

HAL DAUMÉ III — CURRICULUM VITAE

(Last Updated: 13 October 2019)

Contact Information

mailing address:

University of Maryland
Computer Science Department
Iribe Center #4150
College Park, MD 20742 USA

email: me@hal3.name

phone: (301) 405-1073

web: <http://hal3.name/>

pronouns: he/him

land acknowledgment: Lenape+Nacotchtank/Piscataway

Employment History

Professor

July 2018 – Present

Pier Giorgio Perotto Endowed Professor of Computer Science, 2019–2024
Computer Science Department, University of Maryland – College Park
and Language Science Center *and* Institute for Advanced Computer Studies

Senior Principal Researcher

June 2017 – Present

Microsoft Research, New York City

Associate Professor

July 2013 – June 2018

Computer Science Department, University of Maryland – College Park
and Language Science Center *and* Institute for Advanced Computer Studies

Assistant Professor

July 2010 – June 2013

Computer Science Department, University of Maryland – College Park
and University of Maryland Institute for Advanced Computer Studies *and* Department of Linguistics

Assistant Professor

July 2006 – June 2010

School of Computing, University of Utah

Research Intern

Summer 2003

Machine Learning and Applied Statistics Group, Microsoft Research

Education

University of Southern California

August 2006

Ph.D., Computer Science (advisor: Daniel Marcu)

Honors: Outstanding Research Assistant Award, Dean's Doctoral Merit Fellowship

University of Southern California

May 2003

M.S., Computer Science

Carnegie Mellon University

August 1998 – May 2001

B.S., Mathematical Sciences (advisor: Rami Grossberg)

Honors: Research Honors

Awards

- Best paper award ACL 2018: *Learning to Ask Good Questions: Ranking Clarification Questions using Neural Expected Value of Perfect Information* (with Sudha Rao).
- Best paper award NAACL 2016: *Feuding Families and Former Friends: Unsupervised Learning for Dynamic Fictional Relationships* (with Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi and Jordan Boyd-Graber).
- Amazon Research Award 2016.
- Facebook Research Award 2016.
- Best demonstration award NeurIPS 2015: *Interactive Incremental Question Answering* (with Jordan Boyd-Graber, Mohit Iyyer and He He).
- Google Research Award 2014.
- Best paper award CEAS 2011: *Using Classifier Cascades for Scalable E-Mail Classification* (with Jay Pujara and Lise Getoor).
- Google Research Award 2009 (with Suresh Venkatasubramanian).
- Best paper award ECML 2009: *A geometric view of conjugate priors* (with Arvind Agarwal).

Current Advisees

Ph.D. Students:

- Khanh Nguyen (expected graduation ~ 2020)
- Amr Sharaf (expected graduation ~ 2020)
- Kianté Brantley (expected graduation ~ 2022)
- Trista Cao (expected graduation ~ 2023)

Past Advisees

Postdoctoral Researchers:

- Hadi Amiri (Ph.D. National University of Singapore; Postdoc 2013 – 2015; co-advised with Philip Resnik, now Assistant Professor at UMass Lowell)
- Dan Goldwasser (Ph.D. UIUC; Postdoc 2012 – 2014; now Assistant Professor at Purdue University)
- Taesun Moon (Ph.D. UT Austin; Postdoc 2011 – 2013; now Research Scientist at IBM Research)

Ph.D. Students:

- Piyush Rai (Ph.D. 2012; now Assistant Professor at IIT Kanpur)
- Arvind Agarwal (Ph.D. 2012; now Research Scientist at Xerox Labs)
- Amit Goyal (Ph.D. 2013; now Research Scientist at Yahoo!)
- Abhishek Kumar (Ph.D. 2013; now Research Scientist at IBM Research)
- Jagadeesh Jagaralamudi (Ph.D. 2013; now Research Scientist at Google)
- Jiarong Jiang (Ph.D. ~ 2013; now Scientist at TwoSigma)
- He He (Ph.D. 2016; co-advised with Jordan Boyd-Graber, now Assistant Professor at NYU)
won the Larry S. Davis Doctoral Dissertation award 2016
won the CMNS Board of Visitors Outstanding Graduate Student Award in 2016
- Snigdha Chaturvedi (Ph.D. 2016, now Assistant Professor at UNC Chapel Hill)
- Mohit Iyyer (Ph.D. 2017, co-advised with Jordan Boyd-Graber, now Assistant Professor at UMass Amherst)
- Sudha Rao (Ph.D. 2019, now Researcher at Microsoft Research)

Undergraduate and M.S. Students:

- Seth Juarez (M.S. 2009; now at Microsoft Channel 9)
- Scott Alfeld (B.S. 2010; now Assistant Professor at Amherst College)
- Adam Teichert (M.S. 2010; now Assistant Professor at Snow College)
- J. Ozzie Fallick (CS honors student 2011 – 2014)
- Anna Gale (CS honors student 2012 – 2013)
- Daniel Rabinovich (CS honors student 2011 – 2014)
- Cynthia Gan (CS honors student 2014 – 2016)
- Meir Friedenberg (CS honors student 2015 – 2017)
- Harita Kannan (CS honors student 2016 – 2017)
- Ephraim Rothschild (CS honors student 2016 – 2017)

Publications

The fields I publish in—natural language processing and machine learning, for the most part—are conference-driven. I generally publish only in top conferences in these fields (ACL, NAACL, EACL, EMNLP for NLP; and ICML, NeurIPS for ML). Conference papers are generally more competitive than journal papers, but with associated randomness. All these conferences have roughly a 25% acceptance rate per year (it varies somewhat, but those variations are not meaningful, so I have not listed specific acceptance rates). I also do not believe that paper counts or citation counts provide meaningful, unbiased information about the impact of work, so I have not numbered publications nor have I included any approximations of citation counts.

Ph.D. Thesis

- ◇ **Hal Daumé III**. *Practical Structured Learning Techniques for Natural Language Processing*. 2006. Ph.D. Thesis.

Journal Articles

- ◇ Arti Ramesh, Dan Goldwasser, Bert Huang, **Hal Daumé III** and Lise Getoor. *Interpretable Engagement Models for MOOCs using Hinge-loss Markov Random Fields*. 2019. IEEE Transactions on Learning Technologies.
- ◇ Amit Goyal, Ellen Riloff and **Hal Daumé III**. *A Computational Model for Plot Units*. 2013. Computational Intelligence Journal.
- ◇ Ann Irvine, John Morgan, Marine Carpuat, **Hal Daumé III** and Dragos Munteanu. *Measuring Machine Translation Errors in New Domains*. 2013. Transactions of the Association for Computational Linguistics (TACL).
- ◇ Arvind Agarwal and **Hal Daumé III**. *A geometric view of conjugate priors*. 2010. Machine Learning Journal (MLJ).
- ◇ Pu Liu, Qiang Shi, **Hal Daumé III** and Gregory Voth. *A Bayesian Statistics Approach to Multiscale Coarse Graining*. 2009. Journal of Chemical Physics (J.ChPhys).
- ◇ **Hal Daumé III**, John Langford and Daniel Marcu. *Search-based Structured Prediction*. 2009. Machine Learning Journal (MLJ).
- ◇ **Hal Daumé III** and Daniel Marcu. *Domain Adaptation for Statistical Classifiers*. 2006. Journal of Artificial Intelligence Research (JAIR).
- ◇ **Hal Daumé III** and Daniel Marcu. *Induction of Word and Phrase Alignments for Automatic Document Summarization*. 2005. Computational Linguistics (CL).
- ◇ **Hal Daumé III** and Daniel Marcu. *A Bayesian Model for Supervised Clustering with the Dirichlet Process Prior*. 2005. Journal of Machine Learning Research (JMLR).

Conference Papers

- ◇ Khanh Nguyen and **Hal Daumé III**. *Help, Anna! Visual Navigation with Natural Multimodal Assistance via Retrospective Curiosity-Encouraging Imitation Learning*. 2019. Conference on Empirical Methods in Natural Language Processing (EMNLP).

- ◇ Chicheng Zhang, Alekh Agarwal, **Hal Daumé III**, John Langford and Sahand N Negahban. *Warm-starting Contextual Bandits: Robustly Combining Supervised and Bandit Feedback*. 2019. ICML.
- ◇ Sean Welleck, Kianté Brantley, **Hal Daumé III** and Kyunghyun Cho. *Non-Monotonic Sequential Text Generation*. 2019. ICML.
- ◇ Sobhan Miryoosefi, Kianté Brantley, **Hal Daumé III**, Miroslav Dudík and Robert Schapire. *Reinforcement Learning with Convex Constraints*. 2019. NeurIPS.
- ◇ Kenneth Holstein, Jennifer Wortman Vaughan, **Hal Daumé III**, Miroslav Dudík and Hanna Wallach. *Improving fairness in machine learning systems: What do industry practitioners need?.* 2019. CHI.
- ◇ Elissa Redmiles, Lisa Maszkiewicz, Emily Hwang, Dhruv Kuchhal, Everst Liu, Miraida Morales, Denis Peskov, Sudha Rao, Rock Stevens, Kristina Gligorić, Sean Kross, Michelle L. Mazurek and **Hal Daumé III**. *Comparing and Developing Tools to Measure the Readability of Domain-Specific Texts*. 2019. Conference on Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Sudha Rao and **Hal Daumé III**. *Learning to Ask Good Questions: Ranking Clarification Questions using Neural Expected Value of Perfect Information*. 2018. Conference of the Association for Computational Linguistics (ACL). **Best Paper Award**
- ◇ Hoang M Le, Nan Jiang, Alekh Agarwal, Miroslav Dudík, Yisong Yue and **Hal Daumé III**. *Hierarchical Imitation and Reinforcement Learning*. 2018. ICML.
- ◇ **Hal Daumé III**, John Langford and Amr Sharaf. *Residual Loss Prediction: Reinforcement Learning with no Incremental Feedback*. 2018. ICLR.
- ◇ Octavian Suciú, Radu Mărginean, Yiğitcan Kaya, **Hal Daumé III** and Tudor Dumitraş. *When Does Machine Learning FAIL? Generalized Transferability for Evasion and Poisoning Attacks*. 2018. USENIX.
- ◇ Chris Kedzie, Kathleen McKeown and **Hal Daumé III**. *Content Selection in Deep Learning Models of Summarization*. 2018. Conference on Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Amr Sharaf, Shi Feng, Khanh Nguyen, Kianté Brantley and **Hal Daumé III**. *The UMD Neural Machine Translation Systems [at WMT17 Bandit Learning Task]*. 2017. WMT.
- ◇ Khanh Nguyen, **Hal Daumé III** and Jordan Boyd-Graber. *Reinforcement Learning for Bandit Neural Machine Translation with Simulated Human Feedback*. 2017. Conference on Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Akshay Krishnamurthy, Alekh Agarwal, Tzu-Kuo Huang, **Hal Daumé III** and John Langford. *Active Learning for Cost-Sensitive Classification*. 2017. International Conference on Machine Learning (ICML).
- ◇ Sudha Rao, Daniel Marcu, Kevin Knight and **Hal Daumé III**. *Biomedical event extraction using abstract meaning representation*. 2017. BioNLP.
- ◇ Emily Bender, **Hal Daumé III**, Allyson Ettinger and Sudha Rao. *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*. 2017. Conference of the Association for Computational Linguistics (ACL).
- ◇ Mohit Iyyer, Varun Manjunatha, Anupam Guha, Yogarshi Vyas, Jordan Boyd-Graber, **Hal Daumé III** and Larry Davis. *The amazing mysteries of the gutter: Drawing inferences between panels in comic book narratives*. 2017. Computer Vision and Pattern Recognition (CVPR).

- ◇ Snigdha Chaturvedi, Mohit Iyyer and **Hal Daumé III**. *Unsupervised Learning of Evolving Relationships Between Literary Characters*. 2017. National Conference on Artificial Intelligence (AAAI).
- ◇ **Hal Daumé III**, Nikos Karampatziakis, John Langford and Paul Mineiro. *Logarithmic time one-against-some*. 2017. International Conference on Machine Learning (ICML).
- ◇ Kai-Wei Chang, He He, Stéphane Ross, **Hal Daumé III** and John Langford. *A Credit Assignment Compiler for Joint Prediction*. 2016. Advances in Neural Information Processing Systems (NeurIPS).
- ◇ Snigdha Chaturvedi, Shashank Srivastava, **Hal Daumé III** and Chris Dyer. *Modeling Evolving Relationships Between Characters in Literary Novels*. 2016. National Conference on Artificial Intelligence (AAAI).
- ◇ Mohit Iyyer, Varun Manjunatha, Anupam Guha, Yogarshi Vyas, Jordan Boyd-Graber, **Hal Daumé III** and Larry Davis. *The Amazing Mysteries of the Gutter: Drawing Inferences Between Panels in Comic Book Narratives*. 2016. CVPR.
- ◇ He He, Jordan Boyd-Graber and **Hal Daumé III**. *Interpretese vs. Translationese: The Uniqueness of Human Strategies in Simultaneous Interpretation*. 2016. Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).
- ◇ Hadi Amiri and **Hal Daumé III**. *Short Text Representation for Detecting Churn in Microblogs*. 2016. National Conference on Artificial Intelligence (AAAI).
- ◇ Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber and **Hal Daumé III**. *Feuding Families and Former Friends: Unsupervised Learning for Dynamic Fictional Relationships*. 2016. Conference of the North American Chapter of the Association for Computational Linguistics (NAACL). **Best Paper Award**
- ◇ Fabienne Braune, Alexander Fraser, **Hal Daumé III** and Aleš Tamchyna. *A Framework for Discriminative Rule Selection in Hierarchical Moses*. 2016. WMT.
- ◇ Vicente Ordóñez, Xufeng Han, Polina Kuznetsova, Girish Kulkarni, Margaret Mitchell, Kota Yamaguchi, Karl Stratos, Amit Goyal, Jesse Dodge, Alyssa Mensch, **Hal Daumé III**, Alexander C. Berg, Yejin Choi and Tamara L. Berg. *Large scale retrieval and generation of image descriptions*. 2016. IJCV.
- ◇ **Hal Daumé III**, Nikos Karampatziakis, John Langford and Paul Mineiro. *Logarithmic Time One-Against-Some*. 2016. ICML.
- ◇ Hadi Amiri, Philip Resnik, Jordan Boyd-Graber and **Hal Daumé III**. *Learning Text Pair Similarity with Context-sensitive Autoencoders*. 2016. Conference of the Association for Computational Linguistics (ACL).
- ◇ Kathy McKeown, **Hal Daumé III**, Snigdha Chaturvedi, John Paparrizos, Kapil Thadani, Pablo Barrio, Or Biran, Suvarna Bothe, Michael Collins, Kenneth R Fleischmann, Luis Gravano, Rahul Jha, Ben King, Kevin McInerney, Taesun Moon, Arvind Neelakantan, Diarmuid O’Seaghdha, Dragomir Radev, Clay Templeton and Simone Teufel. *Predicting the impact of scientific concepts using full-text features*. 2016. JAIST.
- ◇ He He, Jordan Boyd-Graber, Kevin Kwok and **Hal Daumé III**. *Opponent Modeling in Deep Reinforcement Learning*. 2016. International Conference on Machine Learning (ICML).
- ◇ Snigdha Chaturvedi, Dan Goldwasser and **Hal Daumé III**. *Ask, and Shall You Receive? Understanding Desire Fulfillment in Natural Language Text*. 2016. National Conference on Artificial Intelligence (AAAI).

- ◇ Dong Xu, Rama Chellappa, Trevor Darrell and **Hal Daumé III**. *Guest Editor's Introduction to the Special Issue on Domain Adaptation for Vision Applications*. 2015. IJCV.
- ◇ Sudha Rao, Allyson Ettinger, **Hal Daumé III** and Philip Resnik. *Dialogue focus tracking for zero pronoun resolution*. 2015. Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).
- ◇ Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber and **Hal Daumé III**. *Deep unordered composition rivals syntactic methods for text classification*. 2015. Conference of the Association for Computational Linguistics (ACL).
- ◇ **Hal Daumé III**, Samir Khuller, Manish Purohit and Gregory Sanders. *On Correcting inputs: Inverse Optimization for Online Structured Prediction*. 2015. FSTTCS.
- ◇ Kai-Wei Chang, Akshay Krishnamurthy, Alekh Agarwal, **Hal Daumé III** and John Langford. *Learning to search better than your teacher*. 2015. International Conference on Machine Learning (ICML).
- ◇ Alina Beygelzimer, **Hal Daumé III**, John Langford and Paul Mineiro. *Learning Reductions that Really Work*. 2015. IEEE Proceedings.
- ◇ Naho Orita, Eliana Vornov, Naomi H Feldman and **Hal Daumé III**. *Why discourse affects speakers' choice of referring expressions*. 2015. Conference of the Association for Computational Linguistics (ACL).
- ◇ He He, Alvin Grissom II, Jordan Boyd-Graber and **Hal Daumé III**. *Syntax-based Rewriting for Simultaneous Machine Translation*. 2015. Conference on Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Kai-Wei Chang, **Hal Daumé III**, John Langford and Stéphane Ross. *Efficient programmable learning to search*. 2014. NeurIPS.
- ◇ Snigdha Chaturvedi, Dan Goldwasser and **Hal Daumé III**. *Predicting Instructor Intervention in MOOC Forums*. 2014. Conference of the Association for Computational Linguistics (ACL).
- ◇ Dan Goldwasser and **Hal Daumé III**. *"I Object!" Modeling Latent Pragmatic Effects in Courtroom Dialogues*. 2014. Conference of the European Association for Computational Linguistics (EACL).
- ◇ Arti Ramesh, Dan Goldwasser, Bert Huang, **Hal Daumé III** and Lise Getoor. *Learning Latent Engagement Patterns of Students in Online Courses*. 2014. National Conference on Artificial Intelligence (AAAI).
- ◇ Arti Ramesh, Dan Goldwasser, Bert Huang, **Hal Daumé III** and Lise Getoor. *Uncovering Hidden Engagement Patterns for Predicting Learner Performance in MOOCs*. 2014. Learning at Scale.
- ◇ Aleš Tamchyna, Fabienne Braune, Alexander Fraser, Marine Carpuat, **Hal Daumé III** and Chris Quirk. *Integrating a Discriminative Classifier into Phrase-based and Hierarchical Decoding*. 2014. The Prague Bulletin of Mathematical Linguistics.
- ◇ He He, **Hal Daumé III** and Jason M. Eisner. *Learning to search in branch and bound algorithms*. 2014. NeurIPS.

- ◇ Alvin Grissom II, Jordan Boyd-Graber, He He, John Morgan and **Hal Daumé III**. *Don't Until the Final Verb Wait: Reinforcement Learning for Simultaneous Machine Translation*. 2014. Conference on Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Mohit Iyyer, Jordan Boyd-Graber, Leonardo Claudino, Richard Socher and **Hal Daumé III**. *A Neural Network for Factoid Question Answering over Paragraphs*. 2014. Conference on Empirical Methods in Natural Language Processing (EMNLP).
- ◇ He He, **Hal Daumé III** and Jason Eisner. *Dynamic Feature Selection for Dependency Parsing*. 2013. Conference on Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Yuancheng Luo, Dmitry N. Zotkin, **Hal Daumé III** and Ramani Duraiswami. *Kernel Regression for Head-Related Transfer Function Interpolation and Spectral Extrema Extraction*. 2013. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP).
- ◇ Ann Irvine, Chris Quirk and **Hal Daumé III**. *Monolingual Marginal Matching for Translation Model Adaptation*. 2013. Conference on Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Yuening Hu, Jordan Boyd-Graber, **Hal Daumé III** and Z. Irene Ying. *Binary to Bushy: Bayesian Hierarchical Clustering with the Beta Coalescent*. 2013. Advances in Neural Information Processing Systems (NeurIPS).
- ◇ Jeffrey Ferraro, **Hal Daumé III**, Scott DuVall, Wendy Chapman, Henk Harkema and Peter Haug. *Improving performance of natural language processing part-of-speech tagging on clinical narratives through domain adaptation*. 2013. Journal of the American Medical Informatics Association.
- ◇ Mohammad Rastegari, Jonghyun Choi, Shobeir Fakhraei, **Hal Daumé III** and Larry S. Davis. *Predictable Dual-View Hashing*. 2013. International Conference on Machine Learning (ICML).
- ◇ Junhui Li, Philip Resnik and **Hal Daumé III**. *Modeling Syntactic and Semantic Structures in Hierarchical Phrase-based Translation*. 2013. Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).
- ◇ Snigdha Chaturvedi, **Hal Daumé III** and Taesun Moon. *Discriminatively Enhanced Topic Models*. 2013. International Conference on Data Mining (ICDM).
- ◇ Marine Carpuat, **Hal Daumé III**, Katharine Henry, Ann Irvine, Jagadeesh Jagarlamudi and Rachel Rudinger. *SenseSpotting: Never let your parallel data tie you to an old domain*. 2013. Conference of the Association for Computational Linguistics (ACL).
- ◇ Margaret Mitchell, Jesse Dodge, Amit Goyal, Kota Yamaguchi, Karl Stratos, Xufeng Han, Alyssa Mensch, Alexander C. Berg, Tamara L. Berg and **Hal Daumé III**. *Midge: Generating Image Descriptions From Computer Vision Detections*. 2012. European Chapter of the Association for Computational Linguistics (EACL).
- ◇ Alexandre Passos, Piyush Rai, Jacques Wainer and **Hal Daumé III**. *Flexible Modeling of Latent Task Structures in Multitask Learning*. 2012. International Conference on Machine Learning (ICML).
- ◇ He He, **Hal Daumé III** and Jason Eisner. *Imitation Learning by Coaching*. 2012. Neural Information Processing Systems (NeurIPS).
- ◇ Ching Lik Teo, Yezhou Yang, **Hal Daumé III**, Cornelia Fermüller and Yiannis Aloimonos. *Towards a Watson That Sees: Language-Guided Action Recognition for Robots*. 2012. ICRA.

- ◇ Abhishek Sharma, Abhishek Kumar, **Hal Daumé III** and David Jacobs. *Generalized Multiview Analysis: A Discriminative latent space*. 2012. Computer Vision and Pattern Recognition (CVPR).
- ◇ **Hal Daumé III**, Jeff M. Phillips, Avishek Saha and Suresh Venkatasubramanian. *Efficient Protocols for Distributed Classification and Optimization*. 2012. ALT.
- ◇ Abhishek Kumar and **Hal Daumé III**. *Learning Task Grouping and Overlap in Multi-task Learning*. 2012. International Conference on Machine Learning (ICML).
- ◇ Amit Goyal, **Hal Daumé III** and Graham Cormode. *Sketch Algorithms for Estimating Point Queries in NLP*. 2012. Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Piyush Rai, Abhishek Kumar and **Hal Daumé III**. *Simultaneously Leveraging Output and Task Structures for Multiple-Output Regression*. 2012. Advances in Neural Information Processing Systems (NeurIPS).
- ◇ Karl Stratos, Aneesh Sood, Alyssa Mensch, Xufeng Han, Margaret Mitchell, Kota Yamaguchi, Jesse Dodge, Amit Goyal, **Hal Daumé III**, Alexander C. Berg and Tamara L. Berg. *Understanding and Predicting Importance in Images*. 2012. Computer Vision and Pattern Recognition (CVPR).
- ◇ Jagadeesh Jagarlamudi and **Hal Daumé III**. *Regularized Interlingual Projections: Evaluation on Multilingual Transliteration*. 2012. 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning.
- ◇ Jagadeesh Jagarlamudi and **Hal Daumé III**. *Low-dimensional Discriminative Reranking*. 2012. Conference on North American Chapter of the Association for Computational Linguistics.
- ◇ Amit Goyal, **Hal Daumé III** and Raul Guerra. *Fast Large-Scale Approximate Graph Construction for NLP*. 2012. Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Jordan Boyd-Graber, Brianna Satinoff, He He and **Hal Daumé III**. *Besting the quiz master: crowdsourcing incremental classification games*. 2012. Conference on Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Jagadeesh Jagarlamudi, **Hal Daumé III** and Raghavendra Udupa. *Incorporating Lexical Priors into Topic Models*. 2012. Conference on European Chapter of the Association for Computational Linguistics (EACL).
- ◇ Jiarong Jiang, Adam Teichert, **Hal Daumé III** and Jason Eisner. *Learned Prioritization for Trading Off Accuracy and Speed*. 2012. Advances in Neural Information Processing Systems (NeurIPS).
- ◇ Jesse Dodge, Amit Goyal, Xufeng Han, Alyssa Mensch, Margaret Mitchell, Karl Stratos, Kota Yamaguchi, Yejin Choi, **Hal Daumé III**, Alexander C. Berg and Tamara L. Berg. *Detecting Visual Text*. 2012. North American Chapter of the Association for Computational Linguistics (NAACL).
- ◇ Abhishek Kumar, Alexandru Niculescu-Mizil, Koray Kavukcuoglu and **Hal Daumé III**. *A Binary Classification Framework for Two-Stage Multiple Kernel Learning*. 2012. International Conference on Machine Learning (ICML).
- ◇ Arvind Agarwal and **Hal Daumé III**. *Generative Kernels for Exponential Families*. 2011. Conference on Artificial Intelligence and Statistics (AI-Stats).
- ◇ Abhishek Kumar and **Hal Daumé III**. *A Co-training Approach for Multiview Spectral Clustering*. 2011. International Conference on Machine Learning (ICML).

- ◇ Jiarong Jiang, Piyush Rai and **Hal Daumé III**. *Message-Passing for Approximate MAP Inference with Latent Variables*. 2011. Conference on Neural Information Processing Systems (NeurIPS).
- ◇ Jagadeesh Jagarlamudi, **Hal Daumé III** and Raghavendra Udupa. *From Bilingual Dictionaries to Interlingual Document Representations*. 2011. Association for Computational Linguistics (ACL).
- ◇ Jay Pujara, **Hal Daumé III** and Lise Getoor. *Using Classifier Cascades for Scalable E-Mail Classification*. 2011. CEAS. **Best Paper Award**
- ◇ Anusua Trivedi, Piyush Rai, **Hal Daumé III** and Scott L. DuVall. *Leveraging Social Bookmarks from Partially Tagged Corpus for Improved Webpage Clustering*. 2011. ACM Transactions on Intelligent Systems and Technology.
- ◇ Yezhou Yang, Ching Lik Teo, **Hal Daumé III** and Yiannis Aloimonos. *Corpus-Guided Sentence Generation of Natural Images*. 2011. Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Jagadeesh Jagarlamudi, Raghavendra Udupa, **Hal Daumé III** and Abhijit Bhole. *Improving Bilingual Projections via Sparse Covariance Matrices*. 2011. Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Abhishek Kumar, Piyush Rai and **Hal Daumé III**. *Co-regularized Multi-view Spectral Clustering*. 2011. Conference on Neural Information Processing Systems (NeurIPS).
- ◇ Amit Goyal and **Hal Daumé III**. *Lossy Conservative Update (LCU) sketch: Succinct approximate count storage*. 2011. Conference on Artificial Intelligence (AAAI).
- ◇ Avishek Saha, Piyush Rai, **Hal Daumé III**, Suresh Venkatasubramanian and Scott L. DuVall. *Active Supervised Domain Adaptation*. 2011. European Conference on Machine Learning (ECML).
- ◇ Piyush Rai and **Hal Daumé III**. *Beam Search based MAP Estimates for the Indian Buffet Process*. 2011. International Conference on Machine Learning (ICML).
- ◇ **Hal Daumé III** and Jagadeesh Jagarlamudi. *Domain Adaptation for Machine Translation by Mining Unseen Words*. 2011. Association for Computational Linguistics.
- ◇ Amit Goyal and **Hal Daumé III**. *Approximate Scalable Bounded Space Sketch for Large Data NLP*. 2011. Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Avishek Saha, Piyush Rai, **Hal Daumé III** and Suresh Venkatasubramanian. *Online Learning of Multiple Tasks and Their Relationships*. 2011. Conference on Artificial Intelligence and Statistics (AI-Stats).
- ◇ Arvind Agarwal, Samuel Gerber and **Hal Daumé III**. *Learning Multiple Tasks using Manifold Regularization*. 2010. Conference on Neural Information Processing Systems (NeurIPS).
- ◇ Jagadeesh Jagarlamudi and **Hal Daumé III**. *Extracting Multilingual Topics from Unaligned Corpora*. 2010. European Conference on Information Retrieval (ECIR).
- ◇ Amit Goyal, Ellen Riloff and **Hal Daumé III**. *Automatically Producing Plot Unit Representations for Narrative Text*. 2010. Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Piyush Rai and **Hal Daumé III**. *Infinite Predictor Subspace Models for Multitask Learning*. 2010. Conference on Artificial Intelligence and Statistics (AI-Stats).

- ◇ Jagadeesh Jagarlamudi, Seth Juarez and **Hal Daumé III**. *Kernelized Sorting for Natural Language Processing*. 2010. Conference on Artificial Intelligence (AAAI).
- ◇ Abhishek Kumar, Avishek Saha and **Hal Daumé III**. *A Co-regularization Based Semi-supervised Domain Adaptation*. 2010. Conference on Neural Information Processing Systems (NeurIPS).
- ◇ Arvind Agarwal and **Hal Daumé III**. *Exponential Family Hybrid Semi-Supervised Learning*. 2009. International Joint Conference on Artificial Intelligence (IJCAI). **Best Paper Award**
- ◇ **Hal Daumé III**. *Unsupervised Search-based Structured Prediction*. 2009. International Conference on Machine Learning (ICML).
- ◇ **Hal Daumé III**. *Bayesian Multitask Learning with Latent Hierarchies*. 2009. Conference on Uncertainty in Artificial Intelligence (UAI).
- ◇ **Hal Daumé III**. *Non-Parametric Bayesian Model Areal Linguistics*. 2009. North American Chapter of the Association for Computational Linguistics (NAACL).
- ◇ Piyush Rai, **Hal Daumé III** and Suresh Venkatasubramanian. *Streamed Learning: One-Pass SVMs*. 2009. International Joint Conference on Artificial Intelligence (IJCAI).
- ◇ **Hal Daumé III**. *Markov Random Topic Fields*. 2009. Association for Computational Linguistics (ACL).
- ◇ Amit Goyal, **Hal Daumé III** and Suresh Venkatasubramanian. *Streaming for Large Scale NLP: Language Modeling*. 2009. North American Chapter of the Association for Computational Linguistics (NAACL).
- ◇ Piyush Rai and **Hal Daumé III**. *Multi-Label Prediction via Sparse Infinite CCA*. 2009. Conference on Neural Information Processing Systems (NeurIPS).
- ◇ Piyush Rai and **Hal Daumé III**. *The Infinite Hierarchical Factor Regression Model*. 2008. Conference on Neural Information Processing Systems (NeurIPS).
- ◇ Ulf Hermjakob, Kevin Knight and **Hal Daumé III**. *Name Translation in Statistical Machine Translation: Learning When to Transliterate*. 2008. Conference of the Association for Computational Linguistics (ACL).
- ◇ Percy Liang, **Hal Daumé III** and Dan Klein. *Structure Compilation: Trading Structure for Features*. 2008. International Conference on Machine Learning (ICML).
- ◇ **Hal Daumé III**. *Cross-Task Knowledge-Constrained Self Training*. 2008. Empirical Methods in Natural Language Processing (EMNLP).
- ◇ **Hal Daumé III**. *Frustratingly Easy Domain Adaptation*. 2007. Conference of the Association for Computational Linguistics (ACL). **Test of Time Award Nomination**
- ◇ **Hal Daumé III** and Lyle Campbell. *A Bayesian Model for Discovering Typological Implications*. 2007. Conference of the Association for Computational Linguistics (ACL).
- ◇ **Hal Daumé III**. *Fast search for Dirichlet process mixture models*. 2007. Eleventh International Conference on Artificial Intelligence and Statistics (AISTats).
- ◇ Yee Whye Teh, **Hal Daumé III** and Daniel Roy. *Bayesian Agglomerative Clustering with Coalescents*. 2007. Conference on Neural Information Processing Systems (NeurIPS).

- ◇ Hal Daumé III and Daniel Marcu. *Bayesian Query-Focused Summarization*. 2006. Conference of the Association for Computational Linguistics (ACL).
- ◇ Hal Daumé III and Daniel Marcu. *Learning as Search Optimization: Approximate Large Margin Methods for Structured Prediction*. 2005. International Conference on Machine Learning (ICML).
- ◇ Hal Daumé III and Daniel Marcu. *A Large-Scale Exploration of Effective Global Features for a Joint Entity Detection and Tracking Model*. 2005. Joint Conference on Human Language Technology and Empirical Methods in Natural Language Processing (HLT/EMNLP).
- ◇ Hal Daumé III and Daniel Marcu. *NP Bracketing by Maximum Entropy Tagging and SVM Reranking*. 2004. Empirical Methods in Natural Language Processing.
- ◇ Hal Daumé III and Eric Brill. *Web Search Intent Induction via Automatic Query Reformulation*. 2004. North American Chapter of the Association for Computational Linguistics (NAACL).
- ◇ Hal Daumé III and Daniel Marcu. *A Phrase-Based HMM Approach to Document/Abstract Alignment*. 2004. Empirical Methods in Natural Language Processing (EMNLP).
- ◇ Hal Daumé III, Kevin Knight, Irene Langkilde-Geary, Daniel Marcu and Kenji Yamada. *The Importance of Lexicalized Syntax Models for Natural Language Generation Tasks*. 2002. 2002 International Conference on Natural Language Generation (INLG).
- ◇ Hal Daumé III and Daniel Marcu. *A Noisy-Channel Model for Document Compression*. 2002. 40th Annual Meeting of the Association for Computational Linguistics (ACL).
- ◇ Eric Nyberg and Hal Daumé III. *Integrated Information Management: An Interactive, Extensible Architecture for Information Retrieval*. 2001. 2001 Human Language Technology Conference (HLT).

Reviews, Commentaries and Invited Papers

- ◇ Roger Levy and Hal Daumé III. *Computational methods are invaluable for typology, but the models must match the questions: Commentary on Dunn et al. (2011)*. 2011. Unpublished.
- ◇ Hal Daumé III. *Semi-supervised or Semi-unsupervised?*. 2009. Unpublished.
- ◇ Hal Daumé III. *Book Review: Automatic Summarization (I. Mani)*. 2004. Unpublished.

Workshop Papers

- ◇ Khanh Nguyen and Hal Daumé III. *Global Voices: Crossing Borders in Automatic News Summarization*. 2019. EMNLP Summarization Workshop.
- ◇ Amr Sharaf and Hal Daumé III. *Structured prediction via learning to search under bandit feedback*. 2017. Workshop on Structured Prediction for NLP.
- ◇ Meir Friedenberg, Hadi Amiri, Hal Daumé III and Philip Resnik. *The UMD CLPsych 2016 Shared Task System: Text Representation for Predicting Triage of Forum Posts about Mental Health*. 2016. Workshop on CL for Clinical Psychology.
- ◇ Arti Ramesh, Dan Goldwasser, Bert Huang, Hal Daumé III and Lise Getoor. *Understanding MOOC Discussion Forums using seeded LDA*. 2014. Workshop on Innovative Use of NLP for Building Educational Applications.

- ◇ Arti Ramesh, Dan Goldwasser, Bert Huang, **Hal Daumé III** and Lise Getoor. *Modeling Learner Engagement in MOOCs using Probabilistic Soft Logic*. 2013. NeurIPS Workshop on Data Driven Education.
- ◇ Snigdha Chaturvedi, **Hal Daumé III**, Taesun Moon and Shashank Srivastava. *A Topical Graph Kernel for Link Prediction in Labeled Graphs*. 2013. ICML workshop on Mining and Learning with Graphs (MLG).
- ◇ Dan Goldwasser and **Hal Daumé III**. *Predicting Dialogue Outcomes over Structured Latent Representations*. 2013. NeurIPS Workshop on Output Representation Learning.
- ◇ Jiarong Jiang, Taesun Moon, **Hal Daumé III** and Jason Eisner. *Prioritized Asynchronous Belief Propagation*. 2013. ICML Workshop on Inferning.
- ◇ **Hal Daumé III**, Jeff Phillips, Avishek Saha and Suresh Venkatasubramanian. *Protocols for Learning Classifiers on Distributed Data*. 2012. Workshop on Artificial Intelligence and Statistics (AI-Stats).
- ◇ He He, **Hal Daumé III** and Jason Eisner. *Cost-sensitive Dynamic Feature Selection*. 2012. ICML 2012 Workshop on Interactions between Inference and Learning (Inferning).
- ◇ Ching L. Teo, Yezhou Yang, **Hal Daumé III**, Cornelia Fermüller and Yiannis Aloimonos. *A Corpus-Guided Framework for Robotic Visual Perception*. 2011. AAAI Workshop on Language-Action Tools for Cognitive Artificial Agents.
- ◇ Amit Goyal and **Hal Daumé III**. *Generating Semantic Orientation Lexicon using Large Data and Thesaurus*. 2011. ACL Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA).
- ◇ Amit Goyal, Piyush Rai and **Hal Daumé III**. *Multiple Hash Functions for Learning*. 2011. NeurIPS Big Learning Workshop.
- ◇ Jason Eisner and **Hal Daumé III**. *Speed-Accuracy Tradeoffs in Nondeterministic Inference Algorithms*. 2011. COST: NeurIPS 2011 Workshop on Computational Trade-offs in Statistical Learning.
- ◇ Amit Goyal, Jagadeesh Jagarlamudi, **Hal Daumé III** and Suresh Venkatasubramanian. *Sketch Techniques for Scaling Distributional Similarity to the Web*. 2010. GEometrical Models of Natural Language Semantics Workshop (GEMS) at ACL.
- ◇ Piyush Rai, Avishek Saha, **Hal Daumé III** and Suresh Venkatasubramanian. *Domain Adaptation meets Active Learning*. 2010. HLT/NAACL Workshop on Active Learning for NLP (ALNLP).
- ◇ **Hal Daumé III**, Abhishek Kumar and Avishek Saha. *Frustratingly Easy Semi-Supervised Domain Adaptation*. 2010. Workshop on Domain Adaptation for NLP.
- ◇ Avishek Saha, Piyush Rai, **Hal Daumé III** and Suresh Venkatasubramanian. *Active Online Multitask Learning*. 2010. ICML 2010 Workshop on Budgeted Learning (Budget).
- ◇ Amit Goyal, Ellen Riloff, **Hal Daumé III** and Nathan Gilbert. *Toward Plot Units: Automatic Affect State Analysis*. 2010. HLT/NAACL Workshop on Computational Approaches to Analysis and Generation of Emotion in Text (CAET).
- ◇ Piyush Rai, Anusua Trivedi, **Hal Daumé III** and Scott L. DuVall. *Multiview Clustering with Incomplete Views*. 2010. NeurIPS Workshop on Machine Learning for Social Computing.

- ◇ Piyush Rai and **Hal Daumé III**. *Multitask Learning via Mixture of Linear Subspaces*. 2010. NeurIPS Workshop on Transfer Learning by Learning Rich Generative Models.
- ◇ Abhishek Kumar, Piyush Rai and **Hal Daumé III**. *Co-regularized Spectral Clustering with Multiple Kernels*. 2010. NeurIPS Workshop on New Directions in Multiple Kernel Learning.
- ◇ Anusua Trivedi, Piyush Rai, Scott L. DuVall and **Hal Daumé III**. *Exploiting Tag and Word Correlations for Improved Webpage Clustering*. 2010. CIKM Workshop on Search and Mining User-generated Contents (SMUC).
- ◇ Amit Goyal, Jagadeesh Jagarlamudi, **Hal Daumé III** and Suresh Venkatasubramanian. *Sketching Techniques for Large Scale NLP*. 2010. HLT/NAACL Workshop on the Web as a Corpus (WAC).
- ◇ Piyush Rai and **Hal Daumé III**. *Fast Search for Infinite Latent Feature Models*. 2009. NeurIPS Workshop on Non-parametric Bayes (NP-Bayes).
- ◇ Piyush Rai and **Hal Daumé III**. *Multitask Learning using Nonparametrically Learned Predictor Subspaces*. 2009. NeurIPS Workshop on Learning from Multiple Sources.
- ◇ Adam R. Teichert and **Hal Daumé III**. *Unsupervised Part of Speech Tagging Without a Lexicon*. 2009. NeurIPS Workshop on Grammar Induction, Representation of Language and Language Learning (GIRLLL).
- ◇ Amrish Kapoor, Piyush Rai and **Hal Daumé III**. *Factor Regression Combining Heterogeneous Sources of Information*. 2009. NeurIPS Workshop on Learning From Multiple Sources with Applications to Robotics (LMS).
- ◇ **Hal Daumé III**. *HBC: Hierarchical Bayes Compiler*. 2008. Workshop on Bayesian Inference.
- ◇ Devyani Ghosh, John Carter and **Hal Daumé III**. *Perceptron-based Coherence Predictors*. 2008. 2nd Workshop on Chip Multiprocessor Memory Systems and Interconnects (ICSA).
- ◇ **Hal Daumé III** and Daniel Marcu. *Bayesian Summarization at DUC and a Suggestion for Extrinsic Evaluation*. 2005. Document Understanding Conference (DUC).
- ◇ **Hal Daumé III** and Daniel Marcu. *Bayesian Multi-Document Summarization at MSE*. 2005. Workshop on Multilingual Summarization Evaluation (MSE).
- ◇ **Hal Daumé III**, John Langford and Daniel Marcu. *Search-Based Structured Prediction as Classification*. 2005. NeurIPS Workshop on Advances in Structured Learning for Text and Speech Processing (ASLTSP).
- ◇ **Hal Daumé III** and Daniel Marcu. *A Tree-Position Kernel for Document Compression*. 2004. Fourth Document Understanding Conference (DUC).
- ◇ **Hal Daumé III** and Daniel Marcu. *Generic Sentence Fusion is an Ill-Defined Summarization Task*. 2004. Text Summarization Branches Out Workshop at ACL (TextSum).
- ◇ **Hal Daumé III** and Daniel Marcu. *Supervised clustering with the Dirichlet process*. 2004. NeurIPS Workshop on Learning With Structured Outputs (LwSO).
- ◇ **Hal Daumé III**, Abdesammad Echihabi, Daniel Marcu, Dragos Stefan Munteanu and Radu Soricut. *GLEANS: A Generator of Logical Extracts and Abstracts for Nice Summaries*. 2002. Second Document Understanding Conference (DUC).

Unpublished

- ◇ Amr Sharaf and Hal Daumé III. *Meta-Learning for Contextual Bandit Exploration*. 2019. arxiv.
- ◇ Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III and Kate Crawford. *Datasheets for Datasets*. 2018. arxiv.
- ◇ Hal Daumé III and Katherine Heller. *NeurIPS 2018 Demographics and Inclusion Survey: Summary of Responses*. 2018. NeurIPS (not a normal paper).
- ◇ Hal Daumé III, John Langford and Daniel Marcu. *Searn in Practice*. 2006. Unpublished.
- ◇ Hal Daumé III. *From Zero to Reproducing Kernel Hilbert Spaces in Twelve Pages or Less*. 2004. Unpublished.
- ◇ Hal Daumé III. *Carefully Approximated Bayes Factors for Feature Selection in MaxEnt Models*. 2004. Unpublished.
- ◇ Hal Daumé III. *Notes on CG and LM-BFGS Optimization of Logistic Regression*. 2004. Unpublished.
- ◇ Hal Daumé III. *Yet Another Haskell Tutorial*. 2002. Unpublished.
- ◇ Hal Daumé III. *A Phrase-Based HMM*. 2002. Unpublished.
- ◇ Hal Daumé III. *Asymmetry of Coordination*. 2001. Unpublished.

Publicly Available Software

Much software I have developed recently is on github at <http://github.com/hal3/>. I have also contributed significantly to the vowpal wabbit online training system, originally developed by John Langford: contributions include multiclass classification, structured learning, etc. Below are older software packages.

MEGAM: Optimization software for maximum entropy models; implementation based on conjugate gradient and limited-memory BFGS methods. (Downloaded over 20,000 times.)
<http://hal3.name/megam/>

HBC: Software that compiles statistical models into executable C code. (Downloaded over 6,000 times.)
<http://hal3.name/HBC/>.

TAGCHUNK: Software for performing joint inference over part of speech tags and syntactic chunks. (Downloaded over 400 times.)
<http://hal3.name/TagChunk/>

Teaching Experience

CMSC 421, Artificial Intelligence Fall 2016 (enrollment: 90) Spring 2012 (enrollment: 40)	UMD, CS
CMSC 422, Undergraduate Machine Learning Fall 2015 (enrollment: 90) Spring 2013 (enrollment: 44)	UMD, CS

CMSC 723, Computational Linguistics I Fall 2019 (enrollment: 90) Fall 2012 (enrollment: 51) Fall 2010 (enrollment: 43)	UMD, CS, Ling, iSchool
CMSC 726, Machine Learning Spring 2016 (enrollment: 45) Fall 2011 (enrollment: 47)	UMD, CS
CMSC 828, Reinforcement Learning Fall 2016 (enrollment: 35)	UMD, CS
CMSC 828, Advanced Topics in NLP Fall 2011 (enrollment: 10; co-taught with Naomi Feldman and Jordan Boyd-Graber)	UMD, CS, Ling, iSchool
CS 5300/6300, Artificial Intelligence Spring 2009, Spring 2010 (enrollments: 78, 74) Received Dean’s letter for Excellence in Teaching	U Utah, CS
CS 5350/6350, Machine Learning Spring 2007, Spring 2008, Fall 2008, Fall 2009 (enrollments: 28, 32, 36, 46)	U Utah, CS
CS 5964/6964, Applications of Natural Language Processing Fall 2007 (enrollment: 28)	U Utah, CS
CS 7941, Topics in Machine Learning Every semester, 2007–current (enrollments: about 15) Received Dean’s letter for Excellence in Teaching	U Utah, CS
Various Teaching as a Student USC CSCI 544, Natural Language Processing (Teaching Assistant) (Spring 2005), Outstanding TA Award USC CSCI 544, Natural Language Processing (Invited Lecturer) (Spring 2004) USC CSCI 599, Topics in Statistical Learning (Invited Lecturer) (Spring 2004) USC CSCI 562, Empirical Methods in NLP (Invited Lecturer) (Fall 2003) CMU MATH 229, Set Theory (Grader) (Spring 2000) CMU MATH 115/116, Calculus I & II (Teaching Assistant) (Fall 1999)	USC CS and CMU Math

Invited Talks

- (Meta-)Learning from Interaction** (August 2019) NYU Machine Learning Reading Group
- Beyond demonstrations: Learning behavior from higher-level supervision** (June, 2019) ICML Workshop: Imitation, Intent, and Interaction (I3)
- Out of Order! Flexible neural language generation** (June 2019) NAACL Workshop: Methods for Optimizing and Evaluating Neural Language Generation
- Learning language through interaction** (December 2016, ...) Georgetown, Amazon, USC, GATech, Williams College, UW
- Bias in AI** (November 2016) UMD MCWIC Diversity Summit
- Imitation learning and recurrent neural networks mashup** (December 2015) CIFAR NCAP Workshop
- Algorithms that learn to think on their feet** (October 2015) UC Santa Cruz, Computer Science Seminar
- Interpretese vs Translationese** (August 2015) Shonan Workshop
- Algorithms that learn to think on their feet** (May 2015) Invited Talk at ICLR
- A picture is worth 13.6 words (on average)** (February 2015) MSR NY Tea Talk

- Domain adaptation: the problem of new labels** (December 2014) NeurIPS workshop on Transfer and Multitask Learning
- Algorithms that learn to think of their feet** (October 2014) Columbia University Data Science Institute
- Understanding and adapting statistical models: an exploration in language** (May 2014) University of Toronto, Machine Learning Colloquium
- Predicting linguistic structures accurately and efficiently** (November 2013) Université de Montreal, CS Colloquium
- A picture is worth 13.6 words on average** (November 2013) University de Montreal, LISA Lab Talk Series
- Better! Faster! Stronger! Learning to balance accuracy and efficiency when predicting linguistic structures** (September 2013) National Research Council of Canada, NLP Colloquium
- Imperative Learning** (June 2013) Prediction with Sequential Models workshop at ICML (joint invited talk with John Langford)
- Better! Faster! Stronger! Learning to balance accuracy and efficiency when predicting linguistic structures** (June 2013) Inferning workshop at ICML
- Dynamic Feature Selection** (June 2013) Budgeting workshop at ICML
- Better! Faster! Stronger! Learning to balance accuracy and efficiency when predicting linguistic structures** (January 2013) Microsoft Research, New York
- Complex Predictions Need not be Slow** (September 2012) University of Southern California, CS Colloquium
- Transfer learning in language** (September 2012) Machine Learning for Spoken Language Processing workshop at ICASSP
- NLP for smart people who know nothing about NLP** (August 2012) Course at the Machine Learning Summer School, UC Santa Cruz
- Complex predictions need not be slow** (May 2012) University of Toronto, Statistics Colloquium
- Structured prediction in linguistic domains** (March 2012) University of Utah, CS Colloquium
- Complex prediction need not be slow** (March 2012) Tokyo Institute of Technology, CS Colloquium Series
- Transfer learning in Language** (March 2012) International Workshop on Statistical Machine Learning, Tokyo
- Moving between Tasks and Domains** (November 2011) Temple University, AI Seminar Series
- Structured Prediction Need Not Be Slow** (June 2011) TextGraphs Workshop at NAACL
- Learning Structured Prediction by Demonstration** (June 2011) Laboratory for Telecommunication Science
- Structure and Knowledge in Natural Language Processing** (June 2011) University of Trento, Italy
- Structure and Knowledge in Natural Language Processing** (April 2011) Microsoft Research Tutorial, Redmond, WA
- Structure and Knowledge in Natural Language Processing** (March 2011) Carnegie Mellon University, AI Seminar Series
- Transfer Learning in Language: We've Got a Long Way to Go...** (December, 2010) Workshop on Transfer Learning via Rich Generative Models at NeurIPS

- Linguists get the abstraction, machines get the details** (July, 2010) Linguistics Meets NLP Workshop at ACL
- Structure and Knowledge in Natural Language Processing** (April, 2010) University of Illinois, Computer Science
- Structure and Knowledge in Natural Language Processing** (November, 2009) University of Maryland, Computer Science
- Statistics, Typology and Language Processing** (October, 2009) University of Utah, Department of Linguistics
- { ϵ -,Semi-,Un-}supervised Search-based Structured Prediction** (May, 2009) Yahoo! Research, New York
- Non-parametric Models with Latent Hierarchies** (April, 2009) Columbia University, Department of Statistics
- Domain Adaptation in Natural Language** (March, 2009) Johns Hopkins University, Center for Language and Speech Processing
- Domain Adaptation in Natural Language** (November, 2008) University of Texas at Austin, Computer Science
- Summer School in Computational- and Psycho-Linguistics** (September 2008) University of Edinburgh
- Automatic Document Summarization** (January 2007) Information Sciences Institute
- Search-based Structured Prediction** (March 2007) University of Maryland, College Park
- Bayesian Techniques for Natural Language Processing** (July 2007) Microsoft Research Asia, Beijing, China

Tutorials

- Imitation Learning and its Application to Natural Language Generation** (Dec 2019) NeurIPS Invited Tutorial (with Kyunghyun Cho)
- Imitation Learning** (August 2018) Vector Institute Summer School on Reinforcement Learning
- Advances in Structured Prediction** (June 2015) ICML Tutorial (with John Lanford)
- Hands-on Learning to Search for Joint Prediction** (May 2015) NAACL Tutorial (with John Lanford)
- From Structured Prediction to Inverse Reinforcement Learning** (August 2011) Tutorial at the Conference of the Association for the Advancement of Artificial Intelligence (AAAI)
- Beyond Structured Prediction: Inverse Reinforcement Learning** (July 2011) Tutorial at the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)
- Structured Prediction 101** (May 2011) University of Trento, Italy
- Domain Adaptation** (April 2011) Microsoft Research Tutorial, Redmond, WA
- From Structured Prediction to Inverse Reinforcement Learning** (July 2010) Tutorial at the Conference of the Association for Computational Linguistics (ACL)
- Domain Adaptation** (June 2010) Tutorial at the International Conference in Machine Learning (ICML) Co-presented with John Blitzer (Google)
- Bayesian Techniques for HLT Researchers** (June 2006) Tutorial at the Human Language Technology Conference (HLT/NAACL)

Statistical Learning Theory and SVMs for AI Researchers (February 2006) Tutorial at the ISI Intelligent Systems Division Retreat

NLP for Smart People Who Know Nothing about NLP (January 2006) Toyota Technological Institute at Chicago

Bayesian Methods for Natural Language Processing (March 2005) Tutorial at the Information Sciences Institute, AI seminar

Funding

- Data-Driven Algorithms for Improving Diversity and Fairness in Human-Machine Decision Making** Under Review
NSF RI Medium Proposal. \$1,200,000 over three years (half to UMD). PI: John Dickerson (UMD); Co-I with Jeff Foster and J.P. de Ruiter (Tufts).
- FAI: FAIR-EVAL: Integrating Stakeholder Values in Natural Language Processing Shared Tasks** Under Review
NSF/Amazon Grant Proposal. \$1,244,904 over three years. PI. Co-PIs: Emily Bender and Bernease Herman (UW).
- RI: EAGER: Collaborative Research: Adaptive Heads-up Displays for Simultaneous Interpretation** 2018–2020
NSF-IIS EAGER Grant, \$150k over 18 months. PI: Graham Neubig (CMU); UMD-PI: Jordan Boyd-Graber; UW-PI: Leah Findlater.
- Linguistic Semantics and Discourse from Leaky Distant Supervision** 2016–2020
NSF-IIS Small Grant, \$500k over three years. Sole PI.
- Learning from User Interaction** 2016
Facebook Research Award, \$25k; Sole PI.
- Neural Machine Translation from Weak User Feedback** 2016
Amazon Research Award, \$65k; Sole PI.
- Discrete Algorithms in NLP** 2014
NSF-IIS EAGER Grant, \$75k for one year. Co-PI with Samir Khuller (UMD).
- Learning Integer Linear Programming for Natural Language Processing** 2014
Google Research Award, \$65k; Sole PI.
- NRT-DESE: Flexibility in Language Processes and Technology: Human- and Global-Scale** 2013–2019
NSF NRT-DESE Grant, \$3.0m over five years. PI: Colin Phillips (UMD).
- Bayesian Thinking on Your Feet—Embedding Generative Models in Reinforcement Learning for Sequentially Revealed Data** 2013–2016
NSF-IIS Small Grant, \$550k over three years. PI: Jordan Boyd-Graber (UC Boulder).
- Developing the Next Generation of Community Financial CyberInfrastructure for Monitoring and Modeling Financial Eco-Systems and for Managing Systemic Risk** 2013–2014
NSF CI-PI Infrastructure Grant, \$100k for one year. PI: Louiqa Raschid (UMD).
- Domain Adaptation and Translation** 2012–2013
DARPA CSSG Grant, \$400k over two years. Sole PI.
- Learning the Relationship between the Anatomy and Spatial Hearing** 2011–2012
NSF-IIS Small Grant, \$164k for one year. PI: Ramani Duraiswami (UMD).

Distributed Empirical Language Processing for Human Interaction	2011–2016
DARPA “BOLT” Grant, \$3.3 <i>m</i> for UMD over five years. PI: Doug Oard (UMD).	
Discovering and explaining technical emergence through analysis of the language and structure of scientific publications	2011–2016
IARPA “FUSE” Grant, \$2.8 <i>m</i> for UMD over five years. PI: Kathy McKeown (Columbia University).	
Distributed Domain Adaptation for Millions of Domains	2011
DARPA CSSG Grant, \$92 <i>k</i> over one year. Sole PI.	
Learned Dynamic Prioritization	2010–2013
NSF-IIS Medium Grant, \$496 <i>k</i> for UMD over three years. Co-PI with Jason Eisner (JHU)	
Statistical Linguistic Typology	2009–2012
NSF-IIS Small Grant, \$550 <i>k</i> over three years. Sole PI.	
Graph Algorithms for Large-Data NLP	2009–2010
Google Research Award, \$55 <i>k</i> ; Co-PI with Suresh Venkatasubramanian.	
Language Understanding for Research Papers	2009
U Utah URC Faculty Research & Creative Grant Proposal, \$6 <i>k</i> . Sole PI.	
Advanced Learning Technology for Translation	2009
U Utah Seed Grant, \$35 <i>k</i> . Sole PI.	
Computational Thinking Olympiad: Brainstorming Workshop	2008
NSF-IIS grant, \$20 <i>k</i> over two years. Sole PI.	
Cross-task Learning for Natural Language Processing (REU)	2008
NSF-IIS Research Experience for an Undergraduate Supplement, \$6 <i>k</i> over one year. Sole PI.	
Cross-task Learning for Natural Language Processing	2007–2010
NSF-IIS Small Grant, \$404 <i>k</i> over three years. Sole PI.	

External Service

Coordinating Committee, ACL 2017

Executive Board Chair, NAACL, 2014–2016

Computational Thinking Olympiad Co-program chair of an NSF-sponsored Olympiad in *computational thinking*

Executive Board Member, NAACL, elected

Research blog: **Natural language processing blog** Since 2006, <http://nlpers.blogspot.com>

Workshop Organization

Co-organizer, **Building Linguistically Generalizable NLP Systems**

Workshop at the European Association for Computational Linguistics, 2017

Co-organizer, **Let’s Discuss: Learning Methods for Dialogue**

Workshop at the Conference on Neural Information Processing Systems (NeurIPS) 2016

Co-organizer, **Workshop on Human-Computer Question Answering**

Workshop at the Human Language Technology Conference (HLT/NAACL), 2016

Co-organizer, **Simultaneous Machine Interpretation**

Shonan Workshop, Japan, 2015

Co-organizer, Discrete Algorithms for Machine Learning

Shonan Workshop, Japan, 2013

Co-organizer, Combining Strategies for Reducing the Label Cost

Workshop at the International Conference on Machine Learning (ICML), 2011

Co-organizer, Prior knowledge in text and language processing

Workshop at the International Conference on Machine Learning (ICML), 2008

Co-organizer, Computationally Hard Problems in Speech & Language

Workshop at the Human Language Technology Conference (HLT/NAACL), 2006

Co-organizer, Bayesian Methods in Natural Language Processing

Workshop at the Conference on Neural Information Processing Systems (NeurIPS) 2005

Journal Editorial Roles

Editorial Boards: Transactions of the ACL (2013–2021), Computational Linguistics (2011–2013), Machine Learning Journal (2008–2016), IEEE Transactions on Speech and Language Processing (2008–2012)

Ad-hoc Reviews: Journal of Machine Learning Research, Journal of Computational and Graphical Statistics, Journal of Natural Language Engineering, IEEE’s Transactions on Knowledge and Data Engineering, ACM Transactions on Speech and Language Processing, ACM Transactions on Pattern Analysis and Machine Intelligence, Journal of Artificial Intelligence Research, IBM Systems Journal

Conference Chair Roles

Conference Program Chair: NAACL 2013 (with Katrin Kirchhoff), ICML 2020 (with Aarti Singh)

Diversity and Inclusion Co-Chair: NeurIPS 2018 (with Katherine Heller)

Sponsorship Chair: ICML 2010–2011

Publicity Chair: COLING 2010

Publications Chair: ICML 2009–2010

Publicity Chair: ACL 2008

Conference Area Chair: many times at ICML, IJCAI, NeurIPS, NAACL, ACL, EMNLP, AISTats.

Conference PCs: many times for the above conferences and more

Internal Service

UMD Language Science Executive Committee (2013–present)

MSR Diversity Allies Leadership Committee (2017–present)

UMD CS Diversity Committee (2014–present)

UMD CS Department Council (2015–2016)

UMD CS Artificial Intelligence Field Committee Chair (2013–2016)

UMD Computer Science Hiring Committee (2014, 2019)

Director of Computational Linguistics & Information Processing Lab (2013–2017)

Executive Board, Language Science Center (2013–2017)

UMIACS APT Committee Member (2012)

UMD Computer Science Hiring Committee (2012)

Artificial Intelligence Brown Bag Lunch (2012)

Co-director of Computational Linguistics & Information Processing Lab (2011–2012)

UMD Admissions Committee Member (2010–2010)

Utah Curriculum Committee Chair (2008–2010) Overseen significant advances in math/science requirements, formation of undergraduate tracks, and development of four new undergraduate courses

Utah Masters in Statistics Program, SoC Representative (2008–2010) Representative from the School of Computing to the Masters in Statistics Program

Utah Committee Memberships () Admissions (2007, 2008); Curriculum (2007, 2008); College Council (2007, 2008)

Organizer, Natural Language Seminar Series (2001–2006) Weekly seminar at the Information Sciences Institute