AGASTYA SETH

MS Computer Science | Arizona State University

Tempe, AZ | ☎ +1-602-218-0639 | ✉ agastya.seth@gmail.com | 🌐 agastyaseth.wiki/ | ⭘ github.com/agastyaseth

"Eager to contribute to data science and AI, with a focus on NLP, backed by a solid base in software engineering and machine learning"

EDUCATION

ARIZONA STATE UNIVERSITY (Tempe, AZ) — GPA: 4.0
Masters of Science in Computer Science | Graduate Research Assistant @ ACME Lab — Aug 2023 - May 2025

SHIV NADAR UNIVERSITY (UP, India) — GPA: 3.5
Bachelors of Technology in Electronics and Communication Engineering (Minor in Mathematics) — Aug 2016 - Jun 2020

STANFORD UNIVERSITY (Stanford, CA) — GPA: 4.0
Visiting Undergraduate Student | Data Science & Technology Entrepreneurship — Summer 2018

WORK EXPERIENCE

CADENCE DESIGN SYSTEMS — San Jose, CA
Machine Learning Intern (R&D) - Cerebrus Intelligent Chip Explorer — May 2024 – Aug 2024

- Developed a novel ML approach to predict key performance metrics (TNS and total power) for Verilog RTL designs, bypassing traditional synthesis processes. Automated **feature extraction from ASTs** and implemented **regression models** (XGBoost, MLP, etc.) to provide rapid, accurate design quality assessments. Enhanced design validation with comprehensive unit tests. Overcame challenges like data quality issues and model generalization, significantly improving evaluation efficiency and scalability.
- **Stack:** Python, PyTorch, Transformers, Scikit, Git, TCL, Genus, Verilator, Linux, CSH, Shell

CADENCE DESIGN SYSTEMS — Noida, India
Software Engineer II (R&D) - Quantus Parasitic Extraction — Jul 2020 – Jul 2023

- Developed high-efficiency algorithmic software rooted in **graph theory and geometry processing,** contributing to over 10 significant projects. Improved core libraries and owned & maintained Rccompare tool for multi-million-net design comparisons. Led **AI-driven projects,** enhancing ETA calculator and a novel LLM-based querying system and submitted them at internal conferences.
- **Stack:** C++, Python, Valgrind, Parasoft, ASAN, Perforce, GDB, Perl, Ruby, Linux, Shell, PyTorch, Transformers, Scikit, Git

LIQVID <> SHIV NADAR UNIVERSITY - NLP RESEARCH INITIATIVE ⭘ — Remote (Part-time)
NLP Research Scientist — Aug 2020 – Sep 2021

- Developed an **Automated Essay Scoring (AES)** software for major publishing companies like Wiley with **82% kappa score** through transformer-based deep language models using **DistilBERT** and a novel **ensemble LSTM model** for enhanced explainability & fine-tuning.
- **Stack:** Python, PyTorch, Keras, TensorFlow, AWS EC2, SageMaker, Docker, Kubernetes

VISENZE — Singapore
Data Science Intern — Jan 2020 – Jul 2020

- Managed data-rich projects, improving fashion-attributes models' F1-scores by 15-20% with **augmented data, attention models, & hyperparameter tuning.** Worked on **CNN-based classification & segmentation** models, & innovated a novel pipeline reducing iteration time by 10x with **T-SNE sufficiency-based image parsing.**
- **Stack:** Python, PyTorch, Keras, ONNX, JupyterLab, Docker, BS4, Selenium, JS, HTML/CSS, Git

VIDIMENSIONS — Singapore
Computer Vision Intern — May 2019 – Jul 2019

- Engineered a novel, **lighting-resilient background segmentation model** using distributed **Hough Transform**, achieving **40x faster inference** than deep models for **ARVAS anomaly detection**. Enhanced surveillance systems with an efficient **Faster-RCNN** layer for accurate person and bag detection in camera feeds.
- **Stack:** Python, TensorFlow, Keras, JupyterLab, OpenCV, FastAPI, Git

KEY PROJECTS & PUBLICATIONS

Scalable and Privacy-Preserving CWE Detection and Repair in RTL Designs using LLMs and Graph Neural Networks — Present
ASU ACME Lab

- Leveraging Graph Neural Networks (GNNs) and LLMs to identify and repair Common Weakness Enumeration (CWE) hardware bugs in RTL designs. By analyzing Register Transfer Level (RTL) code and their Abstract Syntax Trees (ASTs), the approach provides a scalable, privacy-preserving solution for robust hardware security.

Wearable Health Monitoring and Signal Processing — Present

ASU ACME Lab
- Collaborating with a startup to develop machine learning models and **adaptive filtering techniques** for enhancing signal fidelity in wearable health monitoring. Analyzing **multimodal sensor data** to mitigate noise and motion artifacts, enabling accurate physiological predictions. Contributing to **algorithm design, signal processing pipelines, and interdisciplinary research** aimed at real-time health monitoring for commercial wearable applications.

Exploring Novel Direct Reward-based Optimization Alignment in LLMs   Present
ASU CogInt NLP Lab
- This project explores a novel alignment technique for LLMs using a reward-based optimization approach. Inspired by **Direct Reward Optimization (DRO)**, the method incorporates a reward model to guide fine-tuning, ensuring the model aligns more effectively with desired outcomes while maintaining robustness across diverse scenarios.

Exploring Reinforcement Learning-Based Alignment in T2V Rectified Flow Models   Present
ASU Vision Language Lab
- Exploring **reinforcement learning** (RL) based **alignment techniques in Text-to-Video (T2V) Rectified Flow Models (RFM)**. By leveraging RL, the goal is to improve the semantic consistency and visual quality of generated videos, ensuring the outputs closely align with the input text prompts, while optimizing the model's performance across diverse scenarios.

Privacy-Preserving ECG Anonymization Using Reinforcement Learning   2024
ASU ACME Lab | UC Irvine Research Collaboration
- Developed a **PPO-based** reinforcement learning framework to **anonymize ECG data**, balancing privacy and utility. Applied both **feature-based and raw data transformations** to minimize patient re-identification while preserving classification accuracy. Explored various feature selection strategies and similarity/distance metrics for optimal RL training. Evaluated privacy effectiveness using **linkage attack methodologies**.

LLM-AID: Leveraging LLMs for Rapid Domain-Specific Accelerator Development 🔗   2024
(In Review) TODAES'25
- This paper introduces LLM-AID, a framework leveraging large LLMs to automate the development of **domain-specific accelerators** (DSAs). It integrates **high-level synthesis (HLS)**, **design space exploration (DSE),** and **symbolic AI** to streamline tasks from generating HDL code to **optimizing power, performance, and area (PPA)**. LLM-AID enhances productivity by reducing manual intervention and enabling iterative refinement, making DSA design more accessible and efficient while addressing challenges like rapid obsolescence and expertise-intensive workflows.

UnSeenTimeQA: Time-sensitive Question Answering Beyond LLMs' Memorization 🔗   2024
ACL 2025
- This paper introduces *UnSeenTimeQA*, a benchmark for time-sensitive question-answering (TSQA) based on synthetic facts, avoiding real-world data contamination. It tests large language models' **temporal reasoning skills** without relying on **pre-trained factual knowledge.** Evaluations reveal that while LLMs handle simpler scenarios well, they struggle with complex event dependencies and parallel timelines, highlighting challenges in **long-range temporal reasoning**.

The Art of Defending: A Systematic Evaluation and Analysis of LLM Defense Strategies on Safety & Over-Defensiveness 🔗   2023
ACL 2024
- This paper introduces the **SODE benchmark**, a dataset and evaluation framework for **assessing safety and over-defensiveness in LLMs**. Through systematic evaluation of various defense strategies, it highlights trade-offs such as the tendency of self-checking techniques to enhance safety while inducing over-defensiveness, and the role of contextual knowledge in compromising safety. The findings aim to guide future improvements in LLM safety mechanisms.

ACHIEVEMENTS

Graduate Research Assistant Scholarship: Granted full tuition waiver as a GRA at ACME Lab under Prof. Krishnendu Chakrabarty
Smart India Hackathon 2019 - Winner: Addressed EV range-anxiety with Tata Motors (Asia's largest Hackathon, 300k participants) 🔗
Google Science Fair 2014 - Regional Finalist: Crafted an Android app for real-time crop prices, empowering Indian farmers 🔗
Shiv Nadar University 70% Merit-Based Scholarship 🔗

ADDITIONAL

Other Publications: (CIC'23) QueryLog: Querying Quantus Output files using LLMs, (CIC'23) Quantus 3DIC Inter-die Extraction (IDX)
Technical Skills: C/C++, Python, SQL, MongoDB, R, JavaScript, Kotlin, Verilog, Linux, Shell, Ruby, Perl, HTML/CSS, JAX, PyTorch, TensorFlow, Caffe, Keras, ONNX, Pinecone, Langchain, AutoML, MLOps, S3, Lambda, SageMaker, EC2
Deep Learning: ViT, T2I, T2V, Diffusion, GANs, LSTMs, BERT, LLMs (LoRA, PEFT, Gradient-based + RL Alignment/RLHF), Transformers, Reinforcement Learning, Retrieval-Augmented Generation (RAG), Agent-based flows, Bayesian Networks,
Software: Jupyter, Valgrind, Android Studio, AWS, Perforce, Git, Docker, Kubernetes, Apache, Spark, Django, MATLAB, Hadoop, Kafka
Certifications/Select MOOCs: Deep Learning Specialization, Full-Stack Web Development, Linear Algebra, Optimization