Belinda Z. Li

✉ bzl@mit.edu — 🌐 belindal.github.io —

Education

Massachusetts Institute of Technology   2020 –
PhD candidate, Electrical Engineering and Computer Science
- **Masters degree** awarded Feb. 2023
- GPA: 5.0/5.0
- Advisor: Jacob Andreas

University of Washington   2015 – 2019
Bachelors of Science, Computer Science, Magna Cum Laude
- Cumulative GPA: 3.95 / 4.0, Major GPA: 3.96 / 4.0
- Advisor: Luke Zettlemoyer

Publications

Peer-Reviewed Publications

[1]   Andi Peng, Andreea Bobu, **Belinda Z. Li**, Theodore Sumers, Ilia Sucholutsky, Nishanth Kumar, Thomas Griffiths, and Julie Shah. "Preference-Conditioned Language-Guided Abstractions". In: *Proceedings of the ACM/IEEE International Conference on Human Robot Interaction (HRI 2024)*. ACM/IEEE. 2024.

[2]   Andi Peng, Ilia Sucholutsky, **Belinda Z. Li**, Theodore Sumers, Thomas Griffiths, Jacob Andreas, and Julie Shah. "Learning with language-guided state abstractions". In: *RSS Workshop on Social Intelligence in Humans and Robots*. 2023.

[4]   **Belinda Z. Li**, Jason Eisner, Adam Pauls, and Sam Thomson. "Toward Interactive Dictation". In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Toronto, Canada: Association for Computational Linguistics, July 2023, pp. 15319–15338.

[7]   **Belinda Z. Li**, Maxwell Nye, and Jacob Andreas. "Language Modeling with Latent Situations". In: *Findings of the Association for Computational Linguistics: ACL 2023*. Toronto, Canada: Association for Computational Linguistics, July 2023, pp. 12556–12571.

[8]   **Belinda Z. Li**, Jane Yu, Madian Khabsa, Luke Zettlemoyer, Alon Halevy, and Jacob Andreas. "Quantifying Adaptability in Pre-trained Language Models with 500 Tasks". In: *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Seattle, United States: Association for Computational Linguistics, July 2022, pp. 4696–4715.

[9]   Qinyuan Ye, **Belinda Z. Li**, Sinong Wang, Benjamin Bolte, Hao Ma, Wen-tau Yih, Xiang Ren, and Madian Khabsa. "On the Influence of Masking Policies in Intermediate Pre-training". In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics, Nov. 2021, pp. 7190–7202.

[10] **Belinda Z. Li**, Maxwell Nye, and Jacob Andreas. "Implicit Representations of Meaning in Neural Language Models". In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Online: Association for Computational Linguistics, Aug. 2021, pp. 1813–1827.

[11] Nayeon Lee, **Belinda Z. Li**, Sinong Wang, Pascale Fung, Hao Ma, Wen-tau Yih, and Madian Khabsa. "On Unifying Misinformation Detection". In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, June 2021, pp. 5479–5485.

[12] **Belinda Z. Li**, Sewon Min, Srinivasan Iyer, Yashar Mehdad, and Wen-tau Yih. "Efficient One-Pass End-to-End Entity Linking for Questions". In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, Nov. 2020, pp. 6433–6441.

[13] Nayeon Lee, **Belinda Z. Li**, Sinong Wang, Wen-tau Yih, Hao Ma, and Madian Khabsa. "Language Models as Fact Checkers?" In: *Proceedings of the Third Workshop on Fact Extraction and VERification (FEVER)*. Association for Computational Linguistics, July 2020, pp. 36–41.

[14] **Belinda Z. Li**, Gabriel Stanovsky, and Luke Zettlemoyer. "Active Learning for Coreference Resolution using Discrete Annotation". In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, July 2020, pp. 8320–8331.

Preprints

[3] **Belinda Z. Li**, Alex Tamkin, Noah Goodman, and Jacob Andreas. *Eliciting Human Preferences with Language Models*. 2023.

[5] Evan Hernandez, **Belinda Z. Li**, and Jacob Andreas. *Measuring and Manipulating Knowledge Representations in Language Models*. 2023.

[6] **Belinda Z. Li**, William Chen, Pratyusha Sharma, and Jacob Andreas. *LaMPP: Language Models as Probabilistic Priors for Perception and Action*. 2023.

[15] Sinong Wang, **Belinda Z. Li**, Madian Khabsa, Han Fang, and Hao Ma. *Linformer: Self-Attention with Linear Complexity*. 2020.

[16] Qinyuan Ye, **Belinda Z. Li**, Sinong Wang, Benjamin Bolte, Hao Ma, Xiang Ren, Wen-tau Yih, and Madian Khabsa. *Studying Strategically: Learning to Mask for Closed-book QA*. 2020.

Industry Experience

Microsoft Semantic Machines   June 2022 – Aug. 2022
Research Intern

Facebook AI   June 2021 – Aug. 2021
Research Intern, AI Integrity Team

Facebook AI   Aug. 2019 – Sep. 2020
Software Engineer, AI Integrity Team (Facebook AI Applied Research)
- Built multi-lingual, multi-modal (image + text) representations for Facebook posts. Shipped to hate speech domain.
https://ai.facebook.com/blog/ai-advances-to-better-detect-hate-speech/

Facebook   June 2018 – Aug. 2018
Software Engineer Intern, Bot Detection/Traffic Quality Team
- Worked on automated bot detection using mouse movements.

Awards

Clare Boothe Luce Graduate Fellowship	2021
National Defence Science and Engineering Graduate Fellowship	2021
National Science Foundation Graduate Research Fellowship *(Declined)*	2021
Ida M. Green Memorial Fellowship	2020 – 2021
Denice Dee Denton Scholars Endowment	2018 – 2019
UW Dean's List	2015 – 2017 (annual), SP18/F18/SP19 (quarterly)
	(every eligible quarter)
National Merit Scholar	2016

Invited Talks

World State in Language Models
MILA   October 2023
ServiceNow   October 2023

Implicit representations of meaning in neural language models
Google N2Formal Team   September 2021

MIT EI Workshop	September 2021
Google DeepMind	July 2021
Primer AI	June 2021

Efficient One-Pass End-to-End Entity Linking for Questions
IE Workshop @ Facebook AI's 2020 NLP Summit   Aug. 2020

Research Mentorship

MEngs / Masters
1. Kunal Handa (Brown BS; Oxford Masters)   July 2023 –
2. Will Chen (MIT MEng; *now PhD @ UC Berkeley*)   May 2022 – Feb. 2023

Undergraduates
1. Nitya Parthasarathy (MIT UROP)   Sep. 2020 – June 2021

Teaching

Massachusetts Institute of Technology   Sep. 2023 –
Teaching Assistant
- Hosting office hours, designing and grading homeworks, leading tutorials
  1. *6.861*: Quantitative Methods for NLP*   Fall 2023

University of Washington   Sep. 2017 – June 2019
Teaching Assistant
- Leading weekly section, designing section material, hosting office hours and review sessions.
  1. *CSE 311: Foundations of Computing I (Discrete Math)*, with Kevin Zatloukal   Spring 2019
  2. *CSE 312: Foundations of Computing II (Probability/Statistics)*, with Martin Tompa   Winter 2019
  3. *CSE 312: Foundations of Computing II (Probability/Statistics)*, with Martin Tompa   Winter 2018
  4. *CSE 331: Software Design and Implementation*, with Kevin Zatloukal   Autumn 2017

Service

Workshop Organizing
- *Wordplay: When Language Meets Games*   ACL 2024

Program Committee – Reviewer
- ACL 2023
- COLING 2022, 2020
- EMNLP 2022, 2021, 2020

Student Volunteer/Scholarship
- ACL 2021, 2020

MIT GAAP Mentor   2020, 2021
- Mentoring individuals from underrepresented groups through the graduate application process.

Miscellany

MIT Graduate Student Union   2021 –
- *Constitution Committee*, drafted GSU's constitution   2023 –
- *Bargaining Committee*, bargained GSU's first contract   2022 – 2023
- *Organizer*   2021 –
Harvard Ballet Company, *Member*   2021 – 2022

Last Updated

December 27, 2023