

Advances in Neural Information Processing Systems 34

35th Conference on Neural Information
Processing Systems (NeurIPS 2021)

Online
6-14 December 2021

Volume 1 of 36

Editors:

M. Ranzato
A. Beygelzimer
Y. Dauphin

P.S. Liang
J. Wortman Vaughan

ISBN: 978-1-7138-4539-3

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© (2021) by individual authors and Neural Information Processing Systems Foundation Inc.
All rights reserved.

Printed by Curran Associates, Inc. (2022)

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

Beyond Value-Function Gaps: Improved Instance-Dependent Regret Bounds for Episodic Reinforcement Learning	1
<i>Christoph Dann, Teodor Vanislavov Marinov, Mehryar Mohri, Julian Zimmert</i>	
Learning One Representation to Optimize All Rewards	13
<i>Yann Ollivier, Ahmed Touati</i>	
Matrix Factorisation and the Interpretation of Geodesic Distance	24
<i>Annie Gray, Patrick Rubin-Delanchy, Nick Whiteley</i>	
UniDoc: Unified Pretraining Framework for Document Understanding.....	39
<i>Nikolaos Barmpalios, Jiuxiang Gu, Rajiv Jain, Jason Kuen, Vlad I Morariu, Ani Nenkova, Tong Sun, Handong Zhao</i>	
Finding Discriminative Filters for Specific Degradations in Blind Super-Resolution	51
<i>Chao Dong, Zhongang Qi, Ying Shan, Xintao Wang, Liangbin Xie</i>	
Counterfactual Explanations Can Be Manipulated.....	62
<i>Anna Hilgard, Himabindu Lakkaraju, Sameer Singh, Dylan Slack</i>	
From Canonical Correlation Analysis to Self-Supervised Graph Neural Networks.....	76
<i>David Wipf, Qitian Wu, Junchi Yan, Philip S Yu, Hengrui Zhang</i>	
BAST: Bayesian Additive Regression Spanning Trees for Complex Constrained Domain.....	90
<i>Zhao Tang Luo, Bani Mallick, Huiyan Sang</i>	
Hyperbolic Busemann Learning with Ideal Prototypes	103
<i>Mina Ghadimi Atigh, Martin Keller-Ressel, Pascal Mettes</i>	
Backward-Compatible Prediction Updates: a Probabilistic Approach.....	116
<i>Peter Gehler, Matthaus Kleindessner, Francesco Locatello, Bernhard Scholkopf, Frederik Trauble, Julius Von Kugelgen</i>	
Truncated Marginal Neural Ratio Estimation.....	129
<i>Alex Cole, Patrick Forre, Gilles Louppe, Benjamin K Miller, Christoph Weniger</i>	
ReAct: Out-Of-distribution Detection with Rectified Activations	144
<i>Chuan Guo, Yixuan Li, Yiyu Sun</i>	
Non-Local Latent Relation Distillation for Self-Adaptive 3D Human Pose Estimation	158
<i>Anirban Chakraborty, Anirudh Jamkhandi, Varun Jampani, Jogendra Nath Kundu, Venkatesh Babu R, Siddharth Seth, Pradyumna Ym</i>	
Fast Training of Neural Lumigraph Representations Using Meta Learning.....	172
<i>Alexander Bergman, Petr Kellnhofer, Gordon Wetzstein</i>	
Analytical Study of Momentum-Based Acceleration Methods in Paradigmatic High-Dimensional Non-Convex Problems	187
<i>Stefano Sarao Mannelli, Pierfrancesco Urbani</i>	

Multimodal Few-Shot Learning with Frozen Language Models.....	200
<i>Serkan Cabi, S. M. Ali Eslami, Felix Hill, Jacob L Menick, Maria Tsimpoukelli, Oriol Vinyals</i>	
Approximating the Permanent with Deep Rejection Sampling.....	213
<i>Juha Harviainen, Mikko Koivisto, Antti Roysko</i>	
Revisiting Model Stitching to Compare Neural Representations	225
<i>Yamini Bansal, Boaz Barak, Preetum Nakkiran</i>	
AugMax: Adversarial Composition of Random Augmentations for Robust Training	237
<i>Anima Anandkumar, Jean Kossaifi, Haotao Wang, Zhangyang Wang, Chaowei Xiao, Zhiding Yu</i>	
Habitat 2.0: Training Home Assistants to Rearrange Their Habitat.....	251
<i>Dhruv Batra, Angel Chang, Devendra Singh Chaplot, Alexander Clegg, Sameer Dharur, Wojciech Galuba, Aaron Gokaslan, Zolt Kira, Vladlen Koltun, Noah Maestre, Oleksandr Maksymets, Jitendra Malik, Franziska Meier, Mustafa Mukadam, Manolis Savva, Andrew Szot, John Turner, Eric Undersander, Vladimír Vondrus, Erik Wijmans, Yili Zhao</i>	
Time Discretization-Invariant Safe Action Repetition for Policy Gradient Methods.....	267
<i>Gunhee Kim, Jaekyeom Kim, Seohong Park</i>	
Meta-Learning Reliable Priors in the Function Space.....	280
<i>Jinfan Chen, Dominique Heyn, Andreas Krause, Jonas Rothfuss</i>	
VoiceMixer: Adversarial Voice Style Mixup.....	294
<i>Hyunseung Chung, Ji-Hoon Kim, Sang-Hoon Lee, Seong-Whan Lee</i>	
Predicting What You Already Know Helps: Provable Self-Supervised Learning.....	309
<i>Jason Lee, Qi Lei, Nikunj Saunshi, Jiacheng Zhuo</i>	
Oracle Complexity in Nonsmooth Nonconvex Optimization.....	324
<i>Guy Kornowski, Ohad Shamir</i>	
CentripetalText: an Efficient Text Instance Representation for Scene Text Detection	335
<i>Jie Chen, Zhouhui Lian, Tao Sheng</i>	
Learning to Select Exogenous Events for Marked Temporal Point Process	347
<i>Gaurav Aggarwal, Abir De, Rishabh Iyer, Ashish Tendulkar, Ping Zhang</i>	
DRIVE: One-Bit Distributed Mean Estimation.....	362
<i>Ran Ben-Basat, Yaniv Ben-Itzhak, Gal Mendelson, Michael Mitzenmacher, Amit Portnoy, Shay Vargaftik</i>	
Learning Space Partitions for Path Planning	378
<i>Brandon Cui, Chris Cummins, Joseph Gonzalez, Dan Klein, Benoit Steiner, Yuandong Tian, Linnan Wang, Kevin Yang, Tianjun Zhang</i>	
Progressive Feature Interaction Search for Deep Sparse Network	392
<i>Chen Gao, Depeng Jin, Yinfeng Li, Yong Li, Quanming Yao</i>	
Local Explanation of Dialogue Response Generation	404
<i>Wenhu Chen, Lise Getoor, Connor Pryor, Yi-Lin Tuan, William Yang Wang</i>	
Scalable Inference in SDEs by Direct Matching of the Fokker–Planck–Kolmogorov Equation	417
<i>Arno Solin, Ella Tamir, Prakhar Verma</i>	

The Complexity of Bayesian Network Learning: Revisiting the Superstructure.....	430
<i>Robert Ganian, Viktoriia Korchemna</i>	
Fast Tucker Rank Reduction for Non-Negative Tensors Using Mean-Field Approximation	443
<i>Kazu Ghalamkari, Mahito Sugiyama</i>	
Learning Stochastic Majority Votes by Minimizing a PAC-Bayes Generalization Bound.....	455
<i>Remi Emonet, Pascal Germain, Benjamin Guedj, Amaury Habrard, Emilie Morvant, Paul Viillard, Valentina Zantedeschi</i>	
Numerical Influence of ReLU'(0) on Backpropagation	468
<i>David Bertoin, Jerome Bolte, Sebastien Gerchinovitz, Edouard Pauwels</i>	
A Contrastive Learning Approach for Training Variational Autoencoder Priors.....	480
<i>Jyoti Aneja, Jan Kautz, Alex Schwing, Arash Vahdat</i>	
What Training Reveals About Neural Network Complexity.....	494
<i>Stefanie Jegelka, Andreas Loukas, Marinos Poiritis</i>	
Class-Agnostic Reconstruction of Dynamic Objects from Videos.....	509
<i>Zhongzheng Ren, Alex Schwing, Xiaoming Zhao</i>	
Unique Sparse Decomposition of Low Rank Matrices.....	523
<i>Xin Bing, Dian Jin, Yuqian Zhang</i>	
Neighborhood Reconstructing Autoencoders.....	536
<i>Hyeokjun Kwon, Yonghyeon Lee, Frank Park</i>	
TopicNet: Semantic Graph-Guided Topic Discovery	547
<i>Bo Chen, Zhibin Duan, Chaojie Wang, Yi Shi Xu, Mingyuan Zhou, Dongsheng Wang</i>	
(Almost) Free Incentivized Exploration from Decentralized Learning Agents.....	560
<i>Cong Shen, Chengshuai Shi, Wei Xiong, Haifeng Xu</i>	
Combining Recurrent, Convolutional, and Continuous-Time Models with Linear State Space Layers.....	572
<i>Tri Dao, Karan Goel, Albert Gu, Isys Johnson, Atri Rudra, Christopher Re, Khaled Saab</i>	
Revisiting Hilbert-Schmidt Information Bottleneck for Adversarial Robustness	586
<i>Jennifer Dy, Stratis Ioannidis, Tong Jian, Aria Masoomi, Zifeng Wang</i>	
T-LoHo: a Bayesian Regularization Model for Structured Sparsity and Smoothness on Graphs	598
<i>Changwoo Lee, Zhao Tang Luo, Huiyan Sang</i>	
The Utility of Explainable AI in Ad Hoc Human-Machine Teaming	610
<i>Muyleng Ghuy, Matthew Gombolay, Reed Jensen, Rohan Paleja, Nadun Ranawaka Arachchige</i>	
Subgoal Search for Complex Reasoning Tasks.....	624
<i>Konrad Czechowski, Lukasz Kucinski, Piotr Milos, Tomasz Odrzygodz, Krzysztof Olejnik, Yuhuai Wu, Michal Zawalski, Marek Zbysinski</i>	

VOLUME 2

MCMC Variational Inference Via Uncorrected Hamiltonian Annealing	639
<i>Justin Domke, Tomas Geffner</i>	

Landmark-RxR: Solving Vision-and-Language Navigation with Fine-Grained Alignment Supervision.....	652
<i>Dong An, Keji He, Yan Huang, Shuanglin Sima, Liang Wang, Qi Wu, Jianhua Yang</i>	
A Winning Hand: Compressing Deep Networks Can Improve Out-Of-Distribution Robustness.....	664
<i>Brian Bartoldson, Shreya Chaganti, James Diffenderfer, Bhavya Kailkhura, Jize Zhang</i>	
On the Importance of Gradients for Detecting Distributional Shifts in the Wild	677
<i>Andrew Geng, Rui Huang, Yixuan Li</i>	
Iterative Methods for Private Synthetic Data: Unifying Framework and New Methods.....	690
<i>Terrance Liu, Giuseppe Vietri, Steven Wu</i>	
Understanding End-To-End Model-Based Reinforcement Learning Methods as Implicit Parameterization	703
<i>Clement Gehring, Jiaoyang Huang, Leslie Kaelbling, Kenji Kawaguchi</i>	
Mirror Langevin Monte Carlo: the Case Under Isoperimetry	715
<i>Qijia Jiang</i>	
Do Different Tracking Tasks Require Different Appearance Models?	726
<i>Luca Bertinetto, Ya-Li Li, Philip Torr, Shengjin Wang, Zhongdao Wang, Hengshuang Zhao</i>	
Towards Robust Vision by Multi-Task Learning on Monkey Visual Cortex.....	739
<i>Santiago Cadena, George Denfield, Arne Nix, Kelli Restivo, Shahd Safarani, Fabian Sinz, Andreas Tolas, Konstantin Willeke</i>	
Arbitrary Conditional Distributions with Energy	752
<i>Junier Oliva, Ryan Strauss</i>	
Learning Domain Invariant Representations in Goal-Conditioned Block MDPs.....	764
<i>Jimmy Ba, Harris Chan, Beining Han, Keiran Paster, Michael Zhang, Chongyi Zheng</i>	
Near-Optimal Multi-Perturbation Experimental Design for Causal Structure Learning	777
<i>Andreas Krause, Scott Sussex, Caroline Uhler</i>	
Fuzzy Clustering with Similarity Queries	789
<i>Wasim Huleihel, Arya Mazumdar, Soumyabrata Pal</i>	
Improving Black-Box Optimization in VAE Latent Space Using Decoder Uncertainty	802
<i>Yarin Gal, Jose Miguel Hernandez-Lobato, Pascal Notin</i>	
Sample Selection for Fair and Robust Training.....	815
<i>Kangwook Lee, Yuji Roh, Changho Suh, Steven Whang</i>	
NeurWIN: Neural Whittle Index Network for Restless Bandits Via Deep RL	828
<i>Santosh Ganji, I-Hong Hou, Ping-Chun Hsieh, Khaled Nakhleh, Srinivas Shakkottai</i>	
Sageflow: Robust Federated Learning Against Both Stragglers and Adversaries.....	840
<i>Minseok Choi, Dong-Jun Han, Jaekyun Moon, Jungwuk Park</i>	
Alias-Free Generative Adversarial Networks.....	852
<i>Timo Aila, Miika Aittala, Janne Hellsten, Erik Harkonen, Tero Karras, Samuli Laine, Jaakko Lehtinen</i>	
Noise2Score: Tweedie’s Approach to Self-Supervised Image Denoising Without Clean Images	864
<i>Kwanyoung Kim, Jong Chul Ye</i>	

Continuous Mean-Covariance Bandits	875
<i>Yihan Du, Zhixuan Fang, Longbo Huang, Siwei Wang</i>	
Dynamic Visual Reasoning by Learning Differentiable Physics Models from Video and Language	887
<i>Zhenfang Chen, Mingyu Ding, Tao Du, Chuang Gan, Ping Luo, Josh Tenenbaum</i>	
Solving Soft Clustering Ensemble Via k-Sparse Discrete Wasserstein Barycenter	900
<i>Hu Ding, Mengying Li, Ruizhe Qin</i>	
Bayesian Adaptation for Covariate Shift	914
<i>Sergey Levine, Aurick Zhou</i>	
Perturb-And-max-product: Sampling and Learning in Discrete Energy-based Models	928
<i>Antoine Dedieu, Dileep George, Miguel Lazaro-Gredilla</i>	
Towards Unifying Behavioral and Response Diversity for Open-Ended Learning in Zero-sum Games	941
<i>Yingfeng Chen, Changjie Fan, Zhipeng Hu, Yujing Hu, Hangtian Jia, Xiangyu Liu, Ying Wen, Yaodong Yang</i>	
Towards Better Understanding of Training Certifiably Robust Models Against Adversarial Examples	953
<i>Jaewook Lee, Sungyoon Lee, Woojin Lee, Jinseong Park</i>	
Mitigating Covariate Shift in Imitation Learning Via Offline Data with Partial Coverage	965
<i>Jonathan Chang, Rahul Kidambi, Dhruv Sreenivas, Wen Sun, Masatoshi Uehara</i>	
Global Filter Networks for Image Classification	980
<i>Jiwen Lu, Yongming Rao, Wenliang Zhao, Jie Zhou, Zheng Zhu</i>	
Catastrophic Data Leakage in Vertical Federated Learning	994
<i>Pin-Yu Chen, Tianyi Chen, Chia-Yi Hsu, Xiao Jin, Chia-Mu Yu</i>	
Fault-Tolerant Federated Reinforcement Learning with Theoretical Guarantee	1007
<i>Zhongxiang Dai, Xiaofeng Fan, Wei Jing, Bryan Kian Hsiang Low, Yining Ma, Cheston Tan</i>	
Compacter: Efficient Low-Rank Hypercomplex Adapter Layers	1022
<i>James Henderson, Rabeeh Karimi Mahabadi, Sebastian Ruder</i>	
Distilling Image Classifiers in Object Detectors	1036
<i>Jose M. Alvarez, Shuxuan Guo, Mathieu Salzmann</i>	
Subgroup Generalization and Fairness of Graph Neural Networks	1048
<i>Junwei Deng, Jiaqi Ma, Qiaozhu Mei</i>	
Scaling Neural Tangent Kernels Via Sketching and Random Features	1062
<i>Haim Avron, Insu Han, Chaewon Kim, Jinwoo Shin, Neta Shoham, Amir Zandieh</i>	
BatchQuant: Quantized-For-all Architecture Search with Robust Quantizer	1074
<i>Haoping Bai, Meng Cao, Ping Huang, Jiulong Shan</i>	
Long Short-Term Transformer for Online Action Detection	1086
<i>Hao Chen, Xinyu Li, Stefano Soatto, Zhuowen Tu, Wei Xia, Yuanjun Xiong, Mingze Xu</i>	
Near Optimal Policy Optimization Via REPS	1100
<i>Peter Bartlett, Jonathan N Lee, Ofir Nachum, Aldo Pacchiano</i>	

Self-Consistent Models and Values.....	1111
<i>Kate Baumli, Greg Farquhar, Angelos Filos, Matteo Hessel, Zita Marinho, David Silver, Hado Van Hasselt</i>	
Learning on Random Balls is Sufficient for Estimating (Some) Graph Parameters.....	1126
<i>Takanori Maehara, Hoang Nt</i>	
Risk-Averse Bayes-Adaptive Reinforcement Learning.....	1142
<i>Nick Hawes, Bruno Lacerda, Marc Rigter</i>	
Iterative Connecting Probability Estimation for Networks	1155
<i>Yang Li, Yichen Qin, Linhan Yu</i>	
Learning to Adapt Via Latent Domains for Adaptive Semantic Segmentation.....	1167
<i>Yang Li, Yunan Liu, Jian Yang, Shanshan Zhang</i>	
Single Layer Predictive Normalized Maximum Likelihood for Out-Of-Distribution Detection.....	1179
<i>Koby Bibas, Meir Feder, Tal Hassner</i>	
Prototypical Cross-Attention Networks for Multiple Object Tracking and Segmentation	1192
<i>Martin Danelljan, Lei Ke, Xia Li, Yu-Wing Tai, Chi-Keung Tang, Fisher Yu</i>	
Algorithmic Instabilities of Accelerated Gradient Descent	1204
<i>Amit Attia, Tomer Koren</i>	
Learning Optimal Predictive Checklists	1215
<i>Marzyeh Ghassemi, Quaid Morris, Berk Ustun, Haoran Zhang</i>	
Finite Sample Analysis of Average-Reward TD Learning and Q-Learning.....	1230
<i>Siva Theja Maguluri, Sheng Zhang, Zhe Zhang</i>	
Generalization Bounds for Graph Embedding Using Negative Sampling: Linear Vs Hyperbolic.....	1243
<i>Marc Cavazza, Atsushi Nitanda, Atsushi Suzuki, Linchuan Xu, Kenji Yamanishi, Jing Wang</i>	
Gradient Starvation: a Learning Proclivity in Neural Networks	1256
<i>Yoshua Bengio, Aaron Courville, Oumar Kaba, Guillaume Lajoie, Mohammad Pezeshki, Doina Precup</i>	
Offline Reinforcement Learning as One Big Sequence Modeling Problem.....	1273
<i>Michael Janner, Sergey Levine, Qiyang Li</i>	
Optimality and Stability in Federated Learning: a Game-Theoretic Approach.....	1287
<i>Kate Donahue, Jon Kleinberg</i>	
Understanding Deflation Process in Over-Parametrized Tensor Decomposition	1299
<i>Rong Ge, Yunwei Ren, Xiang Wang, Mo Zhou</i>	
Privately Learning Subspaces.....	1312
<i>Vikrant Singhal, Thomas Steinke</i>	
On the Value of Interaction and Function Approximation in Imitation Learning	1325
<i>Yanjun Han, Jiantao Jiao, Jingbo Liu, Nived Rajaraman, Kannan Ramchandran, Lin Yang</i>	
Shapeshifter: a Parameter-Efficient Transformer Using Factorized Reshaped Matrices.....	1337
<i>Tom Arodz, Aliakbar Panahi, Seyran Saeedi</i>	

The Adaptive Doubly Robust Estimator and a Paradox Concerning Logging Policy	1351
<i>Masahiro Kato, Kenichiro McAlinn, Shota Yasui</i>	
Regularized Softmax Deep Multi-Agent Q-Learning	1365
<i>Longbo Huang, Ling Pan, Bei Peng, Tabish Rashid, Shimon Whiteson</i>	
Physics-Aware Downsampling with Deep Learning for Scalable Flood Modeling	1378
<i>Zvika Ben-Haim, Niv Giladi, Yossi Matias, Sella Nevo, Daniel Soudry</i>	
Systematic Generalization with Edge Transformers	1390
<i>Dzmitry Bahdanau, Leon Bergen, Timothy O'Donnell</i>	
TransformerFusion: Monocular RGB Scene Reconstruction Using Transformers	1403
<i>Aljaz Bozic, Angela Dai, Matthias Niessner, Pablo Palafox, Justus Thies</i>	
Maximum Likelihood Training of Score-Based Diffusion Models.....	1415
<i>Conor Durkan, Stefano Ermon, Iain Murray, Yang Song</i>	
Global Convergence of Gradient Descent for Asymmetric Low-Rank Matrix Factorization	1429
<i>Simon Du, Tian Ye</i>	
Adaptive Data Augmentation on Temporal Graphs	1440
<i>Siddharth Bhatia, Yujun Cai, Henghui Ding, Bryan Hooi, Yuxuan Liang, Changhu Wang, Yiwei Wang</i>	
Regularized Frank-Wolfe for Dense CRFs: Generalizing Mean Field and Beyond.....	1453
<i>Kartteek Alahari, Đ. Khuê Lê-Huu</i>	
Terra: Imperative-Symbolic Co-Execution of Imperative Deep Learning Programs	1468
<i>Byung-Gon Chun, Eunji Jeong, Geon-Woo Kim, Sehoon Kim, Taebum Kim, Yunmo Koo, Gyeongin Yu</i>	
Uniform Sampling Over Episode Difficulty.....	1481
<i>Sebastien Arnold, Guneet Dhillon, Avinash Ravichandran, Stefano Soatto</i>	

VOLUME 3

Scalable Intervention Target Estimation in Linear Models	1494
<i>Prasanna Sattigeri, Karthikeyan Shanmugam, Ali Tajer, Burak Varici</i>	
Play to Grade: Testing Coding Games as Classifying Markov Decision Process	1506
<i>Emma Brunskill, Allen Nie, Chris Piech</i>	
Distributional Reinforcement Learning for Multi-Dimensional Reward Functions.....	1519
<i>Xiaoyu Chen, Tie-Yan Liu, Tao Qin, Wei Xiong, Pushi Zhang, Li Zhao</i>	
Differentiable Unsupervised Feature Selection Based on a Gated Laplacian.....	1530
<i>Nicolas Casey, Yuval Kluger, Ofir Lindenbaum, Erez Peterfreund, Uri Shaham, Jonathan Svirsky</i>	
Smooth Bilevel Programming for Sparse Regularization.....	1543
<i>Gabriel Peyre, Clarice Poon</i>	
Grounding Representation Similarity Through Statistical Testing.....	1556
<i>Jean-Stanislas Denain, Frances Ding, Jacob Steinhardt</i>	

A Consciousness-Inspired Planning Agent for Model-Based Reinforcement Learning.....	1569
<i>Yoshua Bengio, Zhen Liu, Sitao Luan, Doina Precup, Shuyuan Zhang, Mingde Zhao</i>	
Reward-Free Model-Based Reinforcement Learning with Linear Function Approximation	1582
<i>Quanquan Gu, Weitong Zhang, Dongruo Zhou</i>	
Beltrami Flow and Neural Diffusion on Graphs.....	1594
<i>Michael Bronstein, Benjamin Chamberlain, Francesco Di Giovanni, Xiaowen Dong, Davide Eynard, James Rowbottom</i>	
Think Big, Teach Small: Do Language Models Distil Occam’s Razor?	1610
<i>David Castellano Falcon, Cesar Ferri, Jose Hernandez-Orallo, Gonzalo Jaimovitch-Lopez</i>	
Disentangling Identifiable Features from Noisy Data with Structured Nonlinear ICA	1624
<i>Elisabeth Gassiat, Aapo Hyvarinen, Hermanni Halva, Sylvain Le Corff, Luc Lehericy, Jonathan So, Yongjie Zhu</i>	
Conditionally Parameterized, Discretization-Aware Neural Networks for Mesh-Based Modeling of Physical Systems	1634
<i>Karthikeyan Duraisamy, Aniruddhe Pradhan, Jiayang Xu</i>	
USCO-Solver: Solving Undetermined Stochastic Combinatorial Optimization Problems	1646
<i>Guangmo Tong</i>	
Adaptive Conformal Inference Under Distribution Shift	1660
<i>Emmanuel Candes, Isaac Gibbs</i>	
Periodic Activation Functions Induce Stationarity	1673
<i>Lassi Meronen, Arno Solin, Martin Trapp</i>	
Towards Optimal Strategies for Training Self-Driving Perception Models in Simulation.....	1686
<i>David Acuna, Sanja Fidler, Jonah Philion</i>	
KS-GNN: Keywords Search Over Incomplete Graphs Via Graphs Neural Network.....	1700
<i>Xin Cao, Yixiang Fang, Yu Hao, Yufan Sheng, Wei Wang</i>	
Reconstruction for Powerful Graph Representations	1713
<i>Leonardo Cotta, Christopher Morris, Bruno Ribeiro</i>	
Revealing and Protecting Labels in Distributed Training.....	1727
<i>Francoise Beaufays, Peter Chin, Trung Dang, Rajiv Mathews, Swaroop Ramaswamy, Om Thakkar</i>	
Solving Graph-Based Public Goods Games with Tree Search and Imitation Learning	1739
<i>Victor-Alexandru Darvari, Stephen Hailes, Mirco Musolesi</i>	
Stochastic Optimization of Areas Under Precision-Recall Curves with Provable Convergence.....	1752
<i>Shuiwang Ji, Youzhi Luo, Qi Qi, Zhao Xu, Tianbao Yang</i>	
Transfer Learning of Graph Neural Networks with Ego-Graph Information Maximization.....	1766
<i>Jiawei Han, Haonan Wang, Yidan Xu, Carl Yang, Chao Zhang, Qi Zhu</i>	
You Are Caught Stealing My Winning Lottery Ticket! Making a Lottery Ticket Claim Its Ownership	1780
<i>Tianlong Chen, Xuxi Chen, Zhangyang Wang, Zhenyu Zhang</i>	

Complexity Lower Bounds for Nonconvex-Strongly-Concave Min-Max Optimization	1792
<i>Ali Jadbabaie, Haochuan Li, Yi Tian, Jingzhao Zhang</i>	
Early-Stopped Neural Networks Are Consistent	1805
<i>Ziwei Ji, Justin Li, Matus Telgarsky</i>	
NxMTransformer: Semi-Structured Sparsification for Natural Language Understanding Via ADMM	1818
<i>Yuxiong He, Connor Holmes, Bo Wu, Minjia Zhang</i>	
Reliable Decisions with Threshold Calibration	1831
<i>Alyssa Chen, Stefano Ermon, Roshni Sahoo, Shengjia Zhao</i>	
End-To-End Weak Supervision	1845
<i>Benedikt Boecking, Artur Dubrawski, Salva Ruhling Cachay</i>	
Shift Invariance Can Reduce Adversarial Robustness	1858
<i>Songwei Ge, David Jacobs, Basri Ronen, Vasu Singla</i>	
Wisdom of the Crowd Voting: Truthful Aggregation of Voter Information and Preferences	1872
<i>Grant Schoenebeck, Biaoshuai Tao</i>	
Replay-Guided Adversarial Environment Design	1884
<i>Michael Dennis, Jakob Foerster, Edward Grefenstette, Minqi Jiang, Jack Parker-Holder, Tim Rocktaschel</i>	
There is No Turning Back: a Self-Supervised Approach for Reversibility-Aware Reinforcement Learning	1898
<i>Johan Ferret, Matthieu Geist, Nathan Grinsztajn, Olivier Pietquin, Philippe Preux</i>	
Learning to Execute: Efficient Learning of Universal Plan-Conditioned Policies in Robotics	1912
<i>Danny Driess, Ozgur S. Oguz, Ingmar Schubert, Marc Toussaint</i>	
Self-Diagnosing GAN: Diagnosing Underrepresented Samples in Generative Adversarial Networks	1925
<i>Hye Won Chung, Youngkyu Hong, Haeri Kim, Jinhee Lee</i>	
Online Multi-Armed Bandits with Adaptive Inference	1939
<i>Maria Dimakopoulou, Zhimei Ren, Zhengyuan Zhou</i>	
Efficient Truncated Linear Regression with Unknown Noise Variance	1952
<i>Constantinos Daskalakis, Patroklos Stefanou, Rui Yao, Emmanouil Zampetakis</i>	
Breaking the Dilemma of Medical Image-To-image Translation	1964
<i>Yanle Hu, Detian Huang, Lingke Kong, Chenyu Lian, Qichao Zhou, Zhenjiang Li</i>	
Temporally Abstract Partial Models	1979
<i>Zafarali Ahmed, Gheorghe Comanici, Khimya Khetarpal, Doina Precup</i>	
TransMatcher: Deep Image Matching Through Transformers for Generalizable Person Re- Identification	1992
<i>Shengcai Liao, Ling Shao</i>	
Multi-Objective SPIBB: Seldonian Offline Policy Improvement with Safety Constraints in Finite MDPs	2004
<i>Romain Laroché, Joelle Pineau, Philip Thomas, Harsh Satija</i>	

Is Automated Topic Model Evaluation Broken? the Incoherence of Coherence.....	2018
<i>Jordan Boyd-Graber, Pranav Goel, Andrew Hian-Cheong, Alexander Hoyle, Denis Peskov, Philip Resnik</i>	
INDIGO: GNN-Based Inductive Knowledge Graph Completion Using Pair-Wise Encoding.....	2034
<i>Bernardo Grau, Ian Horrocks, Egor Kostylev, Shuwen Liu</i>	
Do Input Gradients Highlight Discriminative Features?	2046
<i>Prateek Jain, Praneeth Netrapalli, Harshay Shah</i>	
Improving Conditional Coverage Via Orthogonal Quantile Regression	2060
<i>Stephen Bates, Shai Feldman, Yaniv Romano</i>	
Minimizing Polarization and Disagreement in Social Networks Via Link Recommendation.....	2072
<i>Qi Bao, Zhongzhi Zhang, Liwang Zhu</i>	
Adversarial Attacks on Black Box Video Classifiers: Leveraging the Power of Geometric Transformations.....	2085
<i>Abhishek Aich, Salman Asif, Srikanth Krishnamurthy, Shasha Li, Amit Roy-Chowdhury, Chengyu Song, Shitong Zhu</i>	
Optimal Rates for Random Order Online Optimization	2097
<i>Tomer Koren, Yishay Mansour, Uri Sherman</i>	
Discrete-Valued Neural Communication	2109
<i>Anirudh Goyal Alias Parth Goyal, Yoshua Bengio, Kenji Kawaguchi, Alex Lamb, Dianbo Liu, Michael Mozer, Chen Sun</i>	
Skyformer: Remodel Self-Attention with Gaussian Kernel and Nyström Method.....	2122
<i>Yifan Chen, Heng Ji, Yun Yang, Qi Zeng</i>	
TransMIL: Transformer Based Correlated Multiple Instance Learning for Whole Slide Image Classification	2136
<i>Hao Bian, Yang Chen, Xiangyang Ji, Zhuchen Shao, Yifeng Wang, Jian Zhang, Yongbing Zhang</i>	
Multi-View Contrastive Graph Clustering	2148
<i>Zhao Kang, Erlin Pan</i>	
Inverse-Weighted Survival Games.....	2160
<i>Mark Goldstein, Xintian Han, Adler Perotte, Aahlad Puli, Rajesh Ranganath, Thomas Wies</i>	
Generalization Bounds for Meta-Learning Via PAC-Bayes and Uniform Stability.....	2173
<i>Alec Farid, Anirudha Majumdar</i>	
Parallel Bayesian Optimization of Multiple Noisy Objectives with Expected Hypervolume Improvement	2187
<i>Eytan Bakshy, Maximilian Balandat, Samuel Daulton</i>	
Evolution Gym: a Large-Scale Benchmark for Evolving Soft Robots.....	2201
<i>Jagdeep Bhatia, Holly Jackson, Wojciech Matusik, Yunsheng Tian, Jie Xu</i>	
On Calibration and Out-Of-Domain Generalization	2215
<i>Amir Feder, Daniel Greenfeld, Uri Shalit, Yoav Wald</i>	
On the Convergence and Sample Efficiency of Variance-Reduced Policy Gradient Method.....	2228
<i>Chengzhuo Ni, Csaba Szepesvari, Mengdi Wang, Zheng Yu, Junyu Zhang</i>	

Circa: Stochastic ReLUs for Private Deep Learning	2241
<i>Siddharth Garg, Zahra Ghodsi, Nandan Kumar Jha, Brandon Reagen</i>	
Reinforcement Learning in Reward-Mixing MDPs	2253
<i>Constantine Caramanis, Yonathan Efroni, Jeongyeol Kwon, Shie Mannor</i>	
A Gang of Adversarial Bandits	2265
<i>Mark Herbster, Stephen Pasteris, Massimiliano Pontil, Fabio Vitale</i>	
Explaining Hyperparameter Optimization Via Partial Dependence Plots	2280
<i>Bernd Bischl, Giuseppe Casalicchio, Julia Herbinger, Marius Lindauer, Julia Moosbauer</i>	
Robustifying Algorithms of Learning Latent Trees with Vector Variables	2292
<i>Vincent Tan, Fengzhuo Zhang</i>	
Representation Learning on Spatial Networks	2303
<i>Zheng Zhang, Liang Zhao</i>	
Continuous-Time Edge Modelling Using Non-Parametric Point Processes.....	2319
<i>Xuhui Fan, Bin Li, Scott Sisson, Feng Zhou</i>	
Deep Inference of Latent Dynamics with Spatio-Temporal Super-resolution Using Selective Backpropagation Through Time.....	2331
<i>Nauman Ahad, Mark Davenport, Andrea Giovannucci, Harrison A Grier, Matthew Kaufman, Chethan Pandarinath, Andrew Sedler, Feng Zhu</i>	

VOLUME 4

Memory-Efficient Patch-based Inference for Tiny Deep Learning	2346
<i>Han Cai, Wei-Ming Chen, Chuang Gan, Song Han, Ji Lin</i>	
Self-Interpretable Model with Transformation Equivariant Interpretation.....	2359
<i>Xiaoqian Wang, Yipei Wang</i>	
Solving Min-Max Optimization with Hidden Structure Via Gradient Descent Ascent.....	2373
<i>Lampros Flokas, Georgios Piliouras, Emmanouil-Vasileios Vlatakis-Gkaragkounis</i>	
Preserved Central Model for Faster Bidirectional Compression in Distributed Settings	2387
<i>Aymeric Dieuleveut, Constantin Philippenko</i>	
Understanding Instance-Based Interpretability of Variational Auto-Encoders	2400
<i>Kamalika Chaudhuri, Zhifeng Kong</i>	
Voxel-Based 3D Detection and Reconstruction of Multiple Objects from a Single Image.....	2413
<i>Feng Liu, Xiaoming Liu</i>	
Test-Time Classifier Adjustment Module for Model-Agnostic Domain Generalization	2427
<i>Yusuke Iwasawa, Yutaka Matsuo</i>	
Luna: Linear Unified Nested Attention	2441
<i>Xiang Kong, Hao Ma, Xuezhe Ma, Jonathan May, Sinong Wang, Luke Zettlemoyer, Chunting Zhou</i>	
Iterative Causal Discovery in the Possible Presence of Latent Confounders and Selection Bias.....	2454
<i>Yaniv Gurwicz, Shami Nisimov, Gal Novik, Raanan Rohekar</i>	

Hindsight Task Relabelling: Experience Replay for Sparse Reward Meta-RL.....	2466
<i>Pieter Abbeel, Joseph Gonzalez, Charles Packer</i>	
A Bayesian-Symbolic Approach to Reasoning and Learning in Intuitive Physics.....	2478
<i>Dan Gutfreund, Felix Sosa, Akash Srivastava, Charles Sutton, Josh Tenenbaum, Tomer Ullman, Kai Xu</i>	
Associating Objects with Transformers for Video Object Segmentation.....	2491
<i>Yunchao Wei, Yi Yang, Zongxin Yang</i>	
Automatic Symmetry Discovery with Lie Algebra Convolutional Network.....	2503
<i>Nima Dehmamy, Yan Chen Liu, Robin Walters, Dashun Wang, Rose Yu</i>	
Zero Time Waste: Recycling Predictions in Early Exit Neural Networks.....	2516
<i>Klaudia Balazy, Igor T Podolak, Jacek Tabor, Tomasz Trzcinski, Maciej Wolczyk, Bartosz Wojcik, Marek Smieja</i>	
On Model Calibration for Long-Tailed Object Detection and Instance Segmentation.....	2529
<i>Soravit Changpinyo, Wei-Lun Chao, Boqing Gong, Hexiang Hu, Yandong Li, Tai-Yu Pan, Dong Xuan, Cheng Zhang</i>	
ReSSL: Relational Self-Supervised Learning with Weak Augmentation.....	2543
<i>Chen Qian, Xiaogang Wang, Fei Wang, Chang Xu, Shan You, Changshui Zhang, Mingkai Zheng</i>	
Learning to See by Looking at Noise	2556
<i>Manel Baradad Jurjo, Phillip Isola, Antonio Torralba, Tongzhou Wang, Jonas Wulff</i>	
Explicit Loss Asymptotics in the Gradient Descent Training of Neural Networks.....	2570
<i>Maksim Velikanov, Dmitry Yarotsky</i>	
Test-Time Personalization with a Transformer for Human Pose Estimation	2583
<i>Zonglin Di, Nitesh Bharadwaj Gundavarapu, Miao Hao, Yizhuo Li, Xiaolong Wang</i>	
Towards Scalable Unpaired Virtual Try-On Via Patch-Routed Spatially-Adaptive GAN.....	2598
<i>Haoye Dong, Zaiyu Huang, Michael Kampffmeyer, Xiaodan Liang, Zhenyu Xie, Fuwei Zhao</i>	
Bias Out-Of-the-Box: an Empirical Analysis of Intersectional Occupational Biases in Popular Generative Language Models.....	2611
<i>Yuki Asano, Elias Benussi, Frederic Dreyer, Haider Iqbal, Hannah Rose Kirk, Aleksandar Shtedritski, Filippo Volpin, Yennie Jun</i>	
Weisfeiler and Lehman Go Cellular: CW Networks	2625
<i>Cristian Bodnar, Michael Bronstein, Fabrizio Frasca, Pietro Liò, Guido Montufar, Nina Otter, Yuguang Wang</i>	
Learning Conjoint Attentions for Graph Neural Nets.....	2641
<i>L Bai, Tiantian He, Yew Soon Ong</i>	
Hybrid Regret Bounds for Combinatorial Semi-Bandits and Adversarial Linear Bandits	2654
<i>Shinji Ito</i>	
Pay Better Attention to Attention: Head Selection in Multilingual and Multi-Domain Sequence Modeling	2668
<i>Hongyu Gong, Xian Li, Juan Pino, Yun Tang</i>	

Cardinality-Regularized Hawkes-Granger Model.....	2682
<i>Naoki Abe, Tsuyoshi Ide, Georgios Kollias, Dzung Phan</i>	
Aligned Structured Sparsity Learning for Efficient Image Super-Resolution	2695
<i>Yun Fu, Can Qin, Huan Wang, Yulun Zhang</i>	
Why Lottery Ticket Wins? a Theoretical Perspective of Sample Complexity on Sparse Neural Networks	2707
<i>Pin-Yu Chen, Sijia Liu, Meng Wang, Jinjun Xiong, Shuai Zhang</i>	
Constrained Robust Submodular Partitioning	2721
<i>Jeff A Bilmes, Chandrashekar Lavania, Shengjie Wang, Tianyi Zhou</i>	
Online Knapsack with Frequency Predictions.....	2733
<i>Sungjin Im, Ravi Kumar, Mahshid Montazer Qaem, Manish Purohit</i>	
On Component Interactions in Two-Stage Recommender Systems.....	2744
<i>Jiri Hron, Michael Jordan, Niki Kilbertus, Karl Krauth</i>	
Lip to Speech Synthesis with Visual Context Attentional GAN	2758
<i>Joanna Hong, Minsu Kim, Yong Man Ro</i>	
Non-Convex Distributionally Robust Optimization: Non-asymptotic Analysis	2771
<i>Jikai Jin, Liwei Wang, Haiyang Wang, Bohang Zhang</i>	
Goal-Aware Cross-Entropy for Multi-Target Reinforcement Learning	2783
<i>Kibeom Kim, Yoosung Kim, Min Whoo Lee, Minsu Lee, Jehwan Ryu, Byoung-Tak Zhang</i>	
Smooth Normalizing Flows.....	2796
<i>Andreas Kramer, Jonas Kohler, Frank Noe</i>	
MetaAvatar: Learning Animatable Clothed Human Models from Few Depth Images	2810
<i>Andreas Geiger, Qianli Ma, Marko Mihajlovic, Siyu Tang, Shaofei Wang</i>	
Distributed Principal Component Analysis with Limited Communication	2823
<i>Foivos Alimisis, Dan Alistarh, Peter Davies, Bart Vandereycken</i>	
Newton-LESS: Sparsification Without Trade-offs for the Sketched Newton Update.....	2835
<i>Michal Dereziński, Jonathan Lacotte, Michael Mahoney, Mert Pilanci</i>	
Confident Anchor-Induced Multi-Source Free Domain Adaptation	2848
<i>Jiahua Dong, Zhen Fang, Anjin Liu, Tongliang Liu, Gan Sun</i>	
Word2Fun: Modelling Words as Functions for Diachronic Word Representation	2861
<i>Emanuele Di Buccio, Massimo Melucci, Benyou Wang</i>	
Iteratively Reweighted Least Squares for Basis Pursuit with Global Linear Convergence Rate.....	2873
<i>Christian Kummerle, Claudio Mayrink Verdun, Dominik Stoger</i>	
Low-Rank Constraints for Fast Inference in Structured Models	2887
<i>Justin Chiu, Yuntian Deng, Alexander Rush</i>	
Accumulative Poisoning Attacks on Real-Time Data	2899
<i>Yinpeng Dong, Tianyu Pang, Hang Su, Xiao Yang, Jun Zhu</i>	
UCB-Based Algorithms for Multinomial Logistic Regression Bandits	2913
<i>Sanae Amani, Christos Thrampoulidis</i>	

Estimating the Long-Term Effects of Novel Treatments	2925
<i>Keith Battocchi, Eleanor Dillon, Maggie Hei, Greg Lewis, Miruna Oprescu, Vasilis Syrkanis</i>	
Dual Progressive Prototype Network for Generalized Zero-Shot Learning	2936
<i>Xuejin Chen, Houqiang Li, Shaobo Min, Xiaoyan Sun, Chaoqun Wang</i>	
Derivative-Free Policy Optimization for Linear Risk-Sensitive and Robust Control Design: Implicit Regularization and Sample Complexity	2949
<i>Tamer Basar, Bin Hu, Kaiqing Zhang, Xiangyuan Zhang</i>	
G-PATE: Scalable Differentially Private Data Generator Via Private Aggregation of Teacher Discriminators	2965
<i>Carl Gunter, Bhavya Kailkhura, Bo Li, Yunhui Long, Boxin Wang, Zhuolin Yang, Aston Zhang</i>	
On the Existence of the Adversarial Bayes Classifier	2978
<i>Pranjal Awasthi, Natalie Frank, Mehryar Mohri</i>	
Convex-Concave Min-Max Stackelberg Games	2991
<i>Denizalp Goktas, Amy Greenwald</i>	
Misspecified Gaussian Process Bandit Optimization	3004
<i>Ilija Bogunovic, Andreas Krause</i>	
Visual Adversarial Imitation Learning Using Variational Models	3016
<i>Chelsea Finn, Rafael Rafailov, Aravind Rajeswaran, Tianhe Yu</i>	
Object-Aware Regularization for Addressing Causal Confusion in Imitation Learning	3029
<i>Chang Liu, Tie-Yan Liu, Jongjin Park, Tao Qin, Younggyo Seo, Jinwoo Shin, Li Zhao</i>	
Reliable and Trustworthy Machine Learning for Health Using Dataset Shift Detection	3043
<i>Anas Awadalla, Tadayoshi Kohno, Chunjong Park, Shwetak Patel</i>	
Multiclass Boosting and the Cost of Weak Learning	3057
<i>Nataly Brukhim, Elad Hazan, Shay Moran, Indraneel Mukherjee, Robert Schapire</i>	
Partition-Based Formulations for Mixed-Integer Optimization of Trained ReLU Neural Networks	3068
<i>Jan Kronqvist, Ruth Misener, Alexander Thebelt, Calvin Tsay</i>	
Hyperparameter Optimization is Deceiving Us, and How to Stop it.....	3081
<i>A. Feder Cooper, Christopher De Sa, Jessica Forde, Yucheng Lu</i>	
On the Convergence Theory of Debiased Model-Agnostic Meta-Reinforcement Learning	3096
<i>Alireza Fallah, Kristian Georgiev, Aryan Mokhtari, Asuman Ozdaglar</i>	
3D Pose Transfer with Correspondence Learning and Mesh Refinement	3108
<i>Ruibo Li, Guosheng Lin, Fayao Liu, Chaoyue Song, Jiacheng Wei</i>	
Framing RNN as a Kernel Method: a Neural ODE Approach	3121
<i>Gerard Biau, Adeline Fermanian, Pierre Marion, Jean-Philippe Vert</i>	
Contextual Similarity Aggregation with Self-Attention for Visual Re-ranking	3135
<i>Houqiang Li, Jianbo Ouyang, Min Wang, Hui Wu, Wengang Zhou</i>	
Can Information Flows Suggest Targets for Interventions in Neural Circuits?.....	3149
<i>Sanghamitra Dutta, Pulkit Grover, Neil Mehta, Praveen Venkatesh</i>	

AutoBalance: Optimized Loss Functions for Imbalanced Data.....	3163
<i>Jiasi Chen, Mingchen Li, Samet Oymak, Christos Thrampoulidis, Xuechen Zhang</i>	
SyncTwin: Treatment Effect Estimation with Longitudinal Outcomes.....	3178
<i>Ioana Bica, Zhaozhi Qian, Angela Wood, Yao Zhang, Mihaela Van Der Schaar</i>	
Statistical Query Lower Bounds for List-Decodable Linear Regression.....	3191
<i>Ilias Diakonikolas, Daniel Kane, Ankit Pensia, Thanasis Pittas, Alistair Stewart</i>	

VOLUME 5

Unsupervised Motion Representation Learning with Capsule Autoencoders	3205
<i>Mohan Kankanhalli, Xudong Shen, Yongkang Wong, Ziwei Xu</i>	
VigDet: Knowledge Informed Neural Temporal Point Process for Coordination Detection on Social Media.....	3218
<i>Yan Liu, Karishma Sharma, Yizhou Zhang</i>	
An Improved Analysis and Rates for Variance Reduction Under Without-Replacement Sampling Orders	3232
<i>Xinmeng Huang, Xianghui Mao, Wotao Yin, Kun Yuan</i>	
Exploring Forensic Dental Identification with Deep Learning.....	3244
<i>Weikun Han, Lei He, Yuan Liang, Liang Qiu, Yiting Shao, Kun Wang, Chen Wu</i>	
Learning to Generate Realistic Noisy Images Via Pixel-Level Noise-aware Adversarial Training.....	3259
<i>Yuanhao Cai, Xiaowan Hu, Hanspeter Pfister, Haoqian Wang, Donglai Wei, Yulun Zhang</i>	
Multi-Agent Reinforcement Learning for Active Voltage Control on Power Distribution Networks.....	3271
<i>Tim C Green, Yunjie Gu, Wenbin Song, Jianhong Wang, Wangkun Xu</i>	
Looking Beyond Single Images for Contrastive Semantic Segmentation Learning.....	3285
<i>Rene Ranftl, Stephan Richter, Philip Torr, Feihu Zhang</i>	
A Constant Approximation Algorithm for Sequential Random-Order No-Substitution k-Median Clustering	3298
<i>Tom Hess, Michal Moshkovitz, Sivan Sabato</i>	
Dangers of Bayesian Model Averaging Under Covariate Shift.....	3309
<i>Pavel Izmailov, Sanae Lotfi, Patrick Nicholson, Andrew Wilson</i>	
Learning Equilibria in Matching Markets from Bandit Feedback.....	3323
<i>Meena Jagadeesan, Michael Jordan, Jacob Steinhardt, Yixin Wang, Alexander Wei</i>	
Towards Lower Bounds on the Depth of ReLU Neural Networks.....	3336
<i>Amitabh Basu, Marco Di Summa, Christoph Hertrich, Martin Skutella</i>	
The Limitations of Large Width in Neural Networks: a Deep Gaussian Process Perspective	3349
<i>John Cunningham, Geoff Pleiss</i>	
Exact Marginal Prior Distributions of Finite Bayesian Neural Networks	3364
<i>Cengiz Pehlevan, Jacob Zavatone-Veth</i>	
Spatiotemporal Joint Filter Decomposition in 3D Convolutional Neural Networks	3376
<i>Xiuyuan Cheng, Zichen Miao, Qiang Qiu, Ze Wang</i>	

Pooling by Sliced-Wasserstein Embedding.....	3389
<i>Reed Andrews, Joseph F Comer, Heiko Hoffmann, Soheil Kolouri, Navid Naderializadeh</i>	
On the Theory of Reinforcement Learning with Once-Per-Episode Feedback	3401
<i>Peter Bartlett, Niladri Chatterji, Michael Jordan, Aldo Pacchiano</i>	
ResNEsts and DenseNEsts: Block-Based DNN Models with Improved Representation Guarantees	3413
<i>Kuan-Lin Chen, Harinath Garudadri, Ching-Hua Lee, Bhaskar D Rao</i>	
Locally Private Online Change Point Detection.....	3425
<i>Tom Berrett, Yi Yu</i>	
Invariance Principle Meets Information Bottleneck for Out-Of-Distribution Generalization	3438
<i>Kartik Ahuja, Yoshua Bengio, Ethan Caballero, Jean-Christophe Gagnon-Audet, Ioannis Mitliagkas, Irina Rish, Dinghui Zhang</i>	
Repulsive Deep Ensembles Are Bayesian.....	3451
<i>Francesco D'Angelo, Vincent Fortuin</i>	
BayesIMP: Uncertainty Quantification for Causal Data Fusion.....	3466
<i>Siu Lun Chau, Javier Gonzalez, Dino Sejdinovic, Yee Teh, Jean-Francois Ton</i>	
RMM: Reinforced Memory Management for Class-Incremental Learning	3478
<i>Yaoyao Liu, Bernt Schiele, Qianru Sun</i>	
Learning Compact Representations of Neural Networks Using Discriminative Masking (DAM)	3491
<i>Jie Bu, Arka Daw, Anuj Karpatne, M. Maruf</i>	
Neural Auto-Curricula in Two-Player Zero-Sum Games	3504
<i>Xidong Feng, Bo Liu, Stephen McAleer, Oliver Slumbers, Ziyu Wan, Jun Wang, Ying Wen, Yaodong Yang</i>	
ImageBART: Bidirectional Context with Multinomial Diffusion for Autoregressive Image Synthesis.....	3518
<i>Andreas Blattmann, Patrick Esser, Bjorn Ommer, Robin Rombach</i>	
From Global to Local MDI Variable Importances for Random Forests and When They Are Shapley Values.....	3533
<i>Pierre Geurts, Van Anh Huynh-Thu, Gilles Louppe, Antonio Suter, Louis Wehenkel</i>	
Adversarial Robustness of Streaming Algorithms Through Importance Sampling.....	3544
<i>Vladimir Braverman, Avinatan Hassidim, Yossi Matias, Mariano Schain, Sandeep Silwal, Samson Zhou</i>	
Tractable Regularization of Probabilistic Circuits.....	3558
<i>Anji Liu, Guy Van Den Broeck</i>	
On Interaction Between Augmentations and Corruptions in Natural Corruption Robustness.....	3571
<i>Alexander Kirillov, Eric Mintun, Saining Xie</i>	
Dynamic Distillation Network for Cross-Domain Few-Shot Recognition with Unlabeled Data	3584
<i>Chun-Fu (Richard) Chen, Rogerio Feris, Ashraf Islam, Leonid Karlinsky, Rameswar Panda, Richard J. Radke</i>	
Hypergraph Propagation and Community Selection for Objects Retrieval	3596
<i>Guoyuan An, Yuchi Huo, Sung-Eui Yoon</i>	

Deep Learning is Adaptive to Intrinsic Dimensionality of Model Smoothness in Anisotropic Besov Space	3609
<i>Atsushi Nitanda, Taiji Suzuki</i>	
QuPeD: Quantized Personalization Via Distillation with Applications to Federated Learning	3622
<i>Deepesh Data, Suhas Diggavi, Kaan Ozkara, Navjot Singh</i>	
Model Adaptation: Historical Contrastive Learning for Unsupervised Domain Adaptation Without Source Data	3635
<i>Dayan Guan, Jiaxing Huang, Shijian Lu, Aoran Xiao</i>	
The Out-Of-Distribution Problem in Explainability and Search Methods for Feature Importance Explanations	3650
<i>Mohit Bansal, Peter Hase, Harry Xie</i>	
Control Variates for Slate Off-Policy Evaluation	3667
<i>Fernando Amat, Ashok Chandrashekar, Nathan Kallus, Nikos Vlassis</i>	
Stabilizing Deep Q-Learning with ConvNets and Vision Transformers Under Data Augmentation	3680
<i>Nicklas Hansen, Hao Su, Xiaolong Wang</i>	
On Effective Scheduling of Model-Based Reinforcement Learning	3694
<i>Yimin Huang, Hang Lai, Zhenguo Li, Jian Shen, Ruiming Tang, Yong Yu, Weinan Zhang, Xing Zhang</i>	
Removing Inter-Experimental Variability from Functional Data in Systems Neuroscience	3706
<i>Philipp Berens, Benjamin Dunn, Thomas Euler, Katrin Franke, Dominic Gonschorek, Larissa Hofling, David Klindt, Timm Schubert, Klaudia P. Szatko</i>	
Learning Knowledge Graph-Based World Models of Textual Environments	3720
<i>Prithviraj Ammanabrolu, Mark Riedl</i>	
Damped Anderson Mixing for Deep Reinforcement Learning: Acceleration, Convergence, and Stabilization	3732
<i>Bei Jiang, Shangling Jui, Linglong Kong, Yi Liu, Bo Pan, Ke Sun, Yafei Wang, Yingnan Zhao</i>	
Approximate Decomposable Submodular Function Minimization for Cardinality-Based Components	3744
<i>Austin Benson, Jon Kleinberg, Nate Veldt</i>	
Episodic Multi-Agent Reinforcement Learning with Curiosity-driven Exploration	3757
<i>Jiarui Chen, Yingfeng Chen, Changjie Fan, Yang Gao, Jiamin He, Yujing Hu, Jianhao Wang, Chongjie Zhang, Lulu Zheng</i>	
Two Sides of Meta-Learning Evaluation: in Vs. Out of Distribution	3770
<i>Oscar Li, Amrith Setlur, Virginia Smith</i>	
Debiased Visual Question Answering from Feature and Sample Perspectives	3784
<i>Mingkui Tan, Zhiquan Wen, Qi Wu, Qingyao Wu, Guanghui Xu</i>	
Towards a Unified Game-Theoretic View of Adversarial Perturbations and Robustness	3797
<i>Lu Chen, Yiting Chen, Xu Cheng, Jie Ren, Jie Shi, Xin Wang, Yisen Wang, Die Zhang, Quanshi Zhang, Meng Zhou, Zhanpeng Zhou</i>	
On the Out-Of-distribution Generalization of Probabilistic Image Modelling	3811
<i>Steven McDonagh, Andi Zhang, Mingtian Zhang</i>	

Exploiting Local Convergence of Quasi-Newton Methods Globally: Adaptive Sample Size Approach	3824
<i>Qiujiang Jin, Aryan Mokhtari</i>	
PDE-GCN: Novel Architectures for Graph Neural Networks Motivated by Partial Differential Equations	3836
<i>Moshe Eliasof, Eldad Haber, Eran Treister</i>	
Information Directed Reward Learning for Reinforcement Learning	3850
<i>Kamil Ciosek, Andreas Krause, David Lindner, Sebastian Tschiatschek, Matteo Turchetta</i>	
SSMF: Shifting Seasonal Matrix Factorization	3863
<i>Siddharth Bhatia, Bryan Hooi, Koki Kawabata, Rui Liu, Mohit Wadhwa</i>	
Associative Memories Via Predictive Coding	3874
<i>Rafal Bogacz, Simon Frieder, Yujian Hong, Thomas Lukasiewicz, Tommaso Salvatori, Lei Sha, Yuhang Song, Zhenghua Xu</i>	
Robust and Differentially Private Mean Estimation	3887
<i>Sham Kakade, Weihao Kong, Xiyang Liu, Sewoong Oh</i>	
Adaptable Agent Populations Via a Generative Model of Policies	3902
<i>Kenneth Derek, Phillip Isola</i>	
A No-Go Theorem for Robust Acceleration in the Hyperbolic Plane	3914
<i>Linus Hamilton, Ankur Moitra</i>	
Privately Learning Mixtures of Axis-Aligned Gaussians	3925
<i>Ishaq Aden-Ali, Hassan Ashtiani, Christopher Liaw</i>	
Deep Self-Dissimilarities as Powerful Visual Fingerprints	3939
<i>Yuval Bahat, Idan Kligvasser, Tomer Michaeli, Tamar Shaham</i>	
Invariant Causal Imitation Learning for Generalizable Policies	3952
<i>Ioana Bica, Daniel Jarrett, Mihaela Van Der Schaar</i>	
CoAtNet: Marrying Convolution and Attention for All Data Sizes	3965
<i>Zihang Dai, Quoc V Le, Hanxiao Liu, Mingxing Tan</i>	
Mixed Supervised Object Detection by Transferring Mask Prior and Semantic Similarity	3978
<i>Junjie Chen, Yan Liu, Li Niu, Liqing Zhang, Zhijie Zhang</i>	
Celebrating Diversity in Shared Multi-Agent Reinforcement Learning	3991
<i>Li Chenghao, Tonghan Wang, Chengjie Wu, Jun Yang, Chongjie Zhang, Qianchuan Zhao</i>	
Rebounding Bandits for Modeling Satiation Effects	4003
<i>Fatma Kilinc Karzan, Liu Leqi, Zachary Lipton, Alan Montgomery</i>	
Sample Complexity of Tree Search Configuration: Cutting Planes and Beyond	4015
<i>Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, Ellen Vitercik</i>	
IQ-Learn: Inverse soft-Q Learning for Imitation	4028
<i>Shuvam Chakraborty, Chris Cundy, Stefano Ermon, Divyansh Garg, Jiaming Song</i>	
Task-Agnostic Undesirable Feature Deactivation Using Out-of-Distribution Data	4040
<i>Minseok Kim, Jae-Gil Lee, Dongmin Park, Hwanjun Song</i>	

Private Non-Smooth ERM and SCO in Subquadratic Steps	4053
<i>Janardhan Kulkarni, Yin Tat Lee, Daogao Liu</i>	

VOLUME 6

Towards Instance-Optimal Offline Reinforcement Learning with Pessimism	4065
<i>Yu-Xiang Wang, Ming Yin</i>	
Speedy Performance Estimation for Neural Architecture Search.....	4079
<i>Miroslav Fil, Yarin Gal, Clare Lyle, Robin Ru, Lisa Schut, Mark Van Der Wilk</i>	
How Tight Can PAC-Bayes Be in the Small Data Regime?	4093
<i>Wessel Bruinsma, David Burt, Andrew Foong, Richard Turner</i>	
Deep Synoptic Monte-Carlo Planning in Reconnaissance Blind Chess	4106
<i>Gregory Clark</i>	
Dynamic Analysis of Higher-Order Coordination in Neuronal Assemblies Via De-Sparsified Orthogonal Matching Pursuit	4120
<i>Behdash Babadi, Shoutik Mukherjee</i>	
Efficient Training of Retrieval Models Using Negative Cache.....	4134
<i>Ruiqi Guo, Sanjiv Kumar, Erik Lindgren, Sashank Reddi</i>	
Understanding Partial Multi-Label Learning Via Mutual Information	4147
<i>Wei Bao, Xiuwen Gong, Dong Yuan</i>	
Environment Generation for Zero-Shot Compositional Reinforcement Learning	4157
<i>Jongwook Choi, Aleksandra Faust, Izzeddin Gur, Natasha Jaques, Honglak Lee, Yingjie Miao, Manoj Tiwari</i>	
Optimizing Conditional Value-At-Risk of Black-Box Functions.....	4170
<i>Zhongxiang Dai, Patrick Jaillet, Bryan Kian Hsiang Low, Quoc Phong Nguyen</i>	
E(n) Equivariant Normalizing Flows	4181
<i>Fabian Fuchs, Victor Garcia Satorras, Emiel Hoogeboom, Ingmar Posner, Max Welling</i>	
Revitalizing CNN Attention Via Transformers in Self-Supervised Visual Representation Learning	4193
<i>Chongjian Ge, Jianbo Jiao, Youwei Liang, Ping Luo, Yibing Song, Jue Wang</i>	
A Critical Look at the Consistency of Causal Estimation with Deep Latent Variable Models	4207
<i>Pekka Marttinen, Severi Rissanen</i>	
Improving Robustness Using Generated Data.....	4218
<i>Dan Andrei Calian, Sven Gowal, Timothy A Mann, Sylvestre-Alvise Rebuffi, Florian Stimberg, Olivia Wiles</i>	
An Analysis of Constant Step Size SGD in the Non-Convex Regime: Asymptotic Normality and Bias.....	4234
<i>Krishnakumar Balasubramanian, Murat Erdogdu, Stanislav Volgushev, Lu Yu</i>	
Learning to Learn Graph Topologies	4249
<i>Tianyue Cao, Siheng Chen, Xiaowen Dong, Xingyue Pu, Xiaoyun Zhang</i>	
Invertible Tabular GANs: Killing Two Birds with One Stone for Tabular Data Synthesis	4263
<i>Jihoon Cho, Jihyeon Hyeong, Jinsung Jeon, Jaehoon Lee, Noseong Park</i>	

Reducing Collision Checking for Sampling-Based Motion Planning Using Graph Neural Networks.....	4274
<i>Sicun Gao, Chenning Yu</i>	
Sample Complexity Bounds for Active Ranking from Multi-Wise Comparisons.....	4290
<i>Jia Liu, Wenbo Ren, Ness Shroff</i>	
Efficient Bayesian Network Structure Learning Via Local Markov Boundary Search.....	4301
<i>Bryon Aragam, Ming Gao</i>	
Learning Dynamic Graph Representation of Brain Connectome with Spatio-Temporal Attention	4314
<i>Byung-Hoon Kim, Jae-Jin Kim, Jong Chul Ye</i>	
Understanding the Generalization Benefit of Model Invariance from a Data Perspective	4328
<i>Bang An, Furong Huang, Sicheng Zhu</i>	
Improved Variance-Aware Confidence Sets for Linear Bandits and Linear Mixture MDP.....	4342
<i>Simon Du, Xiangyang Ji, Jiaqi Yang, Zihan Zhang</i>	
How Should Pre-Trained Language Models Be Fine-Tuned Towards Adversarial Robustness?	4356
<i>Xinshuai Dong, Min Lin, Anh Tuan Luu, Shuicheng Yan, Hanwang Zhang</i>	
Recursive Bayesian Networks: Generalising and Unifying Probabilistic Context-Free Grammars and Dynamic Bayesian Networks.....	4370
<i>Robert Lieck, Martin Rohrmeier</i>	
EF21: a New, Simpler, Theoretically Better, and Practically Faster Error Feedback.....	4384
<i>Ilyas Fatkhullin, Peter Richtarik, Igor Sokolov</i>	
Mixture Weights Optimisation for Alpha-Divergence Variational Inference	4397
<i>Kamelia Daudel, Randal Douc</i>	
Instance-Dependent Label-noise Learning Under a Structural Causal Model.....	4409
<i>Mingming Gong, Bo Han, Tongliang Liu, Gang Niu, Yu Yao, Kun Zhang</i>	
Combining Human Predictions with Model Probabilities Via Confusion Matrices and Calibration	4421
<i>Gavin Kerrigan, Padhraic Smyth, Mark Steyvers</i>	
LeadCache: Regret-Optimal Caching in Networks	4435
<i>Debjit Paria, Abhishek Sinha</i>	
Probabilistic Attention for Interactive Segmentation	4448
<i>Manjot Bilkhu, Prasad Gabbur, Javier Movellan</i>	
Influence Patterns for Explaining Information Flow in BERT.....	4461
<i>Anupam Datta, Kaiji Lu, Piotr Mardziel, Zifan Wang</i>	
Robust Regression Revisited: Acceleration and Improved Estimation Rates	4475
<i>Arun Jambulapati, Jerry Li, Tselil Schramm, Kevin Tian</i>	
Automatic Unsupervised Outlier Model Selection.....	4489
<i>Leman Akoglu, Ryan Rossi, Yue Zhao</i>	
Pruning Randomly Initialized Neural Networks with Iterative Randomization.....	4503
<i>Daiki Chijiwa, Tomohiro Inoue, Yasutoshi Ida, Kenji Umakoshi, Shin'Ya Yamaguchi</i>	
Probing Inter-Modality: Visual Parsing with Self-Attention for Vision-and-Language Pre-training.....	4514
<i>Jianlong Fu, Yupan Huang, Houqiang Li, Bei Liu, Jiebo Luo, Houwen Peng, Hongwei Xue</i>	

Stability and Generalization of Bilevel Programming in Hyperparameter Optimization	4529
<i>Fan Bao, Chongxuan Li, Guoqiang Wu, Bo Zhang, Jun Zhu</i>	
Regime Switching Bandits	4542
<i>Ningyuan Chen, Xuefeng Gao, Yi Xiong, Xiang Zhou</i>	
MixACM: Mixup-Based Robustness Transfer Via Distillation of Activated Channel Maps.....	4555
<i>Sung-Ho Bae, Zhenguo Li, Jiawei Li, Awais Muhammad, Chuanlong Xie, Fengwei Zhou</i>	
Localization, Convexity, and Star Aggregation.....	4570
<i>Suhas Vijaykumar</i>	
Aligning Silhouette Topology for Self-Adaptive 3D Human Pose Recovery	4582
<i>Varun Jampani, Jogendra Nath Kundu, Ramesha Rakesh Mugaludi, Venkatesh Babu R</i>	
Self-Adaptable Point Processes with Nonparametric Time Decays	4594
<i>Zhimeng Pan, Jeff M Phillips, Zheng Wang, Shandian Zhe</i>	
Offline Meta Reinforcement Learning -- Identifiability Challenges and Effective Data Collection Strategies	4607
<i>Ron Dorfman, Idan Shenfeld, Aviv Tamar</i>	
RoMA: Robust Model Adaptation for Offline Model-Based Optimization	4619
<i>Sungsoo Ahn, Jinwoo Shin, Le Song, Sihyun Yu</i>	
Flexible Option Learning	4632
<i>Martin Klissarov, Doina Precup</i>	
Faster Directional Convergence of Linear Neural Networks Under Spherically Symmetric Data.....	4647
<i>Dachao Lin, Ruoyu Sun, Zhihua Zhang</i>	
Online Facility Location with Multiple Advice.....	4661
<i>Matteo Almanza, Flavio Chierichetti, Silvio Lattanzi, Alessandro Panconesi, Giuseppe Re</i>	
Credit Assignment in Neural Networks Through Deep Feedback Control	4674
<i>Javier Garcia Ordonez, Benjamin F. Grewe, Alexander Meulemans, João Sacramento, Matilde Tristany Farinha, Pau Vilimelis Aceituno</i>	
Robust Online Correlation Clustering	4688
<i>Silvio Lattanzi, Benjamin Moseley, Sergei Vassilvitskii, Yuyan Wang, Rudy Zhou</i>	
Neural Additive Models: Interpretable Machine Learning with Neural Nets.....	4699
<i>Rishabh Agarwal, Rich Caruana, Nicholas Frosst, Geoffrey Hinton, Ben Lengerich, Levi Melnick, Xuezhou Zhang</i>	
Representation Learning for Event-Based Visuomotor Policies	4712
<i>Ashish Kapoor, Sami Mian, Sai Vemprala</i>	
Kernel Functional Optimisation	4725
<i>Arun Kumar Anjanapura Venkatesh, Sunil Gupta, Santu Rana, Alistair Shilton, Svetha Venkatesh</i>	
Generalized Shape Metrics on Neural Representations.....	4738
<i>Simon Kornblith, Erin Kunz, Scott Linderman, Alex H Williams</i>	

Diverse Message Passing for Attribute with Heterophily.....	4751
<i>Xiaochun Cao, Yuanfang Guo, Mengzhe Li, Liyang Liu, Chuan Wang, Liang Yang, Bingxin Niu</i>	
Towards Robust Bisimulation Metric Learning	4764
<i>Tristan Aumentado-Armstrong, Mete Kemertas</i>	
Beyond BatchNorm: Towards a Unified Understanding of Normalization in Deep Learning.....	4778
<i>Robert Dick, Ekdeep S Lubana, Hidenori Tanaka</i>	
Representation Learning Beyond Linear Prediction Functions	4792
<i>Ambuj Tewari, Ziping Xu</i>	
Volume Rendering of Neural Implicit Surfaces	4805
<i>Jiatao Gu, Yoni Kasten, Yaron Lipman, Lior Yariv</i>	
MAUVE: Measuring the Gap Between Neural Text and Human Text Using Divergence Frontiers	4816
<i>Yejin Choi, Zaid Harchaoui, Krishna Pillutla, Swabha Swayamdipta, John Thickstun, Sean Welleck, Rowan Zellers</i>	
Accurately Solving Rod Dynamics with Graph Learning	4829
<i>Jan Bender, Torsten Hadrach, Tassilo Kugelstadt, Dominik L Michels, Wojtek Palubicki, Soeren Pirk, Han Shao</i>	
Limiting Fluctuation and Trajectorial Stability of Multilayer Neural Networks with Mean Field Training	4843
<i>Phan-Minh Nguyen, Huy Tuan Pham</i>	
Medical Dead-Ends and Learning to Identify High-Risk States and Treatments	4856
<i>Mehdi Fatemi, Marzyeh Ghassemi, Taylor Killian, Jayakumar Subramanian</i>	
Overcoming the Convex Barrier for Simplex Inputs	4871
<i>Harkirat Singh Behl, Krishnamurthy Dvijotham, M. Pawan Kumar, Philip Torr</i>	
High-Probability Bounds for Non-Convex Stochastic Optimization with Heavy Tails	4883
<i>Ashok Cutkosky, Harsh Mehta</i>	
Batch Normalization Orthogonalizes Representations in Deep Random Networks.....	4896
<i>Francis Bach, Hadi Daneshmand, Amir Joudaki</i>	
Support Vector Machines and Linear Regression Coincide with Very High-Dimensional Features	4907
<i>Navid Ardehshir, Daniel Hsu, Clayton Sanford</i>	

VOLUME 7

Coupled Segmentation and Edge Learning Via Dynamic Graph Propagation	4919
<i>Anima Anandkumar, Thomas Breuel, Wonmin Byeon, Rui Huang, Jan Kautz, Guilin Liu, Sifei Liu, Zhiding Yu</i>	
Offline RL Without Off-Policy Evaluation	4933
<i>David Brandfonbrener, Joan Bruna, Rajesh Ranganath, Will Whitney</i>	
Continuous Vs. Discrete Optimization of Deep Neural Networks	4947
<i>Nadav Cohen, Omer Elkabetz</i>	

CrypTen: Secure Multi-Party Computation Meets Machine Learning.....	4961
<i>Awni Hannun, Mark Ibrahim, Brian Knott, Shubho Sengupta, Shobha Venkataraman, Laurens Van Der Maaten</i>	
Can Contrastive Learning Avoid Shortcut Solutions?.....	4974
<i>Kayhan Batmanghelich, Stefanie Jegelka, Joshua Robinson, Suvrit Sra, Li Sun, Ke Yu</i>	
See More for Scene: Pairwise Consistency Learning for Scene Classification.....	4987
<i>Gongwei Chen, Shuqiang Jiang, Xinhang Song, Bohan Wang</i>	
Provable Guarantees for Self-Supervised Deep Learning with Spectral Contrastive Loss	5000
<i>Adrien Gaidon, Jeff Z. Haochen, Tengyu Ma, Colin Wei</i>	
Greedy Approximation Algorithms for Active Sequential Hypothesis Testing	5012
<i>Kyra Gan, Su Jia, Andrew Li</i>	
When False Positive is Intolerant: End-To-End Optimization with Low FPR for Multipartite Ranking	5025
<i>Yuan He, Qingming Huang, Peisong Wen, Qianqian Xu, Zhiyong Yang</i>	
Convex Polytope Trees and Its Application to VAE.....	5038
<i>Mohammadreza Armandpour, Ali Sadeghian, Mingyuan Zhou</i>	
The Skellam Mechanism for Differentially Private Federated Learning.....	5052
<i>Naman Agarwal, Peter Kairouz, Ziyu Liu</i>	
Stability and Deviation Optimal Risk Bounds with Convergence Rate $O(1/n)$	5065
<i>Yegor Klochkov, Nikita Zhitovskiy</i>	
SketchGen: Generating Constrained CAD Sketches	5077
<i>Shariq Bhat, Paul Guerrero, Leonidas Guibas, Tom Kelly, Niloy Mitra, Wamiq Para, Peter Wonka</i>	
CLDA: Contrastive Learning for Semi-Supervised Domain Adaptation	5089
<i>Ankit Singh</i>	
Differentially Private N-Gram Extraction	5102
<i>Sivakanth Gopi, Kunho Kim, Janardhan Kulkarni, Sergey Yekhanin</i>	
Capturing Implicit Hierarchical Structure in 3D Biomedical Images with Self-Supervised Hyperbolic Representations.....	5112
<i>Wah Chiu, Jeffrey Gu, Joy Hsu, Gong Wu, Serena Yeung</i>	
Noisy Recurrent Neural Networks	5124
<i>N. Benjamin Erichson, Liam Hodgkinson, Soon Hoe Lim, Michael Mahoney</i>	
Matrix Encoding Networks for Neural Combinatorial Optimization	5138
<i>Jinho Choo, Youngjune Gwon, Yeong-Dae Kwon, Duwon Park, Minah Park, Iljoo Yoon</i>	
When is Unsupervised Disentanglement Possible?	5150
<i>Daniella Horan, Eitan Richardson, Yair Weiss</i>	
Continuous Latent Process Flows.....	5162
<i>Marcus Brubaker, Ruizhi Deng, Andreas Lehrmann, Greg Mori</i>	
Perturbation-Based Regret Analysis of Predictive Control in Linear Time Varying Systems	5174
<i>Yang Hu, Yiheng Lin, Guannan Qu, Guanya Shi, Haoyuan Sun, Adam Wierman</i>	

Dataset Distillation with Infinitely Wide Convolutional Networks.....	5186
<i>Jaehoon Lee, Timothy Nguyen, Roman Novak, Lechao Xiao</i>	
SPANN: Highly-Efficient Billion-scale Approximate Nearest Neighborhood Search	5199
<i>Qi Chen, Mingqin Li, Zengzhong Li, Chuanjie Liu, Haidong Wang, Jingdong Wang, Mao Yang, Bing Zhao</i>	
Distilling Object Detectors with Feature Richness.....	5213
<i>Ming Chang, Tianshi Chen, Yunji Chen, Shaoli Liu, Rui Zhang, Du Zhixing, Xishan Zhang</i>	
Analysis of One-Hidden-layer Neural Networks Via the Resolvent Method.....	5225
<i>Vanessa Piccolo, Dominik Schroder</i>	
Grounding Spatio-Temporal Language with Transformers.....	5236
<i>Katja Hofmann, Tristan Karch, Clement Moulin-Frier, Pierre-Yves Oudeyer, Laetitia Teodorescu</i>	
Learning Where to Learn: Gradient Sparsity in Meta and Continual Learning.....	5250
<i>Massimo Caccia, Seijin Kobayashi, João Sacramento, Simon Schug, Dominic Zhao, Nicolas Zucchet, Johannes Von Oswald</i>	
Domain Invariant Representation Learning with Domain Density Transformations	5264
<i>Atilim Gunes Baydin, Yarin Gal, A. Tuan Nguyen, Toan Tran</i>	
PlayVirtual: Augmenting Cycle-Consistent Virtual Trajectories for Reinforcement Learning.....	5276
<i>Zhibo Chen, Mingxiao Feng, Cuiling Lan, Tao Yu, Wenjun Zeng, Zhizheng Zhang</i>	
Efficient Equivariant Network.....	5290
<i>Yuxuan Chen, Yiming Dong, Lingshen He, Zhouchen Lin, Yisen Wang, Zhengyang Shen</i>	
Unifying Gradient Estimators for Meta-Reinforcement Learning Via Off-Policy Evaluation.....	5303
<i>Tadashi Kozuno, Remi Munos, Mark Rowland, Yunhao Tang, Michal Valko</i>	
Even Your Teacher Needs Guidance: Ground-Truth Targets Dampen Regularization Imposed by Self-Distillation	5316
<i>Lars N Andersen, Kenneth Borup</i>	
Compressing Neural Networks: Towards Determining the Optimal Layer-Wise Decomposition.....	5328
<i>Dan Feldman, Lucas Liebenwein, Alaa Maalouf, Daniela Rus</i>	
Equilibrium and non-Equilibrium Regimes in the Learning of Restricted Boltzmann Machines	5345
<i>Aurelien Decelle, Cyril Furtlehner, Beatriz Seoane</i>	
Imitation with Neural Density Models	5360
<i>Stefano Ermon, Akshat Jindal, Kuno Kim, Jiaming Song, Yang Song, Yanan Sui</i>	
Accurate Point Cloud Registration with Robust Optimal Transport	5373
<i>Ariel H Curiale, Jean Feydy, Peirong Liu, Marc Niethammer, Raul San Jose Estepar, Ruben San Jose Estepar, Zhengyang Shen</i>	
Simple Steps Are All You Need: Frank-Wolfe and Generalized Self-concordant Functions	5390
<i>Mathieu Besancon, Alejandro Carderera, Sebastian Pokutta</i>	
Automatic Data Augmentation for Generalization in Reinforcement Learning.....	5402
<i>Rob Fergus, Maxwell Goldstein, Ilya Kostrikov, Roberta Raileanu, Denis Yarats</i>	

Blending Anti-Aliasing into Vision Transformer	5416
<i>Jiaya Jia, Mu Li, Shengju Qian, Hao Shao, Yi Zhu</i>	
A Trainable Spectral-Spatial Sparse Coding Model for Hyperspectral Image Restoration.....	5430
<i>Theo Bodrito, Jocelyn Chanussot, Julien Mairal, Alexandre Zouaoui</i>	
Posterior Collapse and Latent Variable Non-Identifiability	5443
<i>David Blei, John Cunningham, Yixin Wang</i>	
The Benefits of Implicit Regularization from SGD in Least Squares Problems	5456
<i>Vladimir Braverman, Dean Foster, Quanquan Gu, Sham Kakade, Jingfeng Wu, Difan Zou</i>	
Generalization of Model-Agnostic Meta-Learning Algorithms: Recurring and Unseen Tasks	5469
<i>Alireza Fallah, Aryan Mokhtari, Asuman Ozdaglar</i>	
Factored Policy Gradients: Leveraging Structure for Efficient Learning in MOMDPs	5481
<i>Sumitra Ganesh, Thomas Spooner, Nelson Vadori</i>	
MarioNette: Self-Supervised Sprite Learning	5494
<i>Alexei Efros, Matthew Fisher, Michael Gharbi, Vitor Guizilini, Dmitriy Smirnov, Justin Solomon</i>	
RLlib Flow: Distributed Reinforcement Learning is a Dataflow Problem	5506
<i>Joseph Gonzalez, Eric Liang, Michael Luo, Sven Mika, Ion Stoica, Zhanghao Wu</i>	
Improve Agents Without Retraining: Parallel Tree Search with Off-Policy Correction	5518
<i>Gal Chechik, Gal Dalal, Steven Dalton, Assaf Hallak, Shie Mannor, Iuri Frosio</i>	
Redesigning the Transformer Architecture with Insights from Multi-Particle Dynamical Systems	5531
<i>Soumen Chakrabarti, Tanmoy Chakraborty, Subhabrata Dutta, Tanya Gautam</i>	
Exploring Architectural Ingredients of Adversarially Robust Deep Neural Networks	5545
<i>James Bailey, Sarah Erfani, Quanquan Gu, Hanxun Huang, Xingjun Ma, Yisen Wang</i>	
Center Smoothing: Certified Robustness for Networks with Structured Outputs	5560
<i>Tom Goldstein, Aounon Kumar</i>	
Breaking the Linear Iteration Cost Barrier for Some Well-Known Conditional Gradient Methods Using MaxIP Data-structures	5576
<i>Anshumali Shrivastava, Zhao Song, Zhaozhuo Xu</i>	
Neural Regression, Representational Similarity, Model Zoology & Neural Taskonomy at Scale in Rodent Visual Cortex	5590
<i>George Alvarez, Andrei Barbu, Michael Buice, Colin Conwell, Boris Katz, David Mayo</i>	
A Topological Perspective on Causal Inference.....	5608
<i>Duligur Ibeling, Thomas Icard</i>	
Parameter Inference with Bifurcation Diagrams	5620
<i>Attila Csikasz-Nagy, Neil Dalchau, Gregory Szep</i>	
Scalable Thompson Sampling Using Sparse Gaussian Process Models.....	5631
<i>Artem Artemev, Vincent Dutordoir, Henry Moss, Victor Picheny, Sattar Vakili</i>	
Robust Counterfactual Explanations on Graph Neural Networks	5644
<i>Mohit Bajaj, Lingyang Chu, Peter Cho-Ho Lam, Jian Pei, Lanjun Wang, Zi Yu Xue, Yong Zhang</i>	

Similarity and Matching of Neural Network Representations.....	5656
<i>Adrian Csiszarik, Peter Korosi-Szabo, Akos Matszangosz, Gergely Papp, Daniel Varga</i>	
DOCTOR: a Simple Method for Detecting Misclassification Errors	5669
<i>Daniele Gorla, Federica Granese, Catuscia Palamidessi, Pablo Piantanida, Marco Romanelli</i>	
Contrastive Laplacian Eigenmaps	5682
<i>Peter Koniusz, Ke Sun, Hao Zhu</i>	
Machine Learning Structure Preserving Brackets for Forecasting Irreversible Processes	5696
<i>Kookjin Lee, Panos Stinis, Nathaniel Trask</i>	
On the Variance of the Fisher Information for Deep Learning	5708
<i>Alexander Soen, Ke Sun</i>	
A ² -Net: Learning Attribute-Aware Hash Codes for Large-Scale Fine-Grained Image Retrieval.....	5720
<i>Yang Shen, Xuhao Sun, Xiu-Shen Wei, Jian Yang, Han-Jia Ye</i>	
Shape Registration in the Time of Transformers.....	5731
<i>Luca Cosmo, Riccardo Marin, Simone Melzi, Luca Moschella, Emanuele Rodolà, Giovanni Trappolini</i>	
Brick-By-Brick: Combinatorial Construction with Deep Reinforcement Learning	5745
<i>Minsu Cho, Hyunsoo Chung, Jungtaek Kim, Boris Knyazev, Jinhwi Lee, Jaesik Park, Graham Taylor</i>	
Dissecting the Diffusion Process in Linear Graph Convolutional Networks	5758
<i>Zhouchen Lin, Yifei Wang, Yisen Wang, Jiansheng Yang</i>	
Dynamic Grained Encoder for Vision Transformers.....	5770
<i>Xuming He, Zeming Li, Songtao Liu, Lin Song, Hongbin Sun, Jian Sun, Songyang Zhang, Nanning Zheng</i>	

VOLUME 8

Understanding Negative Samples in Instance Discriminative Self-Supervised Representation Learning	5784
<i>Kento Nozawa, Issei Sato</i>	
On UMAP's True Loss Function	5798
<i>Sebastian Damrich, Fred Hamprecht</i>	
Fast Pure Exploration Via Frank-Wolfe	5810
<i>Alexandre Proutiere, Ruo-Chun Tzeng, Po-An Wang</i>	
iFlow: Numerically Invertible Flows for Efficient Lossless Compression Via a Uniform Coder.....	5822
<i>Ning Kang, Zhenguo Li, Tom Ryder, Shifeng Zhang</i>	
History Aware Multimodal Transformer for Vision-And-Language Navigation.....	5834
<i>Shizhe Chen, Pierre-Louis Guhur, Ivan Laptev, Cordelia Schmid</i>	
Meta Two-Sample Testing: Learning Kernels for Testing with Limited Data.....	5848
<i>Feng Liu, Jie Lu, Danica J. Sutherland, Wenkai Xu</i>	

Process for Adapting Language Models to Society (PALMS) with Values-Targeted Datasets	5861
<i>Christy Dennison, Irene Solaiman</i>	
The Lazy Online Subgradient Algorithm is Universal on Strongly Convex Domains	5874
<i>Daron Anderson, Douglas Leith</i>	
Computer-Aided Design as Language	5885
<i>Sergey Bartunov, Yaroslav Ganin, Ethan Keller, Yujia Li, Stefano Saliceti</i>	
COHESIV: Contrastive Object and Hand Embedding Segmentation in Video	5898
<i>David Fouhey, Richard Higgins, Dandan Shan</i>	
ByPE-VAE: Bayesian Pseudocoresets Exemplar VAE	5910
<i>Qingzhong Ai, Lirong He, Shiyu Liu, Zenglin Xu</i>	
Recovery Analysis for Plug-And-Play Priors Using the Restricted Eigenvalue Condition	5921
<i>Salman Asif, Ulugbek Kamilov, Jiaming Liu, Brendt Wohlberg</i>	
Group Equivariant Subsampling	5934
<i>Hyunjik Kim, Thomas Rainforth, Yee Teh, Jin Xu</i>	
Data Sharing and Compression for Cooperative Networked Control	5947
<i>Jiangnan Cheng, Sandeep Chinchali, Sachin Katti, Marco Pavone, Ao Tang</i>	
Hyperbolic Procrustes Analysis Using Riemannian Geometry	5959
<i>Yuval Kluger, Ya-Wei Eileen Lin, Ronen Talmon</i>	
No Fear of Heterogeneity: Classifier Calibration for Federated Learning with Non-IID Data	5972
<i>Fei Chen, Jiashi Feng, Dapeng Hu, Jian Liang, Mi Luo, Yifan Zhang</i>	
Preconditioned Gradient Descent for Over-Parameterized Nonconvex Matrix Factorization	5985
<i>Salar Fattahi, Jialun Zhang, Richard Y Zhang</i>	
Improving Contrastive Learning on Imbalanced Data Via Open-World Sampling	5997
<i>Tianlong Chen, Ting Chen, Ziyu Jiang, Zhangyang Wang</i>	
Primer: Searching for Efficient Transformers for Language Modeling	6010
<i>Zihang Dai, Quoc V Le, Hanxiao Liu, Wojciech Manke, Noam Shazeer, David So</i>	
Scaling Ensemble Distribution Distillation to Many Classes with Proxy Targets	6023
<i>Mark Gales, Andrey Malinin, Max Ryabinin</i>	
Multi-Person 3D Motion Prediction with Multi-Range Transformers	6036
<i>Medhini Narasimhan, Xiaolong Wang, Jiashun Wang, Huazhe Xu</i>	
STEM: a Stochastic Two-Sided Momentum Algorithm Achieving Near-Optimal Sample and Communication Complexities for Federated Learning	6050
<i>Mingyi Hong, Prashant Khanduri, Jia Liu, Ketan Rajawat, Pranay Sharma, Pramod Varshney, Haibo Yang</i>	
Bubblewrap: Online Tiling and Real-Time Flow Prediction on Neural Manifolds	6062
<i>Anne Draelos, Pranjal Gupta, Na Young Jun, John Pearson, Chaichontat Sriworarat</i>	
The Semi-Random Satisfaction of Voting Axioms	6075
<i>Lirong Xia</i>	

Deep Marching Tetrahedra: a Hybrid Representation for High-Resolution 3D Shape Synthesis	6087
<i>Sanja Fidler, Jun Gao, Ming-Yu Liu, Tianchang Shen, Kangxue Yin</i>	
Learning to Combine Per-Example Solutions for Neural Program Synthesis	6102
<i>Hugo Larochelle, Disha Shrivastava, Daniel Tarlow</i>	
On Success and Simplicity: a Second Look at Transferable Targeted Attacks	6115
<i>Martha Larson, Zhuoran Liu, Zhengyu Zhao</i>	
Provably Efficient, Succinct, and Precise Explanations	6129
<i>Guy Blanc, Jane Lange, Li-Yang Tan</i>	
Refined Learning Bounds for Kernel and Approximate k-Means	6142
<i>Yong Liu</i>	
Learning Causal Semantic Representation for Out-Of-Distribution Prediction	6155
<i>Wei Chen, Tao Li, Chang Liu, Tie-Yan Liu, Tao Qin, Xinwei Sun, Haoyue Tang, Jindong Wang</i>	
A First-Order Primal-dual Method with Adaptivity to Local Smoothness.....	6171
<i>Volkan Cevher, Yura Malitsky, Maria-Luiza Vladarean</i>	
A Theory-Driven Self-Labeling Refinement Method for Contrastive Representation Learning	6183
<i>Steven Chu Hong Hoi, Caiming Xiong, Xiaotong Yuan, Pan Zhou</i>	
Adversarial Robustness with Semi-Infinite Constrained Learning.....	6198
<i>Luiz Chamon, Hamed Hassani, George J. Pappas, Alejandro Ribeiro, Alexander Robey</i>	
Conformal Time-Series Forecasting.....	6216
<i>Ahmed M. Alaa, Kamile Stankeviciute, Mihaela Van Der Schaar</i>	
A 3D Generative Model for Structure-Based Drug Design.....	6229
<i>Jiaqi Guan, Shitong Luo, Jianzhu Ma, Jian Peng</i>	
Bootstrapping the Error of Oja's Algorithm	6240
<i>Robert Lunde, Purnamrita Sarkar, Rachel Ward</i>	
Landscape Analysis of an Improved Power Method for Tensor Decomposition	6253
<i>Joe Kileel, Timo Klock, João M Pereira</i>	
Curriculum Offline Imitating Learning	6266
<i>Minghuan Liu, Tie-Yan Liu, Jian Shen, Zhengyu Yang, Weinan Zhang, Hanyue Zhao, Li Zhao</i>	
Robust Pose Estimation in Crowded Scenes with Direct Pose-Level Inference.....	6278
<i>Gang Hua, Dongkai Wang, Shiliang Zhang</i>	
Ising Model Selection Using ℓ_1 -Regularized Linear Regression: a Statistical Mechanics Analysis.....	6290
<i>Yoshiyuki Kabashima, Xiangming Meng, Tomoyuki Obuchi</i>	
Conformal Prediction Using Conditional Histograms.....	6304
<i>Yaniv Romano, Matteo Sesia</i>	
Contrastive Graph Poisson Networks: Semi-Supervised Learning with Extremely Limited Labels.....	6316
<i>Chen Gong, Liu Liu, Shirui Pan, Sheng Wan, Baosheng Yu, Yibing Zhan</i>	
Collaborative Uncertainty in Multi-Agent Trajectory Forecasting.....	6328
<i>Siheng Chen, Ulrich Neumann, Bohan Tang, Gang Wang, Ya Zhang, Yiqi Zhong</i>	

Network-To-Network Regularization: Enforcing Occam's Razor to Improve Generalization	6341
<i>Rohan Ghosh, Mehul Motani</i>	
Generalized and Discriminative Few-Shot Object Detection Via SVD-Dictionary Enhancement	6353
<i>Cheng Deng, Wei Liu, Aming Wu, Suqi Zhao</i>	
Conditioning Sparse Variational Gaussian Processes for Online Decision-Making	6365
<i>Wesley Maddox, Samuel Stanton, Andrew Wilson</i>	
Spherical Motion Dynamics: Learning Dynamics of Normalized Neural Network Using SGD and Weight Decay	6380
<i>Jian Sun, Ruosi Wan, Xiangyu Zhang, Zhanxing Zhu</i>	
Imitating Deep Learning Dynamics Via Locally Elastic Stochastic Differential Equations	6392
<i>Weijie Su, Hua Wang, Jiayao Zhang</i>	
Probabilistic Forecasting: A Level-Set Approach	6404
<i>Jan Gasthaus, Hilaf Hasson, Tim Januschowski, Bernie Wang</i>	
Roto-Translated Local Coordinate Frames for Interacting Dynamical Systems	6417
<i>Efstathios Gavves, Miltiadis Kofinas, Naveen Nagaraja</i>	
ParK: Sound and Efficient Kernel Ridge Regression by Feature Space Partitions	6430
<i>Daniele Calandriello, Luigi Carratino, Lorenzo Rosasco, Stefano Vigogna</i>	
Scaling Gaussian Processes with Derivative Information Using Variational Inference	6442
<i>David Bindel, Jacob Gardner, Leo Huang, Misha Padidar, Xinran Zhu</i>	
On the Representation of Solutions to Elliptic PDEs in Barron Spaces	6454
<i>Ziang Chen, Jianfeng Lu, Yulong Lu</i>	
A/B Testing for Recommender Systems in a Two-Sided Marketplace	6466
<i>Shaunak Chatterjee, Chun Lo, Preetam Nandy, Divya Venugopalan</i>	
Retiring Adult: New Datasets for Fair Machine Learning	6478
<i>Frances Ding, Moritz Hardt, John Miller, Ludwig Schmidt</i>	
Cardinality Constrained Submodular Maximization for Random Streams	6491
<i>Paul Liu, Aviad Rubinstein, Jan Vondrak, Junyao Zhao</i>	
Self-Instantiated Recurrent Units with Dynamic Soft Recursion	6503
<i>Alvin Chan, Yikang Shen, Yi Tay, Shuai Zhang, Aston Zhang</i>	
Sparse Uncertainty Representation in Deep Learning with Inducing Weights	6515
<i>Martin Kukla, Yingzhen Li, Hippolyt Ritter, Cheng Zhang</i>	
Scalable Inference of Sparsely-Changing Gaussian Markov Random Fields	6529
<i>Salar Fattahi, Andres Gomez</i>	
Grad2Task: Improved Few-Shot Text Classification Using Gradients for Task Representation	6542
<i>Michael Brudno, Frank Rudzicz, Jixuan Wang, Kuan-Chieh Wang</i>	
Learnability of Linear Thresholds from Label Proportions	6555
<i>Rishi Saket</i>	
A Variational Approximate Posterior for the Deep Wishart Process	6567
<i>Laurence Aitchison, Sebastian Ober</i>	

Neural Pseudo-Label Optimism for the Bank Loan Problem	6580
<i>Alex Berg, Edward Chou, Jakob Foerster, Aldo Pacchiano, Shaun Singh</i>	
Visualizing the Emergence of Intermediate Visual Patterns in DNNs	6594
<i>Mingjie Li, Shaobo Wang, Quanshi Zhang</i>	
Learning 3D Dense Correspondence Via Canonical Point Autoencoder	6608
<i>An-Chieh Cheng, Xueting Li, Sifei Liu, Min Sun, Ming-Hsuan Yang</i>	
Speech-T: Transducer for Text to Speech and Beyond	6621
<i>Jiawei Chen, Yichong Leng, Tie-Yan Liu, Tao Qin, Xu Tan, Guihua Wen, Jin Xu</i>	
Multi-Modal Dependency Tree for Video Captioning	6634
<i>Jiebo Luo, Xinxiao Wu, Wentian Zhao</i>	

VOLUME 9

Greedy and Random Quasi-Newton Methods with Faster Explicit Superlinear Convergence	6646
<i>Dachao Lin, Haishan Ye, Zhihua Zhang</i>	
Neural Tangent Kernel Maximum Mean Discrepancy	6658
<i>Xiuyuan Cheng, Yao Xie</i>	
Subgraph Federated Learning with Missing Neighbor Generation	6671
<i>Xiaoxiao Li, Lichao Sun, Carl Yang, Siu Ming Yiu, Ke Zhang</i>	
Bellman-Consistent Pessimism for Offline Reinforcement Learning	6683
<i>Alekh Agarwal, Ching-An Cheng, Nan Jiang, Paul Mineiro, Tengyang Xie</i>	
Can You Learn an Algorithm? Generalizing from Easy to Hard Problems with Recurrent Networks	6695
<i>Eitan Borgnia, Micah Goldblum, Tom Goldstein, Arjun Gupta, Furong Huang, Avi Schwarzschild, Uzi Vishkin</i>	
Sub-Linear Memory: How to Make Performers SLiM	6707
<i>Krzysztof Choromanski, Jared Quincy Davis, Valerii Likhoshesterov, Xingyou Song, Adrian Weller</i>	
Efficient Learning of Discrete-Continuous Computation Graphs	6720
<i>David Friede, Mathias Niepert</i>	
VQ-GNN: a Universal Framework to Scale Up Graph Neural Networks Using Vector Quantization	6733
<i>John Dickerson, Mucong Ding, Tom Goldstein, Furong Huang, Kezhi Kong, Jingling Li, Chen Zhu</i>	
Overcoming Catastrophic Forgetting in Incremental Few-Shot Learning by Finding Flat Minima.....	6747
<i>Jiaxin Chen, Guangyuan Shi, Xiao-Ming Wu, Li-Ming Zhan, Wenlong Zhang</i>	
Functional Neural Networks for Parametric Image Restoration Problems.....	6762
<i>Yanhui Guo, Fangzhou Luo, Xiaolin Wu</i>	
Intrinsic Dimension, Persistent Homology and Generalization in Neural Networks	6776
<i>Tolga Birdal, Leonidas Guibas, Aaron Lou, Umut Simsekli</i>	
GemNet: Universal Directional Graph Neural Networks for Molecules.....	6790
<i>Florian Becker, Johannes Gasteiger, Stephan Gunnemann</i>	

Loss Function Based Second-Order Jensen Inequality and Its Application to Particle Variational Inference	6803
<i>Futoshi Futami, Tomoharu Iwata, Issei Sato, Masashi Sugiyama, Naonori Ueda</i>	
Detecting and Adapting to Irregular Distribution Shifts in Bayesian Online Learning	6816
<i>Alex Boyd, Aodong Li, Stephan Mandt, Padhraic Smyth</i>	
Asynchronous Decentralized SGD with Quantized and Local Updates	6829
<i>Dan Alistarh, Peter Davies, Shigang Li, Giorgi Nadiradze, Amirmojtaba Sabour</i>	
Stochastic Shortest Path: Minimax, Parameter-Free and Towards Horizon-Free Regret	6843
<i>Simon Du, Alessandro Lazaric, Matteo Pirotta, Jean Tarbouriech, Michal Valko, Runlong Zhou</i>	
Nested Counterfactual Identification from Arbitrary Surrogate Experiments	6856
<i>Elias Bareinboim, Juan Correa, Sanghack Lee</i>	
Sim and Real: Better Together	6868
<i>Dotan Di Castro, Shirli Di-Castro, Shie Mannor</i>	
Trustworthy Multimodal Regression with Mixture of Normal-Inverse Gamma Distributions	6881
<i>Huazhu Fu, Zongbo Han, Qinghua Hu, Huan Ma, Changqing Zhang, Joey Tianyi Zhou</i>	
An Empirical Study of Adder Neural Networks for Object Detection	6894
<i>Xinghao Chen, Minjing Dong, Yunhe Wang, Chunjing Xu, Chang Xu</i>	
Does Knowledge Distillation Really Work?	6906
<i>Alexander Alemi, Pavel Izmailov, Polina Kirichenko, Samuel Stanton, Andrew Wilson</i>	
Teachable Reinforcement Learning Via Advice Distillation	6920
<i>Pieter Abbeel, Jacob Andreas, Trevor Darrell, Abhishek Gupta, Olivia Watkins</i>	
Antipodes of Label Differential Privacy: PATE and ALIBI	6934
<i>Mani Malek Esmaeili, Ilya Mironov, Karthik Prasad, Igor Shilov, Florian Tramèr</i>	
Visual Search Asymmetry: Deep Nets and Humans Share Similar Inherent Biases	6946
<i>Shashi Kant Gupta, Gabriel Kreiman, Chia-Chien Wu, Jeremy Wolfe, Mengmi Zhang</i>	
On the Universality of Graph Neural Networks on Large Random Graphs	6960
<i>Alberto Bietti, Nicolas Keriven, Samuel Vaiter</i>	
Inverse Reinforcement Learning in a Continuous State Space with Formal Guarantees	6972
<i>Kevin Bello, Gregory Dexter, Jean Honorio</i>	
Adversarial Attacks on Graph Classifiers Via Bayesian Optimisation	6983
<i>Arno Blaas, Xiaowen Dong, Henry Kenlay, Michael A Osborne, Robin Ru, Xingchen Wan</i>	
Regulating Algorithmic Filtering on Social Media	6997
<i>Sarah Cen, Devavrat Shah</i>	
Argmax Centroid	7012
<i>Chengyue Gong, Qiang Liu, Mao Ye</i>	
Contrastive Learning of Global and Local Video Representations	7025
<i>Daniel McDuff, Yale Song, Zhaoyang Zeng, Shuang Ma</i>	

BooVI: Provably Efficient Bootstrapped Value Iteration.....	7041
<i>Qi Cai, Boyi Liu, Zhaoran Wang, Zhuoran Yang</i>	
Do Wider Neural Networks Really Help Adversarial Robustness?	7054
<i>Deng Cai, Jinghui Chen, Quanquan Gu, Xiaofei He, Boxi Wu</i>	
Exploring the Limits of Out-Of-Distribution Detection	7068
<i>Stanislav Fort, Balaji Lakshminarayanan, Jie Ren</i>	
ABC: Auxiliary Balanced Classifier for Class-Imbalanced Semi-supervised Learning.....	7082
<i>Heeyoung Kim, Hyuck Lee, Seungjae Shin</i>	
BCD Nets: Scalable Variational Approaches for Bayesian Causal Discovery	7095
<i>Chris Cundy, Stefano Ermon, Aditya Grover</i>	
Discovering Dynamic Salient Regions for Spatio-Temporal Graph Neural Networks	7111
<i>Iulia Duta, Marius Leordeanu, Andrei Nicolicioiu</i>	
Information-Constrained Optimization: Can Adaptive Processing of Gradients Help?	7126
<i>Jayadev Acharya, Clement Canonne, Prathamesh Mayekar, Himanshu Tyagi</i>	
Towards Calibrated Model for Long-Tailed Visual Recognition from Prior Perspective	7139
<i>Zenghao Chai, Zhengzhuo Xu, Chun Yuan</i>	
Learning to Draw: Emergent Communication Through Sketching	7153
<i>Jonathon Hare, Daniela Mihai</i>	
Self-Supervised Learning of Event-Based Optical Flow with Spiking Neural Networks	7167
<i>Jesse Hagenaars, Federico Paredes-Valles, Guido De Croon</i>	
On the Value of Infinite Gradients in Variational Autoencoder Models	7180
<i>Bin Dai, Li Wenliang, David Wipf</i>	
Online Robust Reinforcement Learning with Model Uncertainty	7193
<i>Yue Wang, Shaofeng Zou</i>	
Neural View Synthesis and Matching for Semi-Supervised Few-Shot Learning of 3D Pose	7207
<i>Adam Kortylewski, Shenxiao Mei, Angtian Wang, Alan Yuille</i>	
Sharp Impossibility Results for Hyper-Graph Testing	7220
<i>Jiashun Jin, Zheng Tracy Ke, Jiajun Liang</i>	
Evaluating Gradient Inversion Attacks and Defenses in Federated Learning	7232
<i>Sanjeev Arora, Samyak Gupta, Yangsibo Huang, Kai Li, Zhao Song</i>	
Faster Non-Asymptotic Convergence for Double Q-learning	7242
<i>Yingbin Liang, Huaqing Xiong, Lin Zhao</i>	
Towards Tight Communication Lower Bounds for Distributed Optimisation	7254
<i>Dan Alistarh, Janne H. Korhonen</i>	
Fast Multi-Resolution Transformer Fine-tuning for Extreme Multi-label Text Classification.....	7267
<i>Wei-Cheng Chang, Inderjit Dhillon, Hsiang-Fu Yu, Jiong Zhang</i>	
HRFormer: High-Resolution Vision Transformer for Dense Predict	7281
<i>Xilin Chen, Rao Fu, Lang Huang, Weihong Lin, Jingdong Wang, Yuhui Yuan, Chao Zhang</i>	

Manifold Topology Divergence: a Framework for Comparing Data Manifolds	7294
<i>Serguei Barannikov, Evgeny Burnaev, Alexander Filippov, Alexander Korotin, Grigorii Sotnikov, Ekaterina Trimbach, Ilya Trofimov</i>	
Weak-Shot Fine-grained Classification Via Similarity Transfer	7306
<i>Junjie Chen, Liu Liu, Li Niu, Liqing Zhang</i>	
Shape Your Space: a Gaussian Mixture Regularization Approach to Deterministic Autoencoders	7319
<i>Stefan Falkner, Margret Keuper, Amrutha Saseendran, Kathrin Skubch</i>	
An Even More Optimal Stochastic Optimization Algorithm: Minibatching and Interpolation Learning	7333
<i>Nathan Srebro, Blake Woodworth</i>	
Indexed Minimum Empirical Divergence for Unimodal Bandits	7346
<i>Odalric-Ambrym Maillard, Pierre Menard, Hassan Saber</i>	
SOAT: a Scene- and Object-Aware Transformer for Vision-and-Language Navigation	7357
<i>Harsh Agrawal, Dhruv Batra, Stefan Lee, Arjun Majumdar, Abhinav Moudgil</i>	
A Normative and Biologically Plausible Algorithm for Independent Component Analysis	7368
<i>Yanis Bahroun, Dmitri Chklovskii, Anirvan Sengupta</i>	
Regret Bounds for Gaussian-Process Optimization in Large Domains	7385
<i>Andreas Krause, Bernhard Scholkopf, Manuel Wuethrich</i>	
Deeply Shared Filter Bases for Parameter-Efficient Convolutional Neural Networks	7397
<i>Woochul Kang, Daeyeon Kim</i>	
On Optimal Robustness to Adversarial Corruption in Online Decision Problems	7409
<i>Shinji Ito</i>	
Directed Spectrum Measures Improve Latent Network Models of Neural Populations	7421
<i>David Carlson, Kafui Dzirasa, Neil Gallagher</i>	
Uncertainty-Based Offline Reinforcement Learning with Diversified Q-Ensemble	7436
<i>Gaon An, Jang-Hyun Kim, Seungyong Moon, Hyun Oh Song</i>	
Distribution-Free Inference for Regression: Discrete, Continuous, and in Between	7448
<i>Rina Barber, Yonghoon Lee</i>	
Statistical Inference with M-Estimators on Adaptively Collected Data	7460
<i>Lucas Janson, Susan Murphy, Kelly Zhang</i>	
NeuroLKH: Combining Deep Learning Model with Lin-Kernighan-Helsgaun Heuristic for Solving the Traveling Salesman Problem	7472
<i>Zhiguang Cao, Wen Song, Liang Xin, Jie Zhang</i>	
LSH-SMILE: Locality Sensitive Hashing Accelerated Simulation and Learning	7484
<i>Chonghao Sima, Yexiang Xue</i>	
Meta-Learning with an Adaptive Task Scheduler	7497
<i>Chelsea Finn, Defu Lian, Mehrdad Mahdavi, Yu Wang, Ying Wei, Huaxiu Yao, Peilin Zhao</i>	

VOLUME 10

Neural Active Learning with Performance Guarantees	7510
<i>Pranjal Awasthi, Christoph Dann, Claudio Gentile, Ayush Sekhari, Zhilei Wang</i>	
A Gradient Method for Multilevel Optimization.....	7522
<i>Ryo Sato, Akiko Takeda, Mirai Tanaka</i>	
Edge Representation Learning with Hypergraphs	7534
<i>Jinheon Baek, Sung Ju Hwang, Jaehyeong Jo, Minki Kang, Dongki Kim, Seul Lee</i>	
One Question Answering Model for Many Languages with Cross-Lingual Dense Passage Retrieval	7547
<i>Akari Asai, Hanna Hajishirzi, Jungo Kasai, Xinyan Yu</i>	
LEADS: Learning Dynamical Systems that Generalize Across Environments.....	7561
<i>Ibrahim Ayed, Nicolas Baskiotis, Patrick Gallinari, Yuan Yin, Emmanuel De Bezenac</i>	
Stochastic: a Framework for General Stochastic Automatic Differentiation	7574
<i>Emile Krieken, Annette Ten Teije, Jakub Tomczak</i>	
Concentration Inequalities Under sub-Gaussian and Sub-exponential Conditions	7588
<i>Andreas Maurer, Massimiliano Pontil</i>	
Variance-Aware Off-Policy Evaluation with Linear Function Approximation.....	7598
<i>Quanquan Gu, Yifei Min, Tianhao Wang, Dongruo Zhou</i>	
A Provably Efficient Sample Collection Strategy for Reinforcement Learning.....	7611
<i>Alessandro Lazaric, Matteo Pirota, Jean Tarbouriech, Michal Valko</i>	
Improved Regret Bounds for Tracking Experts with Memory	7625
<i>Mark Herbster, James Robinson</i>	
Robustness of Graph Neural Networks at Scale	7637
<i>Aleksandar Bojchevski, Simon Geisler, Stephan Gunnemann, Tobias Schmidt, Daniel Zugner, Hakan Sirin</i>	
Random Noise Defense Against Query-Based Black-Box Attacks.....	7650
<i>Yanbo Fan, Zeyu Qin, Baoyuan Wu, Hongyuan Zha</i>	
SADGA: Structure-Aware Dual Graph Aggregation Network for Text-to-SQL	7664
<i>Ruichu Cai, Zhifeng Hao, Boyan Xu, Jinjie Yuan</i>	
Near-Optimal Offline Reinforcement Learning Via Double Variance Reduction	7677
<i>Yu Bai, Yu-Xiang Wang, Ming Yin</i>	
Joint Modeling of Visual Objects and Relations for Scene Graph Generation.....	7689
<i>Bingbing Ni, Meng Qu, Jian Tang, Minghao Xu</i>	
Going Beyond Linear Transformers with Recurrent Fast Weight Programmers.....	7703
<i>Robert Csordas, Kazuki Irie, Imanol Schlag, Jurgen Schmidhuber</i>	
Reinforced Few-Shot Acquisition Function Learning for Bayesian Optimization.....	7718
<i>Bing-Jing Hsieh, Ping-Chun Hsieh, Xi Liu</i>	
Forster Decomposition and Learning Halfspaces with Noise.....	7732
<i>Ilias Diakonikolas, Daniel Kane, Christos Tzamos</i>	

Cortico-Cerebellar Networks as Decoupling Neural Interfaces	7745
<i>Richard Apps, Ellen Boven, Joseph Pemberton, Rui Ponte Costa</i>	
To the Point: Correspondence-Driven Monocular 3D Category Reconstruction	7760
<i>Iasonas Kokkinos, Filippos Kokkinos</i>	
Proper Value Equivalence	7773
<i>Andre Barreto, Greg Farquhar, Christopher Grimm, David Silver, Satinder Singh</i>	
Challenges and Opportunities in High Dimensional Variational Inference.....	7787
<i>Michael Andersen, Alejandro Catalina, Akash Kumar Dhaka, Jonathan Huggins, Aki Vehtari, Manushi Welandawe</i>	
On the Expressivity of Markov Reward	7799
<i>David Abel, Will Dabney, Anna Harutyunyan, Mark Ho, Michael Littman, Doina Precup, Satinder Singh</i>	
One More Step Towards Reality: Cooperative Bandits with Imperfect Communication.....	7813
<i>Abhimanyu Dubey, Naomi Leonard, Udari Madhushani, Alex Pentland</i>	
Multi-Agent Reinforcement Learning in Stochastic Networked Systems.....	7825
<i>Longbo Huang, Yiheng Lin, Guannan Qu, Adam Wierman</i>	
Neural Scene Flow Prior	7838
<i>Jhony Kaesemodel Pontes, Xueqian Li, Simon Lucey</i>	
The Future is log-Gaussian: ResNets and Their Infinite-depth-and-width Limit at Initialization	7852
<i>Mufan Li, Mihai Nica, Dan Roy</i>	
Grammar-Based Grounded Lexicon Learning	7865
<i>Roger Levy, Jiayuan Mao, Freda Shi, Josh Tenenbaum, Jiajun Wu</i>	
Distributed Deep Learning in Open Collaborations	7879
<i>Alexander Borzunov, Alexey Bukhtiyarov, Michael Diskin, Yacine Jernite, Maxim Kashirin, Iliia Kobelev, Denis Mazur, Gennady Pekhimenko, Dmitry Popov, Dmitry V. Pyrkin, Max Ryabinin, Lucile Saulnier, Anton Sinitsin, Albert Villanova Del Moral, Thomas Wolf, Quentin Lhoest</i>	
Neural Ensemble Search for Uncertainty Estimation and Dataset Shift.....	7898
<i>Thomas Elsken, Chris C Holmes, Frank Hutter, Yee Teh, Sheheryar Zaidi, Arber Zela</i>	
Finding Bipartite Components in Hypergraphs	7912
<i>Peter Macgregor, He Sun</i>	
Hit and Lead Discovery with Explorative RL and Fragment-Based Molecule Generation.....	7924
<i>Doyeong Hwang, Sung Ju Hwang, Seul Lee, Seongok Ryu, Soojung Yang</i>	
Proxy Convexity: A Unified Framework for the Analysis of Neural Networks Trained by Gradient Descent	7937
<i>Spencer Frei, Quanquan Gu</i>	
Covariance-Aware Private Mean Estimation Without Private Covariance Estimation	7950
<i>Gavin Brown, Marco Gaboardi, Adam Smith, Jonathan Ullman, Lydia Zakynthinou</i>	
Label Consistency in Overfitted Generalized k-Means	7965
<i>Arash Amini, Linfan Zhang</i>	

Open-Set Label Noise Can Improve Robustness Against Inherent Label Noise	7978
<i>Bo An, Lue Tao, Hongxin Wei, Renchunzi Xie</i>	
The Complexity of Sparse Tensor PCA	7993
<i>Davin Choo, Tommaso D'Orsi</i>	
Learning to Elect	8006
<i>Cem Anil, Xuchan Bao</i>	
KALE Flow: a Relaxed KL Gradient Flow for Probabilities with Disjoint Support.....	8018
<i>Michael Arbel, Pierre Glaser, Arthur Gretton</i>	
When is Generalizable Reinforcement Learning Tractable?	8032
<i>Yuanzhi Li, Dhruv Malik, Pradeep Ravikumar</i>	
Relational Self-Attention: What's Missing in Attention for Video Understanding	8046
<i>Minsu Cho, Manjin Kim, Suha Kwak, Heeseung Kwon, Chunyu Wang</i>	
Towards Enabling Meta-Learning from Target Models.....	8060
<i>Le Gan, Su Lu, Han-Jia Ye, De-Chuan Zhan</i>	
A Near-Optimal Algorithm for Debiasing Trained Machine Learning Models	8072
<i>Ibrahim Alabdulmohsin, Mario Lucic</i>	
GENESIS-V2: Inferring Unordered Object Representations Without Iterative Refinement.....	8085
<i>Martin Engelcke, Oiwi Parker Jones, Ingmar Posner</i>	
How Data Augmentation Affects Optimization for Linear Regression.....	8095
<i>Boris Hanin, Yi Sun</i>	
An Exact Characterization of the Generalization Error for the Gibbs Algorithm	8106
<i>Gholamali Aminian, Yuheng Bu, Miguel Rodrigues, Laura Toni, Gregory Wornell</i>	
Subgaussian and Differentiable Importance Sampling for Off-Policy Evaluation and Learning	8119
<i>Alberto Maria Metelli, Marcello Restelli, Alessio Russo</i>	
Rethinking Gradient Sparsification as Total Error Minimization.....	8133
<i>Trambak Banerjee, Marco Canini, Aritra Dutta, Panos Kalnis, Ahmed M. Abdelmoniem, Atal Sahu</i>	
Approximate Optimization of Convex Functions with Outlier Noise	8147
<i>Anindya De, Sanjeev Khanna, Huan Li, Mohammadhesam Nikpeysalekde</i>	
Fair Classification with Adversarial Perturbations.....	8158
<i>L. Elisa Celis, Anay Mehrotra, Nisheeth Vishnoi</i>	
Distributed Saddle-Point Problems Under Data Similarity	8172
<i>Aleksandr Beznosikov, Alexander Gasnikov, Alexander Rogozin, Gesualdo Scutari</i>	
Combining Latent Space and Structured Kernels for Bayesian Optimization Over Combinatorial Spaces.....	8185
<i>Aryan Deshwal, Jana Doppa</i>	
Gradual Domain Adaptation Without Indexed Intermediate Domains.....	8201
<i>Wei-Lun Chao, Hong-You Chen</i>	

K-Level Reasoning for Zero-Shot Coordination in Hanabi.....	8215
<i>Brandon Cui, Jakob Foerster, Hengyuan Hu, Luis Pineda</i>	
Learning Markov State Abstractions for Deep Reinforcement Learning	8229
<i>Cameron Allen, Omer Gottesman, George Konidaris, Neev Parikh</i>	
Towards Deeper Deep Reinforcement Learning with Spectral Normalization	8242
<i>Nils Bjorck, Carla Gomes, Kilian Weinberger</i>	
Functionally Regionalized Knowledge Transfer for Low-Resource Drug Discovery	8256
<i>Junzhou Huang, Long-Kai Huang, Zhenhui (Jessie) Li, Ying Wei, Ding Xue, Huaxiu Yao</i>	
Memory-Efficient Approximation Algorithms for Max-k-Cut and Correlation Clustering	8269
<i>Vishnu Narayanan, James Saunderson, Nimita Shinde</i>	
Panoptic 3D Scene Reconstruction from a Single RGB Image	8282
<i>Manuel Dahmert, Angela Dai, Ji Hou, Matthias Niessner</i>	
Measuring Generalization with Optimal Transport	8294
<i>Ching-Yao Chuang, Kristjan Greenewald, Stefanie Jegelka, Youssef Mroueh, Antonio Torralba</i>	
Uniform Concentration Bounds Toward a Unified Framework for Robust Clustering.....	8307
<i>Saptarshi Chakraborty, Swagatam Das, Debolina Paul, Jason Xu</i>	
Learning Signal-Agnostic Manifolds of Neural Fields.....	8320
<i>Katie Collins, Yilun Du, Vincent Sitzmann, Josh Tenenbaum</i>	
Low-Dimensional Structure in the Space of Language Representations is Reflected in Brain Responses	8332
<i>Richard Antonello, Alexander Huth, Javier Turek, Vy Vo</i>	
On the Suboptimality of Thompson Sampling in High Dimensions	8345
<i>Richard Combes, Raymond Zhang</i>	
Learning Debaised and Disentangled Representations for Semantic Segmentation.....	8355
<i>Sanghyeok Chu, Bohyung Han, Dongwan Kim</i>	

VOLUME 11

Diversity Matters When Learning from Ensembles	8367
<i>Juho Lee, Yoonho Lee, Giung Nam, Jongmin Yoon</i>	
Locally Valid and Discriminative Prediction Intervals for Deep Learning Models	8378
<i>Zhen Lin, Jimeng Sun, Shubhendu Trivedi</i>	
Personalized Federated Learning with Gaussian Processes.....	8392
<i>Idan Achituve, Gal Chechik, Ethan Fetaya, Aviv Navon, Aviv Shamsian</i>	
Risk Bounds for Over-Parameterized Maximum Margin Classification on Sub-Gaussian Mixtures	8407
<i>Mikhail Belkin, Yuan Cao, Quanquan Gu</i>	
Implicit SVD for Graph Representation Learning.....	8419
<i>Sami Abu-El-Haija, Valentino Crespi, Aram Galstyan, Hesham Mostafa, Marcel Nassar, Greg Ver Steeg</i>	

Offline Model-Based Adaptable Policy Learning	8432
<i>Xiong-Hui Chen, Qingyang Li, Fan-Ming Luo, Zhiwei Qin, Wenjie Shang, Jieping Ye, Yang Yu</i>	
Multilingual Pre-Training with Universal Dependency Learning	8444
<i>Zuchao Li, Kailai Sun, Hai Zhao</i>	
Parameter-Free HE-Friendly Logistic Regression.....	8457
<i>Junyoung Byun, Jaewook Lee, Woojin Lee</i>	
Active Clustering for Labeling Training Data.....	8469
<i>Quentin Lutz, Alex Scott, Maya Stein, Elie De Panafieu</i>	
Exploring Social Posterior Collapse in Variational Autoencoder for Interaction Modeling	8481
<i>Chen Tang, Masayoshi Tomizuka, Wei Zhan</i>	
Ensembling Graph Predictions for AMR Parsing.....	8495
<i>Ramon Fernandez Astudillo, Thanh Lam Hoang, Yufang Hou, Young-Suk Lee, Vanessa Lopez, Lam Nguyen, Dzong Phan, Gabriele Picco</i>	
On the Interplay Between Data Structure and Loss Function in Classification Problems.....	8506
<i>Giulio Biroli, Marylou Gabrie, Levent Sagun, Stephane D'Ascoli</i>	
Near-Optimal Offline and Streaming Algorithms for Learning Non-Linear Dynamical Systems	8518
<i>Prateek Jain, Suhas Kowshik, Dheeraj Nagaraj, Praneeth Netrapalli</i>	
Mixture Proportion Estimation and PU Learning:A Modern Approach.....	8532
<i>Sivaraman Balakrishnan, Saurabh Garg, Zachary Lipton, Alexander Smola, Yifan Wu</i>	
Escape Saddle Points by a Simple Gradient-Descent Based Algorithm.....	8545
<i>Tongyang Li, Chenyi Zhang</i>	
AC/DC: Alternating Compressed/DeCompressed Training of Deep Neural Networks	8557
<i>Dan Alistarh, Eugenia Iofinova, Alexandra Peste, Adrian Vladu</i>	
HyperSPNs: Compact and Expressive Probabilistic Circuits	8571
<i>Stefano Ermon, Dorsa Sadigh, Andy Shih</i>	
Scaling Vision with Sparse Mixture of Experts.....	8583
<i>Neil Houlsby, Rodolphe Jenatton, Daniel Keysers, Basil Mustafa, Maxim Neumann, Joan Puigcerver, Carlos Riquelme, Andre Susano Pinto</i>	
Two-Sided Fairness in Rankings Via Lorenz Dominance.....	8596
<i>Jamal Atif, Sam Corbett-Davies, Virginie Do, Nicolas Usunier</i>	
Stability & Generalisation of Gradient Descent for Shallow Neural Networks Without the Neural Tangent Kernel	8609
<i>Ilya Kuzborskij, Dominic Richards</i>	
Adversarial Intrinsic Motivation for Reinforcement Learning	8622
<i>Ishan Durugkar, Scott Niekum, Peter Stone, Mauricio Tec</i>	
Machine Learning for Variance Reduction in Online Experiments	8637
<i>Dominic Coey, Matt Goldman, Yongyi Guo, Mikael Konutgan, Wenting Li, Chris Schoener</i>	

L2ight: Enabling On-Chip Learning for Optical Neural Networks Via Efficient In-situ Subspace Optimization.....	8649
<i>Ray Chen, Chenghao Feng, Jiaqi Gu, Zixuan Jiang, David Pan, Hanqing Zhu</i>	
Towards Gradient-Based Bilevel Optimization with Non-convex Followers and Beyond	8662
<i>Risheng Liu, Yaohua Liu, Shangzhi Zeng, Jin Zhang</i>	
Multi-Facet Clustering Variational Autoencoders.....	8676
<i>Fabian Falck, Chris C Holmes, George Nicholson, Matthew Willetts, Christopher Yau, Haoting Zhang</i>	
Synthetic Design: an Optimization Approach to Experimental Design with Synthetic Controls.....	8691
<i>Nick Doudchenko, Khashayar Khosravi, Sebastien Lahaie, Miles Lubin, Vahab Mirrokni, Jean Pouget-Abadie, Jann Spiess, Guido Imbens</i>	
Ranking Policy Decisions.....	8702
<i>Hana Chockler, Daniel Kroening, Hadrien Pouget, Youcheng Sun</i>	
Searching the Search Space of Vision Transformer	8714
<i>Hongyang Chao, Minghao Chen, Jianlong Fu, Haibin Ling, Bei Liu, Bolin Ni, Houwen Peng, Kan Wu</i>	
Relative Stability Toward Diffeomorphisms Indicates Performance in Deep Nets.....	8727
<i>Alessandro Favero, Mario Geiger, Leonardo Petrini, Matthieu Wyart</i>	
Raw Nav-Merge Seismic Data to Subsurface Properties with MLP Based Multi-Modal Information Unscrambler	8740
<i>Anu Chandran, Aditya Desai, Menal Gupta, Anshumali Shrivastava, Antoine Vial-Aussavy, Zhaozhuo Xu</i>	
Inverse Problems Leveraging Pre-Trained Contrastive Representations.....	8753
<i>Alexandros Dimakis, Matt Jordan, Sriram Ravula, Georgios Smyrnis</i>	
The Unbalanced Gromov Wasserstein Distance: Conic Formulation and Relaxation	8766
<i>Gabriel Peyre, Thibault Sejourne, Francois-Xavier Vialard</i>	
Diffusion Models Beat GANs on Image Synthesis	8780
<i>Prafulla Dhariwal, Alexander Nichol</i>	
Learning MDPs from Features: Predict-Then-Optimize for Sequential Decision Making by Reinforcement Learning	8795
<i>Haipeng Chen, Finale Doshi-Velez, Andrew Perrault, Sanket Shah, Milind Tambe, Kai Wang</i>	
A Closer Look at the Worst-Case Behavior of Multi-armed Bandit Algorithms	8807
<i>Anand Kalvit, Assaf Zeevi</i>	
SAPE: Spatially-Adaptive Progressive Encoding for Neural Optimization.....	8820
<i>Daniel Cohen-Or, Raja Giryes, Amir Hertz, Or Perel, Olga Sorkine-Hornung</i>	
A Biased Graph Neural Network Sampler with Near-Optimal Regret.....	8833
<i>Quan Gan, Le Song, David Wipf, Qingru Zhang</i>	
Equilibrium Refinement for the Age of Machines: the One-Sided Quasi-Perfect Equilibrium	8845
<i>Gabriele Farina, Tuomas Sandholm</i>	
Interpreting Representation Quality of DNNs for 3D Point Cloud Processing	8857
<i>Dongrui Liu, Qihan Ren, Wen Shen, Quanshi Zhang</i>	

How Fine-Tuning Allows for Effective Meta-Learning.....	8871
<i>Kurtland Chua, Jason Lee, Qi Lei</i>	
Cooperative Stochastic Bandits with Asynchronous Agents and Constrained Feedback.....	8885
<i>Yu-Zhen Janice Chen, Mohammad Hajiesmaili, John C. S. Lui, Stephen Pasteris, Don Towsley, Lin Yang</i>	
Multiple Descent: Design Your Own Generalization Curve	8898
<i>Mikhail Belkin, Lin Chen, Amin Karbasi, Yifei Min</i>	
On Empirical Risk Minimization with Dependent and Heavy-Tailed Data	8913
<i>Krishnakumar Balasubramanian, Murat Erdogdu, Abhishek Roy</i>	
Gone Fishing: Neural Active Learning with Fisher Embeddings.....	8927
<i>Jordan Ash, Surbhi Goel, Sham Kakade, Akshay Krishnamurthy</i>	
On Riemannian Optimization Over Positive Definite Matrices with the Bures-Wasserstein Geometry	8940
<i>Junbin Gao, Andi Han, Pratik Kumar Jawanpuria, Bamdev Mishra</i>	
Refining Language Models with Compositional Explanations	8954
<i>Ying Chen, Xisen Jin, Xiang Ren, Huihan Yao, Qinyuan Ye</i>	
Going Beyond Linear RL: Sample Efficient Neural Function Approximation	8968
<i>Baihe Huang, Kaixuan Huang, Sham Kakade, Jason Lee, Qi Lei, Runzhe Wang, Jiaqi Yang</i>	
Scalable Neural Data Server: A Data Recommender for Transfer Learning.....	8984
<i>David Acuna, Tianshi Cao, Sasha (Alexandre) Doubov, Sanja Fidler</i>	
What Can Linearized Neural Networks Actually Say About Generalization?	8998
<i>Pascal Frossard, Seyed-Mohsen Moosavi-Dezfooli, Guillermo Ortiz-Jimenez</i>	
CATs: Cost Aggregation Transformers for Visual Correspondence	9011
<i>Seokju Cho, Sunghwan Hong, Sangryul Jeon, Seungryong Kim, Yunsung Lee, Kwanghoon Sohn</i>	
Asynchronous Stochastic Optimization Robust to Arbitrary Delays.....	9024
<i>Alon Cohen, Amit Daniely, Yoel Drori, Tomer Koren, Mariano Schain</i>	
Consistent Non-Parametric Methods for Maximizing Robustness.....	9036
<i>Robi Bhattacharjee, Kamalika Chaudhuri</i>	
Generalizable Multi-Linear Attention Network	9049
<i>Tao Jin, Zhou Zhao</i>	
Labeling Trick: A Theory of Using Graph Neural Networks for Multi-Node Representation Learning	9061
<i>Long Jin, Pan Li, Kai Wang, Yinglong Xia, Muhan Zhang</i>	
SUPER-ADAM: Faster and Universal Framework of Adaptive Gradients.....	9074
<i>Heng Huang, Feihu Huang, Junyi Li</i>	
General Nonlinearities in SO(2)-Equivariant CNNs	9086
<i>Daniel Franzen, Michael Wand</i>	
Denosing Normalizing Flow	9099
<i>Christian Horvat, Jean-Pascal Pfister</i>	

Attention Over Learned Object Embeddings Enables Complex Visual Reasoning	9112
<i>Matt Botvinick, David Ding, Felix Hill, Malcolm Reynolds, Adam Santoro</i>	
Differentially Private Federated Bayesian Optimization with Distributed Exploration	9125
<i>Zhongxiang Dai, Patrick Jaillet, Bryan Kian Hsiang Low</i>	
Differentiable Learning Under Triage	9140
<i>Abir De, Nastaran Okati, Manuel Rodriguez</i>	
A New Theoretical Framework for Fast and Accurate Online Decision-Making	9152
<i>Nicolò Cesa-Bianchi, Tom Cesari, Yishay Mansour, Vianney Perchet</i>	
When Expressivity Meets Trainability: Fewer than N Neurons Can Work.....	9167
<i>Mingyi Hong, Zhi-Quan Luo, Ruoyu Sun, Jiawei Zhang, Yushun Zhang</i>	
Analyzing the Confidentiality of Undistillable Teachers in Knowledge Distillation	9181
<i>Peter Bearel, Yao Fu, Souvik Kundu, Massoud Pedram, Qirui Sun</i>	
High Probability Complexity Bounds for Line Search Based on Stochastic Oracles.....	9193
<i>Billy Jin, Katya Scheinberg, Miaolan Xie</i>	
Pay Attention to MLPs	9204
<i>Zihang Dai, Quoc V Le, Hanxiao Liu, David So</i>	

VOLUME 12

An Image is Worth More than a Thousand Words: Towards Disentanglement in the Wild	9216
<i>Niv Cohen, Aviv Gabbay, Yedid Hoshen</i>	
Dynamics of Stochastic Momentum Methods on Large-Scale, Quadratic Models	9229
<i>Courtney Paquette, Elliot Paquette</i>	
Adversarial Examples in Multi-Layer Random ReLU Networks.....	9241
<i>Peter Bartlett, Sebastien Bubeck, Yeshwanth Cherapanamjeri</i>	
Efficient Statistical Assessment of Neural Network Corruption Robustness	9253
<i>Teddy Furon, Mathias Rousset, Karim Tit</i>	
A Highly-Efficient Group Elastic Net Algorithm with an Application to Function-On-Scalar Regression	9264
<i>Tobia Boschi, Francesca Chiaromonte, Matthew Reimherr</i>	
Hierarchical Clustering: $O(1)$ -Approximation for Well-Clustered Graphs	9278
<i>Bogdan-Adrian Manghiuc, He Sun</i>	
Realistic Evaluation of Transductive Few-Shot Learning	9290
<i>Ismail Ben Ayed, Malik Boudiaf, Pablo Piantanida, Olivier Veilleux</i>	
Qu-ANTI-zation: Exploiting Quantization Artifacts for Achieving Adversarial Outcomes	9303
<i>Tudor Dumitras, Sanghyun Hong, Yigitcan Kaya, Michael-Andrei Panaitescu-Liess</i>	
Differentially Private Stochastic Optimization: New Results in Convex and Non-Convex Settings	9317
<i>Raef Bassily, Cristobal Guzman, Michael Menart</i>	
TacticZero: Learning to Prove Theorems from Scratch with Deep Reinforcement Learning	9330
<i>Amir Dezfouli, Michael Norrish, Christian Walder, Minchao Wu</i>	

Integrating Tree Path in Transformer for Code Representation	9343
<i>Zhi Jin, Ge Li, Han Peng, Wenhan Wang, Yunfei Zhao</i>	
Twins: Revisiting the Design of Spatial Attention in Vision Transformers	9355
<i>Xiangxiang Chu, Haibing Ren, Chunhua Shen, Zhi Tian, Yuqing Wang, Xiaolin Wei, Huaxia Xia, Bo Zhang</i>	
Evaluating State-Of-the-Art Classification Models Against Bayes Optimality	9367
<i>Richard Socher, Ryan Theisen, Lav Varshney, Huan Wang, Caiming Xiong</i>	
Data-Efficient Instance Generation from Instance Discrimination.....	9378
<i>Yujun Shen, Yinghao Xu, Ceyuan Yang, Bolei Zhou</i>	
Reliable Post Hoc Explanations: Modeling Uncertainty in Explainability	9391
<i>Anna Hilgard, Himabindu Lakkaraju, Sameer Singh, Dylan Slack</i>	
Learning Graph Models for Retrosynthesis Prediction	9405
<i>Regina Barzilay, Charlotte Bunne, Connor Coley, Andreas Krause, Vignesh Ram Somnath</i>	
Differentiable Equilibrium Computation with Decision Diagrams for Stackelberg Models of Combinatorial Congestion Games.....	9416
<i>Kengo Nakamura, Shinsaku Sakaue</i>	
Inverse Optimal Control Adapted to the Noise Characteristics of the Human Sensorimotor System	9429
<i>Constantin Rothkopf, Matthias Schultheis, Dominik Straub</i>	
Deep Neural Networks as Point Estimates for Deep Gaussian Processes	9443
<i>Nicolas Durrande, Vincent Dutordoir, Carl Henrik Ek, Zoubin Ghahramani, James Hensman, Mark Van Der Wilk</i>	
Locality Defeats the Curse of Dimensionality in Convolutional Teacher-Student Scenarios	9456
<i>Francesco Cagnetta, Alessandro Favero, Matthieu Wyart</i>	
Causal Identification with Matrix Equations	9468
<i>Elias Bareinboim, Sanghack Lee</i>	
Private and Non-Private Uniformity Testing for Ranking Data	9480
<i>Robert Busa-Fekete, Dimitris Fotakis, Emmanouil Zampetakis</i>	
Model-Based Reinforcement Learning Via Imagination with Derived Memory	9493
<i>Jianguo Chen, Jianye Hao, Shengbo Li, Wulong Liu, Ping Luo, Yao Mu, Bin Wang, Chongjie Zhang, Guangxiang Zhu, Yuzheng Zhuang</i>	
Compositional Transformers for Scene Generation	9506
<i>Dor Arad Hudson, Larry Zitnick</i>	
An Exponential Lower Bound for Linearly Realizable MDP with Constant Suboptimality Gap	9521
<i>Sham Kakade, Ruosong Wang, Yuanhao Wang</i>	
Combating Noise: Semi-Supervised Learning by Region Uncertainty Quantification	9534
<i>Ye Guo, Ya-Li Li, Shengjin Wang, Zhenyu Wang</i>	
Reducing the Covariate Shift by Mirror Samples in Cross Domain Alignment.....	9546
<i>Longjun Cai, Yin Zhao, Minquan Wang</i>	
Permutation-Invariant Variational Autoencoder for Graph-Level Representation Learning	9559
<i>Djork-Arne Clevert, Frank Noe, Robin Winter</i>	

Causal Abstractions of Neural Networks	9574
<i>Atticus Geiger, Thomas Icard, Hanson Lu, Christopher Potts</i>	
Conic Blackwell Algorithm: Parameter-Free Convex-Concave Saddle-Point Solving.....	9587
<i>Julien Grand-Clement, Christian Kroer</i>	
3DP3: 3D Scene Perception Via Probabilistic Programming.....	9600
<i>Marco Cusumano-Towner, Austin Garrett, Matin Ghavamizadeh, Nishad Gothoskar, Dan Gutfreund, Vikash Mansinghka, Falk Pollok, Josh Tenenbaum, Ben Zinberg</i>	
Novel Upper Bounds for the Constrained Most Probable Explanation Task	9613
<i>Vibhav Gogate, Tahrira Rahman, Sara Rouhani</i>	
Why Spectral Normalization Stabilizes GANs: Analysis and Improvements	9625
<i>Giulia Fanti, Zinan Lin, Vyas Sekar</i>	
(Implicit) ² : Implicit Layers for Implicit Representations	9639
<i>Shaojie Bai, Zhichun Huang, J. Zico Kolter</i>	
Mean-Based Best Arm Identification in Contaminated Stochastic Bandits under Reward Contamination	9651
<i>Pin-Yu Chen, Payel Das, Arpan Mukherjee, Ali Tajer</i>	
MADE: Exploration Via Maximizing Deviation from Explored Regions	9663
<i>Joseph Gonzalez, Jiantao Jiao, Paria Rashidinejad, Stuart Russell, Yuandong Tian, Tianjun Zhang</i>	
Variational Automatic Curriculum Learning for Sparse-Reward Cooperative Multi-Agent Problems.....	9681
<i>Jiayu Chen, Huimin Ma, Jiaming Song, Yu Wang, Yi Wu, Yuanfan Xu, Huazhong Yang, Yuanxin Zhang</i>	
Align Before Fuse: Vision and Language Representation Learning with Momentum Distillation	9694
<i>Akhilesh Gotmare, Steven Chu Hong Hoi, Shafiq Joty, Junnan Li, Ramprasaath Selvaraju, Caiming Xiong</i>	
Variational Model Inversion Attacks.....	9706
<i>Yan Fu, Ashish Khisti, Ke Li, Alireza Makhzani, Kuan-Chieh Wang, Richard Zemel</i>	
Graph Neural Networks with Adaptive Residual	9720
<i>Jiayuan Ding, Wei Jin, Xiaorui Liu, Zitao Liu, Yao Ma, Jiliang Tang, Han Xu</i>	
Efficient Active Learning for Gaussian Process Classification by Error Reduction	9734
<i>Francis Alexander, Edward Dougherty, Xiaoning Qian, Byung-Jun Yoon, Guang Zhao</i>	
Non-Asymptotic Analysis for Two Time-Scale TDC with General Smooth Function Approximation.....	9747
<i>Yue Wang, Yi Zhou, Shaofeng Zou</i>	
A Little Robustness Goes a Long Way: Leveraging Robust Features for Targeted Transfer Attacks	9759
<i>Garrett Kenyon, Melanie Mitchell, Jacob Springer</i>	
TriBERT: Human-Centric Audio-visual Representation Learning	9774
<i>Tanzila Rahman, Leonid Sigal, Mengyu Yang</i>	
How Does a Neural Network's Architecture Impact Its Robustness to Noisy Labels?.....	9788
<i>Jimmy Ba, John Dickerson, Jingling Li, Keyulu Xu, Mozhi Zhang</i>	

Calibration and Consistency of Adversarial Surrogate Losses.....	9804
<i>Pranjal Awasthi, Natalie Frank, Anqi Mao, Mehryar Mohri, Yutao Zhong</i>	
The Value of Information When Deciding What to Learn.....	9816
<i>Dilip Arumugam, Benjamin Van Roy</i>	
Co-Adaptation of Algorithmic and Implementational Innovations in Inference-Based Deep Reinforcement Learning.....	9828
<i>Hiroki Furuta, Shixiang (Shane) Gu, Tadashi Kozuno, Yutaka Matsuo, Tatsuya Matsushima</i>	
Can fMRI Reveal the Representation of Syntactic Structure in the Brain?.....	9843
<i>Aniketh Janardhan Reddy, Leila Wehbe</i>	
Robust Implicit Networks Via Non-Euclidean Contractions.....	9857
<i>Francesco Bullo, Alexander Davydov, Saber Jafarpour, Anton Proskurnikov</i>	
A Kernel-Based Test of Independence for Cluster-correlated Data.....	9869
<i>Hongjiao Liu, Anna Plantinga, Michael Wu, Yunhua Xiang</i>	
Efficient Methods for Gaussian Markov Random Fields Under Sparse Linear Constraints.....	9882
<i>David Bolin, Jonas Wallin</i>	
Sparse is Enough in Scaling Transformers.....	9895
<i>Aakanksha Chowdhery, Wojciech Gajewski, Sebastian Jaszczur, Lukasz Kaiser, Jonni Kanerva, Henryk Michalewski, Afroz Mohiuddin</i>	
Sparse Training Via Boosting Pruning Plasticity with Neuroregeneration.....	9908
<i>Zahra Atashgahi, Tianlong Chen, Xiaohan Chen, Huanyu Kou, Shiwei Liu, Decebal Constantin Mocanu, Mykola Pechenizkiy, Li Shen, Zhangyang Wang, Lu Yin</i>	
Low-Fidelity Video Encoder Optimization for Temporal Action Localization.....	9923
<i>Bernard Ghanem, Brais Martinez, Juan Manuel Perez Rua, Mengmeng Xu, Xiatian Zhu</i>	
On Provable Benefits of Depth in Training Graph Convolutional Networks.....	9936
<i>Weilin Cong, Mehrdad Mahdavi, Morteza Ramezani</i>	
Practical Near Neighbor Search Via Group Testing.....	9950
<i>Benjamin Coleman, Joshua Engels, Anshumali Shrivastava</i>	
Baby Intuitions Benchmark (BIB): Discerning the Goals, Preferences, and Actions of Others.....	9963
<i>Moira R Dillon, Kanishk Gandhi, Brenden Lake, Gala Stojnic</i>	
Neural Hybrid Automata: Learning Dynamics with Multiple Modes and Stochastic Transitions.....	9977
<i>Hajime Asama, Sanghyuk Chun, Animesh Garg, Stefano Massaroli, Seong Joon Oh, Jinkyoo Park, Michael Poli, Luca Scimeca, Atsushi Yamashita</i>	
Fast Projection onto the Capped Simplex with Applications to Sparse Regression in Bioinformatics.....	9990
<i>Man Shun Ang, Kun Huang, Nianjun Liu, Jianzhu Ma, Yijie Wang</i>	
The Many Faces of Adversarial Risk.....	10000
<i>Varun Jog, Muni Sreenivas Pydi</i>	
Meta-Adaptive Nonlinear Control: Theory and Algorithms.....	10013
<i>Kamyar Aizzadenesheli, Soon-Jo Chung, Michael O'Connell, Guanya Shi, Yisong Yue</i>	
Compositional Reinforcement Learning from Logical Specifications.....	10026
<i>Rajeev Alur, Suguman Bansal, Osbert Bastani, Kishor Jothimurugan</i>	

Differentiable Quality Diversity	10040
<i>Matthew Fontaine, Stefanos Nikolaidis</i>	
Credit Assignment Through Broadcasting a Global Error Vector.....	10053
<i>L F Abbott, Sueyeon Chung, David Clark</i>	

VOLUME 13

An Online Method for a Class of Distributionally Robust Optimization with Non-Convex Objectives.....	10067
<i>Zhishuai Guo, Rong Jin, Qi Qi, Yi Xu, Tianbao Yang</i>	
A Single Gradient Step Finds Adversarial Examples on Random Two-Layers Neural Networks	10081
<i>Sebastien Bubeck, Yeshwanth Cherapanamjeri, Gauthier Gidel, Remi Tachet Des Combes</i>	
Parameterized Knowledge Transfer for Personalized Federated Learning.....	10092
<i>Song Guo, Xiaosong Ma, Haozhao Wang, Feijie Wu, Wenchao Xu, Jie Zhang</i>	
Contrastively Disentangled Sequential Variational Autoencoder	10105
<i>Junwen Bai, Carla Gomes, Weiran Wang</i>	
Recursive Causal Structure Learning in the Presence of Latent Variables and Selection Bias	10119
<i>Sina Akbari, Amiremad Ghassami, Negar Kiyavash, Ehsan Mokhtarian</i>	
Generalization Error Rates in Kernel Regression: The Crossover from the Noiseless to Noisy Regime	10131
<i>Hugo Cui, Florent Krzakala, Bruno Loureiro, Lenka Zdeborova</i>	
Learning Gaussian Mixtures with Generalized Linear Models: Precise Asymptotics in High-Dimensions.....	10144
<i>Cedric Gerbelot, Florent Krzakala, Bruno Loureiro, Alessandro Pocco, Gabriele Sicuro, Lenka Zdeborova</i>	
Spectral Embedding for Dynamic Networks with Stability Guarantees.....	10158
<i>Ian Gallagher, Andrew Jones, Patrick Rubin-Delanchy</i>	
Infinite Time Horizon Safety of Bayesian Neural Networks.....	10171
<i>Krishnendu Chatterjee, Thomas Henzinger, Mathias Lechner, Đorđe Žikelić</i>	
Towards Understanding Retrosynthesis by Energy-Based Models.....	10186
<i>Bo Dai, Hanjun Dai, Steven Kearnes, Li Li, Ruoxi Sun</i>	
List-Decodable Mean Estimation in Nearly-PCA Time.....	10195
<i>Ilias Diakonikolas, Daniel Kane, Daniel Kongsgaard, Jerry Li, Kevin Tian</i>	
Distributed Zero-Order Optimization Under Adversarial Noise	10209
<i>Arya Akhavan, Massimiliano Pontil, Alexandre Tsybakov</i>	
Reliable Estimation of KL Divergence Using a Discriminator in Reproducing Kernel Hilbert Space	10221
<i>Jennifer Dy, Sandesh Ghimire, Aria Masoomi</i>	
Latent Matters: Learning Deep State-Space Models	10234
<i>Botond Cseke, Alexej Klushyn, Richard Kurle, Maximilian Soelch, Patrick Van Der Smagt</i>	
On the Estimation Bias in Double Q-Learning.....	10246
<i>Jianglun Chen, Beining Han, Hao Hu, Zhizhou Ren, Chongjie Zhang, Guangxiang Zhu</i>	

Mitigating Forgetting in Online Continual Learning with Neuron Calibration	10260
<i>Ping Li, Haiyan Yin, Peng Yang</i>	
Escaping Saddle Points with Compressed SGD	10273
<i>Dmitrii Avdiukhin, Grigory Yaroslavtsev</i>	
Non-Gaussian Gaussian Processes for Few-Shot Regression	10285
<i>Andrzej Bedychaj, Aleksandra Nowak, Massimiliano Patacchiola, Marcin Sendera, Przemyslaw Spurek, Jacek Tabor, Tomasz Trzcinski, Maciej Zieba</i>	
Believe What You See: Implicit Constraint Approach for Offline Multi-Agent Reinforcement Learning	10299
<i>Li Chenghao, Gao Huang, Xiaoteng Ma, Jun Yang, Yiqin Yang, Qiyuan Zhang, Qianchuan Zhao, Zewu Zheng</i>	
Online Learning in Periodic Zero-Sum Games	10313
<i>Tanner Fiez, Georgios Piliouras, Lillian Ratliff, Ryann Sim, Stratis Skoulakis</i>	
K-Net: Towards Unified Image Segmentation	10326
<i>Kai Chen, Chen Change Loy, Jiangmiao Pang, Wenwei Zhang</i>	
Pareto-Optimal Learning-Augmented Algorithms for Online Conversion Problems	10339
<i>Mohammad Hajiesmaili, Russell Lee, Bo Sun, Danny Tsang, Adam Wierman</i>	
Dynaboard: An Evaluation-As-A-Service Platform for Holistic Next-Generation Benchmarking	10351
<i>Kawin Ethayarajh, Somya Jain, Robin Jia, Douwe Kiela, Zhiyi Ma, Christopher Potts, Tristan Thrush, Adina Williams, Ledell Wu</i>	
NTopo: Mesh-Free Topology Optimization Using Implicit Neural Representations	10368
<i>Stelian Coros, Yue Li, Bernhard Thomaszewski, Jonas Zehnder</i>	
Generalization Bounds for (Wasserstein) Robust Optimization	10382
<i>Yang An, Rui Gao</i>	
Faster Matchings Via Learned Duals	10393
<i>Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, Sergei Vassilvitskii</i>	
Online Learning in MDPs with Linear Function Approximation and Bandit Feedback.	10407
<i>Gergely Neu, Julia Olkhovskaya</i>	
Learning Collaborative Policies to Solve NP-Hard Routing Problems	10418
<i>Minsu Kim, Jinkyoo Park, Joungho Kim</i>	
Efficient Mirror Descent Ascent Methods for Nonsmooth Minimax Problems	10431
<i>Heng Huang, Feihu Huang, Xidong Wu</i>	
CO-PILOT: COllaborative Planning and reInforcement Learning on sub-Task Curriculum	10444
<i>Shuang Ao, Jing Jiang, Guodong Long, Qinghua Lu, Tianyi Zhou, Liming Zhu</i>	
Modality-Agnostic Topology Aware Localization	10457
<i>Hanno Ackermann, Daniel Dijkman, Farhad Ghazvinian Zanjani, Ilia Karmanov, Simone Merlin, Fatih Porikli, Max Welling</i>	
Scalable Quasi-Bayesian Inference for Instrumental Variable Regression	10469
<i>Tongzheng Ren, Ziyu Wang, Yuhao Zhou, Jun Zhu</i>	

Kernel Identification Through Transformers.....	10483
<i>Ian Davies, Nicolas Durrande, Vidhi Lalchand, Carl Edward Rasmussen, Fergus Simpson, Alessandro Vullo</i>	
Curriculum Design for Teaching Via Demonstrations: Theory and Applications.....	10496
<i>Rati Devidze, Parameswaran Kamalaruban, Adish Singla, Gaurav Yengera</i>	
Revenue Maximization Via Machine Learning with Noisy Data.....	10510
<i>Ellen Vitercik, Tom Yan</i>	
Exploiting Data Sparsity in Secure Cross-Platform Social Recommendation.....	10524
<i>Chaochao Chen, Jinming Cui, Wang Li, Lingjuan Lyu, Carl Yang</i>	
Parallelizing Thompson Sampling.....	10535
<i>Amin Karbasi, Vahab Mirrokni, Mohammad Shadravan</i>	
Dynamic Causal Bayesian Optimization	10549
<i>Virginia Aglietti, Theodoros Damoulas, Neil Dhir, Javier Gonzalez</i>	
Local Differential Privacy for Regret Minimization in Reinforcement Learning	10561
<i>Evrard Garcelon, Vianney Perchet, Ciara Pike-Burke, Matteo Pirota</i>	
Emergent Discrete Communication in Semantic Spaces.....	10574
<i>Siddharth Agrawal, Dana Hughes, Michael Lewis, Huao Li, Julie A Shah, Katia Sycara, Mycal Tucker</i>	
Drop, Swap, and Generate: a Self-Supervised Approach for Generating Neural Activity	10587
<i>Mehdi Azabou, Max Dabagia, Eva Dyer, Mohammad Gheshlaghi Azar, Keith Hengen, Chi-Heng Lin, Ran Liu, Michal Valko</i>	
Equivariant Manifold Flows.....	10600
<i>Christopher De Sa, Qingxuan Jiang, Isay Katsman, Derek Lim, Ser Nam Lim, Aaron Lou</i>	
Scalable Bayesian GPFA with Automatic Relevance Determination and Discrete Noise Models	10613
<i>Guillaume Hennequin, Kristopher Jensen, Ta-Chu Kao, Jasmine Stone</i>	
Recurrence Along Depth: Deep Convolutional Neural Networks with Recurrent Layer Aggregation	10627
<i>Yanwen Fang, Guodong Li, Jingyu Zhao</i>	
Independent Prototype Propagation for Zero-Shot Compositionality	10641
<i>Doina Bucur, Gertjan Burghouts, Frank Ruis</i>	
Universal Graph Convolutional Networks	10654
<i>Jiawei Han, Dongxiao He, Cuiying Huo, Di Jin, Rui Wang, Xiao Wang, Zhizhi Yu</i>	
Adversarial Feature Desensitization.....	10665
<i>Kartik Ahuja, Pouya Bashivan, Reza Bayat, Mojtaba Faramarzi, Adam Ibrahim, Touraj Laleh, Blake Richards, Irina Rish</i>	
Few-Shot Data-Driven Algorithms for Low Rank Approximation.....	10678
<i>Piotr Indyk, Tal Wagner, David Woodruff</i>	
Neural-PIL: Neural Pre-Integrated Lighting for Reflectance Decomposition.....	10691
<i>Jonathan Barron, Mark Boss, Raphael Braun, Varun Jampani, Hendrik Pa Lensch, Ce Liu</i>	
Asymptotics of the Bootstrap Via Stability with Applications to Inference with Model Selection	10705
<i>Morgane Austern, Vasilis Syrgkanis</i>	

Dynamic Influence Maximization	10718
<i>Binghui Peng</i>	
Risk Monotonicity in Statistical Learning	10732
<i>Zakaria Mhammedi</i>	
Information is Power: Intrinsic Control Via Information Capture	10745
<i>Glen Berseth, John Co-Reyes, Chelsea Finn, Danijar Hafner, Sergey Levine, Nicholas Rhinehart, Jenny Wang</i>	
Extracting Deformation-Aware Local Features by Learning to Deform	10759
<i>Felipe Chamone, Renato Martins, Erickson Nascimento, Guilherme Potje</i>	
Object-Centric Representation Learning with Generative Spatial-Temporal Factorization	10772
<i>Robert Fisher, Wenbin Hu, Nanbo Li, Muhammad Ahmed Raza, Zhaole Sun</i>	
Learning to Simulate Self-Driven Particles System with Coordinated Policy Optimization.....	10784
<i>Ka Ming Hui, Quanyi Li, Chunxiao Liu, Zhenghao Peng, Bolei Zhou</i>	
Gradient-Based Hyperparameter Optimization Over Long Horizons	10798
<i>Paul Micaelli, Amos Storkey</i>	
Stochastic Bias-Reduced Gradient Methods	10810
<i>Hilal Asi, Yair Carmon, Arun Jambulapati, Yujia Jin, Aaron Sidford</i>	
The Causal-Neural Connection: Expressiveness, Learnability, and Inference	10823
<i>Elias Bareinboim, Yoshua Bengio, Kai-Zhan Lee, Kevin Xia</i>	
Validation Free and Replication Robust Volume-Based Data Valuation	10837
<i>Chuan Sheng Foo, Bryan Kian Hsiang Low, Zhaoxuan Wu, Xinyi Xu</i>	
Implicit Finite-Horizon Approximation and Efficient Optimal Algorithms for Stochastic Shortest Path.....	10849
<i>Liyu Chen, Mehdi Jafarnia-Jahromi, Rahul Jain, Haipeng Luo</i>	
A Separation Result Between Data-Oblivious and Data-aware Poisoning Attacks.....	10862
<i>Samuel Deng, Sanjam Garg, Abhradeep Guha Thakurta, Somesh Jha, Saeed Mahloujifar, Mohammad Mahmoody</i>	
Deep Learning Through the Lens of Example Difficulty	10876
<i>Robert Baldock, Hartmut Maennel, Behnam Neyshabur</i>	
R-Drop: Regularized Dropout for Neural Networks	10890
<i>Wei Chen, Juntao Li, Tie-Yan Liu, Qi Meng, Tao Qin, Yue Wang, Lijun Wu, Min Zhang, Xiaobo Liang</i>	
Diversity Enhanced Active Learning with Strictly Proper Scoring Rules	10906
<i>Wray Buntine, Lan Du, Wei Tan</i>	

VOLUME 14

SSUL: Semantic Segmentation with Unknown Label for Exemplar-Based Class-Incremental Learning	10919
<i>Sungmin Cha, Taesup Moon, Youngjoon Yoo, Beomyoung Kim</i>	

Lower and Upper Bounds on the Pseudo-Dimension of Tensor Network Models.....	10931
<i>Behnoush Khavari, Guillaume Rabusseau</i>	
What Makes Multi-Modal Learning Better than Single (Provably)	10944
<i>Xuanyao Chen, Chenzhuang Du, Longbo Huang, Yu Huang, Zihui Xue, Hang Zhao</i>	
Quantifying and Improving Transferability in Domain Generalization	10957
<i>Pascal Poupart, Yaoliang Yu, Guojun Zhang, Han Zhao</i>	
Beyond Pinball Loss: Quantile Methods for Calibrated Uncertainty Quantification	10971
<i>Ian Char, Youngseog Chung, Willie Neiswanger, Jeff Schneider</i>	
Dynamic Inference with Neural Interpreters	10985
<i>Yoshua Bengio, Peter Gehler, Muhammad Waleed Gondal, Shruti Joshi, Francesco Locatello, Nasim Rahaman, Bernhard Scholkopf</i>	
Leveraging Recursive Gumbel-Max Trick for Approximate Inference in Combinatorial Spaces	10999
<i>Artyom Gadetsky, Danil Karpushkin, Denis Rakitin, Kirill Struminsky, Dmitry Vetrov</i>	
Hamiltonian Dynamics with Non-Newtonian Momentum for Rapid Sampling	11012
<i>Aram Galstyan, Greg Ver Steeg</i>	
Dynamic Normalization and Relay for Video Action Recognition.....	11026
<i>Dongqi Cai, Yurong Chen, Anbang Yao</i>	
Robust Visual Reasoning Via Language Guided Neural Module Networks.....	11041
<i>Arjun Akula, Soravit Changpinyo, Varun Jampani, Song-Chun Zhu</i>	
True Few-Shot Learning with Language Models	11054
<i>Kyunghyun Cho, Douwe Kiela, Ethan Perez</i>	
Selective Sampling for Online Best-Arm Identification.....	11071
<i>Romain Camilleri, Maryam Fazel, Lalit Jain, Kevin Jamieson, Zhihan Xiong</i>	
Multi-Task Learning of Order-Consistent Causal Graphs.....	11083
<i>Xinshi Chen, Caleb Ellington, Le Song, Haoran Sun, Eric Xing</i>	
Learning to Iteratively Solve Routing Problems with Dual-Aspect Collaborative Transformer.....	11096
<i>Zhiguang Cao, Zhenghua Chen, Jingwen Li, Yining Ma, Wen Song, Jing Tang, Le Zhang</i>	
Learning Interaction Rules from Multi-Animal Trajectories Via Augmented Behavioral Models	11108
<i>Keisuke Fujii, Emyo Fujioka, Mika Fukushima, Shizuko Hiryu, Kaoru Ide, Yoshinobu Kawahara, Hiroyoshi Kohno, Nozomi Nishiumi, Susumu Takahashi, Naoya Takeishi, Ryoya Tanaka, Kazushi Tsutsui, Ken Yoda</i>	
Differentiable Synthesis of Program Architectures	11123
<i>Guofeng Cui, He Zhu</i>	
Make Sure You're Unsure: a Framework for Verifying Probabilistic Specifications.....	11136
<i>Leonard Berrada, Rudy Bunel, Sumanth Dathathri, Krishnamurthy Dvijotham, Sven Gowal, M. Pawan Kumar, Robert Stanforth, Jonathan Uesato</i>	
Oracle-Efficient Regret Minimization in Factored MDPs with Unknown Structure.....	11148
<i>Yishay Mansour, Aviv Rosenberg</i>	
Linear-Time Probabilistic Solution of Boundary Value Problems	11160
<i>Philipp Hennig, Nicholas Kramer</i>	

Lifelong Domain Adaptation Via Consolidated Internal Distribution.....	11172
<i>Mohammad Rostami</i>	
Counterbalancing Learning and Strategic Incentives in Allocation Markets	11184
<i>Itai Ashlagi, Jamie Kang, Moran Koren, Faidra Monachou</i>	
Controlling Neural Networks with Rule Representations	11196
<i>Sercan Arik, Tomas Pfister, Sungyong Seo, Kihyuk Sohn, Jinsung Yoon, Xiang Zhang</i>	
Making the Most of Your Day: Online Learning for Optimal Allocation of Time	11208
<i>Etienne Boursier, Tristan Garrec, Vianney Perchet, Marco Scarsini</i>	
Federated Reconstruction: Partially Local Federated Learning.....	11220
<i>Zachary Garrett, Sushant Prakash, John Rush, Hakim Sidahmed, Karan Singhal, Shanshan Wu</i>	
Optimal Prediction of Markov Chains with and Without Spectral Gap	11233
<i>YanJun Han, Soham Jana, Yihong Wu</i>	
Subquadratic Overparameterization for Shallow Neural Networks	11247
<i>Volkan Cevher, Armin Eftekhari, Thomas Pethick, Ali Ramezani-Kebrya, Chaehwan Song</i>	
Continuous Doubly Constrained Batch Reinforcement Learning	11260
<i>Kavosh Asadi, Pratik Chaudhari, Rasool Fakoore, Jonas Mueller, Alexander Smola</i>	
Bridging Explicit and Implicit Deep Generative Models Via Neural Stein Estimators.....	11274
<i>Rui Gao, Qitian Wu, Hongyuan Zha</i>	
Score-Based Generative Modeling in Latent Space	11287
<i>Jan Kautz, Karsten Kreis, Arash Vahdat</i>	
Deep Conditional Gaussian Mixture Model for Constrained Clustering.....	11303
<i>Kieran Chin-Cheong, Laura Manduchi, Holger Michel, Julia Vogt, Sven Wellmann</i>	
Bootstrap Your Object Detector Via Mixed Training	11315
<i>Xiang Bai, Yue Cao, Han Hu, Stephen Lin, Yutong Lin, Fangyun Wei, Mengde Xu, Zheng Zhang</i>	
Tensor Decompositions of Higher-Order Correlations by Nonlinear Hebbian Plasticity	11326
<i>Michael Buice, Gabriel Ocker</i>	
Online Adaptation to Label Distribution Shift	11340
<i>Chuan Guo, Yi Su, Kilian Weinberger, Ruihan Wu</i>	
One Explanation is Not Enough: Structured Attention Graphs for Image Classification.....	11352
<i>Alan Fern, Minsuk Kahng, Fuxin Li, Vivswan Shitole, Prasad Tadepalli</i>	
Integrating Expert ODEs into Neural ODEs: Pharmacology and Disease Progression.....	11364
<i>Paul Elbers, Lucas Fleuren, Zhaozhi Qian, William Zame, Mihaela Van Der Schaar</i>	
Shifted Chunk Transformer for Spatio-Temporal Representational Learning.....	11384
<i>Ji Liu, Lv Xun, Sen Yang, Xuefan Zha, Wentao Zhu</i>	
Faster Proximal Algorithms for Matrix Optimization Using Jacobi-Based Eigenvalue Methods.....	11397
<i>Hamza Fawzi, Harry Goulbourne</i>	

Decrypting Cryptic Crosswords: Semantically Complex Wordplay Puzzles as a Target for NLP	11409
<i>Kyle Mahowald, Christopher Potts, Josh Rozner</i>	
An Improved Analysis of Gradient Tracking for Decentralized Machine Learning	11422
<i>Anastasiia Koloskova, Tao Lin, Sebastian Stich</i>	
Entropic Desired Dynamics for Intrinsic Control.....	11436
<i>Kate Baumli, Guillaume Desjardins, Steven Hansen, Nicolas Heess, Volodymyr Mnih, Simon Osindero, David Warde-Farley</i>	
Exploring Cross-Video and Cross-Modality Signals for Weakly-Supervised Audio-Visual Video Parsing.....	11449
<i>Hsin-Ying Lee, Yan-Bo Lin, Yen-Yu Lin, Hung-Yu Tseng, Ming-Hsuan Yang</i>	
Littlestone Classes Are Privately Online Learnable	11462
<i>Noah Golowich, Roi Livni</i>	
Dual Parameterization of Sparse Variational Gaussian Processes.....	11474
<i>Vincent Adam, Paul Chang, Mohammad Emtiyaz Khan, Arno Solin</i>	
Learning to Dehaze with Polarization	11487
<i>Yufei Han, Boxin Shi, Minggui Teng, Chao Xu, Chu Zhou</i>	
Conservative Data Sharing for Multi-Task Offline Reinforcement Learning	11501
<i>Yevgen Chebotar, Chelsea Finn, Karol Hausman, Aviral Kumar, Sergey Levine, Tianhe Yu</i>	
Universal Rate-Distortion-Perception Representations for Lossy Compression	11517
<i>Jun Chen, Ashish Khisti, Jingjing Qian, George Zhang</i>	
What’s a Good Imputation to Predict with Missing Values?	11530
<i>Julie Josse, Marine Le Morvan, Erwan Scornet, Gael Varoquaux</i>	
Replacing Rewards with Examples: Example-Based Policy Search Via Recursive Classification.....	11541
<i>Ben Eysenbach, Sergey Levine, Russ Salakhutdinov</i>	
Hierarchical Skills for Efficient Exploration	11553
<i>Jonas Gehring, Andreas Krause, Gabriel Synnaeve, Nicolas Usunier</i>	
Evidential Softmax for Sparse Multimodal Distributions in Deep Generative Models.....	11565
<i>Phil Chen, Mikhail Itkina, Mykel J Kochenderfer, Ransalu Senanayake</i>	
Submodular + Concave	11577
<i>Moran Feldman, Amin Karbasi, Siddharth Mitra</i>	
DeepGEM: Generalized Expectation-Maximization for Blind Inversion	11592
<i>Katherine Bouman, Jorge Castellanos, Angela Gao, Zachary Ross, Yisong Yue</i>	
Learning to Generate Visual Questions with Noisy Supervision	11604
<i>Zhuoye Ding, Shen Kai, Bo Long, Siliang Tang, Lingfei Wu, Yun Xiao, Yueting Zhuang, Zhen He</i>	
Pure Exploration in Kernel and Neural Bandits	11618
<i>Quanquan Gu, Ruoxi Jiang, Robert Nowak, Rebecca Willett, Dongruo Zhou, Yinglun Zhu</i>	
Numerical Composition of Differential Privacy.....	11631
<i>Sivakanth Gopi, Yin Tat Lee, Lukas Wutschitz</i>	

Coresets for Classification – Simplified and Strengthened	11643
<i>Tung Mai, Cameron Musco, Anup Rao</i>	
Sequential Algorithms for Testing Closeness of Distributions.....	11655
<i>Omar Fawzi, Nicolas Flammarion, Aurelien Garivier, Aadil Oufkir</i>	
Overlapping Spaces for Compact Graph Representations.....	11665
<i>Liudmila Prokhorenkova, Kirill Shevkunov</i>	
Hyperparameter Tuning is All You Need for LISTA.....	11678
<i>Xiaohan Chen, Jialin Liu, Zhangyang Wang, Wotao Yin</i>	
Foundations of Symbolic Languages for Model Interpretability.....	11690
<i>Marcelo Arenas, Pablo Barcelo, Daniel Baez, Jorge Perez, Bernardo Subercaseaux</i>	
Bridging Offline Reinforcement Learning and Imitation Learning: a Tale of Pessimism.....	11702
<i>Jiantao Jiao, Cong Ma, Paria Rashidinejad, Stuart Russell, Banghua Zhu</i>	
Impression Learning: Online Representation Learning with Synaptic Plasticity	11717
<i>Colin Bredenberg, Benjamin Lyo, Cristina Savin, Eero Simoncelli</i>	
How Well Do Feature Visualizations Support Causal Understanding of CNN Activations?.....	11730
<i>Matthias Bethge, Judy Borowski, Wieland Brendel, Robert Geirhos, Thomas Wallis, Roland S. Zimmermann</i>	
Fixes that Fail: Self-Defeating Improvements in Machine-Learning Systems	11745
<i>Chuan Guo, Awni Hannun, Ruihan Wu, Laurens Van Der Maaten</i>	
Coarse-To-fine Animal Pose and Shape Estimation.....	11757
<i>Gim Hee Lee, Chen Li</i>	
Meta-Learning Sparse Implicit Neural Representations.....	11769
<i>Jaeho Lee, Namhoon Lee, Jinwoo Shin, Jihoon Tack</i>	

VOLUME 15

Rethinking Space-Time Networks with Improved Memory Coverage for Efficient Video Object Segmentation.....	11781
<i>Ho Kei Cheng, Yu-Wing Tai, Chi-Keung Tang</i>	
Sparse Spiking Gradient Descent	11795
<i>Dan Goodman, Nicolas Perez-Nieves</i>	
Rethinking Calibration of Deep Neural Networks: Do Not Be Afraid of Overconfidence	11809
<i>Lei Feng, Deng-Bao Wang, Min-Ling Zhang</i>	
Towards Efficient and Effective Adversarial Training.....	11821
<i>Sravanti Addepalli, Arya Baburaj, Venkatesh Babu R, Gaurang Sriramanan</i>	
Intriguing Properties of Contrastive Losses	11834
<i>Ting Chen, Lala Li, Calvin Luo</i>	
Detecting Moments and Highlights in Videos Via Natural Language Queries	11846
<i>Mohit Bansal, Tamara L Berg, Jie Lei</i>	

Stochastic Optimization Under Time Drift: Iterate Averaging, Step-Decay Schedules, and High Probability Guarantees	11859
<i>Joshua Cutler, Dmitriy Drusvyatskiy, Zaid Harchaoui</i>	
Learning Stable Deep Dynamics Models for Partially Observed Or Delayed Dynamical Systems	11870
<i>Florian Dorfler, Andreas Krause, Andreas Schlaginhaufen, Philippe Wenk</i>	
An Uncertainty Principle is a Price of Privacy-Preserving Microdata	11883
<i>John Abowd, Robert Ashmead, Ryan Cumings-Menon, Simson Garfinkel, Daniel Kifer, Philip Leclerc, William Sexton, Ashley Simpson, Christine Task, Pavel Zhuravlev</i>	
Fairness in Ranking Under Uncertainty	11896
<i>Thorsten Joachims, David Kempe, Ashudeep Singh</i>	
Generalized Proximal Policy Optimization with Sample Reuse	11909
<i>Christos G Cassandras, Yannis Paschalidis, James Queeney</i>	
Mosaicking to Distill: Knowledge Distillation from Out-Of-Domain Data	11920
<i>Yifan Bao, Gongfan Fang, Chengchao Shen, Jie Song, Mingli Song, Xinchao Wang, Donglin Xie</i>	
Batch Active Learning at Scale	11933
<i>Gui Citovsky, Giulia Desalvo, Claudio Gentile, Lazaros Karydas, Sanjiv Kumar, Anand Rajagopalan, Afshin Rostamizadeh</i>	
Joint Semantic Mining for Weakly Supervised RGB-D Salient Object Detection	11945
<i>Qi Bi, Wei Ji, Jingjing Li, Huchuan Lu, Yongri Piao, Cheng Yan, Miao Zhang, Li Cheng</i>	
Not All Images Are Worth 16x16 Words: Dynamic Transformers for Efficient Image Recognition	11960
<i>Gao Huang, Rui Huang, Zeyi Huang, Shiji Song, Yulin Wang</i>	
Contrastive Learning for Neural Topic Model	11974
<i>Anh Tuan Luu, Thong Nguyen</i>	
Learning in Two-Player Zero-sum Partially Observable Markov Games with Perfect Recall	11987
<i>Tadashi Kozuno, Remi Munos, Pierre Menard, Michal Valko</i>	
A Geometric Structure of Acceleration and Its Role in Making Gradients Small Fast	11999
<i>Jongmin Lee, Chanwoo Park, Ernest Ryu</i>	
ATISS: Autoregressive Transformers for Indoor Scene Synthesis	12013
<i>Sanja Fidler, Andreas Geiger, Amlan Kar, Karsten Kreis, Despoina Paschalidou, Maria Shugrina</i>	
Generalized Depthwise-Separable Convolutions for Adversarially Robust and Efficient Neural Networks	12027
<i>Hassan Dbouk, Naresh Shanbhag</i>	
A Provably Efficient Model-Free Posterior Sampling Method for Episodic Reinforcement Learning	12040
<i>Christoph Dann, Mehryar Mohri, Tong Zhang, Julian Zimmert</i>	
Fast Federated Learning in the Presence of Arbitrary Device Unavailability	12052
<i>Xinran Gu, Kaixuan Huang, Longbo Huang, Jingzhao Zhang</i>	
On the Structure of Parametric Tournaments with Application to Ranking from Pairwise Comparisons	12065
<i>Arun Rajkumar, Vishnu Veerathu</i>	

SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers	12077
<i>Jose M. Alvarez, Anima Anandkumar, Ping Luo, Wenhai Wang, Enze Xie, Zhiding Yu</i>	
Fairness Via Representation Neutralization	12091
<i>Ahmed Awadallah, Mengnan Du, Xia Hu, Subhabrata Mukherjee, Ruixiang Tang, Guancho Wang</i>	
Residual Relaxation for Multi-View Representation Learning	12104
<i>Zhengyang Geng, Feng Jiang, Chuming Li, Zhouchen Lin, Yifei Wang, Yisen Wang, Jiansheng Yang</i>	
Do Vision Transformers See Like Convolutional Neural Networks?	12116
<i>Alexey Dosovitskiy, Simon Kornblith, Maithra Raghu, Thomas Unterthiner, Chiyuan Zhang</i>	
Optimization-Based Algebraic Multigrid Coarsening Using Reinforcement Learning.....	12129
<i>Scott Maclachlan, Luke Olson, Ali Taghibakhshi, Matthew West</i>	
Delayed Propagation Transformer: a Universal Computation Engine Towards Practical Control in Cyber-Physical Systems	12141
<i>Qiangqiang Guo, Peihao Wang, Zhangyang Wang, Hao Yang, Wenqing Zheng</i>	
Explaining Latent Representations with a Corpus of Examples.....	12154
<i>Jonathan Crabbe, Fergus Imrie, Zhaozhi Qian, Mihaela Van Der Schaar</i>	
Explaining Heterogeneity in Medial Entorhinal Cortex with Task-Driven Neural Networks.....	12167
<i>Alexander Attinger, Malcolm Campbell, Surya Ganguli, Lisa Giocomo, Kiah Hardcastle, Isabel Low, Caitlin S Mallory, Gabriel Mel, Aran Nayebi, Ben Sorscher, Alex H Williams, Dan Yamins</i>	
Beyond Smoothness: Incorporating Low-Rank Analysis into Nonparametric Density Estimation	12180
<i>Antoine Ledent, Robert Vandermeulen</i>	
Multi-View Representation Learning Via Total Correlation Objective	12194
<i>Seunghoon Hong, Hyeongjoo Hwang, Geon-Hyeong Kim, Kee-Eung Kim</i>	
FACMAC: Factored Multi-Agent Centralised Policy Gradients.....	12208
<i>Wendelin Boehmer, Pierre-Alexandre Kamienny, Bei Peng, Tabish Rashid, Christian Schroeder De Witt, Philip Torr, Shimon Whiteson</i>	
EDGE: Explaining Deep Reinforcement Learning Policies	12222
<i>Wenbo Guo, Usman Khan, Xian Wu, Xinyu Xing</i>	
Learning to Assimilate in Chaotic Dynamical Systems	12237
<i>Jed Brown, John McCabe</i>	
Object-Aware Contrastive Learning for Debiased Scene Representation	12251
<i>Hyunwoo Kang, Chun-Liang Li, Sangwoo Mo, Jinwoo Shin, Kihyuk Sohn</i>	
Evaluating Efficient Performance Estimators of Neural Architectures	12265
<i>Wenshuo Li, Shuang Liang, Xuefei Ning, Changcheng Tang, Yu Wang, Huazhong Yang, Zixuan Zhou</i>	
A-NeRF: Articulated Neural Radiance Fields for Learning Human Shape, Appearance, and Pose.....	12278
<i>Helge Rhodin, Shih-Yang Su, Frank Yu, Michael Zollhoefer</i>	
Differential Privacy Over Riemannian Manifolds.....	12292
<i>Karthik Bharath, Matthew Reimherr, Carlos Soto</i>	

How Can Classical Multidimensional Scaling Go Wrong?.....	12304
<i>Anna Gilbert, Benjamin Raichel, Rishi Sonthalia, Greg Van Buskirk</i>	
Modeling Heterogeneous Hierarchies with Relation-Specific Hyperbolic Cones.....	12316
<i>Yushi Bai, Jure Leskovec, Hongyu Ren, Zhitao Ying</i>	
Non-Asymptotic Error Bounds for Bidirectional GANs	12328
<i>Jian Huang, Yuling Jiao, Shiao Liu, Yang Wang, Yunfei Yang</i>	
Confidence-Aware Imitation Learning from Demonstrations with Varying Optimality.....	12340
<i>Zhangjie Cao, Dorsa Sadigh, Yanan Sui, Songyuan Zhang</i>	
Answering Complex Causal Queries with the Maximum Causal Set Effect.....	12351
<i>Zachary Markovich</i>	
Identifiability in Inverse Reinforcement Learning	12362
<i>Haoyang Cao, Samuel Cohen, Lukasz Szpruch</i>	
A Probabilistic State Space Model for Joint Inference from Differential Equations and Data.....	12374
<i>Philipp Hennig, Nicholas Kramer, Jonathan Schmidt</i>	
On Plasticity, Invariance, and Mutually Frozen Weights in Sequential Task Learning	12386
<i>Alessandro Achille, Andrea Censi, Emilio Frazzoli, Julian Zilly</i>	
Provably Efficient Black-Box Action Poisoning Attacks Against Reinforcement Learning	12400
<i>Lifeng Lai, Guanlin Liu</i>	
Fast Approximation of the Sliced-Wasserstein Distance Using Concentration of Random Projections.....	12411
<i>Roland Badeau, Alain Durmus, Pierre E Jacob, Kimia Nadjahi, Umut Simsekli</i>	
Causal Navigation by Continuous-Time Neural Networks	12425
<i>Alexander Amini, Ramin Hasani, Mathias Lechner, Daniela Rus, Charles Vorbach</i>	
Global Convergence of Online Optimization for Nonlinear Model Predictive Control	12441
<i>Sen Na</i>	
Argmax Flows and Multinomial Diffusion: Learning Categorical Distributions	12454
<i>Patrick Forre, Emiel Hoogeboom, Priyank Jaini, Didrik Nielsen, Max Welling</i>	
Learning with User-Level Privacy	12466
<i>Kareem Amin, Satyen Kale, Alex Kulesza, Daniel Levy, Mehryar Mohri, Ziteng Sun, Ananda Theertha Suresh</i>	
Don't Generate Me: Training Differentially Private Generative Models with Sinkhorn Divergence	12480
<i>Alex Bie, Tianshi Cao, Sanja Fidler, Karsten Kreis, Arash Vahdat</i>	
Keeping Your Eye on the Ball: Trajectory Attention in Video Transformers	12493
<i>Yuki Asano, Dylan Campbell, Christoph Feichtenhofer, João Henriques, Florian Metze, Ishan Misra, Mandela Patrick, Andrea Vedaldi</i>	
Variational Bayesian Optimistic Sampling	12507
<i>Tor Lattimore, Brendan O'Donoghue</i>	
Cross-Modal Domain Adaptation for Cost-Efficient Visual Reinforcement Learning	12520
<i>Xiong-Hui Chen, Shengyi Jiang, Feng Xu, Yang Yu, Zongzhang Zhang</i>	

D2C: Diffusion-Decoding Models for Few-Shot Conditional Generation	12533
<i>Stefano Ermon, Chenlin Meng, Abhishek Sinha, Jiaming Song</i>	
Continual Auxiliary Task Learning	12549
<i>Andrew Jacobsen, Raksha Kumaraswamy, Chunlok Lo, Matthew McLeod, Matthew Schlegel, Adam White, Martha White</i>	
Constrained Two-Step Look-Ahead Bayesian Optimization	12563
<i>Peter Frazier, Xiangyu Zhang, Yunxiang Zhang</i>	
Learning with Labeling Induced Abstentions.....	12576
<i>Kareem Amin, Giulia Desalvo, Afshin Rostamizadeh</i>	
SQUALER: Scaling Question Answering by Decoupling Multi-Hop and Logical Reasoning.....	12587
<i>Mattia Atzeni, Jasmina Bogojeska, Andreas Loukas</i>	
Out-Of-Distribution Generalization in Kernel Regression	12600
<i>Blake Bordelon, Abdulkadir Canatar, Cengiz Pehlevan</i>	
FL-WBC: Enhancing Robustness Against Model Poisoning Attacks in Federated Learning from a Client Perspective	12613
<i>Yiran Chen, Louis Divalentin, Amin Hassanzadeh, Ang Li, Hai Li, Jingwei Sun</i>	
Chebyshev-Cantelli PAC-Bayes-Bennett Inequality for the Weighted Majority Vote	12625
<i>Christian Igel, Stephan Lorenzen, Andres Masegosa, Yevgeny Seldin, Yi-Shan Wu</i>	

VOLUME 16

A Multi-Implicit Neural Representation for Fonts	12637
<i>Matthew Fisher, Hailin Jin, Niloy Mitra, Pradyumna Reddy, Zhaowen Wang, Zhifei Zhang</i>	
OctField: Hierarchical Implicit Functions for 3D Modeling	12648
<i>Weikai Chen, Lin Gao, Songrun Liu, Jia-Heng Tang, Bo Wang, Bo Yang, Jie Yang</i>	
The Inductive Bias of Quantum Kernels	12661
<i>Simon Buchholz, Jonas Kubler, Bernhard Scholkopf</i>	
An Exponential Improvement on the Memorization Capacity of Deep Threshold Networks	12674
<i>Amin Karbasi, Dimitris Papailiopoulos, Shashank Rajput, Kartik Sreenivasan</i>	
Pretraining Representations for Data-Efficient Reinforcement Learning.....	12686
<i>Ankesh Anand, Philip Bachman, Laurent Charlin, Aaron Courville, R Devon Hjelm, Michael Noukhovitch, Nitarshan Rajkumar, Max Schwarzer</i>	
Universal Approximation Using Well-Conditioned Normalizing Flows	12700
<i>Holden Lee, Chirag Pabbaraju, Andrej Risteski, Anish Prasad Sevekari</i>	
On the Validity of Modeling SGD with Stochastic Differential Equations (SDEs)	12712
<i>Sanjeev Arora, Zhiyuan Li, Sadhika Malladi</i>	
Proportional Participatory Budgeting with Additive Utilities	12726
<i>Dominik Peters, Grzegorz Pierczynski, Piotr Skowron</i>	
Disentangling the Roles of Curation, Data-Augmentation and the Prior in the Cold Posterior Effect.....	12738
<i>Gregor Bachmann, Thomas Hofmann, Lorenzo Noci, Sebastian Nowozin, Kevin Roth</i>	

Sanity Checks for Lottery Tickets: Does Your Winning Ticket Really Win the Jackpot?	12749
<i>Tianlong Chen, Xiaohan Chen, Xuxi Chen, Ning Liu, Sijia Liu, Xiaolong Ma, Minghai Qin, Xuan Shen, Yanzhi Wang, Zhangyang Wang, Geng Yuan</i>	
Collaborative Causal Discovery with Atomic Interventions	12761
<i>Raghavendra Addanki, Shiva Kasiviswanathan</i>	
Towards Optimally Abstaining from Prediction with OOD Test Examples	12774
<i>Adam Kalai, Varun Kanade</i>	
TokenLearner: Adaptive Space-Time Tokenization for Videos.....	12786
<i>Anelia Angelova, Anurag Arnab, Mostafa Dehghani, Aj Piergiovanni, Michael Ryoo</i>	
Learning in Multi-Stage Decentralized Matching Markets	12798
<i>Xiaowu Dai, Michael Jordan</i>	
Non-Asymptotic Convergence Bounds for Wasserstein Approximation Using Point Clouds	12810
<i>Quentin Merigot, Clement Sarrazin, Filippo Santambrogio</i>	
Understanding Interlocking Dynamics of Cooperative Rationalization	12822
<i>Shiyu Chang, Tommi Jaakkola, Mo Yu, Yang Zhang</i>	
Adversarial Robustness Without Adversarial Training: a Teacher-Guided Curriculum Learning Approach	12836
<i>Sowrya Gali, Vineeth N Balasubramanian, Anindya Sarkar, Anirban Sarkar</i>	
Tactical Optimism and Pessimism for Deep Reinforcement Learning.....	12849
<i>Michael Arbel, Michael Jordan, Ted Moskowitz, Aldo Pacchiano, Jack Parker-Holder</i>	
Towards Hyperparameter-Free Policy Selection for Offline Reinforcement Learning	12864
<i>Nan Jiang, Siyuan Zhang</i>	
FjORD: Fair and Accurate Federated Learning Under Heterogeneous Targets with Ordered Dropout.....	12876
<i>Mario Almeida, Samuel Horvath, Nicholas Lane, Stefanos Laskaridis, Ilias Leontiadis, Stylianos Venieris</i>	
Optimal Uniform OPE and Model-Based Offline Reinforcement Learning in Time-Homogeneous, Reward-Free and Task-Agnostic Settings	12890
<i>Yu-Xiang Wang, Ming Yin</i>	
MixSeq: Connecting Macroscopic Time Series Forecasting with Microscopic Time Series Data	12904
<i>Lei Chen, Ge Jin, Ziqi Liu, Zhiqiang Zhang, Jianyong Zhou, Jun Zhou, Zhibo Zhu</i>	
Pareto Domain Adaptation	12917
<i>Kaixiong Gong, Han Li, Shuang Li, Jian Liang, Chi Harold Liu, Di Liu, Guoren Wang, Fangrui Lv</i>	
Divergence Frontiers for Generative Models: Sample Complexity, Quantization Effects, and Frontier Integrals	12930
<i>Yejin Choi, Zaid Harchaoui, Lang Liu, Sewoong Oh, Krishna Pillutla, Sean Welleck</i>	
Consistency Regularization for Variational Auto-Encoders.....	12943
<i>Adji Bousso Dieng, Samarth Sinha</i>	
Score-Based Generative Neural Networks for Large-Scale Optimal Transport.....	12955
<i>Grady Daniels, Paul Hand, Tyler Maunu</i>	

Interactive Label Cleaning with Example-Based Explanations.....	12966
<i>Andrea Bontempelli, Fausto Giunchiglia, Andrea Passerini, Stefano Teso</i>	
Gradient Descent on Two-Layer Nets: Margin Maximization and Simplicity Bias.....	12978
<i>Sanjeev Arora, Zhiyuan Li, Kaifeng Lyu, Runzhe Wang</i>	
Glance-And-Gaze Vision Transformer.....	12992
<i>Yutong Bai, Yongyi Lu, Wei Shen, Yingda Xia, Qihang Yu, Alan Yuille</i>	
Stochastic L ^q -Convex Function Minimization.....	13004
<i>Javad Lavaei, Haixiang Zhang, Zeyu Zheng</i>	
Self-Supervised GANs with Label Augmentation.....	13019
<i>Qi Cao, Xueqi Cheng, Liang Hou, Huawei Shen</i>	
Shape as Points: a Differentiable Poisson Solver	13032
<i>Andreas Geiger, Chiyu Jiang, Yiyi Liao, Michael Niemeyer, Songyou Peng, Marc Pollefeys</i>	
Outcome-Driven Reinforcement Learning Via Variational Inference	13045
<i>Yarin Gal, Sergey Levine, Rowan McAllister, Vitchyr Pong, Tim G. J. Rudner</i>	
Drawing Robust Scratch Tickets: Subnetworks with Inborn Robustness Are Found Within Randomly Initialized Networks.....	13059
<i>David Cox, Yonggan Fu, Yingyan Lin, Xu Ouyang, Shang Wu, Qixuan Yu, Yang Zhang</i>	
Rectifying the Shortcut Learning of Background for Few-Shot Learning	13073
<i>Xu Luo, Qi Tian, Longhui Wei, Liangjian Wen, Lingxi Xie, Zenglin Xu, Jinrong Yang</i>	
SEAL: Self-Supervised Embodied Active Learning Using Exploration and 3D Consistency	13086
<i>Devendra Singh Chaplot, Murtaza Dalal, Saurabh Gupta, Jitendra Malik, Russ Salakhutdinov</i>	
Sifting Through the Noise: Universal First-Order Methods for Stochastic Variational Inequalities.....	13099
<i>Kimon Antonakopoulos, Volkan Cevher, Ali Kavis, Panayotis Mertikopoulos, Thomas Pethick</i>	
Accommodating Picky Customers: Regret Bound and Exploration Complexity for Multi-Objective Reinforcement Learning.....	13112
<i>Vladimir Braverman, Jingfeng Wu, Lin Yang</i>	
Exact Privacy Guarantees for Markov Chain Implementations of the Exponential Mechanism with Artificial Atoms.....	13125
<i>Matthew Reimherr, Jeremy Seeman, Aleksandra Slavkovic</i>	
The Emergence of Objectness: Learning Zero-Shot Segmentation from Videos.....	13137
<i>Stephen Lin, Runtao Liu, Zhirong Wu, Stella Yu</i>	
Direct Multi-View Multi-person 3D Pose Estimation.....	13153
<i>Yujun Cai, Jiashi Feng, Shuicheng Yan, Jianfeng Zhang, Tao Wang</i>	
MST: Masked Self-Supervised Transformer for Visual Representation	13165
<i>Zhiyang Chen, Rui Deng, Wei Li, Zhaowen Li, Ming Tang, Jinqiao Wang, Liwei Wu, Fan Yang, Chaoyang Zhao, Rui Zhao, Yousong Zhu</i>	
Exploiting Opponents Under Utility Constraints in Sequential Games.....	13177
<i>Martino Bernasconi-De-Luca, Federico Cacciamani, Simone Fioravanti, Nicola Gatti, Alberto Marchesi, Francesco Trovò</i>	

A Compositional Atlas of Tractable Circuit Operations for Probabilistic Inference	13189
<i>Yoojung Choi, Anji Liu, Stefano Teso, Guy Van Den Broeck, Antonio Vergari</i>	
Demystifying and Generalizing BinaryConnect.....	13202
<i>Tim Dockhorn, Vahid Partovi Nia, Eyyub Sari, Yaoliang Yu, Mahdi Zolnouri</i>	
CARMS: Categorical-Antithetic-REINFORCE Multi-Sample Gradient Estimator.....	13217
<i>Alek Dimitriev, Mingyuan Zhou</i>	
Learning to Learn Dense Gaussian Processes for Few-Shot Learning	13230
<i>Zichen Miao, Qiang Qiu, Ze Wang, Xiantong Zhen</i>	
Stochastic Solutions for Linear Inverse Problems Using the Prior Implicit in a Denoiser	13242
<i>Zahra Kadkhodaie, Eero Simoncelli</i>	
Towards Stable and Robust AdderNets.....	13255
<i>Xinghao Chen, Minjing Dong, Yunhe Wang, Chang Xu</i>	
Representing Long-Range Context for Graph Neural Networks with Global Attention.....	13266
<i>Joseph Gonzalez, Paras Jain, Azalia Mirhoseini, Ion Stoica, Matthew Wright, Zhanghao Wu</i>	
Beyond Bandit Feedback in Online Multiclass Classification	13280
<i>Nicolò Cesa-Bianchi, Federico Fusco, Dirk Van Der Hoeven</i>	
Learning Student-Friendly Teacher Networks for Knowledge Distillation.....	13292
<i>Moon-Hyun Cha, Bohyung Han, Daesin Kim, Dae Young Park, Changwook Jeong</i>	
Implicit Transformer Network for Screen Content Image Continuous Super-Resolution.....	13304
<i>Kun Li, Sheng Shen, Jingyu Yang, Huanjing Yue</i>	
Channel Permutations for N:M Sparsity.....	13316
<i>Jeff Pool, Chong Yu</i>	
Curriculum Learning for Vision-And-Language Navigation	13328
<i>Jianqing Fan, Jiajie Peng, Jiwen Zhang, Zhongyu Wei</i>	
Better Algorithms for Individually Fair k-Clustering.....	13340
<i>Deeparnab Chakrabarty, Maryam Negahbani</i>	
Video Instance Segmentation Using Inter-Frame Communication Transformers.....	13352
<i>Miran Heo, Sukjun Hwang, Seon Joo Kim, Seoung Wug Oh</i>	
Progressive Coordinate Transforms for Monocular 3D Object Detection.....	13364
<i>Tong He, Mu Li, Li Wang, Xiangyang Xue, Li Zhang, Zhi Zhang, Yi Zhu</i>	
Structured Reordering for Modeling Latent Alignments in Sequence Transduction.....	13378
<i>Mirella Lapata, Ivan Titov, Bailin Wang</i>	
A Universal Probabilistic Spike Count Model Reveals Ongoing Modulation of Neural Variability	13392
<i>Mate Lengyel, David Liu</i>	
Bellman Eluder Dimension: New Rich Classes of RL Problems, and Sample-Efficient Algorithms	13406
<i>Chi Jin, Qinghua Liu, Sobhan Miryoosefi</i>	
Detecting Anomalous Event Sequences with Temporal Point Processes.....	13419
<i>Jan Gasthaus, Stephan Gunnemann, Tim Januschowski, Oleksandr Shchur, Ali Caner Turkmen</i>	

HNPE: Leveraging Global Parameters for Neural Posterior Estimation	13432
<i>Alexandre Gramfort, Gilles Louppe, Thomas Moreau, Pedro Rodrigues</i>	
Alignment Attention by Matching Key and Query Distributions	13444
<i>Xinjie Fan, Korawat Tanwisuth, Shujian Zhang, Huangjie Zheng, Mingyuan Zhou</i>	
Settling the Variance of Multi-Agent Policy Gradients	13458
<i>Jakub Grudzien Kuba, Linghui Meng, David Mguni, Jun Wang, Muning Wen, Yaodong Yang, Haifeng Zhang, Shangding Gu</i>	
For High-Dimensional Hierarchical Models, Consider Exchangeability of Effects Across Covariates Instead of Across Datasets	13471
<i>Tamara Broderick, Hilary Finucane, Brian Trippe</i>	
Efficient Algorithms for Learning Depth-2 Neural Networks with General ReLU Activations	13485
<i>Pranjal Awasthi, Alex Tang, Aravindan Vijayaraghavan</i>	
VOLUME 17	
Controllable and Compositional Generation with Latent-Space Energy-Based Models	13497
<i>Anima Anandkumar, Weili Nie, Arash Vahdat</i>	
Reverse-Complement Equivariant Networks for DNA Sequences	13511
<i>Vincent Mallet, Jean-Philippe Vert</i>	
Provably Efficient Reinforcement Learning with Linear Function Approximation Under Adaptivity Constraints	13524
<i>Quanquan Gu, Tianhao Wang, Dongruo Zhou</i>	
Nonsmooth Implicit Differentiation for Machine-Learning and Optimization	13537
<i>Jerome Bolte, Tam Le, Edouard Pauwels, Tony Silveti-Falls</i>	
Heuristic-Guided Reinforcement Learning	13550
<i>Ching-An Cheng, Andrey Kolobov, Adith Swaminathan</i>	
Statistical Undecidability in Linear, Non-Gaussian Causal Models in the Presence of Latent Confounders	13564
<i>Konstantin Genin</i>	
A Novel Notion of Barycenter for Probability Distributions Based on Optimal Weak Mass Transport	13575
<i>Elsa Cazelles, Joaquin Fontbona, Felipe Tobar</i>	
Temporal-Attentive Covariance Pooling Networks for Video Recognition	13587
<i>Zilin Gao, Qinghua Hu, Peihua Li, Qilong Wang, Bingbing Zhang</i>	
Revisiting Smoothed Online Learning	13599
<i>Wei Jiang, Shiyin Lu, Tianbao Yang, Lijun Zhang</i>	
Marginalised Gaussian Processes with Nested Sampling	13613
<i>Vidhi Lalchand, Carl Edward Rasmussen, Fergus Simpson</i>	
Provable Benefits of Actor-Critic Methods for Offline Reinforcement Learning	13626
<i>Emma Brunskill, Martin J Wainwright, Andrea Zanette</i>	

Bayesian Bellman Operators	13641
<i>Mattie Fellows, Kristian Hartikainen, Shimon Whiteson</i>	
Uncertainty Calibration for Ensemble-Based Debiasing Methods	13657
<i>Yimeng Chen, Xueqi Cheng, Yanyan Lan, Zhi-Ming Ma, Liang Pang, Ruibin Xiong</i>	
Provably Faster Algorithms for Bilevel Optimization.....	13670
<i>Kaiyi Ji, Yingbin Liang, Junjie Yang</i>	
Neo-GNNs: Neighborhood Overlap-aware Graph Neural Networks for Link Prediction.....	13683
<i>Jaewoo Kang, Hyunwoo Kim, Seoyoon Kim, Junhyun Lee, Seongjun Yun</i>	
Self-Supervised Multi-Object Tracking with Cross-input Consistency	13695
<i>Favyen Bastani, Songtao He, Samuel Madden</i>	
Tree in Tree: from Decision Trees to Decision Graphs.....	13707
<i>Mahsa Shoaran, Bingzhao Zhu</i>	
Test-Time Collective Prediction.....	13719
<i>Stephen Bates, Wenshuo Guo, Michael Jordan, Celestine Mender-Dunner</i>	
A Continuous Mapping for Augmentation Design.....	13732
<i>Puneet Dokania, Ser Nam Lim, Chen Lin, Wanli Ouyang, Keyu Tian, Philip Torr</i>	
Neural Routing by Memory.....	13744
<i>Zhenqiang Li, Zhifeng Li, Wei Liu, Yoichi Sato, Kaipeng Zhang</i>	
GeoMol: Torsional Geometric Generation of Molecular 3D Conformer Ensembles.....	13757
<i>Regina Barzilay, Connor Coley, Octavian Ganea, William Green, Tommi Jaakkola, Klavs Jensen, Lagnajit Pattanaik</i>	
CANITA: Faster Rates for Distributed Convex Optimization with Communication Compression.....	13770
<i>Zhize Li, Peter Richtarik</i>	
Drop-DTW: Aligning Common Signal Between Sequences While Dropping Outliers.....	13782
<i>Konstantinos Derpanis, Mikita Dvornik, Animesh Garg, Isma Hadji, Allan Jepson</i>	
Safe Reinforcement Learning with Natural Language Constraints	13794
<i>Yinlam Chow, Michael Y Hu, Karthik Narasimhan, Peter J Ramadge, Tsung-Yen Yang</i>	
Compositional Modeling of Nonlinear Dynamical Systems with ODE-Based Random Features	13809
<i>Thomas McDonald, Mauricio Alvarez</i>	
Implicit Semantic Response Alignment for Partial Domain Adaptation.....	13820
<i>Zhengming Ding, Hongfu Liu, Wenxiao Xiao</i>	
ToAlign: Task-Oriented Alignment for Unsupervised Domain Adaptation.....	13834
<i>Zhibo Chen, Cuiling Lan, Guoqiang Wei, Wenjun Zeng, Zhizheng Zhang</i>	
Prior-Independent Dynamic Auctions for a Value-maximizing Buyer	13847
<i>Yuan Deng, Hanrui Zhang</i>	
Safe Reinforcement Learning by Imagining the Near Future.....	13859
<i>Yuping Luo, Tengyu Ma, Garrett Thomas</i>	
Contrastive Active Inference	13870
<i>Bart Dhoedt, Pietro Mazzaglia, Tim Verbelen</i>	

Overparameterization Improves Robustness to Covariate Shift in High Dimensions	13883
<i>Ben Adlam, Jeffrey Pennington, Nilesh Tripuraneni</i>	
Logarithmic Regret in Feature-Based Dynamic Pricing.....	13898
<i>Yu-Xiang Wang, Jianyu Xu</i>	
Dimension-Free Empirical Entropy Estimation	13911
<i>Doron Cohen, Aryeh Kontorovich, Aaron Koolyk, Geoffrey Wolfer</i>	
Towards Biologically Plausible Convolutional Networks.....	13924
<i>Peter E Latham, Timothy Lillicrap, Yash Mehta, Roman Pogodin</i>	
DynamicViT: Efficient Vision Transformers with Dynamic Token Sparsification	13937
<i>Cho-Jui Hsieh, Benlin Liu, Jiwen Lu, Yongming Rao, Wenliang Zhao, Jie Zhou</i>	
Learning Transferable Adversarial Perturbations.....	13950
<i>Krishna Kanth Nakka, Mathieu Salzmann</i>	
PortaSpeech: Portable and High-Quality Generative Text-to-Speech.....	13963
<i>Jinglin Liu, Yi Ren, Zhou Zhao</i>	
Exponential Graph is Provably Efficient for Decentralized Deep Training	13975
<i>Yiming Chen, Hanbin Hu, Pan Pan, Wotao Yin, Bicheng Ying, Kun Yuan</i>	
CLIP-It! Language-Guided Video Summarization.....	13988
<i>Trevor Darrell, Medhini Narasimhan, Anna Rohrbach</i>	
Learning Treatment Effects in Panels with General Intervention Patterns.....	14001
<i>Vivek Farias, Andrew Li, Tianyi Peng</i>	
Lossy Compression for Lossless Prediction	14014
<i>Benjamin Bloem-Reddy, Yann Dubois, Chris Maddison, Karen Ullrich</i>	
From Optimality to Robustness: Adaptive Re-Sampling Strategies in Stochastic Bandits	14029
<i>Dorian Baudry, Odalric-Ambrym Maillard, Patrick Saux</i>	
CCVS: Context-Aware Controllable Video Synthesis	14042
<i>Guillaume Le Moing, Jean Ponce, Cordelia Schmid</i>	
An Online Riemannian PCA for Stochastic Canonical Correlation Analysis	14056
<i>Rudrasis Chakraborty, Zihang Meng, Vikas Singh</i>	
Predify: Augmenting Deep Neural Networks with Brain-Inspired Predictive Coding Dynamics.....	14069
<i>B. Ador, Andrea Alamia, Callum Biggs O'May, Bhavin Choksi, Milad Mozafari, Rufin Vanrullen</i>	
Deep Extrapolation for Attribute-Enhanced Generation	14084
<i>Alvin Chan, Ben Krause, Ali Madani, Nikhil Naik</i>	
Generalized DataWeighting Via Class-Level Gradient Manipulation	14097
<i>Xi Chen, Can Chen, Erqun Dong, Dejing Dou, Hao Liu, Xue (Steve) Liu, Shuhao Zheng</i>	
Slow Learning and Fast Inference: Efficient Graph Similarity Computation Via Knowledge Distillation	14110
<i>Yun Fu, Can Qin, Huan Wang, Lichen Wang, Yulun Zhang, Handong Zhao</i>	

Meta Learning Backpropagation and Improving it.....	14122
<i>Louis Kirsch, Jurgen Schmidhuber</i>	
Posterior Meta-Replay for Continual Learning	14135
<i>Maria Cervera, Francesco D'Angelo, Benjamin Ehret, Benjamin F. Grewe, Christian Henning, Seijin Kobayashi, João Sacramento, Regina Traber, Johannes Von Oswald</i>	
Optimizing Reusable Knowledge for Continual Learning Via Metalearning	14150
<i>Julio Hurtado, Alain Raymond, Alvaro Soto</i>	
A Sampling-Based Circuit for Optimal Decision Making.....	14163
<i>Camille Rullan Buxo, Cristina Savin</i>	
Compressed Video Contrastive Learning	14176
<i>Mingyu Ding, Nanyi Fei, Yuqi Huo, Haoyu Lu, Zhiwu Lu, Ping Luo, Ji-Rong Wen</i>	
Uniform-PAC Bounds for Reinforcement Learning with Linear Function Approximation.....	14188
<i>Quanquan Gu, Jiafan He, Dongruo Zhou</i>	
Attention Bottlenecks for Multimodal Fusion	14200
<i>Anurag Arnab, Aren Jansen, Arsha Nagrani, Cordelia Schmid, Chen Sun, Shan Yang</i>	
Convergence of Adaptive Algorithms for Constrained Weakly Convex Optimization.....	14214
<i>Ahmet Alacaoglu, Volkan Cevher, Yura Malitsky</i>	
On the Convergence of Step Decay Step-Size for Stochastic Optimization.....	14226
<i>Mikael Johansson, Sindri Magnússon, Xiaoyu Wang</i>	
BernNet: Learning Arbitrary Graph Spectral Filters Via Bernstein Approximation	14239
<i>Mingguo He, Zengfeng Huang, Zhewei Wei, Hongteng Xu</i>	
Co-Evolution Transformer for Protein Contact Prediction.....	14252
<i>Liang He, Fusong Ju, Tie-Yan Liu, Bin Shao, He Zhang, Nanning Zheng, Jianwei Zhu</i>	
Unsupervised Foreground Extraction Via Deep Region Competition.....	14264
<i>Xiaojian Ma, Ying Nian Wu, Sirui Xie, Peiyu Yu, Song-Chun Zhu, Yixin Zhu</i>	
Leveraging Spatial and Temporal Correlations in Sparsified Mean Estimation.....	14280
<i>Advait Gadhikar, Divyansh Jhunjunwala, Gauri Joshi, Swanand Kadhe, Ankur Mallick</i>	
Last-Iterate Convergence in Extensive-Form Games	14293
<i>Christian Kroer, Chung-Wei Lee, Haipeng Luo</i>	
Class-Incremental Learning Via Dual Augmentation	14306
<i>Zhen Cheng, Cheng-Lin Liu, Xu-Yao Zhang, Fei Zhu</i>	
Robust and Fully-Dynamic Coreset for Continuous-and-Bounded Learning (With Outliers) Problems	14319
<i>Hu Ding, Yiwen Guo, Zixiu Wang</i>	
Rethinking and Reweighting the Univariate Losses for Multi-Label Ranking: Consistency and Generalization	14332
<i>Chongxuan Li, Guoqiang Wu, Kun Xu, Jun Zhu</i>	
Fair Clustering Under a Bounded Cost.....	14345
<i>Brian Brubach, John Dickerson, Seyyed Esmaili, Aravind Srinivasan</i>	

VOLUME 18

Improving Calibration Through the Relationship with Adversarial Robustness	14358
<i>Alex Beutel, Ed Chi, Yao Qin, Xuezhi Wang</i>	
Credal Self-Supervised Learning.....	14370
<i>Eyke Hullermeier, Julian Li</i>	
Spot the Difference: Detection of Topological Changes Via Geometric Alignment.....	14383
<i>Per Steffen Czolbe, Aasa Feragen, Oswin Krause</i>	
Rethinking the Variational Interpretation of Accelerated Optimization Methods.....	14396
<i>Hadi Daneshmand, Antonio Orvieto, Peiyuan Zhang</i>	
Linear and Kernel Classification in the Streaming Model: Improved Bounds for Heavy Hitters	14407
<i>Arvind Mahankali, David Woodruff</i>	
A PAC-Bayes Analysis of Adversarial Robustness	14421
<i>Amaury Habrard, Emilie Morvant, Eric Guillaume Vidot, Paul Viallard</i>	
SE(3)-Equivariant Prediction of Molecular Wavefunctions and Electronic Densities	14434
<i>Mihail Bogojeski, Michael Gastegger, Mario Geiger, Klaus-Robert Muller, Tess Smidt, Oliver Unke</i>	
Modified Frank Wolfe in Probability Space.....	14448
<i>Jose Blanchet, Peter W Glynn, Carson Kent, Jiajin Li</i>	
Bayesian Optimization of Function Networks.....	14463
<i>Raul Astudillo, Peter Frazier</i>	
Look at What I'm Doing: Self-Supervised Spatial Grounding of Narrations in Instructional Videos	14476
<i>Hailin Jin, Bryan Plummer, Bryan Russell, Kate Saenko, Reuben Tan</i>	
RETRIEVE: Coreset Selection for Efficient and Robust Semi-Supervised Learning.....	14488
<i>Feng Chen, Rishabh Iyer, Krishnateja Killamsetty, Xujiang Zhao</i>	
Collaborating with Humans Without Human Data.....	14502
<i>Matt Botvinick, Richard Everett, Edward Hughes, Kevin McKee, Dj Strouse</i>	
Training Feedback Spiking Neural Networks by Implicit Differentiation on the Equilibrium State	14516
<i>Zhouchen Lin, Qingyan Meng, Yisen Wang, Mingqing Xiao, Zongpeng Zhang</i>	
Online Selective Classification with Limited Feedback.....	14529
<i>Ashok Cutkosky, Aditya Gangrade, Anil Kag, Venkatesh Saligrama</i>	
Controlled Text Generation as Continuous Optimization with Multiple Constraints.....	14542
<i>Sachin Kumar, Eric Malmi, Aliaksei Severyn, Yulia Tsvetkov</i>	
S ³ : Sign-Sparse-Shift Reparametrization for Effective Training of Low-bit Shift Networks.....	14555
<i>Xinlin Li, Bang Liu, Wulong Liu, Vahid Partovi Nia, Chunjing Xu, Yaoliang Yu</i>	
Implicit MLE: Backpropagating Through Discrete Exponential Family Distributions.....	14567
<i>Luca Franceschi, Pasquale Minervini, Mathias Niepert</i>	
Scaling Up Continuous-Time Markov Chains Helps Resolve Underspecification	14580
<i>Rebekka Burkholz, Alkis Gotovos, Stefanie Jegelka, John Quackenbush</i>	

Do Neural Optimal Transport Solvers Work? a Continuous Wasserstein-2 Benchmark	14593
<i>Evgeny Burnaev, Alexander Filippov, Aude Genevay, Alexander Korotin, Lingxiao Li, Justin Solomon</i>	
Linear Convergence in Federated Learning: Tackling Client Heterogeneity and Sparse Gradients	14606
<i>Hamed Hassani, Rayana Jaafar, Aritra Mitra, George J. Pappas</i>	
On the Convergence of Prior-Guided Zeroth-Order Optimization Algorithms.....	14620
<i>Shuyu Cheng, Guoqiang Wu, Jun Zhu</i>	
Revisit Multimodal Meta-Learning Through the Lens of Multi-Task Learning	14632
<i>Milad Abdollahzadeh, Ngai-Man (Man) Cheung, Toubia Malekzadeh</i>	
Dynamic Sasvi: Strong Safe Screening for Norm-Regularized Least Squares.....	14645
<i>Makoto Yamada, Hiroaki Yamada</i>	
What Matters for Adversarial Imitation Learning?	14656
<i>Marcin Andrychowicz, Olivier Bachem, Robert Dadashi, Matthieu Geist, Sertan Girgin, Leonard Hussenot, Manu Orsini, Olivier Pietquin, Anton Raichuk, Damien Vincent</i>	
Sequential Causal Imitation Learning with Unobserved Confounders.....	14669
<i>Elias Bareinboim, Daniel Kumor, Junzhe Zhang</i>	
Topic Modeling Revisited: a Document Graph-Based Neural Network Perspective.....	14681
<i>Zheng Dong, Chuan Qin, Dazhong Shen, Chao Wang, Hui Xiong, Hengshu Zhu</i>	
Hard-Attention for Scalable Image Classification.....	14694
<i>Pawel Korus, Nasir Memon, Athanasios Papadopoulos</i>	
Fast Routing Under Uncertainty: Adaptive Learning in Congestion Games Via Exponential Weights.....	14708
<i>Kimon Antonakopoulos, Panayotis Mertikopoulos, Dong Quan Vu</i>	
Profiling Pareto Front with Multi-Objective Stein Variational Gradient Descent.....	14721
<i>Qiang Liu, Xingchao Liu, Xin Tong</i>	
MAP Propagation Algorithm: Faster Learning with a Team of Reinforcement Learning Agents	14734
<i>Stephen Chung</i>	
TransGAN: Two Pure Transformers Can Make One Strong GAN, and that Can Scale Up.....	14745
<i>Shiyu Chang, Yifan Jiang, Zhangyang Wang</i>	
A Central Limit Theorem for Differentially Private Query Answering	14759
<i>Jinshuo Dong, Weijie Su, Linjun Zhang</i>	
Differential Privacy Dynamics of Langevin Diffusion and Noisy Gradient Descent.....	14771
<i>Rishav Chourasia, Reza Shokri, Jiayuan Ye</i>	
Data Driven Semi-Supervised Learning	14782
<i>Maria-Florina Balcan, Dravyansh Sharma</i>	
Meta-Learning Via Learning with Distributed Memory	14795
<i>Sudarshan Babu, Michael Maire, Pedro Savarese</i>	
Physics-Integrated Variational Autoencoders for Robust and Interpretable Generative Modeling.....	14809
<i>Alexandros Kalousis, Naoya Takeishi</i>	

Characterizing the Risk of Fairwashing	14822
<i>Hiromi Arai, Ulrich Aivodji, Sebastien Gambs, Satoshi Hara</i>	
Qimera: Data-Free Quantization with Synthetic Boundary Supporting Samples.....	14835
<i>Kanghyun Choi, Deokki Hong, Youngsok Kim, Jinho Lee, Noseong Park</i>	
Embedding Principle of Loss Landscape of Deep Neural Networks.....	14848
<i>Tao Luo, Zhiqin J Xu, Yaoyu Zhang, Zhongwang Zhang</i>	
Adversarial Reweighting for Partial Domain Adaptation.....	14860
<i>Xiang Gu, Jian Sun, Zongben Xu, Xi Yu, Yan Yang</i>	
M-FAC: Efficient Matrix-Free Approximations of Second-Order Information.....	14873
<i>Dan Alistarh, Elias Frantar, Eldar Kurtic</i>	
Graph Adversarial Self-Supervised Learning.....	14887
<i>Longqi Yang, Wenjing Yang, Liangliang Zhang</i>	
Anti-Backdoor Learning: Training Clean Models on Poisoned Data.....	14900
<i>Nodens Koren, Bo Li, Yige Li, Lingjuan Lyu, Xixiang Lyu, Xingjun Ma</i>	
Locally Most Powerful Bayesian Test for Out-Of-Distribution Detection Using Deep Generative Models.....	14913
<i>Heeyoung Kim, Keunseo Kim, Juncheol Shin</i>	
Stable Neural ODE with Lyapunov-Stable Equilibrium Points for Defending Against Adversarial Attacks.....	14925
<i>Qinxu Ding, Qiyu Kang, Yang Song, Wee Peng Tay</i>	
Robust Compressed Sensing MRI with Deep Generative Priors.....	14938
<i>Marius Arvinte, Giannis Daras, Alexandros Dimakis, Ajil Jalal, Eric Price, Jon Tamir</i>	
H-NeRF: Neural Radiance Fields for Rendering and Temporal Reconstruction of Humans in Motion.....	14955
<i>Thiemo Alldieck, Cristian Sminchisescu, Hongyi Xu</i>	
DOBF: a Deobfuscation Pre-Training Objective for Programming Languages.....	14967
<i>Marie-Anne Lachaux, Guillaume Lample, Baptiste Roziere, Marc Szafraniec</i>	
Detecting Errors and Estimating Accuracy on Unlabeled Data with Self-Training Ensembles.....	14980
<i>Besim Avci, Jiefeng Chen, Somesh Jha, Yingyu Liang, Frederick Liu, Xi Wu</i>	
Exploiting Chain Rule and Bayes' Theorem to Compare Probability Distributions.....	14993
<i>Huangjie Zheng, Mingyuan Zhou</i>	
Actively Identifying Causal Effects with Latent Variables Given Only Response Variable Observable.....	15007
<i>Tian-Zuo Wang, Zhi-Hua Zhou</i>	
Interventional Sum-Product Networks: Causal Inference with Tractable Probabilistic Models	15019
<i>Devendra Dhami, Athresh Karanam, Kristian Kersting, Sriraam Natarajan, Matej Zecevic</i>	
PettingZoo: Gym for Multi-Agent Reinforcement Learning.....	15032
<i>Benjamin Black, Clemens Dieffendahl, Nathaniel Grammel, Ananth Hari, Caroline Horsch, Mario Jayakumar, Yashas Lokesh, Rodrigo Perez-Vicente, Praveen Ravi, Luis S Santos, Ryan Sullivan, J Terry, Niall Williams</i>	

Parametric Complexity Bounds for Approximating PDEs with Neural Networks.....	15044
<i>Zachary Lipton, Tanya Marwah, Andrej Risteski</i>	
Learning-To-learn Non-Convex piecewise-Lipschitz Functions.....	15056
<i>Maria-Florina Balcan, Mikhail Khodak, Dravyansh Sharma, Ameet Talwalkar</i>	
Uncertain Decisions Facilitate Better Preference Learning.....	15070
<i>Cassidy Laidlaw, Stuart Russell</i>	
Decision Transformer: Reinforcement Learning Via Sequence Modeling.....	15084
<i>Pieter Abbeel, Lili Chen, Aditya Grover, Misha Laskin, Kimin Lee, Kevin Lu, Igor Mordatch, Aravind Rajeswaran, Aravind Srinivas</i>	
Probability Paths and the Structure of Predictions Over Time	15098
<i>Sharad Goel, Zhiyuan Jerry Lin, Hao Sheng</i>	
Deep Extended Hazard Models for Survival Analysis	15111
<i>Jonas Mueller, Jane-Ling Wang, Qixian Zhong</i>	
TNASP: a Transformer-Based NAS Predictor with a Self-evolution Framework	15125
<i>Jixiang Li, Ji Liu, Shun Lu, Jianchao Tan, Sen Yang</i>	
Automorphic Equivalence-Aware Graph Neural Network.....	15138
<i>Pan Hui, Yong Li, Fengli Xu, Quanming Yao</i>	
Random Shuffling Beats SGD Only After Many Epochs on Ill-Conditioned Problems	15151
<i>Itay Safran, Ohad Shamir</i>	
Analytic Study of Families of Spurious Minima in Two-Layer ReLU Neural Networks: A Tale of Symmetry II.....	15162
<i>Yossi Arjevani, Michael Field</i>	
CAM-GAN: Continual Adaptation Modules for Generative Adversarial Networks.....	15175
<i>Lawrence Carin, Piyush Rai, P. K. Srijith, Sakshi Varshney, Vinay Kumar Verma</i>	
Structured Dropout Variational Inference for Bayesian Neural Networks	15188
<i>Hung Bui, Nhat Ho, Duong Nguyen, Khai Nguyen, Son Nguyen, Khoat Than</i>	
Neural Relightable Participating Media Rendering.....	15203
<i>Hans-Peter Seidel, Gurprit Singh, Quan Zheng</i>	

VOLUME 19

Efficient Neural Network Training Via Forward and Backward Propagation Sparsification	15216
<i>Zonghao Chen, Shizhe Diao, Tong Zhang, Weizhong Zhang, Xiao Zhou</i>	
Learning to Ground Multi-Agent Communication with Autoencoders.....	15230
<i>Jacob Huh, Phillip Isola, Ser Nam Lim, Toru Lin, Christopher Stauffer</i>	
Large-Scale Wasserstein Gradient Flows	15243
<i>Evgeny Burnaev, Aude Genevay, Alexander Korotin, Lingxiao Li, Petr Mokrov, Justin Solomon</i>	
Who Leads and Who Follows in Strategic Classification?.....	15257
<i>Michael Jordan, Eric Mazumdar, Shankar Sastry, Tijana Zrnic</i>	

Unadversarial Examples: Designing Objects for Robust Vision	15270
<i>Logan Engstrom, Andrew Ilyas, Ashish Kapoor, Aleksander Madry, Hadi Salman, Sai Vemprala</i>	
Deep Jump Learning for Off-Policy Evaluation in Continuous Treatment Settings	15285
<i>Hengrui Cai, Wenbin Lu, Chengchun Shi, Rui Song</i>	
Attention Approximates Sparse Distributed Memory	15301
<i>Trenton Bricken, Cengiz Pehlevan</i>	
Augmented Shortcuts for Vision Transformers.....	15316
<i>Yiping Deng, Kai Han, Yehui Tang, Yunhe Wang, An Xiao, Chang Xu, Chao Xu</i>	
Finding Regions of Heterogeneity in Decision-Making Via Expected Conditional Covariance.....	15328
<i>Saul Blecker, Leora Horwitz, Christina X Ji, Justin Lim, Michael Oberst, David Sontag</i>	
Identifying and Benchmarking Natural Out-Of-Context Prediction Problems.....	15344
<i>David Madras, Richard Zemel</i>	
Label Disentanglement in Partition-Based Extreme Multilabel Classification	15359
<i>Wei-Cheng Chang, Inderjit Dhillon, Cho-Jui Hsieh, Xuanqing Liu, Hsiang-Fu Yu</i>	
Leveraging SE(3) Equivariance for Self-Supervised Category-Level Object Pose Estimation from Point Clouds	15370
<i>A. Abbott, Leonidas Guibas, Xiaolong Li, Shuran Song, He Wang, Yijia Weng, Li Yi</i>	
A Theoretical Analysis of Fine-Tuning with Linear Teachers	15382
<i>Alon Brutzkus, Amir Globerson, Gal Shachaf</i>	
Overinterpretation Reveals Image Classification Model Pathologies	15395
<i>Brandon Carter, David Gifford, Siddhartha Jain, Jonas Mueller</i>	
Neural Circuit Synthesis from Specification Patterns	15408
<i>Bernd Finkbeiner, Christopher Hahn, Markus N Rabe, Frederik Schmitt</i>	
Directional Message Passing on Molecular Graphs Via Synthetic Coordinates	15421
<i>Johannes Gasteiger, Stephan Gunnemann, Chandan Yeshwanth</i>	
Federated Multi-Task Learning Under a Mixture of Distributions.....	15434
<i>Aurelien Bellet, Laetitia Kameni, Othmane Marfoq, Giovanni Neglia, Richard Vidal</i>	
Learning Generative Vision Transformer with Energy-Based Latent Space for Saliency Prediction	15448
<i>Nick Barnes, Ping Li, Jianwen Xie, Jing Zhang</i>	
Regularization in ResNet with Stochastic Depth.....	15464
<i>Fadhel Ayed, Soufiane Hayou</i>	
ResT: an Efficient Transformer for Visual Recognition.....	15475
<i>Yu-Bin Yang, Qinglong Zhang</i>	
Adversarial Examples for k-Nearest Neighbor Classifiers Based on Higher-Order Voronoi Diagrams	15486
<i>Evgenios Kornaropoulos, Chawin Sitawarin, Dawn Song, David Wagner</i>	
Adversarially Robust 3D Point Cloud Recognition Using Self-Supervisions	15498
<i>Anima Anandkumar, Yulong Cao, Christopher B Choy, Zhuoqing Morley Mao, Jiachen Sun, Chaowei Xiao, Zhiding Yu</i>	

Tuning Mixed Input Hyperparameters on the Fly for Efficient Population Based AutoRL.....	15513
<i>Shaan Desai, Vu Nguyen, Jack Parker-Holder, Stephen J Roberts</i>	
Neural Algorithmic Reasoners Are Implicit Planners	15529
<i>Pierre-Luc Bacon, Andreea-Ioana Deac, Ognjen Milinkovic, Mladen Nikolic, Jian Tang, Petar Velickovic</i>	
Self-Supervised Learning with Kernel Dependence Maximization	15543
<i>Arthur Gretton, Yazhe Li, Roman Pogodin, [deadname] Sutherland</i>	
CROCS: Clustering and Retrieval of Cardiac Signals Based on Patient Disease Class, Sex, and Age.....	15557
<i>David Clifton, Dani Kiyasseh, Tingting Zhu</i>	
Representing Hyperbolic Space Accurately Using Multi-Component Floats	15570
<i>Christopher De Sa, Tao Yu</i>	
Dimensionality Reduction for Wasserstein Barycenter.....	15582
<i>Zachary Izzo, Sandeep Silwal, Samson Zhou</i>	
Neural Population Geometry Reveals the Role of Stochasticity in Robust Perception.....	15595
<i>Sueyeon Chung, David Cox, Joel Dapello, James J Dicarlo, Jenelle Feather, Hang Le, Tiago Marques, Josh McDermott</i>	
Unsupervised Learning of Compositional Energy Concepts.....	15608
<i>Yilun Du, Shuang Li, Igor Mordatch, Yash Sharma, Josh Tenenbaum</i>	
Nearly Horizon-Free Offline Reinforcement Learning	15621
<i>Bo Dai, Simon Du, Jialian Li, Tongzheng Ren, Sujay Sanghavi</i>	
Combinatorial Optimization for Panoptic Segmentation: a Fully Differentiable Approach.....	15635
<i>Ahmed Abbas, Paul Swoboda</i>	
Reinforcement Learning with State Observation Costs in Action-Contingent Noiselessly Observable Markov Decision Processes.....	15650
<i>Emma Brunskill, Scott Fleming, Hyunji Alex Nam</i>	
Iterative Amortized Policy Optimization.....	15667
<i>Alessandro Davide Ialongo, Joseph Marino, Alexandre Piche, Yisong Yue</i>	
Revisiting the Calibration of Modern Neural Networks.....	15682
<i>Josip Djolonga, Neil Houlsby, Frances Hubis, Mario Lucic, Matthias Minderer, Rob Romijnders, Dustin Tran, Xiaohua Zhai</i>	
The Decomposition of the Higher-Order Homology Embedding Constructed from the k-Laplacian.....	15695
<i>Yu-Chia Chen, Marina Meila</i>	
Breaking the Moments Condition Barrier: No-Regret Algorithm for Bandits with Super Heavy- Tailed Payoffs.....	15710
<i>Jiayi Huang, Liwei Wang, Lin Yang, Han Zhong</i>	
A Nonparametric Method for Gradual Change Problems with Statistical Guarantees.....	15721
<i>Dan Nicolae, Lizhen Nie</i>	
Nested Graph Neural Networks.....	15734
<i>Pan Li, Muhan Zhang</i>	

Multimodal and Multilingual Embeddings for Large-Scale Speech Mining.....	15748
<i>Paul-Ambroise Duquenne, Hongyu Gong, Holger Schwenk</i>	
Necessary and Sufficient Graphical Conditions for Optimal Adjustment Sets in Causal Graphical Models with Hidden Variables.....	15762
<i>Jakob Runge</i>	
On Blame Attribution for Accountable Multi-Agent Sequential Decision Making	15774
<i>Goran Radanovic, Adish Singla, Stelios Triantafyllou</i>	
FLEX: Unifying Evaluation for Few-Shot NLP.....	15787
<i>Iz Beltagy, Jonathan Bragg, Arman Cohan, Kyle Lo</i>	
A Flow-Based Latent State Generative Model of Neural Population Responses to Natural Images	15801
<i>Mohammad Bashiri, Zhiwei Ding, Zhuokun Ding, Akshay Jagadish, Konstantin-Klemens Lurz, Taliah Muhammad, Fabian Sinz, Andreas Tolias, Edgar Walker</i>	
Learnable Fourier Features for Multi-Dimensional Spatial Positional Encoding.....	15816
<i>Samy Bengio, Cho-Jui Hsieh, Gang Li, Yang Li, Si Si</i>	
Doubly Robust Thompson Sampling with Linear Payoffs.....	15830
<i>Gi-Soo Kim, Wonyoung Kim, Myunghee Cho Paik</i>	
A Computationally Efficient Method for Learning Exponential Family Distributions	15841
<i>Devavrat Shah, Abhin Shah, Gregory Wornell</i>	
Rethinking Neural Operations for Diverse Tasks.....	15855
<i>Tri Dao, Mikhail Khodak, Liam Li, Nicholas Roberts, Christopher Re, Ameet Talwalkar</i>	
Motif-Based Graph Self-Supervised Learning for Molecular Property Prediction	15870
<i>Chee-Kong Lee, Qi Liu, Chengqiang Lu, Hao Wang, Zaixi Zhang</i>	
On Inductive Biases for Heterogeneous Treatment Effect Estimation.....	15883
<i>Alicia Curth, Mihaela Van Der Schaar</i>	
DP-SSL: Towards Robust Semi-supervised Learning with a Few Labeled Samples.....	15895
<i>Jiandong Ding, Yi Xu, Lu Zhang, Shuigeng Zhou</i>	
Transformer in Transformer	15908
<i>Jianyuan Guo, Kai Han, Yunhe Wang, Enhua Wu, Chunjing Xu, An Xiao</i>	
Adversarial Graph Augmentation to Improve Graph Contrastive Learning.....	15920
<i>Cong Hao, Pan Li, Jennifer Neville, Susheel Suresh</i>	
Online Control of Unknown Time-Varying Dynamical Systems.....	15934
<i>Paula Gradu, Elad Hazan, Edgar Minasyan, Max Simchowitz</i>	
Contrastive Reinforcement Learning of Symbolic Reasoning Domains	15946
<i>Wenxin Dong, Noah Goodman, Gabriel Poesia</i>	
Spatial Ensemble: a Novel Model Smoothing Mechanism for Student-Teacher Framework	15957
<i>Tengteng Huang, Yifan Sun, Xun Wang, Haotian Yao, Chi Zhang</i>	
Probabilistic Tensor Decomposition of Neural Population Spiking Activity.....	15969
<i>Sepideh Keshavarzi, Troy Margrie, Maneesh Sahani, Hugo Soulat</i>	

Recurrent Bayesian Classifier Chains for Exact Multi-Label Classification.....	15981
<i>Emmanuel Agu, Luke Buquicchio, Walter Gerych, Tom Hartvigsen, Elke A. Rundensteiner</i>	
Wasserstein Flow Meets Replicator Dynamics: a Mean-Field Analysis of Representation Learning in Actor-Critic	15993
<i>Siyu Chen, Michael Jordan, Zhaoran Wang, Zhuoran Yang, Yufeng Zhang</i>	
Assessing Fairness in the Presence of Missing Data	16007
<i>Qi Long, Yiliang Zhang</i>	
Adversarial Attack Generation Empowered by Min-Max Optimization.....	16020
<i>Pin-Yu Chen, Makan Fardad, Bo Li, Sijia Liu, Jingkang Wang, Jiachen Xu, Tianyun Zhang</i>	
Safe Pontryagin Differentiable Programming	16034
<i>Wanxin Jin, Shaoshuai Mou, George J. Pappas</i>	
Class-Disentanglement and Applications in Adversarial Detection and Defense	16051
<i>Dacheng Tao, Xinmei Tian, Kaiwen Yang, Yonggang Zhang, Tianyi Zhou</i>	
Active 3D Shape Reconstruction from Vision and Touch.....	16064
<i>Roberto Calandra, Michal Drozdal, Jitendra Malik, David Meger, Luis Pineda, Adriana Romero Soriano, Edward Smith</i>	
VOLUME 20	
CAPE: Encoding Relative Positions with Continuous Augmented Positional Embeddings	16079
<i>Ronan Collobert, Tatiana Likhomanenko, Alex Rogozhnikov, Gabriel Synnaeve, Qiantong Xu</i>	
Multi-Armed Bandit Requiring Monotone Arm Sequences	16093
<i>Ningyuan Chen</i>	
Gradient Driven Rewards to Guarantee Fairness in Collaborative Machine Learning.....	16104
<i>Chuan Sheng Foo, Bryan Kian Hsiang Low, Lingjuan Lyu, Xingjun Ma, Chenglin Miao, Xinyi Xu</i>	
Generalizable Imitation Learning from Observation Via Inferring Goal Proximity	16118
<i>Youngwoon Lee, Joseph Lim, Shao-Hua Sun, Andrew Szot</i>	
DualNet: Continual Learning, Fast and Slow.....	16131
<i>Steven Hoi, Chenghao Liu, Quang Pham</i>	
Deformable Butterfly: a Highly Structured and Sparse Linear Transform.....	16145
<i>Graziano Chesi, King Hung Chiu, Rui Lin, Jie Ran, Ngai Wong</i>	
Why Do Pretrained Language Models Help in Downstream Tasks? an Analysis of Head and Prompt Tuning.....	16158
<i>Tengyu Ma, Colin Wei, Sang Michael Xie</i>	
Learning Diverse Policies in MOBA Games Via Macro-Goals	16171
<i>Guangwei Chen, Xueying Du, Qiang Fu, Yiming Gao, Guoan Han, Lanxiao Huang, Zhenjie Lian, Fuhao Qiu, Bei Shi, Liang Wang, Weixuan Wang, Wei Yang, Deheng Ye</i>	
Evaluation of Human-AI Teams for Learned and Rule-Based Agents in Hanabi	16183
<i>Ross Allen, Kimberlee Chang, Edenna Chen, Victor Lopez, Kyle Palko, Jaime Peña, Ho Chit Siu, Yutai Zhou</i>	

Counterfactual Invariance to Spurious Correlations in Text Classification.....	16196
<i>Alexander D'Amour, Jacob Eisenstein, Victor Veitch, Steve Yadlowsky</i>	
Better Safe than Sorry: Preventing Delusive Adversaries with Adversarial Training.....	16209
<i>Songcan Chen, Lei Feng, Sheng-Jun Huang, Lue Tao, Jinfeng Yi</i>	
Determinantal Point Processes Based on Orthogonal Polynomials for Sampling Minibatches in SGD.....	16226
<i>Remi Bardenet, Subhroshekhar Ghosh, Meixia Lin</i>	
Revisiting Contrastive Methods for Unsupervised Learning of Visual Representations.....	16238
<i>Stamatios Georgoulis, Luc V Gool, Wouter Van Gansbeke, Simon Vandenhende</i>	
Neural Analysis and Synthesis: Reconstructing Speech from Self-Supervised Representations	16251
<i>Hyeong-Seok Choi, Hoon Heo, Wansoo Kim, Jie Lee, Juheon Lee, Kyogu Lee</i>	
Auto-Encoding Knowledge Graph for Unsupervised Medical Report Generation	16266
<i>Shen Ge, Fenglin Liu, Xu Sun, Xian Wu, Chenyu You, Sheng Wang</i>	
Diffusion Normalizing Flow	16280
<i>Yongxin Chen, Qinsheng Zhang</i>	
Introspective Distillation for Robust Question Answering.....	16292
<i>Yulei Niu, Hanwang Zhang</i>	
Rethinking the Pruning Criteria for Convolutional Neural Network.....	16305
<i>Zhongzhan Huang, Liang Lin, Ping Luo, Wenqi Shao, Xinjiang Wang</i>	
Adaptive Machine Unlearning.....	16319
<i>Varun Gupta, Christopher Jung, Seth Neel, Aaron Roth, Saeed Sharifi-Malvajerdi, Chris Waites</i>	
EditGAN: High-Precision Semantic Image Editing	16331
<i>Sanja Fidler, Seung Wook Kim, Karsten Kreis, Daiqing Li, Huan Ling, Antonio Torralba</i>	
Deep Molecular Representation Learning Via Fusing Physical and Chemical Information	16346
<i>Lingsheng Cai, Ziyao Li, Guojie Song, Shuwen Yang</i>	
Neural Optimal Feedback Control with Local Learning Rules	16358
<i>Dmitri Chklovskii, Shiva Farashahi, Johannes Friedrich, Alexander Genkin, Siavash Golkar, Anirvan Sengupta</i>	
Reinforcement Learning in Linear MDPs: Constant Regret and Representation Selection	16371
<i>Alessandro Lazaric, Aldo Pacchiano, Matteo Papini, Matteo Pirotta, Marcello Restelli, Andrea Tirinzoni</i>	
Noether Networks: Meta-Learning Useful Conserved Quantities.....	16384
<i>Ferran Alet, Dylan Doblar, Chelsea Finn, Kenji Kawaguchi, Josh Tenenbaum, Allan Zhou</i>	
Uncertainty-Driven Loss for Single Image Super-Resolution.....	16398
<i>Weisheng Dong, Xin Li, Qian Ning, Guangming Shi, Jinjian Wu</i>	
GradInit: Learning to Initialize Neural Networks for Stable and Efficient Training.....	16410
<i>Tom Goldstein, W. Ronny Huang, Kezhi Kong, Renkun Ni, Zheng Xu, Chen Zhu</i>	

Capacity and Bias of Learned Geometric Embeddings for Directed Graphs	16423
<i>Michael Boratko, Kenneth L Clarkson, Andrew McCallum, Nicholas Monath, Luke Vilnis, Dongxu Zhang</i>	
Online Learning of Neural Computations from Sparse Temporal Feedback	16437
<i>Lukas Braun, Tim Vogels</i>	
Self-Supervised Learning with Data Augmentations Provably Isolates Content from Style.....	16451
<i>Michel Besserve, Wieland Brendel, Luigi Gresele, Francesco Locatello, Bernhard Scholkopf, Yash Sharma, Julius Von Kugelgen</i>	
Instance-Conditional Knowledge Distillation for Object Detection.....	16468
<i>Zijian Kang, Jian Sun, Peizhen Zhang, Xiangyu Zhang, Nanning Zheng</i>	
Self-Supervised Representation Learning on Neural Network Weights for Model Characteristic Prediction	16481
<i>Damian Borth, Dimche Kostadinov, Konstantin Schurholt</i>	
Multimodal Virtual Point 3D Detection	16494
<i>Philipp Krahenbuhl, Tianwei Yin, Xingyi Zhou</i>	
On Joint Learning for Solving Placement and Routing in Chip Design	16508
<i>Ruoyu Cheng, Junchi Yan</i>	
Learning with Algorithmic Supervision Via Continuous Relaxations	16520
<i>Christian Borgelt, Oliver Deussen, Hilde Kuehne, Felix Petersen</i>	
Differentiable Multiple Shooting Layers.....	16532
<i>Hajime Asama, Stefano Massaroli, Jinkyoo Park, Michael Poli, Sho Sonoda, Taiji Suzuki, Atsushi Yamashita</i>	
Global-Aware Beam Search for Neural Abstractive Summarization	16545
<i>Kaizhu Huang, Zixun Lan, Ye Ma, Lu Zong</i>	
DROID-SLAM: Deep Visual SLAM for Monocular, Stereo, and RGB-D Cameras	16558
<i>Jia Deng, Zachary Teed</i>	
Few-Shot Object Detection Via Association and Discrimination	16570
<i>Yuhang Cao, Kai Chen, Ying Jin, Dahua Lin, Ziwei Liu, Jiaqi Wang, Tong Wu</i>	
Neural Dubber: Dubbing for Videos According to Scripts	16582
<i>Chenxu Hu, Tingle Li, Qiao Tian, Yuxuan Wang, Wang Yuping, Hang Zhao</i>	
Neural Bootstrapper.....	16596
<i>Hyungjoo Cho, Sungbin Lim, Hyun-Seok Min, Minsuk Shin</i>	
An Axiomatic Theory of Provably-Fair Welfare-Centric Machine Learning	16610
<i>Cyrus Cousins</i>	
HSVA: Hierarchical Semantic-Visual Adaptation for Zero-Shot Learning	16622
<i>Shiming Chen, Hao Li, Yang Liu, Qinmu Peng, Ling Shao, Baigui Sun, Guosen Xie, Xinge You</i>	
Higher Order Kernel Mean Embeddings to Capture Filtrations of Stochastic Processes	16635
<i>Theodoros Damoulas, Blanka Horvath, Maud Lemerrier, Chong Liu, Terry Lyons, Cristopher Salvi</i>	

Low-Rank Subspaces in GANs	16648
<i>Qifeng Chen, Ruili Feng, Yujun Shen, Zheng-Jun Zha, Deli Zhao, Jingren Zhou, Jiapeng Zhu</i>	
Neural Symplectic Form: Learning Hamiltonian Equations on General Coordinate Systems	16659
<i>Yuhan Chen, Takashi Matsubara, Takaharu Yaguchi</i>	
Sample-Efficient Reinforcement Learning is Feasible for Linearly Realizable MDPs with Limited Revisiting	16671
<i>Yuxin Chen, Yuejie Chi, Yuantao Gu, Gen Li, Yuting Wei</i>	
Self-Paced Contrastive Learning for Semi-supervised Medical Image Segmentation with Meta-labels.....	16686
<i>Christian Desrosiers, Marco Pedersoli, Jizong Peng, Ping Wang</i>	
Reverse Engineering Recurrent Neural Networks with Jacobian Switching Linear Dynamical Systems.....	16700
<i>Scott Linderman, Jimmy Smith, David Sussillo</i>	
Learning-Augmented Dynamic Power Management with Multiple States Via New Ski Rental Bounds.....	16714
<i>Antonios Antoniadis, Christian Coester, Marek Elias, Adam Polak, Bertrand Simon</i>	
Learning Equivariant Energy Based Models with Equivariant Stein Variational Gradient Descent.....	16727
<i>Lars Holdijk, Priyank Jaini, Max Welling</i>	
Information Directed Sampling for Sparse Linear Bandits	16738
<i>Wei Deng, Botao Hao, Tor Lattimore</i>	
Linear Convergence of Gradient Methods for Estimating Structured Transition Matrices in High-Dimensional Vector Autoregressive Models.....	16751
<i>Wei Cui, Yulong Liu, Xiao Lv</i>	
Large-Scale Unsupervised Object Discovery	16764
<i>Jean Ponce, Patrick Perez, Cordelia Schmid, Elena Sizikova, Van Huy Vo</i>	
Sparse Steerable Convolutions: an Efficient Learning of SE(3)-Equivariant Features for Estimation and Tracking of Object Poses in 3D Space	16779
<i>Ke Chen, Kui Jia, Hongyang Li, Jiehong Lin, Jiangbo Lu</i>	
Noisy Adaptation Generates Levy Flights in Attractor Neural Networks	16791
<i>Tianhao Chu, Xingsi Dong, Tiejun Huang, Zilong Ji, Si Wu</i>	
On Linear Stability of SGD and Input-Smoothness of Neural Networks.....	16805
<i>Chao Ma, Lexing Ying</i>	
Joint Inference and Input Optimization in Equilibrium Networks	16818
<i>Shaojie Bai, Swaminathan Gurumurthy, J. Zico Kolter, Zachary Manchester</i>	
A Unified Framework for Bandit Multiple Testing.....	16833
<i>Aaditya Ramdas, Ruodu Wang, Ziyu Xu</i>	
Recovering Latent Causal Factor for Generalization to Distributional Shifts	16846
<i>Wei Chen, Chang Liu, Tie-Yan Liu, Tao Qin, Xinwei Sun, Botong Wu, Xiangyu Zheng</i>	
Graph Differentiable Architecture Search with Structure Learning	16860
<i>Yijian Qin, Xin Wang, Zeyang Zhang, Wenwu Zhu</i>	

Designing Counterfactual Generators Using Deep Model Inversion	16873
<i>Akshay Chaudhari, Jia Liang, Vivek Sivaraman Narayanaswamy, Deepta Rajan, Andreas Spanias, Jayaraman Thiagarajan</i>	
A Faster Maximum Cardinality Matching Algorithm with Applications in Machine Learning.....	16885
<i>Nathaniel Lahn, Sharath Raghvendra, Jiacheng Ye</i>	
Dynamic Population-Based Meta-learning for Multi-agent Communication with Natural Language	16899
<i>Abhinav Gupta, Marc Lanctot, Angeliki Lazaridou</i>	
Adversarial Neuron Pruning Purifies Backdoored Deep Models	16913
<i>Yisen Wang, Dongxian Wu</i>	
Towards Robust and Reliable Algorithmic Recourse.....	16926
<i>Shalmali Joshi, Himabindu Lakkaraju, Sohini Upadhyay</i>	

VOLUME 21

Neural Rule-Execution Tracking Machine for Transformer-Based Text Generation.....	16938
<i>Mark Dras, Huang Hu, Daxin Jiang, Mark Johnson, Chongyang Tao, Stephen Wan, Yufei Wang, Can Xu</i>	
Scalable Online Planning Via Reinforcement Learning Fine-Tuning	16951
<i>Brandon Amos, Noam Brown, Arnaud Fickinger, Hengyuan Hu, Stuart Russell</i>	
Adversarial Regression with Doubly Non-Negative Weighting Matrices.....	16964
<i>Jose Blanchet, Tam Le, Truyen Nguyen, Viet Anh Nguyen, Makoto Yamada</i>	
Learned Robust PCA: a Scalable Deep Unfolding Approach for High-Dimensional Outlier Detection	16977
<i>Hanqin Cai, Jialin Liu, Wotao Yin</i>	
Proxy-Normalizing Activations to Match Batch Normalization While Removing Batch Dependence	16990
<i>Zach Eaton-Rosen, Antoine Labatie, Carlo Luschi, Dominic Masters</i>	
Dynamic Bottleneck for Robust Self-Supervised Exploration	17007
<i>Chenjia Bai, Animesh Garg, Lei Han, Jianye Hao, Peng Liu, Zhaoran Wang, Lingxiao Wang</i>	
ProTo: Program-Guided Transformer for Program-Guided Tasks.....	17021
<i>Binghong Chen, Karan Samel, Zelin Zhao, Lee Song</i>	
An Efficient Transfer Learning Framework for Multiagent Reinforcement Learning	17037
<i>Yingfeng Chen, Changjie Fan, Jianye Hao, Yujing Hu, Dong Li, Wulong Liu, Hangyu Mao, Zhaopeng Meng, Hongyao Tang, Weixun Wang, Tianpei Yang, Chengwei Zhang</i>	
Learning to Time-Decode in Spiking Neural Networks Through the Information Bottleneck	17049
<i>Hyeryung Jang, Osvaldo Simeone, Nicolas Skachkovsky</i>	
NEO: Non Equilibrium Sampling on the Orbits of a Deterministic Transform	17060
<i>Arnaud Doucet, Alain Durmus, Yazid Janati El Idrissi, Sylvain Le Corff, Eric Moulines, Charles Ollion, Christian X Robert, Achille Thin</i>	
Relaxing Local Robustness	17072
<i>Matt Fredrikson, Klas Leino</i>	

Tuning Large Neural Networks Via Zero-Shot Hyperparameter Transfer	17084
<i>Igor Babuschkin, Weizhu Chen, David Farhi, Jianfeng Gao, Edward Hu, Xiaodong Liu, Jakub Pachocki, Nick Ryder, Szymon Sidor, Ge Yang</i>	
Statistical Regeneration Guarantees of the Wasserstein Autoencoder with Latent Space Consistency	17098
<i>Anish Chakrabarty, Swagatam Das</i>	
Leveraging the Inductive Bias of Large Language Models for Abstract Textual Reasoning	17111
<i>Christopher Rytting, David Wingate</i>	
Differentiable Simulation of Soft Multi-Body Systems	17123
<i>Vladlen Koltun, Junbang Liang, Ming Lin, Yiling Qiao</i>	
Good Classification Measures and How to Find Them	17136
<i>Martijn Gosgens, Liudmila Prokhorenkova, Aleksey Tikhonov, Anton Zhiyanov</i>	
Distilling Robust and Non-Robust Features in Adversarial Examples by Information Bottleneck.....	17148
<i>Junho Kim, Byung-Kwan Lee, Yong Man Ro</i>	
Vector-Valued Gaussian Processes on Riemannian Manifolds Via Gauge Independent Projected Kernels	17160
<i>Viacheslav Borovitskiy, Marc Deisenroth, Michael Hutchinson, So Takao, Yee Teh, Alexander Terenin</i>	
On the Representation Power of Set Pooling Networks	17170
<i>Christian Bueno, Alan Hylton</i>	
Learning Policies with Zero Or Bounded Constraint Violation for Constrained MDPs	17183
<i>Dileep Kalathil, Panganamala Kumar, Tao Liu, Chao Tian, Ruida Zhou</i>	
A Prototype-Oriented Framework for Unsupervised Domain Adaptation	17194
<i>Bo Chen, Xinjie Fan, Korawat Tanwisuth, Hao Zhang, Shujian Zhang, Huangjie Zheng, Mingyuan Zhou</i>	
Mining the Benefits of Two-Stage and One-Stage HOI Detection	17209
<i>Chen Gao, Xiaobo Li, Yue Liao, Si Liu, Miao Lu, Yongliang Wang, Aixi Zhang</i>	
Discerning Decision-Making Process of Deep Neural Networks with Hierarchical Voting Transformation	17221
<i>Qing He, Chuan Qin, Ying Sun, Hui Xiong, Hengshu Zhu, Fuzhen Zhuang</i>	
Risk-Averse Heteroscedastic Bayesian Optimization	17235
<i>Ilija Bogunovic, Andreas Krause, Anastasia Makarova, Ilnura Usmanova</i>	
Invertible DenseNets with Concatenated LipSwish	17246
<i>Sandjai Bhulai, Yura Perugachi-Diaz, Jakub Tomczak</i>	
Topological Detection of Trojaned Neural Networks	17258
<i>Chao Chen, Mayank Goswami, Hubert Wagner, Yikai Zhang, Songzhu Zheng</i>	
Provably Strict Generalisation Benefit for Invariance in Kernel Methods	17273
<i>Bryn Elesedy</i>	
Formalizing the Generalization-Forgetting Trade-off in Continual Learning	17284
<i>Prasanna Balaprakash, Krishnan Raghavan</i>	

Risk-Aware Transfer in Reinforcement Learning Using Successor Features	17298
<i>Andre Barreto, Michael Gimelfarb, Chi-Guhn Lee, Scott Sanner</i>	
Causal Inference for Event Pairs in Multivariate Point Processes.....	17311
<i>Kristin P Bennett, Debarun Bhattacharjya, Tian Gao, Nicholas Mattei, Xiao Shou, Dharmashankar Subramanian</i>	
Evaluating Model Performance Under Worst-Case Subpopulations.....	17325
<i>Mike Li, Hongseok Namkoong, Shangzhou Xia</i>	
Privately Publishable Per-Instance Privacy	17335
<i>Rachel Redberg, Yu-Xiang Wang</i>	
Understanding the Limits of Unsupervised Domain Adaptation Via Data Poisoning.....	17347
<i>Pin-Yu Chen, Jihun Hamm, Bhavya Kailkhura, Akshay Mehra</i>	
Coresets for Clustering with Missing Values	17360
<i>Vladimir Braverman, Shaofeng Jiang, Robert Krauthgamer, Xuan Wu</i>	
Boosting with Multiple Sources	17373
<i>Corinna Cortes, Mehryar Mohri, Dmitry Storcheus, Ananda Theertha Suresh</i>	
Dynamic Neural Representational Decoders for High-Resolution Semantic Segmentation	17388
<i>Chunhua Shen, Zhi Tian, Bowen Zhang, Yifan Liu</i>	
Dense Keypoints Via Multiview Supervision	17400
<i>Sujoy Ganguly, Hyun Soo Park, Long Sha, Haozheng Yu, Zhixuan Yu</i>	
Scatterbrain: Unifying Sparse and Low-Rank Attention	17413
<i>Beidi Chen, Tri Dao, Atri Rudra, Christopher Re, Zhao Song, Eric Winsor</i>	
PTR: A Benchmark for Part-Based Conceptual, Relational, and Physical Reasoning	17427
<i>Chuang Gan, Yining Hong, Josh Tenenbaum, Antonio Torralba, Li Yi</i>	
Property-Aware Relation Networks for Few-Shot Molecular Property Prediction	17441
<i>Abulikemu Abuduweili, Dejing Dou, Yaqing Wang, Quanming Yao</i>	
Differentially Private Learning with Adaptive Clipping	17455
<i>Galen Andrew, Brendan McMahan, Swaroop Ramaswamy, Om Thakkar</i>	
Can Less Be More? When Increasing-To-Balancing Label Noise Rates Considered Beneficial.....	17467
<i>Yang Liu, Jialu Wang</i>	
Projected GANs Converge Faster.....	17480
<i>Kashyap Chitta, Andreas Geiger, Jens Muller, Axel Sauer</i>	
Generating High-Quality Explanations for Navigation in Partially-Revealed Environments	17493
<i>Gregory Stein</i>	
De-Randomizing MCMC Dynamics with the Diffusion Stein Operator.....	17507
<i>Markus Heinonen, Samuel Kaski, Zheyang Shen</i>	
Sparsely Changing Latent States for Prediction and Planning in Partially Observable Domains	17518
<i>Martin V. Butz, Christian Gumbsch, Georg Martius</i>	
PreferenceNet: Encoding Human Preferences in Auction Design with Deep Learning.....	17532
<i>Michael Curry, John Dickerson, Samuel Dooley, Neehar Peri</i>	

Large-Scale Learning with Fourier Features and Tensor Decompositions.....	17543
<i>Kim Batselier, Frederiek Wesel</i>	
Hash Layers for Large Sparse Models.....	17555
<i>Stephen Roller, Sainbayar Sukhbaatar, Jason Weston, Arthur Szlam</i>	
Sliced Mutual Information: a Scalable Measure of Statistical Dependence	17567
<i>Ziv Goldfeld, Kristjan Greenewald</i>	
Emergent Communication Under Varying Sizes and Connectivities	17579
<i>Jooyeon Kim, Alice Oh</i>	
Deep Bandits Show-Off: Simple and Efficient Exploration with Deep Networks	17592
<i>Mattia Rigotti, Rong Zhu</i>	
Regret Minimization Experience Replay in Off-Policy Reinforcement Learning.....	17604
<i>Shengyi Jiang, Xu-Hui Liu, Jingcheng Pang, Feng Xu, Zhenghai Xue, Yang Yu</i>	
Relative Uncertainty Learning for Facial Expression Recognition	17616
<i>Weihong Deng, Chengrui Wang, Yuhang Zhang</i>	
An Information-Theoretic Approach to Distribution Shifts	17628
<i>Marco Federici, Patrick Forre, Ryota Tomioka</i>	
TRS: Transferability Reduced Ensemble Via Promoting Gradient Diversity and Model Smoothness	17642
<i>Qian Chen, Bo Li, Linyi Li, Benjamin Rubinstein, Xiaojun Xu, Zhuolin Yang, Ce Zhang, Pan Zhou, Shiliang Zuo</i>	
Towards Sample-Optimal Compressive Phase Retrieval with Sparse and Generative Priors	17656
<i>Subhroshekar Ghosh, Zhaoqiang Liu, Jonathan Scarlett</i>	
Moser Flow: Divergence-Based Generative Modeling on Manifolds	17669
<i>Aditya Grover, Yaron Lipman, Maximilian Nickel, Noam Rozen</i>	
Structure-Aware Random Fourier Kernel for Graphs	17681
<i>Jinyuan Fang, Shangsong Liang, Zaiqiao Meng, Qiang Zhang</i>	
Diffusion Schrodinger Bridge with Applications to Score-Based Generative Modeling	17695
<i>Valentin De Bortoli, Arnaud Doucet, Jeremy Heng, James Thornton</i>	
Improving Transferability of Representations Via Augmentation-Aware Self-Supervision.....	17710
<i>Honglak Lee, Hankook Lee, Kibok Lee, Kimin Lee, Jinwoo Shin</i>	
Long-Short Transformer: Efficient Transformers for Language and Vision.....	17723
<i>Anima Anandkumar, Bryan Catanzaro, Tom Goldstein, Wei Ping, Mohammad Shoeybi, Chaowei Xiao, Chen Zhu</i>	
Post-Training Sparsity-Aware Quantization.....	17737
<i>Freddy Gabbay, Samer Kurzum, Gil Shomron, Uri Weiser</i>	
The Implicit Bias of Minima Stability: A View from Function Space.....	17749
<i>Tomer Michaeli, Rotem Mulayoff, Daniel Soudry</i>	
Breaking the Sample Complexity Barrier to Regret-Optimal Model-Free Reinforcement Learning	17762
<i>Yuxin Chen, Yuejie Chi, Yuantao Gu, Gen Li, Laixi Shi</i>	

Robust Auction Design in the Auto-Bidding World	17777
<i>Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, Song Zuo</i>	

VOLUME 22

Weighted Model Estimation for Offline Model-Based Reinforcement Learning.....	17789
<i>Toru Hishinuma, Kei Senda</i>	
Practical, Provably-Correct Interactive Learning in the Realizable Setting: the Power of True Believers	17801
<i>Kevin Jamieson, Julian Katz-Samuels, Blake Mason, Rob Nowak</i>	
Deconditional Downscaling with Gaussian Processes	17813
<i>Shahine Bouabid, Siu Lun Chau, Dino Sejdinovic</i>	
Image Generation Using Continuous Filter Atoms	17826
<i>Seunghyun Hwang, Zichen Miao, Qiang Qiu, Ze Wang</i>	
Latent Equilibrium: Arbitrarily Fast Computation with Arbitrarily Slow Neurons	17839
<i>Benjamin Ellenberger, Paul Haider, Jakob Jordan, Laura Kriener, Mihai A. Petrovici, Walter Senn</i>	
Learning Fast-Inference Bayesian Networks.....	17852
<i>Vaidyanathan Peruvemba Ramaswamy, Stefan Szeider</i>	
Per-Pixel Classification is Not All You Need for Semantic Segmentation	17864
<i>Bowen Cheng, Alexander Kirillov, Alex Schwing</i>	
Deep Markov Factor Analysis: Towards Concurrent Temporal and Spatial Analysis of fMRI Data	17876
<i>Amirreza Farnoosh, Sarah Ostadabbas</i>	
BooVAE: Boosting Approach for Continual Learning of VAE.....	17889
<i>Evgeny Burnaev, Evgenii Egorov, Anna Kuzina</i>	
Handling Long-Tailed Feature Distribution in AdderNets.....	17902
<i>Xinghao Chen, Minjing Dong, Yunhe Wang, Chang Xu</i>	
Pessimism Meets Invariance: Provably Efficient Offline Mean-Field Multi-Agent RL	17913
<i>Minshuo Chen, Yan Li, Zhaoran Wang, Ethan Wang, Zhuoran Yang, Tuo Zhao</i>	
A Law of Iterated Logarithm for Multi-Agent Reinforcement Learning	17927
<i>Bhumesh Kumar, Gagan Chandrashekhkar Thoppe</i>	
MOMA: Multi-Object Multi-Actor Activity Parsing	17939
<i>Ehsan Adeli, Michael Cooper, Siddharth Kapoor, Fei-Fei Li, Yiyun Liang, Zelun Luo, Juan Carlos Niebles, Wanze Xie</i>	
The Pareto Frontier of Model Selection for General Contextual Bandits.....	17956
<i>Teodor Vanislavov Marinov, Julian Zimmert</i>	
Teaching an Active Learner with Contrastive Examples	17968
<i>Yuxin Chen, Adish Singla, Chaoqi Wang</i>	
Structured Denoising Diffusion Models in Discrete State-Spaces	17981
<i>Jacob Austin, Jonathan Ho, Daniel D. Johnson, Daniel Tarlow, Rianne Van Den Berg</i>	

Emergent Communication of Generalizations.....	17994
<i>Noah Goodman, Jesse Mu</i>	
Distributed Machine Learning with Sparse Heterogeneous Data.....	18008
<i>Sahand Negahban, Patrick Rebeschini, Dominic Richards</i>	
Manipulating SGD with Data Ordering Attacks.....	18021
<i>Ross J Anderson, Murat Erdogdu, Dmitry Kazhdan, Nicolas Papernot, I Shumailov, Zakhar Shumaylov, Yiren Zhao</i>	
Graph Posterior Network: Bayesian Predictive Uncertainty for Node Classification	18033
<i>Bertrand Charpentier, Simon Geisler, Stephan Gunnemann, Maximilian Stadler, Daniel Zugner</i>	
Locality Sensitive Teaching	18049
<i>Beidi Chen, Chaojian Li, Yingyan Lin, Weiyang Liu, Anshumali Shrivastava, Le Song, Zhaozhuo Xu</i>	
No-Press Diplomacy from Scratch.....	18063
<i>Anton Bakhtin, Noam Brown, Adam Lerer, David Wu</i>	
Remember What You Want to Forget: Algorithms for Machine Unlearning	18075
<i>Jayadev Acharya, Gautam Kamath, Ayush Sekhari, Ananda Theertha Suresh</i>	
Learning Latent Causal Graphs Via Mixture Oracles	18087
<i>Bryon Aragam, Bohdan Kivva, Goutham Rajendran, Pradeep Ravikumar</i>	
ErrorCompensatedX: Error Compensation for Variance Reduced Algorithms.....	18102
<i>Yao Li, Ji Liu, Hanlin Tang, Ming Yan</i>	
Deep Contextual Video Compression	18114
<i>Bin Li, Jiahao Li, Yan Lu</i>	
On the Frequency Bias of Generative Models.....	18126
<i>Andreas Geiger, Yiyi Liao, Katja Schwarz</i>	
Learning Curves of Generic Features Maps for Realistic Datasets with a Teacher-Student Model.....	18137
<i>Hugo Cui, Cedric Gerbelot, Sebastian Goldt, Florent Krzakala, Bruno Loureiro, Marc Mezard, Lenka Zdeborova</i>	
It Has Potential: Gradient-Driven Denoisers for Convergent Solutions to Inverse Problems	18152
<i>Yochai Blau, Regev Cohen, Daniel Freedman, Ehud Rivlin</i>	
Training Over-Parameterized Models with Non-decomposable Objectives	18165
<i>Aditya Menon, Harikrishna Narasimhan</i>	
Reinforcement Learning for Optimization of Variational Quantum Circuit Architectures	18182
<i>Vedran Dunjko, Wojciech Masarczyk, Mateusz Ostaszewski, Eleanor Scerri, Lea M. Trenkwalder</i>	
Moshpit SGD: Communication-Efficient Decentralized Training on Heterogeneous Unreliable Devices	18195
<i>Eduard Gorbunov, Gennady Pekhimenko, Vsevolod Plokhotnyuk, Max Ryabinin</i>	
IRM---when it Works and When it Doesn't: a Test Case of Natural Language Inference	18212
<i>Yonatan Belinkov, Yana Dranker, He He</i>	

Self-Supervised Learning Disentangled Group Representation as Feature	18225
<i>Jianqiang Huang, Qianru Sun, Tan Wang, Zhongqi Yue, Hanwang Zhang</i>	
SaLKG: Learning from Knowledge Graph Explanations for Commonsense Reasoning	18241
<i>Aaron Chan, Tanishq Gupta, Boyuan Long, Xiang Ren, Soumya Sanyal, Jiashu Xu</i>	
Supervising the Transfer of Reasoning Patterns in VQA	18256
<i>Grigory Antipov, Moez Baccouche, Corentin Kervadec, Madiha Nadri, Christian Wolf</i>	
Conformal Bayesian Computation	18268
<i>Edwin Fong, Chris C Holmes</i>	
A Unified Approach to Fair Online Learning Via Blackwell Approachability	18280
<i>Evgenii Chzhen, Christophe Giraud, Gilles Stoltz</i>	
Training Neural Networks is ER-Complete	18293
<i>Mikkel Abrahamsen, Linda Kleist, Tilmann Miltzow</i>	
Understanding the Under-Coverage Bias in Uncertainty Estimation	18307
<i>Yu Bai, Song Mei, Huan Wang, Caiming Xiong</i>	
Decentralized Q-Learning in Zero-sum Markov Games	18320
<i>Tamer Basar, David Leslie, Asuman Ozdaglar, Muhammed Sayin, Kaiqing Zhang</i>	
Fast Certified Robust Training with Short Warmup.....	18335
<i>Cho-Jui Hsieh, Zhouxing Shi, Yihan Wang, Jinfeng Yi, Huan Zhang</i>	
Vector-Valued Distance and Gyrocalculus on the Space of Symmetric Positive Definite Matrices.....	18350
<i>Federico Lopez, Beatrice Pozzetti, Michael Strube, Steve Trettel, Anna Wienhard</i>	
Improved Transformer for High-Resolution GANs	18367
<i>Ting Chen, Dimitris Metaxas, Han Zhang, Zizhao Zhang, Long Zhao</i>	
Learning High-Precision Bounding Box for Rotated Object Detection Via Kullback-Leibler Divergence.....	18381
<i>Qi Ming, Qi Tian, Wentao Wang, Junchi Yan, Jirui Yang, Xiaojiang Yang, Xue Yang</i>	
On Locality of Local Explanation Models	18395
<i>Karla Diazordaz, Sahra Ghalebikesabi, Chris C Holmes, Lucile Ter-Minassian</i>	
FlexMatch: Boosting Semi-Supervised Learning with Curriculum Pseudo Labeling.....	18408
<i>Wenxin Hou, Manabu Okumura, Takahiro Shinozaki, Hao Wu, Jindong Wang, Yidong Wang, Bowen Zhang</i>	
Relative Flatness and Generalization	18420
<i>Linara Adilova, Mario Boley, Michael Kamp, Henning Petzka, Cristian Sminchisescu</i>	
The Image Local Autoregressive Transformer.....	18433
<i>Chenjie Cao, Yanwei Fu, Yuxin Hong, Xiang Li, Chengrong Wang, Chengming Xu, Xiangyang Xue</i>	
Towards Multi-Grained Explainability for Graph Neural Networks.....	18446
<i>Tat-Seng Chua, Xiangnan He, Xiang Wang, Yingxin Wu, An Zhang</i>	
Behavior from the Void: Unsupervised Active Pre-Training	18459
<i>Pieter Abbeel, Hao Liu</i>	

Autonomous Reinforcement Learning Via Subgoal Curricula.....	18474
<i>Chelsea Finn, Abhishek Gupta, Karol Hausman, Sergey Levine, Archit Sharma</i>	
Statistically and Computationally Efficient Linear Meta-Representation Learning	18487
<i>Prateek Jain, Praneeth Netrapalli, Sewoong Oh, Kiran Thekumparampil</i>	
Decentralized Learning in Online Queuing Systems.....	18501
<i>Etienne Boursier, Vianney Perchet, Flore Sentenac</i>	
Explainable Semantic Space by Grounding Language to Vision with Cross-Modal Contrastive Learning	18513
<i>Minkyu Choi, Kuan Han, Zhongming Liu, Yizhen Zhang</i>	
BulletTrain: Accelerating Robust Neural Network Training Via Boundary Example Mining	18527
<i>Chuan Guo, Weizhe Hua, G. Edward Suh, Yichi Zhang, Zhiru Zhang</i>	
Neural Distance Embeddings for Biological Sequences	18539
<i>Gabriele Corso, Jure Leskovec, Pietro Liò, Michal Pandy, Petar Velickovic, Zhitao Ying</i>	
Fitting Summary Statistics of Neural Data with a Differentiable Spiking Network Simulator	18552
<i>Guillaume Bellec, Johanni Brea, Wulfram Gerstner, Alireza Modirshanechi, Shuqi Wang</i>	
PerSim: Data-Efficient Offline Reinforcement Learning with Heterogeneous Agents Via Personalized Simulators	18564
<i>Anish Agarwal, Abdullah Alomar, Varkey Alumootil, Devavrat Shah, Dennis Shen, Zhi Xu, Cindy Yang</i>	
Online Sign Identification: Minimization of the Number of Errors in Thresholding Bandits	18577
<i>Odalric-Ambrym Maillard, Reda Ouhamma, Vianney Perchet</i>	
All Tokens Matter: Token Labeling for Training Better Vision Transformers	18590
<i>Jiashi Feng, Qibin Hou, Zi-Hang Jiang, Xiaojie Jin, Yujun Shi, Anran Wang, Li Yuan, Daquan Zhou</i>	
Partition and Code: Learning How to Compress Graphs.....	18603
<i>Giorgos Bouritsas, Michael Bronstein, Nikolaos Karalias, Andreas Loukas</i>	
Knowledge-Inspired 3D Scene Graph Prediction in Point Cloud.....	18620
<i>Aimin Hao, Hong Qin, Shoulong Zhang, Shuai Li</i>	
Online Variational Filtering and Parameter Learning	18633
<i>Andrew Campbell, Arnaud Doucet, Thomas Rainforth, Yuyang Shi</i>	

VOLUME 23

Heavy Ball Neural Ordinary Differential Equations	18646
<i>Andrea Bertozzi, Hangjie Ji, Tan Nguyen, Stanley Osher, Vai Suliafu, Bao Wang, Hedi Xia</i>	
Structure Learning in Polynomial Time: Greedy Algorithms, Bregman Information, and Exponential Families	18660
<i>Bryon Aragam, Ming Gao, Bohdan Kivva, Goutham Rajendran</i>	
On the Sample Complexity of Learning Under Geometric Stability	18673
<i>Alberto Bietti, Joan Bruna, Luca Venturi</i>	

SIMILAR: Submodular Information Measures Based Active Learning in Realistic Scenarios.....	18685
<i>Nathan Beck, Rishabh Iyer, Krishnateja Killamsetty, Suraj Kothawade</i>	
Monte Carlo Tree Search with Iteratively Refining State Abstractions	18698
<i>Zaheen Ahmad, Caleb Y Ho, J. Zico Kolter, Samuel Sokota</i>	
Flattening Sharpness for Dynamic Gradient Projection Memory Benefits Continual Learning	18710
<i>Guangyong Chen, Danruo Deng, Jianye Hao, Pheng-Ann Heng, Qiong Wang</i>	
Taxonomizing Local Versus Global Structure in Neural Network Loss Landscapes.....	18722
<i>Joseph Gonzalez, Liam Hodgkinson, Michael Mahoney, Kannan Ramchandran, Ryan Theisen, Yaoqing Yang, Joe Zou</i>	
Learning Models for Actionable Recourse.....	18734
<i>Osbert Bastani, Himabindu Lakkaraju, Alexis Ross</i>	
Efficient and Accurate Gradients for Neural SDEs	18747
<i>James Foster, Patrick Kidger, Xuechen (Chen) Li, Terry Lyons</i>	
EIGNN: Efficient Infinite-Depth Graph Neural Networks.....	18762
<i>Bryan Hooi, Kenji Kawaguchi, Juncheng Liu, Yiwei Wang, Xiaokui Xiao</i>	
Fractal Structure and Generalization Properties of Stochastic Optimization Algorithms	18774
<i>Alexander Camuto, George Deligiannidis, Murat Erdogdu, Mert Gurbuzbalaban, Umut Simsekli, Lingjiong Zhu</i>	
An Infinite-Feature Extension for Bayesian ReLU Nets that Fixes Their Asymptotic Overconfidence	18789
<i>Matthias Hein, Philipp Hennig, Agustinus Kristiadi</i>	
Bandit Phase Retrieval	18801
<i>Botao Hao, Tor Lattimore</i>	
Lower Bounds on Metropolized Sampling Methods for Well-Conditioned Distributions.....	18812
<i>Yin Tat Lee, Ruoqi Shen, Kevin Tian</i>	
Taming Communication and Sample Complexities in Decentralized Policy Evaluation for Cooperative Multi-Agent Reinforcement Learning.....	18825
<i>Jia Liu, Zhuqing Liu, Songtao Lu, Xin Zhang, Zhengyuan Zhu</i>	
Federated Graph Classification Over Non-IID Graphs	18839
<i>Jing Ma, Han Xie, Li Xiong, Carl Yang</i>	
SubTab: Subsetting Features of Tabular Data for Self-Supervised Representation Learning	18853
<i>Lindsay Edwards, Ehsan Hajiramezani, Talip Ucar</i>	
Convergence Rates of Stochastic Gradient Descent Under Infinite Noise Variance.....	18866
<i>Murat Erdogdu, Mert Gurbuzbalaban, Umut Simsekli, Hongjian Wang, Lingjiong Zhu</i>	
Conflict-Averse Gradient Descent for Multi-task Learning	18878
<i>Xiaojie Jin, Bo Liu, Qiang Liu, Xingchao Liu, Peter Stone</i>	
Amortized Synthesis of Constrained Configurations Using a Differentiable Surrogate	18891
<i>Ryan Adams, Szymon Rusinkiewicz, Xingyuan Sun, Tianju Xue</i>	
Efficient First-Order Contextual Bandits: Prediction, Allocation, and Triangular Discrimination	18907
<i>Dylan Foster, Akshay Krishnamurthy</i>	

Distributed Estimation with Multiple Samples Per User: Sharp Rates and Phase Transition	18920
<i>Jayadev Acharya, Clement Canonne, Yuhan Liu, Ziteng Sun, Himanshu Tyagi</i>	
Revisiting Deep Learning Models for Tabular Data	18932
<i>Artem Babenko, Yury Gorishniy, Valentin Khrulkov, Ivan Rubachev</i>	
Backdoor Attack with Imperceptible Input and Latent Modification.....	18944
<i>Khoa Doan, Yingjie Lao, Ping Li</i>	
SOPE: Spectrum of Off-Policy Estimators.....	18958
<i>Yash Chandak, Stephen Giguere, Scott Niekum, Philip Thomas, Christina Yuan</i>	
Label-Imbalanced and Group-Sensitive Classification Under Overparameterization.....	18970
<i>Ganesh Ramachandra Kini, Samet Oymak, Orestis Paraskevas, Christos Thrampoulidis</i>	
Neural Program Generation Modulo Static Analysis	18984
<i>Dipak Chaudhuri, Swarat Chaudhuri, Christopher Jermaine, Rohan Mukherjee, Thomas Reps, Yeming Wen</i>	
Unfolding Taylor's Approximations for Image Restoration	18997
<i>Xueyang Fu, Aiping Liu, Zeyu Xiao, Zhiwei Xiong, Gang Yang, Man Zhou</i>	
Metropolis-Hastings Data Augmentation for Graph Neural Networks	19010
<i>Jung-Woo Ha, Jisu Jeong, Hyunwoo Kim, Kyung-Min Kim, Sihyeon Kim, Seunghun Lee, Hyeonjin Park, Jinyoung Park</i>	
Strategic Behavior is Bliss: Iterative Voting Improves Social Welfare.....	19021
<i>Joshua Kavner, Lirong Xia</i>	
Agnostic Reinforcement Learning with Low-Rank MDPs and Rich Observations	19033
<i>Christoph Dann, Yishay Mansour, Mehryar Mohri, Ayush Sekhari, Karthik Sridharan</i>	
Functional Regularization for Reinforcement Learning Via Learned Fourier Features	19046
<i>Alexander Li, Deepak Pathak</i>	
Adaptive First-Order Methods Revisited: Convex Minimization Without Lipschitz Requirements.....	19056
<i>Kimon Antonakopoulos, Panayotis Mertikopoulos</i>	
Adapting to Function Difficulty and Growth Conditions in Private Optimization.....	19069
<i>Hilal Asi, John Duchi, Daniel Levy</i>	
Support Recovery of Sparse Signals from a Mixture of Linear Measurements	19082
<i>Venkata Gandikota, Arya Mazumdar, Soumyabrata Pal</i>	
Stochastic Gradient Descent-Ascent and Consensus Optimization for Smooth Games: Convergence Analysis Under Expected Co-coercivity	19095
<i>Hugo Berard, Gauthier Gidel, Simon Lacoste-Julien, Nicolas Loizou, Ioannis Mitliagkas</i>	
Tighter Expected Generalization Error Bounds Via Wasserstein Distance.....	19109
<i>German Bassi, Borja Rodríguez Galvez, Mikael Skoglund, Ragnar Thobaben</i>	
Unifying Width-Reduced Methods for Quasi-Self-Concordant Optimization	19122
<i>Deeksha Adil, Brian Bullins, Sushant Sachdeva</i>	
Bridging the Imitation Gap by Adaptive Insubordination	19134
<i>Unnat Jain, Aniruddha Kembhavi, Svetlana Lazebnik, Iou-Jen Liu, Jordi Salvador, Alex Schwing, Luca Weihs</i>	

Adversarial Robustness with Non-Uniform Perturbations	19147
<i>Sergul Aydore, Jeffrey Bickford, Ecenaz Erdemir, Luca Melis</i>	
Container: Context Aggregation Networks	19160
<i>Aniruddha Kembhavi, Hongsheng Li, Jiasen Lu, Roozbeh Mottaghi, Peng Gao</i>	
ConE: Cone Embeddings for Multi-Hop Reasoning Over Knowledge Graphs	19172
<i>Jiajun Chen, Shuiwang Ji, Jie Wang, Feng Wu, Zhanqiu Zhang</i>	
Federated Hyperparameter Tuning: Challenges, Baselines, and Connections to Weight-Sharing	19184
<i>Maria-Florina Balcan, Mikhail Khodak, Liam Li, Tian Li, Virginia Smith, Ameet Talwalkar, Renbo Tu</i>	
Training for the Future: A Simple Gradient Interpolation Loss to Generalize Along Time.....	19198
<i>Abir De, Anshul Nasery, Vihari Piratla, Sunita Sarawagi, Soumyadeep Thakur</i>	
Agent Modelling Under Partial Observability for Deep Reinforcement Learning	19210
<i>Stefano Albrecht, Filippos Christianos, Georgios Papoudakis</i>	
Leveraging Distribution Alignment Via Stein Path for Cross-Domain Cold-Start Recommendation	19223
<i>Chaochao Chen, Weiming Liu, Jiajie Su, Xiaolin Zheng</i>	
Conservative Offline Distributional Reinforcement Learning	19235
<i>Osbert Bastani, Dinesh Jayaraman, Yecheng Ma</i>	
Separation Results Between Fixed-Kernel and Feature-Learning Probability Metrics	19248
<i>Carles Domingo I Enrich, Youssef Mroueh</i>	
Risk Minimization from Adaptively Collected Data: Guarantees for Supervised and Policy Learning	19261
<i>Aurelien Bibaut, Antoine Chambaz, Maria Dimakopoulou, Nathan Kallus, Mark Van Der Laan</i>	
Bayesian Optimization with High-Dimensional Outputs	19274
<i>Eytan Bakshy, Maximilian Balandat, Wesley Maddox, Andrew Wilson</i>	
Finding Optimal Tangent Points for Reducing Distortions of Hard-Label Attacks	19288
<i>Li Chen, Xiangyu Guo, Chen Ma, Yisen Wang, Jun-Hai Yong</i>	
Scalable Diverse Model Selection for Accessible Transfer Learning	19301
<i>Daniel Bolya, Judy Hoffman, Rohit Mittapalli</i>	
Light Field Networks: Neural Scene Representations with Single-Evaluation Rendering	19313
<i>Fredo Durand, Bill Freeman, Semon Rezchikov, Vincent Sitzmann, Josh Tenenbaum</i>	
ViSER: Video-Specific Surface Embeddings for Articulated 3D Shape Reconstruction	19326
<i>Forrester Cole, Varun Jampani, Ce Liu, Deva Ramanan, Deqing Sun, Daniel Vlasic, Gengshan Yang</i>	
Understanding the Effect of Stochasticity in Policy Optimization	19339
<i>Bo Dai, Jincheng Mei, Dale Schuurmans, Csaba Szepesvari, Chenjun Xiao</i>	
Fine-Grained Zero-Shot Learning with DNA as Side Information	19352
<i>Zeynep Akata, Sarkhan Badirli, Mehmet M Dundar, George Mohler, Christine Picard</i>	
Optimal Underdamped Langevin MCMC Method.....	19363
<i>Zhengmian Hu, Heng Huang, Feihu Huang</i>	

Scheduling Jobs with Stochastic Holding Costs.....	19375
<i>Dabeen Lee, Milan Vojnovic</i>	
REMIPS: Physically Consistent 3D Reconstruction of Multiple Interacting People Under Weak Supervision.....	19385
<i>Eduard Bazavan, Mihai Fieraru, Vlad Olaru, Cristian Sminchisescu, Teodor Szente, Mihai Zanfir</i>	
Differentiable Annealed Importance Sampling and the Perils of Gradient Noise.....	19398
<i>Chelsea Finn, Roger Grosse, Kyle Hsu, Jianing Li, Guodong Zhang</i>	
PSD Representations for Effective Probability Models.....	19411
<i>Carlo Ciliberto, Alessandro Rudi</i>	
Exploiting a Zoo of Checkpoints for Unseen Tasks.....	19423
<i>Kenneth Church, Jiaji Huang, Qiang Qiu</i>	
Towards Open-World Feature Extrapolation: An Inductive Graph Learning Approach.....	19435
<i>Qitian Wu, Junchi Yan, Chenxiao Yang</i>	
Adversarial Teacher-Student Representation Learning for Domain Generalization.....	19448
<i>Yuan-Chia Cheng, Zu-Yun Shiao, Yu-Chiang Frank Wang, Fu-En Yang</i>	
Stochastic Bandits with Groups of Similar Arms.....	19461
<i>Odalric-Ambrym Maillard, Fabien Pesquerel, Hassan Saber</i>	
Tracking Without Re-Recognition in Humans and Machines.....	19473
<i>Lakshmi Narasimhan Govindarajan, Junkyung Kim, Drew Linsley, Girik Malik, Ennio Mingolla, Thomas Serre</i>	
Rethinking Conditional GAN Training: An Approach Using Geometrically Structured Latent Manifolds.....	19487
<i>Nick Barnes, Moshir Farazi, Stephen Gould, Salman H Khan, Sameera Ramasinghe</i>	

VOLUME 24

How to Transfer Algorithmic Reasoning Knowledge to Learn New Algorithms?.....	19500
<i>Andreea-Ioana Deac, Jian Tang, Petar Velickovic, Louis-Pascal Xhonneux</i>	
Fast Axiomatic Attribution for Neural Networks.....	19513
<i>Robin Hesse, Stefan Roth, Simone Schaub-Meyer</i>	
OSOA: One-Shot Online Adaptation of Deep Generative Models for Lossless Compression.....	19525
<i>Fabio Maria Carlucci, Zhenguo Li, Chen Zhang, Shifeng Zhang</i>	
Compressive Visual Representations.....	19538
<i>Anurag Arnab, John Canny, Ian Fischer, Sergio Guadarrama, Kuang-Huei Lee</i>	
Multi-Armed Bandits with Bounded Arm-Memory: Near-Optimal Guarantees for Best-Arm Identification and Regret Minimization.....	19553
<i>Arindam Khan, Arnab Maiti, Vishakha Patil</i>	
Grounding Inductive Biases in Natural Images: Invariance Stems from Variations in Data.....	19566
<i>Diane Bouchacourt, Mark Ibrahim, Ari Morcos</i>	

Directed Graph Contrastive Learning.....	19580
<i>Yongxing Dai, Henghui Ding, Xinke Li, Yuxuan Liang, Zekun Tong, Changhu Wang</i>	
Space-Time Mixing Attention for Video Transformer.....	19594
<i>Adrian Bulat, Brais Martinez, Juan Manuel Perez Rua, Swathikiran Sudhakaran, Georgios Tzimiropoulos</i>	
Particle Dual Averaging: Optimization of Mean Field Neural Network with Global Convergence Rate Analysis.....	19608
<i>Atsushi Nitanda, Taiji Suzuki, Denny Wu</i>	
Learning Tree Interpretation from Object Representation for Deep Reinforcement Learning.....	19622
<i>Guiliang Liu, Pascal Poupart, Oliver Schulte, Xiangyu Sun</i>	
Only Train Once: a One-Shot Neural Network Training and Pruning Framework.....	19637
<i>Tianyi Chen, Tianyu Ding, Biyi Fang, Bo Ji, Luming Liang, Yixin Shi, Xiao Tu, Guanyi Wang, Sheng Yi, Zhihui Zhu</i>	
Referring Transformer: a One-Step Approach to Multi-task Visual Grounding.....	19652
<i>Muchen Li, Leonid Sigal</i>	
Decoupling the Depth and Scope of Graph Neural Networks.....	19665
<i>Ren Chen, Long Jin, Rajgopal Kannan, Andrey Malevich, Viktor Prasanna, Ajitesh Srivastava, Yinglong Xia, Hanqing Zeng, Muhan Zhang</i>	
Fast and Memory Efficient Differentially Private-SGD Via JL Projections.....	19680
<i>Zhiqi Bu, Sivakanth Gopi, Janardhan Kulkarni, Yin Tat Lee, Hanwen Shen, Uthaipon Tantipongpipat</i>	
Formalizing Generalization and Adversarial Robustness of Neural Networks to Weight Perturbations.....	19692
<i>Pin-Yu Chen, Chia-Yi Hsu, Yu-Lin Tsai, Chia-Mu Yu</i>	
Pipeline Combinators for Gradual AutoML.....	19705
<i>Guillaume Baudart, Martin Hirzel, Kiran Kate, Parikshit Ram, Avi Shinnar, Jason Tsay</i>	
Boost Neural Networks by Checkpoints.....	19719
<i>Qiao Liu, Hairong Lv, Jinxiang Ou, Feng Wang, Guoyizhe Wei, Xian Wei</i>	
Model Selection for Bayesian Autoencoders.....	19730
<i>Edwin Bonilla, Maurizio Filippone, Pietro Michiardi, Dimitrios Milios, Simone Rossi, Ba-Hien Tran</i>	
Three Operator Splitting with Subgradients, Stochastic Gradients, and Adaptive Learning Rates.....	19743
<i>Alex Gu, Suvrit Sra, Alp Yurtsever</i>	
Knowledge-Adaptation Priors.....	19757
<i>Mohammad Emtiyaz Khan, Siddharth Swaroop</i>	
Provably Efficient Multi-Task Reinforcement Learning with Model Transfer.....	19771
<i>Zhi Wang, Chicheng Zhang</i>	
Predicting Molecular Conformation Via Dynamic Graph Score Matching.....	19784
<i>Shitong Luo, Chence Shi, Jian Tang, Minkai Xu</i>	

When in Doubt: Neural Non-Parametric Uncertainty Quantification for Epidemic Forecasting	19796
<i>Harshavardhan Kamarthi, Lingkai Kong, B. Aditya Prakash, Alexander Rodriguez, Chao Zhang</i>	
Bounds All Around: Training Energy-Based Models with Bidirectional Bounds.....	19808
<i>Jes Frellsen, Zhiyong Gao, Cong Geng, Søren Hauberg, Jia Wang</i>	
CogView: Mastering Text-To-Image Generation Via Transformers	19822
<i>Ming Ding, Wenyi Hong, Junyang Lin, Zhou Shao, Jie Tang, Hongxia Yang, Zhuoyi Yang, Da Yin, Wendi Zheng, Chang Zhou, Xu Zou</i>	
Time-Independent Generalization Bounds for SGLD in Non-convex Settings.....	19836
<i>Tyler Farghly, Patrick Rebeschini</i>	
Nonuniform Negative Sampling and Log Odds Correction with Rare Events Data	19847
<i>Chong Wang, Haiying Wang, Aonan Zhang</i>	
Algorithmic Stability and Generalization of an Unsupervised Feature Selection Algorithm.....	19860
<i>Qiang Cheng, Xinxing Wu</i>	
On Learning Sparse Vectors from Mixture of Responses	19876
<i>Nikita Polyanskii</i>	
Convergence and Alignment of Gradient Descent with Random Backpropagation Weights	19888
<i>John Lafferty, Ganlin Song, Ruitu Xu</i>	
Adder Attention for Vision Transformer.....	19899
<i>Hanting Chen, Lin Li, Han Shu, Jiahao Wang, Yunhe Wang, Yujiu Yang</i>	
Reverse Engineering Learned Optimizers Reveals Known and Novel Mechanisms	19910
<i>Niru Maheswaranathan, Luke Metz, Jascha Sohl-Dickstein, Ruoxi Sun, David Sussillo</i>	
Matching a Desired Causal State Via Shift Interventions	19923
<i>Chandler Squires, Caroline Uhler, Jiaqi Zhang</i>	
Unsupervised Noise Adaptive Speech Enhancement by Discriminator-Constrained Optimal Transport	19935
<i>Hsin-Yi Lin, Xugang Lu, Yu Tsao, Huan-Hsin Tseng</i>	
Optimality of Variational Inference for Stochasticblock Model with Missing Links.....	19947
<i>Solenne Gaucher, Olga Klopp</i>	
Policy Learning Using Weak Supervision.....	19960
<i>Hongyi Guo, Yang Liu, Jingkang Wang, Zhaowei Zhu</i>	
Chasing Sparsity in Vision Transformers: an End-To-End Exploration	19974
<i>Tianlong Chen, Yu Cheng, Zhe Gan, Zhangyang Wang, Lu Yuan, Lei Zhang</i>	
Graphical Models in Heavy-Tailed Markets	19989
<i>Daniel Palomar, Jiayi Ying, Jose Vinicius De Miranda Cardoso</i>	
A Shading-Guided Generative Implicit Model for Shape-Accurate 3D-Aware Image Synthesis.....	20002
<i>Bo Dai, Chen Change Loy, Xingang Pan, Christian Theobalt, Xudong Xu</i>	
XCiT: Cross-Covariance Image Transformers.....	20014
<i>Alaaeldin Ali, Piotr Bojanowski, Mathilde Caron, Matthijs Douze, Herve Jegou, Armand Joulin, Ivan Laptev, Natalia Neverova, Gabriel Synnaeve, Hugo Touvron, Jakob Verbeek</i>	

Row-Clustering of a Point Process-valued Matrix	20028
<i>Yongtao Guan, Huiyan Sang, Ganggang Xu, Lihao Yin</i>	
Fine-Grained Neural Network Explanation by Identifying Input Features with Predictive Information	20040
<i>Azade Farshad, Ashkan Khakzar, Seong Tae Kim, Yawei Li, Nassir Navab, Yang Zhang</i>	
Fast Minimum-Norm Adversarial Attacks Through Adaptive Norm Constraints	20052
<i>Battista Biggio, Wieland Brendel, Maura Pintor, Fabio Roli</i>	
Uncertainty Quantification and Deep Ensembles	20063
<i>Rahul Rahaman, Alexandre Thiery</i>	
Directed Probabilistic Watershed	20076
<i>Sebastian Damrich, Enrique Fita Sanmartin, Fred Hamprecht</i>	
Laplace Redux - Effortless Bayesian Deep Learning	20089
<i>Matthias Bauer, Erik Daxberger, Runa Eschenhagen, Philipp Hennig, Alexander Immer, Agustinus Kristiadi</i>	
Hessian Eigenspectra of More Realistic Nonlinear Models	20104
<i>Zhenyu Liao, Michael Mahoney</i>	
Explicable Reward Design for Reinforcement Learning Agents	20118
<i>Rati Devidze, Parameswaran Kamalaruban, Goran Radanovic, Adish Singla</i>	
A Minimalist Approach to Offline Reinforcement Learning	20132
<i>Scott Fujimoto, Shixiang (Shane) Gu</i>	
SIMONE: View-Invariant, Temporally-Abstracted Object Representations Via Unsupervised Video Decomposition	20146
<i>Matt Botvinick, Chris Burgess, Antonia Creswell, Goker Erdogan, Rishabh Kabra, Alexander Lerchner, Loic Matthey, Daniel Zoran</i>	
Simple Stochastic and Online Gradient Descent Algorithms for Pairwise Learning	20160
<i>Yunwen Lei, Puyu Wang, Zhenhuan Yang, Tianbao Yang, Yiming Ying</i>	
User-Level Differentially Private Learning Via Correlated Sampling	20172
<i>Badi Ghazi, Ravi Kumar, Pasin Manurangsi</i>	
Asynchronous Decentralized Online Learning	20185
<i>Jinjie Gu, Jiyan Jiang, Wenpeng Zhang, Wenwu Zhu</i>	
Multi-Step Budgeted Bayesian Optimization with Unknown Evaluation Costs	20197
<i>Raul Astudillo, Eytan Bakshy, Maximilian Balandat, Peter Frazier, Daniel Jiang</i>	
Model-Based Domain Generalization	20210
<i>Hamed Hassani, George J. Pappas, Alexander Robey</i>	
Alpha-IOU: a Family of Power Intersection Over Union Losses for Bounding Box Regression	20230
<i>James Bailey, Ying Chi, Sarah Erfani, Jiabo He, Xian-Sheng Hua, Xingjun Ma</i>	
Practical Large-Scale Linear Programming Using Primal-Dual Hybrid Gradient	20243
<i>David Applegate, Mateo Diaz, Oliver Hinder, Haihao Lu, Miles Lubin, Brendan O'Donoghue, Warren Schudy</i>	

On the Provable Generalization of Recurrent Neural Networks.....	20258
<i>Xing Cao, Bo Hu, Bo Shen, Lifu Wang</i>	
Differentiable Spline Approximations.....	20270
<i>Aditya Balu, Minsu Cho, Anjana Deva Prasad, Baskar Ganapathysubramanian, Chinmay Hegde, Ameya Joshi, Biswajit Khara, Adarsh Krishnamurthy, Soumik Sarkar</i>	
Rate-Optimal Subspace Estimation on Random Graphs	20283
<i>Ping Li, Cun-Hui Zhang, Fan Zhou, Zhixin Zhou</i>	
Estimating the Unique Information of Continuous Variables	20295
<i>Dar Gilboa, Abdullah Makkeh, Luca Mazzucato, Amin Nejatbakhsh, Ari Pakman, Elad Schneidman, Michael Wibral</i>	
Reliable Causal Discovery with Improved Exact Search and Weaker Assumptions	20308
<i>Ignavier Ng, Kun Zhang, Jiji Zhang, Yujia Zheng</i>	
Node Dependent Local Smoothing for Scalable Graph Learning	20321
<i>Bin Cui, Yang Li, Wen Ouyang, Zeang Sheng, Yangyu Tao, Zhi Yang, Mingyu Yang, Wentao Zhang</i>	
Parallel and Efficient Hierarchical k-Median Clustering.....	20333
<i>Vincent Cohen-Addad, Silvio Lattanzi, Ashkan Norouzi-Fard, Christian Sohler, Ola Svensson</i>	
Human-Adversarial Visual Question Answering	20346
<i>Wojciech Galuba, Vedanuj Goswami, Douwe Kiela, Jose Magana, Devi Parikh, Sasha Sheng, Amanpreet Singh, Tristan Thrush</i>	

VOLUME 25

Across-Animal Odor Decoding by Probabilistic Manifold Alignment	20360
<i>Pedro Herrero-Vidal, Dmitry Rinberg, Cristina Savin</i>	
Excess Capacity and Backdoor Poisoning	20373
<i>Avrim Blum, Naren Manoj</i>	
A Convergence Analysis of Gradient Descent on Graph Neural Networks	20385
<i>Pranjal Awasthi, Abhimanyu Das, Sreenivas Gollapudi</i>	
Differentiable Rendering with Perturbed Optimizers	20398
<i>Justin Carpentier, Ivan Laptev, Quentin Le Lidec, Cordelia Schmid</i>	
BCORLE(λ): an Offline Reinforcement Learning and Evaluation Framework for Coupons Allocation in E-Commerce Market	20410
<i>Dou An, Xueying Li, Bo Tang, Hongyin Tang, Chenyang Xi, Feiyu Xiong, Qingyu Yang, Yang Zhang</i>	
Nested Variational Inference	20423
<i>Babak Esmaeili, Hao Wu, Heiko Zimmermann, Jan-Willem Van De Meent</i>	
Exponential Bellman Equation and Improved Regret Bounds for Risk-Sensitive Reinforcement Learning	20436
<i>Yudong Chen, Yingjie Fei, Zhaoran Wang, Zhuoran Yang</i>	
On Sensitivity of Meta-Learning to Support Data.....	20447
<i>Mayank Agarwal, Yuekai Sun, Mikhail Yurochkin</i>	

On Large-Cohort Training for Federated Learning	20461
<i>Zachary Charles, Zachary Garrett, Zhouyuan Huo, Sergei Shmulyian, Virginia Smith</i>	
Generic Neural Architecture Search Via Regression	20476
<i>Deming Chen, Cong Hao, Pan Li, Yuhong Li, Jinjun Xiong</i>	
The Best of Both Worlds: Stochastic and Adversarial Episodic MDPs with Unknown Transition	20491
<i>Longbo Huang, Tiancheng Jin, Haipeng Luo</i>	
Private Learning Implies Quantum Stability	20503
<i>Srinivasan Arunachalam, Yihui Quek, John A Smolin</i>	
Interesting Object, Curious Agent: Learning Task-Agnostic Exploration.....	20516
<i>Victoria Dean, Abhinav Gupta, Simone Parisi, Deepak Pathak</i>	
SimiGrad: Fine-Grained Adaptive Batching for Large Scale Training Using Gradient Similarity Measurement	20531
<i>Yuxiong He, Heyang Qin, Samyam Rajbhandari, Olatunji Ruwase, Feng Yan, Lei Yang</i>	
Variational Inference for Continuous-Time Switching Dynamical Systems.....	20545
<i>Bastian Alt, Heinz Koeppl, Lukas Kohs</i>	
Implicit Regularization in Matrix Sensing Via Mirror Descent	20558
<i>Patrick Rebeschini, Fan Wu</i>	
STORM+: Fully Adaptive SGD with Recursive Momentum for Nonconvex Optimization.....	20571
<i>Volkan Cevher, Ali Kavis, Kfir Levy</i>	
Skipping the Frame-Level: Event-Based Piano Transcription with Neural Semi-CRFs.....	20583
<i>Frank Cwitkowitz, Zhiyao Duan, Yujia Yan</i>	
Deep Learning on a Data Diet: Finding Important Examples Early in Training	20596
<i>Gintare Karolina Dziugaite, Surya Ganguli, Mansheej Paul</i>	
BNS: Building Network Structures Dynamically for Continual Learning	20608
<i>Wenpeng Hu, Bing Liu, Han Peng, Qi Qin, Dongyan Zhao</i>	
Auditing Black-Box Prediction Models for Data Minimization Compliance	20621
<i>Mark Crovella, Krishna Gummadi, Bashir Rastegarpanah</i>	
Dueling Bandits with Team Comparisons.....	20633
<i>Lee Cohen, Yishay Mansour, Ulrike Schmidt-Kraepelin</i>	
Meta Internal Learning	20645
<i>Raphael Bensadoun, Tomer Galanti, Shir Gur, Lior Wolf</i>	
Uniform Convergence of Interpolators: Gaussian Width, Norm Bounds and Benign Overfitting	20657
<i>Frederic Koehler, Nathan Srebro, [deadname] Sutherland, Lijia Zhou</i>	
Adaptive Wavelet Distillation from Neural Networks Through Interpretations	20669
<i>Wooseok Ha, Francois Lanusse, Chandan Singh, Srigokul Upadhyayula, Bin Yu</i>	
Generative Occupancy Fields for 3D Surface-Aware Image Synthesis	20683
<i>Bo Dai, Dahua Lin, Xingang Pan, Xudong Xu</i>	
Relaxed Marginal Consistency for Differentially Private Query Answering	20696
<i>Ryan McKenna, Jerome Miklau, Siddhant Pradhan, Daniel Sheldon</i>	

Local Policy Search with Bayesian Optimization	20708
<i>Sarah Muller, Sebastian Trimpe, Alexander Von Rohr</i>	
DominoSearch: Find Layer-Wise Fine-grained N:M Sparse Schemes from Dense Neural Networks.....	20721
<i>Henk Corporaal, Hongsheng Li, Andrew Nelson, Sander Stuijk, Wei Sun, Rob Wijnhoven, Aojun Zhou</i>	
Techniques for Symbol Grounding with SATNet	20733
<i>David Rolnick, Xujie Si, Sever Topan</i>	
Object DGCNN: 3D Object Detection Using Dynamic Graphs.....	20745
<i>Justin Solomon, Yue Wang</i>	
Safe Policy Optimization with Local Generalized Linear Function Approximations	20759
<i>Yanan Sui, Akifumi Wachi, Yunyue Wei</i>	
Symplectic Adjoint Method for Exact Gradient of Neural ODE with Minimal Memory	20772
<i>Takashi Matsubara, Yuto Miyatake, Takaharu Yaguchi</i>	
Exponential Separation Between Two Learning Models and Adversarial Robustness	20785
<i>Grzegorz Gluch, Ruediger Urbanke</i>	
The Balancing Principle for Parameter Choice in Distance-Regularized Domain Adaptation	20798
<i>Marius-Constantin Dinu, Hamid Eghbal-Zadeh, Bernhard A. Moser, Bernhard Nessler, Hoan Duc Nguyen, Sergei Pereverzyev, Natalia Shepeleva, Werner Zellinger</i>	
Gaussian Kernel Mixture Network for Single Image Defocus Deblurring.....	20812
<i>Hui Ji, Yuhui Quan, Zicong Wu</i>	
Cockpit: a Practical Debugging Tool for the Training of Deep Neural Networks	20825
<i>Felix Dangel, Philipp Hennig, Frank Schneider</i>	
MEST: Accurate and Fast Memory-Economic Sparse Training Framework on the Edge	20838
<i>Yifan Gong, Chaoyang He, Qing Jin, Zhenglun Kong, Zhengang Li, Xue Lin, Ning Liu, Sijia Liu, Xiaolong Ma, Wei Niu, Minghai Qin, Bin Ren, Siyue Wang, Yanzhi Wang, Geng Yuan, Zheng Zhan</i>	
Precise Characterization of the Prior Predictive Distribution of Deep ReLU Networks	20851
<i>Gregor Bachmann, Thomas Hofmann, Lorenzo Noci, Sebastian Nowozin, Kevin Roth</i>	
RED : Looking for Redundancies for Data-Free Structured Compression of Deep Neural Networks.....	20863
<i>Kevin Bailly, Matthieu Cord, Arnaud Dapogny, Edouard Yvinec</i>	
TestRank: Bringing Order into Unlabeled Test Instances for Deep Learning Tasks.....	20874
<i>Qiuxia Lai, Min Li, Yu Li, Yannan Liu, Qiang Xu</i>	
Large Scale Learning on Non-Homophilous Graphs: New Benchmarks and Strong Simple Methods.....	20887
<i>Omkar Bhalerao, Vaishnavi Gupta, Felix Hohne, Sijia Linda Huang, Xiuyu Li, Derek Lim, Ser Nam Lim</i>	
Reinforcement Learning Based Disease Progression Model for Alzheimer’s Disease	20903
<i>Yurui Cao, Anirudh Choudhary, Ravishankar Iyer, David Jones, Krishnakant Saboo, Gregory Worrell</i>	
Catch-A-Waveform: Learning to Generate Audio from a Single Short Example	20916
<i>Gal Greshler, Tomer Michaeli, Tamar Shaham</i>	

Explanation-Based Data Augmentation for Image Classification	20929
<i>Wynne Hsu, Mong Li Lee, Sandareka Wickramanayake</i>	
Data-Efficient GAN Training Beyond (Just) Augmentations: a Lottery Ticket Perspective.....	20941
<i>Tianlong Chen, Yu Cheng, Zhe Gan, Jingjing Liu, Zhangyang Wang</i>	
When Are Solutions Connected in Deep Networks?	20956
<i>Pierre Brechet, Marco Mondelli, Quynh Nguyen</i>	
TOHAN: a One-Step Approach Towards Few-Shot Hypothesis Adaptation.....	20970
<i>William Cheung, Haoang Chi, Bo Han, James Kwok, Long Lan, Feng Liu, Tongliang Liu, Wenjing Yang</i>	
Learning Graph Cellular Automata	20983
<i>Cesare Alippi, Daniele Grattarola, Lorenzo Livi</i>	
Efficient Online Estimation of Causal Effects by Deciding What to Observe	20995
<i>David Childers, Shantanu Gupta, Zachary Lipton</i>	
Perturbation Theory for the Information Bottleneck	21008
<i>Vudtiwat Ngampruetikorn, David Schwab</i>	
Deconvolutional Networks on Graph Data.....	21019
<i>Hong Cheng, Jia Li, Jiajin Li, Yueting Li, Yang Liu, Jianwei Yu</i>	
Variational Multi-Task Learning with Gumbel-Softmax Priors.....	21031
<i>Ling Shao, Jiayi Shen, Marcel Worring, Xiantong Zhen</i>	
Accelerating Quadratic Optimization with Reinforcement Learning	21043
<i>Goran Banjac, Francesco Borrelli, Ken Goldberg, Joseph Gonzalez, Jeffrey Ichnowski, Paras Jain, Michael Luo, Bartolomeo Stellato, Ion Stoica</i>	
Deep Residual Learning in Spiking Neural Networks.....	21056
<i>Yanqi Chen, Wei Fang, Tiejun Huang, Timothee Masquelier, Yonghong Tian, Zhaofei Yu</i>	
Duplex Sequence-To-Sequence Learning for Reversible Machine Translation.....	21070
<i>Jiajun Chen, Shujian Huang, Lei Li, Jingjing Xu, Zaixiang Zheng, Hao Zhou</i>	
Improved Coresets and Sublinear Algorithms for Power Means in Euclidean Spaces	21085
<i>Vincent Cohen-Addad, David Saulpic, Chris Schwiegelshohn</i>	
Accelerated Sparse Neural Training: a Provable and Efficient Method to Find N:M Transposable Masks	21099
<i>Ron Banner, Brian Chmiel, Itay Hubara, Moshe Isard, Joseph Naor, Daniel Soudry</i>	
Learning and Generalization in RNNs	21112
<i>Navin Goyal, Abhishek Panigrahi</i>	
Improving Visual Quality of Image Synthesis by a Token-Based Generator with Transformers	21125
<i>Hongyang Chao, Jianlong Fu, Jianbo Wang, Huan Yang, Yanhong Zeng</i>	
The Effect of the Intrinsic Dimension on the Generalization of Quadratic Classifiers	21138
<i>Volkan Cevher, Leello Tadesse Dadi, Fabian Latorre, Paul Rolland</i>	
DeepReduce: a Sparse-Tensor Communication Framework for Federated Deep Learning	21150
<i>Aritra Dutta, Panos Kalnis, Kelly Kostopoulou, Xin Li, Alexandros Ntoulas, Hang Xu</i>	

Provably Efficient Causal Reinforcement Learning with Confounded Observational Data.....	21164
<i>Zhaoran Wang, Lingxiao Wang, Zhuoran Yang</i>	
Predicting Deep Neural Network Generalization with Perturbation Response Curves	21176
<i>Pin-Yu Chen, Payel Das, Brian Quanz, Yair Schiff</i>	
Exploiting Domain-Specific Features to Enhance Domain Generalization.....	21189
<i>Manh-Ha Bui, Dinh Phung, Anh Tran, Toan Tran</i>	
Optimal Order Simple Regret for Gaussian Process Bandits	21202
<i>Alberto Bernacchia, Nacime Bouziani, Sepehr Jalali, Da-Shan Shiu, Sattar Vakili</i>	

VOLUME 26

Generalization Guarantee of SGD for Pairwise Learning	21216
<i>Yunwen Lei, Mingrui Liu, Yiming Ying</i>	
Supercharging Imbalanced Data Learning with Energy-Based Contrastive Representation Transfer	21229
<i>Lawrence Carin, Junya Chen, Benjamin Goldstein, Ricardo Henao, Chenyang Tao, Zidi Xiu</i>	
Heavy Ball Momentum for Conditional Gradient	21244
<i>Georgios Giannakis, Bingcong Li, Alireza Sadeghi</i>	
PARP: Prune, Adjust and Re-Prune for Self-Supervised Speech Recognition	21256
<i>Shiyu Chang, Yung-Sung Chuang, David Cox, Jim Glass, Sameer Khurana, Cheng-I Jeff Lai, Yi-Lun Liao, Alexander Liu, Kaizhi Qian, Yang Zhang</i>	
Robust Learning of Optimal Auctions.....	21273
<i>Wenshuo Guo, Michael Jordan, Emmanouil Zampetakis</i>	
Disrupting Deep Uncertainty Estimation Without Harming Accuracy	21285
<i>Ran El-Yaniv, Ido Galil</i>	
SOFT: Softmax-Free Transformer with Linear Complexity	21297
<i>Weiguo Gao, Jiachen Lu, Chunjing Xu, Tao Xiang, Hang Xu, Jinghan Yao, Junge Zhang, Li Zhang, Xiatian Zhu</i>	
Task-Adaptive Neural Network Search with Meta-Contrastive Learning.....	21310
<i>Jinheon Baek, Sung Ju Hwang, Eunyoung Hyung, Wonyong Jeong, Hayeon Lee, Geon Park</i>	
Neural Flows: Efficient Alternative to Neural ODEs	21325
<i>Marin Bilos, Stephan Gunnemann, Tim Januschowski, Syama Sundar Rangapuram, Johanna Sommer</i>	
Multi-Objective Meta Learning.....	21338
<i>Pengxin Guo, Baijiong Lin, Qiao Xiao, Feiyang Ye, Zhixiong Yue, Yu Zhang</i>	
A Self Consistent Theory of Gaussian Processes Captures Feature Learning Effects in Finite CNNs	21352
<i>Gadi Naveh, Zohar Ringel</i>	
Mini-Batch Consistent Slot Set Encoder for Scalable Set Encoding.....	21365
<i>Andreis Bruno, Sung Ju Hwang, Juho Lee, Jeffrey Willette</i>	
Efficient and Local Parallel Random Walks	21375
<i>Michael Kapralov, Silvio Lattanzi, Navid Nouri, Jakab Tardos</i>	

Amortized Variational Inference for Simple Hierarchical Models.....	21388
<i>Abhinav Agrawal, Justin Domke</i>	
Online Matching in Sparse Random Graphs: Non-Asymptotic Performances of Greedy Algorithm	21400
<i>Nathan Noiry, Vianney Perchet, Flore Sentenac</i>	
End-To-end Reconstruction Meets Data-driven Regularization for Inverse Problems	21413
<i>Marcello Carioni, Subhadip Mukherjee, Carola-Bibiane Schonlieb, Ozan Oktm</i>	
An Online Passive-Aggressive Algorithm for Difference-of-squares Classification	21426
<i>Lawrence Saul</i>	
Finite-Sample Analysis of Off-Policy TD-Learning Via Generalized Bellman Operators	21440
<i>Zaiwei Chen, Siva Theja Maguluri, Sanjay Shakkottai, Karthikeyan Shanmugam</i>	
A Bi-Level Framework for Learning to Solve Combinatorial Optimization on Graphs	21453
<i>Zhigang Hua, Gan Liu, Feng Qi, Runzhong Wang, Junchi Yan, Shuang Yang, Xiaokang Yang, Jiayi Zhang, Jun Zhou</i>	
Improved Learning Rates of a Functional Lasso-Type SVM with Sparse Multi-Kernel Representation.....	21467
<i>Yong Liu, Jiankun Liu, Junhui Wang, Shaogao Lv</i>	
When Does Contrastive Learning Preserve Adversarial Robustness from Pretraining to Finetuning?	21480
<i>Pin-Yu Chen, Lijie Fan, Chuang Gan, Sijia Liu, Gaoyuan Zhang</i>	
Learning Transferable Features for Point Cloud Detection Via 3D Contrastive Co-Training	21493
<i>Chao Ma, Chunwei Wang, Yunbo Wang, Hang Xu, Zhen Yang, Chaoqiang Ye, Zeng Yihan</i>	
SILG: the Multi-Domain Symbolic Interactive Language Grounding Benchmark.....	21505
<i>Austin W. Hanjie, Karthik Narasimhan, Sida Wang, Luke Zettlemoyer, Victor Zhong</i>	
A Surrogate Objective Framework for Prediction+Programming with Soft Constraints	21520
<i>Liting Chen, Qingwei Lin, Chuan Luo, Jie Yan, Kai Yan, Dongmei Zhang</i>	
Learning to Predict Trustworthiness with Steep Slope Loss	21533
<i>Mohan Kankanhalli, Yan Luo, Yongkang Wong, Qi Zhao</i>	
On the Periodic Behavior of Neural Network Training with Batch Normalization and Weight Decay.....	21545
<i>Nadezhda Chirkova, Maxim Kodryan, Ekaterina Lobacheva, Andrey Malinin, Dmitry Vetrov</i>	
NeRV: Neural Representations for Videos.....	21557
<i>Hao Chen, Bo He, Ser Nam Lim, Yixuan Ren, Abhinav Shrivastava, Hanyu Wang</i>	
Surrogate Regret Bounds for Polyhedral Losses	21569
<i>Rafael Frongillo, Bo Waggoner</i>	
Last Iterate Convergence of SGD for Least-Squares in the Interpolation Regime.....	21581
<i>Nicolas Flammarion, Loucas Pillaud-Vivien, Aditya Vardhan Varre</i>	
Generative Vs. Discriminative: Rethinking the Meta-Continual Learning	21592
<i>Mohammadamin Banayeeanzade, Hosein Hasani, Rasoul Mirzaiezadeh, Mahdieh Soleymani</i>	
Model, Sample, and Epoch-Wise Descents: Exact Solution of Gradient Flow in the Random Feature Model.....	21605
<i>Antoine Bodin, Nicolas Macris</i>	

Rethinking Graph Transformers with Spectral Attention	21618
<i>Dominique Beaini, Will Hamilton, Devin Kreuzer, Vincent Letourneau, Prudencio Tossou</i>	
Perceptual Score: What Data Modalities Does Your Model Perceive?	21630
<i>Itai Gat, Idan Schwartz, Alex Schwing</i>	
PiRank: Scalable Learning to Rank Via Differentiable Sorting	21644
<i>Bruno Charron, Stefano Ermon, Aditya Grover, Robin Swezey</i>	
Deceive D: Adaptive Pseudo Augmentation for GAN Training with Limited Data	21655
<i>Bo Dai, Liming Jiang, Chen Change Loy, Wayne Wu</i>	
CoFrNets: Interpretable Neural Architecture Inspired by Continued Fractions	21668
<i>Amit Dhurandhar, Tejaswini Pedapati, Isha Puri, Karthikeyan Shanmugam, Kush Varshney, Dennis Wei</i>	
Iterative Teaching by Label Synthesis.....	21681
<i>Weiyang Liu, Zhen Liu, Liam Paull, Bernhard Scholkopf, Hanchen Wang, Adrian Weller</i>	
On Density Estimation with Diffusion Models	21696
<i>Jonathan Ho, Diederik Kingma, Ben Poole, Tim Salimans</i>	
FastCorrect: Fast Error Correction with Edit Alignment for Automatic Speech Recognition	21708
<i>Yichong Leng, Xiangyang Li, Edward Lin, Linquan Liu, Tie-Yan Liu, Renqian Luo, Tao Qin, Xu Tan, Jin Xu, Linchen Zhu</i>	
Integrated Latent Heterogeneity and Invariance Learning in Kernel Space	21720
<i>Peng Cui, Zheyuan Hu, Bo Li, Jiashuo Liu, Zheyang Shen</i>	
Hierarchical Reinforcement Learning with Timed Subgoals	21732
<i>Dieter Buchler, Nico Gurtler, Georg Martius</i>	
Fair Scheduling for Time-Dependent Resources.....	21744
<i>Bo Li, Minming Li, Ruilong Zhang</i>	
SNIPS: Solving Noisy Inverse Problems Stochastically	21757
<i>Michael Elad, Bahjat Kawar, Gregory Vaksman</i>	
Stateful ODE-Nets Using Basis Function Expansions	21770
<i>N. Benjamin Erichson, Liam Hodgkinson, Michael Mahoney, Alejandro Queiruga</i>	
Beyond the Signs: Nonparametric Tensor Completion Via Sign Series	21782
<i>Chanwoo Lee, Miaoyan Wang</i>	
Functional Variational Inference Based on Stochastic Process Generators.....	21795
<i>Jose Miguel Hernandez-Lobato, Chao Ma</i>	
TTT++: When Does Self-Supervised Test-Time Training Fail Or Thrive?	21808
<i>Alexandre Alahi, Baptiste Bellot-Gurlet, Parth Kothari, Yuejiang Liu, Taylor Mordan, Bastien Van Delft</i>	
Double Machine Learning Density Estimation for Local Treatment Effects with Instruments	21821
<i>Elias Bareinboim, Yonghan Jung, Jin Tian</i>	
Dirichlet Energy Constrained Learning for Deep Graph Neural Networks.....	21834
<i>Rui Chen, Soo-Hyun Choi, Xia Hu, Xiao Huang, Li Li, Daochen Zha, Kaixiong Zhou</i>	

Accelerating Robotic Reinforcement Learning Via Parameterized Action Primitives.....	21847
<i>Murtaza Dalal, Deepak Pathak, Russ Salakhutdinov</i>	
Boosted CVaR Classification.....	21860
<i>Chen Dan, J. Zico Kolter, Pradeep Ravikumar, Arun Suggala, Runtian Zhai</i>	
Disentangled Contrastive Learning on Graphs.....	21872
<i>Hang Li, Haoyang Li, Xin Wang, Zehuan Yuan, Ziwei Zhang, Wenwu Zhu</i>	
Widening the Pipeline in Human-Guided Reinforcement Learning with Explanation and Context-Aware Data Augmentation.....	21885
<i>Lin Guan, Suna (Sihang) Guo, Subbarao Kambhampati, Mudit Verma, Ruohan Zhang</i>	
SOLQ: Segmenting Objects by Learning Queries.....	21898
<i>Bin Dong, Tiancai Wang, Yichen Wei, Fangao Zeng, Xiangyu Zhang</i>	
Extending Lagrangian and Hamiltonian Neural Networks with Differentiable Contact Models.....	21910
<i>Amit Chakraborty, Biswadip Dey, Yaofeng Desmond Zhong</i>	
Best-Case Lower Bounds in Online Learning.....	21923
<i>Cristobal Guzman, Nishant Mehta, Ali Mortazavi</i>	
A Comprehensively Tight Analysis of Gradient Descent for PCA.....	21935
<i>Ping Li, Zhiqiang Xu</i>	
On Robust Optimal Transport: Computational Complexity and Barycenter Computation.....	21947
<i>Hung Bui, Nhat Ho, Khang Le, Huy Nguyen, Quang M Nguyen, Tung Pham</i>	
Asymptotically Best Causal Effect Identification with Multi-Armed Bandits.....	21960
<i>Silvia Chiappa, Alan Malek</i>	
Learning Rule Influences Recurrent Network Representations but Not Attractor Structure in Decision-Making Tasks.....	21972
<i>Jonathan Kao, Michael Kleinman, Brandon McMahan</i>	
Few-Shot Segmentation Via Cycle-Consistent Transformer.....	21984
<i>Guoliang Kang, Yunchao Wei, Yi Yang, Gengwei Zhang</i>	
DropGNN: Random Dropouts Increase the Expressiveness of Graph Neural Networks.....	21997
<i>Lukas Faber, Karolis Martinkus, Pal Andras Papp, Roger Wattenhofer</i>	
Photonic Differential Privacy with Direct Feedback Alignment.....	22010
<i>Alessandro Cappelli, Julien Launay, Hamlet Medina, Ruben Ohana, Iacopo Poli, Alain Rakotomamonjy, Liva Ralaivola</i>	
Searching Parameterized AP Loss for Object Detection.....	22021
<i>Tao Chenxin, Gao Huang, Zizhang Li, Yong Liu, Xizhou Zhu, Jifeng Dai</i>	
Fair Exploration Via Axiomatic Bargaining.....	22034
<i>Jackie Baek, Vivek Farias</i>	
Unifying Lower Bounds on Prediction Dimension of Convex Surrogates.....	22046
<i>Jessica Finocchiaro, Rafael Frongillo, Bo Waggoner</i>	
Ultrahyperbolic Neural Networks.....	22058
<i>Marc Law</i>	

VOLUME 27

NeuroMLR: Robust & Reliable Route Recommendation on Road Networks	22070
<i>Vrittika Bagadia, Jayant Jain, Sahil Manchanda, Sayan Ranu</i>	
Risk Bounds and Calibration for a Smart Predict-Then-Optimize Method.....	22083
<i>Paul Grigas, Heyuan Liu</i>	
Three-Dimensional Spike Localization and Improved Motion Correction for Neuropixels Recordings.....	22095
<i>Julien Boussard, Nishchal Dethle, Hyun Dong Lee, Liam Paninski, Erdem Varol</i>	
Semi-Supervised Semantic Segmentation Via Adaptive Equalization Learning.....	22106
<i>Jinshi Cui, Han Hu, Hanzhe Hu, Liwei Wang, Fangyun Wei, Qiwei Ye</i>	
On the Bias-Variance-Cost Tradeoff of Stochastic Optimization	22119
<i>Xin Chen, Niao He, Yifan Hu</i>	
Averaging on the Bures-Wasserstein Manifold: Dimension-free Convergence of Gradient Descent.....	22132
<i>Jason Altschuler, Sinho Chewi, Patrik R Gerber, Austin Stromme</i>	
Reinforcement Learning in Newcomblike Environments	22146
<i>James Bell, Linda Linsefors, Caspar Oesterheld, Joar Skalse</i>	
Comprehensive Knowledge Distillation with Causal Intervention.....	22158
<i>Xiang Deng, Zhongfei Zhang</i>	
Reinforcement Learning with Latent Flow.....	22171
<i>Pieter Abbeel, Yang Gao, Misha Laskin, Aravind Rajeswaran, Wenling Shang, Aravind Srinivas, Xiaofei Wang</i>	
Understanding How Encoder-Decoder Architectures Attend.....	22184
<i>Kyle Aitken, Yuan Cao, Niru Maheswaranathan, Vinay Ramasesh</i>	
Latent Execution for Neural Program Synthesis Beyond Domain-Specific Languages.....	22196
<i>Xinyun Chen, Dawn Song, Yuandong Tian</i>	
Two Steps to Risk Sensitivity.....	22209
<i>Peter Dayan, Christopher Gagne</i>	
DECAF: Generating Fair Synthetic Data Using Causally-Aware Generative Networks	22221
<i>Jeroen Berrevoets, Trent Kyono, Boris Van Breugel, Mihaela Van Der Schaar</i>	
EvoGrad: Efficient Gradient-Based Meta-Learning and Hyperparameter Optimization	22234
<i>Ondrej Bohdal, Timothy Hospedales, Yongxin Yang</i>	
Biological Key-Value Memory Networks.....	22247
<i>Ching Fang, Danil Tyulmankov, Annapurna Vadaparty, Guangyu Robert Yang</i>	
Correlated Stochastic Block Models: Exact Graph Matching with Applications to Recovering Communities	22259
<i>Miklos Racz, Anirudh Sridhar</i>	
Twice Regularized MDPs and the Equivalence Between Robustness and Regularization	22274
<i>Esther Derman, Matthieu Geist, Shie Mannor</i>	

Nearly Minimax Optimal Reinforcement Learning for Discounted MDPs.....	22288
<i>Quanquan Gu, Jiafan He, Dongruo Zhou</i>	
Sparse Deep Learning: a New Framework Immune to Local Traps and Miscalibration.....	22301
<i>Faming Liang, Yan Sun, Wenjun Xiong</i>	
Calibrating Predictions to Decisions: a Novel Approach to Multi-Class Calibration.....	22313
<i>Stefano Ermon, Michael Kim, Tengyu Ma, Roshni Sahoo, Shengjia Zhao</i>	
Lower Bounds and Optimal Algorithms for Smooth and Strongly Convex Decentralized Optimization Over Time-Varying Networks.....	22325
<i>Elnur Gasanov, Alexander Gasnikov, Dmitry Kovalev, Peter Richtarik</i>	
Testing Probabilistic Circuits	22336
<i>Kuldeep S Meel, Yash Pralhad Pote</i>	
Pseudo-Spherical Contrastive Divergence	22348
<i>Stefano Ermon, Jiaming Song, Yang Song, Lantao Yu</i>	
NORESQA: a Framework for Speech Quality Assessment Using Non-Matching References.....	22363
<i>Anurag Kumar, Pranay Manocha, Buye Xu</i>	
AFEC: Active Forgetting of Negative Transfer in Continual Learning	22379
<i>Chenglong Bao, Zhongfan Jia, Qian Li, Kaisheng Ma, Liyuan Wang, Mingtian Zhang, Yi Zhong, Jun Zhu</i>	
Heterogeneous Multi-Player Multi-Armed ABandits: Closing the Gap and Generalization.....	22392
<i>Cong Shen, Chengshuai Shi, Wei Xiong, Jing Yang</i>	
SWAD: Domain Generalization by Seeking Flat Minima	22405
<i>Junbum Cha, Han-Cheol Cho, Sanghyuk Chun, Kyungjae Lee, Yunsung Lee, Seunghyun Park, Sungrae Park</i>	
Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting	22419
<i>Mingsheng Long, Jianmin Wang, Jiehui Xu, Haixu Wu</i>	
Predicting Event Memorability from Contextual Visual Semantics.....	22431
<i>Fen Fang, Joo-Hwee Lim, Ana Molino, Vigneshwaran Subbaraju, Qianli Xu</i>	
Achieving Forgetting Prevention and Knowledge Transfer in Continual Learning	22443
<i>Zixuan Ke, Bing Liu, Nianzu Ma, Lei Shu, Hu Xu</i>	
Bandits with Many Optimal Arms.....	22457
<i>Alexandra Carpentier, James Cheshire, Pierre Menard, Rianne De Heide</i>	
Combiner: Full Attention Transformer with Sparse Computation Cost	22470
<i>Bo Dai, Hanjun Dai, Zihang Dai, Jure Leskovec, Hongyu Ren, Dale Schuurmans, Mengjiao Yang</i>	
Geometry Processing with Neural Fields	22483
<i>Serge Belongie, Bharath Hariharan, Vladlen Koltun, Guandao Yang</i>	
Contextual Recommendations and Low-Regret Cutting-Plane Algorithms	22498
<i>Sreenivas Gollapudi, Guru Guruganesh, Kostas Kollias, Renato Leme, Pasin Manurangsi, Jon Schneider</i>	

Speech Separation Using an Asynchronous Fully Recurrent Convolutional Neural Network.....	22509
<i>Timo Gerkmann, Xiaolin Hu, Jean-Marie Lemerrier, Kai Li, Yi Luo, Weiyi Zhang</i>	
Reinforcement Learning Enhanced Explainer for Graph Neural Networks	22523
<i>Dongsheng Li, Xiang Li, Caihua Shan, Yifei Shen, Yao Zhang</i>	
NAS-Bench-x11 and the Power of Learning Curves.....	22534
<i>Frank Hutter, Yash Savani, Colin White, Shen Yan</i>	
Observation-Free Attacks on Stochastic Bandits.....	22550
<i>Jacob Abernethy, Bhuvesh Kumar, Yinglun Xu</i>	
Learning Disentangled Behavior Embeddings	22562
<i>Maisan Abboud, Shay Achvat, Amir Ghanayim, Shahar Levy, Gal Mishne, Jackie Schiller, Sivan Schwartz, Changhao Shi</i>	
The Sensory Neuron as a Transformer: Permutation-Invariant Neural Networks for Reinforcement Learning	22574
<i>David Ha, Yujin Tang</i>	
Fast Extra Gradient Methods for Smooth Structured Nonconvex-Nonconcave Minimax Problems	22588
<i>Donghwan Kim, Sucheol Lee</i>	
Analysis of Sensing Spectral for Signal Recovery Under a Generalized Linear Model.....	22601
<i>Junjie Ma, Arian Maleki, Ji Xu</i>	
Revisiting ResNets: Improved Training and Scaling Strategies	22614
<i>Irwan Bello, Ekin Dogus Cubuk, Xianzhi Du, William Fedus, Tsung-Yi Lin, Jonathon Shlens, Aravind Srinivas, Barret Zoph</i>	
Sparse Flows: Pruning Continuous-Depth Models.....	22628
<i>Alexander Amini, Ramin Hasani, Lucas Liebenwein, Daniela Rus</i>	
Spectrum-To-Kernel Translation for Accurate Blind Image Super-Resolution	22643
<i>Yun Cao, Shuo Chen, Xiaozhong Ji, Chuming Lin, Tong Lu, Donghao Luo, Ying Tai, Guangpin Tao, Wenzhuo Wang</i>	
On the Rate of Convergence of Regularized Learning in Games: from Bandits and Uncertainty to Optimism and Beyond.....	22655
<i>Angeliki Giannou, Panayotis Mertikopoulos, Emmanouil-Vasileios Vlatakis-Gkaragkounis</i>	
SLAPS: Self-Supervision Improves Structure Learning for Graph Neural Networks.....	22667
<i>Layla El Asri, Bahare Fatemi, Seyed Mehran Kazemi</i>	
Aligning Pretraining for Detection Via Object-Level Contrastive Learning.....	22682
<i>Yue Gao, Han Hu, Stephen Lin, Fangyun Wei, Zhirong Wu</i>	
Double/Debiased Machine Learning for Dynamic Treatment Effects	22695
<i>Greg Lewis, Vasilis Syrgkanis</i>	
Local Disentanglement in Variational Auto-Encoders Using Jacobian L_1 Regularization.....	22708
<i>Daniel Lee, Travers Rhodes</i>	
Design of Experiments for Stochastic Contextual Linear Bandits	22720
<i>Emma Brunskill, Kefan Dong, Jonathan N Lee, Andrea Zanette</i>	

Encoding Spatial Distribution of Convolutional Features for Texture Representation	22732
<i>Zhile Chen, Feng Li, Jinxiu Liang, Yuhui Quan, Yong Xu</i>	
Training Certifiably Robust Neural Networks with Efficient Local Lipschitz Bounds.....	22745
<i>Anima Anandkumar, Yujia Huang, J. Zico Kolter, Yuanyuan Shi, Huan Zhang</i>	
Average-Reward Learning and Planning with Options	22758
<i>Abhishek Naik, Rich Sutton, Yi Wan</i>	
SSAL: Synergizing Between Self-Training and Adversarial Learning for Domain Adaptive Object Detection	22770
<i>Mohsen Ali, Muhammad Haris Khan, Muhammad Akhtar Munir, M. Sarfraz</i>	
Counterexample Guided RL Policy Refinement Using Bayesian Optimization	22783
<i>Pallab Dasgupta, Briti Gangopadhyay</i>	
Stable, Fast and Accurate: Kernelized Attention with Relative Positional Encoding	22795
<i>Tianle Cai, Di He, Guolin Ke, Shanda Li, Tie-Yan Liu, Shengjie Luo, Dinglan Peng, Liwei Wang, Shuxin Zheng</i>	
Learning in Non-Cooperative Configurable Markov Decision Processes.....	22808
<i>Alessandro Concetti, Alberto Maria Metelli, Giorgia Ramponi, Marcello Restelli</i>	
Identification of Partially Observed Linear Causal Models: Graphical Conditions for the Non- Gaussian and Heterogeneous Cases	22822
<i>Jeffrey Adams, Niels Hansen, Kun Zhang</i>	
DIB-R++: Learning to Predict Lighting and Material with a Hybrid Differentiable Renderer	22834
<i>Wenzheng Chen, Sanja Fidler, Clement Fuji Tsang, Jun Gao, Sameh Khamis, Joey Litalien, Or Litany, Zian Wang</i>	
Coresets for Time Series Clustering	22849
<i>Lingxiao Huang, K Sudhir, Nisheeth Vishnoi</i>	
A Variational Perspective on Diffusion-Based Generative Models and Score Matching	22863
<i>Aaron Courville, Chin-Wei Huang, Jae Hyun Lim</i>	
Online Active Learning with Surrogate Loss Functions.....	22877
<i>Giulia Desalvo, Claudio Gentile, Tobias Sommer Thune</i>	
Does Preprocessing Help Training Over-Parameterized Neural Networks?	22890
<i>Zhao Song, Shuo Yang, Ruizhe Zhang</i>	
Causal Influence Detection for Improving Efficiency in Reinforcement Learning.....	22905
<i>Georg Martius, Bernhard Scholkopf, Maximilian Seitzer</i>	
LADA: Look-Ahead Data Acquisition Via Augmentation for Deep Active Learning	22919
<i>Joonho Jang, Yoon-Yeong Kim, Il-Chul Moon, Kyungwoo Song</i>	

VOLUME 28

Policy Optimization in Adversarial MDPs: Improved Exploration Via Dilated Bonuses.....	22931
<i>Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei</i>	
Multiclass Versus Binary Differentially Private PAC Learning.....	22943
<i>Mark Bun, Marco Gaboardi, Satchit Sivakumar</i>	

Adversarially Robust Change Point Detection	22955
<i>Mengchu Li, Yi Yu</i>	
Cycle Self-Training for Domain Adaptation	22968
<i>Hong Liu, Mingsheng Long, Jianmin Wang</i>	
Novel Visual Category Discovery with Dual Ranking Statistics and Mutual Knowledge Distillation	22982
<i>Kai Han, Bingchen Zhao</i>	
Stochastic Anderson Mixing for Nonconvex Stochastic Optimization	22995
<i>Chenglong Bao, Yang Liu, Fuchao Wei</i>	
Sample-Efficient Reinforcement Learning for Linearly-Parameterized MDPs with a Generative Model	23009
<i>Jianqing Fan, Bingyan Wang, Yuling Yan</i>	
NN-Baker: a Neural-network Infused Algorithmic Framework for Optimization Problems on Geometric Intersection Graphs.....	23023
<i>Evan McCarty, Anastasios Sidiropoulos, Yusu Wang, Qi Zhao</i>	
A Note on Sparse Generalized Eigenvalue Problem	23036
<i>Yunfeng Cai, Guanhua Fang, Ping Li</i>	
RMIX: Learning Risk-Sensitive Policies for Cooperative Reinforcement Learning Agents	23049
<i>Bo An, Xu He, Svetlana Obraztsova, Wei Qiu, Zinovi Rabinovich, Rundong Wang, Xinrun Wang, Runsheng Yu</i>	
Optimal Policies Tend to Seek Power	23063
<i>Andrew Critch, Rohin Shah, Logan Smith, Prasad Tadepalli, Alex Turner</i>	
Catalytic Role of Noise and Necessity of Inductive Biases in the Emergence of Compositional Communication	23075
<i>Tomasz Korbak, Pawel Kolodziej, Lukasz Kucinski, Piotr Milos</i>	
PLUR: a Unifying, Graph-Based View of Program Learning, Understanding, and Repair	23089
<i>Zimin Chen, Vincent J Hellendoorn, Pascal Lamblin, Petros Maniatis, Pierre-Antoine Manzagol, Subhodeep Moitra, Daniel Tarlow</i>	
COCO-LM: Correcting and Contrasting Text Sequences for Language Model Pretraining	23102
<i>Payal Bajaj, Paul Bennett, Jiawei Han, Yu Meng, Xia Song, Chenyan Xiong, Saurabh Tiwary</i>	
Minibatch and Momentum Model-Based Methods for Stochastic Weakly Convex Optimization	23115
<i>Qi Deng, Wenzhi Gao</i>	
XDO: a Double Oracle Algorithm for Extensive-Form Games	23128
<i>Pierre Baldi, Roy Fox, Jb Lanier, Stephen McAleer, Kevin A Wang</i>	
Active Assessment of Prediction Services as Accuracy Surface Over Attribute Combinations	23140
<i>Soumen Chakrabarti, Vihari Piratla, Sunita Sarawagi</i>	
A Mechanistic Multi-Area Recurrent Network Model of Decision-making	23152
<i>Chandramouli Chandrasekaran, Jonathan Kao, Michael Kleinman</i>	
Learning to Compose Visual Relations	23166
<i>Yilun Du, Shuang Li, Nan Liu, Josh Tenenbaum, Antonio Torralba</i>	

Identity Testing for Mallows Model.....	23179
<i>Robert Busa-Fekete, Dimitris Fotakis, Balazs Szorenyi, Emmanouil Zampetakis</i>	
Bandits with Knapsacks Beyond the Worst Case.....	23191
<i>Karthik Abinav Sankararaman, Aleksandrs Slivkins</i>	
Closing the Loop in Medical Decision Support by Understanding Clinical Decision-Making: A Case Study on Organ Transplantation.....	23205
<i>Alihan Huyuk, Fergus Imrie, Daniel Jarrett, Yuchao Qin, Alexander Gimson, Mihaela Van Der Schaar</i>	
Change Point Detection Via Multivariate Singular Spectrum Analysis.....	23218
<i>Arwa Alanqary, Abdullah Alomar, Devavrat Shah</i>	
Meta-Learning to Improve Pre-training.....	23231
<i>David Duvenaud, Simon Kornblith, Jonathan Lorraine, Matthew McDermott, Aniruddh Raghu</i>	
Fair Sparse Regression with Clustering: an InveX Relaxation for a Combinatorial Problem.....	23245
<i>Adarsh Barik, Jean Honorio</i>	
Probabilistic Margins for Instance Reweighting in Adversarial Training.....	23258
<i>Chen Gong, Bo Han, Feng Liu, Tongliang Liu, Gang Niu, Masashi Sugiyama, Mingyuan Zhou, Qizhou Wang</i>	
Unbalanced Optimal Transport Through Non-Negative Penalized Linear Regression.....	23270
<i>Laetitia Chapel, Remi Flamary, Cedric Fevotte, Gilles Gasso, Haoran Wu</i>	
The Difficulty of Passive Learning in Deep Reinforcement Learning.....	23283
<i>Pablo Samuel Castro, Will Dabney, Georg Ostrovski</i>	
Intriguing Properties of Vision Transformers.....	23296
<i>Munawar Hayat, Salman H Khan, Muhammad Muzammal Naseer, Kanchana Ranasinghe, Fahad Shahbaz Khan, Ming-Hsuan Yang</i>	
PartialFed: Cross-Domain Personalized Federated Learning Via Partial Initialization.....	23309
<i>Bo Bai, Hongxing Huo, Benyuan Sun, Yi Yang</i>	
Adaptive Diffusion in Graph Neural Networks.....	23321
<i>Ming Ding, Yuxiao Dong, Evgeny Kharlamov, Jie Tang, Jialin Zhao</i>	
Recurrent Submodular Welfare and Matroid Blocking Semi-Bandits.....	23334
<i>Constantine Caramanis, Orestis Papadigenopoulos</i>	
Representer Point Selection Via Local Jacobian Expansion for Post-Hoc Classifier Explanation of Deep Neural Networks and Ensemble Models.....	23347
<i>Scott Sanner, Yi Sui, Ga Wu</i>	
Editing a Classifier by Rewriting Its Prediction Rules.....	23359
<i>David Bau, Mahalaxmi Elango, Aleksander Madry, Shibani Santurkar, Antonio Torralba, Dimitris Tsipras</i>	
How Modular Should Neural Module Networks Be for Systematic Generalization?.....	23374
<i>Xavier Boix, Vanessa D'Amario, Tomotake Sasaki</i>	
Contrast and Mix: Temporal Contrastive Video Domain Adaptation with Background Mixing.....	23386
<i>Abir Das, Rameswar Panda, Kate Saenko, Aadarsh Sahoo, Rutav Shah</i>	

The Flip Side of the Reweighted Coin: Duality of Adaptive Dropout and Regularization	23401
<i>Richard Baraniuk, Hamid Javadi, Daniel Lejeune</i>	
Active Learning of Convex Halfspaces on Graphs	23413
<i>Thomas Gaertner, Maximilian Thiessen</i>	
Differentiable Spike: Rethinking Gradient-Descent for Training Spiking Neural Networks	23426
<i>Shikuang Deng, Shi Gu, Yufei Guo, Yongqing Hai, Yuhang Li, Shanghang Zhang</i>	
Probabilistic Entity Representation Model for Reasoning Over Knowledge Graphs	23440
<i>Narendra Choudhary, Sumeet Katariya, Nikhil Rao, Chandan Reddy, Karthik Subbian</i>	
Black Box Probabilistic Numerics	23452
<i>Philip Breen, Christopher Foley, Toni Karvonen, Chris Oates, Onur Teymur</i>	
Interpolation Can Hurt Robust Generalization Even When There is No Noise	23465
<i>Michael Aerni, Konstantin Donhauser, Reinhard Heckel, Alexandru Tifrea, Fanny Yang</i>	
On the Equivalence Between Neural Network and Support Vector Machine	23478
<i>Yilan Chen, Wei Huang, Lam Nguyen, Tsui-Wei Weng</i>	
Learning Semantic Representations to Verify Hardware Designs.....	23491
<i>David Bieber, C. Richard Ho, Wenjie (Joe) Jiang, Rishabh Singh, Charles Sutton, Shobha Vasudevan, Hamid Shojaei</i>	
Rebooting ACGAN: Auxiliary Classifier GANs with Stable Training	23505
<i>Minsu Cho, Minguk Kang, Jaesik Park, Woohyeon Shim</i>	
Towards a Theoretical Framework of Out-Of-Distribution Generalization	23519
<i>Tianle Cai, Zhenguo Li, Ruichen Li, Liwei Wang, Chuanlong Xie, Haotian Ye</i>	
Slice Sampling Reparameterization Gradients	23532
<i>Ryan Adams, Diana Cai, David Zoltowski</i>	
Multi-Label Learning with Pairwise Relevance Ordering	23545
<i>Sheng-Jun Huang, Ming-Kun Xie</i>	
Sampling with Trustworthy Constraints: a Variational Gradient Framework	23557
<i>Qiang Liu, Xingchao Liu, Xin Tong</i>	
Robust and Decomposable Average Precision for Image Retrieval	23569
<i>Nicolas Audebert, Xavier Bitot, Clement Rambour, Elias Ramzi, Nicolas Thome</i>	
Fast Rates for Prediction with Limited Expert Advice	23582
<i>Gilles Blanchard, El Mehdi Saad</i>	
Probabilistic Transformer for Time Series Analysis	23592
<i>David S Matteson, Binh Tang</i>	
A Hierarchical Reinforcement Learning Based Optimization Framework for Large-Scale Dynamic Pickup and Delivery Problems	23609
<i>Jianye Hao, Xiaotian Hao, Zhigang Li, Xing Liu, Jiawen Lu, Yi Ma, Zhaopeng Meng, Jie Tang, Tong Xialiang, Mingxuan Yuan</i>	
Spatio-Temporal Variational Gaussian Processes	23621
<i>Theodoros Damoulas, Oliver Hamelijnck, Niki Loppi, Arno Solin, William Wilkinson</i>	

MERLOT: Multimodal Neural Script Knowledge Models	23634
<i>Jize Cao, Yejin Choi, Ali Farhadi, Jack Hessel, Ximing Lu, Jae Sung Park, Youngjae Yu, Rowan Zellers</i>	
Fast Approximate Dynamic Programming for Infinite-Horizon Markov Decision Processes	23652
<i>Gyula Max, Peyman Mohajerin Esfahani, Mohamad Amin Sharifi Kolarijani</i>	
Adaptive Risk Minimization: Learning to Adapt to Domain Shift.....	23664
<i>Nikita Dhawan, Chelsea Finn, Abhishek Gupta, Sergey Levine, Henrik Marklund, Marvin Zhang</i>	
Learning State Representations from Random Deep Action-Conditional Predictions	23679
<i>Richard L Lewis, Satinder Singh, Vivek Veeriah, Risto Vuorio, Zeyu Zheng</i>	
Mixability Made Efficient: Fast Online Multiclass Logistic Regression.....	23692
<i>Pierre Gaillard, Remi Jezequel, Alessandro Rudi</i>	
Tracking People with 3D Representations	23703
<i>Angjoo Kanazawa, Jitendra Malik, Georgios Pavlakos, Jathushan Rajasegaran</i>	
Off-Policy Risk Assessment in Contextual Bandits	23714
<i>Kamyar Azizzadenesheli, Audrey Huang, Liu Leqi, Zachary Lipton</i>	
Adaptive Denoising Via GainTuning	23727
<i>Peter Crozier, Carlos Fernandez-Granda, Ramon Manzorro, Sreyas Mohan, Eero Simoncelli, Joshua L Vincent</i>	
Optimal Sketching for Trace Estimation	23741
<i>Shuli Jiang, Hai Pham, David Woodruff, Richard Zhang</i>	
Estimating Multi-Cause Treatment Effects Via Single-Cause Perturbation.....	23754
<i>Alicia Curth, Zhaozhi Qian, Mihaela Van Der Schaar</i>	
Be Confident! Towards Trustworthy Graph Neural Networks Via Confidence Calibration.....	23768
<i>Hongrui Liu, Chuan Shi, Xiao Wang, Cheng Yang</i>	
Learning Riemannian Metric for Disease Progression Modeling.....	23780
<i>Stanley Durrleman, Samuel Gruffaz, Bruno Jedynek, Etienne Maheux, Pierre-Emmanuel Poulet</i>	

VOLUME 29

Bias and Variance of the Bayesian-Mean Decoder	23793
<i>Arthur Prat-Carrabin, Michael Woodford</i>	
MIRACLE: Causally-Aware Imputation Via Learning Missing Data Mechanisms	23806
<i>Alexis Bellot, Trent Kyono, Yao Zhang, Mihaela Van Der Schaar</i>	
Efficient Training of Visual Transformers with Small Datasets.....	23818
<i>Wei Bi, Bruno Lepri, Yahui Liu, Marco Nadai, Enver Sangineto, Nicu Sebe</i>	
Small Random Initialization is Akin to Spectral Learning: Optimization and Generalization Guarantees for Overparameterized Low-Rank Matrix Reconstruction	23831
<i>Mahdi Soltanolkotabi, Dominik Stoger</i>	

Efficient Combination of Rematerialization and Offloading for Training DNNs	23844
<i>Olivier Beaumont, Lionel Eyraud-Dubois, Alena Shilova</i>	
Particle Cloud Generation with Message Passing Generative Adversarial Networks	23858
<i>Javier Duarte, Dimitrios Gunopulos, Raghav Kansal, Breno Orzari, Maurizio Pierini, Hao Su, Thiago Tomei, Mary Touranakou, Jean-Roch Vlimant</i>	
CoFiNet: Reliable Coarse-To-fine Correspondences for Robust PointCloud Registration	23872
<i>Benjamin Busam, Slobodan Ilic, Fu Li, Mahdi Saleh, Hao Yu</i>	
Partial Success in Closing the Gap Between Human and Machine Vision	23885
<i>Matthias Bethge, Wieland Brendel, Robert Geirhos, Benjamin Mitzkus, Kantharaju Narayanappa, Tizian Thieringer, Felix A. Wichmann</i>	
LLC: Accurate, Multi-Purpose Learnt Low-dimensional Binary Codes	23900
<i>Ali Farhadi, Prateek Jain, Sham Kakade, Aditya Kusupati, Jae Sung Park, Krishna Pillutla, Vivek Ramanujan, Raghav Somani, Matthew Wallingford</i>	
Analytic Insights into Structure and Rank of Neural Network Hessian Maps	23914
<i>Gregor Bachmann, Thomas Hofmann, Sidak Pal Singh</i>	
Well-Tuned Simple Nets Excel on Tabular Datasets	23928
<i>Josif Grabocka, Frank Hutter, Arlind Kadra, Marius Lindauer</i>	
POODLE: Improving Few-Shot Learning Via Penalizing Out-of-Distribution Samples.....	23942
<i>Binh-Son Hua, Duong Le, Khoi Nguyen, Khoi Duc Nguyen, Rang Nguyen, Quoc-Huy Tran</i>	
Combinatorial Pure Exploration with Bottleneck Reward Function	23956
<i>Wei Chen, Yihan Du, Yuko Kuroki</i>	
Densely Connected Normalizing Flows	23968
<i>Matej Grbic, Ivan Grubisic, Sinisa Segvic</i>	
Snowflake: Scaling GNNs to High-Dimensional Continuous Control Via Parameter Freezing.....	23983
<i>Charles Blake, Maximilian Igl, Vitaly Kurin, Shimon Whiteson</i>	
Subgame Solving Without Common Knowledge.....	23993
<i>Tuomas Sandholm, Brian Zhang</i>	
Fair Algorithms for Multi-Agent Multi-Armed Bandits.....	24005
<i>Safwan Hossain, Evi Micha, Nisarg Shah</i>	
VAST: Value Function Factorization with Variable Agent Sub-Teams	24018
<i>Philipp Altmann, Lenz Belzner, Thomas Gabor, Claudia Linnhoff-Popien, Thomy Phan, Fabian Ritz</i>	
On the Stochastic Stability of Deep Markov Models	24033
<i>Jan Drgona, Mahantesh Halappanavar, Frank Liu, Sayak Mukherjee, Jiaxin Zhang</i>	
Multiwavelet-Based Operator Learning for Differential Equations	24048
<i>Paul Bogdan, Gaurav Gupta, Xiongye Xiao</i>	
Intermediate Layers Matter in Momentum Contrastive Self Supervised Learning	24063
<i>Aakash Kaku, Narges Razavian, Sahana Upadhy</i>	

An Efficient Pessimistic-Optimistic Algorithm for Stochastic Linear Bandits with General Constraints.....	24075
<i>Bin Li, Xin Liu, Pengyi Shi, Lei Ying</i>	
Efficiently Learning One Hidden Layer ReLU Networks from Queries.....	24087
<i>Sitan Chen, Adam Klivans, Raghu Meka</i>	
Learning Nonparametric Volterra Kernels with Gaussian Processes	24099
<i>Magnus Ross, Michael T Smith, Mauricio Alvarez</i>	
DiBS: Differentiable Bayesian Structure Learning	24111
<i>Andreas Krause, Lars Lorch, Jonas Rothfuss, Bernhard Scholkopf</i>	
Nonparametric Estimation of Continuous DPPs with Kernel Methods.....	24124
<i>Remi Bardenet, Michaël Fanuel</i>	
FINE Samples for Learning with Noisy Labels.....	24137
<i>Sangwook Cho, Jinhwan Choi, Taehyeon Kim, Jongwoo Ko, Se-Young Yun</i>	
Residual2Vec: Debiasing Graph Embedding with Random Graphs	24150
<i>Yong-Yeol Ahn, Isabel Constantino, Sadamori Kojaku, Jisung Yoon</i>	
Benign Overfitting in Multiclass Classification: All Roads Lead to Interpolation.....	24164
<i>Vidya Muthukumar, Christos Thrampoulidis, Ke Wang</i>	
Instance-Dependent Bounds for Zeroth-order Lipschitz Optimization with Error Certificates.....	24180
<i>Francois Bachoc, Tom Cesari, Sebastien Gerchinovitz</i>	
Training Neural Networks with Fixed Sparse Masks	24193
<i>Varun Nair, Colin Raffel, Yi-Lin Sung</i>	
VATT: Transformers for Multimodal Self-Supervised Learning from Raw Video, Audio and Text	24206
<i>Hassan Akbari, Shih-Fu Chang, Wei-Hong Chuang, Yin Cui, Boqing Gong, Rui Qian, Liangzhe Yuan</i>	
Analyzing the Generalization Capability of SGLD Using Properties of Gaussian Channels.....	24222
<i>Flavio Calmon, Rui Gao, Yizhe Huang, Hao Wang</i>	
Learning to Schedule Heuristics in Branch and Bound.....	24235
<i>Antonia Chmiela, Ambros Gleixner, Elias Khalil, Andrea Lodi, Sebastian Pokutta</i>	
On Training Implicit Models	24247
<i>Shaojie Bai, Zhengyang Geng, Zhouchen Lin, Yisen Wang, Xin-Yu Zhang</i>	
MLP-Mixer: an all-MLP Architecture for Vision	24261
<i>Lucas Beyer, Alexey Dosovitskiy, Neil Houlsby, Daniel Keysers, Alexander Kolesnikov, Mario Lucic, Andreas Steiner, Ilya Tolstikhin, Thomas Unterthiner, Jakob Uszkoreit, Jessica Yung, Xiaohua Zhai</i>	
A Framework to Learn with Interpretation.....	24273
<i>Pavlo Mozharovskyi, Jayneel Parekh, Florence D'Alche-Buc</i>	
One Loss for All: Deep Hashing with a Single Cosine Similarity Based Learning Objective	24286
<i>Chee Seng Chan, Jiun Tian Hoe, Kam Woh Ng, Yi-Zhe Song, Tao Xiang, Tianyu Zhang</i>	
Fast and Accurate Randomized Algorithms for Low-Rank Tensor Decompositions.....	24299
<i>Linjian Ma, Edgar Solomonik</i>	

Communication-Efficient SGD: from Local SGD to One-Shot Averaging	24313
<i>Alex Olshevsky, Yannis Paschalidis, Artin Spiridonoff</i>	
Memory Efficient Meta-Learning with Large Images.....	24327
<i>John Bronskill, Katja Hofmann, Daniela Massiceti, Sebastian Nowozin, Massimiliano Patacchiola, Richard Turner</i>	
On the Power of Differentiable Learning Versus PAC and SQ Learning	24340
<i>Emmanuel Abbe, Pritish Kamath, Eran Malach, Colin Sandon, Nathan Srebro</i>	
Can We Globally Optimize Cross-Validation Loss? Quasiconvexity in Ridge Regression.....	24352
<i>Tamara Broderick, Zachary Frangella, Will Stephenson, Madeleine Udell</i>	
Adaptive Proximal Gradient Methods for Structured Neural Networks.....	24365
<i>Aurelie Lozano, Eunho Yang, Jihun Yun</i>	
Discovering and Achieving Goals Via World Models	24379
<i>Kostas Daniilidis, Danijar Hafner, Russell Mendonca, Deepak Pathak, Oleh Rybkin</i>	
Understanding and Improving Early Stopping for Learning with Noisy Labels	24392
<i>Yingbin Bai, Bo Han, Jiatong Li, Tongliang Liu, Yinian Mao, Gang Niu, Erkun Yang, Yanhua Yang</i>	
Distributionally Robust Imitation Learning.....	24404
<i>Mohammad Ali Bashiri, Xinhua Zhang, Brian Ziebart</i>	
On the Power of Edge Independent Graph Models	24418
<i>Sudhanshu Chanpuriya, Cameron Musco, Konstantinos Sotiropoulos, Charalampos Tsourakakis</i>	
Stochastic Online Linear Regression: the Forward Algorithm to Replace Ridge.....	24430
<i>Odalric-Ambrym Maillard, Reda Ouhamma, Vianney Perchet</i>	
Dr Jekyll & Mr Hyde: the Strange Case of Off-Policy Policy Updates	24442
<i>Romain Laroche, Remi Tachet Des Combes</i>	
Understanding Adaptive, Multiscale Temporal Integration in Deep Speech Recognition Systems.....	24455
<i>Menoua Keshishian, Nima Mesgarani, Samuel Norman-Haignere</i>	
VidLanKD: Improving Language Understanding Via Video-Distilled Knowledge Transfer	24468
<i>Mohit Bansal, Jaemin Cho, Hao Tan, Zineng Tang</i>	
Detecting Individual Decision-Making Style: Exploring Behavioral Stylometry in Chess.....	24482
<i>Ashton Anderson, Jon Kleinberg, Reid McIlroy-Young, Siddhartha Sen, Yu Wang</i>	
Coupled Gradient Estimators for Discrete Latent Variables	24498
<i>Zhe Dong, Andriy Mnih, George Tucker</i>	
AutoGEL: an Automated Graph Neural Network with Explicit Link Information	24509
<i>Lei Chen, Shimin Di, Zhili Wang</i>	
RL for Latent MDPs: Regret Guarantees and a Lower Bound.....	24523
<i>Constantine Caramanis, Yonathan Efroni, Jeongyeol Kwon, Shie Mannor</i>	
Adaptive Sampling for Minimax Fair Classification.....	24535
<i>Greg Fields, Mohammad Ghavamzadeh, Tara Javidi, Shubhanshu Shekhar</i>	

Structured in Space, Randomized in Time: Leveraging Dropout in RNNs for Efficient Training.....	24545
<i>Chita Das, Huaipan Jiang, Mahmut Kandemir, Anup Sarma, Sonali Singh, Rui Zhang</i>	
Variational Continual Bayesian Meta-Learning.....	24556
<i>Jinyuan Fang, Shangsong Liang, Zaiqiao Meng, Emine Yilmaz, Qiang Zhang</i>	
Recognizing Vector Graphics Without Rasterization.....	24569
<i>Xuanyi Dong, Xinyang Jiang, Lu Liu, Dongsheng Li, Caihua Shan, Yifei Shen</i>	
On Episodes, Prototypical Networks, and Few-Shot Learning	24581
<i>Luca Bertinetto, Steinar Laenen</i>	
Pointwise Bounds for Distribution Estimation Under Communication Constraints.....	24593
<i>Wei-Ning Chen, Peter Kairouz, Ayfer Ozgur</i>	
CHIP: CHannel Independence-Based Pruning for Compact Neural Networks.....	24604
<i>Saman Aliari Zonouz, Huy Phan, Yang Sui, Yi Xie, Miao Yin, Bo Yuan</i>	
Federated Split Task-Agnostic Vision Transformer for COVID-19 CXR Diagnosis	24617
<i>Boah Kim, Gwanghyun Kim, Jeongsol Kim, Sangjoon Park, Jong Chul Ye</i>	
Active Offline Policy Selection.....	24631
<i>Yutian Chen, Misha Denil, Caglar Gulcehre, Ksenia Konyushova, Daniel Mankowitz, Cosmin Paduraru, Thomas Paine, Nando De Freitas</i>	

VOLUME 30

Unsupervised Representation Transfer for Small Networks: I Believe I Can Distill On-The-Fly	24645
<i>Hee Min Choi, Hyoa Kang, Dokwan Oh</i>	
Understanding Bandits with Graph Feedback	24659
<i>Houshuang Chen, Zengfeng Huang, Shuai Li, Chihao Zhang</i>	
Information-Theoretic Generalization Bounds for Black-box Learning Algorithms	24670
<i>Aram Galstyan, Hrayr Harutyunyan, Maxim Raginsky, Greg Ver Steeg</i>	
Trash Or Treasure? an Interactive Dual-Stream Strategy for Single Image Reflection Separation.....	24683
<i>Xiaojie Guo, Qiming Hu</i>	
Rot-Pro: Modeling Transitivity by Projection in Knowledge Graph Embedding	24695
<i>Lei Huang, Jie Luo, Tengwei Song</i>	
Planning from Pixels in Environments with Combinatorially Hard Search Spaces	24707
<i>Marco Bagatella, Georg Martius, Miroslav Olsak, Michal Rolínek</i>	
PLUGIn: a Simple Algorithm for Inverting Generative Models with Recovery Guarantees	24719
<i>Babhru Joshi, Xiaowei Li, Yaniv Plan, Ozgur Yilmaz</i>	
Modular Gaussian Processes for Transfer Learning.....	24730
<i>Antonio Artes, Pablo Moreno-Muñoz, Mauricio Alvarez</i>	
Neural Human Performer: Learning Generalizable Radiance Fields for Human Performance Rendering	24741
<i>Duygu Ceylan, Henry Fuchs, Dahun Kim, Youngjoong Kwon</i>	

Locally Differentially Private Estimation of Functionals of Discrete Distributions.....	24753
<i>Cristina Butucea, Yann Issartel</i>	
Asymptotics of Representation Learning in Finite Bayesian Neural Networks	24765
<i>Abdulkadir Canatar, Cengiz Pehlevan, Ben Ruben, Jacob Zavatore-Veth</i>	
Adaptive Ensemble Q-Learning: Minimizing Estimation Bias Via Error Feedback.....	24778
<i>Sen Lin, Hang Wang, Junshan Zhang</i>	
Domain Adaptation with Invariant Representation Learning: What Transformations to Learn?	24791
<i>Ruichu Cai, Jaime Carbonell, Mingming Gong, Zijian Li, Petar Stojanov, Kun Zhang</i>	
CSDI: Conditional Score-Based Diffusion Models for Probabilistic Time Series Imputation.....	24804
<i>Stefano Ermon, Jiaming Song, Yang Song, Yusuke Tashiro</i>	
Causal Bandits with Unknown Graph Structure.....	24817
<i>Yangyi Lu, Amirhossein Meisami, Ambuj Tewari</i>	
Piper: Multidimensional Planner for DNN Parallelization	24829
<i>Deepak Narayanan, Amar Phanishayee, Jakub Tarnawski</i>	
Causal Effect Inference for Structured Treatments	24841
<i>Jean Kaddour, Matt Kusner, Qi Liu, Ricardo Silva, Yuchen Zhu</i>	
Efficient Hierarchical Bayesian Inference for Spatio-Temporal Regression Models in Neuroimaging.....	24855
<i>Chang Cai, Yijing Gao, Sanjay Ghosh, Ali Hashemi, Stefan Haufe, Klaus-Robert Muller, Srikantan Nagarajan</i>	
Topological Attention for Time Series Forecasting	24871
<i>Florian Graf, Christoph Hofer, Roland Kwitt, Sebastian Zeng</i>	
Local Signal Adaptivity: Provable Feature Learning in Neural Networks Beyond Kernels	24883
<i>Stefani Karp, Yuanzhi Li, Aarti Singh, Ezra Winston</i>	
IA-RED ² : Interpretability-Aware Redundancy Reduction for Vision Transformers.....	24898
<i>Rogério Feris, Yifan Jiang, Aude Oliva, Bowen Pan, Rameswar Panda, Zhangyang Wang</i>	
Symbolic Regression Via Deep Reinforcement Learning Enhanced Genetic Programming Seeding	24912
<i>Ruben Glatt, Mikel Landajuela, Terrell Mundhenk, Brenden K Petersen, Claudio P Santiago, Daniel Faissol</i>	
Choose a Transformer: Fourier Or Galerkin	24924
<i>Shuhao Cao</i>	
A Causal Lens for Controllable Text Generation	24941
<i>Zhiting Hu, Li Erran Li</i>	
Differentially Private Multi-Armed Bandits in the Shuffle Model.....	24956
<i>Haim Kaplan, Yishay Mansour, Uri Stemmer, Jay Tenenbaum</i>	
Dual Adaptivity: a Universal Algorithm for Minimizing the Adaptive Regret of Convex Functions.....	24968
<i>Wei Jiang, Wei-Wei Tu, Guanghui Wang, Lijun Zhang, Zhi-Hua Zhou</i>	
Learning Hard Optimization Problems: a Data Generation Perspective	24981
<i>Ferdinando Fioretto, James Kotary, Pascal Van Hentenryck</i>	

Canonical Capsules: Self-Supervised Capsules in Canonical Pose	24993
<i>Boyang Deng, Geoffrey Hinton, Sara Sabour, Weiwei Sun, Andrea Tagliasacchi, Soroosh Yazdani, Kwang Moo Yi</i>	
Characterizing Generalization Under Out-Of-Distribution Shifts in Deep Metric Learning.....	25006
<i>Marzyeh Ghassemi, Timo Milbich, Bjorn Ommer, Karsten Roth, Ludwig Schmidt, Samarth Sinha</i>	
Dynamics-Regulated Kinematic Policy for Egocentric Pose Estimation	25019
<i>Ryo Hachiuma, Kris Kitani, Zhengyi Luo, Ye Yuan</i>	
Never Go Full Batch (in Stochastic Convex Optimization)	25033
<i>Idan Amir, Yair Carmon, Tomer Koren, Roi Livni</i>	
Collaborative Learning in the Jungle (Decentralized, Byzantine, Heterogeneous, Asynchronous and Nonconvex Learning).....	25044
<i>El Mahdi El-Mhamdi, Sadegh Farhadkhani, Rachid Guerraoui, Arsany Guirguis, Lê-Nguyên Hoang, Sebastien Rouault</i>	
Not All Low-Pass Filters Are Robust in Graph Convolutional Networks.....	25058
<i>Yatao Bian, Heng Chang, Junzhou Huang, Yu Rong, Xin Wang, Tingyang Xu, Shiji Zhou, Wenwu Zhu</i>	
Counterfactual Maximum Likelihood Estimation for Training Deep Networks	25072
<i>Wenhu Chen, Michael Saxon, William Yang Wang, Xinyi Wang</i>	
Robust Optimization for Multilingual Translation with Imbalanced Data	25086
<i>Hongyu Gong, Xian Li</i>	
A/B/n Testing with Control in the Presence of Subpopulations	25100
<i>Dennis Bohle, Olivier Cappe, Aurelien Garivier, Christina Katsimerou, Wouter Koolen, Yoan Russac</i>	
Using Random Effects to Account for High-Cardinality Categorical Features and Repeated Measures in Deep Neural Networks	25111
<i>Saharon Rosset, Giora Simchoni</i>	
Learning Debaised Representation Via Disentangled Feature Augmentation.....	25123
<i>Jaegul Choo, Eungyeup Kim, Jihyeon Lee, Jungsoo Lee, Juyoung Lee</i>	
Scallop: from Probabilistic Deductive Databases to Scalable Differentiable Reasoning	25134
<i>Binghong Chen, Jiani Huang, Ziyang Li, Mayur Naik, Karan Samel, Xujie Si, Le Song</i>	
Learning to Synthesize Programs as Interpretable and Generalizable Policies	25146
<i>Joseph Lim, Shao-Hua Sun, Dweep Trivedi, Jesse Zhang</i>	
The Functional Specialization of Visual Cortex Emerges from Training Parallel Pathways with Self-Supervised Predictive Learning	25164
<i>Shahab Bakhtiari, Timothy Lillicrap, Patrick Mineault, Christopher Pack, Blake Richards</i>	
Adversarial Training Helps Transfer Learning Via Better Representations	25179
<i>Zhun Deng, Kenji Kawaguchi, Kailas Vodrahalli, Linjun Zhang, James Zou</i>	
Improving Coherence and Consistency in Neural Sequence Models with Dual-System, Neuro-Symbolic Reasoning	25192
<i>Brenden Lake, Maxwell Nye, Josh Tenenbaum, Michael Tessler</i>	

Learning the Optimal Tikhonov Regularizer for Inverse Problems	25205
<i>Giovanni Alberti, Ernesto De Vito, Matti Lassas, Luca Ratti, Matteo Santacesaria</i>	
NovelD: a Simple Yet Effective Exploration Criterion.....	25217
<i>Joseph Gonzalez, Kurt Keutzer, Yuandong Tian, Xiaolong Wang, Yi Wu, Huazhe Xu, Tianjun Zhang</i>	
On Margin-Based Cluster Recovery with Oracle Queries.....	25231
<i>Marco Bressan, Nicolò Cesa-Bianchi, Silvio Lattanzi, Andrea Paudice</i>	
Multi-Scale Representation Learning on Proteins	25244
<i>Charlotte Bunne, Andreas Krause, Vignesh Ram Somnath</i>	
Sparse Quadratic Optimisation Over the Stiefel Manifold with Application to Permutation Synchronisation	25256
<i>Florian Bernard, Daniel Cremers, Johan Thunberg</i>	
Second-Order Neural ODE Optimizer.....	25267
<i>Tianrong Chen, Guan-Horng Liu, Evangelos Theodorou</i>	
Graph Neural Networks with Local Graph Parameters	25280
<i>Pablo Barcelo, Floris Geerts, Juan Reutter, Maksimilian Ryschkov</i>	
Closing the Gap: Tighter Analysis of Alternating Stochastic Gradient Methods for Bilevel Problems.....	25294
<i>Tianyi Chen, Yuejiao Sun, Wotao Yin</i>	
Dense Unsupervised Learning for Video Segmentation.....	25308
<i>Nikita Araslanov, Stefan Roth, Simone Schaub-Meyer</i>	
Charting and Navigating the Space of Solutions for Recurrent Neural Networks.....	25320
<i>Omri Barak, Kabir V Dabholkar, Elia Turner</i>	
Fast Training Method for Stochastic Compositional Optimization Problems	25334
<i>Hongchang Gao, Heng Huang</i>	
Dual-Stream Network for Visual Recognition	25346
<i>Errui Ding, Shumin Han, Teli Ma, Mingyuan Mao, Yan Peng, Baochang Zhang, Renrui Zhang, Honghui Zheng, Peng Gao</i>	
Estimating High Order Gradients of the Data Distribution by Denoising	25359
<i>Stefano Ermon, Wenzhe Li, Chenlin Meng, Yang Song</i>	
Machine Versus Human Attention in Deep Reinforcement Learning Tasks.....	25370
<i>Dana Ballard, Suna (Sihang) Guo, Mary Hayhoe, Bo Liu, Peter Stone, Ruohan Zhang, Yifeng Zhu</i>	
Reusing Combinatorial Structure: Faster Iterative Projections Over Submodular Base Polytopes.....	25386
<i>Swati Gupta, Jai Moondra, Hassan Mortagy</i>	
Constrained Optimization to Train Neural Networks on Critical and Under-Represented Classes.....	25400
<i>Olivio Donati, Ertunc Erdil, Andeas Hotker, Ender Konukoglu, Sara Sangalli</i>	
Collapsed Variational Bounds for Bayesian Neural Networks.....	25412
<i>Andrew Foong, Siddharth Swaroop, Marcin Tomczak, Richard Turner</i>	

Consistent Estimation for PCA and Sparse Regression with Oblivious Outliers	25427
<i>Chih-Hung Liu, Rajai Nasser, Gleb Novikov, David Steurer, Stefan Tiegel, Tommaso D'Orsi</i>	
Offline Constrained Multi-Objective Reinforcement Learning Via Pessimistic Dual Value Iteration.....	25439
<i>Zhaoran Wang, Runzhe Wu, Zhuoran Yang, Yufeng Zhang</i>	
Absolute Neighbour Difference Based Correlation Test for Detecting Heteroscedastic Relationships	25452
<i>Lifeng Zhang</i>	
Batch Multi-Fidelity Bayesian Optimization with Deep Auto-Regressive Networks	25463
<i>Robert Kirby, Shibo Li, Shandian Zhe</i>	
Mastering Atari Games with Limited Data	25476
<i>Pieter Abbeel, Yang Gao, Thanard Kurutach, Shaohuai Liu, Weirui Ye</i>	
Dealing with Misspecification in Fixed-Confidence Linear Top-m Identification.....	25489
<i>Remy Degenne, Clemence Reda, Andrea Tirinzoni</i>	

VOLUME 31

Why Generalization in RL is Difficult: Epistemic POMDPs and Implicit Partial Observability	25502
<i>Ryan Adams, Dibya Ghosh, Aviral Kumar, Sergey Levine, Jad Rahme, Amy Zhang</i>	
Set Prediction in the Latent Space	25516
<i>Ekapol Chuangsuwanich, Tirasan Khandhawit, Burin Naowarat, Chawan Piansaddhayanon, Konpat Preechakul, Sira Sriswasdi</i>	
Best of Both Worlds: Practical and Theoretically Optimal Submodular Maximization in Parallel	25528
<i>Yixin Chen, Tommoy Dey, Alan Kuhnle</i>	
Fine-Grained Generalization Analysis of Inductive Matrix Completion.....	25540
<i>Rodrigo Alves, Marius Kloft, Antoine Ledent, Yunwen Lei</i>	
Learning Frequency Domain Approximation for Binary Neural Networks	25553
<i>Kai Han, Yehui Tang, Yunhe Wang, Chunjing Xu, Chang Xu, Yixing Xu</i>	
Reformulating Zero-Shot Action Recognition for Multi-label Actions.....	25566
<i>Kevin Duarte, Alec Kerrigan, Yogesh Rawat, Mubarak Shah</i>	
Optimal Best-Arm Identification Methods for Tail-Risk Measures	25578
<i>Shubhada Agrawal, Sandeep Juneja, Wouter Koolen</i>	
SyMetric: Measuring the Quality of Learnt Hamiltonian Dynamics Inferred from Vision.....	25591
<i>Aleksandar Botev, Irina Higgins, Andrew Jaegle, Peter Wirsberger</i>	
Learning with Holographic Reduced Representations	25606
<i>Sunil Gandhi, Ashwinkumar Ganesan, Hang Gao, James Holt, Mark McLean, Tim Oates, Edward Raff</i>	
Learning Barrier Certificates: Towards Safe Reinforcement Learning with Zero Training-Time Violations	25621
<i>Yuping Luo, Tengyu Ma</i>	
On the Second-Order Convergence Properties of Random Search Methods	25633
<i>Aurelien Lucchi, Antonio Orvieto, Adamos Solomou</i>	

Noether's Learning Dynamics: Role of Symmetry Breaking in Neural Networks	25646
<i>Daniel Kunin, Hidenori Tanaka</i>	
A Theory of the Distortion-Perception Tradeoff in Wasserstein Space	25661
<i>Dror Freirich, Ron Meir, Tomer Michaeli</i>	
Neural Production Systems	25673
<i>Anirudh Goyal Alias Parth Goyal, Philippe Beaudoin, Yoshua Bengio, Charles Blundell, Aniket Didolkar, Nicolas Heess, Nan Rosemary Ke, Michael Mozer</i>	
Smoothness Matrices Beat Smoothness Constants: Better Communication Compression Techniques for Distributed Optimization	25688
<i>Filip Hanzely, Peter Richtarik, Mher Safaryan</i>	
Increasing Liquid State Machine Performance with Edge-Of-Chaos Dynamics Organized by Astrocyte-modulated Plasticity.....	25703
<i>Vladimir Ivanov, Konstantinos Michmizos</i>	
Fair Sortition Made Transparent.....	25720
<i>Bailey Flanigan, Gregory Kehne, Ariel Procaccia</i>	
A Max-Min Entropy Framework for Reinforcement Learning	25732
<i>Seungyul Han, Youngchul Sung</i>	
Reward is Enough for Convex MDPs	25746
<i>Guillaume Desjardins, Brendan O'Donoghue, Satinder Singh, Tom Zahavy</i>	
Fast Doubly-Adaptive MCMC to Estimate the Gibbs Partition Function with Weak Mixing Time Bounds.....	25760
<i>Cyrus Cousins, Shahrzad Haddadan, Eli Upfal, Yue Zhuang</i>	
Does Enforcing Fairness Mitigate Biases Caused by Subpopulation Shift?.....	25773
<i>Subha Maity, Debarghya Mukherjee, Yuekai Sun, Mikhail Yurochkin</i>	
Implicit Deep Adaptive Design: Policy-Based Experimental Design Without Likelihoods	25785
<i>Adam Foster, Michael Gutmann, Desi R Ivanova, Steven Kleinegesse, Thomas Rainforth</i>	
Sample-Efficient Learning of Stackelberg Equilibria in General-Sum Games	25799
<i>Yu Bai, Chi Jin, Huan Wang, Caiming Xiong</i>	
Non-Approximate Inference for Collective Graphical Models on Path Graphs Via Discrete Difference of Convex Algorithm.....	25812
<i>Yasunori Akagi, Hideaki Kim, Takeshi Kurashima, Naoki Marumo, Hiroyuki Toda</i>	
Implicit Task-Driven Probability Discrepancy Measure for Unsupervised Domain Adaptation	25824
<i>Kaiqi Jiang, Mao Li, Xinhua Zhang</i>	
SBO-RNN: Reformulating Recurrent Neural Networks Via Stochastic Bilevel Optimization.....	25839
<i>Yanhua Li, Guojun Wu, Yun Yue, Ziming Zhang, Haichong Zhang</i>	
Navigating to the Best Policy in Markov Decision Processes	25852
<i>Aymen Al Marjani, Aurelien Garivier, Alexandre Proutiere</i>	
A Faster Decentralized Algorithm for Nonconvex Minimax Problems	25865
<i>Heng Huang, Feihu Huang, Wenhan Xian, Yanfu Zhang</i>	

Generalization Bounds for Meta-Learning: an Information-Theoretic Analysis.....	25878
<i>Qi Chen, Mario Marchand, Changjian Shui</i>	
ReLU Regression with Massart Noise.....	25891
<i>Ilias Diakonikolas, Jong Ho Park, Christos Tzamos</i>	
Identification of the Generalized Condorcet Winner in Multi-Dueling Bandits.....	25904
<i>Viktor Bengs, Bjorn Haddendorst, Eyke Hullermeier</i>	
Robust Inverse Reinforcement Learning Under Transition Dynamics Mismatch.....	25917
<i>Volkan Cevher, Yu-Ting Huang, Parameswaran Kamalaruban, Luca Viano, Adrian Weller</i>	
Re-Ranking for Image Retrieval and Transductive Few-shot Classification.....	25932
<i>Mathieu Aubry, Shell Xu Hu, Xi Shen, Othman Sbai, Yang Xiao</i>	
Post-Processing for Individual Fairness	25944
<i>Debarghya Mukherjee, Felix Petersen, Yuekai Sun, Mikhail Yurochkin</i>	
OpenMatch: Open-Set Semi-supervised Learning with Open-set Consistency Regularization	25956
<i>Donghyun Kim, Kate Saenko, Kuniaki Saito</i>	
End-To-End Training of Multi-Document Reader and Retriever for Open-Domain Question Answering	25968
<i>Chris Dyer, Will Hamilton, Siva Reddy, Devendra Singh, Dani Yogatama</i>	
Fast Algorithms for L_∞ -Constrained S-rectangular Robust MDPs	25982
<i>Bahram Behzadian, Chin Pang Ho, Marek Petrik</i>	
Instance-Optimal Mean Estimation Under Differential Privacy.....	25993
<i>Ziyue Huang, Yuting Liang, Ke Yi</i>	
Look at the Variance! Efficient Black-Box Explanations with Sobol-Based Sensitivity Analysis	26005
<i>Remi Cadene, Mathieu Chalvidal, Matthieu Cord, Thomas Fel, Thomas Serre, David Vigouroux</i>	
PatchGame: Learning to Signal Mid-Level Patches in Referential Games	26015
<i>Anubhav Anubhav, Kamal Gupta, Vinoj Yasanga Jayasundara Magalle Hewa, Abhinav Shrivastava, Gowthami Somepalli, Matthias Zwicker</i>	
Implicit Generative Copulas.....	26028
<i>Mohamed Ghanmi, Tim Janke, Florian Steinke</i>	
Tensor Normal Training for Deep Learning Models.....	26040
<i>Donald Goldfarb, Yi Ren</i>	
Unintended Selection: Persistent Qualification Rate Disparities and Interventions	26053
<i>Yang Liu, Reilly Raab</i>	
Revisiting 3D Object Detection from an Egocentric Perspective.....	26066
<i>Dragomir Anguelov, Boyang Deng, Thomas Funkhouser, Mahyar Najibi, Charles R Qi, Yin Zhou</i>	
Optimizing Information-Theoretical Generalization Bound Via Anisotropic Noise of SGLD	26080
<i>Wei Chen, Tie-Yan Liu, Qi Meng, Bohan Wang, Huishuai Zhang, Jieyu Zhang</i>	
Addressing Algorithmic Disparity and Performance Inconsistency in Federated Learning.....	26091
<i>Sen Cui, Jian Liang, Weishen Pan, Fei Wang, Changshui Zhang</i>	

A Mathematical Framework for Quantifying Transferability in Multi-Source Transfer Learning	26103
<i>Shao-Lun Huang, Xinyi Tong, Xiangxiang Xu, Lizhong Zheng</i>	
Morie Attack (MA): a New Potential Risk of Screen Photos	26117
<i>Ruohao Guo, Dantong Niu, Yisen Wang</i>	
Fast Bayesian Inference for Gaussian Cox Processes Via Path Integral Formulation	26130
<i>Hideaki Kim</i>	
Lattice Partition Recovery with Dyadic CART	26143
<i>Oscar Hernan Madrid Padilla, Alessandro Rinaldo, Yi Yu</i>	
Robust Deep Reinforcement Learning Through Adversarial Loss	26156
<i>Luca Daniel, Alexandre Megretski, Tuomas Oikarinen, Tsui-Wei Weng, Wang Zhang</i>	
Provable Model-Based Nonlinear Bandit and Reinforcement Learning: Shelve Optimism, Embrace Virtual Curvature	26168
<i>Kefan Dong, Tengyu Ma, Jiaqi Yang</i>	
You Only Look at One Sequence: Rethinking Transformer in Vision Through Object Detection	26183
<i>Jiemin Fang, Yuxin Fang, Bencheng Liao, Wenyu Liu, Jianwei Niu, Jiyang Qi, Xinggang Wang, Rui Wu</i>	
Learning to Delegate for Large-Scale Vehicle Routing	26198
<i>Sirui Li, Cathy Wu, Zhongxia Yan</i>	
Effective Meta-Regularization by Kernelized Proximal Regularization	26212
<i>Weisen Jiang, James Kwok, Yu Zhang</i>	
Towards Context-Agnostic Learning Using Synthetic Data	26223
<i>Charles Jin, Martin Rinard</i>	
Minimax Optimal Quantile and Semi-Adversarial Regret Via Root-Logarithmic Regularizers	26237
<i>Blair Bilodeau, Nicolò Campolongo, Jeffrey Negrea, Francesco Orabona, Dan Roy</i>	
Gradient-Free Adversarial Training Against Image Corruption for Learning-based Steering	26250
<i>Tom Goldstein, Weizi Li, Ming Lin, Yu Shen, Manli Shu, Laura Zheng</i>	
Deep Proxy Causal Learning and Its Application to Confounded Bandit Policy Evaluation	26264
<i>Arthur Gretton, Heishiro Kanagawa, Liyuan Xu</i>	
Certifying Robustness to Programmable Data Bias in Decision Trees	26276
<i>Aws Albarghouthi, Loris D'Antoni, Anna Meyer</i>	
TöRF: Time-Of-Flight Radiance Fields for Dynamic Scene View Synthesis	26289
<i>Benjamin Attal, Aaron Gokaslan, Changil Kim, Eliot Laidlaw, Matthew O'Toole, Christian Richardt, James Tompkin</i>	
Sequence-To-Sequence Learning with Latent Neural Grammars	26302
<i>Yoon Kim</i>	
Exploration-Exploitation in Multi-Agent Competition: Convergence with Bounded Rationality	26318
<i>Stefanos Leonardos, Georgios Piliouras, Kelly Spendlove</i>	
Low-Rank Extragradient Method for Nonsmooth and Low-Rank Matrix Optimization Problems	26332
<i>Dan Garber, Atara Kaplan</i>	

Which Mutual-Information Representation Learning Objectives Are Sufficient for Control? 26345
Carlos Florensa, Abhishek Gupta, Sergey Levine, Kate Rakelly

VOLUME 32

A Geometric Perspective Towards Neural Calibration Via Sensitivity Decomposition 26358
Yen-Chang Hsu, Zsolt Kira, Junjiao Tian, Dylan Yung

Towards a Unified Information-Theoretic Framework for Generalization 26370
Gintare Karolina Dziugaite, Mahdi Haghifam, Shay Moran, Dan Roy

Bayesian Decision-Making Under Misspecified Priors with Applications to Meta-learning 26382
Miro Dudik, Daniel Hsu, Akshay Krishnamurthy, Thodoris Lykouris, Robert Schapire, Max Simchowitz, Christopher Tosh

Neural Trees for Learning on Graphs 26395
Luca Carlone, Siyi Hu, Lisa Peng, Rajat Talak

Enabling Fast Differentially Private SGD Via Just-In-Time Compilation and Vectorization 26409
Gautam Kamath, Pranav Subramani, Nicholas Vadivelu

The Effectiveness of Feature Attribution Methods and Its Correlation with Automatic Evaluation Scores 26422
Daeyoung Kim, Anh Nguyen, Giang Nguyen

Coordinated Proximal Policy Optimization 26437
Zifan Wu, Deheng Ye, Chao Yu, Junge Zhang, Hankz Hankui Zhuo, Haiyin Piao

Unbiased Classification Through Bias-Contrastive and Bias-Balanced Learning 26449
Youngkyu Hong, Eunho Yang

Learning from Inside: Self-Driven Siamese Sampling and Reasoning for Video Question Answering 26462
Nan Duan, Lei Ji, Mengfei Li, Lijun Wu, Nong Xiao, Weijiang Yu, Haoteng Zheng

Identification and Estimation of Joint Probabilities of Potential Outcomes in Observational Studies with Covariate Information 26475
Ryusei Shingaki, Manabu Kuroki

Online False Discovery Rate Control for Anomaly Detection in Time Series 26487
Laurent Callot, Tim Januschowski, Baris Kurt, Quentin Rebjock

Pragmatic Image Compression for Human-In-the-Loop Decision-Making 26499
Anca Dragan, Sergey Levine, Sid Reddy

Generalized Linear Bandits with Local Differential Privacy 26511
Yuxuan Han, Zhipeng Liang, Yang Wang, Jiheng Zhang

On the Algorithmic Stability of Adversarial Training 26523
Guang Cheng, Qifan Song, Yue Xing

Width-Based Lookaheads with Learnt Base Policies and Heuristics Over the Atari-2600 Benchmark 26536
Nir Lipovetzky, Stefan O'Toole, Adrian Pearce, Miquel Ramirez

Characterizing Possible Failure Modes in Physics-Informed Neural Networks.....	26548
<i>Amir Gholami, Robert Kirby, Aditi Krishnapriyan, Michael Mahoney, Shandian Zhe</i>	
Artistic Style Transfer with Internal-External Learning and Contrastive Learning.....	26561
<i>Haibo Chen, Ailin Li, Dongming Lu, Zhizhong Wang, Wei Xing, Huiming Zhang, Zhiwen Zuo, Lei Zhao</i>	
Fast Abductive Learning by Similarity-Based Consistency Optimization	26574
<i>Le-Wen Cai, Wang-Zhou Dai, Yu-Xuan Huang, Yuan Jiang, Stephen H Muggleton</i>	
To Beam Or Not to Beam: that is a Question of Cooperation for Language GANs.....	26585
<i>Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, Thomas Scialom, Jacopo Staiano</i>	
Shapley Residuals: Quantifying the Limits of the Shapley Value for Explanations	26598
<i>Sorelle Friedler, Indra Kumar, Carlos Scheidegger, Suresh Venkatasubramanian</i>	
The Elastic Lottery Ticket Hypothesis	26609
<i>Xiaohan Chen, Yu Cheng, Zhe Gan, Jingjing Liu, Shuohang Wang, Zhangyang Wang</i>	
Joint Inference for Neural Network Depth and Dropout Regularization.....	26622
<i>Mohammadmahdi Gilany, Kishan K C, Rui Li</i>	
Tractable Density Estimation on Learned Manifolds with Conformal Embedding Flows.....	26635
<i>Jesse Cresswell, Brendan Ross</i>	
The Limits of Optimal Pricing in the Dark.....	26649
<i>Quinlan Dawkins, Minbiao Han, Haifeng Xu</i>	
No RL, No Simulation: Learning to Navigate Without Navigating	26661
<i>Devendra Singh Chaplot, Abhinav Gupta, Meera Hahn, Mustafa Mukadam, James Rehg, Shubham Tulsiani</i>	
Analogous to Evolutionary Algorithm: Designing a Unified Sequence Model	26674
<i>Shuo Chen, Wenzhou Chen, Feiyue Huang, Jian Li, Yong Liu, Ying Tai, Chengjie Wang, Yabiao Wang, Chao Xu, Jiangning Zhang</i>	
Improving Compositionality of Neural Networks by Decoding Representations to Inputs	26689
<i>Stefano Ermon, Noah Goodman, Mike Wu</i>	
The Hardness Analysis of Thompson Sampling for Combinatorial Semi-Bandits with Greedy Oracle	26701
<i>Wei Chen, Fang Kong, Shuai Li, Yueran Yang</i>	
Universal Semi-Supervised Learning	26714
<i>Chen Gong, Bo Han, Zhuo Huang, Chao Xue, Jian Yang</i>	
Improving Deep Learning Interpretability by Saliency Guided Training.....	26726
<i>Hector Corrada Bravo, Soheil Feizi, Aya Abdelsalam Ismail</i>	
SurvITE: Learning Heterogeneous Treatment Effects from Time-To-Event Data.....	26740
<i>Alicia Curth, Changhee Lee, Mihaela Van Der Schaar</i>	
Optimal Rates for Nonparametric Density Estimation Under Communication Constraints	26754
<i>Jayadev Acharya, Clement Canonne, Aditya Vikram Singh, Himanshu Tyagi</i>	
Rank Overspecified Robust Matrix Recovery: Subgradient Method and Exact Recovery	26767
<i>Yudong Chen, Lijun Ding, Liwei Jiang, Qing Qu, Zhihui Zhu</i>	

Improving Computational Efficiency in Visual Reinforcement Learning Via Stored Embeddings.....	26779
<i>Pieter Abbeel, Lili Chen, Kimin Lee, Aravind Srinivas</i>	
Learning Generalized Gumbel-Max Causal Mechanisms	26792
<i>Tamir Hazan, Daniel D. Johnson, Guy Lorberbom, Chris Maddison, Daniel Tarlow</i>	
Bandit Learning with Delayed Impact of Actions	26804
<i>Chien-Ju Ho, Yang Liu, Wei Tang</i>	
A Stochastic Newton Algorithm for Distributed Convex Optimization.....	26818
<i>Brian Bullins, Kshitij Patel, Ohad Shamir, Nathan Srebro, Blake Woodworth</i>	
Are Transformers More Robust than CNNs?	26831
<i>Yutong Bai, Jieru Mei, Cihang Xie, Alan Yuille</i>	
Towards Sharper Generalization Bounds for Structured Prediction.....	26844
<i>Shaojie Li, Yong Liu</i>	
Automated Discovery of Adaptive Attacks on Adversarial Defenses	26858
<i>Pavol Bielik, Petar Tsankov, Martin Vechev, Chengyuan Yao</i>	
PolarStream: Streaming Object Detection and Segmentation with Polar Pillars.....	26871
<i>Oscar Beijbom, Qi Chen, Sourabh Vora</i>	
Representation Costs of Linear Neural Networks: Analysis and Design	26884
<i>Zhen Dai, Mina Karzand, Nathan Srebro</i>	
Teaching Via Best-Case Counterexamples in the Learning-with-Equivalence-Queries Paradigm	26897
<i>Yuxin Chen, Akash Kumar, Adish Singla</i>	
Distilling Meta Knowledge on Heterogeneous Graph for Illicit Drug Trafficker Detection on Social Media.....	26911
<i>Yiyue Qian, Yanfang (Fa Ye), Chuxu Zhang, Yiming Zhang</i>	
Curriculum Disentangled Recommendation with Noisy Multi-Feedback.....	26924
<i>Yudong Chen, Hong Chen, Xin Wang, Rui Wang, Feng Xia, Ruobing Xie, Wenwu Zhu</i>	
Interpretable Agent Communication from Scratch (with a Generic Visual Processor Emerging on the Side).....	26937
<i>Roberto Dessi, Eugene Kharitonov, Baroni Marco</i>	
MAU: a Motion-Aware Unit for Video Prediction and Beyond	26950
<i>Zheng Chang, Wen Gao, Siwei Ma, Shanshe Wang, Xiang Xinguang, Yan Ye, Xinfeng Zhang</i>	
Successor Feature Landmarks for Long-Horizon Goal-Conditioned Reinforcement Learning	26963
<i>Wilka Carvalho, Jongwook Choi, Christopher Hoang, Honglak Lee, Sungryull Sohn</i>	
Streaming Belief Propagation for Community Detection.....	26976
<i>Mohammadhossein Bateni, Filipe Miguel Goncalves De Almeida, Andre Linhares, Andrea Montanari, Ashkan Norouzi-Fard, Jakab Tardos, Yuchen Wu</i>	
The Staircase Property: How Hierarchical Structure Can Guide Deep Learning	26989
<i>Emmanuel Abbe, Enric Boix-Adsera, Matthew S Brennan, Guy Bresler, Dheeraj Nagaraj</i>	
MagNet: a Neural Network for Directed Graphs.....	27003
<i>Nathan Brugnone, Yixuan He, Matthew Hirn, Michael Perlmutter, Xitong Zhang</i>	

Hardware-Adaptive Efficient Latency Prediction for NAS Via Meta-Learning	27016
<i>Song Chong, Sung Ju Hwang, Hayeon Lee, Sewoong Lee</i>	
Topological Relational Learning on Graphs.....	27029
<i>Yuzhou Chen, Baris Coskunuzer, Yulia Gel</i>	
Learning Theory Can (Sometimes) Explain Generalisation in Graph Neural Networks.....	27043
<i>Leena Chennuru Vankadara, Pascal Esser, Debarghya Ghoshdastidar</i>	
Federated Linear Contextual Bandits	27057
<i>Ruiquan Huang, Cong Shen, Weiqiang Wu, Jing Yang</i>	
Least Square Calibration for Peer Reviews	27069
<i>Xiaohui Bei, Sijun Tan, Jibang Wu, Haifeng Xu</i>	
Scaling Up Exact Neural Network Compression by ReLU Stability	27081
<i>Abhinav Kumar, Srikumar Ramalingam, Thiago Serra, Xin Yu</i>	
Passive Attention in Artificial Neural Networks Predicts Human Visual Selectivity	27094
<i>Ishita Dasgupta, Erin Grant, Tom Griffiths, Nori Jacoby, Thomas Langlois, Haicheng Zhao</i>	
GRIN: Generative Relation and Intention Network for Multi-Agent Trajectory Prediction.....	27107
<i>Tong He, Longyuan Li, Li Wenliang, David Wipf, Tianjun Xiao, Junchi Yan, Jian Yao, Zheng Zhang</i>	
Instance-Dependent Partial Label Learning	27119
<i>Xin Geng, Congyu Qiao, Ning Xu, Min-Ling Zhang</i>	
Deep Learning with Label Differential Privacy	27131
<i>Badi Ghazi, Noah Golowich, Ravi Kumar, Pasin Manurangsi, Chiyuan Zhang</i>	
Semialgebraic Representation of Monotone Deep Equilibrium Models and Applications to Certification.....	27146
<i>Tong Chen, Jean Lasserre, Victor Magron, Edouard Pauwels</i>	
The Role of Global Labels in Few-Shot Classification and How to Infer Them.....	27160
<i>Carlo Ciliberto, Massimiliano Pontil, Ruohan Wang</i>	
NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-View Reconstruction.....	27171
<i>Taku Komura, Lingjie Liu, Yuan Liu, Christian Theobalt, Peng Wang, Wenping Wang</i>	
Improved Guarantees for Offline Stochastic Matching Via New Ordered Contention Resolution Schemes.....	27184
<i>Brian Brubach, Nathaniel Grammel, Will Ma, Aravind Srinivasan</i>	
UFC-BERT: Unifying Multi-Modal Controls for Conditional Image Synthesis.....	27196
<i>Ming Ding, Zhikang Li, Jianxin Ma, Rui Men, Jie Tang, Hongxia Yang, Zhu Zhang, Chang Zhou, Jingren Zhou</i>	
Is Bang-Bang Control All You Need? Solving Continuous Control with Bernoulli Policies	27209
<i>Igor Gilitschenski, Martin Riedmiller, Daniela Rus, Wilko Schwarting, Tim Seyde, Bartolomeo Stellato, Markus Wulfmeier</i>	

VOLUME 33

Improving Generalization in Meta-RL with Imaginary Tasks from Latent Dynamics Mixture.....	27222
<i>Sae-Young Chung, Suyoung Lee</i>	
Localization with Sampling-Argmax	27236
<i>Tong Chen, Jiefeng Li, Yong-Lu Li, Yujing Lou, Cewu Lu, Ruiqi Shi</i>	
Improved Regularization and Robustness for Fine-Tuning in Neural Networks.....	27249
<i>Dongyue Li, Hongyang Zhang</i>	
BARTScore: Evaluating Generated Text as Text Generation	27263
<i>Pengfei Liu, Graham Neubig, Weizhe Yuan</i>	
An Analysis of Ermakov-Zolotukhin Quadrature Using Kernels.....	27278
<i>Ayoub Belhadji</i>	
Towards Understanding Why Lookahead Generalizes Better than SGD and Beyond.....	27290
<i>Jiashi Feng, Shuicheng Yan, Hanshu Yan, Xiaotong Yuan, Pan Zhou</i>	
Online Market Equilibrium with Application to Fair Division	27305
<i>Yuan Gao, Christian Kroer, Alex Peysakhovich</i>	
Dynamic Resolution Network	27319
<i>Kai Han, Zhenzhong Lan, Ying Nie, Yunhe Wang, Enhua Wu, Qiulin Zhang, Mingjian Zhu</i>	
Gauge Equivariant Transformer	27331
<i>Yiming Dong, Lingshen He, Zhouchen Lin, Dacheng Tao, Yisen Wang</i>	
Unsupervised Object-Based Transition Models for 3D Partially Observable Environments.....	27344
<i>Chris Burgess, Antonia Creswell, Rishabh Kabra, Murray Shanahan</i>	
Robust Contrastive Learning Using Negative Samples with Diminished Semantics.....	27356
<i>Songwei Ge, David Jacobs, Chun-Liang Li, Shlok Mishra, Haohan Wang</i>	
General Low-Rank Matrix Optimization: Geometric Analysis and Sharper Bounds.....	27369
<i>Yingjie Bi, Javad Lavaei, Haixiang Zhang</i>	
Flow Network Based Generative Models for Non-Iterative Diverse Candidate Generation.....	27381
<i>Yoshua Bengio, Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup</i>	
Policy Finetuning: Bridging Sample-Efficient Offline and Online Reinforcement Learning	27395
<i>Yu Bai, Nan Jiang, Huan Wang, Tengyang Xie, Caiming Xiong</i>	
Reducing Information Bottleneck for Weakly Supervised Semantic Segmentation	27408
<i>Jooyoung Choi, Jungbeom Lee, Jisoo Mok, Sungroh Yoon</i>	
SGD: the Role of Implicit Regularization, Batch-Size and Multiple-epochs	27422
<i>Satyen Kale, Ayush Sekhari, Karthik Sridharan</i>	
AC-GC: Lossy Activation Compression with Guaranteed Convergence	27434
<i>Tor Aamodt, R David Evans</i>	
Label Noise SGD Provably Prefers Flat Global Minimizers.....	27449
<i>Alex Damian, Jason Lee, Tengyu Ma</i>	

Can We Have it All? on the Trade-Off Between Spatial and Adversarial Robustness of Neural Networks	27462
<i>Amit Deshpande, Sandesh Kamath, Subrahmanyam Kambhampati Venkata, Vineeth N Balasubramanian</i>	
Universal Off-Policy Evaluation	27475
<i>Emma Brunskill, Yash Chandak, Erik Learned-Miller, Scott Niekum, Philip Thomas, Bruno Da Silva</i>	
A Non-Commutative Extension of Lee-Seung's Algorithm for Positive Semidefinite Factorizations	27491
<i>Yong Sheng Soh, Antonios Varvitsiotis</i>	
Efficiently Identifying Task Groupings for Multi-Task Learning	27503
<i>Ehsan Amid, Rohan Anil, Chris Fifty, Chelsea Finn, Tianhe Yu, Zhe Zhao</i>	
Instance-Conditioned GAN	27517
<i>Marlene Careil, Arantxa Casanova, Michal Drozdal, Adriana Romero Soriano, Jakob Verbeek</i>	
DeepSITH: Efficient Learning Via Decomposition of What and When Across Time Scales	27530
<i>Marc Howard, Brandon Jacques, Per B Sederberg, Zoran Tiganj</i>	
A Gaussian Process-Bayesian Bernoulli Mixture Model for Multi-Label Active Learning	27542
<i>Weishi Shi, Dayou Yu, Qi Yu</i>	
Differentially Private Empirical Risk Minimization Under the Fairness Lens	27555
<i>My Dinh, Ferdinando Fioretto, Cuong Tran</i>	
A Unified View of cGANs with and Without Classifiers	27566
<i>Si-An Chen, Chun-Liang Li, Hsuan-Tien Lin</i>	
Online and Offline Reinforcement Learning by Planning with a Learned Model	27580
<i>Ioannis Antonoglou, Mohammadamin Barekatin, Thomas Hubert, Amol Mandhane, Julian Schrittwieser, David Silver</i>	
Stochastic Multi-Armed Bandits with Control Variates	27592
<i>Manjesh Kumar Hanawal, Arun Verma</i>	
Near-Optimal No-Regret Learning in General Games	27604
<i>Constantinos Daskalakis, Maxwell Fishelson, Noah Golowich</i>	
Improving Self-Supervised Learning with Automated Unsupervised Outlier Arbitration	27617
<i>Jingyang Lin, Tao Mei, Yingwei Pan, Yu Wang, Ting Yao, Jingjing Zou</i>	
Improving Anytime Prediction with Parallel Cascaded Networks and a Temporal-Difference Loss	27631
<i>Samy Bengio, Michael Iuzzolino, Michael Mozer</i>	
Identifiable Generative Models for Missing Not at Random Data Imputation	27645
<i>Chao Ma, Cheng Zhang</i>	
DNN-Based Topology Optimisation: Spatial Invariance and Neural Tangent Kernel	27659
<i>Benjamin Dupuis, Arthur Jacot</i>	
Baleen: Robust Multi-Hop Reasoning at Scale Via Condensed Retrieval	27670
<i>Omar Khattab, Christopher Potts, Matei Zaharia</i>	

Local Hyper-Flow Diffusion	27683
<i>Kimon Fountoulakis, Pan Li, Shenghao Yang</i>	
Permuton-Induced Chinese Restaurant Process	27695
<i>Yasuhiro Fujiwara, Akisato Kimura, Masahiro Nakano, Takeshi Yamada, Naonori Ueda</i>	
Faster Algorithms and Constant Lower Bounds for the Worst-Case Expected Error	27709
<i>Jonah Brown-Cohen</i>	
On Learning Domain-Invariant Representations for Transfer Learning with Multiple Sources	27720
<i>Hung Bui, Trung Le, Dinh Phung, Trung Phung, Anh Tran, Toan Tran, Tung-Long Vuong</i>	
You Never Cluster Alone	27734
<i>Jie Qin, Ling Shao, Yuming Shen, Ziyi Shen, Philip Torr, Menghan Wang</i>	
Dynamic COVID Risk Assessment Accounting for Community Virus Exposure from a Spatial-Temporal Transmission Model.....	27747
<i>Yuan Chen, Wenbo Fei, Qinxia Wang, Yuanjia Wang, Donglin Zeng</i>	
Dueling Bandits with Adversarial Sleeping	27761
<i>Pierre Gaillard, Aadirupa Saha</i>	
Beware of the Simulated DAG! Causal Discovery Benchmarks May Be Easy to Game.....	27772
<i>Alexander Reisach, Christof Seiler, Sebastian Weichwald</i>	
Automated Dynamic Mechanism Design.....	27785
<i>Vincent Conitzer, Hanrui Zhang</i>	
A Generative Nonparametric Bayesian Model for Whole Genomes.....	27798
<i>Alan Amin, Debora Marks, Eli N Weinstein</i>	
Robust Predictable Control.....	27813
<i>Ben Eysenbach, Sergey Levine, Russ Salakhutdinov</i>	
Unsupervised Speech Recognition	27826
<i>Michael Auli, Alexei Baevski, Alexis Conneau, Wei-Ning Hsu</i>	
Robustness Between the Worst and Average Case	27840
<i>Anna Bair, J. Zico Kolter, Leslie Rice, Huan Zhang</i>	
Online Learning and Control of Complex Dynamical Systems from Sensory Input.....	27852
<i>Oumayma Bounou, Justin Carpentier, Jean Ponce</i>	
Self-Supervised Bug Detection and Repair	27865
<i>Miltiadis Allamanis, Marc Brockschmidt, Henry Jackson-Flux</i>	
Faster Neural Network Training with Approximate Tensor Operations	27877
<i>Menachem Adelman, Ido Hakimi, Kfir Levy, Mark Silberstein</i>	
Learning Interpretable Decision Rule Sets: a Submodular Optimization Approach	27890
<i>Hongxia Du, Kai He, Liang Sun, Bo Yang, Fan Yang, Jingbang Yang, Linxiao Yang</i>	
Spatial-Temporal Super-Resolution of Satellite Imagery Via Conditional Pixel Synthesis.....	27903
<i>Marshall Burke, Stefano Ermon, Yutong He, Nicholas Lai, David Lobell, Chenlin Meng, Dingjie Wang, William Zhang</i>	

On Memorization in Probabilistic Deep Generative Models.....	27916
<i>Chris Williams, Gerrit Van Den Burg</i>	
You Are the Best Reviewer of Your Own Papers: An Owner-Assisted Scoring Mechanism.....	27929
<i>Weijie Su</i>	
Garment4D: Garment Reconstruction from Point Cloud Sequences.....	27940
<i>Zhongang Cai, Fangzhou Hong, Ziwei Liu, Liang Pan</i>	
Fast Policy Extragradient Methods for Competitive Games with Entropy Regularization	27952
<i>Shicong Cen, Yuejie Chi, Yuting Wei</i>	
Shift-Robust GNNs: Overcoming the Limitations of Localized Graph Training Data	27965
<i>Jiawei Han, Bryan Perozzi, Natalia Ponomareva, Qi Zhu</i>	
RIM: Reliable Influence-Based Active Learning on Graphs.....	27978
<i>Bin Cui, Meng Cao, Ping Huang, Jiulong Shan, Yexin Wang, Zhi Yang, Zhenbang You, Wentao Zhang</i>	
Dynamical Wasserstein Barycenters for Time-Series Modeling.....	27991
<i>Shuchin Aeron, Kevin Cheng, Michael Hughes, Eric L Miller</i>	
RelaySum for Decentralized Deep Learning on Heterogeneous Data.....	28004
<i>Lie He, Martin Jaggi, Sai Praneeth Karimireddy, Anastasiia Koloskova, Tao Lin, Sebastian Stich, Thijs Vogels</i>	
Transformers Generalize DeepSets and Can Be Extended to Graphs & Hypergraphs.....	28016
<i>Seunghoon Hong, Jinwoo Kim, Saeyoon Oh</i>	
No Regrets for Learning the Prior in Bandits	28029
<i>Soumya Basu, Branislav Kveton, Csaba Szepesvari, Manzil Zaheer</i>	
Encoding Robustness to Image Style Via Adversarial Feature Perturbations	28042
<i>Micah Goldblum, Tom Goldstein, Manli Shu, Zuxuan Wu</i>	
Continuized Accelerations of Deterministic and Stochastic Gradient Descents, and of Gossip Algorithms.....	28054
<i>Francis Bach, Raphaël Berthier, Mathieu Even, Nicolas Flammarion, Pierre Gaillard, Hadrien Hendrikx, Laurent Massoulié, Adrien Taylor</i>	
Natural Continual Learning: Success is a Journey, Not (just) a Destination	28067
<i>Alberto Bernacchia, Guillaume Hennequin, Kristopher Jensen, Ta-Chu Kao, Gido Van De Ven</i>	

VOLUME 34

Individual Privacy Accounting Via a Renyi Filter	28080
<i>Vitaly Feldman, Tijana Zrnica</i>	
Post-Training Quantization for Vision Transformer	28092
<i>Wen Gao, Kai Han, Zhenhua Liu, Siwei Ma, Yunhe Wang, Wei Zhang</i>	
Unsupervised Part Discovery from Contrastive Reconstruction	28104
<i>Subhabrata Choudhury, Iro Laina, Christian Rupprecht, Andrea Vedaldi</i>	

ASSANet: an Anisotropic Separable Set Abstraction for Efficient Point Cloud Representation Learning	28119
<i>Bernard Ghanem, Hasan Hammoud, Guohao Li, Guocheng Qian, Ali Thabet</i>	
An Empirical Investigation of Domain Generalization with Empirical Risk Minimizers	28131
<i>David Lopez-Paz, David Schwab, Ramakrishna Vedantam</i>	
Fair Sequential Selection Using Supervised Learning Models.....	28144
<i>Mahed Abroshan, Mohammad Mahdi Khalili, Xueru Zhang</i>	
Towards Sample-Efficient Overparameterized Meta-learning	28156
<i>Maryam Fazel, Ibrahim Gulluk, Adhyyan Narang, Samet Oymak, Yue Sun</i>	
ScaleCert: Scalable Certified Defense Against Adversarial Patches with Sparse Superficial Layers	28169
<i>Xiaobing Chen, Yunji Chen, Zidong Du, Qi Guo, Husheng Han, Xing Hu, Ling Liang, Yanzhi Wang, Kaidi Xu</i>	
Towards Mental Time Travel: a Hierarchical Memory for Reinforcement Learning Agents	28182
<i>Andrea Banino, Stephanie Chan, Felix Hill, Andrew Lampinen</i>	
Beyond Tikhonov: Faster Learning with Self-Concordant Losses, Via Iterative Regularization.....	28196
<i>Gaspard Beugnot, Julien Mairal, Alessandro Rudi</i>	
Variational Bayesian Reinforcement Learning with Regret Bounds.....	28208
<i>Brendan O'Donoghue</i>	
Logarithmic Regret from Sublinear Hints	28222
<i>Aditya Bhaskara, Ashok Cutkosky, Ravi Kumar, Manish Purohit</i>	
Independent Mechanism Analysis, a New Concept?.....	28233
<i>Michel Besserve, Luigi Gresele, Bernhard Scholkopf, Vincent Stimper, Julius Von Kugelgen</i>	
Momentum Centering and Asynchronous Update for Adaptive Gradient Methods.....	28249
<i>Yifan Ding, James Duncan, Nicha Dvornek, Tommy Tang, Sekhar C Tatikonda, Juntang Zhuang</i>	
Robustness Via Uncertainty-Aware Cycle Consistency	28261
<i>Zeynep Akata, Yanbei Chen, Uddeshya Upadhyay</i>	
CBP: Backpropagation with Constraint on Weight Precision Using a pseudo-Lagrange Multiplier Method	28274
<i>Doo Seok Jeong, Guhyun Kim</i>	
On the Sample Complexity of Privately Learning Axis-Aligned Rectangles.....	28286
<i>Menachem Sadigurschi, Uri Stemmer</i>	
Implicit Sparse Regularization: the Impact of Depth and Early Stopping	28298
<i>Chinmay Hegde, Jiangyuan Li, Thanh Nguyen, Ka Wai Wong</i>	
Efficient Generalization with Distributionally Robust Learning	28310
<i>Soumyadip Ghosh, Mark Squillante, Ebisa Wollega</i>	
No-Regret Online Learning Over Riemannian Manifolds.....	28323
<i>Yiguang Hong, Guodong Shi, Zhipeng Tu, Xi Wang, Yingyi Wu</i>	
Landmark-Guided Subgoal Generation in Hierarchical Reinforcement Learning	28336
<i>Junsu Kim, Younggyo Seo, Jinwoo Shin</i>	

Minimax Regret for Stochastic Shortest Path.....	28350
<i>Alon Cohen, Yonathan Efroni, Yishay Mansour, Aviv Rosenberg</i>	
Parametrized Quantum Policies for Reinforcement Learning	28362
<i>Hans Briegel, Vedran Dunjko, Casper Gyurik, Sofiene Jerbi, Simon Marshall</i>	
On Pathologies in KL-Regularized Reinforcement Learning from Expert Demonstrations	28376
<i>Yarin Gal, Cong Lu, Michael A Osborne, Tim G. J. Rudner, Yee Teh</i>	
Conditional Generation Using Polynomial Expansions	28390
<i>Grigorios Chrysos, Markos Georgopoulos, Yannis Panagakis</i>	
Efficient Constrained Sampling Via the mirror-Langevin Algorithm	28405
<i>Kwangjun Ahn, Sinho Chewi</i>	
Adaptive Online Packing-Guided Search for POMDPs	28419
<i>Jianye Hao, Dong Li, Wulong Liu, Chenyang Wu, Guoyu Yang, Yang Yu, Zongzhang Zhang</i>	
Turing Completeness of Bounded-Precision Recurrent Neural Networks	28431
<i>Stephen Chung, Hava Siegelmann</i>	
End-To-end Multi-modal Video Temporal Grounding	28442
<i>Yi-Wen Chen, Yi-Hsuan Tsai, Ming-Hsuan Yang</i>	
How Powerful Are Performance Predictors in Neural Architecture Search?	28454
<i>Frank Hutter, Yang Liu, Robin Ru, Colin White, Arber Zela</i>	
Stylized Dialogue Generation with Multi-Pass Dual Learning	28470
<i>Jinpeng Li, Tie-Yan Liu, Hongda Sun, Yingce Xia, Rui Yan, Dongyan Zhao</i>	
Entropy-Based Adaptive Hamiltonian Monte Carlo	28482
<i>Petros Dellaportas, Marcel Hirt, Michalis Titsias</i>	
Continual World: a Robotic Benchmark for Continual Reinforcement Learning	28496
<i>Lukasz Kucinski, Piotr Milos, Razvan Pascanu, Maciej Wolczyk, Michal Zajac</i>	
Towards Best-Of-All-Worlds Online Learning with Feedback Graphs	28511
<i>Liad Erez, Tomer Koren</i>	
ViTAE: Vision Transformer Advanced by Exploring Intrinsic Inductive Bias	28522
<i>Dacheng Tao, Yufei Xu, Qiming Zhang, Jing Zhang</i>	
Open Rule Induction.....	28536
<i>Xingran Chen, Wanyun Cui</i>	
Post-Contextual-Bandit Inference	28548
<i>Aurelien Bibaut, Antoine Chambaz, Maria Dimakopoulou, Nathan Kallus, Mark Van Der Laan</i>	
Revisiting Discriminator in GAN Compression: a Generator-Discriminator Cooperative Compression Scheme	28560
<i>Fei Chao, Rongrong Ji, Shaojie Li, Xudong Mao, Jie Wu, Xuefeng Xiao</i>	
Asymptotically Exact Error Characterization of Offline Policy Evaluation with Misspecified Linear Models.....	28573
<i>Kohei Miyaguchi</i>	

Topographic VAEs Learn Equivariant Capsules	28585
<i>T. Anderson Keller, Max Welling</i>	
MobILE: Model-Based Imitation Learning from Observation Alone	28598
<i>Jonathan Chang, Rahul Kidambi, Wen Sun</i>	
Few-Round Learning for Federated Learning	28612
<i>Dong-Jun Han, Do-Yeon Kim, Jaekyun Moon, Younghyun Park, Jun Seo</i>	
On Path Integration of Grid Cells: Group Representation and Isotropic Scaling	28623
<i>Ruiqi Gao, Xue-Xin Wei, Ying Nian Wu, Jianwen Xie, Song-Chun Zhu</i>	
Online Convex Optimization with Continuous Switching Constraint	28636
<i>Yuanyu Wan, Guanghui Wang, Tianbao Yang, Lijun Zhang</i>	
Why Do Better Loss Functions Lead to Less Transferable Features?	28648
<i>Ting Chen, Simon Kornblith, Honglak Lee, Mohammad Norouzi</i>	
Breaking the Centralized Barrier for Cross-Device Federated Learning	28663
<i>Martin Jaggi, Satyen Kale, Sai Praneeth Karimireddy, Mehryar Mohri, Sashank Reddi, Sebastian Stich, Ananda Theertha Suresh</i>	
Adversarially Robust Learning for Security-Constrained Optimal Power Flow	28677
<i>Aayushya Agarwal, Neeraj Vijay Bedmutha, Priya Donti, J. Zico Kolter, Larry Pileggi</i>	
Learning a Single Neuron with Bias Using Gradient Descent	28690
<i>Ohad Shamir, Gal Vardi, Gilad Yehudai</i>	
Making a (Counterfactual) Difference One Rationale at a Time	28701
<i>Min Chi, Michael Green, Mitchell Plyler</i>	
3D Siamese Voxel-To-BEV Tracker for Sparse Point Clouds	28714
<i>Mingmei Cheng, Le Hui, Lingpeng Wang, Jin Xie, Jian Yang</i>	
Stateful Strategic Regression	28728
<i>Keegan Harris, Hoda Heidari, Steven Wu</i>	
Self-Attention Between Datapoints: Going Beyond Individual Input-Output Pairs in Deep Learning	28742
<i>Neil Band, Yarin Gal, Aidan Gomez, Jannik Kossen, Clare Lyle, Thomas Rainforth</i>	
Your Head is There to Move You Around: Goal-Driven Models of the Primate Dorsal Pathway	28757
<i>Shahab Bakhtiari, Patrick Mineault, Christopher Pack, Blake Richards</i>	
Achieving Rotational Invariance with Bessel-Convolutional Neural Networks	28772
<i>Adrien Bibal, Valentin Delchevalerie, Benoît Frenay, Alexandre Mayer</i>	
Unsupervised Domain Adaptation with Dynamics-Aware Rewards in Reinforcement Learning	28784
<i>Yachen Kang, Jinxin Liu, Hao Shen, Qiangxing Tian, Donglin Wang</i>	
GraphFormers: GNN-Nested Transformers for Representation Learning on Textual Graph	28798
<i>Sanjay Agrawal, Chaozhuo Li, Defu Lian, Zheng Liu, Amit Singh, Guangzhong Sun, Shitao Xiao, Xing Xie, Junhan Yang</i>	
A Universal Law of Robustness Via Isoperimetry	28811
<i>Sebastien Bubeck, Mark Sellke</i>	

On Contrastive Representations of Stochastic Processes	28823
<i>Adam Foster, Emile Mathieu, Yee Teh</i>	
A Domain-Shrinking Based Bayesian Optimization Algorithm with Order-Optimal Regret Performance.....	28836
<i>Sudeep Salgia, Sattar Vakili, Qing Zhao</i>	
Scalars Are Universal: Equivariant Machine Learning, Structured Like Classical Physics	28848
<i>Ben Blum-Smith, David W Hogg, Kate Storey-Fisher, Soledad Villar, Weichi Yao</i>	
Unsupervised Object-Level Representation Learning from Scene Images	28864
<i>Ziwei Liu, Chen Change Loy, Yew Soon Ong, Jiahao Xie, Xiaohang Zhan</i>	
Do Transformers Really Perform Badly for Graph Representation?.....	28877
<i>Tianle Cai, Di He, Guolin Ke, Tie-Yan Liu, Shengjie Luo, Yanming Shen, Chengxuan Ying, Shuxin Zheng</i>	
Powerpropagation: a Sparsity Inducing Weight Reparameterisation	28889
<i>Siddhant Jayakumar, Peter E Latham, Razvan Pascanu, Jonathan Schwarz, Yee Teh</i>	
Stronger NAS with Weaker Predictors.....	28904
<i>Dongdong Chen, Mei Chen, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Zicheng Liu, Zhangyang Wang, Junru Wu, Ye Yu, Lu Yuan</i>	
Convolutional Normalization: Improving Deep Convolutional Network Robustness and Training	28919
<i>Carlos Fernandez-Granda, Xiao Li, Sheng Liu, Qing Qu, Chong You, Yuexiang Zhai, Zhihui Zhu</i>	
Nearly-Tight and Oblivious Algorithms for Explainable Clustering.....	28929
<i>Buddhima Gamlath, Xinrui Jia, Adam Polak, Ola Svensson</i>	

VOLUME 35

Deep Networks Provably Classify Data on Curves	28940
<i>Sam Buchanan, Dar Gilboa, Tingran Wang, John Wright</i>	
COMBO: Conservative Offline Model-Based Policy Optimization	28954
<i>Chelsea Finn, Aviral Kumar, Sergey Levine, Rafael Rafailov, Aravind Rajeswaran, Tianhe Yu</i>	
Time-Series Generation by Contrastive Imitation	28968
<i>Ioana Bica, Daniel Jarrett, Mihaela Van Der Schaar</i>	
Differentially Private Sampling from Distributions	28983
<i>Sofya Raskhodnikova, Satchit Sivakumar, Adam Smith, Marika Swanberg</i>	
On the Expected Complexity of Maxout Networks.....	28995
<i>Guido Montufar, Hanna Tseran</i>	
Cross-View Geo-localization with Layer-to-Layer Transformer	29009
<i>Xiufan Lu, Hongji Yang, Yingying Zhu</i>	
TAAC: Temporally Abstract Actor-Critic for Continuous Control.....	29021
<i>Wei Xu, Haonan Yu, Haichao Zhang</i>	
Learning Robust Hierarchical Patterns of Human Brain Across Many fMRI Studies	29034
<i>Christos Davatzikos, Dushyant Sahoo</i>	

Global Convergence to Local Minmax Equilibrium in Classes of Nonconvex Zero-Sum Games.....	29049
<i>Evan Faulkner, Tanner Fiez, Eric Mazumdar, Adhyyan Narang, Lillian Ratliff</i>	
Bandit Quickest Change-point Detection	29064
<i>Aditya Gopalan, Braghadeesh Lakshminarayanan, Venkatesh Saligrama</i>	
Can Multi-Label Classification Networks Know What They Don't Know?	29074
<i>Alex Bocchieri, Yixuan Li, Weitang Liu, Haoran Wang</i>	
Balanced Chamfer Distance as a Comprehensive Metric for Point Cloud Completion	29088
<i>Dahua Lin, Ziwei Liu, Liang Pan, Tai Wang, Tong Wu, Junzhe Zhang</i>	
Optimal Gradient-Based Algorithms for Non-concave Bandit Optimization.....	29101
<i>Baihe Huang, Kaixuan Huang, Sham Kakade, Jason Lee, Qi Lei, Runzhe Wang, Jiaqi Yang</i>	
On Optimal Interpolation in Linear Regression	29116
<i>Eduard Oravkin, Patrick Rebeschini</i>	
Differentiable Optimization of Generalized Nondecomposable Functions Using Linear Programs.....	29129
<i>Zihang Meng, Lopamudra Mukherjee, Sathya Ravi, Vikas Singh, Yichao Wu</i>	
Towards Understanding Cooperative Multi-Agent Q-Learning with Value Factorization	29142
<i>Beining Han, Zhizhou Ren, Jianhao Wang, Jianing Ye, Chongjie Zhang</i>	
Margin-Independent Online Multiclass Learning Via Convex Geometry.....	29156
<i>Guru Guruganesh, Allen Liu, Jon Schneider, Joshua Wang</i>	
STEP: Out-Of-Distribution Detection in the Presence of Limited In-Distribution Labeled Data	29168
<i>Zhanzhan Cheng, Lan-Zhe Guo, Yu-Feng Li, Shiliang Pu, Zhi Zhou</i>	
Rényi Differential Privacy of the Subsampled Shuffle Model in Distributed Learning.....	29181
<i>Deepesh Data, Suhas Diggavi, Antonious Girgis</i>	
Gradient-Based Editing of Memory Examples for Online Task-free Continual Learning	29193
<i>Junyi Du, Xisen Jin, Xiang Ren, Arka Sadhu</i>	
Tailoring: Encoding Inductive Biases by Optimizing Unsupervised Objectives at Prediction Time	29206
<i>Ferran Alet, Maria Bauza, Leslie Kaelbling, Kenji Kawaguchi, Nurullah Giray Kuru, Tomas Lozano-Perez</i>	
Implicit Bias of SGD for Diagonal Linear Networks: a Provable Benefit of Stochasticity.....	29218
<i>Nicolas Flammarion, Scott Pesme, Loucas Pillaud-Vivien</i>	
Iterative Teacher-Aware Learning.....	29231
<i>Jeffrey L Chen, Jingdong Gao, Quanquan Gu, Junhong Shen, Ying Nian Wu, Luyao Yuan, Dongruo Zhou, Song-Chun Zhu</i>	
Clockwork Variational Autoencoders	29246
<i>Jimmy Ba, Danijar Hafner, Vaibhav Saxena</i>	
How Does it Sound?.....	29258
<i>Xiulong Liu, Eli Shlizerman, Kun Su</i>	
Stabilizing Dynamical Systems Via Policy Gradient Methods	29274
<i>Juan Perdomo, Max Simchowitz, Jack Umenberger</i>	

Language Models Enable Zero-Shot Prediction of the Effects of Mutations on Protein Function.....	29287
<i>Jason Liu, Joshua Meier, Roshan Rao, Alex Rives, Tom Sercu, Robert Verkuil</i>	
Deep Reinforcement Learning at the Edge of the Statistical Precipice	29304
<i>Rishabh Agarwal, Marc Bellemare, Pablo Samuel Castro, Aaron Courville, Max Schwarzer</i>	
DRONE: Data-Aware Low-rank Compression for Large NLP Models.....	29321
<i>Patrick Chen, Inderjit Dhillon, Cho-Jui Hsieh, Hsiang-Fu Yu</i>	
DSelect-K: Differentiable Selection in the Mixture of Experts with Applications to Multi-Task Learning	29335
<i>Yihua Chen, Ed Chi, Aakanksha Chowdhery, Hussein Hazimeh, Lichan Hong, Rahul Mazumder, Maheswaran Sathiamoorthy, Zhe Zhao</i>	
Mind the Gap: Assessing Temporal Generalization in Neural Language Models.....	29348
<i>Devang Agrawal, Phil Blunsom, Kris Cao, Mai Gimenez, Elena Gribovskaya, Tomas Kocisky, Adhi Kuncoro, Angeliki Lazaridou, Adam Liska, Sebastian Ruder, Tayfun Terzi, Dani Yogatama, Susannah Young, Cyprien De Masson D'Autume</i>	
Heavy Tails in SGD and Compressibility of Overparametrized Neural Networks	29364
<i>Melih Barsbey, Murat Erdogdu, Gaël Richard, Milad Sefidgaran, Umut Simsekli</i>	
Targeted Neural Dynamical Modeling	29379
<i>Matthias Hennig, Cole Hurwitz, Justin Jude, Lee Miller, Matthew Perich, Akash Srivastava, Kai Xu</i>	
Exploiting the Intrinsic Neighborhood Structure for Source-Free Domain Adaptation	29393
<i>Luis Herranz, Shangling Jui, Shiqi Yang, Joost Van De Weijer, Yaxing Wang</i>	
Learning with Noisy Correspondence for Cross-Modal Matching.....	29406
<i>Wenbiao Ding, Zhenyu Huang, Xiao Liu, Guocheng Niu, Xi Peng, Hua Wu, Xinyan Xiao</i>	
Offline Reinforcement Learning with Reverse Model-Based Imagination	29420
<i>Haozhe Jiang, Siyuan Li, Wenzhe Li, Jianhao Wang, Chongjie Zhang, Guangxiang Zhu</i>	
Parameter Prediction for Unseen Deep Architectures	29433
<i>Michal Drozdal, Boris Knyazev, Adriana Romero Soriano, Graham Taylor</i>	
FMMformer: Efficient and Flexible Transformer Via Decomposed Near-Field and Far-Field Attention	29449
<i>Long Chen, Tan Nguyen, Stanley Osher, Vai Suliafu, Bao Wang</i>	
Square Root Principal Component Pursuit: Tuning-Free Noisy Robust Matrix Recovery	29464
<i>John Wright, Jingkai Yan, Junhui Zhang</i>	
Neural Bellman-Ford Networks: a General Graph Neural Network Framework for Link Prediction.....	29476
<i>Jian Tang, Louis-Pascal Xhonneux, Zuobai Zhang, Zhaocheng Zhu</i>	
CorticalFlow: a Diffeomorphic Mesh Transformer Network for Cortical Surface Reconstruction	29491
<i>Pierrick Bourgeat, Clinton Fookes, Jurgen Fripp, Darren Fu, Leo Lebrat, Olivier Salvado, Rodrigo Santa Cruz, Frederic De Gournay</i>	
Bridging the Gap Between Practice and PAC-Bayes Theory in Few-Shot Meta-Learning	29506
<i>Xi Chen, Nan Ding, Sebastian Goodman, Tomer Levinboim, Radu Soricut</i>	
SLOE: a Faster Method for Statistical Inference in High-Dimensional Logistic Regression.....	29517
<i>Alexander D'Amour, Cory Y McLean, Steve Yadlowsky, Taedong Yun</i>	

ELLA: Exploration Through Learned Language Abstraction	29529
<i>Siddharth Karamcheti, Suvir Mirchandani, Dorsa Sadigh</i>	
Learning Distilled Collaboration Graph for Multi-Agent Perception.....	29541
<i>Siheng Chen, Chen Feng, Yiming Li, Shunli Ren, Pengxiang Wu, Wenjun Zhang</i>	
Federated-EM with Heterogeneity Mitigation and Variance Reduction	29553
<i>Aymeric Dieuleveut, Gersende Fort, Eric Moulines, Geneviève Robin</i>	
On the Role of Optimization in Double Descent: a Least Squares Study.....	29567
<i>Ilja Kuzborskij, Razvan Pascanu, Amal Rannen-Triki, Omar Rivasplata, Csaba Szepesvari</i>	
Neural Architecture Dilation for Adversarial Robustness	29578
<i>Yanxi Li, Yunhe Wang, Chang Xu, Zhaohui Yang</i>	
Clustering Effect of Adversarial Robust Models.....	29590
<i>Yang Bai, Yong Jiang, Yisen Wang, Shu-Tao Xia, Xin Yan</i>	
On the Cryptographic Hardness of Learning Single Periodic Neurons	29602
<i>Joan Bruna, Min Jae Song, Ilias Zadik</i>	
PCA Initialization for Approximate Message Passing in Rotationally Invariant Models	29616
<i>Marco Mondelli, Ramji Venkataramanan</i>	
Automatic and Harmless Regularization with Constrained and Lexicographic Optimization: A Dynamic Barrier Approach	29630
<i>Chengyue Gong, Qiang Liu, Xingchao Liu</i>	
Corruption Robust Active Learning	29643
<i>Yifang Chen, Simon Du, Kevin Jamieson</i>	
Metadata-Based Multi-Task Bandits with Bayesian Hierarchical Models	29655
<i>Lin Ge, Rui Song, Runzhe Wan</i>	
Program Synthesis Guided Reinforcement Learning for Partially Observed Environments.....	29669
<i>Osbert Bastani, Jeevana Priya Inala, Yewen Pu, Martin Rinard, Armando Solar-Lezama, Yichen Yang</i>	
Robust Allocations with Diversity Constraints	29684
<i>Lodewijk Gelauff, Ashish Goel, Aleksandra Korolova, Kamesh Munagala, Zeyu Shen</i>	
Activation Sharing with Asymmetric Paths Solves Weight Transport Problem Without Bidirectional Connection.....	29697
<i>Jiwoo Hong, Dongsuk Jeon, Jeongwoo Park, Sunghyeon Woo</i>	
BlendGAN: Implicitly GAN Blending for Arbitrary Stylized Face Generation	29710
<i>Qiang Li, Mingcong Liu, Zekui Qin, Pengfei Wan, Guoxin Zhang, Wen Zheng</i>	
Differentially Private Model Personalization	29723
<i>Abhradeep Guha Thakurta, Prateek Jain, John Rush, Adam Smith, Shuang Song</i>	
Rates of Estimation of Optimal Transport Maps Using Plug-In Estimators Via Barycentric Projections.....	29736
<i>Nabarun Deb, Promit Ghosal, Bodhisattva Sen</i>	
Robust Generalization Despite Distribution Shift Via Minimum Discriminating Information	29754
<i>Andreas Krause, Daniel Kuhn, Tobias Sutter</i>	

Soft Calibration Objectives for Neural Networks..... 29768
*Nicholas Cain, Archit Karandikar, Balaji Lakshminarayanan, Michael Mozer, Becca Roelofs,
Jonathon Shlens, Dustin Tran*

Distributional Gradient Matching for Learning Uncertain Neural Dynamics Models 29780
Florian Dorfler, Andreas Krause, Lenart Treven, Philippe Wenk

VOLUME 36

Shaping Embodied Agent Behavior with Activity-Context Priors from Egocentric Video 29794
Kristen Grauman, Tushar Nagarajan

Adjusting for Autocorrelated Errors in Neural Networks for Time Series 29806
Duane Boning, Chris Lang, Fan-Keng Sun

A Geometric Analysis of Neural Collapse with Unconstrained Features..... 29820
Tianyu Ding, Xiao Li, Qing Qu, Jeremias Sulam, Chong You, Jinxin Zhou, Zhihui Zhu

NeRS: Neural Reflectance Surfaces for Sparse-View 3D Reconstruction in the Wild 29835
Deva Ramanan, Shubham Tulsiani, Gengshan Yang, Jason Zhang

Unleashing the Power of Contrastive Self-Supervised Visual Models Via Contrast-Regularized
Fine-Tuning 29848
Jiashi Feng, Bryan Hooi, Dapeng Hu, Jian Liang, Yifan Zhang

Discovery of Options Via Meta-Learned Subgoals..... 29861
*Matteo Hessel, Iurii Kemaev, Junhyuk Oh, David Silver, Satinder Singh, Vivek Veeriah,
Zhongwen Xu, Tom Zahavy, Hado Van Hasselt*

Near-Optimal Lower Bounds for Convex Optimization for All Orders of Smoothness 29874
Ankit Garg, Robin Kothari, Praneeth Netrapalli, Suhail Sherif

Topology-Imbalance Learning for Semi-Supervised Node Classification 29885
Deli Chen, Peng Li, Yankai Lin, Xuancheng Ren, Xu Sun, Guangxiang Zhao, Jie Zhou

Gradient Inversion with Generative Image Prior..... 29898
Jinwoo Jeon, Jaechang Kim, Kangwook Lee, Sewoong Oh, Jungseul Ok

Beta-CROWN: Efficient Bound Propagation with Per-neuron Split Constraints for Neural Network
Robustness Verification..... 29909
Cho-Jui Hsieh, Suman Jana, J. Zico Kolter, Xue Lin, Shiqi Wang, Kaidi Xu, Huan Zhang

Autobahn: Automorphism-Based Graph Neural Nets 29922
Risi Kondor, Erik Thiede, Wenda Zhou

Data Augmentation Can Improve Robustness..... 29935
*Dan Andrei Calian, Sven Gowal, Timothy A Mann, Sylvestre-Alvise Rebuffi, Florian
Stimberg, Olivia Wiles*

Deep Explicit Duration Switching Models for Time Series 29949
*Abdul Fatir Ansari, Konstantinos Benidis, Tim Januschowski, Richard Kurle, Alexander
Smola, Harold Soh, Ali Caner Turkmen, Bernie Wang*

Shared Independent Component Analysis for Multi-Subject Neuroimaging 29962
Pierre Ablin, Alexandre Gramfort, Aapo Hyvarinen, Hugo Richard, Bertrand Thirion

Shape from Blur: Recovering Textured 3D Shape and Motion of Fast Moving Objects	29972
<i>Vittorio Ferrari, Martin R. Oswald, Marc Pollefeys, Denys Rozumnyi</i>	
Batched Thompson Sampling.....	29984
<i>Cem Kalkanli, Ayfer Ozgur</i>	
Delayed Gradient Averaging: Tolerate the Communication Latency for Federated Learning	29995
<i>Song Han, Hongzhou Lin, Yujun Lin, Yao Lu, Ligeng Zhu</i>	
Focal Attention for Long-Range Interactions in Vision Transformers.....	30008
<i>Xiyang Dai, Jianfeng Gao, Chunyuan Li, Bin Xiao, Jianwei Yang, Lu Yuan, Pengchuan Zhang</i>	
Scalable and Stable Surrogates for Flexible Classifiers with Fairness Constraints	30023
<i>Henry C Bendekgey, Erik Sudderth</i>	
Residual Pathway Priors for Soft Equivariance Constraints.....	30037
<i>Gregory Benton, Marc Finzi, Andrew Wilson</i>	
Optimal Algorithms for Stochastic Contextual Preference Bandits	30050
<i>Aadirupa Saha</i>	
Tight High Probability Bounds for Linear Stochastic Approximation with Fixed Stepsize.....	30063
<i>Alain Durmus, Eric Moulines, Alexey Naumov, Sergey Samsonov, Kevin Scaman, Hoi-To Wai</i>	
Learning Large Neighborhood Search Policy for Integer Programming.....	30075
<i>Zhiguang Cao, Wen Song, Yaoxin Wu, Jie Zhang</i>	
Dynamic Trace Estimation	30088
<i>Prathamesh Dharangutte, Christopher Musco</i>	
Provable Representation Learning for Imitation with Contrastive Fourier Features	30100
<i>Ofir Nachum, Mengjiao Yang</i>	
MICO: Improved Representations Via Sampling-Based State Similarity for Markov Decision Processes	30113
<i>Pablo Samuel Castro, Tyler Kastner, Prakash Panangaden, Mark Rowland</i>	
Counterfactual Explanations in Sequential Decision Making Under Uncertainty.....	30127
<i>Abir De, Manuel Rodriguez, Stratis Tsirtsis</i>	
Streaming Linear System Identification with Reverse Experience Replay	30140
<i>Prateek Jain, Suhas Kowshik, Dheeraj Nagaraj, Praneeth Netrapalli</i>	
SmoothMix: Training Confidence-Calibrated Smoothed Classifiers for Certified Robustness	30153
<i>Jongheon Jeong, Do-Guk Kim, Minkyu Kim, Heung-Chang Lee, Sejun Park, Jinwoo Shin</i>	
Action-Guided 3D Human Motion Prediction.....	30169
<i>Xintong Han, Jian-Fang Hu, Zihang Lin, Jiangxin Sun, Jia Xu, Wei-Shi Zheng</i>	
Meta-Learning the Search Distribution of Black-Box Random Search Based Adversarial Attacks	30181
<i>Matthias Hein, Jan Metzner, Maksym Yatsura</i>	
Validating the Lottery Ticket Hypothesis with Inertial Manifold Theory	30196
<i>Dejing Dou, Jiayin Jin, Ruoming Jin, Ji Liu, Jiaxiang Ren, Lingfei Wu, Zeru Zhang, Zijie Zhang, Xin Zhao, Yang Zhou</i>	

Are My Deep Learning Systems Fair? an Empirical Study of Fixed-Seed Training.....	30211
<i>Jiahao Chen, Zeou Hu, Jungwon Kim, Thibaud Lutellier, Viet Hung Pham, Shangshu Qian, Sameena Shah, Lin Tan, Yaoliang Yu</i>	
Rectangular Flows for Manifold Learning	30228
<i>Anthony Caterini, John Cunningham, Gabriel Loaiza-Ganem, Geoff Pleiss</i>	
On the Generative Utility of Cyclic Conditionals	30242
<i>Chang Liu, Tie-Yan Liu, Tao Qin, Haoyue Tang, Jintao Wang</i>	
Structural Credit Assignment in Neural Networks Using Reinforcement Learning.....	30257
<i>Dhawal Gupta, James Kostas, Gabor Mihucz, Matthew Schlegel, Philip Thomas, Martha White</i>	
A Near-Optimal Algorithm for Stochastic Bilevel Optimization Via Double-Momentum.....	30271
<i>Mingyi Hong, Prashant Khanduri, Hoi-To Wai, Zhaoran Wang, Zhuoran Yang, Siliang Zeng</i>	
Generalized Jensen-Shannon Divergence Loss for Learning with Noisy Labels	30284
<i>Hossein Azizpour, Erik Englesson</i>	
Continual Learning Via Local Module Composition	30298
<i>Massimo Caccia, Laurent Charlin, Oleksiy Ostapenko, Pau Rodriguez</i>	
Model-Based Episodic Memory Induces Dynamic Hybrid Controls	30313
<i>Majid Abdolshah, Thommen Karimpanal George, Hung Le, Truyen Tran, Svetha Venkatesh</i>	
FedDR – Randomized Douglas-Rachford Splitting Algorithms for Nonconvex Federated Composite Optimization.....	30326
<i>Lam Nguyen, Nhan H Pham, Dzung Phan, Quoc Tran Dinh</i>	
Adversarial Examples Make Strong Poisons.....	30339
<i>Ping-Yeh Chiang, Wojciech Czaja, Liam Fowl, Jonas Geiping, Micah Goldblum, Tom Goldstein</i>	
Coresets for Decision Trees of Signals.....	30352
<i>Dan Feldman, Ibrahim Jubran, Ilan I Newman, Ernesto Evgeniy Sanches Shayda</i>	
Local Plasticity Rules Can Learn Deep Representations Using Self-Supervised Contrastive Predictions	30365
<i>Guillaume Bellec, Wulfram Gerstner, Bernd Illing, Jean Ventura</i>	
MobTCast: Leveraging Auxiliary Trajectory Forecasting for Human Mobility Prediction	30380
<i>Nuria Oliver, Yongli Ren, Flora Salim, Hao Xue</i>	
Early Convolutions Help Transformers See Better	30392
<i>Trevor Darrell, Piotr Dollar, Ross Girshick, Eric Mintun, Mannat Singh, Tete Xiao</i>	
Error Compensated Distributed SGD Can Be Accelerated	30401
<i>Xun Qian, Peter Richtarik, Tong Zhang</i>	
InfoGCL: Information-Aware Graph Contrastive Learning.....	30414
<i>Haifeng Chen, Wei Cheng, Dongsheng Luo, Dongkuan Xu, Xiang Zhang</i>	
Meta-Learning for Relative Density-Ratio Estimation.....	30426
<i>Yasuhiro Fujiwara, Tomoharu Iwata, Atsutoshi Kumagai</i>	

Overcoming the Curse of Dimensionality with Laplacian Regularization in Semi-Supervised Learning	30439
<i>Francis Bach, Vivien Cabannes, Loucas Pillaud-Vivien, Alessandro Rudi</i>	
Unlabeled Principal Component Analysis.....	30452
<i>Liangzu Peng, Manolis Tsakiris, Yunzhen Yao</i>	
Causal-BALD: Deep Bayesian Active Learning of Outcomes to Infer Treatment-Effects from Observational Data	30465
<i>Yarin Gal, Andrew Jesson, Andreas Kirsch, Uri Shalit, Panagiotis Tigas, Joost Van Amersfoort</i>	
Scalable Rule-Based Representation Learning for Interpretable Classification	30479
<i>Ning Liu, Jianyong Wang, Zhuo Wang, Wei Zhang</i>	
Bridging Non Co-Occurrence with Unlabeled In-the-wild Data for Incremental Object Detection	30492
<i>Na Dong, Mingli Ding, Gim Hee Lee, Yongqiang Zhang</i>	
A Regression Approach to Learning-Augmented Online Algorithms.....	30504
<i>Keerti Anand, Rong Ge, Amit Kumar, Debmalya Panigrahi</i>	

Author Index