Atharva Anand Joshi

atharvaa@andrew.cmu.edu • 412-954-7615 • https://www.linkedin.com/in/atharva-anand-joshi/ • https://atharva253.github.io/

RESEARCH INTERESTS
Speech Processing: Enhancement and Separation, Deep Learning, Representation Learning

---

EDUCATION

Carnegie Mellon University   Pittsburgh, PA
Master of Science in Electrical and Computer Engineering - Applied Advanced Study (AI/ML Systems)   December 2024
GPA: 4.0/4.0

Birla Institute of Technology and Science, Pilani   Pilani, India
Bachelor of Engineering in Electrical and Electronics Engineering   July 2022
GPA: 9.49/10
GRE: 331/340 (AWA: 4/6), TOEFL: 109/120

SKILLS

Programming: Java, Python3, C, CUDA , MATLAB
Deep Learning: PyTorch, TensorFlow, Keras, AWS and GCP for ML
Frameworks: Kalimba Toolkit, PySpark, WandB, OpenMP, LaTeX, Git

---

EXPERIENCE

Hewlett-Packard Inc., Poly   Austin, TX
Research and Development Intern   May – August 2023
- Developed a deep learning model that can suppress stationary and impulsive background noises directly on headsets
- Optimized the inference pipeline on memory, compute and latency for real-time execution on the QCC5171 DSP
- Introduced an efficient approach to Audio ML Development that would make integration into products 6-8 times faster
- Won the 2023 HP Intern Award for this project by securing 4th position among ~90 engineering interns

American Express, Artificial Intelligence Labs   Gurgaon, India
Analyst - Product Development   July - December 2022
- Researched a blend of Tabular models with Tree-based algorithms for the Credit Default Prediction problem
- Enhanced model performance by leveraging extensively hand-engineered features, along with meta features
- Generated valuable business insights pertaining to features selection and effective aggregation of temporal features
Analyst Intern   January - June 2022
- Proposed a template-based journey that allows users to seamlessly create and deploy their machine learning pipelines
- Developed a framework that facilitates deployment of end-to-end Self Learning pipelines for Sequence Models

Adobe Research, India   Bangalore, India
Research Intern   May - August 2021
- Created rich user representations that can be projected onto edge servers, hence powering faster marketing services
- Performed various experiments around the extent of compression and updatability of the representations generated
- Contributed towards a **patent** as a co-inventor - "Generating Concise and Common User Representations for Edge Systems from Event Sequence Data Stored on Hub Systems" (US20230419339A1 - Filed June 24, 2022)

PUBLICATIONS

- **A. A. Joshi**, H. Settibhaktini, and A. Chintanpalli. Modeling concurrent vowel scores using the time delay neural network and multitask learning. IEEE/ACM Transactions on Audio, Speech, and Language Processing, 30:2452-2459, 2022
- **A. A. Joshi**, P. Bhardwaj, and S. M. Zafaruddin. Terahertz wireless transmissions with maximal ratio combining over fluctuating two-ray fading. IEEE Wireless Communications and Networking Conference (WCNC), pages 1575-1580, 2022

PROJECTS

Machine Learning for Speech Processing Applications
- Contributed an end-to-end reproducible deep learning pipeline for the Kinect-WSJ dataset – a multichannel, reverberated and noisy version of the WSJ0-2mix dataset to the ESPnet Speech Processing Toolkit
- Performed benchmark analysis on this speech separation dataset using the current state-of-the-art models
- Currently investigating the data and model scalability aspects for large speech enhancement and separation models

Proactive Servicing: American Express ML Challenge
- Utilized event sequences and demographic data to predict customer intent at the start of the Ask AmEx chat session
- Employed joint training of Bidirectional GRU with Feedforward Networks
- Attained a validation top-5 accuracy score of 0.768 and hence made it to the top 10 leaderboard out of ~80 teams

High Performance Parallel Implementations for Convolutional Neural Networks
- Provided fast OpenMP and CUDA implementations for various subroutines corresponding to the convolution layer
- Investigated several design choices in detail aiming to optimize speedup over a simple sequential C implementation
- Achieved maximum speedup of 4.23x on the Intel(R) Xeon(R) Silver 4208 CPU and 73.87x on the Nvidia Tesla T4 GPU

Quasi-Newton Methods for Deep Learning
- Replicated the results of the NeurIPS 2020 paper: Practical Quasi-Newton Methods for Training Deep Neural Networks
- Analyzed the algorithm theoretically and evaluated its convergence plots
- Compared the convergence trends of this algorithm against other first and second-order descent methods

Concurrent Vowel Identification Using DNN
- Developed a novel computational model to predict the effect of fundamental frequency difference on the identification scores in a Concurrent Vowel Identification Experiment
- Trained a TDNN augmented with a Multi-task Learning setup on the neuron responses from the Auditory Nerve Model

AWARDS AND RECOGNITION

HP Intern Award – Top 5	2023
American Express Modeling Super Bowl – Top 10 Leaderboard	2022
OP Jindal Engineering and Management Scholarship	2020 and 2021
Institute Merit-Based Scholarship, Birla Institute of Technology and Science, Pilani – Top 2%	2019-2022
Regional Runner-up at TCS IT Wiz, a nationwide inter-school quiz competition	2015

MENTORSHIP AND TEACHING

- **Graduate Teaching Assistant** for 10-605: ML with Large Datasets — Spring 2024
  Instructors: Prof. Ameet Talwalkar and Prof. Geoff Gordon
- **Graduate Teaching Assistant** for 18-661: Introduction to ML for Engineers — Fall 2023
  Instructors Prof. Yuejie Chi and Prof. Beidi Chen
- **Undergraduate Teaching Assistant** for BITS F312: Neural Networks and Fuzzy Logic — Fall 2021
  Instructors: Prof. Surekha Bhanot and Prof. Bijoy Krishna Mukherjee

RELEVANT COURSEWORK

Graduate: Introduction to Deep Learning, Machine Learning for Signal Processing, Deep Generative Modeling, Machine Learning with Large Datasets, Natural Language Processing, Introduction to Machine Learning for Engineers, CUDA Programming, Convex Optimization

Undergraduate: Neural Networks and Fuzzy Logic, Artificial Intelligence, Object Oriented Programming (Java), Digital Signal Processing, Digital Image Processing, Communication Systems, Operating Systems

POSITIONS OF RESPONSIBILITY

Joint Coordinator at Ragamalika, the Classical Music and Dance Club of BITS Pilani, Pilani Campus (2020-2021)
- Actively involved in composing music for semester productions and managing professional concerts
- Avid practitioner and performer of Hindustani Classical Vocal Music for the past fourteen years