

MANDAR JOSHI

✉ mandar90@cs.washington.edu 🌐 <https://homes.cs.washington.edu/~mandar90/>

🎓 <https://scholar.google.com/citations?user=uCJk8mIAAAAJ&hl>

🐙 <https://github.com/mandarjoshi90>

PROFILE

I'm a research scientist at Google DeepMind in Seattle. My interests are centered around agents, tool use, and multimodal learning. I often work on solutions involving self-supervised methods and reinforcement learning.

EDUCATION

Doctor of Philosophy (PhD) Computer Science & Engineering University of Washington, Seattle Advisers: Luke Zettlemoyer and Dan Weld Research: Natural Language Processing	2015 - 2022
Master of Technology Computer Science & Engineering Indian Institute of Technology (IIT) Bombay	2012 - 2014
Bachelor of Technology Computer Science & Engineering National Institute of Technology (NIT) Nagpur	2008 - 2012

EMPLOYMENT

- **Google DeepMind**
Research Scientist, May 2022 - Present
Gemini post-training: Agents, tool-use, reinforcement learning, and multimodality for Gemini.
- **Facebook AI Research (FAIR)**
Visiting Researcher, April 2020 - March 2022
Pre-training and few-shot learning for natural language processing.
- **Google**
Research Intern, June 2019 - January 2020
Integrating background knowledge for reading comprehension.
- **Facebook AI Research (FAIR)**
Student Researcher, October 2018 - June 2019
Pre-training methods for natural language processing.
- **Allen Institute for Artificial Intelligence (AI2)**
Research Intern, September 2017 - December 2017
Interpretable methods for question answering.
- **IBM Research**
Software Engineer, August 2014 - May 2015
Research and software development in natural language processing.

SELECTED PUBLICATIONS

- Gemini Team, Google. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. ArXiv 2507.06261, 2025.
- Shikhar Murty, Christopher Manning, Peter Shaw, **Mandar Joshi**, Kenton Lee. BAGEL: Bootstrapping Agents by Guiding Exploration with Language. ICML 2024.
- Peter Shaw*, **Mandar Joshi***, James Cohan, Jonathan Berant, Panupong Pasupat, Hexiang Hu, Urvashi Khandelwal, Kenton Lee, Kristina Toutanova. From Pixels to UI Actions: Learning to Follow Instructions via Graphical User Interfaces. NeurIPS 2023. * equal contribution
- Google AI Team. PaLI-X: On Scaling up a Multilingual Vision and Language Model. ArXiv 2305.18565, 2023.
- Kenton Lee*, **Mandar Joshi***, Iulia Turc, Hexiang Hu, Fangyu Liu, Julian Eisenschlos, Urvashi Khandelwal, Peter Shaw, Ming-Wei Chang, Kristina Toutanova. Pix2Struct: Screenshot Parsing as Pretraining for Visual Language Understanding. ArXiv 2210.03347, 2022.

- Armen Aghajanyan, Dmytro Okhonko, Mike Lewis, **Mandar Joshi**, Hu Xu, Gargi Ghosh, Luke Zettlemoyer. HTLM: Hyper-Text Pre-Training and Prompting of Language Models. *ArXiv:2107.06955, 2021*.
- Arie Cattan, Alon Eirew, Gabriel Stanovsky, **Mandar Joshi**, Ido Dagan. Cross-document Coreference Resolution over Predicted Mentions. *In ACL 2021 Findings (Short)*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, **Mandar Joshi**, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, Veselin Stoyanov. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *ArXiv:1907.11692, 2019*.
- **Mandar Joshi**, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, Omer Levy. SpanBERT: Improving Pre-training by Representing and Predicting Spans. *In TACL 2019*.
- **Mandar Joshi**, Eunsol Choi, Omer Levy, Daniel S. Weld, Luke Zettlemoyer. pair2vec: Compositional Word-Pair Embeddings for Cross-Sentence Inference. *In NAACL 2019*.
- **Mandar Joshi**, Eunsol Choi, Daniel S. Weld, Luke Zettlemoyer. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. *In ACL 2017*.
- **Mandar Joshi**, Uma Sawant, Soumen Chakrabarti. Knowledge Graph and Corpus Driven Segmentation and Answer Inference for Telegraphic Entity-seeking Queries. *In EMNLP 2014*.

AWARDS AND HIGHLIGHTS

- Invited talk on “Efficient Scalable Pre-training for Natural Language Processing” at **KDD 2020 Deep Learning Day**.
- Awarded **Microsoft Endowment Fellowship** for the academic year 2015-2016.
- **All India Rank 2** in GATE (Graduate Aptitude Test in Engineering) 2012 amongst 150,000 applicants. The test is conducted by IITs for admission into their graduate programs.

SERVICE AND COURSES

- *Service*: Program committee for ACL, EMNLP, and NAACL conferences.
- *Graduate Teaching Assistant*: Natural Language Processing (Autumn 2018), Web Search and Mining (Autumn 2013), Computer Programming (Autumn 2012 and Spring 2013).
- *Selected Graduate Coursework*: Natural Language Processing, Machine Learning, Implementation Techniques in Relational Databases, Advanced Machine Learning, Web Search and Mining, Organization of Web Information, Convex Optimization, Computer Vision.