

Satvik Dixit

satvikdixit@cmu.edu | (412) 430-2900 | linkedin.com/in/satvik-dixit

EDUCATION

Carnegie Mellon University

Pittsburgh, PA

Master of Science in Electrical and Computer Engineering

Dec 2024

- Concentration in AI/ML Systems | GPA: 4.0/4.0
- Coursework: Speech Recognition, Deep Learning, Deep Generative Models, ML for Signal Processing

Indian Institute of Technology (IIT) Delhi

New Delhi, India

Bachelor of Technology in Electrical Engineering

Aug 2023

- Concentration in ML | GPA: 8.6/10.0

EXPERIENCE

Massachusetts Institute of Technology (MIT)

Cambridge, MA

ML Research Intern | Senseable Intelligence Lab

May 2022 - Aug 2023

- Worked on the explainability of deep learning embeddings for speech emotion recognition and submitted findings at Interspeech 2024 as the first author
- Created the data pipeline and benchmarked the performance of acoustic features and self-supervised learning embeddings on the emotion classification task across multilingual datasets
- Developed a novel method to explain the acoustic information encoded in the deep learning embeddings based on the 'probing classifiers' framework

SKILLS

Programming Languages: Python, Java, C++, MATLAB, LaTeX, NumPy, Pandas, Bash/Shell

Frameworks and Tools: Scikit-learn, PyTorch, TensorFlow, Hugging Face, Matplotlib, Git, AWS, CUDA

PROJECTS

Discrete Speech Units for Automatic Speech Recognition

Research Project | Professor Shinji Watanabe, CMU

Sep 2023 - Mar 2024

- Explored using discrete tokens for Automatic Speech Recognition (ASR) on the ML-SUPERB dataset, outperformed traditional spectral features while compressing the size of speech data by 1000X.

Implementing Text and Duration Transducer (TDT) Architecture

Research Project | Professor Shinji Watanabe, CMU + NVIDIA

Jan 2024 - Present

- Implementing the TDT architecture (current state-of-the-art model for ASR) in the ESPNET toolkit.

Room Acoustics Simulation

Research Project | Professor Martin Vetterli, EPFL

June 2021 - Aug 2021

- Worked on developing Pyroomacoustics: an open-source Python package for acoustics simulation. Improved RIR simulation accuracy by 30% by adding 'directivity' functionality to mics and sources.

Adapting LLMs for Mathematical Tasks

Research Project | Professor Bhiksha Raj, CMU

Jan 2024 - Present

- Developing a parameter-efficient memory module (AIR-2) to enable LLMs to do quantitative tasks. Benchmarked the performance of LLMs (CodeLlama-7B, Phi-2-2.7B, and DeepSeek-1.3B) by pre-training and fine-tuning on math word problems, arithmetics, and code generation datasets.

LLM Explainability by Latent Space Exploration

Research Project | Professor Giulia Fanti, CMU

Jan 2024 - Present

- Using latent-space exploration techniques to condition language-model-generations and evaluating using detectors like "Toxicity" or "Fairness" and determine multi-dimensional decision boundaries