

Jennifer Hu

Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
43 Vassar St, Cambridge, MA 02139

Email: jennhu@mit.edu
Web: jennhu.github.io

EDUCATION	Massachusetts Institute of Technology (MIT) , Cambridge, MA, USA	
	Ph.D. in Cognitive Science Thesis: “Neural language models and human linguistic knowledge” Committee: Joshua Tenenbaum (chair), Roger Levy (advisor), Christopher Potts, Evelina Fedorenko	2018 – 2023
	Harvard University , Cambridge, MA, USA	
	B.A. in Mathematics and Linguistics (Magna cum laude with Highest Honors) Secondary Field in Germanic Languages & Literatures	2014 – 2018
RESEARCH EXPERIENCE	MIT Department of Brain and Cognitive Sciences	
	Graduate Researcher Advisor: Roger Levy	2018 – 2023
	Allen Institute for Artificial Intelligence	
	Summer Research Intern, Mosaic Team Advisor: Prithviraj Ammanabrolu	2022
	Stanford University Center for the Study of Language and Information	
	Summer Research Intern Advisor: Christopher Potts	2017
	Harvard University Department of Computer Science	
	Program for Research in Science and Engineering Fellow Advisor: Stuart Shieber	2016
PUBLICATIONS	MANUSCRIPTS	
	<ol style="list-style-type: none">[1] Jennifer Hu and Roger Levy. <i>Prompt-based methods may underestimate large language models’ linguistic generalizations</i>. 2023.[2] Daniel Fried, Nicholas Tomlin, Jennifer Hu, Roma Patel, and Aida Nematzadeh. <i>Pragmatics in grounded language learning: Phenomena, tasks, and modeling approaches</i>. 2022.	
	BOOK CHAPTERS	
	<ol style="list-style-type: none">[1] Ethan Wilcox, Jon Gauthier, Jennifer Hu, Peng Qian, and Roger Levy. “Learning syntactic structures from string input”. <i>Algebraic Structures in Natural Language</i>. Taylor & Francis, 2023.	
	JOURNAL ARTICLES	
	<ol style="list-style-type: none">[1] Jennifer Hu, Roger Levy, Judith Degen, and Sebastian Schuster. “Expectations over unspoken alternatives predict pragmatic inferences”. <i>Transactions of the Association for Computational Linguistics</i> (2023).[2] Jennifer Hu, Hannah Small, Hope Kean, Atsushi Takahashi, Leo Zelekman, Daniel Kleinman, Elizabeth Ryan, Alfonso Nieto-Castañón, Victor Ferreira, and Evelina Fedorenko. “Precision fMRI reveals that the language-selective network supports both phrase-structure building and lexical access during language production”. <i>Cerebral Cortex</i> (2022).	
	CONFERENCE PAPERS	
	<ol style="list-style-type: none">[1] Jennifer Hu, Sammy Floyd, Olessia Jouravlev, Evelina Fedorenko, and Edward Gibson. “A fine-grained comparison of pragmatic language understanding in humans and language models”. <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i>. 2023.	

- [2] Pei Zhou, Andrew Zhu, Jennifer Hu, Jay Pujara, Xiang Ren, Chris Callison-Burch, Yejin Choi, and Prithviraj Ammanabrolu. “An AI Dungeon Master’s Guide: Learning to Converse and Guide with Intents and Theory-of-Mind in Dungeons and Dragons”. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2023.
- [3] Irene Zhou, Jennifer Hu, Roger Levy, and Noga Zaslavsky. “Teasing apart models of pragmatics using optimal reference game design”. *Proceedings of the Cognitive Science Society*. 2022.
- [4] Jennifer Hu, Noga Zaslavsky, and Roger Levy. “Competition from novel features drives scalar inferences in reference games”. *Proceedings of the Cognitive Science Society*. 2021.
- [5] Yiwen Wang, Jennifer Hu, Roger Levy, and Peng Qian. “Controlled evaluation of grammatical knowledge in Mandarin Chinese language models”. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 2021.
- [6] Jon Gauthier, Jennifer Hu, Ethan Wilcox, Peng Qian, and Roger Levy. “SyntaxGym: An online platform for targeted evaluation of language models”. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. 2020, pp. 70–76.
- [7] Jennifer Hu, Sherry Yong Chen, and Roger Levy. “A closer look at the performance of neural language models on reflexive anaphor licensing”. *Proceedings of the Society for Computation in Linguistics*. Vol. 3. 2020, pp. 382–392.
- [8] Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. “A systematic assessment of syntactic generalization in neural language models”. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 2020, pp. 1725–1744.
- [9] Ethan Wilcox, Jon Gauthier, Jennifer Hu, Peng Qian, and Roger Levy. “On the predictive power of neural language models for human real-time comprehension behavior”. *Proceedings of the Cognitive Science Society*. 2020.
- [10] Jennifer Hu, James Traer, and Josh H. McDermott. “Separating object resonance and room reverberation in impact sounds”. *Proceedings of the Cognitive Science Society*. 2019.
- [11] Will Monroe, Jennifer Hu, Andrew Jong, and Christopher Potts. “Generating bilingual pragmatic color references”. *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. 2018, pp. 2155–2165.

WORKSHOP PAPERS

- [1] Jennifer Hu, Roger Levy, and Sebastian Schuster. *Predicting scalar diversity with context-driven uncertainty over alternatives*. ACL Workshop on Cognitive Modeling and Computational Linguistics. 2022.
- [2] Jennifer Hu, Roger Levy, and Noga Zaslavsky. *Scalable pragmatic communication via self-supervision*. ICML Workshop on Self-Supervised Learning for Reasoning and Perception. 2021.

EXTENDED ABSTRACTS

- [1] Jennifer Hu, Roger Levy, and Sebastian Schuster. *Predicting scalar diversity with context-driven expectations*. *Proceedings of the Experimental Pragmatics Conference (XPRAG)*. 2022.
- [2] Yiwen Wang, Jennifer Hu, Roger Levy, and Peng Qian. *Facilitative effect induced by classifier-noun mismatch in Mandarin Chinese*. *The 35th Annual Conference on Human Sentence Processing*. 2022.
- [3] Noga Zaslavsky, Jennifer Hu, and Roger Levy. *A Rate–Distortion view of human pragmatic reasoning*. *Proceedings of the Society for Computation in Linguistics*. 2021.
- [4] Irene Zhou, Jennifer Hu, Roger Levy, and Noga Zaslavsky. *Empirical support for a Rate–Distortion account of pragmatic reasoning*. *Proceedings of the Cognitive Science Society*. Member abstract. 2021.
- [5] Jennifer Hu, Hannah Small, Hope Kean, Atsushi Takahashi, Leo Zekelman, Daniel Kleinman, Elizabeth Ryan, Victor Ferreira, and Evelina Fedorenko. *Distributed and overlapping neural mechanisms for lexical access and syntactic encoding during language production*. *Proceedings of the Society for the Neurobiology of Language*. 2020.
- [6] Ethan Wilcox, Jon Gauthier, Jennifer Hu, Peng Qian, and Roger Levy. *Benchmarking neural networks as models of human language processing*. *Proceedings of the 26th Architectures and Mechanisms for Language Processing Conference*. 2020.

- [7] Ethan Wilcox, Jon Gauthier, Peng Qian, Jennifer Hu, and Roger Levy. *Evaluating the effect of model inductive bias and training data in predicting human reading times. Proceedings of the 33rd Annual CUNY Human Sentence Processing Conference*. 2020.
- [8] Noga Zaslavsky, Jennifer Hu, and Roger Levy. *Emergence of pragmatic reasoning from least-effort optimization. Proceedings of Evolution of Language International Conferences*. 2020.
- [9] Jennifer Hu. *A graph-theoretic approach to comparing typologies in Parallel OT and Harmonic Serialism. Proceedings of the 92nd Annual Meeting of the Linguistic Society of America*. 2018.

AWARDS	National Science Foundation Doctoral Dissertation Research Improvement Grant	2021	
	Computationally-Enabled Integrative Neuroscience Training Program	2019	
	National Science Foundation Graduate Research Fellowship	2019	
	Thomas T. Hoopes Prize	2018	
	Friends of Harvard Mathematics Prize	2018	
	Harvard College Research Program Grant	2017	
	Robert Fletcher Rogers Prize	2017	
	Detur Book Prize	2015	
	John Harvard Scholarship	2015	
TEACHING	TEACHING ASSISTANT		
	Language in the Mind and Brain (9.S52), MIT	2021	
	Computational Psycholinguistics (9.19/9.190), MIT	2020	
	Paradoxes and Infinities (PDOX), Johns Hopkins University Center for Talented Youth	2018	
	Linear Algebra and Real Analysis II (MATH 23B), Harvard	2016	
	Linear Algebra and Real Analysis I (MATH 23A), Harvard	2015	
	Vectors: A Tool for Teaching Algebra, Geometry, and Trigonometry (MATH S-323), Harvard	2015	
	GUEST LECTURER		
	“What do language models know about meaning?” <i>The Science of Intelligence (9.58), MIT</i>	2022	
	“Language understanding in minds and machines” <i>Language, Structure, and Cognition (LING 83), Harvard</i>	2020	
	“Testing synchronous tree-adjointing grammar analyses of linguistic phenomena” <i>Topics in Computational Linguistics (LING 98A), Harvard</i>	2016	
	INVITED TALKS	“Neural language models and human linguistic knowledge” <i>International Interdisciplinary Computational Cognitive Science Summer School</i>	2023
		“Benchmarking neural networks as models of human language processing” <i>Advancing Cognitive Science and AI with Cognitive-AI Benchmarking (CogSci Workshop)</i>	2023
		“Investigating ad-hoc scalar implicatures” <i>Department of Linguistics, University of Tübingen</i>	2022
		“Competition from novel features drives scalar inferences in reference games” <i>Language and Cognition Reading Group, Harvard</i>	2021
“Benchmarking neural networks as models of human language processing” <i>DeepMind</i>		2020	
SERVICE		ICML Workshop on Theory of Mind in Communicating Agents, Co-Organizer	2023
	NeurIPS Meaning in Context Workshop, Co-Organizer	2021	
	MIT School of Science Graduate Student Council, Member	2020 – 2021	
	MIT Women’s Advisory Group, Committee member	2019 – 2021	
	Graduate Women at MIT, Co-Chair	2019 – 2021	
	Harvard University Mather House, Non-Resident Tutor	2018 – 2019	
	Reviewer for ACL; EMNLP; CoNLL; CogSci; Language, Cognition and Neuroscience; Linguistics & Philosophy; Open Mind		

MENTORING	Graduate mentor for MIT-Harvard Women in AI	2021 – present
	Irene Zhou, MIT Undergraduate Research Opportunities Program	2020 – 2022
	Eric Hong, MIT Undergraduate Research Opportunities Program	2019

SKILLS

NATURAL LANGUAGES

English (native), German, Mandarin, Latin

PROGRAMMING LANGUAGES

Python, R, MATLAB, HTML/JavaScript