

Education

University of California, Los Angeles (UCLA), Los Angeles, CA

Ph.D. in Computer Science, Artificial Intelligence Concentration; Advisor: Prof. Song-Chun Zhu Sep 2021

- Dissertation: "Learning How and Why: Causal Learning and Explanation from Physical, Interactive, and Communicative Environments"

M.S. in Computer Science; Advisor: Prof. Song-Chun Zhu

Jun 2017

- Thesis: "Learning Complex Functional Manipulations by Human Demonstration and Fluent Discovery"

University of Dayton, Dayton, OH

B.S. in Computer Engineering; Magna Cum Laude; Advisor: Prof. Tarek Taha

May 2015

- Thesis: "High-Performance Declarative Memory through MapReduce"

Skills

Programming Python, C/C++, TensorFlow, PyTorch, Shell, LaTeX, ROS, Javascript, HTML5, CSS, Node.JS, Java, CUDA
Topics Machine Learning, Graphical Models, Robotics, Cognitive Science, Autonomous Driving, Explainable AI, Reinforcement Learning, Bayesian Networks, Statistical Modeling

Professional and Research Experience

Senior Software Engineer, Cruise Automation, San Francisco, CA

Oct 2021 - Present

- Working at the intersection of machine learning and planning for the Behavior Planning team.
- Using imitation learning, PyTorch and ROS to deploy safer and more capable ML models to cars driving on the road today!

Director and President, Center for AI and Robot Autonomy (CARA), Los Angeles, CA

Mar 2021 - Present

- Define updated mission and vision statements, seek alternative funding, and file paperwork to maintain 501(c)(3) status.

Graduate Student Researcher, Computer Science Department, UCLA, Los Angeles, CA

Sep 2016 - Sep 2021

Humans build generalizable and explainable representations of their environment through interaction, observation, imitation, intervention, and language. The following research produced from a collaborative, team-based lab explores how artificial agents can use these five concepts to learn robust and transferable representations of tasks and environments.

Imitation learning: Training a robot to twist open a medicine bottle

- Captured the complex human hand forces required to open seven different medicine bottles with a tactical glove covered in IMUs (inertial measurement units) and force sensors.
- Constructed robot action planner using a haptic network, And-Or Graph, and the generalized Earley parser.
- Result: Questioned the common structure humans see in the procedure to open any medicine bottle vs. the widely varying forces and action sequences used by the robot for each bottle. This prompted an investigation of abstraction and generalization.

Causal learning: Virtual escape room to examine how humans and AI learn transferable causal representations

- Built virtual "escape rooms" to test causal generalization; while the surface-level features change, all rooms are governed by a common abstract causal structure (series of levers) that describes the required actions to "unlock" the room.
- Ran human subject experiments to verify human learners are capable of discovering the correct abstract causal structure.
- Built hierarchical Bayesian model to achieve similar performance as human learners. This causal model was able to solve the escape room while seven state-of-the-art model-free reinforcement learning algorithms failed at the task.
- Result: Both structural abstraction and feature generalization are critical for transfer learning and generalization.

Explainable AI: How can robots explain their behavior to foster trust from humans?

- Expanded prior imitation learning work by building human-understandable visual interfaces to describe the robot's haptic network and And-Or Graph.
- Showed that robots may need to "think" one way to complete a task, but another way to explain their behavior; the explanations that best fostered trust were not the model components that best aided the robot to achieve the task.
- Result: Consider explainability as a first-class citizen when building AI systems that interact with humans.

Representation learning: The role of language in building generalizable representations

- Created a virtual playground for embodied AI to learn interpretable, common-sense representations of its environment.
- Built a simulated environment using Unreal Engine 4 (UE4) that couples language and vision in a scene graph.
- Devised a dataset consisting of images, language labels, and object segmentation.
- Ongoing: Build representation learning algorithm to use language labels to decompose latent encoding of the environment.

Robotics Research Engineer Intern, Center for AI and Robot Autonomy (CARA), Los Angeles, CA Jun 2018 - Mar 2020

- Worked on representation learning project to couple vision and language (details under UCLA Graduate Student Researcher).

Undergraduate Researcher, Air Force Research Lab (AFRL), Dayton, OH May 2014 - Sep 2015

- Conducted Declarative Memory (semantic knowledge retrieval system) research for the ACT-R cognitive architecture.
- Architected a new declarative memory system using CUDA, thread pools, parsers, and inter-process communication (IPC).
- Continued project work between summers as undergrad thesis research.
- Result: Parallelized declarative retrievals; yielded a 100x speedup over the fastest existing implementation.

Software Engineering Intern, Aviation Department, Garmin International, Olathe, KS May 2013 - Aug 2013

- Automated the testing process for small craft airplane ACARS systems that send timed status messages to ground stations.
- Reduced testing time by 40% and saved hundreds of vendor certification testing hours by optimizing simulation timing protocols and adhering to FAA safety standards.

Teaching Experience

Adjunct Professor, Computer Science Department, Santa Monica College, Santa Monica, CA Jun 2016 - Present

- Teach one-two 45-student classes/quarter: lead lectures, hold office hours and create course materials.
- Instructed 25 courses: Internet Programming (HTML, CSS, JavaScript, MySQL, and PHP), Intro to C, and Intro to C++.

Teaching Assistant, Computer Science Department, UCLA, Los Angeles, CA Sep 2015 - Jun 2016

- Teaching assistant for Introduction to C, Introduction to C++: led discussion hours (~50 students), held office hours.

Teaching Assistant, Electric & Computer Engineering Department, University of Dayton, Dayton, OH Jan 2015 - May 2015

- Teaching assistant for Electronic Devices Lab: aided in lab sessions (~40 students).

Enrichment Workshop Tutor, School of Engineering, University of Dayton, Dayton, OH Sep 2012 - May 2015

- Tutor for first-year engineering students covering calculus, chemistry, and physics.
- Led team of tutors: managed 8 tutors overseeing 40 students.

Summer School Teacher, Cristo Rey Kansas City High School, Kansas City, MO May 2011 - Aug 2012

- Taught four junior and senior student classes at a prep school focused on college placement for underrepresented groups.

Speaking Experience

Speaker, Conferences and Industry Events Aug 2017 - Present

- Presented authored papers at four conferences: AAI 2020, CogSci 2018, CoRL 2017, and IROS 2017.
- Invited to speak at three industry events; Keynote Speaker at the RSS 2018 Causal Imitation Workshop.

Co-Host, Deep End Theory Podcast, UCLA Radio, Los Angeles, CA Jan 2016 - Jan 2020

- Created and hosted an hour-long weekly interview and live performance podcast with 29K+ listens.
- Interviewed 32 electronic music artists including multiple Grammy-nominated guests.
- Edited and released over 70 episodes, DJ-ed 40 live mixes, created branding and marketing material to build fanbase.

Publications

- Four journal publications in Science Robotics, Engineering, and TPDS. Media outlets including IEEE Spectrum, Cosmos Magazine, and The Conversation covered the Science Robotics Feature Article, "A tale of two explanations: Enhancing human trust by explaining robot behavior."
- Nine conference publications for AAI, CVPR, IROS, ICRA, and CogSci on topics including cognitive science, causal learning, imitation learning, robotics, and commonsense reasoning.
- For a complete list of publications, please see my CV, my website (mjedmonds.com) or Google Scholar (ycAoqlAAAAJ).