

Anjali Kaushik

akaush39@asu.edu | [LinkedIn-anjali-kaushik-87312a17b](https://www.linkedin.com/in/anjali-kaushik-87312a17b) | [GitHub-anjalikaushik20](https://github.com/anjalikaushik20) | anjalikaushik20.github.io | +1(623)-2700-615

EDUCATION

Arizona State University: Master of Science, Computer Science, Aug'24 – Present, GPA: 3.72/4

Tempe, AZ, USA

Vellore Institute of Technology: B.Tech., Computer Science and Engineering, Jul'18 – May'22, GPA: 8.99/10

Bhopal, MP, India

EXPERIENCE

Research Assistant

Tempe, AZ, USA

RISE Labs at ASU, *Thesis student under Dr. Wenlong Zhang*

January 2025 – Present

- Reinforcement Learning for Assembly Planning with the Burr Puzzle: Developed an offline reinforcement learning pipeline for complex mechanical assembly, introducing a novel vertex-based action space for collision-free, gravity-aware planning. Improved planning efficiency with pruned n-step lookahead and greedy rollout strategies and am currently reducing the sim-to-real gap by implementing the approach on the ABB IRB 14000 YuMi robot for bimanual manipulation.

Knowledge Discovery & Data Mining Group at ASU

March 2025 – June 2026

- Built a GRPO training environment to study mode collapse in LLM math-reasoning datasets, analyzing output diversity and error patterns to improve dataset robustness for synthetic and real math data via a vLLM based pipeline.
- Developed a simulation-to-decision framework for supply chain optimization by integrating a Mixture Density Network (MDN), an LSTM-based simulator, and an LLM-based decision maker to improve decision quality in supply chain management.

Associate Analytics Engineer

Bengaluru, KA, India

Shell Private Limited Company

August 2022 – August 2024

- Automated complex financial-data workflows using Alteryx and SQL, reducing recurring team workload by 40% and standardizing reporting and exception handling.
- Built a Python-based computer vision pipeline for ROI detection and classification for lube-monitoring image analysis, collaborating with global teams and running experiments on Linux-based systems.
- Automated geo-mechanical experiments with Python, reducing report turnaround time by 50% per lab technician.
- Built a process automation framework for business-critical disaster recovery deployments across 1000+ applications, reducing manual coordination effort by roughly 3 business days per team member.

PUBLICATIONS & PATENTS

DELTA: Variational Disentangled Learning for Privacy Preserving Data Reprogramming

ICDM, 2025

- A two-phase variational learning framework for privacy-preserving feature transformation, combining reinforcement learning with disentangled VAE. Demonstrated robust performance across eight real-world datasets, improving downstream task utility by ~9.3% while reducing sensitive attribute leakage by ~35%.

MITRA (Patent Pending)

- Marine Inspection Tool for Rating and Assessment: A solution in lube monitoring used for image classification with 90% accuracy. The project - Shell Lube-Monitor leverages deep-learning, high-resolution image analysis and a vast engine data archive to provide granular insights and early warnings, helping over 500+ ships daily.

Advanced Deepfake Detection using Inception-ResNet-v2

ICCIS, 2023

- Proposed a deepfake detection pipeline combining MTCNN-based face extraction with Inception-ResNet-v2 and custom dense layers in a sequential model, achieving 91.41% accuracy on the DFDC dataset with an interactive GUI.

PROJECTS

Implementation of Flash Attention

January 2025 – May 2025

- Implemented FlashAttention-style attention optimization for transformer inference (Llama 3) using Triton kernels, reducing latency and memory bandwidth via operator fusion and SoftMax tiling. Built attention modules to minimize HBM transfers and improve memory access efficiency, increasing throughput.

Self-Supervised Dense Point Tracking System

August 2024 – December 2024

- Developed a benchmark for point tracking in videos with induced atmospheric turbulence, enhancing the DINO-Tracker with RAFT-based optical flow refinement. Diagnosed robustness gaps in DINO-Tracker, revealing a 32 % average-Jaccard drop under moderate turbulence and mapping failure modes in occluded, high-blur frames.

Swin Transformer for Vision Tasks

August 2024 – December 2024

- Trained Swin Transformer for classification, localization, and segmentation on ChestXRay14, NODE21, and ChexMask datasets to optimize pipelines and improve performance in medical image analysis. Achieved 72.64% accuracy training from scratch and 81.78% accuracy with ImageNet pretraining on ChestXRay14 classification.

TECHNICAL SKILLS

Languages: Python, C++, SQL; **Frameworks/Libraries:** PyTorch, NumPy, Pandas; **Systems/Tools:** Linux, Git, Docker, Nsight Systems/Compute, Netron, SCALE-Sim

Research Areas: Deep Learning, Computer Vision, Image Processing, Large Language Models, Reinforcement Learning;

Performance/Inference: FlashAttention, vLLM, Transformer Inference

Cloud: AWS (EC2, S3, Lambda, SQS, ECR, CloudWatch, IAM, VPC)