

# ABHILASHA RAVICHANDER

aravicha@mpi-sws.org | [lasharavichander.github.io](https://github.com/lasharavichander)

## RESEARCH INTERESTS

---

Artificial Intelligence, Large Language Models, Data-Centric Interpretability, AI Reliability, Transparency, and Creativity.

## EDUCATION

---

2018–2022	<b>Carnegie Mellon University</b> , School of Computer Science, Ph.D. in Language and Information Technologies, <b>Advisors : Eduard Hovy, Norman Sadeh</b>
2016–2018	<b>Carnegie Mellon University</b> , School of Computer Science M.Sc in Language Technologies. <b>QPA: 4.0</b>

## EMPLOYMENT

---

Nov 2025-Present	<b>Max Planck Institute for Software Systems</b> Tenure-track Faculty
2024-2025	<b>University of Washington</b> , Paul G. Allen School of Computer Science Postdoctoral Researcher <b>Advisor : Yejin Choi</b>
2023-2024	<b>Allen Institute for Artificial Intelligence</b> , Seattle, WA Young Investigator

## ACADEMIC AWARDS AND HONORS

---

- ACL 2025 **Outstanding Paper Award**
- NAACL 2025, Outstanding Paper Award Candidate
- TrustNLP 2025 Workshop **Best Paper Award**
- DAAD AiNet Fellow on Natural Language Processing, 2025
- **Rising Star in Generative AI**, 2024
- ACL 2024 **Best Resource Paper Award**
- ACL 2024 **Best Theme Paper Award**
- NAACL 2024, Outstanding Paper Award Candidate
- MASC-SLL 2024 **Best Paper Award**
- **Rising Star in EECS**, 2022
- SoCal NLP Symposium 2022 **Best Paper Award**
- **Rising Star in Data Science**, 2021
- **Outstanding reviewer**, ACL 2020
- **Outstanding reviewer**, EMNLP 2020
- COLING 2018 **Area Chair Favorite Paper Award**
- **\$100,000** grant from Amazon for the 2016 Alexa Prize proposal.

## PUBLICATIONS

---

1. Mohammad Aflah Khan, Mahsa Amani, Soumi Das, Bishwamitra Ghosh, Qinyuan Wu, Krishna P. Gummadi, Manish Gupta, **Abhilasha Ravichander**  
*In Agents We Trust, but Who Do Agents Trust? Latent Preferences Steer LLM Generations*  
The Fourteenth International Conference on Learning Representations (ICLR 2026). *[long paper]*
2. Keivan Rezaei\*, Mehrdad Saberi\*, Soheil Feizi†, **Abhilasha Ravichander**†  
*Revisiting the Past: Data Unlearning with Model State History*  
The Fourteenth International Conference on Learning Representations (ICLR 2026). *[long paper]*
3. Mohammad Aflah Khan, Krishna P. Gummadi, Manish Gupta, **Abhilasha Ravichander**  
*Fractional Rotation, Full Potential? Investigating Performance and Convergence of Partial RoPE*  
arxiv *[short paper]*

4. **Abhilasha Ravichander\***, Shrusti Ghela\*, David Wadden, Yejin Choi  
*HALoGEN: Fantastic LLM Hallucinations and Where To Find Them*  
2025 Annual Conference of the Association for Computational Linguistics (ACL 2025).  
*[ACL 2025 Outstanding Paper Award] [TrustNLP Workshop Best Paper Award] [long paper]*
5. Yiyu Sun, Yu Gai, Lijie Chen, **Abhilasha Ravichander**, Yejin Choi, Nouha Dziri, Dawn Song  
*Why and How LLMs Hallucinate: Connecting the Dots with Subsequence Associations*  
The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS 2025) *[long paper]*
6. Keivan Rezaei, Khyathi Chandu, Soheil Feizi, Yejin Choi, Faeze Brahman, **Abhilasha Ravichander**  
*RESTOR: Knowledge Recovery through Machine Unlearning*  
*Transactions on Machine Learning Research (TMLR 2025) [long paper]*
7. Benjamin Newman, **Abhilasha Ravichander**, Jaehun Jung, Rui Xin, Hamish Ivison, Yegor Kuznetsov, Pang Wei Koh, Yejin Choi  
*The Curious Case of Factuality Finetuning: Models' Internal Beliefs Can Improve Factuality*  
*Under review. [long paper]*
8. Wenting Zhao, Tanya Goyal, Yu Ying Chiu, Liwei Jiang, Benjamin Newman, **Abhilasha Ravichander**, Khyathi Chandu, Ronan Le Bras, Claire Cardie, Yuntian Deng, Yejin Choi  
*WildHallucinations: Evaluating Long-form Factuality in LLMs with Real-World Entity Queries*  
*Under review. [long paper]*
9. Skyler Hallinan, Jaehun Jung, Melanie Sclar, Ximing Lu, **Abhilasha Ravichander**, Sahana Ramnath, Yejin Choi, Sai Praneeth Karimireddy, Niloofar Miresghallah, Xiang Ren  
*The Surprising Effectiveness of Membership Inference with Simple N-Gram Coverage*  
Conference on Language Modeling, 2025 *[long paper]*
10. Alex Gill, **Abhilasha Ravichander\***, Ana Marasovic  
*What Has Been Lost with Synthetic Evaluation?*  
Findings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP Findings 2023). *[long paper]*
11. **Abhilasha Ravichander**, Jillian Fisher, Taylor Sorensen, Ximing Lu, Maria Antoniak, Bill Yuchen Lin, Niloofar Miresghallah, Chandra Bhagavatula, Yejin Choi  
*Information-Guided Identification of Training Data Imprint in (Proprietary) Large Language Models*  
2025 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2025).  
*[Nominated for Best Paper Award] [long paper]*
12. Nishant Balepur, Feng Gu, **Abhilasha Ravichander**, Shi Feng, Jordan Lee Boyd-Graber, Rachel Rudinger  
*Reverse Question Answering: Can an LLM Write a Question so Hard (or Bad) that it Can't Answer?*  
2025 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2025).  
*[long paper]*
13. Bill Yuchen Lin, Yuntian Deng, Khyathi Chandu, Faeze Brahman, **Abhilasha Ravichander**, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, Yejin Choi  
*WildBench: Benchmarking LLMs with Challenging Tasks from Real Users in the Wild*  
International Conference on Learning Representations (ICLR 2025).  
*[Spotlight] [long paper]*
14. Faeze Brahman+, Sachin Kumar+, **Abhilasha Ravichander\***, Vidhisha Balachandran\*, Pradeep Dasigi\*, Valentina Pyatkin\*, Sarah Wiegrefe\*, Nouha Dziri, Khyathi Chandu, Jack Hessel, Yulia Tsvetkov, Noah A. Smith, Yejin Choi, Hannaneh Hajishirzi  
*The Art of Saying No: Contextual Noncompliance in Language Models*  
NeurIPS 2024, *Datasets and Benchmarks [long paper]*
15. Nishant Balepur, **Abhilasha Ravichander**, Rachel Rudinger  
*Artifacts or Abduction: How Do LLMs Answer Multiple-Choice Questions Without the Question?*  
62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024).  
*[MASC-SLL 2024 Best Paper Award] [long paper]*

16. Groeneveld et al.,  
*OLMo: Accelerating the Science of Language Models*  
62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024).  
*[ACL Best Theme Paper Award] [GeekWire Innovation of the Year Award] [long paper]*
17. Soldaini et al.,  
*Dolma: an Open Corpus of Three Trillion Tokens for Language Model Pretraining Research*  
62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024).  
*[ACL Best Resource Paper Award] [long paper]*
18. Da Yin, Faeze Brahman, **Abhilasha Ravichander**, Khyathi Chandu, Kai-Wei Chang, Yejin Choi, Bill Yuchen Lin  
*Lumos: Learning Agents with Unified Data, Modular Design, and Open-Source LLMs*  
62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024). *[long paper]*
19. Yufei Tian, **Abhilasha Ravichander**, Lianhui Qin, Ronan Le Bras, Raja Marjeh, Nanyun Peng, Yejin Choi, Thomas L Griffiths, Faeze Brahman.  
*MacGyver: Are Large Language Models Creative Problem Solvers?*  
2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2024).  
*[Nominated for Best Paper Award] [long paper]*
20. Yanai Elazar, Akshita Bhagia, Ian Magnusson, **Abhilasha Ravichander**, Dustin Schwenk, Alane Suhr, Evan Pete Walsh, Dirk Groeneveld, Luca Soldaini, Sameer Singh, Hannaneh Hajishirzi, Noah A. Smith, Jesse Dodge  
*What's In My Big Data?*  
International Conference on Learning Representations (ICLR 2024).  
*[Spotlight] [long paper]*
21. Peter West, Ximing Lu, Nouha Dziri, Faeze Brahman, Linjie Li, Jena D. Hwang, Liwei Jiang, Jillian Fisher, **Abhilasha Ravichander**, Khyathi Chandu, Benjamin Newman, Pang Wei Koh, Allyson Ettinger, Yejin Choi  
*The Generative AI Paradox: 'What It Can Create, It May Not Understand'*  
International Conference on Learning Representations (ICLR 2024). *[long paper]*
22. Bill Yuchen Lin, **Abhilasha Ravichander**, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, Yejin Choi  
*The Unlocking Spell on Base LLMs: Rethinking Alignment via In-Context Learning*  
International Conference on Learning Representations (ICLR 2024). *[long paper]*
23. Yuanyuan Feng, **Abhilasha Ravichander**, Yaxing Yao, Shikun Zhang, Rex Chen, Shomir Wilson, Norman Sadeh  
*Understanding How to Inform Blind and Low-Vision Users about Data Privacy through Privacy Question Answering Assistants*  
USENIX Security 2024. *[long paper]*
24. Ximing Lu, Faeze Brahman, Peter West, Jaehun Jang, Khyathi Chandu, **Abhilasha Ravichander**, Lianhui Qin, Prithviraj Ammanabrolu, Liwei Jiang, Sahana Ramnath, Nouha Dziri, Jillian Fisher, Bill Yuchen Lin, Skyler Hallinan, Xiang Ren, Sean Welleck, Yejin Choi  
*Inference-Time Policy Adapters (IPA): Tailoring Extreme-Scale LMs without Fine-tuning*  
Empirical Methods in Natural Language Processing (EMNLP 2023). *[long paper]*
25. **Abhilasha Ravichander**\*, Joe Stacey\*, Marek Rei  
*When and Why Does Bias Mitigation Work?*  
Findings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP Findings 2023). *[long paper]*
26. **Abhilasha Ravichander**, Matt Gardner, and Ana Marasović  
*CONDAQA: A Contrastive Reading Comprehension Dataset for Reasoning about Negation*  
2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022).  
*[SoCal NLP Symposium Best Paper Award] [long paper]*
27. Yuanyuan Feng, **Abhilasha Ravichander**, Shikun Zhang, Yaxing Yao, and Norman Sadeh  
*Exploring and Improving the Accessibility of Data Privacy-related Information for People Who Are Blind and Low-vision*  
7th Workshop on Inclusive Privacy and Security (WIPS 2022). *[long paper]*

28. Yanai Elazar, Nora Kassner, Shauli Ravfogel, Amir Feder, **Abhilasha Ravichander**, Marius Mosbach, Yonatan Belinkov, Hinrich Schütze, Yoav Goldberg  
*Measuring Causal Effects of Data Statistics on Language Model’s Factual Predictions*  
arXiv, 2022. [long paper]
29. Siddhant Arora, Henry Hosseini, Christine Utz, Vinayshekhar Bannihatti Kumar, Tristan O. Dhellemmes, **Abhilasha Ravichander**, Peter Story, Jasmine Mangat, Rex Chen, Martin Degeling, Thomas Norton, Thomas Hupperich, Shomir Wilson and Norman Sadeh  
*A Tale of Two Regulatory Regimes: Creation and Analysis of a Bilingual Privacy Policy Corpus*  
13th Language Resources and Evaluation Conference, (LREC 2022). [long paper]
30. Dheeraj Rajagopal, Aman Madaan, Niket Tandon, Yiming Yang, Shrimai Prabhumoye, **Abhilasha Ravichander**, Peter Clark, Eduard Hovy. *CURIE: An Iterative Querying Approach for Reasoning About Situations*  
First Workshop on Commonsense Representation and Reasoning, (CSRR@ACL 2022) [long paper]
31. **Abhilasha Ravichander**, Yonatan Belinkov, Eduard Hovy.  
*Probing the Probing Paradigm: Does Probing Accuracy Entail Task Relevance?*  
16th Conference of the European Chapter of the Association for Computational Linguistics, (EACL 2021). [long paper]
32. **Abhilasha Ravichander**, Siddharth Dalmia, Maria Ryskina, Florian Metze, Eduard Hovy and Alan Black  
*NoiseQA: Challenge Sets for User-Centric Question Answering*  
16th Conference of the European Chapter of the Association for Computational Linguistics, (EACL 2021). [long paper]
33. Yanai Elazar, Nora Kassner, Shauli Ravfogel, **Abhilasha Ravichander**, Eduard Hovy, Hinrich Schütze, Yoav Goldberg.  
*Measuring and Improving Consistency in Pretrained Language Models*  
Transactions of the Association for Computational Linguistics, (TACL 2021). [long paper]
34. **Abhilasha Ravichander**, Alan W Black, Shomir Wilson, Thomas Norton and Norman Sadeh.  
*Breaking Down Walls of Text: How Can NLP Benefit Consumer Privacy?*  
59th Annual Meeting of the Association for Computational Linguistics, (ACL 2021). [long paper]
35. **Abhilasha Ravichander**, Eduard Hovy, Kaheer Suleman, Adam Trischler, Jackie Chi Kit Cheung .  
*On the Systematicity of Probing Contextualized Word Representations: The Case of Hypernymy in BERT*  
2020 Joint Conference on Lexical and Computational Semantics, (\*SEM 2020).[long paper]
36. **Abhilasha Ravichander\***, Aakanksha Naik\*, Carolyn Rose, Eduard Hovy.  
*EQUATE: A Benchmark Evaluation Framework for Quantitative Reasoning in Natural Language Inference*  
2019 Conference on Computational Natural Language Learning, (CoNLL 2019).[long paper]
37. **Abhilasha Ravichander**, Alan W Black, Shomir Wilson, Thomas Norton and Norman Sadeh  
*Question Answering for Privacy Policies: Combining Computational and Legal Perspectives*  
2019 Conference on Empirical Methods in Natural Language Processing (EMNLP 2019).[long paper]
38. **Abhilasha Ravichander\***, Aakanksha Naik\*, Carolyn Rose, Eduard Hovy  
*Exploring Numeracy in Word Embeddings*  
57th Annual Meeting of the Association for Computational Linguistics (ACL 2019).[short paper]
39. Peter Story, Sebastian Zimmeck, Daniel Smullen, **Abhilasha Ravichander**, Ziqi Wang, Joel Reidenberg, N. Cameron Russell and Norman Sadeh  
*MAPS: Scaling Privacy Compliance Analysis to a Million Apps*  
Proceedings on Privacy Enhancing Technologies (PETS 2019). [long paper]
40. Tom Norton, Joel Reidenberg, Norman Sadeh, **Abhilasha Ravichander**  
*Evaluating How Global Privacy Principles Answer Consumers’ Questions About Mobile App Privacy*,  
4th European Privacy Law Scholars Conference (PLSC 2019).
41. **Abhilasha Ravichander\***, Aakanksha Naik\*, Norman Sadeh, Carolyn Rose, Graham Neubig  
*Stress Test Evaluation for Natural Language Inference*,  
27th International Conference on Computational Linguistics (COLING 2018)  
*[COLING Area Chair Favorite Paper Award]* [long paper]

42. **Abhilasha Ravichander**, Alan Black.  
*An Empirical Study of Self-Disclosure in Spoken Dialogue Systems*  
19th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL 2018). [long paper]
43. **Abhilasha Ravichander\***, Thomas Manzini\*, Matthias Grabmair, Graham Neubig, Eric Nyberg.  
*How Would You Say It? Eliciting Lexically Diverse Data for Supervised Semantic Parsing*  
18th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL 2017). [long paper]
44. Paul Michel\*, **Abhilasha Ravichander\***, Shruti Rijhwani\*.  
*Does the Geometry of Word Embeddings Help Document Classification? A Case Study on Persistent Homology-Based Representations,*  
Workshop on Representation Learning For NLP, Annual Meeting of the Association for Computational Linguistics (ACL 2017). [short paper]
45. Shrimai Prabhumoye\*, Fadi Botros\*, Khyathi Chandu\*, Samridhi Choudhary\*, Esha Keni\*, Chaitanya Malaviya\*, Thomas Manzini\*, Rama Pasumarthi\*, Shivani Poddar\*, **Abhilasha Ravichander\***, Zhou Yu, Alan Black.  
*Building CMU Magnus from User Feedback*, Alexa Prize Proceedings, 2017.

## ACADEMIC SERVICE

---

- **Action Editor** ACL Rolling Review 2024, ACL Rolling Review 2025
- **Area Chair** EMNLP 2022, EMNLP 2023, ACL 2023, \*SEM 2025
- **Reviewer**  
*Journals:* Computational Linguistics, Nature Communications  
*Conferences:* NAACL-HLT 2019, ACL 2020, EMNLP 2020, EACL 2021, ACL 2021, ACL Rolling Review 2021, ACL Rolling Review 2022, COLM 2024, Neurips Datasets and Benchmarks 2024, ICLR 2025, ACL Rolling Review 2025  
*Workshops:* ACL SRW 2020, AACL-IJCNLP SRW 2020, EMNLP-SDP 2020, Neurips HAMLETS workshop 2020, RepL4NLP 2021, AmericasNLP 2021, BlackboxNLP 2021, ACL SRW 2022, Blackbox NLP 2022
- **Organizer**, Workshop on Privacy in Natural Language Processing at ACL 2024 (PrivateNLP 2024)
- **Organizer**, Workshop on Representation Learning for NLP at ACL 2023 (Repl4NLP 2023)
- **Session Chair**, EMNLP 2023
- NAACL DEI **Socio-Cultural Inclusion Chair**, NAACL 2022
- Co-founder, **NLP with Friends** (<https://nlpwithfriends.com/>)
- **Student Volunteer** for DEI, CMU LTI Faculty Hiring Committee
- **Student Volunteer**, CMU LTI Ph.D. Admissions Committee
- **Student Volunteer**, NAACL-HLT, 2019, EMNLP 2019
- **Session Chair**, AAAI Spring Symposium Series, 2019
- CMU AI Research mentor (initiative to mentor under-represented minorities in Computer Science)
- **Research Team Lead**, OurCS 2019 (workshop to introduce undergraduate women to computer science research)
- **Program Committee**, CMU LTI Student Research Symposium 2018.
- **Student Volunteer**, Widening NLP Workshop at NAACL-HLT 2018.

## TEACHING

---

- Teaching Assistant, 11-727 Computational Semantics, Carnegie Mellon University, 2020
- Teaching Assistant, 10-606 Mathematical Foundations for Machine Learning, Carnegie Mellon University, 2018
- Teaching Assistant, 10-607 Computational Foundations for Machine Learning, Carnegie Mellon University, 2018
- Stanford Crowd Course Initiative (MOOC), 2015: Taught modules on recursion and computational complexity.

## INVITED TALKS

---

### DFKI, 2026

#### *Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**AAACL 2025, Workshop on Confabulation, Hallucinations and Overgeneration in Multilingual and Practical Settings**  
*Keynote*

**Indian Symposium on Machine Learning, 2025**

*Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**Microsoft Research, 2025**

*Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**Saarland University, 2025**

*Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**University of Sheffield, 2025**

*Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**Cohere Research Connections, 2025**

*Invited Talk*

Hyper-Personalized Knowledge Systems

**Women in AI Research Podcast, 2025**

**University of California, Santa Cruz, 2025**

*Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**Ploutos AI, 2025**

*Invited Talk*

Fantastic LLM Hallucinations and Where to Find Them

**University of Utah, 2025**

*Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**Technische Universität Darmstadt, 2025**

*Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**University of Mannheim, 2025**

*Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**University of Minnesota, 2025**

*Invited Talk*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**University of Massachusetts, Amherst, 2024**

*Invited Talk, Rising Stars Workshop*

Illuminating Generative AI: Mapping Knowledge in Large Language Models

**EMNLP 2024 Panel (WiNLP workshop)**

*Invited Panelist*

Navigating Research in the Age of LLMs

with Isabelle Augenstein, Mrinmaya Sachan, Sunayana Sitaram, and Lu Wang

**Minds Matter Podcast, 2024**

**Compass Tech Summit, 2023**

*Invited Talk*

How Can NLP Benefit Consumer Privacy?

**University of Massachusetts, Amherst, 2023**

*Invited Talk*

How Do We Get to Transparent Large Language Models?

**National University of Singapore, 2023**

*Invited Talk*

How Do We Get to Transparent Large Language Models?

**University of St. Gallen, 2022**

*Invited Talk*

How Can NLP Benefit Consumer Privacy?

**University of Texas at Austin, 2022**

*Invited Talk, Rising Stars Workshop*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

**Johns Hopkins University, 2022**

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

**Microsoft Research, 2022**

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

**University of Washington, 2022**

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

**Allen Institute for Artificial Intelligence, 2022**

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

**University of Rochester, 2022**

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

**George Mason University, 2022**

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

**Emory University, 2022**

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

**Dair.AI ‘Women in NLP’ seminar, 2022**

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

**TU Munich, 2021**

*Invited Talk*

How Can NLP Benefit Consumer Privacy?

**NLLP Talk Series, 2021**

*Invited Talk*

How Can NLP Benefit Consumer Privacy?

## University of Chicago, 2021

*Invited Talk, Rising Stars Workshop*

User-Centric Question Answering

## Workshop on Search-Oriented Conversational AI, 2021

*Invited Talk*

User-Centric Question Answering

## University of Illinois at Chicago, 2021

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

## University of Bocconi, 2021

*Invited Talk*

Interpreting Neural Model Performance for Robust, Trustworthy NLP

## Language Technologies Institute Seminar at Carnegie Mellon University, 2021

*Invited Panelist*

Can Large Language Models Solve NLP?

with Yonatan Bisk, Sam Bowman, and Colin Raffel

## NAACL 2021 Panel

*Invited Panelist*

Getting Into NLP Research

with William Agnew, Pan Xu, Phu Mon Htut, and Elizabeth Salesky

## 13th International Conference on Data Protection and Artificial Intelligence (CPDP 2020)

*Invited Panelist*

with Cameron Russell, Lokke Moerel, and Antoine Bon

## RESEARCH INTERNSHIPS

---

Jun '21–Aug '21	<b>Allen Institute for Artificial Intelligence</b> , Seattle, WA Research Intern <b>Advisors: Ana Marasovic, Matt Gardner</b>
Jun '19–Aug'19	<b>Microsoft Research</b> , Montreal, QC Research Intern <b>Advisors: Adam Trischler, Kaheer Suleman, Jackie Cheung</b>
Jun '14–Aug '14	<b>Institute of Mathematical Sciences</b> , Chennai, India Visiting Student <b>Advisor: Venkatesh Raman</b>

## SOFTWARE DEVELOPMENT EXPERIENCE

---

Sep '15–Jun '16	<b>Platform Engineer, Sensara Technologies</b> , Bangalore, India Designed a novel algorithm to finely segment advertisement boundaries i.e detect at a frame-level granularity when advertisements begin and end. This work is currently in production at adbreaks.in ( <a href="#">link</a> ) and used at scale to segment advertisements every day. Relevant talk about this work <a href="#">link</a> .
Jan '15–Jul '15	<b>Software Development Engineer Intern, Amazon</b> , Bangalore, India Worked on integrated module to fetch delivery charges such that they are pincode and quantity aware for Amazon China, India and UK.