Mourad Heddaya

mourad@uchicago.edu | mheddaya.com | 425-753-5800

My research focuses on open-ended problems in language where there is no clear ground truth. I study how to evaluate and measure language in these settings, and how to design AI interventions that help people reason and make better decisions.

Education

Ph.D. Student in Computer Science, 2021-,
Expected Graduation May 2026.
University of Chicago, Chicago, IL.
Advisor: Chenhao Tan

B.S. in Informatics, 2015-2019,
University of Washington, Seattle, WA.
Research Supervisor: Noah Smith & Mari Ostendorf

Publications

[Measuring and evaluating language in open-ended settings.]

- *CaseSumm: A Large-Scale Dataset for Long-Context Summarization from U.S. Supreme Court Opinions*. **Mourad Heddaya**, K. MacMillan, Hongyuan Mei, Chenhao Tan, A. Malani.
  NAACL 2025 Findings. Accepted with talk at ALEA 2024.
- *A Century of Inflation Narratives*. **Mourad Heddaya**, Chenhao Tan, R. Voigt, Q. Zeng, A. Zentefis.
  SSRN Working Paper, 2025.
- *Causal Micro-Narratives*. **Mourad Heddaya**, Q. Zeng, R. Voigt, A. Zentefis, Chenhao Tan.
  EMNLP 2024 Workshop on Narrative Understanding.
- *Language of Bargaining*. **Mourad Heddaya**, S. Dworkin, R. Voigt, A. Zentefis, Chenhao Tan.
  ACL 2023 Main Conference.
- *LLM Rationalis? Measuring Bargaining Capabilities of AI Negotiators*. C. Shah, A. Agarwal, K. Garg, **Mourad Heddaya**.
  NeurIPS 2025 Workshop on Multi-Turn Interactions in LLMs.

[Improving language model reasoning.]

- *When Internalization Fails: Finding Better Targets for Reasoning Compression*. **Mourad Heddaya**, R. Wadhawan, M. Roberts, Chenhao Tan.
  Under review at ACL 2026.

Internships

Research Scientist Intern at Abridge, Summer 2025.
San Francisco, CA.

- Worked on reasoning compression for language models.

Applied Scientist at Amazon AWS AI Labs, Summer 2023.
Bedrock Team, JFK 14, New York City, NY.
Mentor: Miguel Ballesteros

- Proposed self-supervised alignment, an efficient method for aligning LLMs to human preferences for summarization and toxicity without RLHF (without RL and with less human feed-

back).
- Allow the model to score its own hypotheses (sampled sentences) and incorporate it as self-feedback in the SFT loop, providing more effective regularization for better alignment.
- Project outcome: delivered internal technical report, documented code, and presentation.

Invited Talks

Freestone Grove Partners,
April 2025
Talk Topic: Causal Micro-Narratives

Max Planck Institute for Research on Collective Goods,
Research Group Engel, February 2025
Talk Topic: NLP In the Legal Domain (summarization, reasoning, etc). Talk to occur in early 2025.

University of Chicago,
Language Evolution Acquisition & Processing Workshop (LEAP), January 2023
Talk Title: Language of Bargaining

Teaching

University of Chicago,
CMSC 25400 – Machine Learning, Winter 2023
CMSC 25300 / 35300 – Mathematical Foundation of Machine Learning, Fall 2022
CMSC 35100 - Natural Language Processing, Winter 2022