Archiki Prasad

archiki.github.io @ archiki@cs.unc.edu github.com/archiki
linkedin.com/in/archiki-prasad Google Scholar

Research Interests

My research goal is to build natural language processing systems that can reason in an efficient, robust, and interpretable manner.

Major Interests: Reasoning, Post-Training, Reward Modeling, Self-Supervised Learning.
Other Interests: Compositional Learning, Robustness, Explainability.

Education

Present – Aug 2021 | **The University of North Carolina, Chapel Hill, USA**
Ph.D. in Computer Science | *Advisor: Mohit Bansal*
Concentration: Natural Language Processing

May 2021 – August 2016 | **Indian Institute of Technology Bombay, Maharashtra, India**
Bachelor + Master of Technology, Major: Electrical Engineering | **GPA:** 9.66/10
Minor: Computer Science and Engineering

Experience

Aug 2025 – May 2025 | **Google DeepMind, Seattle, US**
Student Researcher | *Advisors: Pete Shaw, Kenton Lee, Mandar Joshi*
› Developing metrics to evaluate and measure the utility of chain-of-thought rationales

Aug 2024 – May 2024 | **Fundamental AI Research Labs, Meta, New York City, US**
Research Scientist Intern | *Advisors: Jason Weston, Maryam Fazel-Zarandi*
› Enabling LLMs to learn to reason in an iterative and unsupervised manner

Aug 2023 – May 2023 | **Allen Institute of Artificial Intelligence (AI2), Seattle, US**
Research Intern | *Advisors: Tushar Khot, Ashish Sabharwal, Peter Clark*
› Designed an adaptive task decomposition framework for LLM agents on interactive tasks

Aug 2022 – May 2022 | **Adobe Research, San Jose (Remote), US**
Research Scientist Intern (NLP) | *Advisors: Trung Bui, David Yoon, Franck Dernoncourt*
› Developed a challenging benchmark on extracting question-answer pairs from meeting transcripts

Publications

2026 **Archiki Prasad**, Mandar Joshi, Kenton Lee, Mohit Bansal, Peter Shaw *"Effective Reasoning Chains Reduce Intrinsic Dimensionality"* Arxiv Preprint (Under Review) [PDF]

2026 Jaewoo Lee, **Archiki Prasad**, Justin Chih-Yao Chen, Zaid Khan, Elias Stengel-Eskin, Mohit Bansal *"PRInTS: Reward Modeling for Long-Horizon Information Seeking"* Arxiv Preprint (Under Review) [PDF]

2026 Zaid Khan, **Archiki Prasad**, Elias Stengel-Eskin, Jaemin Cho, Mohit Bansal *"One Life to Learn: Inferring Symbolic World Models for Stochastic Environments from Unguided Exploration"* In Proceedings of the fourteenth International Conference on Learning Representations **(ICLR 2026)**[PDF]

2025 **Archiki Prasad**, Elias Stengel-Eskin, Justin Chih-Yao Chen, Zaid Khan, Mohit Bansal *"Learning to Generate Unit Tests for Automated Debugging"* In Proceedings of the Second Conference of Language Modeling (**COLM 2025**)[PDF]

2025 **Archiki Prasad**, Weizhe Yuan, Richard Yuanzhe Pang, Jing Xu, Maryam Fazel-Zarandi, Mohit Bansal, Sainbayar Sukhbaatar, Jason Weston, Jane Yu *"Self-Consistency Preference Optimization"* In Proceedings of the International Conference on Machine Learning (**ICML 2025**)[PDF]

2025 Swarnadeep Saha, **Archiki Prasad**, Justin Chih-Yao Chen, Peter Hase, Elias Stengel-Eskin, Mohit Bansal *"System-1.x: Learning to Balance Fast and Slow Planning with Language Models"* In Proceedings of the thirteenth International Conference on Learning Representations **(ICLR 2025)** [PDF]

2025 Duy Nguyen*, **Archiki Prasad***, Elias Stengel-Eskin, Mohit Bansal *"LASER: Learning to Adaptively Select Reward Models with Multi-Armed Bandits"* In Proceedings of the Annual Conference on Neural Information Processing Systems (**NeurIPS 2025**) [PDF]

2025 Han Wang, **Archiki Prasad**, Elias Stengel-Eskin, Mohit Bansal *"Retrieval-Augmented Generation with Conflicting Evidence"* In Proceedings of the Second Conference of Language Modeling (**COLM 2025**) [PDF]

2025 Han Wang, **Archiki Prasad**, Elias Stengel-Eskin, Mohit Bansal *"ADACAD: Adaptively Decoding to Balance Conflicts between Contextual and Parametric Knowledge"* In Proceedings of the North American Chapter of the Association for Computational Linguistics (**NAACL 2025**) [PDF]

2024 **Archiki Prasad**, Elias Stengel-Eskin, Mohit Bansal *"Rephrase, Augment, Reason: Visual Grounding of Questions for Vision-Language Models"* In Proceedings of the twelfth International Conference on Learning Representations **(ICLR 2024)** [PDF]

2024 **Archiki Prasad**, Alexander Koller, Mareike Hartmann, Peter Clark, Ashish Sabharwal, Mohit Bansal, Tushar Khot *"ADaPT: As-Needed Decomposition and Planning with Language Models"* In Findings of Conference of the North American Chapter of the Association for Computational Linguistics **(Findings of NAACL 2024)** [PDF]

2024 Elias Stengel-Eskin*, **Archiki Prasad***, Mohit Bansal *"REGAL: Refactoring Programs to Discover Generalizable Abstractions"* In Proceedings of the forty-first International Conference on Machine Learning (**ICML 2024**) [PDF]

2024 Han Wang*, **Archiki Prasad***, Elias Stengel-Eskin*, Mohit Bansal *"Soft Self-Consistency Improves Language Model Agents"* In Proceedings of the Annual Conference of the Association for Computational Linguistics **(ACL 2024)** [PDF]

2024 Justin Chih-Yao Chen, **Archiki Prasad**, Swarnadeep Saha, Elias Stengel-Eskin, Mohit Bansal *"MAGICORE: Multi-Agent, Iterative, Coarse-to-Fine Refinement for Reasoning"* In Proceedings of the Conference on Empirical Methods in Natural Language Processing **(EMNLP 2025)** [PDF]

2023 **Archiki Prasad**, Swarnadeep Saha, Xiang Zhou, Mohit Bansal *"RECEVAL: Evaluating Reasoning Chains via Correctness and Informativeness"* In Proceedings of Conference on Empirical Methods in Natural Language Processing **(EMNLP 2023)** [PDF]

2023 **Archiki Prasad**, Trung Bui, Seunghyun Yoon, Hanieh Deilamsalehy, Franck Dernoncourt, Mohit Bansal *"MEETINGQA: Extractive Question-Answering on Meeting Transcripts"* In Proceedings of the Annual Conference of the Association for Computational Linguistics **(ACL 2023)** [PDF]

2023 **Archiki Prasad**, Peter Hase, Xiang Zhou, Mohit Bansal *"GRIPS: Gradient-free, Edit-based Instruction Search for Prompting Large Language Models"* In Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics **(EACL 2023)** [PDF]

2021 **Archiki Prasad***, Mohammad Ali Rehan*, Shreya Pathak*, Preethi Jyothi *"The Effectiveness of Intermediate-Task Training for Code-Switched Natural Language Understanding"* In Proceedings of the Workshop on Multilingual Representation Learning **(MRL 2021) at EMNLP 2021** [PDF] (Best Paper Honorable Mention)

2021 **Archiki Prasad**, Preethi Jyothi, Rajbabu Velmurugan *"An Investigation of End-to-End Models for Robust Speech Recognition"* In Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing **(ICASSP 2021)** [PDF]

2021 **Archiki Prasad**, Vishal Jain, Sharayu Moharir *"Decentralized Age-of-Information Bandits"* In Proceedings of the IEEE Wireless Communications and Networking Conference **(WCNC 2021)** [PDF]

2020 **Archiki Prasad**, Preethi Jyothi *"How Accents Confound: Probing for Accent Information in End-to-End Speech Recognition Systems"* In Proceedings of the 2020 Annual Conference of the Association for Computational Linguistics **(ACL 2020)** [PDF]

2020 Ayush Chauhan, **Archiki Prasad**, Parth Gupta, Amiredddy Prashanth Reddy, Shiv Kumar Saini *"Time Series Forecasting for Cold-Start Items by Learning from Related Items using Memory Networks"* In Companion Proceedings of the Web Conference 2020 **(WWW 2020)** [PDF]

Honors and Awards

- Recipient of the **2025 Apple Scholars in AI/ML PhD fellowship**, awarded to 21 students worldwide, with full funding for up to 2 years
- IIT Bombay Institute Academic Prize for outstanding performance in the academic year 2019-20
- Secured an all India rank of 144 in JEE-Main 2016 out of roughly 10 million applicants.
- Amongst top 1.2% of all selected candidates (200,000) JEE-Advance 2016.

Invited Talks

Microsoft Turing Speaker Series — Spring 2024
"As-Needed Decomposition and Planning with Language Models" [slides]

Apple LLM Reading Group — Fall 2025
"Less Supervision, Better Reasoning: Bootstrapping from Self-Generated Signals" [slides]

Professional Services

Conference Reviewer

- ICLR 2024-2026, ICML, COLM, NeurIPS, AAAI 2025
- ACL 2022-2025, EMNLP 2021-2025, NAACL 2022-2025 (ACL Rolling Review)

Journal Reviewer

- TACL 2025

References

- Mohit Bansal, John R. Louise S. Parker Professor of CS, UNC Chapel Hill.
- Jason Weston, Senior Director Research Scientist, FAIR, Meta, NYC.
- Tushar Khot, Member of Technical Staff, Microsoft AI, Seattle.