

Varun Panuganti

425-295-1418 | varunp5@uw.edu | <https://www.linkedin.com/in/varun-panuganti/> | <https://github.com/VarunP3000> | portfolio: <https://varunp3000.github.io>

EDUCATION

University of Washington <i>B.S. ACMS: Data Science & Statistics; B.S. Informatics</i> GPA: 3.77 Coursework: CSE 416 (ML), CSE 373 (DS&A), Data Structures, Statistical & Scientific Computing, Data Science Foundations	Seattle, WA <i>Sep 2023 – Jun 2027</i>
---	---

EXPERIENCE

LLM Uncertainty Quantification (UQ) <i>University of Washington</i>	Oct 2024 – Present
<ul style="list-style-type: none">Developing processes to integrate a chain of LLMs for data annotation, allowing users to input a set of prompts, select LLMs, and configure confidence thresholds and token limits.Developing an intuitive React web interface for file uploads using JavaScript, HTML, and CSS, enabling CSV file analysis.Implementing backend processes in Node.js (Express) and Python to manage API calls, LLM workflows, and confidence scoring.	<i>Seattle, WA</i>
Computer Science Instructor <i>ICODE</i>	Jul 2024 – Apr 2025
<ul style="list-style-type: none">Guided students in designing self-driving robots using VEX and sensors, culminating in a robot that consistently navigated a maze.Mentored K-12 students in game development with Python (Pygame) and Unreal Engine, helping them build a 90%-complete racing game with AI opponents and physics simulation.Instructed Java fundamentals (OOP, recursion, algorithms, and data structures), leading students to develop a Spring-based application for data processing and visualization.	<i>Sammamish, WA</i>

COURSE PROJECTS

CSE 416 — CIFAR-10 Image Classification	PyTorch, CNNs, GPU — 2025
<ul style="list-style-type: none">Implemented CNNs (conv→ReLU→pool→dropout) with data augmentation and efficient DataLoaders.Trained on GPU with early stopping and LR scheduling; tracked train/val curves and performed error analysis on misclassifications.	
CSE 416 — House Prices (Tabular ML)	scikit-learn, pandas — 2025
<ul style="list-style-type: none">Prepared data (imputation, encoding, scaling) and splits; established linear baselines and regularized models (Ridge/LASSO).Ran hyperparameter sweeps with validation curves; monitored RMSE to ensure generalization and prevent leakage.Interpreted coefficients and feature effects; summarized error sources and key failure modes with plots/tables.	
STAT 534 — Bayesian Linear Models (Numerics)	C/C++, LAPACK/LAPACKE, GSL — 2025
<ul style="list-style-type: none">Computed closed-form log marginal likelihood for $[1 \mid A]$ via stable solves and log-determinants (all on the log scale).Validated against an R baseline on <code>erdata.txt</code> ($n=158$, $p=51$); reproduced check cases precisely.	
STAT 534 — Logistic Regression Model Search (MC3)	R, AIC/BIC, MCMC — 2025
<ul style="list-style-type: none">Built forward/backward subset selection with robust <code>glm</code> wrappers (convergence/NA handling) on a 60-feature dataset.Implemented MC3 over add/remove-one neighbors with neighbor-count-corrected MH; compared best models across 10 chains.Filtered candidates with rcdd linearity tests to avoid separation/unreliable AIC; documented stability vs. greedy baselines.	
CSE 373 — Shortest Paths Finder	Java, Graphs (Dijkstra/A*) — 2024
<ul style="list-style-type: none">Designed adjacency-list graphs and a binary-heap priority queue; implemented Dijkstra's and A* with admissible heuristics.Handled large graphs via careful PQ updates (reinsert for decrease-key) and parent-pointer reconstruction.Wrote JUnit edge-case tests; benchmarked densities and validated $O((V+E) \log V)$ runtime and memory behavior.	

TECHNICAL SKILLS

Languages: Python, C/C++, Java, R, SQL, JavaScript/TypeScript, HTML/CSS

ML/AI: PyTorch, scikit-learn, pandas, NumPy, SciPy, Matplotlib, Jupyter

Stats/Numerics: BLAS, LAPACK/LAPACKE, GSL; optimization, log-sum-exp, numerical stability

Data/Systems: Git/GitHub, MPI; testing (JUnit/`pytest`), Linux CLI

Web: React, Node.js/Express, Flask