

Noah A. Smith

<http://www.cs.cmu.edu/~nasmith>

PROFESSIONAL EXPERIENCE

University of Washington	2015–	Associate Professor Computer Science and Engineering
Carnegie Mellon University	2011–2015	Associate Professor (tenure awarded 7/1/2014)
	2006–2011	Assistant Professor Language Technologies Institute, School of Computer Science
Microsoft Research	2004	} <i>Research internships</i>
New York University	2002	
Thomson Corporation	2000	
Johns Hopkins University	1999	
U.S. Defense Department	1996, 1997	

EDUCATION

Johns Hopkins University	2006	Ph.D. , Computer Science <i>Thesis topic:</i> unsupervised natural language parsing [144]
	2004	M.S. in Engineering , Computer Science
University of Maryland	2001	B.S. with High Honors in Computer Science
	2001	B.A. with Honors in Linguistics <i>Summa cum laude; honors theses:</i> [145–147]
University of Edinburgh	2000	Visiting student in Linguistics and Artificial Intelligence

SELECTED AWARDS

- Finmeccanica career development chair, endowed “to acknowledge promising teaching and research potential in junior faculty members” at CMU (2011–2014)
- Paper awards: ACL 2009 best paper [70], ICLP 2008 best student paper [77], SMT 2013 retrospective best paper [142], and various nominations [29, 38, 60, 84, 90]
- SAS/International Institute of Forecasters award (with Bryan R. Routledge; 2010)
- Institute for Quantitative Research in Finance (“Q Group”) research award (with Shimon Kogan, Bryan R. Routledge, and Jacob S. Sagi; 2008)
- Fannie and John Hertz Foundation fellowship (2001–2006)

SELECTED PRESS

- [The machines vs. Mitt Romney: How artificial intelligence is parsing political rhetoric](#). *Popular Science's* Zero Moment. September 24, 2014. [[35](#), [161](#)]
- [Addiction and seduction on Yelp: The language of food love](#). NPR's *All Things Considered*. April 16, 2014. [[6](#)]
- [The linguist's mother lode: What Twitter reveals about slang, gender, and no-nose emoticons](#). *Time*. September 9, 2013. [[58](#), [67](#), [121](#)]
- [Do you have a Twitter 'accent'?](#) NPR's *Here and Now*. September 4, 2013. [[121](#)]
- [Main tweet: Researchers dig into the intersection of politics and Twitter](#). *Time's* Swampblog. August 14, 2013. [[67](#), [98](#)]
- [Here's how you can use Twitter to beat the spread on sports betting](#). *Washington Post*, August 12, 2013. [[118](#)]
- [Congress' magic words](#). *Washington Post's* Wonkblog. December 31, 2012. [[45](#)]
- [Revealed: How China censors its social networks](#). *New Scientist*. March 8, 2012. [[13](#)]
- [Twitterology: A new science?](#) *New York Times*. October 30, 2011. [[58](#)]
- [Noah Smith on supercomputer trivia and Twitter dialects](#). CBC's *Spark*. February 15, 2011. [[58](#)]
- [You have an accent even on Twitter](#). NPR's *All Things Considered*. January 18, 2011. [[58](#)]
- [Twitter a decent stand-in for public opinion polls](#). *Ars Technica*. May 11, 2010. [[67](#)]

RESEARCH FUNDING

*Current sponsored research (*lead PI)*

DARPA	Machine learning for adaptable heterogeneous indexing and search (Memex program; PI: J. Schneider)	\$3.55M	2014–16
IARPA	Early model-based event recognition using surrogates (Open Source Indicators program; subcontract to Virginia Tech.)	\$218K	2014–15
NSF	Towards effective web privacy notice and choice: a multi-disciplinary perspective (PI: N. Sadeh; 1330596)	\$3.75M	2013–17
NSF	*Data-driven, computational models for discovery and analysis of framing (Socio-Computational Systems program; co-PIs A. Boydstun, J. H. Gross, P. Resnik; IIS-1211277)	\$750K	2012–15
DARPA	Structured distributed semantics: analysis and filtering of text (Deep Exploration and Filtering of Text program; PI: E. Hovy)	\$5.66M	2012–17
NSF	*CAREER: Flexible learning for natural language processing (IIS-1054319)	\$550K	2011–15
ARO	The linguistic-core approach to structured translation and analysis of low-resource languages (PI: J. Carbonell)	\$6.25M	2010–15

Industry gifts, internal, and other

Amazon	Social measurement from social media (Amazon Web Services credit)	\$30K	2013, 2014
Google	Utility-based language models (with B. R. Routledge)	\$76K	2013
Google	Reading is believing (with W. W. Cohen, T. M. Mitchell, and C. Faloutsos)	\$500K	2012
Amazon	Collecting and analyzing tweets for large-scale political, linguistic, and economic analysis (Amazon Web Services credit)	\$5K	2011
Pittsburgh Super-computing Center	Probabilistic models for recovering latent structure in natural language from text (Extreme Science and Engineering Discovery Environment program; DBS110003 and renewals)		2011–4
Google	Economic data: hard and soft (with B. R. Routledge)	\$68K	2011
CMU	Berkman faculty development award: Text-driven forecasting of mergers (with B. R. Routledge)	\$5K	2010
HP Labs	Understanding political discourse through probabilistic models (with W. W. Cohen)	\$25K	2009
Google	Growing translation resources by modeling Wikipedia	\$75K	2008
IBM	Robust, efficient, and integrable Arabic morphosyntactic disambiguation	\$30K	2007

*Past sponsored research (*lead PI)*

NSF	*An exploratory study on practical approaches for robust NLP tools with integrated annotation languages (co-PI C. Dyer; IIS-1352440)	\$100K	2013–14
NSF	*Big multilinguality for data-driven lexical semantics (BIGDATA program; co-PI C. Dyer; IIS-1251131)	\$250K	2013–14
A.P. Sloan	*Identifying corporate entities and relationships from text (co-PI B. R. Routledge)	\$124K	2013–14
IARPA	*Janus: from the past into the future (Open Source Indicators program; subcontract to Raytheon BBN)	\$233K	2012–14
U.S. Army	*Story creation and inference through Bayesian extraction (subcontract on STTR to Decisive Analytics Corporation)	\$30K	2011–12
DARPA	*Text-driven forecasting of voting behavior (N10AP20042)	\$248K	2010–11
IARPA	*Text-driven forecasting (co-PI B. R. Routledge; N10PC20222)	\$257K	2010–11
Qatar NRF	Improved Arabic natural language processing through semisupervised and cross-lingual learning (co-PI K. Oflazer; NPRP-08-485-1-083)	\$1.05M	2009–12
NSF	*Probabilistic models for structure discovery in text (IIS-0915187)	\$450K	2009–12
NSF	An integrated cluster computing architecture for machine translation (PI: S. Vogel; IIS-0844507)	\$465K	2009–11
NSF	*Scaling up unsupervised grammar induction (IIS-0836431)	\$213K	2008–9

DARPA	*Recombination, aggregation, and visualization of information in newsworthy expressions (co-PIs A. W. Black, R. Hwa, and F. L. Crabbe; NBCH-1080004)	\$500K	2008–10
NSF	*Parsing models and algorithms for morphologically rich languages (IIS-0713265)	\$112K	2007–8
DARPA	*Computer science study panel (phase 1; HR00110110013)	\$100K	2007–8

TEACHING EXPERIENCE

Tutorials

- [Structured Sparsity in Natural Language Processing: Models, Algorithms, and Applications](#)
Three-hour tutorial at EACL 2014 (with Mário Figueiredo, André Martins, and Dani Yogatama) and NAACL 2012 (with Mário Figueiredo and André Martins)
- [Probability and Structure in Natural Language Processing](#)
Invited course at
 - Universität Heidelberg, Germany (2014)
 - International Summer School in Language and Speech Technologies, Tarragona, Spain (2012)
 - IBM Thomas J. Watson Research Center (with Shay Cohen, 2011)
- [Sequence Models](#)
Three-hour lecture at the [Lisbon Machine Learning School](#) (2011, 2012, 2013, 2014)
- [Structured Prediction for Natural Language Processing](#)
Three-hour invited tutorial at ICML 2009

Graduate courses

- [Algorithms for Natural Language Processing](#) Fall 2011–13
Algorithms and formalisms used in NLP and CL (with Alon Lavie and Bob Frederking)
- [Structured Prediction for Language and Other Discrete Data](#) Fall 2011, 2013
Statistical structured prediction models (co-designer & co-instructor with William Cohen, Chris Dyer)
- [Probabilistic Graphical Models](#) Fall 2010
Theory and algorithms for probabilistic graphical models (instructor); Koller & Friedman textbook
- [Language and Statistics II](#) Fall 2006–9
Statistical learning for natural language processing (designer & instructor)

Advanced undergraduate course

- [Natural Language Processing](#) Spring 2008–11, 2013–4
Broad introduction to the field (designer & instructor); Jurafsky & Martin (2nd ed.) textbook
Co-taught with Chris Dyer in 2014

Graduate seminars and labs

- [Laboratory in Natural Language Processing](#) Spring 2013–4
- [Text-Driven Forecasting](#) Fall 2009
- [Advanced Natural Language Processing Seminar](#) Spring 2009–11, Fall 2012–14

Teaching prior to Carnegie Mellon

- [Empirical Research Methods in Computer Science](#) Fall 2005
Department of CS, JHU, short course for undergraduates, co-designed and co-taught with David Smith
- [Computational Genomics: Biological Sequence Modeling](#) Fall 2004
Department of CS, JHU, short course for undergraduates, co-designed and co-taught with Roy Tromble
- [Predicting English](#) Summer 2002–3
Center for Language and Speech Processing, JHU; four-hour laboratory exercise with competitive evaluation, co-designed with Jason Eisner [141], taught by others since 2004
- [Introduction to Programming for Linguists](#) (teaching assistant) Fall 2000
Department of Linguistics, University of Maryland

ADVISING

Ph.D., completed

- [Michael Heilman](#) (2008–2011); NSF graduate fellow, PIER scholar; [51, 65, 66, 104, 134, 135, 137, 140, 155, 170, 173, 178, 195]; research scientist at Educational Testing Service
- [Shay B. Cohen](#) (2006–2011); ICLP 2008 best student paper; Computing Innovation postdoctoral fellowship; [12, 14, 16, 47, 57, 62, 63, 72, 77, 78, 82, 108, 127, 139, 154, 204]; postdoctoral fellow at Columbia University, then lecturer (\approx assistant professor) at the University of Edinburgh
- [Dipanjan Das](#) (2008–2012); ACL 2011 best paper; [9, 42, 43, 47, 53, 61, 64, 69, 79, 104, 106, 133, 153, 173, 176, 191, 195, 198, 205]; research scientist at Google
- [André F. T. Martins](#)¹ (2007–2012); ICTI/Portugal scholar; ACL 2009 best paper; SCS Dissertation Award Honorable Mention; IBM Portugal *Prémio Científico*; [3, 9, 17, 42, 49, 50, 52, 56, 59, 70, 71, 79, 81, 130, 131, 138, 152, 174, 200, 205]; research scientist at Priberam
- [Kevin Gimpel](#) (2006–2012); Sandia-CMU graduate fellowship; [5, 40, 44, 48, 61, 68, 76, 78, 102, 104–106, 125, 126, 142, 151, 192, 195, 198, 201]; assistant research professor at TTI-Chicago
- [Tae Yano](#)² (2007–2013); [45, 75, 98, 107, 128, 136, 150, 193, 202]; engineer at Microsoft
- [Nathan Schneider](#) (2008–2014); [7, 9, 25, 32, 39, 46, 64, 99, 101, 104, 117, 119, 120, 133, 149, 171, 173, 176, 184, 185, 190, 194, 195]; postdoctoral researcher at the University of Edinburgh
- [Brendan O'Connor](#) (2009–2014); [2, 13, 36, 37, 39, 51, 58, 67, 104, 117, 119, 121, 124, 132, 148, 170, 186, 187, 195]; assistant professor at the University of Massachusetts

Ph.D., post-thesis proposal

- [Dani Yogatama](#) (2010–); [8, 30, 31, 41, 51, 98, 104, 116, 163, 170, 195]
- [David Bamman](#) (2011–); [13, 28, 36, 94, 117, 119, 124, 165, 166, 187]
- [Waleed Ammar](#)³ (2011–); [24, 122, 167]

Ph.D., pre-thesis proposal

- [Yanchuan Sim](#) (2011–); [21, 23, 34, 35, 41, 123, 164]
- [Sam Thomson](#) (2012–) [29, 117, 183]

¹Co-advisors: Eric Xing (CMU), Mário Figueiredo and Pedro Aguiar (IST, Universidade Técnica de Lisboa).

²Co-advisor: William Cohen (CMU).

³Co-advisor: Chris Dyer (CMU).

- Swabha Swayamdipta³ (2013–) [25, 117]
- Dallas Card (2013–) [161]
- Jeffrey Flanigan⁴ (2014–); [29, 104, 117, 183, 195]

Post-doctoral

- Chris Dyer (2010–2012); [11, 51, 54, 103, 122, 125, 127, 169, 170]; assistant professor at CMU
- Behrang Mohit (2010–2012), co-supervised by Kemal Oflazer, CMU–Qatar; [46, 99, 101, 171, 190, 194]
- Fei Liu (2013–), co-advised by Norman Sadeh; [26, 95]

M.S.

- Cari (Sisson) Bader (2007–8); engineer at Nuance
- Daniel Mills (2010–11); [104, 195] (on leave)
- Victor Chahuneau³ (2012–3) [6, 11, 33, 38, 40, 100, 188, 189]; pursuing a Ph.D. at Stanford
- Lingpeng Kong (2013–) [25, 162]
- Rohan Ramanath⁵ (2013–) [26, 95]
- Jesse Dodge (2013–) [115, 117]

Graduate student visitors and independent study projects

2014: Zita Marinho; 2013: Daniel Preoțiu-Pietro (Sheffield University), Lingpeng Kong; Minghui Qiu (Singapore Management University) [21, 34]; 2012: Swapna Gottipati (Singapore Management University) [34]; 2011: Dong Nguyen [129]; 2008: Aaron Phillips, Narges Sharif-Razavian, Sourish Chaudhuri [109], Severin Hacker, ThuyLinh Nguyen [60]; 2007: Daniel Rashid [111], Mengqiu Wang [84]; 2006: Thuy Linh Nguyen

Ph.D. thesis committees

Ongoing: Adam Anderson (Harvard) [166], Justin Betteridge, Jonathan Clark [54, 103], Qin Gao, Dan Garrette (U. Texas) [22, 27], Greg Hanneman, Jayant Krishnamurthy, Wang Ling, Thuylinh Nguyen [60], Ankur Parikh; 2014: Anil Nelakanti (Université Pierre et Marie Curie); 2013: Khalid El-Arini, Mahesh Joshi [106], Ramnath Balasubramanyan [67, 110], Vladimir Eidelman (U. Maryland); 2012: Aaron Phillips, Sanjika Hewavitharana; 2011: Ming-Wei Chang (UIUC), David Huggins-Daines, Andreas Zollmann [74, 80]; 2010: Andrew Carlson; 2009: Andrew Arnold; 2008: Ashish Venugopal [74, 80], Ying Joy Zhang

Undergraduate research advisees

Rishav Bhowmick (CMU–Qatar) [46, 157], Desai Chen [9, 64, 127, 133, 156, 176],⁶ Arash Enayati (CMU–Qatar), Mohammad Haque, Dimitry Levin [73], Zack McCord, Tobi Owoputi [39], Benjamin Plaut, Naomi Saphra [119],⁷ Neel Shah [158], Shiladitya Sinha [118, 186], Tal Stramer [172], Matthew Thompson [159], Daniel Tasse [179, 206],⁸ Xiaote Zhu

⁴Co-advisors: Chris Dyer (CMU) and Jaime Carbonell (CMU).

⁵Co-advisor: Norman Sadeh (CMU).

⁶Honorable mention, CRA Outstanding Undergraduate Research Award, 2010. Pursuing a Ph.D. at MIT.

⁷Pursuing a Ph.D. at JHU.

⁸Pursuing a Ph.D. at CMU.

High school research advisee

Lily Scherlis [40]

Full-time technical staff (CMU B.S. graduates)

- Zach Paine (2008); later at Google, Apple
- Philip Gianfortoni (2009–10); entered M.S. in Language Technologies at CMU in 2010
- Bill McDowell (2013–4)
- Michael Mordowanec (2014) [96, 182]

SERVICE

Professional organizations

- Secretary-Treasurer of [SIGDAT](#), the Association for Computational Linguistics Special Interest Group on Linguistic Data and Corpus-Based Approaches to Natural Language Processing (2012–15)

Journals

- Associate editor, *Journal of Artificial Intelligence Research* (2014–17)
- Editorial board, *Transactions of the Association for Computational Linguistics* (2012–)
- Editorial board, *Journal of Artificial Intelligence Research* (2011–14)
- Editorial board, *Computational Linguistics* (2009–11)
- Reviewing: *Artificial Intelligence Journal* (2009), *Journal of Machine Learning Research* (2010, 2008, 2004), *Language and Computation* (2008), *IEEE Intelligent Systems* (2008), *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2011), *Journal of Applied Mathematics and Computer Science* (2014), *Journal of Information Technology and Politics* (2007), *Proceedings of the National Academy of Sciences* (2010, 2007), *Computational Linguistics* (2005), *Language Resources and Evaluation* (formerly *Computers and the Humanities*; 2004)

Conferences

- “Tagging, chunking, syntax, and parsing” area co-chair, NAACL 2015
- Co-organizer (with Claire Cardie, Anne Washington, and John Wilkerson), [NLP Unshared Task in Polinformatics 2014](#), a research competition culminating at the ACL 2014 workshop below [160]
- Co-organizer (with Cristian Danescu-Niculescu-Mizil, Jacob Eisenstein, and Kathy McKeown), [Workshop on Language Technologies and Computational Social Science](#) at ACL 2014
- Co-organizer (with Phil Blunsom and Chris Dyer), [Workshop on Twenty Years of Bitext](#) at EMNLP 2013
- “Social media” area co-chair, ACL 2012
- “Machine learning” area chair, EMNLP 2010
- “Parsing and syntax” area chair, NAACL-HLT 2010
- Workshops co-chair, COLING 2010
- “Parsing and syntax” area co-chair, ACL-IJCNLP 2009
- Student travel awards chair, ICML 2008
- Publications co-chair, ACL-HLT 2008

- Reviewing:

Natural language processing: ACL 2014, ACL 2013, ACL 2011, ACL 2010, ACL 2007, COLING-ACL 2006; COLING 2008; CoNLL 2011, CoNLL 2009, CoNLL 2006; EACL 2014, EACL 2009; EMNLP 2013, EMNLP-CoNLL 2012, EMNLP 2011, EMNLP 2009, EMNLP 2008, EMNLP-CoNLL 2007, EMNLP 2006; HLT-NAACL 2009, HLT-NAACL 2007; IJCNLP 2008; IWPT 2011; NAACL 2013
Machine learning and artificial intelligence: ICML 2013, ICML 2012, ICML 2011, ICML 2008, ICML 2004; IJCAI 2005; NIPS 2010, NIPS 2009, NIPS 2008, NIPS 2007; UAI 2009
Other: WWW 2014, WWW 2011

Funding agencies (reviewing)

- National Science Foundation (2012, 2011, 2010, 2009, 2008, 2007, 2006)
- Qatar National Research Foundation (2012, 2011)
- Fundação para a Ciência e a Tecnologia, Portugal (2008)

University, School, and Departments

- LTI open house organizer (2014)
- LTI MIIS program advisory board (2013–)
- Organized a university-wide seminar series, [Machine Learning and the Social Sciences](#) (2013–14)
- Co-organized a university-wide research workshop, [Machine Learning for the Social Sciences](#) (with George Loewenstein; 2012)
- LTI faculty retreat organizer (2012)
- LTI faculty search committee chair (2012)
- MLD faculty search committee (2011, 2012)
- Faculty liaison to LTI student body (2010–14)
- CMU undergraduate research grant/fellowship selection committee (2010–14)
- SCS Intelligence Seminar organizer (2008–10)
- LTI curriculum committee (2009–14)
- LTI graduate admissions committee (2007–14)
- LTI student research symposium judge (2008, 2007)

PRESENTATIONS

Peer-refereed, non-archival conferences

- American Political Science Association, Washington, DC, 8/2014 (given by Amber Boydstun) [[161](#)]
- Annual Meeting of the Comparative Agendas Project, Konstanz, Germany, 6/2014 (given by Amber Boydstun) [[161](#)]
- American Political Science Association, Chicago, IL, 8/2013 (given by Justin Gross) [[164](#)]
- Digital Humanities, Omaha, NE, 7/2013 (given by David Bamman) [[166](#)]
- Linguistic Society of America, Pittsburgh, PA, 1/2011 (given by Jacob Eisenstein) [[58](#)]
- Association for Public Policy Analysis and Management, Washington, DC, 11/1999 [[181](#)]

Invited talks at workshops, research meetings, and other events

- [Atlanta Computational Social Science Workshop](#), Atlanta, GA, 11/2014 [23, 35]
- [Computational Linguistics in Political Science: What Have You Done for Me Lately?](#), Mannheim, Germany, 10–11/2014 [interdisciplinary research]
- [New Directions in Analyzing Text as Data Conference](#), Northwestern University, 10/2014 (given by Yanchuan Sim) [23]
- [ACL 2014 Workshop on Interactive Language Learning, Visualization, and Interfaces](#), Baltimore, MD, 6/2014 [interdisciplinary research]
- [Carnegie Mellon Center for Innovation and Entrepreneurship: Launch|CMU](#), Mountain View, CA, 5/2013 [13, 45, 51, 58, 67, 73, 98, 121]
- [National Science Foundation CISE CAREER workshop](#), Arlington, TX, 5/2013
- [IARPA FUSE PI meeting](#), College Park, MD, 4/2013 [51]
- [AAAI Spring Symposium on Analyzing Microtext](#), Stanford, CA, 3/2013 [13, 39, 58, 67, 98, 104, 121]
- [World Economic Forum](#), Davos, Switzerland, 1/2013 [45, 58, 67, 73, 106, 121]
- [New Directions in Analyzing Text as Data Conference](#), Harvard University, 10/2012 (given by Tae Yano) [45]
- [Insight 3.0: The Web Seen by its Insiders](#) 7/2012, [panel on the technology industry]
- [Workshop on Multilingual Modeling at ACL](#), 7/2012 [47]
- [Lisbon Machine Learning School](#) evening lecture, 7/2012 [13, 45, 51]
- [Inducing Linguistic Structure Workshop at NAACL](#), 6/2012 [168]
- [Computer Assisted Reporting Conference](#), 2/2012 [13, 51, 58, 67, 106]
- [Tech@State: Real-Time Awareness](#), U.S. State Department, 2/2012 [13, 51, 58, 67, 73, 106]
- [Text to Text Generation Workshop at ACL](#), 6/2011 [65, 66, 69, 84, 138, 155]
- [New Directions in Text Analysis Conference](#), Harvard University, 5/2011 [55, 58, 129]
- [South by Southwest: Interactive](#), Austin, 3/2011 [panel with Philip Resnik: “Using text to predict the real world”]
- [Tracking, Transcribing, and Tagging Government: Building Digital Records for Computational Social Science](#), Center for Advanced Study in the Behavioral Sciences, Stanford University, 6/2010 [67, 73, 75, 106, 107]
- [NAACL-HLT Workshop on #SocialMedia: Computational Linguistics in a World of Social Media](#), 6/2010 [67, 73, 75, 106, 107]
- [NIPS 2009 Workshop on Approximate Learning of Large Scale Graphical Models](#), 12/2009 [68, 70, 71, 76]
- [Hadoop Summit](#) (given by Kevin Gimpel), 6/2009 [“Natural language learning with Hadoop”]
- [NAACL-HLT Workshop on Integer Linear Programming for NLP](#) (with André Martins), 6/2009 [70, 71]
- [ACL Workshop on Mobile Language Processing](#), 6/2008 [panel on mobile NLP]
- [ACL Workshop on Issues in Teaching Computational Linguistics](#), 6/2008 [panel on NLP/CL curriculum]
- [Tokyo Forum on Advanced NLP and Text Mining](#), University of Tokyo, 2/2008 [84, 85]
- CMU “Andrew’s Leap” computer science high school outreach program talk, 7/2007
- DARPA ISAT study group on Engineering Ensemble Effects, 4/2007 [89]
- [HLT-NAACL Doctoral Consortium](#), 6/2006 [panel on the academic/industrial job market]

Invited talks at academic, industry, and government research colloquia

- Universität Heidelberg, 11/2014 [4, 23, 35]
- Max-Planck-Institut SWS, 11/2014 [4, 23, 35]
- University of Washington, 4/2014 [31, 35]
- Jump Trading, 3/2014
- Department of Defense, 7/2013 [13, 45, 51]
- eBay Research Labs, 5/2013 [13, 45, 51]
- Cornell University, 3/2013 [13, 45, 51]
- University of Pennsylvania, 12/2012 [13, 45, 51]
- Universitat Politècnica de Catalunya, 7/2012 [13, 45, 51]
- Korea Advanced Institute of Science and Technology, 7/2012 [13, 45, 51, 58]
- Microsoft Research, 5/2012 [13, 45, 51, 58]
- University of Washington, 5/2012 [42, 49, 52]
- University of Maryland, 3/2012 [42, 49, 52]
- Johns Hopkins University, 3/2012 [42, 49, 52]
- Carnegie Mellon University, Language Technologies Institute 25th Anniversary, 10/2011 [13, 51, 55, 58, 67, 73, 75, 106, 107]
- In-Q-Tel, 10/2011 [51, 67, 73, 106, 107]
- IBM T. J. Watson Research Center, 5/2011 [55, 58, 129]
- Toyota Technological Institute at Chicago, 3/2011 [58, 67, 106, 107, 136]
- BlackRock, 3/2011 [58, 67, 106, 107, 136]
- University of North Carolina at Chapel Hill, Political Science Department, 2/2011 [58, 67, 106, 107, 136]
- Carnegie Mellon University, Language Technologies Institute, 9/2010 [67, 73, 75, 106, 107]
- Twitter, 7/2010 [58, 67, 73, 75, 106, 107]
- Facebook, 6/2010 [58, 67, 73, 75, 106, 107]
- HP Labs, Social Computing Group, 6/2010 [58, 67, 73, 75, 106, 107]
- Stanford University, 6/2010 [58, 67, 73, 106, 107]
- University of California Berkeley, 6/2010, [58, 67, 73, 106, 107]
- University of Illinois at Urbana-Champaign, Department of Computer Science, 5/2010 [67, 73, 75, 106, 107]
- University of Texas at Austin, Department of Computer Science, 4/2010 [67, 73, 75, 106, 107]
- Universidade Técnica de Lisboa, Instituto Superior Técnico, 3/2010 [67, 73, 75, 106, 107]
- University of Edinburgh, Institute for Communicating and Collaborative Systems, 3/2010 [67, 73, 75, 106, 107]
- University of Sheffield, Department of Computer Science, 3/2010 [67, 73, 75, 106, 107]
- University of Maryland, Department of Computer Science, 11/2009 [70, 72, 73, 78, 79, 86, 90]
- Princeton University, Department of Computer Science, 10/2009 [72, 78, 86, 90]
- University of Massachusetts at Amherst, Department of Computer Science, 3/2009 [72, 78, 86, 90]
- Universidade Técnica de Lisboa, Instituto Superior Técnico, 7/2008 [84, 142, 144]
- IBM T. J. Watson Research Center, 2/2008 [82, 84, 85]
- Carnegie Mellon University, Language Technologies Institute, 12/2007 [82, 84, 85]
- Department of Defense, 8/2007 [research overview]
- Google, 4/2006 [job talk]

- Microsoft Research, 4/2006 [job talk]
- Stanford University, Department of Computer Science, 4/2006 [job talk]
- University of Illinois at Urbana-Champaign, Department of Computer Science, 3/2006 [job talk]
- University of Wisconsin–Madison, Department of Computer Sciences, 3/2006 [job talk]
- Carnegie Mellon University, Language Technologies Institute, 3/2006 [job talk]
- University of Maryland, Institute for Advanced Computer Study, 12/2005 [90]
- Brown University, Computer Science Department, 10/2005 [87]
- Carnegie Mellon University, Language Technologies Institute, 9/2005 [87]
- Carnegie Mellon University, Center for Automated Learning and Discovery, 9/2005 [90]
- University of Pittsburgh, Computer Science Department, 9/2005 [90]
- Microsoft Research, 8/2004 [91]
- Microsoft Research, 7/2004 [92]
- University of Maryland, Institute for Advanced Computer Study, 4/2002 [93]

PUBLICATIONS

Book

- [1] Noah A. Smith. *Linguistic Structure Prediction*. Synthesis Lectures on Human Language Technologies. Morgan and Claypool, May 2011.

Journal articles and book chapters

- [2] Jacob Eisenstein, Brendan O'Connor, Noah A. Smith, and Eric P. Xing. [Diffusion of language change in social media](#). *PLoS ONE*, November 2014.
- [3] André F. T. Martins, Mário A. T. Figueiredo, Pedro M. Q. Aguiar, Noah A. Smith, and Eric P. Xing. [AD³: Alternating directions dual decomposition for MAP inference in graphical models](#). *Journal of Machine Learning Research*, 2015.
- [4] David Bamman and Noah A. Smith. [Unsupervised discovery of biographical structure from text](#). *Transactions of the Association for Computational Linguistics*, 2(2014):363–376, October 2014.
- [5] Kevin Gimpel and Noah A. Smith. [Phrase dependency machine translation with quasi-synchronous tree-to-tree features](#). *Computational Linguistics*, 40(2), June 2014.
- [6] Dan Jurafsky, Victor Chahuneau, Bryan R. Routledge, and Noah A. Smith. [Narrative framing of consumer sentiment in online restaurant reviews](#). *First Monday*, 19(4), April 2014.
- [7] Nathan Schneider, Emily Danchik, Chris Dyer, and Noah A. Smith. [Discriminative lexical semantic segmentation with gaps: Running the MWE gamut](#). *Transactions of the Association for Computational Linguistics*, 2:193–206, April 2014.
- [8] Dani Yogatama, Chong Wang, Bryan R. Routledge, Noah A. Smith, and Eric P. Xing. [Dynamic models of streaming text](#). *Transactions of the Association for Computational Linguistics*, 2:181–192, April 2014.
- [9] Dipanjan Das, Desai Chen, André F. T. Martins, Nathan Schneider, and Noah A. Smith. [Frame-semantic parsing](#). *Computational Linguistics*, 40(1):9–56, March 2014.
- [10] Noah A. Smith and André F. T. Martins. [Linguistic structure prediction with the sparseptron](#). *ACM Crossroads*, 19(3):44–48, April 2013.
- [11] Victor Chahuneau, Noah A. Smith, and Chris Dyer. [pycdec: A Python interface to cdec](#). *Prague Bulletin of Mathematical Linguistics*, 98:51–61, October 2012.
- [12] Shay B. Cohen and Noah A. Smith. [Empirical risk minimization for probabilistic grammars: Sample complexity and hardness of learning](#). *Computational Linguistics*, 38(3), September 2012.
- [13] David Bamman, Brendan O'Connor, and Noah A. Smith. [Censorship and content deletion in Chinese social media](#). *First Monday*, 17(3), March 2012.
- [14] Shay B. Cohen, Robert J. Simmons, and Noah A. Smith. [Products of weighted logic programs](#). *Theory and Practice of Logic Programming*, 11(2–3):263–296, January 2011.
- [15] Jason Eisner and Noah A. Smith. [Favor short dependencies: Parsing with soft and hard constraints on dependency length](#). In Harry Bunt, Paola Merlo, and Joakim Nivre, editors, *Trends in Parsing Technology: Dependency Parsing, Domain Adaptation, and Deep Parsing*, volume 43 of *Text, Speech, and Language Technology*, chapter 8, pages 121–150. Springer, January 2011.
- [16] Shay B. Cohen and Noah A. Smith. [Covariance in unsupervised learning of probabilistic grammars](#). *Journal of Machine Learning Research*, 11:3017–3051, November 2010.
- [17] André F. T. Martins, Noah A. Smith, Eric P. Xing, Mário A. T. Figueiredo, and Pedro M. Q. Aguiar. [Nonextensive information theoretic kernels on measures](#). *Journal of Machine Learning Research*, 10:935–975, April 2009.

- [18] Noah A. Smith. [Review of Computational Approaches to Morphology and Syntax](#) by Brian Roark and Richard Sproat. *Computational Linguistics*, 34(3):453–457, September 2008.
- [19] Noah A. Smith and Mark Johnson. [Weighted and probabilistic context-free grammars are equally expressive](#). *Computational Linguistics*, 33(4):477–491, December 2007.
- [20] Philip Resnik and Noah A. Smith. [The Web as a parallel corpus](#). *Computational Linguistics*, 29(3):349–380, September 2003.

Refereed conference publications (full-length)

- [21] Minghui Qiu, Yanchuan Sim, Noah A. Smith, and Jing Jiang. Modeling user arguments, interactions, and attributes for stance prediction in online debate forums. In *Proceedings of the SIAM Conference on Data Mining*, Vancouver, BC, April–May 2015. **SDM 2015.**
- [22] Dan Garrette, Chris Dyer, Jason Baldridge, and Noah A. Smith. [Weakly-supervised grammar-informed Bayesian CCG parser learning](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, Austin, TX, January 2015. **AAAI 2015.**
- [23] Yanchuan Sim, Bryan R. Routledge, and Noah A. Smith. The utility of text: The case of amicus briefs and the Supreme Court. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Austin, TX, January 2015. **AAAI 2015.**
- [24] Waleed Ammar, Chris Dyer, and Noah A. Smith. [Conditional random field autoencoders for unsupervised structured prediction](#). In *Advances in Neural Information Processing Systems* 27, Montréal, Quebec, December 2014. **NIPS 2014.**
Selected for oral presentation (top 5% of accepted papers).
- [25] Lingpeng Kong, Nathan Schneider, Swabha Swayamdipta, Archana Bhatia, Chris Dyer, and Noah A. Smith. [A dependency parser for tweets](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Doha, Qatar, October 2014. **EMNLP 2014.**
- [26] Fei Liu, Rohan Ramanath, Norman Sadeh, and Noah A. Smith. [A step towards usable privacy policy: Automatic alignment of privacy statements](#). In *Proceedings of the International Conference on Computational Linguistics*, Dublin, Ireland, August 2014. **COLING 2014.**
- [27] Dan Garrette, Chris Dyer, Jason Baldridge, and Noah A. Smith. [Weakly-supervised Bayesian learning of a CCG supertagger](#). In *Proceedings of the Conference on Computational Natural Language Learning*, Baltimore, MD, June 2014. **CoNLL 2014.**
- [28] David Bamman, Ted Underwood, and Noah A. Smith. [A Bayesian mixed effects model of literary character](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Baltimore, MD, June 2014. **ACL 2014.**
- [29] Jeffrey Flanigan, Sam Thomson, Jaime Carbonell, Chris Dyer, and Noah A. Smith. [A discriminative graph-based parser for the abstract meaning representation](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Baltimore, MD, June 2014. **ACL 2014.**
Nominated for best paper award.
- [30] Dani Yogatama and Noah A. Smith. [Linguistic structured sparsity in text categorization](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Baltimore, MD, June 2014. **ACL 2014.**
- [31] Dani Yogatama and Noah A. Smith. [Making the most of bag of words: Sentence regularization with alternating direction method of multipliers](#). In *Proceedings of the International Conference on Machine Learning*, Beijing, China, June 2014. **ICML 2014.**
- [32] Nathan Schneider, Spencer Onuffer, Nora Kazour, Emily Danchik, Michael T. Mordowanec, Henrietta Conrad, and Noah A. Smith. [Comprehensive annotation of multiword expressions in a social web](#)

- corpus. In *Proceedings of the Language Resources and Evaluation Conference*, Reykjavik, Iceland, May 2014. **LREC 2014.**
- [33] Victor Chahuneau, Eva Schlinger, Chris Dyer, and Noah A. Smith. [Translating into morphologically rich languages with synthetic phrases](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Seattle, WA, October 2013. **EMNLP 2013.**
- [34] Swapna Gottipati, Minghui Qiu, Yanchuan Sim, Jing Jiang, and Noah A. Smith. [Learning topics and positions from Debatepedia](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Seattle, WA, October 2013. **EMNLP 2013.**
- [35] Yanchuan Sim, Brice D. L. Acree, Justin H. Gross, and Noah A. Smith. [Measuring ideological proportions in political speeches](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Seattle, WA, October 2013. **EMNLP 2013.**
- [36] David Bamman, Brendan O’Connor, and Noah A. Smith. [Learning latent personas of film characters](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Sofia, Bulgaria, August 2013. **ACL 2013.**
- [37] Brendan O’Connor, Brandon Stewart, and Noah A. Smith. [Learning to extract international relations from political context](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Sofia, Bulgaria, August 2013. **ACL 2013.**
- [38] Victor Chahuneau, Noah A. Smith, and Chris Dyer. [Knowledge-rich morphological priors for Bayesian language models](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, Atlanta, GA, June 2013. **NAACL 2013.**
Nominated for best paper award.
- [39] Olutobi Owoputi, Brendan O’Connor, Chris Dyer, Kevin Gimpel, Nathan Schneider, and Noah A. Smith. [Improved part-of-speech tagging for online conversational text with word clusters](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, Atlanta, GA, June 2013. **NAACL 2013.**
- [40] Victor Chahuneau, Kevin Gimpel, Bryan R. Routledge, Lily Scherlis, and Noah A. Smith. [Word salad: Relating food prices and descriptions](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and Natural Language Learning*, Jeju, Korea, July 2012. **EMNLP 2012.**
- [41] Dani Yogatama, Yanchuan Sim, and Noah A. Smith. [A probabilistic model for canonicalizing named entity mentions](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Jeju, Korea, July 2012. **ACL 2012.**
- [42] Dipanjan Das, André F. T. Martins, and Noah A. Smith. [An exact dual decomposition algorithm for shallow semantic parsing with constraints](#). In *Proceedings of the Joint Conference on Lexical and Computational Semantics*, Montréal, Québec, June 2012. ***SEM 2012.**
- [43] Dipanjan Das and Noah A. Smith. [Graph-based lexicon expansion with sparsity-inducing penalties](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, Montréal, Québec, June 2012. **NAACL 2012.**
- [44] Kevin Gimpel and Noah A. Smith. [Structured ramp loss minimization for machine translation](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, Montréal, Québec, June 2012. **NAACL 2012.**
- [45] Tae Yano, Noah A. Smith, and John D. Wilkerson. [Textual predictors of bill survival in Congressional committees](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, pages 793–802, Montréal, Québec, June 2012. **NAACL 2012.**
- [46] Behrang Mohit, Nathan Schneider, Rishav Bhowmick, Kemal Oflazer, and Noah A. Smith. [Recall-oriented learning of named entities in Arabic Wikipedia](#). In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics*, Avignon, France, April 2012. **EACL 2012.**

- [47] Shay B. Cohen, Dipanjan Das, and Noah A. Smith. [Unsupervised structure prediction with non-parallel multilingual guidance](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Edinburgh, UK, July 2011. **EMNLP 2011.**
- [48] Kevin Gimpel and Noah A. Smith. [Quasi-synchronous phrase dependency grammars for machine translation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Edinburgh, UK, July 2011. **EMNLP 2011.**
- [49] André F. T. Martins, Noah A. Smith, Pedro M. Q. Aguiar, and Mário A. T. Figueiredo. [Dual decomposition with many overlapping components](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Edinburgh, UK, July 2011. **EMNLP 2011.**
- [50] André F. T. Martins, Noah A. Smith, Pedro M. Q. Aguiar, and Mário A. T. Figueiredo. [Structured sparsity in structured prediction](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Edinburgh, UK, July 2011. **EMNLP 2011.**
- [51] Dani Yogatama, Michael Heilman, Brendan O'Connor, Chris Dyer, Bryan R. Routledge, and Noah A. Smith. [Predicting a scientific community's response to an article](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Edinburgh, UK, July 2011. **EMNLP 2011.**
- [52] André F. T. Martins, Pedro M. Q. Aguiar, Mário A. T. Figueiredo, Noah A. Smith, and Eric P. Xing. [An augmented Lagrangian approach to constrained MAP inference](#). In *Proceedings of the International Conference on Machine Learning*, Bellevue, WA, June/July 2011. **ICML 2011.**
- [53] Dipanjan Das and Noah A. Smith. [Semi-supervised frame-semantic parsing for unknown predicates](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Portland, OR, June 2011. **ACL 2011.**
- [54] Chris Dyer, Jonathan H. Clark, Alon Lavie, and Noah A. Smith. [Unsupervised word alignment with arbitrary features](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Portland, OR, June 2011. **ACL 2011.**
- [55] Jacob Eisenstein, Noah A. Smith, and Eric P. Xing. [Discovering sociolinguistic associations with structured sparsity](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Portland, OR, June 2011. **ACL 2011.**
- [56] André F. T. Martins, Noah A. Smith, Eric P. Xing, Pedro M. Q. Aguiar, and Mário A. T. Figueiredo. [Online learning of structured predictors with multiple kernels](#). In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, Fort Lauderdale, FL, April 2011. **AISTATS 2011.**
- [57] Shay B. Cohen and Noah A. Smith. [Empirical risk minimization with approximations of probabilistic grammars](#). In *Advances in Neural Information Processing Systems 23*, Vancouver, BC, December 2010. **NIPS 2010.**
- [58] Jacob Eisenstein, Brendan O'Connor, Noah A. Smith, and Eric P. Xing. [A latent variable model for geographic lexical variation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Cambridge, MA, October 2010. **EMNLP 2010.**
- [59] André F. T. Martins, Noah A. Smith, Eric P. Xing, Pedro M. Q. Aguiar, and Mário A. T. Figueiredo. [Turbo parsers: Dependency parsing by approximate variational inference](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Cambridge, MA, October 2010. **EMNLP 2010.**
- [60] ThuyLinh Nguyen, Stephan Vogel, and Noah A. Smith. [Nonparametric word segmentation for machine translation](#). In *Proceedings of the International Conference on Computational Linguistics*, Beijing, China, August 2010. **COLING 2010. Best paper finalist.**
- [61] Kevin Gimpel, Dipanjan Das, and Noah A. Smith. [Distributed asynchronous online learning for natural language processing](#). In *Proceedings of the Conference on Computational Natural Language Learning*, Uppsala, Sweden, July 2010. **CoNLL 2010.**

- [62] Shay B. Cohen and Noah A. Smith. [Viterbi training for PCFGs: Hardness results and competitiveness of uniform initialization](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 1502–1511, Uppsala, Sweden, July 2010. **ACL 2010**.
- [63] Shay B. Cohen, David M. Blei, and Noah A. Smith. [Variational inference for adaptor grammars](#). In *Proceedings of the North American Chapter of the Association for Computational Linguistics Human Language Technologies Conference*, Los Angeles, CA, June 2010. **NAACL 2010**.
- [64] Dipanjan Das, Nathan Schneider, Desai Chen, and Noah A. Smith. [Probabilistic frame-semantic parsing](#). In *Proceedings of the North American Chapter of the Association for Computational Linguistics Human Language Technologies Conference*, Los Angeles, CA, June 2010. **NAACL 2010**.
- [65] Michael Heilman and Noah A. Smith. [Good question! statistical ranking for question generation](#). In *Proceedings of the North American Chapter of the Association for Computational Linguistics Human Language Technologies Conference*, Los Angeles, CA, June 2010. **NAACL 2010**.
- [66] Michael Heilman and Noah A. Smith. [Tree edit models for recognizing textual entailments, paraphrases, and answers to questions](#). In *Proceedings of the North American Chapter of the Association for Computational Linguistics Human Language Technologies Conference*, Los Angeles, CA, June 2010. **NAACL 2010**.
- [67] Brendan O’Connor, Ramnath Balasubramanyan, Bryan R. Routledge, and Noah A. Smith. [From tweets to polls: Linking text sentiment to public opinion time series](#). In *Proceedings of the International AAAI Conference on Weblogs and Social Media*, pages 122–129, Washington, DC, May 2010. **ICWSM 2010**.
- [68] Kevin Gimpel and Noah A. Smith. [Feature-rich translation by quasi-synchronous lattice parsing](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 219–228, Singapore, August 2009. **EMNLP 2009**.
- [69] Dipanjan Das and Noah A. Smith. [Paraphrase identification as probabilistic quasi-synchronous recognition](#). In *Proceedings of the Joint Conference of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, pages 468–476, Singapore, August 2009. **ACL 2009**.
- [70] André F. T. Martins, Noah A. Smith, and Eric P. Xing. [Concise integer linear programming formulations for dependency parsing](#). In *Proceedings of the Joint Conference of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, pages 342–350, Singapore, August 2009. **ACL 2009**. **Best paper award**.
- [71] André F. T. Martins, Noah A. Smith, and Eric P. Xing. [Polyhedral outer approximations with application to natural language parsing](#). In *Proceedings of the International Conference on Machine Learning*, pages 713–720, Montréal, Québec, June 2009. **ICML 2009**.
- [72] Shay B. Cohen and Noah A. Smith. [Shared logistic normal distributions for soft parameter tying in unsupervised grammar induction](#). In *Proceedings of the North American Association for Computational Linguistics Human Language Technologies Conference*, pages 74–82, Boulder, CO, May/June 2009. **NAACL 2009**.
- [73] Shimon Kogan, Dmitry Levin, Bryan R. Routledge, Jacob S. Sagi, and Noah A. Smith. [Predicting risk from financial reports with regression](#). In *Proceedings of the North American Association for Computational Linguistics Human Language Technologies Conference*, pages 272–280, Boulder, CO, May/June 2009. **NAACL 2009**.
- [74] Ashish Venugopal, Andreas Zollmann, Noah A. Smith, and Stephan Vogel. [Preference grammars: Softening syntactic constraints to improve statistical machine translation](#). In *Proceedings of the North American Association for Computational Linguistics Human Language Technologies Conference*, pages 236–244, Boulder, CO, May/June 2009. **NAACL 2009**.

- [75] Tae Yano, William W. Cohen, and Noah A. Smith. [Predicting response to political blog posts with topic models](#). In *Proceedings of the North American Association for Computational Linguistics Human Language Technologies Conference*, pages 477–485, Boulder, CO, May/June 2009. **NAACL 2009**.
- [76] Kevin Gimpel and Noah A. Smith. [Cube summing, approximate inference with non-local features, and dynamic programming without semirings](#). In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics*, pages 157–166, Athens, Greece, March/April 2009. **EACL 2009**.
- [77] Shay B. Cohen, Robert J. Simmons, and Noah A. Smith. [Dynamic programming algorithms as products of weighted logic programs](#). In *Proceedings of the International Conference on Logic Programming*, Udine, Italy, December 2008. **ICLP 2008**.
Best student paper award.
- [78] Shay B. Cohen, Kevin Gimpel, and Noah A. Smith. [Logistic normal priors for unsupervised probabilistic grammar induction](#). In *Advances in Neural Information Processing Systems 21*, pages 321–328, Vancouver, BC, December 2008. **NIPS 2008**.
- [79] André F. T. Martins, Dipanjan Das, Noah A. Smith, and Eric P. Xing. [Stacking dependency parsers](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 157–166, Waikiki, HI, October 2008. **EMNLP 2008**.
- [80] Ashish Venugopal, Andreas Zollmann, Noah A. Smith, and Stephan Vogel. [Wider pipelines: \$N\$ -best alignments and parses in MT training](#). In *Proceedings of the Conference of the Association for Machine Translation in the Americas*, Waikiki, HI, October 2008. **AMTA 2008**.
- [81] André F. T. Martins, Mário A. T. Figueiredo, Pedro M. Q. Aguiar, Noah A. Smith, and Eric P. Xing. [Nonextensive entropic kernels](#). In *Proceedings of the International Conference on Machine Learning*, pages 640–647, Helsinki, Finland, July 2008. **ICML 2008**.
- [82] Shay B. Cohen and Noah A. Smith. [Joint morphological and syntactic disambiguation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 208–217, Prague, Czech Republic, June 2007. **EMNLP-CoNLL 2007**.
- [83] David A. Smith and Noah A. Smith. [Probabilistic models of nonprojective dependency trees](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 132–140, Prague, Czech Republic, June 2007. **EMNLP-CoNLL 2007**.
- [84] Mengqiu Wang, Noah A. Smith, and Teruko Mitamura. [What is the Jeopardy model? a quasi-synchronous grammar for QA](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 22–32, Prague, Czech Republic, June 2007. **EMNLP-CoNLL 2007**.
Nominated for best paper award.
- [85] Noah A. Smith, Douglas L. Vail, and John D. Lafferty. [Computationally efficient M-estimation of log-linear structure models](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 752–759, Prague, Czech Republic, June 2007. **ACL 2007**.
- [86] Noah A. Smith and Jason Eisner. [Annealing structural bias in multilingual weighted grammar induction](#). In *Proceedings of the International Conference on Computational Linguistics and Annual Meeting of the Association for Computational Linguistics*, pages 569–576, Sydney, Australia, July 2006. **COLING-ACL 2006**.
- [87] Jason Eisner and Noah A. Smith. [Parsing with soft and hard constraints on dependency length](#). In *Proceedings of the International Workshop on Parsing Technologies*, pages 30–41, Vancouver, BC, October 2005. **IWPT 2005**.
- [88] Noah A. Smith, David A. Smith, and Roy W. Tromble. [Context-based morphological disambiguation with random fields](#). In *Proceedings of the Human Language Technology Conference and Conference on*

- Empirical Methods in Natural Language Processing*, pages 475–482, Vancouver, BC, October 2005. **EMNLP 2005.**
- [89] Jason Eisner, Eric Goldlust, and Noah A. Smith. [Compiling Comp Ling: Practical weighted dynamic programming and the Dyna language](#). In *Proceedings of the Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 281–290, Vancouver, BC, October 2005. **EMNLP 2005.**
 - [90] Noah A. Smith and Jason Eisner. [Contrastive estimation: Training log-linear models on unlabeled data](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 354–362, Ann Arbor, MI, June 2005. **ACL 2005.**
Nominated for best paper award.
 - [91] David A. Smith and Noah A. Smith. [Bilingual parsing with factored estimation: Using English to parse Korean](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 49–56, Barcelona, Spain, July 2004. **EMNLP 2004.**
 - [92] Noah A. Smith and Jason Eisner. [Annealing techniques for unsupervised statistical language learning](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 487–494, Barcelona, Spain, July 2004. **ACL 2004.**
 - [93] Noah A. Smith. [From words to corpora: Recognizing translation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 95–102, Philadelphia, PA, July 2002. **EMNLP 2002.**

Refereed conference publications (short)

- [94] David Bamman, Chris Dyer, and Noah A. Smith. [Distributed representations of geographically situated language](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Baltimore, MD, June 2014. **ACL 2014.**
- [95] Rohan Ramanath, Fei Liu, Norman Sadeh, and Noah A. Smith. [Unsupervised alignment of privacy policies using hidden Markov models](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Baltimore, MD, June 2014. **ACL 2014.**
- [96] Michael T. Mordowanec, Nathan Schneider, Chris Dyer, and Noah A. Smith. [Simplified dependency annotations with GFL-Web](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics, companion volume*, Baltimore, MD, June 2014. **ACL 2014 demonstration track.**
- [97] André F. T. Martins, Miguel Almeida, and Noah A. Smith. [Turning on the turbo: Fast third-order non-projective turbo parsers](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Sofia, Bulgaria, August 2013. **ACL 2013.**
- [98] Tae Yano, Dani Yogatama, and Noah A. Smith. [A penny for your tweets: Campaign contributions and Capitol Hill microblogs](#). In *Proceedings of the International AAAI Conference on Weblogs and Social Media*, Boston, MA, July 2013. **ICWSM 2013.**
- [99] Nathan Schneider, Behrang Mohit, Chris Dyer, Kemal Oflazer, and Noah A. Smith. [Supersense tagging for Arabic: the MT-in-the-middle attack](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, Atlanta, GA, June 2013. **NAACL 2013.**
- [100] Chris Dyer, Victor Chahuneau, and Noah A. Smith. [A simple, fast, and effective reparameterization of IBM model 2](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, Atlanta, GA, June 2013. **NAACL 2013.**
- [101] Nathan Schneider, Behrang Mohit, Kemal Oflazer, and Noah A. Smith. [Coarse lexical semantic annotation with supersenses: An Arabic case study](#). In *Proceedings of the Annual Meeting of the*

- Association for Computational Linguistics*, Jeju, Korea, July 2012. **ACL 2012.**
- [102] Kevin Gimpel and Noah A. Smith. [Concavity and initialization for unsupervised dependency parsing](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, Montréal, Québec, June 2012. **NAACL 2012.**
- [103] Jonathan H. Clark, Chris Dyer, Alon Lavie, and Noah A. Smith. [Better hypothesis testing for statistical machine translation: Controlling for optimizer instability](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics, companion volume*, Portland, OR, June 2011. **ACL 2011.**
- [104] Kevin Gimpel, Nathan Schneider, Brendan O’Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. [Part-of-speech tagging for Twitter: Annotation, features, and experiments](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics, companion volume*, Portland, OR, June 2011. **ACL 2011.**
- [105] Kevin Gimpel and Noah A. Smith. [Softmax-margin CRFs: Training log-linear models with cost functions](#). In *Proceedings of the North American Chapter of the Association for Computational Linguistics Human Language Technologies Conference*, Los Angeles, CA, June 2010. **NAACL 2010.**
- [106] Mahesh Joshi, Dipanjan Das, Kevin Gimpel, and Noah A. Smith. [Movie reviews and revenues: An experiment in text regression](#). In *Proceedings of the North American Chapter of the Association for Computational Linguistics Human Language Technologies Conference*, Los Angeles, CA, June 2010. **NAACL 2010.**
- [107] Tae Yano and Noah A. Smith. [What’s worthy of comment? content and comment volume in political blogs](#). In *Proceedings of the International AAAI Conference on Weblogs and Social Media*, Washington, DC, May 2010. **ICWSM 2010.**
- [108] Shay B. Cohen and Noah A. Smith. [Variational inference for grammar induction with prior knowledge](#). In *Proceedings of the Joint Conference of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing, companion volume*, pages 1–4, Singapore, August 2009. **ACL 2009.**
- [109] Sourish Chaudhuri, Naman K. Gupta, Noah A. Smith, and Carolyn P. Rosé. [Leveraging structural relations for fluent compressions at multiple compression rates](#). In *Proceedings of the Joint Conference of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing, companion volume*, pages 101–104, Singapore, August 2009. **ACL 2009.**
- [110] Ramnath Balasubramanyan, Frank Lin, William W. Cohen, Matthew Hurst, and Noah A. Smith. [From episodes to sagas: Understanding the news by identifying temporally related story sequences](#). In *Proceedings of the International AAAI Conference on Weblogs and Social Media*, San Jose, CA, May 2009. **ICWSM 2009.**
- [111] Daniel R. Rashid and Noah A. Smith. [Relative keyboard input system](#). In *Proceedings of the International Conference on Intelligent User Interfaces*, pages 397–400, Canary Islands, Spain, January 2008. **IUI 2008.**
- [112] Markus Dreyer, David A. Smith, and Noah A. Smith. [Vine parsing and minimum risk reranking for speed and precision](#). In *Proceedings of the Conference on Natural Language Learning*, pages 201–205, New York, NY, June 2006. **CoNLL 2006.**
- [113] Jason Eisner, Eric Goldlust, and Noah A. Smith. [Dyna: A declarative language for implementing dynamic programs](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics, companion volume*, pages 218–221, Barcelona, Spain, July 2004. **ACL 2004.**
- [114] Noah A. Smith and Michael E. Jahr. [Cairo: An alignment visualization tool](#). In *Proceedings of the Language Resources and Evaluation Conference*, pages 549–552, Athens, Greece, May/June 2000. **LREC 2000.**

- [115] Jesse Dodge, Manaal Faruqui, Sujay K. Jauhar, Chris Dyer, Eduard Hovy, and Noah A. Smith. [Retrofitting word vectors to semantic lexicons](#). In *NIPS Deep Learning and Representation Learning Workshop*, Montréal, Quebec, December 2014.
- [116] Dani Yogatama, Manaal Faruqui, Chris Dyer, and Noah A. Smith. [Learning word representations with hierarchical sparse coding](#). In *NIPS Deep Learning and Representation Learning Workshop*, Montréal, Quebec, December 2014.
- [117] Sam Thomson, Brendan O’Connor, Jeffrey Flanigan, David Bamman, Jesse Dodge, Swabha Swayamdipta, Nathan Schneider, Chris Dyer, and Noah A. Smith. [CMU: Arc-factored, discriminative semantic dependency parsing](#). In *Proceedings of the International (COLING) Workshop on Semantic Evaluations*, Dublin, Ireland, August 2014. **SemEval 2014.**
- [118] Shiladitya Sinha, Chris Dyer, Kevin Gimpel, and Noah A. Smith. [Predicting the NFL using Twitter](#). In *Proceedings of the ECML/PKDD Workshop on (Machine Learning and Data Mining for) Sports Analytics*, Prague, Czech Republic, September 2013.
- [119] Nathan Schneider, Brendan O’Connor, Naomi Saphra, David Bamman, Manaal Faruqui, Noah A. Smith, Chris Dyer, and Jason Baldridge. [A framework for \(under\)specifying dependency syntax without overloading annotators](#). In *Proceedings of the ACL Linguistic Annotation Workshop*, Sofia, Bulgaria, August 2013. **LAW 2013.**
- [120] Nathan Schneider, Chris Dyer, and Noah A. Smith. Exploiting and expanding corpus resources for frame-semantic parsing. April 2013. International FrameNet Workshop.
- [121] Jacob Eisenstein, Brendan O’Connor, Noah A. Smith, and Eric P. Xing. [Mapping the geographical diffusion of new words](#). In *Proceedings of the NIPS Workshop on Social Network and Social Media Analysis: Methods, Models and Applications*, Lake Tahoe, NV, December 2012.
- [122] Waleed Ammar, Chris Dyer, and Noah A. Smith. [Transliteration by sequence labeling with lattice encodings and reranking](#). In *Proceedings of the ACL Named Entities Workshop*, Jeju, Korea, July 2012.
- [123] Yanchuan Sim, Noah A. Smith, and David A. Smith. [Discovering factions in the computational linguistics community](#). In *Proceedings of the ACL Workshop on Rediscovering Fifty Years of Discoveries*, Jeju, Korea, July 2012.
- [124] Brendan O’Connor, David Bamman, and Noah A. Smith. [Computational text analysis for social science: Model complexity and assumptions](#). In *Proceedings of the NIPS Workshop on Computational Social Science and the Wisdom of Crowds*, Sierra Nevada, Spain, December 2011.
- [125] Chris Dyer, Kevin Gimpel, Jonathan H. Clark, and Noah A. Smith. [The CMU-ARK German-English translation system](#). In *Proceedings of the EMNLP Workshop on Statistical Machine Translation*, Edinburgh, UK, July 2011. **SMT 2011.**
- [126] Kevin Gimpel and Noah A. Smith. [Generative models of monolingual and bilingual gappy patterns](#). In *Proceedings of the EMNLP Workshop on Statistical Machine Translation*, Edinburgh, UK, July 2011. **SMT 2011.**
- [127] Desai Chen, Chris Dyer, Shay B. Cohen, and Noah A. Smith. [Unsupervised bilingual POS tagging with Markov random fields](#). In *Proceedings of the EMNLP Workshop on Unsupervised Learning in NLP*, Edinburgh, UK, July 2011. **UNSUP 2011.**
- [128] Jacob Eisenstein, Tae Yano, William W. Cohen, Noah A. Smith, and Eric P. Xing. [Structured databases of named entities from Bayesian nonparametrics](#). In *Proceedings of the EMNLP Workshop on Unsupervised Learning in NLP*, Edinburgh, UK, July 2011. **UNSUP 2011.**
- [129] Dong Nguyen, Noah A. Smith, and Carolyn P. Rosé. [Author age prediction from text using linear regression](#). In *Proceedings of the ACL Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, Portland, OR, June 2011. **LATECH 2011.**

- [130] André F. T. Martins, Noah A. Smith, Eric P. Xing, Pedro M. Q. Aguiar, and Mário A. T. Figueiredo. [Augmenting dual decomposition for MAP inference](#). In *Proceedings of the International Workshop on Optimization for Machine Learning*, Whistler, BC, December 2010. **OPT 2010.**
- [131] André F. T. Martins, Noah A. Smith, Eric P. Xing, Pedro M. Q. Aguiar, and Mário A. T. Figueiredo. [Online multiple kernel learning for structured prediction](#). In *Proceedings of the NIPS Workshop on New Directions in Multiple Kernel Learning*, Whistler, BC, December 2010.
- [132] Brendan O'Connor, Jacob Eisenstein, Eric P. Xing, and Noah A. Smith. [Discovering demographic language variation](#). In *Proceedings of the NIPS Workshop on Machine Learning for Social Computing*, Whistler, BC, December 2010.
- [133] Desai Chen, Nathan Schneider, Dipanjan Das, and Noah A. Smith. [SEMAFOR: Frame argument resolution with log-linear models](#). In *Proceedings of the International (ACL) Workshop on Semantic Evaluations*, Uppsala, Sweden, July 2010. **SemEval 2010.**
- [134] Michael Heilman and Noah A. Smith. [Extracting simplified statements for factual question generation](#). In *Proceedings of the AIED Workshop on Question Generation*, Pittsburgh, PA, June 2010.
- [135] Michael Heilman and Noah A. Smith. [Rating computer-generated questions with Mechanical Turk](#). In *Proceedings of the NAACL-HLT Workshop on Creating Speech and Language Data With Mechanical Turk*, Los Angeles, CA, June 2010.
- [136] Tae Yano, Philip Resnik, and Noah A. Smith. [Shedding \(a thousand points of\) light on biased language](#). In *Proceedings of the NAACL-HLT Workshop on Creating Speech and Language Data With Mechanical Turk*, Los Angeles, CA, June 2010.
- [137] Michael Heilman and Noah A. Smith. [Ranking automatically generated questions as a shared task](#). In *Proceedings of the AIED Workshop on Question Generation*, Brighton, UK, July 2009.
- [138] André F. T. Martins and Noah A. Smith. [Summarization with a joint model for sentence extraction and compression](#). In *Proceedings of the NAACL-HLT Workshop on Integer Linear Programming for Natural Language Processing*, Boulder, CO, June 2009.
- [139] Shay B. Cohen and Noah A. Smith. [The shared logistic normal distribution for grammar induction](#). In *Proceedings of the NIPS Workshop on Speech and Language: Unsupervised Latent-Variable Models*, Whistler, BC, December 2008.
- [140] Noah A. Smith, Michael Heilman, and Rebecca Hwa. [Question generation as a competitive undergraduate course project](#). In *Proceedings of the NSF Workshop on the Question Generation Shared Task and Evaluation Challenge*, Arlington, VA, September 2008.
- [141] Jason Eisner and Noah A. Smith. [Competitive grammar writing](#). In *Proceedings of the ACL Workshop on Issues in Teaching Computational Linguistics*, pages 97–105, Columbus, OH, June 2008.
- [142] Kevin Gimpel and Noah A. Smith. [Rich source-side context for statistical machine translation](#). In *Proceedings of the ACL Workshop on Statistical Machine Translation*, pages 9–17, Columbus, OH, June 2008. **SMT 2008.**
Five-year retrospective best paper award.
- [143] Noah A. Smith and Jason Eisner. [Guiding unsupervised grammar induction using contrastive estimation](#). In *Proceedings of the IJCAI Workshop on Grammatical Inference Applications*, pages 73–82, Edinburgh, UK, July 2005.

Theses

- [144] Noah A. Smith. [Novel Estimation Methods for Unsupervised Discovery of Latent Structure in Natural Language Text](#). Ph.D. thesis, Department of Computer Science, Johns Hopkins University, Baltimore, MD, October 2006. Supervised by Jason Eisner.
- [145] Noah A. Smith. [Detection of translational equivalence](#). Technical report 4253, Department of Com-

- puter Science, University of Maryland College Park, College Park, MD, May 2001. Undergraduate honors thesis, supervised by Philip Resnik.
- [146] Noah A. Smith. [Ellipsis happens, and deletion is how](#). In Andrea Gualmini, Soo-Min Hong, and Mitsue Motomura, editors, *University of Maryland Working Papers in Linguistics*, volume 11, pages 176–191. Department of Linguistics, University of Maryland, November 2001. Undergraduate honors thesis, supervised by Norbert Hornstein.
 - [147] Alison J. Deming, Steven P. Denny, Jessica Exelbert, Anne Italiano, Dan Malinow, Katie E. Praske, Bradley Rhoderick, Noah A. Smith, Amanda Stamper, and Margaret E. Wood. Smart growth: An analysis in three counties, May 2001. University of Maryland College Park, Gemstone team thesis, supervised by Jacqueline Rogers.

OTHER RESEARCH OUTCOMES

Supervised doctoral theses

- [148] Brendan O'Connor. *Statistical Text Analysis for Social Science*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, August 2014.
- [149] Nathan Schneider. *Lexical Semantic Analysis in Natural Language Text*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, June 2014.
- [150] Tae Yano. *Text as Actuator: Text-Driven Response Modeling and Prediction in Politics*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, July 2013.
- [151] Kevin Gimpel. *Discriminative Feature-Rich Modeling for Syntax-Based Machine Translation*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, August 2012.
- [152] André F. T. Martins. *The Geometry of Constrained Structured Prediction: Applications to Inference and Learning of Natural Language Syntax*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, May 2012.

Awarded the IBM Portugal *Prémio Científico* and Honorable Mention for the SCS Dissertation Prize.

- [153] Dipanjan Das. *Semi-Supervised and Latent-Variable Models of Natural Language Semantics*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, April 2012.
- [154] Shay B. Cohen. *Computational Learning of Probabilistic Grammars in the Unsupervised Setting*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, September 2011.
- [155] Michael Heilman. *Automatic Factual Question Generation from Text*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, April 2011.

Supervised undergraduate theses

- [156] Desai Chen. [Unsupervised bilingual POS tagging with Markov random fields](#), May 2011. Computer Science honors thesis, School of Computer Science, Carnegie Mellon University.
- [157] Rishav Bhowmick. [Rich entity type recognition in text](#), May 2010. Computer Science honors thesis, Department of Computer Science, Carnegie Mellon University–Qatar.
- [158] Neel Shah. [Predicting risk from financial reports with supervised topic models](#), May 2010. Computer Science honors thesis, School of Computer Science, Carnegie Mellon University.
- [159] Matthew Thompson. [Möbius: Exploring a new modality for poetry generation](#), April 2009. Linguistics senior thesis, Department of Philosophy, Carnegie Mellon University.

Unpublished technical reports and working papers

- [160] Noah A. Smith, Claire Cardie, Anne L. Washington, and John D. Wilkerson. [Overview of the 2014 NLP unshared task in PoliInformatics](#). In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pages 5–7, Baltimore, MD, June 2014.
- [161] Amber Boydstun, Dallas Card, Justin H. Gross, Philip Resnik, and Noah A. Smith. [Tracking the development of media frames within and across policy issues](#), August 2014.
- [162] Lingpeng Kong and Noah A. Smith. [An empirical comparison of parsing methods for stanford dependencies](#). April 2014.
- [163] Dani Yogatama, Bryan R. Routledge, and Noah A. Smith. [A sparse and adaptive prior for time-dependent model parameters](#). October 2013.
- [164] Justin H. Gross, Brice Acree, Yanchuan Sim, and Noah A. Smith. [Testing the etch-a-sketch hypothesis: A computational analysis of Mitt Romney’s ideological makeover during the 2012 primary vs. general elections](#), August 2013. Presented at the Annual Meeting of the American Political Science Association.
- [165] David Bamman and Noah A. Smith. [New alignment methods for discriminative summarization](#). May 2013.
- [166] David Bamman, Adam Anderson, and Noah A. Smith. [Inferring social rank in an Old Assyrian trade network](#). July 2013. Presented at Digital Humanities.
- [167] Waleed Ammar, Shomir Wilson, Norman Sadeh, and Noah A. Smith. [Automatic categorization of privacy policies: A pilot study](#). Technical Report CMU-LTI-12-019, Carnegie Mellon University, Pittsburgh, PA, December 2012.
- [168] Noah A. Smith. [Adversarial evaluation for models of natural language](#). July 2012.
- [169] Chris Dyer, Noah A. Smith, Graham Morehead, Phil Blunsom, and Abby Levenberg. [The CMU-Oxford translation system for the NIST open machine translation 2012 evaluation](#), May 2012.
- [170] Dani Yogatama, Michael Heilman, Brendan O’Connor, Chris Dyer, Bryan R. Routledge, and Noah A. Smith. [Predicting responses and discovering social factors in scientific literature](#). Technical Report CMU-LTI-11-015, Carnegie Mellon University, Pittsburgh, PA, September 2011.
- [171] Behrang Mohit, Nathan Schneider, Rishav Bhowmick, Kemal Oflazer, and Noah A. Smith. [Recall-oriented learning for named entity recognition in Wikipedia](#). Technical Report CMU-LTI-11-012, Carnegie Mellon University, Pittsburgh, PA, August 2011.
- [172] Tal Stramer, Bryan R. Routledge, and Noah A. Smith. [Predicting FED action from text](#). Technical Report CMU-LTI-11-005, Carnegie Mellon University, Pittsburgh, PA, May 2011.
- [173] Nathan Schneider, Rebecca Hwa, Philip Gianfortoni, Dipanjan Das, Michael Heilman, Alan W. Black, Frederick L. Crabbe, and Noah A. Smith. [Visualizing topical quotations over time to understand news discourse](#). Technical Report CMU-LTI-10-013, Carnegie Mellon University, Pittsburgh, PA, July 2010.
- [174] André F. T. Martins, Kevin Gimpel, Noah A. Smith, Eric P. Xing, Pedro M. Q. Aguiar, and Mário A. T. Figueiredo. [Aggressive online learning of structured classifiers](#). Technical Report CMU-ML-10-109, Carnegie Mellon University, Pittsburgh, PA, June 2010.
- [175] Kevin Gimpel and Noah A. Smith. [Softmax-margin training for structured log-linear models](#). Technical Report CMU-LTI-10-008, Carnegie Mellon University, Pittsburgh, PA, June 2010.
- [176] Dipanjan Das, Nathan Schneider, Desai Chen, and Noah A. Smith. [SEMAFOR 1.0: A probabilistic frame-semantic parser](#). Technical Report CMU-LTI-10-001, Carnegie Mellon University, Pittsburgh, PA, April 2010.
- [177] Noah A. Smith. [Text-driven forecasting](#). March 2010.
- [178] Michael Heilman and Noah A. Smith. [Question generation via overgenerating transformations and ranking](#). Technical Report CMU-LTI-09-013, Carnegie Mellon University, Pittsburgh, PA, June 2009.
- [179] Dan Tasse and Noah A. Smith. [SOUR CREAM: Toward semantic processing of recipes](#). Technical

- Report CMU-LTI-08-005, Carnegie Mellon University, Pittsburgh, PA, May 2008.
- [180] Yaser Al-Onaizan, Jan Cuřin, Michael Jahr, Kevin Knight, John Lafferty, I. Dan Melamed, Noah A. Smith, Franz-Josef Och, David Purdy, and David Yarowsky. [Statistical machine translation](#). CLSP Research Notes 42, Johns Hopkins University, Baltimore, MD, 1999.
 - [181] Margaret E. Wood, Noah A. Smith, Anne Italiano, Jessica Exelbert, Steven Denny, and Konrad Aschenbach. Edgewood Terrace: The decline and revitalization of a mixed-income neighborhood. November 1999. Association for Public Policy Analysis and Management Fall Conference on Global and Comparative Perspectives.

Publicly available datasets and software

- [182] [GFL-WEB](#), developed by Michael T. Mordowanec. Web interface for the Graph Fragment Language for syntactic annotation, 2014, see [96].
- [183] [JAMR](#), developed by Jeffrey Flanigan and Sam Thomson. Semantic parser for the Abstract Meaning Representation, 2014, see [29].
- [184] [COMPREHENSIVE MULTIWORD EXPRESSIONS CORPUS](#), developed by Nathan Schneider. English Web Treebank annotated with multiword expressions, 2014, see [32].
- [185] [AMALGR](#), developed by Nathan Schneider. Multiword expression identification tool, 2014, see [7].
- [186] [NFL TWEET DATASET](#), developed by Shiladitya Sinha, Brendan O’Connor, Chris Dyer, and Kevin Gimpel. NFL game data and identifiers of tweets aligned to games, 2013, see [118].
- [187] [CMU MOVIE SUMMARY CORPUS](#), developed by David Bamman and Brendan O’Connor. Corpus of movie plot summaries and associated metadata, 2013, see [36].
- [188] [GLOBAL VOICES MALAGASY-ENGLISH PARALLEL CORPUS](#), developed by Victor Chahuneau. Corpus of parallel news articles from the Global Voices citizen media project, 2012.
- [189] [WORD SALAD](#), developed by Victor Chahuneau. Corpus of restaurant menus, 2012, see [40].
- [190] [ARABIC WIKIPEDIA SUPERSENSE CORPUS](#), developed by Nathan Schneider and Behrang Mohit. Articles tagged with nominal supersenses, 2012, see [101].
- [191] [AD³](#), developed by André Martins. Approximate MAP decoder, 2012, see [3, 42, 49, 52, 131].
- [192] [RAMPION](#), developed by Kevin Gimpel. Algorithm for training statistical machine translation models based on minimizing structured ramp loss, 2012, see [44].
- [193] [CONGRESSIONAL BILLS CORPUS](#), developed by Tae Yano. Congressional bills and committee outcomes, 2012, see [45].
- [194] [ARABIC WIKIPEDIA NAMED ENTITY CORPUS AND TAGGER](#), developed by Nathan Schneider and Behrang Mohit. Articles tagged with named entities, and a statistical tagger, 2012, see [46].
- [195] [TWITTER PART-OF-SPEECH TAGGING](#), developed by Olutobi Owoputi, Kevin Gimpel, Nathan Schneider, Brendan O’Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, Chris Dyer, and Noah A. Smith. Dataset of tweets manually annotated with part-of-speech tags; part-of-speech tagger trained on this data; a simple browser-based POS tagging annotation interface, 2011, see [39, 104].
- [196] [QUESTION-ANSWER DATA](#), developed by Michael Heilman, Shay Cohen, and Kevin Gimpel. Question-answer pairs generated by undergraduates for the purpose of developing and evaluating question answering systems., 2010, see [140].
- [197] [AMAZON MECHANICAL TURK POLITICAL BIAS DATA](#), developed by Tae Yano. Sentences from political blogs with crowdsourced annotations of political bias, 2010, see [136].
- [198] [MOVIE\\$ CORPUS](#), developed by Mahesh Joshi, Dipanjan Das, and Kevin Gimpel. Collection of pre-release movie reviews, metadata, and opening weekend revenues, 2010, see [106].

- [199] [SEMAFOR](#), developed by Dipanjan Das, Nathan Schneider, and Desai Chen. Frame-semantic parser, 2010, see [64].
- [200] [TURBOPARSER](#), developed by André Martins. Multilingual dependency parser, 2009, see [70].
- [201] [QUIPU](#), developed by Kevin Gimpel. Statistical machine translation system, 2009, see [68].
- [202] [POLITICAL BLOG CORPUS](#), developed by Tae Yano. Text collection from five American political blogs, 2009, see [75].
- [203] [10-K CORPUS](#), developed with three others. Collection of 10-K reports and preceding and following stock return volatility measurements, 2009, see [73].
- [204] [DAGEEM](#), developed by Shay Cohen. Unsupervised dependency grammar induction, 2008, see [78].
- [205] [MSTPARSER](#), [STACKED](#), developed by André Martins and Dipanjan Das. Multilingual dependency parser, 2008, see [79].
- [206] [CURD](#), developed by Dan Tasse. Corpus of semantically annotated recipes, 2008, see [179].
- [207] [DYNA](#), developed with five others. Declarative programming language for weighted dynamic programming, 2004, see [89, 113].
- [208] [EGYPT](#), developed with nine others. Toolkit for statistical machine translation, including GIZA training module and CAIRO word alignment visualizer, 1999, see [114, 180].