

ABHILASHA RAVICHANDER

aravicha@cs.cmu.edu | www.cs.cmu.edu/~aravicha

RESEARCH INTERESTS

Natural Language Understanding, Textual Entailment, Representation Learning, Question Answering, Robustness, Interpretability

EDUCATION

2018–Present	Carnegie Mellon University- School of Computer Science, Doctor of Philosophy in Language and Information Technologies (PhD), Advisors : Eduard Hovy, Norman Sadeh
2016–2018	Carnegie Mellon University- School of Computer Science M.Sc in Language Technologies. QPA: 4.0
2011–2015	P.E.S Institute of Technology- Department of Computer Science and Engineering Bachelor of Engineering in Computer Science, First Class with Distinction

PUBLICATIONS

1. **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]
2. **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]
3. 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]
4. **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]
5. 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*
arXiv 2021 preprint, [long paper]
6. **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]
7. **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]
8. **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]
9. **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]
10. 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]

11. 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).
12. Peter Story, Sebastian Zimmeck, **Abhilasha Ravichander**, Daniel Smullen, Ziqi Wang, Joel Reidenberg, N. Cameron Russell
and Norman Sadeh, *Natural Language Processing for Mobile App Privacy Compliance*,
AAAI Spring Symposium Series, 2019.
13. **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) ****Area Chair Favorite Paper**** [long paper]
14. **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]
15. **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]
16. 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, Association for Computational Linguistics, 2017 (ACL 2017). [short paper]
17. **Abhilasha Ravichander***, Shruti Rijhwani*, Rajat Kulshreshtha*, Chirag Nagpal, Tadas Baltrušaitis, Louis-Phillipe Morency.
Preserving Intermediate Objectives: One Simple Trick to Improve Learning for Hierarchical Models, arXiv, 2017.
18. Shrimai Prabhunoye*, 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 HONORS

- Invited to Rising Stars in Data Science, 2021
- Outstanding reviewer, NAACL 2019, EMNLP 2020, ACL 2020
- NAACL Student Scholarship, 2019
- Area Chair Favorite Paper at COLING, 2018
- Team awarded **100,000\$** stipend by Amazon to compete in the AlexaPrize, 2016.
- Recipient of Graduate Research Fellowship from Carnegie Mellon University, 2016/17
- Best Poster Award, Machine Learning Project Symposium at Carnegie Mellon University
- Travel grant award, Workshop for Women in Machine Learning (NeurIPS 2015)

ACADEMIC SERVICE

- Socio-Cultural Inclusion Chair, NAACL 2022
- Co-organizer, **NLP with Friends** seminar series
- Student Volunteer for DEI, CMU LTI Faculty Hiring Committee
- Reviewer, NAACL-HLT 2019, ACL 2020, ACL SRW 2020, AACL-IJCNLP SRW 2020, EMNLP 2020, EMNLP-SDP 2020,
Neurips HAMLETS workshop 2020, EACL 2021, ACL 2021, AmericasNLP 2021, RepL4NLP 2021, BlackboxNLP 2021, ACL
Rolling Review 2021
- 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.

RESEARCH INTERNSHIPS

Jun '21–Present	Research Intern, Allen Institute of AI , Seattle, WA <i>Advisors: Ana Marasovic, Matt Gardner</i> <i>Project Topic: A Benchmark for Negation in Natural Language Understanding</i>
Jun '19–Aug'19	Research Intern, Microsoft Research , Montreal, QC <i>Advisors: Adam Trischler, Kaheer Suleman</i> <i>Project Topic: Hypernymy in Neural Models of Language</i>
Jun '14–Aug '14	Visiting Student, Institute of Mathematical Sciences , Chennai, India <i>Advisors: Venkatesh Raman</i> I worked on problems in parameterized complexity, with particular interest in $K_{i,j}$ -free graphs. I developed new Fixed Parameter Tractable algorithms for three problems- RED-BLUE DOMINATING SET, CONSTRAINT BIPARTITE DOMINATING SET and THRESHOLD DOMINATING SET in $K_{i,j}$ -free graphs, defined a new Dominating Set variant, the THRESHOLD BIPARTITE DOMINATING SET, and presented a Fixed Parameter Tractable algorithm for it. I proved the $W[1]$ -hardness of the DENSE SUBGRAPH problem in $K_{i,j}$ -free graphs.

SOFTWARE DEVELOPMENT EXPERIENCE

Sep '15–Jun '16	Platform Engineer, Sensara Technologies , 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 (link) and used at scale to segment advertisements every day. Relevant talk about this work link .
Jan '15–Jul '15	Software Development Engineer Intern, Amazon , Bangalore, India Worked on integrated module to fetch delivery charges such that they are pincode and quantity aware for Amazon China, India and UK.

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

- "Understanding Neural Model Performance"
University of Illinois at Chicago, 2021
- "User-Centric Question Answering"
Workshop on Search-Oriented Conversational AI, 2021
University of Chicago, 2021
- "How Can NLP Benefit Consumer Privacy?"
TU Munich, 2021
- Invited Panelist, 'Can Large Language Models Solve NLP?', Language Technologies Institute Seminar at Carnegie Mellon University, with *Yonatan Bisk, Sam Bowman, and Colin Raffel*
- Invited Panelist, NAACL 2021 Panel on 'Getting Into NLP Research', with *William Agnew, Pan Xu, Phu Mon Htut, and Elizabeth Salesky*
- Invited Panelist, 13th International Conference on Data Protection and Artificial Intelligence (CPDP 2020), with *Cameron Russell, Lokke Moerel, and Antoine Bon*

WORKSHOPS AND SUMMER SCHOOLS

- 2018 | “Position” paper, **Young Researchers Round Table on Spoken Dialog Systems**. Melbourne, Australia.
- 2017 | **Machine Learning Summer School**, *Max Planck Institute for Information Systems*. Tübingen, Germany.
- 2014 | **Summer School on Non-Convex Optimization for Machine Learning**, *IIT Bombay*. Bombay, India.
- 2013 | **Summer Student Program in Theoretical Computer Science**, *Institute of Mathematical Sciences*. Chennai, India.
- 2013 | MIT Media Lab Design Innovation Workshop (Bangalore, India)