

# Charles Sutton

School of Informatics  
University of Edinburgh  
Edinburgh, EH8 9AB

Email: [csutton@inf.ed.ac.uk](mailto:csutton@inf.ed.ac.uk)  
Phone: +44 (0)131 651 5634  
<http://homepages.inf.ed.ac.uk/csutton>

## Education

PhD in computer science. 2008. University of Massachusetts Amherst.  
Dissertation: Efficient Training Methods for Conditional Random Fields  
Advisor: Andrew McCallum  
Nominated for ACM Doctoral Dissertation Award

MS in computer science. 2005. University of Massachusetts Amherst.

BA in computer science and philosophy. 1999. St. Mary's College of Maryland.  
*Summa cum laude*

## Professional Positions

**Reader.** (Equivalent rank to Associate Professor in US.) University of Edinburgh. Director of EPSRC Centre of Doctoral Training in Data Science (2014-2022), which fully funds 50 four-year studentships in Data Science over 8 years. August 2014-present.

**Lecturer.** (Equivalent rank to Assistant Professor in US.) University of Edinburgh. September 2009-July 2014.

**Turing Fellow,** Alan Turing Institute, June 2016 - present.

**Postdoctoral Researcher.** University of California, Berkeley. August 2007-August 2009.  
Supervisor: Michael I. Jordan

**Research Assistant.** University of Massachusetts Amherst. September 2003-August 2007.  
Advisor: Andrew McCallum

**Research Intern.** Microsoft Research, Cambridge, UK. Summer 2006.  
Mentor: Tom Minka

**Research Assistant.** University of Massachusetts Amherst. September 2000-August 2003.  
Advisor: Paul R. Cohen

**Instructor.** Center for Talented Youth, Lancaster, PA. Summer 2003 and 2004. Instructor for three-week course in theory of computation to three classes of about fifteen gifted high school students.

**Programmer.** Cypcorp, Inc., Austin, TX. June 1999-August 2000.

**Programmer.** Veridian Systems, California, MD. Summer 1997 and 1998.

## Grants

LUCID: Clearer Software by Integrating Natural Language Analysis into Software Engineering, EPSRC grant EP/P005314/1, 2017-2020, £377k (UoE portion; joint with University College London).

LODE: Unsupervised Machine Learning for Organizing and Exploring Data. Unrestricted gift from Huawei, 2017, £100k.

Reducing the Annotation Tax of Programming Language Types using Machine Learning and Big Code (PI). Microsoft Research PhD Scholarship. 2017-2020.

Statistical Natural Language Processing Methods for Computer Program Source Code (PI), EPSRC grant EP/K024043/1, 2013-2016, £460k.

Fast, Locally Adaptive Inference for Machine Learning in Graphical Models (PI), EPSRC grant EP/J00104X/1, 2011-2014, £100k.

Data Driven Development Tools: Statistical Language Processing for Programming Language Text (PI). Microsoft Research PhD Scholarship. 2012-2015.

EPSRC Centre for Doctoral Training in Data Science (coI and deputy director). EP/L016427/1. £4.5M.

IDEAL: Intelligent Domestic Energy Advice Loop (coI), EPSRC grant EP/K002732/1, 2012-2016, £2.1M.

Data-Driven Methods for a New National Household Energy Survey (coI), EPSRC grant EP/M008223/1, 2015-2018, £750k.

App Guarden: Resilient Application Stores (coI), EPSRC grant EP/K032666/1, 2012-2015, £723k.

Amazon Web Services Research Credit, 2012-2014.

## Students Advised

Kai Xu, PhD student, 2017-present.

Irene Vlasi-Pandi, PhD student, 2017-present.

Lazar Valkov, PhD student, 2016-present.

Ryan Davies, Centre for Doctoral Training (MSc+PhD student), 2016-present.

Simão Eduardo, Centre for Doctoral Training (MSc+PhD student), 2016-present.

Michael-Rafael Karampatsis, Centre for Doctoral Training (MSc+PhD student), 2015-present.

Maria Gorinova, PhD student, 2016-present. Co-supervised with Prof Andrew D Gordon.

Akash Srivastava, PhD student, 2015-present.

Miltiadis Allamanis, PhD, 2012-2016.

Krzysztof Geras, PhD, 2011-2016.

Yichuan Zhang, PhD, 2010-2015.

Daniel Renshaw, secondary supervisor for PhD, University of Edinburgh, 2012-2015.

Joel Lang, secondary supervisor for PhD, University of Edinburgh, 2009-2012.

Ioannis Konstas, secondary supervisor for PhD, University of Edinburgh, 2011-2015.

Greg Coppola, secondary supervisor for PhD, University of Edinburgh, 2011-2014.

MSc thesis supervision, University of Edinburgh: Junyi Yang, Minxue Xia, João Gaurda, Ridho Rahman, Linan Gong, Nikos Katirzis, Stefan Fiott, Madeleine Vazquez, Nadi Awad, Sjoerd Dost, Agnieszka Bomersbach, Sotirios Antonopoulos, Yitong Hu, Maria Alecu, Sergey Dudoladov, Joaquim Castella, Facundo Bellosi, Shihan Wang, Ilja Kuzborskij, Aron Deak, Ouye Xie, Spyros Marketos, Xiaohu Liu, Anca-Elena Leuca, Yichuan Zhang

BSc thesis supervision, University of Edinburgh: Damodar Sojka, Ioannis Karamanlakis, Stefan Sabev (MInf), Romyana Rumenova, Balint Vekerdy (MInf), Jacob Essex, Sean Bowie, Scott Hofmann, Blagovest Velez, Dimitar Dimitrov, Ivan Trendafilov

David Matthews, member of PhD committee, University of Edinburgh, 2011

Michael Auli, internal assessor for third-year PhD review, 2011

Rob Hall, second-year PhD student, University of Massachusetts, 2006-2007, transferred and finished PhD at Carnegie Mellon University

Michael Sindelar, undergraduate, 2006, now PhD student at University of Massachusetts Amherst

David Nachum, undergraduate, UC Berkeley, 2009

Abishek Arun, internal assessor for third-year PhD review, 2010

Robert McWilliam, external assessor for third-year EngD review, 2010

### **Full-Time Researchers Supervised**

Dr Annie Louis, postdoctoral researcher, April 2017-present.

Dr Mingjun Zhong, postdoctoral researcher, July 2013-April 2017.

Dr Jaroslav Fowkes, postdoctoral researcher, October 2013-August 2016.

Pankajan Chanthirasegaran, November 2015-September 2016.

### **PhD Examination**

Sébastien Bratières, external examiner, University of Cambridge, August 2017

Tim Rocktäschel, external examiner, University College London, July 2017

Veselin Raychev, external member of PhD committee, ETH Zurich, August 2016

Thanh Tran, external member of PhD committee, University of Massachusetts, December 2012

Daniele Masato, external examiner, University of Aberdeen, December 2011

Thomas Joyce, internal examiner, November 2015

Jyri Kivinen, internal examiner, November 2013

Dimitrios Milios, internal examiner, October 2013

Sasa Petrovic, internal examiner, December 2012

Andrew Dai, internal examiner, September 2012

Shahzad Asif, internal examiner, April 2012

### **Teaching Experience**

Courses Taught: University of Edinburgh: Data Mining and Exploration (MSc course, Spring 2010, 2011, 2012, 2013). Introduction to Applied Machine Learning (Undergraduate, Autumn 2010 and 2011), Machine Learning and Pattern Recognition (MSc; Spring 2013, Autumn 2013). Introduction to Research in Data Science (Autumn 2014, Autumn 2015).

Instructor, Center for Talented Youth. Summer 2003 and 2004

## Software Packages

GRMM. Created GRMM, an open-source toolkit for inference in graphical models. Implements training for conditional random fields with arbitrary structure and parameter tying, with inference using junction tree, generalized belief propagation, or Gibbs sampling. Available at <http://mallet.cs.umass.edu/grmm/>

MALLET. Major contributor to MALLET, open-source toolkit for document classification and sequence labeling using conditional random fields. <http://mallet.cs.umass.edu/>

## Professional Activities

Director, EPSRC Centre for Doctoral Training in Data Science, January 2016 – August 2017. Served as Deputy Director from April 2014-December 2016.

Member of the Peer Review College, Engineering and Physical Sciences Research Council, 2017–present.

Expert reviewer, Horizon 2020 funding, 2017.

Area Chair, International Conference on Machine Learning (ICML), 2017.

Programme Committee member, ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), 2017.

Co-organizer, NIPS 2016 Workshop on Artificial Intelligence for Data Science, December 2016, Barcelona, Spain.

Co-organizer, Symposium on Communicating Machine Learning, sponsored by the Alan Turing Institute, August 2016.

Member of the Steering Committee, EPSRC Network on Computational Statistics and Machine Learning, 2013-2016. Chaired the Steering Committee 2015-2016.

Organizer, Cross-CDT Workshop on Statistics and Machine Learning, August 2016.

Programme Committee member, Workshop on Natural Language Processing and Software Engineering, Microsoft Research, Seattle WA, USA, November 2016.

Programme Committee member, Workshop on Natural Language Processing and Software Engineering, Microsoft Research, Redmond WA, USA, October 2015.

Area Chair, International Conference on Artificial Intelligence and Statistics (AISTATS 2016)

Co-organizer, Dagstuhl Workshop on “Big Code”, 2015.

Area Chair, Neural Information Processing Systems (NIPS) conference, 2014.

Associate Editor, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2013-2014.

Panel member, ICT Prioritization Panel, Engineering and Physical Sciences Research Council (EPSRC), January 2015.

Production Editor, Journal for Machine Learning Research (JMLR), 2014-present.

Member of the Editorial Board, Journal for Artificial Intelligence Research (JAIR), 2013-2014

Area Chair, Conference of the European Chapter of the Association for Computational Linguistics (EACL 2012).

Area Chair, International Conference on Artificial Intelligence and Statistics (AISTATS 2012)

Local Arrangements Chair, Twenty-Ninth International Conference on Machine Learning (ICML 2012)

Co-organizer, Bayes-250 workshop, held to mark the 250th anniversary of the death of Rev Thomas Bayes, 5-7 Sept 2012.

Fundraising Chair, Twelfth, Thirteenth, and Fourteenth International Conferences on Artificial Intelligence and Statistics (AISTATS 2009, 2010, and 2011).

Co-organizer, Workshop on Joint Inference in Natural Language Processing. Held June 8, 2006, at HLT/NAACL.

Reviewing. *Conferences*: NIPS, ICML, ICLR, UAI, ACL, EMNLP, NAACL, IJCAI, CoNLL, AISTATS, PLDI, OOPSLA. *Journals*: Journal of Machine Learning Research, Machine Learning Journal, Transactions of the Association for Computational Linguistics (TACL), ACM Transactions on Speech and Language Processing, Statistics and Computing. *Funding Agencies*: National Science Foundation, EU Horizon 2020, Engineering and Physical Sciences Research Council, ETH Zurich Postdoctoral Fellowship Program.

PhD selector, School of Informatics, 2012-2014.

MSc Specialist Area Advisor, Learning from Data Specialism, University of Edinburgh, 2010-2012.

MInf Project Coordinator, University of Edinburgh, 2010-2012.

Organizer, Seminar Series for Institute of Adaptive and Neural Computation, University of Edinburgh, 2009-2014.

Organizer, Center for Intelligent Systems Seminar, University of California, Berkeley, 2008-2009.

## Invited Talks

Amazon (Cambridge, UK), August 2017.

Dagstuhl workshop on Machine Learning and Formal Methods, August 2017.

sourced tech talk, Moscow, June 2017.

Facebook London, May 2017.

University of Glasgow, January 2017.

Invited speaker, Google Zurich, August 2016.

International Conference of the ERCIM Working Group on Computational and Methodological Statistics (CMStatistics 2015). Invited to speak in parallel session. London, December 2015.

Keynote speaker, Workshop on Natural Language Processing and Software Engineering, Microsoft Research, Redmond WA, USA, October 2015.

EPFL, Lausanne, Switzerland. November 2015.

University of Washington, Seattle WA, USA, October 2015.

CMStatistics 2015 (parallel session). London, 2015.

Microsoft Research, Redmond WA, USA, May 2015.

University of Massachusetts Amherst, December 2014.

Keynote speaker, 13th International Conference on Intelligent Virtual Agents. August 2013.

Keynote speaker, Workshop on Collective Learning and Inference on Structured Data (CoLISD). Bristol, UK. September 2012.

Gatsby Computational Neuroscience Unit, University College London, November 2012.

Fifth International Conference of the ERCIM Working Group on Computing and Statistics (ERCIM 2012) (parallel session). Oviedo, Spain.

INFORMS 2012 meeting (parallel session). Phoenix, AZ, USA.

Cambridge University, July 2011.

Robert Gordon University, May 2011.

University of Sheffield, March 2011.

University of Newcastle, Department of Statistics, March 2011.

Xerox Research Centre Europe. Grenoble, France. February 2011.

University of Glasgow. Sept 2010.

University College London, Gatsby Computational Neuroscience Institute. May 2010.

Cambridge University. March 2010.

Carnegie Mellon University. April 2007.

Stanford University. April 2007.

## Awards

Nominated for ACM Doctoral Dissertation Award, 2008  
One of two nominations from UMass.

Honorable Mention, NSF Graduate Research Fellowship, 2001

Phi Beta Kappa, 1999

## Publications

### Journals

- [1] Jaroslav Fowkes, Razvan Ranca, Miltiadis Allamanis, Mirella Lapata, and Charles Sutton. Autofolding for source code summarization. *Transactions on Software Engineering*, 2017. In press.
- [2] Weikun Wang, Giuliano Casale, and Charles Sutton. A Bayesian approach to parameter inference in queueing networks. *ACM Transactions on Modeling and Computer Simulation*, 27(1), August 2016.
- [3] Charles Sutton and Andrew McCallum. An Introduction to Conditional Random Fields. *Foundations and Trends in Machine Learning*, 4(4):267–373, 2012. (Over 200 citations on Google Scholar.)
- [4] Charles Sutton and Michael I. Jordan. Bayesian inference in queueing networks. *Annals of Applied Statistics*, 5(1):254–282, 2011.
- [5] Charles Sutton and Andrew McCallum. Piecewise training for structured prediction. *Machine Learning*, 77(2–3):165–194, 2009.

- [6] Charles Sutton, Andrew McCallum, and Khashayar Rohanimanesh. Dynamic conditional random fields: Factorized probabilistic models for labeling and segmenting sequence data. *Journal of Machine Learning Research*, 8:693–723, March 2007. (Over 380 citations on Google Scholar.)
- [7] Charles Sutton. Computers and Octi: Report from the 2001 tournament. *ICGA Journal*, 25(2):105–112, June 2002.

### Refereed Conferences

- [8] Akash Srivastava, Lazar Valkov, Chris Russell, Michael Gutmann, and Charles Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. *Advances in Neural Information Processing Systems (NIPS)*, 2017. To appear.
- [9] Miltiadis Allamanis, Pankajan Chanthirasegaran, Pushmeet Kohli, and Charles Sutton. Learning continuous semantic representations of symbolic expressions. *International Conference on Machine Learning (ICML)*, 2017.
- [10] Akash Srivastava and Charles Sutton. Autoencoding variational inference for topic models. *International Conference on Learning Representations (ICLR)*, 2017.
- [11] Jaroslav Fowkes and Charles Sutton. A Bayesian Network Model for Interesting Itemsets. *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery (ECML-PKDD)*, 2016.
- [12] Krzysztof Geras and Charles Sutton. Composite denoising autoencoders. *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery (ECML-PKDD)*, 2016.
- [13] Jaroslav Fowkes and Charles Sutton. Parameter-free probabilistic API mining across GitHub. *Foundations of Software Engineering (FSE)*, 2016. (27% acceptance rate).
- [14] Jaroslav Fowkes and Charles Sutton. A subsequence interleaving model for sequential pattern mining. *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
- [15] Miltiadis Allamanis, Hao Peng, and Charles Sutton. A Convolutional Attention Network for Extreme Summarization of Source Code. *International Conference in Machine Learning (ICML)*, 2016.
- [16] Wei Chen, David Aspinall, Andrew D Gordon, Charles Sutton, and Igor Muttik. More semantics more robust: Improving android malware classifiers. *ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2016. (25% acceptance rate).
- [17] Jaroslav Fowkes, Pankajan Chanthirasegaran, Razvan Ranca, Miltiadis Allamanis, Mirella Lapata, and Charles Sutton. TASSAL: Autofolding for source code summarization. *International Conference on Software Engineering (ICSE)*, 2016. Demo track (32% acceptance rate).
- [18] Wei Chen, David Aspinall, Andrew Gordon, Charles Sutton, and Igor Muttik. Learning and verifying unwanted behaviours. *Workshop on Hot Issues in Security Principles and Trust (HotSpot 2016)*, 2016.
- [19] Wei Chen, David Aspinall, Andrew Gordon, Charles Sutton, and Igor Muttik. On robust malware classifiers by verifying unwanted behaviours. *International Conference on Integrated Formal Methods*, 2016.

- [20] Mingjun Zhong, Nigel Goddard, and Charles Sutton. Latent Bayesian melding for integrating individual and population models. *Advances in Neural Information Processing Systems (NIPS)*, 2015. Spotlight (4% of submissions).
- [21] Miltiadis Allamanis, Earl T. Barr, Christian Bird, and Charles Sutton. Suggesting accurate method and class names. *Foundations of Software Engineering (FSE)*, 2015.
- [22] Krzysztof Geras and Charles Sutton. Scheduled denoising autoencoders. *International Conference on Representation Learning (ICLR)*, 2015.
- [23] Yichuan Zhang and Charles Sutton. Semi-separable Hamiltonian Monte Carlo for inference in Bayesian hierarchical models. *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [24] Mingjun Zhong, Nigel Goddard, and Charles Sutton. Signal aggregate constraints in additive factorial HMMs, with application to energy disaggregation. *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [25] Miltiadis Allamanis, Earl T Barr, Christian Bird, and Charles Sutton. Learning natural coding conventions. *Symposium on the Foundations of Software Engineering (FSE)*, 2014. Winner, ACM SIGSOFT Distinguished Paper Award.
- [26] Miltos Allamanis and Charles Sutton. Mining idioms from source code. *Symposium on the Foundations of Software Engineering (FSE)*, 2014.
- [27] Quim Castella and Charles Sutton. Word storms: Multiples of word clouds for visual comparison of documents. *International World Wide Web Conference (WWW)*, 2014.
- [28] Miltos Allamanis and Charles Sutton. Mining source code repositories at massive scale using language modeling. *Working Conference on Mining Software Repositories (MSR)*, 2013.
- [29] Miltos Allamanis and Charles Sutton. Why, when, and what: Analyzing stack overflow questions by topic, type, and code. *Working Conference on Mining Software Repositories (MSR)*, 2013.
- [30] Krzysztof Geras and Charles Sutton. Multiple-source cross validation. *International Conference on Machine Learning (ICML)*, 2013.
- [31] Thanh T. L. Tran, Yanlei Diao, Charles Sutton, and Anna Liu. Supporting user-defined functions on uncertain data. *Proceedings of the VLDB Endowment (PVLDB)*, 2013.
- [32] Yichuan Zhang, Charles Sutton, Amos Storkey, and Zoubin Ghahramani. Continuous relaxations for discrete Hamiltonian Monte Carlo. *Advances in Neural Information Processing Systems (NIPS)*, 2012. Spotlight (5% of submissions).
- [33] Yichuan Zhang and Charles Sutton. Quasi-Newton Markov chain Monte Carlo. *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [34] Zhao Cao, Charles Sutton, Yanlei Diao, and Prashant Shenoy. Distributed inference and query processing for RFID tracking and monitoring. *Proceedings of the VLDB Endowment (PVLDB)*, 4(5):326–337, 2011.
- [35] Charles Sutton and Michael I. Jordan. Learning and inference in queueing networks. *Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.
- [36] Peter Bodik, Rean Griffith, Charles Sutton, Armando Fox, Michael I. Jordan, and David A. Patterson. Statistical machine learning makes automatic control practical for internet



- datacenters. *Workshop on Hot Topics in Cloud Computing (HotCloud '09)*, 2009. (Over **130** citations on Google Scholar.)
- [37] Thanh Tran, Charles Sutton, Richard Cocci, Yanming Nie, Yanlei Diao, and Prashant Shenoy. Probabilistic inference over RFID streams in mobile environments. *International Conference on Data Engineering (ICDE)*, 2009. (Over **100** citations on Google Scholar.)
- [38] Yanlei Diao, Boduo Li, Anna Liu, Liping Peng, Charles Sutton, Thanh Tran, and Michael Zink. Capturing data uncertainty in high-volume stream processing. *Conference on Innovative Data Systems Research (CIDR)*, 2009.
- [39] Charles Sutton and Michael I. Jordan. Probabilistic inference in queueing networks. *Workshop on Tackling Computer Systems Problems with Machine Learning Techniques (SYSML)*, 2008.
- [40] Blaine Nelson, Marco Barreno, Fuching Jack Chi, Anthony D. Joseph, Benjamin I. P. Rubinstein, Udam Saini, Charles Sutton, J. D. Tygar, and Kai Xia. Exploiting machine learning to subvert your spam filter. *Proceedings of the First USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, 2008. (Over **65** citations on Google Scholar.)
- [41] Robert Hall, Charles Sutton, and Andrew McCallum. Unsupervised deduplication using cross-field dependencies. *Conference on Knowledge Discovery and Data Mining (KDD)*, 2008.
- [42] Charles Sutton and Andrew McCallum. Piecewise pseudolikelihood for efficient CRF training. *International Conference on Machine Learning (ICML)*, 2007. (Over **75** citations on Google Scholar.)
- [43] Charles Sutton and Andrew McCallum. Improved dynamic schedules for belief propagation. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2007.
- [44] Charles Sutton, Michael Sindelar, and Andrew McCallum. Reducing weight undertraining in structured discriminative learning. *Conference on Human Language Technology and North American Association for Computational Linguistics (HLT-NAACL)*, 2006.
- [45] Chris Pal, Charles Sutton, and Andrew McCallum. Sparse forward-backward using minimum divergence beams for fast training of conditional random fields. *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2006.
- [46] Charles Sutton and Andrew McCallum. Piecewise training of undirected models. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2005. (Over **150** citations on Google Scholar.)
- [47] Charles Sutton and Andrew McCallum. Composition of conditional random fields for transfer learning. *Conference on Human Language Technology and Empirical Methods in Natural Language Processing (HLT-EMNLP)*, 2005.
- [48] Max Welling and Charles Sutton. Learning in Markov random fields with contrastive free energies. *Conference on Artificial Intelligence and Statistics (AISTATS)*, 2005.
- [49] Charles Sutton and Andrew McCallum. Collective segmentation and labeling of distant entities in information extraction. *ICML Workshop on Statistical Relational Learning and Its Connections to Other Fields*, 2004. (Over **130** citations on Google Scholar.)
- [50] Charles Sutton, Khashayar Rohanimanesh, and Andrew McCallum. Dynamic conditional random fields: Factorized probabilistic models for labeling and segmenting sequence data. *International Conference on Machine Learning (ICML)*, 2004.

## Book Chapters

- [51] Blaine Nelson, Marco Barreno, Fuching Jack Chi, Anthony D. Joseph, Benjamin I. P. Rubinstein, Udam Saini, Charles Sutton, J. D. Tygar, and Kai Xia. Misleading learners: Co-opting your spam filter. Jeffrey J. P. Tsai and Philip S. Yu, editors, *Machine Learning in Cyber Trust: Security, Privacy, Reliability*. Springer, 2009.
- [52] Charles Sutton and Andrew McCallum. An introduction to conditional random fields for relational learning. Lise Getoor and Ben Taskar, editors, *Introduction to Statistical Relational Learning*. MIT Press, 2007. (Over 850 citations on Google Scholar.)

## Dissertation

- [53] Charles Sutton. *Efficient Training Methods for Conditional Random Fields*. PhD thesis, University of Massachusetts, 2008.

## Workshops

- [54] Akash Srivastava, James Zou, Ryan P. Adams, and Charles Sutton. Clustering with a Reject Option: Interactive Clustering as Bayesian Prior Elicitation. *Workshop on Human Interpretability in Machine Learning Workshop on Human Interpretability in Machine Learning (co-located with ICML)*, 2016.
- [55] Krzysztof J. Geras, Abdel rahman Mohamed, Rich Caruana, Gregor Urban, Shengjie Wang, Ozlem Aslan, Matthai Philipose, Matthew Richardson, and Charles Sutton. Blending LSTMs into CNNs. *International Conference on Learning Representations (ICLR Workshop)*, 2016.
- [56] Wei Chen, Charles Sutton, Andrew Gordon, David Aspinall, Igor Muttik, and Qi Shen. Compact explanations of why malware is bad. *International Workshop on the Use of AI in Formal Methods (AI4FM)*, 2015.
- [57] Wei Chen, Charles Sutton, David Aspinall, Andrew Gordon, Qi Shen, and Igor Muttik. Verifying anti-security policies learnt from android malware families. *International Seminar on Program Verification, Automated Debugging and Symbolic Computation*, 2015.
- [58] Peter Bodik, Rean Griffith, Charles Sutton, Armando Fox, Michael I. Jordan, and David A. Patterson. Automatic exploration of datacenter performance regimes. *First Workshop on Automated Control for Datacenters and Clouds (ACDC '09)*, 2009.
- [59] Hanna Wallach, Charles Sutton, and Andrew McCallum. Bayesian modeling of dependency trees using hierarchical Pitman-Yor priors. *ICML Workshop on Prior Knowledge for Text and Language Processing*, 2008.
- [60] Peter Bodik, Charles Sutton, Armando Fox, David Patterson, and Michael I. Jordan. Response-time modeling for resource allocation and energy-informed SLAs. *NIPS Workshop on Statistical Learning Techniques for Solving Systems Problems (MLSys 07)*, 2007.
- [61] Charles Sutton, Chris Pal, and Andrew McCallum. Sparse forward-backward for fast training of conditional random fields. *NIPS Workshop on Structured Learning for Text and Speech Processing*, 2005.
- [62] Andrew McCallum and Charles Sutton. Piecewise training with parameter independence diagrams: Comparing globally- and locally-trained linear-chain CRFs. *NIPS Workshop on Learning with Structured Outputs*, 2004.

- [63] Andrew McCallum, Khashayar Rohanimanesh, and Charles Sutton. Dynamic conditional random fields for jointly labeling multiple sequences. *NIPS Workshop on Syntax, Semantics, and Statistics*, December 2003.
- [64] Charles Sutton, Brendan Burns, Clayton Morrison, and Paul R. Cohen. Guided incremental construction of belief networks. *International Symposium on Intelligent Data Analysis*. Springer-Verlag, 2003.
- [65] Paul R. Cohen and Charles Sutton. Very predictive n-grams for space-limited probabilistic models. *International Symposium on Intelligent Data Analysis*. Springer-Verlag, 2003.
- [66] Brendan Burns, Charles Sutton, Clayton Morrison, and Paul R. Cohen. Information theory and representation in associative word learning. *Third International Workshop on Epigenetic Robotics*. Lund University Cognitive Studies, Volume 101, 2003.
- [67] Paul R. Cohen, Charles Sutton, and Brendan Burns. Learning effects of robot actions using temporal associations. *International Conference on Development and Learning (ICDL)*, 2002.

### **Unrefereed Papers**

- [68] Charles Sutton and Tom Minka. Local training and belief propagation. Technical Report TR-2006-121, Microsoft Research, 2006.
- [69] Charles Sutton and Andrew McCallum. Fast, piecewise training for discriminative finite-state and parsing models. Technical Report IR-403, Center for Intelligent Information Retrieval, 2005.
- [70] Charles Sutton. Conditional probabilistic context-free grammars. Master's thesis, University of Massachusetts, 2004.