

Sameer Singh

Education

- 2014 **PhD (Computer Science)**, *University of Massachusetts*, Amherst, MA.
Title: Scaling MCMC Inference and Belief Propagation to Large, Dense Graphical Models
- 2007 **MS (Computer Science)**, *Vanderbilt University*, Nashville, TN.
- 2004 **BE (Electronics & Comm Engg)**, *NSIT (Delhi University)*, New Delhi, India.

Professional Experience

Academic

- July 2016– **Assistant Professor**, *Computer Science*, University of California, Irvine.
- Oct 2016– **Assistant Professor**, *Language Science (courtesy appointment)*, University of California, Irvine.
- Apr 2017– **Assistant Professor**, *Elec Engg & CS (courtesy appointment)*, University of California, Irvine.
- 2013–2016 **Postdoctoral Research Associate**, *Computer Science*, University of Washington, Seattle.

Industry

- June–Sept 2012 **Research Intern**, *Microsoft Research*, Cambridge, UK.
- June–Sept 2010 **Research Intern**, *Google Research*, Mountain View, CA.
- July–Sept 2009 **Applied Research Intern**, *Yahoo! Labs*, Santa Clara, CA.
- May–Dec 2007 **Software Engineer Intern**, *Google Inc*, Pittsburgh, PA.
- May–Aug 2006 **Intern**, *Adv Computing Center for Research and Education (ACCRE)*, *Vanderbilt University*, Nashville, TN.

Teaching

- CS 272 **Statistical Natural Language Processing**, *Graduate Course*, UC Irvine.
Winter 2020, Winter 2019, Winter 2018, Winter 2017
- CS 175 **Projects in AI (in Minecraft)**, *Undergraduate Course*, UC Irvine.
Fall 2020, Spring 2019, Spring 2017, also published as **Sameer Singh**. Minecraft as a Platform for Project-Based Learning in AI. In *AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI)*, 2020
- CS 273A **Machine Learning**, *Graduate Course*, UC Irvine, Fall 2018, Fall 2017.
- CS 178 **Machine Learning and Data Mining**, *Undergraduate Course*, UC Irvine, Fall 2017.
- BANA 290 **ML for Text**, *Graduate course*, *Merage School of Business*, UC Irvine, Spring 2018 (co-instructor).
- Guest Lecturer **LAW 5903: The AI Frontier: Disrupting Legal Services**, *UCI Law School*, Fall/Spring '20, Spring '19.
Econ 229: Big Data, *UCI School of Economics*, Spring 2017.
CSEP 517: Natural Language Processing, *Univ of Washington*, Fall 2015.
CSE 546: Machine Learning, *Univ of Washington*, Fall 2015, Fall 2014.
CS691: Probabilistic Graphical Models, *UMass, Amherst*, Spring 2012.
- Teaching Assistant **Probabilistic Graphical Models**, *UMass, Amherst*, Spring 2011.
Intro to Programming, *UMass, Amherst*, Spring 2008.
Intro to Data Structures, *Vanderbilt*, Fall 2004, Spring 2005.

Awards

Research Awards

- 2020 **Hellman Fellow**, Hellman Family Foundation.
- 2020 **Dean's Mid-Career Award for Excellence in Research**, University of California, Irvine, CA.
- 2015 **DARPA Riser**, DARPA Wait, What? Event, St. Louis, MO.
- 2010-2011 **Yahoo! Key Scientific Challenges (KSC) Award**, in Machine Learning & Statistics.
- 2010-2011 **Accomplishments in Search & Mining Award**, UMass CS Dept and Yahoo!.

Paper Awards

- 2020 **Overall Best Paper Award**, Association of Computational Linguistics (ACL).
Awarded to *Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh*. *Beyond Accuracy: Behavioral Testing of NLP models with CheckList*. In Association for Computational Linguistics (ACL), 2020
- 2020 **Best Paper Runners-Up**, Automated Knowledge Base Completion (AKBC).
Awarded to *Pouya Pezeshkpour, Yifan Tian, and Sameer Singh*. *Revisiting Evaluation of Knowledge Base Completion Models*. In Automated Knowledge Base Construction (AKBC), 2020
- 2019 **Best Paper**, EMNLP Workshop on Machine Reading and Question Answering (MRQA), Hongkong.
Awarded to *Anthony Chen, Gabriel Stanovsky, Sameer Singh, and Matt Gardner*. *Evaluating Question Answering Evaluation*. In Workshop on Machine Reading and Question Answering (MRQA), 2019
- 2019 **Best Demo Paper**, Empirical Methods in Natural Language Processing (EMNLP), Hongkong.
Awarded to *Eric Wallace, Jens Tuyls, Junlin Wang, Sanjaya Subramanian, Matt Gardner, and Sameer Singh*. *AllenNLP Interpret: A Framework for Explaining Predictions of NLP Models*. In Demo at the Empirical Methods in Natural Language Processing (EMNLP), 2019
- 2018 **Honorable Mention for Best Paper**, Annual Meeting of Computational Linguistics (ACL), Melbourne.
Awarded to *Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin*. *Semantically Equivalent Adversarial Rules for Debugging NLP models*. In Association for Computational Linguistics (ACL), 2018
- 2017 **Amazon Best Poster Award**, Southern California ML Symposium, Los Angeles, CA.
Awarded to *Zhengli Zhao, Dheeru Dua, and Sameer Singh*. *Generating Natural Adversarial Examples*. In International Conference on Learning Representations (ICLR), 2018
- 2016 **Audience Appreciation Award**, International Conference of SIGKDD, San Francisco, CA.
Awarded to *Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin*. *"Why Should I Trust You?": Explaining the Predictions of Any Classifier*. In Knowledge Discovery and Data Mining (KDD), August 2016
- 2016 **Best Paper Award**, ICML Workshop on Human Interpretability in Machine Learning, New York, NY.
Awarded to *Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin*. *Model-Agnostic Interpretability of Machine Learning*. In ICML Workshop on Human Interpretability in Machine Learning (WHI), June 2016
- 2015 **Grand Prize Winner**, Yelp Dataset Challenge, Round 4.
Awarded to *Nitish Gupta and Sameer Singh*. *Collective Factorization for Relational Data: An Evaluation on the Yelp Datasets*. Technical report, Yelp Dataset Challenge, Round 4, 2015
- 2014 **Exceptional Submission Award**, ACL Workshop on Semantic Parsing, Baltimore, MD.
Awarded to *Tim Rocktaschel, Sameer Singh, Matko Bosnjak, and Sebastian Riedel*. *Low-dimensional Embeddings of Logic*. In ACL 2014 Workshop on Semantic Parsing (SP14), 2014
- 2004 **Best Undergraduate Project**, Department of Instrumentation & Control (NSIT).

Service & Teaching Awards

- 2019 **DTEI Dean's Honoree Award for Undergraduate Teaching**, University of California, Irvine, CA.
- 2019 **Outstanding Area Chair**, Empirical Methods in Natural Language Processing (EMNLP), Hongkong.
- 2018 **Dean's Award for Excellence in Undergraduate Teaching**, University of California, Irvine, CA.

Professional Activity

Tutorials

- Dec 2020 **Explaining Machine Learning Predictions: State-of-the-art, Challenges, and Opportunities**, *Neural Information Processing Systems (NeurIPS)*.
- Nov 2020 **Interpreting Predictions of NLP Models**, *Empirical Methods in Natural Language Processing (EMNLP)*.
- June 2019 **Adversarial Learning in NLP**, *North-American Assoc of Computational Linguistics (NAACL)*, Minnesota.
- Feb 2018 **Mining Knowledge Graphs from Text**, *Intl Conf on Web Search & Data Mining (WSDM)*, Los Angeles.
- Feb 2017 **Knowledge Graph Construction From Text**, *AAAI Conf on Advances in Artificial Intel*, San Francisco, CA.

Invited Talks (Academic Events)

- May 2020 **Evaluating and Testing NLP Capabilities**, *Google Inc.*, (virtual).
- Feb 2020 **Evaluating and Testing NLP Capabilities**, *Dagstuhl Seminar on Software Engineering for ML (SEML)*, Dagstuhl, Germany.
- Jan 2020 **Evaluating and Testing Question Answering Capabilities**, *AAAI Workshop on Reasoning for Complex QA (RCQA)*, New York, NY.
- Jan 2020 **Explaining and Debugging ML**, *AAAI Workshop on Artificial Intelligence Safety (SafeAI)*, New York, NY.
- Jan 2020 **Explaining and Debugging ML**, *Information Theory and Applications (ITA)*, San Diego, CA.
- Jan 2020 **Explaining and Debugging ML**, *PSB Workshop on AI Ethics and Values in Biomedicine*, Kona, Hawaii.
- Nov 2019 **Inducing Fake, and Real, Information from NLP Models**, *2nd EMNLP Workshop on Fact Extraction and Verification*, Hongkong.
- Sep 2019 **Explaining and Debugging ML**, *Los Alamos National Labs*, Los Alamos, NM.
- Sep 2019 **Adversarial Perturbations for Debugging NLP**, *Southern California NLP Symposium*, Los Angeles, CA.
- May 2019 **Adversarial Perturbations for Debugging NLP**, *ICLR Workshop on Debugging ML*, New Orleans, LA.
- July 2018 **On Knowledge Graph Embeddings, with Application to Generation**, *TTIC Summer Workshop on Collaborative & Knowledge-backed Language Generation*, Chicago, IL.
- July 2018 **Local, Model-Agnostic Explanations of Machine Learning Predictions**, *Contributed Talk at the Joint Statistical Meeting (JSM)*, Vancouver, Canada.
- July 2018 **Questioning Question Answering Answers**, *ACL 2018 Workshop on Machine Reading and Question Answering (MRQA)*, Melbourne, AUS.
- Dec 2017 **Multimodal KB Extraction and Completion**, *NeurIPS Workshop on Automated Knowledge Base Construction (AKBC)*, Long Beach, CA.
- Dec 2017 **That Doesn't Make Sense! A Case Study in Actively Annotating Model Explanations**, *NeurIPS Workshop on Learning with Limited Labeled Data (LLD)*, Long Beach, CA.
- Nov 2016 **"Why Should I Trust You?" Explaining the Predictions of Any Classifier**, *Workshop on Fairness, Accountability, and Transparency in ML (FATML)*, New York, NY.
- Oct 2016 **Interpreting machine learning predictions**, *UCI Symp on Recent Advances in Data Science*, Irvine, CA.
- April 2015 **Declarative ML With Wolfe**, *Dagstuhl Seminar on Probabilistic Programming*, Dagstuhl, Germany.
- Aug 2014 **Large-Scale Entity Resolution**, *COLING AHA! Workshop*, Dublin, Ireland.
- Nov 2013 **Universal Schema for TACKBP**, *NIST*, Gaithersburg, MD.
- Apr 2011 **Large-scale Cross-doc Coreference**, *Machine Reading Project Ph 3 Kickoff*, Seattle, WA.

Invited Talks (Universities)

- Nov 2019 **Explaining and Debugging ML**, *UCI IMBS Seminar*, Irvine, CA.
- Nov 2019 **Explaining and Debugging ML**, *Indian Institute of Technology*, Delhi, India.
- Sep 2019 **Explaining and Debugging ML**, *Pomona College CS Colloquium*, Pomona, CA.
- May 2019 **Adversarial Perturbations for Debugging NLP**, *Stanford NLP Seminar*, Stanford, CA.

- Oct 2018 **Explaining Predictions of Any ML Algorithm**, *UC Riverside Department Colloquium*, Riverside, CA.
- Aug 2018 **Injecting Prior Information and Multiple Modalities into KB Embeddings**, *ExoBrain meeting*, Seoul, Korea.
- Aug 2018 **Local, Model-Agnostic Explanations of Machine Learning Predictions**, *KAIST*, Daejeon, Korea.
- Apr 2018 **Local, Model-Agnostic Explanations for Machine Learning**, *UCSD AI Seminar*, San Diego, CA.
- Apr 2018 **Local, Model-Agnostic Explanations for Machine Learning**, *Caltech IST Symposium*, Pasadena, CA.
- Mar 2017 **Intuitive Interactions with Black-box Machine Learning**, *USC ISI NLP Seminar*, Los Angeles, CA.
- Nov 2016 **Intuitive Explanations & Interactions with Black-box ML**, *UCSD AI Seminar*, San Diego, CA.
- Oct 2016 **Towards Intuitive Explanations and Interactions with Black-box ML**, *UCI CS Seminar*, Irvine, CA.
- Jun 2016 **Interactive Training for Relation Extraction**, *University of Washington AI Seminar*, Seattle, WA.
- Jan 2016 **Interactive Training of Relation Embeddings**, *University of Maryland*, College Park, MD.
- Jan 2016 **Interactive Training of Relation Embeddings**, *Johns Hopkins University*, Baltimore, MD.
- Jan 2015 **Interactive Matrix Factorization**, *UW CSE Department Summit*, Seattle, WA.
- Feb 2013 **Large-Scale Entity Resolution**, *Univ of Pennsylvania*, Philadelphia, PA.
- Nov 2012 **Machines That Read**, *Computer Science Dept Colloquium*, Williams College, MA.

Invited Talks (Non-academic)

- Feb 2020 **Evaluating and Testing NLP Capabilities**, *NEC Research*, Heidelberg, Germany.
- Jan 2020 **Explaining and Debugging ML**, *RE-Work Deep Learning Day*, San Francisco, CA.
- Dec 2019 **Explaining and Debugging ML**, *PyData*, Los Angeles, CA.
- Nov 2019 **Explaining and Debugging ML**, *Adobe*, Noida, India.
- Jul 2019 **Discovering Natural Bugs Using Adversarial Data Perturbations**, *Meetup on Robust AI: Debugging NLP models*, Seattle, WA.
- Jan 2019 **Explaining Decisions of Black-box Machine Learning**, *Financial Conduct Authority*, UK, remote.
- Sep 2018 **Explaining Decisions of Black-box Machine Learning**, *Vale AI Centre of Excellence*, remote.
- Apr 2018 **Oh Shoot! What Now? Technology in a Time of Crisis**, *Filene Research*, <http://innovation.uci.edu/2018/05/filene-research-institute-engages-conversations-on-finance-technology/>.
- Apr 2018 **Explaining Black-Box Machine Learning Predictions**, *FICO World Conference*, Miami, FL.
- Dec 2017 **Explaining Black-Box Machine Learning Predictions**, *H2O.AI World Conference*, San Francisco, CA.
- Jul 2017 **Explaining Black-Box Machine Learning Predictions**, *Orange County ACM Meeting*, Irvine, CA.
- Jun 2017 **Explaining Black-Box Machine Learning Predictions**, *PyData SoCal Meetup*, Los Angeles, CA.
- Jun 2017 **Interpreting machine learning predictions**, *Nutanix .NEXT Conference*, Washington, DC.
- Mar 2017 **Explaining the Predictions of any Classifier**, *Global Data Science Conference*, Santa Clara, CA.
- Dec 2016 **Explaining the Predictions of Any Classifier**, *OC Advanced Analytics Meetup*, Newport Beach, CA.
- Aug 2015 **Wolfe: A Declarative Machine Learning Stack**, *Big Data Scala By the Bay*, San Francisco, CA.
- Dec 2014 **Wolfe: Declarative Machine Learning**, *Montreal Scala Meetup*, Montreal, Canada.

Media Appearances, Articles, and Quotes

- Sep 2018 **Forbes**, *Quoted in article*, <https://www.forbes.com/sites/jasonbloomberg/2018/09/16/dont-trust-artificial-intelligence-time-to-open-the-ai-black-box/>.
- Jul 2017 **Canada TV's Your Morning**, *Interviewed for live television*, <http://www.ctvnews.ca/sci-tech/explainable-ai-the-push-to-make-sure-machines-don-t-learn-to-be-racist-1.3488349>.
- Jul 2017 **The Wrap**, *Quoted in an article on Artificial Intelligence*, <http://www.thewrap.com/why-mark-zuckerberg-and-elon-musks-are-artificial-intelligence-adversaries-draft/>.
- Jun 2017 **Fast Company**, *Quoted in an article on Machine Learning*, <https://www.fastcompany.com/40433959/why-the-military-and-corporate-america-want-to-make-ai-explain-itself>.

Funding

Grants and Sponsored Research

- 2020-2023 **DARPA, RED**, Reverse Engineering of Deceptions, (recommended) \$300 000 (Co-PI with Lowd, UOregon).
- 2020-2022 **NSF, EAGER-SaTC**, Multi-level Attack and Defense Simulation Environment for AI Education and Research, \$300 000 (Co-PI, with PI Li and Co-PI Gago-Masague).
- 2020-2023 **NSF, RI Small**, Post hoc Explanations in the Wild: Exposing Vulnerabilities and Ensuring Robustness, \$450 000 (PI, collaboration with PI Lakkaraju, Harvard).
- 2019-2023 **DARPA, MCS**, Machine Common Sense, \$1 250 000.
- 2019-2022 **DARPA, LwLL**, Learning with Less Labels, \$250 000.
- 2019-2022 **NSF, CCRI**, ENS: ML Democratization via a Linked, Annotated Repository of Datasets, \$1 800 000 (PI, with Co-PIs Smyth and Papadopoulos).
- 2019 **Amazon, Alexa Challenge**, ZotBot, \$250 000.
- 2019 **UCI CORCL, Faculty Research Award**, Answering Complex Questions, \$3 000.
- 2018-2021 **NSF, RI Small**, Modeling Multiple Modalities for Knowledge-Base Construction, \$450 000.
- 2018-2020 **NSF, CRII**, Explaining Decisions of Black-box Models via Input Perturbations, \$175 000.

Gifts and Donations

- Jan '20 - Dec '20 **Allen Institute for Artificial Intelligence, Research Support**, Knowledge Extraction and Reasoning.
- Jun '19 - Jun '18 **Base 11 Foundation and Deloitte, Teaching Support**, Introduction to Data Science.
- Jan '19 - Dec '19 **Allen Institute for Artificial Intelligence, Research Support**, Knowledge Extraction and Reasoning.
- Jun '18 - Dec '18 **Allen Institute for Artificial Intelligence, Research Support**, Knowledge Extraction and Reasoning.
- Jun '17 - Jun '18 **Fair Isaac Corporation, Research Support**, Explaining Machine Learning Predictions.
- Jun '17 - Jun '18 **Adobe Research, Data Science Faculty Award**, Explanations for Machine Learning Predictions.

Other Support

- 2015 **Support Grant for Junior Researchers**, Schloss Dagstuhl-NSF, Dagstuhl, Germany.
- 2014, 2015 **UW CSE Postdoc Research Award**, University of Washington, Seattle, WA.
- 2014 **Amazon Machine Learning Data Grant**, Amazon, Seattle, WA.
- 2012 **Finalist**, Facebook PhD Fellowship, Menlo Park, CA.
- 2010-2011 **Graduate School Fellowship Award**, for continuing students, Univ of Massachusetts.

Professional Service

Conference and Workshop Organization

- member** TAC KBP Scientific Advisory Board, 2020
- chair** 2nd Conference on Automated Knowledge Base Construction (AKBC), 2020
- organizer** Deep Learning Day, Knowledge Discovery and Data Mining (KDD), 2020
- senior area chair** Empirical Methods in Natural Language Processing (EMNLP), 2020
- area chair** Neural Information Processing Systems (NeurIPS), 2020
- area chair** International Conference on Machine Learning, 2020
- organizer** NeurIPS Workshop on Knowledge Representation & Reasoning Meets Machine Learning (KR2ML), 2019
- organizer** 1st AKBC Workshop on Knowledge Bases and Multiple Modalities (KBMM), 2019
- area chair** Neural Information Processing Systems (NeurIPS), 2019
- area chair** Empirical Methods in Natural Language Processing (EMNLP), 2019
- program chair** 1st Conference on Automated Knowledge Base Construction (AKBC), 2019
- area chair** International Conference on Learning Representations (ICLR), 2019

- chair** Southern California NLP Symposium, 2018
- session chair** Annual Conference of the Association for Computational Linguistics (ACL), 2018
- area chair** Neural Information Processing Systems (NeurIPS), 2018
- area chair** Conference on Computational Natural Language Learning (CONLL), 2018
- area chair** Neural Information Processing Systems (NeurIPS), 2017
- area chair** Empirical Methods in Natural Language Processing (EMNLP), 2017
- publication chair** Annual Conference of the Association for Computational Linguistics (ACL), 2017
- organizer** NAACL Workshop on “Automated Knowledge Base Construction”, 2016
- organizer** AAAI Workshop on “Declarative Learning Based Programming”, 2016
- organizer** NeurIPS Workshop on “Machine Learning Systems (LearningSys)”, 2015
- organizer** NeurIPS Workshop on “Automated Knowledge Base Construction”, 2014
- organizer** CIKM Workshop on “Automated Knowledge Base Construction”, 2013
- organizer** NeurIPS Workshop on “Big Learning”, 2011, 2012, 2013
- organizer** ICML Workshop on “Inferning”: Interactions between Inference and Learning, 2012, 2013
- co-chair** North-East Student Colloquium on Artificial Intelligence (NESCAI), 2010

Reviewing/PC Member

- journal** Journal for ML Research (2019-), Journal for AI Research (2016-), Transactions of the ACL (2015-), Big Data (2016-), Computational Linguistics (2016-), ACM Transactions on Intelligent Systems and Technology (2016-), PLOS ONE (2017), Communications of the ACM (2018), Pattern Recognition (2018)
- conferences** EMNLP (2010, 2011, 2012, 2014, 2015), KDD (2011, 2015, 2018), IJCAI (2011, 2015, 2016), NeurIPS (2011, 2012, 2013, 2014, 2015, 2016), EACL (2012, 2014, 2017), ICML (2013, 2014, 2015, 2018), NAACL (2013, 2015, 2016, 2018), UAI (2013, 2018), ACL (2014, 2015, 2016, 2017, 2018), COLING (2014, 2018), CIKM 2014, CONLL 2013, AAAI (2015, 2017), WSDM (2015, 2016), AISTATS 2017, ECML-PKDD (2016, 2017, 2018), WWW 2017, IJCNLP 2017
- workshops** AKBC 2010, NESCAI 2010, OPTMAS 2012, AKBC 2012, OPTMAS 2013, LML 2013, AKBC 2013, StarAI 2014, LTPM 2014, NAMPI 2016, StarAI 2017, DeepStruct 2017, StarAI 2018, EthNLP 2018, CD-MAKE 2018, PyData SoCal 2018
- agencies** NSF CISE (IIS Division): Panelist 2016, The Royal Society Workshop on AI for Social Good: Panelist 2017, NSF CISE (CCF Division): Panelist 2018, NSF CISE (CHS Division): Panelist 2018, NSF CISE (IIS Division): Panelist 2019, NSF CHS (IIS Division): Panelist 2020, NSF CCRI: Ad-hoc reviewer 2019

External Committee Member

- 2019** **PhD Committee**, *Konstantinos Skianis*, Advised by Michalis Vazirgiannis, Ecole Polytechnique.
- 2018** **PhD Committee**, *Marco Tulio Ribeiro*, Advised by Carlos Guestrin, Univ of Washington.
- 2018** **PhD Committee**, *Andrew Chisholm*, Advised by Ben Hachey, University of Sydney.
- 2018** **PhD Committee**, *Linlin Wang*, Advised by Gerard de Melo, Tsinghua University.
- 2017** **PhD Proposal Committee**, *Marco Tulio Ribeiro*, Advised by Carlos Guestrin, Univ of Washington.
- 2015** **PhD Committee**, *Xiao Ling*, Advised by Daniel S. Weld, Univ of Washington.

University Service

Outreach Activities

- course creator** Introduction to Data Science for High-School, 2019-2020
- founding organizer** North American Computational Linguistics Olympiad (NACLO), UC Irvine, 2017-2020
- mentor** High School and Middle School teams, Paul Allen Computing Challenge, 2014
- judge** Science Fair, Hampden Charter School of Science, 2012
- session chair** Massachusetts Statewide Undergraduate Research Conference, 2012

Mentoring

Current PhD Students

- Aug 2019- **Dylan Slack**, *PhD*, Computer Science, UC Irvine.
- Aug 2018- **Anthony Chen**, *PhD*, Computer Science, UC Irvine.
- Sep 2017- **Robert L. Lockwood**, *PhD*, co-advisor: Padhraic Smyth, UC Irvine.
- Aug 2017- **Dheeru Dua**, *PhD*, Computer Science, UC Irvine.
- Apr 2017- **Pouya Pezeshkpour**, *PhD*, EECS, UC Irvine.
- Dec 2016- **Zhengli Zhao**, *PhD*, Computer Science, UC Irvine.

Graduated PhD Students

- 2013-2018 **Marco T. Ribeiro**, *PhD*, co-advisor: Carlos Guestrin, University of Washington.
- 2017-2018 **Forough Arabshahi**, *PhD*, co-advisor: Anima Anandkumar, UC Irvine.

Graduated MS Students

- 2019-2020 **Taylor Shin**, *MS*, Computer Science, UC Irvine.
- 2018-2019 **Chirag Choudhary**, *MS*, Computer Science, UC Irvine.
- 2017-2018 **Ananya**, *MS*, Computer Science, UC Irvine.

Publications

Book Chapters

- BC2 Forest Agostinelli, Guillaume Hocquet, **Sameer Singh**, and Pierre Baldi. From Reinforcement Learning to Deep Reinforcement Learning: An Overview. In *Braverman Readings in Machine Learning: Key Ideas from Inception to Current State*, Springer Press. 2018
- BC1 Jeremy Kubica, **Sameer Singh**, and Daria Sorokina. Parallel Large-scale Feature Selection. In *Scaling Up Machine Learning*, Cambridge University Press. 2011

Journal Publications

- J5 Jihyun Park, Dimitrios Kotzias, Patty Kuo, Robert L. Logan, Kritzia Merced, **Sameer Singh**, Michael Tanana, Efi Karra-Taniskidou, Jennifer Elston Lafata, David C. Atkins, Ming Tai-Seale, Zac E Imel, and Padhraic Smyth. Detecting Conversation Topics in Primary Care Office Visits from Transcripts of Patient-Provider Interactions. *Journal of the American Medical Informatics Association*, TBD, 2019
- J4 Yanbing Bai, Chang Gao, **Sameer Singh**, Magaly Koch, Bruno Adriano, Erick Mas, and Shunichi Koshimura. A Framework of Rapid Regional Tsunami Damage Recognition from Post-event TerraSAR-X Imagery Using Deep Neural Networks. *IEEE Geoscience and Remote Sensing Letters*, PP, 2018
- J3 Xiao Ling, **Sameer Singh**, and Dan Weld. Design Challenges for Entity Linking. *Transactions of the Association for Computational Linguistics (ACL)*, 3, 2015
- J2 D.G. Feitelson, T.O.S. Adeshiyani, D. Balasubramanian, Y. Etsion, G. Madl, E.P. Osses, **Sameer Singh**, K. Suwanmongkol, M. Xie, and S.R. Schach. Fine-Grain Analysis of Common Coupling and its Application to a Linux Case Study. *Journal of Systems and Software (JSS)*, 80, 2007
- J1 S.R. Schach, T.O.S. Adeshiyani, D. Balasubramanian, G. Madl, E.P. Osses, **Sameer Singh**, K. Suwanmongkol, M. Xie, and D.G. Feitelson. Common Coupling and Pointer Variables, with Application to a Linux Case Study. *Software Quality Journal (SQJ)*, 15, 2007

Selective, Peer-Reviewed Conference Papers

- C49 Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and **Sameer Singh**. Beyond Accuracy: Behavioral Testing of NLP models with CheckList. In *Association for Computational Linguistics (ACL)*, 2020 **Overall Best Paper Award**
- C48 Robert L. Logan, Matt Gardner, and **Sameer Singh**. On Importance Sampling-Based Evaluation of Latent Language Models. In *Association for Computational Linguistics (ACL)*, 2020

- C47 Sanjay Subramanian, Ben Bogin, Nitish Gupta, Tomer Wolfson, **Sameer Singh**, Jonathan Berant, and Matt Gardner. Obtaining Faithful Interpretations from Compositional Neural Networks. In *Association for Computational Linguistics (ACL)*, 2020
- C46 Dheeru Dua, **Sameer Singh**, and Matt Gardner. Benefits of Intermediate Annotations in Reading Comprehension. In *Association for Computational Linguistics (ACL)*, 2020
- C45 Ananth Gottumukkala, Dheeru Dua, **Sameer Singh**, and Matt Gardner. Dynamic Sampling Strategies for Multi-Task Reading Comprehension. In *Association for Computational Linguistics (ACL)*, 2020
- C44 Pouya Pezeshkpour, Yifan Tian, and **Sameer Singh**. Revisiting Evaluation of Knowledge Base Completion Models. In *Automated Knowledge Base Construction (AKBC)*, 2020 *Best Paper Runners-Up*
- C43 Dan Barsever, **Sameer Singh**, and Emre Neftci. Building a Better Lie Detector with BERT: The Difference Between Truth and Lies. In *International Joint Conference on Neural Networks (IJCNN)*, 2020
- C42 **Sameer Singh**. Minecraft as a Platform for Project-Based Learning in AI. In *AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI)*, 2020
- C41 Dylan Slack, Sophie Hilgard, Emily Jia, **Sameer Singh**, and Himabindu Lakkaraju. Fooling LIME and SHAP: Adversarial Attacks on Post hoc Explanation Methods. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2020
- C40 Nitish Gupta, Kevin Lin, Dan Roth, **Sameer Singh**, and Matt Gardner. Neural Module Networks for Reasoning over Text. In *International Conference on Learning Representations (ICLR)*, 2020
- C39 Piyush Gupta, Nikaash Puri, Sukriti Verma, Dhruv Kayastha, Shripad Deshmukh, Balaji Krishnamurthy, and **Sameer Singh**. Explain Your Move: Understanding Agent Actions Using Specific and Relevant Feature Attribution. In *International Conference on Learning Representations (ICLR)*, 2020
- C38 Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and **Sameer Singh**. Universal Adversarial Triggers for Attacking and Analyzing NLP. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2019
- C37 Eric Wallace, Yizhong Wang, Sujian Li, **Sameer Singh**, and Matt Gardner. Do NLP Models Know Numbers? Probing Numeracy in Embeddings. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2019
- C36 Matthew E. Peters, Mark Neumann, Robert L. Logan, Roy Schwartz, Vidur Joshi, **Sameer Singh**, and Noah A. Smith. Knowledge Enhanced Contextual Word Representations. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2019
- C35 Robert L. Logan, Nelson F. Liu, Matthew E. Peters, Matt Gardner, and **Sameer Singh**. Barack's Wife Hillary: Using Knowledge Graphs for Fact-Aware Language Modeling. In *Association for Computational Linguistics (ACL)*, 2019
- C34 Marco Tulio Ribeiro, Carlos Guestrin, and **Sameer Singh**. Are Red Roses Red? Evaluating Consistency of Question-Answering Models. In *Association for Computational Linguistics (ACL)*, 2019
- C33 Sewon Min, Eric Wallace, **Sameer Singh**, Matt Gardner, Hannaneh Hajishirzi, and Luke Zettlemoyer. Compositional Questions Do Not Necessitate Multi-hop Reasoning. In *Association for Computational Linguistics (ACL)*, 2019
- C32 Pouya Pezeshkpour, Yifan Tian, and **Sameer Singh**. Investigating Robustness and Interpretability of Link Prediction via Adversarial Modifications. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019
- C31 Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, **Sameer Singh**, and Matt Gardner. DROP: A Reading Comprehension Benchmark Requiring Discrete Reasoning Over Paragraphs. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019
- C30 Ananya Ananya, Nitya Parthasarathi, and **Sameer Singh**. GenderQuant: Quantifying Mention-Level Genderedness. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019
- C29 Jun Seok Kang, Robert L. Logan, Zewei Chu, Yang Chen, Dheeru Dua, Kevin Gimpel, **Sameer Singh**, and Niranjan Balasubramanian. PoMo: Generating Entity-Specific Post-Modifiers in Context. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019
- C28 Pouya Pezeshkpour, Liyan Chen, and **Sameer Singh**. Embedding Multimodal Relational Data for Knowledge Base Completion. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2018

- C27 Marzieh Saeidi, Max Bartolo, Patrick Lewis, **Sameer Singh**, Tim Rocktaschel, Mike Sheldon, Guillaume Bouchard, and Sebastian Riedel. Interpretation of Natural Language Rules in Conversational Machine Reading. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2018
- C26 Marco Tulio Ribeiro, **Sameer Singh**, and Carlos Guestrin. Semantically Equivalent Adversarial Rules for Debugging NLP models. In *Association for Computational Linguistics (ACL)*, 2018 *Honorable Mention for Best Paper Award*
- C25 Zhengli Zhao, Dheeru Dua, and **Sameer Singh**. Generating Natural Adversarial Examples. In *International Conference on Learning Representations (ICLR)*, 2018 *Best Poster Award at SoCal ML Symposium*
- C24 Forough Arabshahi, **Sameer Singh**, and Animashree Anandkumar. Combining Symbolic Expressions and Black-box Function Evaluations for Training Neural Programs. In *International Conference on Learning Representations (ICLR)*, 2018
- C23 Marco Tulio Ribeiro, **Sameer Singh**, and Carlos Guestrin. Anchors: High-Precision Model-Agnostic Explanations. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018
- C22 Nitish Gupta, **Sameer Singh**, and Dan Roth. Entity Linking via Joint Encoding of Types, Descriptions, and Context. In *Empirical Methods in Natural Language Processing (EMNLP)*, September 2017
- C21 Parisa Kordjamshidi, Daniel Khashabi, Christos Christodoulopoulos, Bhargav Mangipudi, **Sameer Singh**, and Dan Roth. Better call Saul: Flexible Programming for Learning and Inference in NLP. In *International Conference on Computational Linguistics (COLING)*, December 2016
- C20 Marco Tulio Ribeiro, **Sameer Singh**, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Knowledge Discovery and Data Mining (KDD)*, August 2016 *Audience Appreciation Award*
- C19 Hannah Rashkin, **Sameer Singh**, and Yejin Choi. Connotation Frames: A Data-Driven Investigation. In *Association for Computational Linguistics (ACL)*, August 2016
- C18 Tim Rocktaschel, **Sameer Singh**, and Sebastian Riedel. Injecting Logical Background Knowledge into Embeddings for Relation Extraction. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2015
- C17 Tianqi Chen, **Sameer Singh**, Ben Taskar, and Carlos Guestrin. Efficient Second-Order Gradient Boosting for Conditional Random Fields. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015
- C16 Ignacio Cano, **Sameer Singh**, and Carlos Guestrin. Distributed Non-Parametric Representations for Vital Filtering: UW at TREC KBA 2014. In *Text REtrieval Conference (TREC): Knowledge-Base Acceleration (KBA) Track*, 2014
- C15 Fabian M. Suchanek, **Sameer Singh**, Sebastian Riedel, and Partha P. Talukdar. AKBC 2013: Third Workshop on Automated Knowledge Base Construction. In *ACM Conference of Information and Knowledge Management (CIKM)*, 2013
- C14 **Sameer Singh** and Thore Graepel. Automated Probabilistic Modeling for Relational Data. In *ACM Conference of Information and Knowledge Management (CIKM)*, 2013
- C13 Jiaping Zheng, Luke Vilnis, **Sameer Singh**, Jinho Choi, and Andrew McCallum. Dynamic Knowledge-Base Alignment for Coreference Resolution. In *Conference on Computational Natural Language Learning (CoNLL)*, 2013
- C12 **Sameer Singh**, Limin Yao, David Belanger, Ari Kobren, Sam Anzaroot, Michael Wick, Alexandre Passos, Harshal Pandya, Jinho Choi, Brian Martin, and Andrew McCallum. Universal Schema for Slot Filling and Cold Start: UMass IESL at TACKBP 2013. In *Text Analysis Conference on Knowledge Base Population (TAC KBP)*, 2013
- C11 Michael Wick, **Sameer Singh**, and Andrew McCallum. A Discriminative Hierarchical Model for Fast Coreference at Large Scale. In *Association for Computational Linguistics (ACL)*, 2012
- C10 **Sameer Singh**, Michael Wick, and Andrew McCallum. Monte Carlo MCMC: Efficient Inference by Approximate Sampling. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2012
- C9 **Sameer Singh**, Amarnag Subramanya, Fernando Pereira, and Andrew McCallum. Large-Scale Cross-Document Coreference Using Distributed Inference and Hierarchical Models. In *Association for Computational Linguistics (ACL)*, 2011
- C8 **Sameer Singh**, Dustin Hillard, and Chris Leggetter. Minimally-Supervised Extraction of Entities from Text Advertisements. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2010

- C7 **Sameer Singh**, Limin Yao, Sebastian Riedel, and Andrew McCallum. Constraint-Driven Rank-Based Learning for Information Extraction. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2010
- C6 **Sameer Singh**, Karl Schultz, and Andrew McCallum. Bi-directional Joint Inference for Entity Resolution and Segmentation using Imperatively-Defined Factor Graphs. In *Machine Learning and Knowledge Discovery in Databases (Lecture Notes in Computer Science) and European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, 2009
- C5 **Sameer Singh**, Jeremy Kubica, Scott E. Larsen, and Daria Sorokina. Parallel Large Scale Feature Selection for Logistic Regression. In *SIAM International Conference on Data Mining (SDM)*, 2009
- C4 Andrew McCallum, Karl Schultz, and **Sameer Singh**. FACTORIE: Probabilistic Programming via Imperatively Defined Factor Graphs. In *Neural Information Processing Systems (NeurIPS)*, 2009
- C3 Michael Wick, Khashyar Rohanimanesh, **Sameer Singh**, and Andrew McCallum. Training Factor Graphs with Reinforcement Learning for Efficient MAP Inference. In *Neural Information Processing Systems (NeurIPS)*, 2009
- C2 **Sameer Singh** and Julie A. Adams. Transfer of Learning for Complex Domains: A Demonstration Using Multiple Robots. In *International Conference on Robotics and Automation (ICRA)*, 2006
- C1 **Sameer Singh**. Finding the shortest path for a mobile robot in an unmapped maze from minimum runs. In *Int Conf on CAD, CAM, Robotics and Autonomous Factories (INCARF)*, 2003

Peer-reviewed Workshops, Demonstrations, and Symposia

- W40 Pouya Pezeshkpour, Zhengli Zhao, and **Sameer Singh**. Data Importance-Based Active Learning for Limited Labels. In *CVPR Workshop on Visual Learning with Limited Labels (VL3)*, 2020
- W39 Sanjay Subramanian, **Sameer Singh**, and Matt Gardner. Analyzing Compositionality of Visual Question Answering. In *NeurIPS Workshop on Visually Grounded Interaction and Language (ViGIL)*, 2019
- W38 Zhengli Zhao, Nicolas Papernot, **Sameer Singh**, Neoklis Polyzotis, and Augustus Odena. Improving Differentially Private Models with Active Learning. In *NeurIPS Workshop on Privacy in Machine Learning (PriML)*, 2019
- W37 Yoshitomo Matsubara, Sabur Baidya, Davide Callegaro, Marco Levorato, and **Sameer Singh**. Distilled Split Deep Neural Networks for Edge-Assisted Real-Time Systems. In *Mobicom Workshop on Hot Topics in Video Analytics and Intelligent Edges*, 2019
- W36 Anthony Chen, Gabriel Stanovsky, **Sameer Singh**, and Matt Gardner. Evaluating Question Answering Evaluation. In *Workshop on Machine Reading and Question Answering (MRQA)*, 2019 *Best Paper Award*
- W35 Dheeru Dua, Ananth Gottumukkala, Alon Talmor, **Sameer Singh**, and Matt Gardner. ORB: An Open Reading Benchmark for Comprehensive Evaluation of Machine Reading Comprehension. In *Workshop on Machine Reading and Question Answering (MRQA)*, 2019
- W34 Eric Wallace, Jens Tuyls, Junlin Wang, Sanjay Subramanian, Matt Gardner, and **Sameer Singh**. AllenNLP Interpret: A Framework for Explaining Predictions of NLP Models. In *Demo at the Empirical Methods in Natural Language Processing (EMNLP)*, 2019 *Best Demo Paper Award*
- W33 Forough Arabshahi, **Sameer Singh**, and Animashree Anandkumar. Towards Solving Differential Equations through Neural Programming. In *ICML Workshop on Neural Abstract Machines and Program Induction (NAMPI)*, 2018
- W32 Zhengli Zhao, Dheeru Dua, and **Sameer Singh**. Generating Natural Adversarial Examples. In *NeurIPS Workshop on Machine Deception*, 2017
- W31 Pouya Pezeshkpour, Liyan Chen, and **Sameer Singh**. Embedding Multimodal Relational Data. In *Workshop on Automated Knowledge Base Construction (AKBC)*, 2017
- W30 Robert L. Logan, Samuel Humeau, and **Sameer Singh**. Multimodal Attribute Extraction. In *Workshop on Automated Knowledge Base Construction (AKBC)*, 2017
- W29 Ananya Ananya and **Sameer Singh**. How Biased Are We? Automated Detection of Gendered Language. In *ACL Workshop on Women and Underrepresented Minorities in NLP (WiNLP)*, August 2017
- W28 Parisa Kordjamshidi, **Sameer Singh**, Daniel Khashabi, Christos Christodoulopoulos, Mark Summons, Saurabh Sinha, and Dan Roth. Relational Learning and Feature Extraction by Querying over Heterogeneous Information Networks. In *International Workshop on Statistical Relational AI (StarAI)*, July 2017

- W27 **Sameer Singh**, Marco Tulio Ribeiro, and Carlos Guestrin. Programs as Black-Box Explanations. In *NeurIPS Workshop on Interpretable Machine Learning in Complex Systems*, November 2016
- W26 Marco Tulio Ribeiro, **Sameer Singh**, and Carlos Guestrin. Nothing Else Matters: Model-Agnostic Explanations By Identifying Prediction Invariance. In *NeurIPS Workshop on Interpretable Machine Learning in Complex Systems*, November 2016
- W25 Marco Tulio Ribeiro, **Sameer Singh**, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Demo at the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, June 2016
- W24 Marco Tulio Ribeiro, **Sameer Singh**, and Carlos Guestrin. Model-Agnostic Interpretability of Machine Learning. In *ICML Workshop on Human Interpretability in Machine Learning (WHI)*, June 2016 **Best Paper Award**
- W23 Marco Tulio Ribeiro, **Sameer Singh**, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *CHI Workshop on Human-Centred Machine Learning (HCML)*, May 2016
- W22 **Sameer Singh** and Sebastian Riedel. Creating Interactive and Visual Educational Resources for AI. In *AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI)*, 2016
- W21 **Sameer Singh**, Tim Rocktaschel, Luke Hewitt, Jason Naradowsky, and Sebastian Riedel. WOLFE: An NLP-friendly Declarative Machine Learning Stack. In *Demo at the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2015
- W20 **Sameer Singh**, Tim Rocktaschel, and Sebastian Riedel. Towards Combined Matrix and Tensor Factorization for Universal Schema Relation Extraction. In *NAACL Workshop on Vector Space Modeling for NLP*, 2015
- W19 Ivan Sanchez, Tim Rocktaschel, Sebastian Riedel, and **Sameer Singh**. Towards Extracting Faithful and Descriptive Representations of Latent Variable Models. In *AAAI Spring Symposium on Knowledge Representation and Reasoning (KRR): Integrating Symbolic and Neural Approaches*, 2015
- W18 Guillaume Bouchard, **Sameer Singh**, and Theo Trouillon. On Approximate Reasoning Capabilities of Low-Rank Vector Spaces. In *AAAI Spring Symposium on Knowledge Representation and Reasoning (KRR): Integrating Symbolic and Neural Approaches*, 2015
- W17 Tim Rocktaschel, **Sameer Singh**, Matko Bosnjak, and Sebastian Riedel. Low-dimensional Embeddings of Logic. In *ACL 2014 Workshop on Semantic Parsing (SP14)*, 2014 **Exceptional Submission Award**
- W16 Sebastian Riedel, **Sameer Singh**, Vivek Srikumar, Tim Rocktaschel, Larysa Visengeriyeva, and Jan Noessner. WOLFE: Strength Reduction and Approximate Programming for Probabilistic Programming. In *International Workshop on Statistical Relational AI (StarAI)*, 2014
- W15 Xiao Ling, **Sameer Singh**, and Dan Weld. Context Representation for Named Entity Linking. In *Pacific Northwest Regional NLP Workshop (NW-NLP)*, 2014
- W14 Mathias Niepert and **Sameer Singh**. Out of Many, One: Unifying Web-Extracted Knowledge Bases. In *Workshop on Automated Knowledge Base Construction (AKBC)*, 2014
- W13 **Sameer Singh**, Sebastian Riedel, Luke Hewitt, and Tim Rocktaschel. Designing an IDE for Probabilistic Programming: Challenges and a Prototype. In *NeurIPS Workshop on Probabilistic Programming*, 2014
- W12 Victoria (Xi) Lin, **Sameer Singh**, Luheng He, Ben Taskar, and Luke Zettlemoyer. Multi-label Learning with Posterior Regularization. In *NeurIPS Workshop on Modern Machine Learning and Natural Language Processing*, 2014
- W11 **Sameer Singh**, Sebastian Riedel, Brian Martin, Jiaping Zheng, and Andrew McCallum. Joint Inference of Entities, Relations, and Coreference. In *CIKM Workshop on Automated Knowledge Base Construction (AKBC)*, 2013
- W10 Michael Wick, **Sameer Singh**, Ari Kobren, and Andrew McCallum. Assessing Confidence of Knowledge Base Content with an Experimental Study in Entity Resolution. In *CIKM Workshop on Automated Knowledge Base Construction (AKBC)*, 2013
- W9 Michael Wick, **Sameer Singh**, Harshal Pandya, and Andrew McCallum. A Joint Model for Discovering and Linking Entities. In *CIKM Workshop on Automated Knowledge Base Construction (AKBC)*, 2013
- W8 **Sameer Singh**, Sebastian Riedel, and Andrew McCallum. Anytime Belief Propagation Using Sparse Domains. In *Neural Information Processing Systems (NeurIPS) Workshop on Resource Efficient Machine Learning*, 2013

- W7 **Sameer Singh**, Michael Wick, and Andrew McCallum. Monte Carlo MCMC: Efficient Inference by Sampling Factors. In *NAACL/HLT Workshop on Automated Knowledge Base Construction (AKBC-WEKEX)*, 2012
- W6 **Sameer Singh** and Thore Graepel. Compiling Relational Database Schemata into Probabilistic Graphical Models. In *NeurIPS Workshop on Probabilistic Programming*, 2012
- W5 **Sameer Singh** and Andrew McCallum. Towards Asynchronous Distributed MCMC Inference for Large Graphical Models. In *Neural Information Processing Systems (NeurIPS) Workshop on Big Learning*, 2011
- W4 **Sameer Singh**, Brian Martin, and Andrew McCallum. Inducing Value Sparsity for Parallel Inference in Tree-shaped Models. In *Neural Information Processing Systems (NeurIPS) Workshop on Computational Trade-offs in Statistical Learning (COST)*, 2011
- W3 **Sameer Singh**, Amarnag Subramanya, Fernando Pereira, and Andrew McCallum. Distributed MAP Inference for Undirected Graphical Models. In *Neural Information Processing Systems (NeurIPS) Workshop on Learning on Cores, Clusters, and Clouds (LCCC)*, 2010
- W2 Andrew McCallum, Khashyar Rohanimanesh, Michael Wick, Karl Schultz, and **Sameer Singh**. FACTORIE: Efficient Probabilistic Programming via Imperative Declarations of Structure, Inference and Learning. In *NeurIPS Workshop on Probabilistic Programming*, 2008
- W1 T. Kichkaylo, C. van Buskirk, **Sameer Singh**, H. Neema, M. Orosz, and R. Neches. Mixed-Initiative Planning for Space Exploration Missions. In *International Conference on Automated Planning and Scheduling Workshop (ICAPS)*, 2007

Online Articles

- O2 **Sameer Singh**. Comment on Semantic Based Adversarial Examples Fool Face Recognition, August 2019. URL <https://syncedreview.com/2019/08/09/semantic-based-adversarial-examples-fool-face-recognition/>
- O1 Marco Tulio Ribeiro, **Sameer Singh**, and Carlos Guestrin. Introduction to Local Interpretable Model-Agnostic Explanations (LIME), August 2016. URL <https://www.oreilly.com/learning/introduction-to-local-interpretable-model-agnostic-explanations-lime>

Invited Papers, Unrefereed Publications

- U8 Igor Burago, Davide Callegaro, Marco Levorato, and **Sameer Singh**. Intelligent Data Filtering in Constrained IoT Systems. In *Asilomar Conference on Signals, Systems, and Computers*, 2017
- U7 Igor Burago, Marco Levorato, and **Sameer Singh**. Semantic Compression for Edge-Assisted Systems. In *Information Theory and Applications (ITA) Workshop*, February 2017
- U6 Nitish Gupta and **Sameer Singh**. Collective Factorization for Relational Data: An Evaluation on the Yelp Datasets. Technical report, Yelp Dataset Challenge, Round 4, 2015 *Grand Prize Winner*
- U5 **Sameer Singh**, Amarnag Subramanya, Fernando Pereira, and Andrew McCallum. Wikilinks: A Large-scale Cross-Document Coreference Corpus Labeled via Links to Wikipedia. Technical report, University of Massachusetts Amherst, CMPSCI UM-CS-2012-015, 2012
- U4 **Sameer Singh**, Gregory Druck, and Andrew McCallum. Constraint-Driven Training of Complex Models Using MCMC. Technical report, University of Massachusetts Amherst, CMPSCI UM-CS-2012-032, 2012
- U3 **Sameer Singh**, Michael Wick, and Andrew McCallum. Distantly Labeling Data for Large Scale Cross-Document Coreference. Technical report, Computing Research Repository (CoRR) eprint arXiv:1005.4298, 2010
- U2 **Sameer Singh**. Option Discovery in Hierarchical Reinforcement Learning for Training Large Factor Graphs for Information Extraction, 2009
- U1 Khashyar Rohanimanesh, Michael Wick, **Sameer Singh**, and Andrew McCallum. Reinforcement Learning for MAP Inference in Large Factor Graphs. Technical report, University of Massachusetts Amherst, CMPSCI UM-CS-2008-040, 2008

Patents

- P2 **Sameer Singh**, Thore Graepel, Lucas J. Bordeaux, and Andrew D. Gordon. Relational database management. Technical report, US Patent Number 0188928, 2014
- P1 **Sameer Singh**, E. S. Larsen, Jeremy Kubica, and Andrew W. Moore. Feature selection for large scale models. Technical report, US Patent Number 8190537, 2008