STEVEN HILLIS

(636) 484-2548
stevenrhillis@gmail.com

EDUCATION

Carnegie Mellon University: School of Computer Science Pittsburgh, PA *December 2019*

Master of Science in **Intelligent Information Systems**

Southeast Missouri State University Cape Girardeau, MO *May 2018*

Bachelor of Science in **Computer Science** and **Economics**

EXPERIENCE

AssemblyAI, Deep Learning Engineer Remote *January 2021 – November 2021*

- Developed and scaled a speaker diarization feature powered by a Fast ResNet-34 speaker embeddings model, delivering reliable count estimation, low word-level DER, and RTFs under 2% for audios 10s to 12hr
- Trained a Conformer-CTC ASR model using intermediate-CTC and auxiliary AED losses, distributed over multiple nodes on enterprise-scale data, with a streaming attention mask for variable-latency inference
- Productionized a fault-tolerant, batchwise inference harness for streaming ASR with tight latency constraints

Hyperia, Machine Learning Engineer Denver, CO *January 2020 – December 2020*

- Produced a speaker diarization feature using a res2net embeddings model trained with margin loss
- Iteratively bootstrapped Jasper 10x5 ASR model on freely available and increasingly diverse data
- Designed a tool converting written text into flawed, speech-like text to produce sufficient conversational data for finetuning a transformer model on error correction and contextualization of ASR output
- Leveraged a prompting approach with GPT-3 followed by a finetuned T5 model for producing structured, well-formed, suitably abstractive transcript summaries with acceptable balance of coverage and fluency

Federal Reserve Bank of Cleveland, Data Science Intern Pittsburgh, PA *Summer 2019*

- Instituted a robust and efficient data pipeline from complex tables to pandas dataframes using MySQL queries
- Constructed a financial lexicon from embeddings, clustered with dynamic KMeans, visualized with PCA, TSNE

PROJECTS

Defenses for Adversarial Attacks on ASR Neural Networks Carnegie Mellon: Capstone *Fall 2019*

- Extended IBM's Adversarial Robustness Toolbox to handle speech models: Listen-Attend-Spell, DeepSpeech2
- Explored novel defenses for underserved speech and text domains against new and existing attacks

Multimodal, Multilingual Grapheme-to-Phoneme Conversion Carnegie Mellon: Course *Spring 2019*

- Introduced state-of-the-art multilingual neural grapheme-to-phoneme model for low-resource languages, leveraging an auxiliary audio modality during training without introducing dependency during inference
- Published: DeepLo 2019

Facial Image Classification and Verification Carnegie Mellon: Course *Spring 2019*

- Performed multiclass classification over augmented facial images using a modified ShuffleNetV2 architecture
- Adapted the model for facial verification, generating cosine distance similarity scores from facial embeddings

Speech to Speech Translation for Unwritten Languages Carnegie Mellon: Research *Spring 2019*

- Evaluated four unsupervised representations of speech data on downstream BLEU to determine optimal intermediate for languages without stable writing systems in the traditional speech-to-speech pipeline
- Published: Interspeech 2019

SKILLS

Programming Languages: Python, Java, SQL, C++, Perl, JavaScript, C#, C

Python Tools: Pytorch, NumPy, Hugging Face, pandas, NLTK, scikit-learn, wandb, Tensorflow, FairScale, Comet

DevOps Tools: Git, Docker, Jenkins, Flask, microservices, CI/CD; AWS S3, EC2, SageMaker; Google Colab, Gradio