

# Chaitanya Malaviya

412-816-7472 | [cmalaviy@seas.upenn.edu](mailto:cmalaviy@seas.upenn.edu) | <https://chaitanyamalaviya.github.io/>

## EDUCATION

---

<b>University of Pennsylvania</b> PhD in Computer and Information Science <i>Advisors: Mark Yatskar and Dan Roth</i>	Sep 2020 – May 2025 (expected)
<b>Carnegie Mellon University</b> Masters of Language Technologies (MLT) <i>Advisor: Graham Neubig</i>	Aug 2016 - Aug 2018
<b>Nanyang Technological University, Singapore</b> Bachelors in Computer Engineering <i>Fully funded through the SIA-NOL Youth Scholarship, Presidents' Research Scholar</i>	Aug 2012 - May 2016

## RESEARCH EXPERIENCE

---

<b>Graduate Research Assistant</b> <b>University of Pennsylvania</b> <ul style="list-style-type: none"><li>• Evaluation &amp; benchmarking of language models in realistic scenarios [1 - 4] [6] [7].</li><li>• Aligning models to diverse users, especially domain experts [1] [2] [4] [5] [6].</li><li>• Post-training and learning from human feedback [5].</li></ul>	Sep 2020 - ongoing <i>Philadelphia, PA</i>
<b>Research Intern</b> <b>Allen Institute for Artificial Intelligence (Ai2)</b> <ul style="list-style-type: none"><li>• Contextualized evaluation of language models [1].</li><li>• Evaluation of scientific question answering.</li></ul>	May 2024 - ongoing <i>Seattle, WA</i>
<b>Student Researcher</b> <b>Google DeepMind</b> <ul style="list-style-type: none"><li>• Evaluating language models on domain-specific long-form writing tasks [2].</li><li>• Gemini 1.5 release evaluations [3].</li></ul>	Sep 2023 - March 2024 <i>New York City, NY</i>
<b>Student Researcher</b> <b>Google Brain</b> <ul style="list-style-type: none"><li>• Evaluating retrieval for queries with implicit constraints [7].</li></ul>	June 2022 - Dec 2022 <i>Seattle, WA</i>
<b>Predoctoral Young Investigator</b> <b>Allen Institute for Artificial Intelligence (Ai2)</b> <ul style="list-style-type: none"><li>• Mosaic Team (Commonsense Reasoning) [12 - 15].</li></ul>	Oct 2018 – Aug 2020 <i>Seattle, WA</i>
<b>Graduate Research Assistant</b> <b>Language Technologies Institute, Carnegie Mellon University</b> <ul style="list-style-type: none"><li>• Morphology, typology, machine translation [16 - 21].</li></ul>	Aug 2016 – Aug 2018 <i>Pittsburgh, PA</i>

## AWARDS

---

- Outstanding Paper Award, ACL 2023
- Departmental PhD Grant, University of Pennsylvania 2020
- Graduate Research Fellowship, Carnegie Mellon University 2016-18
- EMNLP Student Scholarship 2017
- Amazon Alexa Grant to CMU for participation in the Alexa Prize 2017
- SIA-NOL Youth 100% Scholarship for undergraduate studies 2012 –2016
- Presidents’ Research Scholar, School of Computer Science and Engineering, NTU 2014-2015
- Best Industrial Attachment Project, School of Computer Science and Engineering, NTU 2015

## PUBLICATIONS

---

- [1] **Chaitanya Malaviya**, Joseph Chee Chang, Dan Roth, Mohit Iyyer, Mark Yatskar, Kyle Lo. *Contextualized Evaluations: Taking the Guesswork Out of Language Model Evaluations*. Transactions of the Association for Computational Linguistics (TACL) 2024.
- [2] **Chaitanya Malaviya**, Priyanka Agrawal, Kuzman Ganchev, Pranesh Srinivasan, Fantine Huot, Jonathan Berant, Mark Yatskar, Dipanjan Das, Mirella Lapata, Chris Alberti. **DOLOMITES: Domain-specific Long-form Methodical Tasks**. Transactions of the Association for Computational Linguistics (TACL) 2024. ([Website](#))
- [3] Gemini Team, Google. Contributor (Evals): **Chaitanya Malaviya**. *Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context*. Technical Report 2024.
- [4] Ori Yoran, Samuel Amouyal, **Chaitanya Malaviya**, Ben Bogin, Ofir Press, Jonathan Berant. *AssistantBench: Can Web Agents Solve Realistic and Time-consuming Tasks?*. Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP) 2024.
- [5] **Chaitanya Malaviya**, Subin Lee, Dan Roth, Mark Yatskar. *What if you said that differently?: How Explanation Formats Affect Human Feedback Efficacy and User Perception*. Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL) 2024.
- [6] **Chaitanya Malaviya**, Subin Lee, Sihao Chen, Mark Yatskar, Dan Roth. **EXPERTQA: Expert-Curated Questions and Attributed Answers**. Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL) 2024.
- [7] **Chaitanya Malaviya**, Peter Shaw, Ming-Wei Chang, Kenton Lee, Kristina Toutanova. **QUEST: A Dataset of Entity-Seeking Queries with Implicit Set Operations**. Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL) 2023. (**Outstanding paper award**)
- [8] **Chaitanya Malaviya**, Sudeep Bhatia, Mark Yatskar. *Cascading Biases: Investigating the Effect of Heuristic Annotation Strategies on Data and Models*. Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP) 2022.

- [9] Sihao Chen, **Chaitanya Malaviya**, Hagai Taitelbaum, Tal Schuster, Senaka Buthpitiya, Alex Fabrikant. *Reference (In-)Determinacy in Natural Language Inference*. Preprint 2024.
- [10] Qing Lyu, Kumar Shridhar, **Chaitanya Malaviya**, Li Zhang, Yanai Elazar, Niket Tandon, Marianna Apidianaki, Mrinmaya Sachan, Chris Callison-Burch. *Calibrating Large Language Models with Sample Consistency*. AACL 2025.
- [11] Lucy Yuewei Yuan, **Chaitanya Malaviya**, Mark Yatskar. *AmbiCoref: Evaluating Human and Model Sensitivity to Ambiguous Coreference*. Findings of the Conference of the European Chapter of the Association for Computational Linguistics (EACL) 2023.
- [12] Yiben Yang, **Chaitanya Malaviya**, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Ji-Ping Wang, Chandra Bhagavatula, Yejin Choi, Doug Downey. *Generative Data Augmentation for Commonsense Reasoning*. Findings of ACL: EMNLP 2020.
- [13] **Chaitanya Malaviya**, Chandra Bhagavatula, Antoine Bosselut, Yejin Choi. *Commonsense Knowledge Base Completion with Structural and Semantic Context*. Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI) 2020.
- [14] Chandra Bhagavatula, Ronan Le Bras, **Chaitanya Malaviya**, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Scott Wen-tau Yih, Yejin Choi. *Abductive Commonsense Reasoning*. Proceedings of the International Conference on Learning Representations (ICLR) 2020.
- [15] Antoine Bosselut, Hannah Rashkin, Maarten Sap, **Chaitanya Malaviya**, Asli Celikyilmaz and Yejin Choi. *CoMET: Commonsense Transformers for Automatic Knowledge Graph Construction*. Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL) 2019.
- [16] **Chaitanya Malaviya\***, Shijie Wu\*, Ryan Cotterell. *A Simple Joint Model for Improved Contextual Neural Lemmatization*. Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL) 2019.
- [17] **Chaitanya Malaviya**, Matthew R. Gormley, Graham Neubig. *Neural Factor Graph Models for Cross-lingual Morphological Tagging*. Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL) 2018.
- [18] **Chaitanya Malaviya**, Pedro Ferreira, André Martins. *Sparse and Constrained Attention for Neural Machine Translation*. Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL) 2018.
- [19] **Chaitanya Malaviya**, Graham Neubig and Patrick Littell. *Learning Language Representations for Typology Prediction*. Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP) 2017.
- [20] Arya D. McCarthy, Ekaterina Vylomova, Shijie Wu, **Chaitanya Malaviya**, Lawrence Wolf-Sonkin, Garrett Nicolai, Miikka Silfverberg, Sebastian Mielke, Jeffrey Heinz, Ryan Cotterell, Mans Hulden. *The SIGMORPHON 2019 Shared Task: Morphological Analysis in Context and Cross-Lingual Transfer for Inflection*. SIGMORPHON 2019.
- [21] Shrimai Prabhumoye\*, Fadi Botros\*, **Chaitanya Malaviya\***, Zhou Yu, Alan Black et al. *Building CMU Magnus from User Feedback*. In Proceedings of AWS re:INVENT, 2017.

## SERVICE AND OUTREACH

---

- **Area Chair/Action Editor:** ACL Rolling Review (2024-).
- **Reviewer:** ACL (2019-2023), NAACL (2019, 2021), EMNLP (2019-2022), ACL Rolling Review (2021-), COLM (2024), NeurIPS (2021), ICLR (2022, 2023), AAAI (2021), JAIR.
- **Co-organizer:** SIGMORPHON 2019 Shared Task: Crosslinguality and Context in Morphology.
- **Student Volunteer:** ACL 2018, EMNLP 2017, Interspeech 2014.
- **Mentor** (2020, 2021), Graduate Student Mentorship Initiative (GSMI), Científico Latino. Assisted one student per year with their graduate school application.
- **Mentor** for the PhD Applicant Support Program for students from underrepresented groups @ UPenn CIS - 2020: Assisted two students with graduate school applications.
- **Mentor** for the PhD Program @ UPenn CIS - 2022: Mentored one 1st year PhD student.
- **Co-organizer** of CLunch (Computational Linguistics Lunch), University of Pennsylvania, Spring 2022.
- Member of the Diversity, Equity and Inclusion Committee at the Allen Institute for AI.

## MENTORSHIP

---

- Anirudh Bharadwaj, May 2024-, University of Pennsylvania
- Nikhil Kumar, July 2024-, University of Pennsylvania
- Yifei Li, Nov 2023-, University of Pennsylvania
- Subin Lee, May 2023-Oct 2023, University of Pennsylvania
- Lucy Yuewei Yuan, Sep 2022-March 2023, University of Pennsylvania
- Sarah Payne, Katherine Hann, Madeline Kalil, Sep 2021-Dec 2021, University of Pennsylvania

## TEACHING

---

- **Teaching Assistant** at the University of Pennsylvania. CIS 530: Computational Linguistics, Spring 2020 and Fall 2024.

## TALKS

---

- Talk at the Allen Institute for AI, 2024. [DOLOMITES: Domain-Specific Long-Form Methodical Tasks](#).
- Talk at EMNLP, 2024. [DOLOMITES: Domain-Specific Long-Form Methodical Tasks](#).
- Talk at Google DeepMind, 2024. [ExpertQA: Expert-Curated Questions and Attributed Answers](#).
- Talk at NAACL, 2024. [ExpertQA: Expert-Curated Questions and Attributed Answers](#).
- Talk at NAACL, 2024. [How do Explanation Formats Affect Human Feedback Efficacy?](#).
- Talk at ACL 2023. [QUEST: A Dataset of Entity-Seeking Queries with Implicit Set Operations](#).
- Talk at EMNLP 2022. [Cascading Biases: How Biases Transfer from Data to Models](#).
- Talk at NAACL 2019. [Contextual Neural Lemmatization](#).
- Talk at ACL 2018. [Sparse and Constrained Attention for NMT](#).
- Talk at Allen Institute for Artificial Intelligence. [Neural Factor Graph Models](#).