Arushi Rai

arushirai1999@gmail.com arushirai1.github.io arushirai1 arushi-r-568887b3

Education

University of Pittsburgh August 2021 - November 2026 (expected)
Doctorate in Computer Science *GPA: 3.77*
Dissertation: Language Provides Structure for Weakly-Supervised Localization in Object Detection and Sports Coaching
Advisor: Dr. Adriana Kovashka

Illinois Institute of Technology July 2017 - Dec 2020
Bachelor of Science in Computer Science *GPA: 3.76*
Honors: Magna Cum Laude

Experience

Research Scientist April 2026 – Present
Meta

- Developing LLM-based agents for recommender systems and improving inference efficiency of multimodal LLMs and embedding models through speculative decoding.

Graduate Research Assistant August 2021 - Present
University of Pittsburgh

- Dissertation research on using language as structure for weakly-supervised localization, spanning object detection from noisy in-the-wild captions and video understanding for sports coaching (temporal grounding, feedback generation, and figure skating scoring). See publications below for details.

Software Engineering Intern (PhD) May 2025 – October 2025
Meta – Menlo Park, USA

- Designed an efficient test-time alignment method that improves LLM reasoning and adds reasoning capability to previously non-reasoning models, such as backtracking and verification.
- Built a task-adaptive MoE-inspired router for test-time alignment; published in Findings of ACL 2026.

Software Engineering Intern (PhD) May 2024 – August 2024
Meta – Menlo Park, USA

- Built large-scale paired-multimodal dataset (1B+ items) using engagement-based heuristics for semantic similarity, with scalable noise-filtering methods.
- Improved fine-grained representation of vision-language models using contrastive learning, achieving 1.7% absolute gain on fine-grained topic classification and emergent super fine-grained semantic similarities recognition capability on multimodal retrieval.

Machine Learning Engineer Intern (Computer Vision) May 2022 - August 2022
Gather.ai - Pittsburgh, USA

- Improved case counting pipeline by 13% on difficult, irregular cases by targeting intermediate stack segmentation and mean-shift clustering.
- Implemented **instance segmentation approaches** like Panoptic DeepLab and increased robustness to viewpoint shifts through perspective augmentations and fusing depth cues with RGB images.

Computer Vision Research Intern (REU) May 2020 - August 2020
University of Central Florida — Center for Research in Computer Vision *Link to Report*

- Extended self-supervised learning frameworks such as DeepCluster and XDC in Pytorch to handle multimodal data from HowTo100M like video clips and narrations to achieve 5% improvement on long-complex activity recognition.

Peer-Reviewed Publications

Conference

[1] **Arushi Rai** and Adriana Kovashka, "Learning Consistent Temporal Grounding between Related Tasks in Sports Coaching," European Conference on Computer Vision (ECCV), Sep 2026. [pdf]

[2] **Arushi Rai**, Qiang Zhang, Hanqing Zeng, Yunkai Zhang, Dipesh Tamboli, Xiangjun Fan, Zhuokai Zhao, and Lizhu Zhang, "TARo: Token-level Adaptive Routing for LLM Test-time Alignment," Findings of the Association for Computational Linguistics: ACL 2026, July 2026. [pdf]

[3] **Arushi Rai** and Adriana Kovashka, "Generalizing Sports Feedback Generation by Watching Competitions and Reading Books: A Rock Climbing Case Study," IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), March 2026. [pdf]

[4] **Arushi Rai** and Adriana Kovashka, "Rubric-Constrained Figure Skating Scoring," IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), March 2025. Also presented as an oral at the 3rd Workshop on Computer Vision for Winter Sports (CV4WS). [pdf]

[5] **Arushi Rai** and Adriana Kovashka, "VEIL: Vetting Extracted Image Labels from In-the-Wild Captions for Weakly-Supervised Object Detection," The 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL), March 2024. [pdf]

Workshop - Archival, in CVPR Workshops Proceedings

[1] **Arushi Rai***, Kyle Buettner*, and Adriana Kovashka, "Strategies to Leverage Foundational Model Knowledge in Object Affordance Grounding," MMFM2: The 2nd Workshop on What is Next in Multimodal Foundation Models?, held in conjunction with the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2024. [pdf] (*=equal contribution)

[2] **Arushi Rai** and Adriana Kovashka, "Improving Language-Supervised Object Detection With Linguistic Structure Analysis," Workshop on Open-Domain Reasoning Under Multi-Modal Settings (O-DRUM), held in conjunction with the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2023. [pdf]

Talks, Leadership, Service, Teaching Experience

Oral Presentation – Rubric-Constrained Figure Skating Scoring March 2025
3rd Workshop on Computer Vision for Winter Sports (CV4WS) at WACV 2025

Teaching Assistant – University of Pittsburgh 2022 – 2024
Intro to Computer Vision (CS 1674/2074, Fall 2024); Operating Systems (CS 1550, Spring 2023); Systems Software (CS 0449, Fall 2022 & Fall 2023); Algorithms and Data Structures 2 (CS 1501, Fall 2023)

Conference Reviewer – CVPR 2024–2026, ECCV 2024, 2026, AAAI 2024, 2026, BMVC 2026, The Web Conference 2027 2024 – Present
Including Emergency Reviewer for CVPR 2025

Workshop Reviewer – CV4WS @ WACV 2026, GRAIL-V @ CVPR 2026, SURGeLLM @ ACL 2026, DemoDiv @ CVPR 2025 2025 – 2026
Workshops on Winter Sports, Video, LLMs, and Demos

Computer Vision Interest Group Moderator 2021-2022
Machine Learning Collective

ACM@IIT President 2020
Association for Computing Machinery Chapter