Pranjal Aggarwal

Pranjal2041 | Website

Education

2024–Present Ph.D. in Language Technologies, School of Computer Science, Carnegie Mellon University
2019–2024 B.Tech. and M.Tech. in Computer Science, Indian Institute of Technology Delhi (**GPA: 9.4, Department Rank: 4**)

Publications

[1] **Pranjal Aggarwal**, Seungone Kim, Jack Lanchantin, Sean Welleck, Jason Weston, Ilia Kulikov, and Swarnadeep Saha. "OptimalThinkingBench: Evaluating Over and Underthinking in LLMs". In: *Proceedings of the **International Conference on Learning Representations (ICLR), 2026***. 2026. Link.

[2] **Pranjal Aggarwal** and Sean Welleck. "Programming with Pixels: Computer-Use Meets Software Engineering". In: *Proceedings of the **International Conference on Learning Representations (ICLR), 2026***. 2026. Link.

[3] Shreyas Chaudhari*, **Pranjal Aggarwal***, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan, Ameet Deshpande, and Bruno Castro da Silva. "RLHF Deciphered: A Critical Analysis of Reinforcement Learning from Human Feedback for LLMs". In: ***ACM Computing Surveys*** 58.2 (2026). Link.

[4] Weihua Du, **Pranjal Aggarwal**, Sean Welleck, and Yiming Yang. "Agentic-R1: Distilled Dual-Strategy Reasoning". In: *Proceedings of the 2025 Conference on **Empirical Methods in Natural Language Processing***. 2025. Link.

[5] **Pranjal Aggarwal** and Sean Welleck. "L1: Controlling How Long A Reasoning Model Thinks With Reinforcement Learning". In: *Proceedings of the **Conference on Language Modeling (COLM), 2025***. 2025. Link.

[6] **Pranjal Aggarwal**, Bryan Parno, and Sean Welleck. "AlphaVerus: Bootstrapping Formally Verified Code Generation through Self-Improving Translation and Treefinement". In: *Proceedings of the 42nd **International Conference on Machine Learning** (Best Theme Track Paper, LTI SRS)*. 2025. Link.

[7] **Pranjal Aggarwal***, Aman Madaan*, Ankit Anand, Srividya Pranavi Potharaju, Swaroop Mishra, Pei Zhou, Aditya Gupta, Dheeraj Rajagopal, Karthik Kappaganthu, Yiming Yang, Shyam Upadhyay, Mausam, and Manaal Faruqui. "AutoMix: Automatically Mixing Language Models". In: *Proceedings of **Neural Information Processing Systems, 2024***. 2024. Link.

[8] **Pranjal Aggarwal**, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan, and Ameet Deshpande. "GEO: Generative Engine Optimization". In: *Proceedings of the 30th ACM SIGKDD Conference on **Knowledge Discovery and Data Mining***. KDD '24. Barcelona, Spain: Association for Computing Machinery, 2024, pp. 5–16. Link.

[9] **Pranjal Aggarwal**, Aman Madaan, Yiming Yang, and Mausam. "Let's Sample Step by Step: Adaptive-Consistency for Efficient Reasoning and Coding with LLMs". In: *Proceedings of the 2023 Conference on **Empirical Methods in Natural Language Processing (Oral)***. Singapore: Association for Computational Linguistics, Dec. 2023, pp. 12375–12396. Link.

[10] **Pranjal Aggarwal**, Ameet Deshpande, and Karthik Narasimhan. "SemSup-XC: Semantic Supervision for Zero and Few-shot Extreme Classification". In: *Proceedings of the 40th **International Conference on Machine Learning***. ICML'23. Honolulu, Hawaii, USA, 2023. Link.

[11] **Pranjal Aggarwal***, Krithika Rangarajan*, Dhruv Kumar Gupta, Rohan Raju Dhanakshirur, Akhil Baby, Chandan Pal, Arunkumar Gupta, Smriti Hari, Subhashis Banerjee, and Chetan Arora. "Deep learning for detection of iso-dense, obscure masses in mammographically dense breasts." In: ***European radiology*** (2023). Link.

** Equal Contribution*

Preprints

[1] Maxwell Jones, **Pranjal Aggarwal**, Lawrence Keunho Jang, Trung Bui, Franck Dernoncourt, Gang Wu, Jun-Yan Zhu, and Ruslan Salakhutdinov. "Image Editing Software as a Tool: MLLM Agents for Precise Image Editing via API-Driven Control". In: *Under Review*. 2026.

[2] **Pranjal Aggarwal**, Marjan Ghazvininejad, Seungone Kim, Ilia Kulikov, Jack Lanchantin, Xian Li, Tianjian Li, Bo Liu, Graham Neubig, Anaelia Ovalle, Swarnadeep Saha, Sainbayar Sukhbaatar, Sean Welleck, Jason Weston, Chenxi Whitehouse, Adina Williams, Jing Xu, Ping Yu, Weizhe Yuan, Jingyu Zhang, and Wenting Zhao. "Reasoning Over Mathematical Objects: On-Policy Reward Modeling and Test Time Aggregation". In: *arXiv preprint* (2026). Link.

[3] Alex Wilf, **Pranjal Aggarwal**, Bryan Parno, Daniel Fried, Louis-Philippe Morency, Paul Pu Liang, and Sean Welleck. "Propose, Solve, Verify: Self-Play Through Formal Verification". In: *arXiv preprint* (2025). Link.

[4] Wenting Zhao, **Pranjal Aggarwal**, Swarnadeep Saha, Asli Celikyilmaz, Jason Weston, and Ilia Kulikov. "The Majority is not always right: RL training for solution aggregation". In: *arXiv preprint* (2025). Link.

Research Interests

Computer-Use Agents, Test-Time Compute for Reasoning, Formal Verification for Code.

Work Experience

Ph.D. Student, *Carnegie Mellon University, Advisor: Prof. Sean Welleck* Aug'24–Present

- **Computer-Use Agents:** Developing more efficient architectures for computer-use agents and scaling training across softwares, environments, inference compute.
- **Reasoning:** Developing methods for controlling and optimizing test-time compute in reasoning LLMs through reinforcement learning.
- **Formal Verification for Code:** Developing methods for generating formally verified code through self-improving translation and self-play guided by verifier feedback.

Research Scientist Intern @ Meta, *Team Lead: Jason Weston, Mentor: Swarnadeep Saha* May'25–Nov'25

- Researched improving reasoning in LLMs, including evaluating optimal test-time compute allocation, RL-trained solution aggregation, and building benchmarks, training data, and verification methods for graduate-level mathematical and scientific reasoning.

Undergraduate Researcher, *Data Analytics & Intelligence Lab, Advisor: Prof. Mausam* Dec'22–Jun'24

- Focused on improving efficiency of LLM reasoning and code generation through adaptive sampling methods and cost-quality optimization via model routing.

NLP Research Intern, *Princeton University, Advisor: Prof. Karthik Narasimhan* Jun'22–Aug'22

- Researched zero-shot extreme classification, optimization for generative search engines, and reinforcement learning from human feedback in LLMs.

Undergraduate Researcher, *CV Lab, IIT Delhi, Advisor: Prof. Chetan Arora* Jul'21–Mar'22

- Developed deep learning methods for cancer detection in mammographically dense breast images.

Awards

SoftBank Group–Arm Fellowship, Carnegie Mellon University Link

- Selected as one of eight CMU students for the SoftBank Group–Arm Fellowship for PhD research.

Winner, Model Attribution Challenge at SaTML 2023 Link

- Developed the best method for attributing blackbox fine-tuned language models to their corresponding pre-trained models; Oral presentation @ SaTML'23.

First Place, Tower Research Capital Data Challenge Link

- Secured 1st rank in machine learning for market data challenge by Tower Research Capital at IIT Delhi.

Google and Microsoft Research Travel Grants

- Awarded research travel grants from Microsoft and Google Research for academic conferences.

Winner, HackMIT Sponsor Track Link

- Won MIT's HackMIT 2020 in the vehicle crash detection challenge by Cambridge Mobile Telematics.

Press Coverage

Research featured in *The New York Times*, *Wired*, *The Observer*, *VentureBeat*, *Inc.*, *a16z*, *Entrepreneur*, *Digiday*, *Search Engine Land*, *MarkTechPost*, *Interconnects*, and others. *(One representative link per outlet.)*

Academic Honors

- All India Rank 143 in IIT JEE-Advanced 2019 among 1.2 million candidates nationwide.
- Selected among top 50 students nationwide for the Orientation-Cum-Selection Camp (OCSC) for the International Olympiad in Astronomy and Astrophysics.
- Received the IIT Delhi Semester Merit Award (top 7%) in 4 out of 8 semesters.
- Ranked 303rd in KVPY'19, a prestigious science examination by the Indian Institute of Science.

Academic Service

Reviewer: Multiple editions of ACL, EMNLP, ICML, ICLR, NeurIPS, KDD, ECCV
Outstanding Teaching Assistant Award, Natural Language Processing (Graduate Level), Spring 2023–2024
Teaching Assistant, Principles of Artificial Intelligence (Graduate Level), Fall 2023–2024
Academic Mentor, Electrodynamics and Quantum Mechanics, Fall 2020–2021
Volunteer at ACL 2023, ICML 2023, EMNLP 2023