Jason M. Eisner
Curriculum Vitae

CONTACT

Department of Computer Science
Johns Hopkins University
3400 N. Charles St., Hackerman 324C
Baltimore, MD 21218-2691
U.S.A.

Phone: (410) 516-8438
Fax: (410) 516-5050

jason@cs.jhu.edu
http://cs.jhu.edu/~jason

EDUCATION
AND HONORS

Ph.D. in Computer Science　　　　　　　　　　　　　　　　　2001
University of Pennsylvania
　　　Graduate Teaching Award
　　　Thesis: *Smoothing a Probabilistic Lexicon via Syntactic Transformations*
　　　Advisor: Mitch Marcus

B.A./M.A. in Mathematics　　　　　　　　　　　　　　　　　1993
University of Cambridge
　　　First-class honours
　　　Note: Second undergraduate degree.

A.B. in Psychology, Cognitive Science track　　　　　　　　　1990
Harvard University
　　　Summa cum laude
　　　Junior-year election to Phi Beta Kappa
　　　GPA: 3.9/4.0; GRE: 800/800/800
　　　Thesis: *Dynamical-Systems Behavior in Recurrent and Non-Recurrent
　　　　　　Connectionist Nets*

PROFESSIONAL
EXPERIENCE

Johns Hopkins University　　　　　　　　　　　　　　　　July 2014–
Professor of Computer Science
Joint appointment in Cognitive Science (2003–)
Member of the Center for Language and Speech Processing (2000–)
Member of the Mathematical Institute for Data Science (2017–)
Affiliate of the Human Language Technology Center of Excellence (2007–)

Johns Hopkins University　　　　　　　　　　　July 2007–June 2014
Associate Professor of Computer Science

Johns Hopkins University　　　　　　　　　　　July 2000–June 2007
Assistant Professor of Computer Science

University of Rochester　　　　　　　　　　　　Jan. 2000–June 2001
Assistant Professor of Computer Science
Secondary appointment in Linguistics

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–).
– 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* (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), *IEEE Transactions on Audio, Speech, and Language Processing* (2007), *Cognition* (2002), *J. Algorithms* (1997), *Language and Speech* (1999).

Conferences

– Diversity & Inclusion co-chair, NAACL (2018).
– 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). 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); SIG-PHON Workshop on Finite-State Phonology (2000).
- Organizer, NSF international workshop on Probabilistic Representations of Linguistic Meaning (2014).
- Publications chair, ACL (2005).
- Co-chair for HLT/NAACL Tutorial and Workshop Programs (2003).
- Program committee member / reviewer for ACL (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), NIPS (2018, 2017, 2016, 2014, 2013, 2011, 2010, 2007), NLP-LING (2010), SCiL (2018), SIGMORPHON (2018, 2014, 2012, 2010, 2008, 2006, 2004, 2002, 2000, 1998), WiNLP—Widening NLP Workshop (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 (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
- Board member, Women and Underrepresented Minorities in Natural Language Processing (WINLP) (2016–).
- Organizing committee member, SIGDAT (which runs the EMNLP conference) (2007–).
- Problems committee, North American Computational Linguistics Olympiad (2006–).
- President, ACL SIGMORPHON (Computational Morphology and Phonology) (2001–); executive committee (1998–).

Other
- Judge for "NLP Idol" special event at NAACL (2012).
- Mentor for NAACL Student Research Workshop (2013, 2012, 2009).
- NSF proposal reviewer (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
Quora Top Writer	2016, 2017
Best Featured Actor in a Musical DC Metro Theater Arts, Best of 2014	2014
Alumni Association Excellence in Teaching Award Johns Hopkins University, Whiting School of Engineering	2013
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
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

UNIVERSITY SERVICE

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

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

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

Leader of campus-wide Machine Learning Group, (2007–).

Mentor for a few junior faculty members.

Search committees: Bloomberg Distinguished Professorship in Computational Cognitive Science (2014, co-chair), Bloomberg Distinguished Professorship in Computational Healthcare (2014), Director of Admissions (2012), Applied Math and Statistics (2012, 2011), CS machine learning subcommittee (2014, 2013, 2012, 2011, 2010), HLTCOE (2008–).

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

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

Director of graduate studies, Computer Science Department (2002–2013).

Onstage presenter for various Homewood events: admissions open houses (spring 2013, fall 2013, spring 2014), Hopkins Engineering Sampler Seminar (2013), Dissertation Writing Workshop (2013, 2016), QuarkNet (2013), WSE centennial video (2013), TA Training Institute (2005–).

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–?).

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

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

CS/CLSP admissions committee (2001–).

See also "Teaching" section.

GRANTS	Amazon Research Award: Continuous-Time Reinforcement Learning For Personalization (PI, $99K)	2018
	NSF RI-Small: Linguistic Structure in Neural Sequence Models (PI, $395K)	2017–2020
	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

Discrete Latent Variables in NLP: Good, Bad, and Indifferent. Invited talk at ACL Workshop on Relevance of Linguistic Structure in Neural Architectures for NLP, August 2018.

Recovering Syntactic Structure from Surface Features. Pennsylvania State University, January 2018. (invited talk)

Probabilistically Modeling Surface Patterns Using Latent Structure. Symposium on Perceptrons and Syntactic Structures at 60: Computational Modeling of Language, at Society for Computation in Linguistics (SCiL), January 2018. (invited talk)

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. (invited talk)

Tractable deep models of sequential structure. LTI Colloquium, Carnegie Mellon University, March 2017. (invited talk)

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

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

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

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

Probabilistic inference on strings. Triangle Computer Science Distinguished Lecture Series (UNC+Duke+NCSU), November 2015. Similar invited 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, November 2014. (invited talk)

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

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

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

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

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

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, 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

Dingquan Wang and Jason Eisner. Surface statistics of an unknown language indicate how to parse it. *Transactions of the Association for Computational Linguistics (TACL)*, to appear 2018.

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)*, to appear 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

Sebastian 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*, 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, Sebastian 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, Sebastian 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.

Hongyuan Mei and Jason Eisner. The neural Hawkes process: A neurally self-modulating multivariate point process. In *Advances in Neural Information Processing Systems (NeurIPS)*, 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^3)*, 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.

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 Natu-*

ral 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 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)*, 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.

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.

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.

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. ∀-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.

Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, 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.

Refereed Presentations

Jason Eisner and Nathaniel Wesley Filardo. Treating machine learning algorithms as declaratively specified circuits. In *Proceedings of the SysML Conference*, 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, 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

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

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

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: Universal 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, 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.

DISSERTATIONS SUPERVISED

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 available on request. Some of it has been requested often.

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 toolkit in 2012 for robust training of graphical models under approximations.

We have $1\frac{1}{2}$ implementations of the Dyna programming language. My students and I released a Dyna-to-C++ compiler in 2005 for an early, simpler design of Dyna. It was used for both research and teaching. In 2013, we began working hard on a new, more flexible runtime engine for a much more expressive language design. The new system has already been used for teaching. Both versions can be obtained from `http://dyna.org`.

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

– *Machine Learning: Linguistic and Sequence Modeling*   2018
  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.

– *Natural Language Processing*   2001–2004, 2006–2014, 2016–
  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: 60.

– *Declarative Methods* 2005–2014, 2017–

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.

– *Seminar 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: 10–15.

– *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–

- 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 Computational Innovation Fellow;*   2009–2011
       now at Facebook)

Ph.D. dissertation students

Jacob Buckman	2018–
Sebastian Mielke	2017–
Hongyuan Mei *(Bloomberg Fellow)*	2016–
Matthew Francis-Landau	2016–
Chu-Cheng Lin	2015–
Dingquan Wang	2014–
Ryan Cotterell *(Fulbright Scholar, NDSEG Fellow, Facebook Fellow;* *now Asst. Prof., Cambridge Univ.)*	2013–
Tim Vieira	2011–
Nathaniel W. Filardo *(HLTCOE Fellow;* *now postdoctoral fellow, Univ. of Cambridge)*	2007–2017
Nicholas Andrews *(now postdoctoral fellow, JHU HLTCOE)*	2009–2016
Matthew Gormley *(HLTCOE Fellow;* *now Asst. Teaching Prof., CMU MLD)*	2009–2015
Markus Dreyer *(Wolman Fellow;* *now Sr. Machine Learning Scientist, Amazon)*	2003–2011
Zhifei Li (co-adv.) *(later at Google; now CEO/founder of Mobvoi)*	2008–2010
David A. Smith *(NSF Fellow; then Research Prof., U. Mass;* *now Asst. Prof., Northeastern Univ.)*	2002–2010
Roy Tromble *(NDSEG Fellow; now at Google)*	2002–2009
Noah A. Smith *(Hertz Foundation Fellow; then Assoc.* *Prof., CMU; now at University of Washington)*	2001–2006

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

Michael Paul *(NSF Fellow; Microsoft Fellow; Dean's Fellow;* *now Asst. Prof., UC-Boulder)*	2010–2015
Adam Teichert *(now Asst. Prof., Snow College)*	2010–2014
Darcey Riley (a.k.a. Halley Orshan) *(Dean's Fellow; now at Google)*	2012–2015
Jason R. Smith *(switched advisors in 2011; now at Google)*	2006–2011
Omar Zaidan *(now at Microsoft Research)*	2005–2008
John Blatz *(now at Google)*	2004–2008
Eric Goldlust *(Wolman Fellow; Muuss Research Award;* *Hon. Mention for CRA Outstanding Undergraduate Award)*	2004–2005

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

Hao Zhu (visitor from Tsinghua University)	2018–
Guanghui Chen (visitor from Peking University)	2018–
Xiang Lisa Li	2018–
Shijie Wu	2017–2018

Xiaochen Li	2017
Bo Liu	2016–2017
Yuan Jing (Vincent) Yan	2016–2017
Jiawei Huang	2016–
Akshay Srivatsan	2016–
Mozhi Zhang	2015–
Elan Hourticolon-Retzler	2015
Juneki Hong *(Most Valuable Peer Award)*	2013–
Sharon Li *(Outstanding Senior Award)*	2013–2014
Katherine Wu *(Pistritto Fellow, Outstanding Senior Award)*	2012–2013
Jay Feldman *(Pistritto Fellow, Outstanding Senior Award)*	2012–2013
Michael Tontchev *(High school—Grand Prize Winner,*	2010–2012
Baltimore Science Fair)	
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 *(Pistritto Fellow)*	2006–2007
Constantinos Michael	2005–2006
George Shafer	2005–2006
John Graettinger *(Outstanding Research Award)*	2005–2006
Michael Kornbluh *(Pistritto Fellow, Outstanding Senior Award)*	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

December 12, 2018