
Bryan A. Plummer

Computer Vision • Machine Learning • Natural Language Processing

RESEARCH INTERESTS

My research interests fall within the umbrella of artificial intelligence with a focus on multimodal machine learning, efficient neural networks, explainable and fair AI, and robust ML.

ACADEMIC EXPERIENCE

Assistant Professor , Boston University	July 2020 – Present
Research Assistant Professor , Boston University	September 2019 – June 2020
Postdoctoral Associate , Boston University	April 2018 – August 2019
Supervisors: Kate Saenko and Stan Sclaroff	
Graduate Research Assistant , UIUC	August 2013 – April 2018
Supervisor: Svetlana Lazebnik	
Undergraduate Research Assistant , UIUC	September 2011 – December 2012
Supervisor: Derek Hoiem	

EDUCATION

- 2013 – 2018 **University of Illinois at Urbana-Champaign, Urbana, IL**
PHD IN COMPUTER SCIENCE
Thesis: Grounding Natural Language Phrases in Images and Video
Advisor: Svetlana Lazebnik
- 2011 – 2013 **University of Illinois at Urbana-Champaign, Urbana, IL**
BS IN COMPUTER SCIENCE
- 2009 – 2011 **Mesa Community College, Mesa, AZ**
ASSOCIATE IN SCIENCE, ASSOCIATE IN ARTS

GRANTS

Total funding as PI: \$8,603,148 (\$2,990,837 to BU)

- [1] **Meta:** Online Domain Adaptation for Social Media Analysis. PI at Boston University. *Amount Awarded to BU: \$200,000, Dates: 2022-2024*
- [2] **Adobe:** Gift. *Amount Awarded to BU: \$20,000, Dates: 2023*
- [3] **MIT Lincoln Laboratory:** Adaptive Network Weight Generation for Learning Generalizable Representations. PI at Boston University. *Amount Awarded to BU: \$108,225, Dates: 2022-2024*
- [4] **NSF Infrastructure Innovation for Biological Research:** Collaborative Research: Image-based Readouts of Cellular State using Universal Morphology Embeddings. PI at Boston University. *Amount Awarded: \$1,011,711 (\$498,980 to BU), Dates: 2022-2025*
- [5] **DARPA SemaFor:** TONIC: Trusted Online Content. PI & Team Leader Bryan Plummer (phases 2&3, Co-PI in phase 1), co-PI Kate Saenko (phases 2&3, PI in phase 1). *Amount Awarded: \$7,263,212 (\$2,163,632 to BU), Dates: 2020-2025*

CONFERENCE PUBLICATIONS

- [1] S. Wang, A. Liu, and B. A. Plummer. Noise-aware generalization: robustness to in-domain noise and out-of-domain generalization. In *International Conference on Learning Representations (ICLR)*, 2026.
- [2] V. Agrawal, J. Peters, T. N. Thompson, M. V. Sanian, C. Pham, N. Moshkov, A. Kazi, A. Pillai, J. Freeman, B. Kang, S. L. Farhi, E. Fraenkel, R. Stewart, L. Paavolainen, B. A. Plummer, and J. C. Caicedo. Chammi-75: pre-training multi-channel models with heterogeneous microscopy images. In *International Conference on Learning Representations (ICLR)*, 2026.
- [3] C. Pham, J. C. Caicedo, and B. A. Plummer. ChA-MAEViT: Unifying channel-aware masked autoencoders and multi-channel vision transformers for improved cross-channel learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
- [4] W. Qin, A. Burns, B. A. Plummer, and M. Betke. Walk and read less: Improving the efficiency of vision-and-language navigation via tuning-free multimodal token pruning. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2025.
- [5] A. Liu, K. Miller, V. Saligrama, K. Saenko, B. Gong, S.-N. Lim, and B. A. Plummer. Temporal experts averaging for large-scale temporal domain generalization. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2025.

-
- [6] Y. Wang, Z. Zhang, M. Piana, Z. Zhou, P. Gerstoft, and B. A. Plummer. Real, fake, or manipulated? Detecting machine-influenced text. In *Findings of Empirical Methods in Natural Language Processing (EMNLP)*, 2025.
- [7] A. Ray, J. Duan, E. Brown, R. Tan, D. Bashkurova, R. Hendrix, K. Ehsani, A. Kembhavi, B. A. Plummer, R. Krishna, K.-H. Zeng, and K. Saenko. SAT: Dynamic spatial aptitude training for multimodal language models. In *Conference on Language Modeling (COLM)*, 2025.
- [8] M. Qraitem, P. Teterwak, K. Saenko, and B. A. Plummer. Web artifact attacks disrupt vision language models. In *IEEE International Conference on Computer Vision (ICCV)*, 2025.
- [9] N. Li, K. J. Shih, and B. A. Plummer. Enhancing virtual try-on with synthetic pairs and error-aware noise scheduling. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [10] N. Tasnim and B. A. Plummer. RECAST: Reparameterized, compact weight adaptation for sequential tasks. In *International Conference on Learning Representations (ICLR)*, 2025.
- [11] P. Teterwak, K. Saito, T. Tsiligkaridis, B. A. Plummer, and K. Saenko. Is large-scale pretraining the secret to good domain generalization? In *International Conference on Learning Representations (ICLR)*, 2025.
- [12] P. Teterwak, K. Saito, T. Tsiligkaridis, K. Saenko, and B. A. Plummer. ERM++: An improved baseline for domain generalization. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2025.
- [13] C. Pham and B. A. Plummer. Enhancing feature diversity boosts channel-adaptive vision transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [14] S. Wang and B. A. Plummer. LNL+K: Enhancing learning with noisy labels through noise source knowledge integration. In *European Conference on Computer Vision (ECCV)*, 2024.
- [15] M. Qraitem, K. Saenko, and B. A. Plummer. From fake to real: Pretraining on balanced synthetic images to prevent spurious correlations in image recognition. In *European Conference on Computer Vision (ECCV)*, 2024.
- [16] A. Liu, Z. Li, Z. Chen, N. Li, Y. Xu, and B. A. Plummer. Panofree: Tuning-free holistic multi-view image generation with cross-view self-guidance. In *European Conference on Computer Vision (ECCV)*, 2024.
- [17] A. Burns, K. Saenko, and B. A. Plummer. Tell me what’s next: Textual foresight for generic ui representations. In *Findings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [18] Z. Zhang, W. Qin, and B. A. Plummer. Machine-generated text localization. In *Findings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [19] N. Li, Q. Liu, K. K. Singh, Y. Wang, J. Zhang, B. A. Plummer, and Z. Lin. UniHuman: A unified model for editing human images in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [20] R. Tan, X. Sun, P. Hu, J.-h. Wang, H. Deilamsalehy, B. A. Plummer, B. Russell, and K. Saenko. Koala: Key frame-conditioned long video-llm. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [21] C. Pham, B. Liu, Y. Yang, Z. Chen, T. Liu, J. Yuan, B. A. Plummer, Z. Wang, and H. Yang. Let models speak ciphers: Multiagent debate through embeddings. In *International Conference on Learning Representations (ICLR)*, 2024.
- [22] P. Teterwak, S. Nelson, N. Dryden, D. Bashkurova, K. Saenko, and B. A. Plummer. Learning to compose superweights for neural parameter allocation search. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2024.
- [23] C. Pham, P. Teterwak, S. Nelson, and B. A. Plummer. Mixturegrowth: Growing neural networks by recombining learned parameters. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2024.
- [24] Z. Zhang, J. Zheng, J. Z. Fang, and B. A. Plummer. Text-to-image editing by image information removal. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2024.
- [25] Z. Zhang, Y. Gu, B. A. Plummer, X. Miao, J. Liu, and H. Wang. Movie genre classification by language augmentation and shot sampling. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2024.
- [26] A. Burns, K. Srinivasan, J. Ainslie, G. Brown, B. A. Plummer, K. Saenko, J. Ni, and M. Guo. A suite of generative tasks for multi-level multimodal webpage understanding. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2023.
- [27] Z. Zhang, Y. Gu, and B. A. Plummer. Show, write, and retrieve: Entity-aware article generation and retrieval. In *Findings of Empirical Methods in Natural Language Processing (EMNLP)*, 2023.
- [28] Z. Chen, C. Pham, M. Doron, S. Wang, N. Moshkov, B. A. Plummer, and J. C. Caicedo. CHAMMI: A benchmark for channel-adaptive models in microscopy imaging. In *Advances in Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*, 2023.

-
- [29] A. Ray, F. Radenovic, A. Dubey, B. A. Plummer, R. Krishna, and K. Saenko. Cola: A benchmark for compositional text-to-image retrieval. In *Advances in Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*, 2023.
- [30] Z. Zhang, H. He, B. A. Plummer, Z. Liao, and H. Wang. Complex scene image editing by scene graph comprehension. In *British Machine Vision Conference (BMVC)*, 2023.
- [31] N. Li, K. J. Shih, and B. A. Plummer. Collecting the puzzle pieces: Disentangled self-driven human pose transfer by permuting textures. In *IEEE International Conference on Computer Vision (ICCV)*, 2023.
- [32] M. Qraitem, K. Saenko, and B. A. Plummer. Bias mimicking: A simple sampling approach for bias mitigation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [33] R. Tan, A. Ray, A. Burns, B. A. Plummer, J. Salamon, O. Nieto, B. Russell, and K. Saenko. Language-guided audio-visual source separation via trimodal consistency. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [34] N. Li and B. A. Plummer. Supervised attribute information removal and reconstruction for image manipulation. In *European Conference on Computer Vision (ECCV)*, 2022.
- [35] R. Tan, B. A. Plummer, K. Saenko, J. P. Lewis, A. Sud, and T. Leung. Newsstories: Illustrating articles with visual summaries. In *European Conference on Computer Vision (ECCV)*, 2022.
- [36] A. Burns, D. Arsan, S. Agrawal, R. Kumar, K. Saenko, and B. A. Plummer. A dataset for interactive vision-language navigation with unknown command feasibility. In *European Conference on Computer Vision (ECCV)*, 2022.
- [37] B. A. Plummer, N. Dryden, J. Frost, T. Hoefler, and K. Saenko. Neural parameter allocation search. In *International Conference on Learning Representations (ICLR)*, 2022.
- [38] R. Tan, B. A. Plummer, K. Saenko, H. Jin, and B. Russell. Look at what i'm doing: Self-supervised spatial grounding of narrations in instructional videos. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [39] D. Kim, K. Saito, T.-H. Oh, B. A. Plummer, S. Sclaroff, and K. Saenko. CDS: Cross-Domain Self-supervised Pre-training. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [40] S. Mishra, Z. Zhang, Y. Shen, R. Kumar, V. Saligrama, and B. A. Plummer. Effectively leveraging attributes for visual similarity. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [41] R. Tan, H. Xu, K. Saenko, and B. A. Plummer. LoGAN: Latent graph co-attention network for weakly-supervised video moment retrieval. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2021.
- [42] R. Tan, B. A. Plummer, and K. Saenko. Detecting cross-modal inconsistency to defend against neural fake news. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [43] A. Burns, D. Kim, D. Wijaya, K. Saenko, and B. A. Plummer. Learning to scale multilingual representations for vision-language tasks. In *European Conference on Computer Vision (ECCV)*, 2020.
- [44] B. A. Plummer, M. I. Vasileva, V. Petsiuk, K. Saenko, and D. Forsyth. Why do these match? Explaining the behavior of image similarity models. In *European Conference on Computer Vision (ECCV)*, 2020.
- [45] D. Kim, K. Saito, K. Saenko, S. Sclaroff, and B. A. Plummer. MULE: Multimodal universal language embedding. In *AAAI Conference on Artificial Intelligence*, 2020.
- [46] A. Burns, R. Tan, K. Saenko, S. Sclaroff, and B. A. Plummer. Language features matter: Effective language representations for vision-language tasks. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [47] R. Tan, M. I. Vasileva, K. Saenko, and B. A. Plummer. Learning similarity conditions without explicit supervision. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [48] H. Xu, K. He, B. A. Plummer, L. Sigal, S. Sclaroff, and K. Saenko. Multilevel language and vision integration for text-to-clip retrieval. In *AAAI Conference on Artificial Intelligence*, 2019.
- [49] B. A. Plummer, M. H. Kiapour, S. Zheng, and R. Piramuthu. Give me a hint! Navigating Image Databases using Human-in-the-loop Feedback. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2019.
- [50] M. I. Vasileva, B. A. Plummer, K. Dusad, S. Rajpal, R. Kumar, and D. Forsyth. Learning type-aware embeddings for fashion compatibility. In *European Conference on Computer Vision (ECCV)*, 2018.
- [51] B. A. Plummer, P. Kordas, M. H. Kiapour, S. Zheng, R. Piramuthu, and S. Lazebnik. Conditional image-text embedding networks. In *European Conference on Computer Vision (ECCV)*, 2018.
- [52] B. A. Plummer, A. Mallya, C. M. Cervantes, J. Hockenmaier, and S. Lazebnik. Phrase localization and visual relationship detection with comprehensive image-language cues. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [53] B. A. Plummer, M. Brown, and S. Lazebnik. Enhancing video summarization via vision-language embedding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [54] T. Tommasi, A. Mallya, B. A. Plummer, S. Lazebnik, A. C. Berg, and T. L. Berg. Solving visual madlibs with multiple cues. In *British Machine Vision Conference (BMVC)*, 2016.

-
- [55] B. A. Plummer, L. Wang, C. M. Cervantes, J. C. Caicedo, J. Hockenmaier, and S. Lazebnik. Flickr30k Entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.

JOURNAL PUBLICATIONS

- [56] D. Lteif, D. Appapogu, S. A. Bargal, B. A. Plummer, and V. B. Kolachalama. Anatomy-guided, modality-agnostic segmentation of neuroimaging abnormalities. *Human Brain Mapping*, 46(14):e70329, 2025.
- [57] J. Wu, B. M. Trifiro, L. R. Ranker, J. M. Origgi, E. J. Benjamin, R. M. Robertson, A. Bhatnagar, A. C. Stokes, Z. Xuan, D. Wijaya, B. Plummer, J. Cornacchione Ross, J. L. Fetterman, and T. Hong. Health warnings on instagram advertisements for synthetic nicotine e-cigarettes and engagement. *JAMA Network Open*, 7(9):e2434434–e2434434, 2024.
- [58] C. Xue, S. S. Kowshik, D. Lteif, S. Puducheri, V. H. Jasodanand, O. T. Zhou, A. S. Walia, O. B. Guney, J. D. Zhang, S. T. Pham, A. Kaliaev, V. C. Andreu-Arasa, B. C. Dwyer, C. W. Farris, H. Hao, S. Kedar, A. Z. Mian, D. L. Murman, S. A. O’Shea, A. B. Paul, S. Rohatgi, M.-H. Saint-Hilaire, E. A. Sartor, B. N. Setty, J. E. Small, A. Swaminathan, O. Taraschenko, J. Yuan, Y. Zhou, S. Zhu, C. Karjadi, T. F. A. Ang, S. A. Bargal, B. A. Plummer, K. L. Poston, M. Ahangaran, R. Au, and V. B. Kolachalama. Ai-based differential diagnosis of dementia etiologies on multimodal data. *Nature Medicine*, 2024.
- [59] D. Lteif, S. Sreerama, S. A. Bargal, B. A. Plummer, R. Au, and V. B. Kolachalama. Disease-driven domain generalization for neuroimaging-based assessment of alzheimers disease. *Human Brain Mapping*, 45(8):e26707, 2024.
- [60] B. A. Plummer, K. J. Shih, Y. Li, K. Xu, S. Lazebnik, S. Sclaroff, and K. Saenko. Revisiting image-language networks for open-ended phrase detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 44(4):2155–2167, 2022.
- [61] O. Watkins, S. Huang, J. Frost, K. Bhatia, E. Weiner, P. Abbeel, T. Darrell, B. A. Plummer, K. Saenko, and A. Dragan. Explaining robot policies. *Applied AI Letters*, 2021.
- [62] T. Tommasi, A. Mallya, B. A. Plummer, S. Lazebnik, A. C. Berg, and T. L. Berg. Combining multiple cues for visual madlibs question answering. *International Journal of Computer Vision (IJCV)*, 127(1):38–60, 2019.
- [63] B. A. Plummer, L. Wang, C. M. Cervantes, J. C. Caicedo, J. Hockenmaier, and S. Lazebnik. Flickr30k Entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *International Journal of Computer Vision (IJCV)*, 123(1):74–93, 2017.

WORKSHOP PUBLICATIONS

- [64] K. Deng, A. Ray, R. Tan, S. Gabriel, B. A. Plummer, and K. Saenko. Socratis: are large multimodal models emotionally aware? In *IEEE International Conference on Computer Vision Workshop (ICCVW) on Emotionally And Culturally Intelligent AI*, 2023.
- [65] M. Graitem and B. A. Plummer. From coarse to fine-grained concept based discrimination for phrase detection. In *IEEE Conference on Computer Vision and Pattern Recognition Workshop (CVPRW) on Computer Vision in the Wild*, 2023.
- [66] V. Petsiuk, A. E. Siemenn, S. Surbehera, Z. Chin, K. Tyser, G. Hunter, A. Raghavan, Y. Hicke, B. A. Plummer, O. Kerret, T. Buonassisi, K. Saenko, A. Solar-Lezama, and I. Drori. Human evaluation of text-to-image models on a multi-task benchmark. In *Advances in Neural Information Processing Systems (NeurIPS) Workshop on Human Evaluation of Generative Models (HEGM)*, 2022.
- [67] D. Kim, K. Saito, S. Mishra, S. Sclaroff, K. Saenko, and B. A. Plummer. Self-supervised visual attribute learning for fashion compatibility. In *IEEE International Conference on Computer Vision Workshop (ICCVW) on Visual Inductive Priors for Data-Efficient Deep Learning*, 2021.
- [68] D. Kim, T. Lan, C. Zou, N. Xu, B. A. Plummer, S. Sclaroff, J. Eledath, and G. Medioni. Mila: multi-task learning from videos via efficient inter-frame attention. In *IEEE International Conference on Computer Vision Workshop (ICCVW) on Multi-Task Learning in Computer Vision (MTL)*, 2021.
- [69] J. Frost, O. Watkins, E. Weiner, P. Abbeel, T. Darrell, B. A. Plummer, and K. Saenko. Explaining reinforcement learning policies through counterfactual trajectories. In *International Conference on Machine Learning (ICML) Workshop on Human in the Loop Learning (HILL)*, 2021.
- [70] P. D. Tsatsoulis, B. A. Plummer, and D. Forsyth. Visual analogies: a framework for defining aspect categorization. In *European Conference on Computer Vision Workshop (ECCVW) on Transferring and Adapting Source Knowledge in Computer Vision (TASK-CV)*, 2016.

PATENTS

- [1] D. A. Forsyth, R. Kumar, K. Dusad, K. Li, M. I. Vasileva, B. A. Plummer, Y. Shen, and S. Rajpal. Search engine use of neural network regressor for multi-modal fashion recommendations based on visual semantic embeddings. *US Patent 12,131,365*, 2024.

-
- [2] R. Piramuthu, T. S. Keefer, K. C. Crookston, A. S. Rekhi, N. A. K. Nazimudeen, P. Gudipati, S. Lin, J. F. Weigel, F. Zhong, S. Ramesh, M. Kiapour, S. Zheng, A. O. Pereira, R. S. Lanka, M. A. u. Islam, N. A. Whyte, G. Iyengar, and B. A. Plummer. Computer vision and image characteristic search. *US Patent 11,250,487*, 2022.
- [3] R. Piramuthu, T. S. Keefer, A. S. Rekhi, P. Gudipati, M. Kiapour, S. Zheng, A. O. Pereira, R. S. Lanka, M. A. u. Islam, N. A. Whyte, G. Iyengar, and B. A. Plummer. Computer vision for unsuccessful queries and iterative search. *US Patent 11,200,611*, 2021.
- [4] B. M. Fields, S. Roberson, A. Harish, H. Lim, M. J. Waughtel, B. A. Plummer, and P. A. Brown. Driver organization and management for driver's education. *US Patent 10,373,523*, 2019.
- [5] B. A. Plummer, P. A. Brown, B. M. Fields, and S. Roberson. Gaze tracking for a vehicle operator. *US Patent 9,547,798*, 2017.
- [6] B. M. Fields, S. Roberson, A. Harish, H. Lim, M. J. Waughtel, B. A. Plummer, and P. A. Brown. Real-time driver observation and progress monitoring. *US Patent 9,586,591*, 2017.
- [7] B. A. Plummer, P. A. Brown, J. He, B. M. Fields, S. Roberson, S. Cielocha, and J. Peng. System and method for monitoring and reducing vehicle operator impairment. *US Patent 9,758,173*, 2017.
- [8] B. A. Plummer, P. A. Brown, J. He, B. M. Fields, S. Roberson, S. Cielocha, and J. Peng. Systems and methodologies for real-time driver gaze location determination and analysis utilizing computer vision technology. *US Patent 9,275,532*, 2016.
- [9] B. A. Plummer, D. Cross, and N. L. Tofte. Method of estimating damage to a roof. *US Patent 9,262,564*, 2016.
- [10] B. A. Plummer, D. Cross, and N. L. Tofte. Systems and methods for assessing a roof and generating models. *US Patent 9,098,655*, 2015.
- [11] B. A. Plummer, D. Cross, and N. L. Tofte. Systems and methods for assessing property damage. *US Patent 8,756,085*, 2014.
- [12] B. M. Fields, J. He, J. A. Nepomuceno, S. Roberson, B. A. Plummer, K. C. Houdek, and N. Jain. Real-time driver observation and scoring for driver's education. *US Patent 8,876,535*, 2014.
- [13] B. A. Plummer and D. Cross. Systems and methods for assessing a roof. *US Patent 8,874,454*, 2014.

TEACHING

BU FALL 2022-	CS 541: Applied Machine Learning , <i>Instructor</i> , Enrollment: 90
BU SPRING 2021;2025-	CS 591/598: Multimodal Machine Learning , <i>Instructor</i> , Enrollment: 50
BU SPRING 2022;2023	CS 585: Image and Video Computing , <i>Instructor</i> , Enrollment: 120
BU FALL 2020;2021	CS 542: Machine Learning , <i>Instructor</i> , Enrollment: 90-120
UIUC SPRING 2017	CS 543: Computer Vision , <i>Teaching Assistant</i> (Instructor: Derek Hoiem)

MENTORING

PhD Students

Andrea Burns (coadvised, 2023), Reuben Tan (coadvised, 2024), Zhongping Zhang (2024), Nannan Li (2025), Siqi Wang (2025), Diala Lteif (coadvised), Piotr Teterwak (coadvised), Maan Qraitem (coadvised), Arijit Ray (coadvised), Chau Pham, Aoming Liu (coadvised), Keanu Nichols, Nazia Tasnim, Mahir Patel (coadvised), Nicholas Ikechukwu (coadvised)

Other Students

"Elly" Yitong Wang (undergrad, 2025), Margherita Piana (undergrad, 2025), Harsh Khatri (masters, 2024), Nishant Nadkarni (masters, 2023), Divya Appapogu (masters, 2023), Duruvan Saravanan (masters, 2023), Soren Nelson (masters, 2022), Xiaohan Zou (masters, 2022), Manyuan Lu (BA/MS, 2022), Zora Che (undergrad, 2022), Murtadha Al-Barani (masters, 2021), Sanjna Agrawal (undergrad, 2021), "John" Yuanming Chai (undergrad, 2021), Julius Frost (BA/MS, 2021), Farheen Rahman (undergrad, 2020), Tammy Qiu (undergrad, 2019), Paige Kordas (undergrad, 2018)

INDUSTRY/RESEARCH LAB EXPERIENCE

Ebay Research Labs , <i>Research Intern</i>	May 2017 – August 2017
Explored methods of using computer vision for e-commerce tasks including human-in-the-loop approaches.	
Google , <i>Software Engineering Intern, PhD</i>	May 2016 – August 2016
Developed vision-language approaches for tasks using video data.	
A9 , <i>Software Developer Intern</i>	May 2015 – August 2015
Investigated possible alternative approaches to existing algorithms developed by the Visual Search team based on recent advances in image description.	
State Farm Research and Development Center , <i>IT/Systems Intern</i>	August 2014 – May 2015
Advisor to three computer vision projects in development at the center.	
MIT Lincoln Laboratory , <i>Summer Research Program Intern</i>	June 2014 – August 2014

Developed an approach using active learning for object recognition on natural images with an emphasis on finding rare object categories.

State Farm Research and Development Center, IT/Systems Intern November 2012 – August 2013

Built a system to use 3D scans of objects to detect anomalies and created a program to do real time ($\geq 5\text{Hz}$) gaze tracking using low resolution images on computationally limited platforms.

NASA Jet Propulsion Laboratory, JPLSIP Intern June 2012 – August 2012

Created a class that performs image matching using phase correlation to register an image and then align them as well as create an interface for HDF-EOS 5 images.

Aqueous Solutions, LLC, Student SDE Intern August 2011 – January 2013

Rewrote the installer, added automatic updating, and contributed to many other features in applications developed by the company.

NASA Goddard Space Flight Center, SIES Intern June 2011 – August 2011

Conducted a parameter search to produce specific results in a model that simulates snowflake growth and added GPU processing capabilities.

HONORS

† National award ‡ Regional award

- 2023 Best Paper Runner-up for “Text-to-image Editing by Image Information Removal,” CVPR Workshop on AI for Content Creation
- 2023 Hariri Institute Junior Faculty Fellow
- 2021 ICCV Outstanding Reviewer
- 2019 Best Fashion Paper Runner-up for “Learning Similarity Conditions Without Explicit Supervision,” ICCV Workshop on Computer Vision for Fashion, Art and Design
- 2019, 2022 NeurIPS Top Reviewer
- 2018–2021 CVPR Outstanding Reviewer
- 2016 Best Paper Award for “Visual Analogies: A Framework for Defining Aspect Categorization,” ECCV TASK-CV Workshop
- 2015 †NSF Graduate Research Fellowship Honorable Mention
- 2014, 2015 3M Foundation Fellowship
- 2013 UIUC Research Park Most Outstanding Undergraduate Intern Finalist
- 2012 †Barry M. Goldwater Scholar
- 2012 The Illinois Club Make-A-Difference Award
- 2011 †Coca-Cola Academic Team, Gold Scholar
- 2011 ‡All-Arizona Academic Team, First Team
- 2011 †Distinguished Chapter Officer Team, Phi Theta Kappa Honor Society
- 2011 †National Community College Aerospace Scholar
- 2011 †Presidential Volunteer Service Award, Gold
- 2011 ‡Most Distinguished Officer, Arizona Hall of Honor, Phi Theta Kappa Honor Society
- 2011 MCC Excellence in Service Award
- 2011 PTK Five Star Competitive Edge Award
- 2011 MCC Honors Speaker, Academic Achievement Awards Keynote
- 2010, 2011 MCC Honors in Action Award
- 2010 †Leaders of Promise, Phi Theta Kappa Honor Society
- 2010 ‡UCAN Serve Impact Scholar
- 2010 MCC Service Scholar
- 2010 Academic Achievement Award, MCC Psychology Department
- 2009 MCC Foundation Scholarship

OUTREACH AND SERVICE

- 2024 – PRESENT Boston University IVC REU+AI Bootcamp Director
- 2023 – 2024 Boston University AI4All Co-Director
- 2021 Boston University Research in Science & Engineering Program (RISE) Mentor
- 2018 – 2020 Boston University AI4All Project Mentor
- 2016 CS Admissions Committee Graduate Application Reviewer
- 2014 – 2015 CS Grad Peer Mentor
- 2013 – 2018 CS Grad Ambassador
- 2013 – 2015 SigBot and SigArt Project Mentor, ACM UIUC Student Chapter
- 2012 – 2013 SigBot Chair, ACM UIUC Student Chapter
- 2012 – 2013 Student Representative, UIUC CS Department Grievance Committee
- 2011 – 2012 Speakers Committee, 18th Reflections Projections Conference
- 2011 Logistics Committee, 17th Reflections Projections Conference
- 2010 – 2011 VP of Leadership, MCC Chapter of the Phi Theta Kappa Honor Society

- 2009 – 2011 AmeriCorps Volunteer, site: Arizona Science Center
- 2010 VP of Administration, MCC Chapter of the Phi Theta Kappa Honor Society
- 2010 LeaderCorps Member, Arizona Governor's Office
- 2010 Logistics Committee, Arizona National and Community Service Conference
- 2010 STEM Club Advisor, Orangewood Elementary School
- 2009 – 2010 Organizing Committee Member, MCC Beautification Project
- 2009 – 2010 Learn and Serve Camp Team Leader, Arizona Department of Education
- 2009 – 2010 Student Member, MCC Service Learning Advisory Board
- 2009 – 2010 Student Member, MCC Community Strengthening Allies

PROFESSIONAL ACTIVITIES

- EDITORIAL BOARD International Journal of Computer Vision'24-Present
- LEAD AREA CHAIR CVPR'26
- AREA CHAIR/SPC ECCV'20'24; AAAI'22'26; CVPR'23-'25; ACL ARR'23-'25; NeurIPS'23-'25; ICCV'25; ICML'25; WACV'26; ICLR'26
- CONFERENCE REVIEWER CVPR'17-'22; ICCV'17-'23; CHI'18; ECCV'18'22; BMVC'19; WACV'24; NeurIPS'19'20'22; AAAI'20; ICML'20'23; ACL'20-'22; NeurIPS D&B'23'24; EMNLP'23; ICLR'25
- JOURNAL REVIEWER Transactions on Pattern Analysis and Machine Intelligence'15'16'20-'25; International Journal of Computer Vision'18'19'24; Transactions on Multimedia'19
- GRANT REVIEWER NIH'21

Workshop Organization

- 2026 CVPR: Workshop on the Findings of the CVPR, Lead Organizer
- 2025 ICCV: Workshop on the Findings of the ICCV, Lead Organizer
- 2024 ACM SIGIR: 1st Workshop on Multimodal Representation and Retrieval, Committee Member

Invited Talks

- 2025 University of California, Berkeley; Columbia University; University of Wisconsin-Madison; University of Washington; University of Michigan, Ann Arbor; Korea University; University of Maryland, College Park; ECI on Advances in Optics for Biotechnology, Medicine, and Surgery
- 2024 University of Edinburgh; Trinity College Dublin; CVPR Area Chair Workshop; Princeton University, University of Oxford; Stanford University
- 2023 University of Pittsburgh; CVPR Workshop - Computer Vision in the Wild; Boston College
- 2022 Amazon; University of Bristol
- 2021 ICCV Workshop – Closing the Loop between Vision and Language; CVPR Workshop – Media Forensics; University of Illinois at Urbana-Champaign; Zalando; NVIDIA
- 2020 Keynote at Amazon Computer Vision Conference; MIT Lincoln Laboratory; Brown University
- 2019 ICCV Workshop – Linguistics Meets Image and Video Retrieval; Accenture