

AKSHATH RAGHAV RAVIKIRAN

☎ 765-404-8121 ✉ araviki@purdue.edu [in linkedin.com/in/akshathrr](https://www.linkedin.com/in/akshathrr) [github akshathraghav.github.io](https://github.com/akshathraghav)

Education

Purdue University, West Lafayette

Junior Year || August 2022 – December 2025

Bachelor of Science in Computer Engineering

GPA: 3.81

Honors: Eli Shay Electrical Engineering Scholarship (F'23 - S'24), Dean's List and Semester Honors (F'22 - S'24)

Papers

V.Purohit, W.Jiang, **A.R.Ravikiran**, J.Davis. [A Partial Replication of MaskFormer in TensorFlow on TPUs for the TensorFlow Model Garden](#), arXiv Technical Report, 2024

R.Jain, **A.R.Ravikiran**, P.Katti. [Time-Driven Fire Risk Forecasting: Leveraging Historical Trends for Enhanced Seasonal Modeling](#), ResearchGate White Paper, 2023

Experience

TensorFlow Model Developer

August 2023 – Present

Google X Purdue Duality Lab (Prof. James Davis, Purdue-ECE)

West Lafayette, IN

- Re-engineered the [MaskFormer model](#) to natively run on Google Cloud TPUs. Validated module-level precision on funded GPUs/TPUs; responsible for the PR to Model Garden from guideline compliance to documenting results.
- Engineered the **inference module**, including the PQ metrics implementation to work with TF 2.x's task flow. Integrated code for auxiliary losses to address **inability of loss convergence**, and fixed run-time issues across computes.

Data Science Intern

March 2023 – July 2023

Ambee (Climate Intelligence)

Bangalore, India

- Built a global [forest-fire forecasting system](#), from prototype to production, that remains integrated into Ambee's proprietary **API** dashboard. Developed modularized components implemented within an end-to-end **AWS** lifecycle (Spark, Glue, S3) ensuring tri-monthly forecast generation, complimented by robust **ETL pipelines** (Docker).
- Co-authored a white paper outlining unique strategies targeting historical Fire Weather Index, enhancing a **Boosted Multi-Target RF Regressor's** performance to surpass government forecasts (NIFC & CWFIS) in risk classification.

Data Science Lead

December 2022 – April 2023

Lightning Wildfire Lab (Prof. Yuan Wang, Purdue-EAPS X NASA)

West Lafayette, IN

- [Supervised codebase development](#) for short-term wildfire forecasting; Responsible for bundling netCDF data on the basis of spatio-temporal features to package into **LSTM**, **CNN** and **ConvLSTM** deep learning models.
- Automated API requests for large scale fire data collection from USGS and Copernicus; Developed scripts using **Xarray**, **GeoPandas** and **netCDF4** to process Landsat and GeoTIFF data from NASA/NOAA satellites.

Projects

GrammarFlow

February 2024 - April 2024

- [Published a package](#) to **ensure parseability** of LLM outputs into structured dataclasses, improving robustness of AI agents. Designed to address issues with LangChain on local LMs, GrammarFlow achieves parsing success of **99%+** across fine-tuned variants from **Llama**, **Mistral** and **Dolphin** on popular reasoning datasets.
- Introduced **GNBF syntax** generation to interface with efficient token sampling, incorporating regex patterns into decoding algorithms. Enables **complex use-cases**, including multi-grammar and nested sequence generation.

Reproducible AI Software (RAIS)

January 2024 - May 2024

- Awarded **"Outstanding Sophomore in VIP"** for directing efforts to practically evaluate AI/ML research reproducibility. Categorized reproducibility types and quantified reusability with a scoring system.
- Implemented ETL pipelines to analyze over 3000 academic & corporate research projects for validating our [data-driven pipeline](#). Employed **LLM chains integrated within repository file structures** to assess structural cohesion, documentation, community engagement, etc.

Amazon OpenSearch Service

September 2023 - December 2023

- Reworked the workload generation process within the [official macrobenchmarking framework](#), enabling custom features for user-defined workloads. **Defined** documentation for creating performance benchmarks on hosted indices.
- Enhanced extraction efficiency (upto 41.65%) by deploying **multi-process** capabilities for simultaneous data retrieval from clusters, ensuring optimal **throughput** and accelerated performance.

Technical Skills

Languages: Python, C, C++, CUDA, Java, JavaScript, MATLAB

Frameworks: TensorFlow, PyTorch, Keras, Xarray, Matplotlib, GeoPandas, OpenCV2, ONNX, Django, Node.js

Tools: Linux, Git, Docker, DVC, MLFlow, ZenML, Kubernetes, Redis, MongoDB, PostgreSQL, Elasticsearch

Cloud Utilities: Google Cloud Console (Compute), Amazon Web Services (S3, Lambda, Spark)