

Samuel R. Bowman, *Curriculum Vitae*

CONTACT INFORMATION *Email:* bowman@nyu.edu

DATE COMPILED Tuesday 23rd June, 2020

RESEARCH AREA natural language processing, computational linguistics, representation learning

ACADEMIC POSITIONS **New York University**, New York, New York USA

Assistant Professor of Linguistics and Data Science **2016–2020**
Assistant Professor of Linguistics, Data Science, and Computer Science **2020–**
Joint appointment in the Department of Linguistics and the Center for Data Science, associated appointment in the Courant Institute’s Department of Computer Science.

EDUCATION **Stanford University**, Stanford, California USA

Ph.D., Linguistics **2011–2016**
Dissertation title: *Modeling Natural Language Semantics in Learned Representations*
Dissertation committee: Christopher Manning and Christopher Potts (chairs),
Thomas F. Icard, Percy Liang

The University of Chicago, Chicago, Illinois USA

Combined BA/MA, Linguistics **2007–2011**
MA Thesis: *Vowel harmony, opacity, and finite-state OT*
Advisor: Jason Riggle

FUNDING AND AWARDS Google Cloud Platform Research Credit Grant (in kind, retail value \$5,000), 2020
Samuel R. Bowman (PI)

Intuit Gift Research Support (\$75,000)
Samuel R. Bowman (PI)

Moore-Sloan Data Science Environment funding (\$10,000), 2019
Subaward funding through an internal competition in support of a three-day mini-course for NYC-area undergraduates from underrepresented groups.
Samuel R. Bowman (PI), Nikita Nangia, Phu Mon Htut, Alex Wang

DeepMind Gift Research Support (\$3,000), 2019
In support of expenses for the *GLUE* and *SuperGLUE* projects.
Samuel R. Bowman (PI)

*SEM Best Paper Award, 2019
For: Probing What Different NLP Tasks Teach Machines about Function Word Comprehension

Philanthropic research funding (\$2,000,000), 2019
From Eric and Wendy Schmidt, by recommendation of the Schmidt Futures program.
For: Massive Open Annotated Data for Natural Language Understanding
Samuel R. Bowman (PI)

NSF CRII Program (award no. 1850208, \$174,894), 2019
For: CRII: RI: Can Low-Bias Machine Learners Acquire English Grammar? Deep Learning and Linguistic Acceptability
Samuel R. Bowman (PI)

NVIDIA GPU Grant (in kind, retail value \$1,400), 2019 Samuel R. Bowman (PI)

ACL Top Reviewer, 2018
Given to a few percent of the reviewers for the conference in recognition for unusually constructive reviews.

Selected as team leader for the Jelinek Summer Workshop (JSALT), 2018
Funded six-week language technology research event for large teams.
For: General-Purpose Sentence Representation Learning
Samuel R. Bowman (Lead PI), Ellie Pavlick (Co-PI)

Samsung Research (sponsored research, \$2.2 million), 2017–2020
For: Improving Deep Learning using Latent Structure
Kyunghyun Cho (PI), Joan Bruna (PI), Samuel R. Bowman (PI)

Moore-Sloan Data Science Environment Seed Grant (NYU internal, \$25,000), 2017–2018
For: Semi-supervised NLP Techniques for Automated Cybercrime Forum Analysis
Damon McCoy (PI), Samuel R. Bowman (PI)

University Research Challenge Fund (NYU internal seed grant, \$14,000), 2017–2018
For: Unsupervised Sentence Representation Learning
Samuel R. Bowman (PI)

Google Faculty Research Award (\$91,000), 2017–2018
For: A Corpus and Challenge for Cross-Genre Natural Language Understanding
Samuel R. Bowman (PI), Angeliki Lazaridou

Tencent Holdings Faculty Research Award (\$50,000), 2017
Samuel R. Bowman (PI)

NVIDIA GPU Grant (in kind, retail value \$3,600), 2017
Samuel R. Bowman (PI)

EMNLP Best New Data Set or Resource Award, 2015
For: A large annotated corpus for learning natural language inference

Google Faculty Research Award (\$62,000), 2015–2016
For: Representations and Resources for Wide-Coverage Natural Language Inference
Christopher Potts (PI), Christopher D. Manning, Gabor Angeli, Samuel R. Bowman, Kelvin Gu

Donna Schweers and Thomas Geiser Fellowship, 2014–2017
Through the Stanford Interdisciplinary Graduate Fellowship, a competitive internal funding program a with 10% acceptance rate

Stanford University PhD Fellowship, 2011–2014

Google European Doctoral Fellowship in Speech Technology, 2011
Declined: offer was made after I had enrolled in a US PhD program

Phi Beta Kappa, The University of Chicago, 2011

College Honors, The University of Chicago, 2011

Fellowship for the Center for Language and Speech Processing (CLSP) Summer School in Human Language Technology at The Johns Hopkins University, 2010

USA National Merit Scholar, 2007–2011

PEER-REVIEWED
PUBLICATIONS

2020. Anhad Mohananey, Katharina Kann, and Samuel R. Bowman. Self-Training for Un-supervised Parsing with PRPN. *Proceedings of The 16th International Conference on Parsing Technologies (IWPT)*.

2020. Alex Warstadt and Samuel R. Bowman. Do self-supervised neural networks acquire a bias towards structural linguistic generalizations? *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)*.

2020. {Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut}, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel R. Bowman. Intermediate-Task Transfer Learning with Pretrained Language Models: When and Why Does It Work? *Proceedings of the Annual Conference of the Association for Computational Linguistics (ACL)*.

2020. {Yada Pruksachatkun, Phil Yeres}, Haokun Liu, Jason Phang, Phu Mon Htut, Alex Wang, Ian Tenney, and Samuel R. Bowman. jiant: A Software Toolkit for Research on General-Purpose Text Understanding Models. *Proceedings of the Annual Conference of the Association for Computational Linguistics (ACL), demonstration track*.

2020. Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. BLiMP: A Benchmark of Linguistic Minimal Pairs for English. *Transactions of the Association for Computational Linguistics (TACL)*.

2020. Katharina Kann, Samuel R. Bowman, and Kyunghyun Cho. Learning to Learn Morphological Inflection for Resource-Poor Languages. *Proceedings of the AAAI Conference on Artificial Intelligence*.

2019. {Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh}, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems. *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.

2019. Nishant Subramani, Samuel R. Bowman, and Kyunghyun Cho. Can Unconditional Language Models Recover Arbitrary Sentences? *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.

2019. Katharina Kann, Anhad Mohananey, Samuel R. Bowman, and Kyunghyun Cho. Neural Unsupervised Parsing Beyond English. *Proceedings of The Workshop on Deep Learning for Low-Resource NLP (DeepLo)*.

2019. {Alex Warstadt, Yu Cao, Ioana Grosu, Wei Peng, Hagen Blix, Yining Nie, Anna Alsop, Shikha Bordia, Haokun Liu, Alicia Parrish, Sheng-Fu Wang, Jason Phang, Anhad Mohananey, Phu Mon Htut, Paloma Jeretic} and Samuel R. Bowman. Investigating BERTs Knowledge of Language: Five Analysis Methods with NPIs. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

2019. Katharina Kann, Kyunghyun Cho, and Samuel R. Bowman. Towards Realistic Practices In Low-Resource Natural Language Processing: The Development Set. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
2019. Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. Neural Network Acceptability Judgments. *Transactions of the Association for Computational Linguistics (TACL)*.
2019. Alex Wang, Jan Hula, Patrick Xia, Raghavendra Pappagari, R. Thomas McCoy, Roma Patel, Najoung Kim, Ian Tenney, Yinghui Huang, Katherin Yu, Shuning Jin, Berlin Chen, Benjamin Van Durme, Edouard Grave, Ellie Pavlick, and Samuel R. Bowman. How to Get Past Sesame Street: Sentence-Level Pretraining Beyond Language Modeling. *Proceedings of the Annual Conference of the Association for Computational Linguistics (ACL)*.
2019. Nikita Nangia and Samuel R. Bowman. Human vs. Muppet: A Conservative Estimate of Human Performance on the GLUE Benchmark. *Proceedings of the Annual Conference of the Association for Computational Linguistics (ACL)*.
2019. Najoung Kim, Roma Patel, Adam Poliak, Patrick Xia, Alex Wang, Tom McCoy, Ian Tenney, Alexis Ross, Tal Linzen, Benjamin Van Durme, Samuel R. Bowman, and Ellie Pavlick. Probing What Different NLP Tasks Teach Machines about Function Word Comprehension. **Best Paper Award**. *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM)*.
2019. Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. On Measuring Social Biases in Sentence Encoders. *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
2019. Shikha Bordia and Samuel R. Bowman. Identifying and Reducing Gender Bias in Word-Level Language Models. *Proceedings of the NAACL Student Research Workshop*.
2019. Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, Ellie Pavlick. What do you learn from context? Probing for sentence structure in contextualized word representations. *Proceedings of the International Conference on Learning Representations (ICLR)*.
2019. Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
2019. Katharina Kann, Alex Warstadt, Adina Williams, and Samuel R. Bowman. Verb Argument Structure Alternations in Word and Sentence Embeddings. *Proceedings of the Society for Computation in Linguistics (SCiL)*.
2018. Yun Chen, Victor O.K. Li, Kyunghyun Cho, and Samuel R. Bowman. A Stable and Effective Learning Strategy for Trainable Greedy Decoding. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
2018. Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. XNLI: Cross-lingual Sentence Understanding through Inference. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

2018. Phu Mon Htut, Kyunghyun Cho, and Samuel R. Bowman. Grammar Induction with Neural Language Models: An Unusual Replication. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP, short paper)*.
2018. WooJin Chung, Sheng-Fu Wang, and Samuel R. Bowman. The Lifted Matrix-Space Model for Semantic Composition. *Proceedings of the Twenty-Second Conference on Computational Natural Language Learning (CoNLL)*.
2018. Nikita Nangia and Samuel R. Bowman. ListOps: A Diagnostic Dataset for Latent Tree Learning. *Proceedings of the NAACL Student Research Workshop*.
2018. Phu Mon Htut, Samuel R. Bowman, and Kyunghyun Cho. Training a Ranking Function for Open-Domain Question Answering. *Proceedings of the NAACL Student Research Workshop*.
2018. Adina Williams, Nikita Nangia, and Samuel R. Bowman. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
2018. Adina Williams, Andrew Drozdov, and Samuel R. Bowman. Do latent tree learning models identify meaningful structure in sentences? *Transactions of the Association for Computational Linguistics (TACL)*, vol. 6, pp. 253–267.
2018. Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R. Bowman, and Noah A. Smith. Annotation Artifacts in Natural Language Inference Data. *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL, short paper)*.
2018. Yichen Gong and Samuel R. Bowman. Ruminating Reader: Reasoning with Gated Multi-Hop Attention. *Proceedings of the ACL Workshop on Machine Reading for Question Answering*.
2017. Rohan Kshirsagar, Robert Morris, and Samuel R. Bowman. Detecting and Explaining Crisis. *Proceedings of the 2017 Computational Linguistics and Clinical Psychology Workshop*.
2017. Sebastian Brarda, Philip Yeres, and Samuel R. Bowman. Sequential Attention. *Proceedings of the 2nd Workshop on Representation Learning for NLP*.
2016. Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Jozefowicz, and Samy Bengio. Generating Sentences from a Continuous Space. *Proceedings of the Twentieth Conference on Computational Natural Language Learning (CoNLL)*.
2016. Samuel R. Bowman, Jon Gauthier, Abhinav Rastogi, Raghav Gupta, Christopher D. Manning, and Christopher Potts. A Fast Unified Model for Parsing and Sentence Understanding. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*.
2015. Samuel R. Bowman, Christopher D. Manning, and Christopher Potts. Tree-structured composition in neural networks without tree-structured architectures. *Proceedings of the NIPS 2015 Workshop on Cognitive Computation: Integrating Neural and Symbolic Approaches*.
2015. Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. A large annotated corpus for learning natural language inference. **Best New Data Set or Resource Award**. *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

2015. Samuel R. Bowman, Christopher Potts, and Christopher D. Manning. Recursive Neural Networks Can Learn Logical Semantics. *Proceedings of the 3rd Workshop on Continuous Vector Space Models and their Compositionality*.

2015. Samuel R. Bowman, Christopher Potts, and Christopher D. Manning. Learning Distributed Word Representations for Natural Logic Reasoning. *Proceedings of the AAAI Spring Symposium on Knowledge Representation and Reasoning*.

2014. Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, John Bauer, Samuel R. Bowman, and Christopher D. Manning. A gold standard dependency corpus for English. *Proceedings of the 9th International Conference on Language Resources and Evaluation (LREC)*.

2013. Marie-Catherine de Marneffe, Miriam Connor, Natalia Silveira, Samuel R. Bowman, Timothy Dozat, and Christopher D. Manning. More constructions, more genres: Extending Stanford Dependencies. *Proceedings of the 13th International Conference on Dependency Linguistics*.

2012. Samuel R. Bowman and Harshit Chopra. Automatic Animacy Classification. *Proceedings of the NAACL Student Research Workshop*.

2011. Geoffrey Zweig, Les Atlas, Kris Demuynck, Fei Sha, Patrick Nguyen, Dirk van Compernelle, Damianos Karakos, Pascal Clark, Meihong Wang, Gregory Sell, Samuel Thomas, Samuel Bowman and Justine Kao. Speech Recognition with Segmental Conditional Random Fields: A Summary of the JHU CLSP 2010 Summer Workshop. *Proceedings of the 36th International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.

2010. Sam Bowman and Karen Livescu. Modeling pronunciation variation with context-dependent articulatory feature decision trees. *Proceedings of INTERSPEECH*.

PUBLICATIONS
REVIEWED AS
ABSTRACTS

2013. Samuel R. Bowman and Benjamin Lokshin. Idiosyncratic transparency in Kazakh vowel harmony. *Proceedings of the Annual Meetings on Phonology*.

2013. Samuel R. Bowman. Two arguments for vowel harmony by trigger competition. *Proceedings of the 49th Annual Meeting of the Chicago Linguistic Society (CLS)*.

OTHER
PUBLICATIONS

2018. Georgiana Dinu, Miguel Ballesteros, Avirup Sil, Samuel R. Bowman, Wael Hamza, Anders Sogaard, Tahira Naseem, and Yoav Goldberg. *Proceedings of the Workshop on the Relevance of Linguistic Structure in Neural Architectures for NLP*. Association for Computational Linguistics.

2017. Nikita Nangia, Adina Williams, Angeliki Lazaridou, and Samuel R. Bowman. The RepEval 2017 Shared Task: Multi-Genre Natural Language Inference with Sentence Representations. *Proceedings of the 2nd Workshop on Evaluating Vector Space Representations for NLP*.

2017. Samuel R. Bowman, Yoav Goldberg, Felix Hill, Angeliki Lazaridou, Omer Levy, Roi Reichart, and Anders Sogaard. *Proceedings of the 2nd Workshop on Evaluating Vector Space Representations for NLP*. Association for Computational Linguistics.

2017. Vasant Dhar and Samuel R. Bowman. A Perspective on Natural Language Understanding Capability: An Interview with Sam Bowman. *Big Data*. 5(1): 5–11.

INVITED
PRESENTATIONS

2020. Evaluating Recent Progress Toward General-Purpose Language Understanding Models. Google Research Virtual Talk Series. Remote Presentation.

2020. Research with HuggingFace Transformers (software library). Panel co-moderator. International Conference on Learning Representations. Virtual Conference.

2020. Learning Acceptability Judgments from Text Alone. UC Irvine Language Science Colloquium. Remote Presentation.

2020. Learning Acceptability Judgments from Text Alone. Max Planck–NYU Center for Language, Music, and Emotion (CLaME) workshop on Language. New York, NY. (Wrote talk; called in sick; talk was delivered by collaborator Alex Warstadt.)

2019. Evaluating Recent Progress Toward General-Purpose Language Understanding Models. University of Washington. Seattle, WA.

2019. Evaluating Recent Progress Toward General-Purpose Language Understanding Models. Allen Institute for Artificial Intelligence. Seattle, WA.

2019. Learning Acceptability Judgments from Raw Text Alone. The Ohio State University Linguistics Colloquium. Columbus, OH.

2019. Task-Independent Language Understanding. IBM Research AI. Yorktown Heights, NY.

2019. Task-Independent Language Understanding. Cornell Learning Machines Seminar Series. New York, NY.

2019. Panel discussion: A linguist, an NLP engineer, and a psycholinguist walk into a bar... The Workshop on Evaluating Vector-Space Representations for NLP (RepEval). NAACL. Minneapolis, MN.

2019. Task-Independent Sentence Understanding. Joint (joint-joint?) invited talk at the Joint Conference on Lexical and Computational Semantics (*SEM) and the International Workshop on Semantic Evaluation (SemEval). NAACL. Minneapolis, MN.

2019. Task-Independent Sentence Understanding. Grammarly. New York, NY.

2019. Task-Independent Sentence Understanding. University of Pennsylvania NLP Talk Series. Philadelphia, PA.

2019. Task-Independent Sentence Understanding. Toyota Technological Institute at Chicago (TTI-C) Colloquium. Chicago, IL.

2019. What Deep Learning Researchers Talk about when They Talk about Sentence Understanding. Panel “What should linguists know about Natural Language Processing and Machine Learning?” Society for Computation in Linguistics (SCiL). LSA. New York, NY.

2018. Task-Independent Sentence Understanding Models. AI Summit. New York, NY.

2018. Toward Task-Independent Sentence Understanding. Samsung AI Forum. Seoul, Korea.

2018. Toward Task-Independent Sentence Understanding. KAIST Computer Science. Daejeon, Korea.

2018. Toward Task-Independent Sentence Understanding. Seoul National University Linguistics Colloquium. Seoul, Korea.

2018. Evaluating the Semantic and Syntactic Abilities of Neural Network Models. Seoul National University Linguistics Colloquium. Seoul, Korea.

2018. Interviewee. NLP Highlights podcast. Allen Institute for AI. Seattle, WA.

2018. General Purpose Sentence Representation Learning at JSALT 2018. Jelinek Memorial Summer Workshop (JSALT) Closing Recap Talk. Johns Hopkins University. Baltimore, MD.

2018. Sentence Representation Learning: Evaluation and the State of the Art. Lecture at the JHU Summer School on Human Language Technology. Baltimore, MD.

2018. GLUE: Toward Task-Independent Sentence Understanding. Invited talk at the Workshop on New Forms of Generalization in Deep Learning and Natural Language Processing. NAACL. New Orleans, LA.

2018. Two Early Efforts toward Using Deep Learning in Syntax and Semantics. CoAStAL NLP. University of Copenhagen. Copenhagen, Denmark.

2018. Two Early Efforts toward Using Deep Learning in Syntax and Semantics *and* Sentence Understanding with Neural Networks and Natural Language Inference. Invited talk series at Centre of Linguistic Theory and Studies in Probability. University of Gothenburg. Gothenburg, Sweden.

2018. Teaching Neural Networks Compositional Semantics. Joint invited talk at the workshop “Perceptrons and Syntactic Structures at Sixty” and the first annual meeting of the Society for Computation in Linguistics (SCiL). Salt Lake City, UT.

2017. Sentence Understanding with Neural Networks and Natural Language Inference. Language Technologies Institute Colloquium. Carnegie Mellon University. Pittsburgh, PA.

2017. Participant in panel “Current Status and Evolution of AI Voice Assistants.” Samsung Global AI Forum. Samsung 837. New York, NY.

2017. Sentence Understanding with Neural Networks and Natural Language Inference. Lunch Seminar. Insight AI. New York, NY.

2017. Semi-supervised learning and sentence understanding: platitudes and provocations. Google Natural Language Understanding Workshop. Google. New York, NY.

2017. Participant in panel “AI in the Newsroom: Technology and Practical Application.” Artificial Intelligence: Practice and Implications for Journalism. Tow Center for Digital Journalism, Columbia University. New York, NY.

2017. Sentence Understanding with Neural Networks and Natural Language Inference. NLP Speaker Series. Columbia University. New York, NY.

2017. Sentence Understanding with Neural Networks and Natural Language Inference. Linguistics Colloquium. State University of New York. Stony Brook, NY.

2017. Sentence Understanding with Neural Networks and Natural Language Inference. Computational Linguistics and Information Processing (CLIP) Colloquium. University of Maryland. College Park, MD.

2016. Learning neural networks for sentence understanding with the Stanford NLI corpus. Google. New York, NY.

2016. Learning neural networks for sentence understanding with the Stanford NLI corpus. Forum for Artificial Intelligence. University of Texas. Austin, TX.

2016. Learning neural networks for sentence understanding with the Stanford NLI corpus. Facebook AI Research. New York, NY.

2016. Learning neural networks for sentence understanding with the Stanford NLI corpus. South England NLP Meetup. University College London. London, UK.

2016. Modeling Natural Language Semantics with Learned Representations. Nuance Sunnyvale seminar series. Nuance Communications. Sunnyvale, CA.

2016. Modeling Natural Language Semantics with Learned Representations. Berkeley NLP Group weekly meeting. UC Berkeley. Berkeley, CA.

2016. Modeling Natural Language Semantics with Learned Representations. Linguistics–Center for Data Science joint colloquium. New York University. New York, NY.

2016. Modeling Natural Language Semantics with Learned Representations. Linguistics colloquium. Northwestern University. Evanston, IL.

2016. Modeling Natural Language Semantics with Learned Representations. College of Information and Computer Sciences colloquium. University of Massachusetts Amherst. Amherst, MA.

2016. Modeling Natural Language Semantics with Learned Representations. Linguistics–Computer Science joint colloquium. Georgetown University. Washington, DC.

2014. Samuel R. Bowman. Can recursive neural tensor networks learn logical reasoning? Nuance Sunnyvale seminar series. Nuance Communications. Sunnyvale, CA.

2014. Samuel R. Bowman. Can recursive neural networks learn to do natural language inference? 3rd CSLI Workshop on Logic, Rationality & Intelligent Interaction. Center for the Study of Language and Information, Stanford University. Stanford, CA.

2014. Samuel R. Bowman. Transparent vowels in Agreement by Correspondence: Open issues. Conference on Agreement by Correspondence (ABC↔C). UC Berkeley. Berkeley, CA.

2013. Samuel R. Bowman and Benjamin Lokshin. Idiosyncratic transparency in Kazakh vowel harmony. Phorum. UC Berkeley. Berkeley, CA.

2013. Samuel R. Bowman. Two arguments for vowel harmony by trigger competition. The University of Edinburgh Phonology/Phonetics Workshop. Edinburgh, Scotland.

REFEREED
PRESENTATIONS
WITHOUT
ASSOCIATED
PROCEEDINGS
PAPERS

2020. Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang and Samuel R. Bowman. BLiMP: A Benchmark of Linguistic Minimal Pairs for English. Society for Computation in Linguistics (SCiL, non-archival abstract). New Orleans, LA.

2019. Jason Phang and Samuel R. Bowman. Semi-Supervised Learning for Natural Language Understanding: A Negative Result. Natural Language, Dialog and Speech (NDS) Symposium. New York, NY.

2019. Leslie Huang, Kyunghyun Cho, and Samuel R. Bowman. Automatic extraction of polysemous words from contextualized embeddings. Natural Language, Dialog and Speech (NDS) Symposium. New York, NY.
2019. Samuel R. Bowman and Xiaodan Zhu. Deep Learning for Natural Language Inference. **Three-Hour Tutorial**. Conference of the North American Chapter of the Association for Computational Linguistics (NAACL). Minneapolis, MN.
2018. Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. The Workshop on Analyzing and interpreting neural networks for NLP (BlackboxNLP, cross-submission track). Brussels, Belgium.
2018. Kelly W. Zhang and Samuel R. Bowman. Language Modeling Teaches You More Syntax than Translation Does: Lessons Learned Through Auxiliary Task Analysis. The Workshop on Analyzing and interpreting neural networks for NLP (BlackboxNLP, cross-submission track). Brussels, Belgium.
2018. Phu Mon Htut, Kyunghyun Cho, and Samuel R. Bowman. Grammar Induction with Neural Language Models: An Unusual Replication. The Workshop on Analyzing and interpreting neural networks for NLP (BlackboxNLP, non-archival track). Brussels, Belgium.
2018. Yun Chen, Kyunghyun Cho, Samuel R. Bowman, and Victor O.K. Li. Stable and Effective Trainable Greedy Decoding for Sequence to Sequence Learning. International Conference on Learning Representations (ICLR, workshop track). Vancouver, BC.
2018. Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R. Bowman, and Noah A. Smith. Annotation Artifacts in Natural Language Inference Data. Workshop on New Forms of Generalization in Deep Learning and Natural Language Processing (cross-submission track). New Orleans, LA.
2018. Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R. Bowman, and Noah A. Smith. Annotation Artifacts in Natural Language Inference Data. 5th Pacific Northwest Regional NLP Workshop (cross-submission track). Redmond, WA.
2018. Alex Warstadt and Samuel R. Bowman. Acceptability Judgments from a Neural Network. The 92nd Annual Meeting of the Linguistic Society of America (LSA). Salt Lake City, UT.
2017. Andrew Drozdov and Samuel R. Bowman. The Coadaptation Problem when Learning How and What to Compose. The 2nd Workshop on Representation Learning for NLP. Vancouver, BC, Canada.
2016. Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Jozefowicz, and Samy Bengio. Generating Sentences from a Continuous Space. International Conference on Learning Representations (ICLR, workshop track). San Juan, PR.
2015. Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. A large annotated corpus of entailments and contradictions. 8th California Universities Semantics and Pragmatics Workshop (CUSP). Stanford, CA.
2014. Samuel R. Bowman, Christopher Potts, and Christopher D. Manning. Recursive Neural Networks for Learning Logical Semantics. BayLearn Symposium. Berkeley, CA.

2014. Samuel R. Bowman. Can recursive neural tensor networks learn logical reasoning? International Conference on Learning Representations (ICLR), workshop track. Banff, AB, Canada.

2013. Samuel R. Bowman and Benjamin Lokshin. Idiosyncratic transparency in Kazakh vowel harmony. The Ninth Workshop on Altaic Formal Linguistics (WAFL). Ithaca, NY.

2013. Samuel R. Bowman. Two arguments for vowel harmony by trigger competition. The 21st Manchester Phonology Meeting (mfm). Manchester, UK.

2013. Samuel R. Bowman. Seto vowel harmony and neutral vowels. The 87th Annual Meeting of the Linguistic Society of America (LSA). Boston, MA.

2012. Robert Podesva, Annette D'Onofrio, Eric Acton, Samuel R. Bowman, Jeremy Calder, Hsin-Chang Chen, Benjamin Lokshin, and Janneke Van Hofwegen. Linguistic and social effects on perceptions of voice onset time in Korean stops. The 164th Meeting of the Acoustical Society of America (ASA). Kansas City, MO.

TEMPORARY
RESEARCH
POSITIONS

Google, Inc., Mountain View, CA and New York, NY

Visiting Researcher Program
Software Engineering Intern

Summer/Fall 2019
Summers 2012, 2013, 2014, 2015

Johns Hopkins University, Baltimore MD

Team Lead, CLP JSALT Summer Workshop (Sentence Representations Team) **Summer 2018**
Undergraduate Researcher, CLSP Summer Workshop (CRF Speech Team) **Summer 2010**

TEACHING
EXPERIENCE

New York University, New York, New York USA

Instructor **Fall 2020**
Seminar in Semantics: Topic TBA (LING-GA 3340)

Instructor **Spring 2020**
Natural Language Understanding and Computational Semantics (DS-GA 1012/LING-GA 1012)
Machine Learning for Language Understanding (LING-UA 52/DS-UA 203)

Instructor **Spring 2019**
Natural Language Understanding and Computational Semantics (DS-GA 1012/LING-GA 1012, with Katharina Kann)
Seminar in Semantics: Team Project on Linguistic Knowledge in Reusable Sentence Encoders (LING-GA 3340)

Instructor **Fall 2018**
Patterns in Language (LING-UA 6)

Instructor **Spring 2018**
Natural Language Understanding and Computational Semantics (DS-GA 1012/LING-GA 1012)
Seminar in Semantics: Deep Learning in Semantics (LING-GA 3340, with Chris Barker)

Instructor **Fall 2017**
Natural Language Processing with Representation Learning (DS-GA 1011, with Kyunghyun Cho)

Instructor **Fall 2016**
Natural Language Understanding with Distributed Representations (DS-GA 3001-001)
Seminar in Semantics: Artificial Neural Networks (LING-GA 3340)

Stanford University, Stanford, California USA

Teaching Assistant **Spring 2014**
Natural Language Understanding (CS 224U/LINGUIST 188/288)
Instructors: Bill MacCartney (Google, Inc.) and Christopher Potts
Guest lecture: Recursive neural networks for semantic interpretation

Teaching Assistant **Winter 2014**
From Languages to Information (CS 124/LINGUIST 180)
Instructor: Dan Jurafsky
Note: Coursera-based ‘flipped-classroom’ course.

REVIEWING

2018–2020. Standing reviewer pool, Transactions of the Association for Computational Linguistics (TACL).

2020. Ninth Joint Conference on Lexical and Computational Semantics (*SEM).

2020. ACL Student Research Workshop (ACL SRW).

2020. Ad-Hoc Reviewer, Directorate for Computer & Information Science & Engineering (CISE), National Science Foundation (NSF).

2017, 2020. Ad-Hoc Reviewer, Directorate for Social, Behavioral & Economic Sciences (SBE), National Science Foundation (NSF).

2015, 2016, 2017, 2018, 2019, 2020. Conference on Empirical Methods in Natural Language Processing (EMNLP).

2020. International Joint Conferences on Artificial Intelligence (IJCAI).

2019, 2020. The Workshop on Representation Learning for NLP (RepL4NLP).

2018, 2019, 2020. Society for Computation in Linguistics (SCiL) Annual Meeting.

2019. New York Academy of Sciences Natural Language, Dialog & Speech (NDS) Symposium.

2019. BlackboxNLP: Analyzing and interpreting neural networks for NLP.

2016, 2017, 2018, 2019. Annual Meeting of the Association for Computational Linguistics (ACL).

2017, 2019. The Workshop on Evaluating Vector-Space Representations for NLP (RepEval).

2017, 2018, 2019. Conference on Neural Information Processing Systems (NeurIPS).

2018. Computational Linguistics (CL).

2018. International Conference on Computational Linguistics (COLING).

2017, 2018. International Conference on Machine Learning (ICML).

2015, 2016, 2017, 2018. International Conference on Learning Representations (ICLR).

2018. Conference on Computational Natural Language Learning (CoNLL).

2018. Panel reviewer, Directorate for Computer & Information Science & Engineering (CISE), National Science Foundation (NSF).

2017. Journal of Linguistic Issues in Language Technology (LiLT).

2017. AAAI Conference on Artificial Intelligence.

2017. Journal of Artificial Intelligence Research (JAIR).

2015. Journal of Natural Language Engineering.

2014–2015. Semantics and Linguistic Theory (SALT) 25.

2014. AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches.

OTHER EXTERNAL SERVICE

2018–. Co-organizer, GLUE and SuperGLUE benchmark competitions.

2017–. Consultant/Advisor, ASAPP (NLP for customer service; compensated role).

2019, 2020, 2021. Area Chair, International Conference on Learning Representations (ICLR).

2020. Area Chair, Conference on Neural Information Processing Systems (NeurIPS).

2020. Area Chair, Conference on Computational Natural Language Learning (CoNLL).

2020. Steering Committee Member, Workshop on Simple and Efficient Natural Language Processing (SustainLP, hosted at EMNLP).

2020. Senior Area Chair for Textual Inference and Other Areas of Semantics, Annual Meeting of the Association for Computational Linguistics (ACL).

2020. Best Paper Selection Committee Member (general and executive committees), International Conference on Learning Representations (ICLR).

2019, 2020. Steering Committee Member, New York Academy of Sciences Natural Language, Dialog & Speech (NDS) Symposium.

2019. Area Chair for Semantics, Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).

2017–2018. Consultant/Advisor, Koko (NLP for mental health; compensated role).

2018. Faculty Advisor to the Student Research Workshop, Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).

2018. Co-organizer, Workshop on the Relevance of Linguistic Structure in Neural NLP (hosted at ACL).

2018. Area Chair for Formal Semantics, Seventh Joint Conference on Lexical and Computational Semantics (*SEM).

2017. Co-organizer and Shared Task Chair, The Second Workshop on Evaluating Vector Space Representations for NLP (RepEval 2, hosted at EMNLP).

UNIVERSITY AND
DEPARTMENT
SERVICE

2017-. Application Reader, Data Science Fellows program, Center for Data Science.

2017-. Track Advisor and Admissions Reader, NLP track, Data Science MS program, Center for Data Science.

2017-. Co-organizer, NLP and Text as Data Speaker Series (weekly colloquium, with A. Spirling), Center for Data Science.

2016-. Member, PhD Admissions Committee, Center for Data Science.

2016-. Member, Graduate Curriculum Committee, Center for Data Science.

2020. Member, Annual Merit Review Committee, Department of Linguistics.

2020. Track Head for Machine Learning, PhD Admissions Committee, Center for Data Science.

2019-2020. Member, Computational Social Science Faculty Search Committee, Center for Data Science and Department of Linguistics.

2019. Faculty organizer, NYC AI Workshop (three-day mini-course for NYC-area undergraduates from underrepresented groups), Center for Data Science.

2019. Member, Machine Learning Faculty Search Committee, Department of Computer Science and Center for Data Science. (Hired Lerrel Pinto and Andrew Wilson.)

2017. Chair, Machine Learning Faculty Search Committee, Department of Computer Science and Center for Data Science. (Hired Rajesh Ranganath.)

2016-2017. Faculty advisor, Data Future Lab incubator.

2014-2015. Co-founder and organizer, Stanford Natural Logic and Natural Language Inference Reading Group.

2014-2015. Member, Graduate Admissions Committee, Stanford University Department of Linguistics.

2013-2014. Member, Graduate Studies Committee, Stanford University Department of Linguistics.

2013-2014. Social organizer, Stanford University Natural Language Processing Group.

2012-2014. Organizer, Stanford Phonetics and Phonology Workshop.

2012-2013. Corpus TA and librarian, Stanford University Department of Linguistics.

PHD ADVISING

2018-. Jason Phang (PhD, transferred from MS program in 2019, Data Science, advisor)

2017-. Phu Mon Htut (PhD, Data Science, co-advisor with Kyunghyun Cho)

2017-. Alex Wang (PhD, Computer Science, co-advisor with Kyunghyun Cho, NSF GRFP recipient)

2017-. Alex Warstadt (PhD, Linguistics, co-advisor with Chris Barker)

2017-. Nikita Nangia (PhD, transferred from MS program in 2018, Data Science, advisor)

2017-2019. Nishant Subramani (PhD, Computer Science, co-advisor with Kyunghyun Cho)

2017-2018. Adina Williams (PhD, Linguistics, research supervisor alongside dissertation chair Liina Pyllkanen, now at Facebook AI Research)

DISSERTATION
COMMITTEES

2019-. Sheng-Fu Wang (PhD, Linguistics)

2019-. Elman Mansimov (PhD, Computer Science)

2019-. Sébastien Jean (PhD, Computer Science)

2019. Zhouhan Lin (PhD, Informatics, The University of Montreal/Montreal Institute for Learning Algorithms)

2017-2019. Lisheng Fu (PhD, Computer Science)

2017-2018. Rodrigo Nogueira (PhD, Computer Science, proposal only)

2017-2018. Adina Williams (PhD, Linguistics, now at Facebook AI Research)

2017-2018. Xiang Zhang (PhD, Computer Science, now at Element AI)

2017. Yacine Jernite (PhD, Computer Science, now at Facebook AI Research)

2017. Arvind R. Neelakantan (PhD, Computer Science, UMass Amherst, now at Google Brain)

QUALIFYING EXAM
COMMITTEES

2019. Dima Taji (Computer Science, chair)

2019. Rasika Bhalerao (Computer Science, NYU Tandon School of Engineering)

2019. Nishant Subramani (Computer Science)

2019. Alex Wang (Computer Science)

2019. Phu Mon Htut (Data Science, co-chair)

2018. Lingfan Yu (Computer Science)

2018. Elman Mansimov (Computer Science)

2018. Sébastien Jean (Computer Science, chair)

QUALIFYING PAPER 2019. Alex Warstadt (PhD, Linguistics, chair)
 COMMITTEES
 2018. Sheng-Fu Wang (PhD, Linguistics, co-chair with Maria Gouskova)
 2017. WooJin Chung (PhD, Linguistics, chair)

OTHER ADVISING 2019-. Phil Yeres (Full-time research engineer, Data Science, supervisor)
 2019-. Clara Vania (Postdoc, Data Science, supervisor)
 2019-2020. Haokun Liu (MS, Computer Science, research supervisor)
 2019-2020. Yada Pruksachatkun (MS, Data Science, research supervisor)
 2018-2019. Jason Phang (MS, Data Science, research supervisor)
 2019. Leslie Huang (PhD, Data Science, research rotation supervisor)
 2018-2019. Katharina Kann (Postdoc, Data Science, co-supervisor with Kyunghyun Cho)
 2018-2019. Anhad Mohananey (MS, Computer Science, research supervisor)
 2018-2019. Shikha Bordia (MS, Computer Science, research supervisor)
 2016-2018. Yichen Gong (BS, Computer Science, Tandon School of Engineering, research supervisor, now at Horizon Robotics)
 2016-2018. Kelly Zhang (BS, Computer Science, research supervisor, now Harvard PhD student & NSF GRFP recipient)
 2018. Amanpreet Singh (MS, Computer Science, research supervisor, now at Facebook AI Research)
 2018. Thibault Févry (MS, Data Science, research supervisor, now a Google AI Resident)
 2018. Yun Chen (visiting PhD student, University of Hong Kong, co-host with Kyunghyun Cho, now at Huawei Noah's Ark Lab)
 2016-2017. Andrew Drozdov (MS, Computer Science, research supervisor, now U. Mass. Amherst PhD student)
 2017. Haoyue Shi (visiting undergraduate, Computer Science, Peking University, host, now TTI-Chicago PhD student)
 2017. Xiaonan Zhao (MS, Computer Science, research supervisor, now at Amazon)

MEDIA COVERAGE 2019. Google achieves state-of-the-art NLP performance with an enormous language model and data set. VentureBeat.
 OF MY WORK
 2019. Machines Beat Humans on a Reading Test. But Do They Understand? Quanta, reprinted at Wired.

2019. Google AIs ALBERT claims top spot in multiple NLP performance benchmarks. VentureBeat.

2019. Facebook AIs RoBERTa improves Googles BERT pretraining methods. VentureBeat.

2019. AI researchers launch SuperGLUE, a rigorous benchmark for language understanding. VentureBeat.

MEDIA COVERAGE
AS THIRD-PARTY
EXPERT

2019. An AI helped us write this article. Vox Future Perfect.

2019. Computers could write full - and wrong - news stories - researchers. New York Times and elsewhere via Reuters.

2019. Musk-backed OpenAI delays publishing research over ethics concerns. Financial Times.

2019. The AI That Can Write a Fake News Story From a Handful of Words. Bloomberg.

2018. Finally, a Machine That Can Finish Your Sentence. The New York Times.