

Uncertainty in Artificial Intelligence

Proceedings of the Twenty-Seventh Conference (2011)

Uncertainty in Artificial Intelligence

Proceedings of the Twenty-Seventh Conference (2011)

July 14–17, 2011

Barcelona, Spain

Edited by

Fabio G. Cozman

Escola Politecnica

Universidade de Sao Paulo

Sao Paulo, Brazil

Avi Pfeffer

Charles River Analytics

Boston, MA

Conference Chair

Peter Grünwald

Centrum Wiskunde & Informatica

Amsterdam, the Netherlands

Local Arrangements Chair

Lluís Godó

Artificial Intelligence Research

Institute (IIIA-CSIC)

Bellaterra, Spain

AUAI Press

Corvallis, Oregon

Sponsored by

Artificial Intelligence Journal

IBM Research

Google

Microsoft Research

Pascal Network of Excellence

Gobierno de España – Ministerio
de Ciencia e Innovación

The cover design is used with permission from Elsevier.

Production by Copy General Corporation, Sterling, Virginia.

Published by AUAJ Press for
Association for Uncertainty in Artificial Intelligence
<http://auai.org>

Editorial Office:
P.O. Box 866
Corvallis, Oregon 97339
USA

Copyright © 2011 by AUAJ Press
All rights reserved
Printed in the United States of America

No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without the prior written permission of the publisher.

ISBN 978-0-9749039-7-2

Contents

<i>Preface</i>	<i>xi</i>
<i>Acknowledgments</i>	<i>xii</i>
CONTRIBUTED PAPERS	
Graphical Models for Bandit Problems	1
Kareem Amin, Michael Kearns, Umar Syed	
Extended Lifted Inference with Joint Formulas	11
Udi Apsel, Ronen I. Brafman	
Learning is Planning: Near Bayes-optimal Reinforcement Learning via Monte-Carlo Tree Search	19
John Asmuth, Michael Littman	
Solving Cooperative Reliability Games	27
Yoram Bachrach, Reshef Meir, Michal Feldman, Moshe Tennenholtz	
Active Diagnosis via AUC Maximization: An Efficient Approach for Multiple Fault Identification in Large Scale, Noisy Networks	35
Gowtham Bellala, Jason Stanley, Clayton Scott, Suresh K. Bhavnani	
Semi-supervised Learning with Density Based Distances.	43
Avleen S. Bijral, Nathan Ratliff, Nati Srebro	
Deconvolution of Mixing Time Series on a Graph	51
Alexander W. Blocker, Edoardo M. Airolti	
Factored Filtering of Continuous-Time Systems.	61
E. Busra Celikkaya, Christian R. Shelton, William Lam	
Near-Optimal Target Learning With Stochastic Binary Signals	69
Mithun Chakraborty, Sanmay Das, Malik Magdon-Ismael	
Filtered Fictitious Play for Perturbed Observation Potential Games and Decentralised POMDPs.	77
Archie C. Chapman, Simon A. Williamson, Nicholas R. Jennings	
A Framework for Optimizing Paper Matching	86
Laurent Charlin, Richard S. Zemel, Craig Boutilier	
A Temporally Abstracted Viterbi Algorithm	96
Shaunak Chatterjee, Stuart Russell	
Smoothing Proximal Gradient Method for General Structured Sparse Learning	105
Xi Chen, Qihang Lin, Seyoung Kim, Jaime G. Carbonell, Eric P. Xing	
EDML: A Method for Learning Parameters in Bayesian Networks	115
Arthur Choi, Khaled S. Refaat, Adnan Darwiche	
Strictly Proper Mechanisms with Cooperating Players	125
SangIn Chun, Ross D. Shachter	
A Logical Characterization of Constraint-Based Causal Discovery	135
Tom Claassen, Tom Heskes	

Ensembles of Kernel Predictors	145
Corinna Cortes, Mehryar Mohri, Afshin Rostamizadeh	
Bayesian Network Learning with Cutting Planes	153
James Cussens	
Active Learning for Developing Personalized Treatment	161
Kun Deng, Joelle Pineau, Susan Murphy	
Efficient Optimal Learning for Contextual Bandits	169
Miroslav Dudik, Daniel Hsu, Satyen Kale, Nikos Karampatziakis, John Langford, Lev Reyzin, Tong Zhang	
A Unifying Framework for Linearly Solvable Control	179
Krishnamurthy Dvijotham, Emanuel Todorov	
Boosting as a Product of Experts	187
Narayanan U. Edakunni, Gavin Brown, Tim Kovacs	
PAC-Bayesian Policy Evaluation for Reinforcement Learning	195
Mahdi Milani Fard, Joelle Pineau, Csaba Szepesvári	
On the Complexity of Decision Making in Possibilistic Decision Trees	203
Hélène Fargier, Nahla Ben Amor, Wided Guezguez	
Inference in Probabilistic Logic Programs using Weighted CNFs	211
Daan Fierens, Guy Van den Broeck, Ingo Thon, Bernd Gutmann, Luc De Raedt	
Efficient Inference in Markov Control Problems.	221
Thomas Furnston, David Barber	
Dynamic Consistency and Decision Making under Vacuous Belief	230
Phan H. Giang	
Hierarchical Affinity Propagation	238
Inmar E. Givoni, Clement Chung, Brendan J. Frey	
Approximation by Quantization	247
Vibhav Gogate, Pedro Domingos	
Probabilistic Theorem Proving	256
Vibhav Gogate, Pedro Domingos	
Generalized Fisher Score for Feature Selection	266
Quanquan Gu, Zhenhui Li, Jiawei Han	
Active Semi-Supervised Learning using Submodular Functions.	274
Andrew Guillory, Jeff Bilmes	
Bregman Divergence as General Framework to Estimate Unnormalized Statistical Models	283
Michael U. Gutmann, Jun-ichiro Hirayama	
Reasoning about RoboCup Soccer Narratives.	291
Hannaneh Hajishirzi, Julia Hockenmaier, Erik T. Mueller, Eyal Amir	
Suboptimality Bounds for Stochastic Shortest Path Problems	301
Eric A. Hansen	
Sequential Inference for Latent Force Models	311
Jouni Hartikainen, Simo Särkkä	
What Cannot be Learned with Bethe Approximations	319
Uri Heinemann, Amir Globerson	

Portfolio Allocation for Bayesian Optimization	327
Matthew Hoffman, Eric Brochu, Nando de Freitas	
Sum-Product Networks: A New Deep Architecture	337
Hoifung Poon, Pedro Domingos	
Lipschitz Parametrization of Probabilistic Graphical Models	347
Jean Honorio	
Efficient Probabilistic Inference with Partial Ranking Queries.	355
Jonathan Huang, Ashish Kapoor, Carlos Guestrin	
Noisy-OR Models with Latent Confounding	363
Antti Hyttinen, Frederick Eberhardt, Patrik O. Hoyer	
Discovering Causal Structures in Binary Exclusive-OR Skew Acyclic Models	373
Takanori Inazumi, Takashi Washio, Shohei Shimizu, Joe Suzuki, Akihiro Yamamoto, Yoshinobu Kawahara	
Detecting Low-Complexity Unobserved Causes	383
Dominik Janzing, Eleni Sgouritsa, Oliver Stegle, Jonas Peters, Bernhard Schölkopf	
Online Importance Weight Aware Updates.	392
Nikos Karampatziakis, John Langford	
Modeling Social Networks with Node Attributes using the Multiplicative Attribute Graph Model.	400
Myunghwan Kim, Jure Leskovec	
Pitman-Yor Diffusion Trees.	410
David A. Knowles, Zoubin Ghahramani	
Learning Determinantal Point Processes	419
Alex Kulesza, Ben Taskar	
Message-Passing Algorithms for Quadratic Programming Formulations of MAP Estimation	428
Akshat Kumar, Shlomo Zilberstein	
An Efficient Protocol for Negotiation over Combinatorial Domains with Incomplete Information	436
Minyi Li, Quoc Bao Vo, Ryszard Kowalczyk	
Noisy Search with Comparative Feedback	445
Shiau Hong Lim, Peter Auer	
Variational Algorithms for Marginal MAP	453
Qiang Liu, Alexander T. Ihler	
Classification of Sets using Restricted Boltzmann Machines	463
Jérôme Louradour, Hugo Larochelle	
Belief Change with Noisy Sensing in the Situation Calculus	471
Jianbing Ma, Weiru Liu, Paul Miller	
Improving the Scalability of Optimal Bayesian Network Learning with External-Memory Frontier Breadth-First Branch and Bound Search	479
Brandon Malone, Changhe Yuan, Eric A. Hansen, Susan Bridges	
Order-of-Magnitude Influence Diagrams	489
Radu Marinescu, Nic Wilson	
Asymptotic Efficiency of Deterministic Estimators for Discrete Energy-Based Models: Ratio Matching and Pseudolikelihood	497
Benjamin M. Marlin, Nando de Freitas	

Reconstructing Pompeian Households.	506
David Mimno	
Conditional Restricted Boltzmann Machines for Structured Output Prediction	514
Volodymyr Mnih, Hugo Larochelle, Geoffrey E. Hinton	
Compact Mathematical Programs For DEC-MDPs With Structured Agent Interactions.	523
Hala Mostafa, Victor Lesser	
Fractional Moments on Bandit Problems.	531
Ananda Narayanan B, Balaraman Ravindran	
Dynamic Mechanism Design for Markets with Strategic Resources	539
Swaprava Nath, Onno Zoeter, Yadati Narahari, Chris R. Dance	
Multidimensional Counting Grids: Inferring Word Order from Disordered Bags of Words	547
Nebojsa Jojic, Alessandro Perina	
Partial Order MCMC for Structure Discovery in Bayesian Networks.	557
Teppo Niinimäki, Pekka Parviainen, Mikko Koivisto	
A Geometric Traversal Algorithm for Reward-Uncertain MDPs	565
Eunsoo Oh, Kee-Eung Kim	
Iterated Risk Measures for Risk-Sensitive Markov Decision Processes with Discounted Cost	573
Takayuki Osogami	
Price Updating in Combinatorial Prediction Markets with Bayesian Networks	581
David M. Pennock, Lirong Xia	
Identifiability of Causal Graphs using Functional Models	589
Jonas Peters, Joris M. Mooij, Dominik Janzing, Bernhard Schölkopf	
Nonparametric Divergence Estimation with Applications to Machine Learning on Distributions	599
Barnabas Poczos, Liang Xiong, Jeff Schneider	
Compressed Inference for Probabilistic Sequential Models	609
Gungor Polatkan, Oncel Tuzel	
Fast MCMC Sampling for Markov Jump Processes and Continuous Time Bayesian Networks.	619
Vinayak Rao, Yee Whye Teh	
New Probabilistic Bounds on Eigenvalues and Eigenvectors of Random Kernel Matrices	627
Nima Reyhani, Hideitsu Hino, Ricardo Vigário	
Learning with Missing Features	635
Afshin Rostamizadeh, Alekh Agarwal, Peter Bartlett	
Symbolic Dynamic Programming for Discrete and Continuous State MDPs.	643
Scott Sanner, Karina Valdivia Delgado, Leliane Nunes de Barros	
Generalized Fast Approximate Energy Minimization via Graph Cuts: α -Expansion β -Shrink Moves	653
Mark Schmidt, Karteek Alahari	
An Efficient Algorithm for Computing Interventional Distributions in Latent Variable Causal Models	661
Ilya Shpitser, Thomas S. Richardson, James M. Robins	
Graph Cuts is a Max-Product Algorithm	671
Daniel Tarlow, Inmar E. Givoni, Richard S. Zemel, Brendan J. Frey	
Adjustment Criteria in Causal Diagrams: An Algorithmic Perspective	681
Johannes Textor, Maciej Liśkiewicz	

Learning Mixed Graphical Models from Data with p Larger than n	689
Inma Tur, Robert Castelo	
Robust Learning Bayesian Networks for Prior Belief.	698
Maomi Ueno	
Distributed Anytime MAP Inference	708
Joop van de Ven, Fabio Ramos	
A Sequence of Relaxations Constraining Hidden Variable Models	717
Greg Ver Steeg, Aram Galstyan	
The Structure of Signals: Causal Interdependence Models for Games of Incomplete Information	727
Michael P. Wellman, Lu Hong, Scott E. Page	
Generalised Wishart Processes	736
Andrew Gordon Wilson, Zoubin Ghahramani	
Sparse Matrix-Variate Gaussian Process Blockmodels for Network Modeling	745
Feng Yan, Zenglin Xu, Yuan (Alan) Qi	
Hierarchical Maximum Margin Learning for Multi-Class Classification	753
Jian-Bo Yang, Ivor W. Tsang	
Planar Cycle Covering Graphs	761
Julian Yarkony, Alexander T. Ihler, Charless C. Fowlkes	
Tightening MRF Relaxations with Planar Subproblems	770
Julian Yarkony, Ragib Morshed, Alexander T. Ihler, Charless C. Fowlkes	
Rank/Norm Regularization with Closed-Form Solution: Application to Subspace Clustering	778
Yao-Liang Yu, Dale Schuurmans	
Measuring the Hardness of Stochastic Sampling on Bayesian Networks with Deterministic Causalities: the k -Test	786
Haohai Yu, Robert van Engelen	
Risk Bounds for Infinitely Divisible Distribution	796
Chao Zhang, Dacheng Tao	
Kernel-based Conditional Independence Test and Application in Causal Discovery	804
Kun Zhang, Jonas Peters, Dominik Janzing, Bernhard Schölkopf	
Smoothing Multivariate Performance Measures.	814
Xinhua Zhang, Ankan Saha, S. V. N. Vishwanathan	
Belief Propagation by Message Passing in Junction Trees: Computing Each Message Faster Using GPU Parallelization	822
Lu Zheng, Ole Mengshoel, Jike Chong	
Sparse Topical Coding.	831
Jun Zhu, Eric P. Xing	
Testing Whether Linear Equations are Causal: A Free Probability Theory Approach	839
Jakob Zscheischler, Dominik Janzing, Kun Zhang	
ABSTRACTS	
Incentives in Group Decision-Making With Uncertainty and Subjective Beliefs	849
Ruggiero Cavallo	
Learning High-Dimensional DAGs with Latent and Selection Variables	850
Diego Colombo, Marloes H. Maathuis, Markus Kalisch, Thomas S. Richardson	

Characterization and Greedy Learning of Interventional Markov Equivalence Classes of DAGs . . .	851
Alain Hauser, Peter Bühlmann	
Correction for Hidden Confounders in the Genetic Analysis of Gene Expression	852
Jennifer Listgarten, Carl Kadie, Eric E. Schadt, David Heckerman	
Statistical Mechanics of Semi-Supervised Clustering in Sparse Graphs	853
Greg Ver Steeg, Aram Galstyan, Armen Allahverdyan	

Preface

This year’s Conference on Uncertainty in Artificial Intelligence (UAI) is the twenty-seventh of a series of meetings that began as a small workshop in 1985. Over the years, the UAI community has developed techniques that have been widely integrated into general Artificial Intelligence curricula and adopted in many fields outside of Artificial Intelligence. The meeting is now the premier international conference on issues relating to representation and management of uncertainty within the field of Artificial Intelligence. UAI has a wide scope that includes, but is not limited to, modeling, inference, learning and decision making under uncertainty, with an interest both in theory and in applied work. This year UAI also hosted the 8th Bayesian Modelling Applications Workshop, a focused forum for interchange among those interested in real world applications of graphical models and Bayesian networks.

This volume contains all papers presented at UAI 2011, held at the Campus Roger de Lluria of the Universitat Pompeu Fabra, in Barcelona, Spain, July 14–17, 2011. Papers appearing in this volume were subjected to a rigorous review process; 285 regular papers were submitted to UAI this year, and 96 were accepted (24 for plenary presentation and 72 for poster presentation). We are confident that the proceedings, like past UAI Conference Proceedings, will become an important archival reference for the field.

Based on the recommendations of the Program Committee and our own overview of the accepted papers, we selected one paper as recipient of the UAI 2011 Microsoft Best Paper Award and another paper as the recipient of the UAI 2011 Google Best Student Paper Award. These awards were given for outstanding technical contributions. The UAI 2011 Microsoft Best Paper Award was given to the paper *Sum-Product Networks: A New Deep Architecture* by Hoifung Poon and Pedro Domingos, and the UAI 2011 Google Best Student Paper Award was given to the paper *Generalised Wishart Processes* by Andrew Wilson and Zoubin Ghahramani. Two other papers deserve recognition. The runner up for the Best Paper Award was *A Sequence of Relaxations Constraining Hidden Variable Models* by Greg Ver Steeg and Aram Galstyan, and the runner up for the Best Student Paper Award was *Graph Cuts is a Max-Product Algorithm* by Daniel Tarlow, Inmar Givoni, Richard Zemel and Brendan Frey. The journal *Artificial Intelligence* (AIJ) has invited authors of best papers to submit journal versions of their papers and will work to ensure an expedited review process for these papers.

To accommodate the publishing traditions of research fields other than computer science, UAI allowed for the submission of *not-for-proceedings*-papers. Only extended abstracts of the five such papers appear at the end of these proceedings.

In addition to the presentation of technical papers, we were very pleased to have four distinguished invited speakers: James M. Rehg (Georgia Institute of Technology), Hod Lipson (Cornell University), Karl Friston (University College London), and, as Banquet Speaker, Colin Howson (University of Toronto).

UAI 2011 also hosted a tutorial day, on July 14th. Three tutorials were given on this day: Ilya Shpitser (Harvard University) reviewed “Causal Inference: from Effects of Interventions to Learning and Inference with Partial Observability”; Milind Tambe and Christopher Kiekintveld (University of Southern California and University of Texas at El Paso) discussed “Game Theory for Security: Lessons Learned from Deployed Applications”; and Pedro Domingos and Kristian Kersting (University of Washington and Fraunhofer IAIS) presented material on “Combining Logic and Probability: Languages, Algorithms and Applications”.

Fabio G. Cozman and Avi Pfeffer
Program Co-Chairs
Peter Grünwald
Conference Chair


Acknowledgments

The success of a rigorously reviewed conference hinges upon the diligent efforts of many parties. We assembled a Senior Program Committee and a Program Committee with internationally renowned experts in the range of specialties covered by UAI. At least three Program Committee members (or in some cases, auxiliary reviewers) reviewed each paper. Despite the time pressure imposed by rigid conference deadlines, these individuals did their best to supply balanced, fair, and helpful reviews to all authors of submitted papers. Our Program Committee, with the assistance of several auxiliary reviewers, did an outstanding job for which the program chairs and authors alike are quite grateful. The Senior Program Committee acted in a supervisory role, ensuring that paper reviews lived up to UAI standards, moderating discussion among Program Committee members and, ultimately, making recommendations to the program chairs. In some cases Senior Program Committee members wrote supplementary reviews or consulted outside experts to ensure that authors received fair and detailed feedback. We are extremely grateful for their efforts.


A number of people have made significant contributions towards making UAI 2011 possible. We acknowledge and thank:

- For all the help with the local organization of UAI 2011, Lluís Godó (local arrangements chair), Tommaso Flaminio and Enrico Marchioni (local arrangements), and Angel Gil-Estallo (venue arrangements).
- For their help with the financial aspects of the conference, Prakash Shenoy and Ellen Fuller.
- For her leadership in organizing the 8th Bayesian Modelling Applications Workshop, Ann Nicholson.
- For extraordinary help in managing the conference web site and other social media pages, Mark Crowley.
- For much needed help in managing copyright forms, Martijn de Jongh.
- For help during the conference, the student volunteers Joshua Guerin, David Knowles, Minyi Li, Nikos Karampatziakis, Pere Pardo, Amanda Vidal, Andrew Wilson.

Finally, we gratefully acknowledge the generous support provided by our sponsors: Artificial Intelligence Journal, Microsoft Research, Google Inc., IBM Research, Pascal II Network of Excellence.



Microsoft®
Research
IBM Research



Senior Program Committee

Alina Beygelzimer	IBM Research, United States
Max Chickering	Microsoft adCenter, United States
Adnan Darwiche	University of California, Los Angeles, United States
Nando de Freitas	University of British Columbia, Canada
Didier Dubois	Institut de Recherche en Informatique de Toulouse, France
Bob Givan	Purdue University, United States
Russell Greiner	University of Alberta, Canada
Tom Heskes	Radboud University Nijmegen, The Netherlands
Tommi Jaakkola	Massachusetts Institute of Technology, United States
Manfred Jaeger	Aalborg Universitet, Denmark
Jerome Lang	LAMSADE, Universite Paris-Dauphine, France
Helge Langseth	The Norwegian University of Science and Technology, Norway
Kathryn Laskey	George Mason University, United States
Chris Meek	Microsoft Research, United States
Joris Mooij	Radboud University Nijmegen, The Netherlands
Ann Nicholson	Monash University, Australia
David Poole	University of British Columbia, Canada
Thomas Richardson	University of Washington, United States
Teemu Roos	University of Helsinki, Finland
Regis Sabbadin	INRA, France
Teddy Seidenfeld	Carnegie Mellon University, United States
Christian Shelton	University of California, Riverside, United States
Prakash Shenoy	University of Kansas, United States
Ricardo Silva	University College London, United Kingdom
Satinder Singh	University of Michigan, United States
Jin Tian	Iowa State University of Science and Technology, United States
Yair Weiss	The Hebrew University of Jerusalem, Israel
Michael Wellman	University of Michigan, United States
Nevin Zhang	The Hong Kong University of Science and Technology, Hong Kong

Program Committee

Ryan Adams	University of Toronto, Canada
John-Mark Agosta	Imperium, United States
Ayesha Ali	University of Guelph, Canada
Russell Almond	Florida State University, United States
Eyal Amir	University of Illinois Urbana-Champaign, United States
Nahla Ben Amor	Institut Supérieur de Gestion, Tunisia
Dragomir Anguelov	Google, United States
Alessandro Antonucci	IDSIA, Switzerland
Dimitrios Antos	Harvard University, United States
Elias Bareinboim	University of California, Los Angeles, United States
Salem Benferhat	Universite d'Artois, France
Philippe Besnard	Universite Paul Sabatier, France
Concha Bielza	Technical University of Madrid, Spain
Blai Bonet	Universidad Simon Bolivar, Venezuela
Grégory Bonnet	University of Caen Lower-Normandy, France
Remco Bouckaert	University of Auckland, New Zealand
Darius Braziunas	University of Toronto, Canada
Emma Brunskill	University of California Berkeley, United States
Olivier Buffet	LORIA-INRIA, France
Tiberio Caetano	NICTA, Australia
Peter Carbonetto	University of Chicago, United States

Robert Castelo	Universitat Pompeu Fabra, Spain
Nicolò Cesa-Bianchi	Università degli Studi di Milano, Italy
Hei Chan	Institute of Statistical Mathematics, Japan
Yu-han Chang	USC ISI, United States
Tao Chen	EMC Research Lab, China
Yiling Chen	USC ISI, United States
Arthur Choi	University of California, Los Angeles, United States
Tianjiao Chu	University of Pittsburgh, United States
Tom Claassen	Radboud University Nijmegen, The Netherlands
Barry Cobb	Virginia Military Institute, United States
Michael Collins	Columbia University, United States
James Cussens	University of York, United Kingdom
Denver Dash	Intel Labs Pittsburgh, United States
Cassio Polpo de Campos	IDSIA, Switzerland
Luc de Raedt	Katholieke Universiteit Leuven, Belgium
Sebastien Destercke	CIRAD, France
Vanessa Didelez	University of Bristol, United Kingdom
Francisco Javier Díez	UNED, Spain
Arnaud Doucet	University of British Columbia, Canada
Marek Druzdzel	University of Pittsburgh, United States; Bialystok University of Technology, Poland
David Dunson	Duke University, United States
Alain Dutech	LORIA-INRIA, France
Frederick Eberhardt	Washington University in St Louis, United States
Tal El-Hay	IBM Research, Israel
Saher Esmeir	DRW Holdings, United States
Sander Evers	Radboud University Nijmegen, The Netherlands
Vitaly Feldman	IBM Research, United States
Alan Fern	Oregon State University, United States
María Julia Flores	Universidad de Castilla-La Mancha, Spain
Charles Fox	Sheffield University, United Kingdom
Yaakov (Kobi) Gal	Ben-Gurion University of the Negev, Israel
Aram Galstyan	University of Southern California, United States
Roman Garnett	Carnegie Mellon University, United States
Phan Giang	George Mason University, United States
Mark Girolami	University College London, United Kingdom
Amir Globerson	The Hebrew University of Jerusalem, Israel
Lluis Godo	Artificial Intelligence Research Institute - CSIC, Spain
Vibhav Gogate	University of Washington, United States
Judy Goldsmith	University of Kentucky, United States
Manuel Gomez	Universidad de Granada, Spain
Christophe Gonzales	LIP6-UPMC, France
Dilan Gorur	University of California Irvine, United States
Stephen Gould	Australian National University, Australia
Yuhong Guo	Temple University, United States
Maya Gupta	University of Washington, United States
Petr Hajek	Academy of Sciences of the Czech Republic, Czech Republic
Milos Hauskrecht	University of Pittsburgh, United States
Elad Hazan	Technion, Israel
Tamir Hazan	TTI Chicago, United States
David Heckerman	Microsoft Research, United States
Jesse Hoey	University of Waterloo, Canada
Arjen Hommersom	Radboud University Nijmegen, The Netherlands
Eric Horvitz	Microsoft Research, United States
Patrik Hoyer	University of Helsinki, Finland

Bert Huang	Columbia University, United States
Alexander T. Ihler	University of California, Irvine, United States
Dominik Janzing	MPI for Intelligent Systems, Germany
David Jensen	University of Massachusetts Amherst, United States
Nebojsa Jojic	Microsoft Research, United States
Matti Kääriäinen	Nokia Research Center, Finland
Souhila Kaci	Universite d'Artois, France
Leslie Kaelbling	MIT, United States
Satyen Kale	Yahoo! Research, United States
Markus Kalisch	ETH Zurich, Switzerland
Ece Kamar	Microsoft Research, United States
Kalev Kask	University of California Irvine, United States
Sergey Kirshner	Purdue University, United States
Mikko Koivisto	Helsinki Institute for Information Technology, Finland
Petri Kontkanen	HIIT/CoSCo, Finland
Kevin Korb	Monash University, Australia
Andreas Krause	ETH Zurich, Switzerland
Manabu Kuroki	Osaka University, Japan
Branislav Kveton	Intel Labs, United States
Simon Lacoste-Julien	University of Cambridge, United Kingdom
Hugo Larochelle	Université de Sherbrooke, Canada
Pedro Larranaga	Technical University of Madrid, Spain
Jonathan Lawry	University of Bristol, United Kingdom
Jan Lemeire	Vrije Universiteit Brussel, Belgium
Alex Lenkoski	Heidelberg University, Germany
Weiru Liu	Queen's University Belfast, United Kingdom
Dan Lizotte	University of Michigan, United States
Daniel Lowd	University of Oregon, United States
Aurelie Lozano	IBM Research, United States
Thomas Lukasiewicz	Oxford University, United Kingdom
Marloes Maathuis	ETH Zurich, Switzerland
Omid Madani	SRI AI Center, United States
Dmitry Malioutov	DRW Holdings, United States
Subramani Mani	Vanderbilt University, United States
Vikash Mansinghka	Navia Systems Inc. and MIT, United States
Dimitris Margaritis	Amazon.com (A9.com), United States
Benjamin Marlin	University of British Columbia, Canada
Srujana Merugu	IBM Research, United States
Ofer Meshi	The Hebrew University of Jerusalem, Israel
Taneli Mielikäinen	Nokia Research Center, United States
Brian Milch	Google, United States
Thomas Minka	Microsoft Research, United Kingdom
Claire Monteleoni	Columbia University, United States
Martin Neil	Queen Mary University of London, United Kingdom
Uri Nodelman	Stanford University, United States
Michael Osborne	Oxford University, United Kingdom
John Paisley	Princeton University, United States
Hector Palacios	Universidad Carlos III de Madrid, Spain
David Parkes	Harvard University, United States
David Pennoek	Yahoo! Research, United States
Patrice Perny	UPMC, Paris 6, France
Nathalie Peyrard	INRA, France
Cedric Pralet	ONERA, France
Doina Precup	McGill University, Canada

Bob Price	PARC, United States
David Pynadath	USC Institute for Creative Technologies, United States
Erik Quaeghebeur	Ghent University, Belgium
Emmanuel Rachelson	EDF Research and Development
Fabio Tozeto Ramos	University of Sydney, Australia
Roland Ramsahai	University of Cambridge, United Kingdom
Marc'Aurelio Ranzato	University of Toronto, Canada
Silja Renooij	Universiteit Utrecht, The Netherlands
Daniel Roy	University of Cambridge, United Kingdom
Rafael Rumi	University of Almeria, Spain
Daniil Ryabko	INRIA Lille, France
Antonio Salmeron	University of Almeria, Spain
Scott Sanner	NICTA, Australia
Mark Schmidt	INRIA/ENS, France
Dale Schuurmans	University of Alberta, Canada
Ross Shachter	Stanford University, United States
Shai Shalev-Shwartz	The Hebrew University of Jerusalem, Israel
Ohad Shamir	Microsoft Research, United States
Chung-chieh Shan	Rutgers University, United States
Ilya Shpitser	Harvard University, United States
Ajit Singh	University of Washington, United States
Moninder Singh	IBM Research, United States
Mathieu Sinn	University of Waterloo, Canada
Richard Socher	Stanford University, United States
Peter Spirtes	Carnegie Mellon University, United States
Harald Steck	Bell Labs Alcatel-Lucent, United States
Oliver Stegle	Max Planck Institute, Germany
Charles Sutton	University of Edinburgh, United Kingdom
Markus Svensén	Microsoft Research, United Kingdom
Csaba Szepesvari	University of Alberta, Canada
Ben Taskar	University of Pennsylvania, United States
Yee Whye Teh	Gatsby Computational Neuroscience Unit UCL, United Kingdom
Florent Teichteil-Königsbuch	ONERA - The French Aerospace Lab, France
Bo Thiessen	Microsoft Research, United States
Robert Tillman	Carnegie Mellon University, United States
Jo-Anne Ting	University of British Columbia, Canada
Michalis K. Titsias	University of Manchester, United Kingdom
Volker Tresp	Siemens AG, Germany
Ioannis Tsamardinos	University of Crete, Greece
Richard Turner	University of Cambridge, United Kingdom
Charles Twardy	George Mason University, United States
Marco Valtorta	University of South Carolina, United States
Linda van der Gaag	Universiteit Utrecht, The Netherlands
Jirka Vomlel	Academy of Sciences of the Czech Republic, Czech Republic
Yi Wang	National University of Singapore, Singapore
Renata Wassermann	Universidade de São Paulo, Brazil
Max Welling	University of California Irvine, United States
Paul Weng	Universite Paris 6, France
Ydo Wexler	DRW Trading, United States
Wim Wiegerinck	Radboud University Nijmegen, The Netherlands
Nic Wilson	University College Cork, Ireland
David Wingate	MIT, United States
Pierre-Henri Wuillemin	Universite Paris 6, France

Dit-Yan Yeung	The Hong Kong University of Science and Technology, Hong Kong
Chun-Nam Yu	University of Alberta, Canada
Yaming Yu	University of California Irvine, United States
Changhe Yuan	Mississippi State University, United States
Luke Zettlemoyer	University of Washington, United States
Jiji Zhang	Lingnan University, Hong Kong
Kun Zhang	MPI for Biological Cybernetics, Germany
Onno Zoeter	Xerox Research Centre Europe, France

Auxiliary reviewers

Arash Afkanpour, Saeed Amizadeh, Jaesik Choi, José Costa, Fernando Fernandez, John Fisher III, José A. Gámez, Perry Groot, Marek Grzes, Joshua T. Guerin, Bernd Gutmann, Hannaneh Hajishirzi, Yi Huang, Guillaume Infantes, Xueyan Jiang, Vincenzo Lagani, Brandon Malone, Steven Mascaro, Nicholas Mattei, Julian McAuley, Stijn Meganck, Richard Neapolitan, James Neufeld, Maximilian Nickel, Jens D. Nielsen, James Petteerson, Leonard Kin Man Poon, José M. Puerta, Novi Quadrianto, Deepak Ramachandran, Debajyoti Ray, Merel Rietbergen, Jonah Schupbach, Alexander J. Smola, Olivier Spanjaard, Peter Stobbe, Ingo Thon, Hamed Valizadegan, Michal Valko, Martijn van Otterlo, Kush R. Varshney, Marina Velikova, Fabio Vitale, Jingsong Wang, Jun Wang, Martha White, David Whitehouse, Yu Zhang.