

Venkata Datta Adithya Gadhamsetty

helses.com — linkedin.com/in/datta-adithya — github.com/dat-adi

Email: vgadham1@jh.edu

Mobile: +1 (443) 226-7847

Education

- Johns Hopkins University** Baltimore, Maryland
 - Master of Science in Engineering Degree, Computer Science; GPA: 3.61/4.0* *Expected: May 2026*
 - Courses: Machine Learning, Parallel Computing for Data Science, NLP: Self-Supervised Models, Explainable AI*
- Vellore Institute of Technology** Andhra Pradesh, India
 - Bachelor of Technology Degree, Computer Science; GPA: 3.64/4.0* *August 2023*
 - Courses: Data Structures and Algorithms, Advanced Data Analytics, Web Development, Operating Systems*

Skills

- Languages:** Python, Go, C++, Rust, SQL, Bash, Typescript
- SDE Toolkit:** Docker, Git, Linux, AWS, PydanticAI, Langchain
- Machine Learning Toolkit:** Scikit-Learn, PyTorch, Huggingface, Transformers, vLLM

Experience

- Johns Hopkins University** Baltimore, Maryland
 - Research Assistant @ The Burns Lab (Dr. Randal Burns)* *Jan 2025 - Present*
 - LLM Inference Arithmetic:** Performed a trade-off analysis of the Arithmetic Intensity between computing versus offloading transformer attention scores onto AWS data stores for **1m+ token contexts**.
 - Emergent Sparsity in LLMs:** Extracted unstructured sparse matrices that fall between DenseMM and SpMM structures from TEAL, Wanda and SparseGPT. At present, I'm exploring sparsity ablations to detect emergent patterns for MMM adoption in Meta's Llama models. Utilized **Triton, PyTorch** and **CUDA**.
- Bajaj Finserv Direct Limited** Pune, India
 - Software Development Engineer* *Jan 2023 - Jul 2024*
 - Data Migration Platform:** Built the DRIFT unified ETL orchestration framework managing 300+ jobs with real-time observability, reducing operational costs by 66% using **Python, DuckDB, Polars**, and **PyArrow**.
 - DRIFT API:** Architected the development of a high-performance API delivering job lifecycle analytics and improving developer productivity by 20% with **FastAPI** and **PostgreSQL**. This project consolidated job ingestion, management, processing, delivery and notification management into a single interface.
 - Customer360:** Engineered a data mining solution to **automate the generation of over 100+ customer portfolios** for employee dashboards, streamlining operations across business units, using BeanieODM on AWS DocumentDB.
 - Offer Generation Engine:** Designed data pipeline processing personalized offers at scale for 30M+ customers with feature engineering, risk scoring, and ML-driven prioritization using **Parquet, AWS S3**, and **Typer CLI**.
- Open Source Community (Mozilla)** Amaravati, Andhra Pradesh, India
 - President* *May 2021 - May 2022*
 - GitBook Admin Guide:** Designed an administration guide to improve onboarding efficiency for new executive board members by upto 20%. Leveraged **GitHub, CI/CD, GitBook** and **Markdown**.
 - Major League Hacking:** Hosted day-long MLH hackathons at VIT, mentored and guided over 200 students into adopting **Linux, Docker** and the **Rust** toolchain.

Projects

- MapApp (Golang, Preact, TypeScript, Cesium, DuckDB):** Built a large-scale data visualization application empowering real estate agents with geospatial data during HopHacks 2025, earning a spot in the **Top 10**. We interacted with clients who were limited by the Google MyMaps 20k dot limit, so we built our application to surpass it by **supporting up to 1M+ dots on the globe** with minimal latency.
- Synthetic Data-Boosted SLMs (PyTorch, HuggingFace, QLoRA):** Enhanced TinyLlama 1.1B model through **stratified synthetic data sampling** and topic modeling, achieving 10% better performance than OLMoE-1B-7B (8.79 vs 9.71 perplexity) while using 6x fewer parameters in resource-constrained environments.
- Orkanskader (PyTorch, DuckDB, scikit-learn, Shapely, Fiona):** Modeled hurricane impact using **historical Atlantic hurricane data (1979-present)**. Built **Random Forest** for radius prediction (**MAPE 24.2%**) and **LSTM** for trajectory forecasting (**MAPE 3.5%**). Integrated **geospatial analysis** for residential property damage estimation.
- Novel Crawlers (Python, FastAPI, Docker, Redis):** Programmed a web spider to scrape a webnovel with over **14 million words**, deriving statistics and to generate an EPUB for my favorite light novel.

Publications

- A Study on Advanced Networking using Intelligent Systems Applications:** Presented at the I2CT Conference 2023, this study explores the recent advances and implementations of intelligent systems in SDN-based IoT architectures.
- ASM-SDN:** Published in The Journal of Supercomputing, this paper discusses the development of an automated station migration system in a cluster-based heterogeneous SDN environment.

Honors and Awards

- Winter Season of Code:** OSCHub, Best Team Lead Award, December 2020.
- Raman Research Award:** publication of ASM-SDN, October 2023.