

Advances in Neural Information Processing Systems 35

36th Conference on Neural Information
Processing Systems (NeurIPS 2022)

New Orleans, Louisiana, USA
28 November - 9 December 2022

Volume 1 of 50

ISBN: 978-1-7138-7108-8
eISBN: 978-1-7138-7312-9

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© (2022) by individual authors and Neural Information Processing Systems Foundation Inc.
All rights reserved.

Printed with permission by Curran Associates, Inc. (2023)

For permission requests, please contact Neural Information Processing Systems Foundation Inc.
at the address below.

Neural Information Processing Systems Foundation Inc.
1269 Law Street
San Diego, CA 92109
USA

Phone: (858) 453-4100

Fax: (858) 587-0417

<https://neurips.cc/>

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

VOLUME 1

Federated Submodel Optimization for Hot and Cold Data Features	1
<i>Guihai Chen, Yucheng Ding, Yanghe Feng, Chengfei Lyu, Chaoyue Niu, Shaojie Tang, Fan Wu</i>	
On Kernelized Multi-Armed Bandits with Constraints	14
<i>Bo Ji, Xingyu Zhou</i>	
Geometric Order Learning for Rank Estimation	27
<i>Chang-Su Kim, Seon-Ho Lee, Nyeong Ho Shin</i>	
Structured Recognition for Generative Models with Explaining Away.....	40
<i>Neil Burgess, Maneesh Sahani, Hugo Soulat, Changmin Yu</i>	
NAS-Bench-Graph: Benchmarking Graph Neural Architecture Search.....	54
<i>Yijian Qin, Xin Wang, Zeyang Zhang, Ziwei Zhang, Wenwu Zhu</i>	
Fast Bayesian Coresets Via Subsampling and Quasi-Newton Refinement	70
<i>Trevor Campbell, Cian Naik, Judith Rousseau</i>	
What You See is What You Classify: Black Box Attributions	84
<i>Radhakrishna Achanta, Fernando Perez-Cruz, Nathanael Perraudin, Steven Stalder, Michele Volpi</i>	
Adaptive Interest for Emphatic Reinforcement Learning.....	95
<i>Kavosh Asadi, Rasool Fakoor, Taesup Kim, Martin Klissarov, Jonas Mueller, Alexander Smola</i>	
Scaling & Shifting Your Features: A New Baseline for Efficient Model Tuning	109
<i>Jiashi Feng, Dongze Lian, Xinchao Wang, Daquan Zhou</i>	
Zero-Shot Video Question Answering Via Frozen Bidirectional Language Models	124
<i>Ivan Laptev, Antoine Miech, Cordelia Schmid, Josef Sivic, Antoine Yang</i>	
Active Learning with Neural Networks: Insights from Nonparametric Statistics	142
<i>Robert Nowak, Yinglun Zhu</i>	
IM-Loss: Information Maximization Loss for Spiking Neural Networks	156
<i>Yuanpei Chen, Yufei Guo, Xuhui Huang, Xiaode Liu, Zhe Ma, Yinglei Wang, Liwen Zhang</i>	
Using Natural Language and Program Abstractions to Instill Human Inductive Biases in Machines.....	167
<i>Jonathan D Cohen, Carlos G. Correa, Ishita Dasgupta, Nathaniel Daw, Tom Griffiths, Robert Hawkins, Michael Y Hu, Sreejan Kumar, Raja Marjeh, Karthik Narasimhan</i>	
Second Thoughts Are Best: Learning to Re-Align with Human Values from Text Edits	181
<i>Chenyang Jia, Ruibo Liu, Tony Liu, Soroush Vosoughi, Ge Zhang, Ziyu Zhuang</i>	
SatMAE: Pre-Training Transformers for Temporal and Multi-Spectral Satellite Imagery	197
<i>Marshall Burke, Yezhen Cong, Stefano Ermon, Yutong He, Samar Khanna, Patrick Liu, David Lobell, Chenlin Meng, Erik Rozi</i>	

On Sample Optimality in Personalized Collaborative and Federated Learning	212
<i>Mathieu Even, Laurent Massoulié, Kevin Scaman</i>	
Offline Multi-Agent Reinforcement Learning with Knowledge Distillation	226
<i>Phillip Isola, Yen-Chen Lin, Wei-Cheng Tseng, Tsun-Hsuan Johnson Wang</i>	
Decentralized Gossip-Based Stochastic Bilevel Optimization Over Communication Networks	238
<i>Mengdi Wang, Shuoguang Yang, Xuezhou Zhang</i>	
Conditional Meta-Learning of Linear Representations	253
<i>Carlo Ciliberto, Giulia Denevi, Massimiliano Pontil</i>	
Theory and Approximate Solvers for Branched Optimal Transport with Multiple Sources.....	267
<i>Enrique Fita Sanmartín, Fred Hamprecht, Peter Lippmann</i>	
CHIMLE: Conditional Hierarchical IMLE for Multimodal Conditional Image Synthesis	280
<i>Ke Li, Seyed Alireza Moazenipourasil, Shichong Peng</i>	
Active Ranking Without Strong Stochastic Transitivity.....	297
<i>Farzad Farnoud, Quanquan Gu, Tao Jin, Hao Lou, Yue Wu, Pan Xu</i>	
Offline Goal-Conditioned Reinforcement Learning Via f-Advantage Regression	310
<i>Osbert Bastani, Dinesh Jayaraman, Jason Yecheng Ma, Jason Yan</i>	
Rethinking and Improving Robustness of Convolutional Neural Networks: A Shapley Value-Based Approach in Frequency Domain.....	324
<i>Yiting Chen, Qibing Ren, Junchi Yan</i>	
Adversarial Style Augmentation for Domain Generalized Urban-Scene Segmentation.....	338
<i>Gim Hee Lee, Nicu Sebe, Yuyang Zhao, Zhun Zhong</i>	
Fully Sparse 3D Object Detection	351
<i>Lue Fan, Feng Wang, Naiyan Wang, Zhao-Xiang Zhang</i>	
Diffusion Visual Counterfactual Explanations	364
<i>Maximilian Augustin, Valentyn Boreiko, Francesco Croce, Matthias Hein</i>	
Recurrent Video Restoration Transformer with Guided Deformable Attention	378
<i>Jiezhong Cao, Yuchen Fan, Luc V Gool, Simon Green, Eddy Ilg, Jingyun Liang, Rakesh Ranjan, Radu Timofte, Xiaoyu Xiang, Kai Zhang</i>	
A Consolidated Cross-Validation Algorithm for Support Vector Machines Via Data Reduction	394
<i>Boxiang Wang, Archer Yang</i>	
On-Demand Sampling: Learning Optimally from Multiple Distributions	406
<i>Nika Haghtalab, Michael Jordan, Eric Zhao</i>	
Asynchronous SGD Beats Minibatch SGD Under Arbitrary Delays	420
<i>Francis Bach, Mathieu Even, Konstantin Mishchenko, Blake Woodworth</i>	
Coresets for Relational Data and the Applications	434
<i>Jiaxiang Chen, Hu Ding, Ruomin Huang, Qingyuan Yang</i>	
Model-Based Offline Reinforcement Learning with Pessimism-Modulated Dynamics Belief.....	449
<i>Yanhui Geng, Kaiyang Guo, Shao Yunfeng</i>	

Generating Training Data with Language Models: Towards Zero-Shot Language Understanding	462
<i>Jiawei Han, Jiaxin Huang, Yu Meng, Yu Zhang</i>	
Wavelet Score-Based Generative Modeling	478
<i>Simon Coste, Valentin De Bortoli, Florentin Guth, Stephane Mallat</i>	
Robust Binary Models by Pruning Randomly-Initialized Networks	492
<i>Chen Liu, Mathieu Salzmann, Sabine Sússtrunk, Ziqi Zhao</i>	
Why Do Tree-Based Models Still Outperform Deep Learning on Typical Tabular Data?	507
<i>Leo Grinsztajn, Edouard Oyallon, Gael Varoquaux</i>	
Generalizing Consistent Multi-Class Classification with Rejection to Be Compatible with Arbitrary Losses	521
<i>Bo An, Tianchi Cai, Yuzhou Cao, Lei Feng, Lihong Gu, Jinjie Gu, Gang Niu, Masashi Sugiyama</i>	

VOLUME 2

Markovian Interference in Experiments	535
<i>Vivek Farias, Andrew Li, Tianyi Peng, Andrew Zheng</i>	
Identifiability and Generalizability from Multiple Experts in Inverse Reinforcement Learning.....	550
<i>Volkan Cevher, Boris Nikolov, Paul Rolland, Norman Schürhoff, Luca Viano</i>	
Parallel Tempering with a Variational Reference	565
<i>Alexandre Bouchard-Côté, Trevor Campbell, Nikola Surjanovic, Saifuddin Syed</i>	
Provably Efficient Reinforcement Learning in Partially Observable Dynamical Systems.....	578
<i>Nathan Kallus, Jason Lee, Ayush Sekhari, Wen Sun, Masatoshi Uehara</i>	
Off-Policy Evaluation for Episodic Partially Observable Markov Decision Processes Under Non-Parametric Models.....	593
<i>Rui Miao, Zhengling Qi, Xiaoke Zhang</i>	
Efficient Knowledge Distillation from Model Checkpoints.....	607
<i>Gao Huang, Rui Huang, Shiji Song, Chaofei Wang, Qisen Yang</i>	
Decoupled Self-Supervised Learning for Graphs.....	620
<i>Zhengyu Chen, Zhimeng Guo, Suhang Wang, Teng Xiao, Zeyang Zhuang</i>	
Shadow Knowledge Distillation: Bridging Offline and Online Knowledge Transfer	635
<i>Zhe Jin, Lujun Li</i>	
ComENet: Towards Complete and Efficient Message Passing for 3D Molecular Graphs	650
<i>Shuiwang Ji, Yuchao Lin, Haoran Liu, Yi Liu, Limei Wang</i>	
VLMbench: A Compositional Benchmark for Vision-And-Language Manipulation.....	665
<i>Xiaotong Chen, Odest Chadwicke Jenkins, Xin Wang, Kaizhi Zheng</i>	
Tiered Reinforcement Learning: Pessimism in the Face of Uncertainty and Constant Regret.....	679
<i>Wei Chen, Jiawei Huang, Nan Jiang, Tie-Yan Liu, Tao Qin, Li Zhao</i>	
Between Stochastic and Adversarial Online Convex Optimization: Improved Regret Bounds Via Smoothness.....	691
<i>Cristóbal Guzmán, Hedi Hadiji, Sarah Sachs, Tim Van Erven</i>	

Differentially Private Learning Needs Hidden State (Or Much Faster Convergence).....	703
<i>Reza Shokri, Jiayuan Ye</i>	
BR-SNIS: Bias Reduced Self-Normalized Importance Sampling	716
<i>Gabriel Cardoso, Eric Moulines, Jimmy Olsson, Sergey Samsonov, Achille Thin</i>	
Learning to Configure Computer Networks with Neural Algorithmic Reasoning	730
<i>Luca Beurer-Kellner, Laurent Vanbever, Martin Vechev, Petar Velickovic</i>	
Early Stage Convergence and Global Convergence of Training Mildly Parameterized Neural Networks	743
<i>Chao Ma, Mingze Wang</i>	
On Divergence Measures for Bayesian Pseudocoresets	757
<i>Jungwon Choi, Jung-Woo Ha, Balhae Kim, Juho Lee, Seanie Lee, Yoonho Lee</i>	
Unsupervised Learning of Equivariant Structure from Sequences	768
<i>Kenji Fukumizu, Masanori Koyama, Takeru Miyato</i>	
Multi-Class H-Consistency Bounds	782
<i>Pranjal Awasthi, Anqi Mao, Mehryar Mohri, Yutao Zhong</i>	
On the Frequency-Bias of Coordinate-MLPs	796
<i>Simon Lucey, Lachlan E. Macdonald, Sameera Ramasinghe</i>	
DC-BENCH: Dataset Condensation Benchmark	810
<i>Justin Cui, Cho-Jui Hsieh, Si Si, Ruochen Wang</i>	
Mask Matching Transformer for Few-Shot Segmentation	823
<i>Ling Chen, Siyu Jiao, Shant Navasardyan, Honghui Shi, Yunchao Wei, Gengwei Zhang, Yao Zhao</i>	
Queue Up Your Regrets: Achieving the Dynamic Capacity Region of Multiplayer Bandits	837
<i>Nicholas Bambos, Ilai Bistriz</i>	
Differentially Private Covariance Revisited	850
<i>Wei Dong, Yuting Liang, Ke Yi</i>	
Trimmed Maximum Likelihood Estimation for Robust Generalized Linear Model	862
<i>Pranjal Awasthi, Abhimanyu Das, Weihao Kong, Rajat Sen</i>	
Causal Discovery in Linear Latent Variable Models Subject to Measurement Error	874
<i>Amiremad Ghassami, Negar Kiyavash, Mohamed Nafea, Ilya Shpitser, Yuqin Yang, Kun Zhang</i>	
Density-Driven Regularization for Out-Of-Distribution Detection.....	887
<i>Wenjian Huang, Hao Wang, Chengyan Wang, Jiahao Xia, Jianguo Zhang</i>	
Sparsity in Continuous-Depth Neural Networks	901
<i>Hananeh Aliee, Ignacio Ibarra, Niki Kilbertus, Till Richter, Mikhail Solonin, Fabian Theis</i>	
Environment Diversification with Multi-Head Neural Network for Invariant Learning.....	915
<i>Bo-Wei Huang, Chang-Sheng Kao, Keng-Te Liao, Shou-De Lin</i>	
Learning Probabilistic Models from Generator Latent Spaces with Hat EBM	928
<i>Mitch Hill, Jonathan Mitchell, Erik Nijkamp, Bo Pang, Song-Chun Zhu</i>	

Learning Best Combination for Efficient N:M Sparsity.....	941
<i>Fei Chao, Rongrong Ji, Ke Li, Mingbao Lin, Zhihang Lin, Yiting Luo, Yongjian Wu, Yuxin Zhang</i>	
Why Do Artificially Generated Data Help Adversarial Robustness.....	954
<i>Guang Cheng, Qifan Song, Yue Xing</i>	
Neural Surface Reconstruction of Dynamic Scenes with Monocular RGB-D Camera.....	967
<i>Hongrui Cai, Wanquan Feng, Xuetao Feng, Yan Wang, Juyong Zhang</i>	
Global Optimal K-Medoids Clustering of One Million Samples.....	982
<i>Yankai Cao, Kaixun Hua, Jiayang Ren</i>	
Batch Multi-Fidelity Active Learning with Budget Constraints.....	995
<i>Robert Kirby, Shibo Li, Jeff M Phillips, Xin Yu, Shandian Zhe</i>	
UniCLIP: Unified Framework for Contrastive Language-Image Pre-Training.....	1008
<i>Bumsoo Kim, Jongsuk Kim, Junmo Kim, Seung Hwan Kim, Honglak Lee, Janghyeon Lee, Hyounguk Shon</i>	
Efficient Multi-Agent Communication Via Self-Supervised Information Aggregation.....	1020
<i>Feng Chen, Cong Guan, Chenghe Wang, Hao Yin, Yang Yu, Lei Yuan, Zongzhang Zhang</i>	
Accelerated Training of Physics-Informed Neural Networks (PINNs) Using Meshless Discretizations.....	1034
<i>Varun Shankar, Ramansh Sharma</i>	
DOPE: Doubly Optimistic and Pessimistic Exploration for Safe Reinforcement Learning.....	1047
<i>Archana Bura, Jean-Francois Chamberland, Aria Hasanzadezonuzy, Dileep Kalathil, Srinivas Shakkottai</i>	
Improved Regret Analysis for Variance-Adaptive Linear Bandits and Horizon-Free Linear Mixture MDPs.....	1060
<i>Kwang-Sung Jun, Yeoneung Kim, Insoon Yang</i>	
Communication-Efficient Topologies for Decentralized Learning with O(1) Consensus Rate.....	1073
<i>Xexin Jin, Weijian Li, Lei Shi, Zhuoqing Song, Ming Yan, Wotao Yin, Kun Yuan</i>	
Moderate-Fitting as a Natural Backdoor Defender for Pre-Trained Language Models.....	1086
<i>Yangyi Chen, Ganqu Cui, Yangdong Deng, Chong Fu, Ming Gu, Zhiyuan Liu, Yujia Qin, Maosong Sun, Jingang Wang, Wei Wu, Weilin Zhao, Biru Zhu</i>	
Dataset Distillation Via Factorization.....	1100
<i>Songhua Liu, Kai Wang, Xinchao Wang, Xingyi Yang, Jingwen Ye</i>	
Adaptive Sampling for Discovery.....	1114
<i>Eunjae Shim, Ambuj Tewari, Ziping Xu, Paul Zimmerman</i>	
A Large Scale Search Dataset for Unbiased Learning to Rank.....	1127
<i>Xiaokai Chu, Haitao Mao, Jiliang Tang, Shuaiqiang Wang, Wenwen Ye, Dawei Yin, Lixin Zou</i>	
SegNeXt: Rethinking Convolutional Attention Design for Semantic Segmentation.....	1140
<i>Ming-Ming Cheng, Meng-Hao Guo, Qibin Hou, Shi-Min Hu, Zhengning Liu, Cheng-Ze Lu</i>	
Understanding Hyperdimensional Computing for Parallel Single-Pass Learning.....	1157
<i>Christopher De Sa, Tao Yu, Yichi Zhang, Zhiru Zhang</i>	

Syndicated Bandits: A Framework for Auto Tuning Hyper-Parameters in Contextual Bandit Algorithms.....	1170
<i>Qin Ding, Cho-Jui Hsieh, Yue Kang, Thomas Chun Man Lee, Yi-Wei Liu, James Sharpnack</i>	
Benign, Tempered, Or Catastrophic: Toward a Refined Taxonomy of Overfitting	1182
<i>Amirhesam Abedsoltan, Misha Belkin, Neil Mallinar, Preetum Nakkiran, Parthe Pandit, James Simon</i>	
Pre-Trained Adversarial Perturbations.....	1196
<i>Yuanhao Ban, Yinpeng Dong</i>	
An Empirical Study on Disentanglement of Negative-Free Contrastive Learning.....	1210
<i>Jinkun Cao, Yang Gao, Jialei Huang, Ruiqian Nai, Qing Yang</i>	
MABSplit: Faster Forest Training Using Multi-Armed Bandits	1223
<i>Ryan Kang, Jaeyong Lee, Chris Piech, Ilan Shomorony, Sebastian Thrun, Mo Tiwari, Martin Zhang</i>	
Counterfactual Fairness with Partially Known Causal Graph	1238
<i>Mingming Gong, Bo Han, Tongliang Liu, Susan Wei, Kun Zhang, Aoqi Zuo</i>	
Controlled Sparsity Via Constrained Optimization Or: How I Learned to Stop Tuning Penalties and Love Constraints.....	1253
<i>Yoshua Bengio, Akram Erraqabi, Jose Gallego-Posada, Simon Lacoste-Julien, Juan Ramirez</i>	
Algorithms and Hardness for Learning Linear Thresholds from Label Proportions	1267
<i>Rishi Saket</i>	
Predictive Coding Beyond Gaussian Distributions	1280
<i>Thomas Lukasiewicz, Beren Millidge, Luca Pinchetti, Tommaso Salvatori, Yuhang Song, Yordan Yordanov</i>	
Semi-Supervised Active Linear Regression	1294
<i>Pranjal Awasthi, Fnu Devvrit, Nived Rajaraman</i>	
Boosting Barely Robust Learners: A New Perspective on Adversarial Robustness	1307
<i>Avrim Blum, Omar Montasser, Greg Shakhnarovich, Hongyang Zhang</i>	

VOLUME 3

Decision-Focused Learning Without Decision-Making: Learning Locally Optimized Decision Losses.....	1320
<i>Andrew Perrault, Sanket Shah, Milind Tambe, Kai Wang, Bryan Wilder</i>	
Neural Stochastic PDEs: Resolution-Invariant Learning of Continuous Spatiotemporal Dynamics.....	1333
<i>Andris Gerasimovics, Maud Lemercier, Cristopher Salvi</i>	
Okapi: Generalising Better by Making Statistical Matches Match.....	1345
<i>Myles Bartlett, Novi Quadrianto, Sara Romiti, Viktoriia Sharmanska</i>	
Revisiting Heterophily for Graph Neural Networks.....	1362
<i>Xiao-Wen Chang, Chenqing Hua, Qincheng Lu, Sitao Luan, Doina Precup, Shuyuan Zhang, Mingde Zhao, Jiaqi Zhu</i>	
Museformer: Transformer with Fine- And Coarse-Grained Attention for Music Generation	1376
<i>Wei Hu, Tie-Yan Liu, Peiling Lu, Tao Qin, Xu Tan, Rui Wang, Wei Ye, Botao Yu, Shikun Zhang</i>	

Emergent Communication: Generalization and Overfitting in Lewis Games	1389
<i>Emmanuel Dupoux, Jean-Bastien Grill, Paul Michel, Olivier Pietquin, Mathieu Rita, Florian Strub, Corentin Tallec</i>	
Towards Efficient Post-Training Quantization of Pre-Trained Language Models.....	1405
<i>Haoli Bai, Lu Hou, Xin Jiang, Irwin King, Michael R Lyu, Lifeng Shang</i>	
TweetNERD - End to End Entity Linking Benchmark for Tweets	1419
<i>Aria Haghighi, Raheleh Makki, Sneha Mehta, Shubhanshu Mishra, Ali Mollahosseini, Aman Saini</i>	
Riemannian Neural SDE: Learning Stochastic Representations on Manifolds	1434
<i>Hyomin Kim, Junseok Kwon, Kyungjae Lee, Sung Woo Park</i>	
OccGen: Selection of Real-World Multilingual Parallel Data Balanced in Gender Within Occupations	1445
<i>Christine Basta, Marta Costa-Jussà, Oriol Domingo, André Rubungo</i>	
Learning in Observable POMDPs, Without Computationally Intractable Oracles.....	1458
<i>Noah Golowich, Ankur Moitra, Dhruv Rohatgi</i>	
Cross-Image Context for Single Image Inpainting	1474
<i>Tingliang Feng, Wei Feng, Weiqi Li, Di Lin</i>	
Efficient and Effective Augmentation Strategy for Adversarial Training	1488
<i>Sravanti Addepalli, Samyak Jain, Venkatesh Babu R</i>	
Multi-Sample Training for Neural Image Compression	1502
<i>Chenjian Gao, Han Gao, Dailan He, Kunzan Liu, Hongwei Qin, Yan Wang, Tongda Xu</i>	
Adaptive Data Debiasing Through Bounded Exploration	1516
<i>Yang Liu, Parinaz Naghizadeh, Yifan Yang</i>	
Learning to Navigate Wikipedia by Taking Random Walks.....	1529
<i>Arun Ahuja, Sheila Babayan, Ishita Dasgupta, Rob Fergus, Will Grathwohl, Christine Kaeser-Chen, Kenneth Marino, John Schultz, Wendy Shang, Manzil Zaheer</i>	
When Does Return-Conditioned Supervised Learning Work for Offline Reinforcement Learning?	1542
<i>Alberto Bietti, David Brandfonbrener, Joan Bruna, Jacob Buckman, Romain Laroche</i>	
Provable Subspace Identification Under Post-Nonlinear Mixtures	1554
<i>Xiao Fu, Qi Lyu</i>	
S ³ -NeRF: Neural Reflectance Field from Shading and Shadow Under a Single Viewpoint	1568
<i>Chaofeng Chen, Guanying Chen, Zhenfang Chen, Kwan-Yee K. Wong, Wenqi Yang</i>	
AdaFocal: Calibration-Aware Adaptive Focal Loss	1583
<i>Arindam Ghosh, Matthew Gormley, Thomas Schaaf</i>	
PDEBench: An Extensive Benchmark for Scientific Machine Learning	1596
<i>Francesco Alesiani, Raphael Leiteritz, Daniel Mackinlay, Mathias Niepert, Dirk Pflüger, Timothy Praditia, Makoto Takamoto</i>	
Learning Robust Dynamics Through Variational Sparse Gating	1612
<i>Samira Ebrahimi Kahou, Danijar Hafner, Arnav Kumar Jain, Shruti Joshi, Vincent Michalski, Shivakanth Sujit</i>	

Where to Pay Attention in Sparse Training for Feature Selection?	1627
<i>Zahra Atashgahi, Decebal Constantin Mocanu, Mykola Pechenizkiy, Ghada Sokar</i>	
Maximizing Revenue Under Market Shrinkage and Market Uncertainty	1643
<i>Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm</i>	
General Cutting Planes for Bound-Propagation-Based Neural Network Verification	1656
<i>Cho-Jui Hsieh, Suman Jana, J. Zico Kolter, Bo Li, Linyi Li, Shiqi Wang, Kaidi Xu, Huan Zhang</i>	
StrokeRehab: A Benchmark Dataset for Sub-Second Action Identification	1671
<i>Carlos Fernandez-Granda, Aakash Kaku, Kangning Liu, Natasha Pandit, Avinash Parnandi, Haresh Rengaraj Rajamohan, Heidi Schambra, Kannan Venkataramanan, Anita Venkatesan, Audre Wirtanen</i>	
An α -Regret Analysis of Adversarial Bilateral Trade	1685
<i>Yossi Azar, Amos Fiat, Federico Fusco</i>	
LDSA: Learning Dynamic Subtask Assignment in Cooperative Multi-Agent Reinforcement Learning	1698
<i>Xunhan Hu, Houqiang Li, Mingyu Yang, Jian Zhao, Wengang Zhou, Jiangcheng Zhu</i>	
Mildly Conservative Q-Learning for Offline Reinforcement Learning.....	1711
<i>Xiu Li, Zongqing Lu, Jiafei Lyu, Xiaoteng Ma</i>	
Iterative Feature Matching: Toward Provable Domain Generalization with Logarithmic Environments.....	1725
<i>Yining Chen, Tengyu Ma, Andrej Risteski, Elan Rosenfeld, Mark Sellke</i>	
Certifying Robust Graph Classification Under Orthogonal Gromov-Wasserstein Threats	1737
<i>Hongwei Jin, Zishun Yu, Xinhua Zhang</i>	
Functional Ensemble Distillation	1751
<i>Idan Achituve, Ethan Fetaya, Coby Penso</i>	
Use-Case-Grounded Simulations for Explanation Evaluation	1764
<i>Valerie Chen, Nari Johnson, Gregory Plumb, Ameet Talwalkar, Nicholay Topin</i>	
Lethal Dose Conjecture on Data Poisoning.....	1776
<i>Soheil Feizi, Alexander Levine, Wenxiao Wang</i>	
Online Decision Mediation.....	1790
<i>Alihan Hüyük, Daniel Jarrett, Mihaela Van Der Schaar</i>	
Neural-Symbolic Entangled Framework for Complex Query Answering.....	1806
<i>Huajun Chen, Hui Chen, Zezhong Xu, Peng Ye, Wen Zhang</i>	
Reinforcement Learning with Automated Auxiliary Loss Search	1820
<i>Tairan He, Dongsheng Li, Minghuan Liu, Kan Ren, Che Wang, Yuqing Yang, Weinan Zhang, Yuge Zhang</i>	
FinRL-Meta: Market Environments and Benchmarks for Data-Driven Financial Reinforcement Learning	1835
<i>Jiechao Gao, Jian Guo, Xiao-Yang Liu, Jingyang Rui, Zhaoran Wang, Christina Wang, Ziyi Xia, Hongyang Yang, Ming Zhu</i>	

TempEL: Linking Dynamically Evolving and Newly Emerging Entities	1850
<i>Isabelle Augenstein, Johannes Deleu, Thomas Demeester, Chris Develder, Lucie-Aimée Kaffee, Klim Zaporozjets</i>	
M ⁴ I: Multi-Modal Models Membership Inference	1867
<i>Pingyi Hu, Ruoxi Sun, Hu Wang, Zihan Wang, Minhui Xue</i>	
Best of Both Worlds Model Selection	1883
<i>Christoph Dann, Claudio Gentile, Aldo Pacchiano</i>	
The Unreasonable Effectiveness of Fully-Connected Layers for Low-Data Regimes	1896
<i>Guillem Braso, Ismail Elezi, Peter Kocsis, Laura Leal-Taixé, Matthias Niessner, Peter Sükenik</i>	
Augmentations in Hypergraph Contrastive Learning: Fabricated and Generative	1909
<i>Tianlong Chen, Jingrui He, Yang Shen, Zhangyang Wang, Tianxin Wei, Yuning You</i>	
On the Global Convergence Rates of Decentralized Softmax Gradient Play in Markov Potential Games.....	1923
<i>Bo Dai, Na Li, Jincheng Mei, Dale Schuurmans, Runyu Zhang</i>	
Sym-NCO: Leveraging Symmetricity for Neural Combinatorial Optimization.....	1936
<i>Minsu Kim, Jinkyoo Park, Junyoung Park</i>	
Few-Shot Parameter-Efficient Fine-Tuning is Better and Cheaper than In-Context Learning	1950
<i>Mohit Bansal, Tenghao Huang, Haokun Liu, Jay Mohta, Mohammed Muqeeth, Colin Raffel, Derek Tam</i>	
HF-NeuS: Improved Surface Reconstruction Using High-Frequency Details	1966
<i>Ivan Skorokhodov, Yiqun Wang, Peter Wonka</i>	
On the Epistemic Limits of Personalized Prediction	1979
<i>Flavio Calmon, Carol Long, Lucas Monteiro Paes, Berk Ustun</i>	
DeepInteraction: 3D Object Detection Via Modality Interaction	1992
<i>Jiaqi Chen, Wei Li, Zhenwei Miao, Zeyu Yang, Li Zhang, Xiatian Zhu</i>	
Deep Differentiable Logic Gate Networks	2006
<i>Christian Borgelt, Oliver Deussen, Hilde Kuehne, Felix Petersen</i>	
Maximizing and Satisficing in Multi-Armed Bandits with Graph Information	2019
<i>Gautam Dasarathy, Mohit Malu, Nikhil Rao, Parth Thaker</i>	
MoGDE: Boosting Mobile Monocular 3D Object Detection with Ground Depth Estimation	2033
<i>Shan Chang, Minyi Guo, Yunzhe Li, Quan Liu, Yunsong Zhou, Hongzi Zhu</i>	
Causality Preserving Chaotic Transformation and Classification Using Neurochaos Learning	2046
<i>Aditi Kathpalia, Nithin Nagaraj, Harikrishnan N B</i>	
GOOD: A Graph Out-Of-Distribution Benchmark.....	2059
<i>Shurui Gui, Shuiwang Ji, Xiner Li, Limei Wang</i>	
CascadeXML: Rethinking Transformers for End-To-End Multi-Resolution Training in Extreme Multi-Label Classification.....	2074
<i>Rohit Babbar, Atmadeep Banerjee, Siddhant Kharbanda, Erik Schultheis</i>	

VF-PS: How to Select Important Participants in Vertical Federated Learning, Efficiently and Securely?	2088
<i>Lukas Burkhalter, Bolin Ding, Bo Du, Fangcheng Fu, Anwar Hithnawi, Jiawei Jiang, Bo Li, Ce Zhang</i>	

VOLUME 4

PopArt: Efficient Sparse Regression and Experimental Design for Optimal Sparse Linear Bandits	2102
<i>Kyoungseok Jang, Kwang-Sung Jun, Chicheng Zhang</i>	
Augmenting Online Algorithms with ϵ -Accurate Predictions	2115
<i>Anupam Gupta, Debmalya Panigrahi, Bernardo Subercaseaux, Kevin Sun</i>	
Unsupervised Multi-Object Segmentation by Predicting Probable Motion Patterns.....	2128
<i>Subhabrata Choudhury, Laurynas Karazija, Iro Laina, Christian Rupprecht, Andrea Vedaldi</i>	
Incrementality Bidding Via Reinforcement Learning Under Mixed and Delayed Rewards.....	2142
<i>Ashwinkumar Badanidiyuru Varadaraja, Zhe Feng, Tianxi Li, Haifeng Xu</i>	
Masked Generative Adversarial Networks Are Data-Efficient Generation Learners	2154
<i>Kaiwen Cui, Dayan Guan, Jiaxing Huang, Shengcai Liao, Shijian Lu, Aoran Xiao, Eric Xing, Fangneng Zhan</i>	
What You See is What You Get: Principled Deep Learning Via Distributional Generalization	2168
<i>Jaroslav Blasiok, Bogdan Kulynych, Preetum Nakkiran, Yao-Yuan Yang, Yaodong Yu</i>	
Towards Understanding the Condensation of Neural Networks at Initial Training	2184
<i>Tao Luo, Zhou Qixuan, Zhi-Qin Xu, Yaoyu Zhang, Hanxu Zhou</i>	
CoNT: Contrastive Neural Text Generation	2197
<i>Chenxin An, Jiangtao Feng, Xuanjing Huang, Lingpeng Kong, Kai Lv, Xipeng Qiu</i>	
GAPX: Generalized Autoregressive Paraphrase-Identification X.....	2211
<i>Hayden Housen, Renyu Li, Ser Nam Lim, Yifei Zhou</i>	
Scalable Infomin Learning	2226
<i>Yanzhi Chen, Yingzhen Li, Weihao Sun, Adrian Weller</i>	
Learning to Accelerate Partial Differential Equations Via Latent Global Evolution	2240
<i>Jure Leskovec, Takashi Maruyama, Tailin Wu</i>	
Monte Carlo Augmented Actor-Critic for Sparse Reward Deep Reinforcement Learning from Suboptimal Demonstrations	2254
<i>Ashwin Balakrishna, Wyame Benslimane, Daniel Brown, Jules Dedieu, Ken Goldberg, Albert Wilcox</i>	
Not Too Little, Not Too Much: A Theoretical Analysis of Graph (over)smoothing	2268
<i>Nicolas Keriven</i>	
Kernel Similarity Matching with Hebbian Networks	2282
<i>Kyle Luther, Sebastian Seung</i>	
HumanLiker: A Human-Like Object Detector to Model the Manual Labeling Process.....	2294
<i>Ping Guo, Chenglong Liu, Peng Wang, Haoran Wei, Yangguang Zhu</i>	

Scalable Representation Learning in Linear Contextual Bandits with Constant Regret Guarantees.....	2307
<i>Alessandro Lazaric, Matteo Papini, Matteo Pirootta, Andrea Tirinzoni, Ahmed Touati</i>	
DevFly: Bio-Inspired Development of Binary Connections for Locality Preserving Sparse Codes	2320
<i>Rana Alkhoury Maroun, Qinghai Guo, Barbara Webb, Tianqi Wei</i>	
Why Neural Networks Find Simple Solutions: The Many Regularizers of Geometric Complexity	2333
<i>David Barrett, Benoit Dherin, Michael Munn, Mihaela Rosca</i>	
Efficiently Computing Local Lipschitz Constants of Neural Networks Via Bound Propagation	2350
<i>Cho-Jui Hsieh, J. Zico Kolter, Zhouxing Shi, Yihan Wang, Huan Zhang</i>	
A Causal Analysis of Harm	2365
<i>Sander Beckers, Hana Chockler, Joseph Halpern</i>	
Seeing the Forest and the Tree: Building Representations of Both Individual and Collective Dynamics with Transformers.....	2377
<i>Mehdi Azabou, Max Dabagia, Eva Dyer, Ran Liu, Jingyun Xiao</i>	
Semi-Supervised Learning with Decision Trees: Graph Laplacian Tree Alternating Optimization	2392
<i>Miguel A. Carreira-Perpinan, Arman Zharmagambetov</i>	
Riemannian Score-Based Generative Modelling.....	2406
<i>Valentin De Bortoli, Arnaud Doucet, Michael Hutchinson, Emile Mathieu, Yee Whye Teh, James Thornton</i>	
Intra-Agent Speech Permits Zero-Shot Task Acquisition	2423
<i>Joshua Abramson, Federico Carnevale, Petko I Georgiev, Aurelia Guy, Chia-Chun Hung, Timothy Lillicrap, Alistair Muldal, Adam Santoro, Gregory Wayne, Chen Yan</i>	
Free Probability for Predicting the Performance of Feed-Forward Fully Connected Neural Networks	2439
<i>Reda Chhaibi, Tariq Daouda, Ezechiel Kahn</i>	
The Minority Matters: A Diversity-Promoting Collaborative Metric Learning Algorithm	2451
<i>Shilong Bao, Xiaochun Cao, Yuan He, Qingming Huang, Qianqian Xu, Zhiyong Yang</i>	
Open-Ended Reinforcement Learning with Neural Reward Functions	2465
<i>Robert Meier, Asier Mujika</i>	
A Reduction to Binary Approach for Debiasing Multiclass Datasets	2480
<i>Ibrahim Alabdulmohsin, Sanmi Koyejo, Jessica Schrouff</i>	
Obj2Seq: Formatting Objects as Sequences with Class Prompt for Visual Tasks	2494
<i>Zhiyang Chen, Wei Li, Zhaowen Li, Ming Tang, Haixin Wang, Jinqiao Wang, Liwei Wu, Fan Yang, Chaoyang Zhao, Rui Zhao, Yousong Zhu</i>	
Learn to Explain: Multimodal Reasoning Via Thought Chains for Science Question Answering	2507
<i>Kai-Wei Chang, Peter Clark, Ashwin Kalyan, Pan Lu, Swaroop Mishra, Liang Qiu, Oyvind Tafford, Tanglin Xia, Song-Chun Zhu</i>	
Few-Shot Audio-Visual Learning of Environment Acoustics	2522
<i>Ziad Al-Halah, Changan Chen, Kristen Grauman, Sagnik Majumder</i>	
The Phenomenon of Policy Churn.....	2537
<i>Andre Barreto, Georg Ostrovski, John Quan, Tom Schaul</i>	

Molecule Generation by Principal Subgraph Mining and Assembling.....	2550
<i>Wenbing Huang, Xiangzhe Kong, Yang Liu, Zhixing Tan</i>	
Implicit Neural Representations with Levels-Of-Experts.....	2564
<i>Serge Belongie, Zekun Hao, Ming-Yu Liu, Arun Mallya</i>	
Planning for Sample Efficient Imitation Learning.....	2577
<i>Qifeng Chen, Yang Gao, Weirui Ye, Zhao-Heng Yin</i>	
Concept Activation Regions: A Generalized Framework for Concept-Based Explanations.....	2590
<i>Jonathan Crabbé, Mihaela Van Der Schaar</i>	
Towards Safe Reinforcement Learning with a Safety Editor Policy	2608
<i>Wei Xu, Haonan Yu, Haichao Zhang</i>	
Understanding Cross-Domain Few-Shot Learning Based on Domain Similarity and Few-Shot Difficulty	2622
<i>Namgyu Ho, Jin-Hwa Kim, Sungnyun Kim, Jaehoon Oh, Hwanjun Song, Se-Young Yun</i>	
Ambiguous Images with Human Judgments for Robust Visual Event Classification	2637
<i>Reno Kriz, Anqi Liu, Kate Sanders, Benjamin Van Durme</i>	
Sustainable Online Reinforcement Learning for Auto-Bidding	2651
<i>Rongquan Bai, Yusen Huo, Zhiyu Mou, Mingzhou Xie, Jian Xu, Chuan Yu, Bo Zheng</i>	
Uni-Perceiver-MoE: Learning Sparse Generalist Models with Conditional MoEs.....	2664
<i>Jifeng Dai, Hongsheng Li, Xiaogang Wang, Wenhai Wang, Xiaohua Wang, Jinguo Zhu, Xizhou Zhu</i>	
Improved Coresets for Euclidean k-Means	2679
<i>Vincent Cohen-Addad, Kasper Green Larsen, David Saulpic, Chris Schwiegelshohn, Omar Ali Sheikh-Omar</i>	
Accelerated Linearized Laplace Approximation for Bayesian Deep Learning	2695
<i>Zhijie Deng, Feng Zhou, Jun Zhu</i>	
Learning to Reason with Neural Networks: Generalization, Unseen Data and Boolean Measures.....	2709
<i>Emmanuel Abbe, Samy Bengio, Elisabetta Cornacchia, Jon Kleinberg, Aryo Lotfi, Maithra Raghu, Chiyuan Zhang</i>	
Using Partial Monotonicity in Submodular Maximization.....	2723
<i>Moran Feldman, Loay Mualem</i>	
Enhanced Meta Reinforcement Learning Via Demonstrations in Sparse Reward Environments	2737
<i>Sapana Chaudhary, Dileep Kalathil, Jaewon Kim, Desik Rengarajan, Srinivas Shakkottai</i>	
Riemannian Diffusion Models.....	2750
<i>Milad Aghajohari, Joey Bose, Aaron Courville, Chin-Wei Huang, Prakash Panangaden</i>	
Training and Inference on Any-Order Autoregressive Models the Right Way	2762
<i>Stefano Ermon, Dorsa Sadigh, Andy Shih</i>	
Lazy and Fast Greedy MAP Inference for Determinantal Point Process.....	2776
<i>Kaito Fujii, Shinichi Hemmi, Satoru Iwata, Taihei Oki, Shinsaku Sakaue</i>	
Generalization Bounds for Stochastic Gradient Descent Via Localized ϵ -Covers	2790
<i>Murat Erdogdu, Sejun Park, Umut Simsekli</i>	

Semi-Supervised Semantic Segmentation Via Gentle Teaching Assistant	2803
<i>Ying Jin, Dahua Lin, Jiaqi Wang</i>	
Anytime-Valid Inference for Multinomial Count Data.....	2817
<i>Michael Lindon, Alan Malek</i>	
What I Cannot Predict, I Do Not Understand: A Human-Centered Evaluation Framework for Explainability Methods	2832
<i>Remi Cadene, Julien Colin, Thomas Fel, Thomas Serre</i>	
S4ND: Modeling Images and Videos as Multidimensional Signals with State Spaces	2846
<i>Stephen Baccus, Tri Dao, Gordon Downs, Karan Goel, Albert Gu, Eric Nguyen, Christopher Ré, Preey Shah</i>	
Understanding Programmatic Weak Supervision Via Source-Aware Influence Function	2862
<i>Cheng-Yu Hsieh, Alexander Ratner, Haonan Wang, Jieyu Zhang</i>	
Approximation with CNNs in Sobolev Space: With Applications to Classification	2876
<i>Jian Huang, Yuling Jiao, Yuanyuan Lin, Guohao Shen</i>	

VOLUME 5

Sample Complexity of Learning Heuristic Functions for Greedy-Best-First and A* Search.....	2889
<i>Taihei Oki, Shinsaku Sakaue</i>	
TransTab: Learning Transferable Tabular Transformers Across Tables.....	2902
<i>Jimeng Sun, Zifeng Wang</i>	
Perceptual Attacks of No-Reference Image Quality Models with Human-In-The-Loop	2916
<i>Guodong Guo, Dingquan Li, Kede Ma, Xionghuo Min, Xiaokang Yang, Guangtao Zhai, Weixia Zhang</i>	
Sketch-GNN: Scalable Graph Neural Networks with Sublinear Training Complexity	2930
<i>Bang An, Mucong Ding, Furong Huang, Tahseen Rabbani, Evan Wang</i>	
Tree Mover's Distance: Bridging Graph Metrics and Stability of Graph Neural Networks	2944
<i>Ching-Yao Chuang, Stefanie Jegelka</i>	
DivBO: Diversity-Aware CASH for Ensemble Learning.....	2958
<i>Bin Cui, Yang Li, Yupeng Lu, Yu Shen, Yaofeng Tu, Wentao Zhang</i>	
Revisiting Graph Contrastive Learning from the Perspective of Graph Spectrum.....	2972
<i>Deyu Bo, Nian Liu, Jian Pei, Chuan Shi, Xiao Wang</i>	
Functional Indirection Neural Estimator for Better Out-Of-Distribution Generalization.....	2984
<i>Thai Hung Le, Man Ngo, Kha Pham, Truyen Tran</i>	
Combinatorial Bandits with Linear Constraints: Beyond Knapsacks and Fairness.....	2997
<i>Zhixuan Fang, Qingsong Liu, Siwei Wang, Weihang Xu</i>	
Will Bilevel Optimizers Benefit from Loops.....	3011
<i>Kaiyi Ji, Yingbin Liang, Mingrui Liu, Lei Ying</i>	
Combining Explicit and Implicit Regularization for Efficient Learning in Deep Networks	3024
<i>Dan Zhao</i>	

MBW: Multi-View Bootstrapping in the Wild	3039
<i>Tim Clifford, Mosam Dabhi, Ian Fasel, László Jeni, Simon Lucey, Chaoyang Wang</i>	
On a Mallows-Type Model for (Ranked) Choices.....	3052
<i>Yifan Feng, Yuxuan Tang</i>	
(De-)Randomized Smoothing for Decision Stump Ensembles.....	3066
<i>Marc Fischer, Miklós Horváth, Mark Müller, Martin Vechev</i>	
Learning to Break the Loop: Analyzing and Mitigating Repetitions for Neural Text Generation.....	3082
<i>Deng Cai, Jian Li, Huayang Li, Xiaojiang Liu, Jin Xu, Jianhao Yan</i>	
Debiased Machine Learning Without Sample-Splitting for Stable Estimators.....	3096
<i>Morgane Austern, Qizhao Chen, Vasilis Syrgkanis</i>	
Near-Optimal Sample Complexity Bounds for Constrained MDPs	3110
<i>Csaba Szepesvari, Sharan Vaswani, Lin Yang</i>	
Integral Probability Metrics PAC-Bayes Bounds	3123
<i>Ron Amit, Baruch Epstein, Ron Meir, Shay Moran</i>	
Bellman Residual Orthogonalization for Offline Reinforcement Learning.....	3137
<i>Martin J Wainwright, Andrea Zanette</i>	
Quantum Speedups of Optimizing Approximately Convex Functions with Applications to Logarithmic Regret Stochastic Convex Bandits.....	3152
<i>Tongyang Li, Ruizhe Zhang</i>	
Learning Neural Acoustic Fields	3165
<i>Yilun Du, Chuang Gan, Andrew Luo, Michael Tarr, Josh Tenenbaum, Antonio Torralba</i>	
A Universal Error Measure for Input Predictions Applied to Online Graph Problems.....	3178
<i>Giulia Bernardini, Alexander Lindermayr, Alberto Marchetti-Spaccamela, Nicole Megow, Leen Stougie, Michelle Sweering</i>	
Online Reinforcement Learning for Mixed Policy Scopes.....	3191
<i>Elias Bareinboim, Junzhe Zhang</i>	
Self-Explaining Deep Models with Logic Rule Reasoning.....	3203
<i>Meeyoung Cha, Sungwon Han, Seungeon Lee, Xiting Wang, Xing Xie, Xiaoyuan Yi</i>	
XTC: Extreme Compression for Pre-Trained Transformers Made Simple and Efficient.....	3217
<i>Yuxiong He, Conglong Li, Xiaoxia Wu, Zhewei Yao, Minjia Zhang</i>	
Meta-Album: Multi-Domain Meta-Dataset for Few-Shot Image Classification	3232
<i>Dustin Carrión-Ojeda, Sergio Escalera, Isabelle Guyon, Mike Huisman, Felix Mohr, Haozhe Sun, Ihsan Ullah, Jan N. Van Rijn, Joaquin Vanschoren, Phan Anh Vu</i>	
S ³ GC: Scalable Self-Supervised Graph Clustering	3248
<i>Fnu Devvrit, Inderjit Dhillon, Prateek Jain, Aditya Sinha</i>	
Contrastive Neural Ratio Estimation	3262
<i>Patrick Forré, Benjamin K Miller, Christoph Weniger</i>	
An Information-Theoretic Framework for Deep Learning.....	3279
<i>Hong Jun Jeon, Benjamin Van Roy</i>	

Uncoupled Learning Dynamics with $O(\log T)$ Swap Regret in Multiplayer Games.....	3292
<i>Ioannis Anagnostides, Gabriele Farina, Christian Kroer, Chung-Wei Lee, Haipeng Luo, Tuomas Sandholm</i>	
Robust Semi-Supervised Learning When Not All Classes Have Labels.....	3305
<i>Lan-Zhe Guo, Yu-Feng Li, Jie-Jing Shao, Zhi-Fan Wu, Yi-Ge Zhang</i>	
Private Multiparty Perception for Navigation	3318
<i>Mia Chiquier, Hui Lu, Carl Vondrick</i>	
Improving Task-Specific Generalization in Few-Shot Learning Via Adaptive Vicinal Risk Minimization	3329
<i>Long-Kai Huang, Ying Wei</i>	
EgoTaskQA: Understanding Human Tasks in Egocentric Videos	3343
<i>Siyuan Huang, Baoxiong Jia, Ting Lei, Song-Chun Zhu</i>	
C-Mixup: Improving Generalization in Regression	3361
<i>Chelsea Finn, Yiping Wang, Huaxiu Yao, Linjun Zhang, James Zou</i>	
Generalised Mutual Information for Discriminative Clustering.....	3377
<i>Charles Bouveyron, Arnaud Droit, Warith Harchaoui, Mickaël Leclercq, Pierre-Alexandre Mattei, Louis Ohl, Frederic Precioso</i>	
Consistent Interpolating Ensembles Via the Manifold-Hilbert Kernel	3391
<i>Clay Scott, Yutong Wang</i>	
Geo-Neus: Geometry-Consistent Neural Implicit Surfaces Learning for Multi-View Reconstruction	3403
<i>Qiancheng Fu, Yew Soon Ong, Wenbing Tao, Qingshan Xu</i>	
Sublinear Algorithms for Hierarchical Clustering.....	3417
<i>Arpit Agarwal, Sanjeev Khanna, Huan Li, Prathamesh Patil</i>	
Is Sortition Both Representative and Fair?.....	3431
<i>Soroush Ebadian, Gregory Kehne, Evi Micha, Ariel Procaccia, Nisarg Shah</i>	
Beyond Rewards: A Hierarchical Perspective on Offline Multiagent Behavioral Analysis	3444
<i>Yannick Assogba, Lucas Dixon, Andrei Kapishnikov, Been Kim, Shayegan Omidshafiei</i>	
Dynamic Pricing and Assortment Under a Contextual MNL Demand.....	3461
<i>Vineet Goyal, Noemie Perivier</i>	
DGD ² : A Linearly Convergent Distributed Algorithm for High-Dimensional Statistical Recovery	3475
<i>Marie Maros, Gesualdo Scutari</i>	
Pseudo-Riemannian Graph Convolutional Networks.....	3488
<i>Shirui Pan, Nico Potyka, Steffen Staab, Bo Xiong, Chuan Zhou, Shichao Zhu</i>	
CroCo: Self-Supervised Pre-Training for 3D Vision Tasks by Cross-View Completion.....	3502
<i>Leonid Antsfeld, Vaibhav Arora, Romain Brégier, Yohann Cabon, Boris Chidlovskii, Gabriela Csurka, Vincent Leroy, Thomas Lucas, Jerome Revaud, Philippe Weinzaepfel</i>	
Sound and Complete Verification of Polynomial Networks.....	3517
<i>Elias Abad Rocamora, Volkan Cevher, Grigorios Chrysos, Fanghui Liu, Mehmet Fatih Sahin</i>	
Better SGD Using Second-Order Momentum.....	3530
<i>Ashok Cutkosky, Hoang Tran</i>	

Learning Predictions for Algorithms with Predictions	3542
<i>Maria-Florina Balcan, Misha Khodak, Ameet Talwalkar, Sergei Vassilvitskii</i>	
Unsupervised Point Cloud Completion and Segmentation by Generative Adversarial Autoencoding Network.....	3556
<i>Jie Guo, Yanwen Guo, Changfeng Ma, Fei Pan, Chongjun Wang, Yang Yang</i>	
CalFAT: Calibrated Federated Adversarial Training with Label Skewness.....	3569
<i>Chen Chen, Yuchen Liu, Lingjuan Lyu, Xingjun Ma</i>	
Rethinking Generalization in Few-Shot Classification.....	3582
<i>Tom Drummond, Mehrtash Harandi, Markus Hiller, Rongkai Ma</i>	
Stimulative Training of Residual Networks: A Social Psychology Perspective of Loafing	3596
<i>Tao Chen, Baopu Li, Wanli Ouyang, Shengji Tang, Peng Ye</i>	
EGSDE: Unpaired Image-To-Image Translation Via Energy-Guided Stochastic Differential Equations.....	3609
<i>Fan Bao, Chongxuan Li, Min Zhao, Jun Zhu</i>	
Cryptographic Hardness of Learning Halfspaces with Massart Noise	3624
<i>Ilias Diakonikolas, Daniel Kane, Pasin Manurangsi, Lisheng Ren</i>	
Frank-Wolfe-Based Algorithms for Approximating Tyler's M-Estimator	3637
<i>Lior Danon, Dan Garber</i>	
Reinforcement Learning with Non-Exponential Discounting	3649
<i>Heinz Koeppl, Constantin Rothkopf, Matthias Schultheis</i>	
Picking on the Same Person: Does Algorithmic Monoculture Lead to Outcome Homogenization?.....	3663
<i>Rishi Bommasani, Kathleen A. Creel, Dan Jurafsky, Ananya Kumar, Percy Liang</i>	

VOLUME 6

Causal Identification Under Markov Equivalence: Calculus, Algorithm, and Completeness	3679
<i>Elias Bareinboim, Amin Jaber, Adele Ribeiro, Jiji Zhang</i>	
Unravelling the Performance of Physics-Informed Graph Neural Networks for Dynamical Systems.....	3691
<i>Ravinder Bhattoo, Suresh Bishnoi, N M Anoop Krishnan, Gunjan Kumar, Sayan Ranu, Abishek Thangamuthu</i>	
Dynamic Fair Division with Partial Information.....	3703
<i>Gerdus Benade, Daniel Halpern, Alexandros Psomas</i>	
Generalized Variational Inference in Function Spaces: Gaussian Measures Meet Bayesian Deep Learning	3716
<i>Robert Hu, Dino Sejdinovic, Veit David Wild</i>	
Communicating Natural Programs to Humans and Machines.....	3731
<i>Sam Acquaviva, Gabrielle Ecanow, Marta Kryven, Maxwell Nye, Yewen Pu, Theodoros Sechopoulos, Josh Tenenbaum, Michael Tessler, Catherine Wong</i>	
SCAMPS: Synthetics for Camera Measurement of Physiological Signals	3744
<i>Tadas Baltrusaitis, Javier Hernandez, Brian Hill, Jonathan Lester, Xin Liu, Daniel McDuff, Miah Wander</i>	

A Closer Look at Learned Optimization: Stability, Robustness, and Inductive Biases	3758
<i>James Harrison, Luke Metz, Jascha Sohl-Dickstein</i>	
"Lossless" Compression of Deep Neural Networks: A High-Dimensional Neural Tangent Kernel Approach	3774
<i>Yongqi Du, Lingyu Gu, Zhenyu Liao, Shiliang Pu, Robert Qiu, Di Xie, Yuan Zhang</i>	
Privacy of Noisy Stochastic Gradient Descent: More Iterations Without More Privacy Loss	3788
<i>Jason Altschuler, Kunal Talwar</i>	
Theseus: A Library for Differentiable Nonlinear Optimization.....	3801
<i>Brandon Amos, Stuart Anderson, Ricky T. Q. Chen, Daniel Detone, Jing Dong, Taosha Fan, Maurizio Monge, Mustafa Mukadam, Joseph Ortiz, Luis Pineda, Paloma Sodhi, Shobha Venkataraman, Austin Wang</i>	
K-Radar: 4D Radar Object Detection for Autonomous Driving in Various Weather Conditions.....	3819
<i>Seung-Hyun Kong, Dong-Hee Paek, Kevin Tirta Wijaya</i>	
Asymmetric Temperature Scaling Makes Larger Networks Teach Well Again.....	3830
<i>Wen-Shu Fan, Bingshuai Li, Xin-Chun Li, Yinchuan Li, Shaoming Song, Shao Yunfeng, De-Chuan Zhan</i>	
Solving Quantitative Reasoning Problems with Language Models.....	3843
<i>Anders Andreassen, Cem Anil, David Dohan, Ethan Dyer, Guy Gur-Ari, Theo Gutman-Solo, Aitor Lewkowycz, Henryk Michalewski, Vedant Misra, Behnam Neyshabur, Vinay Ramasesh, Imanol Schlag, Ambrose Slone, Yuhuai Wu</i>	
Structural Knowledge Distillation for Object Detection	3858
<i>Marius Cordts, Philip De Rijk, Dariu Gavrilă, Lukas Schneider</i>	
Thompson Sampling Efficiently Learns to Control Diffusion Processes	3871
<i>Mohsen Bayati, Mohamad Kazem Shirani Faradonbeh, Mohamad Sadegh Shirani Faradonbeh</i>	
Discrete Compositional Representations as an Abstraction for Goal Conditioned Reinforcement Learning	3885
<i>Yoshua Bengio, Anirudh Goyal, Riashat Islam, Kenji Kawaguchi, Alex Lamb, Romain Laroche, Xin Li, Remi Tachet Des Combes, Hongyu Zang</i>	
Graph Convolution Network Based Recommender Systems: Learning Guarantee and Item Mixture Powered Strategy.....	3900
<i>Enhong Chen, Leyan Deng, Defu Lian, Chenwang Wu</i>	
Local Metric Learning for Off-Policy Evaluation in Contextual Bandits with Continuous Actions	3913
<i>Yunseon Choi, Wonseok Jeon, Kee-Eung Kim, Byung-Jun Lee, Haanvid Lee, Jongmin Lee, Yung-Kyun Noh</i>	
Flare7K: A Phenomenological Nighttime Flare Removal Dataset	3926
<i>Yuekun Dai, Ruicheng Feng, Chongyi Li, Chen Change Loy, Shangchen Zhou</i>	
USB: A Unified Semi-Supervised Learning Benchmark for Classification	3938
<i>Hao Chen, Yue Fan, Lan-Zhe Guo, Wenxin Hou, Yu-Feng Li, Satoshi Nakamura, Heli Qi, Bhiksha Raj, Marios Savvides, Bernt Schiele, Takahiro Shinozaki, Wang Sun, Ran Tao, Jindong Wang, Renjie Wang, Yidong Wang, Zhen Wu, Xing Xie, Linyi Yang, Wei Ye, Yue Zhang, Zhi Zhou</i>	

Nocturne: A Scalable Driving Benchmark for Bringing Multi-Agent Learning One Step Closer to the Real World.....	3962
<i>Brandon Amos, Jakob Foerster, Nathan Lichtlé, Eugene Vinitsky, Xiaomeng Yang</i>	
Quasi-Newton Methods for Saddle Point Problems.....	3975
<i>Chengchang Liu, Luo Luo</i>	
Self-Supervised Contrastive Pre-Training for Time Series Via Time-Frequency Consistency.....	3988
<i>Theodoros Tsiligkaridis, Xiang Zhang, Ziyuan Zhao, Marinka Zitnik</i>	
Uncalibrated Models Can Improve Human-AI Collaboration.....	4004
<i>Tobias Gerstenberg, Kailas Vodrahalli, James Zou</i>	
Learning Dynamical Systems Via Koopman Operator Regression in Reproducing Kernel Hilbert Spaces.....	4017
<i>Carlo Ciliberto, Vladimir Kostic, Andreas Maurer, Pietro Novelli, Massimiliano Pontil, Lorenzo Rosasco</i>	
Self-Supervised Surround-View Depth Estimation with Volumetric Feature Fusion.....	4032
<i>Junhwa Hur, Seong-Gyun Jeong, Jung-Hee Kim, Tien Phuoc Nguyen</i>	
CARLANE: A Lane Detection Benchmark for Unsupervised Domain Adaptation from Simulation to Multiple Real-World Domains	4046
<i>Julian Gebele, Johann Haselberger, Bonifaz Stuhr</i>	
On Enforcing Better Conditioned Meta-Learning for Rapid Few-Shot Adaptation	4059
<i>Tom Drummond, Mehrtash Harandi, Markus Hiller</i>	
Oracle-Efficient Online Learning for Smoothed Adversaries.....	4072
<i>Nika Haghtalab, Yanjun Han, Abhishek Shetty, Kunhe Yang</i>	
A Policy-Guided Imitation Approach for Offline Reinforcement Learning	4085
<i>Jiang Li, Jianxiong Li, Haoran Xu, Xianyuan Zhan</i>	
Sample-Efficient Learning of Correlated Equilibria in Extensive-Form Games.....	4099
<i>Yu Bai, Song Mei, Ziang Song</i>	
VectorAdam for Rotation Equivariant Geometry Optimization	4111
<i>Alec Jacobson, Selena Zihan Ling, Nicholas Sharp</i>	
Matrix Multiplicative Weights Updates in Quantum Zero-Sum Games: Conservation Laws & Recurrence.....	4123
<i>Rahul Jain, Georgios Piliouras, Ryann Sim</i>	
On the Convergence Theory for Hessian-Free Bilevel Algorithms.....	4136
<i>Kaiyi Ji, Yingbin Liang, Daouda Sow</i>	
Equivariant Networks for Crystal Structures.....	4150
<i>Oumar Kaba, Siamak Ravanbakhsh</i>	
A General Framework for Auditing Differentially Private Machine Learning.....	4165
<i>Francis Ferraro, Maya Fuchs, Tyler Leblond, Fred Lu, Joseph Munoz, Edward Raff, Brian Testa, Elliott Zaresky-Williams</i>	
Generalization Analysis on Learning with a Concurrent Verifier.....	4177
<i>Kengo Nakamura, Masaaki Nishino, Norihito Yasuda</i>	

Spartan: Differentiable Sparsity Via Regularized Transportation.....	4189
<i>Ser Nam Lim, Kai Sheng Tai, Taipeng Tian</i>	
Focal Modulation Networks	4203
<i>Xiyang Dai, Jianfeng Gao, Chunyuan Li, Jianwei Yang</i>	
HSurf-Net: Normal Estimation for 3D Point Clouds by Learning Hyper Surfaces.....	4218
<i>Jin-San Cheng, Yi Fang, Zhizhong Han, Qing Li, Yu-Shen Liu, Cheng Wang</i>	
Robust Streaming PCA.....	4231
<i>Daniel Bienstock, Minchan Jeong, Apurv Shukla, Se-Young Yun</i>	
NeMF: Neural Motion Fields for Kinematic Animation	4244
<i>Chengnan He, Holly Rushmeier, Jun Saito, James Zachary, Yi Zhou</i>	
Global Normalization for Streaming Speech Recognition in a Modular Framework.....	4257
<i>Cyril Allauzen, Michael D Riley, David Rybach, Matt Shannon, Ehsan Variani, Ke Wu</i>	
Resource-Adaptive Federated Learning with All-In-One Neural Composition	4270
<i>Pengfei Guo, Yiqun Mei, Vishal Patel, Mo Zhou</i>	
SoteriaFL: A Unified Framework for Private Federated Learning with Communication Compression.....	4285
<i>Yuejie Chi, Boyue Li, Zhize Li, Haoyu Zhao</i>	
Your Transformer May Not Be as Powerful as You Expect	4301
<i>Di He, Shanda Li, Tie-Yan Liu, Shengjie Luo, Liwei Wang, Shuxin Zheng</i>	
Redundancy-Free Message Passing for Graph Neural Networks.....	4316
<i>Rongqin Chen, Ye Li, Leong Hou U, Shenghui Zhang</i>	
Diffusion-LM Improves Controllable Text Generation	4328
<i>Ishaan Gulrajani, Tatsunori Hashimoto, Xiang Li, Percy Liang, John Thickstun</i>	
Making Sense of Dependence: Efficient Black-Box Explanations Using Dependence Measure.....	4344
<i>Thomas Fel, Paul Novello, David Vigouroux</i>	
Energy-Based Contrastive Learning of Visual Representations	4358
<i>Beomsu Kim, Jong Chul Ye</i>	
Why Robust Generalization in Deep Learning is Difficult: Perspective of Expressive Power	4370
<i>John Hopcroft, Jikai Jin, Binghui Li, Liwei Wang, Han Zhong</i>	
Asynchronous Actor-Critic for Multi-Agent Reinforcement Learning	4385
<i>Christopher Amato, Weihao Tan, Yuchen Xiao</i>	
Polynomial Neural Fields for Subband Decomposition and Manipulation	4401
<i>Jonathan Barron, Serge Belongie, Sagie Benaim, Thomas Funkhouser, Kyle Genova, Bharath Hariharan, Varun Jampani, Guandao Yang</i>	
On the Generalizability and Predictability of Recommender Systems.....	4416
<i>John Dickerson, Sujay Khandagale, Duncan McElfresh, Jonathan Valverde, Colin White</i>	
Optimal Learning Rates for Regularized Conditional Mean Embedding Learning.....	4433
<i>Arthur Gretton, Zhu Li, Dimitri Meunier, Mattes Mollenhauer</i>	

Divert More Attention to Vision-Language Tracking 4446
Heng Fan, Mingzhe Guo, Liping Jing, Zhipeng Zhang

Rethinking Image Restoration for Object Detection 4461
Xiaochun Cao, Wenqi Ren, Shangquan Sun, Tao Wang

VOLUME 7

Optimal Brain Compression: A Framework for Accurate Post-Training Quantization and Pruning 4475
Dan Alistarh, Elias Frantar

Challenging Common Assumptions in Convex Reinforcement Learning 4489
Piersilvio De Bartolomeis, Riccardo De Santi, Mirco Mutti, Marcello Restelli

Association Graph Learning for Multi-Task Classification with Category Shifts 4503
Jiayi Shen, Cees Snoek, Marcel Worring, Zehao Xiao, Xiantong Zhen

Delving into Sequential Patches for Deepfake Detection 4517
Errui Ding, Jiazhi Guan, Zhibin Hong, Chengbin Quan, Jingdong Wang, Youjian Zhao, Hang Zhou

On the Detrimental Effect of Invariances in the Likelihood for Variational Inference 4531
Jan Gasthaus, Ralf Herbrich, Tim Januschowski, Richard Kurle, Yuyang Bernie Wang

Self-Supervised Learning Through Efference Copies 4543
Qinghai Guo, Timoleon Moraitis, Franz Scherr

Precise Learning Curves and Higher-Order Scalings for Dot-Product Kernel Regression 4558
Hong Hu, Yue Lu, Theodor Misiakiewicz, Jeffrey Pennington, Lechao Xiao

A Unified Model for Multi-Class Anomaly Detection 4571
Lei Cui, Xinyi Le, Xin Lu, Yujun Shen, Kai Yang, Zhiyuan You, Yu Zheng

Constrained GPI for Zero-Shot Transfer in Reinforcement Learning 4585
Gunhee Kim, Jaekyeom Kim, Seohong Park

UnfoldML: Cost-Aware and Uncertainty-Based Dynamic 2D Prediction for Multi-Stage Classification 4598
Rishikesan Kamaleswaran, Alind Khare, Glenn Matlin, Monish Ramadoss, Alexey Tumanov, Yanbo Xu, Chao Zhang

Generative Status Estimation and Information Decoupling for Image Rain Removal 4612
Qing Guo, Ping Li, Di Lin, Ruonan Liu, Jia Shen, Miaohui Wang, Xin Wang, Wuyuan Xie, Renjie Zhang

Learning with Little Mixing 4626
Stephen Tu, Ingvar Ziemann

ComGAN: Unsupervised Disentanglement and Segmentation Via Image Composition 4638
Rui Ding, Kehua Guo, Liwei Wang, Zheng Wu, Xiangyuan Zhu

Sparse Winning Tickets Are Data-Efficient Image Recognizers 4652
Tianlong Chen, Xuxi Chen, Mukund Varma T, Subhashini Venugopalan, Zhangyang Wang, Zhenyu Zhang

VAEL: Bridging Variational Autoencoders and Probabilistic Logic Programming.....	4667
<i>Giuseppe Marra, Eleonora Misino, Emanuele Sansone</i>	
The Alignment Property of SGD Noise and How it Helps Select Flat Minima: A Stability Analysis.....	4680
<i>Weijie Su, Mingze Wang, Lei Wu</i>	
EvenNet: Ignoring Odd-Hop Neighbors Improves Robustness of Graph Neural Networks	4694
<i>Bolin Ding, Runlin Lei, Yaliang Li, Zhen Wang, Zhewei Wei</i>	
Differentiable Analog Quantum Computing for Optimization and Control	4707
<i>Jiaqi Leng, Ming Lin, Yuxiang Peng, Yi-Ling Qiao, Xiaodi Wu</i>	
Promising Or Elusive? Unsupervised Object Segmentation from Real-World Single Images.....	4722
<i>Bo Yang, Yafei Yang</i>	
Efficiently Factorizing Boolean Matrices Using Proximal Gradient Descent	4736
<i>Sebastian Dalleiger, Jilles Vreeken</i>	
Learning from Future: A Novel Self-Training Framework for Semantic Segmentation.....	4749
<i>Ye Du, Jingjing Fei, Zehua Fu, Wei Li, Qingjie Liu, Yujun Shen, Haochen Wang, Liwei Wu, Rui Zhao</i>	
A Simple and Provably Efficient Algorithm for Asynchronous Federated Contextual Linear Bandits	4762
<i>Quanquan Gu, Jiafan He, Yifei Min, Tianhao Wang</i>	
How Powerful Are K-Hop Message Passing Graph Neural Networks.....	4776
<i>Yixin Chen, Jiarui Feng, Fuhai Li, Anindya Sarkar, Muhan Zhang</i>	
Dynamic Tensor Product Regression.....	4791
<i>Aravind Reddy, Zhao Song, Lichen Zhang</i>	
Generalization Analysis of Message Passing Neural Networks on Large Random Graphs	4805
<i>Gitta Kutyniok, Yunseok Lee, Ron Levie, Sohir Maskey</i>	
On Infinite Separations Between Simple and Optimal Mechanisms	4818
<i>Alexandros Psomas, Ariel Schwartzman Cohenca, S. Weinberg</i>	
Online Learning and Pricing for Network Revenue Management with Reusable Resources.....	4830
<i>Huiwen Jia, Siqian Shen, Cong Shi</i>	
Learning to Find Proofs and Theorems by Learning to Refine Search Strategies: The Case of Loop Invariant Synthesis	4843
<i>Jonathan Laurent, André Platzer</i>	
Exploitability Minimization in Games and Beyond	4857
<i>Denizalp Goktas, Amy Greenwald</i>	
Where2comm: Communication-Efficient Collaborative Perception Via Spatial Confidence Maps	4874
<i>Siheng Chen, Shaoheng Fang, Yue Hu, Zixing Lei, Yiqi Zhong</i>	
Your Out-Of-Distribution Detection Method is Not Robust!	4887
<i>Mohammad Azizmalayeri, Mohammad Manzuri, Mohammad Hossein Rohban, Arshia Soltani Moakhar, Arman Zarei, Reihaneh Zohrabi</i>	
Discovering and Overcoming Limitations of Noise-Engineered Data-Free Knowledge Distillation.....	4902
<i>Deepak Mishra, Piyush Raikwar</i>	

NaturalProver: Grounded Mathematical Proof Generation with Language Models.....	4913
<i>Yejin Choi, Hannaneh Hajishirzi, Jiacheng Liu, Ximing Lu, Sean Welleck</i>	
Smoothed Online Convex Optimization Based on Discounted-Normal-Predictor.....	4928
<i>Wei Jiang, Tianbao Yang, Jinfeng Yi, Lijun Zhang</i>	
One for All: Simultaneous Metric and Preference Learning Over Multiple Users.....	4943
<i>Gregory Canal, Ramya Korlakai Vinayak, Blake Mason, Robert Nowak</i>	
Few-Shot Fast-Adaptive Anomaly Detection.....	4957
<i>Ser Nam Lim, Tsung-Yu Lin, Ashish Shah, Rui Wang, Ze Wang, Yipin Zhou</i>	
SegViT: Semantic Segmentation with Plain Vision Transformers.....	4971
<i>Xiangxiang Chu, Yifan Liu, Chunhua Shen, Quan Tang, Zhi Tian, Xiaolin Wei, Bowen Zhang</i>	
Unsupervised Learning from Incomplete Measurements for Inverse Problems.....	4983
<i>Dongdong Chen, Mike Davies, Julián Tachella</i>	
Redeeming Intrinsic Rewards Via Constrained Optimization	4996
<i>Pulkit Agrawal, Eric Chen, Zhang-Wei Hong, Joni Pajarinen</i>	
A Unified Evaluation of Textual Backdoor Learning: Frameworks and Benchmarks.....	5009
<i>Yangyi Chen, Ganqu Cui, Bingxiang He, Zhiyuan Liu, Maosong Sun, Lifan Yuan</i>	
MissDAG: Causal Discovery in the Presence of Missing Data with Continuous Additive Noise Models.....	5024
<i>Howard Bondell, Erdun Gao, Mingming Gong, Wei Huang, Tongliang Liu, Ignavier Ng, Li Shen, Kun Zhang</i>	
Escaping Saddle Points with Bias-Variance Reduced Local Perturbed SGD for Communication Efficient Nonconvex Distributed Learning.....	5039
<i>Tomoya Murata, Taiji Suzuki</i>	
AttCAT: Explaining Transformers Via Attentive Class Activation Tokens	5052
<i>Rhongo Jang, Chengyin Li, Xin Li, Deng Pan, Yao Qiang, Dongxiao Zhu</i>	
A Theoretical Study on Solving Continual Learning.....	5065
<i>Zixuan Ke, Gyuhak Kim, Tatsuya Konishi, Bing Liu, Changnan Xiao</i>	
BiMLP: Compact Binary Architectures for Vision Multi-Layer Perceptrons.....	5080
<i>Xinghao Chen, Yunhe Wang, Yixing Xu</i>	
Understanding Deep Neural Function Approximation in Reinforcement Learning Via ϵ -Greedy Exploration	5093
<i>Volkan Cevher, Fanghui Liu, Luca Viano</i>	
Misspecified Phase Retrieval with Generative Priors.....	5109
<i>Jiulong Liu, Zhaoqiang Liu, Xinshao Wang</i>	
Data-Efficient Augmentation for Training Neural Networks	5124
<i>Tian Yu Liu, Baharan Mirzasoleiman</i>	
Divide and Contrast: Source-Free Domain Adaptation Via Adaptive Contrastive Learning.....	5137
<i>Weikai Chen, Hui Cheng, Zhen Li, Guanbin Li, Siyuan Li, Liang Lin, Ziyi Zhang</i>	

Towards Human-Level Bimanual Dexterous Manipulation with Reinforcement Learning	5150
<i>Yuanpei Chen, Hao Dong, Xidong Feng, Jiechuan Jiang, Zongqing Lu, Stephen McAleer, Shengjie Wang, Tianhao Wu, Yaodong Yang, Song-Chun Zhu</i>	
Outlier-Robust Sparse Mean Estimation for Heavy-Tailed Distributions	5164
<i>Ilias Diakonikolas, Daniel Kane, Jasper Lee, Ankit Pensia</i>	
Local-Global MCMC Kernels: The Best of Both Worlds	5178
<i>Alain Durmus, Marylou Gabrié, Evgeny Lagutin, Eric Moulines, Alexey Naumov, Sergey Samsonov</i>	
The Computational and Learning Benefits of Daleian Neural Networks	5194
<i>Adam Haber, Elad Schneidman</i>	
CLIPDraw: Exploring Text-To-Drawing Synthesis Through Language-Image Encoders	5207
<i>Kevin Frans, Lisa Soros, Olaf Witkowski</i>	
Chaotic Dynamics Are Intrinsic to Neural Network Training with SGD	5219
<i>Maximilian Granz, Luis Herrmann, Tim Landgraf</i>	
Efficient and Modular Implicit Differentiation.....	5230
<i>Quentin Berthet, Mathieu Blondel, Marco Cuturi, Roy Frostig, Stephan Hoyer, Felipe Llinares-Lopez, Fabian Pedregosa, Jean-Philippe Vert</i>	
Effective Dimension in Bandit Problems Under Censorship.....	5243
<i>Saurabh Amin, Gauthier Guinet, Patrick Jaillet</i>	

VOLUME 8

Which Explanation Should I Choose? A Function Approximation Perspective to Characterizing Post Hoc Explanations.....	5256
<i>Tessa Han, Himabindu Lakkaraju, Suraj Srinivas</i>	
Accelerating Certified Robustness Training Via Knowledge Transfer	5269
<i>Kevin Eykholt, Amir Rahmati, Pratik Vaishnavi</i>	
MOMA-LRG: Language-Refined Graphs for Multi-Object Multi-Actor Activity Parsing	5282
<i>Ehsan Adeli, Zane Durante, Zhuoyi Huang, Emily Jin, Fei-Fei Li, Linden Li, Lun Yu Li, Ruo Chen Liu, Zelun Luo, Juan Carlos Niebles, Jiajun Wu, Wanze Xie</i>	
MTNeuro: A Benchmark for Evaluating Representations of Brain Structure Across Multiple Levels of Abstraction	5299
<i>Nauman Ahad, Mehdi Azabou, Eva Dyer, William Gray-Roncal, Joy Jackson, Matthew Jin, Erik Johnson, Christopher Liding, Ran Liu, Jorge Quesada, Lakshmi Sathidevi, Carolina Urzay, Jingyun Xiao</i>	
FLamby: Datasets and Benchmarks for Cross-Silo Federated Learning in Realistic Healthcare Settings	5315
<i>Shadi Albarqouni, Mathieu Andreux, Salman Avestimehr, Samy-Safwan Ayed, Aurélien Bellet, Edwige Cyffers, Aymeric Dieuleveut, Felix Grimberg, Chaoyang He, Martin Jaggi, Sai Praneeth Karimireddy, Regis Loeb, Marco Lorenzi, Paul Mangold, Tanguy Marchand, Othmane Marfoq, Erum Mushtaq, Boris Muzellec, Giovanni Neglia, Jean Ogier Du Terrail, Constantin Philippenko, Santiago Silva, Maria Telenczuk, Marc Tommasi</i>	
Foundation Posteriors for Approximate Probabilistic Inference	5335
<i>Noah Goodman, Mike Wu</i>	

Blackbox Attacks Via Surrogate Ensemble Search	5348
<i>Salman Asif, Zikui Cai, Srikanth Krishnamurthy, Amit Roy-Chowdhury, Chengyu Song</i>	
Transition to Linearity of General Neural Networks with Directed Acyclic Graph Architecture.....	5363
<i>Misha Belkin, Chaoyue Liu, Libin Zhu</i>	
A Comprehensive Study on Large-Scale Graph Training: Benchmarking and Rethinking.....	5376
<i>Tianlong Chen, Keyu Duan, Xia Hu, Zirui Liu, Peihao Wang, Zhangyang Wang, Wenqing Zheng, Kaixiong Zhou</i>	
Scale-Invariant Learning by Physics Inversion	5390
<i>Philipp Holl, Vladlen Koltun, Nils Thuerey</i>	
Regret Bounds for Multilabel Classification in Sparse Label Regimes.....	5404
<i>Róbert Busa-Fekete, Heejin Choi, Krzysztof Dembczynski, Claudio Gentile, Henry Reeve, Balazs Szorenyi</i>	
Sample Constrained Treatment Effect Estimation.....	5417
<i>Raghavendra Addanki, David Arbour, Tung Mai, Cameron Musco, Anup Rao</i>	
CATER: Intellectual Property Protection on Text Generation APIs Via Conditional Watermarks	5431
<i>Xuanli He, Ruoxi Jia, Jiwei Li, Lingjuan Lyu, Fangzhao Wu, Qiongfai Xu, Yi Zeng</i>	
On the Symmetries of the Synchronization Problem in Cryo-EM: Multi-Frequency Vector Diffusion Maps on the Projective Plane	5446
<i>Arash Behboodi, Gabriele Cesa, Taco Cohen, Max Welling</i>	
Sketching Based Representations for Robust Image Classification with Provable Guarantees	5459
<i>Nishanth Dikkala, Sankeerth Rao Karingula, Raghu Meka, Jelani Nelson, Rina Panigrahy, Xin Wang</i>	
ResQ: A Residual Q Function-Based Approach for Multi-Agent Reinforcement Learning Value Factorization.....	5471
<i>Yongquan Fu, Jun Liu, WeiQuan Liu, Xinwang Liu, Mengwei Qiu, Siqu Shen, Cheng Wang</i>	
Non-Identifiability and the Blessings of Misspecification in Models of Molecular Fitness	5484
<i>Alan Amin, Jonathan Frazer, Debora Marks, Eli Weinstein</i>	
Conditional Diffusion Process for Inverse Halftoning	5498
<i>Hao Jiang, Yadong Mu</i>	
Towards a Standardised Performance Evaluation Protocol for Cooperative MARL.....	5510
<i>Ruan John De Kock, Roland Dubb, Rihab Gorsane, Omayma Mahjoub, Arnu Pretorius, Siddarth Singh</i>	
Size and Depth of Monotone Neural Networks: Interpolation and Approximation	5522
<i>Dan Mikulincer, Daniel Reichman</i>	
Action-Modulated Midbrain Dopamine Activity Arises from Distributed Control Policies	5535
<i>Jack Lindsey, Ashok Litwin-Kumar</i>	
Latent Hierarchical Causal Structure Discovery with Rank Constraints.....	5549
<i>Clark Glymour, Biwei Huang, Charles Jia Han Low, Feng Xie, Kun Zhang</i>	
Beyond Real-World Benchmark Datasets: An Empirical Study of Node Classification with GNNs.....	5562
<i>Seiji Maekawa, Koki Noda, Makoto Onizuka, Yuya Sasaki</i>	

Doubly-Asynchronous Value Iteration: Making Value Iteration Asynchronous in Actions.....	5575
<i>Richard Sutton, Tian Tian, Kenny Young</i>	
Turbocharging Solution Concepts: Solving NEs, CEs and CCEs with Neural Equilibrium Solvers	5586
<i>Thomas Anthony, Ian Gemp, Siqi Liu, Luke Marris, Andrea Tacchetti, Karl Tuyls</i>	
NOMAD: Nonlinear Manifold Decoders for Operator Learning	5601
<i>Georgios Kissas, George J. Pappas, Paris Perdikaris, Jacob Seidman</i>	
Semi-Supervised Video Salient Object Detection Based on Uncertainty-Guided Pseudo Labels.....	5614
<i>Chenyang Lu, Huchuan Lu, Yongri Piao, Miao Zhang</i>	
Debiased Causal Tree: Heterogeneous Treatment Effects Estimation with Unmeasured Confounding.....	5628
<i>Qing Cui, Linbo Jiang, Longfei Li, Xinyu Li, Caizhi Tang, Huiyuan Wang, Ya-Lin Zhang, Jun Zhou, Feng Zhu</i>	
Asymptotic Properties for Bayesian Neural Network in Besov Space.....	5641
<i>Jaeyong Lee, Kyeongwon Lee</i>	
A PAC-Bayesian Generalization Bound for Equivariant Networks	5654
<i>Arash Behboodi, Gabriele Cesa, Taco Cohen</i>	
Neural Network Architecture Beyond Width and Depth.....	5669
<i>Zuwei Shen, Haizhao Yang, Shijun Zhang</i>	
S-Prompts Learning with Pre-Trained Transformers: An Occam’s Razor for Domain Incremental Learning	5682
<i>Xiaopeng Hong, Zhiwu Huang, Yabin Wang</i>	
OmniVL: One Foundation Model for Image-Language and Video-Language Tasks.....	5696
<i>Dongdong Chen, Yu-Gang Jiang, Ce Liu, Chong Luo, Junke Wang, Zuxuan Wu, Yujia Xie, Lu Yuan, Yucheng Zhao, Luowei Zhou</i>	
Self-Organized Group for Cooperative Multi-Agent Reinforcement Learning.....	5711
<i>Shuncheng He, Xiangyang Ji, Yuhang Jiang, Zhiqiang Lou, Jianzhun Shao, Hongchang Zhang</i>	
(Optimal) Online Bipartite Matching with Degree Information.....	5724
<i>Anders Aamand, Justin Chen, Piotr Indyk</i>	
Fairness in Federated Learning Via Core-Stability.....	5738
<i>Mintong Kang, Bo Li, Linyi Li, Ruta Mehta, Bhaskar Ray Chaudhury</i>	
Log-Polar Space Convolution Layers.....	5751
<i>Bing Su, Ji-Rong Wen</i>	
CoNSoLe: Convex Neural Symbolic Learning	5766
<i>Haoran Li, Hanghang Tong, Yang Weng</i>	
DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps.....	5775
<i>Fan Bao, Jianfei Chen, Chongxuan Li, Cheng Lu, Yuhao Zhou, Jun Zhu</i>	
Hierarchical Channel-Spatial Encoding for Communication-Efficient Collaborative Learning	5788
<i>Song Guo, Tao Guo, Xun Liu, Yi Liu, Zhihao Qu, Zhenda Xu, Jiewei Zhang, Jie Zhang, Qihua Zhou</i>	

Attracting and Dispersing: A Simple Approach for Source-Free Domain Adaptation	5802
<i>Shangling Jui, Joost Van De Weijer, Kai Wang, Yaxing Wang, Shiqi Yang</i>	
SCINet: Time Series Modeling and Forecasting with Sample Convolution and Interaction.....	5816
<i>Muxi Chen, Qiuxia Lai, Minhao Liu, Lingna Ma, Qiang Xu, Zhijian Xu, Ailing Zeng</i>	
Exploration-Guided Reward Shaping for Reinforcement Learning Under Sparse Rewards.....	5829
<i>Rati Devidze, Parameswaran Kamalaruban, Adish Singla</i>	
Active Exploration for Inverse Reinforcement Learning	5843
<i>Andreas Krause, David Lindner, Giorgia Ramponi</i>	
Subspace Recovery from Heterogeneous Data with Non-Isotropic Noise	5854
<i>John Duchi, Vitaly Feldman, Lunjia Hu, Kunal Talwar</i>	
Lifelong Neural Predictive Coding: Learning Cumulatively Online Without Forgetting.....	5867
<i>C Lee Giles, Daniel Kifer, Ankur Mali, Alex Ororbia</i>	
Align Then Fusion: Generalized Large-Scale Multi-View Clustering with Anchor Matching Correspondences	5882
<i>Jiaqi Jin, Suyuan Liu, Xinwang Liu, Wenxuan Tu, Siwei Wang, En Zhu, Xinzhong Zhu</i>	
Geodesic Graph Neural Network for Efficient Graph Representation Learning	5896
<i>Yixin Chen, Lecheng Kong, Muhan Zhang</i>	
Improved Differential Privacy for SGD Via Optimal Private Linear Operators on Adaptive Streams.....	5910
<i>Sergey Denisov, Abhradeep Guha Thakurta, H. Brendan McMahan, John Rush, Adam Smith</i>	
On Privacy and Personalization in Cross-Silo Federated Learning	5925
<i>Shengyuan Hu, Ken Liu, Virginia Smith, Steven Wu</i>	
SHAQ: Incorporating Shapley Value Theory into Multi-Agent Q-Learning.....	5941
<i>Yunjie Gu, Tae-Kyun Kim, Jianhong Wang, Yuan Zhang</i>	
Trajectory Balance: Improved Credit Assignment in GFlowNets	5955
<i>Yoshua Bengio, Emmanuel Bengio, Moksh Jain, Nikolay Malkin, Chen Sun</i>	
Instance-Dependent Near-Optimal Policy Identification in Linear MDPs Via Online Experiment Design.....	5968
<i>Kevin Jamieson, Andrew Wagenmaker</i>	
ProcTHOR: Large-Scale Embodied AI Using Procedural Generation	5982
<i>Matt Deitke, Kiana Ehsani, Winson Han, Alvaro Herrasti, Aniruddha Kembhavi, Eric Kolve, Roozbeh Mottaghi, Jordi Salvador, Eli Vanderbilt, Luca Weihs</i>	
Efficient Sampling on Riemannian Manifolds Via Langevin MCMC.....	5995
<i>Xiang Cheng, Suvrit Sra, Jingzhao Zhang</i>	
OpenFWI: Large-Scale Multi-Structural Benchmark Datasets for Full Waveform Inversion.....	6007
<i>Yinpeng Chen, Chengyuan Deng, Shihang Feng, Yinan Feng, Peng Jin, Youzuo Lin, Hanchen Wang, Qili Zeng, Xitong Zhang</i>	
High-Order Pooling for Graph Neural Networks with Tensor Decomposition	6021
<i>Chenqing Hua, Guillaume Rabusseau, Jian Tang</i>	
Natural Image Synthesis for the Retina with Variational Information Bottleneck Representation.....	6034
<i>Christophe Moser, Demetri Psaltis, Babak Rahmani</i>	

VOLUME 9

Imitating Past Successes Can Be Very Suboptimal	6047
<i>Benjamin Eysenbach, Sergey Levine, Russ Salakhutdinov, Soumith Udatha</i>	
A Communication-Efficient Algorithm with Linear Convergence for Federated Minimax Learning.....	6060
<i>Zhenyu Sun, Ermin Wei</i>	
Dynamic Graph Neural Networks Under Spatio-Temporal Distribution Shift.....	6074
<i>Haoyang Li, Zhou Qin, Xin Wang, Zeyang Zhang, Ziwei Zhang, Wenwu Zhu</i>	
Biologically Inspired Dynamic Thresholds for Spiking Neural Networks	6090
<i>Jianchuan Ding, Yufei Ding, Bo Dong, Felix Heide, Xin Yang, Baocai Yin, Yunduo Zhou</i>	
GraphQNTK: Quantum Neural Tangent Kernel for Graph Data.....	6104
<i>Yehui Tang, Junchi Yan</i>	
Trajectory-Guided Control Prediction for End-To-End Autonomous Driving: A Simple Yet Strong Baseline	6119
<i>Li Chen, Xiaosong Jia, Hongyang Li, Yu Qiao, Penghao Wu, Junchi Yan</i>	
360-MLC: Multi-View Layout Consistency for Self-Training and Hyper-Parameter Tuning.....	6133
<i>Yueh-Cheng Liu, Bolivar Solarte, Min Sun, Yi-Hsuan Tsai, Chin-Hsuan Wu</i>	
Distributed Optimization for Overparameterized Problems: Achieving Optimal Dimension Independent Communication Complexity.....	6147
<i>Mingyi Hong, Bingqing Song, Ioannis Tsaknakis, Hoi-To Wai, Chung-Yiu Yau</i>	
Falsification Before Extrapolation in Causal Effect Estimation.....	6161
<i>Zeshan M Hussain, Michael Oberst, Ming-Chieh Shih, David Sontag</i>	
Surprise Minimizing Multi-Agent Learning with Energy-Based Models	6175
<i>Yuri Lawryshyn, Konstantinos N Plataniotis, Xiao Qi Shi, Karush Suri</i>	
Geodesic Self-Attention for 3D Point Clouds	6190
<i>Mingsong Chen, Zhengyu Li, Xuan Tang, Xihao Wang, Xian Wei, Zihao Xu, Hui Yu</i>	
Test Time Adaptation Via Conjugate Pseudo-Labels.....	6204
<i>Sachin Goyal, J. Zico Kolter, Aditi Raghunathan, Mingjie Sun</i>	
SPD Domain-Specific Batch Normalization to Crack Interpretable Unsupervised Domain Adaptation in EEG	6219
<i>Jun-Ichiro Hirayama, Motoaki Kawanabe, Reinmar Kobler, Qibin Zhao</i>	
AVLEN: Audio-Visual-Language Embodied Navigation in 3D Environments.....	6236
<i>Anoop Cherian, Sudipta Paul, Amit Roy-Chowdhury</i>	
Semantic Uncertainty Intervals for Disentangled Latent Spaces.....	6250
<i>Anastasios Angelopoulos, Stephen Bates, Phillip Isola, Yaniv Romano, Swami Sankaranarayanan</i>	
A Fast Scale-Invariant Algorithm for Non-Negative Least Squares with Non-Negative Data.....	6264
<i>Jelena Diakonikolas, Chenghui Li, Swati Padmanabhan, Chaobing Song</i>	

Self-Supervised Amodal Video Object Segmentation	6278
<i>Yanwei Fu, Tong He, Yuxin Hong, Francesco Locatello, Chiyu Wang, David P Wipf, Tianjun Xiao, Jian Yao, Zheng Zhang</i>	
Beyond the Return: Off-Policy Function Estimation Under User-Specified Error-Measuring Distributions	6292
<i>Audrey Huang, Nan Jiang</i>	
Disentangling Transfer in Continual Reinforcement Learning	6304
<i>Lukasz Kucinski, Piotr Milos, Razvan Pascanu, Maciej Wolczyk, Michal Zajac</i>	
Meta-Learning with Self-Improving Momentum Target	6318
<i>Hankook Lee, Jaeho Lee, Jongjin Park, Jinwoo Shin, Jihoon Tack</i>	
Online Allocation and Learning in the Presence of Strategic Agents	6333
<i>Shipra Agrawal, Steven Yin, Assaf Zeevi</i>	
Learning Contrastive Embedding in Low-Dimensional Space.....	6345
<i>Shuo Chen, Chen Gong, Jun Li, Gang Niu, Masashi Sugiyama, Jian Yang</i>	
Near-Optimal Randomized Exploration for Tabular Markov Decision Processes	6358
<i>Qiwen Cui, Simon Du, Maryam Fazel, Ruoqi Shen, Zhihan Xiong</i>	
Defending Against Adversarial Attacks Via Neural Dynamic System	6372
<i>Xiyuan Li, Weiwei Liu, Zou Xin</i>	
On the Robustness of Graph Neural Diffusion to Topology Perturbations.....	6384
<i>Qiyu Kang, Yang Song, Wee Peng Tay, Sijie Wang, Kai Zhao</i>	
Few-Shot Relational Reasoning Via Connection Subgraph Pretraining.....	6397
<i>Qian Huang, Jure Leskovec, Hongyu Ren</i>	
Equivariant Networks for Zero-Shot Coordination	6410
<i>Jakob Foerster, Darius Muglich, Christian Schroeder De Witt, Elise Van Der Pol, Shimon Whiteson</i>	
Randomized Sketches for Clustering: Fast and Optimal Kernel k-Means	6424
<i>Yong Liu, Dan Meng, Weiping Wang, Rong Yin</i>	
Quantile Constrained Reinforcement Learning: A Reinforcement Learning Framework Constraining Outage Probability	6437
<i>Myungsik Cho, Whiyoung Jung, Jongeui Park, Youngchul Sung</i>	
Procedural Image Programs for Representation Learning.....	6450
<i>Manel Baradad, Richard Chen, Rogerio Feris, Phillip Isola, Antonio Torralba, Tongzhou Wang, Jonas Wulff</i>	
RényiCL: Contrastive Representation Learning with Skew Rényi Divergence	6463
<i>Kyungmin Lee, Jinwoo Shin</i>	
On the Efficient Implementation of High Accuracy Optimality of Profile Maximum Likelihood.....	6478
<i>Moses Charikar, Zhihao Jiang, Kirankumar Shiragur, Aaron Sidford</i>	
Paraphrasing is All You Need for Novel Object Captioning	6492
<i>Wan-Cyuan Fan, Louis-Philippe Morency, Russ Salakhutdinov, Yao-Hung Hubert Tsai, Frank Wang, Cheng-Fu Yang</i>	

Adversarial Robustness is at Odds with Lazy Training.....	6505
<i>Raman Arora, Poorya Mianjy, Enayat Ullah, Yunjuan Wang</i>	
Uncertainty Estimation for Multi-View Data: The Power of Seeing the Whole Picture	6517
<i>Joanna Dipnall, Lan Du, Belinda Gabbe, Myong Chol Jung, He Zhao</i>	
Motion Transformer with Global Intention Localization and Local Movement Refinement.....	6531
<i>Dengxin Dai, Li Jiang, Bernt Schiele, Shaoshuai Shi</i>	
No-Regret Learning in Games with Noisy Feedback: Faster Rates and Adaptivity Via Learning Rate Separation.....	6544
<i>Kimon Antonakopoulos, Volkan Cevher, Yu-Guan Hsieh, Panayotis Mertikopoulos</i>	
First-Order Algorithms for Min-Max Optimization in Geodesic Metric Spaces.....	6557
<i>Michael Jordan, Tianyi Lin, Emmanouil-Vasileios Vlatakis-Gkaragkounis</i>	
Feature-Proxy Transformer for Few-Shot Segmentation.....	6575
<i>Wei Chen, Yifan Sun, Yi Yang, Jian-Wei Zhang</i>	
Conformal Frequency Estimation with Sketched Data.....	6589
<i>Stefano Favaro, Matteo Sesia</i>	
Revisiting Active Sets for Gaussian Process Decoders	6603
<i>Cilie Feldager, Søren Hauberg, Pablo Moreno-Muñoz</i>	
Exact Learning Dynamics of Deep Linear Networks with Prior Knowledge.....	6615
<i>Lukas Braun, Clémentine Dominé, James Fitzgerald, Andrew Saxe</i>	
AMP: Automatically Finding Model Parallel Strategies with Heterogeneity Awareness.....	6630
<i>Dacheng Li, Hongyi Wang, Eric Xing, Hao Zhang</i>	
A Consistently Adaptive Trust-Region Method.....	6640
<i>Fadi Hamad, Oliver Hinder</i>	
Online Neural Sequence Detection with Hierarchical Dirichlet Point Process	6654
<i>Weihan Li, Gang Pan, Yu Qi</i>	
Leveraging Inter-Layer Dependency for Post -Training Quantization	6666
<i>Liang Li, Yuanliu Liu, Changbao Wang, Dandan Zheng</i>	
Emergence of Hierarchical Layers in a Single Sheet of Self-Organizing Spiking Neurons	6680
<i>Paul Bertens, Seong-Whan Lee</i>	
A Gradient Sampling Method with Complexity Guarantees for Lipschitz Functions in High and Low Dimensions.....	6692
<i>Damek Davis, Dmitriy Drusvyatskiy, Yin Tat Lee, Swati Padmanabhan, Guanghao Ye</i>	
CyCLIP: Cyclic Contrastive Language-Image Pretraining	6704
<i>Hritik Bansal, Sumit Bhatia, Shashank Goel, Aditya Grover, Ryan Rossi, Vishwa Vinay</i>	
When Does Dough Become a Bagel? Analyzing the Remaining Mistakes on ImageNet	6720
<i>Benjamin Caine, Sara Fridovich-Keil, Raphael Gontijo Lopes, Rebecca Roelofs, Vijay Vasudevan</i>	
Spatial Pruned Sparse Convolution for Efficient 3D Object Detection.....	6735
<i>Yukang Chen, Jianhui Liu, Xiaojuan Qi, Xiao Tan, Zhuotao Tian, Xiaoqing Ye</i>	

Sampling Without Replacement Leads to Faster Rates in Finite-Sum Minimax Optimization.....	6749
<i>Aniket Das, Michael Muehlebach, Bernhard Schölkopf</i>	
Feature Learning in L_2 -Regularized DNNs: Attraction/Repulsion and Sparsity	6763
<i>Franck Gabriel, Eugene Golikov, Clement Hongler, Arthur Jacot</i>	
Set-Based Meta-Interpolation for Few-Task Meta-Learning.....	6775
<i>Bruno Andreis, Sung Ju Hwang, Kenji Kawaguchi, Juho Lee, Seanie Lee</i>	
Non-Deep Networks.....	6789
<i>Alexey Bochkovskiy, Jia Deng, Ankit Goyal, Vladlen Koltun</i>	
Low-Rank Optimal Transport: Approximation, Statistics and Debiasing	6802
<i>Marco Cuturi, Meyer Scetbon</i>	
Embodied Scene-Aware Human Pose Estimation	6815
<i>Shun Iwase, Kris Kitani, Zhengyi Luo, Ye Yuan</i>	

VOLUME 10

Privacy Induces Robustness: Information-Computation Gaps and Sparse Mean Estimation.....	6829
<i>Kristian Georgiev, Samuel Hopkins</i>	
Batch Bayesian Optimization on Permutations Using the Acquisition Weighted Kernel.....	6843
<i>Roberto Bondesan, Efstratios Gavves, Changyong Oh, Max Welling</i>	
Supervised Training of Conditional Monge Maps.....	6859
<i>Charlotte Bunne, Marco Cuturi, Andreas Krause</i>	
Formalizing Consistency and Coherence of Representation Learning	6873
<i>Luke Dickens, Artur Garcez, Alessandra Russo, Harald Strömfelt</i>	
Positively Weighted Kernel Quadrature Via Subsampling	6886
<i>Satoshi Hayakawa, Terry Lyons, Harald Oberhauser</i>	
Guaranteed Conservation of Momentum for Learning Particle-Based Fluid Dynamics	6901
<i>Vladlen Koltun, Lukas Prantl, Nils Thuerey, Benjamin Ummenhofer</i>	
M4Singer: A Multi-Style, Multi-Singer and Musical Score Provided Mandarin Singing Corpus.....	6914
<i>Xiao Chen, Liqun Deng, Jinzheng He, Rongjie Huang, Ruiqi Li, Jinglin Liu, Yi Ren, Shoutong Wang, Lichao Zhang, Zhou Zhao, Jieming Zhu</i>	
Bayesian Spline Learning for Equation Discovery of Nonlinear Dynamics with Quantified Uncertainty	6927
<i>Daniel Huang, Hao Sun, Luning Sun, Jian-Xun Wang</i>	
Convexity Certificates from Hessians	6941
<i>Joachim Giesen, Julien Klaus, Sören Laue, Niklas Merk, Konstantin Wiedom</i>	
FR: Folded Rationalization with a Unified Encoder	6954
<i>Ruixuan Li, Wei Liu, Jun Wang, Haozhao Wang, Chao Yue, Yuankai Zhang</i>	
Max-Min Off-Policy Actor-Critic Method Focusing on Worst-Case Robustness to Model Misspecification	6967
<i>Youhei Akimoto, Kazuto Fukuchi, Jun Sakuma, Rei Sato, Takumi Tanabe</i>	

One-Inlier is First: Towards Efficient Position Encoding for Point Cloud Registration.....	6982
<i>Zhi Chen, Lin Guo, Wenbing Tao, Fan Yang</i>	
Adaptive Distribution Calibration for Few-Shot Learning with Hierarchical Optimal Transport.....	6996
<i>Dandan Guo, Long Tian, Hongyuan Zha, He Zhao, Mingyuan Zhou</i>	
InsNet: An Efficient, Flexible, and Performant Insertion-Based Text Generation Model	7011
<i>Sidi Lu, Tao Meng, Nanyun Peng</i>	
Weighted Distillation with Unlabeled Examples.....	7024
<i>Cenk Baykal, Fotis Iliopoulos, Vasilis Kontonis, Gaurav Menghani, Khoa Trinh, Erik Vee</i>	
When Does Group Invariant Learning Survive Spurious Correlations?.....	7038
<i>Yimeng Chen, Yanyan Lan, Zhi-Ming Ma, Ruibin Xiong</i>	
SNAKE: Shape-Aware Neural 3D Keypoint Field.....	7052
<i>Xiaoxue Chen, Chuang Gan, Wenbing Huang, Xiaodong Mu, Fuchun Sun, Peixing You, Hao Zhao, Chengliang Zhong, Guyue Zhou</i>	
Single Loop Gaussian Homotopy Method for Non-Convex Optimization	7065
<i>Shinji Ito, Hidenori Iwakiri, Akiko Takeda, Yuhang Wang</i>	
Minimax Optimal Online Imitation Learning Via Replay Estimation.....	7077
<i>J. Bagnell, Sanjiban Choudhury, Jiantao Jiao, Matt Peng, Nived Rajaraman, Kannan Ramchandran, Gokul Swamy, Steven Wu</i>	
Multi-Layer State Evolution Under Random Convolutional Design.....	7089
<i>Max Daniels, Cedric Gerbelot, Florent Krzakala, Lenka Zdeborová</i>	
Mixture-Of-Experts with Expert Choice Routing	7103
<i>Zhifeng Chen, Andrew Dai, Nan Du, Yanping Huang, James Laudon, Quoc V Le, Tao Lei, Hanxiao Liu, Vincent Zhao, Yanqi Zhou</i>	
GULP: A Prediction-Based Metric Between Representations.....	7115
<i>Enric Boix-Adsera, Hannah Lawrence, Philippe Rigollet, George Stepaniants</i>	
On the Convergence of Policy Gradient Methods to Nash Equilibria in General Stochastic Games	7128
<i>Angeliki Giannou, Kyriakos Lotidis, Panayotis Mertikopoulos, Emmanouil-Vasileios Vlatakis-Gkaragkounis</i>	
SAGDA: Achieving $O(\epsilon^{-2})$ Communication Complexity in Federated Min-Max Learning	7142
<i>Jia Liu, Zhuqing Liu, Haibo Yang, Xin Zhang</i>	
ALMA: Hierarchical Learning for Composite Multi-Agent Tasks.....	7155
<i>Robby Costales, Shariq Iqbal, Fei Sha</i>	
Improved Bounds on Neural Complexity for Representing Piecewise Linear Functions	7167
<i>Kuan-Lin Chen, Harinath Garudadri, Bhaskar D Rao</i>	
Assaying Out-Of-Distribution Generalization in Transfer Learning	7181
<i>Thomas Brox, Andrea Dittadi, Peter Gehler, Max Horn, David Kernert, Francesco Locatello, Chris Russell, Bernt Schiele, Bernhard Schölkopf, Carl-Johann Simon-Gabriel, Florian Wenzel, Dominik Zietlow</i>	
M2N: Mesh Movement Networks for PDE Solvers.....	7199
<i>Junqing Chen, Xiang Chen, Junpeng Gao, Matthew Piggott, Zuoqiang Shi, Wenbin Song, Fanglei Sun, Zheng Tian, Joseph G Wallwork, Jun Wang, Mingrui Zhang</i>	

Physics-Informed Implicit Representations of Equilibrium Network Flows.....	7211
<i>Francesco Bullo, Francesco Seccamonte, Kevin D. Smith, Ananthram Swami</i>	
Learning Interface Conditions in Domain Decomposition Solvers.....	7222
<i>Scott Maclachlan, Nicolas Nytko, Luke Olson, Ali Taghibakhshi, Matthew West, Tareq Uz Zaman</i>	
TANKBind: Trigonometry-Aware Neural Networks for Drug-Protein Binding Structure Prediction.....	7236
<i>Chengtao Li, Wei Lu, Jiahua Rao, Qifeng Wu, Jixian Zhang, Shuangjia Zheng</i>	
Hamiltonian Latent Operators for Content and Motion Disentanglement in Image Sequences.....	7250
<i>Asif Khan, Amos Storkey</i>	
Convolutional Neural Networks on Graphs with Chebyshev Approximation, Revisited.....	7264
<i>Mingguo He, Zhewei Wei, Ji-Rong Wen</i>	
A Kernelised Stein Statistic for Assessing Implicit Generative Models.....	7277
<i>Gesine D Reinert, Wenkai Xu</i>	
Fine-Grained Semantically Aligned Vision-Language Pre-Training.....	7290
<i>Xin He, Juncheng Li, Long Qian, Siliang Tang, Qi Tian, Longhui Wei, Lingxi Xie, Linchao Zhu, Yueting Zhuang</i>	
Fast Distance Oracles for Any Symmetric Norm.....	7304
<i>Yichuan Deng, Zhao Song, Omri Weinstein, Ruizhe Zhang</i>	
Outlier-Robust Sparse Estimation Via Non-Convex Optimization.....	7318
<i>Yu Cheng, Ilias Diakonikolas, Rong Ge, Shivam Gupta, Daniel Kane, Mahdi Soltanolkotabi</i>	
Recurrent Convolutional Neural Networks Learn Succinct Learning Algorithms.....	7328
<i>Surbhi Goel, Sham Kakade, Adam Kalai, Cyril Zhang</i>	
Navigating Memory Construction by Global Pseudo-Task Simulation for Continual Learning.....	7342
<i>Yejia Liu, Shaolei Ren, Wang Zhu</i>	
Contact-Aware Human Motion Forecasting.....	7356
<i>Richard I Hartley, Miaomiao Liu, Wei Mao, Mathieu Salzmann</i>	
Spectral Bias in Practice: The Role of Function Frequency in Generalization.....	7368
<i>Sara Fridovich-Keil, Raphael Gontijo Lopes, Rebecca Roelofs</i>	
Out-Of-Distribution Detection with an Adaptive Likelihood Ratio on Informative Hierarchical VAE.....	7383
<i>Bo An, Yewen Li, Tongliang Liu, Xin Miao, Chaojie Wang, Xiaobo Xia</i>	
FOF: Learning Fourier Occupancy Field for Monocular Real-Time Human Reconstruction.....	7397
<i>Qiao Feng, Yu-Kun Lai, Kun Li, Yebin Liu, Jingyu Yang</i>	
Rate-Optimal Online Convex Optimization in Adaptive Linear Control.....	7410
<i>Asaf Benjamin Cassel, Tomer Koren, Alon Peled-Cohen</i>	
RTFormer: Efficient Design for Real-Time Semantic Segmentation with Transformer.....	7423
<i>Errui Ding, Haocheng Feng, Chenhui Gou, Junyu Han, Jian Wang, Jingdong Wang, Qiman Wu</i>	
Staggered Rollout Designs Enable Causal Inference Under Interference Without Network Knowledge.....	7437
<i>Mayleen Cortez, Matthew Eichhorn, Christina Yu</i>	

Censored Quantile Regression Neural Networks for Distribution-Free Survival Analysis	7450
<i>Jong-Hyeon Jeong, Yichen Jia, Tim Pearce, Jun Zhu</i>	
Learning Graph-Embedded Key-Event Back-Tracing for Object Tracking in Event Clouds.....	7462
<i>Junhui Hou, Xianqiang Lyu, Zhiyu Zhu</i>	
Efficient and Near-Optimal Smoothed Online Learning for Generalized Linear Functions	7477
<i>Adam Block, Max Simchowitz</i>	
Systematic Improvement of Neural Network Quantum States Using Lanczos	7490
<i>Hongwei Chen, Adrian Feiguin, Douglas Hendry, Phillip Weinberg</i>	
Grounded Reinforcement Learning: Learning to Win the Game Under Human Commands.....	7504
<i>Huaijie Wang, Yi Wu, Shusheng Xu</i>	
Enhance the Visual Representation Via Discrete Adversarial Training	7520
<i>Yuefeng Chen, Ranjie Duan, Xiaodan Li, Xiaofeng Mao, Gege Qi, Hui Xue', Shaokai Ye, Rong Zhang, Yao Zhu</i>	
Trading off Resource Budgets for Improved Regret Bounds	7534
<i>Damon Falck, Thomas Orton</i>	
Natural Color Fool: Towards Boosting Black-Box Unrestricted Attacks	7546
<i>Yaya Cheng, Lianli Gao, Jingkuan Song, Shengming Yuan, Qilong Zhang</i>	
Old Can Be Gold: Better Gradient Flow Can Make Vanilla-GCNs Great Again	7561
<i>Tianlong Chen, Ying Ding, Ajay Jaiswal, Justin Rousseau, Peihao Wang, Zhangyang Wang</i>	
Egocentric Video-Language Pretraining.....	7575
<i>Chengfei Cai, Dima Damen, Difei Gao, Bernard Ghanem, Wang Hongfa, Weijie Kong, Kevin Qinghong Lin, Wei Liu, Mike Zheng Shou, Mattia Soldan, Rong-Cheng Tu, Jinpeng Wang, Michael Wray, Eric Z. Xu, Rui Yan, Wenzhe Zhao</i>	
Society of Agents: Regret Bounds of Concurrent Thompson Sampling.....	7587
<i>Qinxun Bai, Yan Chen, Maria Dimakopoulou, Perry Dong, Wei Xu, Zhengyuan Zhou</i>	
A Non-Asymptotic Analysis of Non-Parametric Temporal-Difference Learning.....	7599
<i>Francis Bach, Eloïse Berthier, Ziad Kobeissi</i>	

VOLUME 11

CEIP: Combining Explicit and Implicit Priors for Reinforcement Learning with Demonstrations	7614
<i>Alex Schwing, Yu-Xiong Wang, Kai Yan</i>	
Memorization and Optimization in Deep Neural Networks with Minimum Over-Parameterization	7628
<i>Mohammad Hossein Amani, Simone Bombari, Marco Mondelli</i>	
LAMP: Extracting Text from Gradients with Language Model Priors.....	7641
<i>Mislav Balunovic, Dimitar Dimitrov, Nikola Jovanovic, Martin Vechev</i>	
Explainability Via Causal Self-Talk	7655
<i>Junhyung Kim, Neil Rabinowitz, Nicholas Roy</i>	
Continual Learning with Evolving Class Ontologies	7671
<i>Shu Kong, Zhiqiu Lin, Deepak Pathak, Deva Ramanan, Yu-Xiong Wang</i>	

A Gradient Estimator Via L1-Randomization for Online Zero-Order Optimization with Two Point Feedback.....	7685
<i>Arya Akhavan, Evgenii Chzhen, Massimiliano Pontil, Alexandre Tsybakov</i>	
On the SDEs and Scaling Rules for Adaptive Gradient Algorithms.....	7697
<i>Sanjeev Arora, Kaifeng Lyu, Sadhika Malladi, Abhishek Panigrahi</i>	
Improving Neural Ordinary Differential Equations with Nesterov's Accelerated Gradient Method.....	7712
<i>Ho Huu Nghia Nguyen, Tan Nguyen, Stanley Osher, Huyen Vo, Thieu Vo</i>	
VER: Scaling On-Policy RL Leads to the Emergence of Navigation in Embodied Rearrangement.....	7727
<i>Dhruv Batra, Irfan Essa, Erik Wijmans</i>	
Neural Attentive Circuits.....	7741
<i>Nicolas Ballas, Yoshua Bengio, Erran Li Li, Francesco Locatello, Chris Pal, Nasim Rahaman, Bernhard Schölkopf, Martin Weiss</i>	
CLEVRER-Humans: Describing Physical and Causal Events the Human Way.....	7755
<i>Noah Goodman, Jiayuan Mao, Jiajun Wu, Xuelin Yang, Xikun Zhang</i>	
3D Concept Grounding on Neural Fields.....	7769
<i>Yilun Du, Chuang Gan, Yining Hong, Chunru Lin, Josh Tenenbaum</i>	
Evaluating Graph Generative Models with Contrastively Learned Features.....	7783
<i>Kaveh Hassani, Hamed Shirzad, Danica J. Sutherland</i>	
Path Independent Equilibrium Models Can Better Exploit Test-Time Computation.....	7796
<i>Cem Anil, Shaojie Bai, Roger Grosse, J. Zico Kolter, Kaiqu Liang, Ashwini Pokle, Johannes Treutlein, Yuhuai Wu</i>	
Implications of Model Indeterminacy for Explanations of Automated Decisions.....	7810
<i>Ashton Anderson, Marc-Etienne Brunet, Richard Zemel</i>	
Debugging and Explaining Metric Learning Approaches: An Influence Function Based Perspective.....	7824
<i>Jin Song Dong, Yun Lin, Ruofan Liu, Xianglin Yang</i>	
Meta-Reinforcement Learning with Self-Modifying Networks.....	7838
<i>Mathieu Chalvidal, Thomas Serre, Rufin Vanrullen</i>	
FairVFL: A Fair Vertical Federated Learning Framework with Contrastive Adversarial Learning.....	7852
<i>Yongfeng Huang, Hao Liao, Lingjuan Lyu, Tao Qi, Chuhan Wu, Fangzhao Wu, Xing Xie, Tong Xu, Zhongliang Yang</i>	
Incorporating Bias-Aware Margins into Contrastive Loss for Collaborative Filtering.....	7866
<i>Tat-Seng Chua, Wenchang Ma, Xiang Wang, An Zhang</i>	
Multiview Human Body Reconstruction from Uncalibrated Cameras.....	7879
<i>Cem Keskin, Hyun Soo Park, Chengcheng Tang, Luan Tran, Yuanlu Xu, Zhixuan Yu, Linguang Zhang</i>	
MACK: Multimodal Aligned Conceptual Knowledge for Unpaired Image-Text Matching.....	7892
<i>Yan Huang, Liang Wang, Yuming Wang, Yunan Zeng</i>	
Learning Optical Flow from Continuous Spike Streams.....	7905
<i>Xiaopeng Fan, Tiejun Huang, Ruiqin Xiong, Zhaofei Yu, Jing Zhao, Rui Zhao</i>	

Improving 3D-Aware Image Synthesis with a Geometry-Aware Discriminator	7921
<i>Qifeng Chen, Yujun Shen, Zifan Shi, Yinghao Xu, Dit-Yan Yeung, Deli Zhao</i>	
A Consistent and Differentiable L_p Canonical Calibration Error Estimator	7933
<i>Matthew Blaschko, Teodora Popordanoska, Raphael Sayer</i>	
Transform Once: Efficient Operator Learning in Frequency Domain	7947
<i>Federico Berto, Tri Dao, Stefano Ermon, Stefano Massaroli, Jinkyoo Park, Michael Poli, Christopher Ré</i>	
Fused Orthogonal Alternating Least Squares for Tensor Clustering	7960
<i>Dan Nicolae, Jiacheng Wang</i>	
A Solver-Free Framework for Scalable Learning in Neural ILP Architectures	7972
<i>Mausam Singla, Yatin Nandwani, Rishabh Ranjan, Parag Singla</i>	
Distributed Distributionally Robust Optimization with Non-Convex Objectives	7987
<i>Yang Jiao, Dongjin Song, Kai Yang</i>	
Parameter-Free Regret in High Probability with Heavy Tails	8000
<i>Ashok Cutkosky, JiuJia Zhang</i>	
Non-Convex Bilevel Games with Critical Point Selection Maps	8013
<i>Michael Arbel, Julien Mairal</i>	
Consistent Sufficient Explanations and Minimal Local Rules for Explaining the Decision of Any Classifier Or Regressor.....	8027
<i>Nicolas Brunel, Salim I. Amoukou</i>	
High-Dimensional Additive Gaussian Processes Under Monotonicity Constraints	8041
<i>Francois Bachoc, Andrés López-Lopera, Olivier Roustant</i>	
Spherical Channels for Modeling Atomic Interactions.....	8054
<i>Abhishek Das, Adeesh Kolluru, Janice Lan, Muhammed Shuaibi, Anuroop Sriram, Zachary Ulissi, Brandon Wood, Larry Zitnick</i>	
Handcrafted Backdoors in Deep Neural Networks.....	8068
<i>Nicholas Carlini, Sanghyun Hong, Alexey Kurakin</i>	
Touch and Go: Learning from Human-Collected Vision and Touch	8081
<i>Chenyang Ma, Andrew Owens, Fengyu Yang, Wenzhen Yuan, Jiacheng Zhang, Jing Zhu</i>	
SoLar: Sinkhorn Label Refinery for Imbalanced Partial-Label Learning	8104
<i>Gang Chen, Lei Feng, Yixuan Li, Yuren Mao, Haobo Wang, Mingxuan Xia, Junbo Zhao</i>	
Log-Linear-Time Gaussian Processes Using Binary Tree Kernels.....	8118
<i>Michael K. Cohen, Samuel Daulton, Michael A Osborne</i>	
Recovering Private Text in Federated Learning of Language Models	8130
<i>Danqi Chen, Tianyu Gao, Samyak Gupta, Yangsibo Huang, Kai Li, Zexuan Zhong</i>	
INRAS: Implicit Neural Representation for Audio Scenes	8144
<i>Mingfei Chen, Eli Shlizerman, Kun Su</i>	
Non-Monotonic Latent Alignments for CTC-Based Non-Autoregressive Machine Translation.....	8159
<i>Yang Feng, Chenze Shao</i>	

Attention-Based Neural Cellular Automata.....	8174
<i>Derek Nowrouzezahrai, Chris Pal, Mattie Tesfaldet</i>	
Learning Deep Input-Output Stable Dynamics	8187
<i>Ryosuke Kojima, Yuji Okamoto</i>	
Grounded Video Situation Recognition.....	8199
<i>C. V. Jawahar, Zeeshan Khan, Makarand Tapaswi</i>	
Policy Optimization with Advantage Regularization for Long-Term Fairness in Decision Systems	8211
<i>Sicun Gao, Min Kyung Lee, Zhizhen Qin, Eric Yu</i>	
Gradient Descent: The Ultimate Optimizer	8214
<i>Kartik Chandra, Erik Meijer, Jonathan Ragan-Kelley, Audrey Xie</i>	
DAGMA: Learning DAGs Via M-Matrices and a Log-Determinant Acyclicity Characterization.....	8226
<i>Bryon Aragam, Kevin Bello, Pradeep Ravikumar</i>	
Forward-Backward Latent State Inference for Hidden Continuous-Time semi-Markov Chains	8240
<i>Nicolai Engelmann, Heinz Koepl</i>	
LobsDICE: Offline Learning from Observation Via Stationary Distribution Correction Estimation.....	8252
<i>Youngsoo Jang, Geon-Hyeong Kim, Kee-Eung Kim, Jongmin Lee, Hongseok Yang</i>	
Ensemble of Averages: Improving Model Selection and Boosting Performance in Domain Generalization	8265
<i>Devansh Arpit, Huan Wang, Caiming Xiong, Yingbo Zhou</i>	
Is L^2 Physics Informed Loss Always Suitable for Training Physics Informed Neural Network?.....	8278
<i>Di He, Shanda Li, Liwei Wang, Chuwei Wang</i>	
Vision GNN: An Image is Worth Graph of Nodes.....	8291
<i>Jianyuan Guo, Kai Han, Yehui Tang, Yunhe Wang, Enhua Wu</i>	
ELIGN: Expectation Alignment as a Multi-Agent Intrinsic Reward.....	8304
<i>Michael Bernstein, Ranjay Krishna, Fei-Fei Li, Zixian Ma, Rose Wang</i>	
Factored DRO: Factored Distributionally Robust Policies for Contextual Bandits.....	8318
<i>Emma Brunskill, Yash Chandak, Tatsunori Hashimoto, Tong Mu</i>	
Could Giant Pre-Trained Image Models Extract Universal Representations?.....	8332
<i>Yue Cao, Han Hu, Stephen Lin, Yutong Lin, Ze Liu, Zheng Zhang, Nanning Zheng</i>	
A Scalable Deterministic Global Optimization Algorithm for Training Optimal Decision Tree.....	8347
<i>Yankai Cao, Kaixun Hua, Jiayang Ren</i>	
Thor: Wielding Hammers to Integrate Language Models and Automated Theorem Provers.....	8360
<i>Konrad Czechowski, Mateja Jamnik, Albert Qiaochu Jiang, Wenda Li, Piotr Milos, Tomasz Odrzygóźdź, Szymon Tworkowski, Yuhuai Wu</i>	
Rethinking Knowledge Graph Evaluation Under the Open-World Assumption	8374
<i>Zhouchen Lin, Haotong Yang, Muhan Zhang</i>	
KERPLE: Kernelized Relative Positional Embedding for Length Extrapolation	8386
<i>Ta-Chung Chi, Ting-Han Fan, Peter J Ramadge, Alexander Rudnicky</i>	

VOLUME 12

Zonotope Domains for Lagrangian Neural Network Verification	8400
<i>Alex Dimakis, Jonathan Hayase, Matt Jordan, Sewoong Oh</i>	
Neural Basis Models for Interpretability	8414
<i>Abhimanyu Dubey, Dhruv Mahajan, Filip Radenovic</i>	
RecursiveMix: Mixed Learning with History.....	8427
<i>Xiang Li, Renjie Song, Jian Yang, Lingfeng Yang, Borui Zhao</i>	
GAR: Generalized Autoregression for Multi-Fidelity Fusion	8441
<i>Yuxin Wang, Zheng Xing, Wei Xing</i>	
Anonymized Histograms in Intermediate Privacy Models.....	8456
<i>Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi</i>	
Truly Deterministic Policy Optimization	8469
<i>Tim Bretl, Saba Ghaffari, Ehsan Saleh, Matthew West</i>	
Language Models with Image Descriptors Are Strong Few-Shot Video-Language Learners	8483
<i>Mohit Bansal, Shih-Fu Chang, Derek Hoiem, Heng Ji, Jie Lei, Manling Li, Xudong Lin, Shuohang Wang, Zhenhailong Wang, Ruochen Xu, Ziyi Yang, Luowei Zhou, Chenguang Zhu</i>	
3DB: A Framework for Debugging Computer Vision Models	8498
<i>Logan Engstrom, Andrew Ilyas, Ashish Kapoor, Guillaume Leclerc, Aleksander Madry, Hadi Salman, Shibani Santurkar, Sai Vemprala, Vibhav Vineet, Kai Xiao, Greg Yang, Pengchuan Zhang</i>	
Empirical Gateaux Derivatives for Causal Inference	8512
<i>Michael Jordan, Yixin Wang, Angela Zhou</i>	
Calibrated Data-Dependent Constraints with Exact Satisfaction Guarantees.....	8526
<i>Yuekai Sun, Songkai Xue, Mikhail Yurochkin</i>	
Towards Consistency in Adversarial Classification.....	8538
<i>Jamal Atif, Yann Chevaleryre, Raphael Ettetdgui, Laurent Meunier, Rafael Pinot</i>	
Residual Multiplicative Filter Networks for Multiscale Reconstruction	8550
<i>Marcus Brubaker, David Fleet, David Lindell, Shayan Shekarforoush</i>	
WT-MVSNet: Window-Based Transformers for Multi-View Stereo	8564
<i>Yikang Ding, Wensen Feng, Jia Guo, Dihe Huang, Jinli Liao, Shihao Ren, Yoli Shavit, Kai Zhang</i>	
Private Isotonic Regression	8577
<i>Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi</i>	
A Closer Look at Offline RL Agents	8591
<i>Benoit Boulet, Yuwei Fu, Di Wu</i>	
Generalization Bounds for Estimating Causal Effects of Continuous Treatments	8605
<i>Huanhuan Chen, Shengfei Lyu, Xin Wang, Tianhao Wu, Xingyu Wu</i>	
Better Uncertainty Calibration Via Proper Scores for Classification and Beyond	8618
<i>Florian Buettner, Sebastian Gruber</i>	

Video Diffusion Models	8633
<i>William Chan, David Fleet, Alexey Gritsenko, Jonathan Ho, Mohammad Norouzi, Tim Salimans</i>	
Formulating Robustness Against Unforeseen Attacks.....	8647
<i>Sihui Dai, Saeed Mahloujifar, Prateek Mittal</i>	
Single Model Uncertainty Estimation Via Stochastic Data Centering.....	8662
<i>Rushil Anirudh, Timo Bremer, Vivek Sivaraman Narayanaswamy, Jayaraman Thiagarajan</i>	
Autoinverse: Uncertainty Aware Inversion of Neural Networks	8675
<i>Navid Ansari, Vahid Babaei, Hans-Peter Seidel, Nima Vahidi Ferdowsi</i>	
FedPop: A Bayesian Approach for Personalised Federated Learning.....	8687
<i>Alain Durmus, Nikita Kotelevskii, Eric Moulines, Maxime Vono</i>	
Semantic Diffusion Network for Semantic Segmentation.....	8702
<i>Jimin Pi, Haoru Tan, Sitong Wu</i>	
Iterative Structural Inference of Directed Graphs.....	8717
<i>Jun Pang, Aoran Wang</i>	
An Efficient Graph Generative Model for Navigating Ultra-Large Combinatorial Synthesis Libraries	8731
<i>Brandon Anderson, Chaoyi Chang, Pawel Gniewek, Aryan Pedawi, Henry Van Den Bedem</i>	
Learning to Discover and Detect Objects	8746
<i>Ismail Elezi, Vladimir Fomenko, Laura Leal-Taixé, Aljosa Osep, Deva Ramanan</i>	
Simulation-Guided Beam Search for Neural Combinatorial Optimization	8760
<i>Jinho Choo, Youngjune Gwon, André Hottung, Jeongwoo Jae, Jihoon Kim, Yeong-Dae Kwon, Kevin Tierney</i>	
FlowHMM: Flow-Based Continuous Hidden Markov Models	8773
<i>Pawel Lorek, Rafal Nowak, Tomasz Trzcinski, Maciej Zieba</i>	
Near Instance-Optimal PAC Reinforcement Learning for Deterministic MDPs.....	8785
<i>Aymen Al Marjani, Emilie Kaufmann, Andrea Tirinzoni</i>	
VICRegL: Self-Supervised Learning of Local Visual Features.....	8799
<i>Adrien Bardes, Yann Lecun, Jean Ponce</i>	
Alleviating Adversarial Attacks on Variational Autoencoders with MCMC	8811
<i>Anna Kuzina, Jakub Tomczak, Max Welling</i>	
Policy Gradient with Serial Markov Chain Reasoning.....	8824
<i>Oya Celiktutan, Edoardo Cetin</i>	
DTG-SSOD: Dense Teacher Guidance for Semi-Supervised Object Detection.....	8840
<i>Gang Li, Xiang Li, Ding Liang, Yujie Wang, Wu Yichao, Shanshan Zhang</i>	
Human-AI Shared Control Via Policy Dissection	8853
<i>Lan Feng, Quanyi Li, Zhenghao Peng, Haibin Wu, Bolei Zhou</i>	
Improving Generative Adversarial Networks Via Adversarial Learning in Latent Space	8868
<i>Yang Li, Yichuan Mo, Liangliang Shi, Junchi Yan</i>	

ShapeCrafter: A Recursive Text-Conditioned 3D Shape Generation Model	8882
<i>Yiwen Chen, Rao Fu, Daniel Ritchie, Srinath Sridhar, Xiao Zhan</i>	
SoundSpaces 2.0: A Simulation Platform for Visual-Acoustic Learning	8896
<i>Dhruv Batra, Paul Calamia, Changan Chen, Alexander Clegg, Sanchit Garg, Kristen Grauman, Philip Kobernik, Philip Robinson, Carl Schissler</i>	
AutoWS-Bench-101: Benchmarking Automated Weak Supervision with 100 Labels	8912
<i>Dyah Adila, Aws Albarghouthi, Tzu-Heng Huang, Xintong Li, Cheng-Yu Liu, Haotian Ma, Lauren Pick, Nicholas Roberts, Frederic Sala, Spencer Schoenberg</i>	
Towards Disentangling Information Paths with Coded ResNeXt	8926
<i>Apostolos Avranas, Marios Kountouris</i>	
Time-Conditioned Dances with Simplicial Complexes: Zigzag Filtration Curve Based Supra-Hodge Convolution Networks for Time-Series Forecasting	8940
<i>Yuzhou Chen, Yulia Gel, H. Vincent Poor</i>	
Are Defenses for Graph Neural Networks Robust?.....	8954
<i>Aleksandar Bojchevski, Simon Geisler, Stephan Günnemann, Felix Mujkanovic</i>	
GraB: Finding Provably Better Data Permutations than Random Reshuffling	8969
<i>Christopher De Sa, Wentao Guo, Yucheng Lu</i>	
Amortized Proximal Optimization	8982
<i>Juhan Bae, Roger Grosse, Jeff Z. Haochen, Paul Vicol</i>	
A Statistical Online Inference Approach in Averaged Stochastic Approximation	8998
<i>Chuhan Xie, Zhihua Zhang</i>	
SAViT: Structure-Aware Vision Transformer Pruning Via Collaborative Optimization.....	9010
<i>Zheyang Li, Shiliang Pu, Ye Ren, Wenming Tan, Jun Xiao, Zhi Yang, Kai Zhang, Chuanyang Zheng</i>	
Deciding What to Model: Value-Equivalent Sampling for Reinforcement Learning	9024
<i>Dilip Arumugam, Benjamin Van Roy</i>	
Giga-Scale Kernel Matrix-Vector Multiplication on GPU	9045
<i>Siu Lun Chau, Joan Glaunès, Robert Hu, Dino Sejdinovic</i>	
MultiScan: Scalable RGBD Scanning for 3D Environments with Articulated Objects	9058
<i>Angel Chang, Hanxiao Jiang, Yongsen Mao, Manolis Savva, Yiming Zhang</i>	
Are GANs Overkill for NLP?.....	9072
<i>David Alvarez-Melis, Vikas Garg, Adam Kalai</i>	
“Why Not Other Classes?”: Towards Class-Contrastive Back-Propagation Explanations.....	9085
<i>Xiaoqian Wang, Yipei Wang</i>	
Neural Stochastic Control.....	9098
<i>Wei Lin, Jingdong Zhang, Qunxi Zhu</i>	
Constrained Update Projection Approach to Safe Policy Optimization	9111
<i>Juntao Dai, Jiaming Ji, Pengfei Li, Gang Pan, Long Yang, Yaodong Yang, Linrui Zhang, Binbin Zhou</i>	

DetCLIP: Dictionary-Enriched Visual-Concept Paralleled Pre-Training for Open-World Detection	9125
<i>Jianhua Han, Zhenguo Li, Xiaodan Liang, Youpeng Wen, Dan Xu, Hang Xu, Chunjing Xu, Lewei Yao, Wei Zhang</i>	
The Sample Complexity of One-Hidden-Layer Neural Networks	9139
<i>Ohad Shamir, Nati Srebro, Gal Vardi</i>	
Learning from Few Samples: Transformation-Invariant SVMs with Composition and Locality at Multiple Scales	9151
<i>P. R. Kumar, Tao Liu, Xi Liu, Ruida Zhou</i>	
Learning Efficient Vision Transformers Via Fine-Grained Manifold Distillation	9164
<i>Jianyuan Guo, Kai Han, Zhiwei Hao, Han Hu, Ding Jia, Yehui Tang, Yunhe Wang, Chao Zhang</i>	
Equivariant Graph Hierarchy-Based Neural Networks	9176
<i>Jiaqi Han, Wenbing Huang, Yu Rong, Tingyang Xu</i>	
VOLUME 13	
Learning Interacting Dynamical Systems with Latent Gaussian Process ODEs	9188
<i>Melih Kandemir, Barbara Rakitsch, Çagatay Yildiz</i>	
OLIVES Dataset: Ophthalmic Labels for Investigating Visual Eye Semantics.....	9201
<i>Ghassan Alregib, Kiran Kokilepersaud, Yash-Yee Logan, Mohit Prabhushankar, Stephanie Trejo Corona, Charles Wykoff</i>	
Off-Policy Evaluation for Action-Dependent Non-Stationary Environments	9217
<i>Nathaniel Bastian, Emma Brunskill, Yash Chandak, Bruno Da Silva, Shiv Shankar, Philip Thomas</i>	
Fast Mixing of Stochastic Gradient Descent with Normalization and Weight Decay	9233
<i>Zhiyuan Li, Tianhao Wang, Dingli Yu</i>	
OPEN: Orthogonal Propagation with Ego-Network Modeling	9249
<i>Xiaochun Cao, Yuanfang Guo, Dongxiao He, Lina Kang, Mengzhe Li, Bingxin Niu, Chuan Wang, Zhen Wang, Liang Yang, Qiuliang Zhang</i>	
Beyond the Best: Distribution Functional Estimation in Infinite-Armed Bandits	9262
<i>Tavor Baharav, Yanjun Han, Jiantao Jiao, David Tse, Yifei Wang</i>	
Adversarial Training for High-Stakes Reliability	9274
<i>Tim Bauman, Lawrence Chan, Daniel De Haas, Tao Lin, Noa Nabeshima, Seraphina Nix, Adam Scherlis, Peter Schmidt-Nielsen, Buck Shlegeris, Nate Thomas, Benjamin Weinstein-Raun, Daniel Ziegler</i>	
ELEVATER: A Benchmark and Toolkit for Evaluating Language-Augmented Visual Models	9287
<i>Jyoti Aneja, Jianfeng Gao, Houdong Hu, Ping Jin, Yong Jae Lee, Chunyuan Li, Liunian Li, Haotian Liu, Zicheng Liu, Jianwei Yang, Pengchuan Zhang</i>	
Augmented RBMLE-UCB Approach for Adaptive Control of Linear Quadratic Systems	9302
<i>P. R. Kumar, Akshay Mete, Rahul Singh</i>	
Stochastic Window Transformer for Image Restoration.....	9315
<i>Xueyang Fu, Feng Wu, Jie Xiao, Zheng-Jun Zha</i>	

Phase Transition from Clean Training to Adversarial Training	9330
<i>Guang Cheng, Qifan Song, Yue Xing</i>	
PFL-Bench: A Comprehensive Benchmark for Personalized Federated Learning	9344
<i>Daoyuan Chen, Bolin Ding, Dawei Gao, Weirui Kuang, Yaliang Li</i>	
Squeezeformer: An Efficient Transformer for Automatic Speech Recognition.....	9361
<i>Amir Gholami, Kurt Keutzer, Sehoon Kim, Nicholas Lee, Michael Mahoney, Jitendra Malik, Karttikeya Mangalam, Albert Shaw</i>	
Deep Generalized Schrödinger Bridge	9374
<i>Tianrong Chen, Guan-Horng Liu, Oswin So, Evangelos Theodorou</i>	
Characterizing the Ventral Visual Stream with Response-Optimized Neural Encoding Models.....	9389
<i>Keith Jamison, Meenakshi Khosla, Amy Kuceyeski, Mert Sabuncu</i>	
Learning Sparse Features Can Lead to Overfitting in Neural Networks	9403
<i>Francesco Cagnetta, Leonardo Petrini, Eric Vanden-Eijnden, Matthieu Wyart</i>	
Public Wisdom Matters! Discourse-Aware Hyperbolic Fourier Co-Attention for Social Text Classification	9417
<i>Md Shad Akhtar, S M Phaneendra Angara, Tanmoy Chakraborty, Karish Grover</i>	
Harmonizing the Object Recognition Strategies of Deep Neural Networks with Humans	9432
<i>Thomas Fel, Drew Linsley, Ivan F Rodriguez Rodriguez, Thomas Serre</i>	
Learning Mixed Multinomial Logits with Provable Guarantees	9447
<i>Yiqun Hu, David Simchi-Levi, Zhenzhen Yan</i>	
Defining and Characterizing Reward Hacking.....	9460
<i>Nikolaus Howe, Dmitrii Krasheninnikov, David Krueger, Joar Skalse</i>	
Learning Distinct and Representative Modes for Image Captioning.....	9472
<i>Qi Chen, Chaorui Deng, Qi Wu</i>	
Lifting the Information Ratio: An Information-Theoretic Analysis of Thompson Sampling for Contextual Bandits	9486
<i>Gergely Neu, Iuliia Olkhovskaia, Matteo Papini, Ludovic Schwartz</i>	
Is This the Right Neighborhood? Accurate and Query Efficient Model Agnostic Explanations.....	9499
<i>Amit Dhurandhar, Karthikeyan Natesan Ramamurthy, Karthikeyan Shanmugam</i>	
Object Scene Representation Transformer	9512
<i>Daniel Duckworth, Klaus Greff, Leonidas Guibas, Thomas Kipf, Mario Lucic, Aravindh Mahendran, Filip Pavetic, Mehdi S. M. Sajjadi, Sjoerd Van Steenkiste</i>	
Data-Driven Conditional Robust Optimization	9525
<i>Nymisha Bandi, Abhilash Reddy Chenreddy, Erick Delage</i>	
COLD Decoding: Energy-Based Constrained Text Generation with Langevin Dynamics	9538
<i>Yejin Choi, Daniel Khashabi, Lianhui Qin, Sean Welleck</i>	
Matching in Multi-Arm Bandit with Collision	9552
<i>Zhixuan Fang, Siwei Wang, Yirui Zhang</i>	
Multimodal Contrastive Learning with LIMoE: The Language-Image Mixture of Experts	9564
<i>Neil Houlsby, Rodolphe Jenatton, Basil Mustafa, Joan Puigcerver, Carlos Riquelme</i>	

Few-Shot Learning for Feature Selection with Hilbert-Schmidt Independence Criterion.....	9577
<i>Yasuhiro Fujiwara, Yasutoshi Ida, Tomoharu Iwata, Atsutoshi Kumagai</i>	
Statistical Learning and Inverse Problems: A Stochastic Gradient Approach	9591
<i>Yuri Fonseca, Yuri Saporito</i>	
Hyperparameter Sensitivity in Deep Outlier Detection: Analysis and a Scalable Hyper-Ensemble Solution	9603
<i>Leman Akoglu, Xueying Ding, Lingxiao Zhao</i>	
TVLT: Textless Vision-Language Transformer	9617
<i>Mohit Bansal, Jaemin Cho, Yixin Nie, Zineng Tang</i>	
A Deep Reinforcement Learning Framework for Column Generation.....	9633
<i>Amine Aboussalah, Cheng Chi, Elias Khalil, Zoha Sherkat-Masoumi, Juyoung Wang</i>	
Learning to Drop Out: An Adversarial Approach to Training Sequence VAEs	9645
<i>Carl Allen, Joachim M Buhmann, Kushal Jain, Djordje Miladinovic, Max Paulus, Kumar Shridhar</i>	
On the Adversarial Robustness of Mixture of Experts	9660
<i>Pranjal Awasthi, Srinadh Bhojanapalli, Rodolphe Jenatton, Joan Puigcerver, Carlos Riquelme</i>	
Safety Guarantees for Neural Network Dynamic Systems Via Stochastic Barrier Functions	9672
<i>Morteza Lahijanian, Luca Laurenti, Rayan Mazouz, Karan Muvvala, Akash Ratheesh Babu</i>	
Simple and Optimal Greedy Online Contention Resolution Schemes	9687
<i>Vasilis Livanos</i>	
Composition Theorems for Interactive Differential Privacy	9700
<i>Xin Lyu</i>	
On Margins and Generalisation for Voting Classifiers	9713
<i>Felix Biggs, Benjamin Guedj, Valentina Zantedeschi</i>	
Effective Backdoor Defense by Exploiting Sensitivity of Poisoned Samples	9727
<i>Weixin Chen, Haoqian Wang, Baoyuan Wu</i>	
Rethinking the Reverse-Engineering of Trojan Triggers	9738
<i>Hailun Ding, Shiqing Ma, Kai Mei, Zhenting Wang, Juan Zhai</i>	
Antigen-Specific Antibody Design and Optimization with Diffusion-Based Generative Models for Protein Structures	9754
<i>Shitong Luo, Jianzhu Ma, Jian Peng, Xingang Peng, Yufeng Su, Sheng Wang</i>	
Learning Single-Index Models with Shallow Neural Networks.....	9768
<i>Alberto Bietti, Joan Bruna, Clayton Sanford, Min Jae Song</i>	
Information Bottleneck Theory of High-Dimensional Regression: Relevancy, Efficiency and Optimality.....	9784
<i>Vudtiwat Ngampruetikorn, David Schwab</i>	
RainNet: A Large-Scale Imagery Dataset and Benchmark for Spatial Precipitation Downscaling	9797
<i>Xuanhong Chen, Kairui Feng, Naiyuan Liu, Ziang Liu, Yifan Lu, Bingbing Ni, Zhengyan Tong</i>	

Dataset Distillation Using Neural Feature Regression	9813
<i>Jimmy Ba, Ehsan Nezhadarya, Yongchao Zhou</i>	
ZeroC: A Neuro-Symbolic Model for Zero-Shot Concept Recognition and Acquisition at Inference Time.....	9828
<i>Jure Leskovec, Kevin Liu, Rok Susic, Megan Tjandrasuwita, Tailin Wu, Zhengxuan Wu, Xuelin Yang</i>	
Theoretically Better and Numerically Faster Distributed Optimization with Smoothness-Aware Quantization Techniques	9841
<i>Peter Richtarik, Mher Safaryan, Bokun Wang</i>	
Sparse Structure Search for Delta Tuning	9853
<i>Ning Ding, Shengding Hu, Zhiyuan Liu, Maosong Sun, Yadao Wang, Yasheng Wang, Zhen Zhang</i>	
On the Safety of Interpretable Machine Learning: A Maximum Deviation Approach	9866
<i>Elizabeth Daly, Amit Dhurandhar, Rahul Nair, Moninder Singh, Kush Varshney, Dennis Wei</i>	
Non-Stationary Transformers: Exploring the Stationarity in Time Series Forecasting	9881
<i>Yong Liu, Mingsheng Long, Jianmin Wang, Haixu Wu</i>	
Risk-Driven Design of Perception Systems	9894
<i>Anthony Corso, Xin Du, Craig Innes, Sydney Katz, Mykel J Kochenderfer, Subramanian Ramamoorthy</i>	
A Simple Approach to Automated Spectral Clustering	9907
<i>Jicong Fan, Yiheng Tu, Haijun Zhang, Zhao Zhang, Mingbo Zhao</i>	
Joint Entropy Search for Multi-Objective Bayesian Optimization.....	9922
<i>Axel Gandy, Nikolas Kantas, Behrang Shafei, Ben Tu</i>	
Subgroup Robustness Grows on Trees: An Empirical Baseline Investigation	9939
<i>Josh Gardner, Zoran Popovic, Ludwig Schmidt</i>	
Robustness to Unbounded Smoothness of Generalized SignSGD	9955
<i>Michael Crawshaw, Mingrui Liu, Francesco Orabona, Wei Zhang, Zhenxun Zhuang</i>	
GhostNetV2: Enhance Cheap Operation with Long-Range Attention	9969
<i>Jianyuan Guo, Kai Han, Yehui Tang, Yunhe Wang, Chang Xu, Chao Xu</i>	

VOLUME 14

Analyzing Sharpness Along GD Trajectory: Progressive Sharpening and Edge of Stability	9983
<i>Jian Li, Zhouzi Li, Zixuan Wang</i>	
You Never Stop Dancing: Non-Freezing Dance Generation Via Bank-Constrained Manifold Projection	9995
<i>Huang Hu, Jian-Fang Hu, Zhi Jin, Hanjiang Lai, Jiangxin Sun, Chunyu Wang</i>	
Nonnegative Tensor Completion Via Integer Optimization.....	10008
<i>Anil Aswani, Caleb Bugg, Chen Chen</i>	
LION: Latent Point Diffusion Models for 3D Shape Generation.....	10021
<i>Sanja Fidler, Zan Gojcic, Karsten Kreis, Or Litany, Arash Vahdat, Francis Williams, Xiaohui Zeng</i>	

Distribution-Informed Neural Networks for Domain Adaptation Regression	10040
<i>Elizabeth Ainsworth, Kaiyu Guan, Jingrui He, Sheng Wang, Jun Wu</i>	
Momentum Adversarial Distillation: Handling Large Distribution Shifts in Data-Free Knowledge Distillation	10055
<i>Kien Do, Haripriya Harikumar, Thai Hung Le, Dang Nguyen, Dung Nguyen, Santu Rana, Truyen Tran, Svetha Venkatesh</i>	
Hard ImageNet: Segmentations for Objects with Strong Spurious Cues	10068
<i>Soheil Feizi, Mazda Moayeri, Sahil Singla</i>	
VideoMAE: Masked Autoencoders Are Data-Efficient Learners for Self-Supervised Video Pre-Training	10078
<i>Yibing Song, Zhan Tong, Jue Wang, Limin Wang</i>	
Scalable Design of Error-Correcting Output Codes Using Discrete Optimization with Graph Coloring	10094
<i>Saurabh Amin, Samarth Gupta</i>	
A New Family of Generalization Bounds Using Samplewise Evaluated CMI.....	10108
<i>Giuseppe Durisi, Fredrik Hellström</i>	
Maximum-Likelihood Inverse Reinforcement Learning with Finite-Time Guarantees	10122
<i>Alfredo Garcia, Mingyi Hong, Chenliang Li, Siliang Zeng</i>	
Lost in Latent Space: Examining Failures of Disentangled Models at Combinatorial Generalisation.....	10136
<i>Jeffrey Bowers, Casimir Ludwig, Gaurav Malhotra, Milton Montero, Rui Ponte Costa</i>	
MultiGuard: Provably Robust Multi-Label Classification Against Adversarial Examples	10150
<i>Neil Gong, Jinyuan Jia, Wenjie Qu</i>	
On Measuring Excess Capacity in Neural Networks.....	10164
<i>Florian Graf, Roland Kwitt, Marc Niethammer, Bastian Rieck, Sebastian Zeng</i>	
On Leave-One-Out Conditional Mutual Information for Generalization	10179
<i>Alessandro Achille, Suhas Diggavi, Aditya Golatkar, Mohamad Rida Rammal, Stefano Soatto</i>	
Near-Isometric Properties of Kronecker-Structured Random Tensor Embeddings	10191
<i>Qijia Jiang</i>	
MExMI: Pool-Based Active Model Extraction Crossover Membership Inference	10203
<i>Chengfang Fang, Haibo Hu, Jie Shi, Yaxin Xiao, Qingqing Ye, Huadi Zheng</i>	
Parameter-Efficient Masking Networks	10217
<i>Yue Bai, Yun Fu, Xu Ma, Zhiqiang Tao, Huan Wang, Yitian Zhang</i>	
How Sampling Impacts the Robustness of Stochastic Neural Networks.....	10230
<i>Sina Däubener, Asja Fischer</i>	
Toward Robust Spiking Neural Network Against Adversarial Perturbation	10244
<i>Lei Deng, Xing Hu, Ling Liang, Yuan Xie, Kaidi Xu</i>	
GRASP: Navigating Retrosynthetic Planning with Goal-Driven Policy.....	10257
<i>Zhengxing Huang, Kun Kuang, Ying Wei, Fei Wu, Huaxiu Yao, Yemin Yu</i>	
End-To-End Symbolic Regression with Transformers	10269
<i>Francois Charton, Stéphane D'Ascoli, Pierre-Alexandre Kamienny, Guillaume Lample</i>	

Gold-Standard Solutions to the Schrödinger Equation Using Deep Learning: How Much Physics Do We Need?.....	10282
<i>Leon Gerard, Philipp Grohs, Philipp Marquetand, Michael Scherbela</i>	
EcoFormer: Energy-Saving Attention with Linear Complexity	10295
<i>Jianfei Cai, Haoyu He, Jing Liu, Zizheng Pan, Bohan Zhuang</i>	
Wild-Time: A Benchmark of in-The-Wild Distribution Shift Over Time.....	10309
<i>Bochuan Cao, Caroline Choi, Chelsea Finn, Pang Wei Koh, Yoonho Lee, Huaxiu Yao</i>	
Sound and Complete Causal Identification with Latent Variables Given Local Background Knowledge.....	10325
<i>Tian Qin, Tian-Zuo Wang, Zhi-Hua Zhou</i>	
A Unified Diversity Measure for Multiagent Reinforcement Learning.....	10339
<i>Yuan Li, Zongkai Liu, Peng Sun, Zifan Wu, Yaodong Yang, Chao Yu</i>	
HornNet: Efficient High-Order Spatial Interactions with Recursive Gated Convolutions	10353
<i>Ser Nam Lim, Jiwen Lu, Yongming Rao, Yansong Tang, Wenliang Zhao, Jie Zhou</i>	
Learning Tractable Probabilistic Models from Inconsistent Local Estimates	10367
<i>Vibhav Gogate, Shasha Jin, Vasundhara Komaragiri, Tahrira Rahman</i>	
Decomposed Knowledge Distillation for Class-Incremental Semantic Segmentation.....	10380
<i>Donghyeon Baek, Bumsub Ham, Junghyup Lee, Sanghoon Lee, Youngmin Oh</i>	
Minimax Optimal Algorithms for Fixed-Budget Best Arm Identification	10393
<i>Junya Honda, Junpei Komiyama, Taira Tsuchiya</i>	
Reduced Representation of Deformation Fields for Effective Non-Rigid Shape Matching.....	10405
<i>Riccardo Marin, Maks Ovsjanikov, Emanuele Rodolà, Ramana Subramanyam Sundararaman</i>	
BEVFusion: A Simple and Robust LiDAR-Camera Fusion Framework.....	10421
<i>Tingting Liang, Zhiwei Lin, Tao Tang, Zhi Tang, Bing Wang, Yongtao Wang, Zhongyu Xia, Hongwei Xie, Kaicheng Yu</i>	
Automatic Differentiation of Programs with Discrete Randomness.....	10435
<i>Gaurav Arya, Christopher Rackauckas, Frank Schäfer, Moritz Schauer</i>	
A Closer Look at the Adversarial Robustness of Deep Equilibrium Models.....	10448
<i>Yang Liu, Tianyu Pang, Zonghan Yang</i>	
Near-Optimal Private and Scalable k-Clustering.....	10462
<i>Vincent Cohen-Addad, Alessandro Epasto, Vahab Mirrokni, Shyam Narayanan, Peilin Zhong</i>	
NS ³ : Neuro-Symbolic Semantic Code Search.....	10476
<i>Miltiadis Allamanis, Shushan Arakelyan, Luis Garcia, Anna Hakhverdyan, Christophe Hauser, Xiang Ren</i>	
Revisiting Sparse Convolutional Model for Visual Recognition.....	10492
<i>Xili Dai, Xingjian Gao, Shao-Lun Huang, Mingyang Li, Yi Ma, Shengbang Tong, Chong You, Pengyuan Zhai, Zhihui Zhu</i>	
Temporal Latent Bottleneck: Synthesis of Fast and Slow Processing Mechanisms in Sequence Learning	10505
<i>Yoshua Bengio, Aniket Didolkar, Anirudh Goyal, Nitesh Bharadwaj Gundavarapu, Kshitij Gupta, Nan Rosemary Ke, Alex Lamb</i>	

A Near-Optimal Primal-Dual Method for Off-Policy Learning in CMDP	10521
<i>Fan Chen, Zaiwen Wen, Junyu Zhang</i>	
DReS-FL: Dropout-Resilient Secure Federated Learning for Non-IID Clients Via Secret Data Sharing	10533
<i>Songze Li, Jiawei Shao, Yuchang Sun, Jun Zhang</i>	
BackdoorBench: A Comprehensive Benchmark of Backdoor Learning	10546
<i>Hongrui Chen, Chao Shen, Shaokui Wei, Baoyuan Wu, Danni Yuan, Mingda Zhang, Zihao Zhu</i>	
REVIVE: Regional Visual Representation Matters in Knowledge-Based Visual Question Answering	10560
<i>Dongdong Chen, Yuanze Lin, Yujia Xie, Yichong Xu, Lu Yuan, Chenguang Zhu</i>	
FedAvg with Fine Tuning: Local Updates Lead to Representation Learning.....	10572
<i>Liam Collins, Hamed Hassani, Aryan Mokhtari, Sanjay Shakkottai</i>	
Modeling Transitivity and Cyclicity in Directed Graphs Via Binary Code Box Embeddings.....	10587
<i>Michael Boratko, Andrew McCallum, Cameron Musco, Dongxu Zhang</i>	
Generalization Bounds for Gradient Methods Via Discrete and Continuous Prior	10600
<i>Jian Li, Luo Bei, Xuanyuan Luo</i>	
On Deep Generative Models for Approximation and Estimation of Distributions on Manifolds	10615
<i>Minshuo Chen, Biraj Dahal, Alexander Havrilla, Wenjing Liao, Tuo Zhao</i>	
Memory Efficient Continual Learning with Transformers	10629
<i>Cedric Archambeau, Beyza Ermis, Aditya Rawal, Martin Wistuba, Giovanni Zappella</i>	
Counterfactual Neural Temporal Point Process for Estimating Causal Influence of Misinformation on Social Media.....	10643
<i>Defu Cao, Yan Liu, Yizhou Zhang</i>	
Curriculum Reinforcement Learning Using Optimal Transport Via Gradual Domain Adaptation.....	10656
<i>Fei Fang, Peide Huang, Laixi Shi, Mengdi Xu, Ding Zhao, Jiacheng Zhu</i>	
Temporally-Consistent Survival Analysis	10671
<i>Lucas Maystre, Daniel Russo</i>	
Symbolic Distillation for Learned TCP Congestion Control.....	10684
<i>Ang Chen, Kuo-Feng Hsu, S P Sharan, Zhangyang Wang, Jiarong Xing, Wenqing Zheng</i>	
A Probabilistic Graph Coupling View of Dimension Reduction	10696
<i>Julien Chiquet, Thibault Espinasse, Franck Picard, Hugues Van Assel</i>	
Hardness of Noise-Free Learning for Two-Hidden-Layer Neural Networks	10709
<i>Sitan Chen, Aravind Gollakota, Adam Klivans, Raghu Meka</i>	
Convergence Beyond the Over-Parameterized Regime Using Rayleigh Quotients	10725
<i>Marc Lelarge, David A. R. Robin, Kevin Scaman</i>	
Optimistic Posterior Sampling for Reinforcement Learning with Few Samples and Tight Guarantees.....	10737
<i>Denis Belomestny, Daniele Calandriello, Pierre Ménard, Eric Moulines, Remi Munos, Alexey Naumov, Mark Rowland, Daniil Tiapkin, Michal Valko</i>	

On Convergence of FedProx: Local Dissimilarity Invariant Bounds, Non-Smoothness and Beyond..... 10752
Ping Li, Xiaotong Yuan

VOLUME 15

Can Push-Forward Generative Models Fit Multimodal Distributions? 10766
Valentin De Bortoli, Julie Delon, Agnes Desolneux, Antoine Salmona

Improved Imaging by Inx Regularizers with Global Optima Guarantees..... 10780
Neil Bourne, Tingting Mu, Samuel Pinilla, Jeyan Thiyagalingam

The Neural Covariance SDE: Shaped Infinite Depth-And-Width Networks at Initialization 10795
Mufan Li, Mihai Nica, Dan Roy

Rethinking and Scaling Up Graph Contrastive Learning: An Extremely Efficient Approach with
Group Discrimination..... 10809
Vincent CS Lee, Shirui Pan, Philip S Yu, Yu Zheng, Yizhen Zheng

Diverse Weight Averaging for Out-Of-Distribution Generalization 10821
*Matthieu Cord, Patrick Gallinari, Matthieu Kirchmeyer, Thibaud Rahier, Alain
Rakotomamonjy, Alexandre Rame*

Bivariate Causal Discovery for Categorical Data Via Classification with Optimal Label
Permutation 10837
Yang Ni

Phase Transitions in When Feedback is Useful 10849
Lokesh Boominathan, Xaq Pitkow

Stochastic Second-Order Methods Improve Best-Known Sample Complexity of SGD for Gradient-
Dominated Functions..... 10862
Niao He, Negar Kiyavash, Saeed Masiha, Saber Salehkaleybar, Patrick Thiran

Posterior and Computational Uncertainty in Gaussian Processes 10876
John Cunningham, Philipp Hennig, Marvin Pförtner, Geoff Pleiss, Jonathan Wenger

Understanding the Evolution of Linear Regions in Deep Reinforcement Learning 10891
Setareh Cohan, Nam Hee Kim, David Rolnick, Michiel Van De Panne

Causal Discovery in Heterogeneous Environments Under the Sparse Mechanism Shift Hypothesis 10904
Ronan Perry, Bernhard Schölkopf, Julius Von Kügelgen

Towards Practical Control of Singular Values of Convolutional Layers 10918
Ekaterina Bulatova, Anton Obukhov, Maxim Rakhuba, Alexandra Senderovich

Deep Multi-Modal Structural Equations for Causal Effect Estimation with Unstructured Proxies..... 10931
Shachi Deshpande, Volodymyr Kuleshov, Zheng Li, Dhruv Sreenivas, Kaiwen Wang

Generic Bounds on the Approximation Error for Physics-Informed (and) Operator Learning 10945
Tim De Ryck, Siddhartha Mishra

Learning Distributions Generated by Single-Layer ReLU Networks in the Presence of Arbitrary
Outliers 10959
Saikiran Bulusu, M. Cenk Gursoy, Geethu Joseph, Pramod Varshney

GenerSpeech: Towards Style Transfer for Generalizable Out-Of-Domain Text-To-Speech	10970
<i>Chenye Cui, Rongjie Huang, Jinglin Liu, Yi Ren, Zhou Zhao</i>	
On Non-Linear Operators for Geometric Deep Learning.....	10984
<i>Joan Bruna, Jakob Maier, Edouard Oyallon, Grégoire Sergeant-Perthuis</i>	
Momentum Aggregation for Private Non-Convex ERM.....	10996
<i>Ashok Cutkosky, Hoang Tran</i>	
Learning in Congestion Games with Bandit Feedback	11009
<i>Qiwen Cui, Simon Du, Maryam Fazel, Zhihan Xiong</i>	
LOG: Active Model Adaptation for Label-Efficient OOD Generalization.....	11023
<i>Lan-Zhe Guo, Yu-Feng Li, Jie-Jing Shao, Xiao-Wen Yang</i>	
PolarMix: A General Data Augmentation Technique for LiDAR Point Clouds	11035
<i>Kaiwen Cui, Dayan Guan, Jiaying Huang, Shijian Lu, Ling Shao, Aoran Xiao</i>	
Learning from Distributed Users in Contextual Linear Bandits Without Sharing the Context	11049
<i>Christina Fragouli, Osama Hanna, Lin Yang</i>	
Provably Feedback-Efficient Reinforcement Learning Via Active Reward Learning.....	11063
<i>Dingwen Kong, Lin Yang</i>	
Recurrent Memory Transformer.....	11079
<i>Aydar Bulatov, Mikhail Burtsev, Yury Kuratov</i>	
Hierarchical Lattice Layer for Partially Monotone Neural Networks	11092
<i>Takayuki Katsuki, Kohei Miyaguchi, Hiroki Yanagisawa</i>	
Class-Dependent Label-Noise Learning with Cycle-Consistency Regularization	11104
<i>De Cheng, Yuxuan Du, Xinbo Gao, Bo Han, Tongliang Liu, Yixiong Ning, Nannan Wang, Heng Yang</i>	
SeqPATE: Differentially Private Text Generation Via Knowledge Distillation.....	11117
<i>He He, Ziyue Huang, Zhiliang Tian, Yu-Xiang Wang, Nevin L. Zhang, Yingxiu Zhao</i>	
Bayesian Persuasion for Algorithmic Recourse.....	11131
<i>Valerie Chen, Keegan Harris, Hoda Heidari, Joon Kim, Ameet Talwalkar, Steven Wu</i>	
Instance-Based Uncertainty Estimation for Gradient-Boosted Regression Trees.....	11145
<i>Jonathan Brophy, Daniel Lowd</i>	
Transormer: Transformer for Sentence Scoring with Sliding Language Modeling	11160
<i>Yichong Leng, Dongsheng Li, Tao Qin, Kaitao Song, Xu Tan, Yicheng Zou</i>	
Models Out of Line: A Fourier Lens on Distribution Shift Robustness.....	11175
<i>Brian Bartoldson, Timo Bremer, James Diffenderfer, Sara Fridovich-Keil, Bhavya Kailkhura</i>	
Deep Learning Methods for Proximal Inference Via Maximum Moment Restriction	11189
<i>Andrew Beam, David Bellamy, Tom Kolokotronis, Benjamin Kompa, James M Robins</i>	
Nest Your Adaptive Algorithm for Parameter-Agnostic Nonconvex Minimax Optimization	11202
<i>Niao He, Xiang Li, Junchi Yang</i>	
Brownian Noise Reduction: Maximizing Privacy Subject to Accuracy Constraints.....	11217
<i>Aaditya Ramdas, Ryan Rogers, Justin Whitehouse, Steven Wu</i>	

Are AlphaZero-Like Agents Robust to Adversarial Perturbations?.....	11229
<i>Cho-Jui Hsieh, Li-Cheng Lan, Meng-Yu Tsai, I-Chen Wu, Ti-Rong Wu, Huan Zhang</i>	
AnimeSR: Learning Real-World Super-Resolution Models for Animation Videos	11241
<i>Gen Li, Ying Shan, Xintao Wang, Yanze Wu</i>	
On the Spectral Bias of Convolutional Neural Tangent and Gaussian Process Kernels.....	11253
<i>Meirav Galun, Amnon Geifman, David Jacobs, Basri Ronen</i>	
Fairness Transferability Subject to Bounded Distribution Shift.....	11266
<i>Yatong Chen, Yang Liu, Reilly Raab, Jialu Wang</i>	
Improving Self-Supervised Learning by Characterizing Idealized Representations	11279
<i>Yann Dubois, Stefano Ermon, Tatsunori Hashimoto, Percy Liang</i>	
On the Difficulty of Learning Chaotic Dynamics with RNNs.....	11297
<i>Daniel Durstewitz, Jonas Mikhaeil, Zahra Monfared</i>	
SKFlow: Learning Optical Flow with Super Kernels.....	11313
<i>Yuanqi Chen, Guodong Guo, Ge Li, Shangkun Sun, Yu Zhu</i>	
Mingling Foresight with Imagination: Model-Based Cooperative Multi-Agent Reinforcement Learning	11327
<i>Yunpeng Baiia, Guoliang Fan, Dapeng Li, Zhiwei Xu, Yuan Zhan, Bin Zhang</i>	
End-To-End Stochastic Optimization with Energy-Based Model	11341
<i>Jiaming Cui, Rui Feng, Lingkai Kong, B. Aditya Prakash, Chao Zhang, Yuchen Zhuang</i>	
Learning Low-Dimensional Generalizable Natural Features from Retina Using a U-Net	11355
<i>Michael Berry, Elizabeth De Laittre, Benjamin Hoshal, Thierry Mora, Stephanie Palmer, Siwei Wang</i>	
Split-KI and PAC-Bayes-Split-KI Inequalities for Ternary Random Variables	11369
<i>Yevgeny Seldin, Yi-Shan Wu</i>	
Wasserstein K-Means for Clustering Probability Distributions.....	11382
<i>Xiaohui Chen, Yun Yang, Yubo Zhuang</i>	
Learning Long-Term Crop Management Strategies with CyclesGym	11396
<i>Ioannis Athanasiadis, Joachim M Buhmann, Amanda Burton, Luca Corinzia, Juan Herrera, Andreas Krause, Scott Sussex, Matteo Turchetta</i>	
Identification, Amplification and Measurement: A Bridge to Gaussian Differential Privacy.....	11410
<i>Bei Jiang, Linglong Kong, Yi Liu, Ke Sun</i>	
MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields	11423
<i>Ilyes Batatia, Gabor Csanyi, David P Kovacs, Christoph Ortner, Gregor Simm</i>	
Follow-The-Perturbed-Leader for Adversarial Markov Decision Processes with Bandit Feedback	11437
<i>Liyu Chen, Yan Dai, Haipeng Luo</i>	
Improving Multi-Task Generalization Via Regularizing Spurious Correlation	11450
<i>Ed Chi, Lichan Hong, Ziniu Hu, Yizhou Sun, Tiansheng Yao, Xinyang Yi, Zhe Zhao</i>	
Estimating and Explaining Model Performance When Both Covariates and Labels Shift.....	11467
<i>Lingjiao Chen, Matei Zaharia, James Zou</i>	

Tenrec: A Large-Scale Multipurpose Benchmark Dataset for Recommender Systems	11480
<i>Lei Chen, Bo Hu, Beibei Kong, Shujie Li, Yudong Li, Zang Li, Xiaohu Qie, Yu Xu, Min Yang, Chenyun Yu, Fajie Yuan, Guanghu Yuan</i>	
Joint Entropy Search for Maximally-Informed Bayesian Optimization	11494
<i>Frank Hutter, Carl Hvarfner, Luigi Nardi</i>	
Provable General Function Class Representation Learning in Multitask Bandits and MDP	11507
<i>Simon Du, Gao Huang, Rui Lu, Andrew Zhao</i>	
Weighted Mutual Learning with Diversity-Driven Model Compression	11520
<i>David Campos, Chenjuan Guo, Wei Huang, Li Wang, Bin Yang, Miao Zhang</i>	
S2P: State-Conditioned Image Synthesis for Data Augmentation in Offline Reinforcement Learning	11534
<i>Daesol Cho, H. Jin Kim, Dongseok Shim</i>	

VOLUME 16

Neural Collapse with Normalized Features: A Geometric Analysis Over the Riemannian Manifold	11547
<i>Laura Balzano, Qing Qu, Peng Wang, Can Yaras, Zhihui Zhu</i>	
Conformalized Fairness Via Quantile Regression	11561
<i>Lei Ding, Bei Jiang, Linglong Kong, Meichen Liu, Wulong Liu, Dengdeng Yu</i>	
Efficient Methods for Non-Stationary Online Learning	11573
<i>Yan-Feng Xie, Lijun Zhang, Peng Zhao, Zhi-Hua Zhou</i>	
Increasing the Scope as You Learn: Adaptive Bayesian Optimization in Nested Subspaces	11586
<i>Luigi Nardi, Leonard Papenmeier, Matthias Poloczek</i>	
ACIL: Analytic Class-Incremental Learning with Absolute Memorization and Privacy Protection	11602
<i>Zhiping Lin, Kar-Ann Toh, Hongxin Wei, Zhenyu Weng, Renchunzi Xie, Huiping Zhuang</i>	
MsSVT: Mixed-Scale Sparse Voxel Transformer for 3D Object Detection on Point Clouds	11615
<i>Ziyang Bian, Lihe Ding, Shaocong Dong, Jianan Li, Haiyang Wang, Jie Wang, Ying Wang, Tingfa Xu, Xinli Xu</i>	
Deterministic Langevin Monte Carlo with Normalizing Flows for Bayesian Inference	11629
<i>Biwei Dai, Richard Grumitt, Uros Seljak</i>	
Evolution of Neural Tangent Kernels Under Benign and Adversarial Training	11642
<i>Alexander Amini, Ramin Hasani, Noel Loo, Daniela Rus</i>	
Zero-Sum Stochastic Stackelberg Games	11658
<i>Denizalp Goktas, Amy Greenwald, Sadie Zhao</i>	
Evaluating Out-Of-Distribution Performance on Document Image Classifiers	11673
<i>Yutong Ai, David Kuang, Stefan Larson, Kevin Leach, Yi Yang Gordon Lim</i>	
Distributed Learning of Conditional Quantiles in the Reproducing Kernel Hilbert Space	11686
<i>Heng Lian</i>	
Spatial Mixture-Of-Experts	11697
<i>Nikoli Dryden, Torsten Hoefler</i>	

Amortized Mixing Coupling Processes for Clustering.....	11714
<i>Liping Jing, Huafeng Liu</i>	
Hilbert Distillation for Cross-Dimensionality Networks.....	11726
<i>Jiajun Bu, Zhe Liu, Dian Qin, Haishuai Wang, Hongjia Xu, Sheng Zhou</i>	
Provably Efficient Offline Multi-Agent Reinforcement Learning Via Strategy-Wise Bonus.....	11739
<i>Qiwen Cui, Simon Du</i>	
A Best-Of-Both-Worlds Algorithm for Bandits with Delayed Feedback	11752
<i>Saeed Masoudian, Yevgeny Seldin, Julian Zimmert</i>	
LIFT: Language-Interfaced Fine-Tuning for Non-Language Machine Learning Tasks	11763
<i>Tuan Dinh, Michael Gira, Kangwook Lee, Ziqian Lin, Dimitris Papailiopoulos, Shashank Rajput, Jy-Yong Sohn, Yuchen Zeng, Ruisu Zhang</i>	
Capturing Failures of Large Language Models Via Human Cognitive Biases	11785
<i>Erik Jones, Jacob Steinhardt</i>	
Template Based Graph Neural Network with Optimal Transport Distances	11800
<i>Marco Corneli, Nicolas Courty, Rémi Flamary, Titouan Vayer, Cédric Vincent-Cuaz</i>	
Knowledge Distillation Improves Graph Structure Augmentation for Graph Neural Networks.....	11815
<i>Yufei Huang, Stan Z. Li, Haitao Lin, Lirong Wu</i>	
Learning Invariant Graph Representations for Out-Of-Distribution Generalization.....	11828
<i>Haoyang Li, Xin Wang, Ziwei Zhang, Wenwu Zhu</i>	
A Quadrature Rule Combining Control Variates and Adaptive Importance Sampling.....	11842
<i>Rémi Leluc, François Portier, Johan Segers, Aigerim Zhuman</i>	
GAL: Gradient Assisted Learning for Decentralized Multi-Organization Collaborations	11854
<i>Enmao Diao, Jie Ding, Vahid Tarokh</i>	
Direct Advantage Estimation.....	11869
<i>Nico Grtler, Alexander Neitz, Hsiao-Ru Pan, Bernhard Schlkopf</i>	
Honor of Kings Arena: An Environment for Generalization in Competitive Reinforcement Learning	11881
<i>Jingxiao Chen, Minwen Deng, Qiang Fu, Lanxiao Huang, Xiyang Ji, Siqin Li, Liu Linc, Hongyang Qin, Liang Wang, Hua Wei, Wei Yang, Deheng Ye, Yong Yu, Weinan Zhang</i>	
On the Symmetries of Deep Learning Models and Their Internal Representations	11893
<i>Davis Brown, Tegan Emerson, Charles Godfrey, Henry Kvinge</i>	
TabNAS: Rejection Sampling for Neural Architecture Search on Tabular Datasets	11906
<i>Gabriel Bender, Da Huang, Pieter-Jan Kindermans, Quoc V Le, Hanxiao Liu, Yifeng Lu, Madeleine Udell, Chengrun Yang</i>	
Embed and Emulate: Learning to Estimate Parameters of Dynamical Systems with Uncertainty Quantification.....	11918
<i>Ruoxi Jiang, Rebecca Willett</i>	
A Differentiable Semantic Metric Approximation in Probabilistic Embedding for Cross-Modal Retrieval	11934
<i>Lianli Gao, Gongfu Li, Hao Li, Jingkuan Song, Pengpeng Zeng, Haonan Zhang</i>	

Friendly Noise Against Adversarial Noise: A Powerful Defense Against Data Poisoning Attack.....	11947
<i>Tian Yu Liu, Baharan Mirzasoleiman, Yu Yang</i>	
Dict-TTS: Learning to Pronounce with Prior Dictionary Knowledge for Text-To-Speech	11960
<i>Ziyue Jiang, Jinglin Liu, Yi Ren, Zhe Su, Qian Yang, Zhou Zhao, Zhenhui Ye</i>	
AZ-Whiteness Test: A Test for Signal Uncorrelation on Spatio-Temporal Graphs	11975
<i>Cesare Alippi, Daniele Zambon</i>	
ViSioNS: Visual Search in Natural Scenes Benchmark.....	11987
<i>Gaston Bujia, Juan Esteban Kamienkowski, Gonzalo Ruarte, Fermín Travi</i>	
Value Function Decomposition for Iterative Design of Reinforcement Learning Agents	12001
<i>Evan Archer, Alisa Devlic, James Macglashan, Takuma Seno, Craig Sherstan, Peter Stone, Peter Wurman</i>	
HandMeThat: Human-Robot Communication in Physical and Social Environments.....	12014
<i>Jiayuan Mao, Josh Tenenbaum, Yanming Wan</i>	
Task-Agnostic Graph Explanations	12027
<i>Edward Huang, Shuiwang Ji, Sumeet Katariya, Nikhil Rao, Karthik Subbian, Xianfeng Tang, Yaochen Xie</i>	
Embrace the Gap: VAEs Perform Independent Mechanism Analysis	12040
<i>Michel Besserve, Jack Brady, Wieland Brendel, Luigi Gresele, Georg Martius, Patrik Reizinger, Bernhard Schölkopf, Julius Von Kügelgen, Dominik Zietlow</i>	
Dataset Inference for Self-Supervised Models	12058
<i>Franziska Boenisch, Yannis Cattan, Nikita Dhawan, Haonan Duan, Adam Dziedziec, Jonas Guan, Muhammad Ahmad Kaleem, Nicolas Papernot</i>	
Statistically Meaningful Approximation: A Case Study on Approximating Turing Machines with Transformers.....	12071
<i>Yining Chen, Tengyu Ma, Colin Wei</i>	
Improved Feature Distillation Via Projector Ensemble	12084
<i>Yudong Chen, Frank De Hoog, Zi Huang, Jiajun Liu, Sen Wang, Xuwei Xu</i>	
Does GNN Pretraining Help Molecular Representation?	12096
<i>Hanjun Dai, Ruoxi Sun, Adams Wei Yu</i>	
ULNeF: Untangled Layered Neural Fields for Mix-And-Match Virtual Try-On	12110
<i>Dan Casas, Miguel Otaduy, Igor Santesteban, Nils Thuerey</i>	
Introspective Learning : A Two-Stage Approach for Inference in Neural Networks.....	12126
<i>Ghassan Alregib, Mohit Prabhushankar</i>	
Bayesian Active Learning with Fully Bayesian Gaussian Processes.....	12141
<i>Francisco Antunes, Frederik Hüttel, Carlos Lima Azevedo, Francisco Pereira, Christoffer Riis</i>	
E-MAPP: Efficient Multi-Agent Reinforcement Learning with Parallel Program Guidance.....	12154
<i>Can Chang, Ni Mu, Ling Pan, Jiajun Wu, Huazhe Xu</i>	
In Defense of the Unitary Scalarization for Deep Multi-Task Learning.....	12169
<i>Alessandro De Palma, Ilya Kostrikov, Vitaly Kurin, Pawan K Mudigonda, Shimon Whiteson</i>	

Flexible Neural Image Compression Via Code Editing.....	12184
<i>Chenjian Gao, Dailan He, Hongwei Qin, Yan Wang, Tongda Xu</i>	
Homomorphic Matrix Completion.....	12197
<i>Zechu Steven Li, Xiao-Yang Liu, Xiaodong Wang</i>	
Spending Thinking Time Wisely: Accelerating MCTS with Virtual Expansions.....	12211
<i>Pieter Abbeel, Yang Gao, Weirui Ye</i>	
Intrinsic Dimensionality Estimation Using Normalizing Flows.....	12225
<i>Christian Horvat, Jean-Pascal Pfister</i>	
Moment Distributionally Robust Tree Structured Prediction.....	12237
<i>Kevin Gimpel, Yesu Li, Danyal Saeed, Xinhua Zhang, Brian Ziebart</i>	
Minimax Optimal Fixed-Budget Best Arm Identification in Linear Bandits.....	12253
<i>Vincent Tan, Junwen Yang</i>	
Tempo: Accelerating Transformer-Based Model Training Through Memory Footprint Reduction.....	12267
<i>Muralidhar Andoorveedu, Gennady Pekhimenko, Bojian Zheng, Zhanda Zhu</i>	
I2DFormer: Learning Image to Document Attention for Zero-Shot Image Classification.....	12283
<i>Luc V Gool, Muhammad Ferjad Naeem, Federico Tombari, Yongqin Xian</i>	
Benchmarking Heterogeneous Treatment Effect Models Through the Lens of Interpretability.....	12295
<i>Ioana Bica, Jonathan Crabbé, Alicia Curth, Mihaela Van Der Schaar</i>	
AD-DROP: Attribution-Driven Dropout for Robust Language Model Fine-Tuning.....	12310
<i>Jinghao Deng, Shaoliang Nie, Xiaojun Quan, Qifan Wang, Tao Yang</i>	
Reinforced Genetic Algorithm for Structure-Based Drug Design.....	12325
<i>Connor Coley, Tianfan Fu, Wenhao Gao, Jimeng Sun</i>	

VOLUME 17

A Variational Edge Partition Model for Supervised Graph Representation Learning.....	12339
<i>Bo Chen, Yilin He, Chaojie Wang, Hao Zhang, Mingyuan Zhou</i>	
Learning Optimal Flows for Non-Equilibrium Importance Sampling.....	12352
<i>Yu Cao, Eric Vanden-Eijnden</i>	
Decision Trees with Short Explainable Rules.....	12365
<i>Ferdinando Cicalese, Eduardo Laber, Marco Molinaro, Victor Feitosa Souza</i>	
NAS-Bench-360: Benchmarking Neural Architecture Search on Diverse Tasks.....	12380
<i>Misha Khodak, Nicholas Roberts, Frederic Sala, Junhong Shen, Ameet Talwalkar, Renbo Tu</i>	
Sparse Fourier Backpropagation in Cryo-EM Reconstruction.....	12395
<i>Kiarash Jamali, Dari Kimanius, Sjors Scheres</i>	
Modular Flows: Differential Molecular Generation.....	12409
<i>Vikas Garg, Markus Heinonen, Samuel Kaski, Yogesh Verma</i>	
Anonymous Bandits for Multi-User Systems.....	12422
<i>Hossein Esfandiari, Vahab Mirrokni, Jon Schneider</i>	

Unsupervised Object Detection Pretraining with Joint Object Priors Generation and Detector Learning	12435
<i>Lei Bai, Meilin Chen, Wanli Ouyang, Donglian Qi, Shixiang Tang, Yizhou Wang, Yunfeng Yan, Haiyang Yang, Rui Zhao, Feng Zhu</i>	
Invariance Learning in Deep Neural Networks with Differentiable Laplace Approximations.....	12449
<i>Vincent Fortuin, Alexander Immer, Gunnar Rätsch, Tycho Van Der Ouderaa, Mark Van Der Wilk</i>	
QueryPose: Sparse Multi-Person Pose Regression Via Spatial-Aware Part-Level Query	12464
<i>Mingshu He, Lei Jin, Kai Su, Xiaojuan Wang, Yabo Xiao, Dongdong Yu, Zehuan Yuan</i>	
EAGER: Asking and Answering Questions for Automatic Reward Shaping in Language-Guided RL.....	12478
<i>Thomas Carta, Sylvain Lamprier, Pierre-Yves Oudeyer, Olivier Sigaud</i>	
OpenFilter: A Framework to Democratize Research Access to Social Media AR Filters	12491
<i>Francisco Escolano, Francesco Galati, Thomas Hofmann, Nuria Oliver, Bill Psomas, Piera Riccio</i>	
Improving Policy Learning Via Language Dynamics Distillation.....	12504
<i>Edward Grefenstette, Jesse Mu, Tim Rocktäschel, Luke Zettlemoyer, Victor Zhong</i>	
Learning Energy Networks with Generalized Fenchel-Young Losses	12516
<i>Mathieu Blondel, Robert Dadashi, Matthieu Geist, Leonard Hussenot, Felipe Llinares-Lopez</i>	
Interaction-Grounded Learning with Action-Inclusive Feedback	12529
<i>Dylan J Foster, Nan Jiang, John Langford, Paul Mineiro, Lekan Molu, Ida Momennejad, Akanksha Saran, Tengyang Xie</i>	
Making Look-Ahead Active Learning Strategies Feasible with Neural Tangent Kernels	12542
<i>Wonho Bae, Mohamad Amin Mohamadi, Danica J. Sutherland</i>	
The Neural Testbed: Evaluating Joint Predictions.....	12554
<i>Seyed Mohammad Asghari, Vikranth Dwaracherla, Botao Hao, Morteza Ibrahimi, Dieterich Lawson, Xiuyuan Lu, Brendan O'Donoghue, Ian Osband, Benjamin Van Roy, Zheng Wen</i>	
Learning Dense Object Descriptors from Multiple Views for Low-Shot Category Generalization	12566
<i>Zixuan Huang, James Rehg, Stefan Stojanov, Anh Thai</i>	
Monte Carlo Tree Descent for Black-Box Optimization	12581
<i>Sicun Gao, Yaoguang Zhai</i>	
Teacher Forcing Recovers Reward Functions for Text Generation	12594
<i>Yongchang Hao, Yuxin Liu, Lili Mou</i>	
Masked Autoencoding for Scalable and Generalizable Decision Making.....	12608
<i>Pieter Abbeel, Aditya Grover, Fangchen Liu, Hao Liu</i>	
Distributional Reward Estimation for Effective Multi-Agent Deep Reinforcement Learning	12619
<i>Yi Chang, Hechang Chen, Jifeng Hu, Sili Huang, Haiyin Piao, Lichao Sun, Yanchao Sun</i>	
Sampling from Log-Concave Distributions with Infinity-Distance Guarantees.....	12633
<i>Oren Mangoubi, Nisheeth Vishnoi</i>	
ELASTIC: Numerical Reasoning with Adaptive Symbolic Compiler	12647
<i>Yashar Moshfeghi, Jiaxin Zhang</i>	

Training Spiking Neural Networks with Local Tandem Learning.....	12662
<i>Yansong Chua, Haizhou Li, Xinchao Wang, Jibin Wu, Qu Yang, Malu Zhang</i>	
FiLM: Frequency Improved Legendre Memory Model for Long-Term Time Series Forecasting	12677
<i>Rong Jin, Ziqing Ma, Liang Sun, Xue Wang, Qingsong Wen, Tao Yao, Wotao Yin, Tian Zhou</i>	
Differentially Private Linear Sketches: Efficient Implementations and Applications	12691
<i>Divyakant Agrawal, Amr El Abbadi, Dan Qiao, Rachel Redberg, Yu-Xiang Wang, Fuheng Zhao</i>	
Coordinates Are NOT Lonely - Codebook Prior Helps Implicit Neural 3D Representations	12705
<i>Tao Chen, Pei Cheng, Zilong Huang, Wen Liu, Fukun Yin, Gang Yu</i>	
Scalable Neural Video Representations with Learnable Positional Features.....	12718
<i>Subin Kim, Jaeho Lee, Jinwoo Shin, Sihyun Yu</i>	
Data Augmentation MCMC for Bayesian Inference from Privatized Data	12732
<i>Jordan Awan, Ruobin Gong, Nianqiao Ju, Vinayak Rao</i>	
Neural Circuit Architectural Priors for Embodied Control.....	12744
<i>Nikhil Bhattasali, Tatiana Engel, Anthony M Zador</i>	
Bayesian Optimization Over Discrete and Mixed Spaces Via Probabilistic Reparameterization.....	12760
<i>Eytan Bakshy, Maximilian Balandat, Samuel Daulton, David Eriksson, Michael A Osborne, Xingchen Wan</i>	
Learning to Constrain Policy Optimization with Virtual Trust Region	12775
<i>Majid Abdolshah, Kien Do, Sunil Gupta, Thommen Karimpanal George, Thai Hung Le, Dung Nguyen, Svetha Venkatesh</i>	
Verification and Search Algorithms for Causal DAGs	12787
<i>Arnab Bhattacharyya, Davin Choo, Kirankumar Shiragur</i>	
Causally Motivated Multi-Shortcut Identification and Removal.....	12800
<i>Maggie Makar, Jiayun Zheng</i>	
Avalon: A Benchmark for RL Generalization Using Procedurally Generated Worlds	12813
<i>Joshua Albrecht, Abraham Fetterman, Bryden Fogelman, Ellie Kitanidis, Maksis Knutins, Zack Polizzi, Kanjun Qiu, Michael Rosenthal, Nicole Seo, James Simon, Bartosz Wróblewski</i>	
Learning Equivariant Segmentation with Instance-Unique Querying.....	12826
<i>James Liang, Dongfang Liu, Wenguan Wang</i>	
NeuPhysics: Editable Neural Geometry and Physics from Monocular Videos	12841
<i>Alexander Gao, Ming Lin, Yi-Ling Qiao</i>	
Measuring Data Reconstruction Defenses in Collaborative Inference Systems.....	12855
<i>Hongxin Hu, Ziang Li, Ao Ren, Juan Wang, Xiaoyang Xu, Mengda Yang, Wenzhe Yi</i>	
ZARTS: On Zero-Order Optimization for Neural Architecture Search.....	12868
<i>Wenxuan Guo, Jianlin Su, Xiaoxing Wang, Junchi Yan, Xiaokang Yang</i>	
Make Some Noise: Reliable and Efficient Single-Step Adversarial Training	12881
<i>Adel Bibi, Pau De Jorge Aranda, Puneet Dokania, Gregory Rogez, Amartya Sanyal, Philip Torr, Riccardo Volpi</i>	

Structural Pruning Via Latency-Saliency Knapsack.....	12894
<i>Jose M. Alvarez, Jianna Liu, Lei Mao, Pavlo Molchanov, Maying Shen, Hongxu Yin</i>	
Risk Bounds of Multi-Pass SGD for Least Squares in the Interpolation Regime	12909
<i>Vladimir Braverman, Quanquan Gu, Sham Kakade, Jingfeng Wu, Difan Zou</i>	
Decision-Based Black-Box Attack Against Vision Transformers Via Patch-Wise Adversarial Removal	12921
<i>Yahong Han, Xiaohui Kuang, Yucheng Shi, Yu-An Tan</i>	
EfficientFormer: Vision Transformers at MobileNet Speed.....	12934
<i>Georgios Evangelidis, Ju Hu, Yanyu Li, Jian Ren, Sergey Tulyakov, Yanzhi Wang, Yang Wen, Geng Yuan</i>	
Holomorphic Equilibrium Propagation Computes Exact Gradients Through Finite Size Oscillations	12950
<i>Axel Laborieux, Friedemann Zenke</i>	
Learning Substructure Invariance for Out-Of-Distribution Molecular Representations	12964
<i>Xiaosong Jia, Qitian Wu, Junchi Yan, Nianzu Yang, Kaipeng Zeng</i>	
The Dollar Street Dataset: Images Representing the Geographic and Socioeconomic Diversity of the World.....	12979
<i>Cody Coleman, Sudnya Damos, William Gaviria Rojas, Vijay Janapa Reddi, David Kanter, Keertan Kini</i>	
LST: Ladder Side-Tuning for Parameter and Memory Efficient Transfer Learning.....	12991
<i>Mohit Bansal, Jaemin Cho, Yi-Lin Sung</i>	
CGLB: Benchmark Tasks for Continual Graph Learning	13006
<i>Dongjin Song, Dacheng Tao, Xikun Zhang</i>	
Pre-Trained Image Encoder for Generalizable Visual Reinforcement Learning	13022
<i>Yang Gao, Xueqian Wang, Yi Wu, Huazhe Xu, Zhengrong Xue, Bo Yuan, Zhecheng Yuan</i>	
Amortized Inference for Heterogeneous Reconstruction in Cryo-EM	13038
<i>Axel Levy, Julien N. P Martel, Frederic Poitevin, Gordon Wetzstein, Ellen Zhong</i>	
RKHS-SHAP: Shapley Values for Kernel Methods	13050
<i>Siu Lun Chau, Javier González, Robert Hu, Dino Sejdinovic</i>	
Single-Stage Visual Relationship Learning Using Conditional Queries.....	13064
<i>Alakh Desai, Subarna Tripathi, Nuno Vasconcelos, Tz-Ying Wu</i>	
Unlabelled Sample Compression Schemes for Intersection-Closed Classes and Extremal Classes.....	13078
<i>Benjamin Rubinstein, Joachim Rubinstein</i>	
VCT: A Video Compression Transformer.....	13091
<i>Eirikur Agustsson, Sergi Caelles, Sung Jin Hwang, Mario Lucic, Fabian Mentzer, David Minnen, George D Toderici</i>	
Amortized Inference for Causal Structure Learning.....	13104
<i>Andreas Krause, Lars Lorch, Jonas Rothfuss, Bernhard Schölkopf, Scott Sussex</i>	
Emergent Graphical Conventions in a Visual Communication Game.....	13119
<i>Lifeng Fan, Tao Gao, Jungseock Joo, Shuwen Qiu, Sirui Xie, Song-Chun Zhu, Yixin Zhu</i>	

VOLUME 18

Repairing Neural Networks by Leaving the Right Past Behind	13132
<i>Melanie F. Pradier, Yingzhen Li, Aditya Nori, Ryutaro Tanno</i>	
Selective Compression Learning of Latent Representations for Variable-Rate Image Compression	13146
<i>Seyoon Jeong, Munchurl Kim, Jooyoung Lee</i>	
Multi-LexSum: Real-World Summaries of Civil Rights Lawsuits at Multiple Granularities.....	13158
<i>Nathan Dahlberg, Doug Downey, Kyle Lo, Margo Schlanger, Zejiang Shen, Lauren Yu</i>	
Increasing Confidence in Adversarial Robustness Evaluations.....	13174
<i>Wieland Brendel, Nicholas Carlini, Florian Tramer, Roland S. Zimmermann</i>	
Local Bayesian Optimization Via Maximizing Probability of Descent.....	13190
<i>Jacob Gardner, Roman Garnett, Quan Nguyen, Kaiwen Wu</i>	
Staircase Attention for Recurrent Processing of Sequences	13203
<i>Da Ju, Stephen Roller, Sainbayar Sukhbaatar, Jason E Weston</i>	
Causal Inference with Non-IID Data Using Linear Graphical Models	13214
<i>Karthika Mohan, Judea Pearl, Chi Zhang</i>	
Data-Driven Offline Decision-Making Via Invariant Representation Learning.....	13226
<i>Aviral Kumar, Sergey Levine, Han Qi, Yi Su</i>	
Untargeted Backdoor Watermark: Towards Harmless and Stealthy Dataset Copyright Protection.....	13238
<i>Yang Bai, Yong Jiang, Bo Li, Yiming Li, Shu-Tao Xia, Yong Yang</i>	
Learning Symmetric Rules with SATNet	13251
<i>Sangho Lim, Eun-Gyeol Oh, Hongseok Yang</i>	
The Privacy Onion Effect: Memorization is Relative.....	13263
<i>Nicholas Carlini, Matthew Jagielski, Nicolas Papernot, Andreas Terzis, Florian Tramer, Chiyuan Zhang</i>	
Langevin Autoencoders for Learning Deep Latent Variable Models.....	13277
<i>Yusuke Iwasawa, Wataru Kumagai, Yutaka Matsuo, Shohei Taniguchi</i>	
Envy-Free Policy Teaching to Multiple Agents.....	13290
<i>Jiarui Gan, R Majumdar, Goran Radanovic, Adish Singla</i>	
Provably Efficient Model-Free Constrained RL with Linear Function Approximation	13303
<i>Arnob Ghosh, Ness Shroff, Xingyu Zhou</i>	
Towards Optimal Communication Complexity in Distributed Non-Convex Optimization.....	13316
<i>Brian Bullins, Kumar Kshitij Patel, Nati Srebro, Lingxiao Wang, Blake Woodworth</i>	
Eliciting Thinking Hierarchy Without a Prior	13329
<i>Zhihuan Huang, Yuqing Kong, Yunqi Li, Jinzhao Wu, Yubo Zhang</i>	
Fair and Efficient Allocations Without Obvious Manipulations.....	13342
<i>Alexandros Psomas, Paritosh Verma</i>	
Bridge the Gap Between Architecture Spaces Via a Cross-Domain Predictor	13355
<i>Yuqiao Liu, Zeqiong Lv, Yanan Sun, Yehui Tang, Yunhe Wang</i>	

Shield Decentralization for Safe Multi-Agent Reinforcement Learning	13367
<i>Christopher Amato, Daniel Melcer, Stavros Tripakis</i>	
Domain Generalization Without Excess Empirical Risk	13380
<i>Vladlen Koltun, Ozan Sener</i>	
Star Temporal Classification: Sequence Modeling with Partially Labeled Data	13392
<i>Ronan Collobert, Awni Hannun, Vineel Prapat, Gabriel Synnaeve</i>	
Signal Processing for Implicit Neural Representations	13404
<i>Zhiwen Fan, Yifan Jiang, Peihao Wang, Zhangyang Wang, DeJia Xu</i>	
Fault-Aware Neural Code Rankers	13419
<i>Andres Cudas, Mark Encarnación, Jianfeng Gao, Jeevana Priya Inala, Shuvendu Lahiri, Madanlal Musuvathi, Chenglong Wang, Mei Yang</i>	
Prompt Certified Machine Unlearning with Randomized Gradient Smoothing and Quantization	13433
<i>Tianshi Che, Lingjuan Lyu, Zijie Zhang, Xin Zhao, Yang Zhou</i>	
What Makes a "Good" Data Augmentation in Knowledge Distillation - A Statistical Perspective	13456
<i>Yun Fu, Michael Jones, Suhas Lohit, Huan Wang</i>	
Supervising the Multi-Fidelity Race of Hyperparameter Configurations	13470
<i>Josif Grabocka, Arlind Kadra, Martin Wistuba</i>	
Streaming Radiance Fields for 3D Video Synthesis	13485
<i>Lingzhi Li, Li Shen, Zhen Shen, Ping Tan, Zhongshu Wang</i>	
Byzantine-Tolerant Federated Gaussian Process Regression for Streaming Data	13499
<i>Zhenyuan Yuan, Xu Zhang, Minghui Zhu</i>	
Neural Matching Fields: Implicit Representation of Matching Fields for Visual Correspondence	13512
<i>Seokju Cho, Sunghwan Hong, Susung Hong, Sangryul Jeon, Seungryong Kim, Dongbo Min, Jisu Nam</i>	
Consistency of Constrained Spectral Clustering Under Graph Induced Fair Planted Partitions	13527
<i>Ambedkar Dukkipati, Shubham Gupta</i>	
Local Spatiotemporal Representation Learning for Longitudinally-Consistent Neuroimage Analysis	13541
<i>Kelly Botteron, Neel Dey, Guido Gerig, Mengwei Ren, Martin Styner</i>	
Overparameterization from Computational Constraints	13557
<i>Sanjam Garg, Somesh Jha, Saeed Mahloujifar, Mohammad Mahmoody, Mingyuan Wang</i>	
A Unifying Framework of Off-Policy General Value Function Evaluation	13570
<i>Yingbin Liang, Zhaoran Wang, Tengyu Xu, Zhuoran Yang</i>	
Anchor-Changing Regularized Natural Policy Gradient for Multi-Objective Reinforcement Learning	13584
<i>Dileep Kalathil, P. R. Kumar, Tao Liu, Chao Tian, Ruida Zhou</i>	
Do Current Multi-Task Optimization Methods in Deep Learning Even Help?	13597
<i>Orhan Firat, Ankush Garg, Behrooz Ghorbani, Justin Gilmer, Derrick Xin</i>	
TAP-Vid: A Benchmark for Tracking Any Point in a Video	13610
<i>Yusuf Aytar, Joao Carreira, Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens, Lucas Smaira, Yi Yang, Andrew Zisserman</i>	

Alignment-Guided Temporal Attention for Video Action Recognition	13627
<i>Xun Guo, Zhenyang Li, Yan Lu, Yizhou Zhao</i>	
Meta Reinforcement Learning with Finite Training Tasks - A Density Estimation Approach	13640
<i>Gilad Adler, Zohar Rimon, Aviv Tamar</i>	
TotalSelfScan: Learning Full-Body Avatars from Self-Portrait Videos of Faces, Hands, and Bodies.....	13654
<i>Hujun Bao, Junting Dong, Qi Fang, Yudong Guo, Sida Peng, Qing Shuai, Xiaowei Zhou</i>	
DHRL: A Graph-Based Approach for Long-Horizon and Sparse Hierarchical Reinforcement Learning	13668
<i>Inkyu Jang, H. Jin Kim, Jigang Kim, Seungjae Lee</i>	
A Classification of G-Invariant Shallow Neural Networks	13679
<i>Devanshu Agrawal, James Ostrowski</i>	
A Conditional Randomization Test for Sparse Logistic Regression in High-Dimension	13691
<i>Sylvain Arlot, Binh T. Nguyen, Bertrand Thirion</i>	
Biologically-Plausible Determinant Maximization Neural Networks for Blind Separation of Correlated Sources.....	13704
<i>Bariscan Bozkurt, Alper Erdogan, Cengiz Pehlevan</i>	
Generalized One-Shot Domain Adaptation of Generative Adversarial Networks.....	13718
<i>Tiande Guo, Congying Han, Yinglu Liu, Tao Mei, Ting Yao, Zicheng Zhang</i>	
Representing Spatial Trajectories as Distributions	13731
<i>Didac Suris, Carl Vondrick</i>	
EPIC-KITCHENS VISOR Benchmark: Video Segmentations and Object Relations.....	13745
<i>Dima Damen, Ahmad Darkhalil, Sanja Fidler, David Fouhey, Richard Higgins, Amlan Kar, Jian Ma, Dandan Shan, Bin Zhu</i>	
Inverse Design for Fluid-Structure Interactions Using Graph Network Simulators.....	13759
<i>Kelsey Allen, Peter Battaglia, Jessica Hamrick, Tatiana Lopez-Guevara, Tobias Pfaff, Alvaro Sanchez Gonzalez, Kimberly Stachenfeld</i>	
What Makes Graph Neural Networks Miscalibrated?	13775
<i>Daniel Cremers, Hans Hao-Hsun Hsu, Yuesong Shen, Christian Tomani</i>	
Stochastic Adaptive Activation Function.....	13787
<i>Jae Youn Hwang, Haeyun Lee, Kyungsu Lee, Jaeseung Yang</i>	
ActionSense: A Multimodal Dataset and Recording Framework for Human Activities Using Wearable Sensors in a Kitchen Environment	13800
<i>Joseph Delpreto, Michael Foshey, Yunzhu Li, Chao Liu, Yiyue Luo, Wojciech Matusik, Daniela Rus, Antonio Torralba</i>	
Video Compression Dataset and Benchmark of Learning-Based Video-Quality Metrics	13814
<i>Anastasia Antsiferova, Aleksandr Gushchin, Dmitriy Kulikov, Sergey Lavrushkin, Maksim Smirnov, Dmitriy Vatolin</i>	
Scalable Distributional Robustness in a Class of Non-Convex Optimization with Guarantees	13826
<i>Avinandan Bose, Tien Mai, Arunesh Sinha</i>	
Prototypical VoteNet for Few-Shot 3D Point Cloud Object Detection.....	13838
<i>Xiaojuan Qi, Shizhen Zhao</i>	

Learning from a Sample in Online Algorithms	13852
<i>C. J. Argue, Alan Frieze, Anupam Gupta, Christopher Seiler</i>	
Robust Rent Division	13864
<i>Dominik Peters, Ariel Procaccia, David Zhu</i>	
Efficient Dataset Distillation Using Random Feature Approximation	13877
<i>Alexander Amini, Ramin Hasani, Noel Loo, Daniela Rus</i>	
Scalable Sensitivity and Uncertainty Analyses for Causal-Effect Estimates of Continuous-Valued Interventions	13892
<i>Alyson Douglas, Yarin Gal, Andrew Jesson, Peter Manshausen, Nicolai Meinshausen, Uri Shalit, Maëlys Solal, Philip Stier</i>	

VOLUME 19

Sparse Interaction Additive Networks Via Feature Interaction Detection and Sparse Selection	13908
<i>James Enouen, Yan Liu</i>	
Faster Stochastic Algorithms for Minimax Optimization Under Polyak-Lojasiewicz Condition.....	13921
<i>Lesi Chen, Luo Luo, Boyuan Yao</i>	
Kantorovich Strikes Back! Wasserstein GANs Are Not Optimal Transport?	13933
<i>Evgeny Burnaev, Alexander Kolesov, Alexander Korotin</i>	
List-Decodable Sparse Mean Estimation Via Difference-Of-Pairs Filtering.....	13947
<i>Ilias Diakonikolas, Daniel Kane, Sushrut Karmalkar, Ankit Pensia, Thanasis Pittas</i>	
Distributional Convergence of the Sliced Wasserstein Process	13961
<i>Jonathan Niles-Weed, Jiaqi Xi</i>	
VTC-LFC: Vision Transformer Compression with Low-Frequency Components.....	13974
<i>Feng Ding, Hao Li, Hao Luo, Fan Wang, Zhenyu Wang, Pichao Wang</i>	
A Theory of PAC Learnability Under Transformation Invariances	13989
<i>Avrim Blum, Omar Montasser, Han Shao</i>	
PALBERT: Teaching ALBERT to Ponder	14002
<i>Nikita Balagansky, Daniil Gavrilov</i>	
Distributed Methods with Compressed Communication for Solving Variational Inequalities, with Theoretical Guarantees.....	14013
<i>Aleksandr Beznosikov, Michael Diskin, Alexander Gasnikov, Peter Richtarik, Max Ryabinin</i>	
Analyzing Data-Centric Properties for Graph Contrastive Learning.....	14030
<i>Mark Heimann, Danai Koutra, Ekdeep S Lubana, Jayaraman Thiagarajan, Puja Trivedi</i>	
Learning Bipartite Graphs: Heavy Tails and Multiple Components.....	14044
<i>José Vinícius De Miranda Cardoso, Daniel Palomar, Jiayi Ying</i>	
Training Scale-Invariant Neural Networks on the Sphere Can Happen in Three Regimes	14058
<i>Maxim Kodryan, Ekaterina Lobacheva, Maksim Nakhodnov, Dmitry Vetrov</i>	
Exploring the Whole Rashomon Set of Sparse Decision Trees	14071
<i>Zhi Chen, Cynthia Rudin, Margo Seltzer, Takuya Takagi, Rui Xin, Chudi Zhong</i>	

Graph Self-Supervised Learning with Accurate Discrepancy Learning.....	14085
<i>Jinheon Baek, Sung Ju Hwang, Dongki Kim</i>	
Multi-Scale Adaptive Network for Single Image Denoising.....	14099
<i>Yuanbiao Gou, Peng Hu, Jiancheng Lv, Xi Peng, Joey Tianyi Zhou</i>	
A Theory of Weight Distribution-Constrained Learning	14113
<i>Daniel Lee, Haim Sompolinsky, Ben Sorscher, Weishun Zhong</i>	
Hierarchical Normalization for Robust Monocular Depth Estimation	14128
<i>Bin Fu, Chunhua Shen, Billz Wang, Wei Yin, Gang Yu, Chi Zhang</i>	
Deep Compression of Pre-Trained Transformer Models.....	14140
<i>Leland Chang, Chia-Yu Chen, Kaoutar El Maghraoui, Chi-Chun Charlie Liu, Sanchari Sen, Vijayalakshmi Viji Srinivasan, Swagath Venkataramani, Naigang Wang</i>	
Constrained Predictive Coding as a Biologically Plausible Model of the Cortical Hierarchy.....	14155
<i>Yanis Bahroun, Dmitri Chklovskii, Stavash Golkar, Anirvan Sengupta, Tiberiu Tesileanu</i>	
Washing the Unwashable : On the (Im)possibility of Fairwashing Detection.....	14170
<i>Ulrich Aïvodji, Aisha Alaagib, Natalie Dullerud, Sébastien Gambs, Nicolas Papernot, Ali Shahin Shamsabadi, Sierra Wylie, Mohammad Yaghini</i>	
Near-Optimal Collaborative Learning in Bandits.....	14183
<i>Emilie Kaufmann, Clémence Réda, Sattar Vakili</i>	
Hypothesis Testing for Differentially Private Linear Regression	14196
<i>Daniel Alabi, Salil Vadhan</i>	
Bayesian Optimistic Optimization: Optimistic Exploration for Model-Based Reinforcement Learning	14210
<i>Tianci Li, Chenyang Wu, Yang Yu, Zongzhang Zhang</i>	
TokenMixup: Efficient Attention-Guided Token-Level Data Augmentation for Transformers.....	14224
<i>Hyeong Kyu Choi, Joonmyung Choi, Hyunwoo Kim</i>	
Error Analysis of Tensor-Train Cross Approximation	14236
<i>Zhexuan Gong, Alexander Lidiak, Zhen Qin, Gongguo Tang, Michael B Wakin, Zhihui Zhu</i>	
Approximate Secular Equations for the Cubic Regularization Subproblem	14250
<i>Yihang Gao, Michael Ng, Man-Chung Yue</i>	
On the Theoretical Properties of Noise Correlation in Stochastic Optimization	14261
<i>Francis Bach, Hans Kersting, Aurelien Lucchi, Antonio Orvieto, Frank Proske</i>	
Test-Time Prompt Tuning for Zero-Shot Generalization in Vision-Language Models.....	14274
<i>Anima Anandkumar, Tom Goldstein, De-An Huang, Weili Nie, Manli Shu, Chaowei Xiao, Zhiding Yu</i>	
SemMAE: Semantic-Guided Masking for Learning Masked Autoencoders.....	14290
<i>Gang Li, Daqing Liu, Bing Su, Chaoyue Wang, Changwen Zheng, Heliang Zheng</i>	
BiT: Robustly Binarized Multi-Distilled Transformer.....	14303
<i>Raghuraman Krishnamoorthi, Meng Li, Zechun Liu, Yashar Mehdad, Barlas Oguz, Aasish Pappu, Lin Xiao, Scott Yih</i>	

Perturbation Learning Based Anomaly Detection	14317
<i>Jinyu Cai, Jicong Fan</i>	
Knowledge-Aware Bayesian Deep Topic Model.....	14331
<i>Bo Chen, Zhibin Duan, Miaoge Li, Chaojie Wang, Dongsheng Wang, Yi. Shi Xu, Mingyuan Zhou</i>	
SelecMix: Debiased Learning by Contradicting-Pair Sampling.....	14345
<i>Inwoo Hwang, Jin-Hwa Kim, Yunhyeok Kwak, Sangjun Lee, Seong Joon Oh, Damien Teney, Byoung-Tak Zhang</i>	
Are All Frames Equal? Active Sparse Labeling for Video Action Detection	14358
<i>Aayush Rana, Yogesh Rawat</i>	
Unsupervised Domain Adaptation for Semantic Segmentation Using Depth Distribution	14374
<i>Huajun Liu, Quanliang Wu</i>	
P2P: Tuning Pre-Trained Image Models for Point Cloud Analysis with Point-To-Pixel Prompting	14388
<i>Jiwen Lu, Yongming Rao, Ziyi Wang, Xumin Yu, Jie Zhou</i>	
Finding Differences Between Transformers and ConvNets Using Counterfactual Simulation Testing	14403
<i>Sarah Bargal, Nataniel Ruiz, Kate Saenko, Stan Sclaroff, Cihang Xie</i>	
In the Eye of the Beholder: Robust Prediction with Causal User Modeling	14419
<i>Amir Feder, Guy Horowitz, Roi Reichart, Nir Rosenfeld, Yoav Wald</i>	
Variational Inference Via Wasserstein Gradient Flows.....	14434
<i>Francis Bach, Silvère Bonnabel, Sinho Chewi, Marc Lambert, Philippe Rigollet</i>	
projUNN: Efficient Method for Training Deep Networks with Unitary Matrices.....	14448
<i>Randall Balestriero, Bobak Kiani, Yann Lecun, Seth Lloyd</i>	
Reinforcement Learning in a Birth and Death Process: Breaking the Dependence on the State Space	14464
<i>Jonatha Anselmi, Bruno Gaujal, Louis-Sébastien Rebuffi</i>	
Multi-Dataset Training of Transformers for Robust Action Recognition	14475
<i>Junwei Liang, Chunhua Shen, Jun Zhang, Enwei Zhang</i>	
An Embarrassingly Simple Approach to Semi-Supervised Few-Shot Learning	14489
<i>Yuxin Peng, Xiu-Shen Wei, H.-Y. Xu, Faen Zhang, Wei Zhou</i>	
Recipe for a General, Powerful, Scalable Graph Transformer	14501
<i>Dominique Beaini, Vijay Prakash Dwivedi, Michael Galkin, Anh Tuan Luu, Ladislav Rampásek, Guy Wolf</i>	
ALIFE: Adaptive Logit Regularizer and Feature Replay for Incremental Semantic Segmentation	14516
<i>Donghyeon Baek, Bumsub Ham, Youngmin Oh</i>	
Rare Gems: Finding Lottery Tickets at Initialization	14529
<i>Matthew Grinde, Kangwook Lee, Alliot Nagle, Dimitris Papailiopoulos, Jy-Yong Sohn, Kartik Sreenivasan, Hongyi Wang, Eric Xing, Liu Yang</i>	
Fast Vision Transformers with HiLo Attention.....	14541
<i>Jianfei Cai, Zizheng Pan, Bohan Zhuang</i>	

Online Bipartite Matching with Advice: Tight Robustness-Consistency Tradeoffs for the Two-Stage Model	14555
<i>Billy Jin, Will Ma</i>	
Identifying Good Directions to Escape the NTK Regime and Efficiently Learn Low-Degree Plus Sparse Polynomials	14568
<i>Yu Bai, Jason Lee, Eshaan Nichani</i>	
Pure Transformers Are Powerful Graph Learners	14582
<i>Sungjun Cho, Seunghoon Hong, Jinwoo Kim, Honglak Lee, Moontae Lee, Seonwoo Min, Dat Nguyen</i>	
Orthogonal Transformer: An Efficient Vision Transformer Backbone with Token Orthogonalization	14596
<i>Ran He, Huaibo Huang, Xiaoqiang Zhou</i>	
Using Mixup as a Regularizer Can Surprisingly Improve Accuracy & Out-Of-Distribution Robustness.....	14608
<i>Puneet Dokania, Ser Nam Lim, Francesco Pinto, Philip Torr, Harry Yang</i>	
NeurOLight: A Physics-Agnostic Neural Operator Enabling Parametric Photonic Device Simulation	14623
<i>Duane Boning, Ray Chen, Chenghao Feng, Zhengqi Gao, Jiaqi Gu, David Pan, Hanqing Zhu</i>	
Learning the Structure of Large Networked Systems Obeying Conservation Laws	14637
<i>Rajasekhar Anguluri, Gautam Dasarathy, Anirudh Rayas</i>	
FP8 Quantization: The Power of the Exponent	14651
<i>Tijmen Blankevoort, Andrey Kuzmin, Markus Nagel, Jorn Peters, Yuwei Ren, Mart Van Baalen</i>	
Bridging the Gap Between Vision Transformers and Convolutional Neural Networks on Small Datasets	14663
<i>Chuanbin Liu, Zhiying Lu, Hongtao Xie, Yongdong Zhang</i>	
Private Set Generation with Discriminative Information	14678
<i>Dingfan Chen, Mario Fritz, Raouf Kerkouche</i>	
The First Optimal Algorithm for Smooth and Strongly-Convex-Strongly-Concave Minimax Optimization.....	14691
<i>Alexander Gasnikov, Dmitry Kovalev</i>	

VOLUME 20

Provable Defense Against Backdoor Policies in Reinforcement Learning.....	14704
<i>Shubham Bharti, Adish Singla, Xuezhou Zhang, Jerry Zhu</i>	
Diffusion Models as Plug-And-Play Priors	14715
<i>Alexandros Graikos, Nebojsa Jojic, Nikolay Malkin, Dimitris Samaras</i>	
CoupAlign: Coupling Word-Pixel with Sentence-Mask Alignments for Referring Image Segmentation	14729
<i>Wei Ke, Xiaodan Liang, Jianzhuang Liu, Zicheng Zhang, Yi Zhu</i>	
Saliency-Aware Neural Architecture Search	14743
<i>Ramtin Hosseini, Pengtao Xie</i>	

VaiPhy: A Variational Inference Based Algorithm for Phylogeny	14758
<i>Hazal Koptagel, Oskar Kviman, Jens Lagergren, Harald Melin, Negar Safinianaini</i>	
Repeated Augmented Rehearsal: A Simple but Strong Baseline for Online Continual Learning.....	14771
<i>Albert Bifet, Eibe Frank, Yunzhe Jia, Nick Jin Sean Lim, Bernhard Pfahringer, Yaqian Zhang</i>	
You Only Live Once: Single-Life Reinforcement Learning.....	14784
<i>Annie Chen, Chelsea Finn, Sergey Levine, Archit Sharma</i>	
Compressible-Composable NeRF Via Rank-Residual Decomposition	14798
<i>Xiaokang Chen, Jiaxiang Tang, Jingbo Wang, Gang Zeng</i>	
Data-Efficient Pipeline for Offline Reinforcement Learning with Limited Data	14810
<i>Emma Brunskill, Yannis Flet-Berliac, Deon Jordan, Allen Nie, William Steenbergen</i>	
Hardness in Markov Decision Processes: Theory and Practice.....	14824
<i>Michelangelo Conserva, Paulo Rauber</i>	
Injecting Domain Knowledge from Empirical Interatomic Potentials to Neural Networks for Predicting Material Properties	14839
<i>Daniel Karls, George Karypis, Ilya Nikiforov, Zeren Shui, Ellad Tadmor, Mingjian Wen</i>	
Learning Modular Simulations for Homogeneous Systems	14852
<i>Jayesh Gupta, Ashish Kapoor, Sai Vemprala</i>	
On Gap-Dependent Bounds for Offline Reinforcement Learning.....	14865
<i>Qiwen Cui, Simon Du, Xinqi Wang</i>	
Semi-Discrete Normalizing Flows Through Differentiable Tessellation.....	14878
<i>Brandon Amos, Ricky T. Q. Chen, Maximilian Nickel</i>	
Don't Pour Cereal into Coffee: Differentiable Temporal Logic for Temporal Action Segmentation.....	14890
<i>Mohan Kankanhalli, Yogesh Rawat, Mubarak Shah, Yongkang Wong, Ziwei Xu</i>	
Batch-Size Independent Regret Bounds for Combinatorial Semi-Bandits with Probabilistically Triggered Arms Or Independent Arms	14904
<i>Wei Chen, Carlee Joe-Wong, Xutong Liu, John C. S. Lui, Siwei Wang, Jinhang Zuo</i>	
Less-Forgetting Multi-Lingual Fine-Tuning.....	14917
<i>Lu Chen, Nan Duan, Yunjun Gao, Yaobo Liang, Yuren Mao, Haobo Wang, Kai Wang</i>	
Adversarial Attack on Attackers: Post-Process to Mitigate Black-Box Score-Based Query Attacks	14929
<i>Sizhe Chen, Xiaolin Huang, Zhehao Huang, Qinghua Tao, Yingwen Wu, Cihang Xie</i>	
Retaining Knowledge for Learning with Dynamic Definition	14944
<i>Benjamin Coleman, Zichang Liu, Anshumali Shrivastava, Tianyi Zhang</i>	
HUMANISE: Language-Conditioned Human Motion Generation in 3D Scenes	14959
<i>Yixin Chen, Siyuan Huang, Wei Liang, Tengyu Liu, Zan Wang, Yixin Zhu</i>	
Keypoint-Guided Optimal Transport with Applications in Heterogeneous Domain Adaptation.....	14972
<i>Xiang Gu, Jian Sun, Zongben Xu, Yucheng Yang, Wei Zeng</i>	
Sequence-To-Set Generative Models	14986
<i>Longtao Tang, Yu Yang, Ying Zhou</i>	

Near-Optimal Multi-Agent Learning for Safe Coverage Control.....	14998
<i>Andreas Krause, Manish Prajapat, Matteo Turchetta, Melanie Zeilinger</i>	
Network Change Point Localisation Under Local Differential Privacy	15013
<i>Tom Berrett, Mengchu Li, Yi Yu</i>	
Independence Testing for Bounded Degree Bayesian Networks.....	15027
<i>Arnab Bhattacharyya, Clément L Canonne, Qiping Yang</i>	
Beyond Spectral Gap: The Role of the Topology in Decentralized Learning.....	15039
<i>Hadrien Hendrikx, Martin Jaggi, Thijs Vogels</i>	
GriddlyJS: A Web IDE for Reinforcement Learning	15051
<i>Christopher Bamford, Minqi Jiang, Tim Rocktäschel, Mikayel Samvelyan</i>	
Periodic Graph Transformers for Crystal Material Property Prediction	15066
<i>Shuiwang Ji, Yuchao Lin, Yi Liu, Keqiang Yan</i>	
Decentralized, Communication- And Coordination-Free Learning in Structured Matching Markets.....	15081
<i>Chinmay Maheshwari, Eric Mazumdar, Shankar Sastry</i>	
Improving GANs with a Dynamic Discriminator.....	15093
<i>Bo Dai, Yujun Shen, Yinghao Xu, Ceyuan Yang, Deli Zhao, Bolei Zhou</i>	
Deliberated Domain Bridging for Domain Adaptive Semantic Segmentation	15105
<i>Kai Chen, Huaian Chen, Lin Chen, Xin Jin, Yi Jin, Zhixiang Wei, Miao Zheng</i>	
Learning to Share in Networked Multi-Agent Reinforcement Learning.....	15119
<i>Ge Li, Zongqing Lu, Yaowei Wang, Yuxuan Yi</i>	
Scalable and Efficient Non-Adaptive Deterministic Group Testing.....	15132
<i>Dariusz Kowalski, Dominik Pajak</i>	
Structure-Preserving 3D Garment Modeling with Neural Sewing Machines.....	15147
<i>Xipeng Chen, Xiaodan Liang, Liang Lin, Philip Torr, Guangrun Wang, Dizhong Zhu</i>	
Rethinking Variational Inference for Probabilistic Programs with Stochastic Support	15160
<i>Luke Ong, Thomas Rainforth, Tim Reichelt</i>	
Variance Reduced ProxSkip: Algorithm, Theory and Application to Federated Learning	15176
<i>Grigory Malinovsky, Peter Richtarik, Kai Yi</i>	
DreamShard: Generalizable Embedding Table Placement for Recommender Systems	15190
<i>Bhargav Bhushanam, Louis Feng, Xia Hu, Arun Kejariwal, Kwei-Herng Lai, Zirui Liu, Qiaoyu Tan, Yuandong Tian, Daochen Zha</i>	
Order-Invariant Cardinality Estimators Are Differentially Private.....	15204
<i>Charlie Dickens, Justin Thaler, Daniel Ting</i>	
NeuForm: Adaptive Overfitting for Neural Shape Editing.....	15217
<i>Paul Guerrero, Leonidas Guibas, Connor Lin, Niloy Mitra, Gordon Wetzstein</i>	
Inductive Logical Query Answering in Knowledge Graphs.....	15230
<i>Michael Galkin, Hongyu Ren, Jian Tang, Zhaocheng Zhu</i>	
Drawing Out of Distribution with Neuro-Symbolic Generative Models.....	15244
<i>Tuan Anh Le, Yichao Liang, Siddharth N, Josh Tenenbaum</i>	

Unsupervised Object Representation Learning Using Translation and Rotation Group Equivariant VAE	15255
<i>Tristan Bepler, Alireza Nasiri</i>	
PointTAD: Multi-Label Temporal Action Detection with Learnable Query Points.....	15268
<i>Bin Kang, Xintian Shi, Jing Tan, Limin Wang, Xiaotong Zhao</i>	
Unpacking Reward Shaping: Understanding the Benefits of Reward Engineering on Sample Complexity	15281
<i>Abhishek Gupta, Sham Kakade, Sergey Levine, Aldo Pacchiano, Yuexiang Zhai</i>	
LASSIE: Learning Articulated Shapes from Sparse Image Ensemble Via 3D Part Discovery	15296
<i>Wei-Chih Hung, Varun Jampani, Yuanzhen Li, Michael Rubinstein, Ming-Hsuan Yang, Chun-Han Yao</i>	
Retrieval-Augmented Diffusion Models	15309
<i>Andreas Blattmann, Jonas Müller, Kaan Oktay, Björn Ommer, Robin Rombach</i>	
Logical Credal Networks.....	15325
<i>Francisco Barahona, Debarun Bhattacharjya, Tian Gao, Alexander Gray, Radu Marinescu, Haifeng Qian, Ryan Riegel, Pravinda Sahu</i>	
Neural Set Function Extensions: Learning with Discrete Functions in High Dimensions	15338
<i>Stefanie Jegelka, Nikolaos Karalias, Andreas Loukas, Joshua Robinson</i>	
Minimax-Optimal Multi-Agent RL in Markov Games with a Generative Model.....	15353
<i>Yuxin Chen, Yuejie Chi, Gen Li, Yuting Wei</i>	
Bi-Directional Weakly Supervised Knowledge Distillation for Whole Slide Image Classification	15368
<i>Xiaoyuan Luo, Linhao Qu, Zhijian Song, Manning Wang</i>	
Private and Communication-Efficient Algorithms for Entropy Estimation.....	15382
<i>Gecia Bravo-Hermsdorff, Róbert Busa-Fekete, Mohammad Ghavamzadeh, Andres Munoz Medina, Umar Syed</i>	
PKD: General Distillation Framework for Object Detectors Via Pearson Correlation Coefficient.....	15394
<i>Weihan Cao, Anda Cheng, Jian Cheng, Ke Cheng, Jianfei Gao, Yifan Zhang</i>	
Off-Team Learning.....	15407
<i>Brandon Cui, Jakob Foerster, Hengyuan Hu, Andrei Lupu, Samuel Sokota</i>	
NUWA-Infinity: Autoregressive Over Autoregressive Generation for Infinite Visual Synthesis.....	15420
<i>Nan Duan, Yuejian Fang, Zhe Gan, Xiaowei Hu, Jian Liang, Zicheng Liu, Jianfeng Wang, Lijuan Wang, Chenfei Wu</i>	
SoftPatch: Unsupervised Anomaly Detection with Noisy Data.....	15433
<i>Xi Jiang, Yong Liu, Jianlin Liu, Qiang Nie, Chengjie Wang, Jinbao Wang, Kai Wu, Feng Zheng</i>	
Stability Analysis and Generalization Bounds of Adversarial Training	15446
<i>Yanbo Fan, Zhi-Quan Luo, Ruoyu Sun, Jue Wang, Jiancong Xiao</i>	
LasUIE: Unifying Information Extraction with Latent Adaptive Structure-Aware Generative Language Model.....	15460
<i>Tat-Seng Chua, Hao Fei, Fei Li, Bobo Li, Jingye Li, Libo Qin, Shengqiong Wu, Meishan Zhang, Min Zhang</i>	

STaR: Bootstrapping Reasoning with Reasoning 15476
Noah Goodman, Jesse Mu, Yuhuai Wu, Eric Zelikman

VOLUME 21

TPU-KNN: K Nearest Neighbor Search at Peak FLOP/s 15489
Felix Chern, Andy Davis, Ruiqi Guo, Blake Hechtman, Sanjiv Kumar, David Majnemer

Efficient Meta Reinforcement Learning for Preference-Based Fast Adaptation 15502
Yitao Liang, Anji Liu, Jianzhu Ma, Jian Peng, Zhizhou Ren

Weakly Supervised Representation Learning with Sparse Perturbations 15516
Kartik Ahuja, Yoshua Bengio, Jason Hartford

A Multi-Resolution Framework for U-Nets with Applications to Hierarchical VAEs 15529
Dominic Danks, George Deligiannidis, Arnaud Doucet, Fabian Falck, Chris C Holmes, Matthew Willetts, Christopher Williams, Christopher Yau

Watermarking for Out-Of-Distribution Detection 15545
Chen Gong, Bo Han, Feng Liu, Tongliang Liu, Qizhou Wang, Jing Zhang, Yonggang Zhang

K-LITE: Learning Transferable Visual Models with External Knowledge 15558
Trevor Darrell, Zhe Gan, Jianfeng Gao, Xiaowei Hu, Kurt Keutzer, Chunyuan Li, Ce Liu, Anna Rohrbach, Sheng Shen, Lijuan Wang, Yujia Xie, Jianwei Yang, Lu Yuan, Pengchuan Zhang

VeriDark: A Large-Scale Benchmark for Authorship Verification on the Dark Web 15574
Antonio Barbalau, Florin Brad, Radu Tudor Ionescu, Andrei Manolache, Marius Popescu

EHRSQL: A Practical Text-To-SQL Benchmark for Electronic Health Records 15589
Seongsu Bae, Edward Choi, Hyeonji Hwang, Jong-Yeup Kim, Yeonsu Kwon, Gyubok Lee, Minjoon Seo, Woncheol Shin, Seongjun Yang

Where Do Models Go Wrong? Parameter-Space Saliency Maps for Explainability 15602
Eitan Borgnia, Micah Goldblum, Tom Goldstein, Furong Huang, Roman Levin, Manli Shu

Using Embeddings for Causal Estimation of Peer Influence in Social Networks 15616
Irina Cristali, Victor Veitch

Collaborative Learning by Detecting Collaboration Partners 15629
Shu Ding, Wei Wang

Linear Label Ranking with Bounded Noise 15642
Dimitris Fotakis, Alkis Kalavasis, Vasilis Kontonis, Christos Tzamos

Quo Vadis: Is Trajectory Forecasting the Key Towards Long-Term Multi-Object Tracking? 15657
Patrick Dendorfer, Laura Leal-Taixé, Aljosa Osep, Vladimir Yugay

Wasserstein Iterative Networks for Barycenter Estimation 15672
Evgeny Burnaev, Vage Egiazarian, Alexander Korotin, Lingxiao Li

Identifiability of Deep Generative Models Without Auxiliary Information 15687
Bryon Aragam, Bohdan Kivva, Goutham Rajendran, Pradeep Ravikumar

Task Discovery: Finding the Tasks that Neural Networks Generalize on 15702
Andrei Atanov, Andrei Filatov, Ajay Sohmshty, Teresa Yeo, Amir Zamir

Iron: Private Inference on Transformers.....	15718
<i>Hanxiao Chen, Meng Hao, Hongwei Li, Pengzhi Xing, Guowen Xu, Tianwei Zhang</i>	
On Batch Teaching with Sample Complexity Bounded by VCD	15732
<i>Farnam Mansouri, Hans Simon, Adish Singla, Sandra Zilles</i>	
Approaching Quartic Convergence Rates for Quasi-Stochastic Approximation with Application to Gradient-Free Optimization.....	15743
<i>Caio Kalil Lauand, Sean Meyn</i>	
PAC: Assisted Value Factorization with Counterfactual Predictions in Multi-Agent Reinforcement Learning	15757
<i>Vaneet Aggarwal, Tian Lan, Hanhan Zhou</i>	
Translation-Equivariant Representation in Recurrent Networks with a Continuous Manifold of Attractors.....	15770
<i>Si Wu, Ying Nian Wu, Wenhao Zhang</i>	
OpenXAI: Towards a Transparent Evaluation of Model Explanations	15784
<i>Chirag Agarwal, Nari Johnson, Satyapriya Krishna, Himabindu Lakkaraju, Martin Pawelczyk, Isha Puri, Eshika Saxena, Marinka Zitnik</i>	
Better Best of Both Worlds Bounds for Bandits with Switching Costs.....	15800
<i>Idan Amir, Guy Azov, Tomer Koren, Roi Livni</i>	
Uncertainty-Aware Hierarchical Refinement for Incremental Implicitly-Refined Classification.....	15811
<i>Yang Cao, Jian Yang, Kecheng Zheng, Kai Zhu</i>	
The Hessian Screening Rule.....	15823
<i>Johan Larsson, Jonas Wallin</i>	
Sharp Analysis of Stochastic Optimization Under Global Kurdyka-Lojasiewicz Inequality	15836
<i>Jalal Etesami, Ilyas Fatkhullin, Niao He, Negar Kiyavash</i>	
Plan to Predict: Learning an Uncertainty-Foreseeing Model for Model-Based Reinforcement Learning	15849
<i>Chen Chen, Jianye Hao, Zifan Wu, Chao Yu, Hankz Hankui Zhuo</i>	
Co-Modality Graph Contrastive Learning for Imbalanced Node Classification	15862
<i>Yiyue Qian, Qianlong Wen, Yanfang Ye, Chunhui Zhang, Chuxu Zhang, Yiming Zhang</i>	
Efficiency Ordering of Stochastic Gradient Descent.....	15875
<i>Vishwaraj Doshi, Do-Young Eun, Jie Hu</i>	
Muffliato: Peer-To-Peer Privacy Amplification for Decentralized Optimization and Averaging	15889
<i>Aurélien Bellet, Edwige Cyffers, Mathieu Even, Laurent Massoulié</i>	
Mining Multi-Label Samples from Single Positive Labels	15903
<i>Youngin Cho, Jaegul Choo, Mohammad Azam Khan, Daejin Kim</i>	
Sequential Information Design: Learning to Persuade in the Dark	15917
<i>Martino Bernasconi, Matteo Castiglioni, Nicola Gatti, Alberto Marchesi, Francesco Trovò</i>	
AutoML Two-Sample Test	15929
<i>Simon Buchholz, Jonas M. Kübler, Krikamol Muandet, Bernhard Schölkopf, Vincent Stimper</i>	

Cross-Linked Unified Embedding for Cross-Modality Representation Learning	15942
<i>Zhi-Jie Cao, Xia Chenrui, Ge Gao, Sara Mostafavi, Xinming Tu</i>	
Learning Individualized Treatment Rules with Many Treatments: A Supervised Clustering Approach Using Adaptive Fusion.....	15956
<i>Yufeng Liu, Haixu Ma, Donglin Zeng</i>	
Robust Generalized Method of Moments: A Finite Sample Viewpoint.....	15970
<i>Dhruv Rohatgi, Vasilis Syrgkanis</i>	
k-Sliced Mutual Information: A Quantitative Study of Scalability with Dimension	15982
<i>Ziv Goldfeld, Kristjan Greenewald, Theshani Nuradha, Galen Reeves</i>	
What's the Harm? Sharp Bounds on the Fraction Negatively Affected by Treatment.....	15996
<i>Nathan Kallus</i>	
Bessel Equivariant Networks for Inversion of Transmission Effects in Multi-Mode Optical Fibres	16010
<i>Marco Aversa, Daniele Faccio, Simon Mekhail, Joshua Mitton, Roderick Murray-Smith, Miles Padgett</i>	
Training Subset Selection for Weak Supervision.....	16023
<i>Hunter Lang, David Sontag, Aravindan Vijayaraghavan</i>	
Expansion and Shrinkage of Localization for Weakly-Supervised Semantic Segmentation	16037
<i>Zequan Jie, Jinlong Li, Lin Ma, Xu Wang, Xiaolin Wei</i>	
No Free Lunch from Deep Learning in Neuroscience: A Case Study Through Models of the Entorhinal-Hippocampal Circuit	16052
<i>Ila Fiete, Mikail Khona, Rylan Schaeffer</i>	
On the Importance of Gradient Norm in PAC-Bayesian Bounds	16068
<i>Yossi Adi, Itai Gat, Tamir Hazan, Alex Schwing</i>	
RAMBO-RL: Robust Adversarial Model-Based Offline Reinforcement Learning.....	16082
<i>Nick Hawes, Bruno Lacerda, Marc Rigter</i>	
Learning NP-Hard Multi-Agent Assignment Planning Using GNN: Inference on a Random Graph and Provable Auction-Fitted Q-Learning	16098
<i>Hyunwook Kang, Taehwan Kwon, James R. Morrison, Jinkyoo Park</i>	
Improved Techniques for Deterministic l_2 Robustness.....	16110
<i>Soheil Feizi, Sahil Singla</i>	
Normalizing Flows for Knockoff-Free Controlled Feature Selection	16125
<i>Derek Hansen, Brian Manzo, Jeffrey Regier</i>	
ReFactor GNNs: Revisiting Factorisation-Based Models from a Message-Passing Perspective	16138
<i>Yihong Chen, Luca Franceschi, Pasquale Minervini, Pushkar Mishra, Sebastian Riedel, Pontus Lars Erik Saito Stenetorp</i>	
Efficient Architecture Search for Diverse Tasks.....	16151
<i>Misha Khodak, Junhong Shen, Ameet Talwalkar</i>	
Beyond Not-Forgetting: Continual Learning with Backward Knowledge Transfer.....	16165
<i>Deliang Fan, Sen Lin, Li Yang, Junshan Zhang</i>	

Inherently Explainable Reinforcement Learning in Natural Language	16178
<i>Xiangyu Peng, Chen Xing, Prafulla Kumar Choubey, Chien-Sheng Wu, Caiming Xiong</i>	
Conditional Independence Testing with Heteroskedastic Data and Applications to Causal Discovery	16191
<i>Wiebke Günther, Urmi Ninad, Jakob Runge, Jonas Wahl</i>	
On Reinforcement Learning and Distribution Matching for Fine-Tuning Language Models with No Catastrophic Forgetting	16203
<i>Marc Dymetman, Hady Elsahar, Tomasz Korbak, Germán Kruszewski</i>	
Ask4Help: Learning to Leverage an Expert for Embodied Tasks	16221
<i>Jonghyun Choi, Alvaro Herrasti, Aniruddha Kembhavi, Roozbeh Mottaghi, Kunal Pratap Singh, Luca Weihs</i>	
Human-AI Collaborative Bayesian Optimisation	16233
<i>Arun Kumar A V, Santu Rana, Alistair Shilton, Svetha Venkatesh</i>	
Robust Learning Against Relational Adversaries	16246
<i>Mohannad Alhanahnah, Mihai Christodorescu, Somesh Jha, Xiaozhu Meng, Ke Wang, Yizhen Wang</i>	
Active Bayesian Causal Inference	16261
<i>Christian Knoll, Andreas Krause, Lars Lorch, Robert Peharz, Franz Pernkopf, Christian Toth, Julius Von Kügelgen</i>	

VOLUME 22

Understanding and Improving Robustness of Vision Transformers Through Patch-Based Negative Augmentation	16276
<i>Alex Beutel, Ting Chen, Balaji Lakshminarayanan, Yao Qin, Xuezhi Wang, Chiyuan Zhang</i>	
LogiGAN: Learning Logical Reasoning Via Adversarial Pre-Training	16290
<i>Nan Duan, Yan Gao, Jian-Guang Lou, Xinyu Pi, Wanjun Zhong</i>	
Revisiting Non-Parametric Matching Cost Volumes for Robust and Generalizable Stereo Matching	16305
<i>Kelvin Cheng, Christopher Healey, Tianfu Wu</i>	
On the Interpretability of Regularisation for Neural Networks Through Model Gradient Similarity	16319
<i>Viktor Andersson, Rebecka Jörnsten, Balazs Kulcsar, Vincent Szolnoky</i>	
Controllable 3D Face Synthesis with Conditional Generative Occupancy Fields	16331
<i>Zhaoyang Huang, Hongsheng Li, Keqiang Sun, Quan Wang, Shangzhe Wu, Ning Zhang</i>	
FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness	16344
<i>Tri Dao, Stefano Ermon, Dan Fu, Christopher Ré, Atri Rudra</i>	
Locally Hierarchical Auto-Regressive Modeling for Image Generation	16360
<i>Bohyung Han, Chiheon Kim, Saehoon Kim, Doyup Lee, Tackgeun You</i>	
Learning Enhanced Representation for Tabular Data Via Neighborhood Propagation	16373
<i>Kounianhua Du, Quan Gan, Jiarui Jin, Yangkun Wang, David P Wipf, Weinan Zhang, Zheng Zhang, Xilong Zhao, Ruiwen Zhou</i>	
On Global Convergence of ResNets: From Finite to Infinite Width Using Linear Parameterization	16385
<i>Raphaël Barboni, Gabriel Peyré, Francois-Xavier Vialard</i>	

Spherization Layer: Representation Using Only Angles	16398
<i>Kangil Kim, Hoyong Kim</i>	
On the Identifiability of Nonlinear ICA: Sparsity and Beyond	16411
<i>Ignavier Ng, Kun Zhang, Yujia Zheng</i>	
Self-Supervised Visual Representation Learning with Semantic Grouping	16423
<i>Xiaojuan Qi, Xin Wen, Xiangyu Zhang, Bingchen Zhao, Anlin Zheng</i>	
Optimistic Mirror Descent Either Converges to Nash Or to Strong Coarse Correlated Equilibria in Bimatrix Games.....	16439
<i>Ioannis Anagnostides, Gabriele Farina, Ioannis Panageas, Tuomas Sandholm</i>	
Discovered Policy Optimisation	16455
<i>Jakob Foerster, Jakub Kuba, Alistair Letcher, Chris Lu, Luke Metz, Christian Schroeder De Witt</i>	
Robust Testing in High-Dimensional Sparse Models	16469
<i>Clément L Canonne, Anand Jerry George</i>	
Learning Consistency-Aware Unsigned Distance Functions Progressively from Raw Point Clouds.....	16481
<i>Yi Fang, Zhizhong Han, Yu-Shen Liu, Baorui Ma, Junsheng Zhou</i>	
Understanding Square Loss in Training Overparametrized Neural Network Classifiers	16495
<i>Tianyang Hu, Zhenguo Li, Wenjia Wang, Jun Wang</i>	
Multi-Agent Reinforcement Learning is a Sequence Modeling Problem.....	16509
<i>Jakub Kuba, Runji Lin, Jun Wang, Muning Wen, Ying Wen, Yaodong Yang, Weinan Zhang</i>	
Non-Stationary Bandits with Knapsacks	16522
<i>Jiashuo Jiang, Xiaocheng Li, Shang Liu</i>	
Fast Bayesian Inference with Batch Bayesian Quadrature Via Kernel Recombination.....	16533
<i>Masaki Adachi, Satoshi Hayakawa, Martin Jørgensen, Harald Oberhauser, Michael A Osborne</i>	
Are Two Heads the Same as One? Identifying Disparate Treatment in Fair Neural Networks	16548
<i>Krishnaram Kenthapadi, Matthäus Kleindessner, Francesco Locatello, Michael Lohaus, Chris Russell</i>	
Provable Generalization of Overparameterized Meta-Learning Trained with SGD.....	16563
<i>Longbo Huang, Yu Huang, Yingbin Liang</i>	
When Do Flat Minima Optimizers Work?.....	16577
<i>Jean Kaddour, Matt Kusner, Linqing Liu, Ricardo Silva</i>	
Fast Instrument Learning with Faster Rates	16596
<i>Ziyu Wang, Yuhao Zhou, Jun Zhu</i>	
Sparse Gaussian Process Hyperparameters: Optimize Or Integrate?	16612
<i>Wessel Bruinsma, David Burt, Vidhi Lalchand, Carl Edward Rasmussen</i>	
HierSpeech: Bridging the Gap Between Text and Speech by Hierarchical Variational Inference Using Self-Supervised Representations for Speech Synthesis	16624
<i>Min-Jae Hwang, Seung-Bin Kim, Ji-Hyun Lee, Sang-Hoon Lee, Seong-Whan Lee, Eunwoo Song</i>	

Sharing Knowledge for Meta-Learning with Feature Descriptions.....	16637
<i>Tomoharu Iwata, Atsutoshi Kumagai</i>	
A Variant of Anderson Mixing with Minimal Memory Size	16650
<i>Chenglong Bao, Yang Liu, Fuchao Wei, Guangwen Yang</i>	
AdaptFormer: Adapting Vision Transformers for Scalable Visual Recognition.....	16664
<i>Shoufa Chen, Chongjian Ge, Ping Luo, Yibing Song, Zhan Tong, Jue Wang, Jiangliu Wang</i>	
Symmetry Teleportation for Accelerated Optimization.....	16679
<i>Nima Dehmamy, Robin Walters, Rose Yu, Bo Zhao</i>	
Wasserstein Logistic Regression with Mixed Features	16691
<i>Mohammad Reza Belbasi, Martin Haugh, Aras Selvi, Wolfram Wiesemann</i>	
TaiSu: A 166M Large-Scale High-Quality Dataset for Chinese Vision-Language Pre-Training.....	16705
<i>Haoran Chen, Guojing Ge, Yulong Liu, Ru Peng, Guanhuai Qiao, Qi Song, Jinqiao Wang, Lingxiang Wu, Bin Zhu, Guibo Zhu</i>	
Robust Bayesian Regression Via Hard Thresholding	16718
<i>Qingpei Hu, Zhaohui Li, Fan Zheyi</i>	
Trajectory Inference Via Mean-Field Langevin in Path Space.....	16731
<i>Lénaïc Chizat, Matthieu Heitz, Geoffrey Schiebinger, Stephen Zhang</i>	
SwinTrack: A Simple and Strong Baseline for Transformer Tracking	16743
<i>Heng Fan, Liting Lin, Haibin Ling, Yong Xu, Zhipeng Zhang</i>	
Unknown-Aware Domain Adversarial Learning for Open-Set Domain Adaptation	16755
<i>Joonho Jang, Mingi Ji, Il-Chul Moon, Byeonghu Na, Dong Hyeok Shin, Kyungwoo Song</i>	
Learning Chaotic Dynamics in Dissipative Systems	16768
<i>Anima Anandkumar, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Nikola Kovachki, Zongyi Li, Burigede Liu, Miguel Liu-Schiaffini, Andrew Stuart</i>	
Poisson Flow Generative Models	16782
<i>Tommi Jaakkola, Ziming Liu, Max Tegmark, Yilun Xu</i>	
Bounding and Approximating Intersectional Fairness Through Marginal Fairness	16796
<i>Patrick Loiseau, Mathieu Molina</i>	
Semi-Infinitely Constrained Markov Decision Processes	16808
<i>Yang Peng, Wenhao Yang, Zhihua Zhang, Liangyu Zhang</i>	
Latent Planning Via Expansive Tree Search.....	16821
<i>Robert Giesemann, Florian T. Pokorny</i>	
Invertible Monotone Operators for Normalizing Flows	16836
<i>Byeongkeun Ahn, Youngjoon Hong, Chiyoon Kim, Hyunwoo Kim</i>	
Dynamic Sparse Network for Time Series Classification: Learning What to “See”	16849
<i>Shiwei Liu, Decebal Constantin Mocanu, Elena Mocanu, Mykola Pechenizkiy, Boqian Wu, Qiao Xiao, Yu Zhang</i>	
Learning to Sample and Aggregate: Few-Shot Reasoning Over Temporal Knowledge Graphs	16863
<i>Tarek Abdelzaher, Jinning Li, Zheng Li, Shengzhong Liu, Dachun Sun, Ruijie Wang, Bing Yin</i>	

Evaluating Robustness to Dataset Shift Via Parametric Robustness Sets.....	16877
<i>Michael Oberst, David Sontag, Nikolaj Thams</i>	
CogView2: Faster and Better Text-To-Image Generation Via Hierarchical Transformers	16890
<i>Ming Ding, Wenyi Hong, Jie Tang, Wendi Zheng</i>	
Recursive Reasoning in Minimax Games: A Level k Gradient Play Method.....	16903
<i>Zichu Liu, Laca Pavel</i>	
When to Ask for Help: Proactive Interventions in Autonomous Reinforcement Learning	16918
<i>Chelsea Finn, Archit Sharma, Fahim Tajwar, Annie Xie</i>	
Reinforcement Learning with Neural Radiance Fields.....	16931
<i>Danny Driess, Pete Florence, Yunzhu Li, Ingmar Schubert, Marc Toussaint</i>	
Function Classes for Identifiable Nonlinear Independent Component Analysis	16946
<i>Michel Besserve, Simon Buchholz, Bernhard Schölkopf</i>	
Self-Supervised Heterogeneous Graph Pre-Training Based on Structural Clustering.....	16962
<i>Ziyu Guan, Jianbin Huang, Weigang Lu, Zhe Wang, Cai Xu, Yaming Yang, Wei Zhao</i>	
Learning Audio-Visual Dynamics Using Scene Graphs for Audio Source Separation.....	16975
<i>Narendra Ahuja, Moitrya Chatterjee, Anoop Cherian</i>	
On the Sample Complexity of Stabilizing LTI Systems on a Single Trajectory.....	16989
<i>Yang Hu, Guannan Qu, Adam Wierman</i>	
coVariance Neural Networks	17003
<i>Gonzalo Mateos, Corey McMillan, Alejandro Ribeiro, Saurabh Sihag</i>	
Taming Fat-Tailed (“Heavier-Tailed” with Potentially Infinite Variance) Noise in Federated Learning	17017
<i>Jia Liu, Peiwen Qiu, Haibo Yang</i>	
Exploring Figure-Ground Assignment Mechanism in Perceptual Organization	17030
<i>Yang Cao, Zheng-Jun Zha, Wei Zhai, Jing Zhang</i>	
Two-Stream Network for Sign Language Recognition and Translation.....	17043
<i>Yutong Chen, Shujie Liu, Brian Mak, Fangyun Wei, Yu Wu, Ronglai Zuo</i>	

VOLUME 23

VisFIS: Visual Feature Importance Supervision with Right-For-The-Right-Reason Objectives.....	17057
<i>Mohit Bansal, Peter Hase, Zhuofan Ying</i>	
Operative Dimensions in Unconstrained Connectivity of Recurrent Neural Networks	17073
<i>Matthew Cook, Giacomo Indiveri, Sepp Kollmorgen, Renate Krause, Valerio Mante</i>	
Batch Size-Invariance for Policy Optimization.....	17086
<i>Karl Cobbe, Jacob Hilton, John Schulman</i>	
Inference and Sampling for Archimax Copulas.....	17099
<i>Ali Hasan, Yuting Ng, Vahid Tarokh</i>	
Neural Shape Deformation Priors.....	17117
<i>Lev Markhasin, Matthias Niessner, Jiapeng Tang, Justus Thies, Bi Wang</i>	

The Curse of Unrolling: Rate of Differentiating Through Optimization	17133
<i>Quentin Bertrand, Gauthier Gidel, Fabian Pedregosa, Damien Scieur</i>	
A Deep Learning Dataloader with Shared Data Preparation	17146
<i>Chun Cao, Zenan Li, Hanghang Tong, Guochang Wang, Jian Xie, Jingwei Xu, Yuan Yao</i>	
Learning Physics Constrained Dynamics Using Autoencoders	17157
<i>Karthik Narasimhan, Peter J Ramadge, Justinian Rosca, Tsung-Yen Yang</i>	
Variational Model Perturbation for Source-Free Domain Adaptation	17173
<i>Mengmeng Jing, Jingjing Li, Cees Snoek, Xiantong Zhen</i>	
On the Non-Universality of Deep Learning: Quantifying the Cost of Symmetry	17188
<i>Emmanuel Abbe, Enric Boix-Adsera</i>	
Sharper Convergence Guarantees for Asynchronous SGD for Distributed and Federated Learning	17202
<i>Martin Jaggi, Anastasiia Koloskova, Sebastian Stich</i>	
A Unified Framework for Alternating Offline Model Training and Policy Learning	17216
<i>Yihao Feng, Shentao Yang, Shujian Zhang, Mingyuan Zhou</i>	
Adjoint-Aided Inference of Gaussian Process Driven Differential Equations	17233
<i>Mauricio A Álvarez, Engineer Bainomugisha, Paterne Gahungu, Christopher Lanyon, Michael T Smith, Richard Wilkinson</i>	
BOME! Bilevel Optimization Made Easy: A Simple First-Order Approach	17248
<i>Bo Liu, Qiang Liu, Peter Stone, Stephen Wright, Mao Ye</i>	
Mirror Descent with Relative Smoothness in Measure Spaces, with Application to Sinkhorn and EM	17263
<i>Pierre-Cyril Aubin-Frankowski, Anna Korba, Flavien Léger</i>	
Peer Prediction for Learning Agents	17276
<i>Yiling Chen, Shi Feng, Fang-Yi Yu</i>	
Visual Clues: Bridging Vision and Language Foundations for Image Paragraph Captioning	17287
<i>Nguyen Bach, Xiyang Dai, Ce Liu, Yujia Xie, Lu Yuan, Michael Zeng, Luowei Zhou</i>	
APT-36K: A Large-Scale Benchmark for Animal Pose Estimation and Tracking	17301
<i>Long Lan, Dacheng Tao, Yufei Xu, Junjie Yang, Yuxiang Yang, Jing Zhang</i>	
ShuffleMixer: An Efficient ConvNet for Image Super-Resolution	17314
<i>Jinshan Pan, Long Sun, Jinhui Tang</i>	
Neural Topological Ordering for Computation Graphs	17327
<i>Roberto Bondesan, Mukul Gagrani, Wonseok Jeon, Christopher Lott, Corrado Rainone, Harris Teague, Herke Van Hoof, Yang Yang, Piero Zappi, Weiliang Zeng</i>	
Probable Domain Generalization Via Quantile Risk Minimization	17340
<i>Cian Eastwood, Hamed Hassani, George J. Pappas, Alexander Robey, Bernhard Schölkopf, Shashank Singh, Julius Von Kügelgen</i>	
Locating and Editing Factual Associations in GPT	17359
<i>Alex Andonian, David Bau, Yonatan Belinkov, Kevin Meng</i>	

S-PIFu: Integrating Parametric Human Models with PIFu for Single-View Clothed Human Reconstruction.....	17373
<i>Kennard Chan, Guosheng Lin, Weisi Lin, Haiyu Zhao</i>	
A Composable Machine-Learning Approach for Steady-State Simulations on High-Resolution Grids.....	17386
<i>Lalit Ghule, Chris Hill, Jay Pathak, Rishikesh Ranade</i>	
Outlier Suppression: Pushing the Limit of Low-Bit Transformer Language Models.....	17402
<i>Ruihao Gong, Xianglong Liu, Xiuying Wei, Fengwei Yu, Qi Zhang, Shanghang Zhang, Xiangguo Zhang, Yunchen Zhang</i>	
A Single-Timescale Analysis for Stochastic Approximation with Multiple Coupled Sequences.....	17415
<i>Tianyi Chen, Han Shen</i>	
Bayesian Risk Markov Decision Processes.....	17430
<i>Yifan Lin, Yuxuan Ren, Enlu Zhou</i>	
All Politics is Local: Redistricting Via Local Fairness.....	17443
<i>Pankaj Agarwal, Shao-Heng Ko, Kamesh Munagala, Erin Taylor</i>	
Confident Adaptive Language Modeling.....	17456
<i>Dara Bahri, Mostafa Dehghani, Adam Fisch, Jai Gupta, Donald Metzler, Tal Schuster, Yi Tay, Vinh Tran</i>	
Submodular Maximization in Clean Linear Time.....	17473
<i>Moran Feldman, Amin Karbasi, Ehsan Kazemi, Wenxin Li</i>	
Parameters Or Privacy: A Provable Tradeoff Between Overparameterization and Membership Inference.....	17488
<i>Richard Baraniuk, Hamid Javadi, Blake Mason, Jasper Tan</i>	
EF-BV: A Unified Theory of Error Feedback and Variance Reduction Mechanisms for Biased and Unbiased Compression in Distributed Optimization.....	17501
<i>Laurent Condat, Peter Richtarik, Kai Yi</i>	
DataMUX: Data Multiplexing for Neural Networks.....	17515
<i>Carlos Jimenez, Vishvak Murahari, Karthik Narasimhan, Runzhe Yang</i>	
Biologically-Plausible Backpropagation Through Arbitrary Timespans Via Local Neuromodulators.....	17528
<i>Yuhan Helena Liu, Stefan Mihalas, Eric Shea-Brown, Stephen Smith, Uygur Sümbül</i>	
Revisiting Realistic Test-Time Training: Sequential Inference and Adaptation by Anchored Clustering.....	17543
<i>Kui Jia, Yongyi Su, Xun Xu</i>	
Label Noise in Adversarial Training: A Novel Perspective to Study Robust Overfitting.....	17556
<i>Chengyu Dong, Liyuan Liu, Jingbo Shang</i>	
Active Labeling: Streaming Stochastic Gradients.....	17568
<i>Francis Bach, Vivien Cabannes, Vianney Perchet, Alessandro Rudi</i>	
CEBaB: Estimating the Causal Effects of Real-World Concepts on NLP Model Behavior.....	17582
<i>Eldar D Abraham, Karel D'Oosterlinck, Amir Feder, Yair Gat, Atticus Geiger, Christopher Potts, Roi Reichart, Zhengxuan Wu</i>	

TOIST: Task Oriented Instance Segmentation Transformer with Noun-Pronoun Distillation	17597
<i>Xiaoxue Chen, Pengfei Li, Yongliang Shi, Beiwen Tian, Ya-Qin Zhang, Hao Zhao, Guyue Zhou</i>	
Mind the Gap: Understanding the Modality Gap in Multi-Modal Contrastive Representation Learning	17612
<i>Yongchan Kwon, Victor Weixin Liang, Serena Yeung, Yuhui Zhang, James Zou</i>	
α -ReQ : Assessing Representation Quality in Self-Supervised Learning by Measuring Eigenspectrum Decay.....	17626
<i>Kumar K Agrawal, Arna Ghosh, Arnab Kumar Mondal, Blake Richards</i>	
When Combinatorial Thompson Sampling Meets Approximation Regret	17639
<i>Pierre Perrault</i>	
Pruning Has a Disparate Impact on Model Accuracy.....	17652
<i>Ferdinando Fioretto, Jung-Eun Kim, Rakshit Naidu, Cuong Tran</i>	
Sequence Model Imitation Learning with Unobserved Contexts	17665
<i>J. Bagnell, Sanjiban Choudhury, Gokul Swamy, Steven Wu</i>	
Tikhonov Regularization is Optimal Transport Robust Under Martingale Constraints.....	17677
<i>Jose Blanchet, Jiajin Li, Sirui Lin, Viet Anh Nguyen</i>	
Policy Optimization with Linear Temporal Logic Constraints.....	17690
<i>Swarat Chaudhuri, Hoang Le, Cameron Voloshin, Yisong Yue</i>	
Merging Models with Fisher-Weighted Averaging.....	17703
<i>Michael S Matena, Colin Raffel</i>	
Learning to Generate Inversion-Resistant Model Explanations	17717
<i>Sung Ju Hwang, Hoyong Jeong, Suyoung Lee, Soel Son</i>	
Unsupervised Multi-View Object Segmentation Using Radiance Field Propagation.....	17730
<i>Jiabao Chen, Xinhang Liu, Yu-Wing Tai, Chi-Keung Tang, Huai Yu</i>	
Beyond Mahalanobis Distance for Textual OOD Detection.....	17744
<i>Pierre Colombo, Eduardo Dadoalto, Nathan Noiry, Pablo Piantanida, Guillaume Staerman</i>	
FasterRisk: Fast and Accurate Interpretable Risk Scores	17760
<i>Boxuan Li, Jiachang Liu, Cynthia Rudin, Margo Seltzer, Chudi Zhong</i>	
Graph Learning Assisted Multi-Objective Integer Programming.....	17774
<i>Zhiguang Cao, Abhishek Gupta, Mingyan Lin, Wen Song, Yaoxin Wu, Jie Zhang</i>	
Revisiting Sliced Wasserstein on Images: From Vectorization to Convolution.....	17788
<i>Nhat Ho, Khai Nguyen</i>	
SignRFF: Sign Random Fourier Features.....	17802
<i>Ping Li, Xiaoyun Li</i>	
The Role of Baselines in Policy Gradient Optimization.....	17818
<i>Wesley Chung, Bo Dai, Jincheng Mei, Dale Schuurmans, Csaba Szepesvari, Valentin Thomas</i>	
A Rotated Hyperbolic Wrapped Normal Distribution for Hierarchical Representation Learning.....	17831
<i>Seunghyuk Cho, Dongwoo Kim, Juyong Lee, Jaesik Park</i>	

VOLUME 24

Private Graph All-Pairwise-Shortest-Path Distance Release with Improved Error Rate.....	17844
<i>Chenglin Fan, Ping Li, Xiaoyun Li</i>	
Multi-Fidelity Best-Arm Identification	17857
<i>Alberto Maria Metelli, Riccardo Poiani, Marcello Restelli</i>	
SemiFL: Semi-Supervised Federated Learning for Unlabeled Clients with Alternate Training.....	17871
<i>Enmao Diao, Jie Ding, Vahid Tarokh</i>	
RankFeat: Rank-1 Feature Removal for Out-Of-Distribution Detection.....	17885
<i>Nicu Sebe, Yue Song, Wei Wang</i>	
Jump Self-Attention: Capturing High-Order Statistics in Transformers	17899
<i>Jianxin Li, Jieqi Peng, Siyang Xiao, Shanghang Zhang, Shuai Zhang, Haoyi Zhou</i>	
Learning Infinite-Horizon Average-Reward Restless Multi-Action Bandits Via Index Awareness.....	17911
<i>Jian Li, Shufan Wang, Guojun Xiong</i>	
STNDT: Modeling Neural Population Activity with Spatiotemporal Transformers.....	17926
<i>Trung Le, Eli Shlizerman</i>	
ProtoVAE: A Trustworthy Self-Explainable Prototypical Variational Model.....	17940
<i>Ahcène Boubekki, Srishti Gautam, Stine Hansen, Marina Höhne, Robert Jenssen, Michael Kampffmeyer, Suaiba Salahuddin</i>	
If Influence Functions Are the Answer, Then What is the Question?.....	17953
<i>Juhan Bae, Marzyeh Ghassemi, Roger Grosse, Alston Lo, Nathan Ng</i>	
FACT: Learning Governing Abstractions Behind Integer Sequences.....	17968
<i>Peter Belcak, Ard Kastrati, Flavio Schenker, Roger Wattenhofer</i>	
SAPipe: Staleness-Aware Pipeline for Data Parallel DNN Training	17981
<i>Yangrui Chen, Juncheng Gu, Haibin Lin, Meng Ma, Yanghua Peng, Chuan Wu, Cong Xie, Yibo Zhu</i>	
Probing Classifiers Are Unreliable for Concept Removal and Detection.....	17994
<i>Abhinav Kumar, Amit Sharma, Chenhao Tan</i>	
SCL-WC: Cross-Slide Contrastive Learning for Weakly-Supervised Whole-Slide Image Classification	18009
<i>Xiao Han, Junzhou Huang, Ming-Hui Wang, Xiyue Wang, Jinxi Xiang, Sen Yang, Wei Yang, Zhongyi Yang, Jun Zhang, Jing Zhang</i>	
Reproducibility in Optimization: Theoretical Framework and Limits	18022
<i>Kwangjun Ahn, Prateek Jain, Ziwei Ji, Satyen Kale, Praneeth Netrapalli, Gil I Shamir</i>	
Hierarchical Classification at Multiple Operating Points	18034
<i>Jack Valmadre</i>	
Kernel Multimodal Continuous Attention	18046
<i>Walter Dempsey, Alexander Moreno, Supriya Nagesh, James Rehg, Zhenke Wu</i>	
VisCo Grids: Surface Reconstruction with Viscosity and Coarea Grids	18060
<i>Yaron Lipman, Albert Pumarola, Artsiom Sanakoyeu, Ali Thabet, Lior Yariv</i>	

Exact Shape Correspondence Via 2D Graph Convolution	18072
<i>Yongqiang Chen, James Cheng, Barakeel Fansou Kamhoua, Bo Han, Ma Kaili, Bo Li, Han Yang, Lin Zhang</i>	
Posterior Matching for Arbitrary Conditioning	18088
<i>Junier B Oliva, Ryan Strauss</i>	
CARD: Classification and Regression Diffusion Models.....	18100
<i>Xizewen Han, Huangjie Zheng, Mingyuan Zhou</i>	
What Can the Neural Tangent Kernel Tell Us About Adversarial Robustness?.....	18116
<i>Julia Kempe, Nikolaos Tsilivis</i>	
Benefits of Permutation-Equivariance in Auction Mechanisms.....	18131
<i>Fengxiang He, Wenbing Huang, Tian Qin, Dingfeng Shi, Dacheng Tao</i>	
MoCoDA: Model-Based Counterfactual Data Augmentation	18143
<i>Elliot Creager, Animesh Garg, Ajay Mandlekar, Silviu Pitis</i>	
SkinCon: A Skin Disease Dataset Densely Annotated by Domain Experts for Fine-Grained Debugging and Analysis.....	18157
<i>Zhuo Ran Cai, Roxana Daneshjou, Roberto Novoa, Mert Yuksekgonul, James Zou</i>	
A Permutation-Free Kernel Two-Sample Test	18168
<i>Ilmun Kim, Aaditya Ramdas, Shubhanshu Shekhar</i>	
Simple Unsupervised Object-Centric Learning for Complex and Naturalistic Videos.....	18181
<i>Sungjin Ahn, Gautam Singh, Yi-Fu Wu</i>	
Communication-Efficient Distributed Eigenspace Estimation with Arbitrary Node Failures	18197
<i>Vasileios Charisopoulos, Anil Damle</i>	
On Uncertainty, Tempering, and Data Augmentation in Bayesian Classification	18211
<i>Pavel Izmailov, Sanyam Kapoor, Wesley Maddox, Andrew Wilson</i>	
Optimal Binary Classification Beyond Accuracy	18226
<i>Justin Khim, Shashank Singh</i>	
Information-Theoretic GAN Compression with Variational Energy-Based Model.....	18241
<i>Bohyung Han, Eunhee Kang, Minsoo Kang, Sehwan Ki, Hyong Euk Lee, Hyewon Yoo</i>	
Learning versus Refutation in Noninteractive Local Differential Privacy	18256
<i>Alexander Edmonds, Aleksandar Nikolov, Toniann Pitassi</i>	
Why So Pessimistic? Estimating Uncertainties for Offline RL Through Ensembles, and Why Their Independence Matters.....	18267
<i>Kamyar Ghasemipour, Shixiang Shane Gu, Ofir Nachum</i>	
Private Synthetic Data for Multitask Learning and Marginal Queries	18282
<i>Cedric Archambeau, Sergul Aydore, William Brown, Michael Kearns, Aaron Roth, Ankit Siva, Shuai Tang, Giuseppe Vietri, Steven Wu</i>	
Sample-Efficient Reinforcement Learning of Partially Observable Markov Games.....	18296
<i>Chi Jin, Qinghua Liu, Csaba Szepesvari</i>	

Advancing Model Pruning Via Bi-Level Optimization	18309
<i>Tianlong Chen, Mingyi Hong, Sijia Liu, Parikshit Ram, Yanzhi Wang, Yuguang Yao, Yihua Zhang, Pu Zhao</i>	
Optimal Positive Generation Via Latent Transformation for Contrastive Learning	18327
<i>Hong Chang, Xilin Chen, Yinqi Li, Bingpeng Ma, Shiguang Shan</i>	
MineDojo: Building Open-Ended Embodied Agents with Internet-Scale Knowledge.....	18343
<i>Anima Anandkumar, Linxi Fan, De-An Huang, Yunfan Jiang, Ajay Mandlekar, Andrew Tang, Guanzhi Wang, Yuncong Yang, Haoyi Zhu, Yuke Zhu</i>	
The Price of Unfairness in Linear Bandits with Biased Feedback	18363
<i>Alexandra Carpentier, Solenne Gaucher, Christophe Giraud</i>	
DeepFoids: Adaptive Bio-Inspired Fish Simulation with Deep Reinforcement Learning.....	18377
<i>Yuko Ishiwaka, Masaki Nakada, Shun Ogawa, Tadayuki Tone, Donovan Westwater, Xiao Zeng</i>	
Truncated Matrix Power Iteration for Differentiable DAG Learning.....	18390
<i>Ehsan Abbasnejad, Dong Gong, Mingming Gong, Yuhang Liu, Ignavier Ng, Javen Qinfeng Shi, Kun Zhang, Zhen Zhang</i>	
Learning Debaised Classifier with Biased Committee	18403
<i>Sungsoo Ahn, Sehyun Hwang, Nayeong Kim, Suha Kwak, Jaesik Park</i>	
Modeling the Machine Learning Multiverse	18416
<i>Samuel J. Bell, Jesse Dodge, Onno Kampman, Neil Lawrence</i>	
Label-Aware Global Consistency for Multi-Label Learning with Single Positive Labels	18430
<i>Sheng-Jun Huang, Jiahao Xiao, Ming-Kun Xie</i>	
Unifying Voxel-Based Representation with Transformer for 3D Object Detection	18442
<i>Yilun Chen, Jiaya Jia, Yanwei Li, Zeming Li, Xiaojuan Qi, Jian Sun</i>	
UDC: Unified DNAs for Compressible TinyML Models for Neural Processing Units.....	18456
<i>Igor Fedorov, Ramon Matas, Matthew Mattina, Hokchhay Tann, Paul Whatmough, Chuteng Zhou</i>	
Measures of Information Reflect Memorization Patterns.....	18472
<i>Rachit Bansal, Yonatan Belinkov, Danish Pruthi</i>	
Fine-Grained Analysis of Stability and Generalization for Modern Meta Learning Algorithms.....	18487
<i>Jiechao Guan, Yong Liu, Zhiwu Lu</i>	
On Scrambling Phenomena for Randomly Initialized Recurrent Networks.....	18501
<i>Vaggos Chatziafratis, Ioannis Panageas, Clayton Sanford, Stelios Stavroulakis</i>	
Learning to Branch with Tree MDPs.....	18514
<i>Karen Aardal, Feng Chen, Didier Chetelat, Maxime Gasse, Andrea Lodi, Lara Scavuzzo, Neil Yorke-Smith</i>	
Neural Sheaf Diffusion: A Topological Perspective on Heterophily and Oversmoothing in GNNs.....	18527
<i>Cristian Bodnar, Michael Bronstein, Benjamin Chamberlain, Francesco Di Giovanni, Pietro Lió</i>	
PyKT: A Python Library to Benchmark Deep Learning Based Knowledge Tracing Models.....	18542
<i>Jiahao Chen, Shuyan Huang, Qiongqiong Liu, Zitao Liu, Weiqi Luo, Jiliang Tang</i>	

Adversarial Unlearning: Reducing Confidence Along Adversarial Directions	18556
<i>Benjamin Eysenbach, Sergey Levine, Amrith Setlur, Virginia Smith</i>	
How Would the Viewer Feel? Estimating Wellbeing from Video Scenarios	18571
<i>Steven Basart, Jun Shern Chan, David Forsyth, Dan Hendrycks, Mantas Mazeika, Dawn Song, Jacob Steinhardt, Eric Tang, Andy Zou</i>	
On Elimination Strategies for Bandit Fixed-Confidence Identification	18586
<i>Rémy Degenne, Andrea Tirinzoni</i>	
When Adversarial Training Meets Vision Transformers: Recipes from Training to Architecture	18599
<i>Yiwen Guo, Yichuan Mo, Yisen Wang, Yifei Wang, Dongxian Wu</i>	
Escaping from the Barren Plateau Via Gaussian Initializations in Deep Variational Quantum Circuits	18612
<i>Min-Hsiu Hsieh, Liu Liu, Dacheng Tao, Kaining Zhang</i>	
A Sharp NMF Result with Applications in Network Modeling.....	18628
<i>Jiashun Jin</i>	

VOLUME 25

Decoupling Classifier for Boosting Few-Shot Object Detection and Instance Segmentation	18640
<i>Xiaochen Chen, Bin-Bin Gao, Zhongyi Huang, Guannan Jiang, Jinxiang Lai, Jun Liu, Congchong Nie, Chengjie Wang, Xi Wang</i>	
Private Estimation with Public Data.....	18653
<i>Alex Bie, Gautam Kamath, Vikrant Singhal</i>	
Pre-Activation Distributions Expose Backdoor Neurons	18667
<i>Jianze Li, Li Liu, Rongjun Tang, Runkai Zheng</i>	
DaDA: Distortion-Aware Domain Adaptation for Unsupervised Semantic Segmentation.....	18681
<i>Sujin Jang, Joohan Na, Dokwan Oh</i>	
Large-Batch Optimization for Dense Visual Predictions: Training Faster R-CNN in 4.2 Minutes	18694
<i>Liang Chen, Jianming Liang, Yu Liu, Ping Luo, Guanglu Song, Zeyue Xue, Zhuofan Zong</i>	
Most Activation Functions Can Win the Lottery Without Excessive Depth.....	18707
<i>Rebekka Burkholz</i>	
Rapid Model Architecture Adaption for Meta-Learning	18721
<i>Nicolo Fusi, Xitong Gao, Robert Mullins, I Shumailov, Yiren Zhao</i>	
Exploring Through Random Curiosity with General Value Functions.....	18733
<i>Louis Kirsch, Aditya Ramesh, Jürgen Schmidhuber, Sjoerd Van Steenkiste</i>	
Optimal Transport-Based Identity Matching for Identity-Invariant Facial Expression Recognition.....	18749
<i>Daeha Kim, Byung Cheol Song</i>	
Unsupervised Learning Under Latent Label Shift.....	18763
<i>Saurabh Garg, Zachary Lipton, Pranav Mani, Manley Roberts</i>	
SHINE: SubHypergraph Inductive Neural nEtwork.....	18779
<i>Yuan Luo</i>	

Efficient Aggregated Kernel Tests Using Incomplete U-Statistics	18793
<i>Arthur Gretton, Benjamin Guedj, Ilmun Kim, Antonin Schrab</i>	
Influencing Long-Term Behavior in Multiagent Reinforcement Learning.....	18808
<i>Michael Everett, Jakob Foerster, Jonathan How, Dong-Ki Kim, Miao Liu, Matthew Riemer, Chuangchuang Sun, Gerald Tesauro</i>	
Quantized Training of Gradient Boosting Decision Trees	18822
<i>Zhuoming Chen, Guolin Ke, Tie-Yan Liu, Yu Shi, Shuxin Zheng</i>	
Convergent Representations of Computer Programs in Human and Artificial Neural Networks.....	18834
<i>Evelina Fedorenko, Anna Ivanova, Ben Lipkin, Una-May O'Reilly, Shashank Srikant</i>	
QC-StyleGAN - Quality Controllable Image Generation and Manipulation.....	18850
<i>Tan M. Dinh, Dat Viet Thanh Nguyen, Cuong Pham, Anh Tran, Phong Tran The</i>	
Retrieve, Reason, and Refine: Generating Accurate and Faithful Patient Instructions.....	18864
<i>David Clifton, Shen Ge, Fenglin Liu, Zhangdaihong Liu, Xu Sun, Xian Wu, Bang Yang, Yang Yang, Chenyu You</i>	
Data Distributional Properties Drive Emergent In-Context Learning in Transformers	18878
<i>Stephanie Chan, Felix Hill, Andrew Lampinen, James McClelland, Pierre Richemond, Adam Santoro, Aaditya Singh, Jane Wang</i>	
Algorithms that Approximate Data Removal: New Results and Limitations	18892
<i>Vinith Suriyakumar, Ashia Wilson</i>	
LOT: Layer-Wise Orthogonal Training on Improving ℓ_2 Certified Robustness	18904
<i>Bo Li, Linyi Li, Xiaojun Xu</i>	
Lottery Tickets on a Data Diet: Finding Initializations with Sparse Trainable Networks	18916
<i>Gintare Karolina Dziugaite, Jonathan Frankle, Surya Ganguli, Brett Larsen, Mansheej Paul</i>	
Regularized Molecular Conformation Fields	18929
<i>Xuanjing Huang, Lihao Wang, Yiqun Wang, Xiaoqing Zheng, Hao Zhou, Yi Zhou</i>	
Reconstruction on Trees and Low-Degree Polynomials.....	18942
<i>Frederic Koehler, Elchanan Mossel</i>	
Lower Bounds and Nearly Optimal Algorithms in Distributed Learning with Communication Compression.....	18955
<i>Yiming Chen, Xinqiang Huang, Wotao Yin, Kun Yuan</i>	
Memory Safe Computations with XLA Compiler	18970
<i>Yuze An, Artem Artemev, Tilman Roeder, Mark Van Der Wilk</i>	
Towards Theoretically Inspired Neural Initialization Optimization	18983
<i>Zhouchen Lin, Hong Wang, Yibo Yang, Haobo Yuan</i>	
AnimeRun: 2D Animation Visual Correspondence from Open Source 3D Movies.....	18996
<i>Chao Dong, Bo Li, Yuhang Li, Ziwei Liu, Chen Change Loy, Li Siyao</i>	
Uncertainty Estimation Using Riemannian Model Dynamics for Offline Reinforcement Learning.....	19008
<i>Shie Mannor, Guy Tennenholtz</i>	

Improved Convergence Rate of Stochastic Gradient Langevin Dynamics with Variance Reduction and Its Application to Optimization.....	19022
<i>Yuri Kinoshita, Taiji Suzuki</i>	
Practical Adversarial Attacks on Spatiotemporal Traffic Forecasting Models	19035
<i>Wenzhao Jiang, Hao Liu, Fan Liu</i>	
Efficient Learning of Nonlinear Prediction Models with Time-Series Privileged Information	19048
<i>Fredrik Johansson, Bastian Jung</i>	
Layer Freezing & Data Sieving: Missing Pieces of a Generic Framework for Sparse Training	19061
<i>Zhenglun Kong, Sheng Li, Yanyu Li, Jian Ren, Xulong Tang, Sergey Tulyakov, Yanzhi Wang, Geng Yuan</i>	
ETAB: A Benchmark Suite for Visual Representation Learning in Echocardiography	19075
<i>Ahmed M. Alaa, Anthony Philippakis, David Sontag</i>	
Trustworthy Monte Carlo	19087
<i>Juha Harviainen, Petteri Kaski, Mikko Koivisto</i>	
Double Bubble, Toil and Trouble: Enhancing Certified Robustness Through Transitivity	19099
<i>Andrew Cullen, Sarah Erfani, Shijie Liu, Paul Montague, Benjamin Rubinstein</i>	
An α -No-Regret Algorithm for Graphical Bilinear Bandits	19113
<i>Yann Chevaleyre, Igor Colin, Rida Laraki, Geovani Rizk, Albert Thomas</i>	
A Unified Analysis of Federated Learning with Arbitrary Client Participation.....	19124
<i>Mingyue Ji, Shiqiang Wang</i>	
Deconfounded Representation Similarity for Comparison of Neural Networks	19138
<i>Tianyu Cui, Samuel Kaski, Yogesh Kumar, Pekka Marttinen</i>	
Fairness Without Demographics Through Knowledge Distillation	19152
<i>Junyi Chai, Taeuk Jang, Xiaoqian Wang</i>	
Sleeper Agent: Scalable Hidden Trigger Backdoors for Neural Networks Trained from Scratch.....	19165
<i>Rama Chellappa, Liam Fowl, Micah Goldblum, Tom Goldstein, Hossein Souri</i>	
Dynamic Pricing with Monotonicity Constraint Under Unknown Parametric Demand Model	19179
<i>Su Jia, Andrew Li, R Ravi</i>	
A Win-Win Deal: Towards Sparse and Robust Pre-Trained Language Models.....	19189
<i>Yanan Cao, Peng Fu, Jiangnan Li, Zheng Lin, Yuanxin Liu, Fandong Meng, Weiping Wang, Jie Zhou</i>	
Asymptotics of Smoothed Wasserstein Distances in the Small Noise Regime	19203
<i>Yunzi Ding, Jonathan Niles-Weed</i>	
Fine-Tuning Language Models Over Slow Networks Using Activation Quantization with Guarantees	19215
<i>Beidi Chen, Tri Dao, Yongjun He, Christopher Ré, Luka Rimanic, Jue Wang, Binhang Yuan, Ce Zhang</i>	
Pareto Set Learning for Expensive Multi-Objective Optimization.....	19231
<i>Xi Lin, Zhiyuan Yang, Qingfu Zhang, Xiaoyuan Zhang</i>	

Non-Monotonic Resource Utilization in the Bandits with Knapsacks Problem.....	19248
<i>Robert Kleinberg, Raunak Kumar</i>	
Efficient Identification of Informative Features in Simulation-Based Inference.....	19260
<i>Jonas Beck, Philipp Berens, Yves Bernaerts, Michael Deistler, Jakob H Macke</i>	
Agreement-On-The-Line: Predicting the Performance of Neural Networks Under Distribution Shift	19274
<i>Christina Baek, Yiding Jiang, J. Zico Kolter, Aditi Raghunathan</i>	
Large-Scale Differentiable Causal Discovery of Factor Graphs.....	19290
<i>Jan-Christian Huetter, Romain Lopez, Jonathan Pritchard, Aviv Regev</i>	
Diagnosing Failures of Fairness Transfer Across Distribution Shift in Real-World Medical Settings	19304
<i>Ibrahim Alabdulmohsin, Alexander Brown, Christina Chen, Silvia Chiappa, Alexander D'Amour, Awa Dieng, Natalie Harris, Katherine Heller, Alan Karthikesalingam, Sanmi Koyejo, Yuan Liu, Diana Mincu, Vivek Natarajan, Krista Opsahl-Ong, Subhrajit Roy, Eva Schnider, Jessica Schrouff</i>	
Towards Lightweight Black-Box Attack Against Deep Neural Networks.....	19319
<i>Wan Chaoqun, Bo Han, Ya Li, Tongliang Liu, Chenghao Sun, Xinmei Tian, Qizhou Wang, Yonggang Zhang</i>	
Federated Learning from Pre-Trained Models: A Contrastive Learning Approach.....	19332
<i>Jing Jiang, Lu Liu, Guodong Long, Jie Ma, Yue Tan, Tianyi Zhou</i>	
Theoretically Provable Spiking Neural Networks.....	19345
<i>Shao-Qun Zhang, Zhi-Hua Zhou</i>	
Approximate Euclidean Lengths and Distances Beyond Johnson-Lindenstrauss	19357
<i>Mathieu Luisier, Aleksandros Sobczyk</i>	
InsPro: Propagating Instance Query and Proposal for Online Video Instance Segmentation.....	19370
<i>Naiyu Gao, Fei He, Kaiqi Huang, Jian Jia, Yanhu Shan, Haoyang Zhang, Xin Zhao</i>	
Measuring and Reducing Model Update Regression in Structured Prediction for NLP	19384
<i>Deng Cai, Yi-An Lai, Elman Mansimov, Lei Shu, Yixuan Su, Yi Zhang</i>	
Rethinking Lipschitz Neural Networks and Certified Robustness: A Boolean Function Perspective	19398
<i>Di He, Du Jiang, Liwei Wang, Bohang Zhang</i>	
Multivariate Time-Series Forecasting with Temporal Polynomial Graph Neural Networks.....	19414
<i>Wei Chen, Haozhe Feng, Qinxian Liu, Yijing Liu, Zhongwei Wang, Jian-Wei Zhang, Zihan Zhou</i>	

VOLUME 26

Few-Shot Image Generation Via Adaptation-Aware Kernel Modulation	19427
<i>Milad Abdollahzadeh, Keshigeyan Chandrasegaran, Ngai-Man Man Cheung, Yunqing Zhao</i>	
Learning to Follow Instructions in Text-Based Games	19441
<i>Toryn Klassen, Andrew Li, Sheila McIlraith, Scott Sanner, Mathieu Tuli, Pashootan Vaezipoor</i>	
Concentration of Data Encoding in Parameterized Quantum Circuits	19456
<i>Guangxi Li, Xin Wang, Ruilin Ye, Xuanqiang Zhao</i>	

Improving Variational Autoencoders with Density Gap-Based Regularization.....	19470
<i>Jun Bai, Chenghua Lin, Wenge Rong, Yanmeng Wang, Jianfei Zhang</i>	
RISE: Robust Individualized Decision Learning with Sensitive Variables	19484
<i>Zhengling Qi, Christopher Seymour, Xiaoqing Tan, Lu Tang</i>	
Bayesian Clustering of Neural Spiking Activity Using a Mixture of Dynamic Poisson Factor Analyzers.....	19499
<i>Ian H Stevenson, Xiaojing Wang, Ganchao Wei</i>	
Understanding Deep Contrastive Learning Via Coordinate-Wise Optimization	19511
<i>Yuandong Tian</i>	
Beyond Neural Scaling Laws: Beating Power Law Scaling Via Data Pruning	19523
<i>Surya Ganguli, Robert Geirhos, Ari Morcos, Shashank Shekhar, Ben Sorscher</i>	
Neural Estimation of Submodular Functions with Applications to Differentiable Subset Selection.....	19537
<i>Soumen Chakrabarti, Abir De</i>	
Maximum Class Separation as Inductive Bias in One Matrix	19553
<i>Gertjan Burghouts, Rita Cucchiara, Tejaswi Kasarla, Pascal Mettes, Elise Van Der Pol, Max Van Spengler</i>	
Surprising Instabilities in Training Deep Networks and a Theoretical Analysis	19567
<i>Dong Lao, Yuxin Sun, Ganesh Sundaramoorthi, Anthony Yezzi</i>	
WaveBound: Dynamic Error Bounds for Stable Time Series Forecasting	19579
<i>Youngin Cho, Jaegul Choo, Mohammad Azam Khan, Daejin Kim, Dongmin Kim</i>	
Finite-Time Analysis of Adaptive Temporal Difference Learning with Deep Neural Networks.....	19592
<i>Dongsheng Li, Tao Sun, Bao Wang</i>	
Benign Underfitting of Stochastic Gradient Descent.....	19605
<i>Tomer Koren, Roi Livni, Yishay Mansour, Uri Sherman</i>	
A Geometric Perspective on Variational Autoencoders	19618
<i>Stephanie Allasonniere, Clément Chadebec</i>	
A Data-Augmentation is Worth a Thousand Samples: Analytical Moments and Sampling-Free Training	19631
<i>Randall Balestriero, Yann Lecun, Ishan Misra</i>	
Effective Adaptation in Multi-Task Co-Training for Unified Autonomous Driving	19645
<i>Jianhua Han, Xiaodan Liang, Xiwen Liang, Yangxin Wu, Hang Xu, Chunjing Xu</i>	
Redundant Representations Help Generalization in Wide Neural Networks.....	19659
<i>Diego Doimo, Aldo Glielmo, Sebastian Goldt, Alessandro Laio</i>	
Mean Estimation in High-Dimensional Binary Markov Gaussian Mixture Models	19673
<i>Nir Weinberger, Yihan Zhang</i>	
Rate-Distortion Theoretic Bounds on Generalization Error for Distributed Learning	19687
<i>Romain Chor, Milad Sefidgaran, Abdellatif Zaidi</i>	
Online Frank-Wolfe with Arbitrary Delays	19703
<i>Wei-Wei Tu, Yuanyu Wan, Lijun Zhang</i>	

Isometric 3D Adversarial Examples in the Physical World.....	19716
<i>Yinpeng Dong, Xiao-Shan Gao, Yibo Miao, Jun Zhu</i>	
Tight Lower Bounds on Worst-Case Guarantees for Zero-Shot Learning with Attributes	19732
<i>Stephen Bach, Alessio Mazzetto, Cristina Menghini, Eli Uppfal, Andrew Yuan</i>	
Graph Neural Networks with Adaptive Readouts.....	19746
<i>David Buterez, Jon Paul Janet, Steven J Kiddle, Pietro Liò, Dino Oglic</i>	
Robust Imitation of a Few Demonstrations with a Backwards Model	19759
<i>Jung Yeon Park, Lawson Wong</i>	
Communication Efficient Distributed Learning for Kernelized Contextual Bandits.....	19773
<i>Chuanhao Li, Hongning Wang, Huazheng Wang, Mengdi Wang</i>	
Trust Region Policy Optimization with Optimal Transport Discrepancies: Duality and Algorithm for Continuous Actions.....	19786
<i>Florian Dorfler, Nicolas Lanzetti, Giorgia Ramponi, Antonio Terpin, Batuhan Yardim</i>	
ELIAS: End-To-End Learning to Index and Search in Large Output Spaces.....	19798
<i>Patrick Chen, Inderjit Dhillon, Nilesh Gupta, Cho-Jui Hsieh, Hsiang-Fu Yu</i>	
GStarX: Explaining Graph Neural Networks with Structure-Aware Cooperative Games.....	19810
<i>Yozen Liu, Neil Shah, Yizhou Sun, Shichang Zhang</i>	
Bridging the Gap from Asymmetry Tricks to Decorrelation Principles in Non-Contrastive Self- Supervised Learning.....	19824
<i>Kang-Jun Liu, Takayuki Okatani, Masanori Suganuma</i>	
Lipschitz Bandits with Batched Feedback.....	19836
<i>Yasong Feng, Zengfeng Huang, Tianyu Wang</i>	
Improving Barely Supervised Learning by Discriminating Unlabeled Samples with Super-Class.....	19849
<i>Guan Gui, Lei Qi, Yinghuan Shi, Lei Wang, Zhen Zhao, Luping Zhou</i>	
Low-Rank Modular Reinforcement Learning Via Muscle Synergy	19861
<i>Heng Dong, Jiayuan Liu, Tonghan Wang, Chongjie Zhang</i>	
Neural Temporal Walks: Motif-Aware Representation Learning on Continuous-Time Dynamic Graphs	19874
<i>Ming Jin, Yuan-Fang Li, Shirui Pan</i>	
Understanding Benign Overfitting in Gradient-Based Meta Learning.....	19887
<i>Lisha Chen, Tianyi Chen, Songtao Lu</i>	
Generative Neural Articulated Radiance Fields.....	19900
<i>Alexander Bergman, Eric Chan, Petr Kellnhofer, David Lindell, Gordon Wetzstein, Wang Yifan</i>	
AutoMS: Automatic Model Selection for Novelty Detection with Error Rate Control.....	19917
<i>Dejing Dou, Haiyan Jiang, Haojie Ren, Yifan Zhang, Changliang Zou</i>	
GALOIS: Boosting Deep Reinforcement Learning Via Generalizable Logic Synthesis.....	19930
<i>Yushi Cao, Jianye Hao, Yi Li, Zhiming Li, Yang Liu, Tianpei Yang, Hao Zhang, Yan Zheng</i>	

Faster Deep Reinforcement Learning with Slower Online Network.....	19944
<i>Kavosh Asadi, Rasool Fakoor, Omer Gottesman, Taesup Kim, Michael Littman, Alexander Smola</i>	
Learn to Match with No Regret: Reinforcement Learning in Markov Matching Markets	19956
<i>Michael Jordan, Yifei Min, Zhaoran Wang, Tianhao Wang, Ruitu Xu, Zhuoran Yang</i>	
Efficient Frameworks for Generalized Low-Rank Matrix Bandit Problems	19971
<i>Cho-Jui Hsieh, Yue Kang, Thomas Chun Man Lee</i>	
A Projection-Free Algorithm for Constrained Stochastic Multi-Level Composition Optimization	19984
<i>Krishnakumar Balasubramanian, Saeed Ghadimi, Tesi Xiao</i>	
Green Hierarchical Vision Transformer for Masked Image Modeling	19997
<i>Lang Huang, Chen Qian, Fei Wang, Toshihiko Yamasaki, Shan You, Mingkai Zheng</i>	
Simultaneous Missing Value Imputation and Structure Learning with Groups.....	20011
<i>Miltiadis Allamanis, Wenbo Gong, Angus Lamb, Pablo Morales-Alvarez, Nick Pawlowski, Simon Peyton Jones, Simon Woodhead, Cheng Zhang</i>	
Towards Reliable Simulation-Based Inference with Balanced Neural Ratio Estimation	20025
<i>Arnaud Delaunoy, Joeri Hermans, Gilles Louppe, François Rozet, Antoine Wehenkel</i>	
Interaction Modeling with Multiplex Attention.....	20038
<i>Nick Haber, Isaac Kauvar, Mykel J Kochenderfer, Jiachen Li, Fan-Yun Sun, Jiajun Wu, Ruohan Zhang</i>	
Low-Rank Lottery Tickets: Finding Efficient Low-Rank Neural Networks Via Matrix Differential Equations.....	20051
<i>Gianluca Ceruti, Jonas Kusch, Steffen Schotthöfer, Francesco Tudisco, Emanuele Zangrando</i>	
Causality-Driven Hierarchical Structure Discovery for Reinforcement Learning	20064
<i>Ruizhi Chen, Yunji Chen, Zidong Du, Jiaming Guo, Qi Guo, Xing Hu, Ling Li, Shaohui Peng, Ke Tang, Qi Yi, Xishan Zhang, Rui Zhang</i>	
Pay Attention to Your Loss : Understanding Misconceptions About 1-Lipschitz Neural Networks	20077
<i>Louis Béthune, Thibaut Boissin, Corentin Friedrich, Alberto Gonzalez Sanz, Franck Mamalet, Mathieu Serrurier</i>	
Large-Scale Retrieval for Reinforcement Learning.....	20092
<i>Arthur Guez, Peter Humphreys, Timothy Lillicrap, Laurent Sifre, Olivier Tieleman, Theophane Weber</i>	
Gradient Flow Dynamics of Shallow ReLU Networks for Square Loss and Orthogonal Inputs	20105
<i>Etienne Boursier, Nicolas Flammarion, Loucas Pillaud-Vivien</i>	
Weisfeiler and Leman Go Walking: Random Walk Kernels Revisited.....	20119
<i>Nils M. Kriege</i>	
Outsourcing Training Without Uploading Data Via Efficient Collaborative Open-Source Sampling.....	20133
<i>Junyuan Hong, Lingjuan Lyu, Michael Spranger, Jiayu Zhou</i>	
Multi-Agent Dynamic Algorithm Configuration.....	20147
<i>Miqing Li, Chao Qian, Jiacheng Xu, Ke Xue, Yang Yu, Lei Yuan, Zongzhang Zhang</i>	
TaSIL: Taylor Series Imitation Learning	20162
<i>Nikolai Matni, Daniel Pfommer, Stephen Tu, Thomas Zhang</i>	

Experimental Design for Linear Functionals in Reproducing Kernel Hilbert Spaces 20175
Andreas Krause, Mojmir Mutny

Continuous MDP Homomorphisms and Homomorphic Policy Gradient..... 20189
David Mejer, Prakash Panangaden, Doina Precup, Sahand Rezaei-Shoshtari, Rosie Zhao

VOLUME 27

Score-Based Generative Modeling Secretly Minimizes the Wasserstein Distance 20205
Ying Fan, Dohyun Kwon, Kangwook Lee

Uncertainty-Aware Reinforcement Learning for Risk-Sensitive Player Evaluation in Sports Game..... 20218
Guiliang Liu, Yudong Luo, Pascal Poupart, Oliver Schulte

End-To-End Algorithm Synthesis with Recurrent Networks: Extrapolation Without Overthinking 20232
Arpit Bansal, Eitan Borgnia, Zeyad Emam, Micah Goldblum, Tom Goldstein, Furong Huang, Avi Schwarzschild

Smooth Fictitious Play in Stochastic Games with Perturbed Payoffs and Unknown Transitions..... 20243
Lucas Baudin, Rida Laraki

OOD Link Prediction Generalization Capabilities of Message-Passing GNNs in Larger Test Graphs 20257
Gitta Kutyniok, Bruno Ribeiro, Yangze Zhou

Algorithms with Prediction Portfolios..... 20273
Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, Sergei Vassilvitskii

A Unified Hard-Constraint Framework for Solving Geometrically Complex PDEs..... 20287
Ze Cheng, Songming Liu, Hang Su, Chengyang Ying, Hao Zhongkai, Jun Zhu

Instability and Local Minima in GAN Training with Kernel Discriminators 20300
Evan Becker, Alyson Fletcher, Parthe Pandit, Sundeep Rangan

Versatile Multi-Stage Graph Neural Network for Circuit Representation 20313
Jianye Hao, Dong Li, Guojie Song, Shuwen Yang, Zhihao Yang, Yingxueff Zhang, Zhanguang Zhang

Non-Stationary Bandits Under Recharging Payoffs: Improved Planning with Sublinear Regret..... 20325
Constantine Caramanis, Orestis Papadigenopoulos, Sanjay Shakkottai

Optimal and Adaptive Monteiro-Svaiter Acceleration 20338
Yair Carmon, Danielle Hausler, Arun Jambulapati, Yujia Jin, Aaron Sidford

Chroma-VAE: Mitigating Shortcut Learning with Generative Classifiers 20351
Micah Goldblum, Polina Kirichenko, Andrew Wilson, Wanqian Yang

SparCL: Sparse Continual Learning on the Edge 20366
Jennifer Dy, Yifan Gong, Stratis Ioannidis, Tong Jian, Wei Niu, Bin Ren, Yanzhi Wang, Zifeng Wang, Geng Yuan, Zheng Zhan

Adaptively Exploiting d-Separators with Causal Bandits 20381
Blair Bilodeau, Dan Roy, Linbo Wang

Spectrum Random Masking for Generalization in Image-Based Reinforcement Learning 20393
Guangyao Chen, Yangru Huang, Peixi Peng, Yonghong Tian, Yifan Zhao

Contrastive Graph Structure Learning Via Information Bottleneck for Recommendation.....	20407
<i>Jian Liang, Di Liu, Fei Wang, Chunyu Wei</i>	
Pyramid Attention for Source Code Summarization	20421
<i>Lei Chai, Ming Li</i>	
SIREN: Shaping Representations for Detecting Out-Of-Distribution Objects.....	20434
<i>Xuefeng Du, Gabriel Gozum, Yixuan Li, Yifei Ming</i>	
CLOOB: Modern Hopfield Networks with InfoLOOB Outperform CLIP	20450
<i>Angela Bitto, Andreas Fürst, Sepp Hochreiter, Günter Klambauer, Michael Kopp, David Kreil, Johannes Lehner, Hubert Ramsauer, Elisabeth Rumetshofer, Fei Tang, Viet T. Tran</i>	
I2Q: A Fully Decentralized Q-Learning Algorithm.....	20469
<i>Jiechuan Jiang, Zongqing Lu</i>	
Planning to the Information Horizon of BAMDPs Via Epistemic State Abstraction	20482
<i>Dilip Arumugam, Satinder Singh</i>	
AutoST: Towards the Universal Modeling of Spatio-Temporal Sequences	20498
<i>Jianxin Li, Hui Xiong, Shuai Zhang, Haoyi Zhou</i>	
Constrained Langevin Algorithms with L-Mixing External Random Variables.....	20511
<i>Andrew Lamperski, Yuping Zheng</i>	
Language Conditioned Spatial Relation Reasoning for 3D Object Grounding	20522
<i>Shizhe Chen, Pierre-Louis Guhur, Ivan Laptev, Cordelia Schmid, Makarand Tapaswi</i>	
MVP-N: A Dataset and Benchmark for Real-World Multi-View Object Classification.....	20536
<i>Tae Sung Kim, Jinsung Kim, Hyuk-Jae Lee, Jiayue Wang, Ren Wang</i>	
Data Augmentation for Compositional Data: Advancing Predictive Models of the Microbiome	20551
<i>John Cunningham, Elliott Gordon-Rodriguez, Thomas Quinn</i>	
On the Robustness of Deep Clustering Models: Adversarial Attacks and Defenses	20566
<i>Anshuman Chhabra, Prasant Mohapatra, Ashwin Sekhari</i>	
Exploiting the Relationship Between Kendall's Rank Correlation and Cosine Similarity for Attribution Protection.....	20580
<i>Adams Wai Kin Kong, Fan Wang</i>	
Wavelet Feature Maps Compression for Image-To-Image CNNs	20592
<i>Maor Ashkenazi, Shahaf E. Finder, Eran Treister, Yair Zohav</i>	
Biologically Plausible Solutions for Spiking Networks with Efficient Coding.....	20607
<i>Veronika Koren, Stefano Panzeri</i>	
Finding Optimal Arms in Non-Stochastic Combinatorial Bandits with Semi-Bandit Feedback and Finite Budget	20621
<i>Viktor Bengs, Jasmin Brandt, Björn Haddendorst, Eyke Hüllermeier</i>	
Graph Neural Networks Are Dynamic Programmers	20635
<i>Andrew J Dudzik, Petar Velickovic</i>	
Evaluated CMI Bounds for Meta Learning: Tightness and Expressiveness	20648
<i>Giuseppe Durisi, Fredrik Hellström</i>	

SPD: Synergy Pattern Diversifying Oriented Unsupervised Multi-Agent Reinforcement Learning.....	20661
<i>Shuncheng He, Xiangyang Ji, Yuhang Jiang, Jianzhun Shao, Hongchang Zhang</i>	
Self-Aware Personalized Federated Learning.....	20675
<i>Salman Avestimehr, Huili Chen, Jie Ding, Anit Kumar Sahu, Eric W. Tramel, Shuang Wu, Tao Zhang</i>	
Additive MIL: Intrinsically Interpretable Multiple Instance Learning for Pathology	20689
<i>Syed Ashar Javed, Dinkar Juyal, Harshith Padigela, Aaditya Prakash, Amaro Taylor-Weiner, Limin Yu</i>	
Model-Based Imitation Learning for Urban Driving.....	20703
<i>Roberto Cipolla, Gianluca Corrado, Nicolas Griffiths, Corina Gurau, Anthony Hu, Alex Kendall, Zachary Murez, Jamie Shotton, Hudson Yeo</i>	
Online Training Through Time for Spiking Neural Networks.....	20717
<i>Di He, Zhouchen Lin, Qingyan Meng, Mingqing Xiao, Zongpeng Zhang</i>	
SCONE: Surface Coverage Optimization in Unknown Environments by Volumetric Integration.....	20731
<i>Antoine Guedon, Vincent Lepetit, Pascal Monasse</i>	
WebShop: Towards Scalable Real-World Web Interaction with Grounded Language Agents	20744
<i>Howard Chen, Karthik Narasimhan, John Yang, Shunyu Yao</i>	
Boosting Out-Of-Distribution Detection with Typical Features.....	20758
<i>Yaowu Chen, Yuefeng Chen, Xiaodan Li, Xiang Tian, Chuanlong Xie, Hui Xue', Rong Zhang, Bolun Zheng, Yao Zhu</i>	
Dynamic Learning in Large Matching Markets.....	20770
<i>Anand Kalvit, Assaf Zeevi</i>	
Invariant and Transportable Representations for Anti-Causal Domain Shifts	20782
<i>Yibo Jiang, Victor Veitch</i>	
Symplectic Spectrum Gaussian Processes: Learning Hamiltonians from Noisy and Sparse Data	20795
<i>Tomoharu Iwata, Yusuke Tanaka, Naonori Ueda</i>	
Multiclass Learnability Beyond the PAC Framework: Universal Rates and Partial Concept Classes.....	20809
<i>Alkis Kalavasis, Amin Karbasi, Grigoris Velegkas</i>	
MetaTeacher: Coordinating Multi-Model Domain Adaptation for Medical Image Classification	20823
<i>Liuhan Peng, Liang Tian, Zhenbin Wang, Mao Ye, Xiatian Zhu, Yingying Zhu</i>	
Biological Learning of Irreducible Representations of Commuting Transformations.....	20838
<i>Dmitri Chklovskii, Alexander Genkin, Siavash Golkar, David Lipshutz, Tiberiu Tesileanu</i>	
A Fourier Approach to Mixture Learning.....	20850
<i>Kumar Avinava Dubey, Guru Guruganesh, Mingda Qiao, Ankit Rawat, Manzil Zaheer</i>	
Structured Energy Network as a Loss	20862
<i>Purujit Goyal, Jay Yoon Lee, Andrew McCallum, Dhruvesh Patel, Zhiyang Xu, Wenlong Zhao</i>	
Bayesian Inference Via Sparse Hamiltonian Flows.....	20876
<i>Trevor Campbell, Naitong Chen, Zuheng Xu</i>	
SAPA: Similarity-Aware Point Affiliation for Feature Upsampling.....	20889
<i>Zhiguo Cao, Hongtao Fu, Wenzhe Liu, Yuliang Liu, Hao Lu, Zixuan Ye</i>	

Losses Can Be Blessings: Routing Self-Supervised Speech Representations Towards Efficient Multilingual and Multitask Speech Processing	20902
<i>Yonggan Fu, Cheng-I Jeff Lai, Celine Lin, Kaizhi Qian, Zhifan Ye, Zhongzhi Yu, Yang Zhang</i>	
Gradient Methods Provably Converge to Non-Robust Networks	20921
<i>Ohad Shamir, Gal Vardi, Gilad Yehudai</i>	
Diversity Vs. Recognizability: Human-Like Generalization in One-Shot Generative Models	20933
<i>Victor Boutin, Thomas Serre, Lakshya Singhal, Xavier Thomas</i>	
Task-Level Differentially Private Meta Learning.....	20947
<i>Raef Bassily, Xinyu Zhou</i>	
On the Statistical Efficiency of Reward-Free Exploration in Non-Linear RL.....	20960
<i>Alekh Agarwal, Jinglin Chen, Nan Jiang, Akshay Krishnamurthy, Aditya Modi</i>	
Pessimism for Offline Linear Contextual Bandits Using ℓ_p Confidence Sets	20974
<i>Gene Li, Cong Ma, Nati Srebro</i>	
Discrete-Convex-Analysis-Based Framework for Warm-Starting Algorithms with Predictions.....	20988
<i>Taihei Oki, Shinsaku Sakaue</i>	

VOLUME 28

CAESAR: An Embodied Simulator for Generating Multimodal Referring Expression Datasets	21001
<i>Alexi Gladstone, Haley Green, Tariq Iqbal, Md Mofijul Islam, Reza Mirzaiee</i>	
Generalizing Bayesian Optimization with Decision-Theoretic Entropies.....	21016
<i>Stefano Ermon, Chenlin Meng, Willie Neiswanger, Lantao Yu, Shengjia Zhao</i>	
Ordered Subgraph Aggregation Networks.....	21030
<i>Floris Geerts, Christopher Morris, Mathias Niepert, Chendi Qian, Gaurav Rattan</i>	
Neuron with Steady Response Leads to Better Generalization	21046
<i>Xu Chen, Lun Du, Wei Fang, Qiang Fu, Shi Han, Haitao Mao, Dongmei Zhang</i>	
Laplacian Autoencoders for Learning Stochastic Representations.....	21059
<i>Søren Hauberg, Marco Miani, Pablo Moreno-Muñoz, Nicki Skafte, Frederik Warburg</i>	
Alleviating the Sample Selection Bias in Few-Shot Learning by Removing Projection to the Centroid.....	21073
<i>Yanan Li, Xu Luo, Xinglin Pan, Wenjie Pei, Jing Xu, Zenglin Xu</i>	
A Coupled Design of Exploiting Record Similarity for Practical Vertical Federated Learning	21087
<i>Bingsheng He, Qinbin Li, Zhaomin Wu</i>	
Cooperative Distribution Alignment Via JSD Upper Bound.....	21101
<i>Wonwoong Cho, Ziyu Gong, David Inouye</i>	
An Analytical Theory of Curriculum Learning in Teacher-Student Networks	21113
<i>Stefano Mannelli, Luca Saglietti, Andrew Saxe</i>	
RSA: Reducing Semantic Shift from Aggressive Augmentations for Self-Supervised Learning.....	21128
<i>Yingbin Bai, Cheng Deng, Yuxuan Du, Bo Han, Tongliang Liu, Dadong Wang, Zhaoqing Wang, Erkun Yang</i>	

Evaluating Latent Space Robustness and Uncertainty of EEG-ML Models Under Realistic Distribution Shifts	21142
<i>Brent M Berry, Samarth Rawal, Yogatheesan Varatharajah, Neeraj Wagh, Jionghao Wei</i>	
u-HuBERT: Unified Mixed-Modal Speech Pretraining and Zero-Shot Transfer to Unlabeled Modality	21157
<i>Wei-Ning Hsu, Bowen Shi</i>	
Hierarchical Graph Transformer with Adaptive Node Sampling.....	21171
<i>Qingyong Hu, Chee-Kong Lee, Qi Liu, Zaixi Zhang</i>	
Learning Options Via Compression.....	21184
<i>Benjamin Eysenbach, Chelsea Finn, Yiding Jiang, J. Zico Kolter, Evan Liu</i>	
Fixed-Distance Hamiltonian Monte Carlo.....	21200
<i>Sally Cripps, Hadi Mohasel Afshar</i>	
GlanceNets: Interpretable, Leak-Proof Concept-Based Models.....	21212
<i>Emanuele Marconato, Andrea Passerini, Stefano Teso</i>	
3DOS: Towards 3D Open Set Learning - Benchmarking and Understanding Semantic Novelty Detection on Point Clouds.....	21228
<i>Antonio Alliegro, Francesco Cappio Borlino, Tatiana Tommasi</i>	
Masked Prediction: A Parameter Identifiability View	21241
<i>Daniel Hsu, Bingbin Liu, Pradeep Ravikumar, Andrej Risteski</i>	
Self-Supervised Learning of Brain Dynamics from Broad Neuroimaging Data	21255
<i>Russell Poldrack, Christopher Ré, Armin Thomas</i>	
Characterization of Excess Risk for Locally Strongly Convex Population Risk.....	21270
<i>Zhi-Ming Ma, Ruoyu Wang, Mingyang Yi</i>	
A Non-Asymptotic Moreau Envelope Theory for High-Dimensional Generalized Linear Models	21286
<i>Frederic Koehler, Nati Srebro, Pragya Sur, Danica J. Sutherland, Lijia Zhou</i>	
Towards Efficient 3D Object Detection with Knowledge Distillation	21300
<i>Runyu Ding, Xiaojuan Qi, Shaoshuai Shi, Zhe Wang, Jihan Yang</i>	
CodeRL: Mastering Code Generation Through Pretrained Models and Deep Reinforcement Learning	21314
<i>Akhilesh Deepak Gotmare, Steven Chu Hong Hoi, Hung Le, Silvio Savarese, Yue Wang</i>	
Leveraging the Hints: Adaptive Bidding in Repeated First-Price Auctions	21329
<i>Aaron Flores, Yanjun Han, Tsachy Weissman, Wei Zhang, Zhengyuan Zhou</i>	
Sample Efficiency Matters: A Benchmark for Practical Molecular Optimization.....	21342
<i>Connor Coley, Tianfan Fu, Wenhao Gao, Jimeng Sun</i>	
MGNNI: Multiscale Graph Neural Networks with Implicit Layers	21358
<i>Bryan Hooi, Kenji Kawaguchi, Juncheng Liu, Xiaokui Xiao</i>	
UQGAN: A Unified Model for Uncertainty Quantification of Deep Classifiers Trained Via Conditional GANs.....	21371
<i>Gernot Fink, Philipp Oberdiek, Matthias Rottmann</i>	

Audio-Driven Co-Speech Gesture Video Generation.....	21386
<i>Yuanqi Du, Dahua Lin, Xian Liu, Ziwei Liu, Qianyi Wu, Wayne Wu, Hang Zhou</i>	
Concept Embedding Models: Beyond the Accuracy-Explainability Trade-Off.....	21400
<i>Pietro Barbiero, Gabriele Ciravegna, Michelangelo Diligenti, Mateo Espinosa Zarlenga, Francesco Giannini, Mateja Jamnik, Pietro Lió, Giuseppe Marra, Stefano Melacci, Frederic Precioso, Zohreh Shams, Adrian Weller</i>	
DENSE: Data-Free One-Shot Federated Learning.....	21414
<i>Chen Chen, Shouhong Ding, Bo Li, Lingjuan Lyu, Chunhua Shen, Shuang Wu, Chao Wu, Jie Zhang</i>	
Searching for Better Spatio-Temporal Alignment in Few-Shot Action Recognition.....	21429
<i>Yichao Cao, Xiaobo Lu, Xiu Su, Qingfei Tang, Chang Xu, Shan You</i>	
Fine-Tuning Pre-Trained Language Models Effectively by Optimizing Subnetworks Adaptively	21442
<i>Zhi Jin, Ge Li, Jia Li, Haojie Zhang, Zhongjin Zhang, Yuqi Zhu</i>	
Quality Not Quantity: On the Interaction Between Dataset Design and Robustness of CLIP	21455
<i>Gabriel Ilharco, Thao Nguyen, Sewoong Oh, Ludwig Schmidt, Mitchell Wortsman</i>	
Stars: Tera-Scale Graph Building for Clustering and Learning.....	21470
<i>CJ Carey, Jonathan Halcrow, Rajesh Jayaram, Vahab Mirrokni, Warren Schudy, Peilin Zhong</i>	
Score-Based Diffusion Meets Annealed Importance Sampling.....	21482
<i>Arnaud Doucet, Will Grathwohl, Alexander Matthews, Heiko Strathmann</i>	
PaCo: Parameter-Compositional Multi-Task Reinforcement Learning.....	21495
<i>Lingfeng Sun, Masayoshi Tomizuka, Wei Xu, Haichao Zhang</i>	
Entropy-Driven Mixed-Precision Quantization for Deep Network Design.....	21508
<i>Hesen Chen, Ce Ge, Hao Li, Ming Lin, Xiuyu Sun, Zhenhong Sun, Junyan Wang</i>	
Non-Gaussian Tensor Programs	21521
<i>Eugene Golikov, Greg Yang</i>	
Adaptation Accelerating Sampling-Based Bayesian Inference in Attractor Neural Networks.....	21534
<i>Tianhao Chu, Xingsi Dong, Tiejun Huang, Zilong Ji, Si Wu, Wenhao Zhang</i>	
A Contrastive Framework for Neural Text Generation.....	21548
<i>Nigel Collier, Lingpeng Kong, Tian Lan, Yixuan Su, Yan Wang, Dani Yogatama</i>	
Exploring the Latent Space of Autoencoders with Interventional Assays	21562
<i>Stefan Bauer, Michel Besserve, Felix Leeb, Bernhard Schölkopf</i>	
Pythae: Unifying Generative Autoencoders in Python - A Benchmarking Use Case	21575
<i>Stephanie Allasonniere, Clément Chadebec, Louis Vincent</i>	
Adaptive Multi-Stage Density Ratio Estimation for Learning Latent Space Energy-Based Model.....	21590
<i>Tian Han, Zhisheng Xiao</i>	
An Analysis of Ensemble Sampling	21602
<i>Xiuyuan Lu, Chao Qin, Benjamin Van Roy, Zheng Wen</i>	
Fair Wrapping for Black-Box Predictions	21615
<i>Ibrahim Alabdulmohsin, Sanmi Koyejo, Yishay Mansour, Nyalleng Moorosi, Richard Nock, Alexander Soen, Ke Sun, Lexing Xie</i>	

Bridging Central and Local Differential Privacy in Data Acquisition Mechanisms.....	21628
<i>Alireza Fallah, Ali Makhdoumi, Azarakhsh Malekian, Asuman Ozdaglar</i>	
Meta-Learning Dynamics Forecasting Using Task Inference.....	21640
<i>Robin Walters, Rui Wang, Rose Yu</i>	
Multi-Fidelity Monte Carlo: A Pseudo-Marginal Approach.....	21654
<i>Ryan Adams, Diana Cai</i>	
SAPD+: An Accelerated Stochastic Method for Nonconvex-Concave Minimax Problems.....	21668
<i>Necdet Serhat Aybat, Mert Gurbuzbalaban, Xuan Zhang</i>	
Contrastive Adapters for Foundation Model Group Robustness	21682
<i>Christopher Ré, Michael Zhang</i>	
Decoupled Context Processing for Context Augmented Language Modeling	21698
<i>Ruiqi Guo, Sanjiv Kumar, Zonglin Li</i>	
LISA: Learning Interpretable Skill Abstractions from Language.....	21711
<i>Stefano Ermon, Divyansh Garg, Kuno Kim, Jiaming Song, Skanda Vaidyanath</i>	
Accelerated Primal-Dual Gradient Method for Smooth and Convex-Concave Saddle-Point Problems with Bilinear Coupling	21725
<i>Alexander Gasnikov, Dmitry Kovalev, Peter Richtarik</i>	
Diffusion Curvature for Estimating Local Curvature in High Dimensional Data	21738
<i>Ian Adelstein, Dhananjay Bhaskar, Oluwadamilola Fasina, Smita Krishnaswamy, Kincaid Macdonald, Bastian Rieck, Dawson Thomas</i>	
Hidden Progress in Deep Learning: SGD Learns Parities Near the Computational Limit	21750
<i>Boaz Barak, Benjamin Edelman, Surbhi Goel, Sham Kakade, Eran Malach, Cyril Zhang</i>	
One Positive Label is Sufficient: Single-Positive Multi-Label Learning with Label Enhancement.....	21765
<i>Xin Geng, Jiaqi Lv, Congyu Qiao, Ning Xu, Min-Ling Zhang</i>	

VOLUME 29

Communication Acceleration of Local Gradient Methods Via an Accelerated Primal-Dual Algorithm with an Inexact Prox	21777
<i>Dmitry Kovalev, Peter Richtarik, Abdurakhmon Sadiev</i>	
Aligning Individual Brains with Fused Unbalanced Gromov Wasserstein.....	21792
<i>Nicolas Courty, Stanislas Dehaene, Rémi Flamary, Bertrand Thirion, Alexis Thual, Quang Huy Tran, Tatiana Zemsanova</i>	
This is the Way: Designing and Compiling LEPISZCZE, a Comprehensive NLP Benchmark for Polish.....	21805
<i>Lukasz Augustyniak, Roman Bartusiak, Denis Janiak, Arkadiusz Janz, Tomasz Kajdanowicz, Mikolaj Morzy, Maciej Piasecki, Albert Sawczyn, Piotr Szymanski, Adrian Szymczak, Kamil Tagowski, Marcin Watroba</i>	
An In-Depth Study of Stochastic Backpropagation.....	21819
<i>Hao Chen, Jun Fang, Bing Shuai, Joseph Tighe, Zhuowen Tu, Mingze Xu</i>	

Transformer Memory as a Differentiable Search Index	21831
<i>Dara Bahri, William Cohen, Mostafa Dehghani, Jai Gupta, Kai Hui, Harsh Mehta, Donald Metzler, Jianmo Ni, Zhen Qin, Tal Schuster, Yi Tay, Vinh Tran, Zhe Zhao</i>	
Insights into Pre-Training Via Simpler Synthetic Tasks	21844
<i>Felix Li, Percy Liang, Yuhuai Wu</i>	
Last-Iterate Convergence of Optimistic Gradient Method for Monotone Variational Inequalities.....	21858
<i>Gauthier Gidel, Eduard Gorbunov, Adrien Taylor</i>	
3DILG: Irregular Latent Grids for 3D Generative Modeling	21871
<i>Matthias Niessner, Peter Wonka, Biao Zhang</i>	
Policy Optimization for Markov Games: Unified Framework and Faster Convergence.....	21886
<i>Yu Bai, Na Li, Qinghua Liu, Huan Wang, Caiming Xiong, Runyu Zhang</i>	
C2FAR: Coarse-To-Fine Autoregressive Networks for Precise Probabilistic Forecasting	21900
<i>Shane Bergsma, Lei Guo, Javad Rahimipour Anaraki, Tim Zeyl</i>	
METS-CoV: A Dataset of Medical Entity and Targeted Sentiment on COVID-19 Related Tweets.....	21916
<i>Dading Chong, Zhijiang Guo, Yining Hua, Zichang Su, Zhiyang Teng, Zeqiang Wang, Jiageng Wu, Jie Yang, Peilin Zhou</i>	
Respecting Transfer Gap in Knowledge Distillation	21933
<i>Long Chen, Yulei Niu, Hanwang Zhang, Chang Zhou</i>	
Target Alignment in Truncated Kernel Ridge Regression	21948
<i>Arash Amini, Richard Baumgartner, Dai Feng</i>	
Continual Learning in Environments with Polynomial Mixing Times.....	21961
<i>Ignacio Cases, Maximilian Puelma Touzel, Sharath Chandra Raparthy, Matthew Riemer, Irina Rish, Gopeshh Subbaraj</i>	
ENS-10: A Dataset for Post-Processing Ensemble Weather Forecasts.....	21974
<i>Saleh Ashkboos, Tal Ben-Nun, Nikoli Dryden, Peter Dueben, Lukas Gianinazzi, Torsten Hoefler, Langwen Huang, Luca Kummer</i>	
Panchromatic and Multispectral Image Fusion Via Alternating Reverse Filtering Network.....	21988
<i>Danfeng Hong, Jie Huang, Chongyi Li, Chengjun Xie, Keyu Yan, Feng Zhao, Man Zhou</i>	
Unsupervised Cross-Task Generalization Via Retrieval Augmentation	22003
<i>Bill Yuchen Lin, Chris Miller, Xiang Ren, Kangmin Tan, Beiwen Tian</i>	
Natural Gradient Enables Fast Sampling in Spiking Neural Networks	22018
<i>J. Patrick Connor, Paul Masset, Venkatesh Murthy, Cengiz Pehlevan, Jacob Zavatone-Veth</i>	
Adversarial Auto-Augment with Label Preservation: A Representation Learning Principle Guided Approach	22035
<i>Fengxiang He, Furong Huang, Jiahao Su, Yanchao Sun, Dacheng Tao, Xinmei Tian, Kaiwen Yang, Tianyi Zhou</i>	
Coordinate Linear Variance Reduction for Generalized Linear Programming.....	22049
<i>Jelena Diakonikolas, Cheuk Yin Lin, Chaobing Song, Stephen Wright</i>	
Probabilistic Missing Value Imputation for Mixed Categorical and Ordered Data	22064
<i>Alex Townsend, Madeleine Udell, Yuxuan Zhao</i>	

Gaussian Copula Embeddings.....	22078
<i>Chien Lu, Jaakko Peltonen</i>	
SGAM: Building a Virtual 3D World Through Simultaneous Generation and Mapping	22090
<i>Wei-Chiu Ma, Yuan Shen, Shenlong Wang</i>	
Finding Naturally Occurring Physical Backdoors in Image Datasets.....	22103
<i>Emilio Andere, Arjun Nitin Bhagoji, Roma Bhattacharjee, Josephine Passananti, Emily Wenger, Ben Zhao, Heather Zheng</i>	
Unsupervised Representation Learning from Pre-Trained Diffusion Probabilistic Models	22117
<i>Zhijie Lin, Zijian Zhang, Zhou Zhao</i>	
Learning Causally Invariant Representations for Out-Of-Distribution Generalization on Graphs.....	22131
<i>Yatao Bian, Yongqiang Chen, James Cheng, Bo Han, Ma Kaili, Tongliang Liu, Binghui Xie, Han Yang, Yonggang Zhang</i>	
To Update Or Not to Update? Neurons at Equilibrium in Deep Models.....	22149
<i>Andrea Bragagnolo, Marco Grangetto, Enzo Tartaglione</i>	
Non-Convex Online Learning Via Algorithmic Equivalence	22161
<i>Udaya Ghai, Elad Hazan, Zhou Lu</i>	
TA-MoE: Topology-Aware Large Scale Mixture-Of-Expert Training	22173
<i>Chang Chen, Min Li, Zihua Wu, Chao Yang, Dianhai Yu</i>	
A Combinatorial Perspective on the Optimization of Shallow ReLU Networks.....	22187
<i>Michael S Matena, Colin Raffel</i>	
Large Language Models Are Zero-Shot Reasoners.....	22199
<i>Shixiang Shane Gu, Yusuke Iwasawa, Takeshi Kojima, Yutaka Matsuo, Machel Reid</i>	
Trading off Utility, Informativeness, and Complexity in Emergent Communication.....	22214
<i>Roger Levy, Julie Shah, Mycal Tucker, Noga Zaslavsky</i>	
FiLM-Ensemble: Probabilistic Deep Learning Via Feature-Wise Linear Modulation	22229
<i>Alexander Becker, Bernd Bischl, Stefano D'Aronco, Rodrigo Caye Daudt, Hüseyin Anil Gündüz, Mina Rezaei, Konrad Schindler, Mehmet Ozgur Turkoglu, Jan Wegner</i>	
Meta-DMoE: Adapting to Domain Shift by Meta-Distillation from Mixture-Of-Experts.....	22243
<i>Zhixiang Chi, Li Gu, Jin Tang, Yang Wang, Yuanhao Yu, Tao Zhong</i>	
Beyond Time-Average Convergence: Near-Optimal Uncoupled Online Learning Via Clairvoyant Multiplicative Weights Update	22258
<i>Georgios Piliouras, Ryann Sim, Stratis Skoulakis</i>	
Meta-Reward-Net: Implicitly Differentiable Reward Learning for Preference-Based Reinforcement Learning	22270
<i>Fengshuo Bai, Yali Du, Runze Liu, Yaodong Yang</i>	
One-Shot Neural Backdoor Erasing Via Adversarial Weight Masking	22285
<i>Shuwen Chai, Jinghui Chen</i>	
Revisiting Neural Scaling Laws in Language and Vision.....	22300
<i>Ibrahim Alabdulmohsin, Behnam Neyshabur, Xiaohua Zhai</i>	

Efficient Phi-Regret Minimization in Extensive-Form Games Via Online Mirror Descent.....	22313
<i>Yu Bai, Chi Jin, Song Mei, Ziang Song, Tiancheng Yu</i>	
Long Range Graph Benchmark.....	22326
<i>Dominique Beaini, Vijay Prakash Dwivedi, Michael Galkin, Anh Tuan Luu, Ali Parviz, Ladislav Rampásek, Guy Wolf</i>	
On the Inability of Gaussian Process Regression to Optimally Learn Compositional Functions.....	22341
<i>Matteo Giordano, Kolyan Ray, Johannes Schmidt-Hieber</i>	
Active Learning Through a Covering Lens.....	22354
<i>Avihu Dekel, Guy Hacohen, Daphna Weinshall, Ofer Yehuda</i>	
Uplifting Bandits.....	22368
<i>Yu-Guan Hsieh, Shiva Kasiviswanathan, Branislav Kveton</i>	
Training Uncertainty-Aware Classifiers with Conformalized Deep Learning.....	22380
<i>Bat-Sheva Einbinder, Yaniv Romano, Matteo Sesia, Yanfei Zhou</i>	
A General Approximation Lower Bound in L^p Norm, with Applications to Feed-Forward Neural Networks.....	22396
<i>El Mehdi Achour, Armand Foucault, Sébastien Gerchinovitz, François Malgouyres</i>	
EnvPool: A Highly Parallel Reinforcement Learning Environment Execution Engine.....	22409
<i>Shengyi Huang, Yukun Jiang, Min Lin, Bo Liu, Zichen Liu, Ting Luo, Denys Makoviichuk, Viktor Makoviychuk, Yufan Song, Jiayi Weng, Zhongwen Xu, Shuicheng Yan</i>	
Generative Visual Prompt: Unifying Distributional Control of Pre-Trained Generative Models.....	22422
<i>Fernando De La Torre, Saman Motamed, Shaunak Srivastava, Chen Henry Wu</i>	
Implicit Warping for Animation with Image Sets.....	22438
<i>Ming-Yu Liu, Arun Mallya, Ting-Chun Wang</i>	
FNeVR: Neural Volume Rendering for Face Animation.....	22451
<i>Dapeng Chen, Hong Li, Boyu Liu, Jianzhuang Liu, Xuhui Liu, Wei Peng, Bohan Zeng, Baochang Zhang</i>	
Ontologue: Declarative Benchmark Construction for Ontological Multi-Label Classification.....	22463
<i>Bernease Herman, Bill Howe, Sean Yang</i>	
Make an Omelette with Breaking Eggs: Zero-Shot Learning for Novel Attribute Synthesis.....	22477
<i>Tzu-Yin Chao, Pin-Yu Chen, Wei-Chen Chiu, Ching-Chun Huang, Yu-Hsuan Li</i>	
Exploiting Semantic Relations for Glass Surface Detection.....	22490
<i>Rynson Lau, Jiaying Lin, Yuen-Hei Yeung</i>	
Differentially Private Generalized Linear Models Revisited.....	22505
<i>Raman Arora, Raef Bassily, Cristóbal Guzmán, Michael Menart, Enayat Ullah</i>	
GREED: A Neural Framework for Learning Graph Distance Functions.....	22518
<i>Venkatesan Chakaravarthy, Siddharth Grover, Sourav Medya, Rishabh Ranjan, Sayan Ranu, Yogish Sabharwal</i>	
Domain Adaptation Under Open Set Label Shift.....	22531
<i>Sivaraman Balakrishnan, Saurabh Garg, Zachary Lipton</i>	

Efficient Adversarial Training Without Attacking: Worst-Case-Aware Robust Reinforcement Learning 22547
Furong Huang, Yongyuan Liang, Yanchao Sun, Ruijie Zheng

Alleviating "Posterior Collapse" in Deep Topic Models Via Policy Gradient 22562
Bo An, Bo Chen, Zhibin Duan, Yewen Li, Chaojie Wang, Dongsheng Wang, Mingyuan Zhou

VOLUME 30

Tracking Functional Changes in Nonstationary Signals with Evolutionary Ensemble Bayesian Model for Robust Neural Decoding 22576
Gang Pan, Yu Qi, Yueming Wang, Xinyun Zhu

Zeroth-Order Hard-Thresholding: Gradient Error Vs. Expansivity..... 22589
William De Vazelhes, Bin Gu, Huimin Wu, Xiaotong Yuan, Hualin Zhang

Collaborative Linear Bandits with Adversarial Agents: Near-Optimal Regret Bounds..... 22602
Arman Adibi, Hamed Hassani, Aritra Mitra, George J. Pappas

Differentially Private Graph Learning Via Sensitivity-Bounded Personalized PageRank 22617
Alessandro Epasto, Vahab Mirrokni, Bryan Perozzi, Anton Tsitsulin, Peilin Zhong

How Well Do Unsupervised Learning Algorithms Model Human Real-Time and Life-Long Learning?..... 22628
Yoon Bai, James J Dicarolo, Xiaoxuan Jia, Kenneth Norman, Nicholas Turk-Browne, Ziyu Xiang, Dan Yamins, Chengxu Zhuang

Stochastic Multiple Target Sampling Gradient Descent..... 22643
Nhat Ho, Trung Le, Hoang Phan, Dinh Phung, Ngoc Tran, Toan Tran

Towards Out-Of-Distribution Sequential Event Prediction: A Causal Treatment..... 22656
Liang Sun, Qingsong Wen, Qitian Wu, Junchi Yan, Chenxiao Yang, Zhiqiang Zhou

Hybrid Neural Autoencoders for Stimulus Encoding in Visual and Other Sensory Neuroprostheses 22671
Michael Beyeler, Jacob Granley, Lucas Relic

Tractable Function-Space Variational Inference in Bayesian Neural Networks..... 22686
Zonghao Chen, Yarin Gal, Tim G. J. Rudner, Yee Whye Teh

BMU-MoCo: Bidirectional Momentum Update for Continual Video-Language Modeling 22699
Zhao Cao, Nanyi Fei, Yizhao Gao, Hao Jiang, Yijie Li, Haoyu Lu, Zhiwu Lu

Can Hybrid Geometric Scattering Networks Help Solve the Maximum Clique Problem? 22713
Yimeng Min, Michael Perlmutter, Frederik Wenkel, Guy Wolf

Expected Improvement for Contextual Bandits..... 22725
Sunil Gupta, Santu Rana, Long Tran-Thanh, Hung Tran-The, Tuan Truong, Svetha Venkatesh

Redistribution of Weights and Activations for AdderNet Quantization..... 22739
Haikang Diao, Kai Han, Chuanjian Liu, Ying Nie, Yunhe Wang, Enhua Wu

Physically-Based Face Rendering for NIR-VIS Face Recognition 22752
Jiankang Deng, Jungong Han, Alexandros Lattas, Yunqi Miao, Stefanos Zafeiriou

DGraph: A Large-Scale Financial Dataset for Graph Anomaly Detection	22765
<i>Lei Chen, Xuanwen Huang, Michalis Vazirgiannis, Yang Wang, Chunping Wang, Jiarong Xu, Yang Yang, Zhisheng Zhang</i>	
Globally Convergent Policy Search for Output Estimation.....	22778
<i>Juan Perdomo, Max Simchowitz, Russ Tedrake, Jack Umenberger, Kaiqing Zhang</i>	
Escaping Saddle Points for Effective Generalization on Class-Imbalanced Data	22791
<i>Sumukh K Aithal, Mayank Mishra, Venkatesh Babu R, Harsh Rangwani</i>	
Continuous Deep Q-Learning in Optimal Control Problems: Normalized Advantage Functions Analysis	22806
<i>Stepan Martyanov, Anton Plaksin</i>	
Partial Identification of Treatment Effects with Implicit Generative Models.....	22816
<i>Vahid Balazadeh, Rahul Krishnan, Vasilis Syrgkanis</i>	
Robust Model Selection and Nearly-Proper Learning for GMMs.....	22830
<i>Jerry Li, Allen Liu, Ankur Moitra</i>	
A ² : Efficient Automated Attacker for Boosting Adversarial Training	22844
<i>Shiwen Cui, Ming Gu, Yihua Huang, Changhua Meng, Weiqiang Wang, Zhuoer Xu, Zhenzhe Ying, Guanghui Zhu</i>	
Shape, Light, and Material Decomposition from Images Using Monte Carlo Rendering and Denoising	22856
<i>Jon Hasselgren, Nikolai Hofmann, Jacob Munkberg</i>	
Convergence for Score-Based Generative Modeling with Polynomial Complexity	22870
<i>Holden Lee, Jianfeng Lu, Yixin Tan</i>	
Kernel Interpolation with Sparse Grids	22883
<i>Cameron Musco, Daniel Sheldon, Mohit Yadav</i>	
Contrastive Language-Image Pre-Training with Knowledge Graphs.....	22895
<i>Dongchen Han, Gao Huang, Xuran Pan, Shiji Song, Tianzhu Ye</i>	
Reconstructing Training Data from Trained Neural Networks.....	22911
<i>Niv Haim, Michal Irani, Ohad Shamir, Gal Vardi, Gilad Yehudai</i>	
Hierarchical Agglomerative Graph Clustering in Poly-Logarithmic Depth	22925
<i>Laxman Dhulipala, David Eisenstat, Jakub Lacki, Vahab Mirrokni, Jessica Shi</i>	
On-Device Training Under 256KB Memory	22941
<i>Wei-Ming Chen, Chuang Gan, Song Han, Ji Lin, Wei-Chen Wang, Ligeng Zhu</i>	
Behavior Transformers: Cloning k Modes with One Stone.....	22955
<i>Ariuntuya Arty Altanzaya, Zichen Cui, Lerrel Pinto, Nur Muhammad Shafiullah</i>	
Performative Power	22969
<i>Moritz Hardt, Meena Jagadeesan, Celestine Mender-Dünner</i>	
Diagonal State Spaces Are as Effective as Structured State Spaces	22982
<i>Jonathan Berant, Albert Gu, Ankit Gupta</i>	

Deep Fourier Up-Sampling	22995
<i>Jinwei Gu, Jie Huang, Chongyi Li, Chen Change Loy, Deyu Meng, Hu Yu, Feng Zhao, Man Zhou</i>	
Generative Time Series Forecasting with Diffusion, Denoise, and Disentanglement	23009
<i>Dejing Dou, Yan Li, Xinjiang Lu, Yaqing Wang</i>	
The Missing Invariance Principle Found -- The Reciprocal Twin of Invariant Risk Minimization	23023
<i>Avinash Baidya, Dongsung Huh</i>	
Signal Recovery with Non-Expansive Generative Network Priors	23036
<i>Jorio Coccola</i>	
Towards Understanding the Mixture-Of-Experts Layer in Deep Learning	23049
<i>Zixiang Chen, Yihe Deng, Quanquan Gu, Yuanzhi Li, Yue Wu</i>	
Indicators of Attack Failure: Debugging and Improving Optimization of Adversarial Examples.....	23063
<i>Battista Biggio, Nicholas Carlini, Luca Demetrio, Ambra Demontis, Maura Pintor, Fabio Roli, Angelo Sotgiu</i>	
Beyond Accuracy: Generalization Properties of Bio-Plausible Temporal Credit Assignment Rules	23077
<i>Arna Ghosh, Guillaume Lajoie, Yuhan Helena Liu, Blake Richards, Eric Shea-Brown</i>	
Parameter Tuning and Model Selection in Optimal Transport with Semi-Dual Brenier Formulation.....	23098
<i>Adrien Vacher, Francois-Xavier Vialard</i>	
VITA: Video Instance Segmentation Via Object Token Association.....	23109
<i>Miran Heo, Sukjun Hwang, Seon Joo Kim, Joon-Young Lee, Seoung Wug Oh</i>	
Bridging the Gap: Unifying the Training and Evaluation of Neural Network Binary Classifiers	23121
<i>Kate Candon, Deyuan Li, Yofit Milkessa, Nathan Tsoi, Marynel Vázquez</i>	
Truncated Proposals for Scalable and Hassle-Free Simulation-Based Inference	23135
<i>Michael Deistler, Pedro Goncalves, Jakob H Macke</i>	
When to Update Your Model: Constrained Model-Based Reinforcement Learning	23150
<i>Fengxiang He, Wenbing Huang, Tianying Ji, Mingxuan Jing, Yu Luo, Fuchun Sun</i>	
Noise Attention Learning: Enhancing Noise Robustness by Gradient Scaling	23164
<i>Yang Bo, Wenbo He, Yangdi Lu</i>	
Iso-Dream: Isolating and Leveraging Noncontrollable Visual Dynamics in World Models	23178
<i>Minting Pan, Yunbo Wang, Xiaokang Yang, Xiangming Zhu</i>	
PointNeXt: Revisiting PointNet++ with Improved Training and Scaling Strategies.....	23192
<i>Mohamed Elhoseiny, Bernard Ghanem, Hasan Hammoud, Yuchen Li, Jinjie Mai, Houwen Peng, Guocheng Qian</i>	
Quantum Algorithms for Sampling Log-Concave Distributions and Estimating Normalizing Constants	23205
<i>Andrew M. Childs, Tongyang Li, Jin-Peng Liu, Chunhao Wang, Ruizhe Zhang</i>	
Physics-Embedded Neural Networks: Graph Neural PDE Solvers with Mixed Boundary Conditions.....	23218
<i>Masanobu Horie, Naoto Mitsume</i>	
Mismatched No More: Joint Model-Policy Optimization for Model-Based RL	23230
<i>Benjamin Eysenbach, Alexander Khazatsky, Sergey Levine, Russ Salakhutdinov</i>	

Phase Diagram of Stochastic Gradient Descent in High-Dimensional Two-Layer Neural Networks	23244
<i>Florent Krzakala, Bruno Loureiro, Ludovic Stephan, Rodrigo Veiga, Lenka Zdeborová</i>	
Constrained Stochastic Nonconvex Optimization with State-Dependent Markov Data.....	23256
<i>Krishnakumar Balasubramanian, Saeed Ghadimi, Abhishek Roy</i>	
Adapting Self-Supervised Vision Transformers by Probing Attention-Conditioned Masking Consistency	23271
<i>Judy Hoffman, Viraj Prabhu, Aaditya Singh, Sriram Yenamandra</i>	
MaskTune: Mitigating Spurious Correlations by Forcing to Explore	23284
<i>Saeid Asgari, Ali Gholami, Ghassan Hamarneh, Aliasghar Khani, Fereshte Khani, Ali Mahdavi Amiri, Linh Tran</i>	
A Dataset for Efforts Towards Achieving the Sustainable Development Goal of Safe Working Environments.....	23297
<i>Eirik Lund Flogard, Ole Jakob Mengshoel</i>	
Decomposing NeRF for Editing Via Feature Field Distillation.....	23311
<i>Sosuke Kobayashi, Eiichi Matsumoto, Vincent Sitzmann</i>	
Sample-Then-Optimize Batch Neural Thompson Sampling	23331
<i>Zhongxiang Dai, Patrick Jaillet, Bryan Kian Hsiang Low, Yao Shu</i>	
Video-Based Human-Object Interaction Detection from Tubelet Tokens	23345
<i>Xionghuo Min, Wei Shen, Wei Sun, Danyang Tu, Guangtao Zhai</i>	

VOLUME 31

Efficient and Stable Fully Dynamic Facility Location	23358
<i>Sayan Bhattacharya, Silvio Lattanzi, Nikos Parotsidis</i>	
MCVD - Masked Conditional Video Diffusion for Prediction, Generation, and Interpolation	23371
<i>Alexia Jolicoeur-Martineau, Chris Pal, Vikram Voleti</i>	
Addressing Leakage in Concept Bottleneck Models	23386
<i>Finale Doshi-Velez, Marton Havasi, Sonali Parbhoo</i>	
Adaptive Oracle-Efficient Online Learning	23398
<i>Jacob Abernethy, Zihao Hu, Vidya Muthukumar, Guanghui Wang</i>	
MoVQ: Modulating Quantized Vectors for High-Fidelity Image Generation	23412
<i>Jianfei Cai, Dinh Phung, Tung-Long Vuong, Chuanxia Zheng</i>	
Meta-Auto-Decoder for Solving Parametric Partial Differential Equations.....	23426
<i>Lei Chen, Haotian Chu, Bin Dong, Bei Hua, Xiang Huang, Shi Ji, Yang Li, Hongsheng Liu, Min Wang, Zidong Wang, Kang Yang, Zhanhong Ye, Fan Yu</i>	
Sharpness-Aware Training for Free	23439
<i>Jiawei Du, Jiashi Feng, Vincent Tan, Daquan Zhou, Joey Tianyi Zhou</i>	
Generalization Error Bounds on Deep Learning with Markov Datasets	23452
<i>Lan V. Truong</i>	

AirfRANS: High Fidelity Computational Fluid Dynamics Dataset for Approximating Reynolds-Averaged Navier–Stokes Solutions	23463
<i>Florent Bonnet, Paola Cinnella, Patrick Gallinari, Jocelyn Mazari</i>	
Generalization for Multiclass Classification with Overparameterized Linear Models.....	23479
<i>Rahul Arya, Anant Sahai, Vignesh Subramanian</i>	
Inception Transformer	23495
<i>Chenyang Si, Xinchao Wang, Shuicheng Yan, Weihao Yu, Pan Zhou, Yichen Zhou</i>	
ElasticMVS: Learning Elastic Part Representation for Self-Supervised Multi-View Stereopsis	23510
<i>Zheng Cao, Lu Fang, Ruqi Huang, Ruofan Tang, Jing Xiao, Jinzhi Zhang</i>	
Adaptive Stochastic Variance Reduction for Non-Convex Finite-Sum Minimization	23524
<i>Kimon Antonakopoulos, Volkan Cevher, Leello Tadesse Dadi, Ali Kavis, Stratis Skoulakis</i>	
Thinking Outside the Ball: Optimal Learning with Gradient Descent for Generalized Linear Stochastic Convex Optimization	23539
<i>Idan Amir, Roi Livni, Nati Srebro</i>	
Generalization Properties of NAS Under Activation and Skip Connection Search.....	23551
<i>Volkan Cevher, Grigorios Chrysos, Fanghui Liu, Zhenyu Zhu</i>	
Mesoscopic Modeling of Hidden Spiking Neurons.....	23566
<i>Guillaume Bellec, Wulfram Gerstner, Valentin Schmutz, Shuqi Wang</i>	
SageMix: Saliency-Guided Mixup for Point Clouds	23580
<i>Minkyu Jeon, Hyunwoo Kim, Injae Kim, Sanghyeok Lee, Yunyang Xiong</i>	
Denoising Diffusion Restoration Models	23593
<i>Michael Elad, Stefano Ermon, Bahjat Kawar, Jiaming Song</i>	
On Translation and Reconstruction Guarantees of the Cycle-Consistent Generative Adversarial Networks	23607
<i>Anish Chakrabarty, Swagatam Das</i>	
Adversarial Training with Complementary Labels: On the Benefit of Gradually Informative Attacks	23621
<i>Bo Han, Tongliang Liu, Gang Niu, Masashi Sugiyama, Jingfeng Zhang, Jianan Zhou, Jianing Zhu</i>	
Tractable Optimality in Episodic Latent MABs	23634
<i>Constantine Caramanis, Yonathan Efroni, Jeongyeol Kwon, Shie Mannor</i>	
A Characterization of Semi-Supervised Adversarially Robust PAC Learnability	23646
<i>Idan Attias, Steve Hanneke, Yishay Mansour</i>	
Data-IQ: Characterizing Subgroups with Heterogeneous Outcomes in Tabular Data.....	23660
<i>Ioana Bica, Jonathan Crabbé, Nabeel Seedat, Mihaela Van Der Schaar</i>	
Task-Free Continual Learning Via Online Discrepancy Distance Learning.....	23675
<i>Adrian G. Bors, Fei Ye</i>	
BinauralGrad: A Two-Stage Conditional Diffusion Probabilistic Model for Binaural Audio Synthesis.....	23689
<i>Jiawei Chen, Zehua Chen, Junliang Guo, Lei He, Yichong Leng, Xiangyang Li, Haohe Liu, Tie-Yan Liu, Danilo Mandic, Tao Qin, Xu Tan, Sheng Zhao</i>	

ConfLab: A Data Collection Concept, Dataset, and Benchmark for Machine Analysis of Free- Standing Social Interactions in the Wild	23701
<i>Ekin Gedik, Hayley Hung, Ashraf Islam, Chirag Raman, Stephanie Tan, Jose Vargas Quiros</i>	
Flamingo: A Visual Language Model for Few-Shot Learning.....	23716
<i>Jean-Baptiste Alayrac, Iain Barr, Ricardo Barreira, Mikolaj Binkowski, Sebastian Borgeaud, Andy Brock, Serkan Cabi, Jeff Donahue, Zhitao Gong, Tengda Han, Yana Hasson, Karel Lenc, Pauline Luc, Jacob L Menick, Arthur Mensch, Antoine Miech, Katherine Millican, Marianne Monteiro, Aida Nematzadeh, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Sina Samangooei, Sahand Sharifzadeh, Karén Simonyan, Oriol Vinyals, Andrew Zisserman</i>	
Understanding the Eluder Dimension.....	23737
<i>Dylan J Foster, Pritish Kamath, Gene Li, Nati Srebro</i>	
Predictive Querying for Autoregressive Neural Sequence Models	23751
<i>Alex Boyd, Stephan Mandt, Samuel Showalter, Padhraic Smyth</i>	
Learning State-Aware Visual Representations from Audible Interactions.....	23765
<i>Abhinav Gupta, Unnat Jain, Himangi Mittal, Pedro Morgado</i>	
Context-Based Dynamic Pricing with Partially Linear Demand Model.....	23780
<i>Jinzhi Bu, David Simchi-Levi, Chonghuan Wang</i>	
NeuroSchedule: A Novel Effective GNN-Based Scheduling Method for High-Level Synthesis.....	23792
<i>Mingyang Kou, Hailong Yao, Jun Zeng</i>	
Hand-Object Interaction Image Generation.....	23805
<i>Hezhen Hu, Houqiang Li, Weilun Wang, Wengang Zhou</i>	
DISCO: Adversarial Defense with Local Implicit Functions	23818
<i>Chih-Hui Ho, Nuno Vasconcelos</i>	
Distinguishing Discrete and Continuous Behavioral Variability Using Warped Autoregressive HMMs	23838
<i>Julia Costacurta, Sandeep R Datta, Lea Duncker, Winthrop Gillis, Scott Linderman, Jeffrey Markowitz, Blue Sheffer, Caleb Weinreb, Alex Williams</i>	
RORL: Robust Offline Reinforcement Learning Via Conservative Smoothing.....	23851
<i>Chenjia Bai, Lei Han, Xiaoteng Ma, Zhaoran Wang, Rui Yang, Chongjie Zhang</i>	
Optimal Scaling for Locally Balanced Proposals in Discrete Spaces.....	23867
<i>Hanjun Dai, Dale Schuurmans, Haoran Sun</i>	
The Impact of Task Underspecification in Evaluating Deep Reinforcement Learning	23881
<i>Vindula Jayawardana, Sirui Li, Dajiang Suo, Catherine Tang, Cathy Wu</i>	
Zero-Shot 3D Drug Design by Sketching and Generating	23894
<i>Xinyu Dai, Siyu Long, Hao Zhou, Yi Zhou</i>	
Decoupling Knowledge from Memorization: Retrieval-Augmented Prompt Learning	23908
<i>Huajun Chen, Xiang Chen, Shumin Deng, Fei Huang, Lei Li, Xiaozhuan Liang, Luo Si, Chuanqi Tan, Ningyu Zhang</i>	
Polynomial Time Guarantees for the Burer-Monteiro Method.....	23923
<i>Diego Cifuentes, Ankur Moitra</i>	

Optimal Comparator Adaptive Online Learning with Switching Cost	23936
<i>Ashok Cutkosky, Yannis Paschalidis, Zhiyu Zhang</i>	
A Robust Phased Elimination Algorithm for Corruption-Tolerant Gaussian Process Bandits	23951
<i>Ilija Bogunovic, Andreas Krause, Zihan Li, Jonathan Scarlett</i>	
Fair Rank Aggregation	23965
<i>Diptarka Chakraborty, Syamantak Das, Arindam Khan, Aditya Subramanian</i>	
Beyond IID: Data-Driven Decision-Making in Heterogeneous Environments	23979
<i>Omar Besbes, Will Ma, Omar Mouchtaki</i>	
Neur2SP: Neural Two-Stage Stochastic Programming	23992
<i>Merve Bodur, Justin Dumouchelle, Elias Khalil, Rahul Mihir Patel</i>	
SQ Lower Bounds for Learning Single Neurons with Massart Noise	24006
<i>Ilias Diakonikolas, Daniel Kane, Lisheng Ren, Yuxin Sun</i>	
MaskPlace: Fast Chip Placement Via Reinforced Visual Representation Learning	24019
<i>Yao Lai, Ping Luo, Yao Mu</i>	
List-Decodable Sparse Mean Estimation.....	24031
<i>Jie Shen, Shiwei Zeng</i>	
Structure-Aware Image Segmentation with Homotopy Warping.....	24046
<i>Xiaoling Hu</i>	
Positive-Unlabeled Learning Using Random Forests Via Recursive Greedy Risk Minimization.....	24060
<i>Ryan Ko, Abigail Koay, Jonathan Wilton, Miao Xu, Nan Ye</i>	
Extracting Computational Mechanisms from Neural Data Using Low-Rank RNNs	24072
<i>Srdjan Ostojic, Jonathan Pillow, Adrian Valente</i>	
Pluralistic Image Completion with Gaussian Mixture Models.....	24087
<i>Bo Han, Yewen Li, Tongliang Liu, Jie Ren, Xiaobo Xia, Wenhao Yang, Yibing Zhan</i>	
A Fast Post-Training Pruning Framework for Transformers	24101
<i>Amir Gholami, Joseph Hassoun, Kurt Keutzer, Sehoon Kim, Woosuk Kwon, Michael Mahoney</i>	
SurDis: A Surface Discontinuity Dataset for Wearable Technology to Assist Blind Navigation in Urban Environments.....	24117
<i>Kuan Yew Leong, Siew Mooi Lim</i>	
Interventions, Where and How? Experimental Design for Causal Models at Scale.....	24130
<i>Yashas Annadani, Stefan Bauer, Yarin Gal, Andrew Jesson, Bernhard Schölkopf, Panagiotis Tigas</i>	

VOLUME 32

Dual-Discriminative Graph Neural Network for Imbalanced Graph-Level Anomaly Detection	24144
<i>Charu Aggarwal, Leman Akoglu, Hao Peng, Quan Z. Sheng, Jianlin Su, Jia Wu, Shan Xue, Jian Yang, Zhenyu Yang, Ge Zhang, Chuan Zhou</i>	
Posted Pricing and Dynamic Prior-Independent Mechanisms with Value Maximizers.....	24158
<i>Yuan Deng, Vahab Mirrokni, Hanrui Zhang</i>	

Curious Exploration Via Structured World Models Yields Zero-Shot Object Manipulation.....	24170
<i>Sebastian Blaes, Georg Martius, Cansu Sancaktar</i>	
Estimating Noise Transition Matrix with Label Correlations for Noisy Multi-Label Learning	24184
<i>Shiming Ge, Shikun Li, Tongliang Liu, Xiaobo Xia, Yibing Zhan, Hansong Zhang</i>	
Active Learning Polynomial Threshold Functions.....	24199
<i>Omri Ben-Eliezer, Max Hopkins, Chutong Yang, Hantao Yu</i>	
Single-Phase Deep Learning in Cortico-Cortical Networks.....	24213
<i>Will Greedy, Jack Mellor, Joseph Pemberton, Rui Ponte Costa, Heng Wei Zhu</i>	
Domain Generalization by Learning and Removing Domain-Specific Features.....	24226
<i>Fang Chen, Yu Ding, Bin Liang, Shuming Liang, Lei Wang, Yang Wang</i>	
Torsional Diffusion for Molecular Conformer Generation.....	24240
<i>Regina Barzilay, Jeffrey Chang, Gabriele Corso, Tommi Jaakkola, Bowen Jing</i>	
LiteTransformerSearch: Training-Free Neural Architecture Search for Efficient Language Models.....	24254
<i>Sebastien Bubeck, Gustavo De Rosa, Debadeepta Dey, Mojan Javaheripi, Farinaz Koushanfar, Subhabrata Mukherjee, Tomasz Religa, Shital Shah, Caio Cesar Teodoro Mendes</i>	
AgraSSt: Approximate Graph Stein Statistics for Interpretable Assessment of Implicit Graph Generators	24268
<i>Gesine D Reinert, Wenkai Xu</i>	
On the Limitations of Stochastic Pre-Processing Defenses.....	24280
<i>Kassem Fawaz, Yue Gao, Nicolas Papernot, I Shumailov</i>	
Iterative Scene Graph Generation.....	24295
<i>Siddhesh Khandelwal, Leonid Sigal</i>	
Proximal Point Imitation Learning.....	24309
<i>Volkan Cevher, Angeliki Kamoutsi, Igor Krawczuk, Gergely Neu, Luca Viano</i>	
Heatmap Distribution Matching for Human Pose Estimation	24327
<i>Yujun Cai, Lin Geng Foo, Jun Liu, Haoxuan Qu, Li Xu</i>	
Mining Unseen Classes Via Regional Objectness: A Simple Baseline for Incremental Segmentation	24340
<i>Zhiyuan Fang, Guangyu Gao, Jianbo Jiao, Yunchao Wei, Zekang Zhang</i>	
Bezier Gaussian Processes for Tall and Wide Data	24354
<i>Martin Jørgensen, Michael A Osborne</i>	
Smoothed Embeddings for Certified Few-Shot Learning	24367
<i>Olesya Kuznetsova, Ivan Oseledets, Mikhail Pautov, Aleksandr Petiushko, Nurislam Tursynbek</i>	
Finite Sample Analysis of Dynamic Regression Parameter Learning	24380
<i>Koby Crammer, Mark Kozdoba, Shie Mannor, Edward Moroshko</i>	
Group Meritocratic Fairness in Linear Contextual Bandits.....	24392
<i>Arya Akhavan, Leonardo Cella, John Falk, Riccardo Grazi, Massimiliano Pontil</i>	
New Lower Bounds for Private Estimation and a Generalized Fingerprinting Lemma	24405
<i>Gautam Kamath, Argyris Mouzakis, Vikrant Singhal</i>	

Lower Bounds on Randomly Preconditioned Lasso Via Robust Sparse Designs.....	24419
<i>Jonathan Kelner, Frederic Koehler, Raghu Meka, Dhruv Rohatgi</i>	
Model-Based Safe Deep Reinforcement Learning Via a Constrained Proximal Policy Optimization Algorithm	24432
<i>Shalabh Bhatnagar, Ashish K Jayant</i>	
Exact Solutions of a Deep Linear Network	24446
<i>Botao Li, Xiangming Meng, Liu Ziyin</i>	
An Adaptive Kernel Approach to Federated Learning of Heterogeneous Causal Effects	24459
<i>Arnab Bhattacharyya, Young Lee, Tze-Yun Leong, Thanh Vinh Vo</i>	
FeLMi : Few Shot Learning with Hard Mixup.....	24474
<i>Rama Chellappa, Anoop Cherian, Prithviraj Dhar, Aniket Roy, Anshul Shah, Ketul Shah</i>	
EpiGRAF: Rethinking Training of 3D GANs	24487
<i>Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, Peter Wonka</i>	
Towards Trustworthy Automatic Diagnosis Systems by Emulating Doctors' Reasoning with Deep Reinforcement Learning.....	24502
<i>Arsene Fansi Tchango, Joumana Ghosn, Rishab Goel, Gaetan Marceau Caron, Julien Martel, Zhi Wen</i>	
Towards Improving Faithfulness in Abstractive Summarization.....	24516
<i>Xiuying Chen, Xin Gao, Mingzhe Li, Xiangliang Zhang</i>	
ZIN: When and How to Learn Invariance Without Environment Partition?	24529
<i>Peng Cui, Yong Lin, Lu Tan, Shengyu Zhu</i>	
Random Sharpness-Aware Minimization	24543
<i>Xiangning Chen, Minhao Cheng, Cho-Jui Hsieh, Yong Liu, Siqi Mai, Yang You</i>	
Active Surrogate Estimators: An Active Learning Approach to Label-Efficient Model Evaluation	24557
<i>Sebastian Farquhar, Yarin Gal, Jannik Kossen, Thomas Rainforth</i>	
HAPI: A Large-Scale Longitudinal Dataset of Commercial ML API Predictions.....	24571
<i>Lingjiao Chen, Evan Sabri Eyuboglu, Zhihua Jin, Christopher Ré, Matei Zaharia, James Zou</i>	
Near-Optimal Regret Bounds for Multi-Batch Reinforcement Learning	24586
<i>Xiangyang Ji, Yuhang Jiang, Zihan Zhang, Yuan Zhou</i>	
OST: Improving Generalization of DeepFake Detection Via One-Shot Test-Time Training	24597
<i>Liang Chen, Lingqiao Liu, Yibing Song, Jue Wang, Yong Zhang</i>	
The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games.....	24611
<i>Alexandre Bayen, Jiaxuan Gao, Akash Velu, Eugene Vinitsky, Yu Wang, Yi Wu, Chao Yu</i>	
Resolving the Data Ambiguity for Periodic Crystals.....	24625
<i>Vitaliy Kurlin, Daniel Widdowson</i>	
Video PreTraining (VPT): Learning to Act by Watching Unlabeled Online Videos	24639
<i>Ilge Akkaya, Bowen Baker, Jeff Clune, Adrien Ecoffet, Brandon Houghton, Joost Huizinga, Raul Sampedro, Jie Tang, Peter Zhokov</i>	

GLOBEM Dataset: Multi-Year Datasets for Longitudinal Human Behavior Modeling Generalization	24655
<i>Tim Althoff, Jennifer Brown, Anind Dey, Kevin Kuehn, Xin Liu, Jennifer Mankoff, Mike Merrill, Margaret Morris, Paula Nurius, Shwetak Patel, Yiyi Ren, Eve Riskin, Yasaman Sefidgar, Woosuk Seo, Xuhai Xu, Han Zhang</i>	
Shape and Structure Preserving Differential Privacy	24693
<i>Karthik Bharath, Matthew Reimherr, Aleksandra Slavkovic, Carlos Soto</i>	
Transformers Meet Stochastic Block Models: Attention with Data-Adaptive Sparsity and Cost.....	24706
<i>Sungjun Cho, Seunghoon Hong, Jinwoo Kim, Honglak Lee, Moontae Lee, Seonwoo Min</i>	
Characteristics of Harmful Text: Towards Rigorous Benchmarking of Language Models	24720
<i>Sumanth Dathathri, Iason Gabriel, Amelia Glaese, Lisa Anne Hendricks, Po-Sen Huang, Geoffrey Irving, William Isaac, John Mellor, Maribeth Rauh, Jonathan Uesato, Laura Weidinger, Johannes Welbl</i>	
APG: Adaptive Parameter Generation Network for Click-Through Rate Prediction.....	24740
<i>Hongbo Deng, Feng Li, Pengjie Wang, Jian Xu, Bencheng Yan, Kai Zhang, Bo Zheng</i>	
NeoRL: A Near Real-World Benchmark for Offline Reinforcement Learning	24753
<i>Xiong-Hui Chen, Songyi Gao, Zewen Li, Rong-Jun Qin, Yang Yu, Weinan Zhang, Xingyuan Zhang</i>	
Stochastic Halpern Iteration with Variance Reduction for Stochastic Monotone Inclusions.....	24766
<i>Xufeng Cai, Jelena Diakonikolas, Cristóbal Guzmán, Chaobing Song</i>	
SNN-RAT: Robustness-Enhanced Spiking Neural Network Through Regularized Adversarial Training	24780
<i>Tong Bu, Jianhao Ding, Tiejun Huang, Jian Liu, Zhaofei Yu</i>	
The Mechanism of Prediction Head in Non-Contrastive Self-Supervised Learning.....	24794
<i>Yuanzhi Li, Zixin Wen</i>	
Counterfactual Temporal Point Processes	24810
<i>Kimia Noorbakhsh, Manuel Rodriguez</i>	
Chain-Of-Thought Prompting Elicits Reasoning in Large Language Models.....	24824
<i>Maarten Bosma, Ed Chi, Brian Ichter, Quoc V Le, Dale Schuurmans, Xuezhi Wang, Jason Wei, Fei Xia, Denny Zhou</i>	
House of Cans: Covert Transmission of Internal Datasets Via Capacity-Aware Neuron Steganography.....	24838
<i>Xudong Pan, Yifan Yan, Min Yang, Mi Zhang, Shengyao Zhang</i>	
Polynomial-Time Optimal Equilibria with a Mediator in Extensive-Form Games	24851
<i>Tuomas Sandholm, Brian Zhang</i>	
Dungeons and Data: A Large-Scale NetHack Dataset.....	24864
<i>Eric Hambro, Heinrich Küttler, Vegard Mella, Naila Murray, Roberta Raileanu, Tim Rocktäschel, Danielle Rothermel</i>	
Optimal Dynamic Regret in LQR Control.....	24879
<i>Dheeraj Baby, Yu-Xiang Wang</i>	
Gradient Descent is Optimal Under Lower Restricted Secant Inequality and Upper Error Bound.....	24893
<i>Baptiste Goujaud, Charles Guille-Escuret, Adam Ibrahim, Ioannis Mitliagkas</i>	

GenSDF: Two-Stage Learning of Generalizable Signed Distance Functions..... 24905
Gene Chou, Ilya Chugunov, Felix Heide

Forecasting Human Trajectory from Scene History 24920
Xiran Cai, Terrence Chen, Mancheng Meng, Dinggang Shen, Ziyang Wu, Fan Yang, Xiang Zhou

VOLUME 33

Debiasing Graph Neural Networks Via Learning Disentangled Causal Substructure 24934
Shaohua Fan, Yanhu Mo, Chuan Shi, Jian Tang, Xiao Wang

Stochastic Online Learning with Feedback Graphs: Finite-Time and Asymptotic Optimality..... 24947
Teodor Vanislavov Marinov, Mehryar Mohri, Julian Zimmert

Asymptotics of ℓ_2 Regularized Network Embeddings 24960
Andrew Davison

Differentiable Hierarchical and Surrogate Gradient Search for Spiking Neural Networks 24975
Kaiwei Che, Jie Cheng, Qinghai Guo, Luziwei Leng, Jianxing Liao, Qinghu Meng, Jianguo Zhang, Kaixuan Zhang

On Embeddings for Numerical Features in Tabular Deep Learning..... 24991
Artem Babenko, Yury Gorishniy, Ivan Rubachev

Visual Prompting Via Image Inpainting 25005
Amir Bar, Trevor Darrell, Alexei Efros, Yossi Gandelsman, Amir Globerson

MonoSDF: Exploring Monocular Geometric Cues for Neural Implicit Surface Reconstruction..... 25018
Andreas Geiger, Michael Niemeyer, Songyou Peng, Torsten Sattler, Zehao Yu

OpenAUC: Towards AUC-Oriented Open-Set Recognition 25033
Xiaochun Cao, Yuan He, Qingming Huang, Zitai Wang, Qianqian Xu, Zhiyong Yang

Reduction Algorithms for Persistence Diagrams of Networks: CoralTDA and PrunIT 25046
Cuneyt G Akcora, Baris Coskunuzer, Yulia Gel, Murat Kantarcioglu

TreeMoCo: Contrastive Neuron Morphology Representation Learning 25060
Hanbo Chen, Lei He, Daniel Iascone, Lijuan Liu, Hanchuan Peng, Jiawei Yang, Jianhua Yao

Compositional Generalization in Unsupervised Compositional Representation Learning: A Study
on Disentanglement and Emergent Language 25074
Marc Niethammer, Colin Raffel, Zhenlin Xu

Improving Zero-Shot Generalization in Offline Reinforcement Learning Using Generalized
Similarity Functions 25088
Ilya Kostrikov, Bogdan Mazouze, Ofir Nachum, Jonathan Tompson

GAUDI: A Neural Architect for Immersive 3D Scene Generation..... 25102
*Samira Abnar, Miguel Angel Bautista, Zhuoyuan Chen, Afshin Dehghan, Laurent Dinh,
Hanlin Goh, Pengsheng Guo, Joshua Susskind, Walter Talbott, Alexander Toshev, Daniel
Ulbricht, Shuangfei Zhai*

Mask-Based Latent Reconstruction for Reinforcement Learning25117
Zhibo Chen, Cuiling Lan, Yan Lu, Tao Yu, Zhizheng Zhang

Product Ranking for Revenue Maximization with Multiple Purchases.....	25132
<i>Xiaolong Chen, Peng Cui, Bo Li, Renzhe Xu, Xingxuan Zhang, Yafeng Zhang</i>	
One Model to Edit Them All: Free-Form Text-Driven Image Manipulation with Semantic Modulations.....	25146
<i>Qifeng Chen, Xintong Han, Hongyu Liu, Yibing Song, Jue Wang, Chun Yuan, Ziyang Yuan, Yiming Zhu</i>	
Luckiness in Multiscale Online Learning.....	25160
<i>Wouter Koolen, Muriel F. Pérez-Ortiz</i>	
Deep Active Learning by Leveraging Training Dynamics	25171
<i>Jingrui He, Wei Huang, Andrew J Margenot, Hanghang Tong, Haonan Wang, Ziwei Wu</i>	
Learning and Covering Sums of Independent Random Variables with Unbounded Support.....	25185
<i>Alkis Kalavasis, Konstantinos Stavropoulos, Emmanouil Zampetakis</i>	
LGDN: Language-Guided Denoising Network for Video-Language Modeling	25198
<i>Mingyu Ding, Nanyi Fei, Yuqi Huo, Haoyu Lu, Zhiwu Lu</i>	
LieGG: Studying Learned Lie Group Generators.....	25212
<i>Artem Moskalev, Anna Sepiarskaia, Arnold Smeulders, Ivan Sosnovik</i>	
FourierNets Enable the Design of Highly Non-Local Optical Encoders for Computational Imaging	25224
<i>Misha B Ahrens, Michael Broxton, Alex Chen, Diptodip Deb, Zhenfei Jiao, Kaspar Podgorski, Ruth Sims, Srinivas C Turaga</i>	
Benign Overfitting in Two-Layer Convolutional Neural Networks	25237
<i>Misha Belkin, Yuan Cao, Zixiang Chen, Quanquan Gu</i>	
Discovery of Single Independent Latent Variable	25251
<i>Ori Katz, Uri Shaham, Jonathan Svirsky, Ronen Talmon</i>	
Meta-Ticket: Finding Optimal Subnetworks for Few-Shot Learning Within Randomly Initialized Neural Networks.....	25264
<i>Daiki Chijiwa, Yasutoshi Ida, Atsutoshi Kumagai, Shin'Ya Yamaguchi</i>	
LAION-5B: An Open Large-Scale Dataset for Training Next Generation Image-Text Models.....	25278
<i>Romain Beaumont, Mehdi Cherti, Theo Coombes, Katherine Crowson, Cade Gordon, Jenia Jitsev, Robert Kaczmarczyk, Aarush Katta, Srivatsa Kundurthy, Clayton Mullis, Ludwig Schmidt, Patrick Schramowski, Christoph Schuhmann, Richard Vencu, Ross Wightman, Mitchell Wortsman</i>	
Constants of Motion Network	25295
<i>Muhammad Firmansyah Kasim, Yi Heng Lim</i>	
HUMUS-Net: Hybrid Unrolled Multi-Scale Network Architecture for Accelerated MRI Reconstruction.....	25306
<i>Zalan Fabian, Mahdi Soltanolkotabi, Berk Tinaz</i>	
A Damped Newton Method Achieves Global $O(1/k^2)$ and Local Quadratic Convergence Rate	25320
<i>Alexander Gasnikov, Slavomir Hanzely, Dmitry Kamzolov, Dmitry Pasechnyuk, Peter Richtarik, Martin Takac</i>	
Disentangling the Predictive Variance of Deep Ensembles Through the Neural Tangent Kernel	25335
<i>Seijin Kobayashi, Pau Vilimelis Aceituno, Johannes Von Oswald</i>	

High-Dimensional Limit Theorems for SGD: Effective Dynamics and Critical Scaling.....	25349
<i>Gerard Ben Arous, Reza Gheissari, Aukosh Jagannath</i>	
Online Deep Equilibrium Learning for Regularization by Denoising	25363
<i>Weijie Gan, Ulugbek Kamilov, Jiaming Liu, Shirin Shoushtari, Xiaojian Xu</i>	
Semantic Exploration from Language Abstractions and Pretrained Representations.....	25377
<i>Andrea Banino, Stephanie Chan, Felix Hill, Andrew Lampinen, Neil Rabinowitz, Nicholas Roy, Dj Strouse, Allison Tam, Jane Wang</i>	
Earthformer: Exploring Space-Time Transformers for Earth System Forecasting	25390
<i>Zhihan Gao, Mu Li, Xingjian Shi, Hao Wang, Yuyang Bernie Wang, Dit-Yan Yeung, Yi Zhu</i>	
Benchopt: Reproducible, Efficient and Collaborative Optimization Benchmarks	25404
<i>Pierre Ablin, Pierre-Antoine Bannier, Benjamin Charlier, Mathieu Dagr�eou, Cassio F. Dantas, Tom Dupre La Tour, Ghislain Durif, Alexandre Gramfort, Quentin Klopfenstein, En Lai, Johan Larsson, Tanguy Lefort, Beno�t Mal�ezieux, Mathurin Massias, Thomas Moreau, Badr Moufad, Binh T. Nguyen, Alain Rakotomamonjy, Zaccharie Ramzi, Joseph Salmon, Samuel Vaiter</i>	
SketchBoost: Fast Gradient Boosted Decision Tree for Multioutput Problems	25422
<i>Leonid Iosipoi, Anton Vakhrushev</i>	
Deep Attentive Belief Propagation: Integrating Reasoning and Learning for Solving Constraint Optimization Problems.....	25436
<i>Bo An, Yanchen Deng, Shufeng Kong, Caihua Liu</i>	
Conservative Dual Policy Optimization for Efficient Model-Based Reinforcement Learning	25450
<i>Shenao Zhang</i>	
Decentralized Training of Foundation Models in Heterogeneous Environments.....	25464
<i>Beidi Chen, Tri Dao, Jared Davis, Yongjun He, Percy Liang, Christopher R�e, Binhang Yuan, Ce Zhang, Tianyi Zhang</i>	
Cross Aggregation Transformer for Image Restoration.....	25478
<i>Zheng Chen, Jinjin Gu, Linghe Kong, Xin Yuan, Yongbing Zhang, Yulun Zhang</i>	
Neural Payoff Machines: Predicting Fair and Stable Payoff Allocations Among Team Members.....	25491
<i>Yoram Bachrach, Daphne Cornelisse, Tal Kachman, Mateusz Malinowski, Thomas Rood</i>	
Learning Recourse on Instance Environment to Enhance Prediction Accuracy.....	25504
<i>Abir De, Lokesh N, Guntakanti Sai Koushik, Sunita Sarawagi</i>	
Learning to Re-Weight Examples with Optimal Transport for Imbalanced Classification	25517
<i>Dandan Guo, Zhuo Li, Hongyuan Zha, He Zhao, Meixi Zheng, Mingyuan Zhou</i>	
DIMES: A Differentiable Meta Solver for Combinatorial Optimization Problems.....	25531
<i>Ruizhong Qiu, Zhiqing Sun, Yiming Yang</i>	
Confident Approximate Policy Iteration for Efficient Local Planning in q^r -Realizable MDPs.....	25547
<i>Andr�s Gy�rgy, Tadashi Kozuno, Csaba Szepesvari, Gell�rt Weisz</i>	
Finding Correlated Equilibrium of Constrained Markov Game: A Primal-Dual Approach.....	25560
<i>Ziyi Chen, Shaocong Ma, Yi Zhou</i>	
NSNet: A General Neural Probabilistic Framework for Satisfiability Problems.....	25573
<i>Zhaoyu Li, Xujie Si</i>	

Brain Network Transformer.....	25586
<i>Hejie Cui, Wei Dai, Ying Guo, Xuan Kan, Carl Yang, Zilong Zhang</i>	
A Neural Corpus Indexer for Document Retrieval	25600
<i>Qi Chen, Chengmin Chi, Weiwei Deng, Yingyan Hou, Zheng Liu, Ziming Miao, Hao Sun, Haonan Wang, Yujing Wang, Shibin Wu, Yuqing Xia, Xing Xie, Mao Yang, Qi Zhang, Guoshuai Zhao</i>	
Perfect Sampling from Pairwise Comparisons	25615
<i>Dimitris Fotakis, Alkis Kalavasis, Christos Tzamos</i>	
Improved Utility Analysis of Private CountSketch	25631
<i>Rasmus Pagh, Mikkel Thorup</i>	
Optimal Efficiency-Envy Trade-Off Via Optimal Transport	25644
<i>Christian Kroer, Steven Yin</i>	
Non-Linear Coordination Graphs	25655
<i>Yipeng Kang, Tonghan Wang, Xiaoran Wu, Qianlan Yang, Chongjie Zhang</i>	
SafeBench: A Benchmarking Platform for Safety Evaluation of Autonomous Vehicles	25667
<i>Wenhao Ding, Yihan He, Hanjiang Hu, Bo Li, Zuxin Liu, Weijie Lyu, Shuai Wang, Chejian Xu, Ding Zhao</i>	
Improving Diffusion Models for Inverse Problems Using Manifold Constraints	25683
<i>Hyungjin Chung, Dohoon Ryu, Byeongsu Sim, Jong Chul Ye</i>	
Semi-Supervised Vision Transformers at Scale.....	25697
<i>Rahul Bhotika, Zhaowei Cai, Paolo Favaro, Davide Modolo, Avinash Ravichandran, Stefano Soatto, Zhuowen Tu, Manchen Wang</i>	
Fast Bayesian Estimation of Point Process Intensity as Function of Covariates	25711
<i>Taichi Asami, Hideaki Kim, Hiroyuki Toda</i>	

VOLUME 34

Online PAC-Bayes Learning	25725
<i>Benjamin Guedj, Maxime Haddouche</i>	
Deep Model Reassembly	25739
<i>Songhua Liu, Xinchao Wang, Xingyi Yang, Jingwen Ye, Daquan Zhou</i>	
Fast Algorithms for Packing Proportional Fairness and Its Dual.....	25754
<i>Francisco Criado, David Martinez-Rubio, Sebastian Pokutta</i>	
A Closer Look at Prototype Classifier for Few-Shot Image Classification	25767
<i>Mingcheng Hou, Issei Sato</i>	
When Are Offline Two-Player Zero-Sum Markov Games Solvable?.....	25779
<i>Qiwen Cui, Simon Du</i>	

BigBio: A Framework for Data-Centric Biomedical Natural Language Processing	25792
<i>Gabriel Altay, Stephen Bach, Fabio Barth, Benjamin Beilharz, Stella Biderman, Shamik Bose, Nicholas Broad, Samuel Cahyawijaya, Alison Callahan, Jenny Chim, Michael Cullan, Nathan Dahlberg, Debajyoti Datta, Shlok Deshmukh, Moritz Freidank, Jason Fries, Samuele Garda, Théo Gigant, John Giorgi, Jonas Golde, Patrick Haller, Sunny Kang, Sid Kiblawi, Wojciech Kusa, Yanis Labrak, Daniel León Perrián, Robert Martin, Shubhanshu Mishra, Trishala Neeraj, Marianna Nezhurina, Simon Ott, Marc Pàmies, Jose Posada, Matthias Samwald, Mario Sängler, Natasha Seelam, Ayush Singh, Karthik Rangasai Sivaraman, Rosaline Su, Albert Villanova Del Moral, Minh Chien Vu, Bo Wang, Leon Weber</i>	
Branch & Learn for Recursively and Iteratively Solvable Problems in Predict+Optimize	25807
<i>Xinyi Hu, Jasper Lee, Jimmy Lee, Allen Z. Zhong</i>	
Linear TreeShap	25818
<i>Albert Bifet, Jesse Read, Chao Xu, Peng Yu</i>	
Gradient Estimation with Discrete Stein Operators.....	25829
<i>Jessica Hwang, Lester Mackey, Jiaxin Shi, Michalis Titsias, Yuhao Zhou</i>	
Rapidly Mixing Multiple-Try Metropolis Algorithms for Model Selection Problems	25842
<i>Hyunwoong Chang, Changwoo Lee, Zhao Tang Luo, Huiyan Sang, Quan Zhou</i>	
Distributionally Adaptive Meta Reinforcement Learning	25856
<i>Pulkit Agrawal, Anurag Ajay, Dibya Ghosh, Abhishek Gupta, Sergey Levine</i>	
Toward Efficient Robust Training Against Union of ℓ_p Threat Models.....	25870
<i>Soheil Feizi, Maharshi Gor, Gaurang Sriramanan</i>	
Multi-View Subspace Clustering on Topological Manifold	25883
<i>Wentao Feng, Shudong Huang, Jiancheng Lv, Yazhou Ren, Ivor Tsang, Hongjie Wu, Zenglin Xu</i>	
CLEAR: Generative Counterfactual Explanations on Graphs.....	25895
<i>Ruocheng Guo, Jundong Li, Jing Ma, Saumitra Mishra, Aidong Zhang</i>	
Online Agnostic Multiclass Boosting	25908
<i>Vinod Raman, Ambuj Tewari</i>	
A Contrastive Rule for Meta-Learning	25921
<i>João Sacramento, Simon Schug, Johannes Von Oswald, Dominic Zhao, Nicolas Zucchet</i>	
Distinguishing Learning Rules with Brain Machine Interfaces.....	25937
<i>James M Murray, Jacob Portes, Christian Schmid</i>	
Efficient Training of Low-Curvature Neural Networks.....	25951
<i>François Fleuret, Himabindu Lakkaraju, Kyle Matoba, Suraj Srinivas</i>	
Infinite-Fidelity Coregionalization for Physical Simulation.....	25965
<i>Robert Kirby, Shibo Li, Zheng Wang, Shandian Zhe</i>	
Open High-Resolution Satellite Imagery: The WorldStrat Dataset – with Application to Super-Resolution.....	25979
<i>Julien Cornebise, Freddie Kalaitzis, Ivan Oršolic</i>	
Evaluation Beyond Task Performance: Analyzing Concepts in AlphaZero in Hex.....	25992
<i>Jessica Forde, George Konidaris, Michael Littman, Charles Lovering, Ellie Pavlick</i>	

Semi-Supervised Semantic Segmentation with Prototype-Based Consistency Regularization	26007
<i>Qiuchen Bian, Lingqiao Liu, Haiming Xu, Zhen Yang</i>	
Empirical Phase Diagram for Three-Layer Neural Networks with Infinite Width	26021
<i>Zhenyuan Jin, Tao Luo, Zhou Qixuan, Zhi-Qin Xu, Yaoyu Zhang, Hanxu Zhou</i>	
Benchmarking and Analyzing 3D Human Pose and Shape Estimation Beyond Algorithms.....	26034
<i>Zhongang Cai, Ziwei Liu, Hui En Pang, Lei Yang, Tianwei Zhang</i>	
TTOpt: A Maximum Volume Quantized Tensor Train-Based Optimization and Its Application to Reinforcement Learning	26052
<i>Andrei Chertkov, Andrzej S Cichocki, Ivan Oseledets, Anh-Huy Phan, Roman Schutski, Konstantin Sozykin</i>	
Diversified Recommendations for Agents with Adaptive Preferences.....	26066
<i>Arpit Agarwal, William Brown</i>	
A Mixture of Surprises for Unsupervised Reinforcement Learning	26078
<i>Gao Huang, Yangguang Li, Matthieu Lin, Yong-Jin Liu, Andrew Zhao</i>	
Deep Hierarchical Planning from Pixels	26091
<i>Pieter Abbeel, Ian Fischer, Danijar Hafner, Kuang-Huei Lee</i>	
PeRFception: Perception Using Radiance Fields	26105
<i>Anima Anandkumar, Minsu Cho, Chris Choy, Yoonwoo Jeong, Junha Lee, Jaesik Park, Seungjoo Shin</i>	
Efficient Submodular Optimization Under Noise: Local Search is Robust.....	26122
<i>Lingxiao Huang, Yuyi Wang, Chunxue Yang, Huanjian Zhou</i>	
On the Generalization Power of the Overfitted Three-Layer Neural Tangent Kernel Model.....	26135
<i>Peizhong Ju, Xiaojun Lin, Ness Shroff</i>	
Chromatic Correlation Clustering, Revisited	26147
<i>Shuang Cui, Kai Han, He Huang, Jing Tang, Qing Xiu</i>	
Gradient-Free Methods for Deterministic and Stochastic Nonsmooth Nonconvex Optimization.....	26160
<i>Michael Jordan, Tianyi Lin, Zeyu Zheng</i>	
Multilingual Abusive Comment Detection at Scale for Indic Languages.....	26176
<i>Somnath Banerjee, Mithun Das, Vikram Gupta, Binny Mathew, Animesh Mukherjee, Sumegh Roychowdhury, Punyajoy Saha, Hastagiri Prakash Vanchinathan</i>	
Generalized Delayed Feedback Model with Post-Click Information in Recommender Systems.....	26192
<i>Jiaqi Yang, De-Chuan Zhan</i>	
A Communication-Efficient Distributed Gradient Clipping Algorithm for Training Deep Neural Networks	26204
<i>Yunwen Lei, Chunyang Liao, Mingrui Liu, Zhenxun Zhuang</i>	
On Analyzing Generative and Denoising Capabilities of Diffusion-Based Deep Generative Models	26218
<i>Kamil Deja, Anna Kuzina, Jakub Tomczak, Tomasz Trzcinski</i>	
On the Tradeoff Between Robustness and Fairness.....	26230
<i>Weiwei Liu, Xinsong Ma, Zekai Wang</i>	

Learning Expressive Meta-Representations with Mixture of Expert Neural Processes	26242
<i>Herke Van Hoof, Qi Wang</i>	
Learning Physical Dynamics with Subequivariant Graph Neural Networks.....	26256
<i>Chuang Gan, Jiaqi Han, Wenbing Huang, Jiachen Li, Hengbo Ma, Josh Tenenbaum</i>	
DiSC: Differential Spectral Clustering of Features	26269
<i>Ariel Jaffe, Gal Mishne, Ram Dyuthi Sristi</i>	
Learn What Matters: Cross-Domain Imitation Learning with Task-Relevant Embeddings	26283
<i>Tim Franzmeyer, João Henriques, Philip Torr</i>	
UViM: A Unified Modeling Approach for Vision with Learned Guiding Codes	26295
<i>Lucas Beyer, Jeremiah Harmsen, Neil Houlsby, Alexander Kolesnikov, André Susano Pinto, Xiaohua Zhai</i>	
Rotation-Equivariant Conditional Spherical Neural Fields for Learning a Natural Illumination Prior	26309
<i>Bernhard Egger, James Gardner, William Smith</i>	
Proximal Learning with Opponent-Learning Awareness.....	26324
<i>Jakob Foerster, Roger Grosse, Chris Lu, Stephen Zhao</i>	
HyperTree Proof Search for Neural Theorem Proving	26337
<i>Gabriel Ebner, Amaury Hayat, Marie-Anne Lachaux, Timothee Lacroix, Guillaume Lample, Thibaut Lavril, Xavier Martinet, Aurelien Rodriguez</i>	
The Policy-Gradient Placement and Generative Routing Neural Networks for Chip Design	26350
<i>Ruoyu Cheng, Jianye Hao, Yang Li, Xianglong Lyu, Junchi Yan, Junjie Ye</i>	
IMED-RL: Regret Optimal Learning of Ergodic Markov Decision Processes	26363
<i>Odalric-Ambrym Maillard, Fabien Pesquere</i>	
Coresets for Wasserstein Distributionally Robust Optimization Problems	26375
<i>Hu Ding, Jiawei Huang, Ruomin Huang, Wenjie Liu</i>	
SAMURAI: Shape and Material from Unconstrained Real-World Arbitrary Image Collections.....	26389
<i>Jonathan Barron, Mark Boss, Andreas Engelhardt, Varun Jampani, Abhishek Kar, Hendrik Pa Lensch, Yuanzhen Li, Deqing Sun</i>	
Automatic Differentiation of Nonsmooth Iterative Algorithms.....	26404
<i>Jerome Bolte, Edouard Pauwels, Samuel Vaiter</i>	
Wukong: A 100 Million Large-Scale Chinese Cross-Modal Pre-Training Benchmark.....	26418
<i>Jiaxi Gu, Lu Hou, Runhui Huang, Xin Jiang, Xiaodan Liang, Guansong Lu, Xiaojun Meng, Niu Minzhe, Hang Xu, Chunjing Xu, Lewei Yao, Wei Zhang</i>	
Neural Abstractions	26432
<i>Alessandro Abate, Alec Edwards, Mirco Giacobbe</i>	
Few-Shot Non-Parametric Learning with Deep Latent Variable Model.....	26448
<i>Yiqin Dai, Zhiying Jiang, Ming Li, Jimmy Lin, Ji Xin</i>	
ST-Adapter: Parameter-Efficient Image-To-Video Transfer Learning.....	26462
<i>Hongsheng Li, Ziyi Lin, Junting Pan, Jing Shao, Xiatian Zhu</i>	
Distilled Gradient Aggregation: Purify Features for Input Attribution in the Deep Neural Network.....	26478
<i>Jaesik Choi, Giyoung Jeon, Haedong Jeong</i>	

Temporally Disentangled Representation Learning.....	26492
<i>Guangyi Chen, Weiran Yao, Kun Zhang</i>	

VOLUME 35

Can Adversarial Training Be Manipulated by Non-Robust Features?.....	26504
<i>Songcan Chen, Lei Feng, Sheng-Jun Huang, Lue Tao, Hongxin Wei, Jinfeng Yi</i>	
On the Effectiveness of Fine-Tuning Versus Meta-Reinforcement Learning.....	26519
<i>Pieter Abbeel, Stephen James, Mandi Zhao</i>	
Generalizing Goal-Conditioned Reinforcement Learning with Variational Causal Reasoning.....	26532
<i>Wenhao Ding, Bo Li, Haohong Lin, Ding Zhao</i>	
WinoGAViL: Gamified Association Benchmark to Challenge Vision-And-Language Models	26549
<i>Mohit Bansal, Yonatan Bitton, Nitzan Bitton Guetta, Yuval Elovici, Roy Schwartz, Gabriel Stanovsky, Ron Yosef</i>	
Elucidating the Design Space of Diffusion-Based Generative Models	26565
<i>Timo Aila, Miika Aittala, Tero Karras, Samuli Laine</i>	
TREC: Transient Redundancy Elimination-Based Convolution	26578
<i>Xiaoyong Du, Jiawei Guan, Jiesong Liu, Xipeng Shen, Hsin-Hsuan Sung, Ruofan Wu, Feng Zhang</i>	
Chaotic Regularization and Heavy-Tailed Limits for Deterministic Gradient Descent.....	26590
<i>Soon Hoe Lim, Umut Simsekli, Yijun Wan</i>	
DARE: Disentanglement-Augmented Rationale Extraction	26603
<i>Yanqing An, Enhong Chen, Yichao Du, Qi Liu, Li Wang, Linan Yue</i>	
The Gyro-Structure of Some Matrix Manifolds	26618
<i>Xuan Son Nguyen</i>	
SMPL: Simulated Industrial Manufacturing and Process Control Learning Environments.....	26631
<i>Jiayi Cheng, Benjamin Decardi-Nelson, Animesh Garg, Jinfeng Liu, Xiaohong Liu, Matthew Poon, Bo Song, Sile Tao, Xiaozhou Wang, Dengdeng Yu, An Zhang, Mohan Zhang</i>	
CoPur: Certifiably Robust Collaborative Inference Via Feature Purification.....	26645
<i>Sanmi Koyejo, Bo Li, Jing Liu, Chulin Xie</i>	
Toward Understanding Privileged Features Distillation in Learning-To-Rank	26658
<i>Jan Bakus, Holakou Rahmanian, Sujay Sanghavi, Vishwanathan S. V. N., Shuo Yang</i>	
Contrastive and Non-Contrastive Self-Supervised Learning Recover Global and Local Spectral Embedding Methods.....	26671
<i>Randall Balestriero, Yann Lecun</i>	
Subgame Solving in Adversarial Team Games.....	26686
<i>Federico Cacciamani, Luca Carminati, Gabriele Farina, Nicola Gatti, Pierricardo Olivieri, Tuomas Sandholm, Brian Zhang</i>	
A Framework for Bilevel Optimization that Enables Stochastic and Global Variance Reduction Algorithms.....	26698
<i>Pierre Ablin, Mathieu Dagr�eou, Thomas Moreau, Samuel Vaiter</i>	

Predicting Cellular Responses to Novel Drug Perturbations at a Single-Cell Resolution	26711
<i>Simon Boehm, Stephan Günnemann, Leon Hetzel, Niki Kilbertus, Mohammad Lotfollahi, Fabian Theis</i>	
Accelerated Projected Gradient Algorithms for Sparsity Constrained Optimization Problems	26723
<i>Jan Harold Alcantara, Ching-Pei Lee</i>	
The Stability-Efficiency Dilemma: Investigating Sequence Length Warmup for Training GPT Models.....	26736
<i>Yuxiong He, Conglong Li, Minjia Zhang</i>	
Trading off Image Quality for Robustness is Not Necessary with Regularized Deterministic Autoencoders.....	26751
<i>Stefan Falkner, Margret Keuper, Amrutha Saseendran, Kathrin Skubch</i>	
Does Momentum Change the Implicit Regularization on Separable Data?	26764
<i>Wei Chen, Tie-Yan Liu, Zhi-Ming Ma, Qi Meng, Ruoyu Sun, Bohan Wang, Huishuai Zhang</i>	
Generalization Gap in Amortized Inference	26777
<i>David Barber, Peter Hayes, Mingtian Zhang</i>	
Top Two Algorithms Revisited.....	26791
<i>Dorian Baudry, Rémy Degenne, Rianne De Heide, Marc Jourdan, Emilie Kaufmann</i>	
Relation-Constrained Decoding for Text Generation	26804
<i>Xiang Chen, Xiaojun Wan, Zhixian Yang</i>	
Learning General World Models in a Handful of Reward-Free Deployments	26820
<i>Philip Ball, Edward Grefenstette, Aldo Pacchiano, Jack Parker-Holder, S Roberts, Tim Rocktäschel, Oleh Rybkin, Yingchen Xu</i>	
Bringing Image Scene Structure to Video Via Frame-Clip Consistency of Object Tokens	26839
<i>Amir Bar, Elad Ben Avraham, Trevor Darrell, Amir Globerson, Roei Herzig, Leonid Karlinsky, Karttikeya Mangalam, Anna Rohrbach</i>	
Probabilistic Transformer: Modelling Ambiguities and Distributions for RNA Folding and Molecule Design.....	26856
<i>Jörg Franke, Frank Hutter, Frederic Runge</i>	
PulseImpute: A Novel Benchmark Task for Pulsative Physiological Signal Imputation.....	26874
<i>Varol Aydemir, Santosh Kumar, Alexander Moreno, Supriya Nagesh, James Rehg, David Wetter, Maxwell Xu</i>	
Beyond Separability: Analyzing the Linear Transferability of Contrastive Representations to Related Subpopulations.....	26889
<i>Jeff Z. Haochen, Ananya Kumar, Tengyu Ma, Colin Wei</i>	
Precise Regret Bounds for Log-Loss Via a Truncated Bayesian Algorithm.....	26903
<i>Ananth Grama, Mohsen Heidari, Wojciech Szpankowski, Changlong Wu</i>	
What Are the Best Systems? New Perspectives on NLP Benchmarking.....	26915
<i>Stephan Cléménçon, Pierre Colombo, Ekhine Irurozki, Nathan Noiry</i>	
Learning from Label Proportions by Learning with Label Noise.....	26933
<i>Clay Scott, Yutong Wang, Jianxin Zhang</i>	

Dynamics of SGD with Stochastic Polyak Stepsizes: Truly Adaptive Variants and Convergence to Exact Solution	26943
<i>Simon Lacoste-Julien, Nicolas Loizou, Antonio Orvieto</i>	
MORA: Improving Ensemble Robustness Evaluation with Model Reweighting Attack	26955
<i>Xitong Gao, Cheng-Zhong Xu, Yunrui Yu</i>	
Faster and Scalable Algorithms for Densest Subgraph and Decomposition.....	26966
<i>Chandra Chekuri, Elfarouk Harb, Kent Quanrud</i>	
Extrapolation and Spectral Bias of Neural Nets with Hadamard Product: A Polynomial Net Study	26980
<i>Volkan Cevher, Grigorios Chrysos, Fanghui Liu, Yongtao Wu, Zhenyu Zhu</i>	
Cost-Sensitive Self-Training for Optimizing Non-Decomposable Metrics.....	26994
<i>Venkatesh Babu R, Shrinivas Ramasubramanian, Harsh Rangwani, Kato Takashi, Sho Takemori, Yuhei Umeda</i>	
Tensor Wheel Decomposition and Its Tensor Completion Application.....	27008
<i>Liang-Jian Deng, Hong-Xia Dou, Ting-Zhu Huang, Deyu Meng, Zhong-Cheng Wu</i>	
BOND: Benchmarking Unsupervised Outlier Node Detection on Static Attributed Graphs.....	27021
<i>Canyu Chen, George H Chen, Kaize Ding, Xueying Ding, Yingtong Dou, Xiyang Hu, Zhihao Jia, Jundong Li, Kay Liu, Hao Peng, Kai Shu, Lichao Sun, Philip S Yu, Ruitong Zhang, Yue Zhao</i>	
Movement Penalized Bayesian Optimization with Application to Wind Energy Systems.....	27036
<i>Ilija Bogunovic, Andreas Krause, Shyam Sundhar Ramesh, Pier Giuseppe Sessa</i>	
Associating Objects and Their Effects in Video Through Coordination Games.....	27049
<i>Forrester Cole, Tali Dekel, Bill Freeman, Erika Lu, Michael Rubinstein, Weidi Xie, Andrew Zisserman</i>	
Point-M2AE: Multi-Scale Masked Autoencoders for Hierarchical Point Cloud Pre-Training.....	27061
<i>Rongyao Fang, Peng Gao, Ziyu Guo, Hongsheng Li, Yu Qiao, Dong Wang, Renrui Zhang, Bin Zhao</i>	
Exploring Example Influence in Continual Learning.....	27075
<i>Wei Feng, Fan Lyu, Fanhua Shang, Qing Sun, Liang Wan</i>	
Subspace Clustering in High-Dimensions: Phase Transitions & Statistical-To-Computational Gap	27087
<i>Florent Krzakala, Bruno Loureiro, Luca Pesce, Lenka Zdeborová</i>	
Self-Supervised Fair Representation Learning Without Demographics	27100
<i>Junyi Chai, Xiaoqian Wang</i>	
CEDe: A Collection of Expert-Curated Datasets with Atom-Level Entity Annotations for Optical Chemical Structure Recognition.....	27114
<i>Jaegul Choo, Sehui Han, Rodrigo Hormazabal, Ahra Jo, Yeonsik Jo, Seung Hwan Kim, Honglak Lee, Jaewan Lee, Moontae Lee, Soonyoung Lee, Changyoung Park</i>	
How Mask Matters: Towards Theoretical Understandings of Masked Autoencoders	27127
<i>Yisen Wang, Yifei Wang, Qi Zhang</i>	
Improving Out-Of-Distribution Generalization by Adversarial Training with Structured Priors.....	27140
<i>Yisen Wang, Qixun Wang, Yifei Wang, Hong Zhu</i>	

Graphein - A Python Library for Geometric Deep Learning and Network Analysis on Biomolecular Structures and Interaction Networks	27153
<i>Tom Blundell, Yuanqi Du, Dominic Hall, Charles Harris, Kexin Huang, Arian Jamasb, Pietro Lió, Eric Ma, Ramon Viñas Torné</i>	
ZeroQuant: Efficient and Affordable Post-Training Quantization for Large-Scale Transformers.....	27168
<i>Yuxiong He, Conglong Li, Xiaoxia Wu, Zhewei Yao, Reza Yazdani Aminabadi, Minjia Zhang</i>	
Simplified Graph Convolution with Heterophily	27184
<i>Sudhanshu Chanpuriya, Cameron Musco</i>	
Signal Propagation in Transformers: Theoretical Perspectives and the Role of Rank Collapse.....	27198
<i>Sotiris Anagnostidis, Luca Biggio, Aurelien Lucchi, Lorenzo Noci, Antonio Orvieto, Sidak Pal Singh</i>	
Maximum a Posteriori Natural Scene Reconstruction from Retinal Ganglion Cells with Deep Denoiser Priors	27212
<i>Nora Brackbill, E. J. Chichilnisky, Alan Litke, Alexander Sher, Eero Simoncelli, Eric Wu</i>	
Imbalance Trouble: Revisiting Neural-Collapse Geometry.....	27225
<i>Tina Behnia, Ganesh Ramachandra Kini, Christos Thrampoulidis, Vala Vakilian</i>	
ProtoX: Explaining a Reinforcement Learning Agent Via Prototyping.....	27239
<i>Qihang Lin, Ronilo Ragodos, Tong Wang, Xun Zhou</i>	
NOTE: Robust Continual Test-Time Adaptation Against Temporal Correlation.....	27253
<i>Taesik Gong, Jongheon Jeong, Taewon Kim, Yewon Kim, Sung-Ju Lee, Jinwoo Shin</i>	
Margin-Based Few-Shot Class-Incremental Learning with Class-Level Overfitting Mitigation	27267
<i>Ruixuan Li, Yuhua Li, Shanghang Zhang, Yixiong Zou</i>	
First Hitting Diffusion Models for Generating Manifold, Graph and Categorical Data.....	27280
<i>Qiang Liu, Lemeng Wu, Mao Ye</i>	

VOLUME 36

Forecasting Future World Events with Neural Networks	27293
<i>Owain Evans, Dan Hendrycks, Ryan Jia, Joe Kwon, Richard Li, Mantas Mazeika, Dawn Song, Jacob Steinhardt, Tristan Xiao, Andy Zou</i>	
Human-Robotic Prosthesis as Collaborating Agents for Symmetrical Walking.....	27306
<i>Xiang Gao, He Huang, Jennie Si, Brent Wallace, Ruofan Wu, Junmin Zhong</i>	
Support Recovery in Sparse PCA with Incomplete Data.....	27321
<i>Jean Honorio, Hanbyul Lee, Qifan Song</i>	
Exponentially Improving the Complexity of Simulating the Weisfeiler-Lehman Test with Graph Neural Networks.....	27333
<i>Anders Aamand, Justin Chen, Piotr Indyk, Shyam Narayanan, Ronitt Rubinfeld, Nicholas Schiefer, Sandeep Silwal, Tal Wagner</i>	
Zero-Shot Transfer Learning Within a Heterogeneous Graph Via Knowledge Transfer Networks.....	27347
<i>Ziniu Hu, John Palowitch, Bryan Perozzi, Ruslan Salakhutdinov, Minji Yoon, Dustin Zelle</i>	
Nonlinear Sufficient Dimension Reduction with a Stochastic Neural Network.....	27360
<i>Faming Liang, Siqi Liang, Yan Sun</i>	

Autoregressive Perturbations for Data Poisoning.....	27374
<i>Jonas Geiping, Micah Goldblum, Tom Goldstein, David Jacobs, Pedro Sandoval-Segura, Vasu Singla</i>	
NodeFormer: A Scalable Graph Structure Learning Transformer for Node Classification	27387
<i>Zenan Li, David P Wipf, Qitian Wu, Junchi Yan, Wentao Zhao</i>	
Is Integer Arithmetic Enough for Deep Learning Training?.....	27402
<i>Masoud Asgharian, Alireza Ghaffari, Vahid Partovi Nia, Marzieh S. Tahaei, Mohammadreza Tayaranian</i>	
mRI: Multi-Modal 3D Human Pose Estimation Dataset Using mmWave, RGB-D, and Inertial Sensors	27414
<i>Sizhe An, Yin Li, Umit Ogras</i>	
Effectiveness of Vision Transformer for Fast and Accurate Single-Stage Pedestrian Detection	27427
<i>Panagiotis Barmpoutis, Tania Stathaki, Jing Yuan</i>	
ESCADA: Efficient Safety and Context Aware Dose Allocation for Precision Medicine.....	27441
<i>Ahmet Alparslan Celik, Ilker Demirel, Cem Tekin</i>	
Grow and Merge: A Unified Framework for Continuous Categories Discovery.....	27455
<i>Yutong Feng, Yue Gao, Jianwen Jiang, Rong Jin, Mingqian Tang, Hai Wan, Zhi-Fan Wu, Xinwei Zhang, Xibin Zhao</i>	
Category-Level 6D Object Pose Estimation in the Wild: A Semi-Supervised Learning Approach and a New Dataset.....	27469
<i>Xiaolong Wang, Yang Fu</i>	
Change Event Dataset for Discovery from Spatio-Temporal Remote Sensing Imagery	27484
<i>Kavita Bala, Bharath Hariharan, Utkarsh Mall</i>	
Improved Algorithms for Neural Active Learning	27497
<i>Yikun Ban, Arindam Banerjee, Jingrui He, Hanghang Tong, Yuheng Zhang</i>	
Robust Calibration with Multi-Domain Temperature Scaling	27510
<i>Stephen Bates, Michael Jordan, Yi Ma, Yaodong Yu</i>	
Independence Testing-Based Approach to Causal Discovery Under Measurement Error and Linear Non-Gaussian Models	27524
<i>Haoyue Dai, Peter Spirtes, Kun Zhang</i>	
CUP: Critic-Guided Policy Reuse	27537
<i>Siyuan Li, Chongjie Zhang, Jin Zhang</i>	
Local Identifiability of Deep ReLU Neural Networks: The Theory.....	27549
<i>Francois Bachoc, Joachim Bona-Pellissier, François Malgouyres</i>	
DOMINO: Decomposed Mutual Information Optimization for Generalized Context in Meta-Reinforcement Learning	27563
<i>Jianyu Chen, Jianye Hao, Ping Luo, Yao Mu, Fei Ni, Bin Wang, Yuzheng Zhuang</i>	
Generalization Bounds with Minimal Dependency on Hypothesis Class Via Distributionally Robust Optimization.....	27576
<i>Henry Lam, Yibo Zeng</i>	

QUARK: Controllable Text Generation with Reinforced Unlearning.....	27591
<i>Prithviraj Ammanabrolu, Yejin Choi, Jack Hessel, Liwei Jiang, Ximing Lu, Lianhui Qin, Sean Welleck, Peter West</i>	
Safe Opponent-Exploitation Subgame Refinement.....	27610
<i>Yansen Jing, Mingyang Liu, Qihan Liu, Pingzhong Tang, Chengjie Wu, Jun Yang, Chongjie Zhang</i>	
Parameter-Free Dynamic Graph Embedding for Link Prediction	27623
<i>Hansu Gu, Ning Gu, Dongsheng Li, Jiahao Liu, Tun Lu, Peng Zhang</i>	
Towards a Unified Framework for Uncertainty-Aware Nonlinear Variable Selection with Theoretical Guarantees.....	27636
<i>Beau Coker, Brent Coull, Wenying Deng, Jeremiah Liu, Rajarshi Mukherjee</i>	
Non-Markovian Reward Modelling from Trajectory Labels Via Interpretable Multiple Instance Learning	27652
<i>Tom Bewley, Joseph Early, Christine Evers, Sarvapali Ramchurn</i>	
Explaining Preferences with Shapley Values.....	27664
<i>Siu Lun Chau, Jaime Ferrando Huertas, Robert Hu, Dino Sejdinovic</i>	
Structuring Uncertainty for Fine-Grained Sampling in Stochastic Segmentation Networks.....	27678
<i>Jakob Gawlikowski, Julia Niebling, Frank Nussbaum</i>	
Fair Bayes-Optimal Classifiers Under Predictive Parity	27692
<i>Guang Cheng, Edgar Dobriban, Xianli Zeng</i>	
Pre-Train Your Loss: Easy Bayesian Transfer Learning with Informative Priors.....	27706
<i>Micah Goldblum, Sanyam Kapoor, Yann Lecun, Ravid Shwartz-Ziv, Hossein Souri, Andrew Wilson, Chen Zhu</i>	
Unsupervised Adaptation from Repeated Traversals for Autonomous Driving.....	27716
<i>Mark Campbell, Wei-Lun Chao, Bharath Hariharan, Katie Luo, Cheng Perng Phoo, Kilian Weinberger, Yurong You, Travis Zhang</i>	
Training Language Models to Follow Instructions with Human Feedback.....	27730
<i>Sandhini Agarwal, Diogo Almeida, Amanda Askell, Paul Christiano, Jacob Hilton, Xu Jiang, Fraser Kelton, Jan Leike, Ryan Lowe, Luke Miller, Pamela Mishkin, Long Ouyang, Alex Ray, John Schulman, Maddie Simens, Katarina Slama, Carroll Wainwright, Peter Welinder, Jeffrey Wu, Chong Zhang</i>	
Byzantine Spectral Ranking	27745
<i>John Augustine, Arnhav Datar, Arun Rajkumar</i>	
Non-Rigid Point Cloud Registration with Neural Deformation Pyramid.....	27757
<i>Tatsuya Harada, Yang Li</i>	
Provably Tuning the ElasticNet Across Instances	27769
<i>Maria-Florina Balcan, Mikhail Khodak, Dravyansh Sharma, Ameet Talwalkar</i>	
PlasticityNet: Learning to Simulate Metal, Sand, and Snow for Optimization Time Integration.....	27783
<i>Yadi Cao, Chenfanfu Jiang, Minchen Li, Xuan Li, Craig Schroeder, Yin Yang</i>	
Synthetic Model Combination: An Instance-Wise Approach to Unsupervised Ensemble Learning	27797
<i>Alex Chan, Mihaela Van Der Schaar</i>	

Power and Limitations of Single-Qubit Native Quantum Neural Networks	27810
<i>Mujin Li, Xin Wang, Hongshun Yao, Zhan Yu</i>	
Hedging as Reward Augmentation in Probabilistic Graphical Models	27824
<i>Debarun Bhattacharjya, Radu Marinescu</i>	
Debiased, Longitudinal and Coordinated Drug Recommendation Through Multi-Visit Clinic Records.....	27837
<i>Yuhan Chen, Shuqi Li, Hongda Sun, Ji-Rong Wen, Shufang Xie, Rui Yan</i>	
Disentangling Causal Effects from Sets of Interventions in the Presence of Unobserved Confounders	27850
<i>Ciarán Gilligan-Lee, Olivier Jeunen, Mounia Lalmas, Rishabh Mehrotra</i>	
MATE: Benchmarking Multi-Agent Reinforcement Learning in Distributed Target Coverage Control.....	27862
<i>Mickel Liu, Xuehai Pan, Yizhou Wang, Yaodong Yang, Fangwei Zhong, Song-Chun Zhu</i>	
Receding Horizon Inverse Reinforcement Learning	27880
<i>Wei Gao, David Hsu, Yiqing Xu</i>	
A Reparametrization-Invariant Sharpness Measure Based on Information Geometry	27893
<i>Cheongjae Jang, Sungyoon Lee, Yung-Kyun Noh, Frank Park</i>	
Hyper-Representations as Generative Models: Sampling Unseen Neural Network Weights	27906
<i>Damian Borth, Xavier Giró-I-Nieto, Boris Knyazev, Konstantin Schürholt</i>	
Multi-Game Decision Transformers	27921
<i>Ian Fischer, Daniel Freeman, Sergio Guadarrama, Eric Jang, Kuang-Huei Lee, Lisa Lee, Henryk Michalewski, Igor Mordatch, Ofir Nachum, Winnie Xu, Mengjiao Sherry Yang</i>	
Improving Transformer with an Admixture of Attention Heads.....	27937
<i>Hai Do, Nhat Ho, Khai Nguyen, Khuong Duy Nguyen, Tam Nguyen, Tan Nguyen, Stanley Osher, Minh Pham, Vishwanath Saragadam</i>	
Flexible Diffusion Modeling of Long Videos.....	27953
<i>William Harvey, Vaden Masrani, Saeid Naderiparizi, Christian Weilbach, Frank Wood</i>	
Towards Reasonable Budget Allocation in Untargeted Graph Structure Attacks Via Gradient Debias.....	27966
<i>Stan Z. Li, Zicheng Liu, Zihan Liu, Yun Luo, Lirong Wu</i>	
ToDD: Topological Compound Fingerprinting in Computer-Aided Drug Discovery	27978
<i>Yuzhou Chen, Baris Coskunuzer, Andaç Demir, Yulia Gel, Bulent Kiziltan, Ignacio Segovia- Dominguez</i>	
Geo-SIC: Learning Deformable Geometric Shapes in Deep Image Classifiers	27994
<i>Jian Wang, Miaomiao Zhang</i>	
A Survey and Datasheet Repository of Publicly Available US Criminal Justice Datasets	28008
<i>Bradley Butcher, Adrian Weller, Miri Zilka</i>	
Segmenting Moving Objects Via an Object-Centric Layered Representation.....	28023
<i>Junyu Xie, Weidi Xie, Andrew Zisserman</i>	
NAS-Bench-Suite-Zero: Accelerating Research on Zero Cost Proxies.....	28037
<i>Frank Hutter, Arjun Krishnakumar, Mahmoud Safari, Renbo Tu, Colin White, Arber Zela</i>	

On the Strong Correlation Between Model Invariance and Generalization 28052
Weijian Deng, Stephen Gould, Liang Zheng

On Scalable Testing of Samplers..... 28068
Kuldeep S Meel, Yash Pote

VOLUME 37

How and Why to Manipulate Your Own Agent: On the Incentives of Users of Learning Agents..... 28080
Yoav Kolubus, Noam Nisan

LIPS - Learning Industrial Physical Simulation Benchmark Suite..... 28095
Seif Attoui, David Danan, Pavel Dimitrov, Benjamin Donnot, Clement Etienam, Asma Farjallah, Milad Leyli Abadi, Antoine Marot, Jérôme Picault, Mouadh Yagoubi

FlyView: A Bio-Informed Optical Flow Truth Dataset for Visual Navigation Using Panoramic Stereo Vision 28110
Alix Leroy, Graham Taylor

Controllable Text Generation with Neurally-Decomposed Oracle 28125
Kai-Wei Chang, Sidi Lu, Tao Meng, Nanyun Peng

Active Learning Helps Pretrained Models Learn the Intended Task 28140
Salil Deshpande, Noah Goodman, Jesse Mu, Dat Nguyen, Alex Tamkin

Pragmatically Learning from Pedagogical Demonstrations in Multi-Goal Environments 28154
Hugo Caselles-Dupré, Mohamed Chetouani, Olivier Sigaud

Beyond Black Box Densities: Parameter Learning for the Deviated Components..... 28167
Dat Do, Nhat Ho, Xuanlong Nguyen

Statistical, Robustness, and Computational Guarantees for Sliced Wasserstein Distances 28179
Ziv Goldfeld, Kengo Kato, Sloan Nietert, Ritwik Sadhu

Oracle Inequalities for Model Selection in Offline Reinforcement Learning..... 28194
Emma Brunskill, Bo Dai, Jonathan N Lee, Ofir Nachum, George Tucker

Model-Based Opponent Modeling 28208
Haobin Jiang, Jiechuan Jiang, Zongqing Lu, Xiaopeng Yu, Wanpeng Zhang

What is Where by Looking: Weakly-Supervised Open-World Phrase-Grounding Without Text Inputs..... 28222
Tal Shaharabany, Yoad Tewel, Lior Wolf

DeepMed: Semiparametric Causal Mediation Analysis with Debiased Deep Learning..... 28238
Lin Liu, Zhonghua Liu, Siqu Xu

On the Learning Mechanisms in Physical Reasoning 28252
Shiqian Li, Kewen Wu, Chi Zhang, Yixin Zhu

A Continuous Time Framework for Discrete Denoising Models 28266
Joe Benton, Andrew Campbell, Valentin De Bortoli, George Deligiannidis, Arnaud Doucet, Thomas Rainforth

Simple Mechanisms for Welfare Maximization in Rich Advertising Auctions 28280
Gagan Aggarwal, Kshipra Bhawalkar, Aranyak Mehta, Divyarthi Mohan, Alexandros Psomas

Continuously Tempered PDMP Samplers	28293
<i>Augustin Chevallier, Paul Fearnhead, Robert Salomone, Matthew Sutton</i>	
Distributed Influence-Augmented Local Simulators for Parallel MARL in Large Networked Systems.....	28305
<i>Mustafa Mert Çelikok, Jinke He, Frans Oliehoek, Matthijs Spaan, Miguel Suau</i>	
Tight Mutual Information Estimation with Contrastive Fenchel-Legendre Optimization.....	28319
<i>Larry Carin, Junya Chen, Xinwei Deng, Qing Guo, Jing Huang, Fan Li, Chenyang Tao, Dong Wang, Yuewei Yang</i>	
Exploring the Algorithm-Dependent Generalization of AUPRC Optimization with List Stability	28335
<i>Yuan He, Qingming Huang, Peisong Wen, Qianqian Xu, Zhiyong Yang</i>	
LTMD: Learning Improvement of Spiking Neural Networks with Learnable Thresholding Neurons and Moderate Dropout.....	28350
<i>Tee Hiang Cheng, Meng-Hiot Lim, Siqi Wang</i>	
TransBoost: Improving the Best ImageNet Performance Using Deep Transduction	28363
<i>Guy Bar-Shalom, Omer Belhasin, Ran El-Yaniv</i>	
Sparse Probabilistic Circuits Via Pruning and Growing.....	28374
<i>Meihua Dang, Anji Liu, Guy Van Den Broeck</i>	
Adam Can Converge Without Any Modification on Update Rules.....	28386
<i>Congliang Chen, Zhi-Quan Luo, Naichen Shi, Ruoyu Sun, Yushun Zhang</i>	
Nonlinear MCMC for Bayesian Machine Learning	28400
<i>James Vuckovic</i>	
Continual Learning: A Feature Extraction Formalization, an Efficient Algorithm, and Fundamental Obstructions.....	28414
<i>Binghui Peng, Andrej Risteski</i>	
IKEA-Manual: Seeing Shape Assembly Step by Step.....	28428
<i>Chin-Yi Cheng, Jiayuan Mao, Ruocheng Wang, Jiajun Wu, Ran Zhang, Yunzhi Zhang</i>	
M ³ ViT: Mixture-Of-Experts Vision Transformer for Efficient Multi-Task Learning with Model-Accelerator Co-Design	28441
<i>Tianlong Chen, Yu Cheng, Zhiwen Fan, Cong Hao, Ziyu Jiang, Hanxue Liang, Rishov Sarkar, Zhangyang Wang, Kai Zou</i>	
When to Make Exceptions: Exploring Language Models as Accounts of Human Moral Judgment.....	28458
<i>Fernando Gonzalez, Zhijing Jin, Ojasv Kamal, Sydney Levine, Rada Mihalcea, Mrinmaya Sachan, Maarten Sap, Bernhard Schölkopf, Josh Tenenbaum</i>	
Exponential Family Model-Based Reinforcement Learning Via Score Matching.....	28474
<i>Anmol Kabra, Gene Li, Junbo Li, Nati Srebro, Zhaoran Wang, Zhuoran Yang</i>	
Monte Carlo Tree Search Based Variable Selection for High Dimensional Bayesian Optimization.....	28488
<i>Xiaobin Huang, Chao Qian, Lei Song, Ke Xue</i>	
OpenSRH: Optimizing Brain Tumor Surgery Using Intraoperative Stimulated Raman Histology.....	28502
<i>Sandra Camelo-Piragua, Asadur Chowdury, Kyle Conway, Christian Freudiger, Todd Hollon, Xinhai Hou, Cheng Jiang, Akhil Kondepudi, Honglak Lee, Daniel Orringer</i>	

Assistive Teaching of Motor Control Tasks to Humans	28517
<i>Erdem Biyik, Noah Goodman, Suvir Mirchandani, Dorsa Sadigh, Megha Srivastava</i>	
Retrospective Adversarial Replay for Continual Learning.....	28530
<i>Jeff A Bilmes, Lilly Kumari, Shengjie Wang, Tianyi Zhou</i>	
Unsupervised Image-To-Image Translation with Density Changing Regularization	28545
<i>Qirong Ho, Shaoan Xie, Kun Zhang</i>	
CASA: Category-Agnostic Skeletal Animal Reconstruction	28559
<i>Zeyuan Chen, Shaowei Liu, Zhongzheng Ren, Shenlong Wang, Yuefan Wu</i>	
Regret Bounds for Information-Directed Reinforcement Learning.....	28575
<i>Botao Hao, Tor Lattimore</i>	
Adversarial Reprogramming Revisited	28588
<i>Matthias Englert, Ranko Lazic</i>	
Coded Residual Transform for Generalizable Deep Metric Learning	28601
<i>Yigang Cen, Zhihai He, Shichao Kan, Min Li, Yixiong Liang, Jianxin Wang</i>	
When Does Differentially Private Learning Not Suffer in High Dimensions?.....	28616
<i>Abhradeep Guha Thakurta, Tatsunori Hashimoto, Huseyin A. Inan, Janardhan Kulkarni, Yin-Tat Lee, Xuechen Li, Daogao Liu</i>	
Nearly Optimal Best-Of-Both-Worlds Algorithms for Online Learning with Feedback Graphs	28631
<i>Junya Honda, Shinji Ito, Taira Tsuchiya</i>	
Few-Shot Task-Agnostic Neural Architecture Search for Distilling Large Language Models.....	28644
<i>Ahmed Awadallah, Debadeepta Dey, Jianfeng Gao, Xiaodong Liu, Subhabrata Mukherjee, Wenhui Wang, Dongkuan DK Xu, Xiang Zhang</i>	
Universal Rates for Interactive Learning.....	28657
<i>Steve Hanneke, Amin Karbasi, Shay Moran, Grigoris Velegkas</i>	
Learning Active Camera for Multi-Object Navigation.....	28670
<i>Peihao Chen, Chuang Gan, Weiwen Hu, Wenbing Huang, Dongyu Ji, Thomas Li, Kunyang Lin, Mingkui Tan</i>	
Knowledge Distillation: Bad Models Can Be Good Role Models	28683
<i>Gal Kaplun, Eran Malach, Preetum Nakkiran, Shai Shalev-Shwartz</i>	
On Computing Probabilistic Explanations for Decision Trees	28695
<i>Marcelo Arenas, Pablo Barceló, Miguel Romero Orth, Bernardo Subercaseaux</i>	
Masked Autoencoders that Listen.....	28708
<i>Michael Auli, Alexei Baevski, Christoph Feichtenhofer, Wojciech Galuba, Po-Yao Huang, Jun Cheng Li, Florian Metze, Hu Xu</i>	
AUTOMATA: Gradient Based Data Subset Selection for Compute-Efficient Hyper-Parameter Tuning	28721
<i>Guttu Sai Abhishek, Alexandre Evfimievski, Rishabh Iyer, Krishnateja Killamsetty, Aakriti Lnu, Lucian Popa, Ganesh Ramakrishnan</i>	
DeepTOP: Deep Threshold-Optimal Policy for MDPs and RMABs	28734
<i>I-Hong Hou, Khaled Nakhleh</i>	

Is a Modular Architecture Enough?.....	28747
<i>Yoshua Bengio, Guillaume Lajoie, Sarthak Mittal</i>	
Exploration Via Planning for Information About the Optimal Trajectory	28761
<i>Joseph Abbate, Mark Boyer, Ian Char, Rory Conlin, Stefano Ermon, Viraj Mehta, Willie Neiswanger, Jeff Schneider</i>	
Subquadratic Kronecker Regression with Applications to Tensor Decomposition	28776
<i>Matthew Fahrbach, Gang Fu, Mehrdad Ghadiri</i>	
Robust Anytime Learning of Markov Decision Processes	28790
<i>Nils Jansen, David Parker, Thiago D. Simão, Marnix Suilen</i>	
Discovering Design Concepts for CAD Sketches	28803
<i>Hao Pan, Yuezhi Yang</i>	
Learning Distributed and Fair Policies for Network Load Balancing as Markov Potential Game.....	28815
<i>Zihan Ding, Zhiyuan Yao</i>	
Para-CFlows: C^k -Universal Diffeomorphism Approximators as Superior Neural Surrogates	28829
<i>Zhitang Chen, Wenjing Cun, Chang Feng, Yanhui Geng, Junlong Lyu, Zhijie Xu, Chen Yongwei, Shengyu Zhu</i>	
NCP: Neural Correspondence Prior for Effective Unsupervised Shape Matching.....	28842
<i>Souhaib Attaiki, Maks Ovsjanikov</i>	
Efficient Spatially Sparse Inference for Conditional GANs and Diffusion Models	28858
<i>Stefano Ermon, Song Han, Muyang Li, Ji Lin, Chenlin Meng, Jun-Yan Zhu</i>	

VOLUME 38

A Greek Parliament Proceedings Dataset for Computational Linguistics and Political Analysis.....	28874
<i>Konstantina Dritsa, Panos Louridas, Ioannis Pavlopoulos, Aikaterini Thoma</i>	
Multi-Objective Deep Data Generation with Correlated Property Control	28889
<i>Saleh Alkhalifa, Yuanqi Du, Xiaojie Guo, Austin Leitgeb, Xuanyang Lin, Kevin Minbiole, Bo Pan, Ashley Petersen, Amarda Shehu, Shiyu Wang, Yinkai Wang, William M. Wuest, Yanfang Ye, Liang Zhao</i>	
Domain Adaptation Meets Individual Fairness - and They Get Along	28902
<i>Debarghya Mukherjee, Felix Petersen, Yuekai Sun, Mikhail Yurochkin</i>	
An Asymptotically Optimal Batched Algorithm for the Dueling Bandit Problem	28914
<i>Arpit Agarwal, Rohan Ghuge, Viswanath Nagarajan</i>	
Enhanced Bilevel Optimization Via Bregman Distance.....	28928
<i>Shangqian Gao, Heng Huang, Feihu Huang, Junyi Li</i>	
SAVi++: Towards End-To-End Object-Centric Learning from Real-World Videos	28940
<i>Gamaleldin Elsayed, Klaus Greff, Thomas Kipf, Aravindh Mahendran, Michael Mozer, Sjoerd Van Steenkiste</i>	
Reincarnating Reinforcement Learning: Reusing Prior Computation to Accelerate Progress	28955
<i>Rishabh Agarwal, Marc Bellemare, Pablo Samuel Castro, Aaron Courville, Max Schwarzer</i>	

Global Linear and Local Superlinear Convergence of IRLS for Non-Smooth Robust Regression	28972
<i>Christian Kümmerle, Liangzu Peng, Rene Vidal</i>	
Rashomon Capacity: A Metric for Predictive Multiplicity in Classification	28988
<i>Flavio Calmon, Hsiang Hsu</i>	
Inducing Equilibria Via Incentives: Simultaneous Design-And-Play Ensures Global Convergence	29001
<i>Mingyi Hong, Jiayang Li, Boyi Liu, Yu Nie, Hoi-To Wai, Zhaoran Wang, Zhuoran Yang</i>	
Learning with Convolution and Pooling Operations in Kernel Methods	29014
<i>Song Mei, Theodor Misiakiewicz</i>	
In Differential Privacy, There is Truth: On Vote-Histogram Leakage in Ensemble Private Learning.....	29026
<i>David Lie, Nicolas Papernot, Roei Schuster, I Shumailov, Jiaqi Wang</i>	
Estimating Graphical Models for Count Data with Applications to Single-Cell Gene Network.....	29038
<i>Huaying Fang, Junjie Tang, Ruibin Xi, Feiyi Xiao</i>	
Online Minimax Multiobjective Optimization: Multicalibeating and Other Applications	29051
<i>Daniel Lee, Georgy Noarov, Malleh Pai, Aaron Roth</i>	
Roadblocks for Temporarily Disabling Shortcuts and Learning New Knowledge.....	29064
<i>Bin Li, Hanting Li, Hongjing Niu, Feng Zhao</i>	
Random Rank: The One and Only Strategyproof and Proportionally Fair Randomized Facility Location Mechanism	29076
<i>Haris Aziz, Alexander Lam, Mashbat Suzuki, Toby Walsh</i>	
Distributionally Robust Weighted K-Nearest Neighbors.....	29088
<i>Rui Gao, Liyan Xie, Yao Xie, Minghe Zhang, Shixiang Zhu</i>	
A Character-Level Length-Control Algorithm for Non-Autoregressive Sentence Summarization	29101
<i>Puyuan Liu, Lili Mou, Xiang Zhang</i>	
Neural Lyapunov Control of Unknown Nonlinear Systems with Stability Guarantees.....	29113
<i>Hans De Sterck, Jun Liu, Thanin Quartz, Ruikun Zhou</i>	
Minimax Regret for Cascading Bandits	29126
<i>Sujay Sanghavi, Sanjay Shakkottai, R. Srikant, Daniel Vial</i>	
Mean Estimation with User-Level Privacy Under Data Heterogeneity.....	29139
<i>Rachel Cummings, Vitaly Feldman, Audra McMillan, Kunal Talwar</i>	
Propo: A Message Passing Framework for Customizable and Composable Learning Algorithms	29152
<i>Paavo Parmas, Takuma Seno</i>	
A Lower Bound of Hash Codes' Performance.....	29166
<i>Lianli Gao, Yu Lei, Hengtao Shen, Jingkuan Song, Xiaosu Zhu</i>	
Self-Supervised Image Restoration with Blurry and Noisy Pairs.....	29179
<i>Ming Liu, Rongjian Xu, Zifei Yan, Zhilu Zhang, Wangmeng Zuo</i>	
Embracing Consistency: A One-Stage Approach for Spatio-Temporal Video Grounding.....	29192
<i>Yang Jin, Yongzhi Li, Yadong Mu, Zehuan Yuan</i>	
Pitfalls of Epistemic Uncertainty Quantification Through Loss Minimisation	29205
<i>Viktor Bengs, Eyke Hüllermeier, Willem Waegeman</i>	

Pile of Law: Learning Responsible Data Filtering from the Law and a 256GB Open-Source Legal Dataset.....	29217
<i>Neel Guha, Peter Henderson, Daniel Ho, Dan Jurafsky, Mark Krass, Christopher D Manning, Lucia Zheng</i>	
Active-Passive SimStereo - Benchmarking the Cross-Generalization Capabilities of Deep Learning-Based Stereo Methods	29235
<i>Allen Antony, Mohammed Bennamoun, Farid Boussaid, Laurent Jospin, Hamid Laga, Lian Xu</i>	
Back Razor: Memory-Efficient Transfer Learning by Self-Sparsified Backpropagation.....	29248
<i>Xuxi Chen, Xianzhi Du, Xueqin Huang, Ziyu Jiang, Zhangyang Wang, Denny Zhou</i>	
Patching Open-Vocabulary Models by Interpolating Weights.....	29262
<i>Ali Farhadi, Samir Yitzhak Gadre, Hannaneh Hajishirzi, Gabriel Ilharco, Simon Kornblith, Ludwig Schmidt, Shuran Song, Mitchell Wortsman</i>	
Training Stochastic Stabilized Supralinear Networks by Dynamics-Neutral Growth.....	29278
<i>Mate Lengyel, Wayne Soo</i>	
Post-Hoc Estimators for Learning to Defer to an Expert.....	29292
<i>Wittawat Jitkrittum, Sanjiv Kumar, Aditya Menon, Harikrishna Narasimhan, Ankit Rawat</i>	
Learning Generalizable Part-Based Feature Representation for 3D Point Clouds	29305
<i>Xiang Gu, Jian Sun, Xin Wei</i>	
FourierFormer: Transformer Meets Generalized Fourier Integral Theorem.....	29319
<i>Nhat Ho, Khai Nguyen, Tam Nguyen, Tan Nguyen, Stanley Osher, Minh Pham</i>	
Learning Multi-Resolution Functional Maps with Spectral Attention for Robust Shape Matching.....	29336
<i>Nicolas Donati, Lei Li, Maks Ovsjanikov</i>	
Label-Invariant Augmentation for Semi-Supervised Graph Classification	29350
<i>Hongfu Liu, Han Yue, Chunhui Zhang, Chuxu Zhang</i>	
Practical Adversarial Multivald Conformal Prediction	29362
<i>Osbert Bastani, Varun Gupta, Christopher Jung, Georgy Noarov, Ramya Ramalingam, Aaron Roth</i>	
Test-Time Training with Masked Autoencoders.....	29374
<i>Xinlei Chen, Alexei Efros, Yossi Gandelsman, Yu Sun</i>	
Constraining Gaussian Processes to Systems of Linear Ordinary Differential Equations.....	29386
<i>Andreas Besginow, Markus Lange-Hegermann</i>	
So3krates: Equivariant Attention for Interactions on Arbitrary Length-Scales in Molecular Systems.....	29400
<i>Thorben Frank, Klaus-Robert Müller, Oliver Unke</i>	
HyperDomainNet: Universal Domain Adaptation for Generative Adversarial Networks.....	29414
<i>Aibek Alanov, Vadim Titov, Dmitry Vetrov</i>	
Multiagent Q-Learning with Sub-Team Coordination.....	29427
<i>Xiaotie Deng, Jianye Hao, Wenhan Huang, Kai Li, Jun Luo, Hangyu Mao, Kun Shao, Matthew Taylor, Jun Wang, Dongge Wang, Tianze Zhou</i>	
CLiMB: A Continual Learning Benchmark for Vision-And-Language Tasks.....	29440
<i>Ting-Yun Chang, Georgios Chochlakis, Leticia Pinto Alva, Mohammad Rostami, Tejas Srinivasan, Jesse Thomason</i>	

Bidirectional Learning for Offline Infinite-Width Model-Based Optimization.....	29454
<i>Can Chen, Mark Coates, Jie Fu, Xue Steve Liu, Yingxueff Zhang</i>	
Differentially Private Model Compression.....	29468
<i>Arturs Backurs, Huseyin A. Inan, Janardhan Kulkarni, Fatemehsadat Mireshghallah, Lukas Wutschitz</i>	
Dual-Curriculum Contrastive Multi-Instance Learning for Cancer Prognosis Analysis with Whole Slide Images	29484
<i>Zhenyuan Ning, Chao Tu, Yu Zhang</i>	
Meta-Complementing the Semantics of Short Texts in Neural Topic Models.....	29498
<i>Hady Lauw, Delvin Ce Zhang</i>	
Unified Optimal Transport Framework for Universal Domain Adaptation.....	29512
<i>Wanxing Chang, Ye Shi, Hoang Tuan, Jingya Wang</i>	
GBA: A Tuning-Free Approach to Switch Between Synchronous and Asynchronous Training for Recommendation Models.....	29525
<i>Yufeng Cai, Jing Chen, Hongbo Deng, Lin Qu, Kaixu Ren, Yue Song, Wenbo Su, Pengjie Wang, Jian Xu, Huimin Yi, Yuanxing Zhang, Bo Zheng</i>	
TUSK: Task-Agnostic Unsupervised Keypoints	29538
<i>Jan Hosang, Yuhe Jin, Weiwei Sun, Eduard Trulls, Kwang Moo Yi</i>	
Multi-Block Min-Max Bilevel Optimization with Applications in Multi-Task Deep AUC Maximization.....	29552
<i>Quanqi Hu, Tianbao Yang, Yongjian Zhong</i>	
Coresets for Vertical Federated Learning: Regularized Linear Regression and K-Means Clustering.....	29566
<i>Lingxiao Huang, Zhize Li, Jialin Sun, Haoyu Zhao</i>	
Class-Aware Adversarial Transformers for Medical Image Segmentation.....	29582
<i>Sandeep Chinchali, Siyuan Dong, James Duncan, Fenglin Liu, Lawrence Staib, Ufuk Topcu, Chenyu You, Ruihan Zhao</i>	
MCL-GAN: Generative Adversarial Networks with Multiple Specialized Discriminators.....	29597
<i>Jinyoung Choi, Bohyung Han</i>	
Unsupervised Visual Representation Learning Via Mutual Information Regularized Assignment	29610
<i>Sungik Choi, Sae-Young Chung, Hyunwoo Kim, Dong Hoon Lee</i>	
Mind Reader: Reconstructing Complex Images from Brain Activities	29624
<i>Sikun Lin, Ambuj K Singh, Thomas Sprague</i>	
A Few Expert Queries Suffices for Sample-Efficient RL with Resets and Linear Value Approximation.....	29637
<i>Philip Amortila, Dean Foster, Nan Jiang, Dhruv Madeka</i>	

VOLUME 39

Subsidiary Prototype Alignment for Universal Domain Adaptation	29649
<i>Suvaansh Bhambri, Varun Jampani, Akshay R Kulkarni, Jogendra Nath Kundu, Venkatesh Babu R, Hiran Sarkar</i>	

Dynamic Inverse Reinforcement Learning for Characterizing Animal Behavior	29663
<i>Zoe Ashwood, Aditi Jha, Jonathan Pillow</i>	
FedRolex: Model-Heterogeneous Federated Learning with Rolling Sub-Model Extraction.....	29677
<i>Samiul Alam, Luyang Liu, Ming Yan, Mi Zhang</i>	
Multi-Lingual Acquisition on Multimodal Pre-Training for Cross-Modal Retrieval	29691
<i>Anwen Hu, Qin Jin, Liang Zhang</i>	
Manifold Interpolating Optimal-Transport Flows for Trajectory Inference	29705
<i>Oluwadamilola Fasina, Guillaume Huguet, Smita Krishnaswamy, Manik Kuchroo, Daniel Sumner Magruder, Alexander Tong, Guy Wolf</i>	
Depth is More Powerful than Width with Prediction Concatenation in Deep Forest	29719
<i>Yi-Xiao He, Shen-Huan Lyu, Zhi-Hua Zhou</i>	
Logical Activation Functions: Logit-Space Equivalents of Probabilistic Boolean Operators	29733
<i>Jason D'Eon, Robert Earle, Scott Lowe, Sageev Oore, Thomas Trappenberg</i>	
An Investigation into Whitening Loss for Self-Supervised Learning.....	29748
<i>Rao Anwer, Lei Huang, Salman Khan, Fahad Shahbaz Khan, Xi Weng, Lei Zhao</i>	
Geometric Knowledge Distillation: Topology Compression for Graph Neural Networks	29761
<i>Qitian Wu, Junchi Yan, Chenxiao Yang</i>	
A Benchmark for Compositional Visual Reasoning	29776
<i>Julien Colin, Sebastian Musslick, Thomas Serre, Mohit Vaishnav, Aimen Zerroug</i>	
Learning Articulated Rigid Body Dynamics with Lagrangian Graph Neural Network.....	29789
<i>Ravinder Bhattoo, N M Anoop Krishnan, Sayan Ranu</i>	
Myriad: A Real-World Testbed to Bridge Trajectory Optimization and Deep Learning	29801
<i>Pierre-Luc Bacon, Simon Dufort-Labbé, Nikolaus Howe, Nitarshan Rajkumar</i>	
Batch Bayesian Optimisation Via Density-Ratio Estimation with Guarantees.....	29816
<i>Rafael Oliveira, Fabio Ramos, Louis Tiao</i>	
Amplifying Membership Exposure Via Data Poisoning	29830
<i>Yufei Chen, Chao Shen, Yun Shen, Cong Wang, Yang Zhang</i>	
Boosting the Transferability of Adversarial Attacks with Reverse Adversarial Perturbation	29845
<i>Yanbo Fan, Yi Liu, Zeyu Qin, Li Shen, Jue Wang, Baoyuan Wu, Yong Zhang</i>	
Extra-Newton: A First Approach to Noise-Adaptive Accelerated Second-Order Methods	29859
<i>Kimon Antonakopoulos, Volkan Cevher, Ali Kavis</i>	
FETA: Towards Specializing Foundational Models for Expert Task Applications	29873
<i>Amit Alfassy, Assaf Arbelle, Christoph Auer, Michele Dolfi, Rogerio Feris, Oshri Halimi, Sivan Harary, Roei Herzig, Leonid Karlinsky, Rameswar Panda, Kate Saenko, Eli Schwartz, Peter Staar</i>	
Global Convergence of Federated Learning for Mixed Regression	29889
<i>Lili Su, Jiaming Xu, Pengkun Yang</i>	
BayesPCN: A Continually Learnable Predictive Coding Associative Memory.....	29903
<i>Frank Wood, Jinsoo Yoo</i>	

Optimizing Data Collection for Machine Learning.....	29915
<i>Jose M. Alvarez, Sanja Fidler, Marc Law, James Lucas, Rafid Mahmood</i>	
DP-PCA: Statistically Optimal and Differentially Private PCA.....	29929
<i>Prateek Jain, Weihao Kong, Xiyang Liu, Sewoong Oh</i>	
Semantic Probabilistic Layers for Neuro-Symbolic Learning.....	29944
<i>Kareem Ahmed, Kai-Wei Chang, Stefano Teso, Guy Van Den Broeck, Antonio Vergari</i>	
Adapting to Online Label Shift with Provable Guarantees	29960
<i>Yong Bai, Masashi Sugiyama, Yu-Jie Zhang, Peng Zhao, Zhi-Hua Zhou</i>	
CAGroup3D: Class-Aware Grouping for 3D Object Detection on Point Clouds.....	29975
<i>Lihe Ding, Shaocong Dong, Zhenguo Li, Aoxue Li, Jianan Li, Shaoshuai Shi, Liwei Wang, Haiyang Wang</i>	
Revisiting Injective Attacks on Recommender Systems	29989
<i>Lei Chen, Shimin Di, Haoyang Li</i>	
Learning-Based Motion Planning in Dynamic Environments Using GNNs and Temporal Encoding	30003
<i>Jingkai Chen, Chuchu Fan, Sicun Gao, Chenning Yu, Ruipeng Zhang</i>	
Training Compute-Optimal Large Language Models.....	30016
<i>Sebastian Borgeaud, Elena Buchatskaya, Trevor Cai, Aidan Clark, Bogdan Damoc, Diego De Las Casas, Erich Elsen, Aurelia Guy, Lisa Anne Hendricks, Thomas Hennigan, Jordan Hoffmann, Arthur Mensch, Katherine Millican, Eric Noland, Simon Osindero, Jack Rae, Eliza Rutherford, Laurent Sifre, Karén Simonyan, George Van Den Driessche, Oriol Vinyals, Johannes Welbl</i>	
Robust Imitation Via Mirror Descent Inverse Reinforcement Learning.....	30031
<i>Dong-Sig Han, Hyunseo Kim, Hyundo Lee, Jehwan Ryu, Byoung-Tak Zhang</i>	
Characterizing Datapoints Via Second-Split Forgetting.....	30044
<i>Saurabh Garg, J. Zico Kolter, Zachary Lipton, Pratyush Maini</i>	
VoiceBlock: Privacy Through Real-Time Adversarial Attacks with Audio-To-Audio Models	30058
<i>Keshav Bhandari, Andreas Bugler, Max Morrison, Patrick O'Reilly, Bryan Pardo</i>	
A Theoretical View on Sparsely Activated Networks.....	30071
<i>Cenk Baykal, Nishanth Dikkala, Rina Panigrahy, Cyrus Rashtchian, Xin Wang</i>	
Invariance Learning Based on Label Hierarchy	30085
<i>Kenji Fukumizu, Shoji Toyota</i>	
Effects of Data Geometry in Early Deep Learning.....	30099
<i>George Konidaris, Saket Tiwari</i>	
When to Intervene: Learning Optimal Intervention Policies for Critical Events.....	30114
<i>Chiranjib Bhattacharyya, Niranjan Damera Venkata</i>	
Draft-And-Revise: Effective Image Generation with Contextual RQ-Transformer	30127
<i>Minsu Cho, Wook Shin Han, Chiheon Kim, Saehoon Kim, Doyup Lee</i>	
Finite-Sample Maximum Likelihood Estimation of Location.....	30139
<i>Shivam Gupta, Jasper Lee, Eric Price, Paul Valiant</i>	

GENIE: Higher-Order Denoising Diffusion Solvers	30150
<i>Tim Dockhorn, Karsten Kreis, Arash Vahdat</i>	
Provably Adversarially Robust Detection of Out-Of-Distribution Data (Almost) for Free.....	30167
<i>Julian Bitterwolf, Matthias Hein, Alexander Meinke</i>	
Efficient Graph Similarity Computation with Alignment Regularization	30181
<i>Guang Tan, Wei Zhuo</i>	
Tsetlin Machine for Solving Contextual Bandit Problems	30194
<i>Ole-Christoffer Granmo, Raihan Seraj, Jivitesh Sharma</i>	
Prune and Distill: Similar Reformatting of Image Information Along Rat Visual Cortex and Deep Neural Networks.....	30206
<i>Alessandro Laio, Paolo Muratore, Eugenio Piasini, Sina Tafazoli, Davide Zoccolan</i>	
Graph Scattering Beyond Wavelet Shackles.....	30219
<i>Christian Koke, Gitta Kutyniok</i>	
Matryoshka Representation Learning.....	30233
<i>Gantavya Bhatt, Kaifeng Chen, Ali Farhadi, William Howard-Snyder, Prateek Jain, Sham Kakade, Aditya Kusupati, Vivek Ramanujan, Aniket Rege, Aditya Sinha, Matthew Wallingford</i>	
Off-Policy Evaluation with Deficient Support Using Side Information.....	30250
<i>Paolo Cremonesi, Nicolò Felicioni, Maurizio Ferrari Dacrema, Marcello Restelli</i>	
The Nature of Temporal Difference Errors in Multi-Step Distributional Reinforcement Learning	30265
<i>Bernardo Avila Pires, Marc Bellemare, Will Dabney, Remi Munos, Mark Rowland, Yunhao Tang</i>	
GraphDE: A Generative Framework for Debaised Learning and Out-Of-Distribution Detection on Graphs	30277
<i>Zenan Li, Fan Nie, Qitian Wu, Junchi Yan</i>	
Expectation-Maximization Contrastive Learning for Compact Video-And-Language Representations	30291
<i>Jie Chen, David Clifton, Shen Ge, Jinfa Huang, Peng Jin, Fenglin Liu, Guoli Song, Xian Wu</i>	
Accelerating Sparse Convolution with Column Vector-Wise Sparsity.....	30307
<i>Yunji Chen, Zidong Du, Kai Han, Yijun Tan, Yunhe Wang, Jun Yao, Xianzhi Yu, Kang Zhao</i>	
LLM.int8(): 8-Bit Matrix Multiplication for Transformers at Scale.....	30318
<i>Younes Belkada, Tim Dettmers, Mike Lewis, Luke Zettlemoyer</i>	
Posterior Refinement Improves Sample Efficiency in Bayesian Neural Networks.....	30333
<i>Runa Eschenhagen, Philipp Hennig, Agustinus Kristiadi</i>	
Micro and Macro Level Graph Modeling for Graph Variational Auto-Encoders	30347
<i>Ke Li, Parnis Naddaf, Oliver Schulte, Kiarash Zahrimia</i>	
Spectral Bias Outside the Training Set for Deep Networks in the Kernel Regime.....	30362
<i>Benjamin Bowman, Guido Montufar</i>	
The Unreliability of Explanations in Few-Shot Prompting for Textual Reasoning	30378
<i>Greg Durrett, Xi Ye</i>	

Self-Similarity Priors: Neural Collages as Differentiable Fractal Representations	30393
<i>Stefano Ermon, Kuno Kim, Stefano Massaroli, Chenlin Meng, Michael Poli, Winnie Xu</i>	
Scalable Multi-Agent Covering Option Discovery Based on Kronecker Graphs.....	30406
<i>Vaneet Aggarwal, Jiayu Chen, Jingdi Chen, Tian Lan</i>	
An Algorithm for Learning Switched Linear Dynamics from Data	30419
<i>Guillaume Berger, Morteza Lahijanian, Monal Narasimhamurthy, Sriram Sankaranarayanan, Kandai Watanabe</i>	
LECO: Learnable Episodic Count for Task-Specific Intrinsic Reward	30432
<i>Daejin Jo, Jongmin Kim, Sungwoong Kim, Taehwan Kwon, Donghoon Lee, Daniel Nam, Seungeun Rho</i>	

VOLUME 40

Generalised Implicit Neural Representations	30446
<i>Daniele Grattarola, Pierre Vandergheynst</i>	
EZNAS: Evolving Zero-Cost Proxies for Neural Architecture Scoring	30459
<i>Yash Akhauri, Ravishankar Iyer, Nilesh Jain, Juan Munoz</i>	
Personalized Federated Learning Towards Communication Efficiency, Robustness and Fairness	30471
<i>Yuze Han, Xiang Li, Shiyun Lin, Zhihua Zhang</i>	
Local Linear Convergence of Gradient Methods for Subspace Optimization Via Strict Complementarity	30486
<i>Ron Fisher, Dan Garber</i>	
Joint Learning of 2D-3D Weakly Supervised Semantic Segmentation	30499
<i>Hyeokjun Kweon, Kuk-Jin Yoon</i>	
RNNs of RNNs: Recursive Construction of Stable Assemblies of Recurrent Neural Networks.....	30512
<i>Michaela Ennis, Leo Kozachkov, Jean-Jacques Slotine</i>	
Training Spiking Neural Networks with Event-Driven Backpropagation	30528
<i>Wei Fang, Tiejun Huang, Timothée Masquelier, Xiaodong Xie, Zhaofei Yu, Yaoyu Zhu</i>	
Trade-Off Between Payoff and Model Rewards in Shapley-Fair Collaborative Machine Learning	30542
<i>Patrick Jaillet, Bryan Kian Hsiang Low, Quoc Phong Nguyen</i>	
Efficient Non-Parametric Optimizer Search for Diverse Tasks	30554
<i>Minhao Cheng, Cho-Jui Hsieh, Ruochen Wang, Yuanhao Xiong</i>	
DualCoOp: Fast Adaptation to Multi-Label Recognition with Limited Annotations	30569
<i>Ping Hu, Kate Saenko, Ximeng Sun</i>	
What Can Transformers Learn In-Context? a Case Study of Simple Function Classes	30583
<i>Shivam Garg, Percy Liang, Dimitris Tsipras, Gregory Valiant</i>	
Towards Robust Blind Face Restoration with Codebook Lookup Transformer	30599
<i>Kelvin Chan, Chongyi Li, Chen Change Loy, Shangchen Zhou</i>	
Few-Shot Continual Active Learning by a Robot.....	30612
<i>Ali Ayub, Carter Fendley</i>	

Learning Generalized Policy Automata for Relational Stochastic Shortest Path Problems.....	30625
<i>Rushang Karia, Rashmeet Kaur Nayyar, Siddharth Srivastava</i>	
A Stochastic Linearized Augmented Lagrangian Method for Decentralized Bilevel Optimization	30638
<i>Xiaodong Cui, Mingyi Hong, Lior Horesh, Brian Kingsbury, Jia Liu, Songtao Lu, Mark Squillante, Siliang Zeng</i>	
Boosting the Performance of Generic Deep Neural Network Frameworks with Log-Supermodular CRFs.....	30651
<i>Yangxiao Lu, Nicholas Ruoizzi, Hao Xiong</i>	
Don't Roll the Dice, Ask Twice: The Two-Query Distortion of Matching Problems and Beyond	30665
<i>Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Voudouris</i>	
On the Effect of Pre-Training for Transformer in Different Modality on Offline Reinforcement Learning	30678
<i>Shiro Takagi</i>	
Look Where You Look! Saliency-Guided Q-Networks for Generalization in Visual Reinforcement Learning	30693
<i>David Bertoin, Emmanuel Rachelson, Adil Zouitine, Mehdi Zouitine</i>	
Information-Theoretic Safe Exploration with Gaussian Processes.....	30707
<i>Felix Berkenkamp, Alessandro Bottero, Carlos Luis, Jan Peters, Julia Vinogradska</i>	
On the Consistent Estimation of Optimal Receiver Operating Characteristic (ROC) Curve	30720
<i>Renxiong Liu, Yunzhang Zhu</i>	
Online Algorithms for the Santa Claus Problem	30732
<i>Mohammadtaghi Hajiaghayi, Mohammad Khani, Debmalya Panigrahi, Max Springer</i>	
Instance-Based Learning for Knowledge Base Completion.....	30744
<i>Xingran Chen, Wanyun Cui</i>	
Decomposable Non-Smooth Convex Optimization with Nearly-Linear Gradient Oracle Complexity.....	30756
<i>Sally Dong, Haotian Jiang, Yin Tat Lee, Swati Padmanabhan, Guanghao Ye</i>	
On Optimal Learning Under Targeted Data Poisoning.....	30770
<i>Steve Hanneke, Amin Karbasi, Mohammad Mahmoody, Idan Mehalal, Shay Moran</i>	
Generalized Laplacian Eigenmaps	30783
<i>Piotr Koniusz, Hao Zhu</i>	
OGC: Unsupervised 3D Object Segmentation from Rigid Dynamics of Point Clouds.....	30798
<i>Ziyang Song, Bo Yang</i>	
Look More but Care Less in Video Recognition	30813
<i>Yue Bai, Yun Fu, Huan Wang, Yi Xu, Yitian Zhang</i>	
Transferring Pre-Trained Multimodal Representations with Cross-Modal Similarity Matching	30826
<i>Sungik Choi, Dasol Hwang, Byoungjip Kim, Honglak Lee, Moontae Lee</i>	
Get More at Once: Alternating Sparse Training with Gradient Correction	30840
<i>Deliang Fan, Jian Meng, Jae-Sun Seo, Li Yang</i>	
BLOX: Macro Neural Architecture Search Benchmark and Algorithms	30851
<i>Mohamed Abdelfattah, Thomas Chau, Lukasz Dudziak, Nicholas Lane, Hongkai Wen</i>	

Does Self-Supervised Learning Really Improve Reinforcement Learning from Pixels?	30865
<i>Srijan Das, Xiang Li, Michael Ryoo, Jinghuan Shang</i>	
TCT: Convexifying Federated Learning Using Bootstrapped Neural Tangent Kernels	30882
<i>Michael Jordan, Sai Praneeth Karimireddy, Yi Ma, Alexander Wei, Yaodong Yu</i>	
Graph Coloring Via Neural Networks for Haplotype Assembly and Viral Quasispecies Reconstruction.....	30898
<i>Yu Lin, Vaibhav Rajan, Hansheng Xue</i>	
Active Learning of Classifiers with Label and Seed Queries	30911
<i>Marco Bressan, Nicolò Cesa-Bianchi, Silvio Lattanzi, Andrea Paudice, Maximilian Thiessen</i>	
TANGO: Text-Driven Photorealistic and Robust 3D Stylization Via Lighting Decomposition.....	30923
<i>Yongwei Chen, Kui Jia, Jiabao Lei, Chen Rui, Yabin Zhang</i>	
Analyzing Lottery Ticket Hypothesis from PAC-Bayesian Theory Perspective	30937
<i>Keitaro Sakamoto, Issei Sato</i>	
Make Sharpness-Aware Minimization Stronger: A Sparsified Perturbation Approach	30950
<i>Rongrong Ji, Peng Mi, Tianhe Ren, Li Shen, Xiaoshuai Sun, Dacheng Tao, Yiyi Zhou</i>	
From Gradient Flow on Population Loss to Learning with Stochastic Gradient Descent	30963
<i>Christopher De Sa, Satyen Kale, Jason Lee, Ayush Sekhari, Karthik Sridharan</i>	
Distributional Reinforcement Learning for Risk-Sensitive Policies.....	30977
<i>Shiau Hong Lim, Ilyas Malik</i>	
Text Classification with Born's Rule	30990
<i>Alfio Ferrara, Emanuele Guidotti</i>	
Learning a Condensed Frame for Memory-Efficient Video Class-Incremental Learning	31002
<i>Jun Cen, Yixuan Pei, Xueming Qian, Zhiwu Qing, Nong Sang, Mingqian Tang, Xiang Wang, Yaxiong Wang, Shiwei Zhang</i>	
Conformal Prediction with Temporal Quantile Adjustments.....	31017
<i>Zhen Lin, Jimeng Sun, Shubhendu Trivedi</i>	
Learning Viewpoint-Agnostic Visual Representations by Recovering Tokens in 3D Space	31031
<i>Srijan Das, Michael Ryoo, Jinghuan Shang</i>	
Certifying Some Distributional Fairness with Subpopulation Decomposition.....	31045
<i>Mintong Kang, Bo Li, Linyi Li, Yang Liu, Maurice Weber, Ce Zhang</i>	
A Theoretical Understanding of Gradient Bias in Meta-Reinforcement Learning	31059
<i>Xidong Feng, Bo Liu, Luo Mai, Jie Ren, Jun Wang, Yaodong Yang, Haifeng Zhang, Rui Zhu</i>	
Optimal Algorithms for Decentralized Stochastic Variational Inequalities	31073
<i>Aleksandr Beznosikov, Alexander Gasnikov, Dmitry Kovalev, Michael Pershianov, Peter Richtarik, Abdurakhmon Sadiev</i>	
Mirror Descent Maximizes Generalized Margin and Can Be Implemented Efficiently.....	31089
<i>Kwangjun Ahn, Navid Azizan, Haoyuan Sun, Christos Thrampoulidis</i>	
Adversarial Task Up-Sampling for Meta-Learning	31102
<i>Long-Kai Huang, Ying Wei, Yichen Wu</i>	

MAtt: A Manifold Attention Network for EEG Decoding.....	31116
<i>Jing-Lun Chou, Yue-Ting Pan, Chun-Shu Wei</i>	
Geometry-Aware Two-Scale PIFu Representation for Human Reconstruction.....	31130
<i>Hujun Bao, Zheng Dong, Ziheng Duan, Rynson Lau, Ke Xu, Weiwei Xu</i>	
Relational Proxies: Emergent Relationships as Fine-Grained Discriminators.....	31145
<i>Zeynep Akata, Abhra Chaudhuri, Anjan Dutta, Massimiliano Mancini</i>	
Root Cause Analysis of Failures in Microservices Through Causal Discovery	31158
<i>Saurabh Bagchi, Sarthak Chakraborty, Azam Ikram, Murat Kocaoglu, Subrata Mitra, Shiv Saini</i>	
Anticipating Performativity by Predicting from Predictions	31171
<i>Frances Ding, Celestine Mandler-Dünner, Yixin Wang</i>	
Falconn++: A Locality-Sensitive Filtering Approach for Approximate Nearest Neighbor Search.....	31186
<i>Tao Liu, Ninh Pham</i>	
Pre-Trained Language Models for Interactive Decision-Making	31199
<i>Ekin Akyürek, Anima Anandkumar, Jacob Andreas, Tao Chen, Yilun Du, Linxi Fan, De-An Huang, Shuang Li, Igor Mordatch, Chris Paxton, Xavier Puig, Antonio Torralba, Clinton Wang, Yuke Zhu</i>	
Tight Analysis of Extra-Gradient and Optimistic Gradient Methods for Nonconvex Minimax Problems.....	31213
<i>Yuyang Deng, Haochuan Li, Mehrdad Mahdavi, Pouria Mahdavinia</i>	

VOLUME 41

Learning Generalizable Models for Vehicle Routing Problems Via Knowledge Distillation.....	31226
<i>Jieyi Bi, Zhiguang Cao, Yeow Meng Chee, Jinbiao Chen, Yining Ma, Yuan Sun, Jiahai Wang</i>	
Large-Scale Optimization of Partial AUC in a Range of False Positive Rates.....	31239
<i>Qihang Lin, Tianbao Yang, Yao Yao</i>	
The Burer-Monteiro SDP Method Can Fail Even Above the Barvinok-Pataki Bound.....	31254
<i>Liam O'Carroll, Vaidehi Srinivas, Aravindan Vijayaraghavan</i>	
MinVIS: A Minimal Video Instance Segmentation Framework Without Video-Based Training	31265
<i>Anima Anandkumar, De-An Huang, Zhiding Yu</i>	
Supported Policy Optimization for Offline Reinforcement Learning.....	31278
<i>Mingsheng Long, Zihan Qiu, Jianmin Wang, Haixu Wu, Jialong Wu</i>	
Infinite Recommendation Networks: A Data-Centric Approach	31292
<i>Mehak Dhaliwal, Julian McAuley, Noveen Sachdeva, Carole-Jean Wu</i>	
DDXPlus: A New Dataset for Automatic Medical Diagnosis	31306
<i>Arsene Fansi Tchango, Joumana Ghosn, Rishab Goel, Julien Martel, Zhi Wen</i>	
Clipped Stochastic Methods for Variational Inequalities with Heavy-Tailed Noise.....	31319
<i>Marina Danilova, David Dobre, Pavel Dvurechenskii, Alexander Gasnikov, Gauthier Gidel, Eduard Gorbunov</i>	

A Unified Sequence Interface for Vision Tasks.....	31333
<i>Ting Chen, David Fleet, Geoffrey Hinton, Lala Li, Tsung-Yi Lin, Saurabh Saxena</i>	
Expected Frequency Matrices of Elections: Computation, Geometry, and Preference Learning.....	31347
<i>Niclas Boehmer, Robert Bredebeck, Edith Elkind, Piotr Faliszewski, Stanislaw Szufa</i>	
GMMSeg: Gaussian Mixture Based Generative Semantic Segmentation Models.....	31360
<i>Chen Liang, Jiaxu Miao, Wenguan Wang, Yi Yang</i>	
Understanding and Extending Subgraph GNNs by Rethinking Their Symmetries.....	31376
<i>Beatrice Bevilacqua, Michael Bronstein, Fabrizio Frasca, Haggai Maron</i>	
Parametrically Retargetable Decision-Makers Tend to Seek Power.....	31391
<i>Prasad Tadepalli, Alex Turner</i>	
CageNeRF: Cage-Based Neural Radiance Field for Generalized 3D Deformation and Animation.....	31402
<i>Yuhao Cheng, Shanyan Guan, Shengqi Liu, Bowen Pan, Yicong Peng, Yichao Yan, Xiaokang Yang, Guangtao Zhai</i>	
Meta-Query-Net: Resolving Purity-Informativeness Dilemma in Open-Set Active Learning.....	31416
<i>Jihwan Bang, Jae-Gil Lee, Youngjun Lee, Dongmin Park, Yooju Shin, Hwanjun Song</i>	
Dance of SNN and ANN: Solving Binding Problem by Combining Spike Timing and Reconstructive Attention.....	31430
<i>Hui Lin, Luping Shi, Rong Zhao, Hao Zheng</i>	
Unsupervised Learning for Combinatorial Optimization with Principled Objective Relaxation.....	31444
<i>Cong Hao, Pan Li, Haoyu Peter Wang, Nan Wu, Hang Yang</i>	
PAC-Bayes Compression Bounds So Tight that They Can Explain Generalization.....	31459
<i>Marc Finzi, Micah Goldblum, Sanyam Kapoor, Sanae Lotfi, Andres Potapczynski, Andrew Wilson</i>	
BagFlip: A Certified Defense Against Data Poisoning.....	31474
<i>Aws Albarghouthi, Loris D'Antoni, Yuhao Zhang</i>	
Data Augmentation for Efficient Learning from Parametric Experts.....	31484
<i>Alexandre Galashov, Nicolas Heess, Josh Merel</i>	
Symmetry-Induced Disentanglement on Graphs.....	31497
<i>Andre Freitas, Vikas Garg, Giangiacomo Mercatali</i>	
Conformal Off-Policy Prediction in Contextual Bandits.....	31512
<i>Rob Cornish, Arnaud Doucet, Muhammad Faaiz Taufiq, Yee Whye Teh, Jean-Francois Ton</i>	
Black-Box Generalization: Stability of Zeroth-Order Learning.....	31525
<i>Farzin Haddadpour, Dionysis Kalogieras, Amin Karbasi, Konstantinos Nikolakakis</i>	
First Contact: Unsupervised Human-Machine Co-Adaptation Via Mutual Information Maximization.....	31542
<i>Anca Dragan, Sergey Levine, Siddharth Reddy</i>	
HyperMiner: Topic Taxonomy Mining with Hyperbolic Embedding.....	31557
<i>Bo Chen, Zhibin Duan, Ruiying Lu, Dongsheng Wang, Yi Shi Xu, Mingyuan Zhou</i>	
Visual Concepts Tokenization.....	31571
<i>Yan Lu, Yuwang Wang, Tao Yang, Nanning Zheng</i>	

ZooD: Exploiting Model Zoo for Out-Of-Distribution Generalization	31583
<i>Sung-Ho Bae, Qishi Dong, Tianyang Hu, Zhenguo Li, Awais Muhammad, Chuanlong Xie, Yongxin Yang, Fengwei Zhou</i>	
Nearly-Optimal Bounds for Testing Histogram Distributions	31599
<i>Clément L Canonne, Ilias Diakonikolas, Daniel Kane, Sihan Liu</i>	
TGEA 2.0: A Large-Scale Diagnostically Annotated Dataset with Benchmark Tasks for Text Generation of Pretrained Language Models	31612
<i>Huibin Ge, Chuang Liu, Qun Liu, Deyi Xiong, Yulong Zeng, Xiaohu Zhao</i>	
Quantifying Statistical Significance of Neural Network-Based Image Segmentation by Selective Inference	31627
<i>Vo Nguyen Le Duy, Shogo Iwazaki, Ichiro Takeuchi</i>	
A Differentially Private Linear-Time ϵ PTAS for the Minimum Enclosing Ball Problem.....	31640
<i>Bar Mahpud, Or Sheffet</i>	
BEER: Fast $O(1/T)$ Rate for Decentralized Nonconvex Optimization with Communication Compression	31653
<i>Yuejie Chi, Boyue Li, Zhize Li, Peter Richtarik, Haoyu Zhao</i>	
Autoregressive Search Engines: Generating Substrings as Document Identifiers.....	31668
<i>Michele Bevilacqua, Patrick Lewis, Giuseppe Ottaviano, Fabio Petroni, Sebastian Riedel, Scott Yih</i>	
Sampling with Riemannian Hamiltonian Monte Carlo in a Constrained Space.....	31684
<i>Yunbum Kook, Yin-Tat Lee, Ruoqi Shen, Santosh Vempala</i>	
Are All Losses Created Equal: A Neural Collapse Perspective	31697
<i>Xiao Li, Kangning Liu, Sheng Liu, Qing Qu, Chong You, Jinxin Zhou, Zhihui Zhu</i>	
Fair Ranking with Noisy Protected Attributes	31711
<i>Anay Mehrotra, Nisheeth Vishnoi</i>	
Social-Inverse: Inverse Decision-Making of Social Contagion Management with Task Migrations	31726
<i>Guangmo Tong</i>	
Provable Benefit of Multitask Representation Learning in Reinforcement Learning	31741
<i>Yuan Cheng, Songtao Feng, Yingbin Liang, Jing Yang, Hong Zhang</i>	
Modeling Human Exploration Through Resource-Rational Reinforcement Learning	31755
<i>Marcel Binz, Eric Schulz</i>	
Generating Long Videos of Dynamic Scenes	31769
<i>Timo Aila, Miika Aittala, Tim Brooks, Alexei Efros, Janne Hellsten, Tero Karras, Jaakko Lehtinen, Ming-Yu Liu, Ting-Chun Wang</i>	
DigGAN: Discriminator Gradient Gap Regularization for GAN Training with Limited Data	31782
<i>Tiantian Fang, Alex Schwing, Ruoyu Sun</i>	
ORIENT: Submodular Mutual Information Measures for Data Subset Selection Under Distribution Shift	31796
<i>Rishabh Iyer, Athresh Karanam, Krishnateja Killamsetty, Harsha Kokel</i>	

The BigScience ROOTS Corpus: A 1.6TB Composite Multilingual Dataset.....	31809
<i>David Adelani, Christopher Akiki, Khalid Almubarak, Zaid Alyafeai, Loubna Ben Allal, Stella Biderman, Shamik Bose, Jenny Chim, Pierre Colombo, Javier De La Rosa, Francesco De Toni, Manan Dey, Gerard Dupont, Jörg Frohberg, Aaron Gokaslan, Itziar Gonzalez-Dios, Eduardo González Ponferrada, Suzana Ilic, Yacine Jernite, Hugo Laurençon, Violette Lepercq, Teven Le Scao, Quentin Lhoest, Kyle Lo, Shayne Longpre, Sasha Alexandra Luccioni, Maraim Masoud, Angelina McMillan-Major, Margaret Mitchell, Chenghao Mou, Manuel Muñoz, Sebastian Nagel, Huu Nguyen, Olivier Nguyen, Somaieh Nikpoor, Pedro Ortiz Suarez, Suhas Pai, Long Phan, Giada Pistilli, Anna Rogers, Mario Šaško, Lucile Saulnier, Aitor Soroa, Tristan Thrush, Hieu Tran, Daniel Van Strien, Albert Villanova Del Moral, Paulo Villegas, Leandro Von Werra, Minh Chien Vu, Thomas Wang, Leon Weber, Ian Yu, Jian Zhu</i>	
Module-Aware Optimization for Auxiliary Learning	31827
<i>Hong Chen, Chaoyu Guan, Yue Liu, Xin Wang, Yuwei Zhou, Wenwu Zhu</i>	
GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images.....	31841
<i>Wenzheng Chen, Sanja Fidler, Jun Gao, Zan Gojcic, Daiqing Li, Or Litany, Tianchang Shen, Zian Wang, Kangxue Yin</i>	
BYOL-Explore: Exploration by Bootstrapped Prediction.....	31855
<i>Florent Alché, Bernardo Avila Pires, Daniele Calandriello, Mohammad Gheshlaghi Azar, Jean-Bastien Grill, Zhaohan Guo, Remi Munos, Bilal Piot, Miruna Pislari, Alaa Saade, Corentin Tallec, Yunhao Tang, Shantanu Thakoor, Michal Valko</i>	
SizeShiftReg: A Regularization Method for Improving Size-Generalization in Graph Neural Networks	31871
<i>Davide Buffelli, Pietro Lió, Fabio Vandin</i>	
On the Effectiveness of Lipschitz-Driven Rehearsal in Continual Learning.....	31886
<i>Lorenzo Bonicelli, Matteo Boschini, Simone Calderara, Angelo Porrello, Concetto Spampinato</i>	
Interpreting Operation Selection in Differentiable Architecture Search: A Perspective from Influence-Directed Explanations	31902
<i>Wei Huang, Bin Yang, Miao Zhang</i>	
Collaborative Learning of Discrete Distributions Under Heterogeneity and Communication Constraints.....	31915
<i>Edgar Dobriban, Hamed Hassani, Ximmeng Huang, Donghwan Lee</i>	
On the Stability and Scalability of Node Perturbation Learning	31929
<i>Naoki Hiratani, Peter E Latham, Timothy Lillicrap, Yash Mehta</i>	
Unsupervised Learning of Group Invariant and Equivariant Representations	31942
<i>Marco Bertolini, Djork-Arné Clevert, Tuan Le, Frank Noe, Robin Winter</i>	
Factored Adaptation for Non-Stationary Reinforcement Learning.....	31957
<i>Fan Feng, Biwei Huang, Sara Magliacane, Kun Zhang</i>	
MAGNet: Mesh Agnostic Neural PDE Solver	31972
<i>Dan Assouline, Loubna Benabbou, Yoshua Bengio, Oussama Boussif</i>	
TarGF: Learning Target Gradient Field to Rearrange Objects Without Explicit Goal Specification.....	31986
<i>Hao Dong, Mingdong Wu, Yulong Xia, Fangwei Zhong</i>	
Learning to Compare Nodes in Branch and Bound with Graph Neural Networks.....	32000
<i>Didier Chetelat, Abdel Ghani Labassi, Andrea Lodi</i>	

VOLUME 42

Teach Less, Learn More: On the Undistillable Classes in Knowledge Distillation.....	32011
<i>Ning Liu, Xin Liu, Weibin Meng, Zhicai Ou, Jian Tang, Louis Wang, Zhiyuan Xu, Yichen Zhu</i>	
GAGA: Deciphering Age-Path of Generalized Self-Paced Regularizer.....	32025
<i>Bin Gu, Diyang Li, Xingyu Qu, Xiaohan Zhao</i>	
ATD: Augmenting CP Tensor Decomposition by Self Supervision	32039
<i>Cheng Qian, Navjot Singh, Edgar Solomonik, Jimeng Sun, M Westover, Cao Danica Xiao, Chaoqi Yang</i>	
Towards Learning Universal Hyperparameter Optimizers with Transformers.....	32053
<i>Yutian Chen, Nando De Freitas, David Dohan, Arnaud Doucet, Kazuya Kawakami, Greg Kochanski, Chansoo Lee, Sagi Perel, Marc'Aurelio Ranzato, Xingyou Song, Zi Wang, Richard Zhang</i>	
Learning to Reconstruct Missing Data from Spatiotemporal Graphs with Sparse Observations	32069
<i>Cesare Alippi, Andrea Cini, Ivan Marisca</i>	
Robust Graph Structure Learning Via Multiple Statistical Tests	32083
<i>Rong Jin, Ming Lin, Xiuyu Sun, Senzhang Wang, Yaohua Wang, Fangyi Zhang</i>	
Peripheral Vision Transformer.....	32097
<i>Minsu Cho, Chong Luo, Juhong Min, Yucheng Zhao</i>	
Maximum Common Subgraph Guided Graph Retrieval: Late and Early Interaction Networks	32112
<i>Soumen Chakrabarti, Abir De, Indradyumna Roy</i>	
Differentially Private Learning with Margin Guarantees	32127
<i>Raef Bassily, Mehryar Mohri, Ananda Theertha Suresh</i>	
ADBench: Anomaly Detection Benchmark.....	32142
<i>Songqiao Han, Xiyang Hu, Hailiang Huang, Minqi Jiang, Yue Zhao</i>	
GLIF: A Unified Gated Leaky Integrate-And-Fire Neuron for Spiking Neural Networks	32160
<i>Jian Cheng, Fanrong Li, Zitao Mo, Xingting Yao</i>	
HSDF: Hybrid Sign and Distance Field for Modeling Surfaces with Arbitrary Topologies.....	32172
<i>Weikai Chen, Lin Gao, Jintao Li, Xiaoxu Meng, Li Wang, Bo Yang, Jie Yang</i>	
Inverse Game Theory for Stackelberg Games: The Blessing of Bounded Rationality	32186
<i>Fei Fang, Weiran Shen, Jibang Wu, Haifeng Xu</i>	
When Privacy Meets Partial Information: A Refined Analysis of Differentially Private Bandits	32199
<i>Achraf Azize, Debabrota Basu</i>	
Robust Reinforcement Learning Using Offline Data	32211
<i>Mohammad Ghavamzadeh, Dileep Kalathil, Kishan Panaganti, Zaiyan Xu</i>	
Compositional Generalization Through Abstract Representations in Human and Artificial Neural Networks	32225
<i>Michael Cole, Takuya Ito, Tim Klinger, John Murray, Mattia Rigotti, Doug Schultz</i>	
Self-Consistent Dynamical Field Theory of Kernel Evolution in Wide Neural Networks	32240
<i>Blake Bordelon, Cengiz Pehlevan</i>	

Provably Expressive Temporal Graph Networks.....	32257
<i>Vikas Garg, Samuel Kaski, Diego Mesquita, Amauri Souza</i>	
Maximum Likelihood Training of Implicit Nonlinear Diffusion Model	32270
<i>Wanmo Kang, Dongjun Kim, Se Jung Kwon, Dongsoo Lee, Il-Chul Moon, Byeonghu Na</i>	
First is Better than Last for Language Data Influence.....	32285
<i>Frederick Liu, Pradeep Ravikumar, Mukund Sundararajan, Ankur Taly, Chih-Kuan Yeh</i>	
Deep Combinatorial Aggregation.....	32299
<i>Daniel Cremers, Yuesong Shen</i>	
Efficient Coding, Channel Capacity, and the Emergence of Retinal Mosaics	32311
<i>Greg Field, Na Young Jun, John Pearson</i>	
TA-GATES: An Encoding Scheme for Neural Network Architectures	32325
<i>Yiping Deng, Shuang Liang, Xuefei Ning, Changcheng Tang, Yu Wang, Huazhong Yang, Junbo Zhao, Tianchen Zhao, Zixuan Zhou</i>	
ZSON: Zero-Shot Object-Goal Navigation Using Multimodal Goal Embeddings	32340
<i>Gunjan Aggarwal, Dhruv Batra, Bhavika Devnani, Judy Hoffman, Arjun Majumdar</i>	
Autoformalization with Large Language Models.....	32353
<i>Mateja Jamnik, Albert Qiaochu Jiang, Wenda Li, Markus Rabe, Charles Staats, Christian Szegedy, Yuhuai Wu</i>	
Model-Based Lifelong Reinforcement Learning with Bayesian Exploration.....	32369
<i>Haotian Fu, George Konidaris, Michael Littman, Shangqun Yu</i>	
On Efficient Online Imitation Learning Via Classification	32383
<i>Yichen Li, Chicheng Zhang</i>	
Let Images Give You More: Point Cloud Cross-Modal Training for Shape Analysis	32398
<i>Shuguang Cui, Jiantao Gao, Zhen Li, Xu Yan, Heshen Zhan, Ruimao Zhang, Chaoda Zheng</i>	
On Robust Multiclass Learnability.....	32412
<i>Weiwei Liu, Jingyuan Xu</i>	
Debiased Self-Training for Semi-Supervised Learning.....	32424
<i>Baixu Chen, Junguang Jiang, Mingsheng Long, Pengfei Wan, Jianmin Wang, Ximei Wang</i>	
Rethinking Individual Global Max in Cooperative Multi-Agent Reinforcement Learning.....	32438
<i>Yitian Hong, Yaochu Jin, Yang Tang</i>	
Archimedes Meets Privacy: On Privately Estimating Quantiles in High Dimensions Under Minimal Assumptions	32450
<i>Omri Ben-Eliezer, Dan Mikulincer, Ilias Zadik</i>	
Uncovering the Structural Fairness in Graph Contrastive Learning.....	32465
<i>Chuan Shi, Le Song, Ruijia Wang, Xiao Wang</i>	
Estimation of Entropy in Constant Space with Improved Sample Complexity.....	32474
<i>Maryam Aliakbarpour, Andrew McGregor, Jelani Nelson, Erik Waingarten</i>	
Average Sensitivity of Euclidean k-Clustering.....	32487
<i>Shinji Ito, Yuichi Yoshida</i>	

Multi-Block-Single-Probe Variance Reduced Estimator for Coupled Compositional Optimization.....	32499
<i>Wei Jiang, Gang Li, Yibo Wang, Tianbao Yang, Lijun Zhang</i>	
Differentially Private Online-To-Batch for Smooth Losses	32512
<i>Ashok Cutkosky, Hoang Tran, Qinzi Zhang</i>	
Weak-Shot Semantic Segmentation Via Dual Similarity Transfer	32525
<i>Junjie Chen, Li Niu, Chen Qian, Jianlou Si, Liqing Zhang, Siyuan Zhou</i>	
A Multi-Task Benchmark for Korean Legal Language Understanding and Judgement Prediction.....	32537
<i>Kyoungyeon Cho, Wonseok Hwang, Dongjun Lee, Hanuhl Lee, Minjoon Seo</i>	
Synergy-Of-Experts: Collaborate to Improve Adversarial Robustness	32552
<i>Sen Cui, Bo Han, Jian Liang, Masashi Sugiyama, Changshui Zhang, Jingfeng Zhang</i>	
Improved Fine-Tuning by Better Leveraging Pre-Training Data	32568
<i>Antoni Chan, Xiangyang Ji, Rong Jin, Hao Li, Ziquan Liu, Qi Qian, Yi Xu, Yuanhong Xu</i>	
Transferring Fairness Under Distribution Shifts Via Fair Consistency Regularization	32582
<i>Bang An, Zora Che, Mucong Ding, Furong Huang</i>	
OpenOOD: Benchmarking Generalized Out-Of-Distribution Detection.....	32598
<i>Guangyao Chen, Kunyuan Ding, Xuefeng Du, Dan Hendrycks, Bo Li, Yixuan Li, Ziwei Liu, Wenxuan Peng, Yiyao Sun, Haoqi Wang, Pengyun Wang, Jingkang Yang, Wayne Zhang, Kaiyang Zhou, Zitang Zhou, Dejian Zou</i>	
Two-Layer Neural Network on Infinite Dimensional Data: Global Optimization Guarantee in the Mean-Field Regime	32612
<i>Naoki Nishikawa, Atsushi Nitanda, Taiji Suzuki, Denny Wu</i>	
KSD Aggregated Goodness-Of-Fit Test	32624
<i>Arthur Gretton, Benjamin Guedj, Antonin Schrab</i>	
Efficient Risk-Averse Reinforcement Learning.....	32639
<i>Yinlam Chow, Mohammad Ghavamzadeh, Ido Greenberg, Shie Mannor</i>	
D ² NeRF: Self-Supervised Decoupling of Dynamic and Static Objects from a Monocular Video	32653
<i>Forrester Cole, Cengiz Oztireli, Andrea Tagliasacchi, Tianhao Wu, Fangcheng Zhong</i>	
On the Effective Number of Linear Regions in Shallow Univariate ReLU Networks: Convergence Guarantees and Implicit Bias.....	32667
<i>Jason Lee, Itay Safran, Gal Vardi</i>	
Robust φ -Divergence MDPs.....	32680
<i>Chin Pang Ho, Marek Petrik, Wolfram Wiesemann</i>	
Object Representations as Fixed Points: Training Iterative Refinement Algorithms with Implicit Differentiation	32694
<i>Michael Chang, Tom Griffiths, Sergey Levine</i>	
Benefits of Additive Noise in Composing Classes with Bounded Capacity.....	32709
<i>Hassan Ashtiani, Alireza Fathollah Pour</i>	
Multi-Objective Deep Learning with Adaptive Reference Vectors	32723
<i>Weiyu Chen, James Kwok</i>	

Towards Hard-Pose Virtual Try-On Via 3D-Aware Global Correspondence Learning	32736
<i>Qingling Cai, Zaiyu Huang, Michael Kampffmeyer, Hanhui Li, Xiaodan Liang, Zhenyu Xie</i>	
Pushing the Limits of Fairness Impossibility: Who's the Fairest of Them All?.....	32749
<i>Kinjal Basu, Brian Hsu, Rahul Mazumder, Preetam Nandy</i>	
Recall Distortion in Neural Network Pruning and the Undecayed Pruning Algorithm	32762
<i>Mikey Ferguson, Aidan Good, Jiaqi Lin, Thiago Serra, Hannah Sieg, Jerzy Wiecek, Xin Yu, Shandian Zhe</i>	
Generating Multivariate Time Series with COMmon Source COordInated GAN (COSCI-GAN)	32777
<i>Raymond Ng, Jean-Francois Rajotte, Ali Seyfi</i>	
ConfounderGAN: Protecting Image Data Privacy with Causal Confounder.....	32789
<i>Kelu Jiang, Kun Kuang, Furui Liu, Qi Tian, Zhihua Wang, Fei Wu</i>	

VOLUME 43

Global Convergence of Direct Policy Search for State-Feedback H_∞ Robust Control: A Revisit of Nonsmooth Synthesis with Goldstein Subdifferential.....	32801
<i>Xingang Guo, Bin Hu</i>	
In What Ways Are Deep Neural Networks Invariant and How Should We Measure This?.....	32816
<i>Tim Doster, Tegan Emerson, Grayson Jorgenson, Henry Kvinge, Jesse Lew, Scott Vasquez</i>	
Optimal Weak to Strong Learning	32830
<i>Kasper Green Larsen, Martin Ritzert</i>	
A Neural Pre-Conditioning Active Learning Algorithm to Reduce Label Complexity	32842
<i>Soomin Jeon, Kyu-Hwan Jung, Seo Taek Kong, Hong-Seok Lee, Jaewon Lee, Dongbin Na</i>	
AnoShift: A Distribution Shift Benchmark for Unsupervised Anomaly Detection	32854
<i>Florin Brad, Elena Burceanu, Marius Dragoi, Emanuela Haller, Andrei Manolache</i>	
Grounding Aleatoric Uncertainty for Unsupervised Environment Design	32868
<i>Michael Dennis, Jakob Foerster, Edward Grefenstette, Minqi Jiang, Heinrich Küttler, Andrei Lupu, Jack Parker-Holder, Tim Rocktäschel</i>	
PROSPECT: Labeled Tandem Mass Spectrometry Dataset for Machine Learning in Proteomics	32882
<i>Wassim Gabriel, Victor-George Giurcoiu, Omar Shouman, Vitor Sternlicht, Mathias Wilhelm</i>	
VLMO: Unified Vision-Language Pre-Training with Mixture-Of-Modality-Experts.....	32897
<i>Kriti Aggarwal, Hangbo Bao, Li Dong, Qiang Liu, Owais Khan Mohammed, Songhao Piao, Subhojit Som, Wenhui Wang, Furu Wei</i>	
Hub-Pathway: Transfer Learning from a Hub of Pre-Trained Models	32913
<i>Zhangjie Cao, Mingsheng Long, Yang Shu, Jianmin Wang, Ziyang Zhang</i>	
Towards Better Evaluation for Dynamic Link Prediction	32928
<i>Shenyang Huang, Kellin Pelrine, Farimah Poursafaei, Reihaneh Rabbany</i>	
Coarse-To-Fine Vision-Language Pre-Training with Fusion in the Backbone	32942
<i>Zi-Yi Dou, Zhe Gan, Jianfeng Gao, Aishwarya Kamath, Yann Lecun, Linjie Li, Ce Liu, Zicheng Liu, Nanyun Peng, Jianfeng Wang, Lijuan Wang, Pengchuan Zhang</i>	

On the Discrimination Risk of Mean Aggregation Feature Imputation in Graphs	32957
<i>Kai-Wei Chang, Arjun Subramonian, Yizhou Sun</i>	
VRL3: A Data-Driven Framework for Visual Deep Reinforcement Learning	32974
<i>Dongsheng Li, Xufang Luo, Keith Ross, Che Wang</i>	
Efficient and Effective Optimal Transport-Based Biclustering	32989
<i>Chakib Fettaf, Lazhar Labiod, Mohamed Nadif</i>	
Unifying and Boosting Gradient-Based Training-Free Neural Architecture Search.....	33001
<i>Zhongxiang Dai, Bryan Kian Hsiang Low, Yao Shu, Zhaoxuan Wu</i>	
Hyperbolic Embedding Inference for Structured Multi-Label Prediction	33016
<i>Michael Cochez, Mojtaba Nayyeri, Steffen Staab, Bo Xiong</i>	
Approximate Value Equivalence	33029
<i>Andre Barreto, Christopher Grimm, Satinder Singh</i>	
The Power and Limitation of Pretraining-Finetuning for Linear Regression Under Covariate Shift.....	33041
<i>Vladimir Braverman, Quanquan Gu, Sham Kakade, Jingfeng Wu, Difan Zou</i>	
Rank Diminishing in Deep Neural Networks.....	33054
<i>Ruili Feng, Yukun Huang, Michael Jordan, Zheng-Jun Zha, Deli Zhao, Kecheng Zheng</i>	
Single-Pass Streaming Lower Bounds for Multi-Armed Bandits Exploration with Instance-Sensitive Sample Complexity.....	33066
<i>Sepehr Assadi, Chen Wang</i>	
Optimistic Tree Searches for Combinatorial Black-Box Optimization	33080
<i>Haitham Bou Ammar, Antoine Grosnit, Cedric Malherbe, Rasul Tutunov, Jun Wang</i>	
Robust Feature-Level Adversaries Are Interpretability Tools	33093
<i>Stephen Casper, Dylan Hadfield-Menell, Gabriel Kreiman, Max Nadeau</i>	
A Unified Convergence Theorem for Stochastic Optimization Methods	33107
<i>Xiao Li, Andre Milzarek</i>	
New Definitions and Evaluations for Saliency Methods: Staying Intrinsic, Complete and Sound.....	33120
<i>Sanjeev Arora, Arushi Gupta, Kaifeng Lyu, Nikunj Saunshi, Dingli Yu</i>	
Exponential Separations in Symmetric Neural Networks	33134
<i>Joan Bruna, Aaron Zweig</i>	
Randomized Message-Interception Smoothing: Gray-Box Certificates for Graph Neural Networks.....	33146
<i>Aleksandar Bojchevski, Simon Geisler, Stephan Günnemann, Yan Scholten, Jan Schuchardt</i>	
Oscillatory Tracking of Continuous Attractor Neural Networks Account for Phase Precession and Procession of Hippocampal Place Cells	33159
<i>Tianhao Chu, Tiejun Huang, Zilong Ji, Yuanyuan Mi, Si Wu, Wenhao Zhang, Junfeng Zuo</i>	
Finding and Listing Front-Door Adjustment Sets.....	33173
<i>Elias Bareinboim, Hyunchai Jeong, Jin Tian</i>	
MorphTE: Injecting Morphology in Tensorized Embeddings.....	33186
<i>Guobing Gan, Sunzhu Li, Xiuqing Lu, Benyou Wang, Peng Zhang</i>	

Active Learning with Safety Constraints.....	33201
<i>Romain Camilleri, Lalit Jain, Kevin Jamieson, Jamie Morgenstern, Andrew Wagenmaker</i>	
Is One Annotation Enough? - A Data-Centric Image Classification Benchmark for Noisy and Ambiguous Label Estimation.....	33215
<i>Sabine Dippel, Vasco Grossmann, Rainer Kiko, Reinhard Koch, Mariusz Oszust, Matti Pastell, Lars Schmarje, Jenny Stracke, Anna Valros, Nina Volkmann, Claudius Zelenka</i>	
Sobolev Acceleration and Statistical Optimality for Learning Elliptic Equations Via Gradient Descent.....	33233
<i>Jose Blanchet, Yiping Lu, Lexing Ying</i>	
Block-Recurrent Transformers	33248
<i>Ethan Dyer, Delesley Hutchins, Behnam Neyshabur, Imanol Schlag, Yuhuai Wu</i>	
Learning Two-Player Markov Games: Neural Function Approximation and Correlated Equilibrium.....	33262
<i>Quanquan Gu, Michael Jordan, Chris Junchi Li, Dongruo Zhou</i>	
Dense Interspecies Face Embedding	33275
<i>Subin Jeon, Seon Joo Kim, Seonghyeon Nam, Sejong Yang</i>	
Efficient Scheduling of Data Augmentation for Deep Reinforcement Learning.....	33289
<i>Byungchan Ko, Jungseul Ok</i>	
Mix and Reason: Reasoning Over Semantic Topology with Data Mixing for Domain Generalization	33302
<i>Chaoqi Chen, Yue Huang, Feng Liu, Luyao Tang, Yizhou Yu, Gangming Zhao</i>	
Personalized Online Federated Learning with Multiple Kernels.....	33316
<i>Pouya M. Ghari, Yanning Shen</i>	
Point Transformer V2: Grouped Vector Attention and Partition-Based Pooling	33330
<i>Li Jiang, Yixing Lao, Xihui Liu, Xiaoyang Wu, Hengshuang Zhao</i>	
Learning Concept Credible Models for Mitigating Shortcuts	33343
<i>Sarah Jabbour, Maggie Makar, Michael Sjoding, Jiaxuan Wang, Jenna Wiens</i>	
Neural Approximation of Graph Topological Features.....	33357
<i>Chao Chen, Liangcai Gao, Tengfei Ma, Zhi Tang, Yusu Wang, Zuoyu Yan</i>	
MOVE: Unsupervised Movable Object Segmentation and Detection.....	33371
<i>Adam Bielski, Paolo Favaro</i>	
FreGAN: Exploiting Frequency Components for Training GANs Under Limited Data	33387
<i>Ziqiu Chi, Zhe Wang, Mengping Yang, Yanbing Zhang</i>	
LBD: Decouple Relevance and Observation for Individual-Level Unbiased Learning to Rank.....	33400
<i>Mouxiang Chen, Chenghao Liu, Zemin Liu, Jianling Sun</i>	
MetricFormer: A Unified Perspective of Correlation Exploring in Similarity Learning	33414
<i>Cheng Deng, Heng Huang, Jiexi Yan, Erkun Yang</i>	
Toward a Realistic Model of Speech Processing in the Brain with Self-Supervised Learning	33428
<i>Yves Boubenec, Charlotte Caucheteux, Ewan Dunbar, Alexandre Gramfort, Jean-Remi King, Juliette Millet, Pierre Orhan, Christophe Pallier</i>	
Distributed Inverse Constrained Reinforcement Learning for Multi-Agent Systems.....	33444
<i>Shicheng Liu, Minghui Zhu</i>	

Collaborative Decision Making Using Action Suggestions.....	33457
<i>Dylan Asmar, Mykel J Kochenderfer</i>	
Near-Optimal Regret for Adversarial MDP with Delayed Bandit Feedback.....	33469
<i>Tiancheng Jin, Tal Lenczewski, Haipeng Luo, Yishay Mansour, Aviv Rosenberg</i>	
Provably Sample-Efficient RL with Side Information About Latent Dynamics.....	33482
<i>Miro Dudik, Yao Liu, Dipendra Misra, Robert Schapire</i>	
Optimal Gradient Sliding and Its Application to Optimal Distributed Optimization Under Similarity.....	33494
<i>Aleksandr Beznosikov, Ekaterina Borodich, Alexander Gasnikov, Dmitry Kovalev, Gesualdo Scutari</i>	
Universally Expressive Communication in Multi-Agent Reinforcement Learning.....	33508
<i>Thomas D Barrett, Matthew Morris, Arnu Pretorius</i>	
Fast Stochastic Composite Minimization and an Accelerated Frank-Wolfe Algorithm Under Parallelization.....	33523
<i>Francis Bach, Quentin Berthet, Benjamin Dubois-Taine, Adrien Taylor</i>	
Multi-Granularity Cross-Modal Alignment for Generalized Medical Visual Representation Learning	33536
<i>Varut Vardhanabhuti, Fuying Wang, Shujun Wang, Lequan Yu, Yuyin Zhou</i>	
Learning Robust Rule Representations for Abstract Reasoning Via Internal Inferences.....	33550
<i>Xianggen Liu, Site Mo, Sen Song, Likai Tang, Wenbo Zhang</i>	
Turning the Tables: Biased, Imbalanced, Dynamic Tabular Datasets for ML Evaluation	33563
<i>Duarte Alves, Pedro Bizarro, André Cruz, João Gama, Sérgio Jesus, José Pombal, Rita Ribeiro, Pedro Saleiro</i>	
DropCov: A Simple Yet Effective Method for Improving Deep Architectures	33576
<i>Mingze Gao, Qinghua Hu, Peihua Li, Qilong Wang, Jiangtao Xie, Zhaolin Zhang</i>	

VOLUME 44

A Unifying Framework for Online Optimization with Long-Term Constraints	33589
<i>Matteo Castiglioni, Andrea Celli, Nicola Gatti, Alberto Marchesi, Giulia Romano</i>	
The Least-Control Principle for Local Learning at Equilibrium	33603
<i>Seijin Kobayashi, Alexander Meulemans, João Sacramento, Johannes Von Oswald, Nicolas Zucchet</i>	
Optimizing Relevance Maps of Vision Transformers Improves Robustness.....	33618
<i>Hila Chefer, Idan Schwartz, Lior Wolf</i>	
Stability and Generalization of Kernel Clustering: From Single Kernel to Multiple Kernel.....	33633
<i>Jun-Jie Huang, Weixuan Liang, Yong Liu, Jiyuan Liu, Xinwang Liu, Siwei Wang, Yi Zhang, Sihang Zhou, En Zhu</i>	
Deep Ensembles Work, but Are They Necessary?.....	33646
<i>Taiga Abe, Estefany Kelly Buchanan, John Cunningham, Geoff Pleiss, Richard Zemel</i>	
VICE: Variational Interpretable Concept Embeddings.....	33661
<i>Martin N Hebart, Patrick McClure, Lukas Muttenthaler, Francisco Pereira, Robert Vandermeulen, Charles Zheng</i>	

Random Normalization Aggregation for Adversarial Defense.....	33676
<i>Xinghao Chen, Minjing Dong, Yunhe Wang, Chang Xu</i>	
Distributionally Robust Optimization with Data Geometry.....	33689
<i>Peng Cui, Bo Li, Jiashuo Liu, Jiayun Wu</i>	
Near-Optimal Correlation Clustering with Privacy.....	33702
<i>Vincent Cohen-Addad, Chenglin Fan, Silvio Lattanzi, Slobodan Mitrovic, Ashkan Norouzi-Fard, Nikos Parotsidis, Jakub Tarnawski</i>	
Knowledge Distillation from a Stronger Teacher.....	33716
<i>Tao Huang, Chen Qian, Fei Wang, Chang Xu, Shan You</i>	
Optimal Transport of Classifiers to Fairness.....	33728
<i>Maarten Buyl, Tijl De Bie</i>	
Rethinking the Compositionality of Point Clouds Through Regularization in the Hyperbolic Space.....	33741
<i>Enrico Magli, Antonio Montanaro, Diego Valsesia</i>	
ReCo: Retrieve and Co-Segment for Zero-Shot Transfer.....	33754
<i>Samuel Albanie, Gyungin Shin, Weidi Xie</i>	
Monocular Dynamic View Synthesis: A Reality Check.....	33768
<i>Hang Gao, Angjoo Kanazawa, Ruilong Li, Bryan Russell, Shubham Tulsiani</i>	
Bridging the Gap Between Object and Image-Level Representations for Open-Vocabulary Detection.....	33781
<i>Hanoona Rasheed, Salman Khan, Muhammad Uzair Khattak, Muhammad Maaz, Fahad Shahbaz Khan</i>	
A Simple and Optimal Policy Design for Online Learning with Safety Against Heavy-Tailed Risk.....	33795
<i>David Simchi-Levi, Zeyu Zheng, Feng Zhu</i>	
A Boosting Approach to Reinforcement Learning.....	33806
<i>Nataly Brukhim, Elad Hazan, Karan Singh</i>	
Relaxing Equivariance Constraints with Non-Stationary Continuous Filters.....	33818
<i>David W. Romero, Tycho Van Der Ouderaa, Mark Van Der Wilk</i>	
The Franz-Parisi Criterion and Computational Trade-Offs in High Dimensional Statistics.....	33831
<i>Afonso S Bandeira, Ahmed El Alaoui, Samuel Hopkins, Tselil Schramm, Alexander S Wein, Ilias Zadik</i>	
Robust Neural Posterior Estimation and Statistical Model Criticism.....	33845
<i>Mark Beaumont, Patrick Cannon, Matteo Fasiolo, Sebastian Schmon, Daniel Ward</i>	
Why Do We Need Large Batchsizes in Contrastive Learning? A Gradient-Bias Perspective.....	33860
<i>Changyou Chen, Liqun Chen, Yiran Chen, Trishul Chilimbi, Jiali Duan, Son Tran, Yi Xu, Belinda Zeng, Jianyi Zhang</i>	
Randomized Channel Shuffling: Minimal-Overhead Backdoor Attack Detection Without Clean Datasets.....	33876
<i>Ruisi Cai, Tianlong Chen, Xiaohan Chen, Zhangyang Wang, Zhenyu Zhang</i>	
Structural Analysis of Branch-And-Cut and the Learnability of Gomory Mixed Integer Cuts.....	33890
<i>Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, Ellen Vitercik</i>	

Finite-Time Last-Iterate Convergence for Learning in Multi-Player Games.....	33904
<i>Yang Cai, Argyris Oikonomou, Weiqiang Zheng</i>	
When Are Local Queries Useful for Robust Learning?.....	33920
<i>Pascale Gourdeau, Varun Kanade, Marta Kwiatkowska, James Worrell</i>	
SPoVT: Semantic-Prototype Variational Transformer for Dense Point Cloud Semantic Completion.....	33934
<i>Hao-Yu Hsu, Sheng Yu Huang, Frank Wang</i>	
Improving Intrinsic Exploration with Language Abstractions.....	33947
<i>Noah Goodman, Edward Grefenstette, Mingqi Jiang, Jesse Mu, Roberta Raileanu, Tim Rocktäschel, Victor Zhong</i>	
The Trade-Offs of Model Size in Large Recommendation Models : 100GB to 10MB Criteo-Tb DLRM Model.....	33961
<i>Aditya Desai, Anshumali Shrivastava</i>	
Near-Optimal Goal-Oriented Reinforcement Learning in Non-Stationary Environments	33973
<i>Liyu Chen, Haipeng Luo</i>	
A Unified Framework for Deep Symbolic Regression.....	33985
<i>Ignacio Aravena, Ruben Glatt, Mikel Landajuela, Chak Shing Lee, Garrett Mulcahy, Terrell Mundhenk, Brenden K Petersen, Claudio P Santiago, Jiachen Yang</i>	
VoxGRAF: Fast 3D-Aware Image Synthesis with Sparse Voxel Grids	33999
<i>Andreas Geiger, Yiyi Liao, Michael Niemeyer, Axel Sauer, Katja Schwarz</i>	
Sparse Hypergraph Community Detection Thresholds in Stochastic Block Model.....	34012
<i>Syed Zulqarnain Gilani, David Suter, Giang Truong, Erchuan Zhang</i>	
Understanding Non-Linearity in Graph Neural Networks from the Bayesian-Inference Perspective.....	34024
<i>Austin Benson, Junteng Jia, Pan Li, Rongzhe Wei, Haoteng Yin</i>	
Explainable Reinforcement Learning Via Model Transforms	34039
<i>Mira Finkelstein, Sarah Keren, Yoav Kolubus, Nitsan Levy, Lucy Liu, David Parkes, Jeffrey S Rosenschein</i>	
Real-Valued Backpropagation is Unsuitable for Complex-Valued Neural Networks.....	34052
<i>Yuan Jiang, Zhi-Hao Tan, Yi Xie, Zhi-Hua Zhou</i>	
Learning Dynamics of Deep Linear Networks with Multiple Pathways	34064
<i>Michael Buice, Eric Shea-Brown, Jianghong Shi</i>	
Self-Supervised Aggregation of Diverse Experts for Test-Agnostic Long-Tailed Recognition	34077
<i>Jiashi Feng, Lanqing Hong, Bryan Hooi, Yifan Zhang</i>	
Self-Supervised Learning Via Maximum Entropy Coding.....	34091
<i>Ya-Li Li, Xin Liu, Shengjin Wang, Zhongdao Wang</i>	
A Practical, Progressively-Expressive GNN	34106
<i>Leman Akoglu, Neil Shah, Lingxiao Zhao</i>	
On Learning Fairness and Accuracy on Multiple Subgroups	34121
<i>Tal Arbel, Qi Chen, Christian Gagné, Jiaqi Li, Charles Ling, Changjian Shui, Boyu Wang, Gezheng Xu</i>	

Adv-Attribute: Inconspicuous and Transferable Adversarial Attack on Face Recognition	34136
<i>Shouhong Ding, Shuai Jia, Chao Ma, Chunhua Shen, Xiaokang Yang, Taiping Yao, Bangjie Yin</i>	
Understanding Aesthetics with Language: A Photo Critique Dataset for Aesthetic Assessment.....	34148
<i>Luigi Celona, Clara Fernandez Labrador, Daniel Vera Nieto</i>	
Structuring Representations Using Group Invariants	34162
<i>Arnab Kumar Mondal, Siamak Ravanbakhsh, Mehran Shakerinava</i>	
Black-Box Coreset Variational Inference	34175
<i>Theofanis Karaletsos, Dionysis Manousakas, Hippolyt Ritter</i>	
Distilling Representations from GAN Generator Via Squeeze and Span	34188
<i>Hakan Bilen, Xiaotian Cheng, Xiangyang Ji, Chang Liu, Yu Yang</i>	
A Quantitative Geometric Approach to Neural-Network Smoothness	34201
<i>Somesh Jha, Gautam Prakriya, Zi Wang</i>	
AutoMTL: A Programming Framework for Automating Efficient Multi-Task Learning.....	34216
<i>Hui Guan, Xiao Liu, Lijun Zhang</i>	
Log-Concave and Multivariate Canonical Noise Distributions for Differential Privacy.....	34229
<i>Jordan Awan, Jinshuo Dong</i>	
On the Generalization of Learning Algorithms that Do Not Converge	34241
<i>Nisha Chandramoorthy, Khashayar Gatmiry, Stefanie Jegelka, Andreas Loukas</i>	
PALMER: Perception - Action Loop with Memory for Long-Horizon Planning.....	34258
<i>Onur Beker, Mohammad Mohammadi, Amir Zamir</i>	
Leveraging Factored Action Spaces for Efficient Offline Reinforcement Learning in Healthcare	34272
<i>Finale Doshi-Velez, Maggie Makar, Michael Sjoding, Shengpu Tang, Jenna Wiens</i>	
Visual Correspondence-Based Explanations Improve AI Robustness and human-AI Team Accuracy	34287
<i>Anh Nguyen, Giang Nguyen, Mohammad Reza Taesiri</i>	
Invariance-Aware Randomized Smoothing Certificates.....	34302
<i>Stephan Günnemann, Jan Schuchardt</i>	
On Image Segmentation with Noisy Labels: Characterization and Volume Properties of the Optimal Solutions to Accuracy and Dice.....	34321
<i>Henrik Hult, Fredrik Löfman, Marcus Nordstrom, Jonas Söderberg</i>	
Blessing of Depth in Linear Regression: Deeper Models Have Flatter Landscape Around the True Solution	34334
<i>Salar Fattahi, Jianhao Ma</i>	
Fairness Reprogramming.....	34347
<i>Shiyu Chang, Wenqi Fan, Qing Li, Sijia Liu, Guanhua Zhang, Yang Zhang, Yihua Zhang</i>	
WeightedSHAP: Analyzing and Improving Shapley Based Feature Attributions	34363
<i>Yongchan Kwon, James Zou</i>	

VOLUME 45

Temporal Effective Batch Normalization in Spiking Neural Networks	34377
<i>Shiyan Chen, Jianhao Ding, Chaoteng Duan, Tiejun Huang, Zhaofei Yu</i>	
Remember the Past: Distilling Datasets into Addressable Memories for Neural Networks	34391
<i>Zhiwei Deng, Olga Russakovsky</i>	
Robustness Analysis of Video-Language Models Against Visual and Language Perturbations.....	34405
<i>Madeline C. Schiappa, Hamid Palangi, Yogesh Rawat, Vibhav Vineet, Shruti Vyas</i>	
Hyperbolic Feature Augmentation Via Distribution Estimation and Infinite Sampling on Manifolds.....	34421
<i>Zhi Gao, Mehrtash Harandi, Yunde Jia, Yuwei Wu</i>	
PatchComplete: Learning Multi-Resolution Patch Priors for 3D Shape Completion on Unseen Categories.....	34436
<i>Angela Dai, Yinyu Nie, Yuchen Rao</i>	
Q-ViT: Accurate and Fully Quantized Low-Bit Vision Transformer.....	34451
<i>Xianbin Cao, Peng Gao, Guodong Guo, Yanjing Li, Sheng Xu, Baochang Zhang</i>	
Enhancing Safe Exploration Using Safety State Augmentation	34464
<i>Haitham Bou Ammar, Alexander Cowen-Rivers, Aivar Sootla, Jun Wang</i>	
Contrastive Intrinsic Control for Unsupervised Reinforcement Learning.....	34478
<i>Pieter Abbeel, Michael Laskin, Hao Liu, Xue Bin Peng, Aravind Rajeswaran, Denis Yarats</i>	
Distributed Online Convex Optimization with Compressed Communication.....	34492
<i>Yiguang Hong, Guodong Shi, Zhipeng Tu, Lei Wang, Xi Wang, Deming Yuan</i>	
Local Latent Space Bayesian Optimization Over Structured Inputs	34505
<i>John Bradshaw, Jacob Gardner, Haydn Jones, Matt Kusner, Natalie Maus, Juston Moore</i>	
Graph Neural Network Bandits	34519
<i>Ilija Bogunovic, Parnian Kassraie, Andreas Krause</i>	
Concrete Score Matching: Generalized Score Matching for Discrete Data	34532
<i>Kristy Choi, Stefano Ermon, Chenlin Meng, Jiaming Song</i>	
Regularized Gradient Descent Ascent for Two-Player Zero-Sum Markov Games.....	34546
<i>Thinh Doan, Justin Romberg, Sihan Zeng</i>	
Chefs' Random Tables: Non-Trigonometric Random Features	34559
<i>Krzysztof M Choromanski, Kumar Avinava Dubey, Valerii Likhoshesterov, Frederick Liu, Tamas Sarlos, Adrian Weller</i>	
CS-Shapley: Class-Wise Shapley Values for Data Valuation in Classification	34574
<i>Yangfeng Ji, Stephanie Schoch, Haifeng Xu</i>	
Factuality Enhanced Language Models for Open-Ended Text Generation.....	34586
<i>Bryan Catanzaro, Pascale N Fung, Nayeon Lee, Mostofa Patwary, Wei Ping, Mohammad Shoeybi, Peng Xu</i>	
On the Representation Collapse of Sparse Mixture of Experts	34600
<i>Payal Bajaj, Zewen Chi, Damai Dai, Li Dong, Heyan Huang, Shaohan Huang, Shuming Ma, Xian-Ling Mao, Barun Patra, Saksham Singhal, Xia Song, Furu Wei</i>	

Nearly Optimal Algorithms for Linear Contextual Bandits with Adversarial Corruptions	34614
<i>Quanquan Gu, Jiafan He, Tong Zhang, Dongruo Zhou</i>	
Implicit Bias of Gradient Descent on Reparametrized Models: On Equivalence to Mirror Descent	34626
<i>Sanjeev Arora, Jason Lee, Zhiyuan Li, Tianhao Wang</i>	
HYPRO: A Hybridly Normalized Probabilistic Model for Long-Horizon Prediction of Event Sequences	34641
<i>Hongyuan Mei, Xiaoming Shi, Siqiao Xue, James Zhang</i>	
Towards Understanding Grokking: An Effective Theory of Representation Learning	34651
<i>Ouail Kitouni, Ziming Liu, Eric Michaud, Niklas S Nolte, Max Tegmark, Mike Williams</i>	
Asymptotic Behaviors of Projected Stochastic Approximation: A Jump Diffusion Perspective	34664
<i>Yuze Han, Xiang Li, Jiadong Liang, Zhihua Zhang</i>	
Towards Practical Few-Shot Query Sets: Transductive Minimum Description Length Inference	34677
<i>Ismail Ayed, Malik Boudiaf, Emilie Chouzenoux, Ségolène Martin, Jean-Christophe Pesquet</i>	
Understanding the Generalization Benefit of Normalization Layers: Sharpness Reduction	34689
<i>Sanjeev Arora, Zhiyuan Li, Kaifeng Lyu</i>	
Learning Manifold Dimensions with Conditional Variational Autoencoders	34709
<i>Tong He, Yixuan Qiu, David P Wipf, Yijia Zheng</i>	
Multi-Modal Grouping Network for Weakly-Supervised Audio-Visual Video Parsing	34722
<i>Shentong Mo, Yapeng Tian</i>	
Optimal-Er Auctions Through Attention	34734
<i>Ksenia Balabaeva, Igor Filippov, Dmitry Ivanov, Iskander Safiulin</i>	
Bootstrapped Transformer for Offline Reinforcement Learning	34748
<i>Dongsheng Li, Xufang Luo, Kan Ren, Kerong Wang, Weinan Zhang, Hanye Zhao</i>	
How to Talk So AI Will Learn: Instructions, Descriptions, and Autonomy	34762
<i>Tom Griffiths, Dylan Hadfield-Menell, Robert Hawkins, Mark Ho, Theodore Sumers</i>	
Learning on the Edge: Online Learning with Stochastic Feedback Graphs	34776
<i>Nicolò Cesa-Bianchi, Emmanuel Esposito, Federico Fusco, Dirk Van Der Hoeven</i>	
Globally Gated Deep Linear Networks	34789
<i>Qianyi Li, Haim Sompolinsky</i>	
Markov Chain Score Ascent: A Unifying Framework of Variational Inference with Markovian Gradients	34802
<i>Adji Bousso Dieng, Jacob Gardner, Hongseok Kim, Kyurae Kim, Jisu Oh</i>	
Cache-Augmented Inbatch Importance Resampling for Training Recommender Retriever	34817
<i>Enhong Chen, Jin Chen, Yucheng Li, Defu Lian, Baoyun Wang, Kai Zheng</i>	
Doubly Robust Counterfactual Classification	34831
<i>Edward Kennedy, Kwangho Kim, Jose Zubizarreta</i>	
Rethinking Value Function Learning for Generalization in Reinforcement Learning	34846
<i>Junyeong Lee, Seungyong Moon, Hyun Oh Song</i>	

Improving Certified Robustness Via Statistical Learning with Logical Reasoning.....	34859
<i>Heng Guo, Bojan Karlaš, Bo Li, Linyi Li, Ji Liu, Hengzhi Pei, Boxin Wang, Zhuolin Yang, Ce Zhang, Jiawei Zhang, Zhikuan Zhao</i>	
Transformer-Based Working Memory for Multiagent Reinforcement Learning with Action Parsing	34874
<i>Guangyong Chen, Jianye Hao, Xiaotian Hao, Pheng-Ann Heng, Weixun Wang, Yaodong Yang</i>	
Recruitment Strategies that Take a Chance	34887
<i>Gregory Kehne, Ariel Procaccia, Jingyan Wang</i>	
Fully Convolutional One-Stage 3D Object Detection on LiDAR Range Images.....	34899
<i>Xiangxiang Chu, Chunhua Shen, Zhi Tian, Xiaoming Wang, Xiaolin Wei</i>	
Understanding Robust Learning Through the Lens of Representation Similarities	34912
<i>Arjun Nitin Bhagoji, Christian Cianfarani, Prateek Mittal, Vikash Sehwal, Ben Zhao, Heather Zheng</i>	
Chartalist: Labeled Graph Datasets for UTXO and Account-Based Blockchains	34926
<i>Cuneyt G Akcora, Yulia Gel, Murat Kantarcioglu, Kiarash Shamsi, Friedhelm Victor</i>	
Multi-Instance Causal Representation Learning for Instance Label Prediction and Out-Of-Distribution Generalization	34940
<i>Hanwen Deng, Min-Ling Zhang, Weijia Zhang, Xuanhui Zhang</i>	
PerfectDou: Dominating DouDizhu with Perfect Information Distillation	34954
<i>Fei Fang, Weijun Hong, Yue Lin, Minghuan Liu, Guan Yang, Guangjun Zeng, Weinan Zhang</i>	
On the Double Descent of Random Features Models Trained with SGD	34966
<i>Volkan Cevher, Fanghui Liu, Johan Suykens</i>	
SALSA: Attacking Lattice Cryptography with Transformers	34981
<i>Francois Charton, Mingjie Chen, Kristin E. Lauter, Emily Wenger</i>	
Variable-Rate Hierarchical CPC Leads to Acoustic Unit Discovery in Speech.....	34995
<i>Jan Chorowski, Santiago Cuervo, Adrian Lancucki, Ricard Marxer, Pawel Rychlikowski</i>	
Learning to Attack Federated Learning: A Model-Based Reinforcement Learning Attack Framework.....	35007
<i>Henger Li, Xiaolin Sun, Zizhan Zheng</i>	
Learning Neural Set Functions Under the Optimal Subset Oracle	35021
<i>Yatao Bian, Yingzhen Li, Zijing Ou, Qinliang Su, Tingyang Xu, Peilin Zhao</i>	
A Near-Optimal Best-Of-Both-Worlds Algorithm for Online Learning with Feedback Graphs	35035
<i>Nicolò Cesa-Bianchi, Chloé Rouyer, Yevgeny Seldin, Dirk Van Der Hoeven</i>	
Optimal Query Complexities for Dynamic Trace Estimation.....	35049
<i>David Woodruff, Fred Zhang, Richard Zhang</i>	
Learning from Stochastically Revealed Preference.....	35061
<i>John Birge, Xiaocheng Li, Chunlin Sun</i>	
Mutual Information Divergence: A Unified Metric for Multimodal Generative Models	35072
<i>Jin-Hwa Kim, Yunji Kim, Jiyoung Lee, Sang-Woo Lee, Kang Min Yoo</i>	
Delving into Out-Of-Distribution Detection with Vision-Language Representations.....	35087
<i>Ziyang Cai, Jiuxiang Gu, Wei Li, Yixuan Li, Yifei Ming, Yiyun Sun</i>	

OnePose++: Keypoint-Free One-Shot Object Pose Estimation Without CAD Models.....	35103
<i>Hujun Bao, Xingyi He, Di Huang, Jiaming Sun, Yuang Wang, Xiaowei Zhou</i>	
Confidence-Based Reliable Learning Under Dual Noises.....	35116
<i>Peng Cui, Zhijie Deng, Yang Yue, Jun Zhu</i>	
Multitasking Models Are Robust to Structural Failure: A Neural Model for Bilingual Cognitive Reserve	35130
<i>Giannis Daras, Alex Dimakis, Zoi Gkalitsiou, Negin Raoof</i>	
Learning (Very) Simple Generative Models is Hard	35143
<i>Sitan Chen, Jerry Li, Yuanzhi Li</i>	

VOLUME 46

PEER: A Comprehensive and Multi-Task Benchmark for Protein Sequence Understanding	35156
<i>Chang Ma, Runcheng Liu, Jiarui Lu, Jian Tang, Minghao Xu, Yangtian Zhang, Zuobai Zhang, Zhaocheng Zhu</i>	
Bring Your Own Algorithm for Optimal Differentially Private Stochastic Minimax Optimization	35174
<i>Niao He, Sewoong Oh, Kiran Thekumparampil, Liang Zhang</i>	
A Regret-Variance Trade-Off in Online Learning	35188
<i>Nicolò Cesa-Bianchi, Dirk Van Der Hoeven, Nikita Zhivotovskiy</i>	
Alternating Mirror Descent for Constrained Min-Max Games	35201
<i>Georgios Piliouras, Molei Tao, Andre Wibisono</i>	
Deep Counterfactual Estimation with Categorical Background Variables	35213
<i>Edward De Brouwer</i>	
SnAKe: Bayesian Optimization with Pathwise Exploration	35226
<i>Jose Pablo Folch, Robert Lee, Ruth Misener, Behrang Shafei, Calvin Tsay, Mark Van Der Wilk, David Walz, Shiqiang Zhang</i>	
Self-Supervised Learning with an Information Maximization Criterion.....	35240
<i>Sercan Arik, Alper Erdogan, Shadi Hamdan, Serdar Ozsoy, Deniz Yuret</i>	
TwiBot-22: Towards Graph-Based Twitter Bot Detection.....	35254
<i>Yuyang Bai, Zijian Cai, Zilong Chen, Shangbin Feng, Xinshun Feng, Zhenyu Lei, Jundong Li, Yuhan Liu, Minnan Luo, Zihan Ma, Zhaoxuan Tan, Herun Wan, Heng Wang, Hongrui Wang, Ningnan Wang, Yanbo Wang, Shujie Yang, Binchi Zhang, Qingyue Zhang, Wenqian Zhang, Lijing Zheng, Qinghua Zheng</i>	
Listen to Interpret: Post-Hoc Interpretability for Audio Networks with NMF	35270
<i>Florence D'Alché-Buc, Pavlo Mozharovskiy, Jayneel Parekh, Sanjeel Parekh, Gaël Richard</i>	
Model-Based RL with Optimistic Posterior Sampling: Structural Conditions and Sample Complexity	35284
<i>Alekh Agarwal, Tong Zhang</i>	
Deep Architecture Connectivity Matters for Its Convergence: A Fine-Grained Analysis	35298
<i>Wuyang Chen, Xinyu Gong, Boris Hanin, Wei Huang, Zhangyang Wang</i>	
OrdinalCLIP: Learning Rank Prompts for Language-Guided Ordinal Regression	35313
<i>Xiaoke Huang, Wanhua Li, Xiu Li, Jiwen Lu, Yansong Tang, Jie Zhou, Zheng Zhu</i>	

Kernel Memory Networks: A Unifying Framework for Memory Modeling	35326
<i>Johanni Brea, Wulfram Gerstner, Georgios Iatropoulos</i>	
The First Optimal Acceleration of High-Order Methods in Smooth Convex Optimization	35339
<i>Alexander Gasnikov, Dmitry Kovalev</i>	
CCCP is Frank-Wolfe in Disguise	35352
<i>Suvrit Sra, Alp Yurtsever</i>	
Uni[MASK]: Unified Inference in Sequential Decision Problems.....	35365
<i>David Bignell, Micah Carroll, Sam Devlin, Anca Dragan, Raluca Georgescu, Matthew Hausknecht, Katja Hofmann, Jessy Lin, Stephanie Milani, Orr Paradise, Mingfei Sun</i>	
Efficient Active Learning with Abstention	35379
<i>Robert Nowak, Yinglun Zhu</i>	
SInGE: Sparsity Via Integrated Gradients Estimation of Neuron Relevance	35392
<i>Kevin Bailly, Matthieu Cord, Arnaud Dapogny, Edouard Yvinec</i>	
On the Complexity of Adversarial Decision Making	35404
<i>Dylan J Foster, Alexander Rakhlin, Ayush Sekhari, Karthik Sridharan</i>	
MoCapAct: A Multi-Task Dataset for Simulated Humanoid Control.....	35418
<i>Ching-An Cheng, Matthew Hausknecht, Andrey Kolobov, Ricky Loynd, Felipe Vieira Frujeri, Nolan Wagener</i>	
On the Effectiveness of Persistent Homology	35432
<i>Guido Montufar, Nina Otter, Renata Turkes</i>	
An Adaptive Deep RL Method for Non-Stationary Environments with Piecewise Stable Context	35449
<i>Jiayu Chen, Xiaoyu Chen, Wenxue Cheng, Peng Cheng, Tie-Yan Liu, Tao Qin, Yongqiang Xiong, Pushi Zhang, Li Zhao, Yufeng Zheng, Xiangming Zhu</i>	
Expediting Large-Scale Vision Transformer for Dense Prediction Without Fine-Tuning	35462
<i>Henghui Ding, Han Hu, Ding Jia, Weicong Liang, Weihong Lin, Xiao Luo, Yuhui Yuan, Chao Zhang, Zheng Zhang</i>	
Flowification: Everything is a Normalizing Flow	35478
<i>François Fleuret, Tobias Golling, Samuel Klein, Bálint Máté</i>	
A Time-Resolved Theory of Information Encoding in Recurrent Neural Networks	35490
<i>Rainer Engelken, Sven Goedeke</i>	
A Unified Analysis of Mixed Sample Data Augmentation: A Loss Function Perspective.....	35504
<i>Sanghyuk Chun, Chanwoo Park, Sangdoon Yun</i>	
Recursive Reinforcement Learning.....	35519
<i>Ernst Moritz Hahn, Mateo Perez, Sven Schewe, Fabio Somenzi, Ashutosh Trivedi, Dominik Wojtczak</i>	
Non-Linguistic Supervision for Contrastive Learning of Sentence Embeddings.....	35533
<i>Chongyang Gao, Yiren Jian, Soroush Vosoughi</i>	
Towards Video Text Visual Question Answering: Benchmark and Baseline.....	35549
<i>Aobotao Dai, Bingjia Li, Guangze Li, Wanqing Li, Jie Wang, Shijie Xuyang, Xinkun Yu, Zhihang Yu, Lan Zhang, Minyi Zhao, Shuigeng Zhou, Wenjing Zhou</i>	

4D Unsupervised Object Discovery	35563
<i>Yuntao Chen, Yuqi Wang, Zhao-Xiang Zhang</i>	
Faster Linear Algebra for Distance Matrices	35576
<i>Piotr Indyk, Sandeep Silwal</i>	
Contextual Bandits with Knapsacks for a Conversion Model	35590
<i>Zhen Li, Gilles Stoltz</i>	
Contrastive Learning as Goal-Conditioned Reinforcement Learning	35603
<i>Benjamin Eysenbach, Sergey Levine, Russ Salakhutdinov, Tianjun Zhang</i>	
The Pitfalls of Regularization in Off-Policy TD Learning	35621
<i>J. Zico Kolter, Gaurav Manek</i>	
MCMAE: Masked Convolution Meets Masked Autoencoders	35632
<i>Jifeng Dai, Peng Gao, Hongsheng Li, Ziyi Lin, Teli Ma, Yu Qiao</i>	
Robustness to Label Noise Depends on the Shape of the Noise Distribution	35645
<i>Nicolas Hengartner, Michal Kucer, Diane Oyen, Har Simrat Singh</i>	
Fast Neural Kernel Embeddings for General Activations.....	35657
<i>Insu Han, Amin Karbasi, Jaehoon Lee, Roman Novak, Lechao Xiao, Amir Zandieh</i>	
Deep Invariant Networks with Differentiable Augmentation Layers	35672
<i>Alexandre Gramfort, Thomas Moreau, Cédric Rommel</i>	
Factorized-FL: Personalized Federated Learning with Parameter Factorization & Similarity Matching.....	35684
<i>Sung Ju Hwang, Wonyong Jeong</i>	
Reinforcement Learning with a Terminator.....	35696
<i>Gal Chechik, Gal Dalal, Assaf Hallak, Shie Mannor, Nadav Merlis, Uri Shalit, Lior Shani, Guy Tennenholtz</i>	
How Transferable Are Video Representations Based on Synthetic Data?	35710
<i>Rogerio Feris, Souyoung Jin, Leonid Karlinsky, Yo-Whan Kim, Hilde Kuehne, Samarth Mishra, Aude Oliva, Rameswar Panda, Kate Saenko, Venkatesh Saligrama</i>	
Sub-Exponential Time Sum-Of-Squares Lower Bounds for Principal Components Analysis	35724
<i>Aaron Potechin, Goutham Rajendran</i>	
A Multilabel Classification Framework for Approximate Nearest Neighbor Search	35741
<i>Ville Hyvönen, Elias Jääsaari, Teemu Roos</i>	
Learnable Polyphase Sampling for Shift Invariant and Equivariant Convolutional Networks	35755
<i>Minh Do, Teck-Yian Lim, Renan A. Rojas-Gomez, Alex Schwing, Raymond A. Yeh</i>	
Not All Bits Have Equal Value: Heterogeneous Precisions Via Trainable Noise	35769
<i>Yanjing Li, Michael Maire, Pedro Savarese, Xin Yuan</i>	
Tensor Program Optimization with Probabilistic Programs	35783
<i>Tianqi Chen, Siyuan Feng, Bohan Hou, Hongyi Jin, Ruihang Lai, Wuwei Lin, Masahiro Masuda, Junru Shao, Cody Hao Yu, Xiyou Zhou</i>	
Deep Generative Model for Periodic Graphs	35797
<i>Xiaojie Guo, Shiyu Wang, Liang Zhao</i>	

Exploring the Limits of Domain-Adaptive Training for Detoxifying Large-Scale Language Models	35811
<i>Anima Anandkumar, Bryan Catanzaro, Bo Li, Mostofa Patwary, Wei Ping, Mohammad Shoeybi, Boxin Wang, Chaowei Xiao, Peng Xu</i>	
Relational Reasoning Via Set Transformers: Provable Efficiency and Applications to MARL	35825
<i>Boyi Liu, Vincent Tan, Zhaoran Wang, Kaixin Wang, Zhuoran Yang, Fengzhuo Zhang</i>	
Missing Data Imputation and Acquisition with Deep Hierarchical Models and Hamiltonian Monte Carlo	35839
<i>José Miguel Hernández-Lobato, Chao Ma, Ignacio Peis</i>	
Score-Based Generative Models Detect Manifolds.....	35852
<i>Jakiw Pidstrigach</i>	
Distributionally Robust Optimization Via Ball Oracle Acceleration.....	35866
<i>Yair Carmon, Danielle Hausler</i>	
Palm Up: Playing in the Latent Manifold for Unsupervised Pretraining.....	35880
<i>Hao Liu, Volodymyr Mnih, Satinder Singh, Tom Zahavy</i>	
Fair Infinitesimal Jackknife: Mitigating the Influence of Biased Training Data Points Without Refitting.....	35894
<i>Pierre Dognin, Soumya Ghosh, Inkit Padhi, Prasanna Sattigeri, Kush Varshney</i>	
JAWS: Auditing Predictive Uncertainty Under Covariate Shift.....	35907
<i>Anqi Liu, Drew Prinster, Suchi Saria</i>	
DNA: Proximal Policy Optimization with a Dual Network Architecture.....	35921
<i>Matthew Aitchison, Penny Sweetser</i>	
Triangulation Candidates for Bayesian Optimization.....	35933
<i>Robert Gramacy, Annie Sauer, Nathan Wycoff</i>	
VOLUME 47	
Masked Autoencoders as Spatiotemporal Learners.....	35946
<i>Haoqi Fan, Christoph Feichtenhofer, Kaiming He, Yanghao Li</i>	
PyramidCLIP: Hierarchical Feature Alignment for Vision-Language Model Pretraining.....	35959
<i>Yuting Gao, Rongrong Ji, Ke Li, Jinfeng Liu, Chunhua Shen, Zihan Xu, Jun Zhang</i>	
On the Parameterization and Initialization of Diagonal State Space Models	35971
<i>Karan Goel, Albert Gu, Ankit Gupta, Christopher Ré</i>	
Implicit Regularization Or Implicit Conditioning? Exact Risk Trajectories of SGD in High Dimensions	35984
<i>Ben Adlam, Courtney Paquette, Elliot Paquette, Jeffrey Pennington</i>	
Extrapolative Continuous-Time Bayesian Neural Network for Fast Training-Free Test-Time Adaptation	36000
<i>Xiangming Gu, Hengguan Huang, Hongfu Liu, Hao Wang, Ye Wang, Chang Xiao</i>	
Global Convergence and Stability of Stochastic Gradient Descent.....	36014
<i>Vivak Patel, Bowen Tian, Shushu Zhang</i>	

Trap and Replace: Defending Backdoor Attacks by Trapping Them into an Easy-To-Replace Subnetwork.....	36026
<i>Junyuan Hong, Haotao Wang, Zhangyang Wang, Aston Zhang, Jiayu Zhou</i>	
Reinforcement Learning with Logarithmic Regret and Policy Switches.....	36040
<i>Amin Karbasi, Grigoris Velegkas, Zhuoran Yang</i>	
Cluster and Aggregate: Face Recognition with Large Probe Set.....	36054
<i>Anil K Jain, Minchul Kim, Feng Liu, Xiaoming Liu</i>	
GLIPv2: Unifying Localization and Vision-Language Understanding	36067
<i>Yen-Chun Chen, Xiyang Dai, Jianfeng Gao, Xiaowei Hu, Jenq-Neng Hwang, Liunian Li, Lijuan Wang, Lu Yuan, Haotian Zhang, Pengchuan Zhang</i>	
Rethinking Alignment in Video Super-Resolution Transformers	36081
<i>Chao Dong, Jinjin Gu, Shuwei Shi, Xintao Wang, Liangbin Xie, Yujiu Yang</i>	
Robustness in Deep Learning: The Good (width), the Bad (depth), and the Ugly (initialization)	36094
<i>Volkan Cevher, Grigorios Chrysos, Fanghui Liu, Zhenyu Zhu</i>	
Learning to Scaffold: Optimizing Model Explanations for Teaching.....	36108
<i>Patrick Fernandes, André Martins, Graham Neubig, Danish Pruthi, Marcos Treviso</i>	
AutoLink: Self-Supervised Learning of Human Skeletons and Object Outlines by Linking Keypoints	36123
<i>Xingzhe He, Helge Rhodin, Bastian Wandt</i>	
Composite Feature Selection Using Deep Ensembles.....	36142
<i>Fergus Imrie, Pietro Lió, Alexander Norcliffe, Mihaela Van Der Schaar</i>	
Scaling Multimodal Pre-Training Via Cross-Modality Gradient Harmonization	36161
<i>Hassan Akbari, Feng Han, Yi Liang, Zhangyang Wang, Junru Wu, Cong Yu</i>	
Bounded-Regret MPC Via Perturbation Analysis: Prediction Error, Constraints, and Nonlinearity	36174
<i>Yang Hu, Tongxin Li, Yiheng Lin, Guannan Qu, Adam Wierman</i>	
AniFaceGAN: Animatable 3D-Aware Face Image Generation for Video Avatars	36188
<i>Qifeng Chen, Yu Deng, Xin Tong, Fangyun Wei, Yue Wu, Jiaolong Yang</i>	
Addressing Resource Scarcity Across Sign Languages with Multilingual Pretraining and Unified-Vocabulary Datasets	36202
<i>Mitesh Khapra, Pratyush Kumar, Manideep Ladi, Gokul NC, Sumit Negi, Prem Selvaraj</i>	
Detection and Localization of Changes in Conditional Distributions	36216
<i>Dan Nicolae, Lizhen Nie</i>	
Refining Low-Resource Unsupervised Translation by Language Disentanglement of Multilingual Translation Model.....	36230
<i>Ai Ti Aw, Shafiq Joty, Xuan-Phi Nguyen, Kui Wu</i>	
A Mean-Field Game Approach to Cloud Resource Management with Function Approximation	36243
<i>Tamer Basar, Hubertus Franke, Ravishankar Iyer, Zbigniew Kalbarczyk, Weichao Mao, Haoran Qiu, Chen Wang</i>	
Regret Bounds for Risk-Sensitive Reinforcement Learning	36259
<i>Osbert Bastani, Jason Yecheng Ma, Estelle Shen, Wanqiao Xu</i>	

BILCO: An Efficient Algorithm for Joint Alignment of Time Series.....	36270
<i>Misha B Ahrens, Alex Chen, Jing-Xuan Lim, Xuelong Mi, Mengfan Wang, Yizhi Wang, Guoqiang Yu</i>	
Giving Feedback on Interactive Student Programs with Meta-Exploration.....	36282
<i>Emma Brunskill, Chelsea Finn, Evan Liu, Allen Nie, Chris Piech, Moritz Stephan</i>	
Left Heavy Tails and the Effectiveness of the Policy and Value Networks in DNN-Based Best-First Search for Sokoban Planning	36295
<i>Dieqiao Feng, Carla Gomes, Bart Selman</i>	
Nonparametric Uncertainty Quantification for Single Deterministic Neural Network	36308
<i>Aleksandr Artemenkov, Kirill Fedyanin, Alexander Fishkov, Nikita Kotelevskii, Fedor Noskov, Maxim Panov, Aleksandr Petiushko, Artem Shelmanov, Artem Vazhentsev</i>	
Decoupling Features in Hierarchical Propagation for Video Object Segmentation.....	36324
<i>Yi Yang, Zongxin Yang</i>	
Computationally Efficient Horizon-Free Reinforcement Learning for Linear Mixture MDPs	36337
<i>Quanquan Gu, Dongruo Zhou</i>	
Counterfactual Harm	36350
<i>Rory Beard, Jonathan Richens, Daniel H. Thompson</i>	
Chain of Thought Imitation with Procedure Cloning	36366
<i>Pieter Abbeel, Ofir Nachum, Dale Schuurmans, Mengjiao Sherry Yang</i>	
Revisiting Optimal Convergence Rate for Smooth and Non-Convex Stochastic Decentralized Optimization.....	36382
<i>Yiming Chen, Ximmeng Huang, Pan Pan, Kun Yuan, Xiaohan Zhang, Yingya Zhang</i>	
Training with More Confidence: Mitigating Injected and Natural Backdoors During Training.....	36396
<i>Hailun Ding, Shiqing Ma, Zhenting Wang, Juan Zhai</i>	
Out-Of-Distribution Detection Via Conditional Kernel Independence Model	36411
<i>Jingyang Lin, Qing Ling, Tao Mei, Yingwei Pan, Yu Wang, Ting Yao, Jingjing Zou</i>	
Online Convex Optimization with Hard Constraints: Towards the Best of Two Worlds and Beyond.....	36426
<i>Hengquan Guo, Xin Liu, Honghao Wei, Lei Ying</i>	
ResT V2: Simpler, Faster and Stronger	36440
<i>Yu-Bin Yang, Qinglong Zhang</i>	
Object-Category Aware Reinforcement Learning	36453
<i>Yunji Chen, Zidong Du, Jiaming Guo, Qi Guo, Xing Hu, Shaohui Peng, Qi Yi, Xishan Zhang, Rui Zhang</i>	
Learning Partial Equivariances from Data	36466
<i>Sahas Lohit, David W. Romero</i>	
Photorealistic Text-To-Image Diffusion Models with Deep Language Understanding.....	36479
<i>William Chan, Emily Denton, David Fleet, Kamyar Ghasemipour, Raphael Gontijo Lopes, Jonathan Ho, Burcu Karagol Ayan, Lala Li, Mohammad Norouzi, Chitwan Saharia, Tim Salimans, Saurabh Saxena, Jay Whang</i>	
A Simple Decentralized Cross-Entropy Method	36495
<i>Martin Jagersand, Jun Jin, Jun Luo, Dale Schuurmans, Zichen Zhang</i>	

MSDS: A Large-Scale Chinese Signature and Token Digit String Dataset for Handwriting Verification	36507
<i>Jiajia Jiang, Lianwen Jin, Yuliang Liu, Peirong Zhang</i>	
Do Residual Neural Networks Discretize Neural Ordinary Differential Equations?.....	36520
<i>Pierre Ablin, Gabriel Peyré, Michael Sander</i>	
Diffusion-Based Molecule Generation with Informative Prior Bridges	36533
<i>Chengyue Gong, Qiang Liu, Xingchao Liu, Lemeng Wu, Mao Ye</i>	
Generative Multitask Learning Mitigates Target-Causing Confounding.....	36546
<i>Kyunghyun Cho, Krzysztof Geras, Taro Makino</i>	
Revisit Last-Iterate Convergence of mSGD Under Milder Requirement on Step Size	36559
<i>Lang Chen, Difei Cheng, Vijay Gupta, Xingkang He, Ruinan Jin</i>	
Are You Stealing My Model? Sample Correlation for Fingerprinting Deep Neural Networks	36571
<i>Jiyang Guan, Ran He, Jian Liang</i>	
SecureFedYJ: A Safe Feature Gaussianization Protocol for Federated Learning.....	36585
<i>Mathieu Andreux, Constance Béguier, Tanguy Marchand, Boris Muzellec, Jean Ogier Du Terrail</i>	
When to Trust Your Simulator: Dynamics-Aware Hybrid Offline-And-Online Reinforcement Learning	36599
<i>Jianming Hu, Ming Li, Haoyi Niu, Yiwen Qiu, Shubham Sharma, Xianyuan Zhan, Guyue Zhou</i>	
Data-Efficient Structured Pruning Via Submodular Optimization	36613
<i>Marwa El Halabi, Simon Lacoste-Julien, Suraj Srinivas</i>	
Interpolation and Regularization for Causal Learning	36627
<i>Leena Chennuru Vankadara, Debarghya Ghoshdastidar, Ulrike Luxburg, Luca Rendburg</i>	
A Direct Approximation of AIXI Using Logical State Abstractions	36640
<i>Kee Siong Ng, Tianyu Wang, Samuel Yang-Zhao</i>	
Learning Representations Via a Robust Behavioral Metric for Deep Reinforcement Learning.....	36654
<i>Jianda Chen, Sinno Pan</i>	
Finding Second-Order Stationary Points in Nonconvex-Strongly-Concave Minimax Optimization.....	36667
<i>Cheng Chen, Yujun Li, Luo Luo</i>	
Contextual Squeeze-And-Excitation for Efficient Few-Shot Image Classification.....	36680
<i>John Bronskill, Katja Hofmann, Sebastian Nowozin, Massimiliano Patacchiola, Aliaksandra Shysheya, Richard Turner</i>	
Fuzzy Learning Machine.....	36693
<i>Junbiao Cui, Jiye Liang</i>	
Learning Structure from the Ground up - Hierarchical Representation Learning by Chunking.....	36706
<i>Ishita Dasgupta, Noemi Elteto, Eric Schulz, Shuchen Wu</i>	
AMOS: A Large-Scale Abdominal Multi-Organ Benchmark for Versatile Medical Image Segmentation	36722
<i>Haotian Bai, Chongjian Ge, Yuanfeng Ji, Zhen Li, Ping Luo, Wanling Ma, Xiang Wan, Jie Yang, Ruimao Zhang, Lingyan Zhannng, Ye Zhu</i>	

VOLUME 48

The Price of Ignorance: How Much Does it Cost to Forget Noise Structure in Low-Rank Matrix Estimation?.....	36733
<i>Jean Barbier, Tianqi Hou, Marco Mondelli, Manuel Saenz</i>	
Scalable Interpretability Via Polynomials	36748
<i>Abhimanyu Dubey, Dhruv Mahajan, Filip Radenovic</i>	
DeVRF: Fast Deformable Voxel Radiance Fields for Dynamic Scenes	36762
<i>Yan-Pei Cao, Jussi Keppo, Jia-Wei Liu, Weijia Mao, Xiaohu Qie, Ying Shan, Mike Zheng Shou, David Junhao Zhang, Wenqiao Zhang</i>	
Hiding Images in Deep Probabilistic Models	36776
<i>Haoyu Chen, Kede Ma, Zhenxing Qian, Linqi Song, Xinpeng Zhang</i>	
ViewFool: Evaluating the Robustness of Visual Recognition to Adversarial Viewpoints	36789
<i>Yinpeng Dong, Caixin Kang, Shouwei Ruan, Hang Su, Xingxing Wei, Jun Zhu</i>	
Learning Superpoint Graph Cut for 3D Instance Segmentation	36804
<i>Le Hui, Yaqi Shen, Linghua Tang, Jin Xie, Jian Yang</i>	
Estimating the Arc Length of the Optimal ROC Curve and Lower Bounding the Maximal AUC	36818
<i>Song Liu</i>	
Acceleration in Distributed Sparse Regression	36832
<i>Marie Maros, Gesualdo Scutari</i>	
Latency-Aware Spatial-Wise Dynamic Networks	36845
<i>Yizeng Han, Gao Huang, Yifan Pu, Shiji Song, Guangyu Sun, Chenhao Xue, Zhihang Yuan</i>	
Towards Versatile Embodied Navigation	36858
<i>Luc V Gool, Wei Liang, Hanqing Wang, Wenguan Wang</i>	
Explain My Surprise: Learning Efficient Long-Term Memory by Predicting Uncertain Outcomes	36875
<i>Mikhail Burtsev, Nazar Buzun, Leonid Pugachev, Artyom Sorokin</i>	
Polyhistor: Parameter-Efficient Multi-Task Adaptation for Dense Vision Tasks	36889
<i>Zijian He, Zsolt Kira, Yen-Cheng Liu, Chih-Yao Ma, Junjiao Tian</i>	
LAPO: Latent-Variable Advantage-Weighted Policy Optimization for Offline Reinforcement Learning	36902
<i>Liang Bin, Xi Chen, Chelsea Finn, Alex Yuan Gao, Ali Ghadirzadeh, Wenzhe Li, Jianhao Wang, Tianhe Yu, Chongjie Zhang</i>	
GAMA: Generative Adversarial Multi-Object Scene Attacks	36914
<i>Abhishek Aich, Salman Asif, Akash Gupta, Srikanth Krishnamurthy, Amit Roy-Chowdhury, Chengyu Song, Calvin-Khang Ta</i>	
Predicting Label Distribution from Multi-Label Ranking	36931
<i>Xiuyi Jia, Yunan Lu</i>	
Trajectory of Mini-Batch Momentum: Batch Size Saturation and Convergence in High Dimensions	36944
<i>Andrew Cheng, Kiwon Lee, Courtney Paquette, Elliot Paquette</i>	

Transformers from an Optimization Perspective	36958
<i>Zengfeng Huang, David P Wipf, Yongyi Yang</i>	
PDSketch: Integrated Domain Programming, Learning, and Planning	36972
<i>Leslie Kaelbling, Tomás Lozano-Pérez, Jiayuan Mao, Josh Tenenbaum</i>	
Amortized Projection Optimization for Sliced Wasserstein Generative Models	36985
<i>Nhat Ho, Khai Nguyen</i>	
GT-GAN: General Purpose Time Series Synthesis with Generative Adversarial Networks	36999
<i>Seunghyeon Cho, Jinsung Jeon, Jeonghak Kim, Noseong Park, Haryong Song</i>	
Heterogeneous Skill Learning for Multi-Agent Tasks	37011
<i>Yong Dou, Yuan Li, Donghong Liu, Yuntao Liu, Xinhai Xu</i>	
On Margin Maximization in Linear and ReLU Networks	37024
<i>Ohad Shamir, Nati Srebro, Gal Vardi</i>	
Unsupervised Causal Generative Understanding of Images	37037
<i>Titas Anciukevicius, Patrick Fox-Roberts, Paul Henderson, Edward Rosten</i>	
DART: Articulated Hand Model with Diverse Accessories and Rich Textures	37055
<i>Daiheng Gao, Kailin Li, Cewu Lu, Ping Tan, Feng Wang, Yuliang Xiu, Lixin Yang, Bang Zhang, Peng Zhang</i>	
BadPrompt: Backdoor Attacks on Continuous Prompts	37068
<i>Xiangrui Cai, Yuan Xiaojie, Haidong Xu, Sihan Xu, Ying Zhang</i>	
Off-Policy Evaluation with Policy-Dependent Optimization Response	37081
<i>Wenshuo Guo, Michael Jordan, Angela Zhou</i>	
FIRE: Semantic Field of Words Represented as Non-Linear Functions	37095
<i>Xin Du, Kumiko Tanaka-Ishii</i>	
Sampling in Constrained Domains with Orthogonal-Space Variational Gradient Descent	37108
<i>Qiang Liu, Xin Tong, Ruqi Zhang</i>	
Change-Point Detection for Sparse and Dense Functional Data in General Dimensions	37121
<i>Carlos Misael Madrid Padilla, Daren Wang, Yi Yu, Zifeng Zhao</i>	
ClimbQ: Class Imbalanced Quantization Enabling Robustness on Efficient Inferences	37134
<i>Ming-Syan Chen, Ting-An Chen, De-Nian Yang</i>	
Look Around and Refer: 2D Synthetic Semantics Knowledge Distillation for 3D Visual Grounding	37146
<i>Yasmeen Alsaedy, Eslam Bakr, Mohamed Elhoseiny</i>	
Nonstationary Dual Averaging and Online Fair Allocation	37159
<i>Yuan Gao, Christian Kroer, Luofeng Liao</i>	
Incentivizing Combinatorial Bandit Exploration	37173
<i>Xinyan Hu, Dung Ngo, Aleksandrs Slivkins, Steven Wu</i>	
Transfer Learning on Heterogeneous Feature Spaces for Treatment Effects Estimation	37184
<i>Ioana Bica, Mihaela Van Der Schaar</i>	
Is Out-Of-Distribution Detection Learnable?	37199
<i>Jiahua Dong, Zhen Fang, Bo Han, Yixuan Li, Feng Liu, Jie Lu</i>	

DMAP: A Distributed Morphological Attention Policy for Learning to Locomote with a Changing Body	37214
<i>Alberto Silvio Chiappa, Alessandro Marin Vargas, Alexander Mathis</i>	
On the Role of Overparameterization in Off-Policy Temporal Difference Learning with Linear Function Approximation.....	37228
<i>Valentin Thomas</i>	
Whitening Convergence Rate of Coupling-Based Normalizing Flows	37241
<i>Felix Draxler, Ullrich Köthe, Christoph Schnörr</i>	
Neurosymbolic Deep Generative Models for Sequence Data with Relational Constraints	37254
<i>Osbert Bastani, Maxwell Du, Halley Young</i>	
Semi-Supervised Generative Models for Multiagent Trajectories.....	37267
<i>Ulf Brefeld, Dennis Fassmeyer, Pascal Fassmeyer</i>	
PhysGNN: A Physics--Driven Graph Neural Network Based Model for Predicting Soft Tissue Deformation in Image--Guided Neurosurgery	37282
<i>Dennis Giannacopoulos, Yasmin Salehi</i>	
Towards Diverse and Faithful One-Shot Adaption of Generative Adversarial Networks	37297
<i>Jinfeng Bai, Zhilong Ji, Yuxiang Wei, Mingshuai Yao, Yabo Zhang, Wangmeng Zuo</i>	
Deep Bidirectional Language-Knowledge Graph Pretraining	37309
<i>Antoine Bosselut, Jure Leskovec, Percy Liang, Christopher D Manning, Hongyu Ren, Michihiro Yasunaga, Xikun Zhang</i>	
Theoretical Analysis of Deep Neural Networks for Temporally Dependent Observations	37324
<i>Mingliang Ma, Abolfazl Safikhani</i>	
A Theoretical Framework for Inference Learning.....	37335
<i>Nick Alonso, Jeffrey Krichmar, Beren Millidge, Emre O Neftci</i>	
Cross-Modal Learning for Image-Guided Point Cloud Shape Completion.....	37349
<i>Emanuele Aiello, Enrico Magli, Diego Valsesia</i>	
Coreset for Line-Sets Clustering	37363
<i>Dan Feldman, Sagi Lotan, Ernesto Evgeniy Sanches Shayda</i>	
Robust On-Policy Sampling for Data-Efficient Policy Evaluation in Reinforcement Learning.....	37376
<i>Stefano Albrecht, Josiah Hanna, Lukas Schäfer, Duohan Zhang, Rujie Zhong</i>	
Pre-Trained Model Reusability Evaluation for Small-Data Transfer Learning	37389
<i>Yao-Xiang Ding, Xi-Zhu Wu, Zhi-Hua Zhou, Kun Zhou</i>	
Tree Ensemble Kernels for Bayesian Optimization with Known Constraints Over Mixed-Feature Spaces.....	37401
<i>Robert Lee, Ruth Misener, Behrang Shafei, Nathan Sudermann-Merx, Alexander Thebelt, Calvin Tsay, David Walz</i>	
RLIP: Relational Language-Image Pre-Training for Human-Object Interaction Detection	37416
<i>Samuel Albanie, Tao Feng, Ziyuan Huang, Jianwen Jiang, Dong Ni, Mingqian Tang, Hangjie Yuan</i>	
Skills Regularized Task Decomposition for Multi-Task Offline Reinforcement Learning.....	37432
<i>Sangwoo Cho, Honguk Woo, Minjong Yoo</i>	

Contextual Dynamic Pricing with Unknown Noise: Explore-Then-UCB Strategy and Improved Regrets.....	37445
<i>Yufeng Liu, Yiyun Luo, Will Wei Sun</i>	
Adversarially Robust Learning: A Generic Minimax Optimal Learner and Characterization.....	37458
<i>Steve Hanneke, Omar Montasser, Nati Srebro</i>	
The Implicit Delta Method.....	37471
<i>Nathan Kallus, James McInerney</i>	
Singular Value Fine-Tuning: Few-Shot Segmentation Requires Few-Parameters Fine-Tuning.....	37484
<i>Qiang Chen, Jian Cheng, Errui Ding, Haocheng Feng, Junyu Han, Xiangyu He, Zechao Li, Yanpeng Sun, Jian Wang, Jingdong Wang</i>	
On the Relationship Between Variational Inference and Auto-Associative Memory.....	37497
<i>Louis Annabi, Alexandre Pitti, Mathias Quoy</i>	
Annihilation of Spurious Minima in Two-Layer ReLU Networks.....	37510
<i>Yossi Arjevani, Michael Field</i>	

VOLUME 49

A Closer Look at Weakly-Supervised Audio-Visual Source Localization.....	37524
<i>Shentong Mo, Pedro Morgado</i>	
Posterior Collapse of a Linear Latent Variable Model.....	37537
<i>Zihao Wang, Liu Ziyin</i>	
Accelerating SGD for Highly Ill-Conditioned Huge-Scale Online Matrix Completion.....	37549
<i>Hong-Ming Chiu, Jialun Zhang, Richard Y Zhang</i>	
Lifting Weak Supervision to Structured Prediction.....	37563
<i>Frederic Sala, Harit Vishwakarma</i>	
A Lagrangian Duality Approach to Active Learning.....	37575
<i>Juan Elenter, Navid Naderializadeh, Alejandro Ribeiro</i>	
Instance-Optimal PAC Algorithms for Contextual Bandits.....	37590
<i>Lalit Jain, Kevin Jamieson, Zhaoqi Li, Houssam Nassif, Lillian Ratliff</i>	
xView3-SAR: Detecting Dark Fishing Activity Using Synthetic Aperture Radar Imagery.....	37604
<i>Jared Dunnmon, Bryce Goodman, Ritwik Gupta, David Kroodsma, Daniel Kuster, Tsu-Ting Tim Lin, Fernando Paolo, Nirav Patel</i>	
Understanding the Failure of Batch Normalization for Transformers in NLP.....	37617
<i>Lei Huang, Jiayi Wang, Ji Wu</i>	
Exploration Via Elliptical Episodic Bonuses.....	37631
<i>Mikael Henaff, Minqi Jiang, Roberta Raileanu, Tim Rocktäschel</i>	
Efficient and Effective Multi-Task Grouping Via Meta Learning on Task Combinations.....	37647
<i>Jiang Bian, Wei Cao, Xiaozhuang Song, James Yu, Shun Zheng</i>	
Learning Fractional White Noises in Neural Stochastic Differential Equations.....	37660
<i>Jaesik Choi, Thanh Nguyen-Tang, Anh Tong, Toan Tran</i>	

CryptoGCN: Fast and Scalable Homomorphically Encrypted Graph Convolutional Network Inference.....	37676
<i>Quan Gang, Ran Ran, Wei Wang, Wujie Wen, Nuo Xu, Jieming Yin</i>	
UniGAN: Reducing Mode Collapse in GANs Using a Uniform Generator.....	37690
<i>Li Niu, Ziqi Pan, Liqing Zhang</i>	
UMIX: Improving Importance Weighting for Subpopulation Shift Via Uncertainty-Aware Mixup	37704
<i>Yatao Bian, Zongbo Han, Lanqing Li, Zhipeng Liang, Liu Liu, Bingzhe Wu, Fan Yang, Jianhua Yao, Changqing Zhang, Peilin Zhao</i>	
Exploit Reward Shifting in Value-Based Deep-RL: Optimistic Curiosity-Based Exploration and Conservative Exploitation Via Linear Reward Shaping	37719
<i>Jian Guo, Lei Han, Xiaoteng Ma, Hao Sun, Rui Yang, Bolei Zhou</i>	
Stability and Generalization for Markov Chain Stochastic Gradient Methods.....	37735
<i>Yunwen Lei, Puyu Wang, Yiming Ying, Ding-Xuan Zhou</i>	
Degradation-Aware Unfolding Half-Shuffle Transformer for Spectral Compressive Imaging	37749
<i>Yuanhao Cai, Henghui Ding, Luc V Gool, Jing Lin, Radu Timofte, Haoqian Wang, Xin Yuan, Yulun Zhang</i>	
Deep Surrogate Assisted Generation of Environments.....	37762
<i>Varun Bhatt, Matthew Fontaine, Stefanos Nikolaidis, Bryon Tjanaka</i>	
Toward Equation of Motion for Deep Neural Networks: Continuous-Time Gradient Descent and Discretization Error Analysis.....	37778
<i>Taiki Miyagawa</i>	
FLAIR: Federated Learning Annotated Image Repository.....	37792
<i>Filip Granqvist, Congzheng Song, Kunal Talwar</i>	
Self-Supervised Pretraining for Large-Scale Point Clouds.....	37806
<i>Min Bai, Erran Li Li, Zaiwei Zhang</i>	
Vision Transformers Provably Learn Spatial Structure	37822
<i>Samy Jelassi, Yuanzhi Li, Michael Sander</i>	
Unsupervised Learning of Shape Programs with Repeatable Implicit Parts	37837
<i>Boyang Deng, Congyue Deng, Zhengyang Dong, Sumith Kulal, Yonglong Tian, Jiajun Wu</i>	
Detecting Abrupt Changes in Sequential Pairwise Comparison Data	37851
<i>Wanshan Li, Alessandro Rinaldo, Daren Wang</i>	
Rethinking Resolution in the Context of Efficient Video Recognition.....	37865
<i>Qiushan Guo, Yi Jiang, Ping Luo, Chuofan Ma, Xiaojuan Qi, Zehuan Yuan</i>	
The Effects of Regularization and Data Augmentation Are Class Dependent.....	37878
<i>Randall Balestriero, Leon Bottou, Yann Lecun</i>	
Learning to Mitigate AI Collusion on Economic Platforms	37892
<i>Gianluca Brero, Nicolas Lepore, Eric Mibuari, David Parkes</i>	
The Query Complexity of Cake Cutting.....	37905
<i>Simina Branzei, Noam Nisan</i>	

PAC Prediction Sets for Meta-Learning	37920
<i>Osbert Bastani, Edgar Dobriban, Insup Lee, Sangdon Park</i>	
High-Dimensional Asymptotics of Feature Learning: How One Gradient Step Improves the Representation	37932
<i>Jimmy Ba, Murat Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, Greg Yang</i>	
Pruning’s Effect on Generalization Through the Lens of Training and Regularization.....	37947
<i>Michael Carbin, Gintare Karolina Dziugaite, Jonathan Frankle, Tian Jin, Dan Roy</i>	
Cluster Randomized Designs for One-Sided Bipartite Experiments	37962
<i>Jennifer Brennan, Vahab Mirrokni, Jean Pouget-Abadie</i>	
Deep Equilibrium Approaches to Diffusion Models	37975
<i>Zhengyang Geng, J. Zico Kolter, Ashwini Pokle</i>	
Inducing Neural Collapse in Imbalanced Learning: Do We Really Need a Learnable Classifier at the End of Deep Neural Network?	37991
<i>Shixiang Chen, Xiangtai Li, Zhouchen Lin, Dacheng Tao, Liang Xie, Yibo Yang</i>	
Pruning Neural Networks Via Coresets and Convex Geometry: Towards No Assumptions	38003
<i>Alaa Maalouf, Loay Mualem, Murad Tukan</i>	
Intermediate Prototype Mining Transformer for Few-Shot Semantic Segmentation	38020
<i>Junwei Han, Nian Liu, Yuanwei Liu, Xiwen Yao</i>	
Long-Form Video-Language Pre-Training with Multimodal Temporal Contrastive Learning.....	38032
<i>Jianlong Fu, Bei Liu, Ruihua Song, Yuchong Sun, Hongwei Xue, Huan Yang</i>	
A New Dataset for Multilingual Keyphrase Generation.....	38046
<i>Philippe Langlais, Frédéric Piedboeuf</i>	
Model Preserving Compression for Neural Networks.....	38060
<i>Jerry Chee, Anil Damle, Christopher De Sa, Megan Flynn Née Renz</i>	
Neural Conservation Laws: A Divergence-Free Perspective	38075
<i>Ricky T. Q. Chen, Yaron Lipman, Jack Richter-Powell</i>	
Towards Effective Multi-Modal Interchanges in Zero-Resource Sounding Object Localization.....	38089
<i>Haifeng Huang, Haoyuan Li, Chen Zhang, Yang Zhao, Zhou Zhao</i>	
On the Convergence of Stochastic Multi-Objective Gradient Manipulation and Beyond.....	38103
<i>Jinjie Gu, Jiyan Jiang, Wenpeng Zhang, Wenliang Zhong, Shiji Zhou, Wenwu Zhu</i>	
Decentralized Local Stochastic Extra-Gradient for Variational Inequalities	38116
<i>Aleksandr Beznosikov, Pavel Dvurechenskii, Alexander Gasnikov, Anastasiia Koloskova, Valentin Samokhin, Sebastian Stich</i>	
Model Zoos: A Dataset of Diverse Populations of Neural Network Models.....	38134
<i>Damian Borth, Xavier Giró-I-Nieto, Boris Knyazev, Konstantin Schürholt, Diyar Taskiran</i>	
Weakly-Supervised Multi-Granularity Map Learning for Vision-And-Language Navigation	38149
<i>Peihao Chen, Chuang Gan, Dongyu Ji, Thomas Li, Kunyang Lin, Mingkui Tan, Runhao Zeng</i>	
Thinned Random Measures for Sparse Graphs with Overlapping Communities	38162
<i>Michele Guindani, Federica Zoe Ricci, Erik Sudderth</i>	

Fine-Tuning Language Models to Find Agreement Among Humans with Diverse Preferences	38176
<i>John Aslanides, Michiel Bakker, Jan Balaguer, Matt Botvinick, Lucy Campbell-Gillingham, Martin Chadwick, Amelia Glaese, Nat McAleese, Hannah Sheahan, Christopher Summerfield, Michael Tessler</i>	
What is a Good Metric to Study Generalization of Minimax Learners?	38190
<i>Asuman Ozdaglar, Sarath Pattathil, Jiawei Zhang, Kaiqing Zhang</i>	
Sequencer: Deep LSTM for Image Classification.....	38204
<i>Masato Taki, Yuki Tatsunami</i>	
Double Check Your State Before Trusting It: Confidence-Aware Bidirectional Offline Model-Based Imagination.....	38218
<i>Xiu Li, Zongqing Lu, Jiafei Lyu</i>	
Learning on Arbitrary Graph Topologies Via Predictive Coding	38232
<i>Tianyi Bao, Rafal Bogacz, Thomas Lukasiewicz, Beren Millidge, Luca Pinchetti, Tommaso Salvatori, Yuhang Song</i>	
Robustness Disparities in Face Detection.....	38245
<i>John Dickerson, Samuel Dooley, Tom Goldstein, George Z Wei</i>	
Marksman Backdoor: Backdoor Attacks with Arbitrary Target Class.....	38260
<i>Khoa D Doan, Yingjie Lao, Ping Li</i>	
Memorization Without Overfitting: Analyzing the Training Dynamics of Large Language Models	38274
<i>Armen Aghajanyan, Aram Markosyan, Kushal Tirumala, Luke Zettlemoyer</i>	
Bandit Theory and Thompson Sampling-Guided Directed Evolution for Sequence Optimization.....	38291
<i>Le Cong, Chengzhuo Ni, Csaba Szepesvari, Huazheng Wang, Mengdi Wang, Hui Yuan, Xuezhou Zhang</i>	
Scalable and Efficient Training of Large Convolutional Neural Networks with Differential Privacy	38305
<i>Zhiqi Bu, Jialin Mao, Shiyun Xu</i>	

VOLUME 50

Weakly Supervised Causal Representation Learning	38319
<i>Johann Brehmer, Taco Cohen, Pim De Haan, Phillip Lippe</i>	
Zeroth-Order Negative Curvature Finding: Escaping Saddle Points Without Gradients.....	38332
<i>Bin Gu, Huan Xiong, Hualin Zhang</i>	
Exposing and Exploiting Fine-Grained Block Structures for Fast and Accurate Sparse Training.....	38345
<i>Lihan Hu, Peng Jiang, Shihui Song</i>	
DABS 2.0: Improved Datasets and Algorithms for Universal Self-Supervision.....	38358
<i>Gaurab Banerjee, Noah Goodman, Vincent Liu, Mohamed Owda, Shashank Rammoorthy, Alex Tamkin</i>	
Operator Splitting Value Iteration.....	38373
<i>Amir-Massoud Farahmand, Mohammad Ghavamzadeh, Amin Rakhsha, Andrew Wang</i>	
Enhanced Latent Space Blind Model for Real Image Denoising Via Alternative Optimization	38386
<i>Jie Huang, Yizhong Pan, Chao Ren</i>	

Self-Explaining Deviations for Coordination.....	38400
<i>Anton Bakhtin, Brandon Cui, Jakob Foerster, Hengyuan Hu, Andrei Lupu, Samuel Sokota, David Wu</i>	
Communication Efficient Federated Learning for Generalized Linear Bandits	38411
<i>Chuanhao Li, Hongning Wang</i>	
Active Learning for Multiple Target Models.....	38424
<i>Sheng-Jun Huang, Ying-Peng Tang</i>	
Descent Steps of a Relation-Aware Energy Produce Heterogeneous Graph Neural Networks	38436
<i>Hongjoon Ahn, Quan Gan, Taesup Moon, David P Wipf, Yongyi Yang</i>	
Multi-Agent Performative Prediction with Greedy Deployment and Consensus Seeking Agents	38449
<i>Qiang Li, Hoi-To Wai, Chung-Yiu Yau</i>	
Preservation of the Global Knowledge by Not-True Distillation in Federated Learning	38461
<i>Sangmin Bae, Minchan Jeong, Gihun Lee, Yongjin Shin, Se-Young Yun</i>	
Finite-Time Regret of Thompson Sampling Algorithms for Exponential Family Multi-Armed Bandits.....	38475
<i>Anima Anandkumar, Tianyuan Jin, Xiaokui Xiao, Pan Xu</i>	
Graph Reordering for Cache-Efficient Near Neighbor Search.....	38488
<i>Benjamin Coleman, Santiago Segarra, Anshumali Shrivastava, Alexander Smola</i>	
MetaMask: Revisiting Dimensional Confounder for Self-Supervised Learning.....	38501
<i>Jiangmeng Li, Wenyi Mo, Wenwen Qiang, Bing Su, Hui Xiong, Yanan Zhang, Changwen Zheng</i>	
On Feature Learning in the Presence of Spurious Correlations.....	38516
<i>Nate Gruver, Pavel Izmailov, Polina Kirichenko, Andrew Wilson</i>	
Sparse2Dense: Learning to Densify 3D Features for 3D Object Detection.....	38533
<i>Chi-Wing Fu, Xiaowei Hu, Zhengzhe Liu, Tianyu Wang</i>	
Exploring Length Generalization in Large Language Models.....	38546
<i>Anders Andreassen, Cem Anil, Ethan Dyer, Guy Gur-Ari, Aitor Lewkowycz, Vedant Misra, Behnam Neyshabur, Vinay Ramasesh, Ambrose Slone, Yuhuai Wu</i>	
Stability and Generalization Analysis of Gradient Methods for Shallow Neural Networks.....	38557
<i>Rong Jin, Yunwen Lei, Yiming Ying</i>	
ViTPose: Simple Vision Transformer Baselines for Human Pose Estimation.....	38571
<i>Dacheng Tao, Yufei Xu, Jing Zhang, Qiming Zhang</i>	
Re-Analyze Gauss: Bounds for Private Matrix Approximation Via Dyson Brownian Motion	38585
<i>Oren Mangoubi, Nisheeth Vishnoi</i>	
ASPiRe: Adaptive Skill Priors for Reinforcement Learning	38600
<i>Shuran Song, Manuela Veloso, Mengda Xu</i>	
Neural Differential Equations for Learning to Program Neural Nets Through Continuous Learning Rules.....	38614
<i>Francesco Faccio, Kazuki Irie, Jürgen Schmidhuber</i>	

MEMO: Test Time Robustness Via Adaptation and Augmentation.....	38629
<i>Chelsea Finn, Sergey Levine, Marvin Zhang</i>	
Learning-Augmented Algorithms for Online Linear and Semidefinite Programming	38643
<i>Elena Grigorescu, Young-San Lin, Sandeep Silwal, Maoyuan Song, Samson Zhou</i>	
Text-Adaptive Multiple Visual Prototype Matching for Video-Text Retrieval.....	38655
<i>Wenhong Ge, Junwei Liang, Chengzhi Lin, Chunhua Shen, Ancong Wu, Jun Zhang, Wei-Shi Zheng</i>	
Asymptotically Unbiased Instance-Wise Regularized Partial AUC Optimization: Theory and Algorithm	38667
<i>Shilong Bao, Qingming Huang, Huiyang Shao, Qianqian Xu, Zhiyong Yang</i>	
Universality of Group Convolutional Neural Networks Based on Ridgelet Analysis on Groups	38680
<i>Masahiro Ikeda, Isao Ishikawa, Sho Sonoda</i>	
Error Correction Code Transformer	38695
<i>Yoni Choukroun, Lior Wolf</i>	
Towards Improving Calibration in Object Detection Under Domain Shift	38706
<i>Mohsen Ali, Muhammad Haris Khan, Muhammad Akhtar Munir, M. Sarfraz</i>	
Renyi Differential Privacy of Propose-Test-Release and Applications to Private and Robust Machine Learning.....	38719
<i>Ruoxi Jia, Saeed Mahloujifar, Prateek Mittal, Jiachen T. Wang, Shouda Wang</i>	
A Transformer-Based Object Detector with Coarse-Fine Crossing Representations.....	38733
<i>Jianyuan Guo, Kai Han, Zhishan Li, Ying Nie, Yunhe Wang, Lei Xie</i>	
Beyond Adult and COMPAS: Fair Multi-Class Prediction Via Information Projection.....	38747
<i>Wael Alghamdi, Shahab Asoodeh, Flavio Calmon, Hsiang Hsu, Haewon Jeong, Peter Michalak, Hao Wang</i>	
Explicit Tradeoffs Between Adversarial and Natural Distributional Robustness	38761
<i>Kiarash Banihashem, Soheil Feizi, Mazda Moayeri</i>	
Learning Latent Seasonal-Trend Representations for Time Series Forecasting	38775
<i>Goce Trajcevski, Zhiyuan Wang, Xovee Xu, Weifeng Zhang, Ting Zhong, Fan Zhou</i>	
JAHS-Bench-201: A Foundation for Research on Joint Architecture and Hyperparameter Search	38788
<i>Archit Bansal, Frank Hutter, Maciej Janowski, Danny Stoll, Arber Zela</i>	
Capturing Graphs with Hypo-Elliptic Diffusions	38803
<i>Celia Hacker, Darrick Lee, Harald Oberhauser, Csaba Toth</i>	
A Spectral Approach to Item Response Theory.....	38818
<i>Duc Nguyen, Anderson Ye Zhang</i>	
FedSR: A Simple and Effective Domain Generalization Method for Federated Learning	38831
<i>Ser Nam Lim, A. Tuan Nguyen, Philip Torr</i>	
SIXO: Smoothing Inference with Twisted Objectives.....	38844
<i>Dieterich Lawson, Scott Linderman, Allan Raventós, Andrew Warrington</i>	
Explicable Policy Search.....	38859
<i>Ze Gong, Yu Tony Zhang</i>	

Exploring Evolution-Aware & -Free Protein Language Models as Protein Function Predictors	38873
<i>Qiuyang Ding, Mingyang Hu, Fusong Ju, Jin Su, Hui Wang, Fei Yang, Kevin Yang, Fajie Yuan</i>	
Breaking Bad: A Dataset for Geometric Fracture and Reassembly	38885
<i>Yun-Chun Chen, Animesh Garg, Alec Jacobson, Silvia Sellán, Ziyi Wu</i>	
Fair and Optimal Decision Trees: A Dynamic Programming Approach.....	38899
<i>Emir Demirovic, Mathijs De Weerd, Jacobus Van Der Linden</i>	
Recommender Forest for Efficient Retrieval.....	38912
<i>Enhong Chen, Chao Feng, Wuchao Li, Defu Lian, Zheng Liu</i>	
Graph Few-Shot Learning with Task-Specific Structures	38925
<i>Chen Chen, Jundong Li, Song Wang</i>	
DASCO: Dual-Generator Adversarial Support Constrained Offline Reinforcement Learning	38937
<i>Yevgen Chebotar, Aviral Kumar, Sergey Levine, Quan Vuong</i>	
Beyond L1: Faster and Better Sparse Models with Skglm.....	38950
<i>Pierre-Antoine Bannier, Quentin Bertrand, Gauthier Gidel, Quentin Klopfenstein, Mathurin Massias</i>	
You Can't Count on Luck: Why Decision Transformers and RvS Fail in Stochastic Environments	38966
<i>Jimmy Ba, Sheila McIlraith, Keiran Paster</i>	
Cost-Efficient Gaussian Tensor Network Embeddings for Tensor-Structured Inputs.....	38980
<i>Linjian Ma, Edgar Solomonik</i>	
Neural Transmitted Radiance Fields.....	38994
<i>Boxin Shi, Renjie Wan, Chengxuan Zhu</i>	
Geoclidean: Few-Shot Generalization in Euclidean Geometry	39007
<i>Noah Goodman, Joy Hsu, Jiajun Wu</i>	
Enabling Detailed Action Recognition Evaluation Through Video Dataset Augmentation.....	39020
<i>Jihoon Chung, Olga Russakovsky, Yu Wu</i>	
Unsupervised Skill Discovery Via Recurrent Skill Training	39034
<i>Jiayu Chen, Jingyue Gao, Zheyuan Jiang</i>	
Structural Kernel Search Via Bayesian Optimization and Symbolical Optimal Transport.....	39047
<i>Matthias Bitzer, Mona Meister, Christoph Zimmer</i>	
Robust Models Are Less Over-Confident.....	39059
<i>Paul Gavrikov, Julia Grabinski, Janis Keuper, Margret Keuper</i>	
Near-Optimal No-Regret Learning Dynamics for General Convex Games	39076
<i>Ioannis Anagnostides, Gabriele Farina, Christian Kroer, Chung-Wei Lee, Haipeng Luo, Tuomas Sandholm</i>	
OTKGE: Multi-Modal Knowledge Graph Embeddings Via Optimal Transport.....	39090
<i>Xiaochun Cao, Zongsheng Cao, Yuan He, Qingming Huang, Qianqian Xu, Zhiyong Yang</i>	
ComMU: Dataset for Combinatorial Music Generation.....	39103
<i>Sharang Han, Hyeonchan Hwang, Hyolim Kang, Minjoo Ki, Seon Joo Kim, Taehyun Kim, Hyun Lee, Kwanho Park</i>	

Author Index