

UC Berkeley
FHL Vive Center for
Enhanced Reality
College of Engineering



FHL VIVE CENTER
FOR ENHANCED REALITY

ANNUAL REPORT
2024



MISSION

The main goals of the FHL Vive Center for Enhanced Reality are to sponsor critical fundamental research and high-impact applications in the emerging fields of Virtual Reality (VR), Augmented Reality (AR), and Artificial Intelligence (AI), and at the same time serve as the central hub to facilitate the deployment of disruptive VR, AR, and AI technologies across the Berkeley campus for cross-disciplinary research and education.

We aim to achieve these goals by offering seed grants to our faculty, supervising and facilitating student research activities, and fostering external industry partnerships with stakeholders.

TABLE OF CONTENTS

Mission ————— 02

Founders Messages ————— 04

Leadership Messages ————— 06

Leadership ————— 08

Advisory Board ————— 11

Researchers ————— 12

Vive Center Highlights ————— 14

K-12 STEM Outreach ————— 22


Research Highlights ————— 24

Thank you to our Sponsors ————— 35

FOUNDERS LETTER

WENCHI CHEN

Founder



When we helped establish the FHL Vive Center for Enhanced Reality in 2017, our aim was to sponsor fundamental research and high-impact applications in the emerging fields of Virtual Reality, Augmented Reality, and Artificial Intelligence.

We've seen the impact such research can have. Twenty years ago, the first DARPA Grand Challenge autonomous driving competition led to technology now in use by industry. The Indy Autonomous Challenge, the successor of that challenge, now encourages the next generation of students – including UC Berkeley's AI Racing Tech team – to design the next generation of vehicle autonomy by running multi-car driverless races at speeds approaching 200 miles per hour. Behind the teamwork and friendly competition, the students learn valuable lessons about Sim-to-Real technology development.

The Vive Center's Robot Open Autonomous Racing (ROAR) team has successfully built out programming for students at every grade level. High schoolers attend popular summer camps and some become ROAR Ambassadors, hosting STEM



activities for K-12 students in their area. College students without previous experience get acquainted with autonomous systems on the new ROAR GO go-kart launched this year.

We find it's the students themselves we're most proud of as we watch them climb the ladder from engaging with autonomous models and go-karts to operating the most advanced autonomous race cars ever made. It is so rewarding to hear of ROAR campers turned ambassadors who, in no time, have started at UC Berkeley or UCLA. Or first-year students, new to the field, learning to build their own autonomous go-kart. Or engineering students from other schools on the AI Racing Team, inspired to pursue a doctorate at UC Berkeley.

They inspire us too. They are not just tomorrow's leaders. They are leading the way now.

Cher Wang & Wenchi Chen
Founders, FHL Vive Center for Enhanced Reality

LEADERSHIP MESSAGES

S. SHANKAR SASTRY

Founding Director

I am so gratified by the progress the FHL Vive Center and its programs have made this year. The faculty and students have done a remarkable job advancing the research while captivating our imaginations with innovative strides from program to program.

Things got off to a bang right away as the 2024 Autonomous Vehicles (AV-24s) were unveiled at the CES conference in early January. AI Racing Tech, led

by the Robot Open Autonomous Racing Program (ROAR), was among only three teams selected to demo the new AV-24s around Las Vegas Motor Speedway. Now branded blue-and-gold for the University of California, Berkeley, the new racecar was readied all spring and summer for the international Indy Autonomous Challenge at Indianapolis Motor Speedway in September.

The ROAR GO team competed in the national Autonomous Kart Series for the first time this year, coming in second among vastly more experienced go-kart teams. Meanwhile, we continue building student pipelines with ROAR Ambassadors who bring exciting STEM activities into K-12 classrooms.

In the labs, one team is teaching a robotic arm to play ping pong, predicting an opponent's swing and return trajectory. A quadruped robot project on unsupervised object detection takes inspiration from how animals learn to navigate their environments rather than from classic supervised learning.

Thank you to all our sponsors and supporters who make this work possible.





ALLEN YANG

Executive Director

For 2024, I am happy to report that the FHL Vive Center is pioneering a major initiative to bring cost-efficient, private GPT technologies to the Berkeley campus for education and research purposes. The new Teaching Assistant Intelligence (TAI) project has been granted a top Berkeley domain at: tai.berkeley.edu, and now is open for public beta testing both for the general public and Berkeley students.

A key feature for the TAI project is privacy. The small-sized base GPT models adopted for TAI are

carefully curated to allow Berkeley faculty and students to easily spin up their own local GPT services and websites. The project further develops robust embedding and RAG toolkits to allow users to convert their knowledge base and multimedia documents into searchable vector databases. The whole process can be contained within the Berkeley internal network and there is no need to further connect to any third-party APIs or online services.

The team currently is actively reaching out to a diverse range of Berkeley departments and schools, with the aim to offer the TAI service to different courses and instructors with their multidisciplinary background. The deployment within Berkeley will serve as a proving ground for our AI researchers to gather first-hand user experience for the use of GPT technique in the education market.

We are excited that this new system may offer more Berkeley students a capable 24/7 TA office-hour experience to supplement their in-person lectures. We have also open sourced the project, and welcome other institutions to collaborate with us.

LEADERSHIP



S. SHANKAR SASTRY

Founding Director


Thomas M. Siebel
Professor in Computer
Science



ALLEN YANG

Executive Director

Department of Electrical
Engineering and
Computer Sciences



YI MA

Chief Scientist

Director and Chair
Professor of HKU
Musketters Foundation
Institute of Data Science



SHANNON JACKSON

Chief Creative Officer

Cyrus and Michelle Hadidi
Professor of Rhetoric and
of Theater, Dance, and
Performance Studies



BJOERN HARTMANN

Chief Design Officer

Associate Professor of
Electrical Engineering
and Computer Sciences





ADVISORY BOARD



WENCHI CHEN

President and CEO
VIA Technologies, Inc.
Director, HTC



CHER WANG

Co-founder and Chair
HTC Corporation
Chairperson, VIA
Technologies, Inc.




SAID BAKADIR

Senior Director
Product
Management
Qualcomm Inc.



MARK MELTZER

Senior VP, General
Counsel, Chief
Compliance Officer,
Intuitive Surgical, Inc.



MARK LIU

Former Chairman
Taiwan Semiconductor
Manufacturing Co., Ltd.



THOMAS NESBITT

Founding Director
Center for Health and
Technology
UC Davis Health



HARRY SHUM

Former Executive
VP Technology and
Research
Microsoft Corp.

RESEARCHERS



RUZENA BAJCSY

Professor Emerita
Electrical Engineering
and Computer Sciences



FRANCESCO BORRELLI

FANUC Chair in
Mechanical Systems,
Professor in ME



LUISA CALDAS

Director of xR Lab,
Professor
Architecture



LEE FLEMING

Professor
Industrial Engineering
and Operations Research



RICHARD HERNANDEZ

Bloomberg Chair
Associate Professor
Journalism



JACK MCCAULEY

Innovator in Residence at
Jacobs Hall



JAMES O'BRIEN

Professor
Computer Science



REN NG

Professor
Electrical Engineering
and Computer Sciences



KATHRYN QUIGLEY

Senior Digital Producer
Lawrence Hall of Science



KOUSHIL SREENATH

Assistant Professor
Mechanical Engineering



CLAIRE TOMLIN

Charles A. Desoer Chair in
College of Engineering,
Professor of EECS



STELLA YU



Director
ISCI Vision Group
Adjunct Professor

VIVE CENTER HIGHLIGHTS



UC Berkeley now leads U.S. Autonomous Racing Team AI Racing Tech with a newly re-branded vehicle.

AI RACING TECH PARTNERS AND SPONSORS

	point  one	 ByteTrade	 hitcd
	 PRATT MILLER	 BeamNG	
			









AV-24s Unveiled at CES

The next-gen Dallara AV-24 was unveiled as the “most advanced autonomous racecar ever built” at CES 2024 in Las Vegas as the annual convention opened on January 9. The official vehicle of the Indy Autonomous Challenge, the AV-24 – like the AV-21 model that preceded it – has no driver’s seat. Instead, the cabin contains hardware and controls, cameras and sensors, to enable full automation – all packed into a modified version of a Dallara Indy NXT chassis.

Updates and new features include improved vehicle controls, 360-degree long-range lidar perception, 4D radar perception, enhanced GPS/GNSS accuracy, independent brake controls, and redesigned optical, among other advances.

The Robot Open Autonomous Racing (ROAR) program, together with its multi-university partners comprising the AI Racing Tech team, was one of three top teams selected to demo the new models around the track of the Las Vegas Motor Speedway.

After a few demonstration laps, the speedway lights were turned off for “Racing After Dark,” a spectacular display of three new AV-24s, decked out with LED lights in blue, green, and purple, orbiting the track in near-total darkness.



AI Racing Tech Team at the Indy Autonomous Challenge (IAC) 2024 race on Sept. 6. Photo credit: Walt Kuhn for IAC



AI Racing Tech Team check their software stack. Photo credit: Walt Kuhn for IAC



AI Racing Tech Team gets ready to pull the new AV-24. Photo credit: Walt Kuhn for IAC

AI Racing Tech Revs Up for IAC 2024

The AI Racing Tech team led by UC Berkeley in partnership with UC San Diego and Carnegie Mellon University, spent the spring and summer readying its considerable training and talent to compete in the first-ever head-to-head race of the new AV-24s for the international Indy Autonomous Challenge 2024.

With nine other multi-university teams, the blue-and-gold took to the oval at the famed Indianapolis Motor Speedway in Indianapolis, Indiana – the racing capital of the world – on September 6, 2024.




ROAR GO Medals in National Debut

In May, the new ROAR GO autonomous electric go-kart team earned second place in their first national competition, the 2024 Autonomous Kart Series. The event was held in the heart of U.S. motorsports country, at Purdue University's Grand Prix racetrack, just up the road from the Racing Capital of the World in Indianapolis, Indiana. The Berkeley team also received a special design award for Most Stable Platform at the race.

"It certainly exceeded our expectations," says FHL Vive Center Executive Director Allen Yang, and Principal Investigator of Robot Open Autonomous Racing (ROAR), the center's multi-faceted AI motorsports program. "I am particularly proud about our first entry at AKS against far more experienced teams."

As the most reliable performer during the four-day AKS 2024 testing and competition, the ROAR GO team's special design award for Most Stable Platform for their first entry to the national competition "significantly ratifies" the team's chassis design and construction, according to team lead Michael Wu, a software engineer in Silicon Valley and at VIA Technologies embedded with ROAR GO - and also a UC Berkeley alumnus, M.S.'22 EECS.



Michael Wu, Isabelle Hsu, Allen Yang, and Aaron Xie (ROAR GO TEAM) showcase their Kart at AKS 2024.

But the thrill of the race is only a part of what's at stake, as Yang explains.

"As autonomous driving is being hotly pursued in vehicle manufacturing, tech, and transportation industries, AI racing has created a fantastic platform for college students to get hands-on practice and understanding of the deep challenges and critical safety issues," he says. Students working on safety and assurance for autonomous systems today could help inform future technology, from autonomous vehicle platooning to space travel.

”

AS AUTONOMOUS DRIVING IS BEING HOTLY PURSUED IN VEHICLE MANUFACTURING, TECH, TRANSPORTATION INDUSTRIES, AI RACING HAS CREATED A FANTASTIC PLATFORM FOR COLLEGE STUDENTS TO GET HANDS ON PRACTICE AND UNDERSTANDING HTE DEEP CHALLENGES AND CRITICAL SAFETY ISSUES

**ALLEN YANG,
EXECUTIVE DIRECTOR**

”

ROAR programming and the FHL Vive Center are supported by industry sponsors without which the new ROAR GO team’s participation in the AKS national competition would not be possible. The students, faculty, and leadership send thanks to VIA Technologies, Inc., Hitch Interactive, and Qualcomm Ventures for their contributions to the team’s success.



The ROAR GO Team celebrates their accomplishments in Berkeley, CA.

Teaming Up with BeamNG

In March 2024, the FHL Vive Center for Enhanced Reality announced its newest industry partnership with BeamNG, a leading simulation engine software company based in Bremen, Germany.

Among the Vive Center's projects, the Robot Open Autonomous Racing (ROAR) program has the most synergies with BeamNG. ROAR leads the multi-university AI Racing Tech team, which sends faculty and students around the globe to compete with driverless racecars on legendary courses from Indianapolis to Monza.

Such feats require advanced simulation environments testing autonomous software stacks in a real-world setting. This "Sim-to-Real" technology carries implications not only for the future of transportation, but also for other high-stakes fields demanding critical assurance, such as aerospace, defense, and healthcare.

AI Racing Tech is a leading developer of simulation environments, already supporting integration with BeamNG products. This new partnership will formalize a mutually beneficial relationship incorporating educational and research opportunities, software licensing, technical support, and the like.

"We welcome this opportunity to collaborate with a company that shares our passion for autonomous racing and the advanced technologies and challenging science behind it," said ROAR faculty director S. Shankar Sastry.

ABOUT BEAMNG

BeamNG is a global technology company dedicated to the development of a dynamic simulation platform. Its proprietary soft body physics engine is built in-house and enables highly realistic, real-time simulations. Since 2013, BeamNG's software has been successfully integrated into commercial and academic research and design projects, contributing to a variety of applications from autonomous driving systems to driver training programs.



”

THIS PARTNERSHIP IS SET TO BRING TOGETHER BEAMNG'S EXPERTISE IN SIMULATION TECHNOLOGY WITH BERKELEY'S IMPRESSIVE RECORD OF ACADEMIA AND RESEARCH. THROUGH MUTUAL GROWTH AND INNOVATION, WE AIM TO PUSH BOUNDARIES IN VEHICLE DYNAMICS, SAFETY, AI, AND SIMULATION CAPABILITIES TO CONTRIBUTE TO MORE REALISTIC AND EFFICIENT SIMULATION TECHNOLOGIES AND PAVE THE WAY FOR BREAKTHROUGHS IN RESEARCH AND INDUSTRY.

**THOMAS FISHER
CEO OF BEAMNG**

”

”

WE WELCOME THIS OPPORTUNITY TO COLLABORATE WITH A COMPANY THAT SHARES OUR PASSION FOR AUTONOMOUS RACING AND THE ADVANCED TECHNOLOGIES AND CHALLENGING SCIENCE BEHIND IT.

**SHANKAR SASTRY
FOUNDING DIRECTOR**

”

K-12 STEM OUTREACH


ROAR Academy

In 2024, the ROAR Academy summer program for high schoolers transitioned from a pandemic-constrained online program to an enriching, fully in-person experience on the beautiful UC Berkeley campus. This exciting shift gave participants a more immersive and hands-on learning opportunity.

High school students who have demonstrated an aptitude for academic and professional careers in science, technology, engineering and mathematics (STEM) subjects are eligible to apply. Participants become familiar with Python programming and its applications to autonomous driving, in addition to other evolving STEM topics.

97% of students ranked the course 7 and above	85% of students considered the course well organized and had a good experience	93% of students feel encouraged to attend other ROAR-related projects
---	--	---

The Academy has shown considerable growth every year. In 2024, we were limited by the physical space of our two sessions.




ROAR Ambassadors

The Ambassador program was developed to give back to promising high school students who excel in the Academy and participate in the racing simulation competition. It is designed to bring Berkeley ROAR AI Racing, ROAR Research, and ROAR Academy experience to K-12 schools with the Center’s technical and financial support.

Sponsored activities include organizing AI and Autonomous Driving seminars, participating in ROAR AI Racing Tech, providing program feedback, and potentially becoming an active researcher to help grow Berkeley ROAR. Each ROAR ambassador will be teamed up with one ROAR faculty member and one Berkeley graduate student and offered a \$2,000 grant.

Over the past three years, the ROAR Ambassador program has awarded \$20,000 in grant support to 10 ambassadors. This funding has been instrumental in fostering the development of clubs and programs at participants’ respective high schools, aimed at nurturing students’ interest in STEM-related topics.

2024 Ambassadors



Derek Chen
Sophomore
Burlingame High School
(Burlingame, California)



Mark Menaker
Sophomore
University High School
Irvine, California)



Lynna Xu
Sophomore
Sonoma Academy
(Santa Rosa, California)



Christopher Zhang
Sophomore
John Cooper School
(The Woodlands, Texas)

AS OUR AMBASSADOR PROGRAM STEPS INTO ITS FOURTH YEAR, WE STRIVE TO CONTINUE GROWING SCIENCE, TECHNOLOGY, ENGINEERING, AND MATHEMATICS (STEM) ACTIVITIES IN K-12 EDUCATION

RESEARCH HIGHLIGHTS

Projects in the labs at Cory Hall continue to move science forward. A quadruped robot project on unsupervised object detection takes inspiration from the way animals learn to navigate and manipulate their environments rather than from classic supervised learning. Another team is opening up the possibility of using segmentation masks directly for robot control with a simple algorithm. A third team is teaching a robotic arm to play ping pong, with anticipatory prediction of an opponent's swing and return trajectory.



Quadruped Platform for Unsupervised Object Discovery via Interaction


ABSTRACT: Unsupervised Object Detection leverages the principle that distinct objects tend to move independently of each other, hypothesized to be a core inductive bias in human perceptual learning. This methodology allows for the identification of relevant objects without explicit labels, which is more akin to the way animals learn to navigate and manipulate their environments than classic supervised learning. However, most research on unsupervised object detection to date has been done on static datasets of images or short video clips. Inspired by nature's autonomous agents, we aim to move object detection research towards a robotic platform that can interact with and manipulate objects instead of only passively looking at them.

Towards this goal, we developed:

- A task environment for a robot, inspired by the foraging and object-manipulation behaviors seen in mammals, to pick up and move objects to goal locations;
- A central ROS-based algorithm to coordinate an environment engagement routine, deployed on a Quadruped Unmanned Ground Vehicle (UGV), reminiscent of the adaptable and terrain-agnostic movement strategies in animals;
- An extended Large Language Model (LLM) integration to advance associated research in robust, assured, and trustworthy autonomy.

By improving robustness and interpretability of object perception and manipulation, research on object discovery for autonomous systems may impact real-world applications such as search and rescue operations and assistive technology for the differently-abled. By studying and mimicking the cognitive and functional aspects of animal perception and navigation, we can enhance robotic control architectures to achieve effective operations in out-of-factory scenarios, driving forward the evolution of cognitive robotics.


Work by Daniel Barron*, Robin Dumas*, Bear Haon*, Nakul Srikanth*, Kaylene Stocking, Claire Tomlin



Low - Latency Segmentation Propagation for Robotic Control



ABSTRACT: Models that output open-vocabulary instance segmentation (i.e., segmentation masks in an image showing the exact location of arbitrary objects without a fixed set of object classes) are a promising route to more robust and interpretable policies for robots that need to operate in unstructured environments. However, these models often require significant GPU memory and have high latency, which limits their usefulness as inputs to robot control policies. In this work, we propose a simple algorithm for propagating segmentation masks from a high-latency model across video frames. The algorithm achieves low latency and can be run on a low-performance CPU such as those commonly found in edge computing platforms. This opens up the possibility of using segmentation masks directly for robot control.

Work by Catherine Glossop, Charles Gordon, and Bear Häon.



Role of Uncertainty in Anticipatory Trajectory Prediction for a Ping-Pong Playing Robot


ABSTRACT: Robotic interaction in fast-paced environments presents a substantial challenge, particularly in tasks requiring the prediction of dynamic, non-stationary objects for timely and accurate responses. An example of such a task is ping-pong, where the physical limitations of a robot may prevent it from reaching its goal in the time it takes the ball to cross the table. The scene of a ping-pong match contains rich visual information of a player's movement that can allow future game state prediction, with varying degrees of uncertainty. To this aim, we present a visual modeling, prediction, and control system to inform a ping-pong playing robot utilizing visual model uncertainty to allow earlier motion of the robot throughout the game. We present demonstrations and metrics in simulation to show the benefit of incorporating model uncertainty, the limitations of current standard model uncertainty estimators, and the need for more verifiable model uncertainty estimation. The code is publicly available.



Theory of Human Tetrachromatic Color Experience and Printing

ABSTRACT: Genetic studies indicate that more than 50% of women are genetically tetrachromatic, expressing four distinct types of color photoreceptors (cone cells) in the retina. At least one functional tetrachromat has been identified in laboratory tests. We hypothesize that there is a large latent group in the population capable of fundamentally richer color experience, but we are not yet aware of this group because of a lack of tetrachromatic colors in the visual environment. This paper develops theory and engineering practice for fabricating tetrachromatic colors and potentially identifying tetrachromatic color vision in the wild. First, we apply general d -dimensional color theory to derive and compute all the key color structures of human tetrachromacy for the first time, including its 4D space of possible object colors, 3D space of chromaticities, and yielding a predicted 2D sphere of tetrachromatic hues. We compare this predicted hue sphere to the familiar hue circle of trichromatic color, extending the theory to predict how the higher dimensional topology produces an expanded color experience for tetrachromats. Second, we derive the four reflectance functions for the ideal tetrachromatic inkset, analogous to the well-known CMY printing basis for trichromacy. Third, we develop a method for prototyping tetrachromatic printers using a library of fountain pen inks and a multi-pass inkjet printing platform. Fourth, we generalize existing color tests - sensitive hue ordering tests and rapid isochromatic plate screening tests - to higher-dimensional vision, and prototype variants of these tests for identifying and characterizing tetrachromacy in the wild.

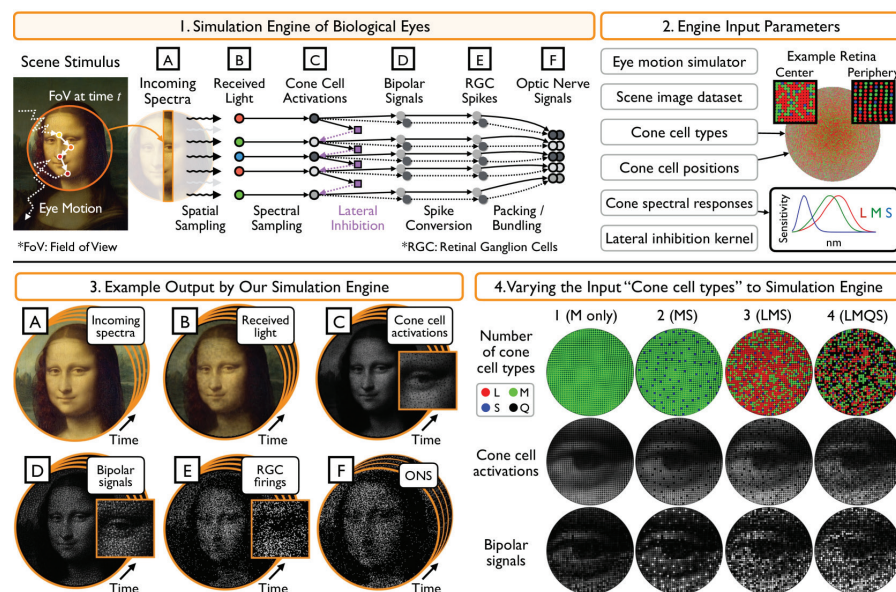
Work by Jessica Lee, Nicholas Jennings, Varun Srivastava, and Ren Ng.



A Computational Framework for Modeling Emergence of Color Vision in the Human Brain

ABSTRACT: It is a mystery how the brain decodes color vision purely from the optic nerve signals it receives, with a core inferential challenge being how it disentangles internal perception with the correct color dimensionality from the unknown encoding properties of the eye. We introduce a computational framework for modeling this emergence of human color vision by simulating both the eye and the cortex. Existing research often overlooks how the cortex develops color vision or represents color space internally, assuming that the color dimensionality is known a priori; however, we argue that the visual cortex has the capability and the challenge of inferring the color dimensionality purely from fluctuations in the optic nerve signals. To validate our theory, we introduce a simulation engine for biological eyes based on established vision science and generate optic nerve signals resulting from looking at natural images. Further, we propose a model of cortical learning based on self-supervised principle and show that this model naturally learns to generate color vision by disentangling retinal invariants from the sensory signals. When the retina contains N types of color photoreceptors, our simulation shows that N -dimensional color vision naturally emerges, verified through formal colorimetry. Using this framework, we also present the first simulation work that successfully boosts the color dimensionality, as observed in gene therapy on squirrel monkeys, and demonstrates the possibility of enhancing human color vision from 3D to 4D.


Work by Atsunobi Kotani and Ren Ng. (arXiv:2408.16916)



SMT-Based Dynamic Multi-Robot Task Allocation

ABSTRACT: Multi-Robot Task Allocation (MRTA) is a problem that arises in many application domains including package delivery, warehouse robotics, and healthcare. In this work, we consider the problem of MRTA for a dynamic stream of tasks with task deadlines and capacitated agents (capacity for more than one simultaneous task). Previous work commonly focuses on the static case, uses specialized algorithms for restrictive task specifications, or lacks guarantees. We propose an approach to Dynamic MRTA for capacitated robots that is based on Satisfiability Modulo Theories (SMT) solving and addresses these concerns. We show our approach is both sound and complete, and that the SMT encoding is general, enabling extension to a broader class of task specifications. We show how to leverage the incremental solving capabilities of SMT solvers, keeping learned information when allocating new tasks arriving online, and to solve non-incrementally, which we provide runtime comparisons of. Additionally, we provide an algorithm to start with a smaller but potentially incomplete encoding that can iteratively be adjusted to the complete encoding. We evaluate our method on a parameterized set of benchmarks encoding multi-robot delivery created from a graph abstraction of a hospital-like environment. The effectiveness of our approach is demonstrated using a range of encodings, including quantifier-free theories of uninterpreted functions and linear or bitvector arithmetic across multiple solvers.

Work by Victoria Tuck, Pei-Wen Chen, Georgios Fainekos, Bardh Hoxha, Hideki Okamoto, S. Shankar Sastry, and Sanjit A. Seshia (arXiv:2403.11737)




A task stream of sets of tasks arrives with monotonically increasing arrival times. Five system sites are shown. Example tasks and robot paths are shown on the right with P, D, and M used to represent the actions succinctly. The result action sequence for the right robot is (M,P,M,D) and (M,P,M,P,M,P,M,D,M,D,M,D). The moves between picks or drops at the same location are used to keep a consistent structure in the plan but take no time.

Incentive-Compatible Vertiport Reservation in Advanced Air Mobility: An Auction-Based Approach

ABSTRACT: The rise of advanced air mobility (AAM) is expected to become a multibillion-dollar industry in the near future. Market-based mechanisms are touted to be an integral part of AAM operations, which comprise heterogeneous operators with private valuations. In this work, we study the problem of designing a mechanism to coordinate the movement of electric vertical take-off and landing (eVTOL) aircraft, operated by multiple operators each having heterogeneous valuations associated with their fleet, between vertiports, while enforcing the arrival, departure, and parking constraints at vertiports. Particularly, we propose an incentive-compatible and individually rational vertiport reservation mechanism that maximizes a social welfare metric, which encapsulates the objective of maximizing the overall valuations of all operators while minimizing the congestion at vertiports. Additionally, we improve the computational tractability of designing the reservation mechanism by proposing a mixed binary linear programming approach that leverages the network flow structure.

Work by Pan-Yang Su, Chinmay Maheshwari, Victoria Tuck, and Shankar Sastry (arXiv:2403.18166)



Schematic representation of the air traffic network with a service provider tasked with coordinating the movement of aircraft of various fleet operators between vertiports in its domain. Each vertiport has a constraint on the number of arriving aircraft, departing aircraft, and parked aircraft.

08 LLM in a Cloudless Day

ABSTRACT: LLM in a Cloudless Day is a Qualcomm-sponsored project, exploring the use of edge AI to facilitate navigation, planning, and future code generation for autonomous robots. Utilizing a Snapdragon System-on-Chip (SoC) integrated into a Samsung Galaxy device, we enable users to interface with an autonomous rover through voice-based interaction with an LLM. Our custom Android application allows for intuitive command generation, telemetry visualization, and a terminal mode for direct interaction with the rover. The LLM processes user inputs to generate function calls to accomplish a desired task. These commands are then sent to and executed by the rover, allowing for high level planning of the rover's actuation. By deploying the LLM on the edge, we allow for real time decision making in cloudless environments, removing the reliance on LLM access via the internet. As the project continues, we aim to extend the LLM's capabilities to include real-time code generation, expanding the rover's autonomous capabilities by adding to the rover's set of executable Robot Operating System (ROS) functions. The project demonstrates a potential implementation of advanced natural language processing with robotics, enabling more accessible human-robot interaction.



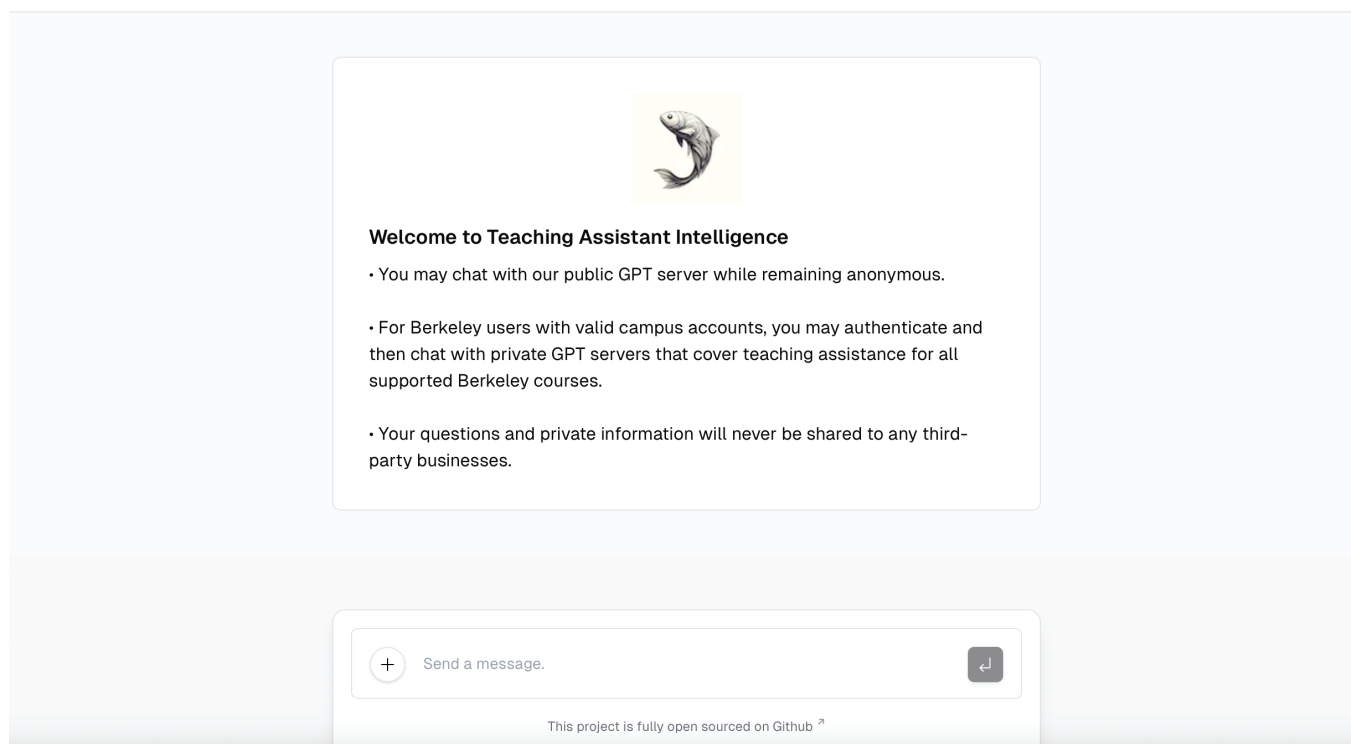
Teaching Assistant Intelligence (TAI) Project

ABSTRACT: The new Teaching Assistant Intelligence (TAI) project at tai.berkeley.edu is now open for public beta testing. A key feature for the TAI project is privacy. The small-sized base GPT models adopted for TAI are carefully curated to allow Berkeley faculty and students to easily spin up their own local GPT services and websites.

The project further develops robust embedding and RAG toolkits to allow users to convert their knowledge base and multimedia documents into searchable vector databases. The whole process can be contained within the Berkeley internal network and there is no need to further connect to any third-party APIs or online services.

The team is reaching out to Berkeley departments and schools, with the aim to offer the TAI service to different courses and instructors.

The deployment within Berkeley will serve as a proving ground for our AI researchers to gather first-hand user experience for the use of GPT technique in the education market.




A glance of [Tai.berkeley.edu](https://tai.berkeley.edu)



THANKS TO OUR SPONSORS





FHL VIVE CENTER

FOR ENHANCED REALITY

WEBSITES

vivecenter.berkeley.edu
roar.berkeley.edu
defi.berkeley.edu

ADDRESS

337 Cory Hall
UC Berkeley
Berkeley, CA 94720

EMAIL

fhlvivecenter@berkeley.edu