Edward Zhang

zedward@engineering.upenn.edu | +1 215 450 6405
Website | zedward23 | Edward Zhang

Research Summary

Ph.D. student at Penn working in computer vision, robotics, multimodal foundation models, and generative learning. Research focuses on perception and synthetic data for specialized, data-scarce, safety-critical environments, with hands-on deployment of learned systems on real robots.

Education

University of Pennsylvania 2024–Present
Ph.D., Computer and Information Science; GRASP Laboratory
Advisors: Eric Eaton and Dan Hashimoto. **NSF Graduate Research Fellow (2024).**

University of Pennsylvania 2019–2023
Penn Engineering and the Wharton School
BSE, Digital Media Design; BS, Entrepreneurship & Innovation. *Magna Cum Laude*; GPA: 3.61/4.0.

Selected Publications and Manuscripts

Limitations of Synthetic Data Generation in Specialized Data-Scarce Domains
Edward Zhang, Marcel Hussing, Tanay Tandon, Shenbagaraj Kannapiran, Jason Hughes, Youkang Wang, Joshua Caswell, Agelos Kratimenos, Yi Fan Li, Milan Manoj, Ethan Sanchez, Sumukh Shrote, Camillo Jose Taylor, Daniel A. Hashimoto, Eric Eaton.
arXiv:2608.13729, August 13, 2026.

A Multi-Robot Platform for Robotic Triage Combining Onboard Sensing and Foundation Models
Jason Hughes, Marcel Hussing*, **Edward Zhang***, Shenbagaraj Kannapiran, Joshua Caswell, Kenneth Chaney, Ruichen Deng, Michaela Feehery, Agelos Kratimenos, Yi Fan Li, Britny Major, Ethan Sanchez, Sumukh Shrote, Youkang Wang, Jeremy Wang, Daudi Zein, Luying Zhang, Ruijun Zhang, Alex Zhou, Tenzi Zhouga, Jeremy Cannon, Zaffer Qasim, Jay Yelon, Fernando Cladera, Kostas Daniilidis, Camillo J. Taylor, Eric Eaton.
*arXiv:2512.08754, December 9, 2025. *Equal contribution.*

Benchmarking Visual Language Models for Automated Trauma Detection in Triage Scenarios
Edward Zhang, Marcel Hussing, Jason Hughes, et al.
Manuscript, available upon request.

Ego-Exo4D: Understanding Skilled Human Activity from First- and Third-Person Perspectives
Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, . . . , **Edward Zhang**, et al.
IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024.

Research Experience

DARPA Triage Challenge — Ph.D. Researcher, Robotics & Multimodal Perception Sep 2024–Present
University of Pennsylvania, GRASP Laboratory; Advisor: Eric Eaton

- Integrate ML and perception components into deployed robotic platforms for automated casualty assessment in mass-casualty scenarios using ROS/ROS2 and Dockerized software stacks.
- Develop body-part and injury perception, trauma classification, and alertness-related perception components; connect multimodal sensing and perception outputs to automated casualty reporting.
- Develop and evaluate VLM-based perception for trauma and injury understanding; coauthor of the multi-robot triage system report and the trauma-detection manuscript.
- Lead and mentor approximately four master's students in data collection, research, and paper preparation.

Synthetic Data for Specialized Data-Scarce Vision
University of Pennsylvania; Ph.D. research with Eric Eaton and Dan Hashimoto

- Led a first-author study benchmarking synthetic data generation in specialized, data-scarce vision, comparing generative augmentation with conventional baselines across five trauma-recognition tasks.
- Evaluated downstream classification on held-out real images with subject-wise train–validation splits, comparing distribution modeling and sample perturbation against non-generative augmentation.
- Analyzed memorization or collapse, distributional drift, and overly canonical synthetic samples using feature-space and nearest-neighbor analyses; no evaluated generative method consistently outperformed the strong non-generative baseline across tasks.
- Ongoing work investigates how to identify visual features that define data sparsity and steer image generation toward those features without memorizing training examples.

Ego-Exo Data for Pose and Segmentation — Researcher Oct 2022–Feb 2024
University of Pennsylvania; Advisor: Jianbo Shi

- Led Penn's data-collection efforts for the Ego-Exo4D collaboration: over 100 participants and 600 hours of footage.
- Investigated semantic segmentation for 3D hand-pose recovery and reconstruction from the collected data.

Multi-Camera Capture for Neural Rendering — Researcher May 2023–Sep 2024
Advisor: Lingjie Liu. Led setup of a general-purpose multi-camera capture apparatus at the SIG Center, working on a pipeline from raw images and video to neural scene representations and investigating semantic information in NeuS2.

SIG Center for Computer Graphics — Researcher Jun 2022–Apr 2023
Advisors: Aline Normoyle and Norman Badler. Conducted funded research on full-body and facial motion capture for sign-language processing and human–computer interaction; presented open-body-fit at MIG 2022.

LLMs for Innovation Management — Pre-Doctoral Student Sep 2023–Jan 2025
Advisor: Valery Yakubovich. Developed patent-retrieval workflows using the USPTO API and language embeddings to assess emerging technologies and market positioning from natural-language interviews.

CEO Transcript Analysis — Research Assistant 2023–2025
Advisor: Claudine Gartenberg. Applied LLMs to label CEO mission statements and analyze their relationship to firm compensation; investigated transcript sentiment.

Game-Based Learning — Wharton Honors Thesis Sep 2022–Apr 2023
Advisor: Ethan Mollick. Developed a Unity game to teach recursion and conducted a literature review and preliminary educational experiments.

Wharton World Research Assistantship Program May 2021–Feb 2022
Advisor: John Zhang. Led analysis of U.S.-listed Chinese stocks using SEC filings, Python, and SQL; presented at the 2022 Wharton Research Symposium.

Technical Skills

Programming	Python, C++, C, CUDA, SQL, C#, R, TypeScript
Machine Learning / Vision	PyTorch, OpenCV, TensorFlow, Hugging Face Transformers, Diffusers, PEFT; vision-language models, generative models, multimodal perception
Robotics / Systems	ROS/ROS2, Docker; robotic perception integration, GPU programming; familiar with the Boston Dynamics Spot API, NVIDIA Jetson platforms, and ZED 2i depth cameras
Graphics	Vulkan, WebGPU, Unity, Unreal Engine, Blender, NVIDIA Nsight

Work Experience and Leadership

Airstrafe Interactive — Software Engineer Jul 2023–Feb 2024
Led a small developer team implementing hardware occlusion culling for arbitrary, non-instanced meshes using Unity's Scriptable Render Pipeline API and Shader Graph, on contract through LatticeWork GT.

LatticeWork GT LLC — Founder Jun 2023–Present
Founded a company for contract software development and internal work in graphics, vision, game engines, data science, and high-performance hardware utilization.

Children's Hospital of Philadelphia — Project Manager Jan–Sep 2023
Managed and designed an interactive assent-form project for an NIH-funded CHOP–UPGRADE collaboration supporting sleep studies involving patients with Down syndrome.

UPGRADE — President Jun 2021–May 2023
Grew the organization from 2 to 100+ members; led 30+ developers across programming, art, sound, and writing to ship *Bionic Traveler* on Steam within one year.

Selected Teaching

Advanced Topics in Machine Perception — Teaching Assistant Sep–Dec 2023
Jianbo Shi. Supported a Ph.D.-level course on deep learning for computer vision.

Modern Datamining — Teaching Assistant Jan–May 2023
Linda Zhao. Taught ML concepts and graded assignments and exams for a Ph.D.-level statistics course using R and Python.

Wharton GPT Workshop — Instructor Aug–Oct 2023
Christian Terwiesch. Designed and taught workshops for faculty and doctoral students on LLM task automation with the OpenAI API.

Additional teaching: Wharton Data Science Academy, Teaching Fellow (2022, 2023); Advanced Rendering, Teaching Assistant (2022); Engineering Summer Academy at Penn, Teaching Assistant (2021).

Additional Publications

open-body-fit: Open-source pipeline for understanding sign language gestures from video

Aline Normoyle, **Edward Zhang**, Norman Badler.

ACM SIGGRAPH Conference on Motion, Interaction and Games (MIG), 2022. Poster presentation.

Selected Invited Talks

An Overview of Neural Radiance Fields — CIS 6600, Penn Apr 2023
Ventures in Visual Computing — Huawei Cloud Mar 2023
What's Up? Gas Prices! — Women in Data Science at Penn Feb 2023

Selected Technical Project

CUDA Path Tracer Code
Implemented a GPU path tracer with arbitrary mesh loading, multiple materials, depth of field, and anti-aliasing; accelerated execution with stream compaction, memory coalescing, and ray caching.