Tejas Gokhale

CONTACT
Email: gokhale@umbc.edu   Website: https://www.tejasgokhale.com
Mail: 1000 Hilltop Circle, ITE 214, Baltimore MD 21250

CURRENT
POSITION

Assistant Professor
Department of Computer Science & Electrical Engineering
University of Maryland, Baltimore County

Affiliated Faculty, UMBC AI Center

Director, Cognitive Vision Group

EDUCATION

Doctor of Philosophy, Arizona State University   2023
School of Computing and Augmented Intelligence
Advisors: Yezhou Yang, Chitta Baral
Thesis: Towards Reliable Semantic Vision

Master of Science, Carnegie Mellon University   2017
Department of Electrical and Computer Engineering
Mentor: Aswin Sankaranarayanan

Bachelor of Engineering (Honours), BITS Pilani   2015
Department of Electrical and Electronics Engineering

EMPLOYMENT
HISTORY

Microsoft Research   Summer 2022
Research Intern, Adaptive Systems and Interaction Group
Mentors: Hamid Palangi, Besa Nushi, Vibhav Vineet, Eric Horvitz

Lawrence Livermore National Laboratory   Summer 2021, 2020
Research Scholar, Machine Intelligence Group
Mentors: Rushil Anirudh, Jay Thiagarajan, Bhavya Kailkhura

Arizona State University
Graduate Research Associate, School of Computing and AI   2018–2023
Graduate Teaching Associate, School of Computing and AI   2018–2020

Snap Research   Summer 2018
Research Intern, Computational Imaging Group
Mentors: Guru Krishnan, Shree Nayar

Carnegie Mellon University   2017–2018
Graduate Student Researcher, Dept. of Electrical and Computer Engineering

ST Microelectronics   Fall 2014
Intern, High Speed Links Group

Steel Authority of India Limited   Summer 2013
Summer Intern, Bhilai Steel Plant

PUBLICATIONS   See my Google Scholar page for recent updates and citation information.
Legend: My graduate advisees are underlined.
Conference proceedings are the de facto form of publication in computer vision, machine learning, natural language processing, and AI. Journals are (to put it mildly) obsolete.

Conference Proceedings

[C1] Nilay Yilmaz, Maitreya Patel, Yiran Lawrence Luo, Tejas Gokhale, Chitta Baral, Suren Jayasuriya, Yezhou Yang. Voila: Evaluation of MLLMs For Perceptual Understanding and Analogical Reasoning. In International Conference on Learning Representations. 2025.
https://arxiv.org/abs/2503.00043   ICLR 2025

[C2] Sourajit Saha, Tejas Gokhale. Improving Shift Invariance in Convolutional Neural Networks with Translation Invariant Polyphase Sampling. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. 2025.
https://arxiv.org/abs/2404.07410   WACV 2025

[C3] Maitreya Patel, Abhiram Kusumba, Sheng Cheng, Changhoon Kim, Tejas Gokhale, Chitta Baral, Yezhou Yang. TripletCLIP: Improving Compositional Reasoning of CLIP via Vision-Language Negatives. In Advances in Neural Information Processing Systems. 2024.
https://arxiv.org/abs/2411.02545   NeurIPS 2024

[C4] Agneet Chatterjee, Gabriela Ben Melech Stan, Estelle Guez Aflalo, Sayak Paul, Dhruba Ghosh, Tejas Gokhale, Ludwig Schmidt, Hannaneh Hajishirzi, Vasudev Lal, Chitta Baral, Yezhou Yang. Getting it Right: Improving Spatial Consistency in Text-to-Image Models. In European conference on computer vision. 2024.
https://arxiv.org/abs/2404.01197   ECCV 2024

[C5] Agneet Chatterjee, Yiran Luo, Tejas Gokhale, Chitta Baral, Yezhou Yang. REVISION: Rendering Tools Enable Spatial Fidelity in Vision-Language Models. In European conference on computer vision. 2024.
https://arxiv.org/abs/2408.02231   ECCV 2024

[C6] Agneet Chatterjee, Tejas Gokhale, Chitta Baral, Yezhou Yang. On the Robustness of Language Guidance for Low-Level Vision Tasks: Findings from Depth Estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2794-2803. 2024.
https://arxiv.org/abs/2404.08540   CVPR 2024

[C7] Maitreya Patel, Tejas Gokhale, Chitta Baral, Yezhou Yang. ConceptBed: Evaluating Concept Learning Abilities of Text-to-Image Diffusion Models. In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, no. 13, pp. 14554-14562. 2024.
https://arxiv.org/abs/2306.04695   AAAI 2024

[C8] Sheng Cheng, Tejas Gokhale, Yezhou Yang. Adversarial Bayesian Augmentation for Single-Source Domain Generalization. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 11400-11410. 2023.
https://arxiv.org/abs/2307.09520   ICCV 2023

[C9] Man Luo, Zhiyuan Fang, Tejas Gokhale, Yezhou Yang, Chitta Baral. End-to-end Knowledge Retrieval for Multi-modal Queries. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 8573-8589. 2023.
https://arxiv.org/abs/2306.00424   ACL 2023

[C10] Tejas Gokhale, Rushil Anirudh, Jayaraman J. Thiagarajan, Bhavya Kailkhura, Chitta Baral, and Yezhou Yang. Improving Diversity with Adversarially Learned Transformations for Domain Generalization. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 434-443. 2023.
https://arxiv.org/abs/2206.07736   WACV 2023

[C11] Maitreya Patel, Tejas Gokhale, Chitta Baral, and Yezhou Yang. CRIPP-VQA: Counterfactual Reasoning about Implicit Physical Properties via Video Question Answering. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, pp. 9856-9870. 2022.
https://arxiv.org/abs/2211.03779   EMNLP 2022

[C12] Tejas Gokhale, Abhishek Chaudhary, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. Semantically Distributed Robust Optimization for Vision-and-Language Inference. In Findings of the Association for Computational Linguistics: ACL 2022, pp. 1493-1513. 2022.
https://arxiv.org/abs/2110.07165   ACL Findings 2022

[C13] Tejas Gokhale, Swaroop Mishra, Man Luo, Bhavdeep Sachdeva, and Chitta Baral. Generalized but not Robust? Comparing the Effects of Data Modification Methods on Out-of-Domain Generalization and Adversarial Robustness. In Findings of the Association for Computational Linguistics: ACL 2022, pp. 2705-2718. 2022.
https://arxiv.org/abs/2203.07653   ACL Findings 2022

[C14] Neeraj Varshney, Pratyay Banerjee, Tejas Gokhale, and Chitta Baral. Unsupervised Natural Language Inference Using PHL Triplet Generation. In Findings of the Association for Computational Linguistics: ACL 2022, pp. 2003-2016. 2022.
https://arxiv.org/abs/2110.08438   ACL Findings 2022

[C15] Yiran Luo, Pratyay Banerjee, Tejas Gokhale, Yezhou Yang, and Chitta Baral. To Find Waldo You Need Contextual Cues: Debiasing Who's Waldo. In 60th Annual Meeting of the Association for Computational Linguistics, ACL 2022, pp. 355-361. 2022.
https://arxiv.org/abs/2203.16682   ACL 2022

[C16] Man Luo, Arindam Mitra, Tejas Gokhale, and Chitta Baral. Improving biomedical information retrieval with neural retrievers. In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, no. 10, pp. 11038-11046. 2022.
https://arxiv.org/abs/2201.07745   AAAI 2022

[C17] Pratyay Banerjee, Tejas Gokhale, Yezhou Yang, and Chitta Baral. Weakly supervised relative spatial reasoning for visual question answering. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 1908-1918. 2021.
https://arxiv.org/abs/2109.01934   ICCV 2021

[C18] Pratyay Banerjee, Tejas Gokhale, Yezhou Yang, and Chitta Baral. WeaQA: Weak Supervision via Captions for Visual Question Answering. In Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, pp. 3420-3435. 2021.
https://arxiv.org/abs/2012.02356   ACL Findings 2021

[C19] Pratyay Banerjee, Tejas Gokhale, and Chitta Baral. Self-Supervised Test-Time Learning for Reading Comprehension. In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 1200-1211. 2021.
https://arxiv.org/abs/2103.11263   NAACL 2021

[C20] Tejas Gokhale, Rushil Anirudh, Bhavya Kailkhura, Jayaraman J. Thiagarajan, Chitta Baral, and

Yezhou Yang. Attribute-guided adversarial training for robustness to natural perturbations. In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, no. 9, pp. 7574-7582. 2021.
https://arxiv.org/abs/2012.01806   AAAI 2021

[C21] Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. MUTANT: A Training Paradigm for Out-of-Distribution Generalization in Visual Question Answering. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 878-892. 2020.
https://arxiv.org/abs/2009.08566   EMNLP 2020

[C22] Zhiyuan Fang, Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. Video2 Commonsense: Generating Commonsense Descriptions to Enrich Video Captioning. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 840-860. 2020.
https://arxiv.org/abs/2003.05162   EMNLP 2020

[C23] Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. Vqa-lol: Visual question answering under the lens of logic. In European conference on computer vision, pp. 379-396. Cham: Springer International Publishing, 2020.
https://arxiv.org/abs/2002.08325   ECCV 2020

Journals and Magazines

[J1] Tejas Gokhale. Towards Robust Visual Understanding: A Paradigm Shift in Computer Vision from Recognition to Reasoning. AI Magazine 1–7. 2024.
https://doi.org/10.1002/aaai.12194   AI Magazine

Peer Reviewed Workshop Papers and Extended Abstracts

[W1] Yiran Luo, Joshua Feinglass, Tejas Gokhale, Kuan-Cheng Lee, Chitta Baral, Yezhou Yang. Grounding Stylistic Domain Generalization with Quantitative Domain Shift Measures and Synthetic Scene Images. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. 2024.
https://arxiv.org/abs/2405.15961   CVPR 2024 Vision Datasets Understanding Workshop

[W2] Tejas Gokhale. Towards Robust Visual Understanding: from Recognition to Reasoning. In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, no. 20, pp. 22665-22665. 2024.
https://ojs.aaai.org/index.php/AAAI/article/view/30281   AAAI New Faculty Highlights

[W3] Tejas Gokhale, Joshua Feinglass, and Yezhou Yang. Covariate Shift Detection via Domain Interpolation Sensitivity. In First Workshop on Interpolation Regularizers and Beyond at NeurIPS 2022.
https://openreview.net/pdf?id=YkPjTHZDdm   NeurIPS 2022 Interpolation Workshop

[W4] Kuldeep Kulkarni, Tejas Gokhale, Rajhans Singh, Pavan Turaga, Aswin C. Sankaranarayanan. Halluci-Net: Scene Completion by Exploiting Object Co-occurrence Relationships. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. 2021.
https://arxiv.org/abs/2004.08614   CVPR 2021 AI for Content Creation Workshop

[W5] Tejas Gokhale, Shailaja Sampat, Zhiyuan Fang, Yezhou Yang, and Chitta Baral. Cooking with blocks: A recipe for visual reasoning on image-pairs. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 5-8. 2019.
https://arxiv.org/abs/1905.12042   CVPR 2019 Vision Meets Cognition Workshop

[W6] Tejas Gokhale. Vision beyond Pixels: Visual Reasoning via Blocksworld Abstractions. In IJCAI, pp. 6436-6437. 2019.
https://www.ijcai.org/Proceedings/2019/0907.pdf   IJCAI 2019 Doctoral Consortium

Technical Reports and Preprints

[P1] Maitreya Patel, Neeraj Varshney, Agneet Chatterjee, Tejas Gokhale, Yezhou Yang, Chitta Baral. Reliability-Checklist: Framework for Comprehensively Evaluating the Reliability of NLP Systems.
https://github.com/Maitreyapatel/reliability-checklist   Tech Report

[P2] Tejas Gokhale, Hamid Palangi, Besmira Nushi, Vibhav Vineet, Eric Horvitz, Ece Kamar, Chitta Baral, and Yezhou Yang. Benchmarking Spatial Relationships in Text-to-Image Generation. 2022.
https://arxiv.org/abs/2212.10015   Tech Report

[P3] Ethan Wisdom, Tejas Gokhale, Chaowei Xiao, and Yezhou Yang. Mole Recruitment: Poisoning of Image Classifiers via Selective Batch Sampling.
https://arxiv.org/abs/2303.17080   Tech Report

Books and Monographs

[B1] *Advances in Multimodal Information Retrieval and Generation*   Springer
Synthesis Lectures on Computer Vision,   ISBN: 978-3-031-57815-1
Man Luo, Tejas Gokhale, Neeraj Varshney, Yezhou Yang, Chitta Baral.   [website]

Chapters:

- Man Luo, Tejas Gokhale, Neeraj Varshney, Yezhou Yang, and Chitta Baral. "Transformer-Driven Models for Language, Vision, and Multimodality." In Advances in Multimodal Information Retrieval and Generation, pp. 11-34. Cham: Springer International Publishing, 2024.
- Man Luo, Tejas Gokhale, Neeraj Varshney, Yezhou Yang, and Chitta Baral. "Multimodal Information Retrieval." In Advances in Multimodal Information Retrieval and Generation, pp. 35-91. Cham: Springer International Publishing, 2024.
- Man Luo, Tejas Gokhale, Neeraj Varshney, Yezhou Yang, and Chitta Baral. "Multimodal Content Generation." In Advances in Multimodal Information Retrieval and Generation, pp. 93-134. Cham: Springer International Publishing, 2024.
- Man Luo, Tejas Gokhale, Neeraj Varshney, Yezhou Yang, and Chitta Baral. "Retrieval Augmented Modeling." In Advances in Multimodal Information Retrieval and Generation, pp. 135-157. Cham: Springer International Publishing, 2024.

Ph.D. Dissertation

[T1] Tejas Gokhale. 2023. Towards Reliable Semantic Vision. Order No. 30426752, Arizona State University. https://www.proquest.com/docview/2813822780

Intellectual Property

[IP1] *Automated Evaluation of Spatial Relationships in Images* [url]   (US Patent App. 18/198,593)

[IP2] *Systems, Methods, and Apparatuses for Implementing Improved Diversity using Adversarially Learned Transformations for Domain Generalization* [url]   (US Patent App. 63/468,653)

"Modular Reasoning using Hybrid Inferential Formalisms"	2024-27
UMBC Center & Institute Departmentally-Engaged Research (CIDER).	($50K)
"Identification of Virga Precipitation Events"	2025-26
UMBC Cybersecurity Institute.	($45K)
Cybersecurity Graduate Fellows Program	2025
UMBC Strategic Awards for Research Transitions (START)	($25K)
"A Framework for Quantifying Typicality of AI-Generated Images"	2024-25
Maryland Procurement Office (via Johns Hopkins University)	(∼$29K)
"Modular Natural Language Understanding" (PI: Frank Ferraro)	2024
UMBC Summer Research Faculty Fellowship (SURFF)	($8K)
"Improving the Continual Learning Ability of Visual Recognition Systems via Targeted Unlearning"	2024
Microsoft Research Accelerate Foundation Models Academic Research	($20K)
Cloud Computing and OpenAI Credits	2024
Google Cloud	(∼$2.5K)
Education Credits	2023-24

TEACHING

Instructor, UMBC

CMSC 472/672 Computer Vision [website]	Fall 2025
CMSC 475/675 Neural Networks [website]	Spring 2025
CMSC 491/691 Robust Machine Learning [website]	Fall 2024
CMSC 491/691 Computer Vision [website]	Spring 2024, Fall 2023
CMSC 898 Pre-Doctoral Candidacy Research	Spring 2025, Fall 2024, Spring 2024
CMSC 799 Master's Thesis Research	Spring 2025, Fall 2024
CMSC 699 Independent Study	Spring 2025, Fall 2024, Spring 2024, Fall 2023
CMSC 499 Independent Study	Spring 2025

Graduate Teaching Associate, Arizona State University

CSE310: Data Structures & Algorithms	Spring 2020
CSE408: Multimedia Information Systems	Spring 2019
CSE110: Introduction to Programming,	Fall 2018

Guest Lecturer, Arizona State University

CSE598, Perception in Robotics	Spring 2022
CSE408, Multimedia Information Systems	Spring 2019

Student Instructor, BITS Pilani Goa Campus

CTE: Advanced Image Processing	Spring 2015

PRESENTATIONS

(Invited Talk), UMBC COEIT Research Day	04/2025
"Cognitive Vision: Concepts, Contexts, and Semantics"	
(Invited Talk), UMBC Information Systems Seminar	11/2024
"Cognitive Vision: Concepts, Contexts, and Semantics"	
(Tutorial), European Conference on Computer Vision	10/2024
"Evaluation and Benchmarking for Text-to-Image Models"	

(Lightning Talk), IARPA Video-LINCS Proposers Day   02/2024
    "Robust Visual Understanding: Knowledge-Guided and Multimodal Reasoning"

(Tutorial), Winter Conference on Applications of Computer Vision   01/2024
    "Challenges with Evaluation of Text-to-Image Models"   [website]

(Invited Talk), PRG Seminar, UMIACS (University of Maryland)   11/2023
    "Robust Visual Understanding in the Multimodal Era"

(Invited Talk) *"Towards Reliable Semantic Vision"*   Spring 2023
    Rochester Institute of Technology (02/23), Binghamton University (03/23),
    University of Maryland Baltimore County (03/23), Indiana University (03/23),
    Case Western Reserve University (03/23), Colorado School of Mines (03/23),
    Temple University (04/2023)

(Tutorial), Winter Conference on Applications of Computer Vision   01/2023
    "Semantic Data Engineering for Robustness Under Multimodal Settings" [website]

(Invited Talk) University of Illinois at Chicago   10/2022
    "Robust Semantic Vision"

(Invited Talk) Microsoft Research   10/2022
    "Benchmarking Spatial Relationships in Text-to-Image Generation"

(Doctoral Consortium) CVPR, New Orleans   06/2022
    "Discovering Transformations for Generalization in Semantic Vision"

(Guest Lecture) Arizona State University CSE 598   03/2022
    "Introduction to Generalization in Semantic Vision"

(Invited Talk) Arizona State University ML Club   09/2021
    "Robust Visual Understanding"

(Doctoral Consortium), IJCAI, Macao   08/2019
    "Vision Beyond Pixels"

(Tutorial) Telluride Neuromorphic Cognition Engineering Workshop,   07/2019
    "Reasoning about Objects and Actions via Block-Play"   [website]

(Invited) Birla Institute of Technology and Science (BITS Pilani)   04/2018
"Deep Learning Methods in Imaging and Computer Vision"

STUDENTS   **PhD**
                • Sourajit Saha   Ph.D. CS [current], UMBC
                • Zhiwei Zhang   Ph.D. CS [current], UMBC
                • Shivanand Kundargi   Ph.D. CS [current], UMBC
                • Jordan Turley   Ph.D. CS [current], UMBC
                • Dylan Lang   Ph.D. CS [current], UMBC

                PhD (as Committee Member)
                • Sheng Cheng (advisor: Yezhou Yang)   Ph.D. CS 2025, ASU [dissertation]
                • Mark Jarzynski (advisor: Marc Olano)   Ph.D. CS [current], UMBC
                • Yiran Luo (advisor: Chitta Baral and Yezhou Yang)   Ph.D. CS [current], ASU

MS Thesis (as Committee Member)
- Naomi Angela Tack (advisor: Don Engel)   M.S. CS 2024, UMBC

Other MS/PhD
- Independent Study: Neel Patel   M.S. CS [current], UMBC
- Independent Study: Shaswati Saha   Ph.D. CS [current], UMBC
- Independent Study: Varun Magotra   M.S. CS [current], UMBC

Undergraduate
- Independent Study: Nicholas Harrell   Spring 2025
- Independent Study: Alexander Shaner   Spring 2025
- Visitors: Tetevi Wilson, Dhanush Bharadwaj   2024-2025
- Visitors: Joey Mule, Luke Parrish   2023-2024
- UMBC CWIT Scholar: Chloe Wood   2024-25
- UMBC CWIT Scholar: Danielle Burton   2023-24

Ph.D. Mentees (at ASU)
- Maitreya Patel   Ph.D. CS [current], ASU
- Agneet Chatterjee   Ph.D. CS [current], ASU
- Nilay Yilmaz   Ph.D. CS [current], ASU

MS (Thesis) Mentees (at ASU)
- Maitreya Patel (see publication [C11])   M.S. CS 2022, ASU [thesis]
- Abhishek Chaudhary (see publication [C12])   M.S. CS 2021, ASU [thesis]

Undergraduate Mentees (at ASU)
- ASU FURI Program: Mertay Dayanc   B.S CS, 2020
- BS CS Capstone: Paul Butler, Jace Lord, Aashwin Ranjan, Sagarika Pannase, William Tith   2019-20

ACADEMIC
SERVICE

National Science Foundation
- Reviewer, GRFP   2025
- Panel, IIS/III   2025

Tutorial Chair, International Conference on Computer Vision (ICCV)   2025

Session Chair
- Winter Applications of Computer Vision (WACV)   2025

Area Chair / Action Editor
- International Conference on Computer Vision (ICCV)   2025
- Advances in Neural Information Processing Systems (NeurIPS)   2024-25
- Winter Applications of Computer Vision (WACV)   2025
- Association for Computational Linguistics (ACL)   2024
- North American Chapter of the ACL (NAACL)   2024
- Empirical Methods in Natural Language Processing (EMNLP)   2024

Reviewer / Program Committee:
- Conference on Computer Vision and Pattern Recognition (CVPR)   2023-25
- International Conference on Computer Vision (ICCV)   2023
- European Conference on Computer Vision (ECCV)   2022-24
- Winter Conference on Applications of Computer Vision (WACV)   2021–24
- International Conference on Machine Learning (ICML)   2023–25
- Advances in Neural Information Processing Systems (NeurIPS)   2022–24
- International Conference on Learning Representations (ICLR)   2022–25
- AAAI Conference on Artificial Intelligence (AAAI)   2021–24
- Conference on Language Models (COLM)   2024
- Association for Computational Linguistics (ACL)   2021–24
- Empirical Methods in Natural Language Processing (EMNLP)   2021–23
- North American Chapter of the ACL (NAACL)   2021–23
- International Conference on Robotics and Automation (ICRA)   2019–2023
- International Conference on Intelligent Robots and Systems (IROS)   2022
- IEEE Robotics and Automation Letter (RA-L)   2020-24
- IEEE Transactions of Pattern Analysis and Machine Intelligence (T-PAMI)   2024-25
- ACM Transactions of Computing for Healthcare   2024
- ACM Computing Surveys   2024
- Springer Machine Vision and Applications (MVAP)   2020
- Springer Book Proposals Reviewer   2024
- Mentor, Undergraduate Student Consortium (AAAI-UC)   AAAI 2024
- Best Student Abstract Award Committee, AAAI   2024

Leadership:
- Director, Cognitive Vision Group   UMBC
- Team Lead, Summer Camp for Applied Language Exploration (SCALE) 2024, JHU Human Language Technology Center of Excellence [Website]   Summer 2024
- Organizer, Tutorial on Responsibly Building Generative Models [Website]   ECCV'24
- Organizer, Tutorial on Reliability of Generative Models in Vision [Website] WACV'24
- Organizer, Workshop on Open-Domain Reasoning under Multi-Modal Settings (ODRUM),   [Website] [YouTube] CVPR'23
- Organizer, Workshop on Open-Domain Retrieval under Multi-Modal Settings (ODRUM),   [Website] [YouTube] CVPR'22
- Organizer, Tutorial on Semantic Data Engineering under Multimodal Settings (SERUM)   [Website] WACV'23
- Organizer, 2021 Frontiers of V&L Seminar Series,   [Website], [YouTube] ASU

UNIVERSITY SERVICE	**University Service (at UMBC):**	
	• UMBC HPCF Governance, Subcommittee for CHIP-GPU	2025-present
	• Course Development, CMSC 475/675: Neural Networks	[Website]
	• Course Development, CMSC 491/691: Robust Machine Learning	[Website]
	• Course Development, CMSC 472/672: Computer Vision	[Website]
	• Regular Graduate Course Proposal, CMSC 672: Computer Vision	Approved
	• Regular Undergraduate Course Proposal, CMSC 472: Computer Vision	Approved
	• PPR Seminar: Advances in Perception, Prediction, and Reasoning	[Website]
	• Graduate Admissions Committee	2023–present
	• Department Publicity Committee	2023–present
	• Faculty Mentor, Center for Women in Technology	2023–present
	• Undergraduate Student Advisor	2023–present
	• Faculty Learning Community, UMBC Faculty Development Center	2024–25
	• Faculty Mentor, UMBC IEEE + Tau Beta Pi Open Lab	Fall 2024
	• Faculty Volunteer, COEIT Ph.D. Open House	2025, 2024
	• CSEE Lightning Talks and Open House	Fall 2023, Fall 2024
	• Reviewer, CSEE Research Day	Spring 2024
	• Reviewer, UMBC ORCA Internal Grants	2025
	• Reviewer, COEIT Cybersecurity Research and Education Proposals	Fall 2024
	• Interviewer, CSEE Faculty Candidates	2024, 2025
	• Interviewer, COEIT Staff Searches	2025
	• Member, Asian and Asian American Faculty Staff Council	2023–present

University Service (at ASU):

- Founder, Summer Vision Reading Group, ASU [Website]
- Course Development, CSE591: Frontier Topics in Vision & Language [YouTube] [website] Spring 2021, ASU
- Volunteer, 2019 Southwest Robotics Symposium, Tempe AZ
- Volunteer, International Conference on Machine Learning 2020, Virtual
- Founding Advisor, ASU Machine Learning Club, ASU
- Award Reviewer, GPSA Teaching Award Reviewer ASU
- Mentor, Graduate Student Mentorship Program, ASU
- Project Mentor, CSE598 - Perception in Robotics, ASU Spring 2022
- Project Mentor, CSE576 - Natural Language Processing, ASU Fall 2018

AWARDS		
	CVPR 2024 VDU Workshop, Best Paper Award	2024
	Research Excellence Award, ASU GPSA	2022
	Outstanding Mentor Award, ASU GPSA	2022
	NeurIPS Top Reviewer	NeurIPS 2022
	CVPR 2022 Doctoral Consortium	CVPR 2022
	ICLR Best Reviewer	ICLR 2022
	SCAI Doctoral Fellowship (ASU),	2022, 2021, 2020
	Engineering Graduate Fellowship, (ASU Engineering)	2023, 2020
	ASU GPSA Travel Award	for WACV 2023
	Graduate College Travel Award (declined)	WACV'23, CVPR'22
	Graduate College Travel Award (accepted)	ICCV'21, EMNLP'20, ECCV'20
	IJCAI 2019 Doctoral Consortium	IJCAI 2019
	Inducted, IEEE Eta Kappa Nu, Sigma Chapter	CMU, 2017

National Talent Scholarship, National Council of Educational Research and Training (Govt. of India) 2007–2015

MEDIA

UMBC team leads research into AI tools that can assess the feasibility of scientific claims
04/2025
UMBC News

Alum inspires next generation of computer vision researchers 10/2024
ASU Full Circle

Frontiers of multimodal learning: A responsible AI approach 09/2023
Microsoft Research Blog

CASC research in ML robustness debuts at AAAI conference 02/2021
News and Press, LLNL Computing

HuggingFace and Intel release a solution for high-fidelity text and image consistency
NetEase (163.com), China 04/2024

REFERENCES   Available upon request