

Muhammad Umar Ramzan

mumarramzanapp@gmail.com | github.com/UmarRamzan | huggingface.co/UmarRamzan

EDUCATION

Lahore University of Management Sciences (LUMS)

Sep 2020 - May 2024

BS - Computer Science (CGPA: 3.52)

Lahore, Pakistan

Thesis: AI-Driven Healthcare Delivery in Pakistan: A Framework for Systemic Improvement [ACM Compass '24]

TECHNICAL SKILLS

Programming Languages: C, C++, Python, Golang, Haskell, Rust

Web/App development: HTML, CSS, JavaScript, React, Flutter, Svelte

ML/AI: Numpy, Pandas, Matplotlib, Transformers, Jupyter Notebook, Scikit, OpenCV, HuggingFace, LangChain

Miscellaneous: MySQL, Git, LaTeX, Scrapy, Postgres, NoSQL, Modal, Replicate, Pinecone

ACADEMIC PROJECTS

Darcheeni: An Intelligent Healthcare Assistant

Jan 2023 - Present

Advisors: Agha Ali Raza (LUMS), Basmaa Ali (Harvard Medical School)

- Conducted a study of 8 outpatient facilities in Lahore, Pakistan to identify gaps within the current healthcare delivery system, examining over 100 clinical encounters in total.
- Designed and deployed a REST API in Python (using Replicate) to transcribe, diarize, and extract relevant information from medical conversations between physicians and patients in real time. The system also employed a fine-tuned LLM (Llama) to provide a preliminary diagnosis that draws from a collection of medical documents ranging back 10 years using Retrieval-Augmented Generation (RAG).
- Created an initial test website (using Svelte) that interfaced with this API to gauge feedback from physicians so as to improve the system.
- Created a pipeline for audio data collected from three clinics in Lahore to be transcribed and stored.
- Fine-tuned the transcription (Whisper / Wav2Vec2) and diarization (PyAnnote) models on context-relevant audio to improve performance on a mixture of English, Urdu, and Punjabi speech using cloud infrastructure provided by Modal.
- Deployed the system within a controlled clinical environment in the outpatient department of a tertiary care hospital, testing on over 40 patients a day.

Audio Deepfake Detection

Sep 2023 - Dec 2023

Advisor: Mobin Javed (LUMS)

- Proposed the study of deepfake audio detection by humans as a research direction.
- Created a dataset containing clips generated using multiple deepfake architectures.
- Created a website (using Svelte) to provide an interface for conducting the study, with data being stored in a Supabase PostgreSQL database.
- Designed and conducted a small-scale user study (70+ participants) in this regard, collecting multiple data points for each user, including accuracy and the time spent on each question.
- Contrasted the results with the detection accuracy of automated Machine Learning approaches to gain insights regarding human behavior in detecting audio deepfakes.

Other Projects

- **Raft Consensus Protocol:** Developed a GO implementation of the Raft consensus protocol in a distributed system. (*Course assignment in Distributed Systems*)
- **Large Clothing Model:** Designed a website (using React) to aggregate clothing items from several websites in Pakistan and to display relevant items to users based on unstructured queries using a fine-tuned language model (Mistral based on the Transformer architecture), text embeddings (OpenAI), reranking (Cohere), a vector database (Pinecone), and prompt engineering.
- **Doctor-Hospital Directory:** Developed a website (using the MERN stack) to connect patients with physicians and hospitals in Pakistan. (*Course project in software engineering*)
- **Urdu Audio Dataset:** Contributed to the collection and creation of an Urdu audio dataset to be used for various downstream ASR tasks. (*Advised by Agha Ali Raza*)
- **Crop Weed Segmentation:** Implemented a U-Net architecture to segment weeds within videos of crop fields in real-time. (*Course project in Computer Vision*)
- **Machine Learning on the Edge:** Conducted an extensive literature review regarding the challenges associated with deploying ML models on edge devices and the solutions currently being explored
- **NLP Projects:** Text generation using Shannon Visualization, Text classification using Naive Bayes, Recreation of the ELIZA chatbot, HMM for part of speech tagging, music generation, and time-series temperature forecasting, Implementing word2vec and neural word embeddings from scratch

EXPERIENCE

Research Assistant, Center for Speech and Language Technologies June 2023 - May 2024
Lahore, Pakistan
Led a research group of 4 students in conducting literature reviews and finding potential research opportunities in the domain of speech recognition.

Self-employed - Private Tutor June 2023 – Sep 2023
Lahore, Pakistan
Tutored 5+ students in the subjects Algorithms, Operating Systems, and Databases

RELEVANT COURSEWORK

Note: Italicised are graduate-level courses

Taken: Algorithms, Data Structures, Databases, Software Engineering, Operating Systems, Network-centric Computing, *Machine Learning, Deep Learning, Learning for Dynamics and Control, Computer Vision, Applied Probability, Distributed Systems, Generative Artificial Intelligence, Network Security*

LANGUAGES

English: C2 (IELTS 8.5/9)

Urdu: Native

Last updated: 22 June, 2024