Tvisha Shah

401-499-7341 tvisha_shah@brown.edu linkedin.com/in/tvisha-v-shah github.com/shahtvisha tvisha.tech

Education

Brown University *MS, Electrical & Computer Engineering* **Sep 2025 – May 2027**
GPA: 4.0/4.0, Coursework: Intelligent Robotics, Sequential Decision Making, Reconfigurable Computing Providence, RI

Mukesh Patel School of Technology, NMIMS *B.Tech, Computer Engineering* **Sep 2021 – May 2025**
CGPA: 3.81/4.0 Mumbai, India

Relevant Work Experience

Brown University / Carney Institute of Brain Sciences *Research Assistant* **Sep 2025 – Present**

- **Developed a multimodal longitudinal modeling framework** over 8 months of passive wearable data (Oura Ring, 49 participants) using regression, mixed-effects, and lagged mixed-effects models to characterize stress–sleep coupling with individualized physiological baselines rather than population-level thresholds.
- **Identified lagged effects of stress physiology on sleep architecture**: reduced HRV predicts multi-night degradation in sleep efficiency and deep sleep; extended via state-space models to infer latent cognitive–physiological stress states. **First-author CCN 2026 acceptance.**

Georgia Institute of Technology *Research Assistant* **Jun – Dec 2024**

- **Built a scalable Economic Entity Extraction pipeline** using long-context LLMs on financial documents, producing a benchmark of **20,139 QA pairs** across 5 complexity levels annotated with Label Studio.
- **Designed an NLP algorithm** for RDF-based knowledge graph ontologies, automating multi-difficulty query generation to **benchmark 16 LLMs** on long-context financial reasoning.

Fractal Analytics *Project Intern* **Apr – Sep 2023**

- **Led end-to-end development** of a no-code workflow orchestrator from system design through production deployment, adopted across **5+ teams** company-wide for real-time operational workflows.
- **Generated $150K+ in FTE savings** by shipping automations that freed 30% of IT technicians' time for higher-value work.

Jai Kisan *Software Development Engineer Intern* **Jan – Mar 2023**

- **Built an error analytics dashboard** surfacing insights from **1M+ payment gateway data points**; drove vendor fixes that achieved a **94% reduction** in payment gateway errors.

Projects

Autonomous City Mapping · Multi-Agent Active SLAM | *PyTorch, Mamba SSM, PPO, OSMnx* **May 2026**

- Designed a cooperative multi-agent RL system in which **3 agents**, each equipped with a **Mamba SSM policy optimised via PPO**, jointly map an unknown urban environment using only onboard 360° LiDAR. Agents share a **Bayesian log-odds occupancy grid** updated in real time via Bresenham ray-casting, with a zone-based lawnmower planner and BFS frontier coverage directing exploration.
- Motivated the choice of **Mamba SSM over transformer attention** by its O(L) vs. O(L^2) sequence complexity; PPO updates propagate gradients through **32-step sequence chunks** enabling long-horizon credit assignment. Trained on OpenStreetMap data from **4 cities** with 50% procedural-map mixing; achieved **90.9% peak region coverage** on Kendall Square, Boston (multi-agent) and **91.0% map accuracy** across 10 held-out episodes (single-agent evaluation).

Self-Improving Robot Policies · Learning from Demonstration | *PyTorch, SAC, HER, Diffusion Policy* **Apr 2026**

- Constructed a complete LfD pipeline on FetchPickAndPlace-v4 to resolve covariate shift and goal distribution gap **without any online expert access**: BC cold-start → iterative **Self-Improvement with Hindsight Experience Replay** → **SAC+HER** online RL → distillation into **Diffusion Policy** and **ACT** with 16-step action chunking for deployment-ready inference.
- Trained SAC+HER to ∼**98% task success** over **1.5M environment steps**; distilled the resulting expert into **Diffusion Policy** and **ACT** architectures via 16-step action chunking, producing deployable inference-time policies that require no access to the RL reward signal.

Zero-Shot 3D Semantic Affordance Mapping | *PyTorch, CLIP, Open3D* **Mar 2026**

- Developed a zero-shot affordance localisation system fusing **CLIP semantic feature lifting** with query-conditioned geometric priors — surface normals, planarity scores, and local curvature — to identify interaction regions in 3D point clouds from free-form natural language, requiring no task-specific training data. Achieved a **+35% IoU improvement** over the CLIP-only baseline on LASO (CVPR 2025); a systematic ablation across **4 affordance categories** quantified each prior's contribution and surfaced novel sensitivity to query phrasing.

Technical Skills

ML / AI: PyTorch, TensorFlow, Scikit-Learn, PPO, SAC, HER, Mamba SSM, Diffusion Policy, ACT, CLIP, ViT, LoRA, RAG, LLMs, VLMs
Robotics & Vision: ROS2, Bayesian SLAM, 3D Point Cloud Processing, Affordance Mapping, Object Detection, Segmentation, OpenCV
Statistical Modeling: State-Space Models, Mixed-Effects Models, Longitudinal Analysis, Time-Series Modeling
Systems & Dev: Python, C, C++, JavaScript, TypeScript, FastAPI, React, Docker, MongoDB, MySQL, Git, Linux