

Tharun Kumar Tiruppali Kalidoss

+1 (310) 961 6434

tharunk@princeton.edu

Senior at Princeton University

tharunkumar.xyz

Education

Princeton University
B.S.E. Computer Science
Aug 2022 - May 2026

Relevant Courses: Theory of Algorithms w/ Dr. Robert Tarjan, Distributed Systems, Real Analysis, Abstract Algebra, Stochastic Systems, Reinforcement Learning, Economics in Computing, Intro ML, Data Structures and Algorithms, Discrete Math, Linear Algebra.

Technical Skills

Python, Typescript, Javascript, Java, SQL, NoSQL, Pytorch, JAX, TensorFlow, Azure, AWS, Mobile Development, Full-stack Development, Swift, Machine Learning, CI/CD, Kubernetes

Relevant Experience

Undergraduate Researcher — Google DeepMind, Princeton, NJ **Sep 2025 – Present**

- Researched and implemented novel State Space Model architectures for efficient long-context language modeling under **Prof. Elad Hazan** in the [Princeton Google DeepMind Lab](#).
- Pretrained 150M, 500M, and 1B parameter LLMs**, conducting ablation studies on various SSM-attention hybrid architectures and ultimately matching/outperforming Transformer models across all settings.
- Built internal PyTorch-compatible packages** enabling reusable, efficient SSM-based modeling for other researchers.

Software Engineering Intern — Amazon Web Services, Seattle, WA **Jun 2025 – Sep 2025**

- Built low-level Python and Go bindings to monitor server hardware health and automate management of a **50,000+ node fleet**, ensuring regional availability by maintaining host redundancy while repairing failed nodes.
- Increased server availability by 15%** and **eliminated 30% of on-call tickets** after shipping to production with CI/CD pipelines
- Received top performance evaluation for contributions to large-scale infrastructure automation.

Member of Technical Staff — UniversalAGI, San Francisco, CA **Jan 2025 – Apr 2025**

- Built LLM agents for Fortune 500 companies and governments. Flying back/forth between Princeton, SF, and Dubai.
- Developed a **state-of-the-art Text-to-SQL LLM agent** achieving **high accuracy on Spider 1.0 benchmark**; successfully deployed the solution **on-premise for the UAE government** w/ Kubernetes on private employee data.
- Developed an **LLM agent** for Apple that builds no-code applications w/ data integrations to external services (e.g. Slack, Gmail).

Research Engineer — Hooglee, San Francisco, CA **May 2024 – Dec 2024**

- Worked with **Sebastian Thrun** to incubate a funded consumer video generation startup, now led by Bichen Wu.
- Engineered synthetic media generation system utilizing wav2lip fine-tuning for audio-driven animation, FLAME 3D facial model integration, and MediaPipe pose estimation, deployed on iOS.
- Developed parameter-efficient LoRA training pipeline with Modal Labs for personalized image generation models, for real-time image synthesis in social media applications

Projects

Learned Diffusion Priors for Sparse-View 3D Gaussian Splatting **Mar 2025**

- Designed a render-repair-retrain pipeline for sparse-view 3D Gaussian Splatting that mitigated novel-view artifacts and measurably improved reconstruction fidelity after retraining.
- Built a ~10k-pair aligned dataset and evaluated diffusion-based and CNN restoration methods, improving PSNR/SSIM, yielding sharper and more stable novel views

Multi-Modal Comic-to-Animation Engine **Nov 2024**

- Built an end-to-end comic-to-anime generation engine using multi-modal agents, Stable Diffusion in-painting, and stable video diffusion with automated lip-syncing and ElevenLabs voice synthesis
- Worked with comic creators with 200k+ readers, landed two animation companies as clients, and got YC interview for this project. (Demo: https://www.youtube.com/watch?v=Hb_hQ3CVrM)

Contextual Bandits for Realtime News Personalization **Oct 2024**

- Implemented a contextual multi-armed bandit over learned user embedding profiles to personalize recommendations for rapid preference adaptation.
- Deployed on iOS app that aggregates posts from multiple news sources, filters low-signal content, and generates concise, sourced summaries via LLM.

RL-based Obstacle Avoidance for Drone Swarms (AGI House Hackathon Winner) **Jul 2024**

- Formulated obstacle avoidance problem as an MDP, and trained a CNN via PPO to teach drones in a swarm to navigate various environments while maintaining swarm formation, used OpenAI gym environment. Won Project Eagle hackathon at AGI House.