

International Conference on Machine Learning (ICML 2023)

Proceedings of Machine Learning Research Volume 202

Honolulu, Hawaii, USA
23 – 29 July 2023

Part 1 of 54

ISBN: 978-1-7138-8918-2

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© (2023) by the Authors
All rights reserved.

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

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

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

Phone: (570) 476-8006
Fax: (570) 476-0860

rasmussen@ptd.net

Additional copies of this publication are available from:

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

TABLE OF CONTENTS

PART 1

Data Structures for Density Estimation	1
<i>Anders Aamand, Alexandr Andoni, Justin Y. Chen, Piotr Indyk, Shyam Narayanan, Sandeep Silwal</i>	
ClusterFuG: Clustering Fully Connected Graphs by Multicut	19
<i>Ahmed Abbas, Paul Swoboda</i>	
Generalization on the Unseen, Logic Reasoning and Degree Curriculum	31
<i>Emmanuel Abbe, Samy Bengio, Aryo Lotfi, Kevin Rizk</i>	
Toward Large Kernel Models	61
<i>Amirhesam Abedsoltan, Mikhail Belkin, Parthe Pandit</i>	
Expertise Trees Resolve Knowledge Limitations in Collective Decision-Making	79
<i>Axel Abels, Tom Lenaerts, Vito Trianni, Ann Nowe</i>	
Comparison of Meta-Learners for Estimating Multi-Valued Treatment Heterogeneous Effects	91
<i>Naoufal Acharki, Ramiro Lugo, Antoine Bertonecello, Josselin Garnier</i>	
BNN-DP: Robustness Certification of Bayesian Neural Networks Via Dynamic Programming	133
<i>Steven Adams, Andrea Patane, Morteza Lahijanian, Luca Laurenti</i>	
SAM Operates Far from Home: Eigenvalue Regularization as a Dynamical Phenomenon	152
<i>Atish Agarwala, Yann Dauphin</i>	
Second-Order Regression Models Exhibit Progressive Sharpening to the Edge of Stability	169
<i>Atish Agarwala, Fabian Pedregosa, Jeffrey Pennington</i>	
Global Optimality of Elman-Type RNNs in the Mean-Field Regime	196
<i>Andrea Agazzi, Jianfeng Lu, Sayan Mukherjee</i>	
SemSup-XC: Semantic Supervision for Zero and Few-Shot Extreme Classification	228
<i>Pranjal Aggarwal, Ameet Deshpande, Karthik R Narasimhan</i>	
Adaptive IMLE for Few-Shot Pretraining-Free Generative Modelling	248
<i>Mehran Aghabozorgi, Shichong Peng, Ke Li</i>	
Scaling Laws for Generative Mixed-Modal Language Models	265
<i>Armen Aghajanyan, Lili Yu, Alexis Conneau, Wei-Ning Hsu, Karen Hambarzumyan, Susan Zhang, Stephen Roller, Naman Goyal, Omer Levy, Luke Zettlemoyer</i>	
Hypothesis Transfer Learning with Surrogate Classification Losses: Generalization Bounds Through Algorithmic Stability	280
<i>Anass Aghbalou, Guillaume Staerman</i>	
Constrained Causal Bayesian Optimization	304
<i>Virginia Aglietti, Alan Malek, Ira Ktena, Silvia Chiappa</i>	
Explaining the Effects of Non-Convergent MCMC in the Training of Energy-Based Models	322
<i>Elisabeth Agoritsas, Giovanni Catania, Aurélien Decelle, Beatriz Seoane</i>	
Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies	337
<i>Gati V Aher, Rosa I. Arriaga, Adam Tauman Kalai</i>	
Interventional Causal Representation Learning	372
<i>Kartik Ahuja, Divyat Mahajan, Yixin Wang, Yoshua Bengio</i>	
Sequential Underspecified Instrument Selection for Cause-Effect Estimation	408
<i>Elisabeth Ailer, Jason Hartford, Niki Kilbertus</i>	
Atari-5: Distilling the Arcade Learning Environment Down to Five Games	421
<i>Matthew Aitchison, Penny Sweetser, Marcus Hutter</i>	
Towards Credible Visual Model Interpretation with Path Attribution	439
<i>Naveed Akhtar, Mohammad A. A. K. Jalwana</i>	

Convergence of First-Order Methods for Constrained Nonconvex Optimization with Dependent Data	458
<i>Ahmet Alacaoglu, Hanbaek Lyu</i>	
Recasting Self-Attention with Holographic Reduced Representations	490
<i>Mohammad Mahmudul Alam, Edward Raff, Stella Biderman, Tim Oates, James Holt</i>	
The Saddle-Point Method in Differential Privacy	508
<i>Wael Alghamdi, Juan Felipe Gomez, Shahab Asoodeh, Flavio Calmon, Oliver Kosut, Lalitha Sankar</i>	
Nonlinear Advantage: Trained Networks Might Not Be as Complex as You Think	529
<i>Christian H. X. Ali Mehmeti-Göpel, Jan Disselhoff</i>	
A Simple Zero-Shot Prompt Weighting Technique to Improve Prompt Ensembling in Text-Image Models	547
<i>James Urquhart Allingham, Jie Ren, Michael W Dusenberry, Xiuye Gu, Yin Cui, Dustin Tran, Jeremiah Zhe Liu, Balaji Lakshminarayanan</i>	
On the Privacy-Robustness-Utility Trilemma in Distributed Learning	569
<i>Youssef Allouah, Rachid Guerraoui, Nirupam Gupta, Rafael Pinot, John Stephan</i>	
Differentially Private Distributed Bayesian Linear Regression with MCMC	627
<i>Baris Alparslan, Sinan Yildirim, Ilker Birbil</i>	
Robust and Scalable Bayesian Online Changepoint Detection	642
<i>Matias Altamirano, Francois-Xavier Briol, Jeremias Knoblauch</i>	

PART 2

Neural Wasserstein Gradient Flows for Discrepancies with Riesz Kernels.....	664
<i>Fabian Altekrüger, Johannes Hertrich, Gabriele Steidl</i>	
Distributed Contextual Linear Bandits with Minimax Optimal Communication Cost.....	691
<i>Sanae Amani, Tor Lattimore, András György, Lin Yang</i>	
A Kernelized Stein Discrepancy for Biological Sequences.....	718
<i>Alan Nawzad Amin, Eli N Weinstein, Debora Susan Marks</i>	
The Optimal Approximation Factors in Misspecified Off-Policy Value Function Estimation	768
<i>Philip Amortila, Nan Jiang, Csaba Szepesvari</i>	
Meta Optimal Transport	791
<i>Brandon Amos, Giulia Luise, Samuel Cohen, Ievgen Redko</i>	
Near-Optimal Φ -Regret Learning in Extensive-Form Games	814
<i>Ioannis Anagnostides, Gabriele Farina, Tuomas Sandholm</i>	
A Modern Look at the Relationship Between Sharpness and Generalization	840
<i>Maksym Andriushchenko, Francesco Croce, Maximilian Müller, Matthias Hein, Nicolas Flammarion</i>	
SGD with Large Step Sizes Learns Sparse Features	903
<i>Maksym Andriushchenko, Aditya Vardhan Varre, Loucas Pillaud-Vivien, Nicolas Flammarion</i>	
Neural Continuous-Discrete State Space Models for Irregularly-Sampled Time Series	926
<i>Abdul Fatir Ansari, Alvin Heng, Andre Lim, Harold Soh</i>	
Paging with Succinct Predictions	952
<i>Antonios Antoniadis, Joan Boyar, Marek Elias, Lene Monrad Favrholt, Ruben Hoeksma, Kim S. Larsen, Adam Polak, Bertrand Simon</i>	
Mixing Predictions for Online Metric Algorithms	969
<i>Antonios Antoniadis, Christian Coester, Marek Elias, Adam Polak, Bertrand Simon</i>	
Exponential Smoothing for Off-Policy Learning	984
<i>Imad Aouali, Victor-Emmanuel Brunel, David Rohde, Anna Korba</i>	
Polynomial Time and Private Learning of Unbounded Gaussian Mixture Models	1018
<i>Jamil Arbas, Hassan Ashtiani, Christopher Liaw</i>	

Principled Acceleration of Iterative Numerical Methods Using Machine Learning.....	1041
<i>Sohei Arisaka, Qianxiao Li</i>	
Faster Rates of Convergence to Stationary Points in Differentially Private Optimization	1060
<i>Raman Arora, Raef Bassily, Tomás González, Cristóbal A Guzmán, Michael Menart, Enayat Ullah</i>	
Prototype-Sample Relation Distillation: Towards Replay-Free Continual Learning.....	1093
<i>Nader Asadi, Mohammadreza Davari, Sudhir Mudur, Rahaf Aljundi, Eugene Belilovsky</i>	
Near-Optimal Algorithms for Private Online Optimization in the Realizable Regime	1107
<i>Hilal Asi, Vitaly Feldman, Tomer Koren, Kunal Talwar</i>	
From Robustness to Privacy and Back	1121
<i>Hilal Asi, Jonathan Ullman, Lydia Zakyntinou</i>	
SGD with AdaGrad Stepsizes: Full Adaptivity with High Probability to Unknown Parameters, Unbounded Gradients and Affine Variance	1147
<i>Amit Attia, Tomer Koren</i>	
Adversarially Robust PAC Learnability of Real-Valued Functions.....	1172
<i>Idan Attias, Steve Hanneke</i>	
Infusing Lattice Symmetry Priors in Attention Mechanisms for Sample-Efficient Abstract Geometric Reasoning	1200
<i>Mattia Atzeni, Mrinmaya Sachan, Andreas Loukas</i>	
Learning to Initiate and Reason in Event-Driven Cascading Processes	1218
<i>Yuval Atzmon, Eli Meir, Shie Mannor, Gal Chechik</i>	
On the Convergence of the MLE as an Estimator of the Learning Rate in the Exp3 Algorithm	1244
<i>Julien Aubert, Luc Lehericy, Patricia Reynaud-Bouret</i>	
Dirichlet Diffusion Score Model for Biological Sequence Generation	1276
<i>Pavel Avdeyev, Chenlai Shi, Yuhao Tan, Kseniia Dudnyk, Jian Zhou</i>	
Gradient Descent Converges Linearly for Logistic Regression on Separable Data	1302
<i>Kyriakos Axiotis, Maxim Sviridenko</i>	
Naive Imputation Implicitly Regularizes High-Dimensional Linear Models.....	1320
<i>Alexis Ayme, Claire Boyer, Aymeric Dieuleveut, Erwan Scornet</i>	
Half-Hop: A Graph Upsampling Approach for Slowing Down Message Passing.....	1341
<i>Mehdi Azabou, Venkataramana Ganesh, Shantanu Thakoor, Chi-Heng Lin, Lakshmi Sathidevi, Ran Liu, Michal Valko, Petar Velickovic, Eva L Dyer</i>	
CLUTR: Curriculum Learning Via Unsupervised Task Representation Learning	1361
<i>Abdus Salam Azad, Izzeddin Gur, Jasper Emhoff, Nathaniel Alexis, Aleksandra Faust, Pieter Abbeel, Ion Stoica</i>	
Personalized Subgraph Federated Learning	1396
<i>Jinheon Baek, Wonyong Jeong, Jiongdoo Jin, Jaehong Yoon, Sung Ju Hwang</i>	
Efficient Self-Supervised Learning with Contextualized Target Representations for Vision, Speech and Language	1416
<i>Alexei Baevski, Arun Babu, Wei-Ning Hsu, Michael Auli</i>	
Efficient Preconditioned Stochastic Gradient Descent for Estimation in Latent Variable Models	1430
<i>Charlotte Baey, Maud Delattre, Estelle Kuhn, Jean-Benoist Leger, Sarah Lemler</i>	

PART 3

Feed Two Birds with One Scone: Exploiting Wild Data for Both Out-Of-Distribution Generalization and Detection	1454
<i>Haoyue Bai, Gregory Canal, Xuefeng Du, Jeongyeol Kwon, Robert D Nowak, Yixuan Li</i>	
Answering Complex Logical Queries on Knowledge Graphs Via Query Computation Tree Optimization.....	1472
<i>Yushi Bai, Xin Lv, Juanzi Li, Lei Hou</i>	

Linear Optimal Partial Transport Embedding.....	1492
<i>Yikun Bai, Ivan Vladimir Medri, Rocio Diaz Martin, Rana Shahroz, Soheil Kolouri</i>	
Implicit Graph Neural Networks: A Monotone Operator Viewpoint.....	1521
<i>Justin Baker, Qingsong Wang, Cory D Hauck, Bao Wang</i>	
Tensor Decompositions Meet Control Theory: Learning General Mixtures of Linear Dynamical Systems.....	1549
<i>Ainesh Bakshi, Allen Liu, Ankur Moitra, Morris Yau</i>	
Block Subsampled Randomized Hadamard Transform for Nyström Approximation on Distributed Architectures	1564
<i>Oleg Balabanov, Matthias Beaupère, Laura Grigori, Victor Lederer</i>	
Efficient Online Reinforcement Learning with Offline Data	1577
<i>Philip J. Ball, Laura Smith, Ilya Kostrikov, Sergey Levine</i>	
Mirror Sinkhorn: Fast Online Optimization on Transport Polytopes	1595
<i>Marin Ballu, Quentin Berthet</i>	
On the Functional Similarity of Robust and Non-Robust Neural Representations	1614
<i>András Balogh, Márk Jelasity</i>	
Robust Budget Pacing with a Single Sample	1636
<i>Santiago R. Balseiro, Rachitesh Kumar, Vahab Mirrokni, Balasubramanian Sivan, Di Wang</i>	
Dynamic Constrained Submodular Optimization with Polylogarithmic Update Time.....	1660
<i>Kiarash Banihashem, Leyla Biabani, Samira Goudarzi, Mohammadtaghi Hajiaghayi, Peyman Jabbarzade, Morteza Monemizadeh</i>	
One Transformer Fits All Distributions in Multi-Modal Diffusion at Scale	1692
<i>Fan Bao, Shen Nie, Kaiwen Xue, Chongxuan Li, Shi Pu, Yaole Wang, Gang Yue, Yue Cao, Hang Su, Jun Zhu</i>	
Optimizing the Collaboration Structure in Cross-Silo Federated Learning.....	1718
<i>Wenxuan Bao, Haohan Wang, Jun Wu, Jingrui He</i>	
MultiDiffusion: Fusing Diffusion Paths for Controlled Image Generation	1737
<i>Omer Bar-Tal, Lior Yariv, Yaron Lipman, Tali Dekel</i>	
Reinforcement Learning with General Utilities: Simpler Variance Reduction and Large State-Action Space.....	1753
<i>Anas Barakat, Ilyas Fatkhullin, Niao He</i>	
Interpretable Neural-Symbolic Concept Reasoning	1801
<i>Pietro Barbiero, Gabriele Ciravegna, Francesco Giannini, Mateo Espinosa Zarlenga, Lucie Charlotte Magister, Alberto Tonda, Pietro Lio, Frederic Precioso, Mateja Jamnik, Giuseppe Marra</i>	
Moccasin: Efficient Tensor Rematerialization for Neural Networks.....	1826
<i>Burak Bartan, Haoming Li, Harris Teague, Christopher Lott, Bistra Dilkina</i>	
User-Level Private Stochastic Convex Optimization with Optimal Rates	1838
<i>Raef Bassily, Ziteng Sun</i>	
A Statistical Perspective on Retrieval-Based Models.....	1852
<i>Soumya Basu, Ankit Singh Rawat, Manzil Zaheer</i>	
Human-Timescale Adaptation in an Open-Ended Task Space.....	1887
<i>Jakob Bauer, Kate Baumli, Feryal Behbahani, Avishkar Bhoopchand, Nathalie Bradley-Schmieg, Michael Chang, Natalie Clay, Adrian Collister, Vibhavari Dasagi, Lucy Gonzalez, Karol Gregor, Edward Hughes, Sheleem Kashem, Maria Loks-Thompson, Hannah Openshaw, Jack Parker-Holder, Shreya Pathak, Nicolas Perez-Nieves, Nemanja Rakicevic, Tim Rocktäschel, Yannick Schroecker, Satinder Singh, Jakub Sygnowski, Karl Tuyls, Sarah York, Alexander Zacherl, Lei M Zhang</i>	
A Kernel Stein Test of Goodness of Fit for Sequential Models.....	1936
<i>Jerome Baum, Heishiro Kanagawa, Arthur Gretton</i>	
Individually Fair Learning with One-Sided Feedback	1954
<i>Yahav Bechavod, Aaron Roth</i>	

Predicting Ordinary Differential Equations with Transformers.....	1978
<i>Sören Becker, Michal Klein, Alexander Neitz, Giambattista Parascandolo, Niki Kilbertus</i>	
Explaining Reinforcement Learning with Shapley Values	2003
<i>Daniel Beechey, Thomas M. S. Smith, Özgür Simsek</i>	
TIDE: Time Derivative Diffusion for Deep Learning on Graphs.....	2015
<i>Maysam Behmanesh, Maximilian Krahn, Maks Ovsjanikov</i>	
Fast as CHITA: Neural Network Pruning with Combinatorial Optimization.....	2031
<i>Riade Benbaki, Wenyu Chen, Xiang Meng, Hussein Hazimeh, Natalia Ponomareva, Zhe Zhao, Rahul Mazumder</i>	
Continuously Parameterized Mixture Models	2050
<i>Christopher M Bender, Yifeng Shi, Marc Niethammer, Junier Oliva</i>	
Controllable Neural Symbolic Regression	2063
<i>Tommaso Bendinelli, Luca Biggio, Pierre-Alexandre Kamienny</i>	
On Second-Order Scoring Rules for Epistemic Uncertainty Quantification	2078
<i>Viktor Bengs, Eyke Hüllermeier, Willem Waegeman</i>	
Certified Robust Neural Networks: Generalization and Corruption Resistance.....	2092
<i>Amine Bennouna, Ryan Lucas, Bart Van Parys</i>	
Gaussian Processes at the Helm(holtz): A More Fluid Model for Ocean Currents.....	2113
<i>Renato Berlinghieri, Brian L. Trippe, David R. Burt, Ryan James Giordano, Kaushik Srinivasan, Tamay Özgökmen, Junfei Xia, Tamara Broderick</i>	
Optimal Rates and Efficient Algorithms for Online Bayesian Persuasion	2164
<i>Martino Bernasconi, Matteo Castiglioni, Andrea Celli, Alberto Marchesi, Francesco Trovò, Nicola Gatti</i>	
Constrained Phi-Equilibria	2184
<i>Martino Bernasconi, Matteo Castiglioni, Alberto Marchesi, Francesco Trovò, Nicola Gatti</i>	
Differentiable and Transportable Structure Learning	2206
<i>Jeroen Berrevoets, Nabeel Seedat, Fergus Imrie, Mihaela Van Der Schaar</i>	
Polyhedral Complex Extraction from ReLU Networks Using Edge Subdivision	2234
<i>Arturs Berzins</i>	
Robust One-Class Classification with Signed Distance Function Using 1-Lipschitz Neural Networks	2245
<i>Louis Béthune, Paul Novello, Guillaume Coiffier, Thibaut Boissin, Mathieu Serrurier, Quentin Vincenot, Andres Troya-Galvis</i>	

PART 4

Neural Algorithmic Reasoning with Causal Regularisation.....	2272
<i>Beatrice Bevilacqua, Kyriacos Nikiforou, Borja Ibarz, Ioana Bica, Michela Paganini, Charles Blundell, Jovana Mitrovic, Petar Velickovic</i>	
Optimally-Weighted Estimators of the Maximum Mean Discrepancy for Likelihood-Free Inference	2289
<i>Ayush Bharti, Masha Naslidnyk, Oscar Key, Samuel Kaski, Francois-Xavier Briol</i>	
Bandit Online Linear Optimization with Hints and Queries	2313
<i>Aditya Bhaskara, Ashok Cutkosky, Ravi Kumar, Manish Purohit</i>	
Improved Online Conformal Prediction Via Strongly Adaptive Online Learning.....	2337
<i>Aadyot Bhatnagar, Huan Wang, Caiming Xiong, Yu Bai</i>	
Data-Copying in Generative Models: A Formal Framework.....	2364
<i>Robi Bhattacharjee, Sanjoy Dasgupta, Kamalika Chaudhuri</i>	
Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling.....	2397
<i>Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, Usvsn Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, Oskar Van Der Wal</i>	

StriderNet: A Graph Reinforcement Learning Approach to Optimize Atomic Structures on Rough Energy Landscapes	2431
<i>Vaibhav Bihani, Sahil Manchanda, Srikanth Sastry, Sayan Ranu, N M Anoop Krishnan</i>	
Modeling Temporal Data as Continuous Functions with Stochastic Process Diffusion	2452
<i>Marin Biloš, Kashif Rasul, Anderson Schneider, Yuriy Nevmyvaka, Stephan Günnemann</i>	
In Or Out? Fixing ImageNet Out-Of-Distribution Detection Evaluation	2471
<i>Julian Bitterwolf, Maximilian Müller, Matthias Hein</i>	
Invariant Slot Attention: Object Discovery with Slot-Centric Reference Frames	2507
<i>Ondrej Biza, Sjoerd Van Steenkiste, Mehdi S. M. Sajjadi, Gamaleldin Fathy Elsayed, Aravindh Mahendran, Thomas Kipf</i>	
Understanding Oversquashing in GNNs Through the Lens of Effective Resistance.....	2528
<i>Mitchell Black, Zhengchao Wan, Amir Nayyeri, Yusu Wang</i>	
Unit Scaling: Out-Of-The-Box Low-Precision Training	2548
<i>Charlie Blake, Douglas Orr, Carlo Luschi</i>	
FLEX: An Adaptive Exploration Algorithm for Nonlinear Systems	2577
<i>Matthieu Blanke, Marc Lelarge</i>	
Not All Strongly Rayleigh Distributions Have Small Probabilistic Generating Circuits	2592
<i>Markus Bläser</i>	
Learning the Dynamics of Sparsely Observed Interacting Systems	2603
<i>Linus Bleistein, Adeline Fermanian, Anne-Sophie Jannot, Agathe Guilloux</i>	
Subset Selection Based on Multiple Rankings in the Presence of Bias: Effectiveness of Fairness Constraints for Multiwinner Voting Score Functions	2641
<i>Niclas Boehmer, L. Elisa Celis, Lingxiao Huang, Anay Mehrotra, Nisheeth K. Vishnoi</i>	
Properties of the Mallows Model Depending on the Number of Alternatives: A Warning for an Experimentalist.....	2689
<i>Niclas Boehmer, Piotr Faliszewski, Sonja Kraiczy</i>	
A Robust Optimisation Perspective on Counterexample-Guided Repair of Neural Networks.....	2712
<i>David Boetius, Stefan Leue, Tobias Sutter</i>	
Beyond the Universal Law of Robustness: Sharper Laws for Random Features and Neural Tangent Kernels	2738
<i>Simone Bombari, Shayan Kiyani, Marco Mondelli</i>	
Sliced-Wasserstein on Symmetric Positive Definite Matrices for M/EEG Signals.....	2777
<i>Clément Bonet, Benoit Malézieux, Alain Rakotomamonjy, Lucas Drumetz, Thomas Moreau, Matthieu Kowalski, Nicolas Courty</i>	
Spherical Fourier Neural Operators: Learning Stable Dynamics on the Sphere	2806
<i>Boris Bonev, Thorsten Kurth, Christian Hundt, Jaideep Pathak, Maximilian Baust, Karthik Kashinath, Anima Anandkumar</i>	
The Regret of Exploration and the Control of Bad Episodes in Reinforcement Learning	2824
<i>Victor Boone, Bruno Gaujal</i>	
Model-Agnostic Measure of Generalization Difficulty	2857
<i>Akhilan Boopathy, Kevin Liu, Jaedong Hwang, Shu Ge, Asaad Mohammedsaleh, Ila R Fiete</i>	
Returning the Favour: When Regression Benefits from Probabilistic Causal Knowledge.....	2885
<i>Shahine Bouabid, Jake Fawkes, Dino Sejdinovic</i>	
In Search for a Generalizable Method for Source Free Domain Adaptation	2914
<i>Malik Boudiaf, Tom Denton, Bart Van Merriënboer, Vincent Dumoulin, Eleni Triantafillou</i>	
Quantum Speedups for Zero-Sum Games Via Improved Dynamic Gibbs Sampling	2932
<i>Adam Bouland, Yosheb M Getachew, Yujia Jin, Aaron Sidford, Kevin Tian</i>	
Diffusion Models as Artists: Are We Closing the Gap Between Humans and Machines?	2953
<i>Victor Boutin, Thomas Fel, Lakshya Singhal, Rishav Mukherji, Akash Nagaraj, Julien Colin, Thomas Serre</i>	
Settling the Reward Hypothesis	3003
<i>Michael Bowling, John D Martin, David Abel, Will Dabney</i>	

ILLUME: Rationalizing Vision-Language Models Through Human Interactions	3021
<i>Manuel Brack, Patrick Schramowski, Björn Deiseroth, Kristian Kersting</i>	
Provably Learning Object-Centric Representations	3038
<i>Jack Brady, Roland S. Zimmermann, Yash Sharma, Bernhard Schölkopf, Julius Von Kügelgen, Wieland Brendel</i>	
Quantifying Human Priors Over Social and Navigation Networks	3063
<i>Gecia Bravo-Hermsdorff</i>	

PART 5

Critical Points and Convergence Analysis of Generative Deep Linear Networks Trained with Bures-Wasserstein Loss	3106
<i>Pierre Bréchet, Katerina Papagiannouli, Jing An, Guido Montufar</i>	
Emergence of Sparse Representations from Noise	3148
<i>Trenton Bricken, Rylan Schaeffer, Bruno Olshausen, Gabriel Kreiman</i>	
Differentially Private Optimization on Large Model at Small Cost	3192
<i>Zhiqi Bu, Yu-Xiang Wang, Sheng Zha, George Karypis</i>	
Machine Learning Force Fields with Data Cost Aware Training	3219
<i>Alexander Bukharin, Tianyi Liu, Shengjie Wang, Simiao Zuo, Weihao Gao, Wen Yan, Tuo Zhao</i>	
Label Differential Privacy and Private Training Data Release	3233
<i>Robert Istvan Busa-Fekete, Andres Munoz Medina, Umar Syed, Sergei Vassilvitskii</i>	
The SSL Interplay: Augmentations, Inductive Bias, and Generalization	3252
<i>Vivien Cabannes, Bobak Kiani, Randall Balestriero, Yann Lecun, Alberto Bietti</i>	
Online Mechanism Design for Information Acquisition	3299
<i>Federico Cacciamani, Matteo Castiglioni, Nicola Gatti</i>	
MyoDex: A Generalizable Prior for Dexterous Manipulation	3327
<i>Vittorio Caggiano, Sudeep Dasari, Vikash Kumar</i>	
What Can Be Learnt with Wide Convolutional Neural Networks?	3347
<i>Francesco Cagnetta, Alessandro Favero, Matthieu Wyart</i>	
Causal Discovery with Latent Confounders Based on Higher-Order Cumulants	3380
<i>Ruichu Cai, Zhiyi Huang, Wei Chen, Zhifeng Hao, Kun Zhang</i>	
On the Connection Between MPNN and Graph Transformer	3408
<i>Chen Cai, Truong Son Hy, Rose Yu, Yusu Wang</i>	
Ske2Grid: Skeleton-To-Grid Representation Learning for Action Recognition	3431
<i>Dongqi Cai, Yangyuxuan Kang, Anbang Yao, Yurong Chen</i>	
Extrapolated Random Tree for Regression	3442
<i>Yuchao Cai, Yuheng Ma, Yiwei Dong, Hanfang Yang</i>	
Cyclic Block Coordinate Descent with Variance Reduction for Composite Nonconvex Optimization	3469
<i>Xufeng Cai, Chaobing Song, Stephen Wright, Jelena Diakonikolas</i>	
Robust Weight Signatures: Gaining Robustness as Easy as Patching Weights?	3495
<i>Ruisi Cai, Zhenyu Zhang, Zhangyang Wang</i>	
Doubly Optimal No-Regret Learning in Monotone Games	3507
<i>Yang Cai, Weiqiang Zheng</i>	
Multi-Agent Learning from Learners	3525
<i>Mine Melodi Caliskan, Francesco Chini, Setareh Maghsudi</i>	
Efficient Learning of Mesh-Based Physical Simulation with Bi-Stride Multi-Scale Graph Neural Network	3541
<i>Yadi Cao, Menglei Chai, Minchen Li, Chenfanfu Jiang</i>	

Variational Sparse Inverse Cholesky Approximation for Latent Gaussian Processes Via Double Kullback-Leibler Minimization	3559
<i>Jian Cao, Myeongjong Kang, Felix Jimenez, Huiyan Sang, Florian Tobias Schaefer, Matthias Katzfuss</i>	
Learning Lightweight Object Detectors Via Multi-Teacher Progressive Distillation	3577
<i>Shengcao Cao, Mengtian Li, James Hays, Deva Ramanan, Yu-Xiong Wang, Liangyan Gui</i>	
One-Sided Matrix Completion from Two Observations Per Row	3599
<i>Steven Cao, Percy Liang, Gregory Valiant</i>	
State and Parameter Learning with PARIS Particle Gibbs	3625
<i>Gabriel Cardoso, Yazid Janati El Idrissi, Sylvain Le Corff, Eric Moulines, Jimmy Olsson</i>	
Grounding Large Language Models in Interactive Environments with Online Reinforcement Learning	3676
<i>Thomas Carta, Clément Romac, Thomas Wolf, Sylvain Lamprier, Olivier Sigaud, Pierre-Yves Oudeyer</i>	
Stein Variational Goal Generation for Adaptive Exploration in Multi-Goal Reinforcement Learning.....	3714
<i>Nicolas Castanet, Olivier Sigaud, Sylvain Lamprier</i>	
Scalable Safe Policy Improvement Via Monte Carlo Tree Search	3732
<i>Alberto Castellini, Federico Bianchi, Edoardo Zorzi, Thiago D. Simão, Alessandro Farinelli, Matthijs T. J. Spaan</i>	
LESS-VFL: Communication-Efficient Feature Selection for Vertical Federated Learning	3757
<i>Timothy Castiglia, Yi Zhou, Shiqiang Wang, Swanand Kadhe, Nathalie Baracaldo, Stacy Patterson</i>	
On the Robustness of Text Vectorizers	3782
<i>Rémi Catellier, Samuel Vaiter, Damien Garreau</i>	
Learning Globally Smooth Functions on Manifolds	3815
<i>Juan Cervino, Luiz F. O. Chamon, Benjamin David Haeffele, Rene Vidal, Alejandro Ribeiro</i>	
Tighter Lower Bounds for Shuffling SGD: Random Permutations and Beyond.....	3855
<i>Jaeyoung Cha, Jaewook Lee, Chulhee Yun</i>	

PART 6

Orthogonality-Enforced Latent Space in Autoencoders: An Approach to Learning Disentangled Representations	3913
<i>Jaehoon Cha, Jeyan Thiyagalingam</i>	
STEERING : Stein Information Directed Exploration for Model-Based Reinforcement Learning	3949
<i>Souradip Chakraborty, Amrit Bedi, Alec Koppel, Mengdi Wang, Furong Huang, Dinesh Manocha</i>	
Thompson Sampling for High-Dimensional Sparse Linear Contextual Bandits	3979
<i>Sunrit Chakraborty, Saptarshi Roy, Ambuj Tewari</i>	
Representations and Exploration for Deep Reinforcement Learning Using Singular Value Decomposition.....	4009
<i>Yash Chandak, Shantanu Thakoor, Zhaohan Daniel Guo, Yunhao Tang, Remi Munos, Will Dabney, Diana L Borsa</i>	
Memory-Based Dual Gaussian Processes for Sequential Learning.....	4035
<i>Paul Edmund Chang, Prakhar Verma, S. T. John, Arno Solin, Mohammad Emtiyaz Khan</i>	
Muse: Text-To-Image Generation Via Masked Generative Transformers	4055
<i>Huiwen Chang, Han Zhang, Jarred Barber, Aaron Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Patrick Murphy, William T. Freeman, Michael Rubinstein, Yuanzhen Li, Dilip Krishnan</i>	
On Investigating the Conservative Property of Score-Based Generative Models	4076
<i>Chen-Hao Chao, Wei-Fang Sun, Bo-Wun Cheng, Chun-Yi Lee</i>	
Robust and Private Stochastic Linear Bandits.....	4096
<i>Vasileios Charisopoulos, Hossein Esfandiari, Vahab Mirrokni</i>	

Streaming Submodular Maximization with Differential Privacy	4116
<i>Anamay Chaturvedi, Huy Nguyen, Thy Dinh Nguyen</i>	
Why Does Throwing Away Data Improve Worst-Group Error?.....	4144
<i>Kamalika Chaudhuri, Kartik Ahuja, Martin Arjovsky, David Lopez-Paz</i>	
Collaborative Multi-Agent Heterogeneous Multi-Armed Bandits	4189
<i>Ronshee Chawla, Daniel Vial, Sanjay Shakkottai, R. Srikant</i>	
Correcting Discount-Factor Mismatch in On-Policy Policy Gradient Methods.....	4218
<i>Fengdi Che, Gautham Vasan, A. Rupam Mahmood</i>	
Fast Federated Machine Unlearning with Nonlinear Functional Theory.....	4241
<i>Tianshi Che, Yang Zhou, Zijie Zhang, Lingjuan Lyu, Ji Liu, Da Yan, Dejing Dou, Jun Huan</i>	
On the Statistical Benefits of Temporal Difference Learning.....	4269
<i>David Cheikhi, Daniel Russo</i>	
Multi-Layer Neural Networks as Trainable Ladders of Hilbert Spaces.....	4294
<i>Zhengdao Chen</i>	
Beyond the Edge of Stability Via Two-Step Gradient Updates	4330
<i>Lei Chen, Joan Bruna</i>	
Trompt: Towards a Better Deep Neural Network for Tabular Data.....	4392
<i>Kuan-Yu Chen, Ping-Han Chiang, Hsin-Rung Chou, Ting-Wei Chen, Tien-Hao Chang</i>	
Differentially Private Stochastic Convex Optimization Under a Quantile Loss Function.....	4435
<i>Du Chen, Geoffrey A. Chua</i>	
Restoration-Degradation Beyond Linear Diffusions: A Non-Asymptotic Analysis for DDIM-Type Samplers	4462
<i>Sitan Chen, Giannis Daras, Alex Dimakis</i>	
Provably Convergent Schrödinger Bridge with Applications to Probabilistic Time Series Imputation	4485
<i>Yu Chen, Wei Deng, Shikai Fang, Fengpei Li, Nicole Tianjiao Yang, Yikai Zhang, Kashif Rasul, Shandian Zhe, Anderson Schneider, Yuriy Nevmyvaka</i>	
ED-Batch: Efficient Automatic Batching of Dynamic Neural Networks Via Learned Finite State Machines	4514
<i>Siyuan Chen, Pratik Pramod Fegade, Tianqi Chen, Phillip Gibbons, Todd Mowry</i>	
Is Learning Summary Statistics Necessary for Likelihood-Free Inference?	4529
<i>Yanzhi Chen, Michael U. Gutmann, Adrian Weller</i>	
Subequivariant Graph Reinforcement Learning in 3D Environments.....	4545
<i>Runfa Chen, Jiaqi Han, Fuchun Sun, Wenbing Huang</i>	
GuardHFL: Privacy Guardian for Heterogeneous Federated Learning	4566
<i>Hanxiao Chen, Meng Hao, Hongwei Li, Kangjie Chen, Guowen Xu, Tianwei Zhang, Xilin Zhang</i>	
Efficient and Degree-Guided Graph Generation Via Discrete Diffusion Modeling	4585
<i>Xiaohui Chen, Jiaying He, Xu Han, Liping Liu</i>	
Evolving Semantic Prototype Improves Generative Zero-Shot Learning	4611
<i>Shiming Chen, Wenjin Hou, Ziming Hong, Xiaohan Ding, Yibing Song, Xinge You, Tongliang Liu, Kun Zhang</i>	
Explore and Exploit the Diverse Knowledge in Model Zoo for Domain Generalization.....	4623
<i>Yimeng Chen, Tianyang Hu, Fengwei Zhou, Zhenguang Li, Zhi-Ming Ma</i>	
Decentralized Stochastic Bilevel Optimization with Improved per-Iteration Complexity	4641
<i>Xuxing Chen, Minhui Huang, Shiqian Ma, Krishna Balasubramanian</i>	
Score Approximation, Estimation and Distribution Recovery of Diffusion Models on Low-Dimensional Data	4672
<i>Minshuo Chen, Kaixuan Huang, Tuo Zhao, Mengdi Wang</i>	

PART 7

Sample Complexity of Probability Divergences Under Group Symmetry	4713
<i>Ziyu Chen, Markos Katsoulakis, Luc Rey-Bellet, Wei Zhu</i>	
Improved Analysis of Score-Based Generative Modeling: User-Friendly Bounds Under Minimal Smoothness Assumptions	4735
<i>Hongrui Chen, Holden Lee, Jianfeng Lu</i>	
Bidirectional Looking with a Novel Double Exponential Moving Average to Adaptive and Non-Adaptive Momentum Optimizers	4764
<i>Yineng Chen, Zuchao Li, Lefei Zhang, Bo Du, Hai Zhao</i>	
HarsanyiNet: Computing Accurate Shapley Values in a Single Forward Propagation.....	4804
<i>Lu Chen, Siyu Lou, Keyan Zhang, Jin Huang, Quanshi Zhang</i>	
Generalized Implicit Follow-The-Regularized-Leader.....	4826
<i>Keyi Chen, Francesco Orabona</i>	
Fisher Information Embedding for Node and Graph Learning.....	4839
<i>Dexiong Chen, Paolo Pellizzoni, Karsten Borgwardt</i>	
Rethinking Visual Reconstruction: Experience-Based Content Completion Guided by Visual Cues	4856
<i>Jiaxuan Chen, Yu Qi, Gang Pan</i>	
Stratified Adversarial Robustness with Rejection	4867
<i>Jiefeng Chen, Jayaram Raghuram, Jihye Choi, Xi Wu, Yingyu Liang, Somesh Jha</i>	
Multi-Task Hierarchical Adversarial Inverse Reinforcement Learning.....	4895
<i>Jiayu Chen, Dipesh Tamboli, Tian Lan, Vaneet Aggarwal</i>	
Model Transferability with Responsive Decision Subjects	4921
<i>Yatong Chen, Zeyu Tang, Kun Zhang, Yang Liu</i>	
Layered State Discovery for Incremental Autonomous Exploration	4953
<i>Liyu Chen, Andrea Tirinzoni, Alessandro Lazaric, Matteo Pirotta</i>	
Optimistic Online Mirror Descent for Bridging Stochastic and Adversarial Online Convex Optimization.....	5002
<i>Sijia Chen, Wei-Wei Tu, Peng Zhao, Lijun Zhang</i>	
Learning to Optimize Differentiable Games	5036
<i>Xuxi Chen, Nelson Vadori, Tianlong Chen, Zhangyang Wang</i>	
Coordinated Dynamic Bidding in Repeated Second-Price Auctions with Budgets.....	5052
<i>Yurong Chen, Qian Wang, Zhijian Duan, Haoran Sun, Zhaohua Chen, Xiang Yan, Xiaotie Deng</i>	
Semi-Offline Reinforcement Learning for Optimized Text Generation	5087
<i>Changyu Chen, Xiting Wang, Yiqiao Jin, Victor Ye Dong, Li Dong, Jie Cao, Yi Liu, Rui Yan</i>	
Lower Bounds for Learning in Revealing POMDPs	5104
<i>Fan Chen, Huan Wang, Caiming Xiong, Song Mei, Yu Bai</i>	
Implicit Neural Spatial Representations for Time-Dependent PDEs.....	5162
<i>Honglin Chen, Rundi Wu, Eitan Grinspun, Changxi Zheng, Peter Yichen Chen</i>	
BEATs: Audio Pre-Training with Acoustic Tokenizers	5178
<i>Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, Wanxiang Che, Xiangzhan Yu, Furu Wei</i>	
Learning to Incentivize Information Acquisition: Proper Scoring Rules Meet Principal-Agent Model	5194
<i>Siyu Chen, Jibang Wu, Yifan Wu, Zhuoran Yang</i>	
Faster Gradient-Free Algorithms for Nonsmooth Nonconvex Stochastic Optimization	5219
<i>Lesi Chen, Jing Xu, Luo Luo</i>	
Efficient Personalized Federated Learning Via Sparse Model-Adaptation	5234
<i>Daoyuan Chen, Liuyi Yao, Dawei Gao, Bolin Ding, Yaliang Li</i>	

A Gromov-Wasserstein Geometric View of Spectrum-Preserving Graph Coarsening	5257
<i>Yifan Chen, Rentian Yao, Yun Yang, Jie Chen</i>	
How to Address Monotonicity for Model Risk Management?	5282
<i>Dangxing Chen, Weicheng Ye</i>	
Sketched Ridgeless Linear Regression: The Role of Downsampling	5296
<i>Xin Chen, Yicheng Zeng, Siyue Yang, Qiang Sun</i>	
Context-Aware Bayesian Network Actor-Critic Methods for Cooperative Multi-Agent Reinforcement Learning	5327
<i>Dingyang Chen, Qi Zhang</i>	
Bidirectional Learning for Offline Model-Based Biological Sequence Design	5351
<i>Can Chen, Yingxue Zhang, Xue Liu, Mark Coates</i>	
Learning to Jump: Thinning and Thickening Latent Counts for Generative Modeling	5367
<i>Tianqi Chen, Mingyuan Zhou</i>	
Lifelong Language Pretraining with Distribution-Specialized Experts	5383
<i>Wuyang Chen, Yanqi Zhou, Nan Du, Yanping Huang, James Laudon, Zhifeng Chen, Claire Cui</i>	
Generalized-Smooth Nonconvex Optimization is as Efficient as Smooth Nonconvex Optimization	5396
<i>Ziyi Chen, Yi Zhou, Yingbin Liang, Zhaosong Lu</i>	
Weakly Supervised Regression with Interval Targets	5428
<i>Xin Cheng, Yuzhou Cao, Ximing Li, Bo An, Lei Feng</i>	
PLay: Parametrically Conditioned Layout Generation Using Latent Diffusion	5449
<i>Chin-Yi Cheng, Forrest Huang, Gang Li, Yang Li</i>	
Identification of the Adversary from a Single Adversarial Example	5472
<i>Minhao Cheng, Rui Min, Haochen Sun, Pin-Yu Chen</i>	
Parallel Online Clustering of Bandits Via Hedonic Game	5485
<i>Xiaotong Cheng, Cheng Pan, Setareh Maghsudi</i>	

PART 8

Mu2 SLAM: Multitask, Multilingual Speech and Language Models	5504
<i>Yong Cheng, Yu Zhang, Melvin Johnson, Wolfgang Macherey, Ankur Bapna</i>	
Understanding the Role of Feedback in Online Learning with Switching Costs	5521
<i>Duo Cheng, Xingyu Zhou, Bo Ji</i>	
Tighter Bounds on the Expressivity of Transformer Encoders	5544
<i>David Chiang, Peter Cholak, Anand Pillay</i>	
Provably Learning Diverse Features in Multi-View Data with Midpoint Mixup	5563
<i>Muthu Chidambaram, Xiang Wang, Chenwei Wu, Rong Ge</i>	
Hiding Data Helps: On the Benefits of Masking for Sparse Coding	5600
<i>Muthu Chidambaram, Chenwei Wu, Yu Cheng, Rong Ge</i>	
PINA: Leveraging Side Information in eXtreme Multi-Label Classification Via Predicted Instance Neighborhood Aggregation	5616
<i>Eli Chien, Jiong Zhang, Cho-Jui Hsieh, Jyun-Yu Jiang, Wei-Cheng Chang, Olgica Milenkovic, Hsiang-Fu Yu</i>	
Tight Certification of Adversarially Trained Neural Networks Via Nonconvex Low-Rank Semidefinite Relaxations	5631
<i>Hong-Ming Chiu, Richard Y. Zhang</i>	
Neural Latent Aligner: Cross-Trial Alignment for Learning Representations of Complex, Naturalistic Neural Data	5661
<i>Cheol Jun Cho, Edward Chang, Gopala Anumanchipalli</i>	
On the Convergence of Federated Averaging with Cyclic Client Participation	5677
<i>Yae Jee Cho, Pranay Sharma, Gauri Joshi, Zheng Xu, Satyen Kale, Tong Zhang</i>	

GREAD: Graph Neural Reaction-Diffusion Networks	5722
<i>Jeongwhan Choi, Seoyoung Hong, Noseong Park, Sung-Bae Cho</i>	
Is Overfitting Necessary for Implicit Video Representation?	5748
<i>Hee Min Choi, Hyoa Kang, Dokwan Oh</i>	
Semi-Parametric Contextual Pricing Algorithm Using Cox Proportional Hazards Model.....	5771
<i>Young-Geun Choi, Gi-Soo Kim, Yunseo Choi, Wooseong Cho, Myunghee Cho Paik, Min-Hwan Oh</i>	
Restoration Based Generative Models.....	5787
<i>Jaemoo Choi, Yesom Park, Myungjoo Kang</i>	
Concept-Based Explanations for Out-Of-Distribution Detectors	5817
<i>Jihye Choi, Jayaram Raghuram, Ryan Feng, Jiefeng Chen, Somesh Jha, Atul Prakash</i>	
Active Causal Structure Learning with Advice	5838
<i>Davin Choo, Themistoklis Gouleakis, Arnab Bhattacharyya</i>	
New Metrics and Search Algorithms for Weighted Causal DAGs	5868
<i>Davin Choo, Kirankumar Shiragur</i>	
Computational Doob H-Transforms for Online Filtering of Discretely Observed Diffusions.....	5904
<i>Nicolas Chopin, Andras Fulop, Jeremy Heng, Alexandre H. Thiery</i>	
Multi-Epoch Matrix Factorization Mechanisms for Private Machine Learning.....	5924
<i>Christopher A. Choquette-Choo, Hugh Brendan McMahan, J Keith Rush, Abhradeep Guha Thakurta</i>	
Taming Graph Kernels with Random Features.....	5964
<i>Krzysztof Marcin Choromanski</i>	
Efficient Graph Field Integrators Meet Point Clouds	5978
<i>Krzysztof Marcin Choromanski, Arijit Sehanobish, Han Lin, Yunfan Zhao, Eli Berger, Tetiana Parshakova, Alvin Pan, David Watkins, Tianyi Zhang, Valerii Likhoshesterov, Somnath Basu Roy Chowdhury, Kumar Avinava Dubey, Deepali Jain, Tamas Sarlos, Snigdha Chaturvedi, Adrian Weller</i>	
ContraBAR: Contrastive Bayes-Adaptive Deep RL	6005
<i>Era Choshen, Aviv Tamar</i>	
Forget Unlearning: Towards True Data-Deletion in Machine Learning	6028
<i>Rishav Chourasia, Neil Shah</i>	
Patch-Level Routing in Mixture-Of-Experts is Provably Sample-Efficient for Convolutional Neural Networks	6074
<i>Mohammed Nowaz Rabbani Chowdhury, Shuai Zhang, Meng Wang, Sijia Liu, Pin-Yu Chen</i>	
What Do CNNs Learn in the First Layer and Why? a Linear Systems Perspective	6115
<i>Rhea Chowers, Yair Weiss</i>	
Unifying Molecular and Textual Representations Via Multi-Task Language Modelling	6140
<i>Dimitrios Christofidellis, Giorgio Giannone, Jannis Born, Ole Winther, Teodoro Laino, Matteo Manica</i>	
Wasserstein Barycenter Matching for Graph Size Generalization of Message Passing Neural Networks	6158
<i>Xu Chu, Yujie Jin, Xin Wang, Shanghang Zhang, Yasha Wang, Wenwu Zhu, Hong Mei</i>	
Shape-Guided Dual-Memory Learning for 3D Anomaly Detection.....	6185
<i>Yu-Min Chu, Chieh Liu, Ting-I Hsieh, Hwann-Tzong Chen, Tyng-Luh Liu</i>	
Multiply Robust Off-Policy Evaluation and Learning Under Truncation by Death	6195
<i>Jianing Chu, Shu Yang, Wenbin Lu</i>	
InfoOT: Information Maximizing Optimal Transport	6228
<i>Ching-Yao Chuang, Stefanie Jegelka, David Alvarez-Melis</i>	
A Toy Model of Universality: Reverse Engineering How Networks Learn Group Operations.....	6243
<i>Bilal Chughtai, Lawrence Chan, Neel Nanda</i>	
Distribution Free Prediction Sets for Node Classification.....	6268
<i>Jase Clarkson</i>	

Sequential Strategic Screening.....	6279
<i>Lee Cohen, Saeed Sharifi -Malvajerdi, Kevin Stangl, Ali Vakilian, Juba Ziani</i>	
Few-Sample Feature Selection Via Feature Manifold Learning.....	6296
<i>David Cohen, Tal Shnitzer, Yuval Kluger, Ronen Talmon</i>	

PART 9

Spatial Implicit Neural Representations for Global-Scale Species Mapping	6320
<i>Elijah Cole, Grant Van Horn, Christian Lange, Alexander Shepard, Patrick Leary, Pietro Perona, Scott Loarie, Oisin Mac Aodha</i>	
K-SHAP: Policy Clustering Algorithm for Anonymous Multi-Agent State-Action Pairs.....	6343
<i>Andrea Coletta, Svitlana Vyetrenko, Tucker Balch</i>	
Inferring Relational Potentials in Interacting Systems	6364
<i>Armand Comas, Yilun Du, Christian Fernandez Lopez, Sandesh Ghimire, Mario Sznajder, Joshua B. Tenenbaum, Octavia Camps</i>	
Task-Specific Experimental Design for Treatment Effect Estimation	6384
<i>Bethany Connolly, Kim Moore, Tobias Schwedes, Alexander Adam, Gary Willis, Ilya Feige, Christopher Frye</i>	
A Mathematical Model for Curriculum Learning for Parities	6402
<i>Elisabetta Cornacchia, Elchanan Mossel</i>	
Learning to Maximize Mutual Information for Dynamic Feature Selection	6424
<i>Ian Connick Covert, Wei Qiu, Mingyu Lu, Na Yoon Kim, Nathan J White, Su-In Lee</i>	
Rethinking Weak Supervision in Helping Contrastive Learning.....	6448
<i>Jingyi Cui, Weiran Huang, Yifei Wang, Yisen Wang</i>	
Bayes-Optimal Learning of Deep Random Networks of Extensive-Width.....	6468
<i>Hugo Cui, Florent Krzakala, Lenka Zdeborova</i>	
A General Representation Learning Framework with Generalization Performance Guarantees.....	6522
<i>Junbiao Cui, Jianqing Liang, Qin Yue, Jiye Liang</i>	
IRNeXt: Rethinking Convolutional Network Design for Image Restoration.....	6545
<i>Yuning Cui, Wenqi Ren, Sining Yang, Xiaochun Cao, Alois Knoll</i>	
Scaling Up Dataset Distillation to ImageNet-1K with Constant Memory.....	6565
<i>Justin Cui, Ruochen Wang, Si Si, Cho-Jui Hsieh</i>	
Learning Dynamic Query Combinations for Transformer-Based Object Detection and Segmentation	6591
<i>Yiming Cui, Linjie Yang, Haichao Yu</i>	
Adaptive Identification of Populations with Treatment Benefit in Clinical Trials: Machine Learning Challenges and Solutions	6603
<i>Alicia Curth, Alihan Hüyük, Mihaela Van Der Schaar</i>	
In Search of Insights, Not Magic Bullets: Towards Demystification of the Model Selection Dilemma in Heterogeneous Treatment Effect Estimation	6623
<i>Alicia Curth, Mihaela Van Der Schaar</i>	
Optimal Stochastic Non-Smooth Non-Convex Optimization Through Online-To-Non-Convex Conversion.....	6643
<i>Ashok Cutkosky, Harsh Mehta, Francesco Orabona</i>	
Monge, Bregman and Occam: Interpretable Optimal Transport in High-Dimensions with Feature-Sparse Maps	6671
<i>Marco Cuturi, Michal Klein, Pierre Ablin</i>	
From Noisy Fixed-Point Iterations to Private ADMM for Centralized and Federated Learning.....	6683
<i>Edwige Cyffers, Aurélien Bellet, Debabrota Basu</i>	
Chameleon: Adapting to Peer Images for Planting Durable Backdoors in Federated Learning.....	6712
<i>Yanbo Dai, Songze Li</i>	

Refined Regret for Adversarial MDPs with Linear Function Approximation	6726
<i>Yan Dai, Haipeng Luo, Chen-Yu Wei, Julian Zimmert</i>	
MultiRobustBench: Benchmarking Robustness Against Multiple Attacks	6760
<i>Sihui Dai, Saeed Mahloujifar, Chong Xiang, Vikash Sehwal, Pin-Yu Chen, Prateek Mittal</i>	
Moderately Distributional Exploration for Domain Generalization	6786
<i>Rui Dai, Yonggang Zhang, Zhen Fang, Bo Han, Xinmei Tian</i>	
Trajectory-Aware Eligibility Traces for Off-Policy Reinforcement Learning	6818
<i>Brett Daley, Martha White, Christopher Amato, Marlos C. Machado</i>	
Efficient Displacement Convex Optimization with Particle Gradient Descent	6836
<i>Hadi Daneshmand, Jason D. Lee, Chi Jin</i>	
Multiple Thinking Achieving Meta-Ability Decoupling for Object Navigation	6855
<i>Ronghao Dang, Lu Chen, Liuyi Wang, Zongtao He, Chengju Liu, Qijun Chen</i>	
Neural Collapse in Deep Linear Networks: From Balanced to Imbalanced Data	6873
<i>Hien Dang, Tho Tran Huu, Stanley Osher, Hung The Tran, Nhat Ho, Tan Minh Nguyen</i>	
Reinforcement Learning Can Be More Efficient with Multiple Rewards	6948
<i>Christoph Dann, Yishay Mansour, Mehryar Mohri</i>	
Best of Both Worlds Policy Optimization	6968
<i>Christoph Dann, Chen-Yu Wei, Julian Zimmert</i>	
Image Generation with Shortest Path Diffusion	7009
<i>Ayan Das, Stathi Fotiadis, Anil Batra, Farhang Nabiei, Fengting Liao, Sattar Vakili, Da-Shan Shiu, Alberto Bernacchia</i>	
Efficient List-Decodable Regression Using Batches	7025
<i>Abhimanyu Das, Ayush Jain, Weihao Kong, Rajat Sen</i>	
Beyond Uniform Lipschitz Condition in Differentially Private Optimization	7066
<i>Rudrajit Das, Satyen Kale, Zheng Xu, Tong Zhang, Sujay Sanghavi</i>	
Understanding Self-Distillation in the Presence of Label Noise	7102
<i>Rudrajit Das, Sujay Sanghavi</i>	

PART 10

Interval Bound Interpolation for Few-Shot Learning with Few Tasks	7141
<i>Shounak Datta, Sankha Subhra Mullick, Anish Chakrabarty, Swagatam Das</i>	
Hypervolume Knowledge Gradient: A Lookahead Approach for Multi-Objective Bayesian Optimization with Partial Information.....	7167
<i>Sam Daulton, Maximilian Balandat, Eytan Bakshy</i>	
Fast Combinatorial Algorithms for Min Max Correlation Clustering	7205
<i>Sami Davies, Benjamin Moseley, Heather Newman</i>	
Predictive Flows for Faster Ford-Fulkerson	7231
<i>Sami Davies, Benjamin Moseley, Sergei Vassilvitskii, Yuyan Wang</i>	
The Persistent Laplacian for Data Science: Evaluating Higher-Order Persistent Spectral Representations of Data.....	7249
<i>Thomas Davies, Zhengchao Wan, Ruben J Sanchez-Garcia</i>	
Mitigating Propagation Failures in Physics-Informed Neural Networks Using Retain-Resample-Release (R3) Sampling	7264
<i>Arka Daw, Jie Bu, Sifan Wang, Paris Perdikaris, Anuj Karpatne</i>	
On the Robustness of Randomized Ensembles to Adversarial Perturbations	7303
<i>Hassan Dbouk, Naresh Shanbhag</i>	
Pre-Computed Memory Or On-The-Fly Encoding? a Hybrid Approach to Retrieval Augmentation Makes the Most of Your Compute.....	7329
<i>Michiel De Jong, Yury Zemlyanskiy, Nicholas Fitzgerald, Joshua Ainslie, Sumit Sanghai, Fei Sha, William W. Cohen</i>	

Continuous Spatiotemporal Transformer.....	7343
<i>Antonio Henrique De Oliveira Fonseca, Emanuele Zappala, Josue Ortega Caro, David Van Dijk</i>	
The Value of Out-Of-Distribution Data.....	7366
<i>Ashwin De Silva, Rahul Ramesh, Carey Priebe, Pratik Chaudhari, Joshua T Vogelstein</i>	
High Fidelity Image Counterfactuals with Probabilistic Causal Models.....	7390
<i>Fabio De Sousa Ribeiro, Tian Xia, Miguel Monteiro, Nick Pawlowski, Ben Glocker</i>	
Learning Noisy OR Bayesian Networks with Max-Product Belief Propagation.....	7426
<i>Antoine Dedieu, Guangyao Zhou, Dileep George, Miguel Lazaro-Gredilla</i>	
Learning-Rate-Free Learning by D-Adaptation	7449
<i>Aaron Defazio, Konstantin Mishchenko</i>	
Scaling Vision Transformers to 22 Billion Parameters.....	7480
<i>Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Peter Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, Rodolphe Jenatton, Lucas Beyer, Michael Tschannen, Anurag Arnab, Xiao Wang, Carlos Riquelme Ruiz, Matthias Minderer, Joan Puigcerver, Utku Evci, Manoj Kumar, Sjoerd Van Steenkiste, Gamaleldin Fathy Elsayed, Aravindh Mahendran, Fisher Yu, Avital Oliver, Fantine Huot, Jasmijn Bastings, Mark Collier, Alexey A. Gritsenko, Vighnesh Birodkar, Cristina Nader Vasconcelos, Yi Tay, Thomas Mensink, Alexander Kolesnikov, Filip Pavetic, Dustin Tran, Thomas Kipf, Mario Lucic, Xiaohua Zhai, Daniel Keysers, Jeremiah J. Harmsen, Neil Houlsby</i>	
Efficient Bound of Lipschitz Constant for Convolutional Layers by Gram Iteration.....	7513
<i>Blaise Delattre, Quentin Barthélemy, Alexandre Araujo, Alexandre Allauzen</i>	
Blossom: An Anytime Algorithm for Computing Optimal Decision Trees	7533
<i>Emir Demirovic, Emmanuel Hebrard, Louis Jean</i>	
Optimizing NOTEARS Objectives Via Topological Swaps.....	7563
<i>Chang Deng, Kevin Bello, Bryon Aragam, Pradeep Kumar Ravikumar</i>	
Uncertainty Estimation by Fisher Information-Based Evidential Deep Learning	7596
<i>Danruo Deng, Guangyong Chen, Yang Yu, Furu Liu, Pheng-Ann Heng</i>	
Multi-Channel Autobidding with Budget and ROI Constraints	7617
<i>Yuan Deng, Negin Golrezaei, Patrick Jaillet, Jason Cheuk Nam Liang, Vahab Mirrokni</i>	
Surrogate Module Learning: Reduce the Gradient Error Accumulation in Training Spiking Neural Networks	7645
<i>Shikuang Deng, Hao Lin, Yuhang Li, Shi Gu</i>	
Confidence and Dispersity Speak: Characterizing Prediction Matrix for Unsupervised Accuracy Estimation.....	7658
<i>Weijian Deng, Yumin Suh, Stephen Gould, Liang Zheng</i>	
Great Models Think Alike: Improving Model Reliability Via Inter-Model Latent Agreement	7675
<i>Ailin Deng, Miao Xiong, Bryan Hooi</i>	
Hyperbolic Image-Text Representations	7694
<i>Karan Desai, Maximilian Nickel, Tanmay Rajpurohit, Justin Johnson, Shanmukha Ramakrishna Vedantam</i>	
Hardware-Aware Compression with Random Operation Access Specific Tile (ROAST) Hashing	7732
<i>Aditya Desai, Keren Zhou, Anshumali Shrivastava</i>	
The Case for 4-Bit Precision: K-Bit Inference Scaling Laws.....	7750
<i>Tim Dettmers, Luke Zettlemoyer</i>	
Fairness in Matching Under Uncertainty.....	7775
<i>Siddhartha Devic, David Kempe, Vatsal Sharan, Aleksandra Korolova</i>	
Efficient Parametric Approximations of Neural Network Function Space Distance.....	7795
<i>Nikita Dhawan, Sicong Huang, Juhan Bae, Roger Baker Grosse</i>	
A Large-Scale Study of Probabilistic Calibration in Neural Network Regression	7813
<i>Victor Dheur, Souhaib Ben Taieb</i>	

Nearly Minimax Optimal Regret for Learning Linear Mixture Stochastic Shortest Path	7837
<i>Qiwei Di, Jiafan He, Dongruo Zhou, Quanquan Gu</i>	
On Over-Squashing in Message Passing Neural Networks: The Impact of Width, Depth, and Topology.....	7865
<i>Francesco Di Giovanni, Lorenzo Giusti, Federico Barbero, Giulia Luise, Pietro Lio, Michael M. Bronstein</i>	
Nearly-Linear Time and Streaming Algorithms for Outlier-Robust PCA	7886
<i>Ilias Diakonikolas, Daniel Kane, Ankit Pensia, Thanasis Pittas</i>	
Near-Optimal Cryptographic Hardness of Agnostically Learning Halfspaces and ReLU Regression Under Gaussian Marginals	7922
<i>Ilias Diakonikolas, Daniel Kane, Lisheng Ren</i>	

PART 11

Improving Graph Generation by Restricting Graph Bandwidth.....	7939
<i>Nathaniel Lee Diamant, Alex M Tseng, Kangway V. Chuang, Tommaso Biancalani, Gabriele Scalia</i>	
Forward-Backward Gaussian Variational Inference Via JKO in the Bures-Wasserstein Space.....	7960
<i>Michael Ziyang Diao, Krishna Balasubramanian, Sinho Chewi, Adil Salim</i>	
Subset-Based Instance Optimality in Private Estimation	7992
<i>Travis Dick, Alex Kulesza, Ziteng Sun, Ananda Theertha Suresh</i>	
Pareto Manifold Learning: Tackling Multiple Tasks Via Ensembles of Single-Task Models.....	8015
<i>Nikolaos Dimitriadis, Pascal Frossard, François Fleuret</i>	
Bayesian Reparameterization of Reward-Conditioned Reinforcement Learning with Energy-Based Models.....	8053
<i>Wenhao Ding, Tong Che, Ding Zhao, Marco Pavone</i>	
DSGD-CECA: Decentralized SGD with Communication-Optimal Exact Consensus Algorithm.....	8067
<i>Lisang Ding, Kexin Jin, Bicheng Ying, Kun Yuan, Wotao Yin</i>	
Open-Vocabulary Universal Image Segmentation with MaskCLIP	8090
<i>Zheng Ding, Jieke Wang, Zhuowen Tu</i>	
Entity Divider with Language Grounding in Multi-Agent Reinforcement Learning	8103
<i>Ziluo Ding, Wanpeng Zhang, Junpeng Yue, Xiangjun Wang, Tiejun Huang, Zongqing Lu</i>	
PixelAsParam: A Gradient View on Diffusion Sampling with Guidance.....	8120
<i>Anh-Dung Dinh, Daochang Liu, Chang Xu</i>	
Second-Order Optimization with Lazy Hessians.....	8138
<i>Nikita Doikov, El Mahdi Chayti, Martin Jaggi</i>	
Polynomial Preconditioning for Gradient Methods.....	8162
<i>Nikita Doikov, Anton Rodomanov</i>	
On Data Manifolds Entailed by Structural Causal Models.....	8188
<i>Ricardo Dominguez-Olmedo, Amir-Hossein Karimi, Georgios Arvanitidis, Bernhard Schölkopf</i>	
Towards Understanding and Reducing Graph Structural Noise for GNNs	8202
<i>Mingze Dong, Yuval Kluger</i>	
SpeedDETR: Speed-Aware Transformers for End-To-End Object Detection	8227
<i>Peiyan Dong, Zhenglun Kong, Xin Meng, Peng Zhang, Hao Tang, Yanzhi Wang, Chih-Hsien Chou</i>	
Understand and Modularize Generator Optimization in ELECTRA-Style Pretraining.....	8244
<i>Chengyu Dong, Liyuan Liu, Hao Cheng, Jingbo Shang, Jianfeng Gao, Xiaodong Liu</i>	
Diversity-Enhancing Generative Network for Few-Shot Hypothesis Adaptation	8260
<i>Ruijiang Dong, Feng Liu, Haoang Chi, Tongliang Liu, Mingming Gong, Gang Niu, Masashi Sugiyama, Bo Han</i>	

PASTA: Pessimistic Assortment Optimization.....	8276
<i>Juncheng Dong, Weibin Mo, Zhengling Qi, Cong Shi, Ethan X Fang, Vahid Tarokh</i>	
Adaptively Weighted Data Augmentation Consistency Regularization for Robust Optimization Under Concept Shift.....	8296
<i>Yijun Dong, Yuege Xie, Rachel Ward</i>	
Does Sparsity Help in Learning Misspecified Linear Bandits?.....	8317
<i>Jialin Dong, Lin Yang</i>	
Symmetry-Aware Robot Design with Structured Subgroups	8334
<i>Heng Dong, Junyu Zhang, Tonghan Wang, Chongjie Zhang</i>	
DoCoFL: Downlink Compression for Cross-Device Federated Learning.....	8356
<i>Ron Dorfman, Shay Vargaftik, Yaniv Ben-Itzhak, Kfir Yehuda Levy</i>	
Meta-Learning the Inductive Bias of Simple Neural Circuits	8389
<i>Will Dorrell, Maria Yuffa, Peter E. Latham</i>	
Self-Repellent Random Walks on General Graphs - Achieving Minimal Sampling Variance Via Nonlinear Markov Chains	8403
<i>Vishwaraj Doshi, Jie Hu, Do Young Eun</i>	
Linear Time GPs for Inferring Latent Trajectories from Neural Spike Trains.....	8424
<i>Matthew Dowling, Yuan Zhao, Il Memming Park</i>	
On the Convergence Rate of Gaussianization with Random Rotations.....	8449
<i>Felix Draxler, Lars Kühmichel, Armand Rousselot, Jens Müller, Christoph Schnoerr, Ullrich Koethe</i>	
PaLM-E: An Embodied Multimodal Language Model	8469
<i>Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, Pete Florence</i>	
Reduce, Reuse, Recycle: Compositional Generation with Energy-Based Diffusion Models and MCMC	8489
<i>Yilun Du, Conor Durkan, Robin Strudel, Joshua B. Tenenbaum, Sander Dieleman, Rob Fergus, Jascha Sohl-Dickstein, Arnaud Doucet, Will Sussman Grathwohl</i>	
Multi-Task Representation Learning for Pure Exploration in Linear Bandits	8511
<i>Yihan Du, Longbo Huang, Wen Sun</i>	
Nonparametric Generative Modeling with Conditional Sliced-Wasserstein Flows.....	8565
<i>Chao Du, Tianbo Li, Tianyu Pang, Shuicheng Yan, Min Lin</i>	
Subsample Ridge Ensembles: Equivalences and Generalized Cross-Validation.....	8585
<i>Jin-Hong Du, Pratik Patil, Arun K. Kuchibhotla</i>	
On Uni-Modal Feature Learning in Supervised Multi-Modal Learning.....	8632
<i>Chenzhuang Du, Jiaye Teng, Tingle Li, Yichen Liu, Tianyuan Yuan, Yue Wang, Yang Yuan, Hang Zhao</i>	
Guiding Pretraining in Reinforcement Learning with Large Language Models	8657
<i>Yuqing Du, Olivia Watkins, Zihan Wang, Cédric Colas, Trevor Darrell, Pieter Abbeel, Abhishek Gupta, Jacob Andreas</i>	
A Flexible Diffusion Model.....	8678
<i>Weitao Du, He Zhang, Tao Yang, Yuanqi Du</i>	
Fast Excess Risk Rates Via Offset Rademacher Complexity	8697
<i>Chenguang Duan, Yuling Jiao, Lican Kang, Xiliang Lu, Jerry Zhijian Yang</i>	
Are Diffusion Models Vulnerable to Membership Inference Attacks?.....	8717
<i>Jinhao Duan, Fei Kong, Shiqi Wang, Xiaoshuang Shi, Kaidi Xu</i>	
Bayesian Progressive Deep Topic Model with Knowledge Informed Textual Data Coarsening Process.....	8731
<i>Zhibin Duan, Xinyang Liu, Yudi Su, Yishi Xu, Bo Chen, Mingyuan Zhou</i>	

PART 12

Are Equivariant Equilibrium Approximators Beneficial?	8747
<i>Zhijian Duan, Yunxuan Ma, Xiaotie Deng</i>	
Evaluating Self-Supervised Learning Via Risk Decomposition	8779
<i>Yann Dubois, Tatsunori Hashimoto, Percy Liang</i>	
Fully Dynamic Submodular Maximization Over Matroids	8821
<i>Paul Duetting, Federico Fusco, Silvio Lattanzi, Ashkan Norouzi-Fard, Morteza Zadimoghaddam</i>	
Optimal No-Regret Learning for One-Sided Lipschitz Functions	8836
<i>Paul Duetting, Guru Guruganesh, Jon Schneider, Joshua Ruizhi Wang</i>	
Integrating Prior Knowledge in Contrastive Learning with Kernel	8851
<i>Benoit Dufumier, Carlo Alberto Barbano, Robin Louiset, Edouard Duchesnay, Pietro Gori</i>	
Q-Flow: Generative Modeling for Differential Equations of Open Quantum Dynamics with Normalizing Flows	8879
<i>Owen M Dugan, Peter Y. Lu, Rumen Dangovski, Di Luo, Marin Soljacic</i>	
Adaptive Whitening in Neural Populations with Gain-Modulating Interneurons	8902
<i>Lyndon Duong, David Lipshutz, David Heeger, Dmitri Chklovskii, Eero P Simoncelli</i>	
Generalization Bounds Using Data-Dependent Fractal Dimensions	8922
<i>Benjamin Dupuis, George Deligiannidis, Umut Simsekli</i>	
Multi-Objective Population Based Training	8969
<i>Arkadiy Dushatskiy, Alexander Chebykin, Tanja Alderliesten, Peter Bosman</i>	
Neural Diffusion Processes	8990
<i>Vincent Dutordoir, Alan Saul, Zoubin Ghahramani, Fergus Simpson</i>	
FAENet: Frame Averaging Equivariant GNN for Materials Modeling	9013
<i>Alexandre Agm Duval, Victor Schmidt, Alex Hernández-Garci'a, Santiago Miret, Fragkiskos D. Malliaros, Yoshua Bengio, David Rolnick</i>	
Blackout Diffusion: Generative Diffusion Models in Discrete-State Spaces	9034
<i>Javier E. Santos, Zachary R. Fox, Nicholas Lubbers, Yen Ting Lin</i>	
The Computational Complexity of Concise Hypersphere Classification	9060
<i>Eduard Eiben, Robert Ganian, Iyad A. Kanj, Sebastian Ordyniak, Stefan Szeider</i>	
E(n) Equivariant Message Passing Simplicial Networks	9071
<i>Floor Eijkelboom, Rob Hesselink, Erik J Bekkers</i>	
Performative Recommendation: Diversifying Content Via Strategic Incentives	9082
<i>Itay Eilat, Nir Rosenfeld</i>	
Hyperparameters in Reinforcement Learning and How to Tune Them	9104
<i>Theresa Eimer, Marius Lindauer, Roberta Raileanu</i>	
Fairness in Streaming Submodular Maximization Over a Matroid Constraint	9150
<i>Marwa El Halabi, Federico Fusco, Ashkan Norouzi-Fard, Jakab Tardos, Jakub Tarnawski</i>	
Difference of Submodular Minimization Via DC Programming	9172
<i>Marwa El Halabi, George Orfanides, Tim Hoheisel</i>	
Graph Positional Encoding Via Random Feature Propagation	9202
<i>Moshe Eliasof, Fabrizio Frasca, Beatrice Bevilacqua, Eran Treister, Gal Chechik, Haggai Maron</i>	
Improving Graph Neural Networks with Learnable Propagation Operators	9224
<i>Moshe Eliasof, Lars Ruthotto, Eran Treister</i>	
Phase Transitions in the Detection of Correlated Databases	9246
<i>Dor Elimelech, Wasim Huleihel</i>	
A New Near-Linear Time Algorithm for K-Nearest Neighbor Search Using a Compressed Cover Tree	9267
<i>Yury Elkin, Vitaliy Kurlin</i>	

Motion Question Answering Via Modular Motion Programs.....	9312
<i>Mark Endo, Joy Hsu, Jiaman Li, Jiajun Wu</i>	
Learning Perturbations to Explain Time Series Predictions	9329
<i>Joseph Enguehard</i>	
Regret Minimization and Convergence to Equilibria in General-Sum Markov Games	9343
<i>Liad Erez, Tal Lincewicz, Uri Sherman, Tomer Koren, Yishay Mansour</i>	
Delayed Bandits: When Do Intermediate Observations Help?.....	9374
<i>Emmanuel Esposito, Saeed Masoudian, Hao Qiu, Dirk Van Der Hoeven, Nicolò Cesa-Bianchi, Yevgeny Seldin</i>	
Scaling Spherical CNNs	9396
<i>Carlos Esteves, Jean-Jacques Slotine, Ameesh Makadia</i>	
Stochastic Gradient Descent Under Markovian Sampling Schemes	9412
<i>Mathieu Even</i>	
Continual Learning in Linear Classification on Separable Data.....	9440
<i>Itay Evron, Edward Moroshko, Gon Buzaglo, Maroun Khriesh, Badea Marjeh, Nathan Srebro, Daniel Soudry</i>	
A Connection Between One-Step RL and Critic Regularization in Reinforcement Learning	9485
<i>Benjamin Eysenbach, Matthieu Geist, Sergey Levine, Ruslan Salakhutdinov</i>	
Neural Status Registers.....	9508
<i>Lukas Faber, Roger Wattenhofer</i>	
Learning Rate Schedules in the Presence of Distribution Shift	9523
<i>Matthew Fahrback, Adel Javanmard, Vahab Mirrokni, Pratik Worah</i>	

PART 13

Predicting Rare Events by Shrinking Towards Proportional Odds	9547
<i>Gregory Faletto, Jacob Bien</i>	
Free-Form Variational Inference for Gaussian Process State-Space Models.....	9603
<i>Xuhui Fan, Edwin V. Bonilla, Terence O’Kane, Scott A Sisson</i>	
Optimizing DDPM Sampling with Shortcut Fine-Tuning	9623
<i>Ying Fan, Kangwook Lee</i>	
LSDS++ : Dual Sampling for Accelerated K-Means++	9640
<i>Chenglin Fan, Ping Li, Xiaoyun Li</i>	
Smart Initial Basis Selection for Linear Programs	9650
<i>Zhenan Fan, Xinglu Wang, Oleksandr Yakovenko, Abdullah Ali Sivas, Owen Ren, Yong Zhang, Zirui Zhou</i>	
General Covariance Data Augmentation for Neural PDE Solvers	9665
<i>Vladimir Fanaskov, Tianchi Yu, Alexander Rudikov, Ivan Oseledets</i>	
The Fast Johnson-Lindenstrauss Transform is Even Faster.....	9689
<i>Ora Nova Fandina, Mikael Møller Høgsgaard, Kasper Green Larsen</i>	
Regression with Label Permutation in Generalized Linear Model.....	9716
<i>Guanhua Fang, Ping Li</i>	
Robust Collaborative Learning with Linear Gradient Overhead.....	9761
<i>Sadegh Farhadkhani, Rachid Guerraoui, Nirupam Gupta, Lê-Nguyễn Hoang, Rafael Pinot, John Stephan</i>	
Neural FIM for Learning Fisher Information Metrics from Point Cloud Data.....	9814
<i>Oluwadamilola Fasina, Guillaume Huguet, Alexander Tong, Yanlei Zhang, Guy Wolf, Maximilian Nickel, Ian Adelman, Smita Krishnaswamy</i>	
Stochastic Policy Gradient Methods: Improved Sample Complexity for Fisher-Non-Degenerate Policies	9827
<i>Ilyas Fatkhullin, Anas Barakat, Anastasia Kireeva, Niao He</i>	

Parallel Neurosymbolic Integration with Concordia	9870
<i>Jonathan Feldstein, Modestas Jurcius, Efthymia Tsamoura</i>	
Why Target Networks Stabilise Temporal Difference Methods	9886
<i>Mattie Fellows, Matthew J. A. Smith, Shimon Whiteson</i>	
Weighted Sampling Without Replacement for Deep Top-K Classification	9910
<i>Dieqiao Feng, Yuanqi Du, Carla P Gomes, Bart Selman</i>	
Improved Online Learning Algorithms for CTR Prediction in Ad Auctions	9921
<i>Zhe Feng, Christopher Liaw, Zixin Zhou</i>	
Fractional Denoising for 3D Molecular Pre-Training.....	9938
<i>Shikun Feng, Yuyan Ni, Yanyan Lan, Zhi-Ming Ma, Wei-Ying Ma</i>	
Improved Algorithms for White-Box Adversarial Streams	9962
<i>Ying Feng, David Woodruff</i>	
Non-Stationary Reinforcement Learning Under General Function Approximation.....	9976
<i>Songtao Feng, Ming Yin, Ruiquan Huang, Yu-Xiang Wang, Jing Yang, Yingbin Liang</i>	
Random Matrix Analysis to Balance Between Supervised and Unsupervised Learning Under the Low Density Separation Assumption	10008
<i>Vasilii Feofanov, Malik Tiomoko, Aladin Virmaux</i>	
SurCo: Learning Linear SURrogates for COmbinatorial Nonlinear Optimization Problems.....	10034
<i>Aaron M Ferber, Taoan Huang, Daochen Zha, Martin Schubert, Benoit Steiner, Bistra Dilkina, Yuandong Tian</i>	
Scaling Laws for Multilingual Neural Machine Translation	10053
<i>Patrick Fernandes, Behrooz Ghorbani, Xavier Garcia, Markus Freitag, Orhan Firat</i>	
Constant Matters: Fine-Grained Error Bound on Differentially Private Continual Observation.....	10072
<i>Hendrik Fichtenberger, Monika Henzinger, Jalaj Upadhyay</i>	
Adapting to Game Trees in Zero-Sum Imperfect Information Games	10093
<i>Côme Fiegel, Pierre Menard, Tadashi Kozuno, Remi Munos, Vianney Perchet, Michal Valko</i>	
User-Defined Event Sampling and Uncertainty Quantification in Diffusion Models for Physical Dynamical Systems	10136
<i>Marc Anton Finzi, Anudhyan Boral, Andrew Gordon Wilson, Fei Sha, Leonardo Zepeda- Nunez</i>	
ACAT: Adversarial Counterfactual Attention for Classification and Detection in Medical Imaging	10153
<i>Alessandro Fontanella, Antreas Antoniou, Wenwen Li, Joanna Wardlaw, Grant Mair, Emanuele Trucco, Amos Storkey</i>	
Explainable Data-Driven Optimization: From Context to Decision and Back Again	10170
<i>Alexandre Forel, Axel Parmentier, Thibaut Vidal</i>	
Hardness of Independent Learning and Sparse Equilibrium Computation in Markov Games	10188
<i>Dylan J Foster, Noah Golowich, Sham M. Kakade</i>	
Disentangled Generative Models for Robust Prediction of System Dynamics	10222
<i>Stathi Fotiadis, Mario Lino Valencia, Shunlong Hu, Stef Garasto, Chris D Cantwell, Anil Anthony Bharath</i>	
Can Forward Gradient Match Backpropagation?	10249
<i>Louis Fournier, Stephane Rivaud, Eugene Belilovsky, Michael Eickenberg, Edouard Oyallon</i>	
Last Switch Dependent Bandits with Monotone Payoff Functions	10265
<i>Ayoub Foussoul, Vineet Goyal, Orestis Papadigenopoulos, Assaf Zeevi</i>	
A Theoretical Analysis of the Learning Dynamics Under Class Imbalance	10285
<i>Emanuele Francazi, Marco Baity-Jesi, Aurelien Lucchi</i>	
SparseGPT: Massive Language Models Can Be Accurately Pruned in One-Shot.....	10323
<i>Elias Frantar, Dan Alistarh</i>	
Learning Temporally AbstractWorld Models Without Online Experimentation	10338
<i>Benjamin Freed, Siddarth Venkatraman, Guillaume Adrien Sartoretti, Jeff Schneider, Howie Choset</i>	

PART 14

A Coupled Flow Approach to Imitation Learning	10357
<i>Gideon Joseph Freund, Elad Sarafian, Sarit Kraus</i>	
Simple Hardware-Efficient Long Convolutions for Sequence Modeling.....	10373
<i>Daniel Y Fu, Elliot L Epstein, Eric Nguyen, Armin W Thomas, Michael Zhang, Tri Dao, Atri Rudra, Christopher Re</i>	
MonoNeRF: Learning Generalizable NeRFs from Monocular Videos Without Camera Poses	10392
<i>Yang Fu, Ishan Misra, Xiaolong Wang</i>	
Go Beyond Imagination: Maximizing Episodic Reachability with World Models	10405
<i>Yao Fu, Run Peng, Honglak Lee</i>	
Specializing Smaller Language Models Towards Multi-Step Reasoning.....	10421
<i>Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, Tushar Khot</i>	
Accelerated Stochastic Optimization Methods Under Quasar-Convexity.....	10431
<i>Qiang Fu, Dongchu Xu, Ashia Camage Wilson</i>	
Meta-Learning Parameterized Skills	10461
<i>Haotian Fu, Shangqun Yu, Saket Tiwari, Michael Littman, George Konidaris</i>	
NeRFool: Uncovering the Vulnerability of Generalizable Neural Radiance Fields Against Adversarial Perturbations	10482
<i>Yonggan Fu, Ye Yuan, Souvik Kundu, Shang Wu, Shunyao Zhang, Yingyan Celine Lin</i>	
Hierarchies of Reward Machines	10494
<i>Daniel Furelos-Blanco, Mark Law, Anders Jonsson, Krystia Broda, Alessandra Russo</i>	
Why Random Pruning is All We Need to Start Sparse.....	10542
<i>Advait Harshal Gadhikar, Sohom Mukherjee, Rebekka Burkholz</i>	
Cell-Free Latent Go-Explore	10571
<i>Quentin Gallouédec, Emmanuel Dellandrea</i>	
Graph Reinforcement Learning for Network Control Via Bi-Level Optimization.....	10587
<i>Daniele Gammelli, James Harrison, Kaidi Yang, Marco Pavone, Filipe Rodrigues, Francisco C. Pereira</i>	
Why is Public Pretraining Necessary for Private Model Training?.....	10611
<i>Arun Ganesh, Mahdi Haghifam, Milad Nasr, Sewoong Oh, Thomas Steinke, Om Thakkar, Abhradeep Guha Thakurta, Lun Wang</i>	
Do Perceptually Aligned Gradients Imply Robustness?.....	10628
<i>Roy Ganz, Bahjat Kawar, Michael Elad</i>	
Solving Linear Programs with Fast Online Learning Algorithms	10649
<i>Wenzhi Gao, Dongdong Ge, Chunlin Sun, Yinyu Ye</i>	
Gradient Descent Finds the Global Optima of Two-Layer Physics-Informed Neural Networks.....	10676
<i>Yihang Gao, Yiqi Gu, Michael Ng</i>	
Generalizing Neural Wave Functions.....	10708
<i>Nicholas Gao, Stephan Günnemann</i>	
On the Impact of Algorithmic Recourse on Social Segregation	10727
<i>Ruijiang Gao, Himabindu Lakkaraju</i>	
DDGR: Continual Learning with Deep Diffusion-Based Generative Replay	10744
<i>Rui Gao, Weiwei Liu</i>	
PAL: Program-Aided Language Models	10764
<i>Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, Graham Neubig</i>	
Out-Of-Domain Robustness Via Targeted Augmentations	10800
<i>Irena Gao, Shiori Sagawa, Pang Wei Koh, Tatsunori Hashimoto, Percy Liang</i>	
Scaling Laws for Reward Model Overoptimization.....	10835
<i>Leo Gao, John Schulman, Jacob Hilton</i>	

The Unreasonable Effectiveness of Few-Shot Learning for Machine Translation	10867
<i>Xavier Garcia, Yamini Bansal, Colin Cherry, George Foster, Maxim Krikun, Melvin Johnson, Orhan Firat</i>	
RLSbench: Domain Adaptation Under Relaxed Label Shift	10879
<i>Saurabh Garg, Nick Erickson, James Sharpnack, Alex Smola, Sivaraman Balakrishnan, Zachary Chase Lipton</i>	
RankMe: Assessing the Downstream Performance of Pretrained Self-Supervised Representations by Their Rank	10929
<i>Quentin Garrido, Randall Balestriero, Laurent Najman, Yann Lecun</i>	
Self-Supervised Learning of Split Invariant Equivariant Representations	10975
<i>Quentin Garrido, Laurent Najman, Yann Lecun</i>	
Federated Heavy Hitter Recovery Under Linear Sketching	10997
<i>Adria Gascon, Peter Kairouz, Ziteng Sun, Ananda Theertha Suresh</i>	
On the Global Convergence of Fitted Q-Iteration with Two-Layer Neural Network Parametrization	11013
<i>Mudit Gaur, Vaneet Aggarwal, Mridul Agarwal</i>	
A Reinforcement Learning Framework for Dynamic Mediation Analysis	11050
<i>Lin Ge, Jitao Wang, Chengchun Shi, Zhenke Wu, Rui Song</i>	
Compositional Score Modeling for Simulation-Based Inference	11098
<i>Tomas Geffner, George Papamakarios, Andriy Mnih</i>	
Cramming: Training a Language Model on a Single GPU in One Day	11117
<i>Jonas Geiping, Tom Goldstein</i>	
Transformers Meet Directed Graphs	11144
<i>Simon Geisler, Yujia Li, Daniel J Mankowitz, Ali Taylan Cemgil, Stephan Günnemann, Cosmin Paduraru</i>	

PART 15

Memory-Based Meta-Learning on Non-Stationary Distributions	11173
<i>Tim Genewein, Gregoire Deletang, Anian Ruoss, Li Kevin Wenliang, Elliot Catt, Vincent Dutoit, Jordi Grau-Moya, Laurent Orseau, Marcus Hutter, Joel Veness</i>	
Towards Reliable Neural Specifications	11196
<i>Chuqin Geng, Nham Le, Xiaojie Xu, Zhaoyue Wang, Arie Gurfinkel, Xujie Si</i>	
Oracles & Followers: Stackelberg Equilibria in Deep Multi-Agent Reinforcement Learning	11213
<i>Matthias Gerstgrasser, David C. Parkes</i>	
Approximately Optimal Core Shapes for Tensor Decompositions	11237
<i>Mehrdad Ghadiri, Matthew Fahrback, Gang Fu, Vahab Mirrokni</i>	
GAT: Guided Adversarial Training with Pareto-Optimal Auxiliary Tasks	11255
<i>Salah Ghamizi, Jingfeng Zhang, Maxime Cordy, Mike Papadakis, Masashi Sugiyama, Yves Le Traon</i>	
On User-Level Private Convex Optimization	11283
<i>Badi Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, Raghu Meka, Chiyuan Zhang</i>	
Contextual Reliability: When Different Features Matter in Different Contexts	11300
<i>Gaurav Rohit Ghosal, Amrith Setlur, Daniel S. Brown, Anca Dragan, Aditi Raghunathan</i>	
Reinforcement Learning from Passive Data Via Latent Intentions	11321
<i>Diby Ghosh, Chethan Anand Bhateja, Sergey Levine</i>	
Harmonic Neural Networks	11340
<i>Atiyo Ghosh, Antonio Andrea Gentile, Mario Dagrada, Chul Lee, Seong-Hyok Sean Kim, Hyukgeun Cha, Yunjun Choi, Dongho Kim, Jeong-Il Kye, Vincent Emanuel Elfving</i>	
Dividing and Conquering a BlackBox to a Mixture of Interpretable Models: Route, Interpret, Repeat	11360
<i>Shantanu Ghosh, Ke Yu, Forough Arabshahi, Kayhan Batmanghelich</i>	

Looped Transformers as Programmable Computers	11398
<i>Angeliki Giannou, Shashank Rajput, Jy-Yong Sohn, Kangwook Lee, Jason D. Lee, Dimitris Papailiopoulos</i>	
Generalized Disparate Impact for Configurable Fairness Solutions in ML.....	11443
<i>Luca Giuliani, Eleonora Misino, Michele Lombardi</i>	
Multicalibration as Boosting for Regression	11459
<i>Ira Globus-Harris, Declan Harrison, Michael Kearns, Aaron Roth, Jessica Sorrell</i>	
Adversarial Robustness of Amortized Bayesian Inference.....	11493
<i>Manuel Gloeckler, Michael Deistler, Jakob H. Macke</i>	
Efficient RL Via Disentangled Environment and Agent Representations.....	11525
<i>Kevin Gmelin, Shikhar Bahl, Russell Mendonca, Deepak Pathak</i>	
Aligning Language Models with Preferences Through F-Divergence Minimization.....	11546
<i>Dongyoung Go, Tomasz Korbak, Germán Kruszewski, Jos Rozen, Nahyeon Ryu, Marc Dymetman</i>	
Robust Consensus in Ranking Data Analysis: Definitions, Properties and Computational Issues	11584
<i>Morgane Goibert, Clément Calauzènes, Ekhine Irurozki, Stephan Cléménçon</i>	
Learning Distributions Over Quantum Measurement Outcomes	11598
<i>Weiyuan Gong, Scott Aaronson</i>	
Convergence of Proximal Point and Extragradient-Based Methods Beyond Monotonicity: The Case of Negative Comonotonicity	11614
<i>Eduard Gorbunov, Adrien Taylor, Samuel Horváth, Gauthier Gidel</i>	
Adaptive Annealed Importance Sampling with Constant Rate Progress	11642
<i>Shirin Goshtasbpour, Victor Cohen, Fernando Perez-Cruz</i>	
Formalizing Preferences Over Runtime Distributions.....	11659
<i>Devon R. Graham, Kevin Leyton-Brown, Tim Roughgarden</i>	
Topological Point Cloud Clustering	11683
<i>Vincent Peter Grande, Michael T Schaub</i>	
On Sampling with Approximate Transport Maps	11698
<i>Louis Grenioux, Alain Oliviero Durmus, Eric Moulines, Marylou Gabrié</i>	
Hidden Symmetries of ReLU Networks.....	11734
<i>Elisenda Grigsby, Kathryn Lindsey, David Rolnick</i>	
EF21-P and Friends: Improved Theoretical Communication Complexity for Distributed Optimization with Bidirectional Compression	11761
<i>Kaja Gruntkowska, Alexander Tyurin, Peter Richtárik</i>	
NerfDiff: Single-Image View Synthesis with NeRF-Guided Distillation from 3D-Aware Diffusion	11808
<i>Jiatao Gu, Alex Trevithick, Kai-En Lin, Joshua M. Susskind, Christian Theobalt, Lingjie Liu, Ravi Ramamoorthi</i>	
DecompDiff: Diffusion Models with Decomposed Priors for Structure-Based Drug Design	11827
<i>Jiaqi Guan, Xiangxin Zhou, Yuwei Yang, Yu Bao, Jian Peng, Jianzhu Ma, Qiang Liu, Liang Wang, Quanquan Gu</i>	
On Excess Mass Behavior in Gaussian Mixture Models with Orlicz-Wasserstein Distances	11847
<i>Aritra Guha, Nhat Ho, Xuanlong Nguyen</i>	
Conformalization of Sparse Generalized Linear Models.....	11871
<i>Etash Kumar Guha, Eugene Ndiaye, Xiaoming Huo</i>	
Privacy-Aware Compression for Federated Learning Through Numerical Mechanism Design.....	11888
<i>Chuan Guo, Kamalika Chaudhuri, Pierre Stock, Michael Rabbat</i>	
Out-Of-Distribution Generalization of Federated Learning Via Implicit Invariant Relationships	11905
<i>Yaming Guo, Kai Guo, Xiaofeng Cao, Tieru Wu, Yi Chang</i>	
FeDXL: Provable Federated Learning for Deep X-Risk Optimization	11934
<i>Zhishuai Guo, Rong Jin, Jiebo Luo, Tianbao Yang</i>	
Provably Efficient Representation Learning with Tractable Planning in Low-Rank POMDP	11967
<i>Jiacheng Guo, Zihao Li, Huazheng Wang, Mengdi Wang, Zhuoran Yang, Xuezhou Zhang</i>	

PART 16

Analyzing Privacy Leakage in Machine Learning Via Multiple Hypothesis Testing: A Lesson from Fano.....	11998
<i>Chuan Guo, Alexandre Sablayrolles, Maziar Sanjabi</i>	
Linkless Link Prediction Via Relational Distillation.....	12012
<i>Zhichun Guo, William Shiao, Shichang Zhang, Yozen Liu, Nitesh V Chawla, Neil Shah, Tong Zhao</i>	
FedBR: Improving Federated Learning on Heterogeneous Data Via Local Learning Bias Reduction	12034
<i>Yongxin Guo, Xiaoying Tang, Tao Lin</i>	
Hierarchical Grammar-Induced Geometry for Data-Efficient Molecular Property Prediction.....	12055
<i>Minghao Guo, Veronika Thost, Samuel W Song, Adithya Balachandran, Payel Das, Jie Chen, Wojciech Matusik</i>	
Graph Neural Networks with Learnable and Optimal Polynomial Bases.....	12077
<i>Yuhe Guo, Zhewei Wei</i>	
LongCoder: A Long-Range Pre-Trained Language Model for Code Completion.....	12098
<i>Daya Guo, Canwen Xu, Nan Duan, Jian Yin, Julian McAuley</i>	
Estimating Heterogeneous Treatment Effects: Mutual Information Bounds and Learning Algorithms.....	12108
<i>Xingzhuo Guo, Yuchen Zhang, Jianmin Wang, Mingsheng Long</i>	
Identifying Useful Learnwares for Heterogeneous Label Spaces.....	12122
<i>Lan-Zhe Guo, Zhi Zhou, Yu-Feng Li, Zhi-Hua Zhou</i>	
High-Dimensional Location Estimation Via Norm Concentration for Subgamma Vectors.....	12132
<i>Shivam Gupta, Jasper C. H. Lee, Eric Price</i>	
GRAFENNE: Learning on Graphs with Heterogeneous and Dynamic Feature Sets	12165
<i>Shubham Gupta, Sahil Manchanda, Sayan Ranu, Srikantha J. Bedathur</i>	
Online Platt Scaling with Calibrating	12182
<i>Chirag Gupta, Aaditya Ramdas</i>	
Multi-Task Structural Learning Using Local Task Similarity Induced Neuron Creation and Removal	12205
<i>Naresh Kumar Gurulingan, Bahram Zonooz, Elahe Arani</i>	
Conditionally Strongly Log-Concave Generative Models	12224
<i>Florentin Guth, Etienne Lempereur, Joan Bruna, Stéphane Mallat</i>	
DRew: Dynamically Rewired Message Passing with Delay	12252
<i>Benjamin Gutteridge, Xiaowen Dong, Michael M. Bronstein, Francesco Di Giovanni</i>	
Kernel Logistic Regression Approximation of an Understandable ReLU Neural Network	12268
<i>Marie Guyomard, Susana Barbosa, Lionel Fillatre</i>	
Conformal Prediction Sets for Graph Neural Networks	12292
<i>Soroush H. Zargarbashi, Simone Antonelli, Aleksandar Bojchevski</i>	
Social Learning Spontaneously Emerges by Searching Optimal Heuristics with Deep Reinforcement Learning.....	12319
<i>Seungwoong Ha, Hawoong Jeong</i>	
Convex Geometry of ReLU-Layers, Injectivity on the Ball and Local Reconstruction.....	12339
<i>Daniel Haider, Martin Ehler, Peter Balazs</i>	
Robust Counterfactual Explanations for Neural Networks with Probabilistic Guarantees.....	12351
<i>Faisal Hamman, Erfaun Noorani, Saumitra Mishra, Daniele Magazzeni, Sanghamitra Dutta</i>	
Wrapped Cauchy Distributed Angular Softmax for Long-Tailed Visual Recognition.....	12368
<i>Boran Han</i>	
On the Impact of Knowledge Distillation for Model Interpretability	12389
<i>Hyeonrok Han, Siwon Kim, Hyun-Soo Choi, Sungroh Yoon</i>	

Alternately Optimized Graph Neural Networks	12411
<i>Haoyu Han, Xiaorui Liu, Haitao Mao, Mohamadali Torkamani, Feng Shi, Victor Lee, Jiliang Tang</i>	
System Identification of Neural Systems: If We Got it Right, Would We Know?.....	12430
<i>Yena Han, Tomaso A Poggio, Brian Cheung</i>	
Total Variation Graph Neural Networks	12445
<i>Jonas Berg Hansen, Filippo Maria Bianchi</i>	
Learning Physical Models that Can Respect Conservation Laws.....	12469
<i>Derek Hansen, Danielle C. Maddix, Shima Alizadeh, Gaurav Gupta, Michael W. Mahoney</i>	
On Pre-Training for Visuo-Motor Control: Revisiting a Learning-From-Scratch Baseline	12511
<i>Nicklas Hansen, Zhecheng Yuan, Yanjie Ze, Tongzhou Mu, Aravind Rajeswaran, Hao Su, Huazhe Xu, Xiaolong Wang</i>	
Leveraging Demonstrations to Improve Online Learning: Quality Matters	12527
<i>Botao Hao, Rahul Jain, Tor Lattimore, Benjamin Van Roy, Zheng Wen</i>	
Coupled Variational Autoencoder.....	12546
<i>Xiaoran Hao, Patrick Shafto</i>	
GNOT: A General Neural Operator Transformer for Operator Learning	12556
<i>Zhongkai Hao, Zhengyi Wang, Hang Su, Chengyang Ying, Yinpeng Dong, Songming Liu, Ze Cheng, Jian Song, Jun Zhu</i>	
Algorithmic Collective Action in Machine Learning	12570
<i>Moritz Hardt, Eric Mazumdar, Celestine Mender-Dünner, Tijana Zrnic</i>	
Gaussian Process Priors for Systems of Linear Partial Differential Equations with Constant Coefficients	12587
<i>Marc Harkonen, Markus Lange-Hegermann, Bogdan Raita</i>	
Theoretical Guarantees of Learning Ensembling Strategies with Applications to Time Series Forecasting	12616
<i>Hilaf Hasson, Danielle C. Maddix, Bernie Wang, Gaurav Gupta, Youngsuk Park</i>	
Global Context Vision Transformers	12633
<i>Ali Hatamizadeh, Hongxu Yin, Greg Heinrich, Jan Kautz, Pavlo Molchanov</i>	
Counterfactual Analysis in Dynamic Latent State Models	12647
<i>Martin B Haugh, Raghav Singal</i>	
Sampling-Based Nyström Approximation and Kernel Quadrature	12678
<i>Satoshi Hayakawa, Harald Oberhauser, Terry Lyons</i>	
Width and Depth Limits Commute in Residual Networks	12700
<i>Soufiane Hayou, Greg Yang</i>	
A Generalization of ViT/MLP-Mixer to Graphs.....	12724
<i>Xiaoxin He, Bryan Hooi, Thomas Laurent, Adam Perold, Yann Lecun, Xavier Bresson</i>	
Domain Adaptation for Time Series Under Feature and Label Shifts	12746
<i>Huan He, Owen Queen, Teddy Koker, Consuelo Cuevas, Theodoros Tsiligkaridis, Marinka Zitnik</i>	
Contrastive Learning Meets Homophily: Two Birds with One Stone	12775
<i>Dongxiao He, Jitao Zhao, Rui Guo, Zhiyong Feng, Di Jin, Yuxiao Huang, Zhen Wang, Weixiong Zhang</i>	

PART 17

Nearly Minimax Optimal Reinforcement Learning for Linear Markov Decision Processes.....	12790
<i>Jiafan He, Heyang Zhao, Dongruo Zhou, Quanquan Gu</i>	
CRISP: Curriculum Based Sequential Neural Decoders for Polar Code Family.....	12823
<i>S Ashwin Hebbar, Viraj Vivek Nadkarni, Ashok Vardhan Makkuva, Suma Bhat, Sewoong Oh, Pramod Viswanath</i>	

Sketch-Flip-Merge: Mergeable Sketches for Private Distinct Counting	12846
<i>Jonathan Hehir, Daniel Ting, Graham Cormode</i>	
Functional Neural Networks: Shift Invariant Models for Functional Data with Applications to EEG Classification	12866
<i>Florian Heinrichs, Mavin Heim, Corinna Weber</i>	
Distance Weighted Supervised Learning for Offline Interaction Data	12882
<i>Joey Hejna, Jensen Gao, Dorsa Sadigh</i>	
Group Equivariant Fourier Neural Operators for Partial Differential Equations	12907
<i>Jacob Helwig, Xuan Zhang, Cong Fu, Jerry Kurtin, Stephan Wojtowytsch, Shuiwang Ji</i>	
Training-Free Neural Active Learning with Initialization-Robustness Guarantees	12931
<i>Apivich Hemachandra, Zhongxiang Dai, Jasraj Singh, See-Kiong Ng, Bryan Kian Hsiang Low</i>	
A Study of Global and Episodic Bonuses for Exploration in Contextual MDPs	12972
<i>Mikael Henaff, Minqi Jiang, Roberta Raileanu</i>	
Robust Camera Pose Refinement for Multi-Resolution Hash Encoding	13000
<i>Hwan Heo, Taekyung Kim, Jiyoung Lee, Jaewon Lee, Soohyun Kim, Hyunwoo J. Kim, Jin-Hwa Kim</i>	
Generalized Teacher Forcing for Learning Chaotic Dynamics	13017
<i>Florian Hess, Zahra Monfared, Manuel Brenner, Daniel Durstewitz</i>	
Causal Modeling of Policy Interventions from Treatment-Outcome Sequences	13050
<i>Çaglar Hizli, S. T. John, Anne Tuulikki Juuti, Tuure Tapani Saarinen, Kirsi Hannele Pietiläinen, Pekka Marttinen</i>	
Monotonicity and Double Descent in Uncertainty Estimation with Gaussian Processes	13085
<i>Liam Hodgkinson, Chris Van Der Heide, Fred Roosta, Michael W. Mahoney</i>	
AdaBoost is Not an Optimal Weak to Strong Learner	13118
<i>Mikael Møller Høgsgaard, Kasper Green Larsen, Martin Ritzert</i>	
Dual Propagation: Accelerating Contrastive Hebbian Learning with Dyadic Neurons	13141
<i>Rasmus Høier, D. Staudt, Christopher Zach</i>	
Multi-Task Off-Policy Learning from Bandit Feedback	13157
<i>Joey Hong, Branislav Kveton, Manzil Zaheer, Sumeet Katariya, Mohammad Ghavamzadeh</i>	
Constrained Optimization Via Exact Augmented Lagrangian and Randomized Iterative Sketching	13174
<i>Ilgee Hong, Sen Na, Michael W. Mahoney, Mladen Kolar</i>	
Revisiting Data-Free Knowledge Distillation with Poisoned Teachers	13199
<i>Junyuan Hong, Yi Zeng, Shuyang Yu, Lingjuan Lyu, Ruoxi Jia, Jiayu Zhou</i>	
Simple Diffusion: End-To-End Diffusion for High Resolution Images	13213
<i>Emiel Hoogeboom, Jonathan Heek, Tim Salimans</i>	
Causal Strategic Classification: A Tale of Two Shifts	13233
<i>Guy Horowitz, Nir Rosenfeld</i>	
Fair and Accurate Decision Making Through Group-Aware Learning	13254
<i>Ramtin Hosseini, Li Zhang, Bhanu Garg, Pengtao Xie</i>	
Approximation Algorithms for Fair Range Clustering	13270
<i>Sedjro Salomon Hotegni, Sepideh Mahabadi, Ali Vakilian</i>	
Decoding Layer Saliency in Language Transformers	13285
<i>Elizabeth Mary Hou, Gregory David Castanon</i>	
PromptBoosting: Black-Box Text Classification with Ten Forward Passes	13309
<i>Bairu Hou, Joe O'Connor, Jacob Andreas, Shiyu Chang, Yang Zhang</i>	
Sparse Learning of Dynamical Systems in RKHS: An Operator-Theoretic Approach	13325
<i>Boya Hou, Sina Sanjari, Nathan Dahlin, Subhonmesh Bose, Umesh Vaidya</i>	
Probably Anytime-Safe Stochastic Combinatorial Semi-Bandits	13353
<i>Yunlong Hou, Vincent Y. F. Tan, Zixin Zhong</i>	
Automatic Data Augmentation Via Invariance-Constrained Learning	13410
<i>Ignacio Hounie, Luiz F. O. Chamon, Alejandro Ribeiro</i>	

Thompson Sampling with Diffusion Generative Prior	13434
<i>Yu-Guan Hsieh, Shiva Kasiviswanathan, Branislav Kveton, Patrick Blöbaum</i>	
Tighter Analysis for ProxSkip	13469
<i>Zhengmian Hu, Heng Huang</i>	
Omnipredictors for Constrained Optimization	13497
<i>Lunjia Hu, Inbal Rachel Livni Navon, Omer Reingold, Chutong Yang</i>	
GFlowNet-EM for Learning Compositional Latent Variable Models	13528
<i>Edward J Hu, Nikolay Malkin, Moksh Jain, Katie E Everett, Alexandros Graikos, Yoshua Bengio</i>	
Blockwise Stochastic Variance-Reduced Methods with Parallel Speedup for Multi-Block Bilevel Optimization.....	13550
<i>Quanqi Hu, Zi-Hao Qiu, Zhishuai Guo, Lijun Zhang, Tianbao Yang</i>	
Language Instructed Reinforcement Learning for Human-AI Coordination.....	13584
<i>Hengyuan Hu, Dorsa Sadigh</i>	

PART 18

Surface Snapping Optimization Layer for Single Image Object Shape Reconstruction.....	13599
<i>Yuan-Ting Hu, Alex Schwing, Raymond A. Yeh</i>	
Learning to Learn from APIs: Black-Box Data-Free Meta-Learning.....	13610
<i>Zixuan Hu, Li Shen, Zhenyi Wang, Baoyuan Wu, Chun Yuan, Dacheng Tao</i>	
For Pre-Trained Vision Models in Motor Control, Not All Policy Learning Methods Are Created Equal.....	13628
<i>Yingdong Hu, Renhao Wang, Li Erran Li, Yang Gao</i>	
Beyond Lipschitz Smoothness: A Tighter Analysis for Nonconvex Optimization	13652
<i>Zhengmian Hu, Xidong Wu, Heng Huang</i>	
Understanding the Impact of Adversarial Robustness on Accuracy Disparity	13679
<i>Yuzheng Hu, Fan Wu, Hongyang Zhang, Han Zhao</i>	
Reinforcement Learning in Low-Rank MDPs with Density Features	13710
<i>Audrey Huang, Jinglin Chen, Nan Jiang</i>	
Composer: Creative and Controllable Image Synthesis with Composable Conditions	13753
<i>Lianghua Huang, Di Chen, Yu Liu, Yujun Shen, Deli Zhao, Jingren Zhou</i>	
Model-Aware Contrastive Learning: Towards Escaping the Dilemmas	13774
<i>Zizheng Huang, Haoxing Chen, Ziqi Wen, Chao Zhang, Huaxiong Li, Bo Wang, Chunlin Chen</i>	
High-Dimensional Clustering onto Hamiltonian Cycle.....	13791
<i>Tianyi Huang, Shenghui Cheng, Stan Z. Li, Zhengjun Zhang</i>	
Banker Online Mirror Descent: A Universal Approach for Delayed Online Bandit Learning	13814
<i>Jiatai Huang, Yan Dai, Longbo Huang</i>	
Fast Algorithms for Distributed k-Clustering with Outliers	13845
<i>Junyu Huang, Qilong Feng, Ziyun Huang, Jinhui Xu, Jianxin Wang</i>	
Searching Large Neighborhoods for Integer Linear Programs with Contrastive Learning	13869
<i>Taoan Huang, Aaron M Ferber, Yuandong Tian, Bistra Dilkina, Benoit Steiner</i>	
On Coresets for Clustering in Small Dimensional Euclidean Spaces.....	13891
<i>Lingxiao Huang, Ruiyuan Huang, Zengfeng Huang, Xuan Wu</i>	
Make-An-Audio: Text-To-Audio Generation with Prompt-Enhanced Diffusion Models	13916
<i>Rongjie Huang, Jiawei Huang, Dongchao Yang, Yi Ren, Luping Liu, Mingze Li, Zhenhui Ye, Jinglin Liu, Xiang Yin, Zhou Zhao</i>	
The Power of Uniform Sampling for k-Median	13933
<i>Lingxiao Huang, Shaofeng H.-C. Jiang, Jianing Lou</i>	
Reparameterized Policy Learning for Multimodal Trajectory Optimization.....	13957
<i>Zhiao Huang, Litian Liang, Zhan Ling, Xuanlin Li, Chuang Gan, Hao Su</i>	

Theoretical Bounds on the Network Community Profile from Low-Rank Semi-Definite Programming.....	13976
<i>Yufan Huang, C. Seshadhri, David F. Gleich</i>	
NeuralStagger: Accelerating Physics-Constrained Neural PDE Solver with Spatial-Temporal Decomposition.....	13993
<i>Xinquan Huang, Wenlei Shi, Qi Meng, Yue Wang, Xiaotian Gao, Jia Zhang, Tie-Yan Liu</i>	
Policy Contrastive Imitation Learning.....	14007
<i>Jialei Huang, Zhao-Heng Yin, Yingdong Hu, Yang Gao</i>	
Are Large Kernels Better Teachers than Transformers for ConvNets?	14023
<i>Tianjin Huang, Lu Yin, Zhenyu Zhang, Li Shen, Meng Fang, Mykola Pechenizkiy, Zhangyang Wang, Shiwei Liu</i>	
Achieving Linear Speedup in Non-IID Federated Bilevel Learning	14039
<i>Minhui Huang, Dewei Zhang, Kaiyi Ji</i>	
Federated Linear Contextual Bandits with User-Level Differential Privacy	14060
<i>Ruiquan Huang, Huanyu Zhang, Luca Melis, Milan Shen, Meisam Hejazinia, Jing Yang</i>	
Straightening Out the Straight-Through Estimator: Overcoming Optimization Challenges in Vector Quantized Networks	14096
<i>Minyoung Huh, Brian Cheung, Pulkit Agrawal, Phillip Isola</i>	
Cut Your Losses with Squentropy	14114
<i>Like Hui, Mikhail Belkin, Stephen Wright</i>	
SOM-CPC: Unsupervised Contrastive Learning with Self-Organizing Maps for Structured Representations of High-Rate Time Series.....	14132
<i>Iris A. M. Huijben, Arthur Andreas Nijdam, Sebastiaan Overeem, Merel M Van Gilst, Ruud Van Sloun</i>	
One-Shot Federated Conformal Prediction	14153
<i>Pierre Humbert, Batiste Le Bars, Aurélien Bellet, Sylvain Arlot</i>	
The Impact of Exploration on Convergence and Performance of Multi-Agent Q-Learning Dynamics.....	14178
<i>Aamal Hussain, Francesco Belardinelli, Dario Paccagnan</i>	
Combinatorial Neural Bandits	14203
<i>Taehyun Hwang, Kyuwook Chai, Min-Hwan Oh</i>	
MAGANet: Achieving Combinatorial Generalization by Modeling a Group Action	14237
<i>Geonho Hwang, Jaewoong Choi, Hyunsoo Cho, Myungjoo Kang</i>	
Information-Theoretic State Space Model for Multi-View Reinforcement Learning.....	14249
<i>Hyeongjoo Hwang, Seokin Seo, Youngsoo Jang, Sungyoon Kim, Geon-Hyeong Kim, Seunghoon Hong, Kee-Eung Kim</i>	
Under-Counted Tensor Completion with Neural Incorporation of Attributes	14283
<i>Shahana Ibrahim, Xiao Fu, Rebecca Hutchinson, Eugene Seo</i>	
On the Identifiability and Estimation of Causal Location-Scale Noise Models	14316
<i>Alexander Immer, Christoph Schultheiss, Julia E Vogt, Bernhard Schölkopf, Peter Bühlmann, Alexander Marx</i>	
Stochastic Marginal Likelihood Gradients Using Neural Tangent Kernels.....	14333
<i>Alexander Immer, Tycho F. A. Van Der Ouderaa, Mark Van Der Wilk, Gunnar Ratsch, Bernhard Schölkopf</i>	
Differentially Private Hierarchical Clustering with Provable Approximation Guarantees	14353
<i>Jacob Imola, Alessandro Epasto, Mohammad Mahdian, Vincent Cohen-Addad, Vahab Mirrokni</i>	
Neural Network Accelerated Implicit Filtering: Integrating Neural Network Surrogates with Provably Convergent Derivative Free Optimization Methods	14376
<i>Brian Irwin, Eldad Haber, Raviv Gal, Avi Ziv</i>	

Principled Offline RL in the Presence of Rich Exogenous Information.....	14390
<i>Riashat Islam, Manan Tomar, Alex Lamb, Yonathan Efroni, Hongyu Zang, Aniket Rajiv Didolkar, Dipendra Misra, Xin Li, Harm Van Seijen, Remi Tachet Des Combes, John Langford</i>	

PART 19

Unveiling the Latent Space Geometry of Push-Forward Generative Models.....	14422
<i>Thibaut Issenhuth, Ugo Tanielian, Jeremie Mary, David Picard</i>	
CO-BED: Information-Theoretic Contextual Optimization Via Bayesian Experimental Design.....	14445
<i>Desi R. Ivanova, Joel Jennings, Tom Rainforth, Cheng Zhang, Adam Foster</i>	
DoG is SGD’s Best Friend: A Parameter-Free Dynamic Step Size Schedule.....	14465
<i>Maor Ivgi, Oliver Hinder, Yair Carmon</i>	
Maximal Initial Learning Rates in Deep ReLU Networks	14500
<i>Gaurav Iyer, Boris Hanin, David Rolnick</i>	
Data-Driven Subgroup Identification for Linear Regression.....	14531
<i>Zachary Izzo, Ruishan Liu, James Zou</i>	
Efficient Training of Language Models Using Few-Shot Learning	14553
<i>Sashank J. Reddi, Sobhan Miryoosefi, Stefani Karp, Shankar Krishnan, Satyen Kale, Seungyeon Kim, Sanjiv Kumar</i>	
Scalable Adaptive Computation for Iterative Generation.....	14569
<i>Allan Jabri, David J. Fleet, Ting Chen</i>	
Unconstrained Online Learning with Unbounded Losses	14590
<i>Andrew Jacobsen, Ashok Cutkosky</i>	
Multi-Objective GFlowNets.....	14631
<i>Moksh Jain, Sharath Chandra Raparthy, Alex Hernández-Garci’a, Jarrid Rector-Brooks, Yoshua Bengio, Santiago Miret, Emmanuel Bengio</i>	
The Price of Differential Privacy Under Continual Observation.....	14654
<i>Palak Jain, Sofya Raskhodnikova, Satchit Sivakumar, Adam Smith</i>	
Graph Ladling: Shockingly Simple Parallel GNN Training Without Intermediate Communication.....	14679
<i>Ajay Kumar Jaiswal, Shiwei Liu, Tianlong Chen, Ying Ding, Zhangyang Wang</i>	
Instant Soup: Cheap Pruning Ensembles in a Single Pass Can Draw Lottery Tickets from Large Models.....	14691
<i>Ajay Kumar Jaiswal, Shiwei Liu, Tianlong Chen, Ying Ding, Zhangyang Wang</i>	
Exploring the Benefits of Training Expert Language Models Over Instruction Tuning.....	14702
<i>Joel Jang, Seungone Kim, Seonghyeon Ye, Doyoung Kim, Lajanugen Logeswaran, Moontae Lee, Kyungjae Lee, Minjoon Seo</i>	
Learning to Boost Training by Periodic Nowcasting Near Future Weights.....	14730
<i>Jinhyeok Jang, Woo-Han Yun, Won Hwa Kim, Youngwoo Yoon, Jaehong Kim, Jaeyeon Lee, Byungok Han</i>	
Unscented Autoencoder.....	14758
<i>Faris Janjos, Lars Rosenbaum, Maxim Dolgov, J. Marius Zoellner</i>	
Curiosity in Hindsight: Intrinsic Exploration in Stochastic Environments.....	14780
<i>Daniel Jarrett, Corentin Tallec, Florent Alché, Thomas Mesnard, Remi Munos, Michal Valko</i>	
BiRT: Bio-Inspired Replay in Vision Transformers for Continual Learning.....	14817
<i>Kishaan Jeeveswaran, Prashant Shivaram Bhat, Bahram Zonooz, Elahe Arani</i>	
Recovering Top-Two Answers and Confusion Probability in Multi-Choice Crowdsourcing.....	14836
<i>Hyeonsu Jeong, Hye Won Chung</i>	
Leveraging Label Non-Uniformity for Node Classification in Graph Neural Networks	14869
<i>Feng Ji, See Hian Lee, Hanyang Meng, Kai Zhao, Jielong Yang, Wee Peng Tay</i>	

Bidirectional Adaptation for Robust Semi-Supervised Learning with Inconsistent Data Distributions	14886
<i>Lin-Han Jia, Lan-Zhe Guo, Zhi Zhou, Jie-Jing Shao, Yuke Xiang, Yu-Feng Li</i>	
Short-Lived High-Volume Bandits	14902
<i>Su Jia, Nishant Oli, Ian Anderson, Paul Duff, Andrew A Li, R. Ravi</i>	
Smooth Non-Stationary Bandits	14930
<i>Su Jia, Qian Xie, Nathan Kallus, Peter I. Frazier</i>	
A Unified Optimization Framework of ANN-SNN Conversion: Towards Optimal Mapping from Activation Values to Firing Rates	14945
<i>Haiyan Jiang, Srinivas Anumasa, Giulia De Masi, Huan Xiong, Bin Gu</i>	
VIMA: Robot Manipulation with Multimodal Prompts	14975
<i>Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, Linxi Fan</i>	
Estimating Causal Effects Using a Multi-Task Deep Ensemble	15023
<i>Ziyang Jiang, Zhuoran Hou, Yiling Liu, Yiman Ren, Keyu Li, David Carlson</i>	
Online Restless Bandits with Unobserved States	15041
<i>Bowen Jiang, Bo Jiang, Jian Li, Tao Lin, Xinbing Wang, Chenghu Zhou</i>	
Detecting Out-Of-Distribution Data Through In-Distribution Class Prior	15067
<i>Xue Jiang, Feng Liu, Zhen Fang, Hong Chen, Tongliang Liu, Feng Zheng, Bo Han</i>	
Towards Stable and Efficient Adversarial Training Against L1 Bounded Adversarial Attacks	15089
<i>Yulun Jiang, Chen Liu, Zhichao Huang, Mathieu Salzmann, Sabine Susstrunk</i>	
Learning Unnormalized Statistical Models Via Compositional Optimization.....	15105
<i>Wei Jiang, Jiayu Qin, Lingyu Wu, Changyou Chen, Tianbao Yang, Lijun Zhang</i>	
Approximate Causal Effect Identification Under Weak Confounding	15125
<i>Wei Jiang, Jiayu Qin, Lingyu Wu, Changyou Chen, Tianbao Yang, Lijun Zhang</i>	
MEWL: Few-Shot Multimodal Word Learning with Referential Uncertainty	15144
<i>Guangyuan Jiang, Manjie Xu, Shiji Xin, Wei Liang, Yujia Peng, Chi Zhang, Yixin Zhu</i>	
NeuralSlice: Neural 3D Triangle Mesh Reconstruction Via Slicing 4D Tetrahedral Meshes.....	15170
<i>Chenbo Jiang, Jie Yang, Shwai He, Yu-Kun Lai, Lin Gao</i>	
Effective Structured Prompting by Meta-Learning and Representative Verbalizer.....	15186
<i>Weisen Jiang, Yu Zhang, James Kwok</i>	
Understanding Incremental Learning of Gradient Descent: A Fine-Grained Analysis of Matrix Sensing	15200
<i>Jikai Jin, Zhiyuan Li, Kaifeng Lyu, Simon Shaolei Du, Jason D. Lee</i>	

PART 20

Thompson Sampling with Less Exploration is Fast and Optimal.....	15239
<i>Tianyuan Jin, Xianglin Yang, Xiaokui Xiao, Pan Xu</i>	
R-U-SURE? Uncertainty-Aware Code Suggestions by Maximizing Utility Across Random User Intents	15262
<i>Daniel D. Johnson, Daniel Tarlow, Christian Walder</i>	
Automatically Auditing Large Language Models Via Discrete Optimization	15307
<i>Erik Jones, Anca Dragan, Aditi Raghunathan, Jacob Steinhardt</i>	
On the Expressive Power of Geometric Graph Neural Networks	15330
<i>Chaitanya K. Joshi, Cristian Bodnar, Simon V Mathis, Taco Cohen, Pietro Lio</i>	
Data-Efficient Contrastive Self-Supervised Learning: Most Beneficial Examples for Supervised Learning Contribute the Least	15356
<i>Siddharth Joshi, Baharan Mirzasoleiman</i>	
Robust Subtask Learning for Compositional Generalization	15371
<i>Kishor Jothimurugan, Steve Hsu, Osbert Bastani, Rajeev Alur</i>	

On Bridging the Gap Between Mean Field and Finite Width Deep Random Multilayer Perceptron with Batch Normalization.....	15388
<i>Amir Joudaki, Hadi Daneshmand, Francis Bach</i>	
FARE: Provably Fair Representation Learning with Practical Certificates.....	15401
<i>Nikola Jovanovic, Mislav Balunovic, Dimitar Iliev Dimitrov, Martin Vechev</i>	
Scaling of Class-Wise Training Losses for Post-Hoc Calibration	15421
<i>Seungjin Jung, Seungmo Seo, Yonghyun Jeong, Jongwon Choi</i>	
Fighting Fire with Fire: Contrastive Debiasing Without Bias-Free Data Via Generative Bias-Transformation	15435
<i>Yeonsung Jung, Hajin Shim, June Yong Yang, Eunho Yang</i>	
Estimating Joint Treatment Effects by Combining Multiple Experiments	15451
<i>Yonghan Jung, Jin Tian, Elias Bareinboim</i>	
The Catalog Problem: Clustering and Ordering Variable-Sized Sets	15528
<i>Mateusz Maria Jurewicz, Graham W. Taylor, Leon Derczynski</i>	
Equivariance with Learned Canonicalization Functions	15546
<i>Sékou-Oumar Kaba, Arnab Kumar Mondal, Yan Zhang, Yoshua Bengio, Siamak Ravanbakhsh</i>	
Biases in Evaluation of Molecular Optimization Methods and Bias Reduction Strategies	15567
<i>Hiroshi Kajino, Kohei Miyaguchi, Takayuki Osogami</i>	
Statistical Indistinguishability of Learning Algorithms.....	15586
<i>Alkis Kalavasis, Amin Karbasi, Shay Moran, Grigoris Velegkas</i>	
Identifying Interpretable Subspaces in Image Representations	15623
<i>Neha Kalibhat, Shweta Bhardwaj, C. Bayan Bruss, Hamed Firooz, Maziar Sanjabi, Soheil Feizi</i>	
Nonlinear Causal Discovery with Latent Confounders	15639
<i>David Kaltenpoth, Jilles Vreeken</i>	
Deep Generative Symbolic Regression with Monte-Carlo-Tree-Search	15655
<i>Pierre-Alexandre Kamienny, Guillaume Lample, Sylvain Lamprier, Marco Virgolin</i>	
One-Vs-The-Rest Loss to Focus on Important Samples in Adversarial Training.....	15669
<i>Sekitoshi Kanai, Shin'ya Yamaguchi, Masanori Yamada, Hiroshi Takahashi, Kentaro Ohno, Yasutoshi Ida</i>	
Large Language Models Struggle to Learn Long-Tail Knowledge	15696
<i>Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, Colin Raffel</i>	
Git-Theta: A Git Extension for Collaborative Development of Machine Learning Models	15708
<i>Nikhil Kandpal, Brian Lester, Mohammed Muqeeth, Anisha Mascarenhas, Monty Evans, Vishal Baskaran, Tenghao Huang, Haokun Liu, Colin Raffel</i>	
A Deep Conjugate Direction Method for Iteratively Solving Linear Systems	15720
<i>Ayano Kaneda, Osman Akar, Jingyu Chen, Victoria Alicia Trevino Kala, David Hyde, Joseph Teran</i>	
Leveraging Proxy of Training Data for Test-Time Adaptation.....	15737
<i>Juwon Kang, Nayeong Kim, Donghyeon Kwon, Jungseul Ok, Suha Kwak</i>	
Beyond Reward: Offline Preference-Guided Policy Optimization.....	15753
<i>Yachen Kang, Diyuhan Shi, Jinxin Liu, Li He, Donglin Wang</i>	
Poisoning Generative Replay in Continual Learning to Promote Forgetting	15769
<i>Siteng Kang, Zhan Shi, Xinhua Zhang</i>	
Node Embedding from Neural Hamiltonian Orbits in Graph Neural Networks	15786
<i>Qiyu Kang, Kai Zhao, Yang Song, Sijie Wang, Wee Peng Tay</i>	
Understanding Gradient Regularization in Deep Learning: Efficient Finite-Difference Computation and Implicit Bias	15809
<i>Ryo Karakida, Tomoumi Takase, Tomohiro Hayase, Kazuki Osawa</i>	
Langevin Thompson Sampling with Logarithmic Communication: Bandits and Reinforcement Learning	15828
<i>Amin Karbasi, Nikki Lijing Kuang, Yian Ma, Siddharth Mitra</i>	

On the Relationship Between Explanation and Prediction: A Causal View	15861
<i>Amir-Hossein Karimi, Krikamol Muandet, Simon Kornblith, Bernhard Schölkopf, Been Kim</i>	
Cocktail Party Attack: Breaking Aggregation-Based Privacy in Federated Learning Using Independent Component Analysis	15884
<i>Sanjay Kariyappa, Chuan Guo, Kiwan Maeng, Wenjie Xiong, G. Edward Suh, Moinuddin K Qureshi, Hsien-Hsin S. Lee</i>	
General Sequential Episodic Memory Model.....	15900
<i>Arjun Karuvally, Terrence Sejnowski, Hava T Siegelmann</i>	
Regression with Sensor Data Containing Incomplete Observations	15911
<i>Takayuki Katsuki, Takayuki Osogami</i>	
Data Representations' Study of Latent Image Manifolds	15928
<i>Ilya Kaufman, Omri Azencot</i>	
Multi-Modal Classifiers for Open-Vocabulary Object Detection	15946
<i>Prannay Kaul, Weidi Xie, Andrew Zisserman</i>	
Learning Mixtures of Markov Chains and MDPs	15970
<i>Chinmaya Kausik, Kevin Tan, Ambuj Tewari</i>	

PART 21

Curious Replay for Model-Based Adaptation.....	16018
<i>Isaac Kauvar, Chris Doyle, Linqi Zhou, Nick Haber</i>	
How Does Information Bottleneck Help Deep Learning?.....	16049
<i>Kenji Kawaguchi, Zhun Deng, Xu Ji, Jiaoyang Huang</i>	
Instrumental Variable Estimation of Average Partial Causal Effects.....	16097
<i>Yuta Kawakami, Manabu Kuroki, Jin Tian</i>	
The Test of Tests: A Framework for Differentially Private Hypothesis Testing	16131
<i>Zeki Kazan, Kaiyan Shi, Adam Groce, Andrew P Bray</i>	
Exact Inference in High-Order Structured Prediction	16152
<i>Chuyang Ke, Jean Honorio</i>	
Neural Wave Machines: Learning Spatiotemporally Structured Representations with Locally Coupled Oscillatory Recurrent Neural Networks	16168
<i>T. Anderson Keller, Max Welling</i>	
Homomorphism AutoEncoder -- Learning Group Structured Representations from Observed Transitions	16190
<i>Hamza Keurti, Hsiao-Ru Pan, Michel Besserve, Benjamin F Grewe, Bernhard Schölkopf</i>	
Rethinking Backdoor Attacks.....	16216
<i>Alaa Khaddaj, Guillaume Leclerc, Aleksandar Makelov, Kristian Georgiev, Hadi Salman, Andrew Ilyas, Aleksander Madry</i>	
PAC Prediction Sets for Large Language Models of Code.....	16237
<i>Adam Khakhar, Stephen Mell, Osbert Bastani</i>	
Accelerated Primal-Dual Methods for Convex-Strongly-Concave Saddle Point Problems	16250
<i>Mohammad Khalafi, Digvijay Boob</i>	
Loss Balancing for Fair Supervised Learning	16271
<i>Mohammad Mahdi Khalili, Xueru Zhang, Mahed Abroshan</i>	
Linearly Constrained Bilevel Optimization: A Smoothed Implicit Gradient Approach	16291
<i>Prashant Khanduri, Ioannis Tsaknakis, Yihua Zhang, Jia Liu, Sijia Liu, Jiawei Zhang, Mingyi Hong</i>	
Emergent Asymmetry of Precision and Recall for Measuring Fidelity and Diversity of Generative Models in High Dimensions	16326
<i>Mahyar Khayatkhoei, Wael Abdalmageed</i>	
Learning-Augmented Private Algorithms for Multiple Quantile Release	16344
<i>Mikhail Khodak, Kareem Amin, Travis Dick, Sergei Vassilvitskii</i>	

CrossSplit: Mitigating Label Noise Memorization Through Data Splitting	16377
<i>Jihye Kim, Aristide Baratin, Yan Zhang, Simon Lacoste-Julien</i>	
Trainability, Expressivity and Interpretability in Gated Neural ODEs	16393
<i>Timothy Doyeon Kim, Tankut Can, Kamesh Krishnamurthy</i>	
SAAL: Sharpness-Aware Active Learning	16424
<i>Yoon-Yeong Kim, Youngjae Cho, Joonho Jang, Byeonghu Na, Yeongmin Kim, Kyungwoo Song, Wanmo Kang, Il-Chul Moon</i>	
Demonstration-Free Autonomous Reinforcement Learning Via Implicit and Bidirectional Curriculum.....	16441
<i>Jigang Kim, Daesol Cho, H. Jin Kim</i>	
Improved Algorithms for Multi-Period Multi-Class Packing Problems with Bandit Feedback	16458
<i>Wonyoung Kim, Garud Iyengar, Assaf Zeevi</i>	
Efficient Latency-Aware CNN Depth Compression Via Two-Stage Dynamic Programming.....	16502
<i>Jinuk Kim, Yeonwoo Jeong, Deokjae Lee, Hyun Oh Song</i>	
Probabilistic Concept Bottleneck Models	16521
<i>Eunji Kim, Dahuin Jung, Sangha Park, Siwon Kim, Sungroh Yoon</i>	
DevFormer: A Symmetric Transformer for Context-Aware Device Placement.....	16541
<i>Haeyeon Kim, Minsu Kim, Federico Berto, Joungho Kim, Jinkyoo Park</i>	
Refining Generative Process with Discriminator Guidance in Score-Based Diffusion Models.....	16567
<i>Dongjun Kim, Yeongmin Kim, Se Jung Kwon, Wanmo Kang, Il-Chul Moon</i>	
Robust Non-Linear Feedback Coding Via Power-Constrained Deep Learning.....	16599
<i>Junghoon Kim, Taejoon Kim, David Love, Christopher Brinton</i>	
LESSON: Learning to Integrate Exploration Strategies for Reinforcement Learning Via an Option Framework.....	16619
<i>Woojun Kim, Jeonghye Kim, Youngchul Sung</i>	
BPipe: Memory-Balanced Pipeline Parallelism for Training Large Language Models.....	16639
<i>Taebum Kim, Hyoungjoo Kim, Gyeong-In Yu, Byung-Gon Chun</i>	
Probabilistic Imputation for Time-Series Classification with Missing Data	16654
<i>Seunghyun Kim, Hyunsu Kim, Eunggu Yun, Hwangrae Lee, Jaehun Lee, Juho Lee</i>	
Variational Curriculum Reinforcement Learning for Unsupervised Discovery of Skills	16668
<i>Seongun Kim, Kyowoon Lee, Jaesik Choi</i>	
Margin-Based Neural Network Watermarking.....	16696
<i>Byungjoo Kim, Suyoung Lee, Seanie Lee, Sooel Son, Sung Ju Hwang</i>	
Regularizing Towards Soft Equivariance Under Mixed Symmetries	16712
<i>Hyunsu Kim, Hyungi Lee, Hongseok Yang, Juho Lee</i>	
Model-Based Offline Reinforcement Learning with Count-Based Conservatism.....	16728
<i>Byeongchan Kim, Min-Hwan Oh</i>	
Transformer-Based Stagewise Decomposition for Large-Scale Multistage Stochastic Optimization	16747
<i>Chanyeong Kim, Jongwoong Park, Hyunglip Bae, Woo Chang Kim</i>	
SurProGenes: Survival Risk-Ordered Representation of Cancer Patients and Genes for the Identification of Prognostic Genes	16771
<i>Junetae Kim, Kyoungsuk Park, Hanseok Jeong, Youngwook Kim, Jeongseon Kim, Sun-Young Kim</i>	
Stable and Consistent Prediction of 3D Characteristic Orientation Via Invariant Residual Learning	16787
<i>Seungwook Kim, Chunghyun Park, Yoonwoo Jeong, Jaesik Park, Minsu Cho</i>	
Prefer to Classify: Improving Text Classifiers Via Auxiliary Preference Learning	16807
<i>Jaehyung Kim, Jinwoo Shin, Dongyeop Kang</i>	

PART 22

An Adaptive Entropy-Regularization Framework for Multi-Agent Reinforcement Learning.....	16829
<i>Woojun Kim, Youngchul Sung</i>	

Practical and Matching Gradient Variance Bounds for Black-Box Variational Bayesian Inference.....	16853
<i>Kyurae Kim, Kaiwen Wu, Jisu Oh, Jacob R. Gardner</i>	
Learnability and Algorithm for Continual Learning.....	16877
<i>Gyuhak Kim, Changnan Xiao, Tatsuya Konishi, Bing Liu</i>	
Unifying Nesterov’s Accelerated Gradient Methods for Convex and Strongly Convex Objective Functions	16897
<i>Junghbin Kim, Insoon Yang</i>	
Denoising MCMC for Accelerating Diffusion-Based Generative Models	16955
<i>Beomsu Kim, Jong Chul Ye</i>	
Structure Learning of Latent Factors Via Clique Search on Correlation Thresholded Graphs.....	16978
<i>Dale Kim, Qing Zhou</i>	
Fair and Robust Estimation of Heterogeneous Treatment Effects for Policy Learning.....	16997
<i>Kwangho Kim, Jose R Zubizarreta</i>	
Proper Losses for Discrete Generative Models	17015
<i>Dhamma Kimpara, Rafael Frongillo, Bo Waggoner</i>	
Controlling Posterior Collapse by an Inverse Lipschitz Constraint on the Decoder Network.....	17041
<i>Yuri Kinoshita, Kenta Oono, Kenji Fukumizu, Yuichi Yoshida, Shin-Ichi Maeda</i>	
A Watermark for Large Language Models	17061
<i>John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, Tom Goldstein</i>	
Probabilistic Contrastive Learning Recovers the Correct Aleatoric Uncertainty of Ambiguous Inputs.....	17085
<i>Michael Kirchhof, Enkelejda Kasneci, Seong Joon Oh</i>	
Training Normalizing Flows from Dependent Data	17105
<i>Matthias Kirchler, Christoph Lippert, Marius Kloft</i>	
IncDSI: Incrementally Updatable Document Retrieval.....	17122
<i>Varsha Kishore, Chao Wan, Justin Lovelace, Yoav Artzi, Kilian Q Weinberger</i>	
Regularization and Variance-Weighted Regression Achieves Minimax Optimality in Linear MDPs: Theory and Practice.....	17135
<i>Toshinori Kitamura, Tadashi Kozuno, Yunhao Tang, Nino Vieillard, Michal Valko, Wenhao Yang, Jincheng Mei, Pierre Menard, Mohammad Gheshlaghi Azar, Remi Munos, Olivier Pietquin, Matthieu Geist, Csaba Szepesvari, Wataru Kumagai, Yutaka Matsuo</i>	
Drug Discovery Under Covariate Shift with Domain-Informed Prior Distributions Over Functions	17176
<i>Leo Klarner, Tim G. J. Rudner, Michael Reutlinger, Torsten Schindler, Garrett M Morris, Charlotte Deane, Yee Whye Teh</i>	
Deep Laplacian-Based Options for Temporally-Extended Exploration	17198
<i>Martin Klissarov, Marlos C. Machado</i>	
Generalized Reductions: Making Any Hierarchical Clustering Fair and Balanced with Low Cost	17218
<i>Marina Knittel, Max Springer, John P Dickerson, Mohammadtaghi Hajiaghayi</i>	
Can We Scale Transformers to Predict Parameters of Diverse ImageNet Models?.....	17243
<i>Boris Knyazev, Doha Hwang, Simon Lacoste-Julien</i>	
Online Learning with Feedback Graphs: The True Shape of Regret.....	17260
<i>Tomáš Kocák, Alexandra Carpentier</i>	
Grounding Language Models to Images for Multimodal Inputs and Outputs	17283
<i>Jing Yu Koh, Ruslan Salakhutdinov, Daniel Fried</i>	
Rigid Body Flows for Sampling Molecular Crystal Structures.....	17301
<i>Jonas Köhler, Michele Invernizzi, Pim De Haan, Frank Noe</i>	
Enabling First-Order Gradient-Based Learning for Equilibrium Computation in Markets.....	17327
<i>Nils Kohring, Fabian Raoul Pieroth, Martin Bichler</i>	
Revisiting Gradient Clipping: Stochastic Bias and Tight Convergence Guarantees	17343
<i>Anastasia Koloskova, Hadrien Hendrikx, Sebastian U Stich</i>	
On Computing Optimal Tree Ensembles.....	17364
<i>Christian Komusiewicz, Pascal Kunz, Frank Sommer, Manuel Sorge</i>	

GOAT: A Global Transformer on Large-Scale Graphs	17375
<i>Kezhi Kong, Jiu hai Chen, John Kirchenbauer, Renkun Ni, C. Bayan Bruss, Tom Goldstein</i>	
Autoregressive Diffusion Model for Graph Generation	17391
<i>Lingkai Kong, Jiaming Cui, Haotian Sun, Yuchen Zhuang, B. Aditya Prakash, Chao Zhang</i>	
End-To-End Full-Atom Antibody Design	17409
<i>Xiangzhe Kong, Wenbing Huang, Yang Liu</i>	
Covariate Balancing Using the Integral Probability Metric for Causal Inference	17430
<i>Insung Kong, Yuha Park, Joonhyuk Jung, Kwonsang Lee, Yongdai Kim</i>	
Masked Bayesian Neural Networks : Theoretical Guarantee and Its Posterior Inference	17462
<i>Insung Kong, Dongyoon Yang, Jongjin Lee, Ilsang Ohn, Gyuseung Baek, Yongdai Kim</i>	
Parameter-Level Soft-Masking for Continual Learning	17492
<i>Tatsuya Konishi, Mori Kurokawa, Chihiro Ono, Zixuan Ke, Gyuhak Kim, Bing Liu</i>	
Pretraining Language Models with Human Preferences	17506
<i>Tomasz Korbak, Kejian Shi, Angelica Chen, Rasika Vinayak Bhalerao, Christopher Buckley, Jason Phang, Samuel R. Bowman, Ethan Perez</i>	
Detecting Adversarial Directions in Deep Reinforcement Learning to Make Robust Decisions	17534
<i>Ezgi Korkmaz, Jonah Brown-Cohen</i>	
Ewald-Based Long-Range Message Passing for Molecular Graphs	17544
<i>Arthur Kosmala, Johannes Gasteiger, Nicholas Gao, Stephan Günnemann</i>	
TabDDPM: Modelling Tabular Data with Diffusion Models	17564
<i>Akim Kotelnikov, Dmitry Baranchuk, Ivan Rubachev, Artem Babenko</i>	
Randomized Schur Complement Views for Graph Contrastive Learning	17580
<i>Vignesh Kothapalli</i>	
Benign Overfitting in Two-Layer ReLU Convolutional Neural Networks.....	17615
<i>Yiwen Kou, Zixiang Chen, Yuanzhou Chen, Quanquan Gu</i>	

PART 23

Variational Mixture of HyperGenerators for Learning Distributions Over Functions.....	17660
<i>Batuhan Koyuncu, Pablo Sanchez Martin, Ignacio Peis, Pablo M. Olmos, Isabel Valera</i>	
Gradient Descent Monotonically Decreases the Sharpness of Gradient Flow Solutions in Scalar Networks and Beyond	17684
<i>Itai Kreisler, Mor Shpigel Nacson, Daniel Soudry, Yair Carmon</i>	
Estimation Beyond Data Reweighting: Kernel Method of Moments	17745
<i>Heiner Kremer, Yassine Nemmour, Bernhard Schölkopf, Jia-Jie Zhu</i>	
Multi-Task Differential Privacy Under Distribution Skew	17784
<i>Walid Krichene, Prateek Jain, Shuang Song, Mukund Sundararajan, Abhradeep Guha Thakurta, Li Zhang</i>	
Towards Bridging the Gaps Between the Right to Explanation and the Right to Be Forgotten	17808
<i>Satyapriya Krishna, Jiaqi Ma, Himabindu Lakkaraju</i>	
Graph Neural Tangent Kernel: Convergence on Large Graphs	17827
<i>Sanjukta Krishnagopal, Luana Ruiz</i>	
Diffusion Models for Black-Box Optimization	17842
<i>Siddarth Krishnamoorthy, Satvik Mehul Mashkaria, Aditya Grover</i>	
Learning to Design Analog Circuits to Meet Threshold Specifications	17858
<i>Dmitrii Krylov, Pooya Khajeh, Junhan Ouyang, Thomas Reeves, Tongkai Liu, Hiba Ajmal, Hamidreza Aghasi, Roy Fox</i>	
Variance Control for Distributional Reinforcement Learning	17874
<i>Qi Kuang, Zhoufan Zhu, Liwen Zhang, Fan Zhou</i>	
Hierarchical Imitation Learning with Vector Quantized Models.....	17896
<i>Kalle Kujanpää, Joni Pajarinen, Alexander Ilin</i>	

SinDDM: A Single Image Denoising Diffusion Model.....	17920
<i>Vladimir Kulikov, Shahar Yadin, Matan Kleiner, Tomer Michaeli</i>	
Towards Explaining Distribution Shifts	17931
<i>Sean Kulinski, David I. Inouye</i>	
Featured Graph Coarsening with Similarity Guarantees	17953
<i>Manoj Kumar, Anurag Sharma, Shashwat Saxena, Sandeep Kumar</i>	
Modeling Dynamic Environments with Scene Graph Memory.....	17976
<i>Andrey Kurenkov, Michael Lingelbach, Tanmay Agarwal, Emily Jin, Chengshu Li, Ruohan Zhang, Li Fei-Fei, Jiajun Wu, Silvio Savarese, Roberto Marti'n-Marti'n</i>	
Tied-Augment: Controlling Representation Similarity Improves Data Augmentation.....	17994
<i>Emirhan Kurtulus, Zichao Li, Yann Dauphin, Ekin Dogus Cubuk</i>	
Cooperation in the Latent Space: The Benefits of Adding Mixture Components in Variational Autoencoders.....	18008
<i>Oskar Kviman, Ricky Molén, Alexandra Hotti, Semih Kurt, Vi'ctor Elvira, Jens Lagergren</i>	
GeCoNeRF: Few-Shot Neural Radiance Fields Via Geometric Consistency.....	18023
<i>Min-Seop Kwak, Jiuhn Song, Seungryong Kim</i>	
Rotation and Translation Invariant Representation Learning with Implicit Neural Representations	18037
<i>Sehyun Kwon, Joo Young Choi, Ernest K. Ryu</i>	
Reward-Mixing MDPs with Few Latent Contexts Are Learnable.....	18057
<i>Jeongyeol Kwon, Yonathan Efroni, Constantine Caramanis, Shie Mannor</i>	
A Fully First-Order Method for Stochastic Bilevel Optimization	18083
<i>Jeongyeol Kwon, Dohyun Kwon, Stephen Wright, Robert D Nowak</i>	
Complexity of Block Coordinate Descent with Proximal Regularization and Applications to Wasserstein CP-Dictionary Learning	18114
<i>Dohyun Kwon, Hanbaek Lyu</i>	
Data-OOB: Out-Of-Bag Estimate as a Simple and Efficient Data Value	18135
<i>Yongchan Kwon, James Zou</i>	
Emergence of Adaptive Circadian Rhythms in Deep Reinforcement Learning	18153
<i>Aqeel Labash, Florian Stelzer, Daniel Majoral, Raul Vicente Zafra</i>	
Synergies Between Disentanglement and Sparsity: Generalization and Identifiability in Multi-Task Learning	18171
<i>Sebastien Lachapelle, Tristan Deleu, Divyat Mahajan, Ioannis Mitliagkas, Yoshua Bengio, Simon Lacoste-Julien, Quentin Bertrand</i>	
Nearly-Optimal Hierarchical Clustering for Well-Clustered Graphs.....	18207
<i>Steinar Laenen, Bogdan Adrian Manghiuc, He Sun</i>	
Hybrid Energy Based Model in the Feature Space for Out-Of-Distribution Detection.....	18250
<i>Marc Lafon, Elias Ramzi, Clément Rambour, Nicolas Thome</i>	
A Theory of Continuous Generative Flow Networks	18269
<i>Salem Lahlou, Tristan Deleu, Pablo Lemos, Dinghui Zhang, Alexandra Volokhova, Alex Hernández-Garci'a, Lena Nehale Ezzine, Yoshua Bengio, Nikolay Malkin</i>	
Automatically Marginalized MCMC in Probabilistic Programming.....	18301
<i>Jinlin Lai, Javier Burroni, Hui Guan, Daniel Sheldon</i>	
DS-1000: A Natural and Reliable Benchmark for Data Science Code Generation	18319
<i>Yuhang Lai, Chengxi Li, Yiming Wang, Tianyi Zhang, Ruiqi Zhong, Luke Zettlemoyer, Wen-Tau Yih, Daniel Fried, Sida Wang, Tao Yu</i>	
ChiPFormer: Transferable Chip Placement Via Offline Decision Transformer.....	18346
<i>Yao Lai, Jinxin Liu, Zhentao Tang, Bin Wang, Jianye Hao, Ping Luo</i>	
FP-Diffusion: Improving Score-Based Diffusion Models by Enforcing the Underlying Score Fokker-Planck Equation	18365
<i>Chieh-Hsin Lai, Yuhta Takida, Naoki Murata, Toshimitsu Uesaka, Yuki Mitsufuji, Stefano Ermon</i>	

Private Statistical Estimation of Many Quantiles	18399
<i>Clément Lalanne, Aurélien Garivier, Rémi Gribonval</i>	
Bootstrap in High Dimension with Low Computation.....	18419
<i>Henry Lam, Zhenyuan Liu</i>	

PART 24

LegendreTron: Uprising Proper Multiclass Loss Learning	18454
<i>Kevin H Lam, Christian Walder, Spiridon Penev, Richard Nock</i>	
Metagenomic Binning Using Connectivity-Constrained Variational Autoencoders.....	18471
<i>Andre Lamurias, Alessandro Tibo, Katja Hose, Mads Albertsen, Thomas Dyhre Nielsen</i>	
Delay-Adapted Policy Optimization and Improved Regret for Adversarial MDP with Delayed Bandit Feedback	18482
<i>Tal Lincewicz, Aviv Rosenberg, Dmitry Sotnikov</i>	
Lottery Tickets in Evolutionary Optimization: On Sparse Backpropagation-Free Trainability	18535
<i>Robert Tjarko Lange, Henning Sprekeler</i>	
On the Occupancy Measure of Non-Markovian Policies in Continuous MDPs.....	18548
<i>Romain Laroche, Remi Tachet Des Combes</i>	
Minimalistic Predictions to Schedule Jobs with Online Precedence Constraints	18563
<i>Alexandra Anna Lassota, Alexander Lindermayr, Nicole Megow, Jens Schlöter</i>	
Speeding Up Bellman Ford Via Minimum Violation Permutations	18584
<i>Silvio Lattanzi, Ola Svensson, Sergei Vassilvitskii</i>	
Who Needs to Know? Minimal Knowledge for Optimal Coordination	18599
<i>Niklas Lauffer, Ameesh Shah, Micah Carroll, Michael D Dennis, Stuart Russell</i>	
Target-Based Surrogates for Stochastic Optimization	18614
<i>Jonathan Wilder Lavington, Sharan Vaswani, Reza Babanezhad Harikandeh, Mark Schmidt, Nicolas Le Roux</i>	
Cluster Explanation Via Polyhedral Descriptions.....	18652
<i>Connor Lawless, Oktay Gunluk</i>	
Pre-Training for Speech Translation: CTC Meets Optimal Transport	18667
<i>Phuong-Hang Le, Hongyu Gong, Changhan Wang, Juan Pino, Benjamin Lecouteux, Didier Schwab</i>	
Bootstrapped Representations in Reinforcement Learning	18686
<i>Charline Le Lan, Stephen Tu, Mark Rowland, Anna Harutyunyan, Rishabh Agarwal, Marc G Bellemare, Will Dabney</i>	
Strategic Classification with Unknown User Manipulations.....	18714
<i>Tosca Lechner, Ruth Urner, Shai Ben-David</i>	
Learning in POMDPs is Sample-Efficient with Hindsight Observability	18733
<i>Jonathan Lee, Alekh Agarwal, Christoph Dann, Tong Zhang</i>	
Towards Deep Attention in Graph Neural Networks: Problems and Remedies	18774
<i>Soo Yong Lee, Fan Chen Bu, Jaemin Yoo, Kijung Shin</i>	
InGram: Inductive Knowledge Graph Embedding Via Relation Graphs.....	18796
<i>Jaejun Lee, Chanyoung Chung, Joyce Jiyoun Whang</i>	
Optimality of Thompson Sampling with Noninformative Priors for Pareto Bandits.....	18810
<i>Jongyeong Lee, Junya Honda, Chao-Kai Chiang, Masashi Sugiyama</i>	
Conditional Graph Information Bottleneck for Molecular Relational Learning	18852
<i>Namkyeong Lee, Dongmin Hyun, Gyoung S. Na, Sungwon Kim, Junseok Lee, Chanyoung Park</i>	
Exploring Chemical Space with Score-Based Out-Of-Distribution Generation.....	18872
<i>Seul Lee, Jaehyeong Jo, Sung Ju Hwang</i>	

Pix2Struct: Screenshot Parsing as Pretraining for Visual Language Understanding	18893
<i>Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvashi Khandelwal, Peter Shaw, Ming-Wei Chang, Kristina Toutanova</i>	
FlexRound: Learnable Rounding Based on Element-Wise Division for Post-Training Quantization	18913
<i>Jung Hyun Lee, Jeonghoon Kim, Se Jung Kwon, Dongsoo Lee</i>	
CoDi: Co-Evolving Contrastive Diffusion Models for Mixed-Type Tabular Synthesis	18940
<i>Chaejeong Lee, Jayoung Kim, Noseong Park</i>	
Minimizing Trajectory Curvature of ODE-Based Generative Models	18957
<i>Sangyun Lee, Beomsu Kim, Jong Chul Ye</i>	
H-Likelihood Approach to Deep Neural Networks with Temporal-Spatial Random Effects for High-Cardinality Categorical Features	18974
<i>Hangbin Lee, Youngjo Lee</i>	
On the Importance of Feature Decorrelation for Unsupervised Representation Learning in Reinforcement Learning	18988
<i>Hojoon Lee, Koanho Lee, Dongyoon Hwang, Hyunho Lee, Byungkun Lee, Jaegul Choo</i>	
HETAL: Efficient Privacy-Preserving Transfer Learning with Homomorphic Encryption	19010
<i>Seewoo Lee, Garam Lee, Jung Woo Kim, Junbum Shin, Mun-Kyu Lee</i>	
QASA: Advanced Question Answering on Scientific Articles	19036
<i>Yoonjoo Lee, Kyungjae Lee, Sunghyun Park, Dasol Hwang, Jaehyeon Kim, Hong-In Lee, Moontae Lee</i>	
Demystifying Disagreement-On-The-Line in High Dimensions	19053
<i>Donghwan Lee, Behrad Moniri, Xinneng Huang, Edgar Dobriban, Hamed Hassani</i>	
On the Correctness of Automatic Differentiation for Neural Networks with Machine-Representable Parameters	19094
<i>Wonyeol Lee, Sejun Park, Alex Aiken</i>	
Implicit Jacobian Regularization Weighted with Impurity of Probability Output	19141
<i>Sungyoon Lee, Jinseong Park, Jaewook Lee</i>	
Unsupervised Skill Discovery for Learning Shared Structures Across Changing Environments	19185
<i>Sang-Hyun Lee, Seung-Woo Seo</i>	
Generalization Analysis for Contrastive Representation Learning	19200
<i>Yunwen Lei, Tianbao Yang, Yiming Ying, Ding-Xuan Zhou</i>	
Learning Control by Iterative Inversion	19228
<i>Gal Leibovich, Guy Jacob, Or Avner, Gal Novik, Aviv Tamar</i>	

PART 25

Sampling-Based Accuracy Testing of Posterior Estimators for General Inference	19256
<i>Pablo Lemos, Adam Coogan, Yashar Hezaveh, Laurence Perreault-Levasseur</i>	
Fast Inference from Transformers Via Speculative Decoding	19274
<i>Yaniv Leviathan, Matan Kalman, Yossi Matias</i>	
Efficient Rate Optimal Regret for Adversarial Contextual MDPs Using Online Function Approximation	19287
<i>Orin Levy, Alon Cohen, Asaf Cassel, Yishay Mansour</i>	
GLOBE-CE: A Translation Based Approach for Global Counterfactual Explanations	19315
<i>Dan Ley, Saumitra Mishra, Daniele Magazzeni</i>	
TIPS: Topologically Important Path Sampling for Anytime Neural Networks	19343
<i>Guihong Li, Kartikeya Bhardwaj, Yuedong Yang, Radu Marculescu</i>	
MAHALO: Unifying Offline Reinforcement Learning and Imitation Learning from Observations	19360
<i>Anqi Li, Byron Boots, Ching-An Cheng</i>	
Internet Explorer: Targeted Representation Learning on the Open Web	19385
<i>Alexander Cong Li, Ellis Langham Brown, Alexei A Efros, Deepak Pathak</i>	

Prototype-Oriented Unsupervised Anomaly Detection for Multivariate Time Series	19407
<i>Yuxin Li, Wenchao Chen, Bo Chen, Dongsheng Wang, Long Tian, Mingyuan Zhou</i>	
Learning Preconditioners for Conjugate Gradient PDE Solvers	19425
<i>Yichen Li, Peter Yichen Chen, Tao Du, Wojciech Matusik</i>	
Parallel Q-Learning: Scaling Off-Policy Reinforcement Learning Under Massively Parallel Simulation	19440
<i>Zechu Li, Tao Chen, Zhang-Wei Hong, Anurag Ajay, Pulkit Agrawal</i>	
Minimum Width of Leaky-ReLU Neural Networks for Uniform Universal Approximation	19460
<i>Li'Ang Li, Yifei Duan, Guanghua Ji, Yongqiang Cai</i>	
FAIRER: Fairness as Decision Rationale Alignment	19471
<i>Tianlin Li, Qing Guo, Aishan Liu, Mengnan Du, Zhiming Li, Yang Liu</i>	
RACE: Improve Multi-Agent Reinforcement Learning with Representation Asymmetry and Collaborative Evolution.....	19490
<i>Pengyi Li, Jianye Hao, Hongyao Tang, Yan Zheng, Xian Fu</i>	
Adversarial Collaborative Learning on Non-IID Features	19504
<i>Qinbin Li, Bingsheng He, Dawn Song</i>	
Near-Optimal Conservative Exploration in Reinforcement Learning Under Episode-Wise Constraints.....	19527
<i>Donghao Li, Ruiquan Huang, Cong Shen, Jing Yang</i>	
Transformers as Algorithms: Generalization and Stability in In-Context Learning	19565
<i>Yingcong Li, Muhammed Emrullah Ildiz, Dimitris Papailiopoulos, Samet Oymak</i>	
Improving Hyperparameter Learning Under Approximate Inference in Gaussian Process Models.....	19595
<i>Rui Li, S. T. John, Arno Solin</i>	
Local Vertex Colouring Graph Neural Networks	19616
<i>Shouheng Li, Dongwoo Kim, Qing Wang</i>	
Analysis of Error Feedback in Federated Non-Convex Optimization with Biased Compression: Fast Convergence and Partial Participation.....	19638
<i>Xiaoyun Li, Ping Li</i>	
How Do Transformers Learn Topic Structure: Towards a Mechanistic Understanding	19689
<i>Yuchen Li, Yuanzhi Li, Andrej Risteski</i>	
BLIP-2: Bootstrapping Language-Image Pre-Training with Frozen Image Encoders and Large Language Models	19730
<i>Junnan Li, Dongxu Li, Silvio Savarese, Steven Hoi</i>	
Nearly Optimal Algorithms with Sublinear Computational Complexity for Online Kernel Regression	19743
<i>Junfan Li, Shizhong Liao</i>	
Revisiting Weighted Aggregation in Federated Learning with Neural Networks.....	19767
<i>Zexi Li, Tao Lin, Xinyi Shang, Chao Wu</i>	
Distribution-Dependent McDiarmid-Type Inequalities for Functions of Unbounded Interaction.....	19789
<i>Shaojie Li, Yong Liu</i>	
Optimal Convergence Rates for Agnostic Nyström Kernel Learning	19811
<i>Jian Li, Yong Liu, Weiping Wang</i>	
Reconstructive Neuron Pruning for Backdoor Defense.....	19837
<i>Yige Li, Xixiang Lyu, Xingjun Ma, Nodens Koren, Lingjuan Lyu, Bo Li, Yu-Gang Jiang</i>	
Meta Learning of Interface Conditions for Multi-Domain Physics-Informed Neural Networks.....	19855
<i>Shibo Li, Michael Penwarden, Yiming Xu, Conor Tillinghast, Akil Narayan, Mike Kirby, Shandian Zhe</i>	
Deep Anomaly Detection Under Labeling Budget Constraints.....	19882
<i>Aodong Li, Chen Qiu, Marius Kloft, Padhraic Smyth, Stephan Mandt, Maja Rudolph</i>	
On the Initialization of Graph Neural Networks	19911
<i>Jiahang Li, Yakun Song, Xiang Song, David Wipf</i>	

Federated Adversarial Learning: A Framework with Convergence Analysis	19932
<i>Xiaoxiao Li, Zhao Song, Jiaming Yang</i>	
How Powerful Are Shallow Neural Networks with Bandlimited Random Weights?	19960
<i>Ming Li, Sho Sonoda, Feilong Cao, Yu Guang Wang, Jiye Liang</i>	
Efficient Quantum Algorithms for Quantum Optimal Control	19982
<i>Xiantao Li, Chunhao Wang</i>	
Low-Switching Policy Gradient with Exploration Via Online Sensitivity Sampling	19995
<i>Yunfan Li, Yiran Wang, Yu Cheng, Lin Yang</i>	
Hierarchical Diffusion for Offline Decision Making	20035
<i>Wenhao Li, Xiangfeng Wang, Bo Jin, Hongyuan Zha</i>	

PART 26

Divide and Conquer Dynamic Programming: An Almost Linear Time Change Point Detection Methodology in High Dimensions	20065
<i>Wanshan Li, Daren Wang, Alessandro Rinaldo</i>	
Architecture-Agnostic Masked Image Modeling -- From ViT Back to CNN	20149
<i>Siyuan Li, Di Wu, Fang Wu, Zelin Zang, Stan Z. Li</i>	
Learning Antidote Data to Individual Unfairness	20168
<i>Peizhao Li, Ethan Xia, Hongfu Liu</i>	
Propensity Matters: Measuring and Enhancing Balancing for Recommendation	20182
<i>Haoxuan Li, Yanghao Xiao, Chunyuan Zheng, Peng Wu, Peng Cui</i>	
GraphCleaner: Detecting Mislabelled Samples in Popular Graph Learning Benchmarks	20195
<i>Yuwen Li, Miao Xiong, Bryan Hooi</i>	
SMURF-THP: Score Matching-Based Uncertainty Quantification for Transformer Hawkes Process	20210
<i>Zichong Li, Yanbo Xu, Simiao Zuo, Haoming Jiang, Chao Zhang, Tuo Zhao, Hongyuan Zha</i>	
Horizon-Free Learning for Markov Decision Processes and Games: Stochastically Bounded Rewards and Improved Bounds	20221
<i>Shengshi Li, Lin Yang</i>	
Transcendental Idealism of Planner: Evaluating Perception from Planning Perspective for Autonomous Driving	20253
<i>Weixin Li, Xiaodong Yang</i>	
Learning for Edge-Weighted Online Bipartite Matching with Robustness Guarantees	20276
<i>Pengfei Li, Jianyi Yang, Shaolei Ren</i>	
FedVS: Straggler-Resilient and Privacy-Preserving Vertical Federated Learning for Split Models	20296
<i>Songze Li, Duanyi Yao, Jin Liu</i>	
Achieving Hierarchy-Free Approximation for Bilevel Programs with Equilibrium Constraints	20312
<i>Jiayang Li, Jing Yu, Boyi Liu, Yu Nie, Zhaoran Wang</i>	
LoSparse: Structured Compression of Large Language Models Based on Low-Rank and Sparse Approximation	20336
<i>Yixiao Li, Yifan Yu, Qingru Zhang, Chen Liang, Pengcheng He, Weizhu Chen, Tuo Zhao</i>	
Nesterov Meets Optimism: Rate-Optimal Separable Minimax Optimization	20351
<i>Chris Junchi Li, Huizhuo Yuan, Gauthier Gidel, Quanquan Gu, Michael Jordan</i>	
Alternating Local Enumeration (TnALE): Solving Tensor Network Structure Search with Fewer Evaluations	20384
<i>Chao Li, Junhua Zeng, Chunmei Li, Cesar F Caiafa, Qibin Zhao</i>	
Understanding the Complexity Gains of Single-Task RL with a Curriculum	20412
<i>Qiyang Li, Yuexiang Zhai, Yi Ma, Sergey Levine</i>	
Does a Neural Network Really Encode Symbolic Concepts?	20452
<i>Mingjie Li, Quanshi Zhang</i>	

Cooperative Open-Ended Learning Framework for Zero-Shot Coordination	20470
<i>Yang Li, Shao Zhang, Jichen Sun, Yali Du, Ying Wen, Xinning Wang, Wei Pan</i>	
Offline Reinforcement Learning with Closed-Form Policy Improvement Operators	20485
<i>Jiachen Li, Edwin Zhang, Ming Yin, Qinxun Bai, Yu-Xiang Wang, William Yang Wang</i>	
Optimal Arms Identification with Knapsacks.....	20529
<i>Shaoang Li, Lan Zhang, Yingqi Yu, Xiangyang Li</i>	
Internally Rewarded Reinforcement Learning	20556
<i>Mengdi Li, Xufeng Zhao, Jae Hee Lee, Cornelius Weber, Stefan Wermter</i>	
Trustworthy Policy Learning Under the Counterfactual No-Harm Criterion	20575
<i>Haoxuan Li, Chunyuan Zheng, Yixiao Cao, Zhi Geng, Yue Liu, Peng Wu</i>	
Structured Cooperative Learning with Graphical Model Priors	20599
<i>Shuangtong Li, Tianyi Zhou, Xinmei Tian, Dacheng Tao</i>	
Low Complexity Homeomorphic Projection to Ensure Neural-Network Solution Feasibility for Optimization Over (Non-)Convex Set.....	20623
<i>Enming Liang, Minghua Chen, Steven Low</i>	
Consistency of Multiple Kernel Clustering	20650
<i>Weixuan Liang, Xinwang Liu, Yong Liu, Chuan Ma, Yunping Zhao, Zhe Liu, En Zhu</i>	
A Distribution Optimization Framework for Confidence Bounds of Risk Measures	20677
<i>Hao Liang, Zhi-Quan Luo</i>	
Accuracy on the Curve: On the Nonlinear Correlation of ML Performance Between Data Subpopulations	20706
<i>Weixin Liang, Yining Mao, Yongchan Kwon, Xinyu Yang, James Zou</i>	
AdaptDiffuser: Diffusion Models as Adaptive Self-Evolving Planners	20725
<i>Zhixuan Liang, Yao Mu, Mingyu Ding, Fei Ni, Masayoshi Tomizuka, Ping Luo</i>	
Learning Compiler Pass Orders Using Coreset and Normalized Value Prediction	20746
<i>Youwei Liang, Kevin Stone, Ali Shameli, Chris Cummins, Mostafa Elhoushi, Jiadong Guo, Benoit Steiner, Xiaomeng Yang, Pengtao Xie, Hugh James Leather, Yuandong Tian</i>	
Adversarial Example Does Good: Preventing Painting Imitation from Diffusion Models Via Adversarial Examples.....	20763
<i>Chumeng Liang, Xiaoyu Wu, Yang Hua, Jiaru Zhang, Yiming Xue, Tao Song, Zhengui Xue, Ruhui Ma, Haibing Guan</i>	
CLUSTSEG: Clustering for Universal Segmentation	20787
<i>James Chenhao Liang, Tianfei Zhou, Dongfang Liu, Wenguan Wang</i>	
Conformal Inference is (almost) Free for Neural Networks Trained with Early Stopping.....	20810
<i>Ziyi Liang, Yanfei Zhou, Matteo Sesia</i>	
Less is More: Task-Aware Layer-Wise Distillation for Language Model Compression	20852
<i>Chen Liang, Simiao Zuo, Qingru Zhang, Pengcheng He, Weizhu Chen, Tuo Zhao</i>	
Statistical Inference and A/B Testing for First-Price Pacing Equilibria.....	20868
<i>Luofeng Liao, Christian Kroer</i>	

PART 27

Supervised Metric Learning to Rank for Retrieval Via Contextual Similarity Optimization	20906
<i>Christopher Liao, Theodoros Tsiligkaridis, Brian Kulis</i>	
Revisiting Domain Randomization Via Relaxed State-Adversarial Policy Optimization	20939
<i>Yun-Hsuan Lien, Ping-Chun Hsieh, Yu-Shuen Wang</i>	
Variational Open-Domain Question Answering	20950
<i>Valentin Liévin, Andreas Geert Motzfeldt, Ida Riis Jensen, Ole Winther</i>	
Generating Novel, Designable, and Diverse Protein Structures by Equivariantly Diffusing Oriented Residue Clouds.....	20978
<i>Yeqing Lin, Mohammed Alquraishi</i>	

Hyperbolic Diffusion Embedding and Distance for Hierarchical Representation Learning.....	21003
<i>Ya-Wei Eileen Lin, Ronald R. Coifman, Gal Mishne, Ronen Talmon</i>	
Simplifying Momentum-Based Positive-Definite Submanifold Optimization with Applications to Deep Learning.....	21026
<i>Wu Lin, Valentin Duruisseaux, Melvin Leok, Frank Nielsen, Mohammad Emtiyaz Khan, Mark Schmidt</i>	
Text Generation with Diffusion Language Models: A Pre-Training Approach with Continuous Paragraph Denoise.....	21051
<i>Zhenghao Lin, Yeyun Gong, Yelong Shen, Tong Wu, Zhihao Fan, Chen Lin, Nan Duan, Weizhu Chen</i>	
Self-Supervised Neural Factor Analysis for Disentangling Utterance-Level Speech Representations	21065
<i>Weiwei Lin, Chenhang He, Man-Wai Mak, Youzhi Tu</i>	
Theory on Forgetting and Generalization of Continual Learning.....	21078
<i>Sen Lin, Peizhong Ju, Yingbin Liang, Ness Shroff</i>	
Accelerated Cyclic Coordinate Dual Averaging with Extrapolation for Composite Convex Optimization.....	21101
<i>Cheuk Yin Lin, Chaobing Song, Jelena Diakonikolas</i>	
Safe Offline Reinforcement Learning with Real-Time Budget Constraints	21127
<i>Qian Lin, Bo Tang, Zifan Wu, Chao Yu, Shangqin Mao, Qianlong Xie, Xingxing Wang, Dong Wang</i>	
Probabilistic Unrolling: Scalable, Inverse-Free Maximum Likelihood Estimation for Latent Gaussian Models	21153
<i>Alexander Lin, Bahareh Tolooshams, Yves Atchade, Demba E. Ba</i>	
Fast Online Value-Maximizing Prediction Sets with Conformal Cost Control	21182
<i>Zhen Lin, Shubhendu Trivedi, Cao Xiao, Jimeng Sun</i>	
Unveiling the Mask of Position-Information Pattern Through the Mist of Image Features	21204
<i>Chieh Hubert Lin, Hung-Yu Tseng, Hsin-Ying Lee, Maneesh Kumar Singh, Ming-Hsuan Yang</i>	
Fair Yet Asymptotically Equal Collaborative Learning.....	21223
<i>Xiaoqiang Lin, Xinyi Xu, See-Kiong Ng, Chuan-Sheng Foo, Bryan Kian Hsiang Low</i>	
Efficient Approximations of Complete Interatomic Potentials for Crystal Property Prediction.....	21260
<i>Yuchao Lin, Keqiang Yan, Youzhi Luo, Yi Liu, Xiaoning Qian, Shuiwang Ji</i>	
Continuation Path Learning for Homotopy Optimization	21288
<i>Xi Lin, Zhiyuan Yang, Xiaoyuan Zhang, Qingfu Zhang</i>	
Speed-Oblivious Online Scheduling: Knowing (Precise) Speeds is Not Necessary	21312
<i>Alexander Lindermayr, Nicole Megow, Martin Rapp</i>	
Graph Mixup with Soft Alignments	21335
<i>Hongyi Ling, Zhimeng Jiang, Meng Liu, Shuiwang Ji, Na Zou</i>	
Deep Graph Representation Learning and Optimization for Influence Maximization.....	21350
<i>Chen Ling, Junji Jiang, Junxiang Wang, My T. Thai, Renhao Xue, James Song, Meikang Qiu, Liang Zhao</i>	
Emergent Agentic Transformer from Chain of Hindsight Experience.....	21362
<i>Hao Liu, Pieter Abbeel</i>	
Shapley Based Residual Decomposition for Instance Analysis.....	21375
<i>Tommy Liu, Amanda S Barnard</i>	
Learning Representations Without Compositional Assumptions.....	21388
<i>Tennison Liu, Jeroen Berrevoets, Zhaozhi Qian, Mihaela Van Der Schaar</i>	
Byzantine-Robust Learning on Heterogeneous Data Via Gradient Splitting.....	21404
<i>Yuchen Liu, Chen Chen, Lingjuan Lyu, Fangzhao Wu, Sai Wu, Gang Chen</i>	
Towards Constituting Mathematical Structures for Learning to Optimize	21426
<i>Jialin Liu, Xiaohan Chen, Zhangyang Wang, Wotao Yin, Hanqin Cai</i>	

AudioLDM: Text-To-Audio Generation with Latent Diffusion Models.....	21450
<i>Haohe Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo Mandic, Wenwu Wang, Mark D Plumbley</i>	
Identifiability of Label Noise Transition Matrix.....	21475
<i>Yang Liu, Hao Cheng, Kun Zhang</i>	
A Group Symmetric Stochastic Differential Equation Model for Molecule Multi-Modal Pretraining	21497
<i>Shengchao Liu, Weitao Du, Zhi-Ming Ma, Hongyu Guo, Jian Tang</i>	
Using Perturbation to Improve Goodness-Of-Fit Tests Based on Kernelized Stein Discrepancy	21527
<i>Xing Liu, Andrew B. Duncan, Axel Gandy</i>	
Cones: Concept Neurons in Diffusion Models for Customized Generation.....	21548
<i>Zhiheng Liu, Ruili Feng, Kai Zhu, Yifei Zhang, Kecheng Zheng, Yu Liu, Deli Zhao, Jingren Zhou, Yang Cao</i>	
Opponent-Limited Online Search for Imperfect Information Games.....	21567
<i>Weiming Liu, Haobo Fu, Qiang Fu, Yang Wei</i>	
Towards Robust and Safe Reinforcement Learning with Benign Off-Policy Data	21586
<i>Zuxin Liu, Zijian Guo, Zhepeng Cen, Huan Zhang, Yihang Yao, Hanjiang Hu, Ding Zhao</i>	
Constrained Decision Transformer for Offline Safe Reinforcement Learning	21611
<i>Zuxin Liu, Zijian Guo, Yihang Yao, Zhepeng Cen, Wenhao Yu, Tingnan Zhang, Ding Zhao</i>	
Understanding and Defending Patched-Based Adversarial Attacks for Vision Transformer.....	21631
<i>Liang Liu, Yanan Guo, Youtao Zhang, Jun Yang</i>	
NUNO: A General Framework for Learning Parametric PDEs with Non-Uniform Data	21658
<i>Songming Liu, Zhongkai Hao, Chengyang Ying, Hang Su, Ze Cheng, Jun Zhu</i>	
Hierarchical Programmatic Reinforcement Learning Via Learning to Compose Programs.....	21672
<i>Guan-Ting Liu, En-Pei Hu, Pu-Jen Cheng, Hung-Yi Lee, Shao-Hua Sun</i>	

PART 28

Online Local Differential Private Quantile Inference Via Self-Normalization.....	21698
<i>Yi Liu, Qirui Hu, Lei Ding, Linglong Kong</i>	
GFlowOut: Dropout with Generative Flow Networks	21715
<i>Dianbo Liu, Moksh Jain, Bonaventure F. P. Dossou, Qianli Shen, Salem Lahlou, Anirudh Goyal, Nikolay Malkin, Chris Chinenye Emezue, Dinghuai Zhang, Nadhir Hassen, Xu Ji, Kenji Kawaguchi, Yoshua Bengio</i>	
2D-Shapley: A Framework for Fragmented Data Valuation.....	21730
<i>Zhihong Liu, Hoang Anh Just, Xiangyu Chang, Xi Chen, Ruoxi Jia</i>	
Causal Structure Learning for Latent Intervened Non-Stationary Data	21756
<i>Chenxi Liu, Kun Kuang</i>	
Structural Re-Weighting Improves Graph Domain Adaptation	21778
<i>Shikun Liu, Tianchun Li, Yongbin Feng, Nhan Tran, Han Zhao, Qiang Qiu, Pan Li</i>	
Dink-Net: Neural Clustering on Large Graphs.....	21794
<i>Yue Liu, Ke Liang, Jun Xia, Sihang Zhou, Xihong Yang, Xinwang Liu, Stan Z. Li</i>	
Oscillation-Free Quantization for Low-Bit Vision Transformers	21813
<i>Shih-Yang Liu, Zechun Liu, Kwang-Ting Cheng</i>	
Understanding the Distillation Process from Deep Generative Models to Tractable Probabilistic Circuits	21825
<i>Xuejie Liu, Anji Liu, Guy Van Den Broeck, Yitao Liang</i>	
Averaged Method of Multipliers for Bi-Level Optimization Without Lower-Level Strong Convexity	21839
<i>Risheng Liu, Yaohua Liu, Wei Yao, Shangzhi Zeng, Jin Zhang</i>	
Graph Switching Dynamical Systems	21867
<i>Yongtuo Liu, Sara Magliacane, Miltiadis Kofinas, Efstratios Gavves</i>	
High Probability Convergence of Stochastic Gradient Methods.....	21884
<i>Zijian Liu, Ta Duy Nguyen, Thien Hang Nguyen, Alina Ene, Huy Nguyen</i>	

OMS-DPM: Optimizing the Model Schedule for Diffusion Probabilistic Models.....	21915
<i>Enshu Liu, Xuefei Ning, Zinan Lin, Huazhong Yang, Yu Wang</i>	
Lazy Agents: A New Perspective on Solving Sparse Reward Problem in Multi-Agent Reinforcement Learning.....	21937
<i>Boyin Liu, Zhiqiang Pu, Yi Pan, Jianqiang Yi, Yanyan Liang, D. Zhang</i>	
RSC: Accelerate Graph Neural Networks Training Via Randomized Sparse Computations.....	21951
<i>Zirui Liu, Chen Shengyuan, Kaixiong Zhou, Daochen Zha, Xiao Huang, Xia Hu</i>	
Algorithms for Bounding Contribution for Histogram Estimation Under User-Level Privacy.....	21969
<i>Yuhan Liu, Ananda Theertha Suresh, Wennan Zhu, Peter Kairouz, Marco Gruteser</i>	
Simple Embodied Language Learning as a Byproduct of Meta-Reinforcement Learning.....	21997
<i>Evan Zheran Liu, Sahaana Suri, Tong Mu, Allan Zhou, Chelsea Finn</i>	
Generating Private Synthetic Data with Genetic Algorithms	22009
<i>Terrance Liu, Jingwu Tang, Giuseppe Vietri, Steven Wu</i>	
FusionRetro: Molecule Representation Fusion Via In-Context Learning for Retrosynthetic Planning.....	22028
<i>Songtao Liu, Zhengkai Tu, Minkai Xu, Zuobai Zhang, Lu Lin, Rex Ying, Jian Tang, Peilin Zhao, Dinghao Wu</i>	
I2 SB: Image-To-Image Schrödinger Bridge.....	22042
<i>Guan-Hong Liu, Arash Vahdat, De-An Huang, Evangelos Theodorou, Weili Nie, Anima Anandkumar</i>	
What Can Online Reinforcement Learning with Function Approximation Benefit from General Coverage Conditions?	22063
<i>Fanghui Liu, Luca Viano, Volkan Cevher</i>	
TR0N: Translator Networks for 0-Shot Plug-And-Play Conditional Generation.....	22092
<i>Zhaoyan Liu, Noël Vouitsis, Satya Krishna Gorti, Jimmy Ba, Gabriel Loaiza-Ganem</i>	
Global Optimization with Parametric Function Approximation.....	22113
<i>Chong Liu, Yu-Xiang Wang</i>	
Deja Vu: Contextual Sparsity for Efficient LLMs at Inference Time.....	22137
<i>Zichang Liu, Jue Wang, Tri Dao, Tianyi Zhou, Binhang Yuan, Zhao Song, Anshumali Shrivastava, Ce Zhang, Yuandong Tian, Christopher Re, Beidi Chen</i>	
Trapdoor Normalization with Irreversible Ownership Verification.....	22177
<i>Hanwen Liu, Zhenyu Weng, Yuesheng Zhu, Yadong Mu</i>	
Same Pre-Training Loss, Better Downstream: Implicit Bias Matters for Language Models	22188
<i>Hong Liu, Sang Michael Xie, Zhiyuan Li, Tengyu Ma</i>	
Taxonomy-Structured Domain Adaptation.....	22215
<i>Tianyi Liu, Zihao Xu, Hao He, Guang-Yuan Hao, Guang-He Lee, Hao Wang</i>	
Dropout Reduces Underfitting.....	22233
<i>Zhuang Liu, Zhiqiu Xu, Joseph Jin, Zhiqiang Shen, Trevor Darrell</i>	
Revisiting Pseudo-Label for Single-Positive Multi-Label Learning	22249
<i>Biao Liu, Ning Xu, Jiaqi Lv, Xin Geng</i>	
Retrosynthetic Planning with Dual Value Networks	22266
<i>Guoqing Liu, Di Xue, Shufang Xie, Yingce Xia, Austin Tripp, Krzysztof Maziarz, Marwin Segler, Tao Qin, Zongzhang Zhang, Tie-Yan Liu</i>	
Online Nonstochastic Control with Adversarial and Static Constraints	22277
<i>Xin Liu, Zixian Yang, Lei Ying</i>	
Optimization for Amortized Inverse Problems.....	22289
<i>Tianci Liu, Tong Yang, Quan Zhang, Qi Lei</i>	
Active Policy Improvement from Multiple Black-Box Oracles.....	22320
<i>Xuefeng Liu, Takuma Yoneda, Chaoqi Wang, Matthew Walter, Yuxin Chen</i>	
Gradient-Based Wang-Landau Algorithm: A Novel Sampler for Output Distribution of Neural Networks Over the Input Space.....	22338
<i>Weitang Liu, Yi-Zhuang You, Ying Wai Li, Jingbo Shang</i>	

VectorMapNet: End-To-End Vectorized HD Map Learning.....	22352
<i>Yicheng Liu, Tianyuan Yuan, Yue Wang, Yilun Wang, Hang Zhao</i>	
Partially Observable Multi-Agent RL with (Quasi-)Efficiency: The Blessing of Information Sharing	22370
<i>Xiangyu Liu, Kaiqing Zhang</i>	
Prometheus: Taming Sample and Communication Complexities in Constrained Decentralized Stochastic Bilevel Learning.....	22420
<i>Zhuqing Liu, Xin Zhang, Prashant Khanduri, Songtao Lu, Jia Liu</i>	
D2Match: Leveraging Deep Learning and Degeneracy for Subgraph Matching	22454
<i>Xuanzhou Liu, Lin Zhang, Jiaqi Sun, Yujie Yang, Haiqin Yang</i>	
Image Shortcut Squeezing: Countering Perturbative Availability Poisons with Compression.....	22473
<i>Zhuoran Liu, Zhengyu Zhao, Martha Larson</i>	

PART 29

Which Invariance Should We Transfer? a Causal Minimax Learning Approach	22488
<i>Mingzhou Liu, Xiangyu Zheng, Xinwei Sun, Fang Fang, Yizhou Wang</i>	
Unsupervised Out-Of-Distribution Detection with Diffusion Inpainting	22528
<i>Zhenzhen Liu, Jin Peng Zhou, Yufan Wang, Kilian Q Weinberger</i>	
NA2 Q: Neural Attention Additive Model for Interpretable Multi-Agent Q-Learning	22539
<i>Zichuan Liu, Yuanyang Zhu, Chunlin Chen</i>	
Contextual Combinatorial Bandits with Probabilistically Triggered Arms	22559
<i>Xutong Liu, Jinhang Zuo, Siwei Wang, John C. S. Lui, Mohammad Hajiesmaili, Adam Wierman, Wei Chen</i>	
Flipping Coins to Estimate Pseudocounts for Exploration in Reinforcement Learning	22594
<i>Sam Lobel, Akhil Bagaria, George Konidaris</i>	
Multi-Symmetry Ensembles: Improving Diversity and Generalization Via Opposing Symmetries	22614
<i>Charlotte Loh, Seungwook Han, Shivchander Sudalairaj, Rumen Dangovski, Kai Xu, Florian Wenzel, Marin Soljacic, Akash Srivastava</i>	
The Flan Collection: Designing Data and Methods for Effective Instruction Tuning	22631
<i>Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, Adam Roberts</i>	
Dataset Distillation with Convexified Implicit Gradients	22649
<i>Noel Loo, Ramin Hasani, Mathias Lechner, Daniela Rus</i>	
Reflected Diffusion Models.....	22675
<i>Aaron Lou, Stefano Ermon</i>	
Never Mind the Metrics---What About the Uncertainty? Visualising Binary Confusion Matrix Metric Distributions to Put Performance in Perspective.....	22702
<i>David Lovell, Dimity Miller, Jaiden Capra, Andrew P. Bradley</i>	
Bilevel Optimization with Coupled Decision-Dependent Distributions.....	22758
<i>Songtao Lu</i>	
Two-Scale Gradient Descent Ascent Dynamics Finds Mixed Nash Equilibria of Continuous Games: A Mean-Field Perspective	22790
<i>Yulong Lu</i>	
STEP: Learning N:M Structured Sparsity Masks from Scratch with Precondition.....	22812
<i>Yucheng Lu, Shivani Agrawal, Suvinay Subramanian, Oleg Rybakov, Christopher De Sa, Amir Yazdanbakhsh</i>	
Contrastive Energy Prediction for Exact Energy-Guided Diffusion Sampling in Offline Reinforcement Learning.....	22825
<i>Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, Jun Zhu</i>	
Exploring the Limits of Model-Targeted Indiscriminate Data Poisoning Attacks.....	22856
<i>Yiwei Lu, Gautam Kamath, Yaoliang Yu</i>	

QAS-Bench: Rethinking Quantum Architecture Search and a Benchmark.....	22880
<i>Xudong Lu, Kaisen Pan, Ge Yan, Jiaming Shan, Wenjie Wu, Junchi Yan</i>	
Learning Dense Correspondences Between Photos and Sketches.....	22899
<i>Xuanchen Lu, Xiaolong Wang, Judith E Fan</i>	
Adversarial Cheap Talk	22917
<i>Chris Lu, Timon Willi, Alistair Letcher, Jakob Nicolaus Foerster</i>	
Federated Conformal Predictors for Distributed Uncertainty Quantification.....	22942
<i>Charles Lu, Yaodong Yu, Sai Praneeth Karimireddy, Michael Jordan, Ramesh Raskar</i>	
Mechanistic Mode Connectivity.....	22965
<i>Ekdeep Singh Lubana, Eric J Bigelow, Robert P. Dick, David Krueger, Hidenori Tanaka</i>	
A Unifying Framework to the Analysis of Interaction Methods Using Synergy Functions	23005
<i>Daniel Lundstrom, Meisam Razaviyayn</i>	
SegCLIP: Patch Aggregation with Learnable Centers for Open-Vocabulary Semantic Segmentation.....	23033
<i>Huaishao Luo, Junwei Bao, Youzheng Wu, Xiaodong He, Tianrui Li</i>	
Image Restoration with Mean-Reverting Stochastic Differential Equations.....	23045
<i>Ziwei Luo, Fredrik K. Gustafsson, Zheng Zhao, Jens Sjölund, Thomas B. Schön</i>	
Dimensionality Reduction for General KDE Mode Finding	23067
<i>Xinyu Luo, Christopher Musco, Cas Widdershoven</i>	
Iterative Approximate Cross-Validation	23083
<i>Yuetian Luo, Zhimei Ren, Rina Barber</i>	
A Closer Look at Few-Shot Classification Again.....	23103
<i>Xu Luo, Hao Wu, Ji Zhang, Lianli Gao, Jing Xu, Jingkuan Song</i>	
HOPE: High-Order Graph ODE for Modeling Interacting Dynamics	23124
<i>Xiao Luo, Jingyang Yuan, Zijie Huang, Huiyu Jiang, Yifang Qin, Wei Ju, Ming Zhang, Yizhou Sun</i>	
Stabilizing GANs' Training with Brownian Motion Controller	23140
<i>Tianjiao Luo, Ziyu Zhu, Jianfei Chen, Jun Zhu</i>	
OCD: Learning to Overfit with Conditional Diffusion Models.....	23157
<i>Shahar Lutati, Lior Wolf</i>	
DiscoBAX: Discovery of Optimal Intervention Sets in Genomic Experiment Design	23170
<i>Clare Lyle, Arash Mehrjou, Pascal Notin, Andrew Jesson, Stefan Bauer, Yarin Gal, Patrick Schwab</i>	
Understanding Plasticity in Neural Networks	23190
<i>Clare Lyle, Zeyu Zheng, Evgenii Nikishin, Bernardo Avila Pires, Razvan Pascanu, Will Dabney</i>	
Bandits with Knapsacks: Advice on Time-Varying Demands	23212
<i>Lixing Lyu, Wang Chi Cheung</i>	
Pairwise Ranking Losses of Click-Through Rates Prediction for Welfare Maximization in Ad Auctions	23239
<i>Boxiang Lyu, Zhe Feng, Zachary Robertson, Sanmi Koyejo</i>	
Which Tricks Are Important for Learning to Rank?.....	23264
<i>Ivan Lyzhin, Aleksei Ustimenko, Andrey Gulin, Liudmila Prokhorenkova</i>	
Learning Neural Constitutive Laws from Motion Observations for Generalizable PDE Dynamics	23279
<i>Pingchuan Ma, Peter Yichen Chen, Bolei Deng, Joshua B. Tenenbaum, Tao Du, Chuang Gan, Wojciech Matusik</i>	

PART 30

LIV: Language-Image Representations and Rewards for Robotic Control	23301
<i>Yecheng Jason Ma, Vikash Kumar, Amy Zhang, Osbert Bastani, Dinesh Jayaraman</i>	

Graph Inductive Biases in Transformers Without Message Passing.....	23321
<i>Liheng Ma, Chen Lin, Derek Lim, Adriana Romero-Soriano, Puneet K. Dokania, Mark Coates, Philip Torr, Ser-Nam Lim</i>	
Learning Signed Distance Functions from Noisy 3D Point Clouds Via Noise to Noise Mapping	23338
<i>Baorui Ma, Yu-Shen Liu, Zhizhong Han</i>	
Learning Intuitive Policies Using Action Features	23358
<i>Mingwei Ma, Jizhou Liu, Samuel Sokota, Max Kleiman-Weiner, Jakob Nicolaus Foerster</i>	
Over-Parametrization Via Lifting for Low-Rank Matrix Sensing: Conversion of Spurious Solutions to Strict Saddle Points	23373
<i>Ziye Ma, Igor Molybog, Javad Lavaei, Somayeh Sojoudi</i>	
Buying Information for Stochastic Optimization	23388
<i>Mingchen Ma, Christos Tzamos</i>	
Generated Graph Detection	23412
<i>Yihan Ma, Zhikun Zhang, Ning Yu, Xinlei He, Michael Backes, Yun Shen, Yang Zhang</i>	
Calibrating Multimodal Learning.....	23429
<i>Huan Ma, Qingyang Zhang, Changqing Zhang, Bingzhe Wu, Huazhu Fu, Joey Tianyi Zhou, Qinghua Hu</i>	
AutoCoreset: An Automatic Practical Coreset Construction Framework.....	23451
<i>Alaa Maalouf, Murad Tukan, Vladimir Braverman, Daniela Rus</i>	
Learning GFlowNets from Partial Episodes for Improved Convergence and Stability.....	23467
<i>Kanika Madan, Jarrid Rector-Brooks, Maksym Korablyov, Emmanuel Bengio, Moksh Jain, Andrei Cristian Nica, Tom Bosc, Yoshua Bengio, Nikolay Malkin</i>	
Applied Online Algorithms with Heterogeneous Predictors	23484
<i>Jessica Maghakian, Russell Lee, Mohammad Hajiesmaili, Jian Li, Ramesh Sitaraman, Zhenhua Liu</i>	
CSP: Self-Supervised Contrastive Spatial Pre-Training for Geospatial-Visual Representations	23498
<i>Gengchen Mai, Ni Lao, Yutong He, Jiaming Song, Stefano Ermon</i>	
Vertical Federated Graph Neural Network for Recommender System.....	23516
<i>Peihua Mai, Yan Pang</i>	
Can Neural Network Memorization Be Localized?	23536
<i>Pratyush Maini, Michael Curtis Mozer, Hanie Sedghi, Zachary Chase Lipton, J Zico Kolter, Chiyuan Zhang</i>	
Fundamental Tradeoffs in Learning with Prior Information.....	23558
<i>Anirudha Majumdar</i>	
Additive Causal Bandits with Unknown Graph	23574
<i>Alan Malek, Virginia Aglietti, Silvia Chiappa</i>	
Weighted Tallying Bandits: Overcoming Intractability Via Repeated Exposure Optimality.....	23590
<i>Dhruv Malik, Conor Igoe, Yuanzhi Li, Aarti Singh</i>	
A Kernel-Based View of Language Model Fine-Tuning	23610
<i>Sadhika Malladi, Alexander Wettig, Dingli Yu, Danqi Chen, Sanjeev Arora</i>	
Performative Reinforcement Learning	23642
<i>Debmalya Mandal, Stelios Triantafyllou, Goran Radanovic</i>	
Differential Privacy Has Bounded Impact on Fairness in Classification.....	23681
<i>Paul Mangold, Michaël Perrot, Aurélien Bellet, Marc Tommasi</i>	
Random Classification Noise Does Not Defeat All Convex Potential Boosters Irrespective of Model Choice	23706
<i>Yishay Mansour, Richard Nock, Robert Williamson</i>	
H-Consistency Bounds for Pairwise Misranking Loss Surrogates	23743
<i>Anqi Mao, Mehryar Mohri, Yutao Zhong</i>	
Cross-Entropy Loss Functions: Theoretical Analysis and Applications	23803
<i>Anqi Mao, Mehryar Mohri, Yutao Zhong</i>	

Supported Trust Region Optimization for Offline Reinforcement Learning	23829
<i>Yixiu Mao, Hongchang Zhang, Chen Chen, Yi Xu, Xiangyang Ji</i>	
Robust Perception Through Equivariance	23852
<i>Chengzhi Mao, Lingyu Zhang, Abhishek Vaibhav Joshi, Junfeng Yang, Hao Wang, Carl Vondrick</i>	
Reliable Measures of Spread in High Dimensional Latent Spaces.....	23871
<i>Anna Marbut, Katy McKinney-Bock, Travis Wheeler</i>	
SRATTA: Sample Re-ATtribution Attack of Secure Aggregation in Federated Learning.....	23886
<i>Tanguy Marchand, Regis Loeb, Ulysse Marteau-Ferey, Jean Ogier Du Terrail, Arthur Pignet</i>	
Neuro-Symbolic Continual Learning: Knowledge, Reasoning Shortcuts and Concept Rehearsal.....	23915
<i>Emanuele Marconato, Gianpaolo Bontempo, Elisa Ficarra, Simone Calderara, Andrea Passerini, Stefano Teso</i>	
Evaluating Unsupervised Denoising Requires Unsupervised Metrics	23937
<i>Adria Marcos Morales, Matan Leibovich, Sreyas Mohan, Joshua Lawrence Vincent, Piyush Haluai, Mai Tan, Peter Crozier, Carlos Fernandez-Granda</i>	
Regions of Reliability in the Evaluation of Multivariate Probabilistic Forecasts.....	23958
<i>Étienne Marcotte, Valentina Zantedeschi, Alexandre Drouin, Nicolas Chapados</i>	
Analyzing Diffusion as Serial Reproduction	24005
<i>Raja Marjieh, Iliia Sucholutsky, Thomas A Langlois, Nori Jacoby, Thomas L. Griffiths</i>	
Quantized Distributed Training of Large Models with Convergence Guarantees	24020
<i>Iliia Markov, Adrian Vladu, Qi Guo, Dan Alistarh</i>	
Efficient Transformed Gaussian Processes for Non-Stationary Dependent Multi-Class Classification	24045
<i>Juan Maroñas, Daniel Hernández-Lobato</i>	
Computational Asymmetries in Robust Classification	24082
<i>Samuele Marro, Michele Lombardi</i>	

PART 31

Neural Network Approximations of PDEs Beyond Linearity: A Representational Perspective	24139
<i>Tanya Marwah, Zachary Chase Lipton, Jianfeng Lu, Andrej Risteski</i>	
Generative Pretraining for Black-Box Optimization	24173
<i>Satvik Mehul Mashkaria, Siddarth Krishnamoorthy, Aditya Grover</i>	
Improved Policy Evaluation for Randomized Trials of Algorithmic Resource Allocation.....	24198
<i>Aditya Mate, Bryan Wilder, Aparna Taneja, Milind Tambe</i>	
Multi-Fidelity Covariance Estimation in the Log-Euclidean Geometry.....	24214
<i>Aimee Maurais, Terrence Alsup, Benjamin Peherstorfer, Youssef Marzouk</i>	
Communication-Constrained Bandits Under Additive Gaussian Noise	24236
<i>Prathamesh Mayekar, Jonathan Scarlett, Vincent Y. F. Tan</i>	
Nonparametric Density Estimation Under Distribution Drift.....	24251
<i>Alessio Mazzetto, Eli Upfal</i>	
PAC-Bayesian Generalization Bounds for Adversarial Generative Models	24271
<i>Sokhna Diarra Mbacke, Florence Clerc, Pascal Germain</i>	
Robustness in Multimodal Learning Under Train-Test Modality Mismatch	24291
<i>Brandon McKinzie, Vaishaal Shankar, Joseph Yitan Cheng, Yinfei Yang, Jonathon Shlens, Alexander T Toshev</i>	
A Model-Free Closeness-Of-Influence Test for Features in Supervised Learning	24304
<i>Mohammad Mehrabi, Ryan A. Rossi</i>	
Stochastic Gradient Succeeds for Bandits	24325
<i>Jincheng Mei, Zixin Zhong, Bo Dai, Alekh Agarwal, Csaba Szepesvari, Dale Schuurmans</i>	
Normalizing Flows for Interventional Density Estimation.....	24361
<i>Valentyn Melnychuk, Dennis Frauen, Stefan Feuerriegel</i>	

Reprogramming Pretrained Language Models for Antibody Sequence Infilling	24398
<i>Igor Melnyk, Vijil Chenthamarakshan, Pin-Yu Chen, Payel Das, Amit Dhurandhar, Inkit Padhi, Devleena Das</i>	
Superhuman Fairness.....	24420
<i>Omid Memarrast, Linh Vu, Brian D Ziebart</i>	
A Model-Based Method for Minimizing CVaR and Beyond.....	24436
<i>Si Yi Meng, Robert M. Gower</i>	
Tuning Language Models as Training Data Generators for Augmentation-Enhanced Few-Shot Learning	24457
<i>Yu Meng, Martin Michalski, Jiaxin Huang, Yu Zhang, Tarek Abdelzaher, Jiawei Han</i>	
On Preemption and Learning in Stochastic Scheduling	24478
<i>Nadav Merlis, Hugo Richard, Flore Sentenac, Corentin Odic, Mathieu Molina, Vianney Perchet</i>	
Quantile Credit Assignment	24517
<i>Thomas Mesnard, Wenqi Chen, Alaa Saade, Yunhao Tang, Mark Rowland, Theophane Weber, Clare Lyle, Audrunas Gruslys, Michal Valko, Will Dabney, Georg Ostrovski, Eric Moulines, Remi Munos</i>	
Is Consensus Acceleration Possible in Decentralized Optimization Over Slowly Time-Varying Networks?.....	24532
<i>Dmitry Metev, Alexander Rogozin, Dmitry Kovalev, Alexander Gasnikov</i>	
Towards Theoretical Understanding of Inverse Reinforcement Learning	24555
<i>Alberto Maria Metelli, Filippo Lazzati, Marcello Restelli</i>	
Quantum Policy Gradient Algorithm with Optimized Action Decoding.....	24592
<i>Nico Meyer, Daniel Scherer, Axel Plinge, Christopher Mutschler, Michael Hartmann</i>	
Training Deep Surrogate Models with Large Scale Online Learning.....	24614
<i>Lucas Thibaut Meyer, Marc Schouler, Robert Alexander Caulk, Alejandro Ribes, Bruno Raffin</i>	
MANSA: Learning Fast and Slow in Multi-Agent Systems	24631
<i>David Henry Mguni, Haojun Chen, Taher Jafferjee, Jianhong Wang, Longfei Yue, Xidong Feng, Stephen Marcus McAleer, Feifei Tong, Jun Wang, Yaodong Yang</i>	
Representation Learning with Multi-Step Inverse Kinematics: An Efficient and Optimal Approach to Rich-Observation RL	24659
<i>Zakaria Mhammedi, Dylan J Foster, Alexander Rakhlin</i>	
Single Point-Based Distributed Zeroth-Order Optimization with a Non-Convex Stochastic Objective Function	24701
<i>Elissa Mhanna, Mohamad Assaad</i>	
Learning Instance-Specific Augmentations by Capturing Local Invariances.....	24720
<i>Ning Miao, Tom Rainforth, Emile Mathieu, Yann Dubois, Yee Whye Teh, Adam Foster, Hyunjik Kim</i>	
Path Neural Networks: Expressive and Accurate Graph Neural Networks.....	24737
<i>Gaspard Michel, Giannis Nikolentzos, Johannes F. Lutzeyer, Michalis Vazirgiannis</i>	
Learning to Acquire Novel Cognitive Tasks with Evolution, Plasticity and Meta-Meta-Learning	24756
<i>Thomas Miconi</i>	
Generative Decoding of Visual Stimuli.....	24775
<i>Eleni Miliotou, Panagiotis Kyriakis, Jason D Hinman, Andrei Irimia, Paul Bogdan</i>	
Cooperative Multi-Agent Reinforcement Learning: Asynchronous Communication and Linear Function Approximation.....	24785
<i>Yifei Min, Jiafan He, Tianhao Wang, Quanquan Gu</i>	
Directed Chain Generative Adversarial Networks.....	24812
<i>Ming Min, Ruimeng Hu, Tomoyuki Ichiba</i>	
An Information-Theoretic Analysis of Nonstationary Bandit Learning	24831
<i>Seungki Min, Daniel Russo</i>	

On the Convergence of Gradient Flow on Multi-Layer Linear Models	24850
<i>Hancheng Min, Rene Vidal, Enrique Mallada</i>	
Optimal Sets and Solution Paths of ReLU Networks	24888
<i>Aaron Mishkin, Mert Pilanci</i>	

PART 32

The Numerical Stability of Hyperbolic Representation Learning	24925
<i>Gal Mishne, Zhengchao Wan, Yusu Wang, Sheng Yang</i>	
DetectGPT: Zero-Shot Machine-Generated Text Detection Using Probability Curvature	24950
<i>Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D Manning, Chelsea Finn</i>	
Diffusion Based Representation Learning.....	24963
<i>Sarthak Mittal, Korbinian Abstreiter, Stefan Bauer, Bernhard Schölkopf, Arash Mehrjou</i>	
Disentangled Multiplex Graph Representation Learning	24983
<i>Yujie Mo, Yajie Lei, Jialie Shen, Xiaoshuang Shi, Heng Tao Shen, Xiaofeng Zhu</i>	
A Unified Audio-Visual Learning Framework for Localization, Separation, and Recognition.....	25006
<i>Shentong Mo, Pedro Morgado</i>	
Pruning Via Sparsity-Indexed ODE: A Continuous Sparsity Viewpoint	25018
<i>Zhanfeng Mo, Haosen Shi, Sinno Jialin Pan</i>	
Text-To-Concept (and Back) Via Cross-Model Alignment	25037
<i>Mazda Moayeri, Keivan Rezaei, Maziar Sanjabi, Soheil Feizi</i>	
A Fast, Well-Founded Approximation to the Empirical Neural Tangent Kernel	25061
<i>Mohamad Amin Mohamadi, Wonho Bae, Danica J. Sutherland</i>	
Special Properties of Gradient Descent with Large Learning Rates.....	25082
<i>Amirkeivan Mohtashami, Martin Jaggi, Sebastian U Stich</i>	
Neural Inverse Operators for Solving PDE Inverse Problems.....	25105
<i>Roberto Molinaro, Yunan Yang, Björn Engquist, Siddhartha Mishra</i>	
Input Uncertainty Propagation Through Trained Neural Networks.....	25140
<i>Paul Monchot, Loic Coquelin, Sébastien Julien Petit, Sébastien Marmin, Erwan Le Pennec, Nicolas Fischer</i>	
Compressing Tabular Data Via Latent Variable Estimation.....	25174
<i>Andrea Montanari, Eric Weiner</i>	
An SDE for Modeling SAM: Theory and Insights	25209
<i>Enea Monzio Compagnoni, Luca Biggio, Antonio Orvieto, Frank Norbert Proske, Hans Kersting, Aurelien Lucchi</i>	
Learning Deductive Reasoning from Synthetic Corpus Based on Formal Logic	25254
<i>Terufumi Morishita, Gaku Morio, Atsuki Yamaguchi, Yasuhiro Sogawa</i>	
WL Meet VC	25275
<i>Christopher Morris, Floris Geerts, Jan Tönshoff, Martin Grohe</i>	
ReLOAD: Reinforcement Learning with Optimistic Ascent-Descent for Last-Iterate Convergence in Constrained MDPs	25303
<i>Ted Moskowitz, Brendan O'Donoghue, Vivek Veeriah, Sebastian Flennerhag, Satinder Singh, Tom Zahavy</i>	
Optimistic Planning by Regularized Dynamic Programming	25337
<i>Antoine Moulin, Gergely Neu</i>	
Neural Signature Kernels as Infinite-Width-Depth-Limits of Controlled ResNets	25358
<i>Nicola Muca Cirone, Maud Lemercier, Cristopher Salvi</i>	
Improving Statistical Fidelity for Neural Image Compression with Implicit Local Likelihood Models.....	25426
<i>Matthew J. Muckley, Alaaeldin El-Nouby, Karen Ullrich, Herve Jegou, Jakob Verbeek</i>	
PFNs4BO: In-Context Learning for Bayesian Optimization.....	25444
<i>Samuel Müller, Matthias Feurer, Noah Hollmann, Frank Hutter</i>	

Achieving High Accuracy with PINNs Via Energy Natural Gradient Descent	25471
<i>Johannes Müller, Marius Zeinhofer</i>	
Uncertain Evidence in Probabilistic Models and Stochastic Simulators	25486
<i>Andreas Munk, Alexander Mead, Frank Wood</i>	
GibbsDDRM: A Partially Collapsed Gibbs Sampler for Solving Blind Inverse Problems with Denoising Diffusion Restoration	25501
<i>Naoki Murata, Koichi Saito, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Yuki Mitsufuji, Stefano Ermon</i>	
DIFF2: Differential Private Optimization Via Gradient Differences for Nonconvex Distributed Learning	25523
<i>Tomoya Murata, Taiji Suzuki</i>	
Efficiently Predicting High Resolution Mass Spectra with Graph Neural Networks	25549
<i>Michael Murphy, Stefanie Jegelka, Ernest Fraenkel, Tobias Kind, David Healey, Thomas Butler</i>	
Dynamical Linear Bandits	25563
<i>Marco Mussi, Alberto Maria Metelli, Marcello Restelli</i>	
Representation-Driven Reinforcement Learning	25588
<i>Ofir Nabati, Guy Tennenholtz, Shie Mannor</i>	
DADAO: Decoupled Accelerated Decentralized Asynchronous Optimization	25604
<i>Adel Nabli, Edouard Oyallon</i>	
Multi-User Reinforcement Learning with Low Rank Rewards	25627
<i>Dheeraj Mysore Nagaraj, Suhas S Kowshik, Naman Agarwal, Praneeth Netrapalli, Prateek Jain</i>	
Statistical Foundations of Prior-Data Fitted Networks	25660
<i>Thomas Nagler</i>	
Do Machine Learning Models Learn Statistical Rules Inferred from Data?	25677
<i>Aaditya Naik, Yinjun Wu, Mayur Naik, Eric Wong</i>	
Sample and Predict Your Latent: Modality-Free Sequential Disentanglement Via Contrastive Estimation	25694
<i>Ilan Naiman, Nimrod Berman, Omri Azencot</i>	
Effectively Using Public Data in Privacy Preserving Machine Learning	25718
<i>Milad Nasr, Saeed Mahloujifar, Xinyu Tang, Prateek Mittal, Amir Houmansadr</i>	

PART 33

Counterfactual Identifiability of Bijective Causal Models	25733
<i>Arash Nasr-Esfahany, Mohammad Alizadeh, Devavrat Shah</i>	
Discovering Object-Centric Generalized Value Functions from Pixels	25755
<i>Somjit Nath, Gopeshh Subbaraj, Khimya Khetarpal, Samira Ebrahimi Kahou</i>	
On Many-Actions Policy Gradient	25769
<i>Michal Nauman, Marek Cygan</i>	
Equivariant Architectures for Learning in Deep Weight Spaces	25790
<i>Aviv Navon, Aviv Shamsian, Idan Achituve, Ethan Fetaya, Gal Chechik, Haggai Maron</i>	
Scalable Multi-Agent Reinforcement Learning Through Intelligent Information Aggregation	25817
<i>Siddharth Nayak, Kenneth Choi, Wenqi Ding, Sydney Dolan, Karthik Gopalakrishnan, Hamsa Balakrishnan</i>	
Geometric Autoencoders - What You See is What You Decode	25834
<i>Philipp Nazari, Sebastian Damrich, Fred A Hamprecht</i>	
Action Matching: Learning Stochastic Dynamics from Samples	25858
<i>Kirill Neklyudov, Rob Brekelmans, Daniel Severo, Alireza Makhzani</i>	

Extending Conformal Prediction to Hidden Markov Models with Exact Validity Via De Finetti's Theorem for Markov Chains	25890
<i>Buddhika Nettasinghe, Samrat Chatterjee, Ramakrishna Tipireddy, Mahantesh M Halappanavar</i>	
ClimaX: A Foundation Model for Weather and Climate	25904
<i>Tung Nguyen, Johannes Brandstetter, Ashish Kapoor, Jayesh K Gupta, Aditya Grover</i>	
Provable Reset-Free Reinforcement Learning by No-Regret Reduction.....	25939
<i>Hoai-An Nguyen, Ching-An Cheng</i>	
Revisiting Over-Smoothing and Over-Squashing Using Ollivier-Ricci Curvature	25956
<i>Khang Nguyen, Nong Minh Hieu, Vinh Duc Nguyen, Nhat Ho, Stanley Osher, Tan Minh Nguyen</i>	
Deep Clustering with Incomplete Noisy Pairwise Annotations: A Geometric Regularization Approach	25980
<i>Tri Nguyen, Shahana Ibrahim, Xiao Fu</i>	
Self-Attention Amortized Distributional Projection Optimization for Sliced Wasserstein Point-Cloud Reconstruction	26008
<i>Khai Nguyen, Dang Nguyen, Nhat Ho</i>	
Building Neural Networks on Matrix Manifolds: A Gyrovector Space Approach	26031
<i>Xuan Son Nguyen, Shuo Yang</i>	
Simple Disentanglement of Style and Content in Visual Representations.....	26063
<i>Lilian Ngweta, Subha Maity, Alex Gittens, Yuekai Sun, Mikhail Yurochkin</i>	
MetaDiffuser: Diffusion Model as Conditional Planner for Offline Meta-RL	26087
<i>Fei Ni, Jianye Hao, Yao Mu, Yifu Yuan, Yan Zheng, Bin Wang, Zhixuan Liang</i>	
LEVER: Learning to Verify Language-To-Code Generation with Execution	26106
<i>Ansong Ni, Srini Iyer, Dragomir Radev, Veselin Stoyanov, Wen-Tau Yih, Sida Wang, Xi Victoria Lin</i>	
Continual Vision-Language Representation Learning with Off-Diagonal Information	26129
<i>Zixuan Ni, Longhui Wei, Siliang Tang, Yueting Zhuang, Qi Tian</i>	
Attributing Image Generative Models Using Latent Fingerprints	26150
<i>Guangyu Nie, Changhoon Kim, Yezhou Yang, Yi Ren</i>	
A Framework for Adapting Offline Algorithms to Solve Combinatorial Multi-Armed Bandit Problems with Bandit Feedback	26166
<i>Guanyu Nie, Yididiya Y. Nadew, Yanhui Zhu, Vaneet Aggarwal, Christopher John Quinn</i>	
SinFusion: Training Diffusion Models on a Single Image Or Video	26199
<i>Yaniv Nikankin, Niv Haim, Michal Irani</i>	
SparseProp: Efficient Sparse Backpropagation for Faster Training of Neural Networks at the Edge	26215
<i>Mahdi Nikdan, Tommaso Pegolotti, Eugenia Iofinova, Eldar Kurtic, Dan Alistarh</i>	
Anti-Exploration by Random Network Distillation.....	26228
<i>Alexander Nikulin, Vladislav Kurenkov, Denis Tarasov, Sergey Kolesnikov</i>	
Input Perturbation Reduces Exposure Bias in Diffusion Models	26245
<i>Mang Ning, Enver Sangineto, Angelo Porrello, Simone Calderara, Rita Cucchiara</i>	
Primal and Dual Analysis of Entropic Fictitious Play for Finite-Sum Problems.....	26266
<i>Atsushi Nitanda, Kazusato Oko, Denny Wu, Nobuhito Takenouchi, Taiji Suzuki</i>	
The Statistical Scope of Multicalibration	26283
<i>Georgy Noarov, Aaron Roth</i>	
Do Embodied Agents Dream of Pixelated Sheep: Embodied Decision Making Using Language Guided World Modelling	26311
<i>Kolby Nottingham, Prithviraj Ammanabrolu, Alane Suhr, Yejin Choi, Hannaneh Hajishirzi, Sameer Singh, Roy Fox</i>	
Gradient-Free Structured Pruning with Unlabeled Data.....	26326
<i>Azade Nova, Hanjun Dai, Dale Schuurmans</i>	

CHiLS: Zero-Shot Image Classification with Hierarchical Label Sets	26342
<i>Zachary Novack, Julian McAuley, Zachary Chase Lipton, Saurabh Garg</i>	
Few-Bit Backward: Quantized Gradients of Activation Functions for Memory Footprint Reduction	26363
<i>Georgii Sergeevich Novikov, Daniel Bershatsky, Julia Gusak, Alex Shonenkov, Denis Valerievich Dimitrov, Ivan Oseledets</i>	
Efficient Exploration Via Epistemic-Risk-Seeking Policy Optimization	26382
<i>Brendan O'Donoghue</i>	
Provable Benefit of Mixup for Finding Optimal Decision Boundaries	26403
<i>Junsoo Oh, Chulhee Yun</i>	
Shedding a PAC-Bayesian Light on Adaptive Sliced-Wasserstein Distances	26451
<i>Ruben Ohana, Kimia Nadjahi, Alain Rakotomamonjy, Liva Ralaivola</i>	
Reasons for the Superiority of Stochastic Estimators Over Deterministic Ones: Robustness, Consistency and Perceptual Quality	26474
<i>Guy Ohayon, Theo Joseph Adrai, Michael Elad, Tomer Michaeli</i>	
On the Within-Group Fairness of Screening Classifiers.....	26495
<i>Nastaran Okati, Stratis Tsirtsis, Manuel Gomez Rodriguez</i>	
Diffusion Models Are Minimax Optimal Distribution Estimators	26517
<i>Kazusato Oko, Shunta Akiyama, Taiji Suzuki</i>	

PART 34

How Many Perturbations Break This Model? Evaluating Robustness Beyond Adversarial Accuracy	26583
<i>Raphael Olivier, Bhiksha Raj</i>	
B-Learner: Quasi-Oracle Bounds on Heterogeneous Causal Effects Under Hidden Confounding	26599
<i>Miruna Oprescu, Jacob Dorn, Marah Ghoummaid, Andrew Jesson, Nathan Kallus, Uri Shalit</i>	
Measuring the Impact of Programming Language Distribution	26619
<i>Gabriel Orlanski, Kefan Xiao, Xavier Garcia, Jeffrey Hui, Joshua Howland, Jonathan Malmaud, Jacob Austin, Rishabh Singh, Michele Catasta</i>	
When Does Privileged Information Explain Away Label Noise?.....	26646
<i>Guillermo Ortiz-Jimenez, Mark Collier, Anant Nawalgaria, Alexander Nicholas D'Amour, Jesse Berent, Rodolphe Jenatton, Efi Kokiopoulou</i>	
Resurrecting Recurrent Neural Networks for Long Sequences	26670
<i>Antonio Orvieto, Samuel L Smith, Albert Gu, Anushan Fernando, Caglar Gulcehre, Razvan Pascanu, Soham De</i>	
Improving Adversarial Robustness Through the Contrastive-Guided Diffusion Process.....	26699
<i>Yidong Ouyang, Liyan Xie, Guang Cheng</i>	
On the Role of Attention in Prompt-Tuning	26724
<i>Samet Oymak, Ankit Singh Rawat, Mahdi Soltanolkotabi, Christos Thrampoulidis</i>	
Revisiting the Linear-Programming Framework for Offline RL with General Function Approximation.....	26769
<i>Asuman E. Ozdaglar, Sarath Pattathil, Jiawei Zhang, Kaiqing Zhang</i>	
Extrapolative Controlled Sequence Generation Via Iterative Refinement.....	26792
<i>Vishakh Padmakumar, Richard Yuanzhe Pang, He He, Ankur P Parikh</i>	
Locally Regularized Neural Differential Equations: Some Black Boxes Were Meant to Remain Closed!	26809
<i>Avik Pal, Alan Edelman, Christopher Vincent Rackauckas</i>	
Controlled Differential Equations on Long Sequences Via Non-Standard Wavelets	26820
<i>Sourav Pal, Zhanpeng Zeng, Sathya N. Ravi, Vikas Singh</i>	
Do the Rewards Justify the Means? Measuring Trade-Offs Between Rewards and Ethical Behavior in the Machiavelli Benchmark.....	26837
<i>Alexander Pan, Jun Shern Chan, Andy Zou, Nathaniel Li, Steven Basart, Thomas Woodside, Hanlin Zhang, Scott Emmons, Dan Hendrycks</i>	

Beyond Homophily: Reconstructing Structure for Graph-Agnostic Clustering.....	26868
<i>Erlin Pan, Zhao Kang</i>	
Better Training of GFlowNets with Local Credit and Incomplete Trajectories.....	26878
<i>Ling Pan, Nikolay Malkin, Dinghuai Zhang, Yoshua Bengio</i>	
A Hybrid Quantum-Classical Approach Based on the Hadamard Transform for the Convolutional Layer.....	26891
<i>Hongyi Pan, Xin Zhu, Salih Furkan Atici, Ahmet Cetin</i>	
Semi Bandit Dynamics in Congestion Games: Convergence to Nash Equilibrium and No-Regret Guarantees.....	26904
<i>Ioannis Panageas, Stratis Skoulakis, Luca Viano, Xiao Wang, Volkan Cevher</i>	
Flash: Concept Drift Adaptation in Federated Learning.....	26931
<i>Kunjai Panchal, Sunav Choudhary, Subrata Mitra, Koyel Mukherjee, Somdeb Sarkhel, Saayan Mitra, Hui Guan</i>	
Learn to Accumulate Evidence from All Training Samples: Theory and Practice	26963
<i>Deep Shankar Pandey, Qi Yu</i>	
Secure Federated Correlation Test and Entropy Estimation	26990
<i>Qi Pang, Lun Wang, Shuai Wang, Wenting Zheng, Dawn Song</i>	
Task-Specific Skill Localization in Fine-Tuned Language Models.....	27011
<i>Abhishek Panigrahi, Nikunj Saunshi, Haoyu Zhao, Sanjeev Arora</i>	
Kernel Sufficient Dimension Reduction and Variable Selection for Compositional Data Via Amalgamation	27034
<i>Junyoung Park, Jeongyoun Ahn, Cheolwoo Park</i>	
Learning Affinity with Hyperbolic Representation for Spatial Propagation.....	27048
<i>Jin-Hwi Park, Jaesung Choe, Inhwan Bae, Hae-Gon Jeon</i>	
TRAK: Attributing Model Behavior at Scale	27074
<i>Sung Min Park, Kristian Georgiev, Andrew Ilyas, Guillaume Leclerc, Aleksander Madry</i>	
Test-Time Style Shifting: Handling Arbitrary Styles in Domain Generalization	27114
<i>Jungwuk Park, Dong-Jun Han, Soyeong Kim, Jaekyun Moon</i>	
Towards Understanding Ensemble Distillation in Federated Learning	27132
<i>Sejun Park, Kihun Hong, Ganguk Hwang</i>	
Learning Controllable Degradation for Real-World Super-Resolution Via Constrained Flows.....	27188
<i>Seobin Park, Dongjin Kim, Sungyong Baik, Tae Hyun Kim</i>	
Differentially Private Sharpness-Aware Training	27204
<i>Jinseong Park, Hoki Kim, Yujin Choi, Jaewook Lee</i>	
Controllability-Aware Unsupervised Skill Discovery	27225
<i>Seohong Park, Kimin Lee, Youngwoon Lee, Pieter Abbeel</i>	
Predictable MDP Abstraction for Unsupervised Model-Based RL	27246
<i>Seohong Park, Sergey Levine</i>	
Neural Stochastic Differential Games for Time-Series Analysis.....	27269
<i>Sungwoo Park, Byoungwoo Park, Moontae Lee, Changhee Lee</i>	
Accelerated Infeasibility Detection of Constrained Optimization and Fixed-Point Iterations	27294
<i>Jisun Park, Ernest K. Ryu</i>	

PART 35

Model-Based Reinforcement Learning with Scalable Composite Policy Gradient Estimators.....	27346
<i>Paavo Parmas, Takuma Seno, Yuma Aoki</i>	
PAC Generalization Via Invariant Representations.....	27378
<i>Advait U Parulekar, Karthikeyan Shanmugam, Sanjay Shakkottai</i>	
Stochastic Gradient Descent-Induced Drift of Representation in a Two-Layer Neural Network.....	27401
<i>Farhad Pashakhanloo, Alexei Koulakov</i>	

Reducing SO(3) Convolutions to SO(2) for Efficient Equivariant GNNs.....	27420
<i>Saro Passaro, C. Lawrence Zitnick</i>	
Federated Online and Bandit Convex Optimization.....	27439
<i>Kumar Kshitij Patel, Lingxiao Wang, Aadirupa Saha, Nathan Srebro</i>	
Brauer’s Group Equivariant Neural Networks	27461
<i>Edward Pearce-Crump</i>	
How Jellyfish Characterise Alternating Group Equivariant Neural Networks	27483
<i>Edward Pearce-Crump</i>	
Can Large Language Models Reason About Program Invariants?	27496
<i>Kexin Pei, David Bieber, Kensen Shi, Charles Sutton, Pengcheng Yin</i>	
Dynamics-Inspired Neuromorphic Visual Representation Learning	27521
<i>Zhengqi Pei, Shuhui Wang</i>	
Feature Directions Matter: Long-Tailed Learning Via Rotated Balanced Representation.....	27542
<i>Gao Peifeng, Qianqian Xu, Peisong Wen, Zhiyong Yang, Huiyang Shao, Qingming Huang</i>	
Fair Neighbor Embedding	27564
<i>Jaakko Peltonen, Wen Xu, Timo Nummenmaa, Jyrki Nummenmaa</i>	
The Ideal Continual Learner: An Agent that Never Forgets	27585
<i>Liangzu Peng, Paris Giampouras, Rene Vidal</i>	
MolDiff: Addressing the Atom-Bond Inconsistency Problem in 3D Molecule Diffusion Generation	27611
<i>Xingang Peng, Jiaqi Guan, Qiang Liu, Jianzhu Ma</i>	
Diagnosis, Feedback, Adaptation: A Human-In-The-Loop Framework for Test-Time Policy Adaptation	27630
<i>Andi Peng, Aviv Netanyahu, Mark K Ho, Tianmin Shu, Andreea Bobu, Julie Shah, Pulkit Agrawal</i>	
Learning Hidden Markov Models When the Locations of Missing Observations Are Unknown	27642
<i>Binyamin Perets, Mark Kozdoba, Shie Mannor</i>	
Estimating the Contamination Factor’s Distribution in Unsupervised Anomaly Detection	27668
<i>Lorenzo Perini, Paul-Christian Bürkner, Arto Klami</i>	
Are Gaussian Data All You Need? the Extents and Limits of Universality in High-Dimensional Generalized Linear Estimation	27680
<i>Luca Pesce, Florent Krzakala, Bruno Loureiro, Ludovic Stephan</i>	
Certifying Ensembles: A General Certification Theory with S-Lipschitzness.....	27709
<i>Aleksandar Petrov, Francisco Eiras, Amartya Sanyal, Philip Torr, Adel Bibi</i>	
The Power of Learned Locally Linear Models for Nonlinear Policy Optimization	27737
<i>Daniel Pfrommer, Max Simchowitz, Tyler Westenbroek, Nikolai Matni, Stephen Tu</i>	
A Scalable Frank-Wolfe-Based Algorithm for the Max-Cut SDP	27822
<i>Chi Bach Pham, Wynita Griggs, James Saunderson</i>	
Attention-Based Recurrence for Multi-Agent Reinforcement Learning Under Stochastic Partial Observability	27840
<i>Thomy Phan, Fabian Ritz, Philipp Altmann, Maximilian Zorn, Jonas Nüßlein, Michael Kölle, Thomas Gabor, Claudia Linnhoff-Popien</i>	
HyperTuning: Toward Adapting Large Language Models Without Back-Propagation.....	27854
<i>Jason Phang, Yi Mao, Pengcheng He, Weizhu Chen</i>	
Linear CNNs Discover the Statistical Structure of the Dataset Using Only the Most Dominant Frequencies.....	27876
<i>Hannah Pinson, Joeri Lenaerts, Vincent Ginis</i>	
Conformal Prediction for Federated Uncertainty Quantification Under Label Shift.....	27907
<i>Vincent Plassier, Mehdi Makni, Aleksandr Rubashevskii, Eric Moulines, Maxim Panov</i>	
Universal Physics-Informed Neural Networks: Symbolic Differential Operator Discovery with Sparse Data.....	27948
<i>Lena Podina, Brydon Eastman, Mohammad Kohandel</i>	

Sequential Kernelized Independence Testing	27957
<i>Aleksandr Podkopaev, Patrick Blöbaum, Shiva Kasiviswanathan, Aaditya Ramdas</i>	
Truncating Trajectories in Monte Carlo Reinforcement Learning.....	27994
<i>Riccardo Poiani, Alberto Maria Metelli, Marcello Restelli</i>	
Hyena Hierarchy: Towards Larger Convolutional Language Models	28043
<i>Michael Poli, Stefano Massaroli, Eric Nguyen, Daniel Y Fu, Tri Dao, Stephen Baccus, Yoshua Bengio, Stefano Ermon, Christopher Re</i>	
Spurious Valleys and Clustering Behavior of Neural Networks.....	28079
<i>Samuele Pollaci</i>	
Multisample Flow Matching: Straightening Flows with Minibatch Couplings.....	28100
<i>Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, Ricky T. Q. Chen</i>	
Minimax Estimation of Discontinuous Optimal Transport Maps: The Semi-Discrete Case	28128
<i>Aram-Alexandre Pooladian, Vincent Divol, Jonathan Niles-Weed</i>	

PART 36

Test-Time Adaptation with Slot-Centric Models	28151
<i>Mihir Prabhudesai, Anirudh Goyal, Sujoy Paul, Sjoerd Van Steenkiste, Mehdi S. M. Sajjadi, Gaurav Aggarwal, Thomas Kipf, Deepak Pathak, Katerina Fragkiadaki</i>	
JAWS-X: Addressing Efficiency Bottlenecks of Conformal Prediction Under Standard and Feedback Covariate Shift	28167
<i>Drew Prinster, Suchi Saria, Anqi Liu</i>	
Equivariant Polynomials for Graph Neural Networks.....	28191
<i>Omri Puny, Derek Lim, Bobak Kiani, Haggai Maron, Yaron Lipman</i>	
Contrast with Reconstruct: Contrastive 3D Representation Learning Guided by Generative Pretraining	28223
<i>Zekun Qi, Runpei Dong, Guofan Fan, Zheng Ge, Xiangyu Zhang, Kaisheng Ma, Li Yi</i>	
An Effective Meaningful Way to Evaluate Survival Models.....	28244
<i>Shi-Ang Qi, Neeraj Kumar, Mahtab Farrokh, Weijie Sun, Li-Hao Kuan, Rajesh Ranganath, Ricardo Henao, Russell Greiner</i>	
Coarse-To-Fine: A Hierarchical Diffusion Model for Molecule Generation in 3D	28277
<i>Bo Qiang, Yuxuan Song, Minkai Xu, Jingjing Gong, Bowen Gao, Hao Zhou, Wei-Ying Ma, Yanyan Lan</i>	
Collaborative Causal Inference with Fair Incentives.....	28300
<i>Rui Qiao, Xinyi Xu, Bryan Kian Hsiang Low</i>	
FREDIS: A Fusion Framework of Refinement and Disambiguation for Unreliable Partial Label Learning	28321
<i>Congyu Qiao, Ning Xu, Jiaqi Lv, Yi Ren, Xin Geng</i>	
Nugget: Neural Agglomerative Embeddings of Text.....	28337
<i>Guanghui Qin, Benjamin Van Durme</i>	
BiBench: Benchmarking and Analyzing Network Binarization.....	28351
<i>Haotong Qin, Mingyuan Zhang, Yifu Ding, Aoyu Li, Zhongang Cai, Ziwei Liu, Fisher Yu, Xianglong Liu</i>	
Not All Semantics Are Created Equal: Contrastive Self-Supervised Learning with Automatic Temperature Individualization.....	28389
<i>Zi-Hao Qiu, Quanqi Hu, Zhuoning Yuan, Denny Zhou, Lijun Zhang, Tianbao Yang</i>	
Shortest Edit Path Crossover: A Theory-Driven Solution to the Permutation Problem in Evolutionary Neural Architecture Search	28422
<i>Xin Qiu, Risto Miikkulainen</i>	
Simple and Fast Group Robustness by Automatic Feature Reweighting	28448
<i>Shikai Qiu, Andres Potapczynski, Pavel Izmailov, Andrew Gordon Wilson</i>	

DRCFS: Doubly Robust Causal Feature Selection.....	28468
<i>Francesco Quinzan, Ashkan Soleymani, Patrick Jaillet, Cristian R. Rojas, Stefan Bauer</i>	
Robust Speech Recognition Via Large-Scale Weak Supervision	28492
<i>Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, Ilya Sutskever</i>	
Shiftable Context: Addressing Training-Inference Context Mismatch in Simultaneous Speech Translation.....	28519
<i>Matthew Raffel, Drew Penney, Lizhong Chen</i>	
Sequential Multi-Dimensional Self-Supervised Learning for Clinical Time Series	28531
<i>Aniruddh Raghu, Payal Chandak, Ridwan Alam, John Guttag, Collin Stultz</i>	
Recovery Bounds on Class-Based Optimal Transport: A Sum-Of-Norms Regularization Framework.....	28549
<i>Arman Rahbar, Ashkan Panahi, Morteza Haghir Chehreghani, Devdatt Dubhashi, Hamid Krim</i>	
Algorithmic Stability of Heavy-Tailed SGD with General Loss Functions	28578
<i>Anant Raj, Lingjiong Zhu, Mert Gurbuzbalaban, Umut Simsekli</i>	
Mastering the Unsupervised Reinforcement Learning Benchmark from Pixels	28598
<i>Sai Rajeswar, Pietro Mazzaglia, Tim Verbelen, Alexandre Piché, Bart Dhoedt, Aaron Courville, Alexandre Lacoste</i>	
SpotEM: Efficient Video Search for Episodic Memory	28618
<i>Santhosh Kumar Ramakrishnan, Ziad Al-Halah, Kristen Grauman</i>	
How Much Does Initialization Affect Generalization?	28637
<i>Sameera Ramasinghe, Lachlan Ewen Macdonald, Moshir Farazi, Hemanth Saratchandran, Simon Lucey</i>	
Model Ratatouille: Recycling Diverse Models for Out-Of-Distribution Generalization.....	28656
<i>Alexandre Rame, Kartik Ahuja, Jianyu Zhang, Matthieu Cord, Leon Bottou, David Lopez-Paz</i>	
A Picture of the Space of Typical Learnable Tasks	28680
<i>Rahul Ramesh, Jialin Mao, Itay Griniasty, Rubing Yang, Han Kheng Teoh, Mark Transtrum, James Sethna, Pratik Chaudhari</i>	
Policy Regularization with Dataset Constraint for Offline Reinforcement Learning	28701
<i>Yuhang Ran, Yi-Chen Li, Fuxiang Zhang, Zongzhang Zhang, Yang Yu</i>	
SpENCNN: Orchestrating Encoding and Sparsity for Fast Homomorphically Encrypted Neural Network Inference.....	28718
<i>Ran Ran, Xinwei Luo, Wei Wang, Tao Liu, Gang Quan, Xiaolin Xu, Caiwen Ding, Wujie Wen</i>	
Feature Learning in Deep Classifiers Through Intermediate Neural Collapse	28729
<i>Akshay Rangamani, Marius Lindegaard, Tomer Galanti, Tomaso A Poggio</i>	
The Unintended Consequences of Discount Regularization: Improving Regularization in Certainty Equivalence Reinforcement Learning	28746
<i>Sarah Rathnam, Sonali Parbhoo, Weiwei Pan, Susan Murphy, Finale Doshi-Velez</i>	
Beam Tree Recursive Cells	28768
<i>Jishnu Ray Chowdhury, Cornelia Caragea</i>	
Monotonic Location Attention for Length Generalization	28792
<i>Jishnu Ray Chowdhury, Cornelia Caragea</i>	
Automated Search for Conjectures on Mathematical Constants Using Analysis of Integer Sequences	28809
<i>Ofir Razon, Yoav Harris, Shahar Gottlieb, Dan Carmon, Ofir David, Ido Kaminer</i>	
Neural Networks Trained with SGD Learn Distributions of Increasing Complexity	28843
<i>Maria Refinetti, Alessandro Ingrosso, Sebastian Goldt</i>	
Simplex Random Features.....	28864
<i>Isaac Reid, Krzysztof Marcin Choromanski, Valerii Likhoshesterov, Adrian Weller</i>	
Bayesian Neural Networks Avoid Encoding Complex and Perturbation-Sensitive Concepts	28889
<i>Qihan Ren, Huiqi Deng, Yunuo Chen, Siyu Lou, Quanshi Zhang</i>	
Escaping Saddle Points in Zeroth-Order Optimization: The Power of Two-Point Estimators	28914
<i>Zhaolin Ren, Yujie Tang, Na Li</i>	

PART 37

Dimension-Independent Certified Neural Network Watermarks Via Mollifier Smoothing.....	28976
<i>Jiaxiang Ren, Yang Zhou, Jiayin Jin, Lingjuan Lyu, Da Yan</i>	
Feature Programming for Multivariate Time Series Prediction.....	29009
<i>Alex Daniel Reneau, Jerry Yao-Chieh Hu, Ammar Gilani, Han Liu</i>	
Run-Off Election: Improved Provable Defense Against Data Poisoning Attacks	29030
<i>Keivan Rezaei, Kiarash Banihashem, Atoosa Chegini, Soheil Feizi</i>	
Learning Control-Oriented Dynamical Structure from Data	29051
<i>Spencer M. Richards, Jean-Jacques Slotine, Navid Azizan, Marco Pavone</i>	
The Edge of Orthogonality: A Simple View of What Makes BYOL Tick.....	29063
<i>Pierre Harvey Richemond, Allison Tam, Yunhao Tang, Florian Strub, Bilal Piot, Felix Hill</i>	
Multi-Agent Best Arm Identification with Private Communications	29082
<i>Alexandre Rio, Merwan Barlier, Igor Colin, Marta Soare</i>	
A Two-Stage Active Learning Algorithm for k-Nearest Neighbors.....	29103
<i>Nicholas Rittler, Kamalika Chaudhuri</i>	
Lowering the Pre-Training Tax for Gradient-Based Subset Training: A Lightweight Distributed Pre- Training Toolkit	29130
<i>Yeonju Ro, Zhangyang Wang, Vijay Chidambaram, Aditya Akella</i>	
The Role of Entropy and Reconstruction in Multi-View Self-Supervised Learning	29143
<i>Borja Rodri'guez Gálvez, Arno Blaas, Pau Rodriguez, Adam Golinski, Xavier Suau, Jason Ramapuram, Dan Busbridge, Luca Zappella</i>	
RLang: A Declarative Language for Describing Partial World Knowledge to Reinforcement Learning Agents	29161
<i>Rafael Rodriguez-Sanchez, Benjamin Adin Spiegel, Jennifer Wang, Roma Patel, Stefanie Tellex, George Konidaris</i>	
Improving Fair Training Under Correlation Shifts	29179
<i>Yuji Roh, Kangwook Lee, Steven Euijong Whang, Changho Suh</i>	
The Statistical Benefits of Quantile Temporal-Difference Learning for Value Estimation	29210
<i>Mark Rowland, Yunhao Tang, Clare Lyle, Remi Munos, Marc G Bellemare, Will Dabney</i>	
Robust Satisficing MDPs	29232
<i>Haolin Ruan, Siyu Zhou, Zhi Chen, Chin Pang Ho</i>	
Infinite Action Contextual Bandits with Reusable Data Exhaust	29259
<i>Mark Rucker, Yinglun Zhu, Paul Mineiro</i>	
Function-Space Regularization in Neural Networks: A Probabilistic Perspective	29275
<i>Tim G. J. Rudner, Sanyam Kapoor, Shikai Qiu, Andrew Gordon Wilson</i>	
A New PHO-Rmula for Improved Performance of Semi-Structured Networks.....	29291
<i>David Rügamer</i>	
Geometric Clifford Algebra Networks	29306
<i>David Ruhe, Jayesh K Gupta, Steven De Keninck, Max Welling, Johannes Brandstetter</i>	
Constrained Monotonic Neural Networks	29338
<i>Davor Runje, Sharath M Shankaranarayana</i>	
Differential Privacy, Linguistic Fairness, and Training Data Influence: Impossibility and Possibility Theorems for Multilingual Language Models	29354
<i>Phillip Rust, Anders Søgaard</i>	
Intrinsic Sliced Wasserstein Distances for Comparing Collections of Probability Distributions on Manifolds and Graphs	29388
<i>Raif M. Rustamov, Subhabrata Majumdar</i>	
SWARM Parallelism: Training Large Models Can Be Surprisingly Communication-Efficient.....	29416
<i>Max Ryabinin, Tim Dettmers, Michael Diskin, Alexander Borzunov</i>	

Hiera: A Hierarchical Vision Transformer Without the Bells-And-Whistles	29441
<i>Chaitanya Ryali, Yuan-Ting Hu, Daniel Bolya, Chen Wei, Haoqi Fan, Po-Yao Huang, Vaibhav Aggarwal, Arkabandhu Chowdhury, Omid Poursaeed, Judy Hoffman, Jitendra Malik, Yanghao Li, Christoph Feichtenhofer</i>	
End-To-End Learning for Stochastic Optimization: A Bayesian Perspective	29455
<i>Yves Rychener, Daniel Kuhn, Tobias Sutter</i>	
Sequential Monte Carlo Learning for Time Series Structure Discovery	29473
<i>Feras Saad, Brian Patton, Matthew Douglas Hoffman, Rif A. Saurous, Vikash Mansinghka</i>	
Active Ranking of Experts Based on Their Performances in Many Tasks	29490
<i>El Mehdi Saad, Nicolas Verzelen, Alexandra Carpentier</i>	
Sample Complexity Bounds for Learning High-Dimensional Simplices in Noisy Regimes.....	29514
<i>Seyed Amir Hossein Saberi, Amir Najafi, Abolfazl Motahari, Babak Khalaj</i>	
Global Selection of Contrastive Batches Via Optimization on Sample Permutations	29542
<i>Vin Sachidananda, Ziyi Yang, Chenguang Zhu</i>	
High-Probability Bounds for Stochastic Optimization and Variational Inequalities: The Case of Unbounded Variance	29563
<i>Abdurakhmon Sadiev, Marina Danilova, Eduard Gorbunov, Samuel Horváth, Gauthier Gidel, Pavel Dvurechensky, Alexander Gasnikov, Peter Richtárik</i>	
End-To-End Differentiable Clustering with Associative Memories	29649
<i>Bishwajit Saha, Dmitry Krotov, Mohammed J Zaki, Parikshit Ram</i>	
Learning to Suggest Breaks: Sustainable Optimization of Long-Term User Engagement	29671
<i>Eden Saig, Nir Rosenfeld</i>	
Multi-Class Graph Clustering Via Approximated Effective p-Resistance	29697
<i>Shota Saito, Mark Herbster</i>	
Off-Policy Evaluation for Large Action Spaces Via Conjunct Effect Modeling	29734
<i>Yuta Saito, Qingyang Ren, Thorsten Joachims</i>	

PART 38

Rethinking Warm-Starts with Predictions: Learning Predictions Close to Sets of Optimal Solutions for Faster L-/L _q -Convex Function Minimization	29760
<i>Shinsaku Sakaue, Taihei Oki</i>	
PAC-Bayesian Offline Contextual Bandits with Guarantees.....	29777
<i>Othmane Sakhi, Pierre Alquier, Nicolas Chopin</i>	
Provably and Practically Efficient Neural Contextual Bandits.....	29800
<i>Sudeep Salgia</i>	
Distributed Linear Bandits Under Communication Constraints	29845
<i>Sudeep Salgia, Qing Zhao</i>	
Optimizing Hyperparameters with Conformal Quantile Regression.....	29876
<i>David Salinas, Jacek Golebiowski, Aaron Klein, Matthias Seeger, Cedric Archambeau</i>	
Raising the Cost of Malicious AI-Powered Image Editing.....	29894
<i>Hadi Salman, Alaa Khaddaj, Guillaume Leclerc, Andrew Ilyas, Aleksander Madry</i>	
Fast, Differentiable and Sparse Top-K: A Convex Analysis Perspective.....	29919
<i>Michael Eli Sander, Joan Puigcerver, Josip Djolonga, Gabriel Peyré, Mathieu Blondel</i>	
TAN Without a Burn: Scaling Laws of DP-SGD	29937
<i>Tom Sander, Pierre Stock, Alexandre Sablayrolles</i>	
Discrete Continuous Optimization Framework for Simultaneous Clustering and Training in Mixture Models.....	29950
<i>Parth Vipul Sangani, Arjun Shashank Kashettiwar, Pritish Chakraborty, Bhuvan Reddy Gangula, Durga S, Ganesh Ramakrishnan, Rishabh K Iyer, Abir De</i>	

Whose Opinions Do Language Models Reflect?.....	29971
<i>Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, Tatsunori Hashimoto</i>	
Streaming Active Learning with Deep Neural Networks	30005
<i>Akanksha Saran, Safoora Yousefi, Akshay Krishnamurthy, John Langford, Jordan T. Ash</i>	
Random Teachers Are Good Teachers	30022
<i>Felix Sarnthein, Gregor Bachmann, Sotiris Anagnostidis, Thomas Hofmann</i>	
Posterior Sampling for Deep Reinforcement Learning	30042
<i>Remo Sasso, Michelangelo Conserva, Paulo Rauber</i>	
Graph Neural Networks Can Recover the Hidden Features Solely from the Graph Structure	30062
<i>Ryoma Sato</i>	
Existence and Estimation of Critical Batch Size for Training Generative Adversarial Networks with Two Time-Scale Update Rule.....	30080
<i>Naoki Sato, Hideaki Iiduka</i>	
StyleGAN-T: Unlocking the Power of GANs for Fast Large-Scale Text-To-Image Synthesis	30105
<i>Axel Sauer, Tero Karras, Samuli Laine, Andreas Geiger, Timo Aila</i>	
Facial Expression Recognition with Adaptive Frame Rate Based on Multiple Testing Correction.....	30119
<i>Andrey Savchenko</i>	
Off-Policy Average Reward Actor-Critic with Deterministic Policy Search	30130
<i>Naman Saxena, Subhojyoti Khastagir, Shishir Kolathaya, Shalabh Bhatnagar</i>	
Gibbsian Polar Slice Sampling	30204
<i>Philip Schär, Michael Habeck, Daniel Rudolf</i>	
Identifiability and Generalizability in Constrained Inverse Reinforcement Learning	30224
<i>Andreas Schlaginhaufen, Maryam Kamgarpour</i>	
Learning Expressive Priors for Generalization and Uncertainty Estimation in Neural Networks.....	30252
<i>Dominik Schnaus, Jongseok Lee, Daniel Cremers, Rudolph Triebel</i>	
Deterministic Equivalent and Error Universality of Deep Random Features Learning	30285
<i>Dominik Schröder, Hugo Cui, Daniil Dmitriev, Bruno Loureiro</i>	
The Acquisition of Physical Knowledge in Generative Neural Networks.....	30321
<i>Luca M. Schulze Buschoff, Eric Schulz, Marcel Binz</i>	
Modality-Agnostic Variational Compression of Implicit Neural Representations	30342
<i>Jonathan Richard Schwarz, Jihoon Tack, Yee Whye Teh, Jaeho Lee, Jinwoo Shin</i>	
Bigger, Better, Faster: Human-Level Atari with Human-Level Efficiency	30365
<i>Max Schwarzer, Johan Samir Obando Ceron, Aaron Courville, Marc G Bellemare, Rishabh Agarwal, Pablo Samuel Castro</i>	
Dissecting the Effects of SGD Noise in Distinct Regimes of Deep Learning	30381
<i>Antonio Sclocchi, Mario Geiger, Matthieu Wyart</i>	
A Fast Optimistic Method for Monotone Variational Inequalities.....	30406
<i>Michael Sedlmayer, Dang-Khoa Nguyen, Radu Ioan Bot</i>	
Double-Weighting for Covariate Shift Adaptation	30439
<i>José I. Segovia-Martín, Santiago Mazuelas, Anqi Liu</i>	
Enhancing Activity Prediction Models in Drug Discovery with the Ability to Understand Human Language	30458
<i>Philipp Seidl, Andreu Vall, Sepp Hochreiter, Günter Klambauer</i>	
Variational Autoencoding Neural Operators.....	30491
<i>Jacob H Seidman, Georgios Kissas, George J. Pappas, Paris Perdikaris</i>	
Neural Markov Jump Processes	30523
<i>Patrick Seifner, Ramses J Sanchez</i>	

PART 39

Bayesian Online Change Point Detection with Hilbert Space Approximate Student-T Process	30553
<i>Jeremy Sellier, Petros Dellaportas</i>	
Incentivizing Exploration with Linear Contexts and Combinatorial Actions	30570
<i>Mark Sellke</i>	
Explainability as Statistical Inference	30584
<i>Hugo Henri Joseph Senetaire, Damien Garreau, Jes Frellsen, Pierre-Alexandre Mattei</i>	
Multi-View Masked World Models for Visual Robotic Manipulation.....	30613
<i>Younggyo Seo, Junsu Kim, Stephen James, Kimin Lee, Jinwoo Shin, Pieter Abbeel</i>	
One-Shot Compression of Large Edge-Exchangeable Graphs Using Bits-Back Coding	30633
<i>Daniel Severo, James Townsend, Ashish J Khisti, Alireza Makhzani</i>	
ModelDiff: A Framework for Comparing Learning Algorithms	30646
<i>Harshay Shah, Sung Min Park, Andrew Ilyas, Aleksander Madry</i>	
Auxiliary Learning as an Asymmetric Bargaining Game.....	30689
<i>Aviv Shamsian, Aviv Navon, Neta Glazer, Kenji Kawaguchi, Gal Chechik, Ethan Fetaya</i>	
Synthetic Prompting: Generating Chain-Of-Thought Demonstrations for Large Language Models	30706
<i>Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, Weizhu Chen</i>	
Complementary Attention for Multi-Agent Reinforcement Learning	30776
<i>Jianzhun Shao, Hongchang Zhang, Yun Qu, Chang Liu, Shuncheng He, Yuhang Jiang, Xiangyang Ji</i>	
Regularization-Free Diffeomorphic Temporal Alignment Nets.....	30794
<i>Ron Shapira Weber, Oren Freifeld</i>	
Toward Efficient Gradient-Based Value Estimation.....	30827
<i>Arsalan Sharifnassab, Richard S. Sutton</i>	
Coin Sampling: Gradient-Based Bayesian Inference Without Learning Rates	30850
<i>Louis Sharrock, Christopher Nemeth</i>	
On Kinetic Optimal Probability Paths for Generative Models	30883
<i>Neta Shaul, Ricky T. Q. Chen, Maximilian Nickel, Matthew Le, Yaron Lipman</i>	
Sequential Changepoint Detection Via Backward Confidence Sequences.....	30908
<i>Shubhanshu Shekhar, Aaditya Ramdas</i>	
Cold Analysis of Rao-Blackwellized Straight-Through Gumbel-Softmax Gradient Estimator	30931
<i>Alexander Shekhovtsov</i>	
Towards Understanding and Improving GFlowNet Training	30956
<i>Max W Shen, Emmanuel Bengio, Ehsan Hajiramezanali, Andreas Loukas, Kyunghyun Cho, Tommaso Biancalani</i>	
On Balancing Bias and Variance in Unsupervised Multi-Source-Free Domain Adaptation.....	30976
<i>Maohao Shen, Yuheng Bu, Gregory W. Wornell</i>	
On Penalty-Based Bilevel Gradient Descent Method.....	30992
<i>Han Shen, Tianyi Chen</i>	
Non-Autoregressive Conditional Diffusion Models for Time Series Prediction	31016
<i>Lifeng Shen, James Kwok</i>	
Cross-Modal Fine-Tuning: Align Then Refine.....	31030
<i>Junhong Shen, Liam Li, Lucio M. Dery, Corey Staten, Mikhail Khodak, Graham Neubig, Ameet Talwalkar</i>	
Auxiliary Modality Learning with Generalized Curriculum Distillation	31057
<i>Yu Shen, Xijun Wang, Peng Gao, Ming Lin</i>	
TGRL: An Algorithm for Teacher Guided Reinforcement Learning.....	31077
<i>Idan Shenfeld, Zhang-Wei Hong, Aviv Tamar, Pulkit Agrawal</i>	

FlexGen: High-Throughput Generative Inference of Large Language Models with a Single GPU	31094
<i>Ying Sheng, Lianmin Zheng, Binhang Yuan, Zhuohan Li, Max Ryabinin, Beidi Chen, Percy Liang, Christopher Re, Ion Stoica, Ce Zhang</i>	
Improved Regret for Efficient Online Reinforcement Learning with Linear Function Approximation.....	31117
<i>Uri Sherman, Tomer Koren, Yishay Mansour</i>	
Fundamental Limits of Two-Layer Autoencoders, and Achieving Them with Gradient Methods	31151
<i>Aleksandr Shevchenko, Kevin Kögler, Hamed Hassani, Marco Mondelli</i>	
Large Language Models Can Be Easily Distracted by Irrelevant Context	31210
<i>Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H. Chi, Nathanael Schärli, Denny Zhou</i>	
Everyone’s Preference Changes Differently: A Weighted Multi-Interest Model for Retrieval.....	31228
<i>Hui Shi, Yupeng Gu, Yitong Zhou, Bo Zhao, Sicun Gao, Jishen Zhao</i>	
A Near-Optimal Algorithm for Safe Reinforcement Learning Under Instantaneous Hard Constraints.....	31243
<i>Ming Shi, Yingbin Liang, Ness Shroff</i>	
Improving the Model Consistency of Decentralized Federated Learning	31269
<i>Yifan Shi, Li Shen, Kang Wei, Yan Sun, Bo Yuan, Xueqian Wang, Dacheng Tao</i>	
UPop: Unified and Progressive Pruning for Compressing Vision-Language Transformers.....	31292
<i>Dachuan Shi, Chaofan Tao, Ying Jin, Zhendong Yang, Chun Yuan, Jiaqi Wang</i>	
Sequence Modeling with Multiresolution Convolutional Memory	31312
<i>Jiaxin Shi, Ke Alexander Wang, Emily Fox</i>	
Statistical Inference on Multi-Armed Bandits with Delayed Feedback	31328
<i>Lei Shi, Jingshen Wang, Tianhao Wu</i>	

PART 40

Provably Efficient Offline Reinforcement Learning with Perturbed Data Sources.....	31353
<i>Chengshuai Shi, Wei Xiong, Cong Shen, Jing Yang</i>	
On the Complexity of Bayesian Generalization	31389
<i>Yu-Zhe Shi, Manjie Xu, John E. Hopcroft, Kun He, Joshua B. Tenenbaum, Song-Chun Zhu, Ying Nian Wu, Wenjuan Han, Yixin Zhu</i>	
Understanding and Generalizing Contrastive Learning from the Inverse Optimal Transport Perspective.....	31408
<i>Liangliang Shi, Gu Zhang, Haoyu Zhen, Jintao Fan, Junchi Yan</i>	
Long Horizon Temperature Scaling.....	31422
<i>Andy Shih, Dorsa Sadigh, Stefano Ermon</i>	
Gradient Descent in Neural Networks as Sequential Learning in Reproducing Kernel Banach Space.....	31435
<i>Alistair Shilton, Sunil Gupta, Santu Rana, Svetha Venkatesh</i>	
SNeRL: Semantic-Aware Neural Radiance Fields for Reinforcement Learning.....	31489
<i>Dongseok Shim, Seungjae Lee, H. Jin Kim</i>	
A Closer Look at the Intervention Procedure of Concept Bottleneck Models.....	31504
<i>Sungbin Shin, Yohan Jo, Sungsoo Ahn, Namhoon Lee</i>	
MetricGAN-OKD: Multi-Metric Optimization of MetricGAN Via Online Knowledge Distillation for Speech Enhancement	31521
<i>Wooseok Shin, Byung Hoon Lee, Jin Sob Kim, Hyun Joon Park, Sung Won Han</i>	
Improved Learning-Augmented Algorithms for the Multi-Option Ski Rental Problem Via Best-Possible Competitive Analysis	31539
<i>Yongho Shin, Changyeol Lee, Gukryeol Lee, Hyung-Chan An</i>	
One-Shot Imitation in a Non-Stationary Environment Via Multi-Modal Skill.....	31562
<i>Sangwoo Shin, Daehee Lee, Minjong Yoo, Woo Kyung Kim, Honguk Woo</i>	

Context Consistency Regularization for Label Sparsity in Time Series	31579
<i>Yooju Shin, Susik Yoon, Hwanjun Song, Dongmin Park, Byunghyun Kim, Jae-Gil Lee, Byung Suk Lee</i>	
Generative Causal Representation Learning for Out-Of-Distribution Motion Forecasting	31596
<i>Shayan Shirahmad Gale Bagi, Zahra Gharaee, Oliver Schulte, Mark Crowley</i>	
Exphormer: Sparse Transformers for Graphs	31613
<i>Hamed Shirzad, Ameya Velingker, Balaji Venkatachalam, Danica J. Sutherland, Ali Kemal Sinop</i>	
Synthetic Data for Model Selection.....	31633
<i>Alon Shoshan, Nadav Bhonker, Igor Kviatkovsky, Matan Fintz, Gerard Medioni</i>	
Probabilistic Attention-To-Influence Neural Models for Event Sequences.....	31657
<i>Xiao Shou, Debarun Bhattacharjya, Tian Gao, Dharmashankar Subramanian, Oktie Hassanzadeh, Kristin Bennett</i>	
Causal Bounds in Quasi-Markovian Graphs	31675
<i>Madhumitha Shridharan, Garud Iyengar</i>	
Repository-Level Prompt Generation for Large Language Models of Code.....	31693
<i>Disha Shrivastava, Hugo Larochelle, Daniel Tarlow</i>	
CLIPood: Generalizing CLIP to Out-Of-Distributions.....	31716
<i>Yang Shu, Xingzhuo Guo, Jialong Wu, Ximei Wang, Jianmin Wang, Mingsheng Long</i>	
Semi-Autoregressive Energy Flows: Exploring Likelihood-Free Training of Normalizing Flows	31732
<i>Phillip Si, Zeyi Chen, Subham Sekhar Sahoo, Yair Schiff, Volodymyr Kuleshov</i>	
Unearthing InSights into Mars: Unsupervised Source Separation with Limited Data	31754
<i>Ali Siahkoochi, Rudy Morel, Maarten V. De Hoop, Erwan Allys, Gregory Sainton, Taichi Kawamura</i>	
Quantitative Universal Approximation Bounds for Deep Belief Networks.....	31773
<i>Julian Sieber, Johann Gehring</i>	
Pricing Experimental Design: Causal Effect, Expected Revenue and Tail Risk.....	31788
<i>David Simchi-Levi, Chonghuan Wang</i>	
Statistical Learning Under Heterogeneous Distribution Shift	31800
<i>Max Simchowitz, Anurag Ajay, Pulkit Agrawal, Akshay Krishnamurthy</i>	
On the Stepwise Nature of Self-Supervised Learning	31852
<i>James B Simon, Maksis Knutins, Liu Ziyin, Daniel Geisz, Abraham J Fetterman, Joshua Albrecht</i>	
Hindsight Learning for MDPs with Exogenous Inputs	31877
<i>Sean R. Sinclair, Felipe Vieira Frujeri, Ching-An Cheng, Luke Marshall, Hugo De Oliveira Barbalho, Jingling Li, Jennifer Neville, Ishai Menache, Adith Swaminathan</i>	
Text-To-4D Dynamic Scene Generation.....	31915
<i>Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, Yaniv Taigman</i>	
The Hessian Perspective into the Nature of Convolutional Neural Networks.....	31930
<i>Sidak Pal Singh, Thomas Hofmann, Bernhard Schölkopf</i>	
When Do Minimax-Fair Learning and Empirical Risk Minimization Coincide?.....	31969
<i>Harvineet Singh, Matthäus Kleindessner, Volkan Cevher, Rumi Chunara, Chris Russell</i>	
Differentiable Simulations for Enhanced Sampling of Rare Events.....	31990
<i>Martin Sipka, Johannes C. B. Dietschreit, Lukáš Grajciar, Rafael Gomez-Bombarelli</i>	
Preprocessors Matter! Realistic Decision-Based Attacks on Machine Learning Systems.....	32008
<i>Chawin Sitawarin, Florian Tramèr, Nicholas Carlini</i>	
Invariance in Policy Optimisation and Partial Identifiability in Reward Learning.....	32033
<i>Joar Max Viktor Skalse, Matthew Farrugia-Roberts, Stuart Russell, Alessandro Abate, Adam Gleave</i>	

A Game-Theoretic Framework for Managing Risk in Multi-Agent Systems.....	32059
<i>Oliver Slumbers, David Henry Mguni, Stefano B Blumberg, Stephen Marcus McAleer, Yaodong Yang, Jun Wang</i>	
On the Effectiveness of Offline RL for Dialogue Response Generation	32088
<i>Paloma Sodhi, Felix Wu, Ethan R. Elenberg, Kilian Q Weinberger, Ryan McDonald</i>	
Fair Densities Via Boosting the Sufficient Statistics of Exponential Families	32105
<i>Alexander Soen, Hisham Husain, Richard Nock</i>	
The Dormant Neuron Phenomenon in Deep Reinforcement Learning.....	32145
<i>Ghada Sokar, Rishabh Agarwal, Pablo Samuel Castro, Utku Evci</i>	

PART 41

Abstracting Imperfect Information Away from Two-Player Zero-Sum Games.....	32169
<i>Samuel Sokota, Ryan D’Orazio, Chun Kai Ling, David J Wu, J Zico Kolter, Noam Brown</i>	
Meta-SAGE: Scale Meta-Learning Scheduled Adaptation with Guided Exploration for Mitigating Scale Shift on Combinatorial Optimization.....	32194
<i>Jiwoo Son, Minsu Kim, Hyeonah Kim, Jinkyoo Park</i>	
Consistency Models	32211
<i>Yang Song, Prafulla Dhariwal, Mark Chen, Ilya Sutskever</i>	
LipsNet: A Smooth and Robust Neural Network with Adaptive Lipschitz Constant for High Accuracy Optimal Control	32253
<i>Xujie Song, Jingliang Duan, Wenxuan Wang, Shengbo Eben Li, Chen Chen, Bo Cheng, Bo Zhang, Junqing Wei, Xiaoming Simon Wang</i>	
Deep Perturbation Learning: Enhancing the Network Performance Via Image Perturbations	32273
<i>Zifan Song, Xiao Gong, Guosheng Hu, Cairong Zhao</i>	
Latent Traversals in Generative Models as Potential Flows.....	32288
<i>Yue Song, T. Anderson Keller, Nicu Sebe, Max Welling</i>	
FedAvg Converges to Zero Training Loss Linearly for Overparameterized Multi-Layer Neural Networks	32304
<i>Bingqing Song, Prashant Khanduri, Xinwei Zhang, Jinfeng Yi, Mingyi Hong</i>	
RGE: A Repulsive Graph Rectification for Node Classification Via Influence.....	32331
<i>Jaeyun Song, Sungyub Kim, Eunho Yang</i>	
Importance Weighted Expectation-Maximization for Protein Sequence Design.....	32349
<i>Zhenqiao Song, Lei Li</i>	
Sketching for First Order Method: Efficient Algorithm for Low-Bandwidth Channel and Vulnerability.....	32365
<i>Zhao Song, Yitan Wang, Zheng Yu, Lichen Zhang</i>	
Sketching Meets Differential Privacy: Fast Algorithm for Dynamic Kronecker Projection Maintenance	32418
<i>Zhao Song, Xin Yang, Yuanyuan Yang, Lichen Zhang</i>	
A Nearly-Optimal Bound for Fast Regression with ℓ_∞ Guarantee.....	32463
<i>Zhao Song, Mingquan Ye, Junze Yin, Lichen Zhang</i>	
Loss-Guided Diffusion Models for Plug-And-Play Controllable Generation	32483
<i>Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin Chen, Arash Vahdat</i>	
Differentiable Tree Operations Promote Compositional Generalization	32499
<i>Paul Soulos, Edward J Hu, Kate McCurdy, Yunmo Chen, Roland Fernandez, Paul Smolensky, Jianfeng Gao</i>	
Are Labels Informative in Semi-Supervised Learning? Estimating and Leveraging the Missing-Data Mechanism.....	32521
<i>Aude Sportisse, Hugo Schmutz, Olivier Humbert, Charles Bouveyron, Pierre-Alexandre Mattei</i>	

Linear Causal Disentanglement Via Interventions	32540
<i>Chandler Squires, Anna Seigal, Salil S Bhate, Caroline Uhler</i>	
Generating Language Corrections for Teaching Physical Control Tasks.....	32561
<i>Megha Srivastava, Noah Goodman, Dorsa Sadigh</i>	
FaDIn: Fast Discretized Inference for Hawkes Processes with General Parametric Kernels	32575
<i>Guillaume Staerman, Cédric Allain, Alexandre Gramfort, Thomas Moreau</i>	
Partial Optimality in Cubic Correlation Clustering	32598
<i>David Stein, Silvia Di Gregorio, Bjoern Andres</i>	
MODEL: Memory Optimizations for Deep Learning	32618
<i>Benoit Steiner, Mostafa Elhoushi, Jacob Kahn, James Hegarty</i>	
Improving Expert Predictions with Conformal Prediction	32633
<i>Eleni Straitouri, Lequn Wang, Nastaran Okati, Manuel Gomez Rodriguez</i>	
Lookahead When it Matters: Adaptive Non-Causal Transformers for Streaming Neural Transducers	32654
<i>Grant Strimel, Yi Xie, Brian John King, Martin Radfar, Ariya Rastrow, Athanasios Mouchtaris</i>	
Kernel QuantTree	32677
<i>Diego Stucchi, Paolo Rizzo, Nicolò Folloni, Giacomo Boracchi</i>	
Topologically Faithful Image Segmentation Via Induced Matching of Persistence Barcodes	32698
<i>Nico Stucki, Johannes C. Paetzold, Suprosanna Shit, Bjoern Menze, Ulrich Bauer</i>	
Towards Robust Graph Incremental Learning on Evolving Graphs	32728
<i>Junwei Su, Dijan Zou, Zijun Zhang, Chuan Wu</i>	
DUET: 2D Structured and Approximately Equivariant Representations.....	32749
<i>Xavier Suau, Federico Danieli, T. Anderson Keller, Arno Blaas, Chen Huang, Jason Ramapuram, Dan Busbridge, Luca Zappella</i>	
Long-Tailed Recognition by Mutual Information Maximization Between Latent Features and Ground-Truth Labels	32770
<i>Min-Kook Suh, Seung-Woo Seo</i>	
Adversarial Learning of Distributional Reinforcement Learning.....	32783
<i>Yang Sui, Yukun Huang, Hongtu Zhu, Fan Zhou</i>	
Distilling Internet-Scale Vision-Language Models into Embodied Agents.....	32797
<i>Theodore Sumers, Kenneth Marino, Arun Ahuja, Rob Fergus, Ishita Dasgupta</i>	
Vector-Valued Control Variates	32819
<i>Zhuo Sun, Alessandro Barp, Francois-Xavier Briol</i>	
MetaModulation: Learning Variational Feature Hierarchies for Few-Shot Learning with Fewer Tasks.....	32847
<i>Wenfang Sun, Yingjun Du, Xiantong Zhen, Fan Wang, Ling Wang, Cees G. M. Snoek</i>	
Revisiting Sampling for Combinatorial Optimization	32859
<i>Haoran Sun, Katayoon Goshvadi, Azade Nova, Dale Schuurmans, Hanjun Dai</i>	
What Makes Entities Similar? a Similarity Flooding Perspective for Multi-Sourced Knowledge Graph Embeddings	32875
<i>Zequn Sun, Jiacheng Huang, Xiaozhou Xu, Qijin Chen, Weijun Ren, Wei Hu</i>	
Maximum Optimality Margin: A Unified Approach for Contextual Linear Programming and Inverse Linear Programming	32886
<i>Chunlin Sun, Shang Liu, Xiaocheng Li</i>	
Tensor Gaussian Process with Contraction for Multi-Channel Imaging Analysis.....	32913
<i>Hu Sun, Ward Manchester, Meng Jin, Yang Liu, Yang Chen</i>	
MABe22: A Multi-Species Multi-Task Benchmark for Learned Representations of Behavior.....	32936
<i>Jennifer J. Sun, Markus Marks, Andrew Wesley Ulmer, Dipam Chakraborty, Brian Geuther, Edward Hayes, Heng Jia, Vivek Kumar, Sebastian Oleszko, Zachary Partridge, Milan Peelman, Alice Robie, Catherine E Schretter, Keith Sheppard, Chao Sun, Param Uttarwar, Julian Morgan Wagner, Erik Werner, Joseph Parker, Pietro Perona, Yisong Yue, Kristin Branson, Ann Kennedy</i>	

PART 42

Dynamic Regularized Sharpness Aware Minimization in Federated Learning: Approaching Global Consistency and Smooth Landscape	32991
<i>Yan Sun, Li Shen, Shixiang Chen, Liang Ding, Dacheng Tao</i>	
When and How Does Known Class Help Discover Unknown Ones? Provable Understanding Through Spectral Analysis	33014
<i>Yiyou Sun, Zhenmei Shi, Yingyu Liang, Yixuan Li</i>	
Learning Prescriptive ReLU Networks	33044
<i>Wei Sun, Asterios Tsiourvas</i>	
All in a Row: Compressed Convolution Networks for Graphs	33061
<i>Junshu Sun, Shuhui Wang, Xinzhe Han, Zhe Xue, Qingming Huang</i>	
Momentum Ensures Convergence of SIGNSGD Under Weaker Assumptions	33077
<i>Tao Sun, Qingsong Wang, Dongsheng Li, Bao Wang</i>	
A Critical Revisit of Adversarial Robustness in 3D Point Cloud Recognition with Diffusion-Driven Purification	33100
<i>Jiachen Sun, Jiong Xiao Wang, Weili Nie, Zhiding Yu, Zhuoqing Mao, Chaowei Xiao</i>	
SDDM: Score-Decomposed Diffusion Models on Manifolds for Unpaired Image-To-Image Translation	33115
<i>Shikun Sun, Longhui Wei, Junliang Xing, Jia Jia, Qi Tian</i>	
A Neural PDE Solver with Temporal Stencil Modeling	33135
<i>Zhiqing Sun, Yiming Yang, Shinjae Yoo</i>	
Feature Expansion for Graph Neural Networks	33156
<i>Jiaqi Sun, Lin Zhang, Guangyi Chen, Peng Xu, Kun Zhang, Yujiu Yang</i>	
Model-Bellman Inconsistency for Model-Based Offline Reinforcement Learning	33177
<i>Yihao Sun, Jiaji Zhang, Chengxing Jia, Haoxin Lin, Junyin Ye, Yang Yu</i>	
Inflow, Outflow, and Reciprocity in Machine Learning	33195
<i>Mukund Sundararajan, Walid Krichene</i>	
When Personalization Harms Performance: Reconsidering the Use of Group Attributes in Prediction	33209
<i>Vinith Menon Suriyakumar, Marzyeh Ghassemi, Berk Ustun</i>	
Tuning Computer Vision Models with Task Rewards	33229
<i>André Susano Pinto, Alexander Kolesnikov, Yuge Shi, Lucas Beyer, Xiaohua Zhai</i>	
Beyond Exponentially Fast Mixing in Average-Reward Reinforcement Learning Via Multi-Level Monte Carlo Actor-Critic	33240
<i>Wesley A Suttle, Amrit Bedi, Bhrij Patel, Brian M. Sadler, Alec Koppel, Dinesh Manocha</i>	
Tight and Fast Generalization Error Bound of Graph Embedding in Metric Space	33268
<i>Atsushi Suzuki, Atsushi Nitanda, Taiji Suzuki, Jing Wang, Feng Tian, Kenji Yamanishi</i>	
Proximal Causal Learning of Conditional Average Treatment Effects	33285
<i>Erik Sverdrup, Yifan Cui</i>	
Inverse Reinforcement Learning Without Reinforcement Learning	33299
<i>Gokul Swamy, David Wu, Sanjiban Choudhury, Drew Bagnell, Steven Wu</i>	
Von Mises Mixture Distributions for Molecular Conformation Generation	33319
<i>Kirk Swanson, Jake Lawrence Williams, Eric M Jonas</i>	
Optimal Randomized Multilevel Monte Carlo for Repeatedly Nested Expectations	33343
<i>Yasa Syed, Guanyang Wang</i>	
Adaptive Coordination in Social Embodied Rearrangement	33365
<i>Andrew Szot, Unnat Jain, Dhruv Batra, Zsolt Kira, Ruta Desai, Akshara Rai</i>	

MG-GNN: Multigrid Graph Neural Networks for Learning Multilevel Domain Decomposition Methods	33381
<i>Ali Taghibakhshi, Nicolas Nytko, Tareq Uz Zaman, Scott Maclachlan, Luke Olson, Matthew West</i>	
Learning Mixtures of Gaussians with Censored Data	33396
<i>Wai Ming Tai, Bryon Aragam</i>	
Approximation and Estimation Ability of Transformers for Sequence-To-Sequence Functions with Infinite Dimensional Input	33416
<i>Shokichi Takakura, Taiji Suzuki</i>	
Learning Neural PDE Solvers with Parameter-Guided Channel Attention	33448
<i>Makoto Takamoto, Francesco Alesiani, Mathias Niepert</i>	
Contextual Conservative Interleaving Bandits	33468
<i>Kei Takemura</i>	
Randomized Gaussian Process Upper Confidence Bound with Tighter Bayesian Regret Bounds.....	33490
<i>Shion Takeno, Yu Inatsu, Masayuki Karasuyama</i>	
Towards Practical Preferential Bayesian Optimization with Skew Gaussian Processes.....	33516
<i>Shion Takeno, Masahiro Nomura, Masayuki Karasuyama</i>	
Robust Explanation for Free Or at the Cost of Faithfulness.....	33534
<i>Zeren Tan, Yang Tian</i>	
Provably Invariant Learning Without Domain Information	33563
<i>Xiaoyu Tan, Lin Yong, Shengyu Zhu, Chao Qu, Xihe Qiu, Xu Yinghui, Peng Cui, Yuan Qi</i>	
Auto-Differentiation of Relational Computations for Very Large Scale Machine Learning	33581
<i>Yuxin Tang, Zhimin Ding, Dimitrije Jankov, Binhang Yuan, Daniel Bourgeois, Chris Jermaine</i>	
Regret-Minimizing Double Oracle for Extensive-Form Games.....	33599
<i>Xiaohang Tang, Le Cong Dinh, Stephen Marcus McAleer, Yaodong Yang</i>	
From Perception to Programs: Regularize, Overparameterize, and Amortize	33616
<i>Hao Tang, Kevin Ellis</i>	
Understanding Self-Predictive Learning for Reinforcement Learning	33632
<i>Yunhao Tang, Zhaohan Daniel Guo, Pierre Harvey Richemond, Bernardo Avila Pires, Yash Chandak, Remi Munos, Mark Rowland, Mohammad Gheshlaghi Azar, Charline Le Lan, Clare Lyle, András György, Shantanu Thakoor, Will Dabney, Bilal Piot, Daniele Calandriello, Michal Valko</i>	
DoMo-AC: Doubly Multi-Step Off-Policy Actor-Critic Algorithm	33657
<i>Yunhao Tang, Tadashi Kozuno, Mark Rowland, Anna Harutyunyan, Remi Munos, Bernardo Avila Pires, Michal Valko</i>	
Towards Understanding Generalization of Graph Neural Networks	33674
<i>Huayi Tang, Yong Liu</i>	
Towards a Better Understanding of Representation Dynamics Under TD-Learning	33720
<i>Yunhao Tang, Remi Munos</i>	
VA-Learning as a More Efficient Alternative to Q-Learning	33739
<i>Yunhao Tang, Remi Munos, Mark Rowland, Michal Valko</i>	
Defects of Convolutional Decoder Networks in Frequency Representation	33758
<i>Ling Tang, Wen Shen, Zhanpeng Zhou, Yuefeng Chen, Quanshi Zhang</i>	

PART 43

Difference-In-Differences Meets Tree-Based Methods: Heterogeneous Treatment Effects Estimation with Unmeasured Confounding.....	33792
<i>Caizhi Tang, Huiyuan Wang, Xinyu Li, Qing Cui, Longfei Li, Jun Zhou</i>	
End-To-End Training of Deep Boltzmann Machines by Unbiased Contrastive Divergence with Local Mode Initialization	33804
<i>Shohei Taniguchi, Masahiro Suzuki, Yusuke Iwasawa, Yutaka Matsuo</i>	

POUF: Prompt-Oriented Unsupervised Fine-Tuning for Large Pre-Trained Models.....	33816
<i>Korawat Tanwisuth, Shujian Zhang, Huangjie Zheng, Pengcheng He, Mingyuan Zhou</i>	
Dual Focal Loss for Calibration	33833
<i>Linwei Tao, Minjing Dong, Chang Xu</i>	
Abstract-To-Executable Trajectory Translation for One-Shot Task Generalization	33850
<i>Stone Tao, Xiaochen Li, Tongzhou Mu, Zhiao Huang, Yuzhe Qin, Hao Su</i>	
Data Feedback Loops: Model-Driven Amplification of Dataset Biases.....	33883
<i>Rohan Taori, Tatsunori Hashimoto</i>	
Deep Regression Unlearning.....	33921
<i>Ayush Kumar Tarun, Vikram Singh Chundawat, Murari Mandal, Mohan Kankanhalli</i>	
How to Trust Your Diffusion Model: A Convex Optimization Approach to Conformal Risk Control	33940
<i>Jacopo Teneggi, Matthew Tivnan, Web Stayman, Jeremias Sulam</i>	
Concurrent Shuffle Differential Privacy Under Continual Observation.....	33961
<i>Jay Tenenbaum, Haim Kaplan, Yishay Mansour, Uri Stemmer</i>	
Finding Generalization Measures by Contrasting Signal and Noise	33983
<i>Jiaye Teng, Bohang Zhang, Ruichen Li, Haowei He, Yequan Wang, Yan Tian, Yang Yuan</i>	
Reinforcement Learning with History Dependent Dynamic Contexts	34011
<i>Guy Tennenholtz, Nadav Merlis, Lior Shani, Martin Mladenov, Craig Boutilier</i>	
PWSHAP: A Path-Wise Explanation Model for Targeted Variables	34054
<i>Lucile Ter-Minassian, Oscar Clivio, Karla Diazordaz, Robin J. Evans, Christopher C. Holmes</i>	
On the Estimation of Gaussian Mixture Copula Models	34090
<i>Ashutosh Tewari</i>	
Target-Aware Generative Augmentations for Single-Shot Adaptation	34105
<i>Kowshik Thopalli, Rakshith Subramanyam, Pavan K. Turaga, Jayaraman J. Thiagarajan</i>	
ELSA: Efficient Label Shift Adaptation Through the Lens of Semiparametric Models	34120
<i>Qinglong Tian, Xin Zhang, Jiwei Zhao</i>	
Spherical Inducing Features for Orthogonally-Decoupled Gaussian Processes	34143
<i>Louis C. Tiao, Vincent Dutordoir, Victor Picheny</i>	
Fast Rates for Maximum Entropy Exploration.....	34161
<i>Daniil Tiapkin, Denis Belomestny, Daniele Calandriello, Eric Moulines, Remi Munos, Alexey Naumov, Pierre Perrault, Yunhao Tang, Michal Valko, Pierre Menard</i>	
Margin-Based Sampling in High Dimensions: When Being Active is Less Efficient than Staying Passive.....	34222
<i>Alexandru Tifrea, Jacob Clarysse, Fanny Yang</i>	
Differentiable Multi-Target Causal Bayesian Experimental Design.....	34263
<i>Panagiotis Tigas, Yashas Annadani, Desi R. Ivanova, Andrew Jesson, Yarin Gal, Adam Foster, Stefan Bauer</i>	
PCA-Based Multi-Task Learning: A Random Matrix Approach	34280
<i>Malik Tiomoko, Romain Couillet, Frederic Pascal</i>	
Perturbation Analysis of Neural Collapse.....	34301
<i>Tom Tirer, Haoxiang Huang, Jonathan Niles-Weed</i>	
Overcoming Simplicity Bias in Deep Networks Using a Feature Sieve.....	34330
<i>Rishabh Tiwari, Pradeep Shenoy</i>	
Beyond In-Domain Scenarios: Robust Density-Aware Calibration.....	34344
<i>Christian Tomani, Futa Kai Waseda, Yuesong Shen, Daniel Cremers</i>	
Distribution Free Domain Generalization	34369
<i>Peifeng Tong, Wu Su, He Li, Jialin Ding, Zhan Haoxiang, Song Xi Chen</i>	
Extending Kernel PCA Through Dualization: Sparsity, Robustness and Fast Algorithms	34379
<i>Francesco Tonin, Alex Lambert, Panagiotis Patrinos, Johan Suykens</i>	
Robust Weak Supervision with Variational Auto-Encoders.....	34394
<i>Francesco Tonolini, Nikolaos Aletras, Yunlong Jiao, Gabriella Kazai</i>	

Fully Bayesian Autoencoders with Latent Sparse Gaussian Processes	34409
<i>Ba-Hien Tran, Babak Shahbaba, Stephan Mandt, Maurizio Filippone</i>	
Discrete Key-Value Bottleneck	34431
<i>Frederik Träuble, Anirudh Goyal, Nasim Rahaman, Michael Curtis Mozer, Kenji Kawaguchi, Yoshua Bengio, Bernhard Schölkopf</i>	
Mimetic Initialization of Self-Attention Layers	34456
<i>Asher Trockman, J Zico Kolter</i>	
Representer Point Selection for Explaining Regularized High-Dimensional Models	34469
<i>Che-Ping Tsai, Jiong Zhang, Hsiang-Fu Yu, Eli Chien, Cho-Jui Hsieh, Pradeep Kumar Ravikumar</i>	
Expected Gradients of Maxout Networks and Consequences to Parameter Initialization.....	34491
<i>Hanna Tseran, Guido Montufar</i>	
Provable Data Subset Selection for Efficient Neural Networks Training.....	34533
<i>Murad Tukan, Samson Zhou, Alaa Maalouf, Daniela Rus, Vladimir Braverman, Dan Feldman</i>	
Jump-Start Reinforcement Learning	34556
<i>Ikechukwu Uchendu, Ted Xiao, Yao Lu, Banghua Zhu, Mengyuan Yan, Joséphine Simon, Matthew Bennice, Chuyuan Fu, Cong Ma, Jiantao Jiao, Sergey Levine, Karol Hausman</i>	
Submodular Order Functions and Assortment Optimization.....	34584
<i>Rajan Udwani</i>	

PART 44

Computationally Efficient PAC RL in POMDPs with Latent Determinism and Conditional Embeddings	34615
<i>Masatoshi Uehara, Ayush Sekhari, Jason D. Lee, Nathan Kallus, Wen Sun</i>	
From Adaptive Query Release to Machine Unlearning	34642
<i>Enayat Ullah, Raman Arora</i>	
Private Federated Learning with Autotuned Compression	34668
<i>Enayat Ullah, Christopher A. Choquette-Choo, Peter Kairouz, Sewoong Oh</i>	
The Monge Gap: A Regularizer to Learn All Transport Maps.....	34709
<i>Théo Uscidda, Marco Cuturi</i>	
Semi-Dual Unbalanced Quadratic Optimal Transport: Fast Statistical Rates and Convergent Algorithm.	34734
<i>Adrien Vacher, François-Xavier Vialard</i>	
Random Grid Neural Processes for Parametric Partial Differential Equations	34759
<i>Arnaud Vadeboncoeur, Ieva Kazlauskaitė, Yanni Papandreou, Fehmi Cirak, Mark Girolami, Omer Deniz Akyildiz</i>	
Delayed Feedback in Kernel Bandits	34779
<i>Sattar Vakili, Danyal Ahmed, Alberto Bernacchia, Ciara Pike-Burke</i>	
Synthetic Data, Real Errors: How (Not) to Publish and Use Synthetic Data	34793
<i>Boris Van Breugel, Zhaozhi Qian, Mihaela Van Der Schaar</i>	
Trading-Off Payments and Accuracy in Online Classification with Paid Stochastic Experts	34809
<i>Dirk Van Der Hoeven, Ciara Pike-Burke, Hao Qiu, Nicolò Cesa-Bianchi</i>	
Causal Isotonic Calibration for Heterogeneous Treatment Effects.....	34831
<i>Lars Van Der Laan, Ernesto Ulloa-Perez, Marco Carone, Alex Luedtke</i>	
Accounting for Informative Sampling When Learning to Forecast Treatment Outcomes Over Time	34855
<i>Toon Vanderschueren, Alicia Curth, Wouter Verbeke, Mihaela Van Der Schaar</i>	
Best Arm Identification in Multi-Agent Multi-Armed Bandits	34875
<i>Filippo Vannella, Alexandre Proutiere, Jaeseong Jeong</i>	
Conditional Tree Matching for Inference-Time Adaptation of Tree Prediction Models.....	34908
<i>Harshit Varma, Abhijeet Awasthi, Sunita Sarawagi</i>	

Optimal LP Rounding and Linear-Time Approximation Algorithms for Clustering Edge-Colored Hypergraphs	34924
<i>Nate Veldt</i>	
Fast $(1+\epsilon)$ -Approximation Algorithms for Binary Matrix Factorization	34952
<i>Ameya Velingker, Maximilian Vötsch, David Woodruff, Samson Zhou</i>	
The Virtues of Laziness in Model-Based RL: A Unified Objective and Algorithms	34978
<i>Anirudh Vemula, Yuda Song, Aarti Singh, Drew Bagnell, Sanjiban Choudhury</i>	
Learning the Right Layers a Data-Driven Layer-Aggregation Strategy for Semi-Supervised Learning on Multilayer Graphs	35006
<i>Sara Venturini, Andrea Cristofari, Francesco Rinaldi, Francesco Tudisco</i>	
Multi-Environment Pretraining Enables Transfer to Action Limited Datasets	35024
<i>David Venuto, Sherry Yang, Pieter Abbeel, Doina Precup, Igor Mordatch, Ofir Nachum</i>	
AbODE: Ab Initio Antibody Design Using Conjoined ODEs	35037
<i>Yogesh Verma, Markus Heinonen, Vikas Garg</i>	
TabLeak: Tabular Data Leakage in Federated Learning	35051
<i>Mark Vero, Mislav Balunovic, Dimitar Iliev Dimitrov, Martin Vechev</i>	
Low-Variance Gradient Estimation in Unrolled Computation Graphs with ES-Single	35084
<i>Paul Vicol</i>	
Arithmetic Sampling: Parallel Diverse Decoding for Large Language Models	35120
<i>Luke Vilnis, Yury Zemlyanskiy, Patrick Murray, Alexandre Tachard Passos, Sumit Sanghai</i>	
Eventual Discounting Temporal Logic Counterfactual Experience Replay	35137
<i>Cameron Voloshin, Abhinav Verma, Yisong Yue</i>	
Transformers Learn In-Context by Gradient Descent	35151
<i>Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, Joao Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, Max Vladymyrov</i>	
Topological Singularity Detection at Multiple Scales	35175
<i>Julius Von Rohrscheidt, Bastian Rieck</i>	
Improving ℓ_1 -Certified Robustness Via Randomized Smoothing by Leveraging Box Constraints	35198
<i>Vaclav Voracek, Matthias Hein</i>	
Vector Quantized Wasserstein Auto-Encoder	35223
<i>Long Tung Vuong, Trung Le, He Zhao, Chuanxia Zheng, Mehrtash Harandi, Jianfei Cai, Dinh Phung</i>	
Competitive Gradient Optimization	35243
<i>Abhijeet Vyas, Brian Bullins, Kamyar Azizzadenesheli</i>	
On Provable Copyright Protection for Generative Models	35277
<i>Nikhil Vyas, Sham M. Kakade, Boaz Barak</i>	
Leveraging Offline Data in Online Reinforcement Learning	35300
<i>Andrew Wagenmaker, Aldo Pacchiano</i>	
Fast Private Kernel Density Estimation Via Locality Sensitive Quantization	35339
<i>Tal Wagner, Yonatan Naamad, Nina Mishra</i>	
Investigating the Role of Model-Based Learning in Exploration and Transfer	35368
<i>Jacob C Walker, Eszter Vértés, Yazhe Li, Gabriel Dulac-Arnold, Ankesh Anand, Theophane Weber, Jessica B Hamrick</i>	
UPSCALE: Unconstrained Channel Pruning	35384
<i>Alvin Wan, Hanxiang Hao, Kaushik Patnaik, Yueyang Xu, Omer Hadad, David Güera, Zhile Ren, Qi Shan</i>	

PART 45

Poisoning Language Models During Instruction Tuning	35413
<i>Alexander Wan, Eric Wallace, Sheng Shen, Dan Klein</i>	

SeMAIL: Eliminating Distractors in Visual Imitation Via Separated Models	35426
<i>Shenghua Wan, Yucen Wang, Minghao Shao, Ruying Chen, De-Chuan Zhan</i>	
Multiplier Bootstrap-Based Exploration.....	35444
<i>Runzhe Wan, Haoyu Wei, Branislav Kveton, Rui Song</i>	
Bandit Multi-Linear DR-Submodular Maximization and Its Applications on Adversarial Submodular Bandits	35491
<i>Zongqi Wan, Jialin Zhang, Wei Chen, Xiaoming Sun, Zhijie Zhang</i>	
Tight Regret Bounds for Single-Pass Streaming Multi-Armed Bandits	35525
<i>Chen Wang</i>	
Improved Active Multi-Task Representation Learning Via Lasso.....	35548
<i>Yiping Wang, Yifang Chen, Kevin Jamieson, Simon Shaolei Du</i>	
Tilted Sparse Additive Models	35579
<i>Yingjie Wang, Hong Chen, Weifeng Liu, Fengxiang He, Tieliang Gong, Youcheng Fu, Dacheng Tao</i>	
From Hypergraph Energy Functions to Hypergraph Neural Networks	35605
<i>Yuxin Wang, Quan Gan, Xipeng Qiu, Xuanjing Huang, David Wipf</i>	
A Closer Look at Self-Supervised Lightweight Vision Transformers	35624
<i>Shaoru Wang, Jin Gao, Zeming Li, Xiaoqin Zhang, Weiming Hu</i>	
PreNAS: Preferred One-Shot Learning Towards Efficient Neural Architecture Search.....	35642
<i>Haibin Wang, Ce Ge, Hesun Chen, Xiuyu Sun</i>	
Adversarial Policies Beat Superhuman Go AIs	35655
<i>Tony Tong Wang, Adam Gleave, Tom Tseng, Kellin Pelrine, Nora Belrose, Joseph Miller, Michael D Dennis, Yawen Duan, Viktor Pogrebniak, Sergey Levine, Stuart Russell</i>	
On Regularization and Inference with Label Constraints.....	35740
<i>Kaifu Wang, Hangfeng He, Tin D. Nguyen, Piyush Kumar, Dan Roth</i>	
Policy Gradient in Robust MDPs with Global Convergence Guarantee	35763
<i>Qiu hao Wang, Chin Pang Ho, Marek Petrik</i>	
Adaptive Smoothing Gradient Learning for Spiking Neural Networks	35798
<i>Ziming Wang, Runhao Jiang, Shuang Lian, Rui Yan, Huajin Tang</i>	
CircuitNet: A Generic Neural Network to Realize Universal Circuit Motif Modeling	35817
<i>Yansen Wang, Xinyang Jiang, Kan Ren, Caihua Shan, Xufang Luo, Dongqi Han, Kaitao Song, Yifei Shen, Dongsheng Li</i>	
Generalized Polyak Step Size for First Order Optimization with Momentum	35836
<i>Xiaoyu Wang, Mikael Johansson, Tong Zhang</i>	
Near-Minimax-Optimal Risk-Sensitive Reinforcement Learning with CVaR	35864
<i>Kaiwen Wang, Nathan Kallus, Wen Sun</i>	
FedHPO-Bench: A Benchmark Suite for Federated Hyperparameter Optimization.....	35908
<i>Zhen Wang, Weirui Kuang, Ce Zhang, Bolin Ding, Yaliang Li</i>	
A/B Testing in Network Data with Covariate-Adaptive Randomization.....	35949
<i>Jialu Wang, Ping Li, Feifang Hu</i>	
Learning Belief Representations for Partially Observable Deep RL.....	35970
<i>Andrew Wang, Andrew C Li, Toryn Q. Klassen, Rodrigo Toro Icarte, Sheila A. McIlraith</i>	
Warm-Start Actor-Critic: From Approximation Error to Sub-Optimality Gap.....	35989
<i>Hang Wang, Sen Lin, Junshan Zhang</i>	
Slot-VAE: Object-Centric Scene Generation with Slot Attention	36020
<i>Yanbo Wang, Letao Liu, Justin Dauwels</i>	
DIVISION: Memory Efficient Training Via Dual Activation Precision.....	36036
<i>Guanchu Wang, Zirui Liu, Zhimeng Jiang, Ninghao Liu, Na Zou, Xia Hu</i>	
CocktailSGD: Fine-Tuning Foundation Models Over 500Mbps Networks	36058
<i>Jue Wang, Yucheng Lu, Binhang Yuan, Beidi Chen, Percy Liang, Christopher De Sa, Christopher Re, Ce Zhang</i>	

Magneto: A Foundation Transformer.....	36077
<i>Hongyu Wang, Shuming Ma, Shaohan Huang, Li Dong, Wenhui Wang, Zhiliang Peng, Yu Wu, Payal Bajaj, Saksham Singhal, Alon Benhaim, Barun Patra, Zhun Liu, Vishrav Chaudhary, Xia Song, Furu Wei</i>	
Direct Parameterization of Lipschitz-Bounded Deep Networks.....	36093
<i>Ruigang Wang, Ian Manchester</i>	
Tighter Information-Theoretic Generalization Bounds from Supersamples.....	36111
<i>Ziqiao Wang, Yongyi Mao</i>	
NP-SemiSeg: When Neural Processes Meet Semi-Supervised Semantic Segmentation.....	36138
<i>Jianfeng Wang, Daniela Massiceti, Xiaolin Hu, Vladimir Pavlovic, Thomas Lukasiewicz</i>	
GC-Flow: A Graph-Based Flow Network for Effective Clustering.....	36157
<i>Tianchun Wang, Farzaneh Mirzazadeh, Xiang Zhang, Jie Chen</i>	
Curriculum Co-Disentangled Representation Learning Across Multiple Environments for Social Recommendation.....	36174
<i>Xin Wang, Zirui Pan, Yuwei Zhou, Hong Chen, Chendi Ge, Wenwu Zhu</i>	
Data Efficient Neural Scaling Law Via Model Reusing.....	36193
<i>Peihao Wang, Rameswar Panda, Zhangyang Wang</i>	
Deep Temporal Sets with Evidential Reinforced Attentions for Unique Behavioral Pattern Discovery	36205
<i>Dingrong Wang, Deep Shankar Pandey, Krishna Prasad Neupane, Zhiwei Yu, Ervine Zheng, Zhi Zheng, Qi Yu</i>	

PART 46

Active Learning Based Structural Inference.....	36224
<i>Aoran Wang, Jun Pang</i>	
Better Diffusion Models Further Improve Adversarial Training	36246
<i>Zekai Wang, Tianyu Pang, Chao Du, Min Lin, Weiwei Liu, Shuicheng Yan</i>	
Polarity is All You Need to Learn and Transfer Faster	36264
<i>Qingyang Wang, Michael Alan Powell, Eric W Bridgeford, Ali Geisa, Joshua T Vogelstein</i>	
Projected Tensor Power Method for Hypergraph Community Recovery	36285
<i>Jinxin Wang, Yuen-Man Pun, Xiaolu Wang, Peng Wang, Anthony Man-Cho So</i>	
Estimating Possible Causal Effects with Latent Variables Via Adjustment.....	36308
<i>Tian-Zuo Wang, Tian Qin, Zhi-Hua Zhou</i>	
InfoDiffusion: Representation Learning Using Information Maximizing Diffusion Models	36336
<i>Yingheng Wang, Yair Schiff, Aaron Gokaslan, Weishen Pan, Fei Wang, Christopher De Sa, Volodymyr Kuleshov</i>	
A Robust Test for the Stationarity Assumption in Sequential Decision Making	36355
<i>Jitao Wang, Chengchun Shi, Zhenke Wu</i>	
GEAR: A GPU-Centric Experience Replay System for Large Reinforcement Learning Models	36380
<i>Hanjing Wang, Man-Kit Sit, Congjie He, Ying Wen, Weinan Zhang, Jun Wang, Yaodong Yang, Luo Mai</i>	
Effective and Efficient Structural Inference with Reservoir Computing.....	36391
<i>Aoran Wang, Tsz Pan Tong, Jun Pang</i>	
Optimal Goal-Reaching Reinforcement Learning Via Quasimetric Learning.....	36411
<i>Tongzhou Wang, Antonio Torralba, Phillip Isola, Amy Zhang</i>	
Model-Free Robust Average-Reward Reinforcement Learning	36431
<i>Yue Wang, Alvaro Velasquez, George K. Atia, Ashley Prater-Bennette, Shaofeng Zou</i>	
Live in the Moment: Learning Dynamics Model Adapted to Evolving Policy	36470
<i>Xiyao Wang, Wichayaporn Wongkamjan, Ruonan Jia, Furong Huang</i>	
Learning to Bid in Repeated First-Price Auctions with Budgets	36494
<i>Qian Wang, Zongjun Yang, Xiaotie Deng, Yuqing Kong</i>	

Network Effects in Performative Prediction Games.....	36514
<i>Xiaolu Wang, Chung-Yiu Yau, Hoi To Wai</i>	
Robustly Learning a Single Neuron Via Sharpness.....	36541
<i>Puqian Wang, Nikos Zarifis, Ilias Diakonikolas, Jelena Diakonikolas</i>	
DualHSIC: HSIC-Bottleneck and Alignment for Continual Learning	36578
<i>Zifeng Wang, Zheng Zhan, Yifan Gong, Yucai Shao, Stratis Ioannidis, Yanzhi Wang, Jennifer Dy</i>	
Enforcing Hard Constraints with Soft Barriers: Safe Reinforcement Learning in Unknown Stochastic Environments	36593
<i>Yixuan Wang, Simon Sinong Zhan, Ruo Chen Jiao, Zhilu Wang, Wanxin Jin, Zhuoran Yang, Zhaoran Wang, Chao Huang, Qi Zhu</i>	
LinSATNet: The Positive Linear Satisfiability Neural Networks.....	36605
<i>Runzhong Wang, Yunhao Zhang, Ziao Guo, Tianyi Chen, Xiaokang Yang, Junchi Yan</i>	
Offline Meta Reinforcement Learning with In-Distribution Online Adaptation	36626
<i>Jianhao Wang, Jin Zhang, Haozhe Jiang, Junyu Zhang, Liwei Wang, Chongjie Zhang</i>	
Reachability-Aware Laplacian Representation in Reinforcement Learning.....	36670
<i>Kaixin Wang, Kuangqi Zhou, Jiashi Feng, Bryan Hooi, Xinchao Wang</i>	
PPG Reloaded: An Empirical Study on What Matters in Phasic Policy Gradient.....	36694
<i>Kaixin Wang, Daquan Zhou, Jiashi Feng, Shie Mannor</i>	
On Heterogeneous Treatment Effects in Heterogeneous Causal Graphs.....	36714
<i>Richard A Watson, Hengrui Cai, Xinming An, Samuel McLean, Rui Song</i>	
Nonparametric Extensions of Randomized Response for Private Confidence Sets	36748
<i>Ian Waudby-Smith, Steven Wu, Aaditya Ramdas</i>	
Global Optimality for Euclidean CCCP Under Riemannian Convexity.....	36790
<i>Melanie Weber, Suvrit Sra</i>	
A Universal Unbiased Method for Classification from Aggregate Observations	36804
<i>Zixi Wei, Lei Feng, Bo Han, Tongliang Liu, Gang Niu, Xiaofeng Zhu, Heng Tao Shen</i>	
NTK-Approximating MLP Fusion for Efficient Language Model Fine-Tuning	36821
<i>Tianxin Wei, Zeming Guo, Yifan Chen, Jingrui He</i>	
Boosting Graph Contrastive Learning Via Graph Contrastive Saliency.....	36839
<i>Chunyu Wei, Yu Wang, Bing Bai, Kai Ni, David Brady, Lu Fang</i>	
Set-Membership Belief State-Based Reinforcement Learning for POMDPs.....	36856
<i>Wei Wei, Lijun Zhang, Lin Li, Huizhong Song, Jiye Liang</i>	
Mitigating Memorization of Noisy Labels by Clipping the Model Prediction.....	36868
<i>Hongxin Wei, Huiping Zhuang, Renchunzi Xie, Lei Feng, Gang Niu, Bo An, Yixuan Li</i>	
Graphically Structured Diffusion Models.....	36887
<i>Christian Dietrich Weilbach, William Harvey, Frank Wood</i>	
Expectation-Complete Graph Representations with Homomorphisms	36910
<i>Pascal Welke, Maximilian Thiessen, Fabian Jögl, Thomas Gärtner</i>	
A Conditional Normalizing Flow for Accelerated Multi-Coil MR Imaging.....	36926
<i>Jeffrey Wen, Rizwan Ahmad, Philip Schniter</i>	
Optimizing Mode Connectivity for Class Incremental Learning.....	36940
<i>Haitao Wen, Haoyang Cheng, Heqian Qiu, Lanxiao Wang, Lili Pan, Hongliang Li</i>	
Towards Learning Geometric Eigen-Lengths Crucial for Fitting Tasks	36958
<i>Yijia Weng, Kaichun Mo, Ruoxi Shi, Yanchao Yang, Leonidas Guibas</i>	
Open-VCLIP: Transforming CLIP to an Open-Vocabulary Video Model Via Interpolated Weight Optimization.....	36978
<i>ZeJia Weng, Xitong Yang, Ang Li, Zuxuan Wu, Yu-Gang Jiang</i>	
Fully-Adaptive Composition in Differential Privacy	36990
<i>Justin Whitehouse, Aaditya Ramdas, Ryan Rogers, Steven Wu</i>	

Scalable Set Encoding with Universal Mini-Batch Consistency and Unbiased Full Set Gradient Approximation.....	37008
<i>Jeffrey Willette, Seanie Lee, Bruno Andreis, Kenji Kawaguchi, Juho Lee, Sung Ju Hwang</i>	

PART 47

Flexible Phase Dynamics for Bio-Plausible Contrastive Learning.....	37042
<i>Ezekiel Williams, Colin Bredenberg, Guillaume Lajoie</i>	
Approximate Stein Classes for Truncated Density Estimation.....	37066
<i>Daniel James Williams, Song Liu</i>	
Theoretical Behavior of XAI Methods in the Presence of Suppressor Variables	37091
<i>Rick Wilming, Leo Kieslich, Benedict Clark, Stefan Hauße</i>	
Marginalization is Not Marginal: No Bad VAE Local Minima When Learning Optimal Sparse Representations	37108
<i>David Wipf</i>	
Uncertainty Estimation for Molecules: Desiderata and Methods	37133
<i>Tom Wollschläger, Nicholas Gao, Bertrand Charpentier, Mohamed Amine Ketata, Stephan Günemann</i>	
The Blessing of Heterogeneity in Federated Q-Learning: Linear Speedup and Beyond.....	37157
<i>Jiin Woo, Gauri Joshi, Yuejie Chi</i>	
Learning Deep Time-Index Models for Time Series Forecasting.....	37217
<i>Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, Steven Hoi</i>	
Sharper Bounds for ℓ_p Sensitivity Sampling.....	37238
<i>David Woodruff, Taisuke Yasuda</i>	
Two Losses Are Better than One: Faster Optimization Using a Cheaper Proxy.....	37273
<i>Blake Woodworth, Konstantin Mishchenko, Francis Bach</i>	
SEGA: Structural Entropy Guided Anchor View for Graph Contrastive Learning	37293
<i>Junran Wu, Xueyuan Chen, Bowen Shi, Shangzhe Li, Ke Xu</i>	
Causal Proxy Models for Concept-Based Model Explanations.....	37313
<i>Zhengxuan Wu, Karel D'Oosterlinck, Atticus Geiger, Amir Zur, Christopher Potts</i>	
Effective Neural Topic Modeling with Embedding Clustering Regularization	37335
<i>Xiaobao Wu, Xinshuai Dong, Thong Thanh Nguyen, Anh Tuan Luu</i>	
Adaptive Compositional Continual Meta-Learning	37358
<i>Bin Wu, Jinyuan Fang, Xiangxiang Zeng, Shangsong Liang, Qiang Zhang</i>	
Anchor Sampling for Federated Learning with Partial Client Participation.....	37379
<i>Feijie Wu, Song Guo, Zhihao Qu, Shiqi He, Ziming Liu, Jing Gao</i>	
Solving High-Dimensional PDEs with Latent Spectral Models.....	37417
<i>Haixu Wu, Tengge Hu, Huakun Luo, Jianmin Wang, Mingsheng Long</i>	
A Law of Robustness Beyond Isoperimetry	37439
<i>Yihan Wu, Heng Huang, Hongyang Zhang</i>	
Uncovering Adversarial Risks of Test-Time Adaptation	37456
<i>Tong Wu, Feiran Jia, Xiangyu Qi, Jiachen T. Wang, Vikash Sehwal, Saeed Mahloujifar, Prateek Mittal</i>	
Stable Estimation of Heterogeneous Treatment Effects	37496
<i>Anpeng Wu, Kun Kuang, Ruoxuan Xiong, Bo Li, Fei Wu</i>	
Rethinking Explaining Graph Neural Networks Via Non-Parametric Subgraph Matching.....	37511
<i>Fang Wu, Siyuan Li, Xurui Jin, Yinghui Jiang, Dragomir Radev, Zhangming Niu, Stan Z. Li</i>	
Understanding Int4 Quantization for Language Models: Latency Speedup, Composability, and Failure Cases	37524
<i>Xiaoxia Wu, Cheng Li, Reza Yazdani Aminabadi, Zhewei Yao, Yuxiong He</i>	
Towards Understanding Generalization of Macro-AUC in Multi-Label Learning.....	37540
<i>Guoqiang Wu, Chongxuan Li, Yilong Yin</i>	

Quantifying the Knowledge in GNNs for Reliable Distillation into MLPs.....	37571
<i>Lirong Wu, Haitao Lin, Yufei Huang, Stan Z. Li</i>	
Delay-Agnostic Asynchronous Coordinate Update Algorithm.....	37582
<i>Xuyang Wu, Changxin Liu, Sindri Magnússon, Mikael Johansson</i>	
Masked Trajectory Models for Prediction, Representation, and Control	37607
<i>Philipp Wu, Arjun Majumdar, Kevin Stone, Yixin Lin, Igor Mordatch, Pieter Abbeel, Aravind Rajeswaran</i>	
Disentangled Multi-Fidelity Deep Bayesian Active Learning	37624
<i>Dongxia Wu, Ruijia Niu, Matteo Chinazzi, Yian Ma, Rose Yu</i>	
Tight Data Access Bounds for Private Top-K Selection.....	37635
<i>Hao Wu, Olga Ohrimenko, Anthony Wirth</i>	
The Implicit Regularization of Dynamical Stability in Stochastic Gradient Descent.....	37656
<i>Lei Wu, Weijie J Su</i>	
Distributional Offline Policy Evaluation with Predictive Error Guarantees.....	37685
<i>Runzhe Wu, Masatoshi Uehara, Wen Sun</i>	
π -Tuning: Transferring Multimodal Foundation Models with Optimal Multi-Task Interpolation.....	37713
<i>Chengyue Wu, Teng Wang, Yixiao Ge, Zeyu Lu, Ruisong Zhou, Ying Shan, Ping Luo</i>	
Learning Functional Distributions with Private Labels	37728
<i>Changlong Wu, Yifan Wang, Ananth Grama, Wojciech Szpankowski</i>	
QuantumDARTS: Differentiable Quantum Architecture Search for Variational Quantum Algorithms	37745
<i>Wenjie Wu, Ge Yan, Xudong Lu, Kaisen Pan, Junchi Yan</i>	
Discover and Cure: Concept-Aware Mitigation of Spurious Correlation	37765
<i>Shirley Wu, Mert Yuksekgonul, Linjun Zhang, James Zou</i>	
On the Training Instability of Shuffling SGD with Batch Normalization	37787
<i>David Xing Wu, Chulhee Yun, Suvrit Sra</i>	

PART 48

DugMatting: Decomposed-Uncertainty-Guided Matting	37846
<i>Jiawei Wu, Changqing Zhang, Zuoyong Li, Huazhu Fu, Xi Peng, Joey Tianyi Zhou</i>	
Personalized Federated Learning Under Mixture of Distributions.....	37860
<i>Yue Wu, Shuaicheng Zhang, Wenchao Yu, Yanchi Liu, Quanquan Gu, Dawei Zhou, Haifeng Chen, Wei Cheng</i>	
Differentially Private Episodic Reinforcement Learning with Heavy-Tailed Rewards.....	37880
<i>Yulian Wu, Xingyu Zhou, Sayak Ray Chowdhury, Di Wang</i>	
Finite-Sample Analysis of Learning High-Dimensional Single ReLU Neuron.....	37919
<i>Jingfeng Wu, Difan Zou, Zixiang Chen, Vladimir Braverman, Quanquan Gu, Sham M. Kakade</i>	
Understanding Backdoor Attacks Through the Adaptability Hypothesis	37952
<i>Xun Xian, Ganghua Wang, Jayanth Srinivasa, Ashish Kundu, Xuan Bi, Mingyi Hong, Jie Ding</i>	
Fair and Optimal Classification Via Post-Processing	37977
<i>Ruicheng Xian, Lang Yin, Han Zhao</i>	
UMD: Unsupervised Model Detection for X2X Backdoor Attacks	38013
<i>Zhen Xiang, Zidi Xiong, Bo Li</i>	
Random Shuffle Transformer for Image Restoration	38039
<i>Jie Xiao, Xueyang Fu, Man Zhou, Hongjian Liu, Zheng-Jun Zha</i>	
Communication-Efficient Federated Hypergradient Computation Via Aggregated Iterative Differentiation	38059
<i>Peiyao Xiao, Kaiyi Ji</i>	
SmoothQuant: Accurate and Efficient Post-Training Quantization for Large Language Models.....	38087
<i>Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, Song Han</i>	

On the Forward Invariance of Neural ODEs	38100
<i>Wei Xiao, Tsun-Hsuan Wang, Ramin Hasani, Mathias Lechner, Yutong Ban, Chuang Gan, Daniela Rus</i>	
COMCAT: Towards Efficient Compression and Customization of Attention-Based Vision Models	38125
<i>Jinqi Xiao, Miao Yin, Yu Gong, Xiao Zang, Jian Ren, Bo Yuan</i>	
Improving Bi-Level Optimization Based Methods with Inspiration from Humans' Classroom Study Techniques	38137
<i>Pengtao Xie</i>	
Future-Conditioned Unsupervised Pretraining for Decision Transformer	38187
<i>Zhihui Xie, Zichuan Lin, Deheng Ye, Qiang Fu, Yang Wei, Shuai Li</i>	
DeSRA: Detect and Delete the Artifacts of GAN-Based Real-World Super-Resolution Models.....	38204
<i>Liangbin Xie, Xintao Wang, Xiangyu Chen, Gen Li, Ying Shan, Jiantao Zhou, Chao Dong</i>	
Semiparametrically Efficient Off-Policy Evaluation in Linear Markov Decision Processes	38227
<i>Chuhan Xie, Wenhao Yang, Zhihua Zhang</i>	
A Critical View of Vision-Based Long-Term Dynamics Prediction Under Environment Misalignment.....	38258
<i>Hanchen Xie, Jiageng Zhu, Mahyar Khayatkhoei, Jiazhi Li, Mohamed E. Hussein, Wael Abdalmageed</i>	
Controlling Type Confounding in Ad Hoc Teamwork with Instance-Wise Teammate Feedback Rectification	38272
<i>Dong Xing, Pengjie Gu, Qian Zheng, Xinrun Wang, Shanqi Liu, Longtao Zheng, Bo An, Gang Pan</i>	
Universal Morphology Control Via Contextual Modulation	38286
<i>Zheng Xiong, Jacob Beck, Shimon Whiteson</i>	
Relevant Walk Search for Explaining Graph Neural Networks.....	38301
<i>Ping Xiong, Thomas Schnake, Michael Gastegger, Grégoire Montavon, Klaus Robert Muller, Shinichi Nakajima</i>	
Why Do Nearest Neighbor Language Models Work?	38325
<i>Frank F. Xu, Uri Alon, Graham Neubig</i>	
MixFlows: Principled Variational Inference Via Mixed Flows	38342
<i>Zuheng Xu, Naitong Chen, Trevor Campbell</i>	
Bit Allocation Using Optimization	38377
<i>Tongda Xu, Han Gao, Chenjian Gao, Yuanyuan Wang, Dailan He, Jinyong Pi, Jixiang Luo, Ziyu Zhu, Mao Ye, Hongwei Qin, Yan Wang, Jingjing Liu, Ya-Qin Zhang</i>	
Regret Bounds for Markov Decision Processes with Recursive Optimized Certainty Equivalents	38400
<i>Wenhao Xu, Xuefeng Gao, Xuedong He</i>	
Probabilistic Categorical Adversarial Attack and Adversarial Training	38428
<i>Han Xu, Pengfei He, Jie Ren, Yuxuan Wan, Zitao Liu, Hui Liu, Jiliang Tang</i>	
Hierarchical Neural Coding for Controllable CAD Model Generation.....	38443
<i>Xiang Xu, Pradeep Kumar Jayaraman, Joseph George Lambourne, Karl D. D. Willis, Yasutaka Furukawa</i>	
Efficient Sequence Transduction by Jointly Predicting Tokens and Durations	38462
<i>Hainan Xu, Fei Jia, Somshubra Majumdar, He Huang, Shinji Watanabe, Boris Ginsburg</i>	
Constrained Efficient Global Optimization of Expensive Black-Box Functions	38485
<i>Wenjie Xu, Yuning Jiang, Bratislav Svetozarevic, Colin Jones</i>	
Pareto Regret Analyses in Multi-Objective Multi-Armed Bandit	38499
<i>Mengfan Xu, Diego Klabjan</i>	
Diverse and Faithful Knowledge-Grounded Dialogue Generation Via Sequential Posterior Inference	38518
<i>Yan Xu, Deqian Kong, Dehong Xu, Ziwei Ji, Bo Pang, Pascale Fung, Ying Nian Wu</i>	
Quantifying the Variability Collapse of Neural Networks.....	38535
<i>Jing Xu, Haoxiong Liu</i>	

Progressive Purification for Instance-Dependent Partial Label Learning.....	38551
<i>Ning Xu, Biao Liu, Jiaqi Lv, Congyu Qiao, Xin Geng</i>	
PFGM++: Unlocking the Potential of Physics-Inspired Generative Models.....	38566
<i>Yilun Xu, Ziming Liu, Yonglong Tian, Shangyuan Tong, Max Tegmark, Tommi Jaakkola</i>	
Geometric Latent Diffusion Models for 3D Molecule Generation.....	38592
<i>Minkai Xu, Alexander S Powers, Ron O. Dror, Stefano Ermon, Jure Leskovec</i>	
The Power of Preconditioning in Overparameterized Low-Rank Matrix Sensing	38611
<i>Xingyu Xu, Yandi Shen, Yuejie Chi, Cong Ma</i>	

PART 49

Fascinating Supervisory Signals and Where to Find Them: Deep Anomaly Detection with Scale Learning	38655
<i>Hongzuo Xu, Yijie Wang, Juhui Wei, Songlei Jian, Yizhou Li, Ning Liu</i>	
Competing for Shareable Arms in Multi-Player Multi-Armed Bandits.....	38674
<i>Renzhe Xu, Haotian Wang, Xingxuan Zhang, Bo Li, Peng Cui</i>	
Sequential Predictive Conformal Inference for Time Series	38707
<i>Chen Xu, Yao Xie</i>	
MPLUG-2: A Modularized Multi-Modal Foundation Model Across Text, Image and Video	38728
<i>Haiyang Xu, Qinghao Ye, Ming Yan, Yaya Shi, Jiabo Ye, Yuanhong Xu, Chenliang Li, Bin Bi, Qi Qian, Wei Wang, Guohai Xu, Ji Zhang, Songfang Huang, Fei Huang, Jingren Zhou</i>	
ProtST: Multi-Modality Learning of Protein Sequences and Biomedical Texts.....	38749
<i>Minghao Xu, Xinyu Yuan, Santiago Miret, Jian Tang</i>	
Bayesian Design Principles for Frequentist Sequential Learning.....	38768
<i>Yunbei Xu, Assaf Zeevi</i>	
SLAMB: Accelerated Large Batch Training with Sparse Communication	38801
<i>Hang Xu, Wenxuan Zhang, Jiawei Fei, Yuzhe Wu, Tingwen Xie, Jun Huang, Yuchen Xie, Mohamed Elhoseiny, Panos Kalnis</i>	
Do Not Train It: A Linear Neural Architecture Search of Graph Neural Networks.....	38826
<i>Peng Xu, Lin Zhang, Xuanzhou Liu, Jiaqi Sun, Yue Zhao, Haiqin Yang, Bei Yu</i>	
An Instrumental Variable Approach to Confounded Off-Policy Evaluation.....	38848
<i>Yang Xu, Jin Zhu, Chengchun Shi, Shikai Luo, Rui Song</i>	
Near-Optimal Quantum Coreset Construction Algorithms for Clustering	38881
<i>Yecheng Xue, Xiaoyu Chen, Tongyang Li, Shaofeng H.-C. Jiang</i>	
A Study on Transformer Configuration and Training Objective.....	38913
<i>Fuzhao Xue, Jianghai Chen, Aixin Sun, Xiaozhe Ren, Zangwei Zheng, Xiaoxin He, Yongming Chen, Xin Jiang, Yang You</i>	
LazyGNN: Large-Scale Graph Neural Networks Via Lazy Propagation	38926
<i>Rui Xue, Haoyu Han, Mohamadali Torkamani, Jian Pei, Xiaorui Liu</i>	
Which Features Are Learnt by Contrastive Learning? on the Role of Simplicity Bias in Class Collapse and Feature Suppression.....	38938
<i>Yihao Xue, Siddharth Joshi, Eric Gan, Pin-Yu Chen, Baharan Mirzasoleiman</i>	
Adaptive Computation with Elastic Input Sequence	38971
<i>Fuzhao Xue, Valerii Likhoshesterov, Anurag Arnab, Neil Houlsby, Mostafa Dehghani, Yang You</i>	
Q-Learning Decision Transformer: Leveraging Dynamic Programming for Conditional Sequence Modelling in Offline RL.....	38989
<i>Taku Yamagata, Ahmed Khalil, Raul Santos-Rodriguez</i>	
Quantum Ridgelet Transform: Winning Lottery Ticket of Neural Networks with Quantum Computation	39008
<i>Hayata Yamasaki, Sathyawageeswar Subramanian, Satoshi Hayakawa, Sho Sonoda</i>	

Compressed Decentralized Proximal Stochastic Gradient Method for Nonconvex Composite Problems with Heterogeneous Data	39035
<i>Yonggui Yan, Jie Chen, Pin-Yu Chen, Xiaodong Cui, Songtao Lu, Yangyang Xu</i>	
Temporally Consistent Transformers for Video Generation	39062
<i>Wilson Yan, Danijar Hafner, Stephen James, Pieter Abbeel</i>	
Distortion and Uncertainty Aware Loss for Panoramic Depth Completion.....	39099
<i>Zhiqiang Yan, Xiang Li, Kun Wang, Shuo Chen, Jun Li, Jian Yang</i>	
Self-Interpretable Time Series Prediction with Counterfactual Explanations	39110
<i>Jingquan Yan, Hao Wang</i>	
Quantum 3D Graph Learning with Applications to Molecule Embedding.....	39126
<i>Ge Yan, Huaijin Wu, Junchi Yan</i>	
Fast Rates in Time-Varying Strongly Monotone Games	39138
<i>Yu-Hu Yan, Peng Zhao, Zhi-Hua Zhou</i>	
Proper Scoring Rules for Survival Analysis	39165
<i>Hiroki Yanagisawa</i>	
Behavior Contrastive Learning for Unsupervised Skill Discovery	39183
<i>Rushuai Yang, Chenjia Bai, Hongyi Guo, Siyuan Li, Bin Zhao, Zhen Wang, Peng Liu, Xuelong Li</i>	
Nested Elimination: A Simple Algorithm for Best-Item Identification from Choice-Based Feedback	39205
<i>Junwen Yang, Yifan Feng</i>	
Towards Better Graph Representation Learning with Parameterized Decomposition & Filtering.....	39234
<i>Mingqi Yang, Wenjie Feng, Yanming Shen, Bryan Hooi</i>	
Weighted Flow Diffusion for Local Graph Clustering with Node Attributes: An Algorithm and Statistical Guarantees	39252
<i>Shenghao Yang, Kimon Fountoulakis</i>	
Chemically Transferable Generative Backmapping of Coarse-Grained Proteins.....	39277
<i>Soojung Yang, Rafael Gomez-Bombarelli</i>	
Data Poisoning Attacks Against Multimodal Encoders.....	39299
<i>Ziqing Yang, Xinlei He, Zheng Li, Michael Backes, Mathias Humbert, Pascal Berrang, Yang Zhang</i>	
Towards Sustainable Learning: Coresets for Data-Efficient Deep Learning.....	39314
<i>Yu Yang, Hao Kang, Baharan Mirzasoleiman</i>	
Improving Adversarial Robustness by Putting More Regularizations on Less Robust Samples.....	39331
<i>Dongyoon Yang, Insung Kong, Yongdai Kim</i>	
Improving Adversarial Robustness of Deep Equilibrium Models with Explicit Regulations Along the Neural Dynamics	39349
<i>Zonghan Yang, Peng Li, Tianyu Pang, Yang Liu</i>	
Mitigating Spurious Correlations in Multi-Modal Models During Fine-Tuning.....	39365
<i>Yu Yang, Besmira Nushi, Hamid Palangi, Baharan Mirzasoleiman</i>	
A Theory of Representation Learning Gives a Deep Generalisation of Kernel Methods.....	39380
<i>Adam X. Yang, Maxime Robeyns, Edward Milsom, Ben Anson, Nandi Schoots, Laurence Aitchison</i>	
Efficient Algorithms for Exact Graph Matching on Correlated Stochastic Block Models with Constant Correlation.....	39416
<i>Joonhyuk Yang, Dongpil Shin, Hye Won Chung</i>	

PART 50

Are Neurons Actually Collapsed? on the Fine-Grained Structure in Neural Representations.....	39453
<i>Yongyi Yang, Jacob Steinhardt, Wei Hu</i>	
Generative Adversarial Symmetry Discovery	39488
<i>Jianke Yang, Robin Walters, Nima Dehmamy, Rose Yu</i>	

Boosting Offline Reinforcement Learning with Action Preference Query	39509
<i>Qisen Yang, Shenzhi Wang, Matthieu Gaetan Lin, Shiji Song, Gao Huang</i>	
Towards Controlled Data Augmentations for Active Learning	39524
<i>Jianan Yang, Haobo Wang, Sai Wu, Gang Chen, Junbo Zhao</i>	
What is Essential for Unseen Goal Generalization of Offline Goal-Conditioned RL?.....	39543
<i>Rui Yang, Lin Yong, Xiaoteng Ma, Hao Hu, Chongjie Zhang, Tong Zhang</i>	
Neural Prediction Errors Enable Analogical Visual Reasoning in Human Standard Intelligence Tests.....	39572
<i>Lingxiao Yang, Hongzhi You, Zonglei Zhen, Dahui Wang, Xiaohong Wan, Xiaohua Xie, Ru-Yuan Zhang</i>	
Change is Hard: A Closer Look at Subpopulation Shift	39584
<i>Yuzhe Yang, Haoran Zhang, Dina Katabi, Marzyeh Ghassemi</i>	
Continual Task Allocation in Meta-Policy Network Via Sparse Prompting	39623
<i>Yijun Yang, Tianyi Zhou, Jing Jiang, Guodong Long, Yuhui Shi</i>	
Hyperbolic Representation Learning: Revisiting and Advancing.....	39639
<i>Menglin Yang, Min Zhou, Rex Ying, Yankai Chen, Irwin King</i>	
Which is Better for Learning with Noisy Labels: The Semi-Supervised Method Or Modeling Label Noise?.....	39660
<i>Yu Yao, Mingming Gong, Yuxuan Du, Jun Yu, Bo Han, Kun Zhang, Tongliang Liu</i>	
How Bad is Top-K Recommendation Under Competing Content Creators?.....	39674
<i>Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, Haifeng Xu</i>	
MultiAdam: Parameter-Wise Scale-Invariant Optimizer for Multiscale Training of Physics-Informed Neural Networks.....	39702
<i>Jiacheng Yao, Chang Su, Zhongkai Hao, Songming Liu, Hang Su, Jun Zhu</i>	
Policy Mirror Ascent for Efficient and Independent Learning in Mean Field Games	39722
<i>Batuhan Yardim, Semih Cayci, Matthieu Geist, Niao He</i>	
Retrieval-Augmented Multimodal Language Modeling.....	39755
<i>Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Richard James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, Wen-Tau Yih</i>	
On the Power of Pre-Training for Generalization in RL: Provable Benefits and Hardness.....	39770
<i>Haotian Ye, Xiaoyu Chen, Liwei Wang, Simon Shaolei Du</i>	
Personalized Federated Learning with Inferred Collaboration Graphs	39801
<i>Rui Ye, Zhenyang Ni, Fangzhao Wu, Siheng Chen, Yanfeng Wang</i>	
Compositional Exemplars for In-Context Learning	39818
<i>Jiacheng Ye, Zhiyong Wu, Jiangtao Feng, Tao Yu, Lingpeng Kong</i>	
Corruption-Robust Algorithms with Uncertainty Weighting for Nonlinear Contextual Bandits and Markov Decision Processes.....	39834
<i>Chenlu Ye, Wei Xiong, Quanquan Gu, Tong Zhang</i>	
GNN&GBDT-Guided Fast Optimizing Framework for Large-Scale Integer Programming.....	39864
<i>Huigen Ye, Hua Xu, Hongyan Wang, Chengming Wang, Yu Jiang</i>	
FedDisco: Federated Learning with Discrepancy-Aware Collaboration	39879
<i>Rui Ye, Mingkai Xu, Jianyu Wang, Chenxin Xu, Siheng Chen, Yanfeng Wang</i>	
Towards Quantum Machine Learning for Constrained Combinatorial Optimization: A Quantum QAP Solver.....	39903
<i>Xinyu Ye, Ge Yan, Junchi Yan</i>	
Temporal Label Smoothing for Early Event Prediction	39913
<i>Hugo Yèche, Alizée Pace, Gunnar Ratsch, Rita Kuznetsova</i>	
From Temporal to Contemporaneous Iterative Causal Discovery in the Presence of Latent Confounders	39939
<i>Raanan Yehezkel Rohekar, Shami Nisimov, Yaniv Gurwicz, Gal Novik</i>	
Doubly Adversarial Federated Bandits.....	39951
<i>Jialin Yi, Milan Vojnovic</i>	

Online Prototype Alignment for Few-Shot Policy Transfer	39968
<i>Qi Yi, Rui Zhang, Shaohui Peng, Jiaming Guo, Yunkai Gao, Kaizhao Yuan, Ruizhi Chen, Siming Lan, Xing Hu, Zidong Du, Xishan Zhang, Qi Guo, Yunji Chen</i>	
MonoFlow: Rethinking Divergence GANs Via the Perspective of Wasserstein Gradient Flows	39984
<i>Mingxuan Yi, Zhanxing Zhu, Song Liu</i>	
SE(3) Diffusion Model with Application to Protein Backbone Generation	40001
<i>Jason Yin, Brian L. Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, Tommi Jaakkola</i>	
CoCo: A Coupled Contrastive Framework for Unsupervised Domain Adaptive Graph Classification.....	40040
<i>Nan Yin, Li Shen, Mengzhu Wang, Long Lan, Zeyu Ma, Chong Chen, Xian-Sheng Hua, Xiao Luo</i>	
Adaptive Estimation of Graphical Models Under Total Positivity	40054
<i>Jiayi Ying, José Vini'cius De Miranda Cardoso, Daniel P. Palomar</i>	
Improving Visual Prompt Tuning for Self-Supervised Vision Transformers.....	40075
<i>Seungryoung Yoo, Eunji Kim, Dahuin Jung, Jungbeom Lee, Sungroh Yoon</i>	
End-To-End Multi-Object Detection with a Regularized Mixture Model	40093
<i>Jaeyoung Yoo, Hojun Lee, Seunghyeon Seo, Inseop Chung, Nojun Kwak</i>	
EM-Network: Oracle Guided Self-Distillation for Sequence Learning.....	40111
<i>Ji Won Yoon, Sunghwan Ahn, Hyeonseung Lee, Minchan Kim, Seok Min Kim, Nam Soo Kim</i>	
Continual Learners Are Incremental Model Generalizers	40129
<i>Jaehong Yoon, Sung Ju Hwang, Yue Cao</i>	
An Investigation into Pre-Training Object-Centric Representations for Reinforcement Learning.....	40147
<i>Jaesik Yoon, Yi-Fu Wu, Heechul Bae, Sungjin Ahn</i>	
Graph Generative Model for Benchmarking Graph Neural Networks	40175
<i>Minji Yoon, Yue Wu, John Palowitch, Bryan Perozzi, Russ Salakhutdinov</i>	
Analyzing Convergence in Quantum Neural Networks: Deviations from Neural Tangent Kernels.....	40199
<i>Xuchen You, Shouvanik Chakrabarti, Boyang Chen, Xiaodi Wu</i>	
Entropy-Driven Unsupervised Keypoint Representation Learning in Videos	40225
<i>Ali Younes, Simone Schaub-Meyer, Georgia Chalvatzaki</i>	

PART 51

The Benefits of Model-Based Generalization in Reinforcement Learning	40254
<i>Kenny John Young, Aditya Ramesh, Louis Kirsch, Jürgen Schmidhuber</i>	
COLA: Orchestrating Error Coding and Learning for Robust Neural Network Inference Against Hardware Defects	40277
<i>Anlan Yu, Ning Lyu, Jieming Yin, Zhiyuan Yan, Wujie Wen</i>	
Delving into Noisy Label Detection with Clean Data	40290
<i>Chenglin Yu, Xinsong Ma, Weiwei Liu</i>	
Bag of Tricks for Training Data Extraction from Language Models.....	40306
<i>Weichen Yu, Tianyu Pang, Qian Liu, Chao Du, Bingyi Kang, Yan Huang, Min Lin, Shuicheng Yan</i>	
Discover-Then-Rank Unlabeled Support Vectors in the Dual Space for Multi-Class Active Learning.....	40321
<i>Dayou Yu, Weishi Shi, Qi Yu</i>	
Long-Term Rhythmic Video Soundtracker.....	40339
<i>Jiashuo Yu, Yaohui Wang, Xinyuan Chen, Xiao Sun, Yu Qiao</i>	
Adversarial Parameter Attack on Deep Neural Networks	40354
<i>Lijia Yu, Yihan Wang, Xiao-Shan Gao</i>	
CodeIPPrompt: Intellectual Property Infringement Assessment of Code Language Models	40373
<i>Zhiyuan Yu, Yuhao Wu, Ning Zhang, Chenguang Wang, Yevgeniy Vorobeychik, Chaowei Xiao</i>	
SeedGNN: Graph Neural Network for Supervised Seeded Graph Matching	40390
<i>Liren Yu, Jiaming Xu, Xiaojun Lin</i>	

Efficient and Equivariant Graph Networks for Predicting Quantum Hamiltonian	40412
<i>Haiyang Yu, Zhao Xu, Xiaofeng Qian, Xiaoning Qian, Shuiwang Ji</i>	
On the Global Convergence of Risk-Averse Policy Gradient Methods with Expected Conditional Risk Measures	40425
<i>Xian Yu, Lei Ying</i>	
Actor-Critic Alignment for Offline-To-Online Reinforcement Learning	40452
<i>Zishun Yu, Xinhua Zhang</i>	
Master-ASR: Achieving Multilingual Scalability and Low-Resource Adaptation in ASR with Modular Learning	40475
<i>Zhongzhi Yu, Yang Zhang, Kaizhi Qian, Cheng Wan, Yonggan Fu, Yongan Zhang, Yingyan Celine Lin</i>	
Coordinate Descent Methods for Fractional Minimization	40488
<i>Ganzhao Yuan</i>	
On the Power of Foundation Models	40519
<i>Yang Yuan</i>	
Automatic Intrinsic Reward Shaping for Exploration in Deep Reinforcement Learning	40531
<i>Mingqi Yuan, Bo Li, Xin Jin, Wenjun Zeng</i>	
Traversing Between Modes in Function Space for Fast Ensembling	40555
<i>Eunggu Yun, Hyungi Lee, Giung Nam, Juho Lee</i>	
Conformal Prediction with Missing Values	40578
<i>Margaux Zaffran, Aymeric Dieuleveut, Julie Josse, Yaniv Romano</i>	
KDEformer: Accelerating Transformers Via Kernel Density Estimation	40605
<i>Amir Zandieh, Insu Han, Majid Daliri, Amin Karbasi</i>	
Bayesian Estimation of Differential Privacy	40624
<i>Santiago Zanella-Beguelin, Lukas Wutschitz, Shruti Tople, Ahmed Salem, Victor Rühle, Andrew Paverd, Mohammad Naseri, Boris Köpf, Daniel Jones</i>	
When is Realizability Sufficient for Off-Policy Reinforcement Learning?	40637
<i>Andrea Zanette</i>	
On Distribution Dependent Sub-Logarithmic Query Time of Learned Indexing	40669
<i>Sepanta Zeighami, Cyrus Shahabi</i>	
Sequential Counterfactual Risk Minimization	40681
<i>Houssam Zenati, Eustache Diemert, Matthieu Martin, Julien Mairal, Pierre Gaillard</i>	
LookupFFN: Making Transformers Compute-Lite for CPU Inference	40707
<i>Zhanpeng Zeng, Michael Davies, Pranav Pulijala, Karthikeyan Sankaralingam, Vikas Singh</i>	
Attribute-Efficient PAC Learning of Low-Degree Polynomial Threshold Functions with Nasty Noise	40719
<i>Shiwei Zeng, Jie Shen</i>	
Generative Graph Dictionary Learning	40749
<i>Zhichen Zeng, Ruike Zhu, Yinglong Xia, Hanqing Zeng, Hanghang Tong</i>	
Stabilizing Transformer Training by Preventing Attention Entropy Collapse	40770
<i>Shuangfei Zhai, Tatiana Likhomanenko, Etai Littwin, Dan Busbridge, Jason Ramapuram, Yizhe Zhang, Jiatao Gu, Joshua M. Susskind</i>	
Offline Learning in Markov Games with General Function Approximation	40804
<i>Yuheng Zhang, Yu Bai, Nan Jiang</i>	
Learning Useful Representations for Shifting Tasks and Distributions	40830
<i>Jianyu Zhang, Leon Bottou</i>	
Nonparametric Iterative Machine Teaching	40851
<i>Chen Zhang, Xiaofeng Cao, Weiyang Liu, Ivor Tsang, James Kwok</i>	
Matrix Estimation for Individual Fairness	40871
<i>Cindy Zhang, Sarah Huiyi Cen, Devavrat Shah</i>	
Graph Contrastive Backdoor Attacks	40888
<i>Hangfan Zhang, Jinghui Chen, Lu Lin, Jinyuan Jia, Dinghao Wu</i>	

Effective Minkowski Dimension of Deep Nonparametric Regression: Function Approximation and Statistical Theories	40911
<i>Zixuan Zhang, Minshuo Chen, Mengdi Wang, Wenjing Liao, Tuo Zhao</i>	
Tractable Control for Autoregressive Language Generation	40932
<i>Honghua Zhang, Meihua Dang, Nanyun Peng, Guy Van Den Broeck</i>	
CataBEEM: Integrating Latent Interaction Categories in Node-Wise Community Detection Models for Network Data.....	40946
<i>Yuhua Zhang, Walter H. Dempsey</i>	
Rethink DARTS Search Space and Renovate a New Benchmark	40976
<i>Jiuling Zhang, Zhiming Ding</i>	
Team Belief DAG: Generalizing the Sequence Form to Team Games for Fast Computation of Correlated Team Max-Min Equilibria Via Regret Minimization.....	40996
<i>Brian Hu Zhang, Gabriele Farina, Tuomas Sandholm</i>	
A Complete Expressiveness Hierarchy for Subgraph GNNs Via Subgraph Weisfeiler-Lehman Tests.....	41019
<i>Bohang Zhang, Guhao Feng, Yiheng Du, Di He, Liwei Wang</i>	

PART 52

Crafting Training Degradation Distribution for the Accuracy-Generalization Trade-Off in Real-World Super-Resolution	41078
<i>Ruofan Zhang, Jinjin Gu, Haoyu Chen, Chao Dong, Yulun Zhang, Wenming Yang</i>	
Prompting Large Language Model for Machine Translation: A Case Study	41092
<i>Biao Zhang, Barry Haddow, Alexandra Birch</i>	
On the Interplay Between Misspecification and Sub-Optimality Gap in Linear Contextual Bandits.....	41111
<i>Weitong Zhang, Jiafan He, Zhiyuan Fan, Quanquan Gu</i>	
When Sparsity Meets Contrastive Models: Less Graph Data Can Bring Better Class-Balanced Representations	41133
<i>Chunhui Zhang, Chao Huang, Yijun Tian, Qianlong Wen, Zhongyu Ouyang, Youhuan Li, Yanfang Ye, Chuxu Zhang</i>	
Spatial-Temporal Graph Learning with Adversarial Contrastive Adaptation	41151
<i>Qianru Zhang, Chao Huang, Lianghao Xia, Zheng Wang, Siu Ming Yiu, Ruihua Han</i>	
Towards Coherent Image Inpainting Using Denoising Diffusion Implicit Models	41164
<i>Guanhua Zhang, Jiabao Ji, Yang Zhang, Mo Yu, Tommi Jaakkola, Shiyu Chang</i>	
CAB: Comprehensive Attention Benchmarking on Long Sequence Modeling.....	41194
<i>Jun Zhang, Shuyang Jiang, Jiangtao Feng, Lin Zheng, Lingpeng Kong</i>	
Adaptive Barrier Smoothing for First-Order Policy Gradient with Contact Dynamics.....	41219
<i>Shenao Zhang, Wanxin Jin, Zhaoran Wang</i>	
One-Step Estimator for Permuted Sparse Recovery	41244
<i>Hang Zhang, Ping Li</i>	
Quantum Lower Bounds for Finding Stationary Points of Nonconvex Functions	41268
<i>Chenyi Zhang, Tongyang Li</i>	
Improving Medical Predictions by Irregular Multimodal Electronic Health Records Modeling	41300
<i>Xinlu Zhang, Shiyang Li, Zhiyu Chen, Xifeng Yan, Linda Ruth Petzold</i>	
FedCR: Personalized Federated Learning Based on Across-Client Common Representation with Conditional Mutual Information Regularization	41314
<i>Hao Zhang, Chenglin Li, Wenrui Dai, Junni Zou, Hongkai Xiong</i>	
On the Optimality of Misspecified Kernel Ridge Regression	41331
<i>Haobo Zhang, Yicheng Li, Weihao Lu, Qian Lin</i>	
Fed-CBS: A Heterogeneity-Aware Client Sampling Mechanism for Federated Learning Via Class-Imbalance Reduction	41354
<i>Jianyi Zhang, Ang Li, Minxue Tang, Jingwei Sun, Xiang Chen, Fan Zhang, Changyou Chen, Yiran Chen, Hai Li</i>	

Learning Subpocket Prototypes for Generalizable Structure-Based Drug Design	41382
<i>Zaixi Zhang, Qi Liu</i>	
No One Idles: Efficient Heterogeneous Federated Learning with Parallel Edge and Server Computation	41399
<i>Feilong Zhang, Xianming Liu, Shiyi Lin, Gang Wu, Xiong Zhou, Junjun Jiang, Xiangyang Ji</i>	
The Wisdom of Hindsight Makes Language Models Better Instruction Followers.....	41414
<i>Tianjun Zhang, Fangchen Liu, Justin Wong, Pieter Abbeel, Joseph E. Gonzalez</i>	
Detecting Adversarial Data by Probing Multiple Perturbations Using Expected Perturbation Score	41429
<i>Shuhai Zhang, Feng Liu, Jiahao Yang, Yifan Yang, Changsheng Li, Bo Han, Mingkui Tan</i>	
On Enhancing Expressive Power Via Compositions of Single Fixed-Size ReLU Network.....	41452
<i>Shijun Zhang, Jianfeng Lu, Hongkai Zhao</i>	
Bi-Directional Masks for Efficient N:M Sparse Training.....	41488
<i>Yuxin Zhang, Yiting Luo, Mingbao Lin, Yunshan Zhong, Jingjing Xie, Fei Chao, Rongrong Ji</i>	
Towards Unbiased Training in Federated Open-World Semi-Supervised Learning.....	41498
<i>Jie Zhang, Xiaosong Ma, Song Guo, Wenchao Xu</i>	
Interactive Object Placement with Reinforcement Learning.....	41510
<i>Shengping Zhang, Quanling Meng, Qinglin Liu, Liqiang Nie, Bineng Zhong, Xiaopeng Fan, Rongrong Ji</i>	
Optimal Shrinkage for Distributed Second-Order Optimization	41523
<i>Fangzhao Zhang, Mert Pilanci</i>	
"Why Did the Model Fail?": Attributing Model Performance Changes to Distribution Shifts	41550
<i>Haoran Zhang, Harvineet Singh, Marzyeh Ghassemi, Shalmali Joshi</i>	
Learning Regions of Interest for Bayesian Optimization with Adaptive Level-Set Estimation	41579
<i>Fengxue Zhang, Jialin Song, James C Bowden, Alexander Ladd, Yisong Yue, Thomas Desautels, Yuxin Chen</i>	
A Category-Theoretical Meta-Analysis of Definitions of Disentanglement.....	41596
<i>Yivan Zhang, Masashi Sugiyama</i>	
On the Convergence of SARSA with Linear Function Approximation.....	41613
<i>Shangtong Zhang, Remi Tachet Des Combes, Romain Laroche</i>	
AdaNPC: Exploring Non-Parametric Classifier for Test-Time Adaptation.....	41647
<i>Yifan Zhang, Xue Wang, Kexin Jin, Kun Yuan, Zhang Zhang, Liang Wang, Rong Jin, Tieniu Tan</i>	
On the Generalization of Multi-Modal Contrastive Learning	41677
<i>Qi Zhang, Yifei Wang, Yisen Wang</i>	
ConCerNet: A Contrastive Learning Based Framework for Automated Conservation Law Discovery and Trustworthy Dynamical System Prediction.....	41694
<i>Wang Zhang, Tsui-Wei Weng, Subhro Das, Alexandre Megretski, Luca Daniel, Lam M. Nguyen</i>	
Towards Trustworthy Explanation: On Causal Rationalization.....	41715
<i>Wenbo Zhang, Tong Wu, Yunlong Wang, Yong Cai, Hengrui Cai</i>	
Demystifying Uneven Vulnerability of Link Stealing Attacks Against Graph Neural Networks.....	41737
<i>He Zhang, Bang Wu, Shuo Wang, Xiangwen Yang, Minhui Xue, Shirui Pan, Xingliang Yuan</i>	
Provable Dynamic Fusion for Low-Quality Multimodal Data	41753
<i>Qingyang Zhang, Haitao Wu, Changqing Zhang, Qinghua Hu, Huazhu Fu, Joey Tianyi Zhou, Xi Peng</i>	
ReDi: Efficient Learning-Free Diffusion Inference Via Trajectory Retrieval	41770
<i>Kexun Zhang, Xianjun Yang, William Yang Wang, Lei Li</i>	
Nearly Optimal Competitive Ratio for Online Allocation Problems with Two-Sided Resource Constraints and Finite Requests	41786
<i>Qixin Zhang, Wenbing Ye, Zaiyi Chen, Haoyuan Hu, Enhong Chen, Yu Yang</i>	
Do You Remember? Overcoming Catastrophic Forgetting for Fake Audio Detection.....	41819
<i>Xiaohui Zhang, Jiangyan Yi, Jianhua Tao, Chenglong Wang, Chu Yuan Zhang</i>	

Coder Reviewer Reranking for Code Generation.....	41832
<i>Tianyi Zhang, Tao Yu, Tatsunori Hashimoto, Mike Lewis, Wen-Tau Yih, Daniel Fried, Sida Wang</i>	
DP-Fast MH: Private, Fast, and Accurate Metropolis-Hastings for Large-Scale Bayesian Inference.....	41847
<i>Wanrong Zhang, Ruqi Zhang</i>	
Nearly-Tight Bounds for Deep Kernel Learning.....	41861
<i>Yifan Zhang, Min-Ling Zhang</i>	

PART 53

OpenFE: Automated Feature Generation with Expert-Level Performance.....	41880
<i>Tianping Zhang, Zheyu Aqa Zhang, Zhiyuan Fan, Haoyan Luo, Fengyuan Liu, Qian Liu, Wei Cao, Li Jian</i>	
Optimal Horizon-Free Reward-Free Exploration for Linear Mixture MDPs.....	41902
<i>Junkai Zhang, Weitong Zhang, Quanquan Gu</i>	
Unlocking Slot Attention by Changing Optimal Transport Costs.....	41931
<i>Yan Zhang, David W. Zhang, Simon Lacoste-Julien, Gertjan J. Burghouts, Cees G. M. Snoek</i>	
Towards a Persistence Diagram that is Robust to Noise and Varied Densities.....	41952
<i>Hang Zhang, Kaifeng Zhang, Kai Ming Ting, Ye Zhu</i>	
Robust Situational Reinforcement Learning in Face of Context Disturbances.....	41973
<i>Jinpeng Zhang, Yufeng Zheng, Chuheng Zhang, Li Zhao, Lei Song, Yuan Zhou, Jiang Bian</i>	
Patch-Level Contrastive Learning Via Positional Query for Visual Pre-Training.....	41990
<i>Shaofeng Zhang, Qiang Zhou, Zhibin Wang, Fan Wang, Junchi Yan</i>	
Men Also Do Laundry: Multi-Attribute Bias Amplification.....	42000
<i>Dora Zhao, Jerone Andrews, Alice Xiang</i>	
Rockmate: An Efficient, Fast, Automatic and Generic Tool for Re-Materialization in PyTorch.....	42018
<i>Xunyi Zhao, Théotime Le Hellard, Lionel Eyraud-Dubois, Julia Gusak, Olivier Beaumont</i>	
Revisiting Structured Variational Autoencoders.....	42046
<i>Yixiu Zhao, Scott Linderman</i>	
On Pitfalls of Test-Time Adaptation.....	42058
<i>Hao Zhao, Yuejiang Liu, Alexandre Alahi, Tao Lin</i>	
Addressing Budget Allocation and Revenue Allocation in Data Market Environments Using an Adaptive Sampling Algorithm.....	42081
<i>Boxin Zhao, Boxiang Lyu, Raul Castro Fernandez, Mladen Kolar</i>	
X-Paste: Revisiting Scalable Copy-Paste for Instance Segmentation Using CLIP and StableDiffusion.....	42098
<i>Hanqing Zhao, Dianmo Sheng, Jianmin Bao, Dongdong Chen, Dong Chen, Fang Wen, Lu Yuan, Ce Liu, Wenbo Zhou, Qi Chu, Weiming Zhang, Nenghai Yu</i>	
Revisiting Simple Regret: Fast Rates for Returning a Good Arm.....	42110
<i>Yao Zhao, Connor Stephens, Csaba Szepesvari, Kwang-Sung Jun</i>	
Transformed Distribution Matching for Missing Value Imputation.....	42159
<i>He Zhao, Ke Sun, Amir Dezfouli, Edwin V. Bonilla</i>	
Protecting Language Generation Models Via Invisible Watermarking.....	42187
<i>Xuandong Zhao, Yu-Xiang Wang, Lei Li</i>	
Local Optimization Achieves Global Optimality in Multi-Agent Reinforcement Learning.....	42200
<i>Yulai Zhao, Zhuoran Yang, Zhaoran Wang, Jason D. Lee</i>	
Simplified Temporal Consistency Reinforcement Learning.....	42227
<i>Yi Zhao, Wenshuai Zhao, Rinu Boney, Juho Kannala, Joni Pajarinen</i>	
RLEG: Vision-Language Representation Learning with Diffusion-Based Embedding Generation.....	42247
<i>Liming Zhao, Kecheng Zheng, Yun Zheng, Deli Zhao, Jingren Zhou</i>	

Optimal Online Generalized Linear Regression with Stochastic Noise and Its Application to Heteroscedastic Bandits.....	42259
<i>Heyang Zhao, Dongruo Zhou, Jiafan He, Quanquan Gu</i>	
Does Continual Learning Equally Forget All Parameters?.....	42280
<i>Haiyan Zhao, Tianyi Zhou, Guodong Long, Jing Jiang, Chengqi Zhang</i>	
Online Learning in Stackelberg Games with an Omniscient Follower.....	42304
<i>Geng Zhao, Banghua Zhu, Jiantao Jiao, Michael Jordan</i>	
Structure-Informed Language Models Are Protein Designers.....	42317
<i>Zaixiang Zheng, Yifan Deng, Dongyu Xue, Yi Zhou, Fei Ye, Quanquan Gu</i>	
Semi-Supervised Offline Reinforcement Learning with Action-Free Trajectories.....	42339
<i>Qinqing Zheng, Mikael Henaff, Brandon Amos, Aditya Grover</i>	
Improved Techniques for Maximum Likelihood Estimation for Diffusion ODEs	42363
<i>Kaiwen Zheng, Cheng Lu, Jianfei Chen, Jun Zhu</i>	
Fast Sampling of Diffusion Models Via Operator Learning.....	42390
<i>Hongkai Zheng, Weili Nie, Arash Vahdat, Kamyar Azizzadenesheli, Anima Anandkumar</i>	
Outline, Then Details: Syntactically Guided Coarse-To-Fine Code Generation	42403
<i>Wenqing Zheng, S P Sharan, Ajay Kumar Jaiswal, Kevin Wang, Yihan Xi, Dejia Xu, Zhangyang Wang</i>	
Revisiting Discriminative Vs. Generative Classifiers: Theory and Implications	42420
<i>Chenyu Zheng, Guoqiang Wu, Fan Bao, Yue Cao, Chongxuan Li, Jun Zhu</i>	
Evidential Interactive Learning for Medical Image Captioning	42478
<i>Ervine Zheng, Qi Yu</i>	
Finding the Missing-Half: Graph Complementary Learning for Homophily-Prone and Heterophily-Prone Graphs	42492
<i>Yizhen Zheng, He Zhang, Vincent Lee, Yu Zheng, Xiao Wang, Shirui Pan</i>	
Multi-Agent Online Scheduling: MMS Allocations for Indivisible Items.....	42506
<i>Shengwei Zhou, Rufan Bai, Xiaowei Wu</i>	
Eliminating Adversarial Noise Via Information Discard and Robust Representation Restoration.....	42517
<i>Dawei Zhou, Yukun Chen, Nannan Wang, Decheng Liu, Xinbo Gao, Tongliang Liu</i>	
Brainformers: Trading Simplicity for Efficiency.....	42531
<i>Yanqi Zhou, Nan Du, Yanping Huang, Daiyi Peng, Chang Lan, Da Huang, Siamak Shakeri, David So, Andrew M. Dai, Yifeng Lu, Zhifeng Chen, Quoc V Le, Claire Cui, James Laudon, Jeff Dean</i>	
Implicit Regularization Leads to Benign Overfitting for Sparse Linear Regression	42543
<i>Mo Zhou, Rong Ge</i>	
ODS: Test-Time Adaptation in the Presence of Open-World Data Shift.....	42574
<i>Zhi Zhou, Lan-Zhe Guo, Lin-Han Jia, Dingchu Zhang, Yu-Feng Li</i>	
Fourmer: An Efficient Global Modeling Paradigm for Image Restoration	42589
<i>Man Zhou, Jie Huang, Chun-Le Guo, Chongyi Li</i>	
Controlled Text Generation with Natural Language Instructions	42602
<i>Wangchunshu Zhou, Yuchen Eleanor Jiang, Ethan Wilcox, Ryan Cotterell, Mrinmaya Sachan</i>	
NNSplitter: An Active Defense Solution for DNN Model Via Automated Weight Obfuscation.....	42614
<i>Tong Zhou, Yukui Luo, Shaolei Ren, Xiaolin Xu</i>	
Deep Latent State Space Models for Time-Series Generation.....	42625
<i>Linqi Zhou, Michael Poli, Winnie Xu, Stefano Massaroli, Stefano Ermon</i>	
SlotGAT: Slot-Based Message Passing for Heterogeneous Graphs.....	42644
<i>Ziang Zhou, Jieming Shi, Renchi Yang, Yuanhang Zou, Qing Li</i>	
Fast Online Node Labeling for Very Large Graphs.....	42658
<i>Baojian Zhou, Yifan Sun, Reza Babanezhad Harikandeh</i>	

PART 54

Horizon-Free and Variance-Dependent Reinforcement Learning for Latent Markov Decision Processes	42698
<i>Runlong Zhou, Ruosong Wang, Simon Shaolei Du</i>	
Phase-Aware Adversarial Defense for Improving Adversarial Robustness	42724
<i>Dawei Zhou, Nannan Wang, Heng Yang, Xinbo Gao, Tongliang Liu</i>	
From Relational Pooling to Subgraph GNNs: A Universal Framework for More Expressive Graph Neural Networks.....	42742
<i>Cai Zhou, Xiyuan Wang, Muhan Zhang</i>	
Towards Omni-Generalizable Neural Methods for Vehicle Routing Problems	42769
<i>Jianan Zhou, Yaoxin Wu, Wen Song, Zhiguang Cao, Jie Zhang</i>	
A Three-Regime Model of Network Pruning	42790
<i>Yefan Zhou, Yaoqing Yang, Arin Chang, Michael W. Mahoney</i>	
Learning to Decouple Complex Systems	42810
<i>Zihan Zhou, Tianshu Yu</i>	
ESC: Exploration with Soft Commonsense Constraints for Zero-Shot Object Navigation.....	42829
<i>Kaiwen Zhou, Kaizhi Zheng, Connor Pryor, Yilin Shen, Hongxia Jin, Lise Getoor, Xin Eric Wang</i>	
On Strengthening and Defending Graph Reconstruction Attack with Markov Chain Approximation.....	42843
<i>Zhanke Zhou, Chenyu Zhou, Xuan Li, Jiangchao Yao, Quanming Yao, Bo Han</i>	
Sharp Variance-Dependent Bounds in Reinforcement Learning: Best of Both Worlds in Stochastic and Deterministic Environments	42878
<i>Runlong Zhou, Zhang Zihan, Simon Shaolei Du</i>	
Learning Unforeseen Robustness from Out-Of-Distribution Data Using Equivariant Domain Translator.....	42915
<i>Sicheng Zhu, Bang An, Furong Huang, Sanghyun Hong</i>	
Markovian Gaussian Process Variational Autoencoders.....	42938
<i>Harrison Zhu, Carles Balsells-Rodas, Yingzhen Li</i>	
Mixture Proportion Estimation Beyond Irreducibility.....	42962
<i>Yilun Zhu, Aaron Fjeldsted, Darren Holland, George Landon, Azaree Lintereur, Clayton Scott</i>	
Exploring Model Dynamics for Accumulative Poisoning Discovery.....	42983
<i>Jianing Zhu, Xiawei Guo, Jiangchao Yao, Chao Du, Li He, Shuo Yuan, Tongliang Liu, Liang Wang, Bo Han</i>	
Decentralized SGD and Average-Direction SAM Are Asymptotically Equivalent	43005
<i>Tongtian Zhu, Fengxiang He, Kaixuan Chen, Mingli Song, Dacheng Tao</i>	
Principled Reinforcement Learning with Human Feedback from Pairwise Or K-Wise Comparisons	43037
<i>Banghua Zhu, Michael Jordan, Jiantao Jiao</i>	
Unleashing Mask: Explore the Intrinsic Out-Of-Distribution Detection Capability	43068
<i>Jianing Zhu, Hengzhuang Li, Jiangchao Yao, Tongliang Liu, Jianliang Xu, Bo Han</i>	
Benign Overfitting in Deep Neural Networks Under Lazy Training.....	43105
<i>Zhenyu Zhu, Fanghui Liu, Grigorios Chrysos, Francesco Locatello, Volkan Cevher</i>	
Interpolation for Robust Learning: Data Augmentation on Wasserstein Geodesics	43129
<i>Jiacheng Zhu, Jielin Qiu, Aritra Guha, Zhuolin Yang, Xuanlong Nguyen, Bo Li, Ding Zhao</i>	
LeadFL: Client Self-Defense Against Model Poisoning in Federated Learning	43158
<i>Chaoyi Zhu, Stefanie Roos, Lydia Y. Chen</i>	
XTab: Cross-Table Pretraining for Tabular Transformers	43181
<i>Bingzhao Zhu, Xingjian Shi, Nick Erickson, Mu Li, George Karypis, Mahsa Shoaran</i>	
Provable Multi-Instance Deep AUC Maximization with Stochastic Pooling.....	43205
<i>Dixian Zhu, Bokun Wang, Zhi Chen, Yaxing Wang, Milan Sonka, Xiaodong Wu, Tianbao Yang</i>	

Surrogate Model Extension (SME): A Fast and Accurate Weight Update Attack on Federated Learning	43228
<i>Junyi Zhu, Ruicong Yao, Matthew B. Blaschko</i>	
Weak Proxies Are Sufficient and Preferable for Fairness with Missing Sensitive Attributes	43258
<i>Zhaowei Zhu, Yuanshun Yao, Jiankai Sun, Hang Li, Yang Liu</i>	
Label Distributionally Robust Losses for Multi-Class Classification: Consistency, Robustness and Adaptivity	43289
<i>Dixian Zhu, Yiming Ying, Tianbao Yang</i>	
Likelihood Adjusted Semidefinite Programs for Clustering Heterogeneous Data	43326
<i>Yubo Zhuang, Xiaohui Chen, Yun Yang</i>	
Are Random Decompositions All We Need in High Dimensional Bayesian Optimisation?	43347
<i>Juliusz Krzysztof Ziomek, Haitham Bou Ammar</i>	
Revisiting Bellman Errors for Offline Model Selection	43369
<i>Joshua P Zitovsky, Daniel De Marchi, Rishabh Agarwal, Michael Rene Kosorok</i>	
Spred: Solving L1 Penalty with SGD	43407
<i>Liu Ziyin, Zihao Wang</i>	
The Benefits of Mixup for Feature Learning	43423
<i>Difan Zou, Yuan Cao, Yuanzhi Li, Quanquan Gu</i>	

Author Index