Akshath Raghav Ravikiran

765-404-8121 araviki@purdue.edu linkedin.com/in/akshathrr akshathraghav.github.io

Education

Purdue University, West Lafayette **Cumulative GPA: 3.63**

Master of Science in Electrical and Computer Engineering *January 2026 – December 2026*

Bachelor of Science in Computer Engineering *August 2022 – December 2025*

Coursework: Programmable Accelerator Architectures, AI Hardware Design, Computer Design & Prototyping ASIC Design, Microprocessors & Microcontrollers, Advanced C Programming

Honors: Eli Shay Electrical Engineering Scholarship, Purdue-OUR Scholar, UIUC HDR Fellowship

Awards: ECE Senior Design Award, ISF DUIRI Award, Outstanding Sophomore in VIP

Organizations: Purdue SoCET (Team Lead), ML@Purdue (Officer), Semiconductor Student Alliance (Board)

Technical Skills

Proficient in Python, Java, C, SQL, and Shell scripting. Familiar with Unix-like systems.
Experienced in ML (CV, NLP, DL), Cloud Technologies (AWS, GCP), and DevOps tools (Linux, Docker, k8s).
Solid understanding of CPU/GPU micro-architecture, and GPU programming with Triton and CUDA.
ARM microcontroller programming and interfacing, and usage of peripherals including timers, interrupts, ADC, DAC.

Experience

SWE Intern - Agile Toolchain (FG-240) **May 2025 – Present**

BMW Group IT *Munich, Bavaria, Germany*

- Building autonomous agents (AI4SoftwareQuality) to crawl enterprise web apps, auto-generate QA Tests and remotely execute them for in-browser-testing and regression validation.
- Deployed platform connectors for ATC Offerings into GAIA using SCP (AWS) utilities. Integrating multi-agent orchestration into JoyCode Frontend in collaboration with South Africa Hub.
- Engineering domain-specific hybrid-RAG (Github/Confluence) for teams to integrate searchable knowledge graphs, exposed via MCP, into their IDEs and workflows.
- Developed custom GAIA Applications to automate AI4Testing use-cases for QA Testing Productivity.

Teaching / Research Assistant **July 2024 – Present**

Purdue SoCET (PI: Prof. Mark Johnson, Purdue-ECE) *West Lafayette, IN, US*

- Designed and functionally verified the lockup-free D-Cache, 2MB scratchpad and systolic array controllers that are integrated within the AMP1 Tensor Core. Implementing optimized PyTorch backends for memory-access-aware kernels.
- Enhanced the AFTx07 RISCV core with the Zicond Ext. for macro-fusion of conditional logic/arithmetic sequences.
- Helped develop semiconductor-design-specific curriculum under the CASCADE Apprenticeship Program w/ Synopsys.

Research Assistant **September 2024 – March 2025**

I-GUIDE (NSF Award 2118329 - PI: Dr. F. Baig, UIUC) *Urbana-Champaign, IL, US*

- Enhanced the Distributed Aging Dams Analysis system by distributing HPC workflows on **GCP Dataproc** to scale geospatial risk mapping across 345 dams in the US. Work was accepted and presented at IGUIDE Forum 2025.
- Parallelized spatial correlation metrics using **Apache Sedona** for identifying vulnerable populations at flood risk.
- Accelerated raster ingestion and indexing workflows using **Apache Spark** for modular & scalable job distribution.

ML Research Intern **August 2023 – April 2024**

Google X Duality Lab (PI: Prof. James Davis, Purdue-ECE) *West Lafayette, IN, US*

- Re-engineered the **MaskFormer** segmentation model to natively run on Google-funded ($100K) GCP TPUs.
- Contributed to a technical white paper, providing implementation guidance for TPU-focused HW/SW co-design.
- Integrated auxiliary losses across modules and performed hyperparam. tuning to increase Panoptic Quality by 25%.
- Performed distributed model training on RCAC's Gilbreth cluster (A100 GPUs), using SLURM and a self-scripted MLOps workflow to design and schedule experiments, helping track model improvements.

Data Science Intern **March 2023 – July 2023**

Ambee (Datair Technology Pvt. Ltd.) *Bangalore, India*

- Built a global forest-fire forecasting system that remains integrated into Ambee's proprietary **MongoDB** dashboard.
- Co-authored a white paper outlining our decisions that enhanced a Boosted Random Forest Regressor's performance to match government forecasts like NIFC (USA) & CWFIS (Canada) with 90%+ accuracies.
- Configured **AWS Glue** to automate the ETL workflows that handle the post-data-retrieval jobs, **Apache Spark** to parallelize transformations and ECS to containerize training scripts.

Projects

BoilerNet – Compute-Enabled Mini-NAS *Jan 2025 - April 2025*

- Designed and assembled custom PCBs, within 3D-printed enclosures, to form a NAS with hot-swappable compute blades & memory slots. Received Purdue ECE's Senior Design Award for Spring '25.
- Single-ported TCP transmission with in-network concurrency interleaves OPs across a 2-way-switched bus topology.
- Provides a MLQ-based priority scheduler to allow for out-of-ordering atomic operations and Layer-7 load-balancing.
- Supports INT8/FP16 quantized MobileNetV2 models, leveraging SPIRAM-aware memory allocation and DMA-friendly I/O pipelines through TFLite Micro.

RISCV Five-Stage Multicore Processor *February 2025 - April 2025*

- Pipelined processsor implementing branch prediction, forwarding/hazard logic and dual-core cache coherency.
- Memory controller arbitrates read/write memory accesses to external memory and supports variable-latency access.
- Two-way associative data cache and direct-mapped instruction cache implement write-back and LRU policies.
- Utilizes functional and contrained random verification for verifying functionality, simulated using ModelSim.
- Synthesizable through Quartus Prime targeting the DE2-115 development board, with **60MHz target frequency**.

tinySpeech – Speech Recognition on Edge Devices *July 2024 - December 2024*

- Reproduced TinySpeech word-recognition models for the VSDSquadronMini board; open-sourced weights and scripts.
- Achieved 91% precision benchmarks without access to DarwinAI's proprietary source code.
- Custom Quantization-Aware Training pipelines allow for training in 2/4/8 bit-widths and PerTensor/PerWeight scales.
- Developed an embedded C-based inference engine, optimized for 8-bit precision targeting RISCV-EC architecture.

Reproducible AI Software (RAIS) *January 2024 - May 2024*

- Awarded "Outstanding Sophomore in VIP" by Purdue for **leading** team to evaluate ML research reproducibility.
- Automated Google Cloud pipelines to scrape & analyze over 3K+ research projects for validating our decision-tree.
- Presented our reusability scoring system and reproducibility categories at the OUR Expo.

GrammarFlow *February 2024 - April 2024*

- Published a package to **ensure parseability** of LLM outputs into structured dataclasses for agentic workflows.
- Introduced **GNBF** grammar generation to incorporate regex patterns into token sampling algorithms for *llama.cpp*.
- Achieves parsing success of **99%+** across variants of **Llama, Mistral and Dolphin** on popular reasoning datasets.