



AT&T Bell Labs, Murray Hill, NJ Artificial Intelligence Research Department	summers 1989–1992
Microsoft Corporation, Seattle, WA Programmer	summer 1988
IBM Research Center, Yorktown Heights, NY Consultant	1987–1988

PROFESSIONAL  
ACTIVITIES

President, ACL SIGMORPHON (Computational Morphology and Phonology) (2001–); executive committee (1998–).

Graduate program chair, CS Department, Johns Hopkins University (2002–).

Problems committee, North American Computational Linguistics Olympiad (2006–).

Editorial board member, “Language and Computation” corner, *Journal of Logic and Computation* (2011–).

Guest editorial board member, special issue on Statistical Learning of Natural Language Structured Input and Output, *Journal of Natural Language Engineering* (2011).

Editorial board member, *Research in Language and Computation* (2006–2010).

Program chair, [EMNLP-CoNLL](#) (2007).

Program chair, [SIGMORPHON Workshop](#) (2008, with Jeff Heinz); SIGPHON Workshop on Finite-State Phonology (2000).

Chair of Best Paper Award committee, EMNLP (2010).

Program area chair or co-chair for NAACL-HLT (2012, morphology/phonology; 2006, syntax/grammar/morphology), EMNLP (2009, phonology/morphology/tagging/chunking/segmentation; 2006, machine learning), EACL (2006, phonology/morphology/finite-state/tagging/segmentation).

Publications chair, ACL (2005).

Co-chair for HLT/NAACL Tutorial and Workshop Programs (2003).

Invited panelist, ACL Workshop on Multiword Expressions (2011).

Advisory committee, ACL SIGDAT (Linguistic Data and Corpus-Based Approaches to NLP) (2007–).

Advisory committee member, Zooniverse project of the Adler Planetarium and JHU Space Telescope Science Institute (2009–).

Advisory board member, cogito.org (2005–).

Advisory board member, AMS GAANN Fellowship Program (2009–).

Special issue editor, *Cognitive Science* (2002).

Journal reviewer, *ACM Transactions on Speech and Language Processing* (2008, 2005), *Computational Linguistics* (2011, 2010, 2007, 2005, 2004, 2002, 2001, 2000), *Cognition* (2002), *IEEE Transactions on Audio, Speech, and Language Processing* (2007), *J. Algorithms* (1997), *J. Logic and Computation* (2011), *Language and Speech* (1999), *Phonology* (2008), *Research on Language and Computation* (2008, 2007).

Program committee for ACL (2012, 2011, 2010, 2009, 2007, 2005, 2004, 2003, 2002), COLING (2008), EACL (2012, 2006, 2003, 1999), EMNLP (2011, 2010, 2009, 2008, 2007, 2006, 2003, 2002), FSMNLP (2005, 2001), ICGI (2012), ICML (2004), IJCAI (2007), IWPT (2009), MITWPL (1999), NAACL (2012, 2010, 2009, 2006), NLP-LING (2010), SIGMORPHON (2012, 2010, 2008, 2006, 2004, 2002, 2000, 1998), TeachCL/TNLP (2008, 2005, 2002), Formal Approaches to Slavic Linguistics 8 (1999), International Conference on Dependency Linguistics (2011), ACL Workshop on Unsupervised Learning (2011); reviewer for AISTATS (2010), ICFP (2008), NIPS (2011, 2010, 2007).

One of the “best reviewer” awards for EMNLP (2010).

NSF panel or *ad hoc* reviewer (2010, 2009, 2008, 2004, 2003).

Board of reviewers, *Handbook of Natural Language Processing* (2008–2009).

Advisor to DARPA seedling in Adaptive Interactive Representations (2008–2009).

Member, EU/NSF joint working group: “ePhilology: Emerging Language Technologies and the Rediscovery of the Past” (2002).

FELLOWSHIPS AND AWARDS	Finalist for best paper award (ACL, EMNLP-CoNLL, ACL, EMNLP)	2009, 2007, 2005, 2002
	Robert B. Pond, Sr. Excellence in Teaching Award Johns Hopkins University, Whiting School of Engineering	2005
	NSF Graduate Research Fellowship (computer science)	1993–1996
	Herchel Smith Harvard Scholarship (mathematics)	1991–1993
	Fulbright Scholarship (creative writing), South Africa	1990–1991
	Harvard National Scholarship	1986–1990
GRANTS	NSF IGERT: Very Large-scale Multilingual Human Language Technology: Empirical, Analytical and Computational Foundations (PI, leading a group of $\approx 25$ faculty) [submitted 2011]	exp. 2012–2017
	NSF RI-Medium: Learned Dynamic Prioritization (PI, \$900K)	2010–2013

JHU Framework for the Future: Initiative in Computational Learning (PI, \$50K, leading a group of 30+ faculty)	2009–
NSF: Computing Innovation Fellows Program (postdoc mentor, \$140K)	2009–2011
NSF RI: Cross-Cutting Research Workshops in Intelligent Information Systems (co-PI, \$647K plus substantial additional funding from other agencies and corporations)	2007–2010
JHU WSE-APL Partnership Fund: Learning with Less (PI, \$68K)	2006–2007
NSF PIRE: Investigation of Meaning Representations in Language Understanding for Machine Translation Systems (co-PI, \$2.5M)	2005–2010
NSF CAREER: Finite-State Machine Learning on Strings and Sequences (PI, \$500K)	2004–2010
NSF ITR: Weighted Dynamic Programming for Statistical Natural Language Processing (PI, \$425K)	2003–2007
ONR MURI: Improving Statistical Translation Models Via Text Analyzers Trained From Parallel Corpora (co-PI, \$4.3M)	2001–2006
NSF ITR/IM+PE+SY: Summer Workshops on Human Language Technology (co-PI, \$2.35M)	2001–2006

## PUBLICATIONS AND PRESENTATIONS

### Invited talks

A non-parametric Bayesian approach to inflectional morphology. JHU Applied Math & Statistics Dept., December 2011; University of Maryland, November 2011; ICML/ACL/ICSA Symposium on Machine Learning in Speech and Language Processing, June 2011; Workshop on Machine Translation and Morphologically-Rich Languages, January 2011.

Toward unsupervised web scraping. DIRE Meeting, May 2011.

A weighted deductive language for declaratively specifying (some) algorithms. University of Bielefeld, July 2010.

Using dynamic programming to help search for reorderings. University of Bielefeld, July 2010.

Variational inference over structured variables for linguistic modeling. University of Edinburgh, May 2010.

A weighted deductive language for declaratively specifying (some) algorithms. University of Edinburgh, May 2010.

Constraint interaction, probabilistic models, and approximate inference. Chicago Linguistic Society Annual Conference, April 2010.

Extending logic programming to support modern statistical AI. Datalog 2.0 Workshop, March 2010.

Weighted deduction as an abstraction level for AI. Joint conference on Statistical Relational Learning + Inductive Logic Programming + Mining and Learning with Graphs, July 2009. (keynote talk)

Joint models with missing data for semi-supervised learning. NAACL Workshop on Semi-Supervised Learning for Natural Language Processing, June 2009. (keynote talk)

Dependency parsing by belief propagation. Boulder Workshop on Dependency Parsing, June 2009; Temple University, November 2008.

Shuffling non-constituents. Second Workshop on Syntax and Structure in Statistical Translation, ACL-08: HLT, June 2008. (keynote talk)

The Dyna language. CMU and Google, May 2008; MIT, November 2006; IBM Yorktown Heights, May 2006; Microsoft Research, August 2005; University of Washington, August 2005.

Searching for optimal permutations with very large-scale neighborhoods. JHU Applied Math & Statistics Dept., November 2007.

Discovering syntactic deep structure via Bayesian statistics. U. of Maryland, May 2007.

Bootstrapping without the boot. MITRE Corporation, August 2006; IPAM Document Space Workshop, January 2006.

Parameterized finite-state machines and their training. U. of Saarland, Germany, March 2004; AT&T Research, October 2002.

Inferring transformations. Mathematics of Language Conference (MoL8), Bloomington, June 2003.

Learning natural-language grammars using a Bayesian prior. Rochester Institute of Technology, May 2000; Johns Hopkins University, February 2000; UCLA Linguistics Department, June 1999; Stanford, 1999; U. of Texas at Austin, 1999; U. of Toronto, 1999; U. of Rochester, 1999; U. of Chicago, 1999.

Doing OT in a straitjacket. Johns Hopkins Cognitive Science Dept., 2002; U. of Rochester Linguistics Dept., 2000; UCLA Linguistics Dept., 1999; Stanford Linguistics Dept., 1999; MIT Linguistics Dept., 1997.

### Journal Articles

John Eng and Jason M. Eisner. [Radiology report entry with automatic phrase completion driven by language modeling](#). *Radiographics*, 24(5):1493–1501, 2004.

Jason Eisner. [Discovering syntactic deep structure via Bayesian statistics](#). *Cognitive Science*, 26(3):255–268, 2002.

Jason Eisner. [Introduction to the special section on linguistically apt statistical methods](#). *Cognitive Science*, 26(3):235–237, 2002.

## Book Chapters

- Jason Eisner and Nathaniel W. Filardo. [Dyna: Extending Datalog for modern AI](#). In Tim Furche, Georg Gottlob, Giovanni Grasso, Oege de Moor, and Andrew Sellers, editors, *Datalog 2.0*, Lecture Notes in Computer Science. Springer, 2011. 40 pages.
- Antti-Veikko Rosti, Eugene Matusov, Jason Smith, Necip Ayan, Jason Eisner, Damianos Karakos, Sanjeev Khudanpur, Gregor Leusch, Zhifei Li, Spyros Matsoukas, Hermann Ney, Richard Schwartz, B. Zhang, and J. Zheng. [Confusion network decoding for MT system combination](#). In *Handbook of Natural Language Processing and Machine Translation*, pages 333–361. Springer, 2011.
- 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*, chapter 8, pages 121–150. Springer, 2010.
- Jason Eisner. [Bilexical grammars and their cubic-time parsing algorithms](#). In Harry Bunt and Anton Nijholt, editors, *Advances in Probabilistic and Other Parsing Technologies*, pages 29–62. Kluwer Academic Publishers, 2000.

## Book Reviews

- Jason Eisner. [Review of Optimality Theory by René Kager](#). *Computational Linguistics*, 26(2):286–290, 2000.

## Ph.D. Thesis

- Jason Eisner. [Smoothing a Probabilistic Lexicon via Syntactic Transformations](#). PhD thesis, University of Pennsylvania, 2001. 318 pages.

## Refereed Conference and Workshop Proceedings

- Jason Eisner and Hal Daumé III. [Learning speed-accuracy tradeoffs in nondeterministic inference algorithms](#). In *COST: NIPS 2011 Workshop on Computational Trade-offs in Statistical Learning*, 2011. 5 pages.
- Veselin Stoyanov and Jason Eisner. [Learning cost-aware, loss-aware approximate inference policies for probabilistic graphical models](#). In *COST: NIPS 2011 Workshop on Computational Trade-offs in Statistical Learning*, 2011. 5 pages.
- Nicholas Andrews and Jason Eisner. [Transformation process priors](#). In *NIPS 2011 Workshop on Bayesian Nonparametrics: Hope or Hype?*, 2011. Extended abstract (3 pages).
- Matthew R. Gormley, Mark Dredze, Benjamin Van Durme, and Jason Eisner. [Shared components topic models with application to selectional preference](#). In *NIPS 2011 Workshop on Learning Semantics*, 2011. Extended abstract (3 pages).
- Markus Dreyer and Jason Eisner. [Discovering morphological paradigms from plain text using a Dirichlet process mixture model](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 616–627, 2011. Supplementary material (9 pages) also available.

- Zhifei Li, Jason Eisner, Ziyuan Wang, Sanjeev Khudanpur, and Brian Roark. [Minimum imputed risk: Unsupervised discriminative training for machine translation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 920–929, 2011.
- Veselin Stoyanov, Alexander Ropson, and Jason Eisner. [Empirical risk minimization of graphical model parameters given approximate inference, decoding, and model structure](#). In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS), JMLR Workshop and Conference Proceedings*, pages 15:725–733, 2011. Supplementary material (4 pages) also available.
- Zhifei Li, Ziyuan Wang, Sanjeev Khudanpur, and Jason Eisner. [Unsupervised discriminative language model training for machine translation using simulated confusion sets](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (COLING)*, pages 656–664, 2010.
- Zhifei Li and Jason Eisner. [First- and second-order expectation semirings with applications to minimum-risk training on translation forests](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 40–51, 2009.
- Markus Dreyer and Jason Eisner. [Graphical models over multiple strings](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 101–110, 2009.
- David A. Smith and Jason Eisner. [Parser adaptation and projection with quasi-synchronous grammar features](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 822–831, 2009.
- Roy Tromble and Jason Eisner. [Learning linear ordering problems for better translation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1007–1016, 2009.
- Zhifei Li, Jason Eisner, and Sanjeev Khudanpur. [Variational decoding for statistical machine translation](#). In *Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 593–601, 2009.
- James Mayfield, David Alexander, Bonnie Dorr, Jason Eisner, Tamer Elsayed, Tim Finin, Clay Fink, Marjorie Freedman, Nikesh Garera, Paul McNamee, Saif Mohammad, Douglas Oard, Christine Piatko, Asad Sayeed, Zareen Syed, Ralph Weischedel, Tan Xu, and David Yarowsky. [Cross-document coreference resolution: A key technology for learning by reading](#). In *Proceedings of the AAAI 2009 Spring Symposium on Learning by Reading and Learning to Read*, 2009.
- Omar F. Zaidan, Jason Eisner, and Christine Piatko. [Machine learning with annotator rationales to reduce annotation cost](#). In *Proceedings of the NIPS\*2008 Workshop on Cost Sensitive Learning*, 2008. 10 pages.
- David A. Smith and Jason Eisner. [Dependency parsing by belief propagation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 145–156, 2008.

- Omar F. Zaidan and Jason Eisner. [Modeling annotators: A generative approach to learning from annotator rationales](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 31–40, 2008.
- Markus Dreyer, Jason R. Smith, and Jason Eisner. [Latent-variable modeling of string transductions with finite-state methods](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1080–1089, 2008.
- Jason Eisner and Noah A. Smith. [Competitive grammar writing](#). In *Proceedings of the Third Workshop on Issues in Teaching Computational Linguistics*, pages 97–105, 2008.
- Damianos Karakos, Jason Eisner, Sanjeev Khudanpur, and Markus Dreyer. [Machine translation system combination using ITG-based alignments](#). In *Proceedings of ACL-08: HLT, Short Papers*, pages 81–84, 2008.
- David A. Smith and Jason Eisner. [Bootstrapping feature-rich dependency parsers with entropic priors](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 667–677, 2007.
- Omar Zaidan, Jason Eisner, and Christine Piatko. [Using “annotator rationales” to improve machine learning for text categorization](#). In *Human Language Technologies: Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, pages 260–267, 2007.
- Damianos Karakos, Jason Eisner, Sanjeev Khudanpur, and Carey E. Priebe. [Cross-instance tuning of unsupervised document clustering algorithms](#). In *Human Language Technologies: Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, pages 252–259, 2007.
- Damianos Karakos, Sanjeev Khudanpur, Jason Eisner, and Carey E. Priebe. [Iterative denoising using Jensen-Renyí divergences with an application to unsupervised document categorization](#). In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2007.
- Jason Eisner and John Blatz. [Program transformations for optimization of parsing algorithms and other weighted logic programs](#). In Shuly Wintner, editor, *Proceedings of FG 2006: The 11th Conference on Formal Grammar*, pages 45–85. CSLI Publications, 2007.
- Joshua Mason, Kathryn Watkins, Jason Eisner, and Adam Stubblefield. [A natural-language approach to automated cryptanalysis of two-time pads](#). In *Proceedings of the ACM Conference on Computer and Communications Security (ACM CCS)*, pages 235–244, 2006.
- Jason Eisner and John Blatz. [Program transformations for optimization of parsing algorithms and other weighted logic programs](#). In Paola Monachesi, Gerald Penn, Giorgio Satta, and Shuly Wintner, editors, *Pre-proceedings of the 11th Conference on Formal Grammar (FG-2006)*, pages 39–59, 2006. Superseded by longer post-proceedings version.
- Markus Dreyer and Jason Eisner. [Better informed training of latent syntactic features](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 317–326, 2006.

- David A. Smith and Jason Eisner. [Minimum-risk annealing for training log-linear models](#). In *Proceedings of the International Conference on Computational Linguistics and the Association for Computational Linguistics (COLING-ACL), Companion Volume*, pages 787–794, 2006.
- Noah A. Smith and Jason Eisner. [Annealing structural bias in multilingual weighted grammar induction](#). In *Proceedings of the International Conference on Computational Linguistics and the Association for Computational Linguistics (COLING-ACL)*, pages 569–576, 2006.
- Jason Eisner and Roy W. Tromble. [Local search with very large-scale neighborhoods for optimal permutations in machine translation](#). In *Proceedings of the HLT-NAACL Workshop on Computationally Hard Problems and Joint Inference in Speech and Language Processing*, pages 57–75, 2006.
- David A. Smith and Jason Eisner. [Quasi-synchronous grammars: Alignment by soft projection of syntactic dependencies](#). In *Proceedings of the HLT-NAACL Workshop on Statistical Machine Translation*, pages 23–30, 2006.
- Roy W. Tromble and Jason Eisner. [A fast finite-state relaxation method for enforcing global constraints on sequence decoding](#). In *Proceedings of the Human Language Technology Conference of the North American Association for Computational Linguistics (HLT-NAACL)*, pages 423–430, 2006.
- Jason Eisner and Noah A. Smith. [Parsing with soft and hard constraints on dependency length](#). In *Proceedings of the International Workshop on Parsing Technologies (IWPT)*, pages 30–41, 2005.
- Jason Eisner and Damianos Karakos. [Bootstrapping without the boot](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT-EMNLP)*, pages 395–402, 2005.
- Jason Eisner, Eric Goldlust, and Noah A. Smith. [Compiling comp ling: Weighted dynamic programming and the Dyna language](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT-EMNLP)*, pages 281–290, 2005.
- Noah A. Smith and Jason Eisner. [Guiding unsupervised grammar induction using contrastive estimation](#). In *International Joint Conference on Artificial Intelligence (IJCAI) Workshop on Grammatical Inference Applications*, pages 73–82, 2005.
- Noah A. Smith and Jason Eisner. [Contrastive estimation: Training log-linear models on unlabeled data](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 354–362, 2005.
- André Kempe, Jean-Marc Champarnaud, Jason Eisner, Franck Guingne, and Florent Nicart. [A class of rational  \$n\$ -WFSM auto-intersections](#). In *Proceedings of the Tenth International Conference on Implementation and Application of Automata (CIAA-2005)*, number 3845 in Lecture Notes in Computer Science, pages 189–200. Springer-Verlag, 2005.
- André Kempe, Jean-Marc Champarnaud, and Jason Eisner. [A note on join and auto-intersection of  \$n\$ -ary rational relations](#). In Loek Cleophas and Bruce Watson, editors,

- Proceedings of the Eindhoven FASTAR Days (Computer Science Technical Report 04-40)*, pages 64–78. Department of Mathematics and Computer Science, Technische Universiteit Eindhoven, Netherlands, 2004.
- Jason Eisner, Eric Goldlust, and Noah A. Smith. [Dyna: A declarative language for implementing dynamic programs](#). In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL), Companion Volume*, pages 218–221, 2004.
- Noah A. Smith and Jason Eisner. [Annealing techniques for unsupervised statistical language learning](#). In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 486–493, 2004.
- Jason Eisner. [Learning non-isomorphic tree mappings for machine translation](#). In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics (ACL), Companion Volume*, pages 205–208, 2003.
- Jason Eisner. [Simpler and more general minimization for weighted finite-state automata](#). In *Proceedings of the Joint Meeting of the Human Language Technology Conference and the North American Chapter of the Association for Computational Linguistics (HLT-NAACL)*, pages 64–71, 2003.
- Jason Eisner. [Parameter estimation for probabilistic finite-state transducers](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1–8, 2002.
- Jason Eisner. [Comprehension and compilation in Optimality Theory](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 56–63, 2002.
- Jason Eisner. [An interactive spreadsheet for teaching the forward-backward algorithm](#). In Dragomir Radev and Chris Brew, editors, *Proceedings of the ACL Workshop on Effective Tools and Methodologies for Teaching NLP and CL*, pages 10–18, 2002.
- Jason Eisner. [Transformational priors over grammars](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 63–70, 2002.
- Jason Eisner. [Expectation semirings: Flexible EM for finite-state transducers](#). In Gertjan van Noord, editor, *Proceedings of the ESSLLI Workshop on Finite-State Methods in Natural Language Processing (FSMNL)*, 2001. Extended abstract (5 pages).
- Jason Eisner. [Easy and hard constraint ranking in Optimality Theory: Algorithms and complexity](#). In Jason Eisner, Lauri Karttunen, and Alain Thériault, editors, *Finite-State Phonology: Proceedings of the 5th Workshop of the ACL Special Interest Group in Computational Phonology (SIGPHON)*, pages 22–33, 2000.
- Jason Eisner. [Directional constraint evaluation in Optimality Theory](#). In *Proceedings of the 18th International Conference on Computational Linguistics (COLING 2000)*, pages 257–263, 2000.
- Jason Eisner and Giorgio Satta. [A faster parsing algorithm for lexicalized tree-adjoining grammars](#). In *Proceedings of the 5th Workshop on Tree-Adjoining Grammars and Related Formalisms (TAG+5)*, pages 14–19, 2000.

- Jason Eisner and Giorgio Satta. [Efficient parsing for bilexical context-free grammars and head-automaton grammars](#). In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 457–464, 1999.
- Jason Eisner. [FOOTFORM decomposed: Using primitive constraints in OT](#). In Benjamin Bruening, editor, *Proceedings of SCIL VIII*, number 31 in MIT Working Papers in Linguistics, pages 115–143, 1998.
- Jason Eisner. [Bilexical grammars and a cubic-time probabilistic parser](#). In *Proceedings of the 5th International Workshop on Parsing Technologies (IWPT)*, pages 54–65, 1997.
- Jason Eisner. [Efficient generation in primitive Optimality Theory](#). In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 313–320, 1997.
- Jason Eisner. [Three new probabilistic models for dependency parsing: An exploration](#). In *Proceedings of the 16th International Conference on Computational Linguistics (COLING-96)*, pages 340–345, 1996.
- Jason Eisner. [Efficient normal-form parsing for combinatory categorial grammar](#). In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 79–86, 1996.
- Breck Baldwin, Jeff Reynar, Mike Collins, Jason Eisner, Adwait Ratnaparkhi, Joseph Rosenzweig, Anoop Sarkar, and Srinivas. [Description of the University of Pennsylvania entry in the MUC-6 competition](#). In *Proceedings of the Sixth Message Understanding Conference*, pages 177–191, 1995.
- Jason Eisner.  [\$\forall\$ -less in Wonderland? Revisiting \*any\*](#). In Janet Fuller, Ho Han, and David Parkinson, editors, *Proceedings of ESCOL 11 (October 1994)*, pages 92–103. DMLL Publications, 1995.
- Mark A. Jones and Jason M. Eisner. [A probabilistic parser applied to software testing documents](#). In *Proceedings of National Conference on Artificial Intelligence (AAAI-92)*, pages 322–328, 1992.
- Mark A. Jones and Jason M. Eisner. [A probabilistic parser and its application](#). In Carl Weir, editor, *Statistically-Based Natural Language Processing Techniques: Papers from the 1992 Workshop*, pages 20–27. Menlo Park: AAAI Press, 1992. Technical Report W-92-01.

### Refereed Presentations

- Jason Eisner. [Dyna: A non-probabilistic programming language for probabilistic AI](#). Extended abstract for talk at the NIPS\*2008 Workshop on Probabilistic Programming, 2008.
- Jason Eisner, Michael Kornbluh, Gordon Woodhull, Raymond Buse, Samuel Huang, Constantinos Michael, and George Shafer. [Visual navigation through large directed graphs and hypergraphs](#). In *Proceedings of the IEEE Symposium on Information Visualization (InfoVis'06), Poster/Demo Session*, pages 116–117, 2006.
- Jason Eisner. [What constraints should OT allow?](#) Talk handout available online (22 pages), Linguistic Society of America (LSA), Chicago, 1997.

## Invited Papers

Damianos Karakos, Sanjeev Khudanpur, Jason Eisner, and Carey E. Priebe. [Unsupervised classification via decision trees: An information-theoretic perspective](#). In *Proceedings of the 2005 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 5, pages 1081–1084, 2005. Invited talk.

## Technical Reports

Jia Cui and Jason Eisner. [Finite-state Dirichlet allocation: Learned priors on finite-state models](#). Technical Report 53, Center for Language and Speech Processing, Johns Hopkins University, 2006. 18 pages.

Jan Hajič, Martin Čmejrek, Bonnie Dorr, Yuan Ding, Jason Eisner, Daniel Gildea, Terry Koo, Kristen Parton, Gerald Penn, Dragomir Radev, and Owen Rambow. [Natural language generation in the context of machine translation](#). Technical report, Center for Language and Speech Processing, Johns Hopkins University, 2004. Final report from 2002 CLSP summer workshop (87 pages).

Jason Eisner. [State-of-the-art algorithms for minimum spanning trees: A tutorial discussion](#). Manuscript available online (78 pages), University of Pennsylvania, 1997.

Jason Eisner. [An empirical comparison of probability models for dependency grammar](#). Technical Report IRCS-96-11, Institute for Research in Cognitive Science, Univ. of Pennsylvania, 1996. Available online (18 pages).

Jason Eisner. [Indirect STV election: A voting system for South Africa](#). White paper, University of Cape Town, 1991. Available online (16 pages).

Jason Eisner. [Dynamical-systems behavior in recurrent and non-recurrent connectionist nets](#). Undergraduate honors thesis, Harvard University, 1990. Available online (57 pages).

## Outreach (general audience)

Jason Eisner. [The science of language: Computational linguistics](#). *Imagine Magazine*, 7(4):14–15, 2000.

Jason Eisner. [Cognitive science and the search for intelligence](#). Invited paper presented to the Socratic Society, University of Cape Town, South Africa, 1991. Available online (24 pages).

## Edited Volumes

Jason Eisner and Jeffrey Heinz, editors. [Proceedings of the Tenth Meeting of the ACL Special Interest Group on Computational Morphology and Phonology](#). Association for Computational Linguistics, 2008. 57 pages.

Jason Eisner, editor. [Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning \(EMNLP-CoNLL\)](#). Association for Computational Linguistics, 2007. 1220 pages.

Jason Eisner, Lauri Karttunen, and Alain Thériault, editors. *Finite-State Phonology: Proceedings of the 5th Workshop of the ACL Special Interest Group in Computational Phonology (SIGPHON)*. Association for Computational Linguistics, 2007. 67 pages.

## PATENTS

Frederick S.M. Herz, Lyle H. Ungar, Jason M. Eisner, and Walter Paul Labys. [Stock market prediction using natural language processing](#). Patent pending, filed 2002.

Frederick S. M. Herz, Jonathan Smith, Paul Labys, and Jason Michael Eisner. [Method of combining shared buffers of continuous digital media data with media delivery scheduling](#). Patent pending, filed 2001.

Frederick S. M. Herz, Walter Paul Labys, David C. Parkes, Sampath Kannan, and Jason M. Eisner. [Secure data interchange](#). Patent pending, filed 2000.

Frederick Herz, Jason Eisner, Lyle Ungar, Walter Paul Labys, Bernie Roemmele, and Jon Hayward. [System for the automatic determination of customized prices and promotions](#). Patent pending, filed 1998.

Jeffrey C. Reynar, Fred Herz, Jason Eisner, and Lyle Ungar. [A Lempel-Ziv data compression technique utilizing a dictionary pre-filled with frequent letter combinations, words and/or phrases](#). U.S. Patent #5,951,623 issued 9/14/1999, filed 1996.

Frederick S. M. Herz, Jason M. Eisner, and Lyle H. Ungar. [System for generation of object profiles for a system for customized electronic identification of desirable objects](#). U.S. Patent #5,835,087 issued 11/10/1998, filed 1995.

Frederick S. M. Herz, Jason M. Eisner, Lyle H. Ungar, and Mitchell P. Marcus. [System for generation of user profiles for a system for customized electronic identification of desirable objects](#). U.S. Patent #5,754,939 issued 5/19/1998, filed 1995.

Frederick S. M. Herz, Jason Eisner, and Marcos Salganicoff. [Pseudonymous server for system for customized electronic identification of desirable objects](#). U.S. Patent #5,754,938 issued 5/19/1998, filed 1995.

Frederick S. M. Herz, Jason M. Eisner, Jonathan M. Smith, and Steven L. Salzberg. [System for customized electronic identification of desirable objects](#). Patent pending, filed 1995.

## TEACHING

Department of Computer Science, Johns Hopkins University

Robert B. Pond, Sr. Excellence in Teaching Award, Whiting School of Engineering, 2005

- *Natural Language Processing* 2001–2004, 2006–  
A mixed graduate-undergraduate class that teaches a synthesis of statistical models, formal grammars, and linguistic theory, with associated algorithms. It is reputed to be one of the most challenging classes in the Computer Science department, requiring both rigor and intellectual flexibility.  
Faculty at several other universities have asked to use the extensive online course materials.  
Enrollment: about 30.

- *Declarative Methods*
2005–
- A new course for juniors, seniors, and graduate students. It surveys computational problems that tend to pop up frequently in different guises (e.g., constraint satisfaction); the specification languages used to describe instances of these problems; general toolkits for solving these instances; and the algorithms run by these toolkits.

Enrollment: about 35.
- *Seminar in Natural Language Processing*
every semester
- A weekly reading and discussion group, exploring important current research in natural language processing and potentially relevant material from related fields. Topics are chosen by the group; each lasts 3–4 weeks.

Attendance: 5–15.
- *Totally Random*
2004, 2005
- A 4-class discussion unit about random numbers and the uses of randomness in computer science. Part of the department’s new freshman experience course.

Enrollment: 8-10.
- *Data Structures*
2003, 2004
- A sophomore-level class, the third and last in the programming sequence for majors. Covers basic data structures and algorithms, basic analysis of algorithms, and object-oriented programming style. Online ”warmups” and highly interactive classes stimulated the students to come up with designs and variations. The challenging weekly assignments often used real-world data.

Faculty at a dozen other universities have asked to use the course materials.

Enrollment: about 50.
- *Finite-State Methods in Natural Language Processing*
2000–2001
- A graduate class on semiring-weighted finite-state transducers. Covers theory and practice, including the theory of formal power series, the extended regular expression calculus, and a range of applications to natural language. Rigorous assignments.

Attendance: about 20.
- *Statistical Language Learning*
2002
- A graduate class about past and present research that has attempted to induce the structure of language from raw data such as text. Lectures are intermixed with reading and discussion of the primary literature.

Attendance: about 10.
- Lecturer, 1st Lisbon Machine Learning School
2011
- Lectures on generative Bayesian modeling.
- Lecturer, NAACL Summer School in Human Language Technology
summers 2002–
- Lectures on NLP and applied probability.
  - *Laboratory exercise* in competitive grammar writing.
- Speaker, TA Training Institute, Johns Hopkins University
2005–

- Large fall lectures: “Preparing for the first day in sciences and engineering.”
- Small spring workshop (some years): “Starting the semester off right in engineering and sciences.”

Department of Computer Science, University of Rochester

- *Statistical Learning of Natural Language* 2000
- *Graduate Problem Seminar* 2000  
 Boot camp for new Ph.D. students. Students learn research skills by teaming up to tackle a series of open-ended engineering problems that touch on research in the department. (I made them build systems for face orientation detection, distributed calendar management, and information retrieval.) Several written and oral presentations are required and receive extensive feedback. The class also includes career advice, familiarization with departmental resources, presentations by other faculty, and a final research project.  
 Enrollment: 10.

Department of Computer Science, University of Pennsylvania  
 Graduate Teaching Award, 1995

- TA in *Introduction to Programming* 1994–1995

ADVISING

**Postdoctoral fellows**

Veselin Stoyanov (*NSF Computing Innovation Fellow*) 2009–2011

**Ph.D. research students**

Tim Vieira 2011–  
 Michael Paul (*NSF Fellow; Dean’s Fellow*) 2010–  
 Adam Teichert 2010–  
 Matthew Gormley (*HLTCOE Fellow*) 2009–  
 Nicholas Andrews 2009–  
 Zhifei Li (co-adv.) (*now at Google*) 2008–2010  
 Nathaniel W. Filardo (*HLTCOE Fellow*) 2007–  
 Jason A. Smith 2006–2011  
 Omar Zaidan 2005–2008  
 John Blatz 2004–2008  
 Eric Goldlust (*Wolman Fellow; Muuss Research Award; Hon. Mention for CRA Outstanding Undergraduate Award*) 2004–2005  
 Markus Dreyer (*Wolman Fellow*) 2003–2011  
 David A. Smith (*NSF Fellow; now Asst. Res. Prof., UMass*) 2002–2010  
 Roy Tromble (*NSDEG Fellow; now at Google*) 2002–2009  
 Noah A. Smith (*Hertz Foundation Fellow; now Assoc. Prof., CMU*) 2001–2006

**H.S./B.S./M.S.E. research students**

Katherine Wu 2012–

Jay Feldman	2012–
Michael Tontchev ( <i>High school—Grand Prize Winner, Baltimore Science Fair</i> )	2010–
Alex Ropson	2010–2011
Ashish Sharma	2009–2010
Wren Thornton	2008
Ian Nowland	2008–2009
Jay Van Der Wall	2007–2009
Asheesh Laroia	2006–2007
Samuel Huang ( <i>Pistritto Fellow</i> )	2006–2007
Constantinos Michael	2005–2006
George Shafer	2005–2006
John Graettinger ( <i>Outstanding Research Award</i> )	2005–2006
Michael Kornbluh ( <i>Pistritto Fellow, Outstanding Senior Award</i> )	2004–2005
Chalaporn Hathaidharm	2002–2004

**Ph.D. thesis committees (not as advisor)**

Jason Smith	exp. 2013
Will Headden (Brown University)	2011
Jia Cui	2008
Mark Thober	2007
Gaja Jarosz Snover	2006
Jonathan Allen	2006
Gideon Mann	2006
Charles Schafer	2006
Myroslava Dzikovska (Univ. of Rochester)	2004
Silviu Cucerzan	2003
Radu “Hans” Florian	2002
Jun Wu	2002
Richard Wicentowski	2002
Grace Ngai	2000

February 3, 2012