Mehul Maheshwari

mahesmeh002@gmail.com | 763-600-3041 | LinkedIn | Personal Website | San Diego, Cali

EDUCATION

University of California, San Diego - *Masters in Science, Computer Science, Specialization in AI*　　Sep 2024 - May 2026
- **GPA 4.0;** Relevant Coursework: ML Algorithms, Statistical Natural Language Processing, Web Mining and Recommender Systems, Advanced Algorithms, Working currently in Q Lab (RL+MAS) and in CSG Lab (MAS)

University of Wisconsin, Madison - *Bachelors in Science, Major in Computer Science, Minor in Math.*　　Sep 2020 - May 2024
- Achieved **Comprehensive Honors**, **accomplished by 1.5% of graduates** by earning distinctions in both the CS Major and through the College Letters & Sciences
- Relevant Coursework: Artificial Neural Networks, Artificial Intelligence

SKILLS

Languages - Python, SQL, React, Java, JavaScript / Typescript, HTML / HTML5, CSS, C#, C, C++, webGL, Matlab

Frameworks & Libraries - Tensorflow, PyTorch, Surprise, Hugging Face, Cornac, Juypter, Kaggle, React, Next.js, Bootstrap,

Tools - AWS Tool: Bedrock, EC2, Route53, Amplify, S3, OpenSearch; Github / Github Actions, Gunicorn, Nginx, SQL/ PostgreSQL / MySQL / SQLite, Unity 3D, PowerBI, Snowflake, Docker, Blender, Claude Sonnet 3

Domains - Machine Learning, Multi-Agent Systems, Natural Language Processing, RAG, Reinforcement Learning, Full Stack

EXPERIENCE & RESEARCH

Q Lab, Research Assistant @ UCSD

Advised by PhD Isadora White, under Professor Lianhui Qin　　Dec 2024-**Present**
- Creating RL benchmarking tools for multi-agent collaboration in an embodied open-world environment.

Coding for Social Good Lab, Research Assistant @ UCSD

Collaborating with PhD students Nazanin Sabri and Jude Rayan on Audio Simulacra　　Dec 2024-**Present**
- Designed and optimized a Multi-Agent System (MAS) leveraging Hugging Face and LLaMA to train for real-time conversational audio simulations. Used quantized models and llama-cpp-python to optimize computation.
- Collaborated with a team to integrate SoTA MAS techniques, reducing resource costs for facilitator training.

Collaborative Computing Group, Research Assistant @ UW-Madison

Under Professor Jacob Thebault-Spieker　　Feb 2024 - Oct 2024
- Published research as **first author** at the Association for the Advancement of Artificial Intelligence Conference on Human Computation and Crowdsourcing.
- Developed a full-stack platform with React and Django to streamline experiment creation for crowdsourcing researchers. It will be used in conjunction with ongoing research to analyze dimensions of bias with five star rating systems.
- Migrated existing database from SQLite to PostgreSQL after identifying shortcomings in the production environment.

OPTUM @ UnitedHealth Group, Technology Development Intern.

　　Internship No. 2 (return offer)　　May 2023 - August 2023
- Saved over $300,000 by collaborating with a team of twelve professionals to develop a full-stack in-house software system using React, Spring Boot, Docker, and Kubernetes, which eliminated dependency on third-party software.

　　Internship No. 1　　May 2022 - August 2022
- Modeled an artificial intelligence integrating microsoft azure to help doctors diagnose patient illnesses.
- Automated backlog reporting and data analysis in Java and Javascript, reducing reporting time to investment stakeholders.

University of Minnesota, Research.

Worked under Professor Victoria Interrante. Mentored by Ph.D. student Chris Curry.　　August 2018 - June 2019
　　under Professor Victoria Interrante
- Presented "Enhancing Spatial Perception and Presence in Immersive Virtual Environments," which was **awarded** by the Minnesota Department of Education: Scholars of Distinction - Meritorious.
- Created a 3D virtual environment in Unity which investigated the relationship between the field of view and cybersickness.

PROJECTS

AWS X VCT Hackathon: Esports Manager Challenge 2024 - Submission**,** Report　　Oct. 2024 - Nov. 2024
- Finalist - **top 25** out of 3299 participants by creating a custom chunking strategy for retrieval augmented generation (RAG).
- Implemented RAG for our large language model (LLM) to provide accurate results quickly, and used advanced prompting techniques for quality responses. Used Claude Sonnet through Bedrock, and an ETL in Python.
- Parsed game logs, player profiles, and other unstructured data, along with web-scraped content, to create a relational database to be used by a LLM. I also developed the front facing chatbot to implement the capabilities of the LLM.

Can Item Metadata Combat the Cold Start Problem? - Report　　August 2024
- **Modeled new recommender systems**, beating baseline models by using metadata on the Steam User-Game dataset.
- Our novel SVD + KNN approach beats existing models such as Factorization Machines and SVD++ implemented on Surprise

Sign Language Learner - Report　　August 2024
- Developed a CNN model using TensorFlow for sign language detection in a team; achieved 99.5% accuracy without overfitting, surpassing baseline models that achieved 96% accuracy on a Kaggle dataset parsed with Pandas.
- Created by experimenting with CNN model architectures, and optimizing a ResNet50-based model using transfer learning

Transformer Encoder / Decoder - Report　　Nov 2024
- Developed a fully functional transformer model with attention without the use of external libraries for transformer blocks. Used self-attention mechanisms for word prediction. Also investigates AliBi encodings.