

Greg Durrett

CONTACT INFORMATION *E-mail:* gdurrett@cs.utexas.edu 2317 Speedway
Website: <http://cs.utexas.edu/~gdurrett/> The University of Texas at Austin
Phone: (607) 342-8276 Austin, TX 78712 USA

POSITIONS **The University of Texas at Austin**, September 2017 –
Assistant Professor

EDUCATION **University of California, Berkeley**
Ph.D. in Computer Science, August 2016
Thesis: *Identifying and Resolving Entities in Text* (Advisor: Dan Klein)

Massachusetts Institute of Technology
B.S. in Computer Science and B.S. in Mathematics, minor in Music, June 2010

PUBLICATIONS Yasumasa Onoe, Michael Boratko, Andrew McCallum, and **Greg Durrett**. 2021. Modeling Fine-Grained Entity Types with Box Embeddings. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Shih-ting Lin, Nathanael Chambers, and **Greg Durrett**. 2021. Conditional Generation of Temporally-ordered Event Sequences. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Jiacheng Xu and Greg Durrett. 2021. Dissecting Generation Modes for Abstractive Summarization Models. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Mahnaz Koupaee, **Greg Durrett**, Nathanael Chambers, and Niranjan Balasubramanian. 2021. Don't Let Discourse Confine Your Model: Sequence Perturbations for Improved Event Language Models. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Jifan Chen and **Greg Durrett**. 2021. Robust Question Answering Through Sub-part Alignment. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Tanya Goyal and **Greg Durrett**. 2021. Annotating and Modeling Fine-grained Factuality in Summarization. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Elisa Ferracane, **Greg Durrett**, Junyi Jessy Li and Katrin Erk. 2021. Did they answer? Subjective acts and intents in conversational discourse. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Qiaochu Chen, Aaron Lamoreaux, Xinyu Wang, **Greg Durrett**, Osbert Bastani and Isil Dillig. 2021. Web Question Answering with Neurosymbolic Program Synthesis. *Proceedings of the Conference on Programming Language Design and Implementation (PLDI)*.

Shrey Desai, Jiacheng Xu, and **Greg Durrett**. 2020. Compressive Summarization with Plausibility and Salience Modeling. *Proceedings of the Conference on Empirical Methods for Natural Language*

Processing (EMNLP).

Jiacheng Xu, Shrey Desai, and **Greg Durrett**. 2020. Understanding Neural Abstractive Summarization Models via Uncertainty. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*.

Shrey Desai and **Greg Durrett**. 2020. Calibration of Pre-trained Transformers. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*.

Wei-Jen Ko, Te-Yuan Chen, Yiyang Huang, **Greg Durrett**, and Junyi Jessy Li. 2020. Inquisitive Question Generation for High Level Text Comprehension. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*.

Tanya Goyal and **Greg Durrett**. 2020. Evaluating Factuality in Generation with Dependency-level Entailment. *Findings of the Conference on Empirical Methods for Natural Language Processing (Findings of EMNLP)*.

Yasumasa Onoe and **Greg Durrett**. 2020. Interpretable Entity Representations through Large-Scale Typing. *Findings of the Conference on Empirical Methods in Natural Language Processing (Findings of EMNLP)*.

Kaj Bostrom and **Greg Durrett**. 2020. Byte Pair Encoding is Suboptimal for Language Model Pretraining. *Findings of the Conference on Empirical Methods in Natural Language Processing (Findings of EMNLP)*.

Xi Ye, Qiaochu Chen, Xinyu Wang, Isil Dillig, and **Greg Durrett**. Sketch-Driven Regular Expression Generation from Natural Language and Examples. *Transactions of the Association for Computational Linguistics (ACL)*.

Tanya Goyal and **Greg Durrett**. 2020. Neural Syntactic Preordering for Controlled Paraphrase Generation. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Xi Ye, Qiaochu Chen, Isil Dillig, and **Greg Durrett**. 2020. Benchmarking Multimodal Regex Synthesis with Complex Structures. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Qiaochu Chen, Xinyu Wang, Xi Ye, **Greg Durrett**, and Isil Dillig. 2020. Multi-Modal Synthesis of Regular Expressions. *Proceedings of the Conference on Programming Language Design and Implementation (PLDI)*.

Jiayi Wei, Maruth Goyal, **Greg Durrett**, and Isil Dillig. 2020. LambdaNet: Probabilistic Type Inference using Graph Neural Networks. *Proceedings of the International Conference on Learning Representations (ICLR)*.

Yasumasa Onoe and **Greg Durrett**. 2020. Fine-Grained Entity Typing for Domain-Independent Entity Linking. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.

Jiacheng Xu and **Greg Durrett**. 2019. Neural Extractive Text Summarization with Syntactic Compression. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*.

Aditya Gupta and **Greg Durrett**. 2019. Effective Use of Transformer Networks for Entity Tracking. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*.

Su Wang, **Greg Durrett**, and Katrin Erk. 2019. Query-focused Scenario Construction. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*.

Tanya Goyal and **Greg Durrett**. 2019. Embedding time expressions for deep temporal ordering models. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Elisa Ferracane, **Greg Durrett**, Junyi Jessy Li, and Katrin Erk. 2019. Evaluating Discourse in Structured Text Representations. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Yasumasa Onoe and **Greg Durrett**. 2019. Learning to Denoise Distantly-Labeled Data for Entity Typing. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Jifan Chen and **Greg Durrett**. 2019. Understanding dataset design choices for multi-hop reasoning. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Wei-Jen Ko, **Greg Durrett**, and Junyi Jessy Li. 2019. Linguistically-Informed Specificity and Semantic Plausibility for Dialog Generation. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Aditya Gupta and **Greg Durrett**. Tracking Discrete and Continuous Entity State for Process Understanding. *NAACL 2019 Workshop on Structured Prediction for NLP (oral)*.

Wei-Jen Ko, **Greg Durrett**, and Junyi Jessy Li. 2019. Domain Agnostic Real-Valued Specificity Prediction. *Proceedings of the AAAI Conference on Artificial Intelligence*.

Jiacheng Xu and **Greg Durrett**. 2018. Spherical Latent Spaces for Stable Variational Autoencoders. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*.

David Mueller and **Greg Durrett**. 2018. Effective Use of Context in Noisy Entity Linking. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*.

Su Wang, Eric Holgate, **Greg Durrett**, and Katrin Erk. 2018. Picking Apart Story Salads. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*.

Su Wang, **Greg Durrett**, and Katrin Erk. 2018. Modeling Semantic Plausibility by Injecting World Knowledge. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Greg Durrett, Jonathan Kummerfeld, Taylor Berg-Kirkpatrick, Rebecca Portnoff, Sadia Afroz, Damon McCoy, Kirill Levchenko, and Vern Paxson. 2017. Identifying Products in Online Cybercrime Marketplaces: A Dataset for Fine-grained Domain Adaptation. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*

Rebecca Portnoff, Sadia Afroz, **Greg Durrett**, Jonathan Kummerfeld, Taylor Berg-Kirkpatrick, Damon McCoy, Kirill Levchenko, and Vern Paxson. 2017. Tools for Automated Analysis of Cybercriminal Markets. *Proceedings of the International World Wide Web Conference (WWW)*.

Greg Durrett, Taylor Berg-Kirkpatrick, and Dan Klein. 2016. Learning-Based Single-Document Summarization with Compression and Anaphoricity Constraints. *Proceedings of the 54th Annual*

Meeting of the Association for Computational Linguistics (ACL).

Matthew Francis-Landau, **Greg Durrett**, and Dan Klein. 2016. Capturing Semantic Similarity for Entity Linking with Convolutional Neural Networks. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).*

Greg Durrett and Dan Klein. 2015. Neural CRF Parsing. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics (ACL).*

James Ferguson, **Greg Durrett** and Dan Klein. 2015. Disfluency Detection with a Semi-Markov Model and Prosodic Features. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).*

Greg Durrett and Dan Klein. 2014. A Joint Model for Entity Analysis: Coreference, Typing, and Linking. *Transactions of the Association for Computational Linguistics (TACL).*

David Hall, **Greg Durrett**, and Dan Klein. 2014. Less Grammar, More Features. *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL).*

Greg Durrett and Dan Klein. 2013. Easy Victories and Uphill Battles in Coreference Resolution. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*
Best Paper Finalist.

Greg Durrett, David Hall, and Dan Klein. 2013. Decentralized Entity-Level Modeling for Coreference Resolution. *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL).*

Taylor Berg-Kirkpatrick, **Greg Durrett**, and Dan Klein. 2013. Unsupervised Transcription of Historical Documents. *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*

Greg Durrett and John DeNero. 2013. Supervised Learning of Complete Morphological Paradigms. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).*

Greg Durrett, Adam Pauls, and Dan Klein. 2012. Syntactic Transfer Using a Bilingual Lexicon. *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP).*

Greg Durrett and Dan Klein. 2011. An Empirical Investigation of Discounting in Cross-Domain Language Models. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (ACL)*

Greg Durrett, Frank Neumann, and Una-May O'Reilly. 2011. Computational Complexity Analysis of Simple Genetic Programming On Two Problems Modeling Isolated Program Semantics. *Foundations of Genetic Algorithms.*

Greg Durrett, Muriel Médard, and Una-May O'Reilly. 2010. A Genetic Algorithm to Minimize Chromatic Entropy. *Proceedings of the European Conference on Evolutionary Computation in Combinatorial Optimization (EvoCOP).*

Accept rates for most NLP conferences can be found at: https://aclweb.org/aclwiki/Conference_acceptance_rates
Findings of EMNLP 2020 represents the next 15% of papers after the top 22% appearing in the main conference.

GRANTS, AWARDS, AND FELLOWSHIPS

- Amazon Research Award (Alexa Fairness in AI), 2021

- Salesforce Research Deep Learning Grant, 2020
 - Walmart Labs, Sponsored Research Project, 2020
 - DARPA KAIROS (with Niranjan Balasubramanian (primary), Minh Hoai Nguyen, Frank Ferraro, Nate Chambers, Katrin Erk, Ray Mooney), 2019-2023
 - Gift from Arm, Inc., 2019
 - NSF Medium: Towards Transparent and Effective Program Analysis for Detecting Mobile Malware (with Isil Dillig (primary), Yu Feng), 2019-2022
 - Texas Health Catalyst Award, 2019
 - NSF Small: Handholding Supervision for Latent Variable Models in NLP (sole author), 2018-2021
 - NSF Medium: Computer-Aided Programming for Data Science (with Isil Dillig (primary), Ruben Martins), 2018-2021
 - Bloomberg Data Science Research Grant, 2017-2018
-
- Facebook Fellowship in Natural Language Processing, 2014-2015
 - Chair's Special Award for service to the Berkeley EECS Department, May 2014
 - Best Paper Finalist at EMNLP 2013, October 2013
 - Outstanding Graduate Student Instructor Award, April 2012
 - National Science Foundation Graduate Research Fellowship, April 2010
 - National Merit Scholar, June 2006
 - Robert C. Byrd Scholar, June 2006

TEACHING

Materials: Teaching materials (slides and course materials) from the following classes publicly released with the following publication:

Greg Durrett, Jifan Chen, Shrey Desai, Tanya Goyal, Lucas Kabela, Yasumasa Onoe and Jiacheng Xu. 2021. Contemporary NLP Modeling in Six Comprehensive Programming Assignments *Proceedings of the Fifth Workshop on Teaching NLP*.

Select materials used at UC Berkeley, University of Washington, Princeton, Ohio State University, Ohio University, Georgia Tech, and ETH Zurich.

CS388: Natural Language Processing (graduate)

Instructor, UT Austin, January 2021 – May 2021; September 2019 – December 2019; August 2018 – December 2018.

Additional online version January 2021 – May 2021

CS378: Natural Language Processing (undergraduate)

Instructor, UT Austin, September 2020 – December 2020; January 2020 – May 2020; January 2019 – May 2019

CS395T: Natural Language Processing (Masters of Computer Science Online program)

Instructor, UT Austin, January 2021 – May 2021

CS395T: Structured Models for NLP (graduate)

Instructor, UT Austin, August 2017 – December 2017

CS288: Statistical Natural Language Processing (graduate)

Teaching Assistant, UC Berkeley, August 2014 – December 2014

CS188: Artificial Intelligence (undergraduate)

Teaching Assistant, UC Berkeley, August 2011 – December 2011

INVITED TALKS

- University of Pennsylvania, March 1, 2021
- UT Austin Institute for Machine Learning Seminar Series, February 8, 2021
- UC Santa Barbara, January 14, 2021

- Microsoft Research, December 9, 2020
- Georgia Tech, October 23, 2020
- Salesforce Research, July 21, 2020
- Carnegie Mellon University, September 27, 2019
- WeCNLP Invited Talk on Low-Resource NLP, September 6, 2019
- Arm, March 16, 2019
- UT Random Structures Seminar, February 27, 2019
- UT Turing Scholars Talk, October 9, 2018
- Bloomberg, July 10, 2018
- Toyota Technological Institute at Chicago, June 25, 2018
- ICES Babuska Forum, UT Austin, March 2, 2018
- Berkeley AI Research Seminar, October 21, 2016
- The University of Texas at Austin, April 14, 2016
- University of California, Irvine, March 16, 2016
- New York University, March 9, 2016
- Carnegie Mellon University, March 7, 2016
- Princeton University, March 1, 2016
- University of California, San Diego, February 22, 2016
- Google New York (NLP group), February 8, 2016
- Microsoft Research, February 4, 2016
- ICSI Berkeley, November 25, 2014
- Facebook, November 20, 2014
- Google Mountain View (SAFT group), February 27, 2014

PROFESSIONAL
SERVICE

- Co-organizer (one of eight) for the ACL 2021 workshop on Natural Language Processing for Programming
- Co-organizer (one of seven) for the EMNLP 2019 workshop on Deep Learning for Low-Resource NLP (DeepLo)
- Publication Chair for EMNLP 2021
- Faculty Advisor (one of two) for the NAACL 2019 Student Research Workshop
- Area Chair for AKBC 2021 (Information Extraction), ACL 2021 (Sentence-level Semantics), NAACL 2021 (Language Generation), CoNLL 2020 (Applications and Resources), EMNLP 2020 (Information Extraction), CoNLL 2019 (Tagging/Chunking/Syntax/Parsing)
- Submission mentor for the ACL 2021 Student Research Workshop (1 paper), NAACL 2021 Student Research Workshop (1 paper), ACL 2020 Student Research Workshop (2 papers), ACL 2019 Student Research Workshop (1 paper), NAACL 2019 Student Research Workshop (1 paper)
- Program Committee member (reviewer): NeurIPS 2021, ICLR 2020, NeurIPS 2020, AKBC 2020, ICML 2020, ACL 2020, EMNLP 2019, NeurIPS 2019, ACL 2019, ICML 2019, IEEE Security and Privacy 2019 (ad-hoc), NAACL 2019, EMNLP 2018, ACL 2018, ICML 2018, GenDeep 2018, NAACL 2018, AAAI 2018 (ad-hoc), EMNLP 2017, New Frontiers in Summarization 2017, CoNLL 2017, ACL 2017, ACL 2017 Demo, EMNLP 2016, COLING 2016, CoNLL 2016, ACL 2016, NAACL 2016, EMNLP 2015, ACL 2015, NAACL 2015, EMNLP 2014, ACL 2014, STAR-SEM 2014, EMNLP 2013
- Conference Service: Mentoring group leader (NAACL 2019, ACL 2020), discussion group organizer (NAACL 2019, ACL 2020), session chair (ACL 2015, EMNLP 2020)
- Outstanding reviewer award for ICLR 2021, ICML 2020 (top 33% of reviewers), NeurIPS 2019 (top 50% of reviewers), EMNLP 2019 (top 1/6 of reviewers), EMNLP 2018 (top 7% / top 100 reviewers), ACL 2018, NAACL 2015 (top 25% of reviewers)
- Regular reviewer for TACL since 2014 (>20 papers reviewed)
- Other journal reviewing: CL (1 paper), JAIR (1 paper), Computer Speech & Language (1 paper)

PHD THESES
SUPERVISED

- Su Wang (PhD, 2020, co-supervised with Katrin Erk): Exploring Distributional Semantics in Lexical Representations and Narrative Modeling

UNDERGRADUATE
AND MS THESES

- Shih-Ting Lin (MS, 2020): Conditional Generation of Temporally-Ordered Event Sequences
- Shrey Desai (BS, 2020): Compressive Summarization with Plausibility and Saliency Modeling
- Xinyu (Lucy) Zhao (BS, 2020): Learning Temporal Relations Across Domains
- Akshay Kumar Gupta (BS, 2020): Topic-Focused Extractive Summarization
- Victoria Lestari (MS, 2018): Building Effective Representations for Domain Adaptation in Coreference Resolution

OTHER PROJECT
SUPERVISION

- Liyan Tang (BS, 2021)
- Isabel Cachola (BS, 2019)
- Dhruv Rajan (BS, 2019)
- Aditya Gupta (MS, 2019)
- Tanya Goyal (MS, 2019)
- David Mueller (BS, 2018)

DEPARTMENTAL
SERVICE

- Undergraduate Studies Committee (departmental), UT Austin Computer Science, 2020-2021
- Linguistics faculty hiring committee (computational position), UT Austin, 2020-2021
- Admissions Committee member, UT Austin, 2016-2017, 2017-2018, 2018-2019, 2019-2020
- Graduate course requirements committee member (departmental), 2017-2018
- Coordinator for the UT Austin AI retreat, Fall 2017
- Student coordinator for the 2014 UC Berkeley Computer Science Graduate Visit Weekend

INDUSTRY
EXPERIENCE

Semantic Machines, Berkeley, CA

Senior Research Scientist, June 2016 – June 2017

Google, Mountain View, CA

Software Engineering Intern–Research, May 2012 – August 2012

Charles River Analytics, Cambridge, MA

Software Engineering Intern, June 2008 – August 2008

ATC–NY, Ithaca, NY

Intern, June 2007 – August 2007