of Implicit Neural Representations.....	59435
<i>Chen Zhang, Steven Tin Sui Luo, Jason Chun Lok Li, Yik-Chung Wu, Ngai Wong</i>	
Pessimism Meets Risk: Risk-Sensitive Offline Reinforcement Learning	59459
<i>Dake Zhang, Boxiang Lyu, Shuang Qiu, Mladen Kolar, Tong Zhang</i>	
Sparse-To-Dense Multimodal Image Registration Via Multi-Task Learning	59490
<i>Kaining Zhang, Jiayi Ma</i>	
Amend to Alignment: Decoupled Prompt Tuning for Mitigating Spurious Correlation in Vision- Language Models	59505
<i>Jie Zhang, Xiaosong Ma, Song Guo, Peng Li, Wenchao Xu, Xueyang Tang, Zicong Hong</i>	
In-Context Principle Learning from Mistakes	59520
<i>Tianjun Zhang, Aman Madaan, Luyu Gao, Steven Zheng, Swaroop Mishra, Yiming Yang, Niket Tandon, Uri Alon</i>	
Sequential Asynchronous Action Coordination in Multi-Agent Systems: A Stackelberg Decision Transformer Approach.....	59559
<i>Bin Zhang, Hangyu Mao, Lijuan Li, Zhiwei Xu, Dapeng Li, Rui Zhao, Guoliang Fan</i>	
Sparsest Models Elude Pruning: An Exposé of Pruning’s Current Capabilities.....	59576
<i>Stephen Zhang, Vardan Papyan</i>	
A Federated Stochastic Multi-Level Compositional Minimax Algorithm for Deep AUC Maximization.....	59601
<i>Xinwen Zhang, Ali Payani, Myungjin Lee, Richard Souvenir, Hongchang Gao</i>	
Riemannian Preconditioned LoRA for Fine-Tuning Foundation Models.....	59641
<i>Fangzhao Zhang, Mert Pilanci</i>	

PART 72

How Language Model Hallucinations Can Snowball.....	59670
<i>Muru Zhang, Ofir Press, William Merrill, Alisa Liu, Noah A. Smith</i>	
Enhancing Storage and Computational Efficiency in Federated Multimodal Learning for Large- Scale Models	59685
<i>Zixin Zhang, Fan Qi, Changsheng Xu</i>	
Online Resource Allocation with Non-Stationary Customers	59700
<i>Xiaoyue Zhang, Hanzhang Qin, Mabel C. Chou</i>	

Flexible Residual Binarization for Image Super-Resolution	59731
<i>Yulun Zhang, Haotong Qin, Zixiang Zhao, Xianglong Liu, Martin Danelljan, Fisher Yu</i>	
Debiased Offline Representation Learning for Fast Online Adaptation in Non-Stationary Dynamics.....	59741
<i>Xinyu Zhang, Wenjie Qiu, Yi-Chen Li, Lei Yuan, Chengxing Jia, Zongzhang Zhang, Yang Yu</i>	
Provable Representation with Efficient Planning for Partially Observable Reinforcement Learning.....	59759
<i>Hongming Zhang, Tongzheng Ren, Chenjun Xiao, Dale Schuurmans, Bo Dai</i>	
Fair Risk Control: A Generalized Framework for Calibrating Multi-Group Fairness Risks	59783
<i>Lujing Zhang, Aaron Roth, Linjun Zhang</i>	
Online Matching with Stochastic Rewards: Provable Better Bound Via Adversarial Reinforcement Learning	59806
<i>Qiankun Zhang, Aocheng Shen, Boyu Zhang, Hanrui Jiang, Bingqian Du</i>	
Trustworthy Alignment of Retrieval-Augmented Large Language Models Via Reinforcement Learning	59827
<i>Zongmeng Zhang, Yufeng Shi, Jinhua Zhu, Wengang Zhou, Xiang Qi, Peng Zhang, Houqiang Li</i>	
Learning Low-Dimensional Latent Dynamics from High-Dimensional Observations: Non-Asymptotics and Lower Bounds.....	59851
<i>Yuyang Zhang, Shahriar Talebi, Na Li</i>	
Matrix Information Theory for Self-Supervised Learning	59897
<i>Yifan Zhang, Zhiquan Tan, Jingqin Yang, Weiran Huang, Yang Yuan</i>	
Switchable Decision: Dynamic Neural Generation Networks.....	59919
<i>Shujian Zhang, Korawat Tanwisuth, Chengyue Gong, Pengcheng He, Mingyuan Zhou</i>	
Interpreting and Improving Large Language Models in Arithmetic Calculation.....	59932
<i>Wei Zhang, Chaoqun Wan, Yonggang Zhang, Yiu-Ming Cheung, Xinmei Tian, Xu Shen, Jieping Ye</i>	
Directly Denoising Diffusion Models	59951
<i>Dan Zhang, Jingjing Wang, Feng Luo</i>	
Disentangled Continual Graph Neural Architecture Search with Invariant Modular Supernet.....	59975
<i>Zeyang Zhang, Xin Wang, Yijian Qin, Hong Chen, Ziwei Zhang, Xu Chu, Wenwu Zhu</i>	
GroupCover: A Secure, Efficient and Scalable Inference Framework for On-Device Model Protection Based on TEEs	59992
<i>Zheng Zhang, Na Wang, Ziqi Zhang, Yao Zhang, Tianyi Zhang, Jianwei Liu, Ye Wu</i>	
Candidate Pseudolabel Learning: Enhancing Vision-Language Models by Prompt Tuning with Unlabeled Data	60004
<i>Jiahao Zhang, Qi Wei, Feng Liu, Lei Feng</i>	
EquiPocket: An E(3)-Equivariant Geometric Graph Neural Network for Ligand Binding Site Prediction	60021
<i>Yang Zhang, Zhewei Wei, Ye Yuan, Chongxuan Li, Wenbing Huang</i>	
Exploring the Benefit of Activation Sparsity in Pre-Training	60040
<i>Zhengyan Zhang, Chaojun Xiao, Qiujieli Qin, Yankai Lin, Zhiyuan Zeng, Xu Han, Zhiyuan Liu, Ruobing Xie, Maosong Sun, Jie Zhou</i>	
Causal Representation Learning from Multiple Distributions: A General Setting	60057
<i>Kun Zhang, Shaoan Xie, Ignavier Ng, Yujia Zheng</i>	
FESSNC: Fast Exponentially Stable and Safe Neural Controller.....	60076
<i>Jingdong Zhang, Luan Yang, Qunxi Zhu, Wei Lin</i>	
Rethinking Guidance Information to Utilize Unlabeled Samples: A Label Encoding Perspective	60099
<i>Yulong Zhang, Yuan Yao, Shuhao Chen, Pengrong Jin, Yu Zhang, Jian Jin, Jiangang Lu</i>	

NExT-Chat: An LMM for Chat, Detection and Segmentation	60116
<i>Ao Zhang, Yuan Yao, Wei Ji, Zhiyuan Liu, Tat-Seng Chua</i>	
Minimax Optimality of Score-Based Diffusion Models: Beyond the Density Lower Bound Assumptions	60134
<i>Kaihong Zhang, Caitlyn H. Yin, Feng Liang, Jingbo Liu</i>	
Irregular Multivariate Time Series Forecasting: A Transformable Patching Graph Neural Networks Approach	60179
<i>Weijia Zhang, Chenlong Yin, Hao Liu, Xiaofang Zhou, Hui Xiong</i>	
Two Heads Are Better than One: Boosting Graph Sparse Training Via Semantic and Topological Awareness	60197
<i>Guibin Zhang, Yanwei Yue, Kun Wang, Junfeng Fang, Yongduo Sui, Kai Wang, Yuxuan Liang, Dawei Cheng, Shirui Pan, Tianlong Chen</i>	
Generalization Analysis for Multi-Label Learning.....	60220
<i>Yi-Fan Zhang, Min-Ling Zhang</i>	
Quantum Algorithms and Lower Bounds for Finite-Sum Optimization.....	60244
<i>Yexin Zhang, Chenyi Zhang, Cong Fang, Liwei Wang, Tongyang Li</i>	
PARDEN, Can You Repeat That? Defending Against Jailbreaks Via Repetition	60271
<i>Ziyang Zhang, Qizhen Zhang, Jakob Foerster</i>	
MLIP: Efficient Multi-Perspective Language-Image Pretraining with Exhaustive Data Utilization	60288
<i>Yu Zhang, Qi Zhang, Zixuan Gong, Yiwei Shi, Yepeng Liu, Duoqian Miao, Yang Liu, Ke Liu, Kun Yi, Wei Fan, Liang Hu, Changwei Wang</i>	
Lightweight Image Super-Resolution Via Flexible Meta Pruning.....	60305
<i>Yulun Zhang, Kai Zhang, Luc Van Gool, Martin Danelljan, Fisher Yu</i>	
Offline Training of Language Model Agents with Functions as Learnable Weights	60315
<i>Shaokun Zhang, Jiayu Zhang, Jiale Liu, Linxin Song, Chi Wang, Ranjay Krishna, Qingyun Wu</i>	
Efficient Contextual Bandits with Uninformed Feedback Graphs	60336
<i>Mengxiao Zhang, Yuheng Zhang, Haipeng Luo, Paul Mineiro</i>	
Efficient Denoising Diffusion Via Probabilistic Masking	60359
<i>Weizhong Zhang, Zhiwei Zhang, Renjie Pi, Zhongming Jin, Yuan Gao, Jieping Ye, Kani Chen</i>	
Navigating Complexity: Toward Lossless Graph Condensation Via Expanding Window Matching	60379
<i>Yuchen Zhang, Tianle Zhang, Kai Wang, Ziyao Guo, Yuxuan Liang, Xavier Bresson, Wei Jin, Yang You</i>	
Confronting Reward Overoptimization for Diffusion Models: A Perspective of Inductive and Primacy Biases	60396
<i>Ziyi Zhang, Sen Zhang, Yibing Zhan, Yong Luo, Yonggang Wen, Dacheng Tao</i>	
Uncertainty-Aware Reward-Free Exploration with General Function Approximation	60414
<i>Junkai Zhang, Weitong Zhang, Dongruo Zhou, Quanquan Gu</i>	
Feature Contamination: Neural Networks Learn Uncorrelated Features and Fail to Generalize	60446
<i>Tianren Zhang, Chujie Zhao, Guanyu Chen, Yizhou Jiang, Feng Chen</i>	
On the Expressive Power of Spectral Invariant Graph Neural Networks.....	60496
<i>Bohang Zhang, Lingxiao Zhao, Haggai Maron</i>	

PART 73

Position: Social Environment Design Should Be Further Developed for AI-Based Policy-Making.....	60527
<i>Edwin Zhang, Sadie Zhao, Tonghan Wang, Safwan Hossain, Henry Gasztowtt, Stephan Zheng, David C. Parkes, Milind Tambe, Yiling Chen</i>	

Neural Jump-Diffusion Temporal Point Processes	60541
<i>Shuai Zhang, Chuan Zhou, Yang Liu, Peng Zhang, Xixun Lin, Zhi-Ming Ma</i>	
The Emergence of Reproducibility and Consistency in Diffusion Models	60558
<i>Huijie Zhang, Jinfan Zhou, Yifu Lu, Minzhe Guo, Peng Wang, Liyue Shen, Qing Qu</i>	
Transferable Facial Privacy Protection Against Blind Face Restoration Via Domain-Consistent Adversarial Obfuscation	60591
<i>Kui Zhang, Hang Zhou, Jie Zhang, Wenbo Zhou, Weiming Zhang, Nenghai Yu</i>	
Fast Text-To-3D-Aware Face Generation and Manipulation Via Direct Cross-Modal Mapping and Geometric Regularization.....	60605
<i>Jinlu Zhang, Yiyi Zhou, Qiancheng Zheng, Xiaoxiong Du, Gen Luo, Jun Peng, Xiaoshuai Sun, Rongrong Ji</i>	
Accelerating Iterative Retrieval-Augmented Language Model Serving with Speculation.....	60626
<i>Zhihao Zhang, Alan Zhu, Lijie Yang, Yihua Xu, Lanting Li, Phitchaya Mangpo Phothilimthana, Zhihao Jia</i>	
Position: Measure Dataset Diversity, Don't Just Claim it.....	60644
<i>Dora Zhao, Jerone T. A. Andrews, Orestis Papakyriakopoulos, Alice Xiang</i>	
Long is More for Alignment:A Simple but Tough-To-Beat Baseline for Instruction Fine-Tuning.....	60674
<i>Hao Zhao, Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion</i>	
Probabilistic Inference in Language Models Via Twisted Sequential Monte Carlo	60704
<i>Stephen Zhao, Rob Breckelmans, Alireza Makhzani, Roger Grosse</i>	
Image Fusion Via Vision-Language Model.....	60749
<i>Zixiang Zhao, Lilun Deng, Haowen Bai, Yukun Cui, Zhipeng Zhang, Yulun Zhang, Haotong Qin, Dongdong Chen, Jianshe Zhang, Peng Wang, Luc Van Gool</i>	
Learning and Forgetting Unsafe Examples in Large Language Models.....	60766
<i>Jiachen Zhao, Zhun Deng, David Madras, James Zou, Mengye Ren</i>	
VideoPrism: A Foundational Visual Encoder for Video Understanding	60785
<i>Long Zhao, Nitesh B. Gundavarapu, Liangzhe Yuan, Hao Zhou, Shen Yan, Jennifer J. Sun, Luke Friedman, Rui Qian, Tobias Weyand, Yue Zhao, Rachel Hornung, Florian Schroff, Ming-Hsuan Yang, David A. Ross, Huisheng Wang, Hartwig Adam, Mikhail Sirotenko, Ting Liu, Boqing Gong</i>	
APT: Adaptive Pruning and Tuning Pretrained Language Models for Efficient Training and Inference.....	60812
<i>Bowen Zhao, Hannaneh Hajishirzi, Qingqing Cao</i>	
Subgoal-Based Demonstration Learning for Formal Theorem Proving	60832
<i>Xueliang Zhao, Wenda Li, Lingpeng Kong</i>	
Absolute Policy Optimization: Enhancing Lower Probability Bound of Performance with High Confidence.....	60866
<i>Weiye Zhao, Feihan Li, Yifan Sun, Rui Chen, Tianhao Wei, Changliu Liu</i>	
Spider: A Unified Framework for Context-Dependent Concept Segmentation.....	60906
<i>Xiaoqi Zhao, Youwei Pang, Wei Ji, Baicheng Sheng, Jiaming Zuo, Lihe Zhang, Huchuan Lu</i>	
Rethinking Adversarial Robustness in the Context of the Right to Be Forgotten.....	60927
<i>Chenxu Zhao, Wei Qian, Yangyi Li, Aobo Chen, Mengdi Huai</i>	
Quantum Implicit Neural Representations	60940
<i>Jiaming Zhao, Wenbo Qiao, Peng Zhang, Hui Gao</i>	
Is Inverse Reinforcement Learning Harder than Standard Reinforcement Learning? A Theoretical Perspective.....	60957
<i>Lei Zhao, Mengdi Wang, Yu Bai</i>	

A Statistical Theory of Regularization-Based Continual Learning.....	61021
<i>Xuyang Zhao, Huiyuan Wang, Weiran Huang, Wei Lin</i>	
Two Fists, One Heart: Multi-Objective Optimization Based Strategy Fusion for Long-Tailed Learning	61040
<i>Zhe Zhao, Pengkun Wang, Haibin Wen, Wei Xu, Song Lai, Qingfu Zhang, Yang Wang</i>	
Gradient-Based Visual Explanation for Transformer-Based CLIP	61072
<i>Chenyang Zhao, Kun Wang, Xingyu Zeng, Rui Zhao, Antoni B. Chan</i>	
CompeteAI: Understanding the Competition Dynamics of Large Language Model-Based Agents.....	61092
<i>Qinlin Zhao, Jindong Wang, Yixuan Zhang, Yiqiao Jin, Kaijie Zhu, Hao Chen, Xing Xie</i>	
Defense Against Backdoor Attack on Pre-Trained Language Models Via Head Pruning and Attention Normalization.....	61108
<i>Xingyi Zhao, Depeng Xu, Shuhan Yuan</i>	
Galore: Memory-Efficient LLM Training by Gradient Low-Rank Projection.....	61121
<i>Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, Yuandong Tian</i>	
When Will Gradient Regularization Be Harmful?.....	61144
<i>Yang Zhao, Hao Zhang, Xiuyuan Hu</i>	
LangCell: Language-Cell Pre-Training for Cell Identity Understanding	61159
<i>Suyuan Zhao, Jiahuan Zhang, Yushuai Wu, Yizhen Luo, Zaiqing Nie</i>	
Optimistic Multi-Agent Policy Gradient	61186
<i>Wenshuai Zhao, Yi Zhao, Zhiyuan Li, Juho Kannala, Joni Pajarinen</i>	
Clifford-Steerable Convolutional Neural Networks	61203
<i>Maksim Zhdanov, David Ruhe, Maurice Weiler, Ana Lucic, Johannes Brandstetter, Patrick Forré</i>	
3D-VLA: A 3D Vision-Language-Action Generative World Model	61229
<i>Haoyu Zhen, Xiaowen Qiu, Peihao Chen, Jincheng Yang, Xin Yan, Yilun Du, Yining Hong, Chuang Gan</i>	
How Does Goal Relabeling Improve Sample Efficiency?.....	61246
<i>Sirui Zheng, Chenjia Bai, Zhuoran Yang, Zhaoran Wang</i>	
PRISE: LLM-Style Sequence Compression for Learning Temporal Action Abstractions in Control	61267
<i>Ruijie Zheng, Ching-An Cheng, Hal Daumé III, Furong Huang, Andrey Kolobov</i>	
Exploiting Negative Samples: A Catalyst for Cohort Discovery in Healthcare Analytics.....	61287
<i>Kaiping Zheng, Horng-Ruey Chua, Melanie Herschel, H. V. Jagadish, Beng Chin Ooi, James Wei Luen Yip</i>	
Constrained Exploration Via Reflected Replica Exchange Stochastic Gradient Langevin Dynamics	61321
<i>Haoyang Zheng, Hengrong Du, Qi Feng, Wei Deng, Guang Lin</i>	

PART 74

GPT-4V(ision) is a Generalist Web Agent, if Grounded.....	61349
<i>Boyuan Zheng, Boyu Gou, Jihyung Kil, Huan Sun, Yu Su</i>	
Characteristic Guidance: Non-Linear Correction for Diffusion Model at Large Guidance Scale	61386
<i>Candi Zheng, Yuan Lan</i>	
Premier-TACO is a Few-Shot Policy Learner: Pretraining Multitask Representation Via Temporal Action-Driven Contrastive Loss	61413
<i>Ruijie Zheng, Yongyuan Liang, Xiyao Wang, Shuang Ma, Hal Daumé III, Huazhe Xu, John Langford, Praveen Palanisamy, Kalyan Shankar Basu, Furong Huang</i>	

ESM All-Atom: Multi-Scale Protein Language Model for Unified Molecular Modeling.....	61432
<i>Kangjie Zheng, Siyu Long, Tianyu Lu, Junwei Yang, Xinyu Dai, Ming Zhang, Zaiqing Nie, Wei-Ying Ma, Hao Zhou</i>	
Bat: Learning to Reason About Spatial Sounds with Large Language Models.....	61454
<i>Zhisheng Zheng, Puyuan Peng, Ziyang Ma, Xie Chen, Eunsol Choi, David Harwath</i>	
Conformal Predictions Under Markovian Data.....	61470
<i>Frédéric Zheng, Alexandre Proutiere</i>	
Detecting and Identifying Selection Structure in Sequential Data.....	61498
<i>Yujia Zheng, Zeyu Tang, Yiwen Qiu, Bernhard Schölkopf, Kun Zhang</i>	
ContPhy: Continuum Physical Concept Learning and Reasoning from Videos.....	61526
<i>Zhicheng Zheng, Xin Yan, Zhenfang Chen, Jingzhou Wang, Qin Zhi Eddie Lim, Joshua B. Tenenbaum, Chuang Gan</i>	
DPN: Decoupling Partition and Navigation for Neural Solvers of Min-Max Vehicle Routing Problems.....	61559
<i>Zhi Zheng, Shunyu Yao, Zhenkun Wang, Xialiang Tong, Mingxuan Yuan, Ke Tang</i>	
On Prompt-Driven Safeguarding for Large Language Models.....	61593
<i>Chujie Zheng, Fan Yin, Hao Zhou, Fandong Meng, Jie Zhou, Kai-Wei Chang, Minlie Huang, Nanyun Peng</i>	
Self-Infilling Code Generation.....	61614
<i>Lin Zheng, Jianbo Yuan, Zhi Zhang, Hongxia Yang, Lingpeng Kong</i>	
Multi-Layer Rehearsal Feature Augmentation for Class-Incremental Learning.....	61649
<i>Bowen Zheng, Da-Wei Zhou, Han-Jia Ye, De-Chuan Zhan</i>	
ERQ: Error Reduction for Post-Training Quantization of Vision Transformers.....	61664
<i>Yunshan Zhong, Jiawei Hu, You Huang, Yuxin Zhang, Rongrong Ji</i>	
Provably Efficient Exploration in Quantum Reinforcement Learning with Logarithmic Worst-Case Regret.....	61681
<i>Han Zhong, Jiachen Hu, Yecheng Xue, Tongyang Li, Liwei Wang</i>	
Towards Efficient Training and Evaluation of Robust Models Against ℓ_0 Bounded Adversarial Perturbations.....	61708
<i>Xuyang Zhong, Yixiao Huang, Chen Liu</i>	
GNNs Also Deserve Editing, and They Need it More than Once.....	61727
<i>Shaochen (Henry) Zhong, Duy Le, Zirui Liu, Zhimeng Jiang, Andrew Ye, Jiamu Zhang, Jiayi Yuan, Kaixiong Zhou, Zhaozhuo Xu, Jing Ma, Shuai Xu, Vipin Chaudhary, Xia Hu</i>	
Causal-IQA: Towards the Generalization of Image Quality Assessment Based on Causal Inference.....	61747
<i>Yan Zhong, Xingyu Wu, Li Zhang, Chenxi Yang, Tingting Jiang</i>	
Jacobian Regularizer-Based Neural Granger Causality.....	61763
<i>Wanqi Zhou, Shuanghao Bai, Shujian Yu, Qibin Zhao, Badong Chen</i>	
Differentially Private Worst-Group Risk Minimization.....	61783
<i>Xinyu Zhou, Raef Bassily</i>	
MVMoE: Multi-Task Vehicle Routing Solver with Mixture-Of-Experts.....	61804
<i>Jianan Zhou, Zhiguang Cao, Yaoxin Wu, Wen Song, Yining Ma, Jie Zhang, Chi Xu</i>	
MD Tree: A Model-Diagnostic Tree Grown on Loss Landscape.....	61825
<i>Yefan Zhou, Jianlong Chen, Qinxue Cao, Konstantin Schürholt, Yaoqing Yang</i>	
On the Emergence of Cross-Task Linearity in Pretraining-Finetuning Paradigm.....	61854
<i>Zhanpeng Zhou, Zijun Chen, Yilan Chen, Bo Zhang, Junchi Yan</i>	
RoboDreamer: Learning Compositional World Models for Robot Imagination.....	61885
<i>Siyuan Zhou, Yilun Du, Jiaben Chen, Yandong Li, Dit-Yan Yeung, Chuang Gan</i>	

MultiMax: Sparse and Multi-Modal Attention Learning.....	61897
<i>Yuxuan Zhou, Mario Fritz, Margret Keuper</i>	
Finite-Time Convergence and Sample Complexity of Actor-Critic Multi-Objective Reinforcement Learning	61913
<i>Tianchen Zhou, FNU Hairi, Haibo Yang, Jia Liu, Tian Tong, Fan Yang, Michinari Momma, Yan Gao</i>	
Attack-Free Evaluating and Enhancing Adversarial Robustness on Categorical Data.....	61934
<i>Yujun Zhou, Yufei Han, Haomin Zhuang, Hongyan Bao, Xiangliang Zhang</i>	
Pedestrian Attribute Recognition as Label-Balanced Multi-Label Learning.....	61964
<i>Yibo Zhou, Hai-Miao Hu, Yirong Xiang, Xiaokang Zhang, Haotian Wu</i>	
Bounding the Excess Risk for Linear Models Trained on Marginal-Preserving, Differentially-Private, Synthetic Data	61979
<i>Yvonne Zhou, Mingyu Liang, Ivan Brugere, Danial Dervovic, Antigoni Polychroniadou, Min Wu, Dana Dachman-Soled</i>	
Conformalized Adaptive Forecasting of Heterogeneous Trajectories	62002
<i>Yanfei Zhou, Lars Lindemann, Matteo Sesia</i>	
Sequential Kernel Goodness-Of-Fit Testing.....	62057
<i>Zhengyu Zhou, Weiwei Liu</i>	
RAUCA: A Novel Physical Adversarial Attack on Vehicle Detectors Via Robust and Accurate Camouflage Generation.....	62076
<i>Jiawei Zhou, Linye Lyu, Daojing He, Yu Li</i>	
CurBench: Curriculum Learning Benchmark.....	62088
<i>Yuwei Zhou, Zirui Pan, Xin Wang, Hong Chen, Haoyang Li, Yanwen Huang, Zhixiao Xiong, Fangzhou Xiong, Peiyang Xu, Shengnan Liu, Wenwu Zhu</i>	
GALA3D: Towards Text-To-3D Complex Scene Generation Via Layout-Guided Generative Gaussian Splatting.....	62108
<i>Xiaoyu Zhou, Xingjian Ran, Yajiao Xiong, Jinlin He, Zhiwei Lin, Yongtao Wang, Deqing Sun, Ming-Hsuan Yang</i>	
Stabilizing Policy Gradients for Stochastic Differential Equations Via Consistency with Perturbation Process	62119
<i>Xiangxin Zhou, Liang Wang, Yichi Zhou</i>	
Language Agent Tree Search Unifies Reasoning, Acting, and Planning in Language Models	62138
<i>Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, Yu-Xiong Wang</i>	
DeCoOp: Robust Prompt Tuning with Out-Of-Distribution Detection	62161
<i>Zhi Zhou, Ming Yang, Jiang-Xin Shi, Lan-Zhe Guo, Yu-Feng Li</i>	
ArCHer: Training Language Model Agents Via Hierarchical Multi-Turn RL.....	62178
<i>Yifei Zhou, Andrea Zanette, Jiayi Pan, Sergey Levine, Aviral Kumar</i>	

PART 75

Graphon Mean Field Games with a Representative Player: Analysis and Learning Algorithm	62210
<i>Fuzhong Zhou, Chenyu Zhang, Xu Chen, Xuan Di</i>	
Effective Federated Graph Matching.....	62257
<i>Yang Zhou, Zijie Zhang, Zeru Zhang, Lingjuan Lyu, Wei-Shinn Ku</i>	
Exploring Training on Heterogeneous Data with Mixture of Low-Rank Adapters.....	62294
<i>Yuhang Zhou, Zihua Zhao, Siyuan Du, Haolin Li, Jiangchao Yao, Ya Zhang, Yanfeng Wang</i>	

Score Identity Distillation: Exponentially Fast Distillation of Pretrained Diffusion Models for One-Step Generation	62307
<i>Mingyuan Zhou, Huangjie Zheng, Zhendong Wang, Mingzhang Yin, Hai Huang</i>	
Iterative Search Attribution for Deep Neural Networks	62332
<i>Zhiyu Zhu, Huaming Chen, Xinyi Wang, Jiayu Zhang, Zhibo Jin, Jason Xue, Jun Shen</i>	
Generative Active Learning for Long-Tailed Instance Segmentation.....	62349
<i>Muzhi Zhu, Chengxiang Fan, Hao Chen, Yang Liu, Weian Mao, Xiaogang Xu, Chunhua Shen</i>	
Asymmetry in Low-Rank Adapters of Foundation Models.....	62369
<i>Jiacheng Zhu, Kristjan Greenewald, Kimia Nadjahi, Haitz Sáez De Ocariz Borde, Rickard Brüel Gabriellsson, Leshem Choshen, Marzyeh Ghassemi, Mikhail Yurochkin, Justin Solomon</i>	
Improving Open-Ended Text Generation Via Adaptive Decoding	62386
<i>Wenhong Zhu, Hongkun Hao, Zhiwei He, Yiming Ai, Rui Wang</i>	
Iterative Data Smoothing: Mitigating Reward Overfitting and Overoptimization in RLHF	62405
<i>Banghua Zhu, Michael I. Jordan, Jiantao Jiao</i>	
Vision Mamba: Efficient Visual Representation Learning with Bidirectional State Space Model.....	62429
<i>Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, Xinggang Wang</i>	
Switched Flow Matching: Eliminating Singularities Via Switching ODEs.....	62443
<i>Qunxi Zhu, Wei Lin</i>	
Catapults in SGD: Spikes in the Training Loss and Their Impact on Generalization Through Feature Learning	62476
<i>Libin Zhu, Chaoyue Liu, Adityanarayanan Radhakrishnan, Mikhail Belkin</i>	
Toward Availability Attacks in 3D Point Clouds.....	62510
<i>Yifan Zhu, Yibo Miao, Yinpeng Dong, Xiao-Shan Gao</i>	
Antibody Design Using a Score-Based Diffusion Model Guided by Evolutionary, Physical and Geometric Constraints	62531
<i>Tian Zhu, Milong Ren, Haicang Zhang</i>	
Online Learning in Betting Markets: Profit Versus Prediction.....	62549
<i>Haiqing Zhu, Alexander Soen, Yun Kuen Cheung, Lexing Xie</i>	
Model Tailor: Mitigating Catastrophic Forgetting in Multi-Modal Large Language Models	62581
<i>Didi Zhu, Zhongyi Sun, Zexi Li, Tao Shen, Ke Yan, Shouhong Ding, Chao Wu, Kun Kuang</i>	
Dynamic Evaluation of Large Language Models by Meta Probing Agents	62599
<i>Kaijie Zhu, Jindong Wang, Qinlin Zhao, Ruochen Xu, Xing Xie</i>	
CRoFT: Robust Fine-Tuning with Concurrent Optimization for OOD Generalization and Open-Set OOD Detection.....	62618
<i>Lin Zhu, Yifeng Yang, Qinying Gu, Xinbing Wang, Chenghu Zhou, Nanyang Ye</i>	
Language Models Represent Beliefs of Self and Others	62638
<i>Wentao Zhu, Zhining Zhang, Yizhou Wang</i>	
Stealthy Imitation: Reward-Guided Environment-Free Policy Stealing.....	62682
<i>Zhixiong Zhuang, Maria-Irina Nicolae, Mario Fritz</i>	
Reinformer: Max-Return Sequence Modeling for Offline RL	62707
<i>Zifeng Zhuang, Dengyun Peng, Jinxin Liu, Ziqi Zhang, Donglin Wang</i>	
COALA: A Practical and Vision-Centric Federated Learning Platform.....	62723
<i>Weiming Zhuang, Jian Xu, Chen Chen, Jingtao Li, Lingjuan Lyu</i>	
GPTSwarm: Language Agents as Optimizable Graphs	62743
<i>Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbullin, Jürgen Schmidhuber</i>	

Towards Efficient Spiking Transformer: A Token Sparsification Framework for Training and Inference Acceleration.....	62768
<i>Zhengyang Zhuge, Peisong Wang, Xingting Yao, Jian Cheng</i>	
Sharp Rates in Dependent Learning Theory: Avoiding Sample Size Deflation for the Square Loss.....	62779
<i>Ingvar Ziemann, Stephen Tu, George J. Pappas, Nikolai Matni</i>	
Converting Transformers to Polynomial Form for Secure Inference Over Homomorphic Encryption.....	62803
<i>Itamar Zimmerman, Moran Baruch, Nir Drucker, Gilad Ezov, Omri Soceanu, Lior Wolf</i>	
Viewing Transformers Through the Lens of Long Convolutions Layers	62815
<i>Itamar Zimmerman, Lior Wolf</i>	
Emergence of In-Context Reinforcement Learning from Noise Distillation.....	62832
<i>Ilya Zisman, Vladislav Kurenkov, Alexander Nikulin, Viacheslav Sini, Sergey Kolesnikov</i>	
Symmetry Induces Structure and Constraint of Learning.....	62847
<i>Liu Ziyin</i>	
Safety Fine-Tuning at (Almost) No Cost: A Baseline for Vision Large Language Models	62867
<i>Yongshuo Zong, Ondrej Bohdal, Tingyang Yu, Yongxin Yang, Timothy Hospedales</i>	
Fool Your (Vision And) Language Model with Embarrassingly Simple Permutations	62892
<i>Yongshuo Zong, Tingyang Yu, Ruchika Chavhan, Bingchen Zhao, Timothy Hospedales</i>	
Leveraging Attractor Dynamics in Spatial Navigation for Better Language Parsing	62914
<i>Xiaolong Zou, Xingxing Cao, Xiaojiao Yang, Bo Hong</i>	
Hybrid ² Neural ODE Causal Modeling and an Application to Glycemic Response	62934
<i>Bob Junyi Zou, Matthew E. Levine, Dessi P. Zaharieva, Ramesh Johari, Emily B. Fox</i>	
Compositional Few-Shot Class-Incremental Learning.....	62964
<i>Yixiong Zou, Shanghang Zhang, Haichen Zhou, Yuhua Li, Ruixuan Li</i>	
BiE: Bi-Exponent Block Floating-Point for Large Language Models Quantization	62978
<i>Lancheng Zou, Wenqian Zhao, Shuo Yin, Chen Bai, Qi Sun, Bei Yu</i>	
Active Statistical Inference.....	62993
<i>Tijana Zrnic, Emmanuel J. Candès</i>	

Author Index