Alan Zhou

azhou23@jhu.edu | atzhou8.github.io | Baltimore, MD

Education

Johns Hopkins University — Baltimore, MD
PhD in Cognitive Science — *Aug. 2022 - Present*

University of California, Berkeley — Berkeley, CA
B.A. in Computer Science and Cognitive Science — *Aug. 2017 - Dec. 2021*

Publications

Submitted

- (Submitted) Gašper Beguš, Thomas Lu, **Alan Zhou**, Peter Wu, and Gopala K. Anumanchipalli. Ciwagan: Articulatory information exchange. **arXiv**

Peer-reviewed Journals and Conferences

** denotes equal contribution*

- (2023) Gašper Beguš*, **Alan Zhou***, Peter Wu, and Gopala K Anumanchipalli. Articulation gan: Unsupervised modeling of articulatory learning. In *ICASSP 2023 IEEE International Conference on Acoustics, Speech and Signal Processing*. **PDF** Video
- (2023) Gašper Beguš, **Alan Zhou**, and Christina Zhao. Encoding of speech in convolutional layers and the brain stem based on language experience. *Scientific Reports*. **PDF**
- (2022) Gašper Beguš and **Alan Zhou**. Interpreting intermediate convolutional layers of generative cnns trained on waveforms. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30. **PDF**
- (2022) Gasper Begus and **Alan Zhou**. Modeling speech recognition and synthesis simultaneously: Encoding and decoding lexical and sublexical semantic information into speech with no direct access to speech data. In *Proc. Interspeech 2022*. **PDF**
- (2022) Gašper Beguš and **Alan Zhou**. Interpreting intermediate convolutional layers in unsupervised acoustic word classification. In *ICASSP 2022 IEEE International Conference on Acoustics, Speech and Signal Processing*. **PDF**

Teaching

At Johns Hopkins

- **Bayesian Inference**, *Teaching Assistant* — Fall 2023
- **Neuroscience: Cognitive**, *Teaching Assistant* — Spring 2022

At UC Berkeley

- **Deep Learning and Phonology**, *Guest Lecturer* — Fall 2021
  (Gave a guest lecture about high-performance computing to linguistics graduate students)
- **Data Structures**, *Academic Intern* — Spring 2018
  (Helped students in lab sections and office hours)

Experience

Berkeley Speech and Computation Lab — Berkeley, CA
Undergraduate Research Assistant | PI: Gašper Beguš — November 2020 to December 2021

- Probed intermediate representations of speech in generative adversarial networks
- Compared intermediate representations in GANs with the auditory brainstem response via latent vector recovery of recorded stimuli

Berkeley Division of Data Science — Berkeley, CA
Research Apprentice | Mentor: Taka'aki Taira — January 2019 to January 2020

- Recovered underlying stress fields from earthquake data using weighted least squares
- Created scripts to calculate and visualize information about the faulting regime, stress orientation, and confidence level of stress fields across Northern California

Projects

F-ZERO Reinforcement Learning
A reinforcement learning agent trained to play the SNES racing game F-ZERO
(GitHub)

- Utilized socket programming to allow an emulator with Lua scripting capabilities to interface with Python and PyTorch
- Used deep Q-learning to create an agent capable of racing in a 3D environment given only screen input

Markov Bot
A Discord bot that creates Markov chains out of user messages in order to simulate text.
(GitHub)

- Developed a means to construct Markov chains for individual users, and to generate novel sentences using constructed chains

Skills

Programming Languages:	Python, Java, C, MATLAB, R, Lua, SQL
Tools/Technologies:	PyTorch, Tensorflow, Keras, Slurm, matplotlib Jupyter, Git, Gradle/Maven
Natural Languages:	English (fluent), Mandarin (conversational)