

Vaishnavi Gupta

(347) 437-1831

vg222@cornell.edu

vaishnavigupta.com | [github](#) | [linkedin](#)

EDUCATION

Cornell University, MS

Computer Science

August 2022 - May 2024

Cornell University, BA

Computer Science + Math

August 2019 - May 2022

GPA: 4.00, Summa Cum Laude

ORGANIZATIONS

Cornell Design & Tech Initiative
Hortus Form (Horticulture Club)

COURSEWORK

Graduate Algorithms
Distributed Systems
Reinforcement Learning
Large Scale Machine Learning
Probability, Hons. Real Analysis
Compilers, Operating Systems
Graduate Network Theory
Hons. OOP, Discrete Math
Computational Genetics
Algebra, Linear Algebra

SKILLS AND TECHNOLOGIES

Software Design + ML

C++ Python (incl PyTorch)

Scala Kotlin Java

Web Development

Backend and frontend experience

JS/Typescript (React, Node.js)

SQL HTML5 CSS

Kubernetes Bazel

AWARDS

Compilers Bakeoff Winner:

Awarded the best compiler for CS 4120 as a team of 3

Cornell Prize for Teaching

Excellence: \$500 reward, given to 1 out of graduating class

Grace Hopper Full Funding

Indian National Olympiad in Informatics Finalist

EXPERIENCE

Databricks - Software Engineer Intern

May 2023 - August 2023

- › Developed an interactive Python debugger for Databricks Notebooks, working end-to-end in Python, Scala, and Typescript.
- › Improved over comparable offerings like Jupyter, by emphasizing features like concurrent user support and persistence through long-running jobs. Also made open source contributions to ipykernel.

Meta AI - Software Engineer Intern

June 2022 - August 2022

- › Made vision transformers for video machine learning tasks more compute-efficient. Implemented adaptive token sampling modules from research papers in PyTorch, and incorporated into downstream tasks like Reels understanding.
- › New trained models showed a 25% decrease in transformer inference time latency and flops, with a 10% final decrease in the production model.

Instagram - Software Engineer Intern

June 2021 - August 2021

- › Worked on infrastructure in the Instagram Suggested Users team, to increase efficiency and recommendation quality. Used C++ and Python, as well as SQL, Hive, and A/B testing frameworks for statistical analysis.
- › Optimizing cache refresh scheduling based on user activity prediction led to a 30% decrease in CPU cycles during peak hours, and an overall 10% decrease.

Cornell University - Student Researcher

January 2021 - Present

- › As a member of [CUAI](#) ↗, I'm working on problems related to improving the performance of Graph Neural Networks on non-homophilous graph structures.
- › Also exploring online learning algorithms, advised by Prof. Robert Kleinberg.

Cornell University - Head Teaching Assistant

August 2020 - Present

I lead a weekly recitation section, handle the autograder, and help write review materials for **CS 4820 - Algorithms** and **CS 2802 - Hons. Discrete Math**.

PUBLICATIONS

- › Large Scale Learning on Non-Homophilous Graphs: New Benchmarks and Strong Simple Methods (*NeurIPS '21*) [PDF](#) ↗
- › Non-Stochastic CDF Estimation Using Threshold Queries (*SODA '23*) [PDF](#) ↗

PROJECTS

PhyloML - [Demo](#) ↗ [Source](#) ↗

- › A phylogenetic tree generator library in OCaml, demoed via a React frontend.
- › Implemented Bayesian inference MCMC sampling algorithms for heuristic based sequence alignment. Also wrote an XML lexer and parser.