

Allyson Ettinger

Visiting Research Scientist, Allen Institute for AI

Contact

Email: allysone@allenai.org

Website: <https://aetting.github.io/>

Appointments

2023–present	Visiting Research Scientist, Allen Institute for AI
2022–2023	Assistant Professor, Dept. of Computer Science, The University of Chicago
2019–2023	Assistant Professor, Dept. of Linguistics, The University of Chicago
2019–2023	Courtesy Faculty, Toyota Technological Institute at Chicago
2018–2019	Research Assistant Professor, Toyota Technological Institute at Chicago

Education

University of Maryland, College Park

PhD - Linguistics, 2013–2018

With Certificate in Neuroscience and Cognitive Science

Dissertation: *Relating lexical and syntactic processes in language: Bridging research in humans and machines.*

Brandeis University

BA, Linguistics/Psychology, Summa Cum Laude, May 2010

Relevant work experience

2011–2013	Lab manager/research assistant, KIT/NYU MEG Research Lab (conducting research in cognitive neuroscience of language)
-----------	--

Publications

- | | |
|------|--|
| 2023 | West, P., Lu, X., Dziri, N., Brahman, F., Li, L., ... Ettinger, A. , Choi, Y. The Generative AI Paradox: “What It Can Create, It May Not Understand”. <i>Proceedings of the Eleventh International Conference on Learning Representations (ICLR 2023)</i> . |
| 2023 | Yang, C., Ettinger, A. Can You Follow Me? Testing Situational Understanding for Chat-GPT. <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023)</i> . |
| 2023 | Ettinger, A. , Hwang, J.D., Pyatkin, V., Bhagavatula, C., Choi, Y. “You Are An Expert Linguistic Annotator”: Limits of LLMs as Analyzers of Abstract Meaning Representation. <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> . |
| 2023 | Dziri, N., Lu, X., Sclar, M. Li, X.L., Jiang, L., Lin, B.Y, West, P., ... Ettinger, A. , Harchaoui, Z., Choi, Y. Faith and Fate: Limits of Transformers on Compositionality. <i>Proceedings of the Thirty-seventh Annual Conference on Neural Information Processing Systems (NeurIPS 2023)</i> . |

- 2023 Li, J., Yu, L., **Ettinger, A.** Counterfactual reasoning: Testing language models’ understanding of hypothetical scenarios. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023)*.
- 2023 Misra, K., Taylor Rayz, J., **Ettinger, A.** COMPS: Conceptual Minimal Pair Sentences for testing Robust Property Knowledge and its Inheritance in Pre-trained Language Models. *Proceedings of The 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. (PDF). **Recipient of Best Paper Award.**
- 2023 Li, J., **Ettinger, A.** Heuristic interpretation as rational inference: A computational model of the N400 and P600 in language processing. *Cognition* 233. (PDF)
- 2022 Li, J., Yu, L., **Ettinger, A.** Counterfactual reasoning: Do language models need world knowledge for causal understanding? *Proceedings of the Workshop on neuro Causal and Symbolic AI (nCSI) at NeurIPS*. (PDF)
- 2022 Kim, S., Yu, L., **Ettinger, A.** “No, they did not”: Dialogue response dynamics in pre-trained language models. *Proceedings of the 29th International Conference on Computational Linguistics (COLING)*. (PDF)
- 2022 Misra, K., Taylor Rayz, J., **Ettinger, A.** A Property Induction Framework for Neural Language Models. *Proceedings of the 44th Annual Conference of the Cognitive Science Society*. (PDF)
- 2021 Pandia, L., **Ettinger, A.** Sorting through the noise: Testing robustness of information processing in pre-trained language models. *Proceedings of The 2021 Conference on Empirical Methods in Natural Language Processing*. (PDF)
- 2021 Pandia, L., Cong, Y., **Ettinger, A.** Pragmatic competence of pre-trained language models through the lens of discourse connectives. *Proceedings of the 2021 SIGNLL Conference on Computational Natural Language Learning*. (PDF)
- 2021 Wu, Q., **Ettinger, A.** Variation and generality in encoding of syntactic anomaly information in sentence embeddings. *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*. (PDF)
- 2021 Yu, L., **Ettinger, A.** On the Interplay Between Fine-tuning and Composition in Transformers. *Findings of ACL: ACL 2021*. (PDF)
- 2021 Misra, K., **Ettinger, A.**, Taylor Rayz, J. Do language models learn typicality judgments from text? *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*. (PDF)
- 2020 Yu, L., **Ettinger, A.** Assessing Phrasal Representation and Composition in Transformers. *Proceedings of The 2020 Conference on Empirical Methods in Natural Language Processing*. (PDF)
- 2020 Misra, K., **Ettinger, A.**, Taylor Rayz, J. Exploring BERT’s Sensitivity to Lexical Cues using Tests from Semantic Priming. *Findings of ACL: EMNLP 2020*. (PDF)
- 2020 Toshniwal, S., Wiseman, S., **Ettinger, A.**, Gimpel, K., Livescu, K. Learning to Ignore: Long Document Coreference with Bounded Memory Neural Networks. *Proceedings of The 2020 Conference on Empirical Methods in Natural Language Processing*. (PDF)
- 2020 **Ettinger, A.** What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics*. (PDF)
- 2020 Klafka, J., & **Ettinger, A.** Spying on your neighbors: Fine-grained probing of contextual embeddings for information about surrounding words. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. (PDF)
- 2020 Toshniwal, S., **Ettinger, A.**, Gimpel, K., Livescu, K. PeTra: A Sparsely Supervised Memory Model for People Tracking. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. (PDF)

- 2018 **Ettinger, A.**, Elgohary, A., Phillips, C., & Resnik, P. Assessing Composition in Sentence Vector Representations. *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*. (PDF)
- 2017 **Ettinger, A.**, Rao, S., Daumé III, H., & Bender, E. M. Towards Linguistically Generalizable NLP Systems: A Workshop and Shared Task. *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*. (PDF)
- 2016 **Ettinger, A.**, Elgohary, A., & Resnik, P. Probing for semantic evidence of composition by means of simple classification tasks. *Proceedings of the First Workshop on Evaluating Vector Space Representations for NLP, ACL 2016*. **Recipient of Best Proposal Award**. (PDF)
- 2016 **Ettinger, A.**, Resnik, P., & Carpuat, M. Retrofitting sense-specific word vectors using parallel text. *Proceedings of NAACL-HLT 2016*. (PDF)
- 2016 **Ettinger, A.**, Feldman, N.H., Resnik, P., & Phillips, C. Modeling N400 amplitude using vector space models of word representation. *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. (PDF)
- 2016 **Ettinger, A.** & Linzen, T. Evaluating vector space models using human semantic priming results. *Proceedings of the First Workshop on Evaluating Vector Space Representations for NLP, ACL 2016*. (PDF)
- 2015 Rao, S., **Ettinger, A.**, Daumé III, H., & Resnik, P. (2015). Dialogue focus tracking for zero pronoun resolution. *Proceedings of NAACL-HLT 2015*, 495-503. (PDF)
- 2015 **Ettinger, A.** & Malamud, S. Mandarin utterance-final particle ba in the conversational scoreboard. *Proceedings of Sinn und Bedeutung 2014*. (PDF)
- 2014 **Ettinger, A.**, Linzen, T., & Marantz, A. The role of morphology in phoneme prediction: Evidence from MEG. *Brain and Language 129*, 14-23. (PDF)

Grants

- 2019–2021 NSF EAGER grant: *World Modeling for Natural Language Understanding*.
(with Kevin Gimpel, Toyota Technological Institute at Chicago)
- 2014–2017 NSF Graduate Research Fellowship

Presentations

- 2023 Invited talk, University of Sheffield NLP Group, University of Sheffield, Sheffield, England.
- 2023 Invited talk, Transdisciplinary Institute in Applied Data Science (TRIADS) seminar series, Washington University in St. Louis, St. Louis, MO.
- 2022 Invited keynote talk, SIGNLL Conference on Computational Natural Language Learning (CoNLL), Abu Dhabi, UAE.
- 2022 Invited talk, ILFC monthly online seminar organized by GdR LIFT, French National Centre for Scientific Research.
- 2022 Invited talk, University of Texas Austin workshop series in formal and computational semantics, Austin, TX.
- 2022 Invited talk, McGill Linguistics Department colloquium, Montreal, Quebec (virtual).
- 2022 Invited talk, University of California, Irvine Summer School on Computational Cognitive Modeling for Language, Irvine, CA.
- 2022 Invited keynote talk, 11th Joint Conference on Lexical and Computational Semantics (*SEM), NAACL 2022, Seattle, WA.

- 2022 Invited speaker and panelist, “The Challenge of Compositionality for AI” workshop, (virtual).
- 2022 Invited talk, Microsoft Cognitive Services Research Group Distinguished Talk Series, Seattle, WA (virtual).
- 2022 Invited keynote talk, Third Workshop on Knowledge Extraction and Integration for Deep Learning Architectures (DeeLIO), ACL 2022, Dublin, Ireland.
- 2022 “*Understanding*” and prediction: *Controlled examinations of meaning sensitivity in pre-trained models*. Invited talk, University of Maryland CLIP Lab Colloquium, College Park, MD.
- 2022 “*Understanding*” and prediction: *Controlled examinations of meaning sensitivity in pre-trained models*. Invited talk, Notre Dame NLP Seminar, Notre Dame, IN (virtual).
- 2022 “*Understanding*” and prediction: *Controlled examinations of meaning sensitivity in pre-trained models*. Invited talk, University of Pennsylvania CLunch, Philadelphia, PA.
- 2022 *Foundations of ambiguity resolution: Assessing meaning extraction in deep learning models*. Invited talk, l’AMORE Mini-workshop on Linguistic Ambiguity and Deep Learning, Barcelona, Spain.
- 2022 *Diagnosing meaning extraction in pre-trained language models*. Invited talk, Carnegie Mellon University brAIIn Seminar, Pittsburgh, PA (virtual).
- 2022 “*Understanding*” and prediction: *Controlled examinations of meaning sensitivity in pre-trained models*. Invited talk, Stanford NLP Seminar, Stanford, CA (virtual).
- 2021 “*Understanding*” and prediction: *Controlled examinations of meaning extraction in natural language processing models*. Invited talk, Ohio State University Linguistics Department Colloquium, Columbus, OH (virtual).
- 2021 “*Understanding*” and prediction: *Controlled examinations of meaning sensitivity in pre-trained models*. Invited talk, van Schijndel research group, Cornell University, Ithaca, NY (virtual).
- 2021 “*Understanding*” and prediction: *Controlled examinations of meaning sensitivity in pre-trained NLP models*. Invited talk, University of Chicago MACSS Computational Social Science Workshop, Chicago, IL (virtual).
- 2021 Invited panelist, *How can findings about the brain improve AI systems?*, panel at Brain2AI workshop, International Conference on Learning Representations (virtual).
- 2021 “*Understanding*” and prediction: *Controlled examinations of meaning sensitivity in pre-trained models*. Invited talk, New York University NLP/Text-as-Data speaker series, New York, NY (virtual).
- 2021 *Pre-trained language models evaluated relative to human cognition*. Invited talk, Dongguk University English Literature and Language Dept. Colloquium, Seoul, South Korea (virtual).
- 2021 *A noisy channel model of N400 and P600 effects in sentence processing*. (with Jiaxuan Li) 34th Annual CUNY Conference on Human Sentence Processing.
- 2020 *Word and phrase representations in machines and in humans: comprehension, context, and composition*. Invited talk, Massachusetts Institute of Technology CompLang discussion group, Cambridge, MA (virtual).
- 2020 *Holding AI to a human standard: Applying cognitive science to evaluation of NLP models*. Invited talk, Massachusetts Institute of Technology Computational Psycholinguistics Lab, Cambridge, MA (virtual).
- 2020 “*Understanding*” and prediction in language: *perspectives from AI and human cognition*. Invited talk, Georgia Institute of Technology Workshop on Language, Technology, and Society, Atlanta, GA (virtual).
- 2020 “*Understanding*” and prediction in language: *perspectives from AI and human cognition*. Invited talk, Northwestern University Department of Linguistics Colloquium, Evanston, IL.

- 2020 (postponed due to COVID-19) *“Understanding” and prediction in language: perspectives from AI and human cognition*. Invited talk, University of Wisconsin-Milwaukee Department of Linguistics Colloquium Series, Milwaukee, WI.
- 2019 *NLP representations from a perspective of human cognition*. Invited talk, Johns Hopkins University Center for Language and Speech Processing Lecture Series, Baltimore, MD.
- 2019 Invited panelist, *A linguist, an NLP engineer, and a psycholinguist walk into a bar?*, interdisciplinary panel at RepEval Workshop, Minneapolis, MN.
- 2019 *Challenges in NLP evaluation (and how linguists can help)*. Invited panel talk, Society for Computation in Linguistics Annual Meeting, New York, NY.
- 2019 *Vector space models*. Invited tutorial, Society for Computation in Linguistics Annual Meeting, New York, NY.
- 2018 *Bridging language research in humans and machines: toward a productive symbiosis*. Invited talk, CUNY Graduate Center, New York, NY.
- 2018 *Opening the black box: leveraging insights from linguistics and cognitive neuroscience to assess the linguistic capacities of NLP systems*. Invited presentation, University of Chicago Department of Linguistics, Chicago, IL.
- 2018 *Bridging language research in humans and machines: toward a productive symbiosis*. Invited talk, University of Chicago Department of Linguistics, Chicago, IL.
- 2018 *Bridging NLP and brain science to improve natural language understanding*. Invited talk, Toyota Technological Institute, Chicago, IL.
- 2017 *Building and evaluating meaning representations for NLP*. Invited talk, Tohoku University, Sendai, Japan.
- 2017 *Drawing cognitive insights from NLP: Assessing the usefulness of engineering models in modeling human language processing*. Invited talk, Tohoku University, Sendai, Japan.
- 2013 *Mandarin utterance-final particle ba in the conversational scoreboard*. [with Sophia Malamud]. Presented at 19th International Congress of Linguists, Geneva, Switzerland, and Linguistic Society of America Annual Meeting, Boston, MA.
- 2012 *Mandarin utterance-final particle ba and conversational goals*. [with Sophia Malamud]. Presented at NYU Semantics Discussion Group Meeting, New York, NY.
- 2010 *Linguistic construction of gender and generations in Hmong-American communities*. [with Mai Youa Moua, James Stanford]. Presented at Linguistic Society of America Annual Meeting, Baltimore, MD.

Thesis Advising

- PhD in Computer Science (in progress): Chih-chan Tien, Lalchand Pandia, Chenghao Yang.
- PhD in Linguistics (in progress): Sanghee Kim, Daniel Lam.
- PhD in Natural Language Understanding (in progress): Kanishka Misra, Purdue University.
- Xiaomeng Ma, CUNY, PhD in Linguistics. Thesis title: *Evaluating Neural Networks as Cognitive Models for Learning Quasi-regularities in Languages*
- Karim Lasri, Ecole normale supérieure, PhD in Language Science. Thesis title: *Linguistic Generalization in Transformer-Based Neural Language Models*
- Laura Aina, Universitat Pompeu Fabra, PhD in Translation and Language Sciences. Thesis title: *A Deep Learning Perspective on Linguistic Ambiguity*.

- Lang Yu, PhD in Computer Science. Thesis title: *Analyzing and Improving Compositionality in Neural Language Models*.
- Daniel Edmiston, PhD in Linguistics. Thesis title: *Identifying Linguistic Structures in Continuous Spaces*.
- Jiaxuan Li, Master’s Program in Computational Social Science. Thesis title: *A Noisy Channel Model of N400 and P600 Effects in Sentence Comprehension*.
- Wanitchaya (Mint) Poonpatanapricha, Master’s Program in Computational Social Science. Thesis title: *Use of emojis as markers of politeness*.
- Ethan Liu, Master’s Program in Statistics. Thesis title: *Causal Relations and Feature Extraction in Contextualized Sentence Representation*.
- Zhou Xing, Master’s Program in Computational Social Science. Thesis title: *Visually Grounded Multilingual Syntax Acquisition*.
- Alexander (Sasha) Tyan, Master’s Program in Computational Social Science. Thesis title: *Using Topic Modelling to Detect Bias Between News Sources*.
- Michael Hanna, Bachelor’s in Linguistics and Computer Science. Thesis title: *Evaluating the Interpretability of Variational Autoencoders’ Latent Space Representations of Sentences*
- Hanson Lu, Bachelor’s in Linguistics and Computer Science. Thesis title: *Applying the Rational Speech Acts Model to Coreference Resolution*

Service

- **Organizer**, “The Big Picture: Crafting a Research Narrative”, Workshop at EMNLP 2023
- **Chair**, Society for Computation in Linguistics, 2021-2022
- **Program committee chair**, Society for Computation in Linguistics, 2020-2021
- **Action editor**, ACL Rolling Review, 2021-2022
- **Senior area chair**, North American Association for Computational Linguistics Annual Meeting, 2022
- **Area chair**, Association for Computational Linguistics Annual Meeting, 2020–2023; Conference on Empirical Methods in Natural Language Processing, 2022; North American Association for Computational Linguistics Annual Meeting, 2021; SIGNLL Conference on Computational Natural Language Learning, 2021
- **Organizer**, Workshop on Building Linguistically Generalizable NLP Systems, EMNLP 2017
- **Organizer**, Student Research Workshop, ACL 2017
- **Chair**, Research Training Committee, UMD Language Science Center (2016–2017)
- **Chair**, Annual Winter Storm graduate training workshop, UMD Language Science Center (2014–2016)

Reviewing

- ACL, EMNLP, ICLR, ICML, NAACL, NeurIPS, EACL, COLING
- Workshop on Analyzing and Interpreting Neural Networks for NLP (BlackBoxNLP)
- Society for Computation in Linguistics (SCiL)
- Workshop on Evaluating Vector Space Representations for NLP (RepEval)
- *Cognition* (journal)
- *Cognitive Science* (journal)
- *Transactions of the Association for Computational Linguistics* (journal)
- *Computational Linguistics* (journal)
- *Psychological Review* (journal)
- *Journal of Natural Language Engineering* (journal)
- Swiss National Science Foundation (grant review)

Natural Languages

- **English:** native
- **Mandarin Chinese:** fluent
- **Spanish:** proficient
- **Korean, French, German, Modern Standard Arabic:** reading/communicative knowledge

Honors and Awards

- Best Paper Award, European Chapter of the Association for Computational Linguistics (EACL) Annual Meeting, 2023
- Best Proposal Award, Workshop on Evaluating Vector Space Representations for NLP, Association for Computational Linguistics (ACL) Annual Meeting 2016
- National Science Foundation Graduate Research Fellowship recipient, 2014
- Member Phi Beta Kappa Academic Honor Society, Brandeis University Chapter
- Justice Louis D. Brandeis full-tuition scholarship, Brandeis University, 2006-2010
- National Presidential Scholar, U.S. Presidential Scholars Program, 2006

Teaching

- **Computational Linguistics, undergrad / grad (2020-2023)**
 - Interdisciplinary introductory course in computational cognitive modeling and natural language processing (cross-listed in linguistics and computer science). Single-quarter course taught separately for undergraduate and graduate students.
- **Computational Linguistics I (2019-2020)**
 - Interdisciplinary, mixed undergraduate/graduate level course in computational cognitive modeling and natural language processing (cross-listed in linguistics and computer science)
- **Computational Linguistics II (2019-2020)**
 - Advanced continuation of Computational Linguistics I
- **Seminar in Computational Linguistics**
 - Graduate seminar course focusing examining approaches to meaning in computational linguistics—particularly efforts to capture meaning and “understanding” in artificial intelligence
- **Language and the Human (2019-present)**
 - Humanities course in the University of Chicago Undergraduate Core Curriculum, structured around discussion and writing on topics involving Language and Identity