39th Conference on Uncertainty in Artificial Intelligence (UAI 2023)

Proceedings of Machine Learning Research Volume 216

Pittsburgh, Pennsylvania, USA
31 July - 4 August 2023

Part 1 of 4

ISBN: 978-1-7138-7806-3

Printed from e-media with permission by:

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

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

Copyright© (2023) by Association For Uncertainty in Artificial Intelligence (AUAI)
All rights reserved.

Printed by Curran Associates, Inc. (2024)

For permission requests, please contact Association For Uncertainty in Artificial Intelligence (AUAI) at the address below.

Association For Uncertainty in Artificial Intelligence (AUAI)
Department of Computer Science
University of British Columbia
Vancouver BC Canada V6T 1Z4

Phone: (703) 787-8044
Fax: (703) 478-0829

www.auai.org

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

SubMix: Learning to Mix Graph Sampling Heuristics 1
S. Abu-El-Haija, J.V. Dillon, B. Fatemi, K. Axiotis, N. Bulut, J. Gasteiger, B. Perozzi, M. Bateni

Guided Deep Kernel Learning 11
I. Achituve, G. Chechik, E. Fetaya

Bounding the Optimal Value Function in Compositional Reinforcement Learning 22
J. Adamczyk, V. Makarenko, A. Arriojas, S. Tiomkin, R.V. Kulkarni

A Decoder Suffices for Query-Adaptive Variational Inference 33
S. Agarwal, G. Hope, A. Younis, E.B. Sudderth

FLASH: Automating Federated Learning Using CASH 45
M.I.I. Alam, K. Kar, T. Salonidis, H. Samulowitz

Transfer Learning for Individual Treatment Effect Estimation 56
A. Aloui, J. Dong, C.P. Le, V. Tarokh

Robust Gaussian Process Regression with the Trimmed Marginal Likelihood 67
D. Andrade, A. Takeda

Two Sides of Miscalibration: Identifying Over and Under-Confidence Prediction for Network Calibration 77
S. Ao, S. Rueger, A. Siddharthan

Quantifying Lottery Tickets Under Label Noise: Accuracy, Calibration, and Complexity 88
V. Arora, D. Irto, S. Goldt, G. Sanguinetti

Bayesian Inference Approach for Entropy Regularized Reinforcement Learning with Stochastic Dynamics 99
A. Arriojas, J. Adamczyk, S. Tiomkin, R.V. Kulkarni

Neural Tangent Kernel at Initialization: LinearWidth Suffices 110
A. Banerjee, P. Cisneros-Velarde, L. Zhu, M. Belkin

Do We Become Wiser with Time? on Causal Equivalence with Tiered Background Knowledge 119
C.W. Bang, V. Didelez

Sample Boosting Algorithm (SamBA) - An Interpretable Greedy Ensemble Classifier Based on Local Expertise for Fat Data 130
B. Bauvin, C. Capponi, F. Clerc, P. Germain, S. Koco, J. Corbeil

Learning Choice Functions with Gaussian Processes 141
A. Benavoli, D. Azzimonti, D. Piga

Inference of a Rumor's Source in the Independent Cascade Model 152
P. Berenbrink, M. Hahn-Klimroth, D. Kaaser, L. Krieg, M. Rau

Best Arm Identification in Rare Events 163
A. Bhattacharjee, S. Vijayan, S. Juneja

BeliefPPG: Uncertainty-Aware Heart Rate Estimation from PPG Signals Via Belief Propagation 173
V. Bieri, P. Streli, B.U. Demirel, C. Holz

Amortized Inference for Gaussian Process Hyperparameters of Structured Kernels 184
M. Bitzer, M. Meister, C. Zimmer

Correcting for Selection Bias and Missing Response in Regression Using Privileged Information 195
P. Boeken, N. de Kroon, M. de Jong, J.M. Mooij, O. Zoeter

Efficient Learning of Minimax Risk Classifiers in High Dimensions 206
K. Bondugula, S. Mazuelas, A. Perez

Approximating Probabilistic Explanations Via Supermodular Minimization 216
L. Bounia, F. Koriche

Inference for Mark-Censored Temporal Point Processes 226
A. Boyd, Y. Chang, S. Mandt, P. Smyth

Testing Conventional Wisdom (of the Crowd) 237
N. Burrell, G. Schoenebeck

Overcoming Language Priors for Visual Question Answering Via Loss Rebalancing Label and Global Context 249
R. Cao, Z. Li

Scaling Integer Arithmetic in Probabilistic Programs 260
W.X. Cao, P. Garg, R. Tjoa, S. Holtzen, T. Millstein, G. Van den Broeck

Human Control: Definitions and Algorithms 271
R. Carey, T. Everitt

Scalable Nonparametric Bayesian Learning for Dynamic Velocity Fields 282
S. Chakraborty, A. Guha, R. Lei, X.L. Nguyen

Learning in Online MDPs: Is There a Price for Handling the Communicating Case? 293
G. Chandrasekaran, A. Tewari

Modified Retrace for Off-Policy Temporal Difference Learning 303
X. Chen, X. Ma, Y. Li, G. Yang, S. Yang, Y. Gao

Benign Overfitting in Adversarially Robust Linear Classification 313
J. Chen, Y. Cao, Q. Gu

An Effective Negotiating Agent Framework Based on Deep Offline Reinforcement Learning 324
S. Chen, J. Zhao, G. Weiss, R. Su, K. Lei

MFA: Multi-Layer Feature-Aware Attack for Object Detection 336
W. Chen, Y. Zhang, Z. Li, Y. Wang

Differential Privacy in Cooperative Multiagent Planning 347
B. Chen, C. Hawkins, M.O. Karabag, C. Neary, M. Hale, U. Topcu

Causal InferenceWith Outcome-Dependent Missingness and Self-Censoring 358
J.M. Chen, D. Malinsky, R. Bhattacharya

Detection of Short-Term Temporal Dependencies in Hawkes Processes with Heterogeneous Background Dynamics 369
Y. Chen, F. Li, A. Schneider, Y. Nevmyvaka, A. Amarasingham, H. Lam

Enhancing Treatment Effect Estimation: A Model Robust Approach Integrating Randomized Experiments and External Controls Using the Double Penalty Integration Estimator 381
Y. Cheng, L. Wu, S. Yang

Adaptivity Complexity for Causal Graph Discovery 391
D. Choo, K. Shiragur

Combinatorial Categorized Bandits with Expert Rankings 403
S.R. Chowdhurry, G. Sinha, N. Natarajan, A. Sharma

Parity Calibration 413
Y. Chung, A. Rumack, C. Gupta

Finite-Sample Guarantees for Nash Q-Learning with Linear Function Approximation 424
P. Cisneros-Velarde, S. Koyejo

Establishing Markov Equivalence in Cyclic Directed Graphs 433
T. Claassen, J.M. Mooij

Expectation Consistency for Calibration of Neural Networks 443
L. Clarte, B. Lourelro, F. Krzakala, L. Zdeborova

Human-In-The-Loop Mixup 454
K.M. Collins, U. Bhatt, W. Liu, V. Piratla, I. Sucholutsky, B. Love, A. Weller

Learning to Reason About Contextual Knowledge for Planning Under Uncertainty 465
C. Cui, S. Amiri, Y. Ding, X. Zhan, S. Zhang

Blackbox Optimization of Unimodal Functions 476
A. Cutkosky, A. Das, W. Kong, C. Lee, R. Sen

Conditional Abstraction Trees for Sample-Efficient Reinforcement Learning 485
M. Dadvar, R.K. Nayyar, S. Srivastava

Improvable Gap Balancing for Multi-Task Learning 496
Y. Dai, N. Fei, Z. Lu

CrysMMNet: Multimodal Representation for Crystal Property Prediction 507
K. Das, P. Goyal, S.C. Lee, S. Bhattacharjee, N. Ganguly

Dirichlet Proportions Model for Hierarchically Coherent Probabilistic Forecasting 518
A. Das, W. Kong, B. Paria, R. Sen

Neural Probabilistic Logic Programming in Discrete-Continuous Domains 529
L. De Smet, P.Z. Dos Martires, R. Manhaeve, G. Marra, A. Kimmig, L. De Raedt

Studying the Effect of GNN Spatial Convolutions on the Embedding Space's Geometry 539
C. Donnat, S.W. Jeong

Deep Gaussian Mixture Ensembles 549
Y. El-Laham, N. Dalmasso, E. Fons, S. Vyetrenko

Personalized Federated Domain Adaptation for Item-To-Item Recommendation 560
Z. Fan, H. Ding, A. Deoras, T.N. Hoang

Generating Synthetic Datasets by Interpolating Along Generalized Geodesics 571
J. Fan, D. Alvarez-Melis

Logit-Based Ensemble Distribution Distillation for Robust Autoregressive Sequence Uncertainties 582
Y. Fathullah, G. Xia, M.J.F. Gales

On the Role of Model Uncertainties in Bayesian Optimization 592
J. Foldager, M. Jordahn, L.K. Hansen, M.R. Andersen

Does Momentum Help in Stochastic Optimization? a Sample Complexity Analysis 602
S. Ganesh, R. Deb, G. Thoppe, A. Budhiraja

Time-Conditioned Generative Modeling of Object-Centric Representations for Video Decomposition and Prediction 613
C. Gao, B. Li

Information Theoretic Clustering Via Divergence Maximization Among Clusters 624
S. Garg, M. Dalirrooyfard, A. Schneider, Y. Adler, Y. Nevmyvaka, Y. Chen, F. Li, G. Cecchi

In- Or Out-Of-Distribution Detection Via Dual Divergence Estimation 635
S. Garg, S. Dutta, M. Dalirrooyfard, A. Schneider, Y. Nevmyvaka

A Data-Driven State Aggregation Approach for Dynamic Discrete Choice Models 647
S. Geng, H. Nassif, C.A. Manzanares

Quasi-Bayesian Nonparametric Density Estimation Via Autoregressive Predictive Updates 658
S. Ghalebikesabi, C. Holmes, E. Fong, B. Lehmann

PART 2

Copula-Based Deep Survival Models for Dependent Censoring 669
A.H.G. Foomani, M. Cooper, R. Greiner, R.G. Krishnan

Probabilistically Robust Conformal Prediction 681
S. Ghosh, Y. Shi, T. Belkhouja, Y. Yan, J.R. Doppa, B. Jones

Fast and Scalable Score-Based Kernel Calibration Tests 691
P. Glaser, D. Widmann, F. Lindsten, A. Gretton

Vacant Holes for Unsupervised Detection of the Outliers in Compact Latent Representation 701
M. Glazunov, A. Zarras

Piecewise Deterministic Markov Processes for Bayesian Neural Networks 712
E. Goan, D. Perrin, K. Mengersen, C. Fookes

A ScalableWalsh-Hadamard Regularizer to Overcome the Low-Degree Spectral Bias of Neural Networks 723
A. Gorji, A. Amrollahi, A. Krause

Stochastic Graphical Bandits with Heavy-Tailed Rewards 734
Y. Gou, J. Yi, L. Zhang

Concurrent Misclassification and Out-Of-Distribution Detection for Semantic Segmentation Via Energy-Based Normalizing Flow 745
D. Gudovskiy, T. Okuno, Y. Nakata

Functional Causal Bayesian Optimization 756
L. Gultchin, V. Aglietti, A. Bellot, S. Chiappa

Causal Discovery for Time Series from Multiple Datasets with Latent Contexts 766
W. Gunther, U. Ninad, J. Runge

Sufficient Identification Conditions and Semiparametric Estimation Under Missing Not at Random Mechanisms 777
A. Guo, J. Zhao, R. Nabi

Interpretable Differencing of Machine Learning Models 788
S. Haldar, D. Saha, D. Wei, R. Nair, E.M. Daly

Differentiable User Models 798
A. Hamalainen, M.M. Celikok, S. Kaski

On the Convergence of Continual Learning with Adaptive Methods 809
S. Han, Y. Kim, T. Cho, J. Lee

Revisiting Bayesian Network Learning with Small Vertex Cover 819
J. Harviainen, M. Koivisto

On Inference and Learning with Probabilistic Generating Circuits 829
J. Harviainen, V.P. Ramaswamy, M. Koivisto

Inference and Sampling of Point Processes from Diffusion Excursions 839
A. Hasan, Y. Chen, Y. Ng, M. Abdelghani, A. Schneider, V. Tarokh

Loosely Consistent Emphatic Temporal-Difference Learning 849
J. He, F. Che, Y. Wan, A.R. Mahmood

Scalable and Robust Tensor Ring Decomposition for Large-Scale Data 860
Y. He, G.K. Atia

Massively Parallel ReweightedWake-Sleep 870
T. Heap, G. Leech, L. Aitchison

Increasing Effect Sizes of Pairwise Conditional Independence Tests Between Random Vectors 879
T. Hochsprung, J. Wahl, A. Gerhardus, U. Ninad, J. Runge

Optimistic Thompson Sampling-Based Algorithms for Episodic Reinforcement Learning 890
B. Hu, T.H. Zhang, N. Hegde, M. Schmidt

ASTRA: Understanding the Practical Impact of Robustness for Probabilistic Programs 900
Z. Huang, S. Dutta, S. Misailovic

A Near-Optimal High-Probability Swap-Regret Upper Bound for Multi-Agent Bandits in Unknown General-Sum Games 911
Z. Huang, J. Pan

Posterior Sampling-Based Online Learning for the Stochastic Shortest Path Model 922
M. Jafarnia-Jahromi, L. Chen, R. Jain, H. Luo

Investigating a Generalization of Probabilistic Material Implication and Bayesian Conditionals 932
M. Jahn, M. Scheutz

Robust Statistical Comparison of Random Variables with Locally Varying Scale of Measurement 941
C. Jansen, G. Schollmeyer, H. Blocher, J. Rodemann, T. Augustin

Noisy Adversarial Representation Learning for Effective and Efficient Image Obfuscation 953
J. Jeong, M. Cho, P. Benz, T.H. Kim

Incentivising Diffusion While Preserving Differential Privacy 963
F. Jia, M. Zhang, J. Liu, B. Khoussainov

Content Sharing Design for SocialWelfare in Networked Disclosure Game 973
F. Jia, C. Qiu, S. Rajtmajer, A. Squicciarini

Multi-View Graph Contrastive Learning for Solving Vehicle Routing Problems 984
Y. Jiang, Z. Cao, Y. Wu, J. Zhang

Bayesian Inference for Vertex-Series-Parallel Partial Orders 995
C. Jiang, G.K. Nicholls, J.E. Lee

Nyström M-Hilbert-Schmidt Independence Criterion 1005
F. Kalinke, Z. Szabo

Causal Discovery with Hidden Confounders Using the Algorithmic Markov Condition 1016
D. Kaltenpoth, J. Vreeken

Heavy-Tailed Linear Bandit with Huber Regression 1027
M. Kang, G.S. Kim

Fed-LAMB: Layer-Wise and Dimension-Wise Locally Adaptive Federated Learning 1037
B. Karimi, P. Li, X. Li

Assessing the Impact of Context Inference Error and Partial Observability on RL Methods for Just-In-Time Adaptive Interventions 1047
K. Karine, P. Klasnja, S.A. Murphy, B.M. Marlin

How to Use Dropout Correctly on Residual Networks with Batch Normalization 1058
B.J. Kim, H. Choi, H. Jang, D. Lee, S.W. Kim

Phase-Shifted Adversarial Training 1068
Y. Kim, S. Kim, I. Seo, B. Shin

On Identifiability of Conditional Causal Effects 1078
Y. Kivva, J. Etesami, N. Kiyavash

Causal Effect Estimation from Observational and Interventional Data Through MatrixWeighted Linear Estimators 1087
K.R. Kladny, J. von Kugelgen, B. Scholkopf, M. Muehlebach

Universal Graph Contrastive Learning with a Novel Laplacian Perturbation 1098
T. Ko, Y. Choi, C.K. Kim

Molecule Design by Latent Space Energy-Based Modeling and Gradual Distribution Shifting 1109
D. Kong, B. Pang, T. Han, Y.N. Wu

Reward-Machine-Guided, Self-Paced Reinforcement Learning 1121
C. Koprulu, U. Topcu

Risk-Aware Curriculum Generation for Heavy-Tailed Task Distributions 1132
C. Koprulu, T.D. Simao, N. Jansen, U. Topcu

Differentially Private Synthetic Data Using KD-Trees 1143
E. Kreacic, N. Nouri, V.K. Potluru, T. Balch, M. Veloso

Optimal Budget Allocation for Crowdsourcing Labels for Graphs 1154
A. Kulkarni, M. Chakraborty, S. Xie, Q. Li

Fixed-Budget Best-Arm Identification with Heterogeneous Reward Variances 1164
A. Lalitha, K. Kalantari, Y. Ma, A. Deoras, B. Kveton

Variable Importance Matching for Causal Inference 1174
Q. Lanners, H. Parikh, A. Volfovsky, C. Rudin, D. Page

Towards Better Certified Segmentation Via Diffusion Models 1185
O. Laousy, A. Araujo, G. Chassagnon, M.P. Revel, S. Garg, F. Khorrami, M. Vakalopoulou

Finding Invariant Predictors Efficiently Via Causal Structure 1196
K. Lee, M.M. Rahman, M. Kocaoglu

When Are Post-Hoc Conceptual Explanations Identifiable? 1207
T. Leemann, M. Kirchhof, Y. Rong, E. Kasneci, G. Kasneci

Memory Mechanism for Unsupervised Anomaly Detection 1219
J. Li, Y. Chen, Y. Xing

Nonconvex Stochastic Scaled Gradient Descent and Generalized Eigenvector Problems 1230
C.J. Li, M.I. Jordan

Gaussian Process Surrogate Models for Neural Networks 1241
M.Y. Li, E. Grant, T.L. Griffiths

CUE: An Uncertainty Interpretation Framework for Text Classifiers Built on Pre-Trained Language Models 1253
J. Li, Z. Sun, B. Liang, L. Gui, Y. He

BISCUIT: Causal Representation Learning from Binary Interactions 1263
P. Lippe, S. Magliacane, S. Lowe, Y.M. Asano, T. Cohen, E. Gavves

Accelerating Voting by Quantum Computation 1274
A. Liu, Q. Han, L. Xia, N. Yu

Residual-Based Error Bound for Physics-Informed Neural Networks 1284
S. Liu, X. Huang, P. Protopapas

No-Regret Linear Bandits Beyond Realizability 1294
C. Liu, M. Yin, Y.X. Wang

Benefits of Monotonicity in Safe Exploration with Gaussian Processes 1304
A. Losalka, J. Scarlett

Practical Privacy-Preserving Gaussian Process Regression Via Secret Sharing 1315
J. Luo, Y. Zhang, J. Zhang, S. Qin, H. Wang, Y. Yu, Z. Xu

PART 3

DeepGD3: Unknown-Aware Deep Generative/Discriminative Hybrid Defect 1326
C.W. Ma, Y. Liu

Federated Learning of Models Pre-Trained on Different Features with Consensus Graphs..... 1336
T. Ma, T.N. Hoang, J. Chen

Random Reshuffling with Variance Reduction: New Analysis and Better Rates 1347
G. Malinovsky, A. Sailanbayev, P. Richtarik

The Shrinkage-Delinkage Trade-Off: An Analysis of Factorized Gaussian Approximations for Variational Inference 1358
C.C. Margossian, L.K. Saul

Partial Identification of Dose Responses with Hidden Confounders..... 1368
M.G. Marmarelis, E. Haddad, A. Jesson, N. Jahanshad, A. Galstyan, G.V. Steeg

Knowledge Intensive Learning of Cutset Networks..... 1380
S. Mathur, V. Gogate, S. Natarajan

TCE: A Test-Based Approach to Measuring Calibration Error..... 1390
T. Matsubara, N. Tax, R. Mudd, I. Guy

Causal Information Splitting: Engineering Proxy Features for Robustness to Distribution Shifts..... 1401
B. Mazaheri, A. Mastakouri, D. Janzing, M. Hardt

KrADagrad: Kronecker Approximation-Domination Gradient Preconditioned Stochastic Optimization..... 1412
J. Mei, A. Moreno, L. Walters

On the Relation Between Policy Improvement and Off-Policy Minimum-Variance Policy Evaluation..... 1423
A.M. Metelli, S. Meta, M. Restelli

On Minimizing the Impact of Dataset Shifts on Actionable Explanations 1434
A.P. Meyer, D. Ley, S. Srinivas, H. Lakkaraju

Adaptive Conditional Quantile Neural Processes..... 1445
P. Mohseni, N. Duffield, B. Mallick, A. Hasanzadeh

Composing Efficient, Robust Tests for Policy Selection..... 1456
D. Morrill, T.J. Walsh, D. Hernandez, P.R. Wurman, P. Stone

On Testability and Goodness of Fit Tests in Missing Data Models 1467
R. Nabi, R. Bhattacharya

Active Metric Learning and Classification Using Similarity Queries 1478
N. Nadagouda, A. Xu, M.A. Davenport

Two-Phase Attacks in Security Games..... 1489
A. Nagorko, P. Ciosmak, T. Michalak

Keep-Alive Caching for the Hawkes Process..... 1499
S. Narayana, I.A. Kash

Simple Transferability Estimation for Regression Tasks 1510
C.N. Nguyen, P. Tran, L.S.T. Ho, V. Dinh, A.T. Tran, T. Hassner, C.V. Nguyen

Probabilistic Multi-Dimensional Classification 1522
V.L. Nguyen, Y. Yang, C. de Campos

Efficient Failure Pattern Identification of Predictive Algorithms 1534
B. Nguyen, V.A. Nguyen

Size-Constrained k-Submodular Maximization in Near-Linear Time 1545
G. Nie, Y. Zhu, Y.Y. Nadew, S. Basu, A. Paven, C.J. Quinn

An Improved Variational Approximate Posterior for the Deep Wishart Process 1555
S.W. Ober, B. Anson, E. Milsom, L. Aitchison

Incentivizing Honest Performative Predictions with Proper Scoring Rules 1564
C. Oesterheld, J. Treutlein, E. Cooper, R. Hudson

Graph Classification Gaussian Processes Via Spectral Features 1575
F.L. Opolka, Y.C. Zhi, P. Lio, X. Dong

Approximate Thompson Sampling Via Epistemic Neural Networks 1586
I. Osband, Z. Wen, S.M. Asghari, V. Dwaracherla, M. Ibrahimi, X. Lu, B. Van Roy

Structure-Aware Robustness Certificates for Graph Classification 1596
P. Osselin, H. Kenlay, X. Dong

Bayesian Numerical Integration with Neural Networks 1606
K. Ott, M. Tiemann, P. Hennig, F.X. Briol

Maximizing Submodular Functions Under Submodular Constraints 1618
M.R. Padmanabhan, Y. Zhu, S. Basu, A. Pavan

Stochastic Generative Flow Networks 1628
L. Pan, D. Zhang, M. Jain, L. Huang, Y. Bengio

Multi-View Independent Component Analysis with Shared and Individual Sources 1639
T. Pandeva, P. Forre

Copula for Instance-Wise Feature Selection and Ranking 1651
H. Peng, G. Fang, P. Li

Boosting AND/OR-Based Computational Protein Design: Dynamic Heuristics and Generalizable UFO 1662
B. Pezeshki, R. Marinescu, A. Ihler, R. Dechter

Exact Count of Boundary Pieces of ReLU Classifiers: Towards the Proper Complexity Measure for Classification 1673
P. Piwek, A. Klukowski, T. Hu

Split, Count, and Share: A Differentially Private Set Intersection Cardinality Estimation Protocol 1684
M. Purcell, Y. Li, K.S. Ng

JANA: Jointly Amortized Neural Approximation of Complex Bayesian Models 1695
S.T. Radev, M. Schmitt, V. Pratz, U. Picchini, U. Kothe, P.C. Burkner

USIM-DAL: Uncertainty-Aware Statistical Image Modeling-Based Dense Active Learning for Super-Resolution 1707
V. Rangnekar, U. Upadhyay, Z. Akata, B. Banerjee

Contrastive Learning for Supervised Graph Matching 1718
G. Ratnayaka, Q. Wang, Y. Li

Validation of Composite Systems by Discrepancy Propagation 1730
D. Reeb, K. Patel, K. Barsim, M. Schiegg, S. Gerwinn

Inference for Probabilistic Dependency Graphs 1741
O.E. Richardson, J.Y. Halpern, C. De Sa

The Past Does Matter: Correlation of Subsequent States in Trajectory Predictions of Gaussian Process Models 1752
S. Ridderbusch, S. Ober-Blobaum, P. Goulart

Approximately Bayes-Optimal Pseudo-Label Selection 1762
J. Rodemann, J. Goschenhofer, E. Dorigatti, T. Nagler, T. Augustin

Hallucinated Adversarial Control for Conservative Offline Policy Evaluation 1774
J. Rothfuss, B. Sukhija, T. Birchler, P. Kassrai, A. Krause

Semi-Supervised Learning of Partial Differential Operators and Dynamical Flows 1785
M. Rotman, A. Dekel, R.I. Ber, L. Wolf, Y. Oz

Is the Volume of a Credal Set a Good Measure for Epistemic Uncertainty? 1795
Y. Sale, M. Caprio, E. Hullermeier

Heteroskedastic Geospatial Tracking with Distributed Camera Networks 1805
C. Samplawski, S. Fang, Z. Wang, D. Ganesan, M. Srivastava, B.M. Marlin

Online Heavy-Tailed Change-Point Detection 1815
A. Sankararaman, B. Narayanaswamy

Learning Good Interventions in Causal Graphs Via Covering 1827
A. Sawarni, R. Madhavan, G. Sinha, S. Barman

Local Message Passing on Frustrated Systems 1837
L. Schmid, J. Brenk, L. Schmalen

Lifelong Bandit Optimization: No Prior and No Regret 1847
F. Schur, P. Kassraie, J. Rothfuss, A. Krause

Learning from Low Rank Tensor Data: A Random Tensor Theory Perspective 1858
M.E.A. Seddik, M. Tiomoko, A. Decurnige, M. Panov, M. Guillaud

MDPose: Real-Time Multi-Person Pose Estimation Via Mixture Density Model 1868
S. Seo, J. Yoo, J. Hwang, N. Kwak

Mnemonist: Locating Model Parameters that Memorize Training Examples 1879
A.S. Shamsabadi, J. Hayes, B. Balle, A. Weller

Implicit Training of Inference Network Models for Structured Prediction 1889
S. Shankar

Efficiently Learning the Graph for Semi-Supervised Learning 1900
D. Sharma, M. Jones

Counting Background Knowledge Consistent Markov Equivalent Directed Acyclic Graphs 1911
V.S. Sharma

SymNet 3.0: Exploiting Long-Range Influences in Learning Generalized Neural Policies for Relational MDPs 1921
V. Sharma, D. Arora, Mausam, P. Singla

Risk-Limiting Financial Audits viaWeighted Sampling Without Replacement 1932
S. Shekhar, Z. Xu, Z. Lipton, P. Liang, A. Ramdas

Learning Nonlinear Causal Effects Via Kernel Anchor Regression 1942
W. Shi, W. Xu

A Bayesian Approach for Bandit Online Optimization with Switching Cost 1953
Z. Shi, J. Tan, F.F. Li

Probabilistic Flow Circuits: Towards Unified Deep Models for Tractable Probabilistic Inference 1964
S. Sidheekh, K. Kersting, S. Natarajan

On the Limitations of Markovian Rewards to Express Multi-Objective, Risk-Sensitive, and Modal Tasks 1974
J. Skalse, A. Abate

PART 4

Aligned Diffusion Schrödinger Bridges 1985
V.R. Somnath, M.R. Martinez, M. Pariset, A. Krause, Y.P. Hsieh, C. Bunne

ViBid: Linear Vision Transformer with Bidirectional Normalization 1996
J. Song, H.C. Lee

Fast Heterogeneous Federated Learning with Hybrid Client Selection 2006
D. Song, G. Shen, D. Gao, I. Yang, X. Zhou, S. Pan, W. Lou, F. Zhou

Solving Multi-Model MDPs by Coordinate Ascent and Dynamic Programming 2016
X. Su, M. Petrik

Differentially Private Stochastic Convex Optimization in (Non)-Euclidean Space Revisited 2026
J. Su, C. Zhao, D. Wang

On the Informativeness of Supervision Signals 2036
I. Sucholutsky, R.M. Battleday, K.M. Collins, R. Marjieh, J.C. Peterson, P. Singh, U. Bhatt, N. Jacoby, A. Weller, T.L. Griffiths

Meta-Learning Control Variates: Variance Reduction with Limited Data 2047
Z. Sun, C.J. Oates, F.X. Briol

Pandering in a (Flexible) Representative Democracy 2058
X. Sun, J. Masur, B. Abramowitz, N. Mattei, Z. Zheng

Locally Regularized Sparse Graph by Fast Proximal Gradient Descent 2069
D. Sun, Y. Yang

Why Out-Of-Distribution Detection Experiments Are Not Reliable - Subtle Experimental Details Muddle the OOD Detector Rankings 2078
K. Szyc, T. Walkowiak, H. Maciejewski

Exploiting Inferential Structure in Neural Processes 2089
D. Tailor, M.E. Khan, E. Nalisnick

Two-Stage Kernel Bayesian Optimization in High Dimensions 2099
J. Tan, N. Nayman

Low-Rank Matrix Recovery with Unknown Correspondence 2111
Z. Tang, T.H. Chang, X. Ye, H. Zha

Fairness-Aware Class Imbalanced Learning on Multiple Subgroups 2123
D.A. Tarzanagh, B. Hou, B. Tong, Q. Long, L. Shen

SPDF: Sparse Pre-Training and Dense Fine-Tuning for Large Language Models 2134
V. Thangarasa, A. Gupta, W. Marshall, T. Li, K. Leong, D. DeCoste, S. Lie, S. Saxena

Bandits with Costly Reward Observations 2147
A.D. Tucker, C. Biddulph, C. Wang, T. Joachims

Probabilistic Circuits that Know What They Don't Know 2157
F. Ventola, S. Braun, Z. Yu, M. Mundt, K. Kersting

A Policy Gradient Approach for Optimization of Smooth Risk Measures 2168
N. Vijayan, L.A. Prashanth

Birds of an Odd Feather: Guaranteed Out-Of-Distribution (OOD) Novel Category Detection 2179
Y. Wald, S. Saria

Exploration for Free: How Does Reward Heterogeneity Improve Regret in Cooperative Multi-Agent Bandits? 2192
X. Wang, L. Yang, Y.Z.J. Chen, X. Liu, M. Hajiesmaili, D. Towsley, J.C.S. Lui

Efficient Privacy-Preserving Stochastic Nonconvex Optimization 2203
L. Wang, B. Jayaraman, D. Evans, Q. Gu

Diversity-Enhanced Probabilistic Ensemble for Uncertainty Estimation 2214
H. Wang, Q. Ji

A Trajectory isWorth Three Sentences: Multimodal Transformer for Offline Reinforcement Learning 2226
Y. Wang, M. Xu, L. Shi, Y. Chi

Robust Distillation forWorst-Class Performance: On the Interplay Between Teacher and Student Objectives 2237
S. Wang, H. Narasimhan, Y. Zhou, S. Hooker, M. Lukasik, A.K. Menon

A Constrained Bayesian Approach to Out-Of-Distribution Prediction 2248
Z. Wang, B. Yuan, J. Lu, B. Ding, Y. Shao, Q. Wu, J. Zhu

On the Role of Generalization in Transferability of Adversarial Examples 2259
Y. Wang, F. Farnia

Bidirectional Attention as a Mixture of Continuous Word Experts 2271
K.C. Wibisono, Y. Wang

Quantifying Aleatoric and Epistemic Uncertainty in Machine Learning: Are Conditional Entropy and Mutual Information Appropriate Measures? 2282
L. Wimmer, Y. Sale, P. Hofman, B. Bischl, E. Hullermeier

Learning to Invert: Simple Adaptive Attacks for Gradient Inversion in Federated Learning 2293
R. Wu, X. Chen, C. Guo, K.Q. Weinberger

Uniform-PAC Guarantees for Model-Based RL with Bounded Eluder Dimension 2304
Y. Wu, J. He, Q. Gu

Robust Quickest Change Detection for Unnormalized Models 2314
S. Wu, E. Diao, T. Banerjee, J. Ding, V. Tarokh

A One-Sample Decentralized Proximal Algorithm for Non-Convex Stochastic Composite Optimization 2324
T. Xiao, X. Chen, K. Balasubramanian, S. Ghadimi

Two-Stage Holistic and Contrastive Explanation of Image Classification 2335
W. Xie, X.H. Li, Z. Lin, L.K.M. Poon, C.C. Cao, N.L. Zhang

Conformal Risk Control for Ordinal Classification 2346
Y. Xu, W. Guo, Z. Wei

E(2)-Equivariant Vision Transformer 2356
R. Xu, K. Yang, K. Liu, F. He

Provably Efficient Adversarial Imitation Learning with Unknown Transitions 2367
T. Xu, Z. Li, Y. Yu, Z.Q. Luo

Pessimistic Model Selection for Offline Deep Reinforcement Learning 2379
C.H. Yang, Z. Qi, Y. Cui, P.Y. Chen

Mixture of Normalizing Flows for European Option Pricing 2390
Y. Yang, T.M. Hospedales

Multi-Modal Differentiable Unsupervised Feature Selection 2400
J. Yang, O. Lindenbaum, Y. Kluger, A. Jaffe

MMEL: A Joint Learning Framework for Multi-Mention Entity Linking 2411
C. Yang, B. He, Y. Wu, C. Xing, L. He, C. Ma

Mitigating Transformer Overconfidence Via Lipschitz Regularization 2422
W. Ye, Y. Ma, X. Cao, K. Tang

Towards Physically Reliable Molecular Representation Learning 2433
S. Yi, Y. Cho, J. Sul, S.W. Ko, S.K. Kim, J. Choo, H. Yoon, J. Lee

Monte-Carlo Search for an Equilibrium in Dec-POMDPs 2444
Y. You, V. Thomas, F. Colas, O. Buffet

Online Estimation of Similarity Matrices with Incomplete Data 2454
F. Yu, Y. Zeng, J. Mao, W. Li

Fast Teammate Adaptation in the Presence of Sudden Policy Change 2465
Z. Zhang, L. Yuan, L. Li, K. Xue, C. Jia, C. Guan, C. Qian, Y. Yu

Energy-Based Predictive Representations for Partially Observed Reinforcement Learning 2477
T. Zhang, T. Ren, C. Xiao, W. Xiao, J.E. Gonzalez, D. Schuurmans, B. Dai

Provably Efficient Representation Selection in Low-Rank Markov Decision Processes: From Online to Offline RL 2488
W. Zhang, J. He, D. Zhou, A. Zhang, Q. Gu

Graph Self-Supervised Learning Via Proximity Divergence Minimization 2498
T. Zhang, Z. Dai, Z. Xu, A. Shrivastava

Greed is Good: Correspondence Recovery for Unlabeled Linear Regression 2509
H. Zhang, P. Li

Conditional Counterfactual Causal Effect for Individual Attribution 2519
R. Zhao, L. Zhang, S. Zhu, Z. Lu, Z. Dong, C. Zhang, J. Xu, Z. Geng, Y. He

Conditionally Optimistic Exploration for Cooperative Deep Multi-Agent Reinforcement Learning 2529
X. Zhao, Y. Pan, C. Xiao, S. Chandar, J. Rajendran

RDM-DC: Poisoning Resilient Dataset Condensation with Robust Distribution Matching 2541
T. Zheng, B. Li

Learning Robust Representation for Reinforcement Learning with Distractions by Reward Sequence Prediction 2551
Q. Zhou, J. Wang, Q. Liu, Y. Kuang, W. Zhou, H. Li

Convergence Rates for Localized Actor-Critic in Networked Markov Potential Games 2563
Z. Zhou, Z. Chen, Y. Lin, A. Wierman

AUC Maximization in Imbalanced Lifelong Learning 2574
X. Zhu, J. Hao, Y. Guo, M. Liu

“Private Prediction Strikes Back!” Private Kernelized Nearest Neighbors with Individual Rényi Filter 2586
Y. Zhu, X. Zhao, C. Guo, Y.X. Wang

MixupE: Understanding and Improving Mixup from Directional Derivative Perspective 2597
Y. Zou, V. Verma, S. Mittal, W.H. Tang, H. Pham, J. Kannala, Y. Bengio, A. Solin, K. Kawaguchi

Regularized Online DR-Submodular Optimization 2608
P. Zuo, Y. Wang, S. Tang

KrADagrad: Kronecker Approximation-Domination Gradient Preconditioned Stochastic Optimization 2618
J. Mei, A. Moreno, L. Walters

KrADagrad: Kronecker Approximation-Domination Gradient Preconditioned Stochastic Optimization (Supplementary Material) 2629
J. Mei, A. Moreno, L. Walters

Author Index