Yifan Zhou

yifanzhou.com
(412)251-9000
yzhou298@asu.edu
linkedin.com/in/yifan-zhou-639974191/

SUMMARY

I am a PhD student at Arizona State University, advised by Dr. Heni Ben Amor. My research focuses on **Robot Learning Machine Learning**, and **LLM Agent**. My publications can be seen in CoRL, COLM, IROS, ICRA and so on. I'm expected to graduate in Aug 2025, and I'm looking for machine learning related *research and engineering positions*.

EDUCATION

Arizona State University   **Aug 2021 – Aug 2025**
Ph.D. Computer Science, GPA: 4.0   Tempe, AZ
Advisor: Dr. Heni Ben Amor

Carnegie Mellon University   **Aug 2019 – May 2021**
M.S. Artificial Intelligence and Innovation, GPA: 3.92   Pittsburgh, PA
Advisor: Dr. Wenzhen Yuan

Southwest Jiaotong University   **Sep 2015 – Jun 2019**
B.E. Computer Science and Technology, GPA: 3.65   Chengdu, China
Advisor: Dr. Xiao Wu

WORK EXPERIENCE

Google LLC   **May 2024 – Aug 2024**
Software Engineering Intern   Sunnyvale, CA

- Proposed and implemented a holistic data quality assurance mechanism for the **SFT** data of a series of **LLM**s, for the product *Gemini for Google Workspace*.
- The proposed mechanism includes 3 parts: *rule-based checkers*, *LLM-based checkers* and *AutoEval-based checkers*, which identify data format issues (e.g., missing critical fields), semantic issues (e.g., the mismatch between model final response and user query), and overall quality issues (e.g., the overall impression of the final response) respectively.
- The implemented checkers detected problems up to 20.1% of the training datapoints. This helps the data curation team to improve the data quality, and in turn helps the training of *Gemini for Google Workspace* LLMs.

LinkedIn Corporation   **May 2020 – Aug 2020**
Machine Learning – Artificial Intelligence Engineering Intern   Sunnyvale, CA

- Created a **BERT** based LinkedIn member profile modeling approach, namely EntityBERT.
- Proposed 2 objectives for the training of EntityBERT: *masked-entity-modeling* and *profile-relation-prediction*.
- The trained model was able to predict missing fields of given LinkedIn user profiles, where the top 1 accuracy was increased by over 10% compared to previous baselines.
- EntityBERT provides the team with a user profile modeling methodology and creates an entity recommendation approach for downstream teams.

PUBLICATIONS

Selected

[1] **Y. Zhou**, X. Liu, Q. Vuong, et al. "AutoMA: Automated Modular Attention enables Context-Rich Imitation Learning using Foundation Models". *Under Review.*

[2] Guan*, **Y. Zhou***, D. Liu, et al. "'Task Success' is not Enough: Investigating the Use of Video-Language Models as Behavior Critics for Catching Undesirable Agent Behaviors". *2024 Conference On Language Modeling* ***(COLM 2024)***.

[3] **Y. Zhou**, S. Sonawani, M. Phielipp, et al. "Learning Modular Language-Conditioned Robot Policies through Attention." ***Autonomous Robots (2023)****: 1-21. In Collection: Large Language Models in Robotics.*

[4] **Y. Zhou**, S. Sonawani, M. Phielipp, et al. "Modularity through Attention: Efficient Training and Transfer of Language-Conditioned Policies for Robot Manipulation." *6th Conference on Robot Learning **(CoRL 2022)***.

[5] **Y. Zhou**, R. Jiang, X. Wu, et al. "BranchGAN: Unsupervised Mutual Image-to-Image Transfer with a Single Encoder and Dual Decoders." *IEEE Transactions on Multimedia **(TMM 2019)** 21, no. 12: 3136-3149.*

Peer Reviewed

[6] X. Liu, **Y. Zhou**, F. Weigend, et al. "Diff-Control: A Stateful Diffusion-based Policy for Imitation Learning". *2024 International Conference on Intelligent Robots and Systems **(IROS 2024)***.

[7] Q. Vuong, et al. "Open X-Embodiment: Robotic Learning Datasets and RT-X Models". *2024 IEEE International Conference on Robotics and Automation **(ICRA 2024 Best Conference Paper Award)***.

[8] X. Liu, **Y. Zhou**, S. Ikemoto, et al. "α-MDF: An Attention-based Multimodal Differentiable Filter for Robot State Estimation" *7th Conference on Robot Learning **(CoRL 2023)***.

[9] S. Sonawani, **Y. Zhou**, H. Ben Amor. "Projecting Robot Intentions Through Visual Cues: Static vs. Dynamic Signaling." *2023 International Conference on Intelligent Robots and Systems **(IROS 2023)***.

[10] X. Liu, G. Clark, J. Campbell, **Y. Zhou** and H. Ben Amor. "Enhancing State Estimation in Robots: A Data-Driven Approach with Differentiable Ensemble Kalman Filters." *2023 International Conference on Intelligent Robots and Systems **(IROS 2023)***.

RESEARCH EXPERIENCE

Arizona State University, Interactive Robotics Lab   **Aug 2021 – Present**
Research Assistant | Advisor: Dr. Heni Ben Amor   Tempe, AZ

Robot Manipulation by Language Conditioned Imitation Learning   *Aug 2021 – Present*
- Proposed a novel method, namely Hierarchical Modularity, for training **language-conditioned manipulation policies**, which allows for efficient training and rapid transfer across different types of robots.
- Conducted evaluations on table-top manipulation tasks, demonstrating the flexibility of the proposed method, which supports expanding itself to new tasks.
- Created an automated pipeline for **synthesis of 3D models** of novel objects using **diffusion models** and **DPT models**, which helps efficient data collection in simulation.
- Utilized **VLMs** to annotate imitation learning datasets for rich contexts, leading to improved performance in policy models.

Robot Imitation Learning using Diffusion Models   *Nov 2023 – May 2024*
- Proposed Diff-Control, a robot policy model which leverages **Diffusion Models** to learn actions. **ControlNet** is integrated, providing extra priors in the form of state history.
- Evaluated Diff-Control's performance on a series of manipulation tasks, including table-top pick and place and kitchen scenario manipulation, finding it outperforming baselines by over 10%.

Identifying Undesired Agent Behaviors using Vision-Language Models (VLMs)   *Dec 2023 – Feb 2024*
- Constructed a benchmark that contains diverse cases of undesirable robot policies.
- Investigated the feasibility of **VLM critics** on the proposed benchmark, finding that GPT-4V can identify undesirable robot behaviors (recall rate of 69%), but with hallucinated information (precision rate of 62%).
- Demonstrated **policy refinement** on a real UR5 robot in five household scenarios, wherein a Code-as-Policies agent updates the policy according to VLM critiques on the rollouts.

Multi-Modal Robot State Estimation using Attention Mechanism   *Feb 2023 – Aug 2023*

- Proposed an attention-based Multimodal Differentiable Filter (α-MDF) for **robot state estimation**, which utilizes **attention mechanisms** to learn **multimodal latent representations** and generate learnable gains that combine multiple input modalities.

Projecting Robot Intention for Human Robot Interaction   *Sep 2022 – Feb 2023*
- Studied the effect of projecting different visual signals to a human partner in a **mixed-reality human-robot collaboration** setup, finding that the mixture of multiple visual signals result in significant advantages, i.e., increased task efficiency and reduced cognitive load.
- Introduced an information theoretic analysis, **Transfer Entropy**, to numerically quantify the degree of information transfer between visual signals and human behavior.

Mixed-Reality Extrinsic Parameter Auto-Correction Using Imitation Learning   *Aug 2022 – Sep 2022*
- Adopted image-based **imitation learning** for **extrinsic parameter correction** in a projector-camera stereo setup, by training a policy to iteratively correct the offset between a QR code and a projected pattern.

Carnegie Mellon University, RoboTouch Lab   **Nov 2019 – May 2020**
Research Student | Advisor: Dr. Wenzhen Yuan   Pittsburgh, PA
- Predicted hardness of objects touched by a **GelSight tactile sensor**, where the sensor signal is a time sequence of images, by training a combination of **CNN**s and **RNN**s.
- Improved R^2 by 0.097 (achieved 0.884) on novel shape dataset by adopting pretrained models, weight freezing and curriculum learning.

Southwest Jiaotong University, Virtual Reality and Multimedia Lab   **Sep 2017 – Jun 2019**
Undergraduate Research Student | Advisor: Dr. Xiao Wu   Chengdu, China
- Proposed BranchGAN, an unsupervised end-to-end **Generative Adversarial Network (GAN)** for mutual **image-to-image transfer** between two domains.
- Introduced a novel backbone architecture with one single encoder and dual decoders to capture the cross-domain distributions and generate images in both domains. Together with a set of 3 comprehensive training objects, the method is able to outperform baselines with great margin on different benchmarks.

SKILLS

Tools & Libraries
> PyTorch, TensorFlow, Jax, MuJoCo, OpenCV, MATLAB, Git, Docker, Kubernetes, Django

Languages
> Python, Java, HTML, JavaScript, SQL, C/C++