Aditeya Prajapati

aditeyak@usc.edu +1 (213) 883-4625 Aditeya Prajapati AD1PRAJAPATI

Education

University of Southern California *Jan 2026 – Dec 2027*
M.S. in Computer Science(Artificial Intelligence) *Los Angeles, CA*

- *Relevant Courses:* Analysis of Algorithm (CSCI 570), Foundations of Artificial Intelligence (CSCI 561), Web Technologies (CSCI 571)

University of Mumbai *Aug 2020 – May 2024*
B.Tech in Information Technology *Mumbai, India*

Research Interests

- **Policy Routing & Multi-Policy Evaluation:** Dynamically composing multiple VLA and control policies during execution via offline RL, and designing efficient evaluation pipelines for tabletop and real-world manipulation.
- **Embodied AI & Generalist Robot Policies:** Scaling and optimizing Vision-Language-Action (VLA) models for open-world, real-world robotic manipulation.
- **Dynamic Policy Composition & Routing:** Leveraging offline reinforcement learning to dynamically orchestrate and switch between heterogeneous control policies at runtime.
- **Robustness Under Distribution Shift:** Designing failure detection mechanisms, risk probes, and conformal prediction frameworks to handle deployment-time visual, environmental, and linguistic variations.
- **Data Augmentation & Sim-to-Real Transfer:** Utilizing contrast set perturbations and simulation data to calibrate robust safety probes for real-world robotic rollouts.

Academic Works

[1] **Motion Complements Appearance: Pose-Augmented Self-Supervised Fusion for Isolated Sign Language Recognition**
A. Prajapati, J. Thomason
In preparation. Target: CVPR 2026.

[2] **SWAP: Stepwise Action Policy Routing for Vision-Language-Action Models**
M. Das, **A. Prajapati**, A. Anwar, J. Thomason
Submitted to CoRL 2026.

[3] **SAFECAST: Robust Failure Detection for VLA Policies with Contrast-Set Training and Calibration**
H. B. Rajaprakash, **A. Prajapati**, R. Xue, A. Anwar, J. Thomason
Submitted to CoRL 2026.

[4] **Topic Modeling and Opinion Analysis on Hyperloop Twitter Data Using Latent Dirichlet Allocation**
K. Pandit, **A. Prajapati**, S. Jain, S. Fernandes, A. R. Shandliya, M. Saeed
To appear in Proceedings of the Helmholtz AI Conference (HAICON 2026).

[5] **Sentiment Analysis of Indian Hyperloop Tweets for Opinion Inversion Prediction**
K. Pandit, S. Jain, P. Sharma, **A. Prajapati**, U. Chaurasiya, M. Saeed
To appear in Proceedings of the Helmholtz AI Conference (HAICON 2026).

[6] **A Resilient Network Blueprint for Wireless Communication in Hyperloop Systems: Addressing Challenges with 5G and Wi-Fi 6**

A. Shandilya, V. Dongre, **A. Prajapati**, K. Pandit, M. Thakkar, H. Jain, P. Rawool, H. Kasturiwale
Accepted for publication in SN Computer Science, Springer Nature.

[7] **Assessment of the Societal Acceptance of Hyperloop System**
A. Prajapati, K. Pandit, S. Jain, A. Maurya, S. Fernandes, S. Singh, P. Sharma, A. Sharma, M. Thakkar, U. Chaurasiya, H. Jain, A. Shandilya
Socio Economic Research Submission for European Hyperloop Week (EHW), Zurich, Switzerland (2024).

[8] **Case Studies: Responsible IIoT Implementations: Smart Bin Waste Prediction Using Machine Learning**
A. Prajapati, S. Uchil, V. Patwa, N. Kapadia, R. G. Mehta
In Industrial Internet of Things for Responsible Technology, CRC Press, Taylor & Francis Group, pp. 54–64 (2025).

Research Experience

University of Southern California *Jan 2026 – Present*
Research Assistant — GLAMOR Lab, PI: Prof. Jesse Thomason *Los Angeles, CA*

- Designing **SWAP (Stepwise Action Policy Routing)**, a framework for dynamically composing multiple VLA policies during execution. Formulated policy routing as an **offline reinforcement learning** problem; trained a routing critic that selects the most appropriate policy (π_0, OpenVLA, or keypoint-based controllers) at each decision step given the current observation.
- Evaluated SWAP on real-world DROID manipulation tasks and Libero simulation; achieved absolute improvements in real-world task success up to **33%** while reducing successful-trajectory action-step length by **28.3%** versus fixed-policy and routing baselines. Built the multi-policy evaluation infrastructure for tabletop manipulation on the Libero benchmark, with real-world deployment on the **Franka Panda** arm.
- Engineered **SAFECAST**, leveraging contrast-set perturbations and conformal prediction to significantly improve hidden-state probe failure detection under deployment-time distribution shifts. Co-designed a state-of-the-art dual-stream (MediaPipe Conformer + self-supervised transformer) **Sign Language Recognition** model achieving a +6% accuracy jump on the ASL-Citizen benchmark.
- Trained **DualLip**, an acoustic-grounded speech-recognition system using HuBERT distillation and contrastive loss on LRS2 via SLURM; containerized and deployed these keypoint-finetuned Auto-AVSR visual-speech models on **Google Cloud Platform (GCP Vertex AI, GKE, Cloud Storage)** for real-world clinical application by patients at the **Keck School of Medicine**.
- Developed a real-time policy selection algorithm using **Implicit Q-Learning (IQL)** for adaptive, optimal operator decision-making during wildfire evacuation dialogues; benchmarked against random and fixed-policy baselines, achieving higher success rates across diverse resident profiles.

University of Southern California *July 2026 – Present*
Research Assistant — LIRA Lab, PI: Prof. Erdem Biyik *Los Angeles, CA*

- Proposed **LLM-regularized social attention** for trajectory forecasting: distilled an LLM's per-agent *influence* estimates into the model's learned attention via a **KL prior** (offline-cached Qwen2.5, zero inference-time cost).
- Cut displacement error (**ADE**) **−7%** at low data on Argoverse 2 vs. no-prior and geometric baselines; aligned attention with interaction-critical agents **3–4**× more than distance heuristics (**17%→64%**).
- Showed via **prompt ablations** that framing dominates model scale (*trajectory-influence* vs. *visual-salience*); hardened LLM decoding from **48%→<1%** parse-failure; validated across Argoverse 2 and ETH/UCY.

Industry Experience

FYND (Shopsense Retail) *Jun 2024 – Dec 2025*
Machine Learning Researcher – PixelBin *Mumbai, India*

- **Generative AI Systems:** Co-engineered an enterprise **AI Photoshoot Automation** system utilizing BirefNet-based apparel segmentation and Flux Fill LoRA, enabling synthetic human model generation with prompt-conditioned control over pose, ethnicity, and camera viewpoint via **ControlNet**, reducing physical model-generation overhead costs by **40%**.
- **Scale & Cloud Pipeline Engineering:** Built an **AI Background Generator** pipeline fusing EraseBG with specialized Flux LoRA inpainting to process over **1M+** production images; developed a lightweight **Flask API** for local microservice testing before deploying the production pipeline end-to-end on **Google Cloud Platform (GCP)**, optimizing image rendering speeds.
- **Computer Vision Architecture:** Deployed U2Net for pixel-perfect background removal, achieving an **F1-score of 0.94**, and architected a product-specific view detection classifier using structured **Chain-of-Thought (CoT) prompting** to anchor LLMs as deterministic vision classifiers with **91.5%** production accuracy.

Reliance Industries *Dec 2022 – Jan 2023*
Data Analyst Intern – Petroleum E&P *Navi Mumbai, India*

- **Time-Series Forecasting:** Analyzed **500GB+** of high-frequency live telemetry sensor streams from **15+** nodes at the KG D6 offshore facility to build predictive LSTM deep learning architectures, minimizing prediction error (**RMSE**) by **14.8%** for regional sales estimation.
- **Applied NLP Deployment:** Developed and deployed an internal conversational AI assistant utilizing TensorFlow and TextBlob text processing modules; trained the network on **10K+** structured operational FAQs to automate user queries and optimize support throughput by **35%**.

Projects

Data-Driven Socio-Technical & Viability Analysis for Hyperloop Networks *EHW 2024*

- **Automated Data Extraction & Infrastructure:** Architected a high-throughput, cross-platform pipeline using UiPath Robotic Process Automation (RPA) and custom Apify scrapers to bypass rate-limits, curating a clean corpus of **10,422** hyperloop-specific text entries spanning 2013–2024. Engineered automated pre-processing wrappers for part-of-speech (POS) dependency parsing using spaCy and NLTK.
- **Latent Topic Modeling & NLP Semantics:** Deployed a **Latent Dirichlet Allocation (LDA)** pipeline with Gensim and pyLDAvis to mathematically map latent clusters of public concern, achieving a peak **Topic Coherence score of 0.55**. Integrated multilingual **BERT models** and VADER for deep polarity evaluations, designing contrastive text-pair probes to isolate and filter pragmatic linguistic noise like sarcasm.
- **Opinion Inversion Predictive Engine:** Formulated *Opinion Inversion* (asymmetric ideological divergence between origin posts and subsequent replies) as a high-dimensional binary classification problem. Mitigated severe minority-class imbalance using **ADASYN** (Adaptive Synthetic Sampling) and trained a **Random Forest Classifier** on dense engagement feature matrices, reaching **84.9% cross-validation accuracy** and an **ROC-AUC score of 0.75**.
- **Parametric Socio-Economic Simulation:** Fused the derived NLP sentiment metrics with industrial transport economics to run parametric structural cost simulations over a proposed **120 km industrial corridor**. Executed comprehensive financial benchmarking against mass transit networks (Mumbai Metro, IndiGo Airlines, Indian Railways) to isolate realistic operational expenditures (OPEX) and calculate break-even horizons under volatile capital costs (CAPEX).

Technical Skills

Programming & Core: Python, C++, SQL, Bash/Shell, CUDA, Linux, Git

Frameworks & Deep Learning: PyTorch, TensorFlow, Hugging Face Transformers, Scikit-learn, torchaudio, torchvision, Distributed Training

Generative AI & Large Models: Large Language Models (LLMs), Vision-Language Models (VLMs), Foundation Models, Fine-Tuning (LoRA/QLoRA), Diffusion Models, Prompt Engineering (Chain-of-Thought), OpenVLA, Flux

Computer Vision & Audio: Computer Vision, Object Detection, Image Segmentation, Automatic Speech Recognition (ASR), Multi-Modal Fusion, OpenCV, MediaPipe, DINOv2

Robotics & Reinforcement Learning: Reinforcement Learning (RL), Offline RL, Robot Manipulation, Motion Planning, Robot Simulation, Libero, DROID, Franka Control Interface (FCI)

Cloud, MLOps & Infrastructure: Google Cloud Platform (GCP), Vertex AI, Kubernetes (GKE), Docker, Containerization, SLURM (GPU Cluster Management), Flask APIs, RESTful Microservices, CI/CD

NLP & Data Engineering: Natural Language Processing (NLP), Topic Modeling (LDA), Text Analytics (spaCy, NLTK), Data Pipelines (ETL), Robotic Process Automation (RPA/UiPath), Data Structures, Pandas, NumPy

Leadership & Extracurricular Activities

- **Global Competition Leadership (European Hyperloop Week):** Directed the socio-technical research division to secure **2nd Place globally** (out of 24 international teams) at EHW 2024 in Zurich, Switzerland. Previously presented the physical Control & Telemetry Subsystems at EHW 2023.
- **Peer Reviewer:** Served as an invited academic reviewer for the *International Journal of Communication Systems* (2024), evaluating novel optimization and machine learning routing algorithms for IoT networks.
- **Amazon ML Summer School Scholar:** Selected for an exclusive cohort mentored directly by Amazon Data Scientists (2023); completed intensive coursework in Deep Neural Networks, Probabilistic Graphical Models, and Unsupervised/Supervised Learning.
- **Government Tech Delegate:** Represented Thakur College of Engineering and Technology at a national conference hosted by the **Ministry of Electronics and Information Technology (MeitY)** in New Delhi, focusing on strategic resource management for large-scale technology initiatives (2023).