

# Dhananjay Ashok

PhD Candidate at the University of Southern California

[dhananjay.ashok99@gmail.com](mailto:dhananjay.ashok99@gmail.com) | [LinkedIn](#), [Website](#), [Google Scholar](#) | (347) 723 2785

**FOCUS:** Self Improving Language-Model Based Agents in Multimodal Environments. Approaches include RL-tuning, synthetic data generation, retrieval augmented generation and curiosity-based exploration

## EDUCATION

**Ph.D. in Computer Science,**  
**University of Southern California**  
(Ongoing)

Self Improvement of Multimodal Agents in Complex Environments

Advisors: [Jonathan May](#), [Jesse Thomason](#)

**M.Sc. in Machine Learning,**  
**Carnegie Mellon University**  
Distribution Shift and LLMs for

Science with: [Zack Lipton](#),

[Barnabas Poczos](#)

**B.Sc. CS and Econ, University of Toronto**

Robotic Control and Neurosymbolic-AI with:

[Animesh Garg](#), [Vijay Ganesh](#)

## AWARDS

- Annenberg Fellowship, USC
- Valerie Brooks Scholarship
- William Kingston Scholarship

## SKILLS

- Algorithms, Data Structures
- Python, C/C++, Bash,
- PyTorch, TensorFlow, Deep Learning
- HuggingFace, Accelerate, DeepSpeed, Natural Language Processing
- Multi-GPU Parallelization and Quantization of LLMs
- Fine-tuning + LoRA, Tuning LLMs via Reinforcement Learning (GRPO, PPO)
- Independent research

## DEBATE

First speaker from a developing country to be judged [Best Speaker](#) at the World School Debating Championship

## INDUSTRY EXPERIENCE

**Applied Research Intern, Capital One** (Upcoming: June 2026 - Aug 2026)

- Investigating approaches to Language Model-based World Models
- Creating a hybrid approach to improve web, coding and embodied agents

**Applied Science Intern, Amazon Core Search** (May 2025 - Aug 2025)

- Devised a state-of-the-art zero-shot dense retrieval algorithm
- Applied method to internal data, operating at an Amazon Marketplace scale

**Machine Learning Research Engineer, Apple Inc.** (May 2023 - Aug 2023)

- Developed systems for automated understanding and processing of log files
- Implemented MultiAgent RL Solutions to 6G Cellular Networking Problems

## RESEARCH EXPERIENCE

- **GLAMOR LAB**, Prof. Jesse Thomason (Jan 2025-Current)  
Researching Self Improvement and Multimodal Grounding of Multimodal Language Model-based Agents, focusing on Curiosity Based Exploration

- **CUTELABNAME**, Prof. Jonathan May (May 2024-Current)  
Investigating [Factual Grounding](#) of Language Models with an emphasis on Synthetic Data Generation and Retrieval Augmented Generation

- **AutonLab**, Prof. Barnabas Poczos (Sep 2022- Dec 2024)  
Researched [Scientific Error Correction](#), developing a method that outperformed GPT3 despite having only 0.1% as many parameters

**ACMI Lab**, Prof. Zachary Chase Lipton (Sep 2022 - Dec 2024)

- Created a State-of-the-art [Few Shot NER](#) System using LLMs
- Developed a principled [Distribution Shift](#) detection and mitigation method

**Vector Institute**, Prof. Animesh Garg (Sep 2019 - May 2022)

- Applied methods from [causal discovery](#) for [robotic manipulation and control](#)

## SELECTED FIRST AUTHOR PUBLICATIONS

[A Little Human Data Goes A Long Way](#): ACL 2025

- Demonstrated that performance declines associated with replacing human generated data with synthetic data is most chronic only after crossing 90% replacement.
- Showed that the best way to use synthetic data is in conjunction with humans

[Language Models Can Predict Their Own Behavior](#): NeurIPS 2025

- Established that the internal states of LLMs can robustly predict how they will behave on particular inputs and developed an algorithm to extract precise signals.
- Used these signals to construct precise and trustworthy early warning system for jailbreaking, alignment failures, low confidence responses, reasoning gaps etc.

[Can VLMs Recall Factual Associations From Visual References?](#) EMNLP 2025

- Curated a controlled benchmark to isolate and establish the failure of Vision Language Models to recall factual information from visual representations.
- Created a diagnostic system to alert users in cases where the VLM has failed to resolve visual entities, leading to a 64% performance improvement

## PUBLICATIONS

- Ashok, Dhananjay, Ruth-Ann Armstrong, and Jonathan May. "Seamless Deception: Larger Language Models Are Better Knowledge Concealers." *arXiv preprint arXiv:2603.14672* (2026).
- Ashok, Dhananjay, et al. "A Representation Sharpening Framework for Zero Shot Dense Retrieval." *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2026.
- Ashok, Dhananjay, et al. "Can VLMs Recall Factual Associations From Visual References?." *Findings of the Association for Computational Linguistics: EMNLP 2025*. 2025.
- Ashok, Dhananjay, and Jonathan May. "Language Models Can Predict Their Own Behavior." *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Ashok, Dhananjay, and Jonathan May. "A little human data goes a long way." *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 2025.
- Ashok, Dhananjay, and Barnabas Poczos. "Controllable text generation in the instruction-tuning era." *arXiv preprint arXiv:2405.01490* (2024).
- Ashok, Dhananjay, and Zachary C. Lipton. "Promptner: Prompting for named entity recognition." *arXiv preprint arXiv:2305.15444* (2023).
- Ashok, Dhananjay, et al. "The student becomes the master: Outperforming GPT3 on scientific factual error correction." *Findings of the Association for Computational Linguistics: EMNLP 2023*. 2023.
- Barker, Matthew, et al. "Feedbacklogs: Recording and incorporating stakeholder feedback into machine learning pipelines." *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 2023.
- Ashok, Dhananjay, et al. "A solver+ gradient descent training algorithm for deep neural networks." *arXiv preprint arXiv:2207.03264* (2022).
- Ashok, Dhananjay, et al. "Logic guided genetic algorithms (student abstract)." *Proceedings of the AAAI conference on artificial intelligence*. Vol. 35. No. 18. 2021.

## GRANTS AND FELLOWSHIPS

- CoefficientGiving Grant: Awarded for research on Language Model Alignment
- Annenberg Fellowship, University of Southern California
- Valerie Brooks Scholarship, University of Toronto
- William Kingston Scholarship, University of Toronto
- Deans Scholarship, University of Toronto