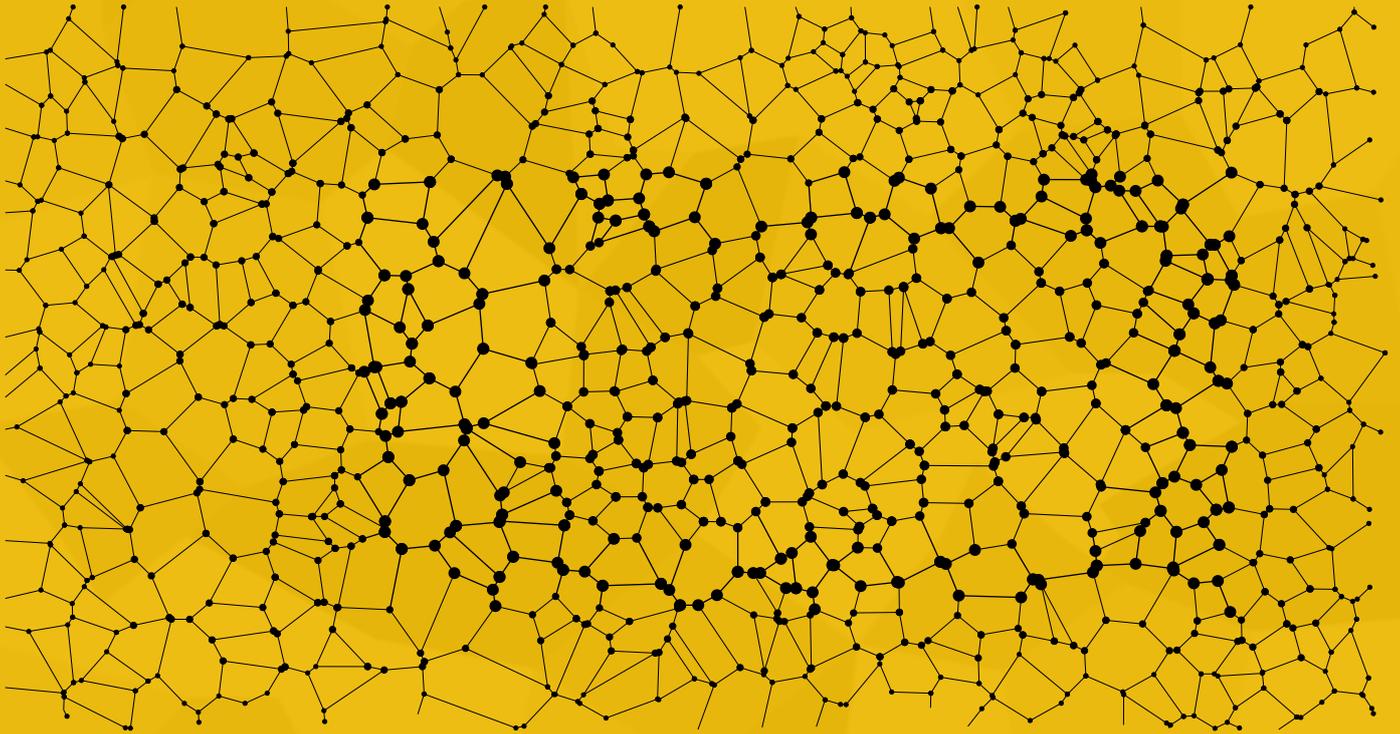


# Uncertainty in Artificial Intelligence





# Uncertainty in Artificial Intelligence

Proceedings of the Thirty-Third Conference (2017)

August 11-15, 2017, Sydney, Australia

**Edited by**

Gal Elidan, Google Inc. and The Hebrew University, Israel  
Kristian Kersting, TU Darmstadt, Germany

**General Chair**

Alexander Ihler, University of California Irvine, USA

**Sponsored by**

Google Inc., Microsoft Research, Amazon Web Services, Artificial  
Intelligence Journal, Disney Research, Uber, NextAI

**AUAI Press Corvallis, Oregon**

# Contents

Preface	5
Organizing Committee	6
Acknowledgments	7
Sponsors	13
Best Paper Awards	14
Tutorials	15
Conference Program	17
Invited Speakers	26
Workshop Program	29
Paper Abstracts	35
Notes	66
Maps	74

---

# Preface

---

The Conference on Uncertainty in Artificial Intelligence (UAI) is the premier international conference on research related to representation, inference, learning and decision making in the presence of uncertainty within the field of Artificial Intelligence. This volume contains the schedule and abstracts of all papers that were accepted for the 33rd UAI Conference, held in Sydney, Australia, from August 11 to 15, 2017.

Papers appearing in this volume were subjected to a rigorous review process. A total of 282 papers were reviewed by at least 3 reviewers each. Of these, 87 papers were accepted, for an acceptance rate of close to 31%. We are very grateful to the program committee and senior program committee members for their diligent efforts. We are confident that the proceedings, like past UAI conference proceedings, will become an important archival reference for the field.

We are pleased to announce that the Best Paper Award is awarded to Shayan Doroudi, Philip S. Thomas, and Emma Brunskill for their paper “Importance Sampling for Fair Policy Selection”. The Best Student Paper Award is awarded to David Buchman and David Poole for their paper “Why Rules are Complex: Real-Valued Probabilistic Logic Programs are not Fully Expressive”. Finally, a Best Student Paper Honorable Mention is given to Daniel Seita, Xinlei Pan, Haoyu Chen, and John Canny for their paper “An Efficient Minibatch Acceptance Test for Metropolis-Hastings”. We are grateful to the members of the best paper committee Tom Heskes (chair), Pascal Poupart, and Sriraam Natarajan.

In addition to the presentation of technical papers, we are very pleased to have five distinguished invited speakers at UAI 2017: Amir Globerson (Tel-Aviv University), Katherine Heller (Duke University), Leslie Pack Kaelbling (MIT), Christopher Ré (Stanford University), and Terry Speed (Walter and Eliza Hall Institute of Medical Research).

The UAI 2017 tutorials program, chaired by Ricardo Silva, consists of four tutorials: “Methods and Models for Large-scale Optimization” by John Duchi (Stanford University), “Representing and Comparing Probabilities with (and without) Kernels” by Arthur Gretton (University College London), “Deep Generative Models” by Shakir Mohamed and Danilo Rezende (DeepMind), and an invited tutorial “Machine Learning and Counterfactual Reasoning for ‘Personalized’ Decision-Making in Healthcare” by Suchi Saria and Hossein Soleimani (Johns Hopkins University).

UAI 2017 also hosts three workshops, chaired by Guy Van den Broeck: “Statistical Relational AI (StarAI)” organized by Angelika Kimmig, David Poole, Jay Pujara, and Tim Rocktäschel. “Causality: Learning, Inference, and Decision-Making”, organized by Elias Bareinboim, Kun Zhang, Caroline Uhler, Jiji Zhang, and Dominik Janzing. “Bayesian Modelling Applications”, organized by John Mark Agosta and Tomas Singliar. UAI 2017 will also hold, for the first time, a training event where participants will learn how to code papers with TensorFlow, MXNet, PyTorch, CNTK and more. The event will include six training sessions: Generative Models, Neural Abstract Machines, Reinforcement Learning, LSTMs and Text, Speech and Time series, and Deep Recommender Systems. The session will be led instructors Alex Dimakis, Chris Fregly, and Nikolaos Vasiloglou.

*Gal Elidan and Kristian Kersting (Program Co-Chairs),  
Alexander Ihler (General Chair)*

---

# Organizing Committee

---

## **General Chair**

Alexander Ihler, University of California Irvine, USA

## **Program Chairs**

Gal Elidan, Google Inc. and The Hebrew University, Israel  
Kristian Kersting, TU Darmstadt, Germany

## **Best Paper Award Committee**

Tom Heskes (chair), Radboud University, Netherlands  
Pascal Pupart, University of Waterloo, Canada  
Sriraam Natarajan, Indiana University, USA

## **Workshops Chair**

Guy Van den Broeck, University of California Los Angeles, USA

## **Tutorial Chair**

Ricardo Silva, University College London, UK

## **Sponsoring and Social Media Chair**

Nikolaos Vasiloglou, Ismion, USA

## **Local Arrangements Chair**

Ann Nicholson, Monash University, Australia

## **Proceedings Chair**

Martin Mladenov, TU Dortmund, Germany

## **Webmaster and OpenReview Chair**

Alejandro Molina, TU Dortmund, Germany

---

# Acknowledgments

---

The success of UAI depends greatly on the efforts of many individuals who volunteer their time to provide expert and detailed reviews of submitted papers. In particular, the Program Committee and Senior Program Committee for UAI 2017 were responsible for generating reviews and recommendations for the 282 submissions to the conference. Each submitted paper was reviewed by at least 3 members of the Program Committee. The Senior Program Committee then assessed the individual reviews for each paper, moderated discussion among Program Committee members if needed, and generated meta-reviews and recommendations for the program chairs. We are extremely grateful for the efforts of all of the individuals listed below.

## Senior Program Committee

Alan Fern	Oregon State University
Amir Globerson	Tel Aviv University
Ann Nicholson	Monash University
Ayesha Ali	University of Guelph
Benjamin Marlin	UMass Amherst
Brian Milch	Google
Cassio de Campos	Queen's University Belfast
Chris Meek	Microsoft Research
Claire Monteleoni	George Washington University
Craig Boutilier	Google
Dhruv Batra	Virginia Tech
Fabio Cuzzolin	Oxford Brookes University
Francisco Diez	UNED
Guy Van den Broek	UCLA
Helene Fargier	Institut de Recherche en Informatique de Toulouse
Jennifer Dy	Northeastern University
Jin Tian	Iowa State University
Jun Zhu	Tsinghua University
Kamalika Chaudhuri	University of California
Marloes Maathuis	ETH Zurich
Michel Besserve	Max Planck Institute
Nevin Zhang	Hong Kong University of Science and Technology
Nuria M. Oliver	Vodafone & Data-Pop Alliance
Ofer Meshi	Google
Pascal Poupart	University of Waterloo
Peter Grnwald	Centrum voor Wiskunde en Informatica
Peter Lucas	Radboud University Nijmegen
Raquel Urtasun	University of Toronto
Rina Dechter	UC-Irvine
Ryan P. Adams	Harvard
Stefano Ermon	Stanford
Thomas Richardson	University of Washington
Tom Heskes	Radboud University Nijmegen
Vibhav Gogate	The University of Texas at Dallas
Wray Buntine	Monash University
Zoltan Szabo	École Polytechnique

## Program Committee

Alessandro Antonucci	IDSIA
Alessio Benavoli	IDSIA
Alex Kirillov	TU Dresden
Alex Schwing	University of Illinois at Urbana-Champaign
Alexandre Bouchard-Cote	UBC
Ali Ghodsi	University of Waterloo
Ami Wiesel	Hebrew University
Amit Gruber	Yahoo!
Anders Madsen	Hugin Expert
Andrey Kolobov	Microsoft Research
Antonio Salmeron	Universidad de Almeria
Antti Hyttinen	University of Helsinki
April Liu	The Hong Kong University of Science and Technology
Aram Galstyan	Information Sciences Institute
Aritanan Gruber	University of Sao Paulo
Arjen Hommersom	University of Nijmegen
Arthur Choi	UCLA
Arvind Agarwal	Xerox
Ashish Sabharwal	Allen Institute for Artificial Intelligence
Aurelie Lozano	IBM
Balaji Lakshminarayanan	UCL
Bart Selman	Cornell
Ben London	Amazon
Bob Price	PARC
Bozhena Bidyuk	Google
Brandon Malone	University of Helsinki
Branislav Kveton	Adobe Research
Brian Ruttenger	Charles River Analytics
Carlo Berzuni	University of Manchester
Changhe Yuan	City University of New York
Changyou Chen	Duke University
Chong Wang	Carnegie Mellon University
Christoph Dann	CMU
Christophe Gonzales	LIP6-UPMC
Clark Glymour	CMU
Cory Butz	University of Regina
Daniel Lowd	University of Oregon
Daniel Malinsky	CMU
David Arbour	University of Massachusetts Amherst
David Balduzzi	Victoria University Wellington
David Jensen	University of Massachusetts Amherst
David Page	UW Madison
David Pennock	Microsoft Research
David Pynadath	USC Institute for Creative Technologies
Debarun Bhattacharjya	IBM Research
Denis Mauá	University of Sao Paulo
Dragomir Anguelov	Google
Dustin Tran	Columbia
Elad Eban	Google
Elias Bareinboim	Purdue University
Erik Cambria	Nanyang Technological University

Ethan Zhao	Monash University
Eunho Yang	KAIST
Eyal Amir	University of Illinois at Urbana-Champaign
Fabio Stella	University of Milan
Farhad Anaraki	University of Colorado
Frederick Eberhardt	Caltech
Georg P Schoenherr	CMU
Gerardo Simari	Universidad Nacional del Sur in Bahia Blanca & CONICET
Greg Ver Steeg	Information Sciences Institute
Hang Su	Tsinghua University
Hong Chang	Institute of Computing Technology
Hung Bui	Adobe
Ilya Shpitser	Johns Hopkins Univeristy
Ioannis Tsamardinos	University of Crete
Jaesik Choi	Ulsan National Institute of Science and Technology
James Cussens	University of York
James Foulds	UC San Diego
Jan Gasthaus	University College London
Janardhan Rao Doppa	Washington State University
Jay Pujara	UC Santa Cruz
Jesse Hoey	University of Waterloo
Jesse Hostetler	Oregon State University
Jian Peng	UIUC
Jiji Zhang	Lingnan University
Jirka Vomlel	Institute of Information Theory and Automation
Joe Suzuki	Osaka University
Johannes Textor	Utrecht University
John Paisley	Columbia University
John Mark Agosta	Toyota Information Technology Center
Jonas Vlasselaer	KU Leuven
Jose Peña	Linköping University
Julian McAuley	UC San Diego
Junxiang Chen	Northeastern University
Justin Domke	University of Massachusetts Amherst
Kacper Chwialkowski	UCL / Gatsby Unit
Kar Wai Lim	ANU
Karthika Mohan	UCLA
Kevin Small	Amazon
Kim-Leng Poh	National University of Singapore
Kun Zhang	CMU
L. Enrique Sucar	INAOE
Lan Du	Monash
Le Song	Georgia Tech
Leonard Poon	Hong Kong Institute of Education
Linbo Wang	University of Washington
Lirong Xia	Rensselaer Polytechnic Institute
Luis Godo	Artificial Intelligence Research Institute
M. Julia Flores	University of Castilla - La Mancha (UCLM)
Mahesh Mohan	George Washington University
Malik Magdon-Ismail	Rensselaer Polytechnic Institute
Manuel Gomez-Olmedo	Universidad de Granada
Manuel Luque	UNED
Marco Valtorta	University of South Carolina

Marcus Hutter	Australian National University
Maria Vanina Martinez	University of Oxford
Mark Crowley	University of Waterloo
Martin Zinkevich	Google
Mathias Niepert	NEC
Mathieu Sinn	IBM Research
Maya Gupta	Google
Michael Lyu	Chinese University of Hong Kong
Mikko Koivisto	University of Helsinki
Mingming Gong	University of Technology Sydney
Minjie Xu	Tsinghua University
Minos Garofalakis	Technical University of Crete
Mohammad Emtiyaz Khan	RIKEN
Nahla Ben Amor	ISG Tunis
Nan Ye	National University of Singapore
Narges Razavian	New York University
Nicolas Drougard	ONERA - The French Aerospace Lab
Nimar Arora	Oracle
Ole Mengshoel	Carnegie Mellon University
Olivier Buffet	LORIA-INRIA
Paola Vicard	Universit degli Studi Roma Tre
Perry Groot	Radboud University Nijmegen
Philipp Geiger	Max Planck Institute for Intelligent Systems
Philipp Hennig	Max Planck Institute
Philippe Leray	University of Nantes
Piyush Rai	IIT Kanpur
Preetam Nandy	ETH
Qi Lou	UCI
Quanquan Gu	University of Virginia
Rafael Rumi	Almeria University
Reza Babanezhad Harikandeh	The University of British Columbia
Robert Castelo	Universitat Pompeu Fabra
Robin Evans	University of Oxford
Roger Baker Grosse	University of Toronto
Roland Ramsahai	University of Cambridge
Roman Garnett	Washington University in St. Louis
Ru He	Amazon
Russell Almond	Florida State University
Scott Sanner	University of Toronto
Sebastien Destercke	CNRS
Shang-Tse Chen	Georgia Tech
Shuai Li	Cambridge University
Silja Renooij	Universiteit Utrecht
Sriraam Natarajan	Indiana University
Steffen Michels	Radboud University Nijmegen
Stephen Bach	Stanford
Su-In Lee	University of Washington
Suriya Gunasekar	TTI
Tameem Adel	University of Manchester
Taneli Mielikainen	Yahoo
Teemu Roos	Helsinki Institute for Information Technology
Tengfei Liu	The Hong Kong University
Thomas Walsh	MIT

Tianjiao Chu	University of Pittsburgh
Tiberio Caetano	Ambiata & University of Sydney
Tom Claassen	Radboud University Nijmegen
Tomas Singliar	Microsoft
Tomi Silander	Xerox Research Centre Europe
Truong-Huy Dinh Nguyen	Texas A&M University
Twan van Laarhoven	Radboud University Nijmegen
Uri Shalit	New York University
Vicenç Gómez	Universitat Pompeu Fabra
Vincent Tan	NUS
Wannes Meert	KU Leuven
Wei Ping	Baidu
William Cheung	Hong Kong Baptist University
Wu-Jun Li	Nanjing University
Xiaotong Yuan	Nanjing University of Information Science & Technology
Yan Liu	USC
Yetian Chen	Amazon
Yexiang Xue	Cornell University
Yi Zhen	Georgia Institute of Technology
Yifeng Zeng	Teesside University
Yingyu Liang	Princeton University
Yoni Halpern	New York University
Yu Zhang	Hong Kong Baptist University
Yuhong Guo	Temple University
Yutian Chen	Google
Zied Eloudi	ISG Tunis

## Additional Acknowledgments

A number of other people have made significant contributions towards making UAI 2017 possible. We acknowledge and thank:

- The OpenReview team:

Andrew McCallum	University of Massachusetts Amherst
Melisa Bok	University of Massachusetts Amherst
Michael Spector	University of Massachusetts Amherst

- The following student scholarship volunteers:

Ouali Abdelkader	CNRS & Universit de Caen Normandie
David Arbour	UMass Amherst
Jessa Bekker	KU Leuven
Jagmohan Chauhan	UNSW & Data61 (CSIRO)
Bence Cserna	University of New Hampshire
Shayan Doroudi	Carnegie Mellon University
Jun Han	Dartmouth College
Chait Manapragada	Monash University
Philippe Morere	The University of Sydney
Swetasudha Panda	Vanderbilt University
Niranjani Prasad	Princeton University
Ehsan Shareghi	Monash University

Benjamin van Niekerk	University of the Witwatersrand
Claire Vernade	Télécom ParisTech
Jiasen Yang	Purdue University
Qinyi Zhang	University of Oxford
He Zhao	Monash University
Fujin Zhu	University of Technology Sydney

---

# Sponsors

---

We gratefully acknowledge the generous support of our sponsors, including support for best paper awards and travel scholarships. Without our sponsors' support it would not be feasible to organize a conference such as UAI 2017 without charging much higher registration fees.

## Gold Sponsors



## Bronze Sponsors



## Startup Sponsor



## Media Sponsor



---

# Best Paper Awards

---

## **Best Paper Award**

*Importance Sampling for Fair Policy Selection*

Shayan Doroudi, Philip S. Thomas, and Emma Brunskill

## **Best Student Paper Award**

*Why Rules are Complex: Real-Valued Probabilistic Logic Programs are not Fully Expressive*

David Buchman and David Poole

## **Best Student Paper Honorable Mention**

*An Efficient Minibatch Acceptance Test for Metropolis-Hastings*

Daniel Seita, Xinlei Pan, Haoyu Chen, and John Canny

---

# Tutorials

---

**Friday August 11th**  
**Room C3.6**

- 8:45–10:15**    Methods and Models for Large-scale Optimization  
*John C. Duchi, Stanford University*
- 10:15–10:35**    **Coffee Break**
- 10:35–12:05**    Representing and Comparing Probabilities with (and without) Kernels  
*Arthur Gretton, University College London*
- 12:05–14:05**    **Lunch Break**
- 14:05–15:35**    Deep Generative Models  
*Shakir Mohamed and Danilo Rezende, DeepMind*
- 15:35–16:00**    **Break**
- 16:00–17:30**    Machine Learning in Healthcare  
*Suchi Saria and Hossein Soleimani, Johns Hopkins University*

## **Tutorial: Methods and Models for Large-scale Optimization**

*John C. Duchi*

In this tutorial, we cover a biased view of recent movements in optimization and its intersection with machine learning and statistics, with an eye toward applying these algorithms. Broadly, we believe these advances lie in two major areas: the first is in large-scale stochastic optimization problems, where the growing size in dataset - both in sample size and dimension - due to the amount of information we track and collect has led to substantial research in the optimization literature. The second is a more recent trend toward non-convex optimization problems. While, as Rockafellar notes, the justifiable dogma in optimization is that "the great watershed in optimization isn't between linearity and nonlinearity, but convexity and nonconvexity," a number of problems in statistics that are nominally non-convex are in fact still (with high probability) easy to solve using simple gradient-based and other methods. Unifying both of these trends is that the inherent randomness in the problems we solve - sampling randomness from collecting data - lends itself to simple optimization methods with convergence guarantees that are, at least, good enough. Our tutorial addresses these trends in three parts: first, by covering methods for large-scale stochastic optimization, second, by discussing optimization models that arise in large-scale problems, and third, by showing how non-convex problems arising from various nice stochastic problems are in practice (and in theory) relatively easy.

## **Tutorial: Representing and Comparing Probabilities with (and without) Kernels**

*Arthur Gretton*

The purpose of this tutorial is to provide an introduction to distribution embeddings and their applications, with a focus on recent tools (from 2014-2017) developed to represent and compare probability distributions. The first part of the talk will focus entirely on the representation of probability distributions, using both kernel methods and explicit features. The focus will be on designing kernels or features to make two distributions as distinguishable as possible. The second part of the talk will focus on more sophisticated applications of distribution representations: model criticism (using Stein’s method to test against a parametric model), learning from probabilities as inputs, testing for independence and testing multi-way interaction. I might also touch on the problem of testing where the inputs are time series.

## **Tutorial: Deep Generative Models**

*Shakir Mohamed and Danilo Rezende*

This tutorial will be a review of recent advances in deep generative models. Generative models have a long history at UAI and recent methods have combined the generality of probabilistic reasoning with the scalability of deep learning to develop learning algorithms that have been applied to a wide variety of problems giving state-of-the-art results in image generation, text-to-speech synthesis, and image captioning, amongst many others. Advances in deep generative models are at the forefront of deep learning research because of the promise they offer for allowing data-efficient learning, and for model-based reinforcement learning. At the end of this tutorial, audience member will have a full understanding of the latest advances in generative modelling covering three of the active types of models: Markov models, latent variable models and implicit models, and how these models can be scaled to high-dimensional data. The tutorial will expose many questions that remain in this area, and for which there remains a great deal of opportunity from members of the UAI community.

## **Tutorial: Machine Learning and Counterfactual Reasoning for “Personalized” Decision-Making in Healthcare**

*Suchi Saria*

Electronic health records and high throughput measurement technologies are changing the practice of healthcare to become more algorithmic and data-driven. This offers an exciting opportunity for statistical machine learning techniques to impact healthcare.

The aim of this tutorial is to introduce you to challenges and techniques for developing “personalized decision-making” tools in medicine. Example topics covered will include: scalable joint models for forecasting risk from “messy” clinical streams, estimating individualized treatment response in populations with heterogeneous treatment effects, and examples of counterfactual reasoning or “what-if” analysis for decision-making. We will also cover example data sources and describe ongoing national initiatives that provide a way for you to get involved.

Target audience: The majority of this tutorial will be targeted at an audience with basic machine learning knowledge. No background in medicine or health care is needed. We will make our slides and any relevant documents accessible after the talk.

---

# 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  
*Ming Jin, Andreas Damianou, Pieter Abbeel, Costas Spanos*

**10:00** 267: Holographic Feature Representations of Deep Networks  
*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  
*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  
*U.N. Niranjan, Arun Rajkumar, Theja Tulabandhula*

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

**11:50** 255: How Good Are My Predictions? Efficiently Approximating Precision-Recall Curves for Massive Datasets  
*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  
*Christian Knoll and Franz Pernkopf*
- 15:20** 259: Improving Optimization-Based Approximate Inference by Clamping Variables  
*Junyao Zhao, Josip Djolonga, Sebastian Tschiatschek, Andreas Krause*
- 15:40** 109: Approximation Complexity of Maximum A Posteriori Inference in Sum-Product Networks  
*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  
*Yitao Liang, Jessa Bekker, Guy Van den Broeck*
- 16:40** 278: A Probabilistic Framework for Multi-Label Learning with Unseen Labels  
*Abhilash Gaure and Piyush Rai*
- 17:00** 297: Hybrid Deep Discriminative/Generative Models for Semi-Supervised Learning  
*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  
*Lorenzo Bisi, Giuseppe De Nittis, Francesco Trovò, Marcello Restelli, Nicola Gatti*
- 105: On the Complexity of Nash Equilibrium Reoptimization  
*Andrea Celli, Alberto Marchesi, Nicola Gatti*
- 26: Shortest Path under Uncertainty: Exploration versus Exploitation  
*Zhan Wei Lim, David Hsu, Wee Sun Lee*
- 35: Learning with Confident Examples: Rank Pruning for Robust Classification with Noisy Labels  
*Curtis G. Northcutt, Tailin Wu, Isaac L. Chuang*
- 37: Monte-Carlo Tree Search using Batch Value of Perfect Information  
*Shahaf S. Shperberg, Solomon Eyal Shimony, Ariel Felner*
- 45: Submodular Variational Inference for Network Reconstruction  
*Lin Chen, Forrest W. Crawford, Amin Karbasi*
- 70: Bayesian Inference of Log Determinants  
*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  
*Stephen Mussmann, Daniel Levy, Stefano Ermon*

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

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

**10:10** 147: Decoupling Homophily and Reciprocity with Latent Space Network Models  
*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  
*Benjamin van Niekerk, Andreas Damianou, Benjamin Rosman*

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

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

**12:00** 225: Importance Sampling for Fair Policy Selection  
*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  
*Xixian Chen, Irwin King, Michael R. Lyu*

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

- 49: Towards Conditional Independence Test for Relational Data  
*Sanghack Lee and Vasant Honavar*
- 50: AutoGP: Exploring the Capabilities and Limitations of Gaussian Process Models  
*Karl Krauth, Edwin V. Bonilla, Kurt Cutajar, Maurizio Filippone*
- 65: A Fast Algorithm for Matrix Eigen-decomposition  
*Zhiqiang Xu, Yiping Ke, Xin Gao*
- 78: Branch and Bound for Regular Bayesian Network Structure Learning  
*Joe Suzuki, Jun Kawahara*
- 81: Effective sketching methods for value function approximation  
*Yangchen Pan, Erfan Sadeqi Azer, Martha White*
- 86: Stochastic L-BFGS Revisited: Improved Convergence Rates 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  
*Matt Barnes, Artur Dubrawski*
- 93: Data-Dependent Sparsity for Subspace Clustering  
*Bo Xin, Yizhou Wang, Wen Gao and David Wipf*
- 132: Weighted Model Counting With Function Symbols  
*Vaishak Belle*
- 136: Triply Stochastic Gradients on Multiple Kernel Learning  
*Xiang Li, Bin Gu, Shuang Ao, Huaimin Wang, Charles X. Ling*
- 141: Coupling Adaptive Batch Sizes with Learning Rates  
*Lukas Balles, Javier Romero, Philipp Hennig*
- 163: Composing Inference Algorithms as Program Transformations  
*Robert Zinkov, Chung-chieh Shan*
- 197: Iterative Decomposition Guided Variable Neighborhood Search for Graphical Model Energy Minimization  
*Abdelkader Ouali, David Allouche, Simon de Givry, Samir Loudni, Yahia Lebbah, Lakhdar Loukil*
- 207: Fair Optimal Stopping Policy for Matching with Mediator  
*Yang Liu*
- 242: Exact Inference for Relational Graphical Models with Interpreted Functions: Lifted Probabilistic Inference Modulo Theories  
*Rodrigo de Salvo Braz, Ciaran O'Reilly*
- 244: Neighborhood Regularized  $\ell^1$ -Graph  
*Yingzhen Yang, Jiashi Feng, Jiahui Yu, Jianchao Yang, Thomas S. Huang*
- 250: Feature-to-Feature Regression for a Two-Step Conditional Independence Test  
*Qinyi Zhang, Sarah Filippi, Seth Flaxman, Dino Sejdinovic*
- 277: Algebraic Equivalence Class Selection for Linear Structural Equation Models  
*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  
*Hossein Soleimani, Adarsh Subbaswamy, Suchi Saria*
- 9:50** 120: Interpreting and Using CPDAGs With Background Knowledge  
*Emilija Perkovic, Markus Kalisch, Marloes H. Maathuis*
- 10:10** 11: Causal Consistency of Structural Equation Models  
*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  
*Jun Han, Qiang Liu*
- 11:20** 289: Continuously Tempered Hamiltonian Monte Carlo  
*Matthew M. Graham, Amos J. Storkey*
- 11:40** 69: Balanced Mini-batch Sampling for SGD Using Determinantal Point Processes  
*Cheng Zhang, Hedvig Kjellstrom, Stephan Mandt*
- 12:00** 54: An Efficient Minibatch Acceptance Test for Metropolis-Hastings  
*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  
*Claire Vernade, Olivier Cappé, Vianney Perchet*
- 14:30** 171: A Practical Method for Solving Contextual Bandit Problems Using Decision Trees  
*Ryan McNellis, Adam N. Elmachtoub, Sechan Oh, Marek Petrik*
- 14:50** 180: Analysis of Thompson Sampling for Stochastic Sleeping Bandits  
*Aritra Chatterjee, Ganesh Ghalme, 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  
*Chunlai Zhou, Fabio Cuzzolin*
- 64: Complexity of Solving Decision Trees with Skew-Symmetric Bilinear Utility  
*Hugo Gilbert, Olivier Spanjaard*
- 72: Stochastic Segmentation Trees for Multiple Ground Truths  
*Jake Snell and Richard S. Zemel*
- 83: Efficient Online Learning for Optimizing Value of Information:  
Theory and Application to Interactive Troubleshooting  
*Yuxin Chen, Jean-Michel Renders, Morteza Haghiri Chehreghani, Andreas Krause*
- 97: Counting Markov Equivalence Classes by Number of Immoralities  
*Adityanarayanan Radhakrishnan, Liam Solus, Caroline Uhler*
- 130: Real-Time Resource Allocation for Tracking Systems  
*Yash Satsangi, Shimon Whiteson, Frans Oliehoek, Henri Bouma*
- 134: Synthesis of Strategies in Influence Diagrams  
*Manuel Luque, Manuel Arias, Francisco Javier Díez*
- 135: Embedding Senses via Dictionary Bootstrapping  
*Byungkon Kang, Kyung-Ah Sohn*
- 138: Importance Sampled Stochastic Optimization for Variational Inference  
*Joseph Sakaya, Arto Klami*
- 155: Multi-dueling Bandits with Dependent Arms  
*Yanan Sui, Vincent Zhuang, Joel W. Burdick, Yisong Yue*
- 185: Convex-constrained Sparse Additive Modeling and Its Extensions  
*Junming Yin, Yaoliang Yu*
- 239: Stein Variational Policy Gradient  
*Yang Liu, Prajit Ramachandran, Qiang Liu, Jian Peng*
- 269: Causal Discovery from Temporally Aggregated Time Series  
*Mingming Gong, Kun Zhang, Bernhard Schölkopf, Clark Glymour, Dacheng Tao*
- 280: Efficient solutions for Stochastic Shortest Path Problems with Dead Ends  
*Felipe Trevizan, Florent Teichteil-Knigsbuch, Sylvie Thiebaux*
- 281: Probabilistic Program Abstractions  
*Steven Holtzen, Todd Millstein, Guy Van den Broeck*
- 286: Communication-Efficient Distributed Primal-Dual Algorithm  
for Saddle Point Problem  
*Yaodong Yu, Sulin Liu, Sinno Jialin Pan*
- 298: Robust Model Equivalence using Stochastic Bisimulation  
for N-Agent Interactive DIDs  
*Muthukumaran Chandrasekaran, Junhuan Zhang, Prashant Doshi, Yifeng Zeng*
- 306: Adversarial Sets for Regularising Neural Link Predictors  
*Pasquale Minervini, Thomas Demeester, Tim Rocktäschel, Sebastian Riedel*

**15:40 - 16:00 Coffee break**

**16:00 - 16:45 Business Meeting**

**16:00 - 18:00 Poster Session 3**

Will include all papers presented on Monday.

---

# Invited Speakers

---

## Keynote Lecture: Learning and Inference with Expectations

*Amir Globerson, School of Computer Science, Tel Aviv University*

Probabilistic models can be characterized using expectations of functions of random variables. This can in turn be used to devise methods for learning representations, performing probabilistic inference, and discriminative learning. The talk will survey such methods, highlighting both theoretical results and their practical implications.

### Biographical details

Prof. Globerson received his BSc in computer science and physics in 1997 from the Hebrew University, and his PhD in computational neuroscience from the Hebrew University in 2006. After his PhD, he was a postdoctoral fellow at the University of Toronto and a Rothschild postdoctoral fellow at MIT. He joined the Hebrew University school of computer science in 2008, and moved to the Tel Aviv University School of Computer Science in 2015. Prof. Globerson's research interests include machine learning, probabilistic inference, convex optimization, neural computation and natural language processing. He is an associate editor for the Journal of Machine Learning Research, and the Associate Editor in Chief for the IEEE Transactions on Pattern Analysis and Machine Intelligence. His work has received several prizes including five paper awards (two at NIPS, two at UAI, and one at ICML), as well as one runner up for best paper at ICML. His research has been supported by several grants and awards from ISF, BSF, GIF, Intel, HP, and Google. In 2015 he was a visiting scientist at Google Mountain View.

## Keynote Lecture: Machine Learning for Healthcare Data

*Katherine A. Heller, Department of Statistical Science, Duke University*

We will discuss multiple ways in which healthcare data is acquired and machine learning methods are currently being introduced into clinical settings. This will include: 1) Modeling disease trends and other predictions, including joint predictions of multiple conditions, from electronic health record (EHR) data using Gaussian processes. 2) Predicting surgical complications and transfer learning methods for combining databases 3) Using mobile apps and integrated sensors for improving the granularity of recorded health data for chronic conditions and 4) The combination of mobile app and social network information in order to predict the spread of contagious disease. Current work in these areas will be presented and the future of machine learning contributions to the field will be discussed.

### Biographical details

Katherine is an Assistant Professor at Duke University, in the Department of Statistical Science and at the Center for Cognitive Neuroscience. Prior to joining Duke she was an NSF Postdoctoral Fellow, in the Computational Cognitive Science group at MIT, and an EPSRC Postdoctoral Fellow at the University of Cambridge. Her Ph.D. is from the Gatsby Unit, where her advisor was Zoubin Ghahramani. Katherine's research interests lie in the fields of machine learning and Bayesian statistics. Specifically, she develops new methods and models to discover

latent structure in data, including cluster structure, using Bayesian nonparametrics, hierarchical Bayes, techniques for Bayesian model comparison, and other Bayesian statistical methods. She applies these methods to problems in the brain and cognitive sciences, where she strives to model human behavior, including human categorization and human social interactions.

## **Keynote Lecture: Intelligent Robots in an Uncertain World**

*Leslie Pack Kaelbling, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology*

The fields of AI and robotics have made great improvements in many individual subfields, including in motion planning, symbolic planning, reasoning under uncertainty, perception, and learning. Our goal is to develop an integrated approach to solving very large problems that are hopelessly intractable to solve optimally. We make a number of approximations during planning, including serializing subtasks, factoring distributions, and determinizing stochastic dynamics, but regain robustness and effectiveness through a continuous state-estimation and replanning process. I will describe our application of these ideas to an end-to-end mobile manipulation system, as well as ideas for current and future work on improving correctness and efficiency through learning.

### **Biographical details**

Leslie is a Professor at MIT. She has an undergraduate degree in Philosophy and a PhD in Computer Science from Stanford, and was previously on the faculty at Brown University. She was the founding editor-in-chief of the Journal of Machine Learning Research. Her goal is to make intelligent robots, although she is not, herself, a robot.

## **Keynote Lecture: Snorkel: Beyond Hand-labeled Data**

*Christopher Ré, Department of Computer Science, Stanford University*

This talk describes Snorkel, whose goal is to make routine machine learning tasks dramatically easier. Snorkel focuses on a key bottleneck in the development of machine learning systems: the lack of large training datasets. In Snorkel, a user implicitly defines large training sets by writing simple programs that label data, instead of tediously hand-labeling individual data items. Snorkel then combines these programs using a generative model to produce probabilistic training labels. This talk will cover the underlying theory, including methods to learn both the parameters and structure of generative models without labeled data, and new convergence guarantees for Gibbs sampling and accelerated non-convex optimization. Additionally we'll describe our preliminary evidence that the Snorkel approach may allow a broader set of users to train machine learning models more easily than previous approaches. Snorkel and its predecessor DeepDive are in daily use by scientists in areas including genomics and drug repurposing, by a number of companies involved in various forms of search, and by law enforcement in the fight against human trafficking. DeepDive and Snorkel are open source on github and available from `DeepDive.Stanford.edu` and `Snorkel.Stanford.edu`. Technical blog posts are available at `Snorkel.Stanford.edu`.

### **Biographical details**

Christopher (Chris) Ré is an assistant professor in the Department of Computer Science at Stanford University in the InfoLab who is affiliated with the Statistical Machine Learning

Group, Pervasive Parallelism Lab, and Stanford AI Lab. His work's goal is to enable users and developers to build applications that more deeply understand and exploit data. His contributions span database theory, database systems, and machine learning, and his work has won best paper at a premier venue in each area, respectively, at PODS 2012, SIGMOD 2014, and ICML 2016. In addition, work from his group has been incorporated into major scientific and humanitarian efforts, including the IceCube neutrino detector, PaleoDeepDive and MEMEX in the fight against human trafficking, and into commercial products from major web and enterprise companies. He received a SIGMOD Dissertation Award in 2010, an NSF CAREER Award in 2011, an Alfred P. Sloan Fellowship in 2013, a Moore Data Driven Investigator Award in 2014, the VLDB early Career Award in 2015, the MacArthur Foundation Fellowship in 2015, and an Okawa Research Grant in 2016.

## **Keynote Lecture: Two Current Analysis Challenges: Single Cell Omics and Nanopore Long-read Sequence Data**

*Terry Speed, Bioinformatics Division, Walter and Eliza Hall Institute of Medical Research*

Abstract. Two of the most exciting recent developments in my field are single cell omics assays, where omics here currently encompasses genomics, transcriptomics and epigenomics, and the long-read DNA sequencing capability introduced by Oxford Nanopore Technologies. Each leads to lots of interesting data permitting biomedical researchers to address questions previously beyond their reach. Each offers numerous opportunities and challenges to those who wish to, or must dive into the data to help the researchers answer these new questions. And each seems likely to drive theoretical developments in areas of interest to participants in UAI 2017. I've been lucky enough to be associated with research projects in these two areas. Building on my experience with them, I'll discuss some of the opportunities, challenges and areas for theoretical development that I see.

### **Biographical details**

Terry Speed completed a BSc (Hons) in mathematics and statistics at the University of Melbourne and a PhD in mathematics and Dip Ed at Monash University. He has held appointments at the University of Sheffield, U.K., the University of Western Australia in Perth, and the University of California at Berkeley, and with the CSIRO in Canberra. In 1997 he took up an appointment with the Walter & Eliza Hall Institute of Medical Research, where he is now an Honorary Fellow and lab head in the Bioinformatics Division. His research interests lie in the application of statistics and bioinformatics to genetics and genomics, and related fields such as proteomics, metabolomics and epigenomics, with a focus on cancer and epigenetics.

---

# Workshop Program

---

**Wednesday August**

**Rooms E5.2, E5.3, E5.4**

## **Fourteenth Annual Bayesian Modeling Applications Workshop**

- 9:00** Invited talk: Probabilistic reasoning with complex heterogeneous observations and applications in geology and medicine  
*David Poole*
- 9:45** A Validation Tool for Conditional Probability Tables (CPT) for Physics Playground  
*Russell Almond, Seyfullah Tingir, Xi Lu, Chen Sun, Seyedahmad Rahimi*
- 10:10** A Delphi-based Knowledge Engineering Approach to Bayesian Networks  
*Steven Mascaro, Ann Nicholson, Shreshth Thakur, Kevin Korb*
- 10:35** **Coffee Break**
- 10:50** Tutorial: OpenMarkov, an open-source tool for probabilistic graphical models  
*Manuel Arias, Francisco Javier Dez, Manuel Luque, Jorge Pérez-Martín*
- 11:35** PROFET: Construction and Inference of DBNs Based on Mathematical Models  
*Hamda Ajmal, Michael Madden, Catherine Enright*
- 12:00** IOOBN: a Modeling Tool using Object Oriented Bayesian Networks with Inheritance  
*Md. Samiullah, Thao Xuan Hoang, David Albrecht, Ann Nicholson, Kevin Korb*
- 12:30** **Lunch Break**
- 14:00** Tabular views of Bayesian networks  
*Russell Almond*
- 14:25** DiLeMMa - Distributed Learning with Markov Chain Monte Carlo Algorithms  
*Ali Zaidi*
- 14:50** **Quo vadis: Bayesian models in the age of “deep everything”  
Community forum**
- 15:10** **End of Workshop**

## 7th Causality Workshop: Learning, Inference, and Decision-Making

8:45 Welcome & opening remarks

9:00 Invited Talk (TBA)

9:55 **Whats hot in sister conferences (UAI, ICML, IJCAI)**

Interpreting and using CPDAGs with background knowledge (UAI-17)

*Emilija Perkovic, Markus Kalisch, Marloes H. Maathuis*

Structure Learning of Linear Gaussian Structural Equation Models  
with Weak Edges (UAI-17)

*Marco Eigenmann, Preetam Nandy, Marloes Maathuis*

10:35 **Break & Posters**

10:55 **Contributed Talks**

Causal Consistency of Structural Equation Models

*Paul K. Rubenstein, Sebastian Weichwald, Stephan Bongers,  
Joris M. Mooij, Dominik Janzing, Moritz Grosse-Wentrup, Bernhard Schoelkopf*

Causal Discovery in the Presence of Measurement Noise: Identifiability Conditions

*Kun Zhang, Mingming Gong, Joseph Ramsey, Kayhan Batmanghelich,  
Peter Spirtes, Clark Glymour*

SAT-Based Causal Discovery under Weaker Assumptions

*Zhalama, Jiji Zhang, Frederick Eberhardt, Wolfgang Mayer*

Learning Dynamic Structure from Undersampled Data

*John W. Cook, David Danks, Sergey M. Plis*

12:15 **Lunch & Posters**

14:00 Invited Talk (TBA)

14:55 **Whats hot in sister conferences (UAI, ICML, IJCAI) (TBA)**

16:00 **Break & Posters**

17:00 **Contributed Talks**

Algebraic Equivalence Class Selection for Linear Structural Equation Models

*Thijs van Ommen, Joris M. Mooij*

Counting Markov Equivalence Classes by Number of Immoralities

*Liam Thomas Solus, Caroline Uhler, Adit Radhakrishnan*

Probabilistic Active Learning of Functions in Structural Causal Models

*Paul K Rubenstein, Bernhard Schoelkopf, Ilya Tolstikhin and Philipp Hennig*

Overcoming Impediments to Estimation with Missing-data

*Karthika Mohan, Judea Pearl*

18:20 **Panel & Closing Discussion**

## Seventh International Workshop on Statistical Relational AI

**9:00** Welcome and introduction

**9:10** Invited Talk: TBA

**10:10** Poster Spotlights (2-minute)

**10:30** Break/Poster Session

**11:30** Invited Talk: TBA

**12:30** Lunch Break

**14:00** Contributed Talks

**15:00** Poster Session

# Fourteenth Annual Bayesian Modeling Applications Workshop

<http://bmaw2017.azurewebsites.net/>

BMAW 2017 aims to foster discussion on the challenges of building working applications of probabilistic methods while considering stakeholders, user interaction, tools, knowledge elicitation, learning, validation, system integration, and deployment. This year we especially appreciate the degree to which the practicality of sharing application development experience comes down to sharing reusable software. The workshop will correspondingly feature real-world Bayesian applications, where teams will explain what tools they adopted to facilitate their work. We will focus the discussion on best practices and tools for sharing data, code and models in easily reproducible ways.

**Invited Speakers** David Poole, UBC

## Organizers

John-Mark Agosta, Microsoft

Tomas Singliar, Microsoft

## Program Committee

Fabio Cozman, U Sao Paulo

José Antonio Gémez Martn, UCLM

Feng Chen, SUNY Albany

Kathryn B. Laskey, GMU

Paulo Cesar Costa, GMU

Marek J Druzdzal, Pitt

Helge Langseth, NTNU

Dennis Buede, Innovative Decisions

Russell Almond, FSU

Rommel Novaes Carvalho, Office of Comptroller General of Brazil

Francisco Javier Diez, UNED Spain

Ann Nicholson, Monash U

V Anne Smith, Andrews U

Maria Julia Flores Gallego, UCLM

Ole Mengshoel, CMU

Thomas Dyhre Nielsen, Aalborg U

Enrique Sucar, INAOE

## 7th Causality Workshop: Learning, Inference, and Decision-Making

<https://www.cs.purdue.edu/homes/eb/causal-uai17/>

key ingredient in building AI systems that are autonomous and can act efficiently in complex and uncertain environments. It's also important to the process of scientific discovery since it underpins how explanations are constructed and the scientific method.

Not surprisingly, the tasks of learning and reasoning with causal-effect relationships have attracted great interest in the artificial intelligence and machine learning communities. This effort has led to a very general theoretical and algorithmic understanding of what causality means and under what conditions it can be inferred. These results have started to percolate through more applied fields that generate the bulk of the data currently available, ranging from genetics to medicine, from psychology to economics.

This one-day workshop will explore causal inference in a broad sense through a set of invited talks, open problems sessions, presentations, and a poster session. In this workshop, we will focus on the foundational side of causality on the one hand, and challenges presented by practical applications on the other. By and large, we welcome contributions from all areas relating to the study of causality.

### Organizers

Elias Bareinboim, Purdue (Chair)

Kun Zhang, CMU

Caroline Uhler, MIT

Jiji Zhang, Lingnan University

Dominik Janzing, MPI

# Seventh International Workshop on Statistical Relational AI

<http://www.starai.org/2017/>

The purpose of the Statistical Relational AI (StarAI) workshop is to bring together researchers and practitioners from two fields: logical (or relational) AI and probabilistic (or statistical) AI. These fields share many key features and often solve similar problems and tasks. Until recently, however, research in them has progressed independently with little or no interaction. The fields often use different terminology for the same concepts and, as a result, keeping-up and understanding the results in the other field is cumbersome, thus slowing down research. Our long term goal is to change this by achieving a synergy between logical and statistical AI.

## Organizers

Angelika Kimmig (Cardiff)

David Poole (UBC)

Jay Pujara (UCSC)

Tim Rocktschel (Oxford)

## Program Committee

Hendrik Blockeel (KU Leuven)

Jaesik Choi (UNIST)

Vitor Santos Costa (Universidade do Porto)

Rajarashi Das (UMass)

Lucas Sterckx (Ghent)

Richard Evans (Google DeepMind)

Vibhav Gogate (UT Dallas)

Kristian Kersting (TU Darmstadt)

Daniel Lowd (University of Oregon)

Pasquale Minervini (University College London)

Mathias Niepert (NEC Labs)

Sriram Natarajan (Indiana)

Scott Sanner (Oregon State University)

Sameer Singh (University of Washington)

Dhanya Sridhar (UC Santa Cruz)

Guy Van den Broeck (UCLA)

Johannes Welbl (University College London)

---

# Paper Abstracts

---

## 6: Near-Orthogonality Regularization in Kernel Methods

*Pengtao Xie, Barnabas Poczos, Eric Xing*

Kernel methods perform nonlinear learning in high-dimensional reproducing kernel Hilbert spaces (RKHSs). Even though their large model-capacity leads to high representational power, it also incurs substantial risk of overfitting. To alleviate this problem, we propose a new regularization approach, near-orthogonality regularization, which encourages the RKHS functions to be close to being orthogonal. This effectively imposes a structural constraint over the function space, which reduces model complexity and can improve generalization performance. Besides, encouraging orthogonality reduces the redundancy among functions, which hence can reduce model size without compromising modeling power and better capture infrequent patterns in the data. Here, we define a family of orthogonality-promoting regularizers by encouraging the Gram matrix of the RKHS functions to be close to an identity matrix where the closeness is measured by Bregman matrix divergences. We apply these regularizers to two kernel methods, and develop an efficient ADMM-based algorithm to solve the regularized optimization problems. We analyze how near-orthogonality affects the generalization performance of kernel methods. Our results suggest that the closer the functions are to being orthogonal, the smaller the generalization error is. Experiments demonstrate the efficacy of near-orthogonality regularization in kernel methods.

## 11: Causal Consistency of Structural Equation Models

*Paul K. Rubenstein\*, Sebastian Weichwald\*, Stephan Bongers, Joris M. Mooij, Dominik Janzing, Moritz Grosse-Wentrup, Bernhard Schölkopf*

Complex systems can be modelled at various levels of detail. Ideally, causal models of the same system should be consistent with one another in the sense that they agree in their predictions of the effects of interventions. We formalise this notion of consistency in the case of Structural Equation Models (SEMs) by introducing exact transformations between SEMs. This provides a general language to consider, for instance, the different levels of description in the following three scenarios: (a) models with large numbers of variables versus models in which the ‘irrelevant’ or unobservable variables have been marginalised out; (b) micro-level models versus macro-level models in which the macro-variables are aggregate features of the micro-variables; (c) dynamical time series models versus models of their stationary behaviour. Our analysis stresses the importance of well specified interventions in the causal modelling process and sheds light on the interpretation of cyclic SEMs.

## 12: FROSH: FasteR Online Sketching Hashing

*Xixian Chen, Irwin King, Michael R. Lyu*

Many hashing methods, especially those that are in the data-dependent category with good learning accuracy, are still inefficient when dealing with three critical problems in modern data analysis. First, data usually come in a streaming fashion, but most of the existing hashing methods are batch-based models. Second, when data become huge, the extensive computational time, large space requirement, and multiple passes to load the data into memory will be prohibitive. Third, data often lack sufficient label information. Although the recently proposed Online Sketching Hashing (OSH) is promising to alleviate all three issues mentioned above, its training procedure still suffers from a high time complexity. In this paper, we propose a FasteR Online Sketching Hashing (FROSH) method to make the training process faster. Compared with OSH, we leverage fast transform to sketch data more compactly. Particularly, we derive independent transformations to guarantee the sketching accuracy, and design a novel implementation to make such transformations applicable to online data sketching without increasing the space cost. We rigorously prove that our method can yield a comparable learning accuracy with a lower time complexity and an equal space cost compared with OSH. Finally, extensive experiments on synthetic and real-world datasets demonstrate the excellent performance of our method.

## 16: Self-Discrepancy Conditional Independence Test

*Sanghack Lee, Vasant Honavar*

Tests of conditional independence (CI) of random variables play an important role in machine learning and causal inference. Of particular interest are kernel-based CI tests which allow us to test for independence among random variables with complex distribution functions. The efficacy of a CI test is measured in terms of its power and its calibratedness. We show that the Kernel CI Permutation Test (KCIPT) suffers from a loss of calibratedness as its power is increased by increasing the number of bootstraps. To address this limitation, we propose a novel CI test, called Self-Discrepancy Conditional Independence Test (SDCIT). SDCIT uses a test statistic that is a modified unbiased estimate of maximum mean discrepancy (MMD), the largest difference in the means of features of the given sample and its permuted counterpart in the kernel-induced Hilbert space. We present results of experiments that demonstrate SDCIT is, relative to the other methods: (i) competitive in terms of its power and calibratedness, outperforming other methods when the number of conditioning variables is large; (ii) more robust with respect to the choice of the kernel function; and (iii) competitive in run time.

## 26: Shortest Path under Uncertainty: Exploration versus Exploitation

*Zhan Wei Lim, David Hsu, Wee Sun Lee*

In the Canadian Traveler Problem (CTP), a traveler seeks a shortest path to a destination through a road network, but unknown to the traveler, some roads may be blocked. This paper studies the Bayesian CTP (BCTP), in which road states are correlated with known prior probabilities and the traveler can infer the states of an unseen road from past observations of other correlated roads. As generalized shortest-path problems, CTP and BCTP have important practical applications. We show that BCTP is NP-complete and give a polynomial-time approximation algorithm, Hedged Shortest Path under Determinization (HSPD), which approximates an optimal solution with a polylogarithmic factor. Preliminary experiments show promising results. HSPD outperforms a widely used greedy algorithm and a state-of-the-art

UCT-based search algorithm for CTP, especially when significant exploration is required.

## 32: Provable Inductive Robust PCA via Iterative Hard Thresholding

*U.N. Niranjan, Arun Rajkumar, Theja Tulabandhula*

The robust PCA problem, wherein, given an input data matrix that is the superposition of a low-rank matrix and a sparse matrix, we aim to separate out the low-rank and sparse components, is a well-studied problem in machine learning. One natural question that arises is that, as in the inductive setting, if features are provided as input as well, can we hope to do better? Answering this in the affirmative, the main goal of this paper is to study the robust PCA problem while incorporating feature information. In contrast to previous works in which recovery guarantees are based on the convex relaxation of the problem, we propose a simple iterative algorithm based on hard-thresholding of appropriate residuals. Under weaker assumptions than previous works, we prove the global convergence of our iterative procedure; moreover, it admits a much faster convergence rate and lesser computational complexity per iteration. In practice, through systematic synthetic and real data simulations, we confirm our theoretical findings regarding improvements obtained by using feature information.

## 35: Learning with Confident Examples: Rank Pruning for Robust Classification with Noisy Labels

*Curtis G. Northcutt, Tailin Wu, Isaac L. Chuang*

$\tilde{P}\tilde{N}$  learning is the problem of binary classification when training examples may be mislabeled (flipped) uniformly with noise rate  $\rho_1$  for positive examples and  $\rho_0$  for negative examples. We propose Rank Pruning (RP) to solve  $\tilde{P}\tilde{N}$  learning and the open problem of estimating the noise rates. Unlike prior solutions, RP is efficient and general, requiring  $\mathcal{O}(T)$  for any unrestricted choice of probabilistic classifier with  $T$  fitting time. We prove RP achieves consistent noise estimation and equivalent empirical risk as learning with uncorrupted labels in ideal conditions, and derive closed-form solutions when conditions are non-ideal. RP achieves state-of-the-art noise estimation and F1, error, and AUC-PR for both MNIST and CIFAR datasets, regardless of the amount of noise. To highlight, RP with a CNN classifier can predict if a MNIST digit is a “1” or “not 1” with only 0.25% error, and 0.46% error across all digits, even when 50% of positive examples are mislabeled and 50% of observed positive labels are mislabeled negative examples.

## 37: Monte-Carlo Tree Search using Batch Value of Perfect Information

*Shahaf S. Shperberg, Solomon Eyal Shimony, Ariel Felner*

This paper focuses on the selection phase of Monte-Carlo Tree Search (MCTS). We define batch value of perfect information (BVPI) in game trees as a generalization of value of computation as proposed by Russell and Wefald, and use it for selecting nodes to sample in MCTS. We show that computing the BVPI is NP-hard, but it can be approximated in polynomial time. In addition, we propose methods that intelligently find sets of fringe nodes with high BVPI, and quickly select nodes to sample from these sets. We apply our new BVPI methods to partial

game trees, both in a stand-alone set of tests, and as a component of a full MCTS algorithm. Empirical results show that our BVPI methods outperform existing node-selection methods for MCTS in different scenarios.

## 45: Submodular Variational Inference for Network Reconstruction

*Lin Chen, Forrest W. Crawford, Amin Karbasi*

In real-world and online social networks, individuals receive and transmit information in real time. Cascading information transmissions (e.g. phone calls, text messages, social media posts) may be understood as a realization of a diffusion process operating on the network, and its branching path can be represented by a directed tree. The process only traverses and thus reveals a limited portion of the edges. The network reconstruction/inference problem is to infer the unrevealed connections. Most existing approaches derive a likelihood and attempt to find the network topology maximizing the likelihood, a problem that is highly intractable. In this paper, we focus on the network reconstruction problem for a broad class of real-world diffusion processes, exemplified by a network diffusion scheme called respondent-driven sampling (RDS). We prove that under realistic and general models of network diffusion, the posterior distribution of an observed RDS realization is a Bayesian log-submodular model. We then propose VINE (Variational Inference for Network rEconstruction), a novel, accurate, and computationally efficient variational inference algorithm, for the network reconstruction problem under this model. Crucially, we do not assume any particular probabilistic model for the underlying network. VINE recovers any connected graph with high accuracy as shown by our experimental results on real-life networks.

## 48: Inverse Reinforcement Learning via Deep Gaussian Process

*Ming Jin, Andreas Damianou, Pieter Abbeel, Costas Spanos*

We propose a new approach to inverse reinforcement learning (IRL) based on the deep Gaussian process (deep GP) model, which is capable of learning complicated reward structures with few demonstrations. Our model stacks multiple latent GP layers to learn abstract representations of the state feature space, which is linked to the demonstrations through the Maximum Entropy learning framework. Incorporating the IRL engine into the nonlinear latent structure renders existing deep GP inference approaches intractable. To tackle this, we develop a non-standard variational approximation framework which extends previous inference schemes. This allows for approximate Bayesian treatment of the feature space and guards against overfitting. Carrying out representation and inverse reinforcement learning simultaneously within our model outperforms state-of-the-art approaches, as we demonstrate with experiments on standard benchmarks (“object world”, “highway driving”) and a new benchmark (“binary world”).

## 49: A Kernel Conditional Independence Test for Relational Data

*Sanghack Lee, Vasant Honavar*

Conditional independence (CI) tests play a central role in statistical inference, machine learning, and causal discovery. Most existing CI tests assume that the samples are independently and identically distributed (i.i.d.). However, this assumption often does not hold in the case of relational data. We define Relational Conditional Independence (RCI), a generalization of CI

to the relational setting. We show how, under a set of structural assumptions, we can test for RCI by reducing the task of testing for RCI on non-i.i.d. data to the problem of testing for CI on several data sets each of which consists of i.i.d. samples. We develop Kernel Relational CI test (KRCIT), a nonparametric test as a practical approach to testing for RCI by relaxing the structural assumptions used in our analysis of RCI. We describe results of experiments with synthetic relational data that show the benefits of KRCIT relative to traditional CI tests that don't account for the non-i.i.d. nature of relational data.

## 50: AutoGP: Exploring the Capabilities and Limitations of Gaussian Process Models

*Karl Krauth, Edwin V. Bonilla, Kurt Cutajar, Maurizio Filippone*

We investigate the capabilities and limitations of Gaussian process (GP) models by jointly exploring three complementary directions: (i) scalable and statistically efficient inference; (ii) flexible kernels; and (iii) objective functions for hyperparameter learning alternative to the marginal likelihood. Our approach outperforms all previous GP methods on the MNIST dataset; performs comparatively to kernel-based methods using the RECTANGLES-IMAGE dataset; and breaks the 1% error-rate barrier in GP models on the MNIST8M dataset, while showing unprecedented scalability (8 million observations) in GP classification. Overall, our approach represents a significant breakthrough in kernel methods and GP models, bridging the gap between deep learning and kernel machines.

## 54: An Efficient Minibatch Acceptance Test for Metropolis-Hastings

*Daniel Seita, Xinlei Pan, Haoyu Chen, John Canny*

We present a novel Metropolis-Hastings method for large datasets that uses small expected-size minibatches of data. Previous work on reducing the cost of Metropolis-Hastings tests yield variable data consumed per sample, with only constant factor reductions versus using the full dataset for each sample. Here we present a method that can be tuned to provide arbitrarily small batch sizes, by adjusting either proposal step size or temperature. Our test uses the noise-tolerant Barker acceptance test with a novel additive correction variable. The resulting test has similar cost to a normal SGD update. Our experiments demonstrate several order-of-magnitude speedups over previous work.

## 56: The total belief theorem

*Chunlai Zhou, Fabio Cuzzolin*

In this paper, motivated by the treatment of conditional constraints in the data association problem, we state and prove the generalisation of the law of total probability to belief functions, as finite random sets. Our results apply to the case in which Dempsters conditioning is employed. We show that the solution to the resulting total belief problem is in general not unique, whereas it is unique when the a-priori belief function is Bayesian. Examples and case studies underpin the theoretical contributions. Finally, our results are compared to previous related work on the generalisation of Jeffreys rule by Spies and Smets.

## 62: Near-Optimal Interdiction of Factored MDPs

*Swetasudha Panda, Yevgeniy Vorobeychik*

Stackelberg games have been widely used to model interactions between attackers and defenders in a broad array of security domains. One related approach involves plan interdiction, whereby a defender chooses a subset of actions to block (remove), and the attacker constructs an optimal plan in response. In previous work, this approach has been introduced in the context of Markov decision processes (MDPs). The key challenge, however, is that the state space of MDPs grows exponentially in the number of state variables. We propose a novel scalable MDP interdiction framework which makes use of factored representation of state, using Fourier representation for representing a value function over a boolean space. We demonstrate that our approach is significantly more scalable, while resulting in near-optimal interdiction decisions.

## 64: Complexity of Solving Decision Trees with Skew-Symmetric Bilinear Utility

*Hugo Gilbert, Olivier Spanjaard*

We study the complexity of solving decision trees with a Skew-Symmetric Bilinear (SSB) utility function. The SSB model is an extension of Expected Utility (EU) with enhanced descriptive possibilities. Unlike EU, the optimality principle does not hold for SSB, which makes its optimization trickier. We show that determining an SSB optimal plan is NP-hard if one only considers deterministic plans while it is polynomial time if one allows randomized plans. With the Weighted EU model (a special case of SSB), the problem becomes polynomial in both settings. Our numerical tests show the operability of the methods.

## 65: A Fast Stochastic Riemannian Eigensolver

*Zhiqiang Xu, Yiping Ke, Xin Gao*

We propose a fast stochastic Riemannian gradient eigensolver for a real and symmetric matrix, and prove its local, eigengap-dependent and linear convergence. The fast convergence is brought by deploying the variance reduction technique which was originally developed for the Euclidean strongly convex problems. In this paper, this technique is generalized to Riemannian manifolds for solving the geodesically non-convex problem of finding a group of top eigenvectors of such a matrix. We first propose the general variance reduction form of the stochastic Riemannian gradient, giving rise to the stochastic variance reduced Riemannian gradient method (SVRRG). It turns out that the operation of vector transport is necessary in addition to using Riemannian gradients and retraction operations. We then specialize it to the problem in question resulting in our SVRRG-EIGS algorithm. We are among the first to propose and analyze the generalization of the stochastic variance reduced gradient (SVRG) to Riemannian manifolds. As an extension of the linearly convergent VR-PCA, it is significant and nontrivial for the proposed algorithm to theoretically achieve a further speedup and empirically make a difference, due to our respect to the inherent geometry of the problem.

## 69: Determinantal Point Processes for Mini-Batch Diversification

*Cheng Zhang, Hedvig Kjellstrom, Stephan Mandt*

We study a mini-batch diversification scheme for stochastic gradient descent (SGD). While classical SGD relies on uniformly sampling data points to form a mini-batch, we propose a non-uniform sampling scheme based on the Determinantal Point Process (DPP). The DPP relies on a similarity measure between data points and gives low probabilities to mini-batches which contain redundant data, and higher probabilities to mini-batches with more diverse data. This simultaneously balances the data and leads to stochastic gradients with lower variance. We term this approach Diversified Mini-Batch SGD (DM-SGD). We show that regular SGD and a biased version of stratified sampling emerge as special cases. Furthermore, DM-SGD generalizes stratified sampling to cases where no discrete features exist to bin the data into groups. We show experimentally that our method results more interpretable and diverse features in unsupervised setups, and in better classification accuracies in supervised setups.

## 70: Bayesian Inference of Log Determinants

*J Fitzsimons, K Cutajar, M Filippone, M Osborne, S Roberts*

The log-determinant of a kernel matrix appears in a variety of machine learning problems, ranging from determinantal point processes and generalized Markov random fields, through to the training of Gaussian processes. Exact calculation of this term is often intractable when the size of the kernel matrix exceeds a few thousands. In the spirit of probabilistic numerics, we reinterpret the problem of computing the log-determinant as a Bayesian inference problem. In particular, we combine prior knowledge in the form of bounds from matrix theory and evidence derived from stochastic trace estimation to obtain probabilistic estimates for the log-determinant and its associated uncertainty within a given computational budget. Beyond its novelty and theoretic appeal, the performance of our proposal is competitive with state-of-the-art approaches to approximating the log-determinant, while also quantifying the uncertainty due to budget-constrained evidence.

## 72: Stochastic Segmentation Trees for Multiple Ground Truths

*Jake Snell, Richard S. Zemel*

A strong machine learning system should aim to produce the full range of valid interpretations rather than a single mode in tasks involving inherent ambiguity. This is particularly true for image segmentation, in which there are many sensible ways to partition an image into regions. We formulate a tree-structured probabilistic model, the stochastic segmentation tree, that represents a distribution over segmentations of a given image. We train this model by optimizing a novel objective that quantifies the degree of match between statistics of the model and ground truth segmentations. Our method allows learning of both the parameters in the tree and the structure itself. We demonstrate on two datasets, including the challenging Berkeley Segmentation Dataset, that our model is able to successfully capture the range of ground truths and to produce novel plausible segmentations beyond those found in the data.

## 75: Fast Amortized Inference and Learning in Log-linear Models with Randomly Perturbed Nearest Neighbor Search

*Stephen Mussmann, Daniel Levy, Stefano Ermon*

Inference in log-linear models scales linearly in the size of output space in the worst-case. This is often a bottleneck in natural language processing and computer vision tasks when the output space is feasibly enumerable but very large. We propose a method to perform inference in log-linear models with sublinear amortized cost. Our idea hinges on using Gumbel random variable perturbations and a pre-computed Maximum Inner Product Search data structure to access the most-likely elements in sublinear amortized time. Our method yields provable runtime and accuracy guarantees. Further, we present empirical experiments on ImageNet and Word Embeddings showing significant speedups for sampling, inference, and learning in log-linear models.

## 78: Branch and Bound for Regular Bayesian Network Structure Learning

*Joe Suzuki, Jun Kawahara*

We consider efficient Bayesian network structure learning (BNSL) using branch and bound. Thus far, it is known that efficient pruning rules have been found for minimizing the description length (MDL) rather than for maximizing the posterior probability such as the Bayesian Dirichlet equivalent uniform (BDeu) that has been used most often. In any BNSL, given data, the simplest structure should be chosen if structures are equally fit to the data. We show that quotient scores satisfy such regularity while the BDeu does not. In this paper, by utilizing the regularity, we propose a pruning rule for maximizing the posterior probability based on the quotient scores, and prove that the bound is tighter than existing one. We examine by experiments that the proposed bound is applied significantly more often than existing one for the BDeu as well. Besides, we prove that the MDL satisfies regularity and that the proposed bound behaves like that of the MDL, and observe that the proposed bound is almost as efficient as one for the MDL. Our novel insight is that BNSL can be efficient by utilizing regularity for constructing a pruning rule.

## 81: Effective sketching methods for value function approximation

*Yangchen Pan, Erfan Sadeqi Azer, Martha White*

High-dimensional representations, such as radial basis function networks or tile coding, are common choices for policy evaluation in reinforcement learning. Learning with such high-dimensional representations, however, can be expensive, particularly for matrix methods, such as least-squares temporal difference learning or quasi-Newton methods that approximate matrix step-sizes. In this work, we explore the utility of sketching for these two classes of algorithms. We highlight issues with sketching the high-dimensional features directly, which can incur significant bias. As a remedy, we demonstrate how to use sketching more sparingly, with only a left-sided sketch, that can still enable significant computational gains and the use of these matrix-based learning algorithms that are less sensitive to parameters. We empirically investigate these algorithms, in four domains with a variety of representations. Our aim is to provide insights into effective use of sketching in practice.

## 83: Efficient Online Learning for Optimizing Value of Information: Theory and Application to Interactive Troubleshooting

*Yuxin Chen, Jean-Michel Renders, Morteza Haghir Chehreghani, Andreas Krause*

We consider the optimal value of information (VoI) problem, where the goal is to sequentially select a set of tests with a minimal cost, so that one can efficiently make the best decision based on the observed outcomes. Existing algorithms are either heuristics with no guarantees, or scale poorly (with exponential run time in terms of the number of available tests). Moreover, these methods assume a known distribution over the test outcomes, which is often not the case in practice.

We propose an efficient sampling-based online learning framework to address the above issues. First, assuming the distribution over hypotheses is known, we propose a dynamic hypothesis enumeration strategy, which allows efficient information gathering with strong theoretical guarantees. We show that with sufficient amount of samples, one can identify a near-optimal decision with high probability. Second, when the parameters of the hypotheses distribution are unknown, we propose an algorithm which learns the parameters progressively via posterior sampling in an online fashion. We further establish a rigorous bound on the expected regret. We demonstrate the effectiveness of our approach on a real-world interactive troubleshooting application and show that one can efficiently make high-quality decisions with low cost.

## 86: Stochastic L-BFGS Revisited: Improved Convergence Rates and Practical Acceleration Strategies

*Renbo Zhao, William B. Haskell, Vincent Y. F. Tan*

We revisit the stochastic limited-memory BFGS (L-BFGS) algorithm. By proposing a new framework for analyzing convergence, we theoretically improve the (linear) convergence rates and computational complexities of the stochastic L-BFGS algorithms in previous works. In addition, we propose several practical acceleration strategies to speed up the empirical performance of such algorithms. We also provide theoretical analyses for most of the strategies. Experiments on large-scale logistic and ridge regression problems demonstrate that our proposed strategies yield significant improvements *vis-à-vis* competing state-of-the-art algorithms.

## 87: The Binomial Block Bootstrap Estimator for Evaluating Loss on Dependent Clusters

*Matt Barnes, Artur Dubrawski*

In this paper, we study the non-IID learning setting where samples exhibit dependency within latent clusters. Our goal is to estimate a learner’s loss on new clusters, an extension of the out-of-bag error. Previously developed cross-validation estimators are well suited to the case where the clustering of observed data is known *a priori*. However, as is often the case in real world problems, we are only given a noisy approximation of this clustering, likely the result of some clustering algorithm. This subtle yet potentially significant issue afflicts domains ranging from image classification to medical diagnostics, where naive cross-validation is an optimistically biased estimator. We present a novel bootstrap technique and corresponding cross-validation method that, somewhat counterintuitively, injects additional dependency to asymptotically

recover the loss in the independent setting.

## 93: Data-Dependent Sparsity for Subspace Clustering

*Bo Xin, Yizhou Wang, Wen Gao and David Wipf*

Subspace clustering is the process of assigning subspace memberships to a set of unlabeled data points assumed to have been drawn from the union of an unknown number of low-dimensional subspaces, possibly interlaced with outliers or other data corruptions. By exploiting the fact that each inlier point has a sparse representation with respect to a dictionary formed by all the other points, an  $\ell_1$  regularized sparse subspace clustering (SSC) method has recently shown state-of-the-art robustness and practical extensibility in a variety of applications. But there remain important lingering weaknesses. In particular, the  $\ell_1$  norm solution is highly sensitive, often in a detrimental direction, to the very types of data structures that motivate interest in subspace clustering to begin with, sometimes leading to poor segmentation accuracy. However, as an alternative source of sparsity, we argue that a certain data-dependent, non-convex penalty function can compensate for dictionary structure in a way that is especially germane to subspace clustering problems. For example, we demonstrate that this proposal displays a form of invariance to feature-space transformations and affine translations that commonly disrupt existing methods, and moreover, in important settings we reveal that its performance quality is lower bounded by the  $\ell_1$  solution. Finally, we provide empirical comparisons on popular benchmarks that corroborate our theoretical findings and demonstrate superior performance when compared to recent state-of-the-art models.

## 97: Counting Markov Equivalence Classes by Number of Immoralities

*Adityanarayanan Radhakrishnan, Liam Solus, Caroline Uhler*

Two directed acyclic graphs (DAGs) are called Markov equivalent if and only if they have the same underlying undirected graph (i.e. skeleton) and the same set of immoralities. When using observational data alone and typical identifiability assumptions, such as faithfulness, a DAG model can only be determined up to Markov equivalence. Therefore, it is desirable to understand the size and number of Markov equivalence classes (MECs) combinatorially. In this paper, we address this enumerative question using a pair of generating functions that encode the number and size of MECs on a skeleton  $G$ , and in doing so we connect this problem to classical problems in combinatorial optimization. The first generating function is a graph polynomial that counts the number of MECs on  $G$  by their number of immoralities. Using connections to the independent set problem, we show that computing a DAG on  $G$  with the maximum possible number of immoralities is NP-hard. The second generating function counts the MECs on  $G$  according to their size. Via computer enumeration, we show that this generating function is distinct for every connected graph on  $p$  nodes for all  $p \leq 10$ .

## 100: Regret Minimization Algorithms for the Follower's Behaviour Identification in Leadership Games

*Lorenzo Bisi, Giuseppe De Nittis, Francesco Trov'ò, Marcello Restelli, Nicola Gatti*

We study for the first time, a leadership game in which one agent, acting as leader, faces another

agent, acting as follower, whose behaviour is not known a priori by the leader, being one among a set of possible behavioural profiles. The main motivation is that in real-world applications the common game-theoretical assumption of perfect rationality is rarely met, and any specific assumption on bounded rationality models, if wrong, could lead to a significant loss for the leader. The question we pose is whether and how the leader can learn the behavioural profile of a follower in leadership games. This is a natural online identification problem: in fact, the leader aims at identifying the followers behavioural profile to exploit at best the potential non-rationality of the opponent, while minimizing the regret due to the initial lack of information. We propose two algorithms based on different approaches and we provide a regret analysis. Furthermore, we experimentally evaluate the pseudo-regret of the algorithms in concrete leadership games, showing that our algorithms outperform the online learning algorithms available in the state of the art.

## 105: On the Complexity of Nash Equilibrium Reoptimization

*Andrea Celli, Alberto Marchesi, Nicola Gatti*

We provide, to the best of our knowledge, the first study about reoptimization complexity of game-theoretical solutions. In a reoptimization problem, we are given an instance, its optimal solution, and a local modification, and we are asked to find the exact or an approximate solution to the modified instance. Reoptimization is crucial whenever an instance needs to be solved repeatedly and, at each repetition, its parameters may slightly change. In this paper, we focus on Nash equilibrium, being the central game-theoretical solution. We study the reoptimization of Nash equilibria satisfying some properties (i.e., maximizing/minimizing the social welfare, the utility of a player or the support size) for some different local modifications of the game (i.e., modification of a payoff or addition/removal of an action), showing that such problems are NP-hard. Furthermore, we assess the approximation complexity of the aforementioned problems, showing that it matches the complexity of the original (non-reoptimization) problems. Finally, we show that, when finding a Nash equilibrium is thought as an optimization problem, reoptimization is useful for finding approximate solutions. Specifically, it allows one to find  $\epsilon$ -Nash equilibria with smaller  $\epsilon$  than that of the solutions returned by the best known approximation algorithms.

## 106: Supervised Restricted Boltzmann Machines

*Tu Dinh Nguyen, Dinh Phung, Viet Huynh, Trung Le*

We propose in this paper the supervised restricted Boltzmann machine (sRBM), a unified framework which combines the versatility of RBM to simultaneously learn the data representation and to perform supervised learning (i.e., a nonlinear classifier or a nonlinear regressor). Unlike the current state-of-the-art classification formulation proposed for RBM in (Larochelle et al., 2012), our model is a hybrid probabilistic graphical model consisting of a distinguished generative component for data representation and a discriminative component for prediction. While the work of (Larochelle et al., 2012) typically incurs no extra difficulty in inference compared with a standard RBM, our discriminative component, modeled as a directed graphical model, renders MCMC-based inference (e.g., Gibbs sampler) very slow and unpractical for use. To this end, we further develop scalable variational inference for the proposed sRBM for both classification and regression cases. Extensive experiments on real-world datasets show that our sRBM achieves better predictive performance than baseline methods. At the same time,

our proposed framework yields learned representations which are more discriminative, hence interpretable, than those of its counterparts. Besides, our method is probabilistic and capable of generating meaningful data conditioning on specific classes a topic which is of current great interest in deep learning aiming at data generation.

## 108: Safe Semi-Supervised Learning of Sum-Product Networks

*Martin Trapp, Tamas Madl, Robert Peharz, Franz Pernkopf, Robert Trapp*

In several domains obtaining class annotations is expensive while at the same time unlabelled data are abundant. While most semi-supervised approaches enforce restrictive assumptions on the data distribution, recent work has managed to learn semi-supervised models in a non-restrictive regime. However, so far such approaches have only been proposed for linear models. In this work, we introduce semi-supervised parameter learning for Sum-Product Networks (SPNs). SPNs are deep probabilistic models admitting inference in linear time in number of network edges. Our approach has several advantages, as it (1) allows generative and discriminative semi-supervised learning, (2) guarantees that adding unlabelled data can increase, but not degrade, the performance (safe), and (3) is computationally efficient and does not enforce restrictive assumptions on the data distribution. We show on a variety of data sets that safe semi-supervised learning with SPNs is competitive compared to state-of-the-art and can lead to a better generative and discriminative objective value than a purely supervised approach.

## 109: Approximation Complexity of Maximum A Posteriori Inference in Sum-Product Networks

*Diarmaid Conaty, Denis D. Maua, Cassio P. de Campos*

We discuss the computational complexity of approximating maximum a posteriori inference in sum-product networks. We first show NP-hardness in trees of height two by a reduction from maximum independent set; this implies non-approximability within a sublinear factor. We show that this is a tight bound, as we can find an approximation within a linear factor in networks of height two. We then show that, in trees of height three, it is NP-hard to approximate the problem within a factor  $2^{f(n)}$  for any sublinear function  $f$  of the size of the input  $n$ . Again, this bound is tight, as we prove that the usual max-product algorithm finds (in any network) approximations within factor  $2^{c \cdot n}$  for some constant  $c < 1$ . Last, we present a simple algorithm, and show that it provably produces solutions at least as good as, and potentially much better than, the max-product algorithm. We empirically analyze the proposed algorithm against max-product using synthetic and real-world data.

## 120: Interpreting and using CPDAGs with background knowledge

*Emilija Perkovic, Markus Kalisch, Marloes H. Maathuis*

We develop terminology and methods for working with maximally oriented partially directed acyclic graphs (maximal PDAGs). Maximal PDAGs arise from imposing restrictions on a Markov equivalence class of directed acyclic graphs, or equivalently on its graphical representation as a completed partially directed acyclic graph (CPDAG), for example when adding background knowledge about certain edge orientations. Although maximal PDAGs often arise

in practice, causal methods have been mostly developed for CPDAGs. In this paper, we extend such methodology to maximal PDAGs. In particular, we develop methodology to read off possible ancestral relationships, we introduce a graphical criterion for covariate adjustment to estimate total causal effects, and we adapt the IDA and joint-IDA frameworks to estimate multi-sets of possible causal effects. We also present a simulation study that illustrates the gain in identifiability of total causal effects as the background knowledge increases. All methods are implemented in the R package `pcalg`.

## 130: Real-Time Resource Allocation for Tracking Systems

*Yash Satsangi, Shimon Whiteson, Frans A. Oliehoek, Henri Bouma*

Automated tracking is key to many computer vision applications. However, many tracking systems struggle to perform in real-time due to the high computational cost of detecting people, especially in ultra high resolution images. We propose a new algorithm called PartiMax that greatly reduces this cost by applying the person detector only to the relevant parts of the image. PartiMax exploits information in the particle filter to select  $k$  of the  $n$  candidate pixel boxes in the image. We prove that PartiMax is guaranteed to make a near-optimal selection with error bounds that are independent of the problem size. Furthermore, empirical results on a real-life dataset show that our system runs in real-time by processing only 10% of the pixel boxes in the image while still retaining 80% of the original tracking performance achieved when processing all pixel boxes.

## 132: Weighted Model Counting With Function Symbols

*Vaishak Belle*

Probabilistic relational languages lift the syntax of relational logic for the specification of large-scale probabilistic graphical models, often admitting concise descriptions for interacting random variables over classes, hierarchies and constraints. The emergence of weighted model counting as an effective and general approach to probabilistic inference has further allowed practitioners to reason about heterogeneous representations, such as Markov logic networks and ProbLog programs, by encoding them as a logical theory. However, much of this work has been limited to an essentially propositional setting: the logical model is understood in terms of ground formulas over a fixed and finite domain; no infinite domains, and certainly no function symbols (other than constants). On the one hand, this is not surprising, because such features are very problematic from a decidability viewpoint, but on the other, they turn out to be very attractive from the point of view of machine learning applications when there is uncertainty about the existence and identity of objects. In this paper, we reconsider the problem of probabilistic reasoning in a logical language with function symbols, and establish some key results that permit effective algorithms.

## 134: Synthesis of strategies in influence diagrams

*Manuel Luque, Manuel Arias, Francisco Javier Díez*

Influence diagrams (IDs) are a powerful tool for representing and solving decision problems under uncertainty. The objective of evaluating an ID is to compute the expected utility and

an optimal strategy, which consists of a policy for each decision. Every policy is usually represented as a table containing a column for each decision scenario, i.e., for each configuration of the variables on which it depends. The no-forgetting assumption, which implies that the decision maker always remembers all past observations and decisions, makes the policies grow exponentially with the number of variables in the ID. For human experts it is very difficult to understand the strategy contained in huge policy tables, not only for their size, but also because the vast majority of columns correspond to suboptimal or impossible scenarios and are hence irrelevant. This makes it difficult to extract the rules of action, to debug the model, and to convince the experts that the recommendations of the ID are reasonable. In this paper, we propose a method that presents the strategy in the form of a compact tree. It has been implemented in OpenMarkov, an open-source software tool for probabilistic graphical models. This facility was essential when evaluating an influence diagram for the mediastinal staging of non-small cell lung cancer; the optimal strategy, whose biggest policy table contained more than 15,000 columns, was synthesized into a tree of only 5 leaves.

### **135: Embedding Senses via Dictionary Bootstrapping**

*Byungkon Kang, Kyung-Ah Sohn*

This paper addresses the problem of embedding senses of a plain word according to its context. Natural language is inherently ambiguous, due to the presence of many multi-sensed words. Such ambiguity might have undesirable influence over many text-mining tools, including word embedding. Traditional word embedding techniques have focused on identifying which words tend to co-occur with one another in order to derive close embeddings for such words. However, the effectiveness of this approach is largely susceptible to the validity and neutrality of the training corpus. To address this problem, we propose to use the dictionary as the authoritative corpus for computing the word embeddings. The basic idea is to simultaneously embed the definition sentence while disambiguating words. Since dictionaries list a set of disambiguated senses, the proposed procedure yields sense embeddings which exhibit semantic characteristics comparable to plain word embeddings.

### **136: Triply Stochastic Gradients on Multiple Kernel Learning**

*Xiang Li, Bin Gu, Shuang Ao, Huaimin Wang, Charles X. Ling*

Multiple Kernel Learning (MKL) is highly useful for learning complex data with multiple cues or representations. However, MKL is known to have poor scalability because of the expensive kernel computation. Dai et al. (2014) proposed to use a doubly Stochastic Gradient Descent algorithm (doubly SGD) to greatly improve the scalability of kernel methods. However, the algorithm is not suitable for MKL because it cannot learn the kernel weights. In this paper, we provide a novel extension to doubly SGD for MKL so that both the decision functions and the kernel weights can be learned simultaneously. To achieve this, we develop the triply Stochastic Gradient Descent (triply SGD) algorithm which involves three sources of randomness – the data points, the random features, and the kernels, which was not considered in previous work. We prove that our algorithm enjoys similar convergence rate as that of doubly SGD. Comparing to several traditional MKL solutions, we show that our method has faster convergence speed and achieved better accuracy. Most importantly, our method makes it possible to learn MKL problems with millions of data points on a normal desktop PC.

## 137: Stochastic Bandit Models for Delayed Conversions

*Claire Vernade, Olivier Cappé, Vianney Perchet*

Online advertising and product recommendation are important domains of applications for multi-armed bandit methods. In these fields, the reward that is immediately available is most often only a proxy for the actual outcome of interest, which we refer to as a *conversion*. For instance, in web advertising, clicks can be observed within a few seconds after an ad display but the corresponding sale –if any– will take hours, if not days to happen. This paper proposes and investigates a new stochastic multi-armed bandit model in the framework proposed by Chapelle (2014) –based on empirical studies in the field of web advertising– in which each action may trigger a future reward that will then happen with a stochastic delay. We assume that the probability of conversion associated with each action is unknown while the distribution of the conversion delay is known, distinguishing between the (idealized) case where the conversion events may be observed whatever their delay and the more realistic setting in which late conversions are censored. We provide performance lower bounds as well as two simple but efficient algorithms based on the UCB and KLUCB frameworks. The latter algorithm, which is preferable when conversion rates are low, is based on a Poissonization argument, of independent interest in other settings where aggregation of Bernoulli observations with different success probabilities is required.

## 138: Importance Sampled Stochastic Optimization for Variational Inference

*Joseph Sakaya, Arto Klami*

Variational inference approximates the posterior distribution of a probabilistic model with a parameterized density by maximizing a lower bound for the model evidence. Modern solutions fit a flexible approximation with stochastic gradient descent, using Monte Carlo approximation for the gradients. This enables variational inference for arbitrary differentiable probabilistic models, and consequently makes variational inference feasible for probabilistic programming languages. In this work we develop more efficient inference algorithms for the task by considering importance sampling estimates for the gradients. We show how the gradient with respect to the approximation parameters can often be evaluated efficiently without needing to re-compute gradients of the model itself, and then proceed to derive practical algorithms that use importance sampled estimates to speed up computation. We present importance sampled stochastic gradient descent that outperforms standard stochastic gradient descent by a clear margin for a range of models, and provide a justifiable variant of stochastic average gradients for variational inference.

## 141: Coupling Adaptive Batch Sizes with Learning Rates

*Lukas Balles, Javier Romero, Philipp Hennig*

Mini-batch stochastic gradient descent and variants thereof have become standard for large-scale empirical risk minimization like the training of neural networks. These methods are usually used with a constant batch size chosen by simple empirical inspection. The batch size significantly influences the behavior of the stochastic optimization algorithm, though, since it determines the variance of the gradient estimates. This variance also changes over the optimization process;

when using a constant batch size, stability and convergence is thus often enforced by means of a (manually tuned) decreasing learning rate schedule.

We propose a practical method for dynamic batch size adaptation. It estimates the variance of the stochastic gradients and adapts the batch size to decrease the variance proportionally to the value of the objective function, removing the need for the aforementioned learning rate decrease. In contrast to recent related work, our algorithm couples the batch size to the learning rate, directly reflecting the known relationship between the two. On popular image classification benchmarks, our batch size adaptation yields faster optimization convergence, while simultaneously simplifying learning rate tuning. A TensorFlow implementation is available.

## **142: Green Generative Modeling: Recycling Dirty Data using Recurrent Variational Autoencoders**

*Yu Wang, Bin Dai, Gang Hua, John Aston, David Wipf*

This paper explores two useful modifications of the recent variational autoencoder (VAE), a popular deep generative modeling framework that dresses traditional autoencoders with probabilistic attire. The first involves a specially-tailored form of conditioning that allows us to simplify the VAE decoder structure while simultaneously introducing robustness to outliers. In a related vein, a second, complementary alteration is proposed to further build invariance to contaminated or dirty samples via a data augmentation process that amounts to recycling. In brief, to the extent that the VAE is legitimately a representative generative model, then each output from the decoder should closely resemble an authentic sample, which can then be resubmitted as a novel input ad infinitum. Moreover, this can be accomplished via special recurrent connections without the need for additional parameters to be trained. We evaluate these proposals on multiple practical outlier-removal and generative modeling tasks, demonstrating considerable improvements over existing algorithms.

## **146: Approximate Evidential Reasoning Using Local Conditioning and Conditional Belief Functions**

*Van Nguyen*

We propose a new message-passing belief propagation method that approximates belief updating on evidential networks with conditional belief functions. By means of local conditioning, the method is able to propagate beliefs on the original multiply-connected network structure using local computations, facilitating reasoning in a distributed and dynamic context. Further, by use of conditional belief functions in the form of partially defined plausibility and basic plausibility assignment functions, belief updating can be efficiently approximated using only partial information of the belief functions involved. Experiments show that the method produces results with high degree of accuracy whilst achieving a significant decrease in computational and space complexity (compared to exact methods).

## **147: Decoupling Homophily and Reciprocity with Latent Space Network Models**

*Jiasen Yang, Vinayak Rao, Jennifer Neville*

Networks form useful representations of data arising in various physical and social domains. In this work, we consider dynamic networks such as communication networks in which links connecting pairs of nodes appear over continuous time. We adopt a point process-based approach, and study latent space models which embed the nodes into Euclidean space. We propose models to capture two different aspects of dynamic network data: (i) communication occurs at a higher rate between individuals with similar features (homophily), and (ii) individuals tend to reciprocate communications from other nodes, but in a manner that varies across individuals. Our framework marries ideas from point process models, including Poisson and Hawkes processes, with ideas from latent space models of static networks. We evaluate our models over a range of tasks on real-world datasets and show that a dual latent space model, which accounts for heterogeneity in both reciprocity and homophily, significantly improves performance for both static and dynamic link prediction.

## 152: Differentially Private Variational Inference for Non-conjugate Models

*Joonas Jälkö, Onur Dikmen, Antti Honkela*

Many machine learning applications are based on data collected from people, such as their tastes and behaviour as well as biological traits and genetic data. Regardless of how important the application might be, one has to make sure individuals' identities or the privacy of the data are not compromised in the analysis. Differential privacy constitutes a powerful framework that prevents breaching of data subject privacy from the output of a computation. Differentially private versions of many important Bayesian inference methods have been proposed, but there is a lack of an efficient unified approach applicable to arbitrary models. In this contribution, we propose a differentially private variational inference method with a very wide applicability. It is built on top of doubly stochastic variational inference, a recent advance which provides a variational solution to a large class of models. We add differential privacy into doubly stochastic variational inference by clipping and perturbing the gradients. The algorithm is made more efficient through privacy amplification from subsampling. We demonstrate the method can reach an accuracy close to non-private level under reasonably strong privacy guarantees, clearly improving over previous sampling-based alternatives especially in the strong privacy regime.

## 155: Multi-dueling Bandits with Dependent Arms

*Yanan Sui, Vincent Zhuang, Joel W. Burdick, Yisong Yue*

The dueling bandits problem is an online learning framework for learning from pairwise preference feedback, and is particularly well-suited for modeling settings that elicit subjective or implicit human feedback. In this paper, we study the problem of multi-dueling bandits with dependent arms, which extends the original dueling bandits setting by simultaneously dueling multiple arms as well as modeling dependencies between arms. These extensions capture key characteristics found in many real-world applications, and allow for the opportunity to develop significantly more efficient algorithms than were possible in the original setting. We propose the SelfSparring algorithm, which reduces the multi-dueling bandits problem to a conventional bandit setting that can be solved using a stochastic bandit algorithm such as Thompson Sampling, and can naturally model dependencies using a Gaussian process prior. We present a no-regret analysis for multi-dueling setting, and demonstrate the effectiveness of our algorithm empirically on a wide range of simulation settings.

## **160: On Loopy Belief Propagation – Local Stability Analysis for Non-Vanishing Fields**

*Christian Knoll, Franz Pernkopf*

In this work we obtain all fixed points of belief propagation and perform a local stability analysis. We consider pairwise interactions of binary random variables and investigate the influence of non-vanishing fields and finite-size graphs on the performance of belief propagation; local stability is heavily influenced by these properties. We show why non-vanishing fields help to achieve convergence and increase the accuracy of belief propagation. We further explain the close connections between the underlying graph structure, the existence of multiple solutions, and the capability of belief propagation (with damping) to converge. Finally we provide insights into why finite-size graphs behave better than infinite-size graphs.

## **162: Value Directed Exploration in Multi-Armed Bandits with Structured Priors**

*Bence Cserna, Marek Petrik, Reazul Hasan Russel, Wheeler Ruml*

Multi-armed bandits are a quintessential machine learning problem requiring balancing exploration with exploitation. While there has been progress in developing algorithms with strong theoretical guarantees, there has been less focus on practical near-optimal finite-time performance. In this paper, we propose an algorithm for Bayesian multi-armed bandits that utilizes approximate value functions. Building on previous work on UCB and Gittins index, we introduce linearly-separable value functions that capture the benefit of exploration when choosing the next arm to pull. Our algorithm enjoys a sub-linear performance guarantee and our simulation results confirm its strength in problems with structured priors. The simplicity and generality of our approach makes it a strong candidate for use in more complex multi-armed bandit problems.

## **163: Composing inference algorithms as program transformations**

*Robert Zinkov, Chung-chieh Shan*

Probabilistic inference procedures are usually coded painstakingly from scratch, for each target model and each inference algorithm. We reduce this effort by generating inference procedures from models automatically. We make this code generation modular by decomposing inference algorithms into reusable program-to-program transformations. These transformations perform exact inference as well as generate probabilistic programs that compute expectations, densities, and MCMC samples. The resulting inference procedures are about as accurate and fast as other probabilistic programming systems on real-world problems.

## **171: A Practical Method for Solving Contextual Bandit Problems Using Decision Trees**

*Ryan McNellis, Adam N. Elmachoub, Sechan Oh, Marek Petrik*

Many efficient algorithms with strong theoretical guarantees have been proposed for the con-

textual multi-armed bandit problem. However, applying these algorithms in practice can be difficult because they require domain expertise to build appropriate features and to tune their parameters. We propose a new method for the contextual bandit problem that is simple, practical, and can be applied with little or no domain expertise. Our algorithm relies on decision trees to model the context-reward relationship. Decision trees are non-parametric, interpretable, and work well without hand-crafted features. To guide the exploration-exploitation trade-off, we use a bootstrapping approach which abstracts Thompson sampling to non-Bayesian settings. We also discuss several computational heuristics and demonstrate the performance of our method on several datasets.

## **173: Computing Nonvacuous Generalization Bounds for Deep (Stochastic) Neural Networks with Many More Parameters than Training Data**

*Gintare Karolina Dziugaite, Daniel M. Roy*

One of the defining properties of deep learning is that models are chosen to have many more parameters than available training data. In light of this capacity for overfitting, it is remarkable that simple algorithms like SGD reliably return solutions with low test error. One roadblock to explaining these phenomena in terms of implicit regularization, structural properties of the solution, and/or easiness of the data is that many learning bounds are quantitatively vacuous when applied to networks learned by SGD in this "deep learning" regime. Logically, in order to explain generalization, we need nonvacuous bounds. We return to an idea by Langford and Caruana (2001), who used PAC-Bayes bounds to compute nonvacuous numerical bounds on generalization error for stochastic two-layer two-hidden-unit neural networks via a sensitivity analysis. By optimizing the PAC-Bayes bound directly, we are able to extend their approach and obtain nonvacuous generalization bounds for deep stochastic neural network classifiers with millions of parameters trained on only tens of thousands of examples. We connect our findings to recent and old work on flat minima and MDL-based explanations of generalization.

## **180: Analysis of Thompson Sampling for Stochastic Sleeping Bandits**

*Aritra Chatterjee, Ganesh Ghalme, Shweta Jain, Rohit Vaish, Y. Narahari*

We study a variant of the stochastic multi-armed bandit problem where the set of available arms varies arbitrarily with time (also known as the sleeping bandit problem). We focus on the Thompson Sampling algorithm and consider a regret notion defined with respect to the best available arm. Our main result is an  $O(\log T)$  regret bound for Thompson Sampling, which generalizes a similar bound known for this algorithm from the classical bandit setting. Our bound also matches (up to constants) the best-known lower bound for the sleeping bandit problem. We show via simulations that Thompson Sampling outperforms the UCB-style AUER algorithm for the sleeping bandit problem.

## **185: Convex-constrained Sparse Additive Modeling and Its Extensions**

*Junming Yin, Yaoliang Yu*

Sparse additive modeling is a class of effective methods for performing high-dimensional non-parametric regression. In this work we show how shape constraints such as convexity/concavity and their extensions, can be integrated into additive models. The proposed sparse difference of convex additive models (SDCAM) can estimate most continuous functions without any a priori smoothness assumption. Motivated by a characterization of difference of convex functions, our method incorporates a natural regularization functional to avoid overfitting and to reduce model complexity. Computationally, we develop an efficient backfitting algorithm with linear per-iteration complexity. Experiments on both synthetic and real data verify that our method is competitive against state-of-the-art sparse additive models, with improved performance in most scenarios.

## 188: Learning Approximately Objective Priors

*Eric Nalisnick, Padhraic Smyth*

Informative Bayesian priors are often difficult to elicit, and when this is the case, modelers usually turn to noninformative or objective priors. However, objective priors such as the Jeffreys and reference priors are not tractable to derive for many models of interest. We address this issue by proposing techniques for learning reference prior approximations: we select a parametric family and optimize a black-box lower bound on the reference prior objective to find the member of the family that serves as a good approximation. We experimentally demonstrate the method's effectiveness by recovering Jeffreys priors and learning the Variational Autoencoder's reference prior.

## 197: Iterative Decomposition Guided Variable Neighborhood Search for Graphical Model Energy Minimization

*Abdelkader Ouali, David Allouche, Simon de Givry, Samir Loudni, Yahia Lebbah, Francisco Eckhardt, Lakhdar Loukil*

Graphical models factorize a global probability distribution/energy function as the product/sum of local functions. A major inference task, known as MAP in Markov Random Fields and MPE in Bayesian Networks, is to find a global assignment of all the variables with maximum a posteriori probability/minimum energy. A usual distinction on MAP solving methods is complete/incomplete, i.e. the ability to prove optimality or not. Most complete methods rely on tree search, while incomplete methods rely on local search. Among them, we study Variable Neighborhood Search (VNS) for graphical models. In this paper, we propose an iterative approach above VNS which uses (partial) tree search inside its local neighborhood exploration. The resulting hybrid method offers a good compromise between completeness and anytime behavior than existing tree search methods while still being competitive for proving optimality. We further propose a parallel version of our method improving its anytime behavior on difficult instances coming from a large graphical model benchmark. Last we experiment on the challenging minimum energy problem found in Computational Protein Design, showing the practical benefit of our parallel version.

## 206: Learning to Draw Samples with Amortized Stein Variational Gradient Descent

*Yihao Feng, Dilin Wang, Qiang Liu*

We propose a simple algorithm to train stochastic neural networks to draw samples from given target distributions for probabilistic inference. Our method is based on iteratively adjusting the neural network parameters so that the output changes along a Stein variational gradient direction (Liu & Wang, 2016) that maximally decreases the KL divergence with the target distribution. Our method works for any target distribution specified by their unnormalized density function, and can train any black-box architectures that are differentiable in terms of the parameters we want to adapt. We demonstrate our method with a number of applications, including variational autoencoder (VAE) with expressive encoders to model complex latent space structures, and hyper-parameter learning of MCMC samplers that allows Bayesian inference to adaptively improve itself when seeing more data.

## 207: Fair Optimal Stopping Policy for Matching with Mediator

*Yang Liu*

In this paper we study an optimal stopping policy for a multi-agent delegated sequential matching system with fairness constraints. We consider a setting where a mediator/decision maker matches a sequence of arriving assignments to multiple groups of agents, with agents being grouped according to certain sensitive attributes that needs to be protected. The decision maker aims to maximize total rewards that can be collected from above matching process (from all groups), while making the matching fair among groups. We discuss two types of fairness constraints: (i) each group has a certain expected deadline before which the match needs to happen; (ii) each group would like to have a guaranteed share of average reward from the matching. We present the exact characterization of fair optimal strategies. Example is provided to demonstrate the computation efficiency of our solution.

## 209: A Reinforcement Learning Approach to Weaning of Mechanical Ventilation in Intensive Care Units

*Niranjani Prasad, Li-Fang Cheng, Corey Chivers, Michael Draugelis, Barbara E Engelhardt*

The management of invasive mechanical ventilation, and the regulation of sedation and analgesia during ventilation, constitutes a major part of the care of patients admitted to intensive care units. Both prolonged dependence on mechanical ventilation and premature extubation are associated with increased risk of complications and higher hospital costs, but clinical opinion on the best protocol for weaning patients off the ventilator varies. This work looks to develop a decision support tool that uses available information to predict time to extubation readiness and recommend a personalized regime of sedation dosage and ventilator support. To this end, we employ off-policy reinforcement learning algorithms to determine the best action at a given patient state from sub-optimal historical ICU data. We compare treatment policies from fitted Q-iteration with extremely randomized trees and with feedforward neural networks, and demonstrate that the policies learnt show promise in recommending weaning protocols with improved outcomes, in terms of minimizing the rate of reintubation and closer regulation of physiological stability.

## 217: Stein Variational Adaptive Importance Sampling

*Jun Han, Qiang Liu*

We propose a novel adaptive importance sampling algorithm which incorporates Stein variational gradient descent algorithm (SVGD) with importance sampling (IS). Our algorithm leverages the nonparametric transforms defined in SVGD to iteratively decrease the KL divergence between our importance proposal and the target distribution. The advantages of this idea are two folds: first, it enables our algorithm to inherit more theoretical interpretability of IS than SVGD; second, we do not restrict the choice of our importance proposal from predefined distribution families as in traditional adaptive IS. Empirical experiments demonstrate that our algorithm performs better than the traditional adaptive IS. In addition, our algorithm achieves comparable results with Hamiltonian annealing importance sampling when employed to estimate the partition functions of graphical models and evaluate the log-likelihoods of deep generative models.

## 220: A Tractable Probabilistic Model for Subset Selection

*Yujia Shen, Arthur Choi, Adnan Darwiche*

Subset selection tasks, such as top-k ranking, induce datasets where examples have cardinalities that are known a priori. In this paper, we propose a tractable probabilistic model for subset selection and show how it can be learned from data. Our proposed model is interpretable and subsumes a previously introduced model based on logistic regression. We show how the parameters of our model can be estimated in closed form given complete data, and propose an algorithm for learning its structure in an interpretable space. We highlight the intuitive structures that we learn via case studies. We finally show how our proposed model can be viewed as an instance of the recently proposed Probabilistic Sentential Decision Diagram.

## 225: Importance Sampling for Fair Policy Selection

*Shayan Doroudi, Philip Thomas, Emma Brusnicki*

We consider the problem of off-policy policy selection in reinforcement learning: using historical data generated from running one policy to compare two or more policies. We show that approaches based on importance sampling can be unfair—they can select the worse of the two policies more often than not. We give two examples where the unfairness of importance sampling could be practically concerning. We then present sufficient conditions to theoretically guarantee fairness and a related notion of safety. Finally, we provide a practical importance sampling-based estimator to help mitigate one of the systematic sources of unfairness resulting from using importance sampling for policy selection.

## 229: Structure Learning of Linear Gaussian Structural Equation Models with Weak Edges

*Marco Eigenmann, Preetam Nandy, Marloes Maathuis*

We consider structure learning of linear Gaussian structural equation models with weak edges. Since the presence of weak edges can lead to a loss of edge orientations in the true underlying

CPDAG, we define a new graphical object that can contain more edge orientations. We show that this object can be recovered from observational data under a type of strong faithfulness assumption. We present a new algorithm for this purpose, called aggregated greedy equivalence search (AGES), that aggregates the solution path of the greedy equivalence search (GES) algorithm for varying values of the penalty parameter. We prove consistency of AGES and demonstrate its performance in a simulation study and on single cell data from Sachs et al. (2005). The algorithm will be made available in the R-package pcalg.

## 234: SAT-Based Causal Discovery under Weaker Assumptions

*Zhalama, Jiji Zhang, Frederick Eberhardt, Wolfgang Mayer*

Using the flexibility of recently developed methods for causal discovery based on Boolean satisfiability (SAT) solvers, we encode a variety of assumptions that weaken the Faithfulness assumption. The encoding results in a number of SAT-based algorithms whose asymptotic correctness relies on weaker conditions than are standardly assumed. This implementation of a whole set of assumptions in the same platform enables us to systematically explore the effect of weakening the Faithfulness assumption on causal discovery. An important effect, suggested by simulation results, is that adopting weaker assumptions greatly alleviates the problem of conflicting constraints and substantially shortens solving time. As a result, SAT-based causal discovery is potentially more scalable under weaker assumptions.

## 237: Learning to Acquire Information

*Yewen Pu, Leslie Pack Kaelbling, Armando Solar-Lezama*

We consider the problem of diagnosis where a set of simple observations are used to infer a potentially complex hidden hypothesis. Finding the optimal subset of observations is intractable in general, thus we focus on the problem of active diagnosis, where the agent selects the next most-informative observation based on the results of previous observations. We show that under the assumption of uniform observation entropy, one can build an *implication model* which directly predicts the outcome of the potential next observation conditioned on the results of past observations, and selects the observation with the maximum entropy. This approach enjoys reduced computation complexity by bypassing the complicated hypothesis space, and can be trained on observation data alone, learning how to query without knowledge of the hidden hypothesis.

## 239: Stein Variational Policy Gradient

*Yang Liu, Prajit Ramachandran, Qiang Liu, Jian Peng*

Policy gradient methods have been successfully applied to many complex reinforcement learning problems. The major drawbacks of policy gradient methods include high variability, slow convergence, and exploration inefficiency. We introduce a maximum entropy policy optimization framework which explicitly encourages parameter exploration. We show that this optimization framework can be reduced to a Bayesian inference problem and then propose a novel Stein variational policy gradient method (SVPG) that utilizes existing policy gradient methods and a repulsive functional to generate a set of diverse but well-behaved policies. On continuous control problems, we find that the SVPG implementations of REINFORCE and advantage actor-critic

algorithms greatly boost the performance of their original versions, in terms of both data efficiency and average return. SVPG is also robust against the initialization variation and can be easily implemented in a parallel manner.

## **242: Exact Inference for Relational Graphical Models with Interpreted Functions: Lifted Probabilistic Inference Modulo Theories**

*Rodrigo de Salvo Braz, Ciaran O'Reilly*

Probabilistic Inference Modulo Theories (PIMT) is a recent framework that expands exact inference on graphical models to use richer languages that include arithmetic, equalities, and inequalities on both integers and real numbers. In this paper, we expand PIMT to a lifted version that also processes random functions and relations. This enhancement is achieved by adapting Inversion, a method from Lifted First-Order Probabilistic Inference literature, to also be modulo theories. This results in the first algorithm for exact probabilistic inference that efficiently and simultaneously exploits random relations and functions, arithmetic, equalities and inequalities.

## **244: Neighborhood Regularized $\ell^1$ -Graph**

*Yingzhen Yang, Jiashi Feng, Jiahui Yu, Jianchao Yang, Pushmeet Kohli, Thomas S. Huang*

$\ell^1$ -Graph, which learns a sparse graph over the data by sparse representation, has been demonstrated to be effective in clustering especially for high dimensional data. Although it achieves compelling performance, the sparse graph generated by  $\ell^1$ -Graph ignores the geometric information of the data by sparse representation for each datum separately. To obtain a sparse graph that is aligned to the underlying manifold structure of the data, we propose the novel Neighborhood Regularized  $\ell^1$ -Graph (NR $\ell^1$ -Graph). NR $\ell^1$ -Graph learns sparse graph with locally consistent neighborhood by encouraging nearby data to have similar neighbors in the constructed sparse graph. We present the optimization algorithm of NR $\ell^1$ -Graph with theoretical guarantee on the convergence and the gap between the sub-optimal solution and the globally optimal solution in each step of the coordinate descent, which is essential for the overall optimization of NR $\ell^1$ -Graph. Its provable accelerated version, NR $\ell^1$ -Graph by Random Projection (NR $\ell^1$ -Graph-RP) that employs randomized data matrix decomposition, is also presented to improve the efficiency of the optimization of NR $\ell^1$ -Graph. Experimental results on various real data sets demonstrate the effectiveness of both NR $\ell^1$ -Graph and NR $\ell^1$ -Graph-RP.

## **250: Feature-to-Feature Regression for a Two-Step Conditional Independence Test**

*Qinyi Zhang, Sarah Filippi, Seth Flaxman, Dino Sejdinovic*

The algorithms for causal discovery and more broadly for learning the structure of graphical models require well calibrated and consistent conditional independence (CI) tests. We revisit the CI tests which are based on two-step procedures and involve regression with subsequent (unconditional) independence test (RESIT) on regression residuals and investigate the assumptions under which these tests operate. In particular, we demonstrate that when going beyond simple functional relationships with additive noise, such tests can lead to an inflated number

of false discoveries. We study the relationship of these tests with those based on dependence measures using reproducing kernel Hilbert spaces (RKHS) and propose an extension of RESIT which uses RKHS-valued regression. The resulting test inherits the simple two-step testing procedure of RESIT, while giving correct Type I control and competitive power. When used as a component of the PC algorithm, the proposed test is more robust to the case where hidden variables induce a switching behaviour in the associations present in the data.

## **255: How Good Are My Predictions? Efficiently Approximating Precision-Recall Curves for Massive Datasets**

*Ashish Sabharwal, Hanie Sedghi*

Large scale machine learning produces massive datasets whose items are often associated with a confidence level and can thus be ranked. However, computing the precision of these resources requires human annotation, which is often prohibitively expensive and is therefore skipped. We consider the problem of cost-effectively approximating precision-recall (PR) or ROC curves for such systems. Our novel approach, called PAULA, provides theoretically guaranteed lower and upper bounds on the underlying precision function while relying on only  $O(\log N)$  annotations for a resource with  $N$  items. This contrasts favorably with  $\Theta(\sqrt{N \log N})$  annotations needed by commonly used sampling based methods. Our key insight is to capitalize on a natural monotonicity property of the underlying confidence-based ranking. PAULA provides tight bounds for PR curves using, e.g., only 17K annotations for resources with 200K items and 48K annotations for resources with 2B items. We use PAULA to evaluate a subset of the much utilized PPDB paraphrase database and a recent Science knowledge base.

## **259: Improving Optimization-Based Approximate Inference by Clamping Variables**

*Junyao Zhao, Josip Djolonga, Sebastian Tschiatschek, Andreas Krause*

While central to the application of probabilistic models to discrete data, the problem of marginal inference is in general intractable and efficient approximation schemes need to exploit the problem structure. Recently, there have been efforts to develop inference techniques that do not necessarily make factorization assumptions about the distribution, but rather exploit the fact that sometimes there exist efficient algorithms for finding the MAP configuration. In this paper, we theoretically prove that for discrete multi-label models the bounds on the partition function obtained by two of these approaches, Perturb-and-MAP and the bound from the infinite Renyi divergence, can be only improved by clamping any subset of the variables. For the case of log-supermodular models we provide a more detailed analysis and develop a set of efficient strategies for choosing the order in which the variables should be clamped. Finally, we present a number of numerical experiments showcasing the improvements obtained by the proposed methods on several models.

## **263: Why Rules are Complex: Real-Valued Probabilistic Logic Programs are not Fully Expressive**

*David Buchman, David Poole*

This paper explores what can and cannot be represented by probabilistic logic programs (PLPs). Propositional PLPs can represent any distribution because they can be acyclic. For relational domains with fixed populations, the probabilistic parameters can be derived as the solutions to polynomial equations. Unfortunately, sometimes they only have complex-valued solutions. Thus PLPs, even with arbitrarily real-valued parameters, cannot represent all distributions. Moreover, they cannot approximate all distributions. Allowing the parameters to be complex numbers, we present a natural truly-cyclic canonical representation that with probability 1 can represent all distributions for (propositional or) relational domains with fixed populations, and, unlike standard representations, has no redundant parameters.

## **266: Treatment-Response Models for Counterfactual Reasoning with Continuous-time, Continuous-valued Interventions**

*Hossein Soleimani, Adarsh Subbaswamy, Suchi Saria*

Treatment effects can be estimated from observational data as the difference in potential outcomes. In this paper, we address the challenge of estimating the potential outcome when treatment-dose levels can vary continuously over time. Further, the outcome variable may not be measured at a regular frequency. Our proposed solution represents the treatment response curves using linear time-invariant dynamical systems—this provides a flexible means for modeling response over time to highly variable dose curves. Moreover, for multivariate data, the proposed method: uncovers shared structure in treatment response and the baseline across multiple markers; and, flexibly models challenging correlation structure both across and within signals over time. For this, we build upon the framework of multiple-output Gaussian Processes. On simulated and a challenging clinical dataset, we show significant gains in accuracy over state-of-the-art models.

## **267: Holographic Feature Representations of Deep Networks**

*Martin A. Zinkevich, Alex Davies, Dale Schuurmans*

It is often asserted that deep networks learn “features”, traditionally expressed by the activations of intermediate nodes. We explore an alternative concept by defining features as partial derivatives of model output with respect to model parameters—extending a simple yet powerful idea from generalized linear models. The resulting features are not equivalent to node activations, and we show that they can induce a holographic representation of the complete model: the network’s output on given data can be exactly replicated by a simple linear model over such features extracted from any ordered cut. We demonstrate useful advantages for this feature representation over standard representations based on node activations.

## **269: Causal Discovery from Temporally Aggregated Time Series**

*Mingming Gong, Kun Zhang, Bernhard Schölkopf, Clark Glymour, Dacheng Tao*

Discovering causal structure of a dynamical system from observed time series is a traditional and important problem. In many practical applications, observed data are obtained by applying subsampling or temporally aggregation to the original causal processes, making it difficult to discover the underlying causal relations. Subsampling refers to the procedure that for every  $k$

consecutive observations, one is kept, the rest being skipped, and recently some advances have been made in causal discovery from such data. With temporal aggregation, the local averages or sums of  $k$  consecutive, non-overlapping observations in the causal process are computed as new observations, and causal discovery from such data is even harder. In this paper, we investigate how to recover causal relations at the original causal frequency from temporally aggregated data. Assuming the time series at the causal frequency follows a vector autoregressive (VAR) model, we show that the causal structure at the causal frequency is identifiable from aggregated time series if the noise terms are independent and non-Gaussian and some other technical conditions hold. We then present an estimation method based on non-Gaussian state-space modeling and evaluate its performance on both synthetic and real data.

## **276: Online Constrained Model-based Reinforcement Learning**

*Benjamin van Niekerk, Andreas Damianou, Benjamin Rosman*

Applying reinforcement learning to robotic systems poses a number of challenging problems. A key requirement is the ability to handle continuous state and action spaces while remaining within a limited time and resource budget. Additionally, for safe operation, the system must make robust decisions under hard constraints. To address these challenges, we propose a model based approach that combines Gaussian Process regression and Receding Horizon Control. Using sparse spectrum Gaussian Processes, we extend previous work by updating the dynamics model incrementally from a stream of sensory data. This results in an agent that can learn and plan in real-time under non-linear constraints. We test our approach on a cart pole swing-up environment and demonstrate the benefits of online learning on an autonomous racing task. The environments dynamics are learned from limited training data and can be reused in new task instances without retraining.

## **277: Algebraic Equivalence of Linear Structural Equation Models**

*Thijs van Ommen, Joris M. Mooij*

Despite their popularity, many questions about the algebraic constraints imposed by linear structural equation models remain open problems. For causal discovery, two of these problems are especially important: the enumeration of the constraints imposed by a model, and deciding whether two graphs define the same statistical model. We show how the half-trek criterion can be used to make progress in both of these problems. We apply our theoretical results to a small-scale model selection problem, and find that taking the additional algebraic constraints into account may lead to significant improvements in model selection accuracy.

## **278: A Probabilistic Framework for Zero-Shot Multi-Label Learning**

*Abhilash Gaure, Aishwarya Gupta, Vinay Kumar Verma, Piyush Rai*

We present a probabilistic framework for multi-label learning for the setting when the test data may require predicting labels that were not available at training time (i.e., the zero-shot learning setting). We develop a probabilistic model that leverages the co-occurrence statistics of the labels via a joint generative model for the label matrix (which denotes the label presence/absence for each training example) and for the label co-occurrence matrix (which denotes how many times a pair of labels co-occurs with each other). In addition to handling the unseen

labels at test time, leveraging the co-occurrence information may also help in the standard multi-label learning setting, especially if the number of training examples is very small and/or the label matrix of training examples has a large fraction of missing entries. Our experimental results demonstrate the efficacy of our model in handling unseen labels.

## **280: Efficient solutions for Stochastic Shortest Path Problems with Dead Ends**

*Felipe Trevizan, Florent Teichteil-Königsbuch, Sylvie Thiebaux*

Many planning problems require maximizing the probability of goal satisfaction as well as minimizing the expected cost to reach the goal. To model and solve such problems, there has been several attempts at extending stochastic shortest path problems (SSPs) to deal with dead-ends and optimize a dual optimization criterion. Unfortunately these extensions lack either theoretical robustness or practical efficiency. We study a new, perhaps even more natural optimization criterion capturing these problems, the Min-Cost given Max-Prob (MCMP) criterion. This criterion leads to the minimum expected cost policy among those with maximum success probability, and accurately accounts for the cost and risk of reaching dead-ends. Moreover, it lends itself to efficient solution methods that build on recent heuristic search algorithms for the dual representation of stochastic shortest paths problems. Our experiments show up to one order of magnitude speed-up over the state of the art.

## **281: Probabilistic Program Abstractions**

*Steven Holtzen, Todd Millstein, Guy Van den Broeck*

Abstraction is a fundamental tool for reasoning about complex systems. Program abstraction has been utilized to great effect for analyzing deterministic programs. At the heart of program abstraction is the relationship between a concrete program, which is difficult to analyze, and an abstract program, which is more tractable. Program abstractions, however, are typically not probabilistic. We generalize non-deterministic program abstractions to probabilistic program abstractions by explicitly quantifying the non-deterministic choices. Our framework upgrades key definitions and properties of abstractions to the probabilistic context. We also discuss preliminary ideas for performing inference on probabilistic abstractions and general probabilistic programs.

## **286: Communication-Efficient Distributed Primal-Dual Algorithm for Saddle Point Problems**

*Yaodong Yu, Sulin Liu, Sinno Jialin Pan*

Primal-dual algorithms, which are proposed to solve reformulated convex-concave saddle point problems, have been proven to be effective for solving a generic class of convex optimization problems, especially when the problems are ill-conditioned. However, the saddle point problem still lacks a distributed optimization framework where primal-dual algorithms can be employed. In this paper, we propose a novel communication-efficient distributed optimization framework to solve the convex-concave saddle point problem based on primal-dual methods. We carefully design local subproblems and a central problem such that our proposed distributed optimiza-

tion framework is communication-efficient. We provide a convergence analysis of our proposed algorithm, and extend it to address non-smooth and non-strongly convex loss functions. We conduct extensive experiments on several real-world datasets to demonstrate competitive performance of the proposed method, especially on ill-conditioned problems.

## **289: Continuously tempered Hamiltonian Monte Carlo**

*Matthew M Graham, Amos J Storkey*

Hamiltonian Monte Carlo (HMC) is a powerful Markov chain Monte Carlo (MCMC) method for performing approximate inference in complex probabilistic models of continuous variables. In common with many MCMC methods, however, the standard HMC approach performs poorly in distributions with multiple isolated modes. We present a method for augmenting the Hamiltonian system with an extra continuous temperature control variable which allows the dynamic to bridge between sampling a complex target distribution and a simpler unimodal base distribution. This augmentation both helps improve mixing in multimodal targets and allows the normalisation constant of the target distribution to be estimated. The method is simple to implement within existing HMC code, requiring only a standard leapfrog integrator. We demonstrate experimentally that the method is competitive with annealed importance sampling and simulating tempering methods at sampling from challenging multimodal distributions and estimating their normalising constants.

## **291: Learning the Structure of Probabilistic Sentential Decision Diagrams**

*Yitao Liang, Jessa Bekker, Guy Van den Broeck*

The probabilistic sentential decision diagram (PSDD) was recently introduced as a tractable representation of probability distributions that are subject to logical constraints. Meanwhile, efforts in tractable learning achieved great success inducing complex joint distributions from data without constraints, while guaranteeing efficient exact probabilistic inference; for instance by learning arithmetic circuits (ACs) or sum-product networks (SPNs). This paper studies the efficacy of PSDDs for the standard tractable learning task without constraints and develops the first PSDD structure learning algorithm, called LearnPSDD. Experiments on standard benchmarks show competitive performance, despite the fact that PSDDs are more tractable and more restrictive than their alternatives. LearnPSDD compares favorably to SPNs, particularly in terms of model size, which is a proxy for tractability. We report state-of-the-art likelihood results on six datasets. Moreover, LearnPSDD retains the ability to learn PSDD structures in probability spaces subject to logical constraints, which is beyond the reach of other representations.

## **292: Interpreting Lion Behaviour with Nonparametric Probabilistic Programs**

*Neil Dhir, Matthijs Vákár, Matthew Wijers, Andrew Markham, Frank Wood, Paul Trethowan, Byron Du Preez, Andrew Loveridge, David MacDonald*

We consider the problem of unsupervised learning of meaningful behavioural segments of high-

dimensional time-series observations, collected from a pride of African lions. We demonstrate, by way of a probabilistic programming system (PPS), a methodology which allows for quick iteration over models and Bayesian inferences, which enables us to learn meaningful behavioural segments. We introduce a new Bayesian nonparametric (BNP) state-space model, which extends the hierarchical Dirichlet process (HDP) hidden Markov model (HMM) with an explicit BNP treatment of duration distributions, to deal with different levels of granularity of the latent behavioural space of the lions. The ease with which this is done exemplifies the flexibility that a PPS gives a scientist. Furthermore, we combine this approach with unsupervised feature learning, using variational autoencoders

## **297: Deep Hybrid Models: Bridging Discriminative and Generative Approaches**

*Volodymyr Kuleshov, Stefano Ermon*

Most methods in machine learning are described as either discriminative or generative. The former often attain higher predictive accuracy, while the latter are more strongly regularized and can deal with missing data. Here, we propose a new framework to combine a broad class of discriminative and generative models, interpolating between the two extremes with a multiconditional likelihood objective. Unlike previous approaches, we couple the two components through shared latent variables, and train using recent advances in variational inference. Instantiating our framework with modern deep architectures gives rise to deep hybrid models, a highly flexible family that generalizes several existing models and is effective in the semi-supervised setting, where it results in improvements over the state of the art on the SVHN dataset.

## **298: Robust Model Equivalence using Stochastic Bisimulation for N-Agent Interactive DIDs**

*Muthukumaran Chandrasekaran, Junhuan Zhang, Prashant Doshi, Yifeng Zeng*

I-DIDs suffer disproportionately from the curse of dimensionality dominated by the exponential growth in the number of models over time. Previous methods for scaling I-DIDs identify notions of equivalence between models, such as behavioral equivalence (BE). But, this requires that the models be solved first. Also, model space compression across agents has not been previously investigated. We present a way to compress the space of models across agents, possibly with different frames, and do so without having to solve them first, using stochastic bisimulation. We test our approach on two non-cooperative partially observable domains with up to 20 agents.

## **306: Adversarial Sets for Regularising Neural Link Predictors**

*Pasquale Minervini, Thomas Demeester, Tim Rocktäschel, Sebastian Riedel*

In adversarial training, a set of models learn together by pursuing competing goals, usually defined on single data instances. However, in relational learning and other non-i.i.d domains, goals can also be defined over sets of instances. For example, a link predictor for the is-a relation needs to be consistent with the transitivity property: if  $\text{IS-A}(x_1, x_2)$  and  $\text{IS-A}(x_2, x_3)$  hold,  $\text{IS-A}(x_1, x_3)$  needs to hold as well. Here we use such assumptions for deriving an inconsistency

loss, measuring the degree to which the model violates the assumptions on an adversarially-generated set of examples. The training objective is defined as a minimax problem, where an adversary finds the most offending adversarial examples by maximising the inconsistency loss, and the model is trained by jointly minimising a supervised loss and the inconsistency loss on the adversarial examples. This yields the first method that can use function-free Horn clauses (as in Datalog) to regularise any neural link predictor, with complexity independent of the domain size. We show that for several link prediction models, the optimisation problem faced by the adversary has efficient closed-form solutions. Experiments on link prediction benchmarks indicate that given suitable prior knowledge, our method can significantly improve neural link predictors on all relevant metrics.

---

# Notes

---













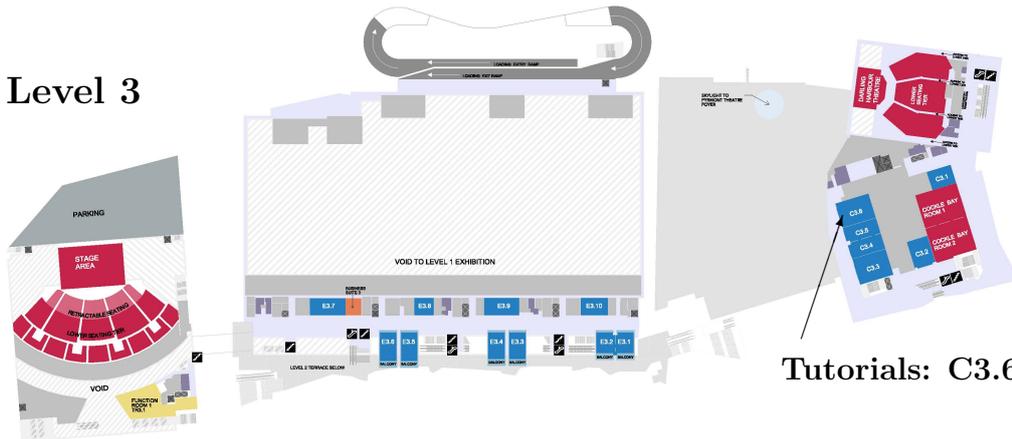


# Maps

## Sydney International Convention Centre

<http://www.iccsydney.com.au/plan-an-event/floor-plans>

### Level 3



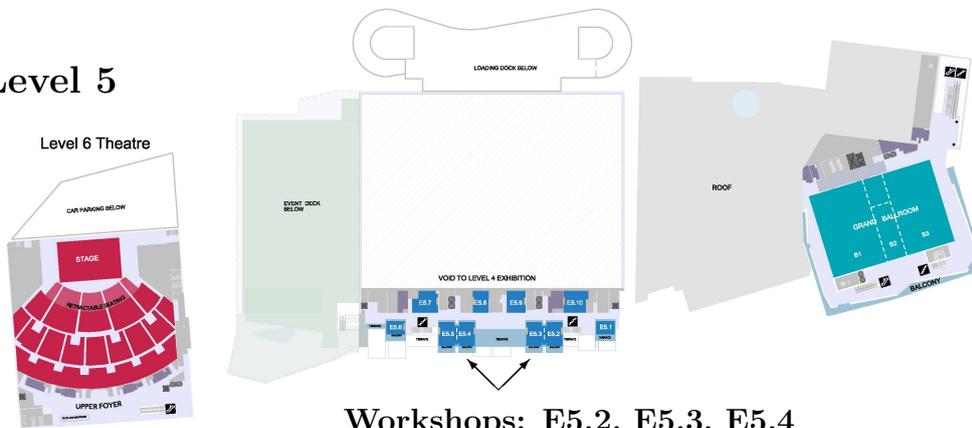
Tutorials: C3.6

### Level 4



Plenary: C4.8  
Posters: C4.9

### Level 5



Workshops: E5.2, E5.3, E5.4

# Sydney ICC Local Area

[http://www.iccsydney.com.au/~media/ICC/Images/Maps/ICCS\\_Transport\\_Map\\_1700x976.ashx](http://www.iccsydney.com.au/~media/ICC/Images/Maps/ICCS_Transport_Map_1700x976.ashx)





