

Sahil Shah

sahilshah379@gmail.com \diamond sahilshah379.github.io

EDUCATION

The University of Texas at Austin

Masters of Science in Electrical and Computer Engineering

Bachelor of Science in Electrical and Computer Engineering, Minor in Robotics

Austin, TX

May 2025 – May 2026

Aug. 2021 – May 2025

PUBLICATIONS

Shah, S., Sharan, S. P.*, Goel, H.*, Choi, M., Munir, M., Pasula, M., Marculescu, R., & Chinchali, S. (2026). NeuS-QA: Grounding Long-Form Video Understanding in Temporal Logic and Neuro-Symbolic Reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*. [Paper]

Shah, S., Goel, H., Narasimhan, S. S., Choi, M., Sharan, S. P., Akcin, O., & Chinchali, S. (2025). A Challenge to Build Neuro-Symbolic Video Agents. In *International Conference on Neuro-symbolic Systems*. [Paper]

Sharan, S. P.*, Choi, M.*, **Shah, S.**, Goel, H., Omama, M., & Chinchali, S. (2025). Neuro-Symbolic Evaluation of Text-to-Video Models using Formal Verification. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. [Paper]

Choi, M., Goel, H.*, Omama, M.*, Yang, Y., **Shah, S.**, & Chinchali, S. (2024). Towards Neuro-Symbolic Video Understanding. In *European Conference on Computer Vision*. **Oral Presentation** [Paper]

Choi, M.*, Sharan, S. P.*, Goel, H., **Shah, S.**, & Chinchali, S. (2025). We'll Fix it in Post: Improving Text-to-Video Generation with Neuro-Symbolic Feedback. *arXiv preprint arXiv:2504.17180*; *Under review at ICLR 2026*. [Paper]

Choi, M.*, Yang, Y.*, Bhatt, N. P.*, Gupta, K., **Shah, S.**, Rai, A., Fridovich-Keil, D., Topcu, U., & Chinchali, S. (2025). Real-Time Privacy Preservation for Robot Visual Perception. In *Transactions on Machine Learning Research*. [Paper]

Munir, M., Goel, H., Wei, X., Choi, M., **Shah, S.**, Bhardwaj, K., Whatmough, P., Chinchali, S., & Marculescu, R. (2025). ObjectAlign: Neuro-Symbolic Object Consistency Verification and Correction. *arXiv preprint arXiv:2511.18701*; *Under review at CVPR 2026*. [Paper]

Yu, F.*, **Shah, S.***, Mittal, Y., Bessler, P., Mohsin, A., Geng, J., Vats, A., & Soloveichik, D. (2025). COFFEE: a High-Performance Approach to Convex Optimization for Thermodynamic Equilibrium Computations. In *2025 Systems and Information Engineering Design Symposium*. **Best Paper** [Paper]

RESEARCH EXPERIENCE

Graduate Research Assistant

Swarm Lab

Aug. 2022 – Present

Austin, TX

- Designed *NeuS-QA*, a system that integrates temporal logic with large vision-language models to answer long-form video queries; achieves consistent improvements of **8–12%** accuracy on temporally complex categories compared to state-of-the-art baselines; **accepted at AAAI 2026**
- Proposed a neuro-symbolic framework for video agents—reactive systems that analyze video streams to reason about events and make decisions—capable of temporal reasoning and structured event understanding, addressing the limitations of current deep learning vision systems; **presented at NeuS 2025**
- Designed *NeuS-V*, a synthetic video evaluation metric that assesses text-to-video alignment by formally checking a video automaton against a desired logic specification; *NeuS-V* demonstrates a higher correlation to human evaluations by over **5 times**; **presented at CVPR 2025**
- Introduced *NSVS-TL*, a method that leverages object and action recognition models to identify scenes of interests given a video stream alongside a linear temporal logic specification, improving the F1 score of complex event identification by **9-15%** on state-of-the-art datasets; **presented at ECCV 2024 as an oral presentation**
- Developed *NeuS-E*, a zero-training video editing and refinement pipeline that improves text-to-video generation by using neuro-symbolic feedback to detect and correct semantic and temporal inconsistencies; enhances prompt alignment by nearly **40%** across diverse text-to-video models and datasets; **submitted to ICLR 2026**
- Developed *PCVS*, a method to conceal sensitive objects in real-time video; models video streams as Markov Chains to predict the sensitivity of entities and uses conformal prediction to ensure over **95%** of streams are deemed safe under logical privacy constraints; **accepted at TMLR**

Founding Engineer

May 2025 – Present

Coargus

Austin, TX

- Built an AI-native IDE for generative video with a model-agnostic orchestration layer, timeline-driven scene graph editor, reusable asset conditioning, and git-style version control to support modular and iterative video generation
- Designed a personalized asset library using LoRA-based conditioning to generate consistent, reusable entities within a single video, enabling continuity of actions and styles across scenes in multi-scene workflows
- Developed timeline-driven scene editing with segment-level re-generation and a metric-guided enhancement loop that evaluates temporal adherence, entity consistency, and motion smoothness to iteratively optimize video quality

INDUSTRY EXPERIENCE

Architecture Intern

May 2024 – Aug. 2024

NVIDIA

Santa Clara, CA

- Proposed a system to model the key interchange for Confidential Computing in the Blackwell architecture
- Implemented 3 commands that manage session keys for an active SPDM session in the PCIe's functional model
- Simulated and tested the startup sequence and trust establishment for next-gen GPUs on Unix using Bash

Software Engineer Intern

Aug. 2023 – Dec. 2023

Tesla

Palo Alto, CA

- Achieved a 22% increase in space efficiency by using IPC to manage shared data between cores in the RTOS
- Designed UDS routines in C++ to facilitate the transfer and storage of HV system data and speed up diagnostics
- Conducted extensive testing in Linux, using Docker to emulate the Cybertruck Battery Management System

Hardware Development Engineer Intern

May 2023 – Aug. 2023

Amazon Web Services

Seattle, WA

- Developed an instrumentation system through running an EC2 SSM agent on a Raspberry Pi to validate servers
- Streamed real-time SMBus data from a microcontroller and the motherboard, ensuring data accuracy within 6%
- Engineered a robust API for efficient lookup table querying, streamlining tests for AWS Outposts servers

Embedded System Intern

May 2022 – Aug. 2022

GXC

Austin, TX

- Wrote a Python CLI on Linux to manage the operations for the data stored on a portable edge device
- Employed REST API to develop a version control tool for Jenkins to optimize product releases and save time
- Verified hardware involving I²C data transfer using an oscilloscope and increased EEPROM read speeds by 450%

STUDENT INVOLVEMENT

Texas Robotics | *Student Ambassador*

Feb. 2024 – Present

- Provided mini-lectures on ongoing research across multiple labs, introducing prospective students to robotics
- Led tours of the lab facility, offering insight on projects and facilitating introductions between students and faculty
- Demonstrated active research threads, enhancing students' understanding of the practical applications of robotics

Longhorn Racing Electric | *Controls Engineer*

Aug. 2021 – June 2023

- Placed 7th overall in the Formula SAE Electric competition in June 2022 and competed in the 2023 season
- Mentored and trained students throughout the year, teaching them how to design large-scale software for a vehicle
- Constructed a high-level system architecture of the electrical control signals across multiple devices on the car
- Developed core software for the Vehicle Control Unit that tracks and updates the state space of the vehicle

IEEE | *Member*

Aug. 2021 – Present

- Coordinated events designed to connect students and provide academic assistance within Texas Engineering
- Mentored underclassmen by offering academic support and advice on navigating university life

TECHNICAL SKILLS

Languages: Python, C/C++, Rust, Java, Lua, Bash

Frameworks: PyTorch, OpenCV, NumPy, CUDA, ROS, CARLA, vLLM, Git, Docker, Linux, LaTeX