

Douglas C Downey

Electrical Engineering and Computer Science Dept.
Northwestern University, Ford 3-345
2133 Sheridan Road
Evanston, IL 60208

ddowney@eecs.northwestern.edu
(847)491-3710
<http://www.cs.northwestern.edu/~ddowney>

Research Interests	Natural Language Processing, Machine Learning, Artificial Intelligence	
Education	University of Washington Ph.D. Computer Science and Engineering, 2008 Dissertation Title: <i>Redundancy in Web-scale Information Extraction: Probabilistic Model and Experimental Results</i> M.S. Computer Science and Engineering, 2004	September 2002 – December 2008
	Case Western Reserve University B.S./M.S. in Computer Science Minors in Mathematics and Economics	August 1996 – May 2000
Professional Experience	Northwestern University , Associate Professor Department of Electrical Engineering and Computer Science	September 2014 – present
	Northwestern University , Assistant Professor Department of Electrical Engineering and Computer Science	September 2008 – 2014
	University of Washington , Research Assistant <i>Advisor: Oren Etzioni</i>	June 2003 – September 2008
	Microsoft Corporation , Research Intern <i>Mentor: Susan Dumais Manager: Eric Horvitz</i>	March 2006 – July 2006
	Intel Corporation , Internet Software Engineer	October 2000 – September 2002
	Case Western Reserve University , Research Assistant <i>Advisor: Randall D. Beer</i>	October 1998 – September 2000
Honors and Awards	NSF CAREER Award, 2014 Northwestern EECS "Teacher of the Year," 2012 US SECRET Security Clearance, 2012-2016 DARPA Computer Science Study Panel, 2011 Microsoft New Faculty Fellowship, 2010 Microsoft Research PhD Fellowship sponsored by Microsoft Live Labs, 2006 IJCAI Distinguished Paper Award, 2005 NSF Graduate Fellowship, 2001 NDSEG Graduate Fellowship (awarded), 2001 United States Presidential Scholar, 1996 Albert W. Smith scholarship, 1996	

Teaching Experience	<p>Northwestern University, Assistant/Associate Professor Machine Learning (EECS 349) Intro to Artificial Intelligence (EECS 348) Statistical Language Models (EECS 395/495) Probabilistic Graphical Models (EECS 474) Web Information Retrieval and Extraction (EECS 395/495) Fundamentals of Computer Programming (EECS 111)</p>
Professional Service	<p>Area Chair/Senior Program Committee: EMNLP 2018,2014; AAI 2018; IJCAI 2016; NAACL-HLT 2013; EMNLP-CoNLL 2012 Action Editor: TACL, 2018-present Program Committees: AAAI 2010-2013, 2018; AISTATS 2017,2018; ACL 2010-2017; CoNLL 2014-2015; COLING 2010,2014,2016, 2018; EMNLP 2010,2013,2015-2017; ICLR 2018; ICML 2010,2013,2014,2018; ICWSM 2014,2017; IJCAI 2009,2011,2013,2017; IJCNLP 2017; KDD 2015; NAACL-HLT 2012,2013,2015,2016,2018; NIPS 2016,2017,2018; SIGIR 2017,2018; WSDM 2015-2017 (outstanding reviewer award 2015); WWW 2011-2015 Reviewer for: TACL (standing reviewer), 2013-2017 Journal of Artificial Intelligence Research, 2008-2009, 2013 Transactions on Knowledge and Data Engineering, 2010-2013 National Center for Women & IT Aspirations in Computing, 2009-2012, 2015 NSF Division of Information and Intelligent Systems, 2009, 2011-2012, 2014-2018 AI Journal, 2008-2010 HLT/NAACL, 2009-2010 VLDB Journal, 2009-2010 User Modeling, 2007 Data Mining and Knowledge Discovery Journal, 2007 AAAI, 2005 ACM Transactions on Internet Technology, 2004</p>
University Service	<p>CS Undergraduate Curriculum Committee, 2009-present EECS Computing Facilities Committee, 2009-present EECS Distinguished Speaker Committee, 2009-present McCormick School of Engineering Freshman Advisor, 2008-2013 Cognitive Science Program Committee, 2012-2014</p>
Publications	<ol style="list-style-type: none"> 1. Yiben Yang, Larry Birnbaum, Jiping Wang, Doug Downey. Extracting Commonsense Properties from Embeddings with Limited Human Guidance. Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL), 2018. 2. Thanapon Noraset, David Demeter, Doug Downey. Controlling Global Statistics in Recurrent Neural Network Text Generation. Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), 2018. 3. Zheng Yuan, Doug Downey. OTyper : A Neural Architecture for Open Named Entity Typing. Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), 2018. 4. Ammar, W., Groeneveld, D., Bhagavatula, C., Beltagy, I., Crawford, M., Downey, D., Dunkelberger, J., Elgohary, A., Feldman, S., Ha, V. and Kinney, R., 2018. Construction of the Literature Graph in Semantic Scholar. Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL –

Industry Track), 2018

5. Jared Fernandez and Doug Downey. Sampling Informative Training Data for RNN Language Models. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL – student paper)*, 2018.
6. Thanapon Noraset, Chen Liang, Larry Birnbaum, Doug Downey. Definition Modeling: Learning to Define Word Embeddings in Natural Language. *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
7. Jared Fernandez, Zhaocheng Yu, Doug Downey. VecShare: A Framework for Sharing Word Representation Vectors. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.
8. Yuji Mo, Stephen Scott, Doug Downey. Learning Hierarchically Decomposable Concepts with Active Over-Labeling. *International Conference on Data Mining (ICDM)*, 2016.
9. Tiantian Zhu, Hongyu Gao, Yi Yang, Kai Bu, Yan Chen, Doug Downey, Kathy Lee, Alok Choudhary. Beating the Artificial Chaos: Fighting OSN Spam Using Its Own Templates. *IEEE/ACM Transactions on Networking*, 2016.
10. Chandra Sekhar Bhagavatula, Thanapon Noraset, and Doug Downey. TabEL: Entity Linking with Web Tables. *International Semantic Web Conference (ISWC)*, 2015.
11. Yi Yang, Doug Downey, and Jordan Boyd-Graber. Efficient Methods for Incorporating Knowledge into Topic Models. *Empirical Methods in Natural Language Processing (EMNLP)*, 2015.
12. Doug Downey, Chandra Sekhar Bhagavatula, Yi Yang. Efficient Methods for Inferring Large Sparse Topic Hierarchies. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2015.
13. Chandra Sekhar Bhagavatula, Thanapon Noraset, Doug Downey. TextJoiner: On-demand Information Extraction with Multi-Pattern Queries. (AKBC workshop, NIPS 2014).
14. Mohammed Alam, Doug Downey. Analyzing the Content Emphasis of Web Search Engines. *Proceedings of the 37th Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR)*, 2014.
15. Thanapon Noraset, Chandra Sekhar Bhagavatula, Doug Downey. Adding High-Precision Links to Wikipedia. *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
16. Fei Huang, Arun Ahuja, Doug Downey, Yi Yang, Yuhong Guo, and Alexander Yates. Learning Representations for Weakly Supervised Natural Language Processing Tasks. *Computational Linguistics*, 2014.
17. Yi Yang, Shimei Pan, Doug Downey, Kunpeng Zhang. Active Learning with Constrained Topic Model (ILLVI Workshop, ACL), 2014.
18. Hongyu Gao, Yi Yang, Kai Bu*, Yan Chen, Doug Downey, Kathy Lee, Alok Choudhary. Spam ain't as diverse as it seems: throttling OSN spam with templates underneath. *Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC)*, 2014.
19. Michael Lucas and Doug Downey. Scaling Semi-supervised Naïve Bayes

- with Feature Marginals. *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*, 2013.
20. Yi Yang, Alex Yates, Doug Downey. Overcoming the Memory Bottleneck in Distributed Training of Latent Variable Models of Text. *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT, short paper)*, 2013.
 21. Bhagavatula, C.S., Noraset, T., and Downey, D. Methods for Exploring and Mining Tables on Wikipedia. *Proceedings of ACM SIGKDD Interactive Data Exploration and Analytics (IDEA)*, 2013.
 22. Downey, D., Bhagavatula, C.S., and Yates, A. Using Natural Language to Integrate, Evaluate, and Optimize Extracted Knowledge Bases. *Proceedings of Automated Knowledge Base Construction: The 3rd Workshop on Knowledge Extraction at CIKM (AKBC)*, 2013.
 23. Raymond V. Burke, Keith D. Allen, Monica R. Howard, Doug Downey, Michael G. Matz, Scott L. Bowen. Tablet-based Video Modeling and Prompting in the Workplace for Individuals with Autism. *Journal of Vocational Rehabilitation*, 38(1): 1-14, 2013.
 24. Kunpeng Zhang, Doug Downey, Zhengzhang Chen, Yusheng Xie, Yu Cheng, Ankit Agrawal, Wei-keng Liao, Alok Choudhary. A probabilistic graphical model for brand reputation assessment in social networks. *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2013.
 25. Brent Hecht, Samuel H. Carton, Mahmood Quaderi, Johannes Schöning, Martin Raubal, Darren Gergle, and Doug Downey. Explanatory Semantic Relatedness and Explicit Spatialization for Exploratory Search. *Proceedings of the ACM Conference on Research and Development in Information Retrieval (SIGIR)*, 2012.
 26. Kunpeng Zhang, Yusheng Xie, Yu Cheng, Daniel Honbo, Doug Downey, Ankit Agrawal, Wei-keng Liao, Alok Choudhary. Sentiment Identification by Incorporating Syntax, Semantics and Context Information. *Proceedings of the ACM Conference on Research and Development in Information Retrieval (SIGIR)*, 2012.
 27. Lev Ratinov, Doug Downey, Dan Roth, Mike Anderson. Local and Global Algorithms for Disambiguation to Wikipedia. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-HLT)*, 2011.
 28. Fei Huang, Alex Yates, Arun Ahuja, and Doug Downey. Language Models as Representations for Weakly-Supervised NLP Tasks. *Conference on Natural Language Learning (CoNLL)*, 2011.
 29. Doug Downey, Oren Etzioni, Stephen Soderland. Analysis of a Probabilistic Model of Redundancy in Unsupervised Information Extraction. *Artificial Intelligence*, 174(11): 726-748, 2010.
 30. Arun Ahuja and Doug Downey. Improved Extraction Assessment through Better Language Models. *Proceedings of Human Language Technologies: The 11th Annual Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL)*, 2010.
 31. Doug Downey, Arun Ahuja, Mike Anderson. Learning to Integrate Relational Databases with Wikipedia. *Proceedings of the IJCAI-09 Workshop on User-Contributed Knowledge and Artificial Intelligence (WIKIAI)*, 2009.

32. Doug Downey, Oren Etzioni. Look Ma, No Hands: Analyzing the Monotonic Feature Abstraction for Text Classification. In D. Koller, Y. Bengio, D. Schuurmans and L. Bottou (Eds.), *Advances in Neural Information Processing Systems (NIPS) 21*, 2008.
33. Doug Downey, Susan Dumais, Dan Liebling and Eric Horvitz. Understanding the relationship between searchers' queries and information goals. *Proceedings of the 17th Conference on Information and Knowledge Management (CIKM)*, 2008
34. Alan Ritter, Doug Downey, Stephen Soderland, and Oren Etzioni. It's a Contradiction -- No, It's Not: A Case Study using Functional Relations. *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2008
35. Doug Downey, Stefan Schoenmackers, Oren Etzioni. Sparse Information Extraction: Unsupervised Language Models to the Rescue. *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics (ACL)*, 2007.
36. Doug Downey, Matthew Broadhead, Oren Etzioni. Locating Complex Named Entities in Web Text. *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.
37. Doug Downey, Susan Dumais, Eric Horvitz. Models of Searching and Browsing: Languages, Studies, and Applications. *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.
38. Doug Downey, Susan Dumais, and Eric Horvitz. Heads and Tails: Studies of Web Search with Common and Rare Queries. *Proceedings of the 30th Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR)*, 2007.
39. Doug Downey, Oren Etzioni, and Stephen Soderland. A Probabilistic Model of Redundancy in Information Extraction. *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI)*, 2005. **Distinguished Paper Award.**
40. Michael J. Cafarella, Doug Downey, Oren Etzioni, Stephen Soderland. KnowItNow: Fast, Scalable Information Extraction from the Web. *Proceedings of the 2005 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2005.
41. Oren Etzioni, Michael Cafarella, Doug Downey, Ana-Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S. Weld, and Alexander Yates. Unsupervised Named-Entity Extraction from the Web: An Experimental Study. *Artificial Intelligence* 165:91-134, 2005.
42. Doug Downey, Oren Etzioni, Daniel S. Weld, and Stephen Soderland. Learning Text Patterns for Web Information Extraction and Assessment. *Proceedings of the AAI-04 Workshop on Adaptive Text Extraction and Mining*, 2004.
43. Oren Etzioni, Michael Cafarella, Doug Downey, Ana-Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S. Weld, and Alexander Yates. Methods for Domain-Independent Information Extraction from the Web: An Experimental Comparison. *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI)*, 2004.
44. Oren Etzioni, Michael Cafarella, Doug Downey, Stanley Kok, Ana-Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S. Weld, and Alexander Yates. Web-scale Information Extraction in KnowItAll. *Proceedings of the Thirteenth International World Wide Web Conference*

(WWW), 2004.

45. Andrew C. Slocum, Douglas C. Downey, Randall D. Beer. Further Experiments in the Evolution of Minimally Cognitive Behavior: From Perceiving Affordances to Selective Attention. *Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior*, 2000.

External
Funding

NSF CCF-1733860: AitF: Mechanism Design and Machine Learning for Peer Grading Total Award: \$700,000. NU Share: \$700,000. Role: Co-PI	2017-2021
Allen Institute for Artificial Intelligence: Parsing Scientific Metadata with Neural Networks Total Award: \$23,508. NU Share: \$23,508. Role: Sole PI	2017-2018
Tencent Corporation: Semantic Constraints in Conversational Text Generation Total Award: \$30,000. NU Share: \$30,000. Role: PI	2017-2018
NSF IIS-1065397: CAREER: Web Information Extraction: Integration and Scaling Total Award: \$550,000. NU Share: \$550,000. Role: Sole PI	2014-2017
Allen Institute for Artificial Intelligence: Open Named Entity Recognition from Scientific Documents Total Award: \$154,431. NU Share: \$154,431. Role: Sole PI	2011-2015
NSF IIS-1065397: Learning Representations of Language for Domain Adaptation Total Award: \$897,982. NU Share: \$200,000. Role: PI	2010-2014
DARPA D11AP00268: Web-scale Active Learning Total Award: \$429,293. NU Share: \$429,293. Role: Sole PI	2010-2013
NSF IIS-1016754, Active Learning of Language Models for Information Extraction Total Award: \$183,736. NU Share: \$183,739. Role: Sole PI	2011-2015
NSF CNS-1064595, Auditing Internet Content for Credibility, Fairness, and Privacy Total Award: \$826,000. NU Share: \$550,000. Role: Co-PI	2011-2015
Microsoft New Faculty Fellowship Total Award: \$200,000. NU Share: \$200,000. Role: Sole PI	2010