

33rd Conference on Uncertainty in Artificial Intelligence 2017

Sydney, Australia
11-15 August 2017

Volume 1 of 2

Editors:

Gal Elidan

Kristian Kersting

ISBN: 978-1-5108-4779-8

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© (2017) by Association For Uncertainty in Artificial Intelligence (AUAI)
All rights reserved.

Printed by Curran Associates, Inc. (2017)

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

Conference Program

All the conference talks will be held in **Room C4.8** on the fourth level of the Sydney International Convention Centre. Poster sessions and coffee breaks will be located in **Room C4.9**.

Saturday August 12th

8:30 Welcome

Keynote Talk

8:40 Intelligent Robots in an Uncertain World
Leslie Pack Kaelbling, MIT

Session 1: Deep Models

9:40 48: Inverse Reinforcement Learning via Deep Gaussian Process 1
Ming Jin, Andreas Damianou, Pieter Abbeel, Costas Spanos

10:00 267: Holographic Feature Representations of Deep Networks 18
Martin A. Zinkevich, Alex Davies, Dale Schuurmans

10:20 173: Computing Nonvacuous Generalization Bounds for Deep (Stochastic) Neural Networks with Many More Parameters than Training Data 43
Gintare Karolina Dziugaite and Daniel M. Roy

10:40 - 11:10 Coffee break

Session 2: Machine Learning

11:10 32: Provable Inductive Robust PCA via Iterative Hard Thresholding 58
U.N. Niranjan, Arun Rajkumar, Theja Tulabandhula

11:30 6: Near-Orthogonality Regularization in Kernel Methods 71
Pengtao Xie, Barnabas Poczos, Eric Xing

11:50 255: How Good Are My Predictions? Efficiently Approximating Precision-Recall Curves for Massive Datasets 88
Ashish Sabharwal and Hanie Sedghi

12:10 - 14:00 Lunch break

Keynote Talk

14:00 Learning and Inference with Expectations
Amir Globerson, University of Tel Aviv

Session 3: Inference

- 15:00** 160: On Loopy Belief Propagation – Local Stability Analysis for Non-Vanishing Fields 100
Christian Knoll and Franz Pernkopf
- 15:20** 259: Improving Optimization-Based Approximate Inference by Clamping Variables 110
Junyao Zhao, Josip Djolonga, Sebastian Tschiatschek, Andreas Krause
- 15:40** 109: Approximation Complexity of Maximum A Posteriori Inference in Sum-Product Networks 124
Diarmaid Conaty, Cassio P. de Campos, Denis D. Maua

16:00 - 16:20 Coffee break

Session 4: Learning

- 16:20** 291: Learning the Structure of Probabilistic Sentential Decision Diagrams 134
Yitao Liang, Jessa Bekker, Guy Van den Broeck
- 16:40** 278: A Probabilistic Framework for Multi-Label Learning with Unseen Labels 146
Abhilash Gaure and Piyush Rai
- 17:00** 297: Hybrid Deep Discriminative/Generative Models for Semi-Supervised Learning 156
Volodymyr Kuleshov and Stefano Ermon

Poster Spotlights 1

17:20 - 17:50 one minute poster spotlight presentations of papers:

- 100: Regret Minimization Algorithms for the Followers Behaviour Identification in Leadership Games 166
Lorenzo Bisi, Giuseppe De Nittis, Francesco Trovò, Marcello Restelli, Nicola Gatti
- 105: On the Complexity of Nash Equilibrium Reoptimization 183
Andrea Celli, Alberto Marchesi, Nicola Gatti
- 26: Shortest Path under Uncertainty: Exploration versus Exploitation 193
Zhan Wei Lim, David Hsu, Wee Sun Lee
- 35: Learning with Confident Examples: Rank Pruning for Robust Classification with Noisy Labels 203
Curtis G. Northcutt, Tailin Wu, Isaac L. Chuang
- 37: Monte-Carlo Tree Search using Batch Value of Perfect Information 228
Shahaf S. Shperberg, Solomon Eyal Shimony, Ariel Felner
- 45: Submodular Variational Inference for Network Reconstruction 238
Lin Chen, Forrest W. Crawford, Amin Karbasi
- 70: Bayesian Inference of Log Determinants 253
J. Fitzsimons, K. Cutajar, M. Filippone, M. Osborne, S. Roberts
- 75: Fast Amortized Inference and Learning in Log-linear Models with Randomly Perturbed Nearest Neighbor Search 264
Stephen Mussmann, Daniel Levy, Stefano Ermon

- 106: Supervised Restricted Boltzmann Machines 279
Tu Dinh Nguyen, Dinh Phung, Viet Huynh, Trung Le
- 108: Safe Semi-Supervised Learning of Sum-Product Networks 289
Martin Trapp, Tamas Madl, Robert Peharz, Franz Pernkopf, Robert Trapp
- 142: Green Generative Modeling: Recycling Dirty Data
using Recurrent Variational Autoencoders 299
Yu Wang, Bin Dai, Gang Hua, John Aston, David Wipf
- 146: Approximate Evidential Reasoning Using Local Conditioning
and Conditional Belief Functions 322
Van Nguyen
- 152: Differentially Private Variational Inference for Non-conjugate Models 332
Joonas Jälkö, Antti Honkela, Onur Dikmen
- 162: Value Directed Exploration in Multi-Armed Bandits with Structured Priors 342
Bence Cserna, Marek Petrik, Reazul Hasan Russel, Wheeler Ruml
- 188: Learning Approximately Objective Priors 355
Eric Nalisnick and Padhraic Smyth
- 206: Learning to Draw Samples with Amortized Stein Variational Gradient Descent 365
Yihao Feng, Dilin Wang, Qiang Liu
- 220: A Tractable Probabilistic Model for Subset Selection 377
Yujia Shen, Arthur Choi, Adnan Darwiche
- 229: Structure Learning of Linear Gaussian Structural Equation Models
with Weak Edges 391
Marco Eigenmann, Preetam Nandy, Marloes Maathuis
- 234: SAT-Based Causal Discovery under Weaker Assumptions 409
Zhalama, Jiji Zhang, Frederick Eberhardt, Wolfgang Mayer
- 237: Learning to Acquire Information 419
Yewen Pu, Leslie Pack Kaelbling, Armando Solar-Lezama

17:50 - 19:50 Poster Session 1

Will include all papers presented on Saturday.

Sunday August 13th

Keynote Talk

8:30 Snorkel: Beyond Hand-labeled Data
Christopher Ré, Stanford

Session 5: Representations

9:30 263: Why Rules are Complex: Real-Valued Probabilistic Logic Programs are not Fully Expressive 429
David Buchman and David Poole

9:50 292: Interpreting Lion Behaviour as Probabilistic Programs 439
Neil Dhir, Matthijs Vakar, Matthew Wijers, Andrew Markham, Frank Wood

10:10 147: Decoupling Homophily and Reciprocity with Latent Space Network Models 456
Jiasen Yang, Vinayak Rao, Jennifer Neville

10:30 - 11:00 Coffee break

Session 6: Reinforcement Learning

11:00 276: Online Constrained Model-based Reinforcement Learning 470
Benjamin van Niekerk, Andreas Damianou, Benjamin Rosman

11:20 209: A Reinforcement Learning Approach to Weaning of Mechanical Ventilation in Intensive Care Units 480
Niranjani Prasad, Li-Fang Cheng, Corey Chivers, Michael Draugelis, Barbara E. Engelhardt

11:40 62: Near-Optimal Interdiction of Factored MDPs 490
Swetasudha Panda and Yevgeniy Vorobeychik

12:00 225: Importance Sampling for Fair Policy Selection 500
Shayan Doroudi, Philip Thomas, Emma Brunskill

12:20 - 14:10 Lunch break

Keynote Talk

14:10 - 15:10 Machine Learning for Healthcare Data
Katherine Heller, Duke University

Poster Spotlights 2

15:10pm - 15:40 one minute poster spotlight presentations of papers:

12: FROSH: FasteR Online Sketching Hashing 515
Xixian Chen, Irwin King, Michael R. Lyu

16: Self-Discrepancy Conditional Independence Test 528
Sanghack Lee and Vasant Honavar

- 49: Towards Conditional Independence Test for Relational Data 539
Sanghack Lee and Vasant Honavar
- 50: AutoGP: Exploring the Capabilities and Limitations
of Gaussian Process Models 552
Karl Krauth, Edwin V. Bonilla, Kurt Cutajar, Maurizio Filippone
- 65: A Fast Algorithm for Matrix Eigen-decomposition 565
Zhiqiang Xu, Yiping Ke, Xin Gao
- 78: Branch and Bound for Regular Bayesian Network Structure Learning 581
Joe Suzuki, Jun Kawahara
- 81: Effective sketching methods for value function approximation 593
Yangchen Pan, Erfan Sadeqi Azer, Martha White
- 86: Stochastic L-BFGS Revisited: Improved Convergence Rates 606
and Practical Acceleration Strategies
Renbo Zhao, William B. Haskell, Vincent Y. F. Tan
- 87: The Binomial Block Bootstrap Estimator for Evaluating Loss
on Dependent Clusters 623
Matt Barnes, Artur Dubrawski
- 93: Data-Dependent Sparsity for Subspace Clustering 639
Bo Xin, Yizhou Wang, Wen Gao and David Wipf
- 132: Weighted Model Counting With Function Symbols 654
Vaishak Belle
- 136: Triply Stochastic Gradients on Multiple Kernel Learning 664
Xiang Li, Bin Gu, Shuang Ao, Huaimin Wang, Charles X. Ling
- 141: Coupling Adaptive Batch Sizes with Learning Rates 675
Lukas Balles, Javier Romero, Philipp Hennig
- 163: Composing Inference Algorithms as Program Transformations 686
Robert Zinkov, Chung-chieh Shan
- 197: Iterative Decomposition Guided Variable Neighborhood Search
for Graphical Model Energy Minimization 696
*Abdelkader Ouali, David Allouche, Simon de Givry, Samir Loudni,
Yahia Lebbah, Lakhdar Loukil*
- 207: Fair Optimal Stopping Policy for Matching with Mediator 710
Yang Liu
- 242: Exact Inference for Relational Graphical Models with Interpreted Functions:
Lifted Probabilistic Inference Modulo Theories 723
Rodrigo de Salvo Braz, Ciaran O'Reilly
- 244: Neighborhood Regularized ℓ^1 -Graph 736
Yingzhen Yang, Jiashi Feng, Jiahui Yu, Jianchao Yang, Thomas S. Huang
- 250: Feature-to-Feature Regression for a Two-Step Conditional Independence Test 751
Qinyi Zhang, Sarah Filippi, Seth Flaxman, Dino Sejdinovic
- 277: Algebraic Equivalence Class Selection for Linear Structural Equation Models 763
Thijs van Ommen, Joris M. Mooij

15:40 - 16:00 Coffee break

16:00 - 18:00 Poster Session 2

Will include all papers presented on Sunday.

18:45 - Banquet Boarding

Banquet Talk: 15 minutes on Artificial Intelligence and Statistical Models
Terry Speed, Walter and Eliza Hall Institute of Medical Research

Monday August 14th

Keynote Talk

- 8:30** Two current analysis challenges: Single Cell Omics and Nanopore Long-read Sequence Data
Terry Speed, Walter and Eliza Hall Institute of Medical Research

Session 7: Causality

- 9:30** 266: Learning Treatment-Response Models from Multivariate Longitudinal Data 780
Hossein Soleimani, Adarsh Subbaswamy, Suchi Saria
- 9:50** 120: Interpreting and Using CPDAGs With Background Knowledge 791
Emilija Perkovic, Markus Kalisch, Marloes H. Maathuis
- 10:10** 11: Causal Consistency of Structural Equation Models 808
Paul K. Rubenstein, Sebastian Weichwald, Stephan Bongers, Joris M. Mooij, Dominik Janzing, Moritz Grosse-Wentrup, Bernhard Schölkopf

10:30 - 11:00 Coffee break

Session 8: Sampling

- 11:00** 217: Stein Variational Adaptive Importance Sampling 823
Jun Han, Qiang Liu
- 11:20** 289: Continuously Tempered Hamiltonian Monte Carlo 833
Matthew M. Graham, Amos J. Storkey
- 11:40** 69: Balanced Mini-batch Sampling for SGD Using Determinantal Point Processes 849
Cheng Zhang, Hedvig Kjellstrom, Stephan Mandt
- 12:00** 54: An Efficient Minibatch Acceptance Test for Metropolis-Hastings 862
Daniel Seita, Xinlei Pan, Haoyu Chen, John Canny

12:20 - 14:10 Lunch break

Session 9: Bandits

- 14:10** 137: Stochastic Bandit Models for Delayed Conversions 879
Claire Vernade, Olivier Cappé, Vianney Perchet
- 14:30** 171: A Practical Method for Solving Contextual Bandit Problems Using Decision Trees 897
Ryan McNellis, Adam N. Elmachtoub, Sechan Oh, Marek Petrik
- 14:50** 180: Analysis of Thompson Sampling for Stochastic Sleeping Bandits 916
Aritra Chatterjee, Ganesh Ghalmé, Shweta Jain, Rohit Vaish, Y. Narahari

Poster Spotlights 3

15:10 - 15:40 one minute poster spotlight presentations of papers:

- 56: The Total Belief Theorem 929
Chunlai Zhou, Fabio Cuzzolin
- 64: Complexity of Solving Decision Trees with Skew-Symmetric Bilinear Utility 939
Hugo Gilbert, Olivier Spanjaard
- 72: Stochastic Segmentation Trees for Multiple Ground Truths 949
Jake Snell and Richard S. Zemel
- 83: Efficient Online Learning for Optimizing Value of Information:
Theory and Application to Interactive Troubleshooting 966
Yuxin Chen, Jean-Michel Renders, Morteza Haghir Chehreghani, Andreas Krause
- 97: Counting Markov Equivalence Classes by Number of Immoralities 984
Adityanarayanan Radhakrishnan, Liam Solus, Caroline Uhler
- 130: Real-Time Resource Allocation for Tracking Systems 994
Yash Satsangi, Shimon Whiteson, Frans Oliehoek, Henri Bouma
- 134: Synthesis of Strategies in Influence Diagrams 1004
Manuel Luque, Manuel Arias, Francisco Javier Díez
- 135: Embedding Senses via Dictionary Bootstrapping 1013
Byungkon Kang, Kyung-Ah Sohn
- 138: Importance Sampled Stochastic Optimization for Variational Inference 1023
Joseph Sakaya, Arto Klami
- 155: Multi-dueling Bandits with Dependent Arms 1033
Yanan Sui, Vincent Zhuang, Joel W. Burdick, Yisong Yue
- 185: Convex-constrained Sparse Additive Modeling and Its Extensions 1047
Junming Yin, Yaoliang Yu
- 239: Stein Variational Policy Gradient 1057
Yang Liu, Prajit Ramachandran, Qiang Liu, Jian Peng
- 269: Causal Discovery from Temporally Aggregated Time Series 1066
Mingming Gong, Kun Zhang, Bernhard Schölkopf, Clark Glymour, Dacheng Tao
- 280: Efficient solutions for Stochastic Shortest Path Problems with Dead Ends 1076
Felipe Trevizan, Florent Teichteil-Knigsbuch, Sylvie Thiebaux
- 281: Probabilistic Program Abstractions 1086
Steven Holtzen, Todd Millstein, Guy Van den Broeck
- 286: Communication-Efficient Distributed Primal-Dual Algorithm
for Saddle Point Problem 1096
Yaodong Yu, Sulin Liu, Sinno Jialin Pan
- 298: Robust Model Equivalence using Stochastic Bisimulation
for N-Agent Interactive DIDs 1113
Muthukumar Chandrasekaran, Junhuan Zhang, Prashant Doshi, Yifeng Zeng
- 306: Adversarial Sets for Regularising Neural Link Predictors 1123
Pasquale Minervini, Thomas Demeester, Tim Rocktäschel, Sebastian Riedel