

## Curriculum Vitae

1

Johns Hopkins University Assistant Professor of Computer Science	July 2000–June 2007
University of Rochester Assistant Professor of Computer Science Secondary appointment in Linguistics	Jan. 2000–June 2001
iReactor Inc., Philadelphia, PA Consultant	1994–2004
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

**Journals**

- Associate editor, *Journal of Artificial Intelligence Research* (2012–2018).
- Action editor, *Transactions of the Association for Computational Linguistics* (2012–2018).
- Organizing committee, Jelinek Summer Workshop on Speech and Language Technology (JSALT)—an 8-week annual event (2001–)
- Editorial board member, *Linguistic Issues in Language Technology* (2014–2022).
- Editorial board member, *Research in Language and Computation* (2006–2011).
- Guest editorial board member (special issue), *Journal of Natural Language Engineering* (2011).
- Special issue editor, *Cognitive Science* (2002).
- Journal reviewer, *J. Applied Logic* (2016), *IEEE Trans. on Pattern Matching and Machine Intelligence* (2014, 2012), *J. Machine Learning Research* (2012), *IEEE Signal Processing Letters* (2012), *Foundations and Trends* (2012), *J. Logic and Computation* (2012, 2011), *Computational Linguistics* (2024, 2021, 2017, 2016, 2011, 2010, 2007, 2005, 2004, 2002, 2001, 2000), *Phonology* (2010, 2008), *Research on Language and Computation* (2008, 2007), *ACM Trans. on Speech and Language Processing* (2007, 2005), *Communications of the ACM* (2025), *IEEE Transactions on Audio, Speech, and Language Processing* (2007), *Cognition* (2002), *J. Algorithms* (1997), *Language and Speech* (1999).

**Conferences**

- Diversity & Inclusion co-chair, NAACL (2019).

- Program area chair or co-chair for ACL (2017, morphology/phonology; 2014, machine learning; 2013, morphology/phonology), COLING (2014, machine learning); 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), NeurIPS (2025). Guest editor for PNAS (2018).
- Chair of Best Paper Award committee, EMNLP (2010). Member of Best Short Paper Award committee, NAACL (2012).
- Program chair, [EMNLP-CoNLL](#) (2007).
- Program chair, [SIGMORPHON Workshop](#) (2008, with Jeff Heinz); SIGPHON Workshop on Finite-State Phonology (2000).
- Organizer, NSF-funded international workshop on [Probabilistic Representations of Linguistic Meaning](#) (2014).
- Publications chair, ACL (2005).
- ICML Workshop Review Committee (2021); co-chair for HLT/NAACL Tutorial and Workshop Programs (2003).
- Organizing committee, ACL Workshop on Deep Learning and Formal Languages (2019).
- Program committee member / reviewer for ACL (2025, 2023, 2020, 2017, 2016, 2014, 2013, 2012, 2011, 2010, 2009, 2007, 2005, 2004, 2003, 2002), AISTATS (2010), COLING (2014, 2012, 2008), EAAI (2013), EACL (2012, 2006, 2003, 1999), EMNLP (2018, 2017, 2016, 2014, 2013, 2012, 2011, 2010, 2009, 2008, 2007, 2006, 2003, 2002 – “best reviewer” awards in 2018, 2010), FSMNLP (2005, 2001), ICGI (2012), ICFP (2008), ICLR (2017), ICML (2019, 2016, 2004), IJCAI (2007), IWPT (2009), MITWPL (1999), NAACL (2016, 2015, 2013, 2012, 2010, 2009, 2006), NeurIPS (2021, 2020, 2019, 2018, 2017, 2016, 2014, 2013, 2011, 2010, 2007), NLP-LING (2010), SCiL (2020, 2018), SIGMORPHON (2023, 2019, 2018, 2014, 2012, 2010, 2008, 2006, 2004, 2002, 2000, 1998), WiNLP—Widening NLP Workshop (2023, 2021, 2019, 2018, 2017), ACL Workshop on Statistical NLP and Weighted Automata (StatFSM) (2016), ACL Workshop on Relevance of Linguistic Structure in Neural Architectures for NLP (RELNLP) (2018), ACL Workshop on Unsupervised Learning (2011), ACL Workshops on Teaching NLP and CL (2024, 2013, 2008, 2005, 2002), CVPR Workshop on Structured Prediction (2013), ICML Workshop on Prediction with Sequential Models (2013), International Conference on Dependency Linguistics (2011), Workshop on Formal Approaches to Slavic Linguistics 8 (1999).

#### **Long-Term Committees**

- ACL Doctoral Dissertation Award Committee (2025, 2026).
- Board member, Widening NLP (WiNLP) (2016–2025).

- Organizing committee member, [SIGDAT](#) (which runs the EMNLP conference) (2007–).
- Working group on policies for anonymity, ACL (2017, 2023).
- Co-chair of Test-of-Time Paper Award Committee, ACL (2022, 2023).
- Problems committee, North American Computational Linguistics Olympiad (2006–2013).
- President, ACL SIGMORPHON (Computational Morphology and Phonology) (2001–2021); executive committee (1998–2001).

#### Other

- Judge for “NLP Idol” special event at NAACL (2012). Panel moderator for NeurIPS Workshop on Interpretability and Robustness in Audio, Speech, and Language (2018).
- Mentor for NAACL Student Research Workshop (2013, 2012, 2009).
- NSF proposal reviewer (2022, 2021, 2019, 2018, 2017, 2014, 2013, 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

Fellow of the Association of Computational Linguistics	2018
Outstanding Paper Award (EMNLP)	2024
Outstanding Paper Award (ACL)	2022
Best Short Paper Award (NAACL)	2021
Best Paper Award (EMNLP-IJCNLP)	2019
Best Long Paper Award (ACL)	2017
Finalist for 5-year retrospective Best Paper Award (SMT Workshop)	2011
Finalist for Best Paper Award (ACL, EMNLP-CoNLL, ACL, EMNLP)	2009, 2007, 2005, 2002
Quora Top Writer	2016, 2017
Best Featured Actor in a Musical DC Metro Theater Arts, Best of 2014	2014
Johns Hopkins Alumni Association <a href="#">Excellence in Teaching Award</a>	2013, 2025
Robert B. Pond, Sr. <a href="#">Excellence in Teaching Award</a> 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

#### UNIVERSITY SERVICE

AI-X Steering Committee (2019–2023).

Common Question Committee (2019–).

Faculty member, [Center for Language and Speech Processing](#) (2000–).

Executive committee, [Mathematical Institute for Data Science](#) (2017–).

Curriculum committee, Computer Science Department (2016–2019 and 2023–, as chair 2017–2019).

Director of Graduate Studies, Computer Science Department (2002–2013).

Onstage presenter for various Homewood events: admissions open houses and parents’ weekends (spring 2013, fall 2013, spring 2014, fall 2014, fall 2017, spring 2018, fall 2018, fall 2019, spring 2023, spring 2024), ‘Hopkins Engineering Sampler Seminar (2013), Dissertation Writing Workshop (2013, 2016), QuarkNet (2013), WSE centennial video (2013), TA Training Institute (2005–2019).

Mentor for multiple junior faculty members.

Service on tenure and promotion committees.

Search committees: Computer Science lecturer (2019), Bloomberg Distinguished Professorship in Computational Cognitive Science (2014, co-chair), Bloomberg Distinguished Professorship in Computational Healthcare (2014), CS machine learning subcommittee (2010–2014 and thereafter), Director of Admissions (2012), Applied Math and Statistics (2012, 2011), HLTCOE (2008–2019).

Chair of education planning committee, Computer Science Department (2013–2014).

Excellence in Teaching awards committee, Whiting School of Engineering (2014).

Judge for the HopHacks hackathon (2014, 2013).

PI of two NSF IGERT proposals involving 25+ faculty across campus (2011, 2010).

External review committees for Cognitive Science and PBS departments (2011).

Advisory committee member, Zooniverse project of the Adler Planetarium and JHU Space Telescope Science Institute (2009–2012).

Advisory board member, AMS GAANN Fellowship Program (2009–2012).

Advisory board member, cogito.org (2005–2016).

CS/CLSP admissions committee (2001–).

*See also “Teaching” section.*

GRANTS	DOE: Theseus: A Computational Science Foundation Model (co-PI, \$7M)	2024–2027
	JSALT: AI-Curated Democratic Discourse (PI)	summer 2024
	NSF RI-Small: Linguistic Structure in Neural Sequence Models (PI, \$395K)	2017–2020
	Amazon Research Award: Continuous-Time Reinforcement Learning For Personalization (PI, \$99K)	2018
	Google Faculty Research Award: Neural Modeling of String-to-String Transductions (PI, \$79K)	2017
	NVIDIA Corporation GPU Grant	2016
	XPS:FULL:Collaborative Research: Parallel and Distributed Circuit Programming for Structured Prediction (PI, \$825K)	2016–2019
	Facebook Faculty Gift (PI, \$50K)	2016
	NSF RI-Small: CompCog: Modeling Latent Discrete Knowledge Across Utterances (PI, \$500K)	2014–2017
	Science of Learning: Innovative Technology for Personalized Foreign Language Education (co-PI, \$200K)	2014–2016
	NSF RI-Medium: Learned Dynamic Prioritization (PI, \$900K)	2010–2014
	NSF PIRE: Investigation of Meaning Representations in Language Understanding for Machine Translation Systems (co-PI, \$2.5M)	2005–2014
	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)	2007–2012
	JHU WSE-APL Partnership Fund: Learning with Less (PI, \$68K)	2006–2007
	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

Work may be browsed by topic at <http://cs.jhu.edu/~jason/papers>.

### Invited talks

- Propositions as parameters. Simons Institute for the Theory of Computing, September 2024.
- Putting planning and reasoning inside language models. Toyota Technological Institute at Chicago, May 2023.
- Dataflow for dialog. 4th Annual Conference on Automated Knowledge Base Construction, November 2022.
- Imposing structure on sequence models. NLPCC (keynote talk), September 2022.
- Is “Introduction to NLP” just “Machine Learning: NLP Edition”? Teaching NLP Workshop at NAACL, June 2021.
- Task-oriented Dialogue as Dataflow Synthesis: A deeper dive. (With Jacob Andreas, Dan Klein, and Percy Liang.) Workshop on Interactive and Executable Semantic Parsing, EMNLP, November 2020.
- Event sequence modeling with the neural Hawkes process. Max Planck Institute for Mathematics in the Sciences (MiS), Leipzig, July 2020; Institute for Advanced Study, Princeton, August 2020; ByteDance, May 2021; George Mason University, October 2021.
- What are natural languages optimized for? NeurIPS Workshop on Emergent Communication (invited talk), December 2019.
- Sequential structure meets hierarchical structure. Carnegie Mellon University, November 2019.
- Improving cross-linguistic robustness by training on synthetic languages. ACL Workshop on Typology for Polyglot NLP (invited talk), August 2019.
- Modeling event streams. NAACL Workshop on Structured Prediction for NLP (invited talk), June 2019.
- [Discrete latent variables in NLP: good, bad, and indifferent](#). ACL Workshop on Relevance of Linguistic Structure in Neural Architectures for NLP (invited talk), August 2018.
- [Recovering syntactic structure from surface features](#). Pennsylvania State University, January 2018. Similar talks at Saarland University, January 2019; Georgetown University (Distinguished Lecture Series), April 2019; Brown University (Eugene Charniak Memorial Workshop), November 2024.

[Probabilistically modeling surface patterns using latent structure](#). Symposium on Perceptrons and Syntactic Structures at 60: Computational Modeling of Language (invited talk), at Society for Computation in Linguistics (SCiL), January 2018.

Strategies for discovering underlying linguistic structure. 8th International Joint Conference on Natural Language Processing (IJCNLP), November 2017. ([invited talk](#))

Probabilistic modeling of natural language. Data Science Seminar, Johns Hopkins University, March 2017.

Tractable deep models of sequential structure. LTI Colloquium, Carnegie Mellon University, March 2017. Similar talk at NeurIPS Workshop on Interpretability and Robustness in Audio, Speech, and Language (invited talk), December 2018.

Methods to discover underlying linguistic structure. NLP and Text-as-Data Speaker Series, New York University, March 2017. Similar invited talk at Duke University, October 2018.

Probabilistic models of related strings. ACL Workshop on Statistical NLP and Weighted Automata (StatFSM) (keynote talk), August 2016.

[Gradually learning to read a foreign language: Adaptive partial machine translation](#). Science of Learning Symposium at Johns Hopkins University ([keynote talk](#)), January 2016. Similar talks at Mount St. Mary's College, April 2016; ACL Workshop on Cognitive Modeling for Language Learning (keynote talk), August 2016; Duolingo, November 2019.

[Graphical models over string-valued random variables](#). IEEE ASRU ([keynote talk](#)), December 2015.

Probabilistic inference on strings. [Triangle Computer Science Distinguished Lecture Series](#) (UNC+Duke+NCSU), November 2015. Similar talks at Hebrew University of Jerusalem, April 2015; Bar-Ilan University, April 2015; Microsoft Research, May 2015; University of Illinois at Urbana-Champaign, June 2015; Georgetown University, November 2015.

A non-probabilistic language for probabilistic AI. [Dagstuhl Workshop on Challenges and Trends in Probabilistic Programming](#), April 2015.

Weighted deduction for analyzing natural language and other data. AAAI Fall Symposium: Natural Language Access to Big Data ([invited talk](#)), November 2014.

Learning phonology from surface data: Reconstructing the inputs. ACL Joint Workshop on Morphology/Phonology/Phonetics and Finite-State Methods ([keynote talk](#)), June 2014.

Open problems in computational phonology and morphology. ACL Joint Workshop on Morphology/Phonology/Phonetics and Finite-State Methods (invited panelist), June 2014.

[Deep learning of recursive structure: Grammar induction](#). 1st International Conference on Learning Representations ([keynote talk](#)), May 2013.

Grammar induction: Beyond local search. International Conference on Grammatical Inference (keynote talk), September 2012.

[Learning approximate inference policies for fast prediction](#). ICML Workshop on Infering: Interactions between Inference and Learning (keynote talk), July 2012.



Dyna: A language for propagating and combining information. Workshop on Architectures for Uncertainty in Knowledge at Scale (AUKS), February 2012.

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.

Invited panelist, ACL Workshop on Multiword Expressions, June 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 (keynote talk), July 2009.

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

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

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

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

Jessy Lin, Nicholas Tomlin, Jacob Andreas, and Jason Eisner. [Decision-oriented dialogue for human-AI collaboration](#). *Transactions of the Association for Computational Linguistics (TACL)*, June 2024.

Jason Eisner. [Time-and-space-efficient weighted deduction](#). *Transactions of the Association for Computational Linguistics (TACL)*, 11:960–973, August 2023.

Tim Vieira, Ryan Cotterell, and Jason Eisner. [Automating the analysis of parsing algorithms \(and other dynamic programs\)](#). *Transactions of the Association for Computational Linguistics*, 2022. Accepted for publication.

Semantic Machines, Jacob Andreas, John Bufe, David Burkett, Charles Chen, Josh Clausman, Jean Crawford, Kate Crim, Jordan DeLoach, Leah Dorner, Jason Eisner, Hao Fang, Alan Guo, David Hall, Kristin Hayes, Kellie Hill, Diana Ho, Wendy Iwaszuk, Smriti Jha, Dan Klein, Jayant Krishnamurthy, Theo Lanman, Percy Liang, Christopher H. Lin, Ilya Lintsbakh, Andy McGovern, Aleksandr Nisnevich, Adam Pauls, Dmitriy Petters, Brent Read, Dan Roth, Subhro Roy, Jesse Rusak, Beth Short, Div Slomin, Ben Snyder, Stephon Striplin, Yu Su, Zachary Tellman, Sam Thomson, Andrei Vorobev, Izabela Witoszko, Jason Wolfe, Abby Wray, Yuchen Zhang, and Alexander Zotov. [Task-oriented dialogue as dataflow synthesis](#). *Transactions of the Association for Computational Linguistics*, 8:556–571, September 2020.

Xiang Lisa Li, Dingquan Wang, and Jason Eisner. [A generative model for punctuation in dependency trees](#). *Transactions of the Association for Computational Linguistics (TACL)*, 7:357–373, March 2019.

Ryan Cotterell, Christo Kirov, Mans Hulden, and Jason Eisner. [On the complexity and typology of inflectional morphological systems](#). *Transactions of the Association for Computational Linguistics (TACL)*, 7:327–342, March 2019.

Dingquan Wang and Jason Eisner. [Surface statistics of an unknown language indicate how to parse it](#). *Transactions of the Association for Computational Linguistics (TACL)*, 6:667–685, December 2018.

- Dingquan Wang and Jason Eisner. [Fine-grained prediction of syntactic typology: Discovering latent structure with supervised learning](#). *Transactions of the Association for Computational Linguistics (TACL)*, 5:147–161, June 2017.
- Tim Vieira and Jason Eisner. [Learning to prune: Exploring the frontier of fast and accurate parsing](#). *Transactions of the Association for Computational Linguistics (TACL)*, 5:263–278, August 2017.
- Dingquan Wang and Jason Eisner. [The Galactic Dependencies treebanks: Getting more data by synthesizing new languages](#). *Transactions of the Association for Computational Linguistics (TACL)*, 4:491–505, September 2016.
- Matthew R. Gormley, Mark Dredze, and Jason Eisner. [Approximation-aware dependency parsing by belief propagation](#). *Transactions of the Association for Computational Linguistics (TACL)*, 3:489–501, August 2015.
- Ryan Cotterell, Nanyun Peng, and Jason Eisner. [Modeling word forms using latent underlying morphs and phonology](#). *Transactions of the Association for Computational Linguistics (TACL)*, 3:433–447, August 2015.
- Francisco Sánchez-Vega, Jason Eisner, Laurent Younes, and Donald Geman. [Learning multivariate distributions by competitive assembly of marginals](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 99, April 2012. 14 pages plus 32-page supplement.
- 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, May-June 2002.
- Jason Eisner. [Introduction to the special section on linguistically apt statistical methods](#). *Cognitive Science*, 26(3):235–237, May-June 2002.

## Book Chapters

- Jason Eisner and Nathaniel W. Filardo. [Dyna: Extending Datalog for modern AI](#). In Oege de Moor, Georg Gottlob, Tim Furche, and Andrew Sellers, editors, *Datalog Reloaded*, volume 6702 of *Lecture Notes in Computer Science*, pages 181–220. Springer, 2011. Longer version available as tech report.
- 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, October 2000.

### Book Reviews

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

### Ph.D. Thesis

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

### Refereed Conference and Workshop Proceedings

Benjamin Lipkin, Benjamin LeBrun, Jacob Hoover Vigly, João Loula, David R. MacIver, Li Du, Jason Eisner, Ryan Cotterell, Vikash Mansinghka, Timothy J. O’Donnell, Alexander K. Lew, and Tim Vieira. [Fast controlled generation from language models with adaptive weighted rejection sampling](#). In *Proceedings of the Conference on Language Modeling (COLM)*, Montreal, October 2025.

Nishant Subramani, Jason Eisner, Justin Svegliato, Benjamin Van Durme, Yu Su, and Sam Thomson. [MICE for CATs: Model-internal confidence estimation for calibrating agents with tools](#). In *Proceedings of the North American Conference on Computational Linguistics*, Albuquerque, April 2025.

João Loula, Benjamin LeBrun, Li Du, Ben Lipkin, Clemente Pasti, Gabriel Grand, Tianyu Liu, Yahya Emara, Marjorie Freedman, Jason Eisner, Ryan Cotterell, Vikash Mansinghka, Alexander K. Lew, Tim Vieira, and Timothy J. O’Donnell. [Syntactic and semantic control of large language models via sequential Monte Carlo](#). In *Proceedings of the International Conference on Learning Representations*, Singapore, April 2025.

Yunmo Chen, Tongfei Chen, Harsh Jhamtani, Patrick Xia, Richard Shin, Jason Eisner, and Benjamin Van Durme. [Learning to retrieve iteratively for in-context learning](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7156–7168, November 2024.

Kaj Bostrom, Harsh Jhamtani, Hao Fang, Sam Thomson, Richard Shin, Patrick Xia, Benjamin Van Durme, Jason Eisner, and Jacob Andreas. [Language-to-code translation with a single labeled example](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8101–8112, November 2024.

Helia Hashemi, Jason Eisner, Corby Rosset, Benjamin Van Durme, and Chris Kedzie. [LLM-Rubric: A multidimensional, calibrated approach to automated evaluation of natural language texts](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 13806–13834, August 2024.

Boshi Wang, Hao Fang, Jason Eisner, Benjamin Van Durme, and Yu Su. [LLMs in the Imaginarium: Tool learning through simulated trial and error](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 10583–10604, August 2024.

- Sky CH-Wang, Benjamin Van Durme, Jason Eisner, and Chris Kedzie. [Do androids know they're only dreaming of electric sheep?](#) In *Findings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4401–4420, August 2024.
- Giovanni Monea, Maxime Peyrard, Martin Josifoski, Vishrav Chaudhary, Jason Eisner, Emre Kiciman, Hamid Palangi, Barun Patra, and Robert West. [A glitch in the Matrix? Locating and detecting language model grounding with Fakepedia.](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 6828–6844, August 2024.
- Li Du, Jason Eisner, Holden Lee, and Ryan Cotterell. [When is a language process a language model?](#) In *Findings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 11083–11094, August 2024.
- Li Du, Afra Amini, Lucas Torroba Hennigen, Xinyan Velocity Yu, Holden Lee, Jason Eisner, and Ryan Cotterell. [Principled gradient-based MCMC for conditional sampling of text.](#) In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, July 2024.
- Nikita Moghe, Patrick Xia, Jacob Andreas, Jason Eisner, Benjamin Van Durme, and Harsh Jhamtani. [Interpreting user requests in the context of natural language standing instructions.](#) In *Findings of the North American Conference on Computational Linguistics (NAACL)*, volume arXiv:2311.09796, pages 4043–4060, June 2024.
- Weiting Tan, Chu-Cheng Lin, and Jason Eisner. [Structure-aware path inference for neural finite state transducers.](#) In *Proceedings of the NeurIPS 2023 Workshop “I Can’t Believe It’s Not Better: Failure Modes in the Age of Foundation Models”*, December 2023.
- Subhro Roy, Sam Thomson, Tongfei Chen, Richard Shin, Adam Pauls, Jason Eisner, and Benjamin Van Durme. [BenchCLAMP: A benchmark for evaluating language models on syntactic and semantic parsing.](#) In *Proceedings of the Thirty-Seventh Conference on Neural Information Processing Systems*, December 2023. Datasets and Benchmarks Track.
- Ruiqi Zhong, Charlie Snell, Dan Klein, and Jason Eisner. [Non-programmers can label programs indirectly via active examples: A case study with text-to-SQL.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5126–5152, December 2023.
- Jie Chi, Brian Lu, Jason Eisner, Peter Bell, Preethi Jyothi, and Ahmed M. Ali. [Unsupervised code-switched text generation from parallel text.](#) In *Proceedings of INTERSPEECH*, Dublin, August 2023.
- Hao Fang, Anusha Balakrishnan, Harsh Jhamtani, John Bufe, Jean Crawford, Jayant Krishnamurthy, Adam Pauls, Jason Eisner, Jacob Andreas, and Dan Klein. [The whole truth and nothing but the truth: Faithful and controllable dialogue response generation with dataflow transduction and constrained decoding.](#) In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5682–5700, July 2023.
- Belinda Z. Li, Jason Eisner, Adam Pauls, and Sam Thomson. [Toward interactive dictation.](#) In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 15319–15338, July 2023.

- Fatemehsadat Mireshghallah, Yu Su, Tatsunori Hashimoto, Jason Eisner, and Richard Shin. [Privacy-preserving domain adaptation of semantic parsers](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 4950–4970, July 2023.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. [Contrastive decoding: Open-ended text generation as optimization](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 12286–12312, July 2023.
- Li Du, Lucas Torroba Hennigen, Tiago Pimentel, Clara Meister, Jason Eisner, and Ryan Cotterell. [A measure-theoretic characterization of tight language models](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 9744–9770, July 2023.
- Andreas Opedal, Ran Zmigrod, Tim Vieira, Ryan Cotterell, and Jason Eisner. [Efficient semiring-weighted Earley parsing](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 3687–3713, July 2023.
- Anej Svete, Benjamin Dayan, Ryan Cotterell, Tim Vieira, and Jason Eisner. [Acyclic weighted finite-state automata with failure transitions](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8289–8305, Abu Dhabi, December 2022.
- Elias Stengel-Eskin, Emmanouil Antonios Platanios, Adam Pauls, Sam Thomson, Hao Fang, Benjamin Van Durme, Jason Eisner, and Yu Su. [When more data hurts: A troubling quirk in developing broad-coverage natural language understanding systems](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11473–11487, Abu Dhabi, December 2022.
- Jiawei Zhou, Jason Eisner, Michael Newman, Emmanouil Anthony Platanios, and Sam Thomson. [Online semantic parsing for latency reduction in task-oriented dialogue](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 1554–1576, Dublin, May 2022.
- Ryan Cotterell and Jason Eisner. [A functionalist account of vowel system typology](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, Dublin, May 2022. Paper was accepted, but we withdrew it in order to add more experiments and analysis before publication.
- Chenghao Yang, Hongyuan Mei, and Jason Eisner. [Transformer embeddings of irregularly spaced events and their participants](#). In *Proceedings of the Tenth International Conference on Learning Representations (ICLR)*, April 2022. 9 pages plus appendices.
- Tim Vieira, Ryan Cotterell, and Jason Eisner. [Searching for more efficient dynamic programs](#). In *Findings of EMNLP’21*, pages 3812–3830, Punta Cana, November 2021.
- Richard Shin, Christopher H. Lin, Sam Thomson, Charles Chen, Subhro Roy, Emmanouil Antonios Platanios, Adam Pauls, Dan Klein, Jason Eisner, and Benjamin Van Durme. [Constrained language models yield few-shot semantic parsers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7699–7715, Punta Cana, November 2021.



- Chu-Cheng Lin, Aaron Jaech, Xin Li, Matt Gormley, and Jason Eisner. [Limitations of autoregressive models and their alternatives](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 5147–5173, Online, June 2021.
- Guanghui Qin and Jason Eisner. [Learning how to ask: Querying LMs with mixtures of soft prompts](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 5203–5212, Online, June 2021. Best Short Paper Award.
- Hongyuan Mei, Tom Wan, and Jason Eisner. [Noise-contrastive estimation for multivariate point processes](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5204–5214, December 2020.
- Hongyuan Mei, Guanghui Qin, Minjie Xu, and Jason Eisner. [Neural Datalog through time: Informed temporal modeling via logical specification](#). In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, July 2020.
- Elizabeth Salesky, Eleanor Chodroff, Tiago Pimentel, Matthew Wiesner, Ryan Cotterell, Alan W. Black, and Jason Eisner. [A corpus for large-scale phonetic typology](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2388–2397, July 2020.
- Matthew Francis-Landau, Tim Vieira, and Jason Eisner. [Evaluation of logic programs with built-ins and aggregation: A calculus for bag relations](#). In *13th International Workshop on Rewriting Logic and Its Applications*, pages 49–63, April 2020. Extended version (27 pages) available on arXiv, October 2020.
- Xiang Lisa Li and Jason Eisner. [Specializing word embeddings \(for parsing\) by information bottleneck](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing*, pages 2744–2754, Hong Kong, November 2019. Best Paper Award.
- Adithya Renduchintala, Philipp Koehn, and Jason Eisner. [Spelling-aware construction of macaronic texts for teaching foreign-language vocabulary](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing*, pages 6439–6444, Hong Kong, November 2019.
- Adithya Renduchintala, Philipp Koehn, and Jason Eisner. [Simple construction of mixed-language texts for vocabulary learning](#). In *Proceedings of the 14th Workshop on Innovative Use of NLP for Building Educational Applications (BEA)*, pages 369–379, Florence, August 2019.
- Sabrina J. Mielke, Ryan Cotterell, Kyle Gorman, Brian Roark, and Jason Eisner. [What kind of language is hard to language-model?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4975–4989, Florence, July 2019.
- Hongyuan Mei, Guanghui Qin, and Jason Eisner. [Imputing missing events in continuous-time event streams](#). In *Proceedings of the 36th International Conference on Machine Learning*, Long Beach, California, June 2019.

- Chu-Cheng Lin, Hao Zhu, Matthew Gormley, and Jason Eisner. [Neural finite-state transducers: Beyond rational relations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 272–283, Minneapolis, June 2019.
- Ekaterina Vylomova, Ryan Cotterell, Tim Baldwin, Trevor Cohn, and Jason Eisner. [Contextualization of morphological inflection](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 2018–2024, Minneapolis, June 2019.
- Sabrina J. Mielke and Jason Eisner. [Spell once, summon anywhere: A two-level open-vocabulary language model](#). In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, pages 6843–6850, Honolulu, January 2019.
- Dingquan Wang and Jason Eisner. [Synthetic data made to order: The case of parsing](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1325–1337, Brussels, November 2018.
- Ryan Cotterell, Christo Kirov, Sabrina J. Mielke, and Jason Eisner. [Unsupervised disambiguation of syncretism in inflected lexicons](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 548–553, New Orleans, June 2018.
- Ryan Cotterell, Sabrina J. Mielke, Jason Eisner, and Brian Roark. [Are all languages equally hard to language-model?](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 536–541, New Orleans, June 2018.
- Chu-Cheng Lin and Jason Eisner. [Neural particle smoothing for sampling from conditional sequence models](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 929–941, New Orleans, June 2018.
- Ryan Cotterell and Jason Eisner. [A deep generative model of vowel formant typology](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 37–46, New Orleans, June 2018.
- Christo Kirov, Ryan Cotterell, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sabrina J. Mielke, Arya D. McCarthy, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. [UniMorph 2.0: Universal morphology](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May 2018.
- Hongyuan Mei and Jason Eisner. [The neural Hawkes process: A neurally self-modulating multivariate point process](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6754–6764, Long Beach, CA, December 2017. First version December 2016 as [arXiv:1612.09328v1](#).
- Adithya Renduchintala, Philipp Koehn, and Jason Eisner. [Knowledge tracing in sequential learning of inflected vocabulary](#). In *Proceedings of the Conference on Natural Language*



- Learning (CoNLL)*, pages 238–247, Vancouver, August 2017.
- Tim Vieira, Matthew Francis-Landau, Nathaniel Wesley Filardo, Farzad Khorasani, and Jason Eisner. [Dyna: Toward a self-optimizing declarative language for machine learning applications](#). In *Proceedings of the First ACM SIGPLAN Workshop on Machine Learning and Programming Languages (MAPL)*, pages 8–17, Barcelona, June 2017. ACM.
- Ryan Cotterell and Jason Eisner. [Probabilistic typology: Deep generative models of vowel inventories](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 1182–1192, Vancouver, August 2017. Best Long Paper Award.
- Nicholas Andrews, Mark Dredze, Benjamin Van Durme, and Jason Eisner. [Bayesian modeling of lexical resources for low-resource settings](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 1029–1039, Vancouver, August 2017.
- Ryan Cotterell, Adam Poliak, Benjamin Van Durme, and Jason Eisner. [Explaining and generalizing skip-gram through exponential family principal component analysis](#). In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics: Human Language Technologies (EACL)*, pages 175–181, Valencia, Spain, April 2017.
- Matthew Francis-Landau, Bing Xue, Jason Eisner, and Vivek Sarkar. [Fine-grained parallelism in probabilistic parsing with Habanero Java](#). In *Proceedings of the Sixth Workshop on Irregular Applications: Architectures and Algorithms (IA<sup>3</sup>)*, pages 78–81, Salt Lake City, November 2016. IEEE Press.
- Jason Eisner. [Inside-outside and forward-backward algorithms are just backprop](#). In *Proceedings of the EMNLP Workshop on Structured Prediction for NLP*, pages 1–17, Austin, TX, November 2016.
- Tim Vieira, Ryan Cotterell, and Jason Eisner. [Speed-accuracy tradeoffs in tagging with variable-order CRFs and structured sparsity](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1973–1978, Austin, TX, November 2016.
- Rebecca Knowles, Adithya Renduchintala, Philipp Koehn, and Jason Eisner. [Analyzing learner understanding of novel L2 vocabulary](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning (CoNLL)*, pages 126–135, Berlin, August 2016.
- Adithya Renduchintala, Rebecca Knowles, Philipp Koehn, and Jason Eisner. [Creating interactive macaronic interfaces for language learning](#). In *Proceedings of ACL-2016 System Demonstrations*, pages 133–138, Berlin, August 2016.
- Adithya Renduchintala, Rebecca Knowles, Philipp Koehn, and Jason Eisner. [User modeling in language learning with macaronic texts](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1859–1869, Berlin, August 2016.
- Ryan Cotterell, Hinrich Schütze, and Jason Eisner. [Morphological smoothing and extrapolation of word embeddings](#). In *Proceedings of the 54th Annual Meeting of the Association*

- for *Computational Linguistics (ACL)*, pages 1651–1660, Berlin, August 2016. Supplementary material (4 pages) also available.
- Nathaniel Wesley Filardo and Jason Eisner. [Rigid tree automata with isolation](#). In *Proceedings of the Fourth International Workshop on Trends in Tree Automata and Tree Transducers (TTATT)*, Seoul, August 2016. 7 pages.
- Pushpendre Rastogi, Ryan Cotterell, and Jason Eisner. [Weighting finite-state transductions with neural context](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 623–633, San Diego, June 2016. 11 pages. Supplementary material (1 page) also available.
- Nanyun Peng, Ryan Cotterell, and Jason Eisner. [Dual decomposition inference for graphical models over strings](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–927, Lisbon, September 2015.
- Ryan Cotterell and Jason Eisner. [Penalized expectation propagation for graphical models over strings](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 932–942, Denver, June 2015. Supplementary material (11 pages) also available.
- He He, Hal Daumé III, and Jason Eisner. [Learning to search in branch-and-bound algorithms](#). In *Advances in Neural Information Processing Systems 27 (NeurIPS)*, pages 3293–3301, Montreal, December 2014.
- Juneki Hong and Jason Eisner. [Deriving multi-headed projective dependency parses from link grammar parses](#). In *13th International Workshop on Treebanks and Linguistic Theories (TLT)*, Tübingen, December 2014. 5 pages plus appendices.
- Nicholas Andrews, Jason Eisner, and Mark Dredze. [Robust entity clustering via phylogenetic inference](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 775–785, Baltimore, June 2014.
- Ryan Cotterell, Nanyun Peng, and Jason Eisner. [Stochastic contextual edit distance and probabilistic FSTs](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 625–630, Baltimore, June 2014.
- He He, Hal Daumé III, and Jason Eisner. [Dynamic feature selection for dependency parsing](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1455–1464, Seattle, October 2013.
- Francis Ferraro and Jason Eisner. [A virtual manipulative for learning log-linear models](#). In *Proceedings of the Fourth Workshop on Teaching NLP and CL*, pages 66–76, Sofia, Bulgaria, August 2013.
- Patrick Littell, Lori Levin, Jason Eisner, and Dragomir Radev. [Introducing computational concepts in a linguistics olympiad](#). In *Proceedings of the Fourth Workshop on Teaching NLP and CL*, pages 18–26, Sofia, Bulgaria, August 2013.

- Matthew Gormley and Jason Eisner. [Nonconvex global optimization for latent-variable models](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 444–454, Sofia, Bulgaria, August 2013.
- Jiarong Jiang, Taesun Moon, Hal Daumé III, and Jason Eisner. [Prioritized asynchronous belief propagation](#). In *ICML Workshop on Inferning: Interactions between Inference and Learning*, Atlanta, June 2013. 5 pages.
- Jiarong Jiang, Adam Teichert, Hal Daumé III, and Jason Eisner. [Learned prioritization for trading off accuracy and speed](#). In *Advances in Neural Information Processing Systems 25 (NeurIPS)*, pages 1331–1339, Lake Tahoe, NV, December 2012.
- He He, Hal Daumé III, and Jason Eisner. [Imitation learning by coaching](#). In *Advances in Neural Information Processing Systems 25 (NeurIPS)*, pages 3149–3157, Lake Tahoe, NV, December 2012.
- Veselin Stoyanov and Jason Eisner. [Easy-first coreference resolution](#). In *Proceedings of the 24th International Conference on Computational Linguistics (COLING)*, pages 2519–2534, Mumbai, December 2012.
- Nathaniel Wesley Filardo and Jason Eisner. [A flexible solver for finite arithmetic circuits](#). In Agostino Dovier and Vítor Santos Costa, editors, *Technical Communications of the 28th International Conference on Logic Programming (ICLP)*, volume 17 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 425–438, Budapest, September 2012.
- Nicholas Andrews, Jason Eisner, and Mark Dredze. [Name phylogeny: A generative model of string variation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 344–355, Jeju, Korea, July 2012.
- Jiarong Jiang, Adam Teichert, Hal Daumé III, and Jason Eisner. [Learned prioritization for trading off accuracy and speed](#). In *ICML Workshop on Inferning: Interactions between Inference and Learning*, Edinburgh, June 2012. 7 pages.
- He He, Hal Daumé III, and Jason Eisner. [Cost-sensitive dynamic feature selection](#). In *ICML Workshop on Inferning: Interactions between Inference and Learning*, Edinburgh, June 2012. 6 pages.
- Veselin Stoyanov and Jason Eisner. [Fast and accurate prediction via evidence-specific MRF structure](#). In *ICML Workshop on Inferning: Interactions between Inference and Learning*, Edinburgh, June 2012. 6 pages.
- Matthew R. Gormley, Mark Dredze, Benjamin Van Durme, and Jason Eisner. [Shared components topic models](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 783–792, Montreal, June 2012.
- Michael Paul and Jason Eisner. [Implicitly intersecting weighted automata using dual decomposition](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 232–242, Montreal, June 2012.

- Jason Smith and Jason Eisner. [Unsupervised learning on an approximate corpus](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 131–141, Montreal, June 2012.
- Veselin Stoyanov and Jason Eisner. [Minimum-risk training of approximate CRF-based NLP systems](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 120–130, Montreal, June 2012.
- Jason Eisner and Hal Daumé III. [Learning speed-accuracy tradeoffs in nondeterministic inference algorithms](#). In *COST: NeurIPS Workshop on Computational Trade-offs in Statistical Learning*, Sierra Nevada, Spain, December 2011. 5 pages.
- Veselin Stoyanov and Jason Eisner. [Learning cost-aware, loss-aware approximate inference policies for probabilistic graphical models](#). In *COST: NeurIPS Workshop on Computational Trade-offs in Statistical Learning*, Sierra Nevada, Spain, December 2011. 5 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, Edinburgh, July 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, Edinburgh, July 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)*, volume 15 of *JMLR Workshop and Conference Proceedings*, pages 725–733, Fort Lauderdale, April 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, Beijing, August 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, Singapore, August 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, Singapore, August 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, Singapore, August 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, Singapore, August 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, Singapore, August 2009. Nominated for Best Paper Award.
- 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*, Stanford, March 2009. AAAI Technical Report SS-09-07.
- Omar F. Zaidan, Jason Eisner, and Christine Piatko. [Machine learning with annotator rationales to reduce annotation cost](#). In *Proceedings of the NeurIPS\*2008 Workshop on Cost Sensitive Learning*, Whistler, BC, December 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, Honolulu, October 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, Honolulu, October 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, Honolulu, October 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, Columbus, Ohio, June 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, Columbus, Ohio, June 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, Prague, June 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, Rochester, NY, April 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, Rochester, NY, April 2007.
- Damianos Karakos, Sanjeev Khudanpur, Jason Eisner, and Carey E. Priebe. [Iterative de-noising using Jensen-Renyí divergences with an application to unsupervised document categorization](#). In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Honolulu, April 2007. 4 pages.
- 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, Alexandria, VA, October 2006.
- 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, Sydney, July 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, Sydney, July 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, Sydney, July 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, New York, June 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, New York, June 2006. Nominated for 5-year retrospective Best Paper award.
- 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, New York, June 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, Vancouver, October 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, Vancouver, October 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, Vancouver, October 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, Edinburgh, July 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, Ann Arbor, Michigan, June 2005. Nominated for Best Paper Award.
- 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, Sophia Antipolis, France, June 2005. Springer-Verlag.
- 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, December 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, Barcelona, July 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, Barcelona, July 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, Sapporo, July 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, Edmonton, May 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, Philadelphia, July 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, Philadelphia, July 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, Philadelphia, July 2002.
- Jason Eisner. [Transformational priors over grammars](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 63–70, Philadelphia, July 2002. Nominated for Best Paper Award.
- 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 (FSMNLP)*, Helsinki, August 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, Luxembourg, August 2000.
- Jason Eisner. [Directional constraint evaluation in Optimality Theory](#). In *Proceedings of the 18th International Conference on Computational Linguistics (COLING 2000)*, pages 257–263, Saarbrücken, Germany, August 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, Paris, May 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, University of Maryland, June 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, Cambridge, MA, 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, MIT, Cambridge, MA, September 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, Madrid, July 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, Copenhagen, August 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, Santa Cruz, June 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, Maryland, October 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, Ithaca, NY, 1995. DMLL Publications.
- 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, San Jose, July 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, July 1992. Technical Report WS-92-01.

## Refereed Presentations

- Ryan Cotterell, Sabrina J. Mielke, Jason Eisner, and Brian Roark. [Are all languages equally hard to language-model?](#) *Proceedings of the Society for Computation in Linguistics (SCiL)*, 2(47):361–362, January 2019.
- Jason Eisner and Nathaniel Wesley Filardo. [Treating machine learning algorithms as declaratively specified circuits](#). In *Proceedings of the Conference on Systems and Machine Learning (SysML)*, Palo Alto, February 2018.
- Dingquan Wang and Jason Eisner. [Predicting fine-grained syntactic typology from surface features](#). In *Proceedings of the Society for Computation in Linguistics (SCiL)*, volume 1, Salt Lake City, January 2018.
- Ryan Cotterell, Christo Kirov, Mans Hulden, and Jason Eisner. [Quantifying the trade-off between two types of morphological complexity](#). In *Proceedings of the Society for Computation in Linguistics (SCiL)*, volume 1, pages 209–210, Salt Lake City, January 2018.
- Nicholas Andrews and Jason Eisner. [Transformation process priors](#). In *NeurIPS Workshop on Bayesian Nonparametrics: Hope or Hype?*, Sierra Nevada, Spain, December 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 *NeurIPS Workshop on Learning Semantics*, Sierra Nevada, Spain, December 2011. Extended abstract (3 pages).
- Jason Eisner. [Dyna: A non-probabilistic programming language for probabilistic AI](#). Extended abstract for talk at the NeurIPS\*2008 Workshop on Probabilistic Programming, December 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, Baltimore, October 2006.

Jason Eisner. [What constraints should OT allow?](#) Talk handout available online (22 pages), Linguistic Society of America (LSA), January 1997.

### Invited Papers

Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vyloмова, Arya D. McCarthy, Katharina Kann, Sabrina Mielke, Garrett Nicolai, Miikka Silfverberg, David Yarowsky, Jason Eisner, and Mans Hulden. [The CoNLL–SIGMORPHON 2018 shared task: Universal morphological reinflection](#). In *Proceedings of the CoNLL SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 1–27, Brussels, October 2018.

Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vyloмова, Patrick Xia, Manaal Faruqui, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. [CoNLL-SIGMORPHON 2017 shared task: Universal morphological reinflection in 52 languages](#). In *Proceedings of the Conference on Natural Language Learning: CoNLL-SIGMORPHON Shared Task System Descriptions*, pages 1–30, Vancouver, August 2017.

Ryan Cotterell, Christo Kirov, John Sylak-Glassman, David Yarowsky, Jason Eisner, and Mans Hulden. [The SIGMORPHON 2016 shared task—morphological reinflection](#). In *Proceedings of the 14th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 10–22, Berlin, August 2016. Supplementary material (4 pages) also available.

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, Philadelphia, March 2005. Invited talk.

### Technical Reports (some under review as conference/journal submissions)

Shepard Xia, Brian Lu, and Jason Eisner. [Let’s think var-by-var: Large language models enable \*ad hoc\* probabilistic reasoning](#). *Computing Research Repository*, arXiv:2412.02081, December 2024.

Kumar Shridhar, Harsh Jhamtani, Hao Fang, Benjamin Van Durme, Jason Eisner, and Patrick Xia. [SCREWS: A modular framework for reasoning with revisions](#). *Computing Research Repository*, arXiv:2309.13075, September 2023.

Li Du, Hongyuan Mei, and Jason Eisner. [Autoregressive modeling with lookahead attention](#). *Computing Research Repository*, arXiv:2305.12272, May 2023.

Clemente Pasti, Andreas Opedal, Tiago Pimentel, Tim Vieira, Jason Eisner, and Ryan Cotterell. [On the intersection of context-free and regular languages](#). *Computing Research Repository*, arXiv:2209.06809, September 2022.

- Ryan Cotterell, Christo Kirov, Mans Hulden, and Jason Eisner. [On the diachronic stability of irregularity in inflectional morphology](#). *Computing Research Repository*, arXiv:1804.08262, April 2018.
- 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, April 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, Baltimore, March 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, June 1991. Available online (16 pages).
- Jason Eisner. [Dynamical-systems behavior in recurrent and non-recurrent connectionist nets](#). Undergraduate honors thesis, Harvard University, April 1990. Available online (57 pages).

## General Audience

- Leon Derczynski, Jason Eisner, Yoav Goldberg, Iryna Gurevych, Lillian Lee, Joakim Nivre, Brian Roark, Thamar Solorio, Yue Zhang, Chengqing Zong, and Diyi Yang. [Report of the ACL committee on the anonymity policy](#). Report available on the wiki of the Association for Computational Linguistics, December 2023.
- Jason Eisner, Jennifer Foster, Iryna Guryvech, Marti Hearst, Heng Ji, Lillian Lee, Christopher Manning, Paola Merlo, Yusuke Miyao, Joakim Nivre, Amanda Stent, and Ming Zhou. [ACL policies and guidelines for submission, review and citation](#). Report available on the wiki of the Association for Computational Linguistics, October 2017.
- Jason Eisner. [The science of language: Computational linguistics](#). *Imagine Magazine*, 7(4):14–15, March 2000.
- Jason Eisner. [Cognitive science and the search for intelligence](#). Invited paper presented to the Socratic Society, University of Cape Town, South Africa, May 1991. Available online (24 pages).

## Edited Volumes

- Mans Hulden, Ryan Cotterell, Jason Eisner, Manaal Faruqui, Christo Kirov, Sandra Kübler, John Sylak-Glassman, Ekaterina Vylomova, Géraldine Walther, Patrick Xia, and David Yarowsky, editors. [Proceedings of the CoNLL SIGMORPHON 2017 Shared Task: Univer-](#)

- sal Morphological Reinflection*. Association for Computational Linguistics, Vancouver, August 2017. 125 pages.
- 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, Columbus, Ohio, June 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, Prague, June 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, Luxembourg, August 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*. U.S. Patent #8,285,619 issued 10/9/2012, 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*. U.S. Patent 7,630,986, issued 2009, 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*. U.S. Patent Application 20,010/014,868, 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, 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, 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, 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, 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.

## DISSERTATIONS SUPERVISED

- Matthew Francis-Landau. *Declarative Programming via Term Rewriting*. PhD thesis, Johns Hopkins University, February 2024.
- Sabrina J. Mielke. *Building and Evaluating Open-Vocabulary Language Models*. PhD thesis, Johns Hopkins University, October 2023.
- Tim Vieira. *Automating the Analysis and Improvement of Dynamic Programming Algorithms with Applications to Natural Language Processing*. PhD thesis, Johns Hopkins University, July 2023.
- Chu-Cheng Lin. *On Expressiveness, Inference, and Parameter Estimation of Discrete Sequence Models*. PhD thesis, Johns Hopkins University, October 2022.
- Ryan Cotterell. *The Probabilistic Typology of Vowel Systems*. PhD thesis, Johns Hopkins University, July 2021.
- Hongyuan Mei. *Neural Probabilistic Methods for Event Sequence Modeling*. PhD thesis, Johns Hopkins University, July 2021.
- Dingquan Wang. *Supervised Training on Synthetic Languages: A Novel Framework for Un-supervised Parsing*. PhD thesis, Johns Hopkins University, October 2019.
- Nathaniel Wesley Filardo. *Dyna 2: Towards a General Weighted Logic Language*. PhD thesis, Johns Hopkins University, October 2017.
- Nicholas Oliver Andrews. *Generative Non-Markov Models for Information Extraction*. PhD thesis, Johns Hopkins University, February 2016.
- Matthew R. Gormley. *Graphical Models with Structured Factors, Neural Factors, and Approximation-Aware Training*. PhD thesis, Johns Hopkins University, October 2015.
- Markus Dreyer. *A Non-Parametric Model for the Discovery of Inflectional Paradigms from Plain Text Using Graphical Models over Strings*. PhD thesis, Johns Hopkins University, Baltimore, MD, April 2011.
- David A. Smith. *Efficient Inference for Trees and Alignments: Modeling Monolingual and Bilingual Syntax with Hard and Soft Constraints and Latent Variables*. PhD thesis, Johns Hopkins University, Baltimore, MD, October 2010.
- Zhifei Li. *Discriminative Training and Variational Decoding in Machine Translation Via Novel Algorithms for Weighted Hypergraphs*. PhD thesis, Johns Hopkins University, Baltimore, MD, April 2010.
- Roy Tromble. *Search and Learning for the Linear Ordering Problem with an Application to Machine Translation*. PhD thesis, Johns Hopkins University, Baltimore, MD, April 2009.
- Noah A. Smith. *Novel Estimation Methods for Unsupervised Discovery of Latent Structure in Natural Language Text*. PhD thesis, Johns Hopkins University, Baltimore, MD, October 2006.

## OTHER SCHOLARLY AND TECHNICAL OUTPUT

In addition to the items below, code and data associated with my papers are often available

from <http://cs.jhu.edu/~jason/papers>, or on request. Some older releases are listed below.

My student Matt Gormley in 2014 released [Pacaya](#), a general graphical models package with support for structured factors.

Veselin Stoyanov and I released the [ERMA](#) toolkit in 2012 for robust training of graphical models under approximations.

We have 4 partial implementations of the [Dyna programming language](#). My students and I released a Dyna-to-C++ compiler in 2005 for the semiring-weighted subset of Dyna. It was used for both research and teaching. Tim Vieira's 2022 dissertation has an interpreter for this subset that includes program transformations and static analysis. In 2013, we prototyped a new, more flexible runtime engine for the full design, and it was used for teaching. Matthew Francis-Landau has been developing a new type of runtime engine, [Dyna-R](#), based on rewriting of relational expressions (R-exprs).

I provide various tutorial materials at <http://cs.jhu.edu/~jason/tutorials>. These include interactive software visualizations for log-linear models (2013), hidden Markov models, and Gaussian mixture models, as well as software for experimenting with probabilistic context-free grammars. They include other writings, slides, and videos as well. My course slides have also been used at a number of other institutions.

Several undergraduates and I released the [Dynasty](#) hypergraph browser in 2006. We are designing a more powerful successor for use with version 2 of Dyna.

Philipp Koehn and I produced the [aclpub](#) publications package that has been used since 2005 for the ACL, NAACL, EMNLP, EACL, COLING, IJCNLP, and CoNLL conferences (among others), as well as their associated workshops.

## TEACHING

Department of Computer Science, Johns Hopkins University

Alumni Association Excellence in Teaching Award, Whiting School of Engineering, 2013

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

- [Natural Language Processing](#) 2001–, except for 2005, 2015, 2019  
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: 70.
- [CS Freshman Experience](#) 2020, 2021  
A weekly seminar to introduce a group of new majors to a range of high-level ideas in computer science.  
Enrollment: 10.
- [Current Topics in Natural Language Processing](#) every semester 2001–



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.

Enrollment: 10.

- *Machine Learning: Linguistic and Sequence Modeling* 2018, 2019  
A graduate class that teaches formal methods for mathematical modeling of real-world sequence data, including finite-state methods, Bayesian nonparametrics, deep learning, and reinforcement learning. Examples are drawn primarily from natural language processing.  
Enrollment: 30.
- *Declarative Methods* 2005–2017, except 2015  
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: 15–35.
- *Current Topics in Machine Learning* 2013, 2014  
A weekly reading and discussion group with a participatory format, jointly led by three faculty.  
Enrollment: about 20.
- *Puzzles, Graphs, and NLP* 2008, 2011  
A 3-class unit to expose freshman CS majors to natural language processing. Discusses various contest puzzles and applied NLP problems, and connects many of them to computational questions about directed graphs.  
Enrollment: About 20.
- *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.
- Tutorials at conferences
- *Structured Belief Propagation for Natural Language Processing* 2014, 2015  
(with Matt Gormley, at ACL 2014 and ACL 2015)
- NAACL Professor (named professorship) at the *Linguistic Institute* 2013
- *Computational linguistics course* at the biennial summer school of the Linguistic Society of America.
- Lecturer, 1st *Lisbon Machine Learning School* 2011
- Lectures on generative Bayesian modeling.
- Lecturer, NAACL Summer School in Human Language Technology summers 2002–2012, 2016
- Lectures on NLP and applied probability.
- *Laboratory exercise* in competitive grammar writing.
- Speaker, TA Training Institute, Johns Hopkins University 2005–2018
- 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 ( <i>NSF Computational Innovation Fellow; now at Facebook</i> )	2009–2011
--	-----------

#### Ph.D. dissertation students

Fei Fang	2025–
Li (Leo) Du	2021–
Brian (Zhichu) Lu	2020–
Sabrina Mielke	2017–2023
Hongyuan Mei ( <i>Bloomberg Fellow; Jelinek Memorial Award; research prof at TTIC; now prof at Purdue</i> )	2016–2021
Matthew Francis-Landau	2016–2024
Chu-Cheng Lin ( <i>now at Google</i> )	2015–2022
Dingquan Wang ( <i>now at Google</i> )	2014–2019
Ryan Cotterell ( <i>Fulbright Scholar; NDSEG Fellow; Facebook Fellow; Jelinek Memorial Award; now prof at ETH Zurich</i> )	2013–2021
Tim Vieira ( <i>now postdoc at ETH Zurich</i> )	2011–2023
Nathaniel W. Filardo ( <i>HLTCOE Fellow; postdoc at U. Cambridge; now at Microsoft Research</i> )	2007–2017
Nicholas Andrews ( <i>now research scientist, JHU HLTCOE</i> )	2009–2016
Matthew Gormley ( <i>HLTCOE Fellow; now teaching prof, CMU MLD</i> )	2009–2015
Markus Dreyer ( <i>Wolman Fellow; now Sr. Machine Learning Scientist, Amazon</i> )	2003–2011
Zhifei Li (co-adv.) ( <i>later at Google; now CEO/founder of Mobvoi</i> )	2008–2010
David A. Smith ( <i>NSF Fellow; research prof at Mass; now prof at Northeastern</i> )	2002–2010
Roy Tromble ( <i>NDSEG Fellow; now at Google</i> )	2002–2009
Noah A. Smith ( <i>Hertz Foundation Fellow; prof at CMU; now prof at U. Washington</i> )	2001–2006

#### Other former Ph.D. students (period as primary advisee)

Jacob Buckman ( <i>now at MILA</i> )	2018–2019
Michael Paul ( <i>NSF Fellow; Microsoft Fellow; Dean's Fellow; now Asst. Prof., UC-Boulder</i> )	2010–2015
Adam Teichert ( <i>now Asst. Prof., Snow College</i> )	2010–2014
Darcey Riley ( <i>Dean's Fellow; moved to Google; now at Notre Dame</i> )	2012–2015
Jason R. Smith ( <i>switched advisors in 2011; now at Google</i> )	2006–2011
Omar Zaidan ( <i>now at Microsoft Research</i> )	2005–2008
John Blatz ( <i>now at Google</i> )	2004–2008
Eric Goldlust ( <i>Wolman Fellow; Muuss Research Award; Hon. Mention for CRA Outstanding Undergraduate Award</i> )	2004–2005

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

Haojun Shi	2025–
Prabhav Singh	2025–
Kenan Rustamov	2025–
Rishi More	2025–
Fengting Yuchi ( <i>visitor from Shanghai Jiao Tong</i> )	2025–
Reva Hirave ( <i>visitor from George Mason</i> )	2024–
Kateryna Morhun ( <i>visitor from MIT</i> )	2024–
Yuqi Li	2024–
Tom Wang	2024–
T.J. Bai	2024–
Weitao Shepard Xia	2024–
Tsugunobu Miyake ( <i>visitor from Bucknell;</i> <i>Hon. Mention for CRA Outstanding Undergraduate Award</i> )	2024
Steven Tan	2021–2023
Qi He	2021–2023
Chenghao Yang ( <i>visitor from Columbia</i> )	2020–2022
Xiang Lisa Li ( <i>CRA Outstanding Undergraduate Award;</i> <i>now at Stanford</i> )	2018–
Guanghui Qin ( <i>visitor from Peking University, then Ph.D. at JHU</i> )	2018–2019
Hao Zhu ( <i>visitor from Tsinghua University, then Ph.D. at CMU</i> )	2018–2019
Shijie Wu	2017–2018
Xiaochen Li	2017
Bo Liu	2016–2017
Yuan Jing (Vincent) Yan	2016–2017
Jiawei Huang	2016–2020
Nikita Srivatsan	2016–2018
Mozhi Zhang	2015–2018
Elan Hourticolon-Retzler	2015
Juneki Hong ( <i>Most Valuable Peer Award</i> )	2013–2018
Sharon Li ( <i>Outstanding Senior Award</i> )	2013–2014
Katherine Wu ( <i>Pistrutto Fellow, Outstanding Senior Award</i> )	2012–2013
Jay Feldman ( <i>Pistrutto Fellow, Outstanding Senior Award</i> )	2012–2013
Michael Tontchev ( <i>High school—Grand Prize Winner;</i> <i>Baltimore Science Fair</i> )	2010–2012
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>Pistrutto Fellow</i> )	2006–2007
Constantinos Michael	2005–2006
George Shafer	2005–2006
John Graettinger ( <i>Outstanding Research Award</i> )	2005–2006
Michael Kornbluh ( <i>Pistrutto Fellow, Outstanding Senior Award</i> )	2004–2005
Chalaporn Hathaidharm	2002–2004

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

Benjamin Börschinger (Macquarie U.)	exp. 2015
Yonatan Bisk (UIUC)	2015
Jiarong Jiang (U. of Maryland)	2014
Francisco Sánchez-Vega (JHU Applied Math & Stats Dept.)	2012
Will Headden (Brown University)	2011
Jia Cui	2008
Mark Thober	2007
Gaja Jarosz Snover (JHU Cognitive Science Dept.)	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

October 6, 2025