Joshua Zhanson

425.281.7055 | josh@jzhanson.com | jzhanson.com | linkedin.com/in/josh-zhanson

Education

Carnegie Mellon University School of Computer Science — *Pittsburgh, PA*
Master of Language Technologies — *August 2022*
QPA: 3.80/4.00 Advisor: Yonatan Bisk Supported by **NSF Graduate Research Fellowship**

Carnegie Mellon University School of Computer Science — *Pittsburgh, PA*
Bachelor of Science in Computer Science, Minor in Machine Learning — *May 2020*
QPA: 3.95/4.00 Dean's List: Fall 2016 - Spring 2019 College & University Honors
Senior thesis: *Investigating and Robustifying Proximal Policy Optimization*
Advised by Emilio Parisotto, Adarsh Prasad, and Ruslan Salakhutdinov

Research Projects

Learning Visual Representations through Embodied Interaction Exploration — *August 2020 - Present*

- Created Find One and Interaction Exploration environments in **Python** built on AI2THOR interactive embodied household robotics simulator to explore visual representation learning with embodiment
- Designed customizable ResNet visual encoders and decoders and LSTM policy model architectures in **Pytorch** for control with pixel inputs and outputs and designed a self-supervised policy and visual pretraining task
- Implemented custom variants of reinforcement learning algorithms Advantage Actor-Critic and Proximal Policy Optimization with hogwild asynchronous multiprocess training to allow running 8+ parallel environments
- Built multiprocess autoencoder baseline, supervised topline, and visual probe experiment pipeline to evaluate quality of learned representations on datasets with 2M+ images generated from different heuristic agent policies in AI2THOR simulator

On Proximal Policy Optimization's Heavy-tailed Gradients — *August 2019 - May 2020*

- Integrated gradient estimators from robust statistics into Advantage Actor-Critic and Proximal Policy Optimization deep reinforcement learning algorithms in **Python** and **Pytorch**
- Evaluated effect of different optimization heuristics on heavy-tailedness of policy gradient and likelihood ratio distributions throughout a training epoch using alpha-index estimator from robust statistics
- Discovered severe heavy-tailedness in off-policy gradients (alpha ≈ 1.0) and likelihood ratios (alpha ≈ 1.2) taken on same batch of data, prompting a reevaluation of the policy gradient reinforcement learning paradigm
- Accepted to **ICML 2021**

Proprioceptive Spatial Representations for Generalized Locomotion — *June 2018 - July 2019*

- Developed JSONWalker environment for robot locomotion and GUI editor with **Python** to allow users to easily construct complex robot bodies in box2d physics simulator
- Wrote scripts in **Python** to randomly construct robot bodies and create datasets of 300 unique robot bodies
- Trained **PyTorch** convolutional models for control using a grid-based proprioceptive robot body state, outperforming baseline models by 20% success rate and solved 9% more unseen robot body configurations
- Accepted to **Workshop on Structure & Priors in Reinforcement Learning** at **ICLR 2019**

Employment

Merit International, Inc. (formerly Sigma Accolade, Inc.) — *Millbrae, CA*
Software Engineer Intern — *May 2018 - August 2018*

- Implemented a feature to allow organizations to prevent duplicate certifications issued to the same user by adding **React** components in **JavaScript** linked to the **Scala** backend with as-you-type **GraphQL** mutations and queries and **Cats** type abstractions for error handling and threading back-end errors to frontend UI

Skills

Languages: Python ~ C/C++ ~ Javascript ~ Scala ~ Standard ML ~ Java ~ Bash ~ Swift
Technologies: Pytorch ~ Tensorflow/Keras ~ OpenCV ~ Numpy ~ Pandas ~ Matplotlib/Seaborn ~ Docker ~ Git ~ React ~ GraphQL ~Typelevel.Cats