

Huzaifah Bin Wasim

huzaifah.wasim241@gmail.com | +1 (678)-790-4416 | Atlanta, GA 30309

EDUCATION

Master of Science, Computer Science (Machine Learning) (GPA 4.0/4.0) Aug 2024 – May 2026
Georgia Institute of Technology Atlanta, GA

- Specialization in Machine Learning.

Bachelor of Electrical Engineering (Conferred) Sep 2018 – Jun 2022
National University of Sciences and Technology (NUST) Islamabad, Pakistan

- Elected specialization stream in Artificial Intelligence.

PUBLICATIONS

- A. Wajid, N. Kamal, M. Sharjeel, R. M. Sheikh, H. B. Wasim, M. H. Ali, W. Hussain, F. Zaffar. "A First Look at Private Communications in Video Games using Visual Features." Proceedings on Privacy Enhancing Technologies (PoPETs), 2021(3), pp. 433–452.

PROFESSIONAL EXPERIENCE

Graduate Research Assistant Jan 2025 – Aug 2025
Georgia Institute of Technology Shenzhen, China

- Conducted research on urban mobility optimization and traffic management using machine learning and predictive analytics, contributing to a 15% improvement in prediction accuracy for next-generation smart city simulations.
- Developed and deployed SUMO-based simulation models with advanced statistical analysis and optimization algorithms, reducing average simulation runtime by 30% while maintaining prediction fidelity across traffic scenarios.

Machine Learning Engineer Jan 2024 – Jul 2024
Adept Technology Solutions Islamabad, Pakistan

- Developed advanced AI and data science algorithms for ASCEND, a synthetic data generation and anonymization platform, leveraging Python and prompt engineering to optimize LLM-based data augmentation workflows, increasing synthetic data quality scores by 25%.
- Led a full codebase refactor to implement modular, scalable software architecture with CI/CD pipelines and cloud platforms (AWS) using FastAPI for API endpoints, reducing deployment time from hours to under 15 minutes.
- Leveraged PySpark to implement proprietary statistical analysis algorithms and used document processing libraries to preprocess raw data at scale, processing 500K+ documents per pipeline run.

Research Assistant Feb 2020 – Jul 2024
Robotics and Machine Intelligence Lab (ROMI), NUST Islamabad, Pakistan

- Co-authored "A First Look at Private Communications in Video Games using Visual Features," published in Proceedings on Privacy Enhancing Technologies (PoPETs) 2021, performing end-to-end ML workflow including data labeling, feature engineering, model training, and hyperparameter tuning.
- Developed the full software stack for a surveillance area coverage optimization project using computer vision, 3D reconstruction, and blind spot identification algorithms, improving area coverage detection by 20%.

Research Intern Jun 2020 – Apr 2021
TUKL-NCAI Deep Learning Lab Islamabad, Pakistan

- Conducted research on predictive and generative algorithms using Transformers, LSTMs, GANs, and Deep Q-Networks.
- Applied research to develop, train, and compare short-term residential electrical load forecasting models.

PROJECTS

Parallel Interleave for Maximal End Component Decomposition Jan 2026
C++, Python, Storm, perf, Valgrind

- Reimplemented the Interleave algorithm for Maximal End Component decomposition in the Storm probabilistic model checker with parallel execution, achieving 2–5x speedup over sequential interleave across the QVBS benchmark suite.
- Built a profiling pipeline using perf and Valgrind with C++ symbol demangling, flamegraph generation, and automated visualization for performance analysis.
- Developed CLI tools for batch profiling (perf record with DWARF call-graphs at 997 Hz, uprobe instrumentation) and automated visualization (flamegraph SVG, filtered text reports).
- Created wrapper functions to resolve C++ template symbol visibility issues that prevented profilers from detecting Interleave entry/exit points.

Legal Knowledge Graph RAG System for Immigration Law

Nov 2025

Python, Neo4j, Graphiti, Gemini, Pydantic, LangChain

- Designed and deployed an end-to-end legal RAG assistant for immigration case law, integrating a Neo4j-based Graphiti knowledge graph containing 264K nodes and 506K relationships with Gemini 2.5 Flash.
- Implemented Pydantic-driven structured extraction to transform unstructured court opinions into a type-safe knowledge graph, attaining a faithfulness score of 0.770 and outperforming NaiveRAG and CBR-RAG by 15–17% with significant reduction in output variance.
- Built a scalable CourtListener ingestion pipeline with rate-limited concurrency and hybrid retrieval, blending vector similarity search with Cypher-based graph traversal to facilitate multi-hop reasoning across precedents and statutes.

PopSign: Generative AI Storybooks for ASL Learning

Oct 2025

Python, Gemini, Sentence-BERT, spaCy, Embedding-Gemma

- Developed a constrained generative AI pipeline with Gemini to author ASL children's stories, using prompt engineering and a custom validation framework to enforce lexical constraints and reduce error rate to 0.
- Created an NLP evaluation system using Sentence-BERT, spaCy, and readability metrics to automatically classify story difficulty across Fountas-Pinnell levels A–Z, enabling data-driven curriculum design.
- Engineered an embedding-based semantic vocabulary sampler with Embedding-Gemma to generate contextually relevant word sets for personalized, level-appropriate story generation.

Surveillance Area Blind Spot Detection

MATLAB, VisualSFM, LIME, Computer Vision

- Led the research project to identify blind spots in a surveillance camera field of view.
- Used computer vision to reconstruct 3D surveilled areas from images and videos.
- Explored explainable AI methods (LIME) to improve the fairness of the model.

SLAM + DQN for TurtleBot3 Navigation

ROS, Ubuntu, TensorFlow, SLAM, Deep Reinforcement Learning

- Combined SLAM with a Deep Q-Network in ROS to improve performance of autonomous TurtleBot3 navigation.

APF + DQN for TurtleBot3 Navigation

ROS, TensorFlow, Python, Gazebo

- Accelerated training of a DQN by integrating an Artificial Potential Field (APF) algorithm for autonomous navigation of TurtleBot3 in Gazebo.

Predictive Object Tracking

Python, Computer Vision, State Estimation

- Implemented a Kalman filter to detect objects and predict their positions to facilitate autonomous vehicle navigation.

Robot Localization

Python, Probabilistic Inference, Particle Filters

- Designed and implemented a Particle Filter with a custom fuzzing algorithm for autonomous vehicle localization in a map.

Robot Path Planning

Python, Search Algorithms, Dynamic Programming

- Implemented A* search and Value Iteration to enable autonomous robot navigation in a warehouse.

PID Controller

Python, Control Systems, Robotics

- Designed and tuned a PID controller to control a quadcopter.

SpaceX Starship Lander

TensorFlow, Unreal Engine 5, OpenAI, RL

- Designed a model of SpaceX's Starship in Unreal Engine 5.
- Designed and deployed a deep reinforcement learning, multi-agent control system using TensorFlow and OpenAI API to safely land the rocket on its landing pad.

Smart Garbage Bin

Jetson Nano, Computer Vision, TensorFlow, LIME, RL

- Used transfer learning to retrain MobileNet v2 on a custom garbage dataset.
- Employed OpenAI API to integrate the classifier with a deep RL agent while using LIME to explain the classifier's output.
- Deployed on Jetson Nano.

Fine-tuned French GPT-3 Chatbot

OpenAI API, GPT, Flask, Python

- Developed a custom webpage for GPT-3-based French language response generation. Users could retrain the model on their own dataset.

Sign Language Translation Using OAK-D

Computer Vision, OpenCV, OAK-D, DepthAI

- Developed an OAK-D-based sign language translation tool.

Short-term Residential Load Forecasting

TensorFlow

- Built and tested multiple predictive models based on Transformers, LSTMs, and reinforcement learning for short-term energy consumption predictions.

Task Scheduling in Cloud Computing

Python, Graph Algorithms

- Implemented a topological sorting algorithm for efficient staging and execution of tasks on cloud computing workflows and pipelines.

Optimal Seam Finding

Python, Dynamic Programming

- Implemented an algorithm to find low-cost seams in 2D images to enable resizing without distorting perspective and relative object size.

Finding C and G Rich Regions in DNA Sequences

Python, Divide and Conquer

- Implemented a modified maximum sub-array sum algorithm to locate consecutive CpG nucleotide islands in mammalian DNA.

Scalable RISC-V Multicore Processor

Xilinx Vivado, OpenPiton, N-tiny, FPGAs

- Interfaced a custom-built 32-bit RISC-V (N-tiny) with the OpenPiton framework to build a scalable multicore processor.
- Implemented 4-way 4-word set-associative caches with LRU policy, Avalon-to-Pico interconnect, and AXI-4 Lite interface.
- Deployed on Genesys2 FPGA.

Speech Recognition on Edge

Linux, Kaldi, FPGAs

- Designed and built a custom hybrid speech recognition model using a combination of Hidden Markov Models and Gaussian Mixture Models.

TECHNICAL SKILLS

Languages: Python, Java, C, C++, Bash, MATLAB, SQL

ML / DL Frameworks: PyTorch, TensorFlow, Scikit-Learn, Stable-Baselines3, Optuna, Hugging Face, LangChain, Graphiti, Gradio, Streamlit

Reinforcement Learning: OpenAI API, Stable Baselines, Gymnasium, OpenSpiel, Deep Q-Networks

Computer Vision & Robotics: OpenCV, ROS, ROS2, Gazebo, SLAM, Kalman Filters, Particle Filters, OAK-D, DepthAI, LIME, 3D Reconstruction

Generative AI & NLP: LLMs, RAG, Graph RAG, Prompt Engineering, Generative AI, NLP, Sentence-BERT, spaCy, Large Language Model APIs, Model Context Protocol

Specializations: Machine Learning, Deep Learning, Computer Vision, Reinforcement Learning, Data Analysis, AI Safety, Knowledge Graphs, Data Structures, Algorithms

Data & Databases: NumPy, Pandas, SciPy, PySpark, Hadoop, PostgreSQL, MySQL, SQLite, Neo4j, ChromaDB, FAISS, Vector Stores

Cloud & DevOps: AWS, Docker, Kubernetes, CI/CD Pipelines, Git, FastAPI, Pydantic

Hardware & Edge: Jetson Nano, FPGAs (Xilinx Vivado, Genesys2), OAK-D Camera

Formal Methods: Storm Model Checker, perf, Valgrind, Probabilistic Model Checking

AWARDS

Finalist, OpenCV AI Competition — Awarded 10 OAK-D cameras for original project idea (Sign Language Translation)

2021