Bhavya Minesh Shah

Phone: +1 (623)-332-9509 | Email: bshah43@asu.edu | LinkedIn | GitHub | Research | Portfolio

EDUCATION

MS in Robotics and Autonomous Systems (AI)	August 2024 – May 2026	**Arizona State University**	**GPA: 3.4 / 4.0**

Relevant Courses: Artificial Intelligence, Machine Learning, Deep Learning, Deployment and Experimentation of Robotic Systems, Perception in Robotics, Generative AI, Robotic Systems.

B.Tech in Computer Science and Engineering	July 2019 – July 2023	**Vellore Institute of Technology**	**GPA: 3.12 / 4.0**

Relevant Courses: Artificial Intelligence, Machine Learning, Deep Learning, Image Processing, Natural Language Processing, Parallel and Distributed Computing, Microcontroller and Microprocessor Interfacing.

PROFESSIONAL EXPERIENCE

Graduate Intern, AI Researcher	September 2025 – November 2025	**Translational Genomics Research Institute**	Phoenix, Arizona

- Ongoing co-op.

Helios Scholar, AI Researcher	May 2025 – August 2025	**Translational Genomics Research Institute**	Phoenix, Arizona

- Worked on Large Language Models (LLMs) and Vision Language Models (VLMs) based Tumor Boards.
- Scheduled jobs on High Performance Computing clusters using SLURM job scheduling and software level libraries like HuggingFace transformers and accelerate that enabled remote training and leveraged multi-GPU capabilities of the cluster.
- Used metrics like Rouge and BLEU score to measure the correctness of the diagnosis generated by medically pre-trained and instruct tuned LLMs like MedGemma, MedLlama 3, Qwen 3, etc given a medical case.
- Trained and tested multimodal and multi-input models using CLIP image encoder that can take in multiple images and texts like MRIs, X-ray reports and CT scans along with the patient histories to solve cases.
- Implemented advanced resampling techniques like Attention Pooling, Perceiver Resampling and Unified Embedding Decoder Architecture to reduce the number of input image tokens to the decoder-only LLM.
- Quantized and optimized model training and testing using LoRA, quantization and using mixed precision with HuggingFace PEFT and BitsAndBytes.
- Presented my findings and work in the form of a presentation and scientific poster.

Artificial Intelligence Engineer	August 2023 – May 2024	**Accurate Industrial Controls Pvt Ltd**	Pune, India

- Being the first of employees in the AI department, I setup procedures and practices of use, hardware like multi-GPU servers, network connections and software like CUDA, CuDNN, WSL2, GPU compatible versions of libraries like PyTorch and TensorFlow 2.0, etc.
- Trained, optimized and deployed open-source models like YOLOv8 for object detection and WaSR for semantic segmentation using multi-GPU finetuning for perception module of autonomous boat project achieving a mAP at 0.5 of 96.4% and mAP at 0.95 of 93.2%.
- Designed the B-Star algorithm, improving the efficiency of the Autonomous Vehicle's path-planning system, achieving a 300% speed increase over A-Star Algorithm for real-time navigation. It is designed to balance between exploration of an unseen environment and quickly finding a path to the goal state which can be tweaked using a single parameter.
- For generalized anomaly detection pipeline, I optimized open-source models like YOLOv8, PatchCore, EfficientAD and EasyOCR by changing the backend from PyTorch to ONNX and TensorRT and deployed them using Nvidia Triton Server running as a Docker container over the host network, achieving real-time performance and scalability, addressing the issue of remote non-parallelized processing request.
- Collaborated with web development team for the integration of the anomaly detection project with a web application created using MongoDB, Express.js, React.js, Node.js to view real-time streams from anomaly detection pipeline using RTSP protocol.
- Enhanced data fusion accuracy and prevented data loss in multi-sensor systems by replacing microcontroller clocks with UNIX timestamps, addressing synchronization issues and improving real-time integration reliability.

Machine Learning Intern	September 2022 – March 2024	**Swasthya AI**	Mumbai, India

- Developed an NLP pipeline to clean and process unstructured medical data, automating data extraction from OPD records and prescription notes, reducing manual effort, and improving efficiency.
- Addressed inaccuracies caused by medical abbreviations and inconsistencies by standardizing terminology and using BERT and LLMs for Named Entity Recognition, improving model accuracy and reducing false positives by 80%.
- Deployed machine learning models on AWS and Azure, optimizing performance for cloud environments to handle large datasets efficiently and ensure scalability. Also used open-source OCRs like PyTesseract and cloud-based OCRs to extract data from medical image documents.

TECHNICAL SKILLS

Programming Languages: Python, C/C++, Java, JavaScript, MATLAB, PHP, HTML, CSS, JQuery, AJAX, Bootstrap, TypeScript, Shell Scripting
Database Technologies: SQL, NoSQL, MongoDB, PostgreSQL, Redis.
Frameworks: Django, Flask, FastAPI, LangChain, ROS Melodic and Noetic, ROS2 Humble, TensorFlow, PyTorch, Keras, MediaPipe, Selenium, Nvidia Triton Inference Server and Client, ONNX, TensorRT, HuggingFace Transformers, Accelerate, PEFT, Numpy, OpenCV, LangChain, Ollama, Nvidia Rapids, CuML.
Simulation/Modelling: Gazebo, MATLAB Simulink, RViz.
Development Tools: Git, Github, CI/CD Pipelines, Jupyter Notebook, Google Colab, VS Code, Pycharm, Docker
Cloud: AWS, Microsoft Azure, Google Cloud Platform, SLURM.
Hardware: Nvidia GPU, Intel CPU, Raspberri Pi, Arduino Uno, ESP32, Nvidia Jetson Nano, Nvidia Jetson Orin, Intel NUC.

PROJECTS

Maze Solver Robot	November 2024 - December 2024

- This project is about solving an observable environment like a maze (on 2D paper plane) viewed fully through an overhead camera that helps detect boundaries of the maze and plan a path. A robot arm with a laser pointer for an end effector is used to trace the path planned. URDF files are used to create digital twins and simulations and calculate forward and inverse kinematics of the robotic arm helping in testing in a simulated environment rather than on expensive real robotic hardware.

Attention-based Transformers	June 2024 - July 2024

- This is a set of 3 Transformer models incorporating Attention mechanism for tasks like Natural Language Translation, Image Captioning and Text Summarization. I deployed the Image Captioning Transformer on a web application locally using Flask for the backend and TensorFlow 2.0 for serving the model. Flask takes in POST request with an image as payload and passes it to the TensorFlow model of the Transformer that generates a caption which is passes as a response to the request.

BStar Path Planning Algorithm	February 2024 - April 2024

- Developed 'BStar', an enhanced path-planning algorithm that significantly outperforms the traditional A* algorithm in speed. Designed and implemented a comprehensive framework to demonstrate its application and simulate real-world scenarios. Incorporated both static and dynamic obstacles, allowing customization to simulate various environments using open-sourced Python libraries.

Warehouse Simulation for Turlebot3 Patrolling Robot	March 2024 - April 2024

- This is a simulation I developed for RAS 598 course using Gazebo and integrated it with ROS2 Humble and Turtlebot3. I designed the world physics using SDF files and Gazebo Ignition and wrote custom nodes and launch files to spawn the Turtlebot3 in the simulated environment and navigate it to a specified goal. The simulated world publishes stream via cameras in the warehouse which on detection of suspicious activity request the Turtlebot3 to check the area out.

EXTRACURRICULAR ACTIVITIES

Team Leader	March 2025 – May 2025	**T.O.M at ASU**	Arizona State University

- Designed a Fall-Alert Ring (custom assistive devices) for Miss Arizona Wheelchair 2025 as she had issues with picking up fallen objects.
- Integrated the Bluetooth module of Arduino Nano with the one on her phone to send real-time SOS alerts to nominated contacts.

Director of Design	August 2025 – December 2025	**T.O.M at ASU**	Arizona State University

- Guided 4 teams in created assistive devices for corresponding need-knowers from ideation to deployment phase.
- Overlooked all the technical events and workshops.

Technical Officer & Mentor	August 2025 – December 2025	**AI Society**	Arizona State University

- Gave a technical presentation on Image Segmentation covering Semantic, Instance, Panoptic and Referred and Reasoning Segmentation.
- Mentored 2 teams on technical projects related to Robotics, Reinforcement Learning, LLMs and Computer Vision.