Sidharth Swamy

sidharthswamy1@gmail.com sidharthswamy.com linkedin.com/in/sidharth-swamy

Education

Stony Brook University **Stony Brook, NY**
Accelerated M.S. in Applied Math & Statistics (Operations Research) *Expected May 2027*
B.S. in Applied Math & Statistics, Minor in Computer Science *Expected May 2027*

Summary

Educated in classical data analysis and optimization fundamentals. Experienced collaborating with an international team in a fast-paced and lean startup in energy and automotive. Passionate about applying data-driven methods to real-world challenges.

Coursework: Analysis of Algorithms, Operations Research, Graph Theory, Computational Geometry, Machine Learning

Experience

Data Science Intern
Zeno - EV Startup *April 2025 – September 2025*

- Built a scalable **data pipeline** for **Modbus** and monitoring devices at battery swap stations, ingesting via **NNG sockets** and streaming over **UDP/TCP with zlib** compression into **Parquet** files on **AWS S3**.
- Built an interactive **analytics layer** (DuckDB + Voila) to detect outages and forecast station-level performance, providing inputs for operations research models in fleet reliability, capacity planning, and energy optimization.
- Designed an ingestion system for **CAN** logs over UDP into **MF4** files, supporting reliability and predictive maintenance analysis.
- Automated build pipelines for firmware flashing tools on Windows with **Github Actions**.

Projects

MNIST Convolutional Neural Network | *Python* **November 2024**

- Built a CNN from scratch with no external libraries to classify **60,000** handwritten digits, achieving **92% accuracy** on test data despite hardware constraints.
- Designed a **modular architecture** for dynamic configuration of layers, neurons, filters, activations, and learning rates.
- Implemented core operations including forward/backpropagation and batch gradient descent, optimizing for efficiency in low-resource settings.

Logistic Regression for Microchip QA & Student Admissions | *Python, NumPy* **April 2025**

- Built logistic regression classifiers for pass/fail prediction, including gradient descent training and accuracy evaluation.
- Added L2-regularized logistic regression and polynomial feature mapping to learn non-linear decision boundaries for the microchip QA dataset.
- Visualized data and decision boundaries with Matplotlib to debug models and compare linear vs. regularized non-linear classifiers.

Multiple Regression & Gene–Environment Interaction Modeling | *Python, R, scikit-learn* **May 2025**

- Applied Bonferroni screening, stepwise regression, and Lasso to isolate significant variables
- Evaluated interaction and nonlinear terms; selected a parsimonious main-effects model based on BIC and diagnostics.
- Positioned findings in the G×E literature context to highlight interpretability and limitations.

Linear Regression Profit Predictor | *Python, NumPy* **March 2025**

- Plotted predictions and forecasted city-level profits on proposed business expansions.
- Built a modular pipeline for data loading, cost computation, and optimization via batch gradient descent with manual gradients.
- Trained model using hand-coded partial derivatives of the MSE function over 1,500 iterations with convergence tracking.

Boeing Database Application and Web Scrapers | *SQL, Python, Microsoft Access* **November 2024**

- Designed an **SQL** database with **45,000+ records** on Boeing aircraft, flights, airports, and incidents.
- Built **Microsoft Access** forms, queries, and reports to generate insights and streamline aviation data analysis.
- Wrote Python scrapers to extract data from **10+ web sources**, reducing manual entry workload.

Skills

Languages: Python, Java, R, SQL
Developer Tools: Windows Subsystem for Linux (WSL), Gurobi
Technologies/Frameworks: Git, Google Cloud Platform (GCP), Microsoft Access