Sahil Shah

sahilshah379@gmail.com ◇ sahilshah379.github.io

Education

University of Illinois Urbana-Champaign — Urbana, IL
Doctor of Philosophy in Computer Science — *Aug. 2026 – May 2030*
- Advisor: Prof. Heng Ji

The University of Texas at Austin — Austin, TX
Masters of Science in Electrical and Computer Engineering — *May 2025 – May 2026*
- Advisor: Prof. Sandeep Chinchali

Bachelor of Science in Electrical and Computer Engineering, Minor in Robotics — *Aug. 2021 – May 2025*

Publications

Shah, S.*, Sharan, S. P.*, Goel, H.*, Pasula, M., Hebbalae, A., Choi, M., & Chinchali, S. (2026). CrossView: Can Vision-Language Models Reason Across Cameras?. In *European Conference on Computer Vision*. [Paper]

Goel, H.*, Sharan, S. P.*, **Shah, S.**, Choi, M., An, J., Grauman, K., & Chinchali, S. (2026). Incentivizing Vision Language Models to Search for Long Video Question Answering. In *European Conference on Computer Vision*. **Spotlight Presentation** [Paper]

Munir, M., Goel, H., Wei, X., Choi, M., **Shah, S.**, Bhardwaj, K., Whatmough, P., Chinchali, S., & Marculescu, R. (2026). ObjectAlign: Neuro-symbolic Object Consistency Verification and Correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. [Paper]

Liang, S., **Shah, S.**, Zhou, C., Sharan, S. P., Goel, H., Sanyal, A., Chinchali, S., & Datta, G. (2026). RT-NeuS: Towards Real-Time Neuro-Symbolic Video Understanding via Adaptive Temporal Verification. In *International Conference on Neuro-symbolic Systems 2026*. [Paper]

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]

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]

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 2025*. [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*. [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.*, 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. *Transactions on Machine Learning Research*. [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]

Research Experience

Graduate Research Assistant — Aug. 2026 – Present
Blender Lab — *Urbana, IL*
- Investigating hierarchical skill learning for robots, building reusable action ontologies that enable adaptive skill transfer across functionally similar tasks in visually distinct domains

Graduate Research Assistant Aug. 2022 – May 2026

Swarm Lab *Austin, TX*

- Proposed grounding long-video understanding in neuro-symbolic reasoning, combining temporal logic and model checking to formally isolate query-relevant video segments for more precise and interpretable VLM reasoning
- Investigated RL post-training for multi-turn agentic video search in VLMs, using verifiable per-step rewards that check each retrieved clip for query-relevant visual evidence instead of only rewarding final-answer accuracy
- Proposed a multi-camera video QA dataset using scene graphs to generate spatiotemporally grounded questions

Industry Experience

Research Scientist Intern May 2026 – Aug. 2026

ByteDance *San Jose, CA*

- Built a program synthesis system for evaluating Seedance 2.0 outputs, where an LLM writes detector programs over composable video skills (OCR, ASR, optical flow, lip sync) to reproduce an annotator's verdict and complaint
- Integrated a meta-learning loop that synthesizes new video skills when existing detectors cannot explain flagged clips, raising annotator agreement from 0.25 to 0.41 as each cycle expands the set of detectable failure modes
- Introduced an evolving video defect taxonomy that adapts across model releases, achieving 89% agreement

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

Technical Skills

Languages: Python, C/C++, Rust, Java, Lua, Bash
Frameworks: PyTorch, OpenCV, verl, NumPy, ROS, CARLA, vLLM, Git, Docker, Linux, LaTeX