

Adithya Vasudev

U.S. Citizen | [linkedin.com/in/adigvasu](https://www.linkedin.com/in/adigvasu) | github.com/adithya-gv

EDUCATION

Georgia Institute of Technology

Atlanta, GA

Masters of Science in Computer Science (GPA 4.0)

January 2024 - December 2024

Concentration: Machine Learning

Coursework: Machine Learning Theory, Natural Language Processing, Efficient Machine Learning, Parallel and Distributed Computing, Systems for Machine Learning, Deep Learning

Bachelors of Science in Computer Science, Graduated w/ Highest Honors (GPA 3.91)

August 2020 - December 2023

Concentrations: Artificial Intelligence, Theoretical Computer Science

EXPERIENCE

Amazon Web Services

Sunnyvale, CA

Software Engineering Intern (AppStudio Generative AI Team)

May 2024 - July 2024

- Developed chatbot and RAG knowledge base to perform intelligent document Q&A and retrieval using AWS and GraphQL.
- Created automatic evaluation framework to integrate chatbot performance evaluation into end-to-end testing pipeline.
- Rolled out agent in a limited environment, and achieved 90% response accuracy on a query set from real users.
- Designed an LLM API and UX toolkit using React & GraphQL to streamline the architecture of AppStudio's LLM agents.
- Deployed toolkit as an internal library, reducing production latency by 10% and build times by 25%.

Software Engineering Intern (AppStudio Backend Team)

May 2023 - July 2023

- Used TypeScript, React, Recoil, and AWS S3 to design and implement a version control system for an application builder.
- Implemented snapshotting, reversion, and sharing functionality to address a highly requested customer use case.
- Decreased ticket resolution times by over 10% by deploying feature as an internal development tool.

ServiceNow

San Diego, CA

Software Engineering Intern (MetricBase Team)

May 2022 - August 2022

- Used JavaScript and REST APIs to design a dashboard for customers that displayed database health and usage efficiency.
- Reduced time delay in server-client metric transmission by 10% by adding HTTP database endpoints for metrics using Java.
- Designed metric, algorithm, and associated HTTP endpoint to measure data stream usage efficiency using Apache Tomcat.
- Reduced time and memory load for query computation by 25% by adding distributed sampling for diagnostic queries.

Software Engineering Intern (Performance Tools Team)

May 2021 - August 2021

- Developed a tool to simulate server load during feature tests using JavaScript to improve test automation suites.
- Rewrote an internal performance testing framework to remove a feature addition and security block.

RESEARCH & PUBLICATIONS

The EarlyBird Gets the WORM: Heuristically Accelerating EarlyBird Convergence | *Sole Author*

- Accepted to the Efficient Natural Language and Speech Processing Workshop at NeurIPS 2024.
- Researched goal-based pruning techniques for vision models and transformers at Dr. Celine Lin's EIC Lab @ Georgia Tech.
- Reduced convergence time of the EarlyBird algorithm by nearly 10% by heuristically masking gradients during training.

Shotluck Holmes: LLMs for Shot-Level Video Summarization | *Co-Author*

- Accepted to the Workshop on Multimodal Generation and Retrieval at ACM MM 2024.
- Trained large language-vision models to summarize multi-shot video given audio and visual information and annotations.
- Achieved results comparable to or exceeding state of the art with models up to 75% smaller.

PROJECTS

Neuraphonic (hackGT X) | *Python, PyTorch, Scikit-Learn, Google Cloud, Twilio, MATLAB, Praat*

- Developed phone call system to detect signs of Parkinson's disease using a signal processing and machine learning pipeline.
- Trained Vision Transformer and Random Forest ensemble model, with prediction accuracy of 80% over processed audio files.
- Won 2nd place best overall project (\$3000 cash prize) out of 189 total teams.
- Project selected for the Create-X Startup Launch program, awarding \$35,000 in funding for project expansion.

TECHNICAL SKILLS

Languages: Java, Python, C, JavaScript, TypeScript, MATLAB, SQL, LaTeX, C++ (basic)

Frameworks/Libraries: PyTorch, Transformers, Tensorflow, Ray, PyTorch Geometric, DGL, NumPy, Pandas, Matplotlib, Scikit Learn, Keras, OpenMPI, REST, GraphQL, Postman, FastAPI, Dash, Flash, MySQL, Node.js, React, Recoil, Express, Docker, Google Cloud, Google Colab, AWS, ROS

Communication: English (Native), Spanish (Professional), Tamil (Elementary)

Awards: Eagle Scout (Boy Scouts of America), 2nd Place Overall Best Project (hackGT X)