

Samuel R. Bowman, *Curriculum Vitae*

CONTACT INFORMATION *Email:* bowman@nyu.edu

DATE COMPILED Saturday 28th July, 2018

RESEARCH INTERESTS natural language processing, representation learning, computational semantics

POSITIONS HELD **New York University**, New York, New York USA

Assistant Professor **2016–**
Department of Linguistics and Center for Data Science (primary joint appointment)
Computer Science Department, Courant Institute of Mathematical Sciences (affiliate, 2017–)

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

PhD, Linguistics (Stanford NLP Group) **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

The Johns Hopkins University, Baltimore, Maryland USA

Summer School on Human Language Technology **2010**
Center for Language and Speech Processing (CLSP)

FUNDING AND AWARDS Selected as team leader for the Jelinek Summer Workshop (JSALT), 2018
Fully-funded six-week language technology research event for teams of 8–12.
For: General-Purpose Sentence Representation Learning

Samsung DMC Research Center Agreement (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 seed grant, \$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

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

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-HLT 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)*.

REFEREED
CONFERENCE
PRESENTATIONS

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.

INVITED
PRESENTATIONS

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

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 \leftrightarrow 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.

OTHER
PUBLICATIONS

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 Søgaard. 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.

INTERSHIPS

Google, Inc., Mountain View, California USA

Software Engineering Intern

Summers 2012, 2013, 2014, 2015

TEACHING EXPERIENCE

New York University, New York, New York USA

Instructor

Spring 2019

Natural Language Understanding and Computational Semantics (DS-GA 1012/LING-GA 1012)
Seminar in Semantics: TBA (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.

PROFESSIONAL SERVICE

2018–2019. Area Chair, International Conference on Learning Representations (ICLR).

2018–. Reviewer, International Conference on Computational Linguistics (COLING).

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

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

2017-. Reviewer, Conference on Neural Information Processing Systems (NIPS).

2017-. Reviewer, International Conference on Machine Learning (ICML).

2016-. Reviewer, Annual Meeting of the Association for Computational Linguistics (ACL).

2015-. Reviewer, International Conference on Learning Representations (ICLR).

2015-. Reviewer, Conference on Empirical Methods in Natural Language Processing (EMNLP).

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

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

2018. Co-organizer and reviewer, 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).

2017-2018. Faculty Advisor to the Student Research Workshop, Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).

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

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

2017. Reviewer, AAAI Conference on Artificial Intelligence.

2017. Reviewer, Journal of Artificial Intelligence Research (JAIR).

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

2016. Reviewer, The First Workshop on Evaluating Vector-Space Representations for NLP (RepEval, hosted at ACL).

2015. Reviewer, Journal of Natural Language Engineering.

2014-2015. Reviewer, Semantics and Linguistic Theory (SALT) 25.

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

2017-. Application Reader, Moore-Sloan Data Science Fellows program, New York University Center for Data Science.

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

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

UNIVERSITY AND
DEPARTMENT
SERVICE

2016–. Faculty advisor, Data Future Lab incubator, New York University.

2016–. Member, PhD Admissions Committee, New York University Center for Data Science.

2016–. Member, Graduate Curriculum Committee, New York University Center for Data Science.

2017. Chair, Machine Learning Faculty Search Committee, Department of Computer Science and Center for Data Science, New York University.

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.

DISSERTATION
COMMITTEES

2017–. Rodrigo Nogueira (PhD, Computer Science)

2017–. Lisheng Fu (PhD, Computer Science)

2017–2018. Adina Williams (PhD, Linguistics)

2017–2018. Xiang Zhang (PhD, Computer Science)

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

2018. Elman Mansimov (PhD, Computer Science)

2018. Sébastien Jean (PhD, Computer Science, chair)

QUALIFYING PAPER
COMMITTEES

2018. Sheng-Fu Wang (PhD, Linguistics, co-chair with Maria Gouskova)

2017. WooJin Chung (PhD, Linguistics, chair)

OTHER ADVISING

2018–. Katharina Kann (Postdoc, Data Science, co-supervisor with Kyunghyun Cho)

2017–. Anhad Mohananey (MS, Computer Science, research supervisor)

2018–. Amanpreet Singh (MS, Computer Science, research supervisor)

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

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

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

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

2016-. Nikita Nangia (PhD, transferred from MS program in 2018, Data Science, advisor)

2018. Yun Chen (visiting PhD student, University of Hong Kong, co-host with Kyunghyun Cho)

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)

2017. Haoyue Shi (visiting BS student, Computer Science, Peking University, host, now TTI-Chicago PhD student)

2017. Xiaonan Zhao (MS, Computer Science, research supervisor, now at Amazon)

2016–2017. Andrew Drozdov (MS, Computer Science, research supervisor, now U. Mass. Amherst PhD student)