

Lawrence Jang

CONTACT

EMAIL: ljang@cs.cmu.edu WEBSITE: lawrencekjang.github.io

RESEARCH INTERESTS

Aligning human preferences and instilling human capabilities within LLM agents; Reproducible evaluation and benchmarking LLM agents; LLM agents for automating everyday human computer tasks

EDUCATION

Carnegie Mellon University, Pittsburgh, PA, United States

M.S. Machine Learning, GPA: 4.00

Jan 2024 - Dec 2024

- Advisor(s): Dr. Ruslan Salakhutdinov
- Relevant courses: Advanced Deep Learning, Probabilistic Graphical Models, Theoretical and Empirical Foundations of Modern Machine Learning, Probability and Mathematical Statistics, Advanced Machine Learning Theory

Carnegie Mellon University, Pittsburgh, PA, United States

B.Sc. Statistics and Machine Learning, GPA: 3.73

Aug 2019 - Dec 2023

- Research advisor(s): Dr. Ron Yurko
- Relevant courses: Convex Optimization, Distributed Machine Learning with Large Datasets, Deep Reinforcement Learning, Non-Parametric Learning, Modern Regression, Multilevel and Hierarchical Modeling

PUBLICATIONS

VideoWebArena: Evaluating Long Context Multimodal Agents with Video Understanding Web Tasks

Lawrence Jang, Yinheng Li, Charles Ding, Justin Lin, Paul Liang, Dan Zhao, Rogerio Bonatti, Kazuhito Koishida (*work done while at Microsoft*)

Published in the Workshop on Open-World Agents at NeurIPS 2024

Under review, ICLR 2025

Windows Agent Arena: Evaluating Multi-Modal OS Agents at Scale

Rogerio Bonatti, Dan Zhao, Francesco Bonacci, Dillon Dupont, Sara Abdali, Yinheng Li, Yadong Lu, Justin Wagle, Kazuhito Koishida, Arthur Bucker, Lawrence Jang, Zack Hui (*work done while at Microsoft*)

Published in the Workshop on Open-World Agents at NeurIPS 2024

Published in the Towards Safe and Trustworthy Agents Workshop at NeurIPS 2024 (Oral)

Under review, ICLR 2025

ICAL: Continual Learning of Multimodal Agents by Transforming Trajectories into Actionable Insights

Gabriel Sarch, Lawrence Jang, Michael J. Tarr, William W. Cohen, Kenneth Marino, Katerina Fragkiadaki

Published in NeurIPS 2024 (Spotlight)

MMoE: Enhancing Multimodal Models with Mixtures of Multimodal Interaction Experts

Haofei Yu*, Zhengyang Qi*, Lawrence Jang*, Ruslan Salakhutdinov, Louis-Philippe Morency, Paul Pu Liang

Published in EMNLP 2024

VisualWebArena: Evaluating Multimodal Agents on Realistic Visual Web Tasks

Jing Yu Koh, Robert Lo*, Lawrence Jang*, Vikram Duvvur*, Ming Chong Lim*, Po-Yu Huang*, Graham Neubig, Shuyan Zhou, Ruslan Salakhutdinov, Daniel Fried

Published in ACL 2024

Published in the LLM Agents Workshop at ICLR 2024

ACADEMIC RESEARCH EXPERIENCE

Machine Learning Department, Carnegie Mellon University

Graduate Research Assistant

Jan 2024 - present

- Conducted research under Dr. Ruslan Salakhutdinov on multimodal agent benchmarks and agentic frameworks.
- Conducted research under Dr. Ruslan Salakhutdinov on training a mixture-of-experts inspired multimodal model based upon multimodal interactions.
- Conducted research under Dr. Katerina Fragkiadaki on retrieval augmented multimodal agents with self-generating in-context memories.

PROFESSIONAL EXPERIENCE	<p>Microsoft <i>Research Intern</i> May 2024 - Aug 2024</p> <ul style="list-style-type: none"> • Helped developed and publish Windows Agent Arena, a scalable OS platform for testing and benchmarking of multi-modal AI agents. • Created VideoWebArena, a video-based benchmark intended to test performance of long-context multimodal agents. <p>Chicago Cubs <i>Research and Development Intern</i> May 2022 - Aug 2022</p> <ul style="list-style-type: none"> • Modeled multivariate joint batted ball distributions with copula models to automate a synthetic data generation of a player's batted ball profile. <p>Major League Baseball Commissioner's Office <i>Data Science Intern</i> May 2021 - Aug 2021</p> <ul style="list-style-type: none"> • Conducted data analysis on spin-rate trends in investigation of foreign substance usage. • Created new high-sensor camera datasets and blog writeups regarding new MLB data innovations. <p>FormFree <i>Data Analytics Intern</i> May 2020 - Aug 2020</p> <ul style="list-style-type: none"> • Identified and manually labeled raw data logs of financial transactions to improve NLP model. • Created and visualized company financial datasets.
HONORS AND AWARDS	<p>Carnegie Mellon University Machine Learning Department TA of the Year 2024</p> <p>Carnegie Mellon University Statistics Department Undergraduate TA of the Year 2023</p> <p>Dean's List (High Honors) 2021, 2022, 2023</p>
TEACHING EXPERIENCE	<p>Carnegie Mellon University <i>Teaching Assistant</i> Aug 2022 - present</p> <ul style="list-style-type: none"> • 10-703 Deep Reinforcement Learning (Fall '22, Spring '23, Fall '23, Fall '24) • 10-708 Probabilistic Graphical Models (Fall '23, Spring '24) • 11-777 Multimodal Machine Learning (Fall '24) • 36-402 Modern Regression (Spring '23, Spring '24) • 36-401 Modern Regression (Fall '23)
VOLUNTEERING AND SERVICE	<p>Moderator for London Machine Learning Meetup: LLaVA 2024</p> <p>CMU MSML Admissions Committee 2023</p> <p>CMU Paths to AI Research Mentor 2024-Present</p> <p>Founder of CMU Korean-American Scientists and Engineers Association Chapter 2022-Present</p> <p>CMU Club Basketball Head Coach 2022-Present</p>
TECHNICAL SKILLS	<p>PROGRAMMING: Python, R, SQL</p> <p>MACHINE LEARNING FRAMEWORKS: PyTorch, Tensorflow, HuggingFace Transformers, scikit-learn, wandb, numpy, pandas, matplotlib, ggplot2, tidyverse</p>