

Tashvik Dhamija

+1 (412) 512-7540 | [linkedin.com/in/tashvikdhamija](https://www.linkedin.com/in/tashvikdhamija) | tdhamija@cs.cmu.edu | <https://tashvikdhamija.github.io/>

EDUCATION

Carnegie Mellon University

Pittsburgh, PA, USA | Present

Master of Science in Artificial Intelligence and Innovation - School of Computer Science

GPA: 4.08/4.0

Coursework: Computer Systems, Intro to Machine Learning (PhD), Advance Deep Learning (PhD), Computer Vision, Advance NLP (PhD), AI Engineering, Learning for 3D Vision, Deep Learning Systems, Generative AI

Delhi Technological University

Delhi, India | May 2023

Bachelor of Technology in Electronics and Communications Engineering

GPA: 9.14/10.0

Coursework: Deep Learning, Computer Vision, Data Structures, Algorithms, Object Oriented Programming, DBMS

EXPERIENCE

Amazon AWS Annapurna Labs

Cupertino, CA, USA | May 2026 – Aug 2026

ML Intern - Neuron Scalable Training Team

- Enabled diffusion and flow-matching model training on AWS Trainium hardware by onboarding state-of-the-art models like Stable Diffusion, FLUX, DiT, and Wan onto Neuron, Annapurna's native PyTorch backend for Trainium.
- Implemented and optimized PyTorch ATen operators for Neuron backend, extending native operator coverage and efficiency.
- Validated long-horizon training dynamics between Neuron and CUDA across text-to-image, text-to-video, and diffusion language models, ensuring convergence, numerical stability, and quality parity.

Inria Startup Studio - Centre at Université Côte d'azur | [Link to Website](#)

Valbonne, France | Oct 2023 – Sep 2024

AI Entrepreneur and Researcher

- Architected a generative AI pipeline leveraging diffusion transformers and cross-modal alignment to synthesize audio-driven video avatars with high temporal coherence and realistic lip-sync, optimized for real-time generation.
- Led a 5-member team to launch MVPs for video generation ([Product Hunt](#)) and automated presentation creation with avatars.
- Formulated go-to-market strategy via business plans and financial models. Showcased the MVPs at VivaTech (Paris) and Bits & Pretzels (Munich) to attract partners and investors to generate interest in 20+ leads.

SwitchOn

Bangalore, India | June 2023 – Sep 2023

Computer Vision Engineer

- Engineered real-time one-shot anomaly detection pipelines using embedding similarity in memory-bank architectures. Also, led the design of the transformer-based OCR module into the anomaly detection stack.
- Achieved more than 99% accuracy with less than 1% false positives in defect identification of goods on manufacturing lines.
- Performed remote rollout and on-device validation of models on Linux embedded hardware to ensure robust edge inference.

PUBLICATIONS & PROJECTS

- Looking In Through Removal: Interpretability Analysis of Physical Priors in IL Policies** (Submitted in CoRL): Introduced a causal interpretability framework to identify which physics properties govern distinct stages of manipulation task execution.
- Dimitra: Audio-driven diffusion model for expressive talking head generation** ([IJCV](#)): Innovated a diffusion transformer for talking head generation, leveraging audio-to-3D latent embeddings to render accurate lip-sync videos.
- Generalizable Deepfake Detection based on Visual Speech Recognition (VSR)** ([IEEE T-BIOM](#)): Engineered a zero-shot model harnessing representation learning on VSR latent space, achieving 30% detection improvement.
- Semantic Segmentation through Transfused Convolution and Transformer** ([Applied Intelligence](#)): Designed hybrid U-Net + Vision Transformer to capture local and global context, improving nuclei and skin lesion image segmentation by 10%.
- Local Distributional Smoothing for Noise-invariant Fingerprint Restoration** ([BIOSIG](#) - Oral Presentation): Developed novel adversarial training algorithm for GAN based denoising models, boosting fingerprint matching by 20%.
- Cross Domain Consistency Fingerprint Denoising** ([IEEE Sensor Letters](#)): Proposed novel unsupervised loss exploiting unlabeled data to improve fingerprint identification by 19%.
- Amazon ML Challenge** (Nationwide-India, ~5000 teams, Secured 2nd Place, [Leaderboard: Team Fishes](#)): Engineered an ensemble of fine-tuned BERT and RoBERTa models with clustering-based quantization to predict product dimensions.

SKILLS

Technical: Python, C/C++, SQL, Bash, Linux, Git, AWS, Conda, Jupyter, Docker, HPCs, GPUs, Distributed & Parallel Compute.

Machine Learning: PyTorch-CUDA, Tensorflow, Keras, Pandas, Numpy, Scikit-Learn, OpenCV, Transformers, Generative AI, Diffusion, GANs, LLMs, Deep Learning, Neural Networks, Computer Vision, Natural Language Processing, HuggingFace.

ACHIEVEMENTS

- Published 5+ peer-reviewed articles with over 128 citations, an h-index of 4, and an i10-index of 2 ([Full Pub list - Scholar](#)).