Esha Pahwa

Pittsburgh, PA | +1-(412)-954-8044 | epahwa@cs.cmu.edu | Google Scholar
eshapahwa.github.io | linkedin.com/in/esha-pahwa/ | eshapahwa

Education

Carnegie Mellon University — Pittsburgh, PA
Master of Science in Machine Learning — Dec 2025
- Primary Courses: Advance Intro to Machine Learning, Multimodal ML, Visual Learning and Recognition, Deep Reinforcement Learning
- **Teaching Assistant for Advance Deep Learning (Masters and Ph.D. level)** under the guidance of Dr. Ruslan Salakhutdinov

Birla Institute of Technology and Science, Pilani — Pilani, India
B.E. in Computer Science, M.Sc. in Chemistry — Jun 2023

Experience

Corvic AI (Enterprise AI - 12M seed) — May 2025 – August 2025
Machine Learning Intern — *Mountain View, CA*
- Added **citation generation** capabilities to LLM responses, linking outputs to grounded sources within enterprise data to improve trust and auditability.
- Built an **LLM-driven orchestration layer** that determines the next operations to perform by interpreting user intent, retrieving relevant documents, and invoking **Python code generation** for execution inside a secure, **Pyodide-based sandbox** to safely run untrusted code, followed by downstream tool invocation.
- Enabled end-to-end support for diverse **document types**, including ingestion, parsing, and transforming them into **multimodal agents** capable of generating contextualized space and agent responses.

Katerina Fragkiadaki's Research Lab — Jan 2025 – May 2025
Research Assitant — *Pittsburgh, PA*
- Designed and implemented **Chain-of-Decoders (CoD)**, a multimodal chain-of-thought framework extending **Pi0** and **OpenVLA** architectures by conditioning inference on latent subgoals across language, visual, and sensor inputs to generate **interpretable, real-time robotic policies**.
- Conducted ablation studies on RLBench using **CoD** variants (e.g., CoD-subgoals, CoD-mask), and introduced a **masked learning strategy** for partial annotations, improving robustness and reducing token usage in multimodal CoT inference.

Adobe Inc. — Jul 2023 – Aug 2024
Member of Technical Staff - I (Generative AI Team) — *Noida, India*
- **Led agentic backend integration** of Adobe's AI tools for automated email, SMS & push notifications, streamlining workflows. Recognized with the **Adobe Spot Award** for innovation.
- **Optimized LLM token management** for email generation systems, dynamically allocating tokens to **reduce operational costs by 27%** and minimize throttling risks.
- **Developed scalable APIs**, including an automated thumbnail generation service using Puppeteer, which **improved rendering speeds by 70% for high-traffic webpages**, and integrated with **Azure Functions for enhanced security** and reliability.
- Designed and implemented a **multilingual API for a prompt library** for accurate **prompt engineering**, catering to solution-specific needs, streamlining content generation across diverse market segments.

Research Intern (Media and Data Science Research Labs - Adobe) — May 2022 - Aug 2022
- Pioneered a **shared encoder–decoder** architecture using Adobe XDM with a **privacy-preserving mapping model**, achieving **93.11% accuracy** and **0.994 ROC** in cross-dataset predictions, and **98.8% segment prediction accuracy** with **0.978 cosine similarity**.

Google Research — Jan 2023 - Jul 2023
Research Associate (Machine Learning and Optimization Team) | Supervisor: Dr. Prateek Jain — *Bengaluru, India*
- Contributed to **product retrieval in Google Shopping Ads** by processing user queries and conducting in-depth analysis to identify opportunities for system optimization. Thesis Link
- Developed and optimized a **Dual Encoder Information Retrieval (IR)** model using **pre-trained transformers** to enhance product retrieval accuracy from large-scale datasets.
- Migrated model training to **JAX with data parallelism across TPUs**, significantly **reducing training time to 14 minutes for a 100K product ID dataset**, demonstrating expertise in scalable machine learning systems.

Computer Vision Center — Aug 2022 - Dec 2022
Researcher (Generative Models) | Supervisor: Dr. Luis Herranz — *Barcelona, Spain*
- Developed ProjectedGAN++ using **EfficientNet and autoencoders** to make networks adaptable to test datasets, **improving image synthesis** on small datasets. Thesis Link
- **Applied transfer learning and self-supervised learning** to enhance GAN performance, improving FID scores - **25%-50% reduction across various datasets** - while reducing overfitting through cross-channel and cross-scale mixing.

Selected Publications

- **SITCOM: Scaling Inference-Time COMpute for VLAs** Paper [NeurIPS 2025 Workshop]
- **MedSkip: Medical Report Generation using Skip Connections and Integrated Attention** Paper | Poster [ICCVW 2021]
- **LVRNet: Lightweight Image Restoration for Aerial Images Under Low Visibility** Paper | Project Page | Code [AAAI 2023]
- **Conditional RGBT Fusion for Effective Crowd Counting** Paper | Poster | Code [IEEE ICIP 2022]
- **DroneAttention: Sparse weighted temporal attention for drone-camera based activity recognition** Paper [Neural Net Journal]

Individual Projects

- **EduMate**: Agentic AI tutor for helping students executing code, providing quizzes, solving doubts and creating flashcards Code
- **Pytorch Hackathon**: Built a web application for diabetic retinopathy detection in eye images. (Rank 25/5000) Blog | Code

Achievements and Other Experience

- **Finalist in Google Deepmind AI Agents Hackathon 2025**
- **Adobe Women-in-Tech Scholarship 2022 (6 selected across India)**, Grace Hopper Celebration Scholarship 2021.
- **Reviewer at: ICCV'23** and **ECCV'22** Workshops.
- **Founder of Computer Vision Research Society**, BITS Pilani with the aim of promoting research culture at my college.

Technical Skills

Languages: Python, Java, SQL, C, HTML/CSS, Rust
Libraries: Tensorflow, OpenCV, Keras, Pytorch, JAX, Numpy, Pandas, Scikit-learn, Matplotlib, Seaborn, Scipy
GenAI Tools: LangChain, OpenAI
Databases: MySQL, PostgreSQL, Pinecone
Cloud Platforms: Azure, AWS
DevOps Tools: Docker, Git