Joshua Zhanson

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

Employment

Microsoft *Mountain View, CA*
Data & Applied Scientist 2 | Microsoft Designer *October 2022 - Present*

- **Template retrieval system**: Led improvement of template retrieval search using a joint transformer + TF-IDF + LLM architecture, increasing **precision by 93%** and **recall by 115%**. Evaluated using offline crowdlabeling and online A/B flighting.
- **Diffusion design search**: Built a scalable embedding-based retrieval system with a large multimodal language model component, increasing design success rate by **9.13%**, seen rate by **3.78%**, and kept rate by **16.83%**.
- **Invitations, banners, posters:** Designed and iterated on LLM prompting architecture and evaluation pipeline for diffusion-based artifact creation, including text replacement with localization capabilities.
- **Automatic LLM prompt optimization**: Applied internal prompt optimization tool to several Designer artifacts, creating new LLM prompting layers, facilitating prompt iteration and model upgrades, and achieving **76.47%** improvement over baseline for invitations artifact and reducing problematic generations by **47%**.
- **Visual creation in M365 Copilot Chat:** Led integration of visual creation capabilities into M365 Copilot Chat and Microsoft Visual Creator Copilot extension with **5K DAU**.
- **Infographics prompting**: Created icon swap flow and reduced prompt length from **7.245 tokens to 388 tokens** and error rate from **4.18% to 0%**. Evaluated prompt on **85 languages**, ensuring broad coverage.

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 - July 2022*

- 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*

- Accepted to **ICML 2021**

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

- Accepted to **Workshop on Structure & Priors in Reinforcement Learning** at **ICLR 2019**

Skills

Languages: Python ~ C# ~ C/C++
Technologies: LLM prompting ~ Pytorch ~ Tensorflow/Keras ~ OpenCV ~ Numpy ~ Pandas ~ Docker ~ Git

Last updated May 17, 2025