Mark **Edmonds**

Staff AI/ML Applied Scientist · Foundation Models for Autonomous Driving

mark@mjedmonds.com | mjedmonds.com | | | |

Summary

Staff AI/ML applied scientist building foundation models for autonomous driving. At General Motors, I develop end-to-end driving models trained with imitation and self-play reinforcement learning. At Cruise, I co-developed DriveGPT (ICML 2025), which demonstrated data and model scaling laws for autoregressive driving behavior models, and distilled large models for on-vehicle deployment. PhD in Computer Science from UCLA on causal learning, imitation learning, and explainable AI, with work published in Science Robotics, ICML, CVPR, and AAAI.

Experience

Staff AI/ML Applied Scientist, General Motors, Remote — Feb 2026 - Present

- Develop end-to-end driving models trained with imitation and self-play reinforcement learning.
- Joined GM through its acquisition of Cruise.

Senior Applied Scientist II, Cruise, Remote — Apr 2024 - Feb 2026

- Co-developed DriveGPT, scaling autoregressive behavior models for driving to 1.4B parameters and 120M training segments (∼50x more data than prior published models); demonstrated data and model scaling laws for driving (ICML 2025).
- Applied imitation and reinforcement learning to trajectory generation for trajectory planning.
- Distilled large behavior models into compact models for on-vehicle deployment.
- Co-inventor on a patent application for autoregressive models for autonomous agents (US 2026/0131822 A1).

Senior Applied Scientist, Cruise, Remote — Oct 2021 - Apr 2024

- Worked at the intersection of machine learning and motion planning on the Behavior Planning team, starting with maneuver planning.
- Applied imitation learning to trajectory generation for behavior planning; used PyTorch and ROS to deploy safer, more capable ML models to vehicles driving on public roads.

Director and President, Center for AI and Robot Autonomy (CARA), Los Angeles, CA — Mar 2021 - May 2025

- Defined updated mission and vision statements, sought alternative funding, and filed 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.

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

- Built a simulated environment in Unreal Engine 4 that couples vision and language in a scene graph for embodied AI research.

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.

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"

Journal Publications

[4] L. Yuan*, X. Gao*, Z. Zheng*, **M. Edmonds**, Y.N. Wu, F. Rossano, H. Lu, Y. Zhu, S.C. Zhu
"In situ bidirectional human-robot value alignment," *Science Robotics 2022.*

[3] Y. Zhu, T.Gao, L. Feng, S. Huang, **M. Edmonds**, H. Liu, F. Gao, C. Zhang, S. Qi, Y.N. Wu, J. Tenenbaum, S.C. Zhu
"Dark, Beyond Deep: A Paradigm Shift to Cognitive AI with Humanlike Common Sense," *Engineering 2020.*

[2] **M. Edmonds***, F. Gao*, H. Liu*, X. Xie*, S. Qi, B. Rothrock, Y. Zhu, Y.N. Wu, H. Lu, S.C. Zhu
"A tale of two explanations: Enhancing human trust by explaining robot behavior," *Science Robotics 2019.*

[1] **M. Edmonds**, T. Atahary, S. Douglass, T. Taha.
"Hardware Accelerated Semantic Declarative Memory Systems through CUDA and MapReduce," *TPDS 2018.*

(* indicates equal contribution)

Conference Publications

[10] X. Huang, E.M. Wolff, P. Vernaza, T. Phan-Minh, H. Chen, D.S. Hayden, **M. Edmonds**, B. Pierce, X. Chen, P.E. Jacob, X. Chen, C. Tairbekov, P. Agarwal, T. Gao, Y. Chai, S.S. Srinivasa
"DriveGPT: Scaling Autoregressive Behavior Models for Driving," *ICML 2025.*

[9] C. Zhang, B. Jia, **M. Edmonds**, S.C. Zhu, Y. Zhu
"ACRE: Abstract Causal REasoning Beyond Covariation," *CVPR 2021.*

[8] **M. Edmonds**, X. Ma, S. Qi, Y. Zhu, H. Lu, S.C. Zhu ***Oral Pres.***
"Theory-based Causal Transfer: Integrating Instance-level Induction and Abstract-level Structure Learning," *AAAI 2020.*

[7] **M. Edmonds**, S. Qi, Y. Zhu, J. Kubricht, S.C. Zhu, H. Lu.
"Decomposing Human Causal Learning: Bottom-up Associative Learning and Top-down Schema Reasoning," *CogSci 2019.*

[6] **M. Edmonds***, J. Kubricht*, Colin Summers, Y. Zhu, B. Rothrock, S.C. Zhu, H. Lu. ***Oral Pres.***
"Human Causal Transfer: Challenges for Deep Reinforcement Learning," *CogSci 2018.*

[5] X. Xie*, H. Liu*, **M. Edmonds**, F. Gao, S. Qi, Y. Zhu, B. Rothrock, S.C. Zhu.
"Unsupervised Learning of Hierarchical Models for Hand-Object Interactions," *ICRA 2018.*

[4] **M. Edmonds***, F. Gao*, X. Xie, H. Liu, S. Qi, Y. Zhu, B. Rothrock, & S.C. Zhu. ***Oral Pres.***
"Feeling the Force: Integrating Force and Pose for Fluent Discovery through Imitation Learning to Open Medicine Bottles," *IROS 2017.*

[3] H. Liu*, X. Xie*, M. Millar*, **M. Edmonds**, F. Gao, Y. Zhu, V. Santos, B. Rothrock, & S.C. Zhu. ***Oral Pres.***
"A Glove-based System for Studying Hand-Object Manipulation via Pose and Force Sensing," *IROS 2017.*

[2] **M. Edmonds**, T. Atahary, T. Taha, & S. Douglass.
"High Performance Declarative Memory Systems through MapReduce," *SNPD 2015.*

[1] D. Prince, **M. Edmonds**, A. Sutter, M. Cusumano, W. Lu, & V. Asari.
"Brain Machine Interface using Emotiv EPOC to control Robai Cyton Robotic Arm," *NAECON 2015.*

(* indicates equal contribution)

Honors, Awards & Recognition

Patent: Autoregressive Models for Autonomous Agents, US 2026/0131822 A1 (application published), USPTO 2026
Press: People prefer robots to explain themselves – and a brief summary doesn't cut it, The Conversation, Online 2020
Press: A Robot That Explains Its Actions Is a First Step Towards AI We Can (Maybe) Trust, IEEE Spectrum, Online 2019
NSF Doctoral Consortium, IROS 2017, Vancouver, BC 2017
The Anthony Horvath and Elmer Steger Award of Excellence, University of Dayton, Dayton, OH 2015
Eta Kappa Nu IEEE Honor Society, Member, Dayton, OH 2014
Tau Beta Pi Engineering Honor Society, Member, Dayton, OH 2014
Eagle Scout, Boy Scouts of America, Kansas City, MO 2011

Invited Talks

Causal Transfer: Challenges for Causal Learning and Reinforcement Learning White Mountain, NH
Student speaker at the ONR MURI Meeting Sep 2018

Causal Imitation: The Necessity of Integrating Observations and Interventions Pittsburgh, PA
Keynote Speaker at the RSS Causal Imitation Workshop Jun 2018

Feeling the Force: Integrating Force and Pose for Imitation Learning Los Angeles, CA
Student speaker at the ONR MURI Meeting Aug 2017

Conference Presentations

Theory-based Causal Transfer: Integrating Instance-level Induction and Abstract-level Structure Learning New York, NY
AAAI 2020 Feb 2020

Human Causal Transfer: Challenges for Deep Reinforcement Learning Madison, WI
CogSci 2018 Jul 2018

Feeling the Force: Integrating Force and Pose for Imitation Learning Mountain View, CA
CoRL 2017 Lightning Talk Nov 2017

Feeling the Force: Integrating Force and Pose for Imitation Learning Vancouver, BC
IROS 2017 Sep 2017

Skills

Tools Python, C++, PyTorch, TensorFlow, CUDA, ROS

Methods Foundation models, Transformers, Scaling laws, Imitation learning, Reinforcement learning (incl. self-play), Model distillation, Behavior prediction and motion planning, Autonomous driving, Robotics, Causal learning, Explainable AI, Bayesian models

Teaching Experience

Adjunct Professor, Computer Science Department, Santa Monica College, Santa Monica, CA Jun 2016 - Sep 2021

- Taught one-two 45-student classes/quarter: led lectures, held office hours and created 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.

Passion Projects

Co-Founder and Community Steward, The Hearth, Boston, MA May 2025 - Present

- Co-founded a collectively governed social club and co-stewarded art, event, and community space, now 130 members.
- Lead governance, finances, and digital infrastructure, including the governance model and membership system.
- Took the financial model from zero to one, reaching financial sustainability within the first month.

Artist and Co-Creator, NikuAI, Cambridge, MA 2023

- Junkyard RAVE: an interactive sculpture for performing music by steering the latent space of six custom-trained RAVE (Realtime Audio Variational autoEncoder) models; steel levers and switches read by Arduinos drive the models through Max MSP. Built from recycled junkyard steel and shown at Arts Around the Lab at MIT.
- MINDVAULT: an AI-assisted short film combining live action with Stable Diffusion, Deforum, and a Dreambooth-fine-tuned actress model, created at the MIT AI for Filmmaking Hackathon. Featured in the 2023 Student World Impact Film Festival (Best Acting; second place for music and AI effects); selected for the 2023 Vancouver Independent International and Lady Filmmakers festivals.

Creator, The Orb

- Built a custom audio guestbook for friends' weddings that records guests' messages, using prompts to guide each guest to reflect on the couple's past, present, or future.
- Adapted it into an interactive demonstration educating audiences on the applications and risks of deepfakes.

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.