

# Nelson F. Liu

---

CONTACT INFORMATION	<i>Email:</i> nfliu@cs.stanford.edu <a href="http://www.cs.stanford.edu/~nfliu">www.cs.stanford.edu/~nfliu</a>	
EDUCATION	<b>Stanford University</b> , Stanford, California Ph.D. Candidate, Computer Science Department Advisor: Percy Liang	2019 - Present
	<b>University of Washington</b> , Seattle, Washington B.S. with Distinction in Computer Science, B.A. in Linguistics Advisor: Noah A. Smith	2015 - 2019
RESEARCH AND INDUSTRY EXPERIENCE	<b>Samaya AI</b> <i>Intern</i>	Summer 2023
	<b>Google Research</b> <i>Student Researcher</i> , with Kenton Lee and Kristina Toutanova	Summer 2022
	<b>University of Washington, Computer Science and Engineering</b> <i>Undergraduate Research Assistant</i> , with Noah A. Smith	2015 - 2019
	<b>Allen Institute for Artificial Intelligence (AI2)</b> <i>Research Intern</i> , with Matt Gardner and Noah A. Smith	2018 - 2019
	<b>USC Information Sciences Institute, Natural Language Group</b> <i>Research Intern</i> , with Kevin Knight and Jonathan May	Summer 2017
	<b>Allen Institute for Artificial Intelligence (AI2)</b> <i>Research Intern</i> , with Matt Gardner	Winter 2017
	<b>scikit-learn</b> <i>Google Summer of Code Developer</i>	Summer 2016
FELLOWSHIPS, AWARDS & HONORS	Best Paper Award, ACL 2022 Workshop on Representation Learning for NLP [7] NSF Graduate Research Fellowship (Natural Language Processing), 2019 Outstanding Undergraduate Honors Thesis Award, University of Washington CSE, 2019 Outstanding Computer Science Senior Award, University of Washington CSE, 2019 Finalist, CRA Outstanding Undergraduate Researcher Award, 2019 Best Paper Award, ACL 2018 Workshop on Representation Learning for NLP [15] Barry M. Goldwater Scholarship, 2018 Washington Research Foundation Fellowship, (2017, 2018) Phi Beta Kappa, 2017 Mary Gates Research Scholarship, 2016	
PUBLICATIONS	[1] <i>Lost in the Middle: How Language Models Use Long Contexts</i> <b>Nelson F. Liu</b> , Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. <i>Transactions of the Association for Computational Linguistics (TACL)</i> . July 2023. [2] <i>Anchor Prediction: Automatic Refinement of Internet Links</i> <b>Nelson F. Liu</b> , Kenton Lee, and Kristina Toutanova. arXiv:2305.14337. May 2023. [3] <i>Evaluating Verifiability in Generative Search Engines</i> <b>Nelson F. Liu</b> , Tianyi Zhang, and Percy Liang. <i>Findings of EMNLP</i> , 2023.	

- [4] *Are Sample-Efficient NLP Models More Robust?*  
**Nelson F. Liu**, Ananya Kumar, Percy Liang, and Robin Jia.  
*Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- [5] *Do Question Answering Modeling Improvements Hold Across Benchmarks?*  
**Nelson F. Liu**, Tony Lee, Robin Jia, and Percy Liang.  
*Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- [6] *Dyna-bAbI: unlocking bAbI’s potential with dynamic synthetic benchmarking*  
Ronen Tamari, Kyle Richardson, Aviad Sar-Shalom, Noam Kahlon, **Nelson F. Liu**, Reut Tsarfaty, and Dafna Shahaf.  
*Joint Conference on Lexical and Computational Semantics (\*SEM)*, 2022.
- [7] *Identifying the Limits of Cross-Domain Knowledge Transfer for Pretrained Models*  
Zhengxuan Wu, **Nelson F. Liu**, and Christopher Potts.  
*ACL Workshop on Representation Learning for NLP (RepL4NLP)*, 2022.  
**Best Paper Award**
- [8] *Making Heads and Tails of Models with Marginal Calibration for Sparse Tagsets*  
Michael Kranzlein, **Nelson F. Liu**, and Nathan Schneider.  
*Findings of EMNLP*, 2021.
- [9] *Lexical Semantic Recognition*  
**Nelson F. Liu**, Daniel Hershcovich, Michael Kranzlein, and Nathan Schneider.  
*ACL Workshop on Multiword Expressions (MWE)*, 2021.
- [10] *Evaluating NLP Models via Contrast Sets*  
Matt Gardner, Yoav Artzi, Victoria Basmova, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hanna Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, **Nelson F. Liu**, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou.  
*Findings of EMNLP*, 2020.
- [11] *Quoref: A Reading Comprehension Dataset with Questions Requiring Coreferential Reasoning*  
Pradeep Dasigi, **Nelson F. Liu**, Ana Marasović, Noah A. Smith and Matt Gardner.  
*Conference on Empirical Methods in Natural Language Processing & International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.
- [12] *Barack’s Wife Hillary: Using Knowledge Graphs for Fact-Aware Language Modeling*  
Robert L. Logan IV, **Nelson F. Liu**, Matthew E. Peters, Matt Gardner and Sameer Singh.  
*Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [13] *Linguistic Knowledge and Transferability of Contextual Representations*  
**Nelson F. Liu**, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith.  
*North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019.
- [14] *Inoculation by Fine-Tuning: A Method for Analyzing Challenge Datasets*  
**Nelson F. Liu**, Roy Schwartz, and Noah A. Smith.  
*North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019.
- [15] *LSTMs Exploit Linguistic Attributes of Data*  
**Nelson F. Liu**, Omer Levy, Roy Schwartz, Chenhao Tan and Noah A. Smith.  
*ACL Workshop on Representation Learning for NLP (RepL4NLP)*, 2018.  
**Best Paper Award.**
- [16] *AllenNLP: A Deep Semantic Natural Language Processing Platform*  
Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, **Nelson F. Liu**, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer.  
*ACL Workshop for Natural Language Processing Open Source Software (NLP-OSS)*, 2018.
- [17] *Augmenting Statistical Machine Translation with Subword Translation of Out-of-Vocabulary Words*  
**Nelson F. Liu**, Jonathan May, Michael Pust, and Kevin Knight.  
arXiv:1808.05700. August 2018.
- [18] *Discovering Phonesthemes with Sparse Regularization*  
**Nelson F. Liu**, Gina-Anne Levow, and Noah A. Smith.  
*NAACL Workshop on Subword and Character Level Models in NLP (SCLeM)*, 2018.

[19] *Crowdsourcing Multiple Choice Science Questions*  
 Johannes Welbl, **Nelson F. Liu**, and Matt Gardner.  
*EMNLP Workshop on Noisy User-generated Text (W-NUT)*, 2017.

PROFESSIONAL  
SERVICE

**Workshop Organization:**

- NAACL 2021 Student Research Workshop
- EMNLP 2020 Workshop for Natural Language Processing Open Source Software

**Program Committee Member:** (\* outstanding reviewer)

- Standing reviewer pool, ACL Rolling Review: 2021, 2022
- Standing reviewer pool, Transactions on Machine Learning Research (TMLR): 2022
- Annual Meeting of the Association for Computational Linguistics (ACL): 2019, 2020\*, 2021, 2023
- Conference on Empirical Methods in Natural Language Processing (EMNLP): 2019\*, 2020\*, 2021
- Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL): 2021
- Conference of the European Chapter of the Association for Computational Linguistics: 2021, 2023
- Annual Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (AACL): 2020, 2022
- International Conference on Computational Linguistics (COLING): 2020
- AAAI Conference on Artificial Intelligence (AAAI): 2020, 2021
- International Conference on Machine Learning (ICML): 2020, 2021
- International Conference on Learning Representations (ICLR): 2021, 2022, 2023
- Conference on Neural Information Processing Systems (NeurIPS): 2021, 2022, 2023
- Conference on Computational Natural Language Learning (CoNLL): 2019, 2020, 2021, 2022, 2023
- Joint Conference on Lexical and Computational Semantics (\*SEM): 2020
- ACL Student Research Workshop: 2019
- Workshop on Analyzing and Interpreting Neural Networks for NLP (BlackboxNLP): 2018, 2019, 2020, 2021, 2023
- NeurIPS Workshop on Distribution Shifts: Connecting Methods and Applications: 2021
- ICML Workshop on Principles of Distribution Shift: 2022
- Workshop on Methods for Optimizing and Evaluating Neural Language Generation: 2019

**Student Volunteer:**

- Conference on Empirical Methods in Natural Language Processing (EMNLP): 2019

DEPARTMENTAL  
SERVICE

Stanford NLP Undergraduate Research Committee: 2020, 2021  
 Stanford CS Undergraduate Mentor: 2020, 2021  
 Stanford AI Lab Undergraduate Mentor: 2019  
 Stanford AI Lab Blog Editor: 2019, 2020, 2021  
 Stanford CS TGIF Organizer: 2022

TEACHING  
EXPERIENCE

**CS224N: NLP with Deep Learning**, Stanford University Winter 2024  
*Head Teaching Assistant*

**CSE447: Natural Language Processing**, University of Washington Winter 2019  
*Teaching Assistant*

**CSE190B: CSE Direct Admit Seminar**, University of Washington Autumn 2018  
*Teaching Assistant*

**CSE481N: Natural Language Processing Capstone**, University of Washington Spring 2018  
*Teaching Assistant*

**CSE447: Natural Language Processing**, University of Washington Winter 2018  
*Teaching Assistant*

DATE COMPILED December 21, 2023