

35th Conference on Uncertainty in Artificial Intelligence (UAI 2019)

Tel Aviv, Israel
22 - 25 July 2019

Volume 1 of 3

ISBN: 978-1-5108-9156-2

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© (2019) by Association For Uncertainty in Artificial Intelligence (AUAI)
All rights reserved.

Printed by Curran Associates, Inc. (2019)

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

VOLUME 1

PERSONALIZED PEER TRUTH SERUM FOR ELICITING MULTI-ATTRIBUTE PERSONAL DATA	1
<i>Naman Goel, Boi Faltings</i>	
CONDITIONAL EXPECTATION PROPAGATION	11
<i>Zheng Wang, Shandian Zhe</i>	
A SPARSE REPRESENTATION-BASED APPROACH TO LINEAR REGRESSION WITH PARTIALLY SHUFFLED LABELS	28
<i>Martin Slawski, Mostafa Rahmani, Ping Li</i>	
ON FAST CONVERGENCE OF PROXIMAL ALGORITHMS FOR SQRT-LASSO OPTIMIZATION: DON'T WORRY ABOUT ITS NONSMOOTH LOSS FUNCTION	39
<i>Xingguo Li, Haoming Jiang, Jarvis Haupt, Raman Arora, Han Liu, Mingyi Hong, Tuo Zhao</i>	
CORRELATED LEARNING FOR AGGREGATION SYSTEMS	80
<i>Tanvi Verma, Pradeep Varakantham</i>	
CAUSAL CALCULUS IN THE PRESENCE OF CYCLES, LATENT CONFOUNDERS AND SELECTION BIAS	91
<i>Patrick Forré, Joris M. Mooij</i>	
VARIATIONAL REGRET BOUNDS FOR REINFORCEMENT LEARNING	108
<i>Ronald Ortner, Pratik Gajane, Peter Auer</i>	
RECOMMENDATION FROM RAW DATA WITH ADAPTIVE COMPOUND POISSON FACTORIZATION	118
<i>Olivier Gouvert, Thomas Oberlin, Cédric Févotte</i>	
ONE-SHOT INFERENCE IN MARKOV RANDOM FIELDS	131
<i>Hao Xiong, Yuanzhen Guo, Yibo Yang, Nicholas Ruozi</i>	
TRULY PROXIMAL POLICY OPTIMIZATION	147
<i>Yuhui Wang, Hao He, Xiaoyang Tan</i>	
LEARNING FACTORED MARKOV DECISION PROCESSES WITH UNAWARENESS	157
<i>Craig Innes, Alex Lascarides</i>	
EXPRESSIVE PRIORS IN BAYESIAN NEURAL NETWORKS: KERNEL COMBINATIONS AND PERIODIC FUNCTIONS	173
<i>Tim Pearce, Russell Tsuchida, Mohamed Zaki, Alexandra Brintrup, Andy Neely</i>	
COUNTDOWN REGRESSION: SHARP AND CALIBRATED SURVIVAL PREDICTIONS	184
<i>Anand Avati, Tony Duan, Sharon Zhou, Ken Jung, Nigam H. Shah, Andrew Ng</i>	
REDUCING EXPLORATION OF DYING ARMS IN MORTAL BANDITS	199
<i>Stefano Tracà, Weiyu Yan, Cynthia Rudin</i>	
COMPARING EM WITH GD IN MIXTURE MODELS OF TWO COMPONENTS	220
<i>Guojun Zhang, Pascal Poupart, George Trimonias</i>	
EFFICIENT SEARCH-BASED WEIGHTED MODEL INTEGRATION	235
<i>Zhe Zeng, Guy Van den Broeck</i>	
CAUSAL DISCOVERY WITH GENERAL NON-LINEAR RELATIONSHIPS USING NON-LINEAR ICA	249
<i>Ricardo Pio Monti, Kun Zhang, Aapo Hyvarinen</i>	
BUBBLERANK: SAFE ONLINE LEARNING TO RE-RANK VIA IMPLICIT CLICK FEEDBACK	270
<i>Chang Li, Branislav Kveton, Tor Lattimore, Ilya Markov, Maarten de Rijke, Csaba Szepesvari, Masrour Zoghi</i>	
COORDINATING USERS OF SHARED FACILITIES VIA DATA-DRIVEN PREDICTIVE ASSISTANTS AND GAME THEORY	286
<i>Philipp Geiger, Michel Besserve, Justus Winkelmann, Claudius Proissl, Bernhard Schoelkopf</i>	
THE INCOMPLETE ROSETTA STONE PROBLEM: IDENTIFIABILITY RESULTS FOR MULTI-VIEW NONLINEAR ICA	296
<i>Luigi Gresele, Paul Rubenstein, Arash Mehrjou, Francesco Locatello, Bernhard Schoelkopf</i>	
RANDOM CLIQUE COVERS FOR GRAPHS WITH LOCAL DENSITY AND GLOBAL SPARSITY	314
<i>Sinead A. Williamson, Mauricio Tec</i>	
RANDOMIZED ITERATIVE ALGORITHMS FOR FISHER DISCRIMINANT ANALYSIS	325
<i>Agniva Chowdhury, Jiasen Yang, Petros Drineas</i>	

DYNAMIC TRIP-VEHICLE DISPATCH WITH SCHEDULED AND ON-DEMAND REQUESTS	344
<i>Taoan Huang, Bohui Fang, Xiaohui Bei, Fei Fang</i>	
FALL OF EMPIRES: BREAKING BYZANTINE-TOLERANT SGD BY INNER PRODUCT MANIPULATION	359
<i>Cong Xie, Oluwasanmi Koyejo, Indranil Gupta</i>	
ADAPTIVE HASHING FOR MODEL COUNTING	369
<i>Jonathan Kuck, Tri Dao, Shenjia Zhao, Burak Bartan, Ashish Sabharwal, Stefano Ermon</i>	
TOWARDS A BETTER UNDERSTANDING AND REGULARIZATION OF GAN TRAINING DYNAMICS	381
<i>Weili Nie, Ankit Patel</i>	
DOMAIN GENERALIZATION VIA MULTIDOMAIN DISCRIMINANT ANALYSIS	406
<i>Shoubo Hu, Kun Zhang, Zhitang Chen, Laiwan Chan</i>	
EFFICIENT PLANNING UNDER UNCERTAINTY WITH INCREMENTAL REFINEMENT	429
<i>Juan Carlos Saborío, Joachim Hertzberg</i>	
CUBIC REGULARIZATION WITH MOMENTUM FOR NONCONVEX OPTIMIZATION	439
<i>Zhe Wang, Yi Zhou, Yingbin Liang, Guanghui Lan</i>	
STABILITY OF LINEAR STRUCTURAL EQUATION MODELS OF CAUSAL INFERENCE	449
<i>Karthik Abinav Sankararaman, Anand Louis, Navin Goyal</i>	
RANDOM SUM-PRODUCT NETWORKS: A SIMPLE AND EFFECTIVE APPROACH TO PROBABILISTIC DEEP LEARNING	460
<i>Robert Peharz, Antonio Vergari, Karl Stelzner, Alejandro Molina, Martin Trapp, Xiaoting Shao, Kristian Kersting, Zoubin Ghahramani</i>	
TOWARDS ROBUST RELATIONAL CAUSAL DISCOVERY	478
<i>Sanghack Lee, Vasant Honavar</i>	
THE ROLE OF MEMORY IN STOCHASTIC OPTIMIZATION	493
<i>Antonio Orvieto, Jonas Kohler, Aurelien Lucchi</i>	
RANDOM SEARCH AND REPRODUCIBILITY FOR NEURAL ARCHITECTURE SEARCH	529
<i>Liam Li, Ameet Talwalkar</i>	
JOINT NONPARAMETRIC PRECISION MATRIX ESTIMATION WITH CONFOUNDING	546
<i>Sinong Geng, Mladen Kolar, Oluwasanmi Koyejo</i>	
GENERAL IDENTIFIABILITY WITH ARBITRARY SURROGATE EXPERIMENTS	557
<i>Sanghack Lee, Juan D. Correa, Elias Bareinboim</i>	
DIFFERENTIABLE PROBABILISTIC MODELS OF SCIENTIFIC IMAGING WITH THE FOURIER SLICE THEOREM	568
<i>Karen Ullrich, Rianne van den Berg, Marcus A. Brubaker, David Fleet, Max Welling</i>	
APPROXIMATE INFERENCE IN STRUCTURED INSTANCES WITH NOISY CATEGORICAL OBSERVATIONS	581
<i>Alireza Heidari, Ihab F. Ilyas, Theodoros Rekatsinas</i>	

VOLUME 2

RANDOMIZED VALUE FUNCTIONS VIA MULTIPLICATIVE NORMALIZING FLOWS	610
<i>Ahmed Touati, Harsh Satija, Joshua Romoff, Joelle Pineau, Pascal Vincent</i>	
A FAST PROXIMAL POINT METHOD FOR COMPUTING EXACT WASSERSTEIN DISTANCE	621
<i>Yujia Xie, Xiangfeng Wang, Ruijia Wang, Hongyuan Zha</i>	
NEURAL DYNAMICS DISCOVERY VIA GAUSSIAN PROCESS RECURRENT NEURAL NETWORKS	642
<i>Qi She, Anqi Wu</i>	
FISHER-BURES ADVERSARY GRAPH CONVOLUTIONAL NETWORKS	659
<i>Ke Sun, Piotr Koniusz, Zhen Wang</i>	
ϵ-BMC: A BAYESIAN ENSEMBLE APPROACH TO EPSILON-GREEDY EXPLORATION IN MODEL-FREE REINFORCEMENT LEARNING	670
<i>Michael Gimelfarb, Scott Sanner, Chi-Guhn Lee</i>	
PERIODIC KERNEL APPROXIMATION BY INDEX SET FOURIER SERIES FEATURES	681
<i>Anthony Tompkins, Fabio Ramos</i>	
EFFICIENT NEURAL NETWORK VERIFICATION WITH EXACTNESS CHARACTERIZATION	692
<i>Krishnamurthy (Dj) Dvijotham, Robert Stanforth, Sven Gowal, Sven Gowal, Chongli Qin, Soham De, Pushmeet Kohli</i>	
AUGMENTING AND TUNING KNOWLEDGE GRAPH EMBEDDINGS	707
<i>Robert Bamler, Farnood Salehi, Stephan Mandt</i>	

A TIGHTER ANALYSIS OF RANDOMISED POLICY ITERATION	720
<i>Meet Taraviya, Shivaram Kalyanakrishnan</i>	
PERTURBED-HISTORY EXPLORATION IN STOCHASTIC LINEAR BANDITS	734
<i>Branislav Kveton, Csaba Szepesvari, Mohammad Ghavamzadeh, Craig Boutilier</i>	
AN IMPROVED CONVERGENCE ANALYSIS OF STOCHASTIC VARIANCE-REDUCED POLICY GRADIENT	745
<i>Pan Xu, Felicia Gao, Quanquan Gu</i>	
DEEP MIXTURE OF EXPERTS VIA SHALLOW EMBEDDING	756
<i>Xin Wang, Fisher Yu, Lisa Dunlap, Yi-An Ma, Ruth Wang, Azalia Mirhoseini, Trevor Darrell, Joseph E. Gonzalez</i>	
SAMPLING-FREE VARIATIONAL INFERENCE OF BAYESIAN NEURAL NETWORKS BY VARIANCE BACKPROPAGATION	770
<i>Manuel Hausmann, Fred A. Hamprecht, Melih Kandemir</i>	
SLICED SCORE MATCHING: A SCALABLE APPROACH TO DENSITY AND SCORE ESTIMATION	781
<i>Yang Song, Sahaj Garg, Jiaxin Shi, Stefano Ermon</i>	
BEYOND STRUCTURAL CAUSAL MODELS: CAUSAL CONSTRAINTS MODELS	812
<i>Tineke Blom, Stephan Bongers, Joris M. Mooij</i>	
BE GREEDY: HOW CHROMATIC NUMBER MEETS REGRET MINIMIZATION IN GRAPH BANDITS	828
<i>Aadirupa Saha, Shreyas Sheshadri, Chiranjib Bhattacharyya</i>	
APPROXIMATE CAUSAL ABSTRACTIONS	848
<i>Sander Beckers, Frederick Eberhardt, Joseph Y. Halpern</i>	
THE SENSITIVITY OF COUNTERFACTUAL FAIRNESS TO UNMEASURED CONFOUNDING	858
<i>Niki Kilbertus, Philip J. Ball, Matt J. Kusner, Adrian Weller, Ricardo Silva</i>	
BELIEF PROPAGATION: ACCURATE MARGINALS OR ACCURATE PARTITION FUNCTION -- WHERE IS THE DIFFERENCE?	870
<i>Christian Knoll, Franz Pernkopf</i>	
FINDING MINIMAL D-SEPARATORS IN LINEAR TIME AND APPLICATIONS	886
<i>Benito van der Zander, Maciej Liskiewicz</i>	
A BAYESIAN APPROACH TO ROBUST REINFORCEMENT LEARNING	898
<i>Esther Derman, Daniel J. Mankowitz, Timothy A. Mann, Shie Mannor</i>	
ADAPTIVITY AND OPTIMALITY: A UNIVERSAL ALGORITHM FOR ONLINE CONVEX OPTIMIZATION	913
<i>Guanghui Wang, Shiyin Lu, Lijun Zhang</i>	
EVACUATE OR NOT? A POMDP MODEL OF THE DECISION MAKING OF INDIVIDUALS IN HURRICANE EVACUATION ZONES	925
<i>Adithya Raam Sankar, Prashant Doshi, Adam Goodie</i>	
PROBABILISTIC PROGRAMMING FOR BIRTH-DEATH MODELS OF EVOLUTION USING AN ALIVE PARTICLE FILTER WITH DELAYED SAMPLING	937
<i>Jan Kudlicka, Lawrence M. Murray, Fredrik Ronquist, Thomas B. Schön</i>	
VARIATIONAL SPARSE CODING	955
<i>Francesco Tonolini, Bjørn Sand Jensen, Roderick Murray-Smith</i>	
LEARNING WITH NON-CONVEX TRUNCATED LOSSES BY SGD	973
<i>Yi Xu, Shenghuo Zhu, Sen Yang, Chi Zhang, Rong Jin, Tianbao Yang</i>	
ACTIVE MULTI-INFORMATION SOURCE BAYESIAN QUADRATURE	991
<i>Alexandra Gessner, Javier Gonzalez, Maren Mahsereci</i>	
CASCADING LINEAR SUBMODULAR BANDITS: ACCOUNTING FOR POSITION BIAS AND DIVERSITY IN ONLINE LEARNING TO RANK	1006
<i>Gaurush Hiranandani, Harvineet Singh, Prakhar Gupta, Iftikhar Ahamath Burhanuddin, Zheng Wen, Branislav Kveton</i>	
SINKHORN AUTOENCODERS	1024
<i>Giorgio Patrini, Rianne van den Berg, Patrick Forré, Marcello Carioni, Samarth Bhargav, Max Welling, Tim Genewein, Frank Nielsen</i>	
HOW TO EXPLOIT STRUCTURE WHILE SOLVING WEIGHTED MODEL INTEGRATION PROBLEMS	1040
<i>Samuel Kolb, Pedro Zuidberg Dos Martires, Luc De Raedt</i>	
MULTI-CLASS GAUSSIAN PROCESS CLASSIFICATION MADE CONJUGATE: EFFICIENT INFERENCE VIA DATA AUGMENTATION	1052
<i>Théo Galy-Fajou, Florian Wenzel, Christian Donner, Manfred Opper</i>	
A FLEXIBLE FRAMEWORK FOR MULTI-OBJECTIVE BAYESIAN OPTIMIZATION USING RANDOM SCALARIZATIONS	1065
<i>Biswajit Paria, Kirthevasan Kandasamy, Barnabás Póczos</i>	

EFFICIENT MULTITASK FEATURE AND RELATIONSHIP LEARNING	1076
<i>Han Zhao, Otilia Stretcu, Alexander J. Smola, Geoffrey J. Gordon</i>	
PRACTICAL MULTI-FIDELITY BAYESIAN OPTIMIZATION FOR HYPERPARAMETER TUNING	1093
<i>Jian Wu, Saul Toscano-Palmerin, Peter I. Frazier, Andrew Gordon Wilson</i>	
ADAPTIVELY TRUNCATING BACKPROPAGATION THROUGH TIME TO CONTROL GRADIENT BIAS	1108
<i>Christopher Aicher, Nicholas J. Foti, Emily B. Fox</i>	
SAMPLING-FREE UNCERTAINTY ESTIMATION IN GATED RECURRENT UNITS WITH APPLICATIONS TO NORMATIVE MODELING IN NEUROIMAGING	1126
<i>Seong Jae Hwang, Ronak R. Mehta, Hyunwoo J. Kim, Sterling C. Johnson, Vikas Singh</i>	
ONLINE FACTORIZATION AND PARTITION OF COMPLEX NETWORKS BY RANDOM WALK	1137
<i>Lin F. Yang, Zheng Yu, Vladimir Braverman, Tuo Zhao, Mengdi Wang</i>	
ON DENSIFICATION FOR MINWISE HASHING	1148
<i>Tung Mai, Anup Rao, Matt Kapilevich, Ryan Rossi, Yasin Abbsi Yadkori, Ritwik Sinha</i>	

VOLUME 3

N-GCN: MULTI-SCALE GRAPH CONVOLUTION FOR SEMI-SUPERVISED NODE CLASSIFICATION	1158
<i>Sami Abu-El-Haija, Amol Kapoor, Bryan Perozzi, Joonseok Lee</i>	
PROBLEM-DEPENDENT REGRET BOUNDS FOR ONLINE LEARNING WITH FEEDBACK GRAPHS	1169
<i>Bingshan Hu, Nishant A. Mehta, Jianping Pan</i>	
WASSERSTEIN FAIR CLASSIFICATION	1190
<i>Ray Jiang, Aldo Pacchiano, Tom Stepleton, Heinrich Jiang, Silvia Chiappa</i>	
VARIATIONAL TRAINING FOR LARGE-SCALE NOISY-OR BAYESIAN NETWORKS	1205
<i>Geng Ji, Dehua Cheng, Huazhong Ning, Changhe Yuan, Hanning Zhou, Liang Xiong, Erik Sudderth</i>	
GUARANTEED SCALABLE LEARNING OF LATENT TREE MODELS	1216
<i>Furong Huang, Niranjan UN, Ioakeim Perros, Robert Chen, Jimeng Sun, Anima Anandkumar</i>	
ON FIRST-ORDER BOUNDS, VARIANCE AND GAP-DEPENDENT BOUNDS FOR ADVERSARIAL BANDITS	1238
<i>Roman Pogodin, Tor Lattimore</i>	
NOISE CONTRASTIVE PRIORS FOR FUNCTIONAL UNCERTAINTY	1252
<i>Danijar Hafner, Dustin Tran, Timothy Lillicrap, Alex Irpan, James Davidson</i>	
FAKE IT TILL YOU MAKE IT: LEARNING-COMPATIBLE PERFORMANCE SUPPORT	1264
<i>Jonathan Bragg, Emma Brunskill</i>	
LITERAL OR PEDAGOGIC HUMAN? ANALYZING HUMAN MODEL MISSPECIFICATION IN OBJECTIVE LEARNING	1274
<i>Smitha Milli, Anca D. Dragan</i>	
CONVERGENCE ANALYSIS OF GRADIENT-BASED LEARNING IN CONTINUOUS GAMES	1284
<i>Benjamin Chasnov, Lillian J. Ratliff, Eric Mazumdar, Sam Burden</i>	
END-TO-END TRAINING OF DEEP PROBABILISTIC CCA ON PAIRED BIOMEDICAL OBSERVATIONS	1298
<i>Gregory Gundersen, Bianca Dumitrascu, Jordan T. Ash, Barbara E. Engelhardt</i>	
APPROXIMATE RELATIVE VALUE LEARNING FOR AVERAGE-REWARD CONTINUOUS STATE MDPS	1309
<i>Hiteshi Sharma, Mehdi Jafarnia-Jahromi, Rahul Jain</i>	
EXACT SAMPLING OF DIRECTED ACYCLIC GRAPHS FROM MODULAR DISTRIBUTIONS	1318
<i>Topi Talvitie, Aleksis Vuoksenmaa, Mikko Koivisto</i>	
INTERVENING ON NETWORK TIES	1328
<i>Eli Sherman, Ilya Shpitser</i>	
GENERATING AND SAMPLING ORBITS FOR LIFTED PROBABILISTIC INFERENCE	1346
<i>Steven Holtzen, Todd Millstein, Guy Van den Broeck</i>	
REAL-TIME ROBOTIC SEARCH USING STRUCTURAL SPATIAL POINT PROCESSES	1356
<i>Olov Andersson, Per Sidén, Johan Dahlin, Patrick Doherty, Mattias Villani</i>	
SOCIAL REINFORCEMENT LEARNING TO COMBAT FAKE NEWS SPREAD	1369
<i>Mahak Goindani, Jennifer Neville</i>	
P30: POLICY-ON POLICY-OFF POLICY OPTIMIZATION	1385
<i>Rasool Fakoor, Pratik Chaudhari, Alexander J. Smola</i>	

CAUSAL INFERENCE UNDER INTERFERENCE AND NETWORK UNCERTAINTY	1401
<i>Rohit Bhattacharya, Daniel Malinsky, Ilya Shpitser</i>	
REVISITING REWEIGHTED WAKE-SLEEP FOR MODELS WITH STOCHASTIC CONTROL FLOW	1417
<i>Tuan Anh Le, Adam R. Kosiorok, N. Siddharth, Yee Whye Teh, Frank Wood</i>	
LEARNABILITY FOR THE INFORMATION BOTTLENECK	1428
<i>Tailin Wu, Ian Fischer, Isaac Chuang, Max Tegmark</i>	
LEARNING BELIEF REPRESENTATIONS FOR IMITATION LEARNING IN POMDPS	1459
<i>Tanmay Gangwani, Joel Lehman, Qiang Liu, Jian Peng</i>	
OBJECT CONDITIONING FOR CAUSAL INFERENCE	1473
<i>David Jensen, Javier Burroni, Matthew Rattigan</i>	
CCMI : CLASSIFIER BASED CONDITIONAL MUTUAL INFORMATION ESTIMATION	1487
<i>Sudipto Mukherjee, Himanshu Asnani, Sreeram Kannan</i>	
EMPIRICAL MECHANISM DESIGN: DESIGNING MECHANISMS FROM DATA	1504
<i>Enrique Areyan Viqueira, Cyrus Cousins, Yasser Mohammad, Amy Greenwald</i>	
ON THE RELATIONSHIP BETWEEN SATISFIABILITY AND MARKOV DECISION PROCESSES	1515
<i>Ricardo Salmon, Pascal Poupart</i>	
INTERPRETABLE ALMOST MATCHING EXACTLY WITH INSTRUMENTAL VARIABLES	1528
<i>M.Ussaid Awan, Yameng Liu, Marco Morucci, Sudeepa Roy, Cynthia Rudin, Alexander Volfovsky</i>	
LOW FREQUENCY ADVERSARIAL PERTURBATION	1544
<i>Chuan Guo, Jared S. Frank, Kilian Q. Weinberger</i>	
MARKOV LOGIC NETWORKS FOR KNOWLEDGE BASE COMPLETION: A THEORETICAL ANALYSIS UNDER THE MCAR ASSUMPTION	1555
<i>Ondrej Kuzelka, Jesse Davis</i>	
IDENTIFICATION IN MISSING DATA MODELS REPRESENTED BY DIRECTED ACYCLIC GRAPHS	1567
<i>Rohit Bhattacharya, Razieh Nabi, Ilya Shpitser, James M. Robins</i>	
A WEIGHTED MINI-BUCKET BOUND FOR SOLVING INFLUENCE DIAGRAM	1584
<i>Junkyu Lee, Radu Marinescu, Alexander Iher, Rina Dechter</i>	
SUBSPACE INFERENCE FOR BAYESIAN DEEP LEARNING	1594
<i>Pavel Izmailov, Wesley J. Maddox, Polina Kirichenko, Timur Garipov, Dmitry Vetrov, Andrew Gordon Wilson</i>	
OFF-POLICY POLICY GRADIENT WITH STATIONARY DISTRIBUTION CORRECTION	1613
<i>Yao Liu, Adith Swaminathan, Alekh Agarwal, Emma Brunskill</i>	
CO-TRAINING FOR POLICY LEARNING	1630
<i>Jialin Song, Ravi Lanka, Yisong Yue, Masahiro Ono</i>	
VARIATIONAL INFERENCE OF PENALIZED REGRESSION WITH SUBMODULAR FUNCTIONS	1644
<i>Koh Takeuchi, Yuichi Yoshida, Yoshinobu Kawahara</i>	
PROBABILITY DISTILLATION: A CAVEAT AND ALTERNATIVES	1654
<i>Chin-Wei Huang, Faruk Ahmed, Kundan Kumar, Alexandre Lacoste, Aaron Courville</i>	
BAYESIAN OPTIMIZATION WITH BINARY AUXILIARY INFORMATION	1666
<i>Yehong Zhang, Zhongxiang Dai, Bryan Kian Hsiang Low</i>	
ON OPEN-UNIVERSE CAUSAL REASONING	1685
<i>Duligur Ibeling, Thomas Icard</i>	
EMBARRASSINGLY PARALLEL MCMC USING DEEP INVERTIBLE TRANSFORMATIONS	1696
<i>Diego Mesquita, Paul Blomstedt, Samuel Kaski</i>	
FAST PROXIMAL GRADIENT DESCENT FOR A CLASS OF NON-CONVEX AND NON-SMOOTH SPARSE LEARNING PROBLEMS	1705
<i>Yingzhen Yang, Jiahui Yu</i>	
BLOCK NEURAL AUTOREGRESSIVE FLOW	1723
<i>Nicola De Cao, Wilker Aziz, Ivan Titov</i>	
EXCLUSIVITY GRAPH APPROACH TO INSTRUMENTAL INEQUALITIES	1736
<i>Davide Poderini, Rafael Chaves, Iris Agresti, Gonzalo Carvacho, Fabio Sciarrino</i>	
Author Index	