

KYLE ROBERT ZENTNER

Mountain View, California
krzentner@gmail.com (619) 694-8313 <https://zentner.io>

Education

Ph.D. Computer Science

Thesis Topic: Leveraging Cross-Task Transfer in Sequential Decision Problems

University of Southern California, Los Angeles, CA

Graduated **May 2024**

GPA 3.80

B.S. Electrical Engineering and Computer Science

University of California at Berkeley, Berkeley, CA

Graduated **December 2015**

GPA 3.584

Publications and Research Projects

Conditionally Combining Robot Skills Using Large Language Models *Presentation at ICRA 2024*

K.R. Zentner, Ryan C. Julian, Brian Ichter, Gaurav S. Sukhatme
Introduced the Language-World benchmark, which adds basic language functionality to the Meta-World benchmark. Demonstrated strong one-shot transfer across 50 tasks in Language-World via Plan Conditioned Behavioral Cloning (PCBC), and Cross-Task Co-Learning.

Guaranteed Trust Region Optimization via Two-Phase KL Penalization

K.R. Zentner*, Ujjwal Puri*, Zhehui Huang, Gaurav S. Sukhatme
Introduced Fixup Policy Optimization (FixPO), a trust-region optimization method that efficiently guarantees constrained policy updates by combining a proximal update rule with occasional “fixup” steps.

Efficient Multi-Task Learning via Iterated Single-Task Transfer *Presentation at IROS 2022*

K.R. Zentner, Ujjwal Puri, Yulun Zhang, Ryan Julian, Gaurav S. Sukhatme
Investigated the effectiveness of Transfer RL on the Meta-World benchmark, discovering an asymmetric transfer pattern. Introduced Value Function Warmup and Inference Transfer, allowing simple Transfer Learning using on-policy RL.

Generating Behaviorally Diverse Policies with Latent Diffusion Models *Poster at NeurIPS 2023*

Shashank Hegde*, Sumeet Batra*, K.R. Zentner, Gaurav S. Sukhatme
Distilled a Quality-Diversity RL policy archive using a (Language Conditioned) Latent Diffusion Graph Hyper-Network, achieving a compression ratio of 13x and allowing flexible sequencing of behaviors using language.

A Simple Approach to Continual Learning by Transferring Skill Parameters

K.R. Zentner*, Ryan C. Julian*, Ujjwal Puri, Yulun Zhang, Gaurav S. Sukhatme
Introduces a basic framework for Transfer Reinforcement Learning and related experiments on the Meta-World benchmark.

Towards Exploiting Geometry and Time for Fast Off-Distribution Adaptation in Multi-Task Robot Learning *NeurIPS 2020 Workshops*

K.R. Zentner*, Ryan C. Julian*, Ujjwal Puri, Yulun Zhang, Gaurav S. Sukhatme
Demonstrates how to achieve cross-task transfer of robust behavior to a new goal with simple neural architectures and a single demonstration.

Garage RL Toolkit *Maintainer 2018 - 2021*

A toolkit for Meta- and Multi-Task RL research.

Meta-World *Lead Maintainer 2021-2023*

Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Avnish Narayan, Hayden Shively,
Adithya Bellathur, Karol Hausman, Chelsea Finn, Sergey Levine
A Meta and Multi-Task RL benchmark with 50 unique tasks with randomized goals, tuned dense reward functions, and scripted policies for generating demonstrations. I introduced the use of fuzzy-logic for writing reward functions, and was lead maintainer from 2021 through early 2023.

Language-World

A language and robotics benchmark based on Meta-World

doexp

A simple experiment runner, which allows programatically describing large numbers of experiments, distributing experiment across computers, and gracefully handling experiment crashes. The reason I can re-run 1000 experiments per week.

UBRL

A small RL framework designed for training large models.

Work Experience

Research Engineer, Robot Learning, 1X Technologies

September 2024 - Present

Helped use deep reinforcement learning for bipedal locomotion. Created simulated environments for evaluating deep learning based humanoid robotic policies. Maintained various internal tools.

Research Assistant, Robotics Embedded Systems Laboratory

Fall 2018-Spring 2024

Collaborated with other PhD students to publish several papers. Helped develop and maintain the Garage RL toolkit. Was awarded the George Bekey service award for work in the lab.

Software Engineer, Google Search

Fall 2016 - Summer 2018

Implemented fine-grained caching of JavaScript modules that decreased bandwidth used on JavaScript by 50%, while measurably decreasing latency.

Worked in a team implementing infrastructure for Google Web Search.

Researched and implemented Persistent XSS prevention techniques.

Intern at X (formerly Google X / Google Robotics)

Spring 2016

Worked on a team to create a variety of software tools used for robotics research. Implemented a tool for manipulating and visualizing virtual objects in a motion planning environment. Worked with another team member to create a dashboard for visualizing sensor data, including proprioception, object detection, depth data, and video feeds. Designed a novel voxel compression technique.

Engineering Coordinator, Pioneers in Engineering

Spring 2014 - Spring 2015

Managed 40 undergraduate volunteer staff on various robotics related projects.

Assisted in fundraising over 10% of the organization's budget.

Coordinated purchasing, supply chain, and production of multiple prototype robots.

Shutdown or pivoted projects which weren't contributing to the organization's immediate mission.

Provided technical oversight to ensure a successful robotics competition with over 200 high school students.

Staff, Pioneers in Engineering

Fall 2012 - Fall 2015

Design and built a modular system for coordinating multiple robots in a competitive setting.

Implemented controls in the Grizzly Bear, an open hardware motor controller.

Taught other staff how to use tools such as build systems, continuous integration, version control, and EDA.

Placed new staff into existing projects.

Helped lead a team developing an open hardware realtime robotics control system.

Intern, Mozilla

Summer 2015

Implemented the 2015 draft of CSS Containment in the Gecko browser engine.

Student Researcher with Dr. Rastislav Bodik at U.C. Berkeley

Fall 2014

Experimented with using the Rosette solver-aided programming language to perform finite state machine composition.

Research Assistant with Dr. Dan Garcia at U.C. Berkeley

Fall 2013 - Spring 2014

Worked with a group to improve a visual programming system, **Snap!**, used in an introductory computer science course at U.C. Berkeley. Worked with Dr. Dan Garcia to implement a guided video tutorial of **Snap!**.

Intern, Qualcomm

Summer 2014

Created a build tool which used pycparser to parse arbitrary C headers and generate RPC instructions suitable for test instrumentation.

Created a Perforce frontend in wxPython which integrated with a proprietary bug tracking system.

Worked in a team to add database integration and a web frontend to a test automation system.

Intern, Wind River*Summer 2013*

Implemented support for directories in the default development filesystem (netDrv) of the VxWorks RTOS.
Improved VxWorks concurrency support by implementing POSIX thread barriers.

Intern, San Diego Supercomputer Center R.E.H.S.*Summer 2010 & Summer 2011*

Manually and algorithmically generated 3D content using Maya and Python.
Assisted students learning how to use Maya.

Assistant System Administrator, River Valley Charter School*Fall 2010 - Summer 2012*

Upgraded and maintained a high school's computer network.

Research Areas of Interests

Transfer Learning.

Reinforcement Learning.

Offline-Reinforcement Learning.

Imitation Learning.

Large Language Models.

Inductive Logic Programming.

Personal Projects

hot-restart*Summer 2024 - Present*

Lightweight implementation of "edit-and-continue" style debugging, allowing for partial program reloads from inside a debugger session with minimal overhead.

DeltaLog*Spring 2017 - Summer 2018*

Designed and implemented a dialect of Datalog capable of computing fuzzy truth values for facts.
Implemented reverse mode automatic differentiation of DeltaLog programs, such that the weights of rules in DeltaLog can be differentiated with respect to the outputs.

EBDDIN*Summer 2017*

Replicated "Evolving Binary Decision Diagrams with Implicit Neutrality". Implemented a general purpose Evolutionary Strategies library in Rust.

lua-live*Fall 2014*

Modified the Lua virtual machine to allow safely updating the bytecode of a program while the program is executing. Designed a novel technique where an updated copy of the program is partially executed and lazily copied into the primary program.

Angelic*Summer 2013 - Fall 2013*

Designed a memory safe, statically typed scripting language for realtime, memory constrained environments.
Implemented an efficient, language specific virtual machine in C. Partially implemented a compiler in JavaScript, including a parser, type inferencer, and code generator.

NDL3*Summer 2014*

Designed and implemented a datagram transport protocol similar to RDP. NDL3 has a very small code size and no dependencies (not even a memory allocator), which allows it to be used across environments as diverse as a web browser and a microcontroller.

Angel-Stage*Fall 2012*

Worked with a small team to implement a 2D real-time-strategy game for AI players.

Presentations

"CSS Containment - Leave my divs alone!", Mozilla

July 30, 2015

"Efficient multi-task learning via iterated single-task transfer", IROS

2022

Honors

Robotics Embedded Systems Laboratory: George Bekey Service Award
Eta Kappa Nu (HKN) Mu Chapter: Candidate of the Semester Award

2020
Spring 2015