Physical Intelligence: Foundation and Systems (ESE 6510)

Instructor:
Antonio Loquercio

Fall 2025
University of Pennsylvania

Today

What is this class about?

- What are you getting into?
- What are you getting out of it?
- What are the expectations?

Why is now a good time to study robot learning?

How we thought
the future would be.

The Jetsons, ca. 1960

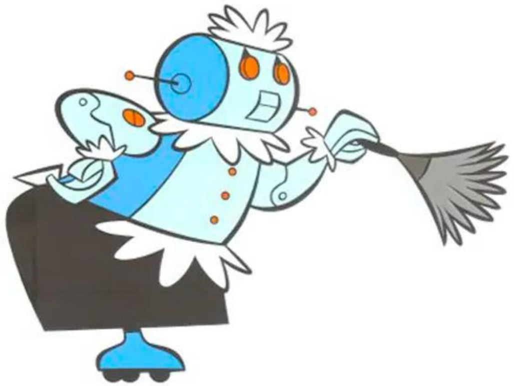

How the future
is going to be.

Moravec’s Paradox

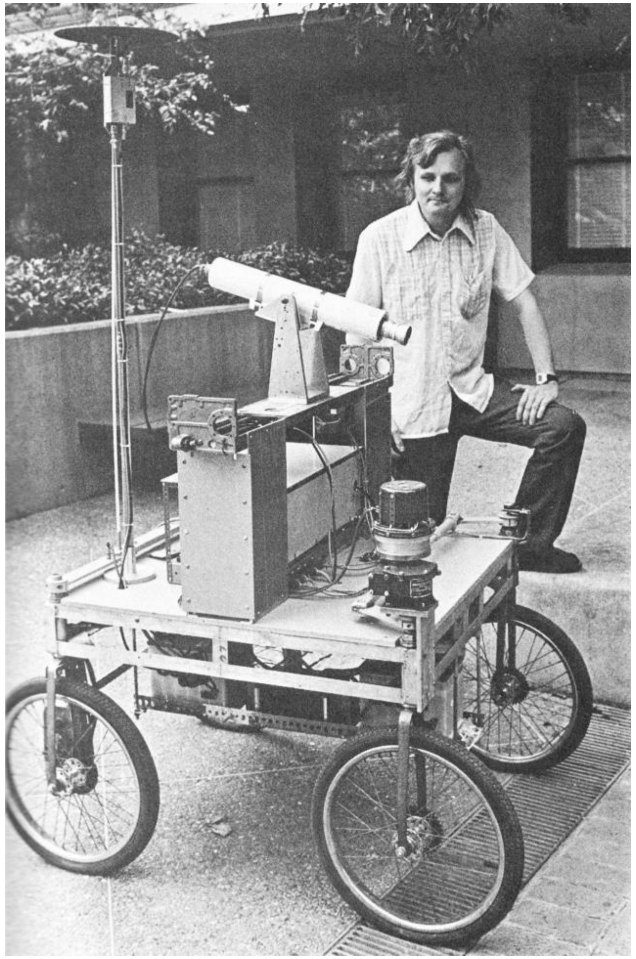

It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility.

Mind Children, H. Moravec, 1988

Hans Moravec's PhD Thesis

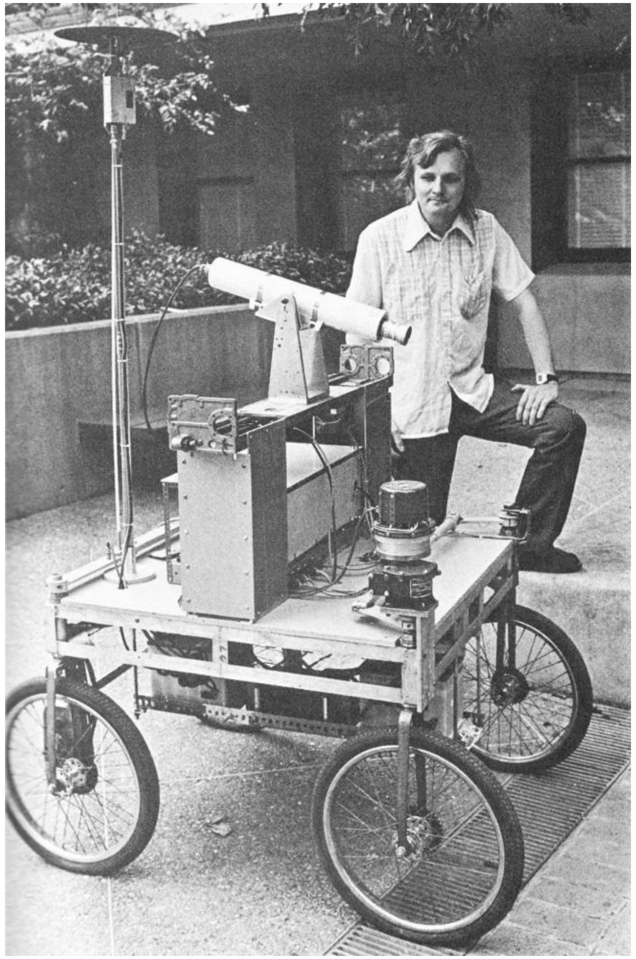

During this period a number of incoming students were assigned to the "cart project'. Each correctly perceived the situation within a year, and went on to something else. The cart's reputation as a serious piece of research apparatus, never too high, sank to new depths.

Why care about **Physical** Intelligence?

"What's so special about robotics?
Surely, the same 'foundational model' should work for everything?"

- In the long term, likely yes
- I don't think we are there yet
    - We don't have enough data
    - We don't have the right data
    - We don't know the right objective(s)
- Even with infinite data, there are challenges that we don't yet know how to solve:
    - Safety (Hallucinations have significant consequences in robotics)
    - Computational efficiency
    - Stuff breaks
    - I am certainly very biased! 😁

Why care about Physical **Intelligence**?

"What if I don't care about this wishy-washy learning stuff?
I just want to make my robot go!"

Small Reason:

If your task is well defined, other stuff might work better (e.g., Roomba)

Bigger Reason:

We have not yet figured out how to build robots that are as adaptable as humans. Learning seems promising (look at CV + NLP).

Biggest Reason:

Science! Emergence of Intelligence is probably the biggest mystery in the universe. We can't help ourselves but try to model it!

Overreaching Theme

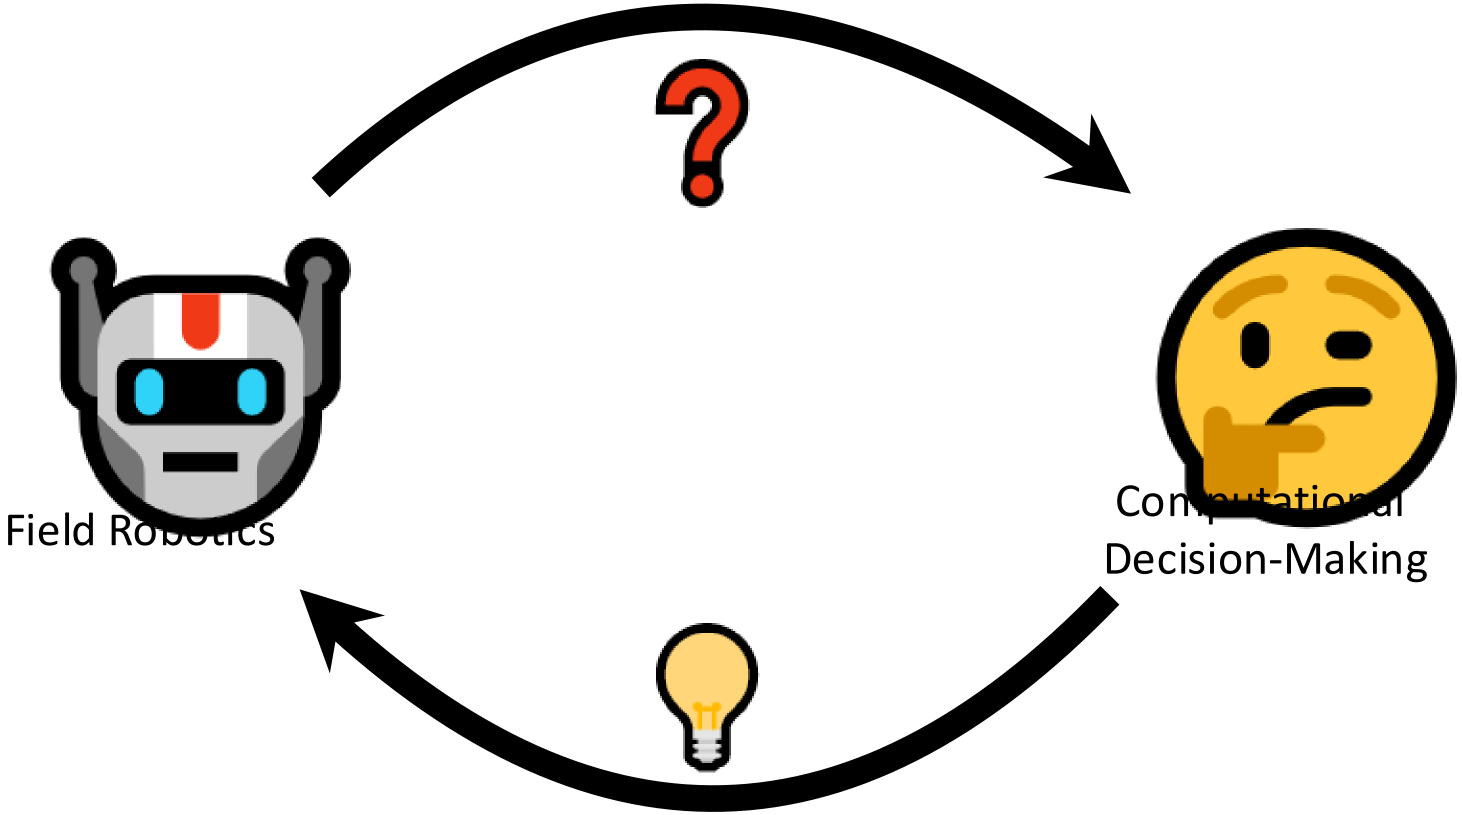

Why care about the **Real World**?

"Why bother with the real-world?
I can study the exact the same problem in simulations."

- We can't yet model everything
  - If we could, we'd be able to predict the future!
  - Unknown unknowns are often what makes a difference.
  - Need to deal with uncertainty
- Our models are limited by human cognition
  - Might be constraining if we're looking for super-human performance
- You are as smart as the environment requires
  - Can't reset or undo things
  - Intelligence ≈ Reward Hacking
  - Everything works given enough constraints

Dunning-Kruger Effect in Robotics

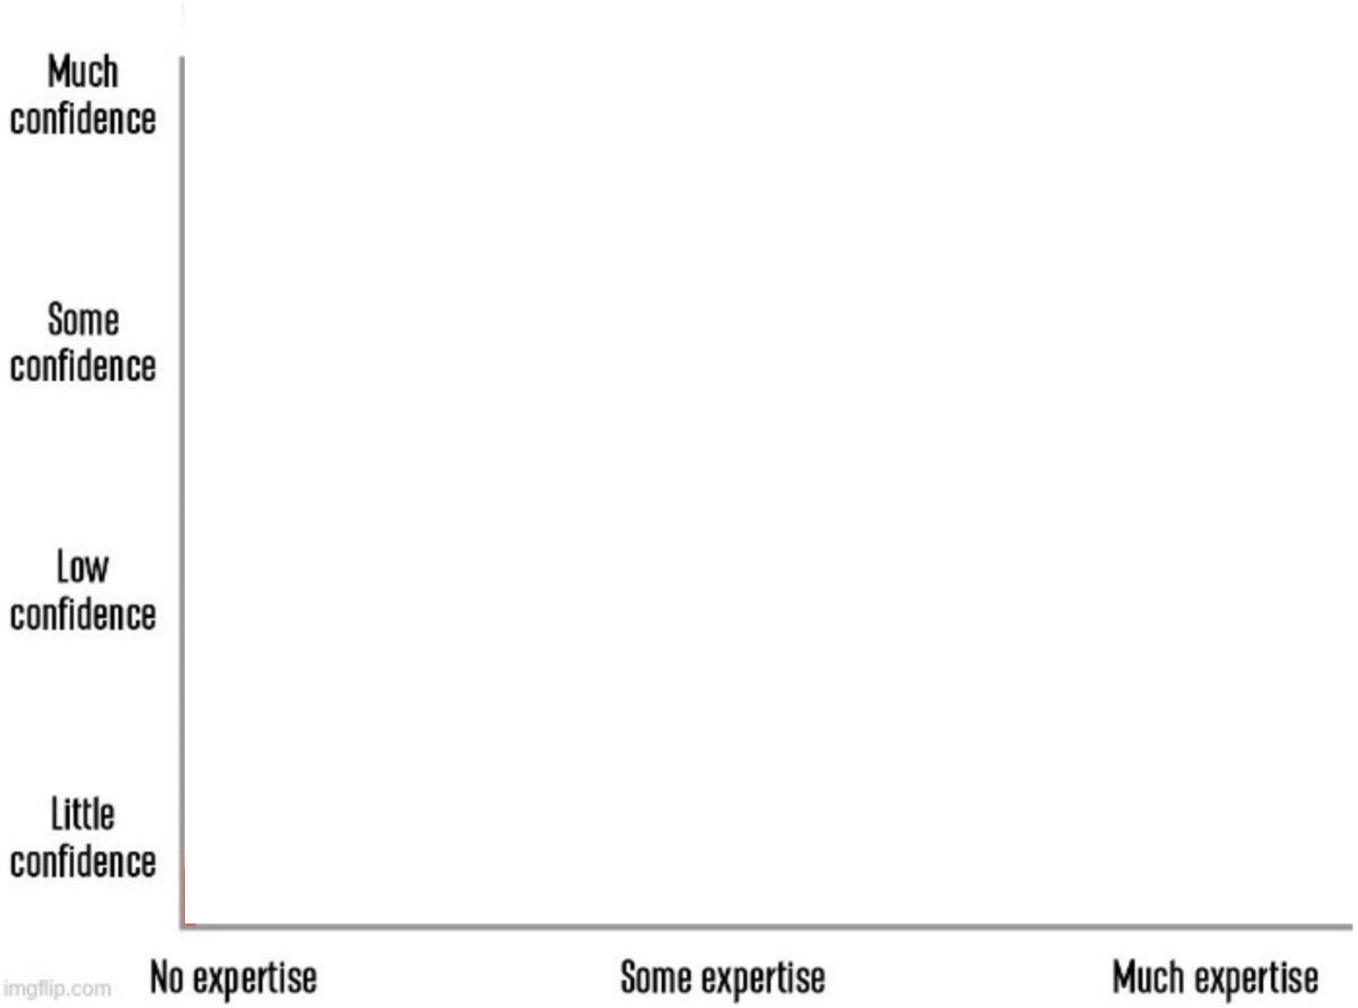

What's different about this class?

- We will see foundational concepts together with fancy systems
  - Study theoretical concepts, then come back to the nitty gritty of making something work
- Understand that hardware is as important as software
  - Contributions are sometimes not what the authors of a paper intended to
- Understand that fancy results are almost never the result of "one" idea
  - But rather, the collection of many small improvements

Article | Open access | Published: 30 August 2023

Champion-level drone racing using deep reinforcement learning

Elia Kaufmann, Leonard Bauersfeld, Antonio Loquercio, Matthias Müller, Vladlen Koltun & Davide Scaramuzza

Nature **620**, 982–987 (2023) | Cite this article

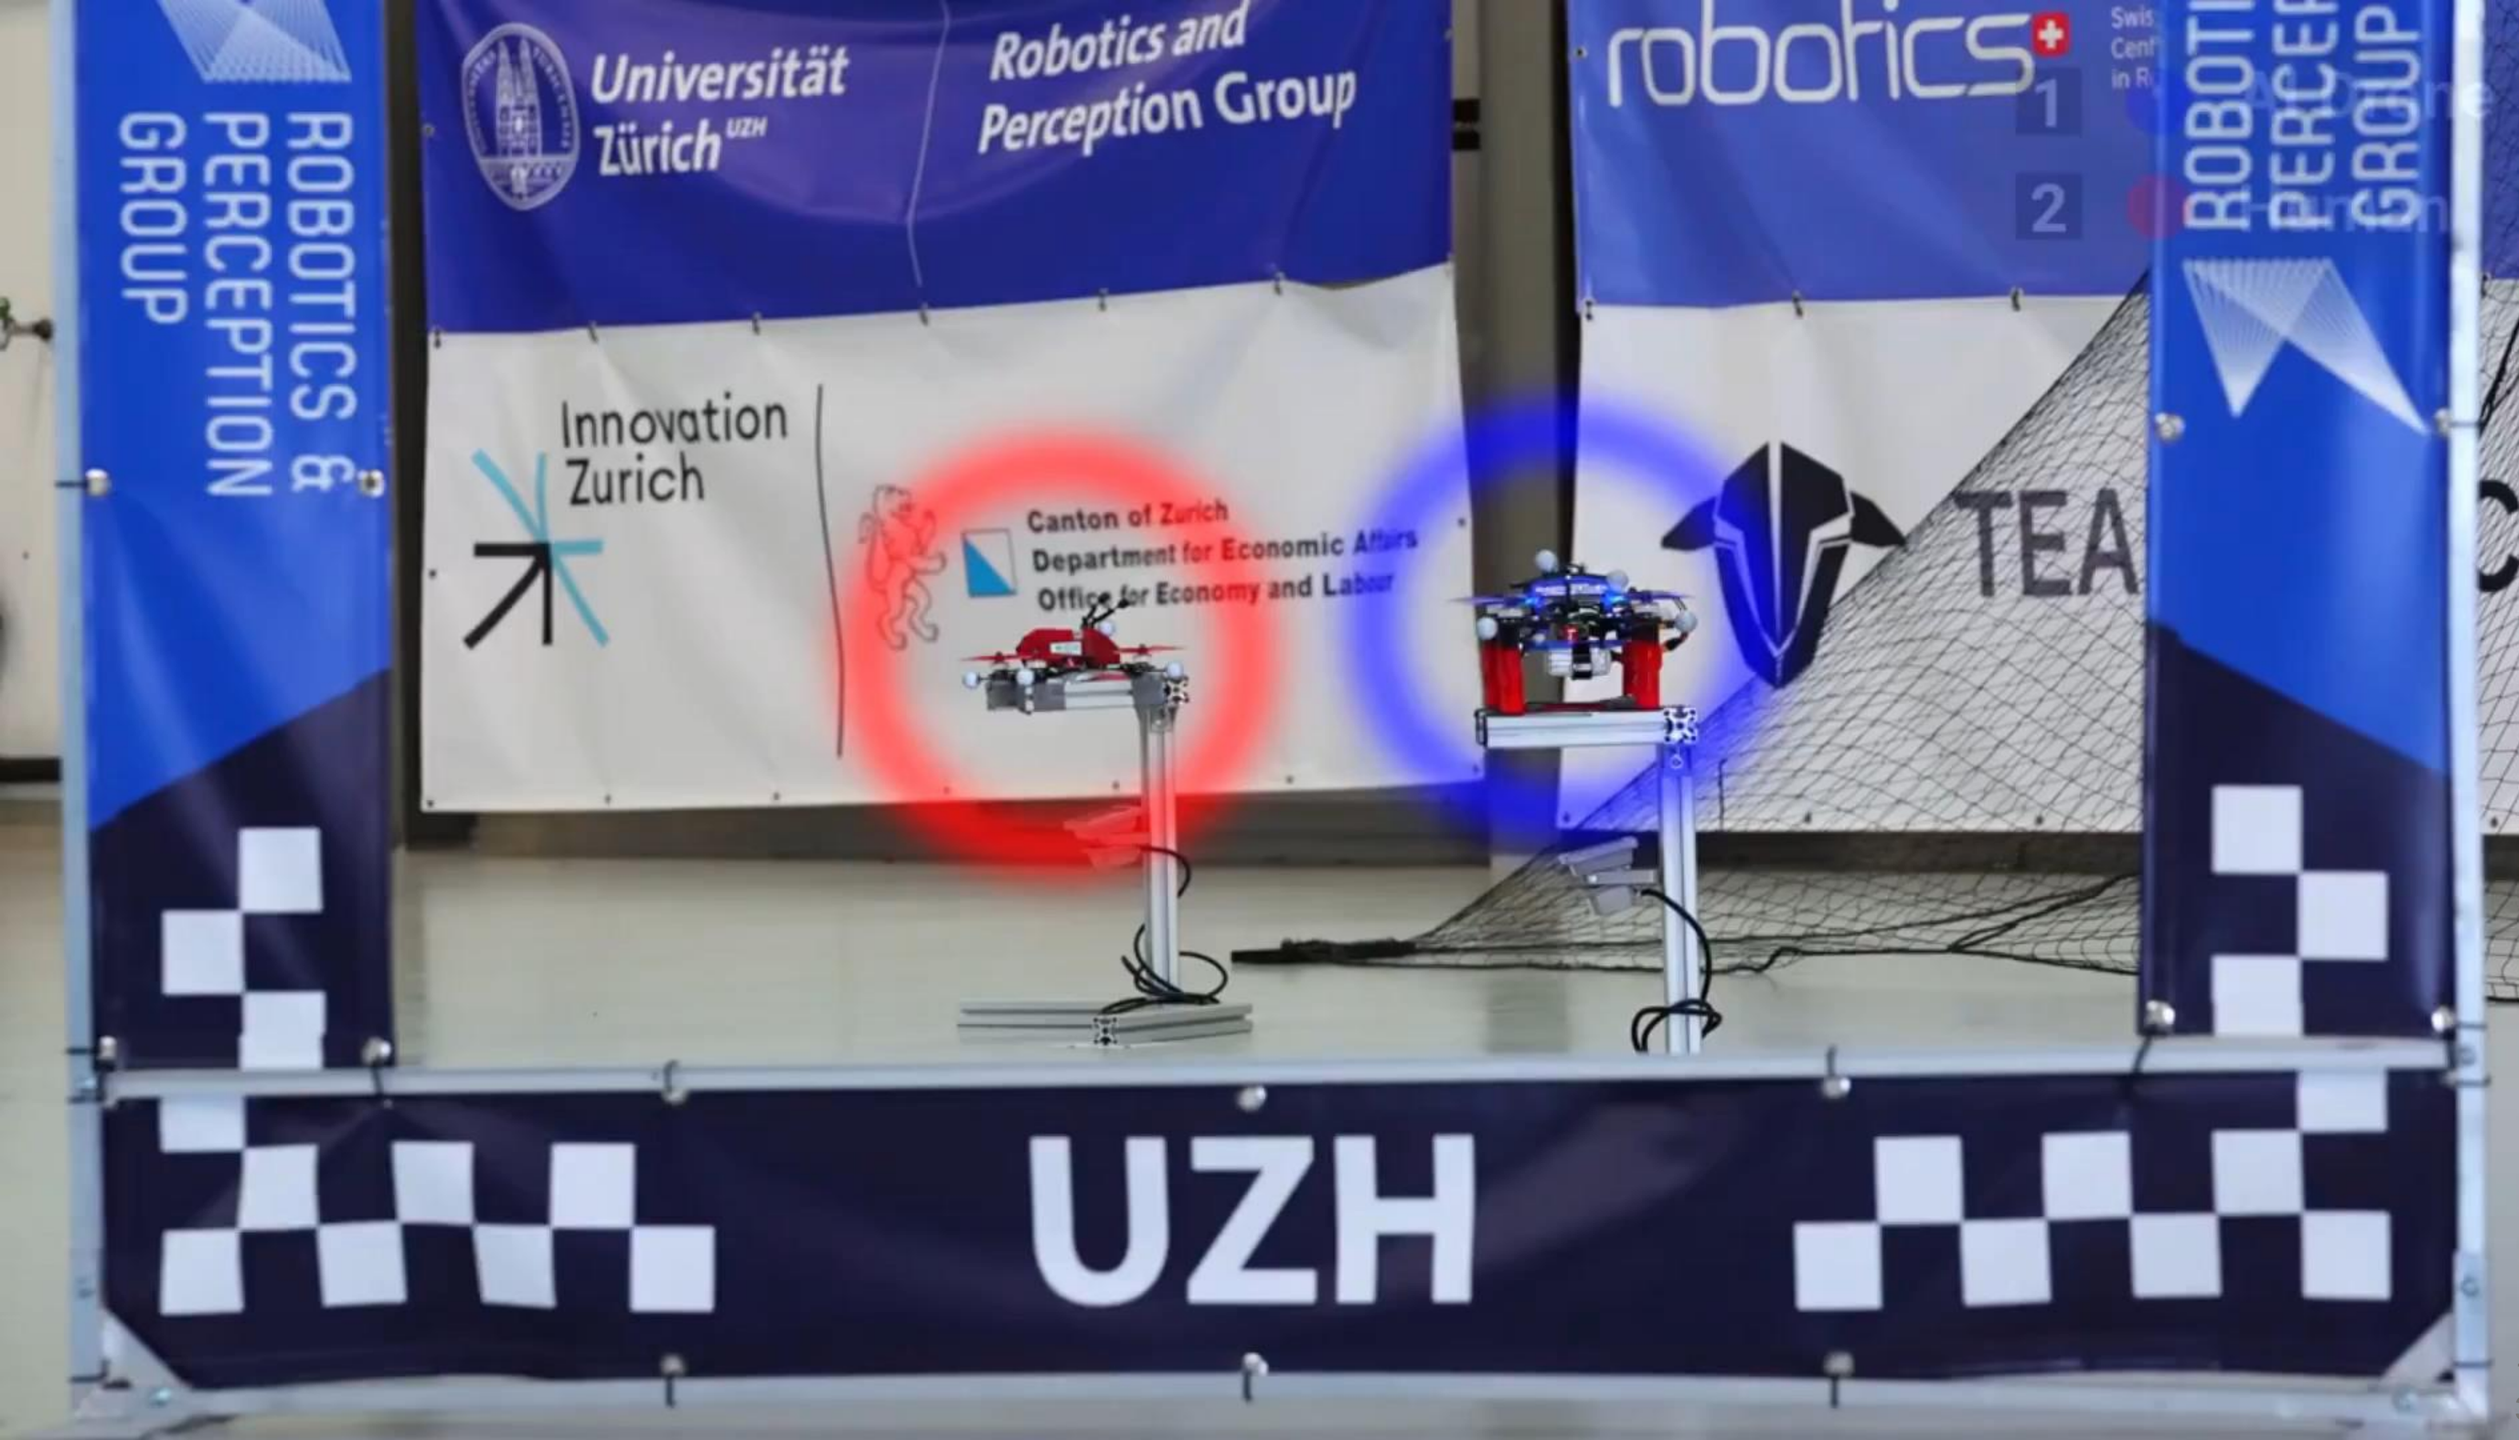
ROBOTICS & PERCEPTION GROUP
Universität Zürich UZH
Robotics and Perception Group
robotics
1
2
Innovation Zurich
Canton of Zurich
Department for Economic Affairs
Office for Economy and Labour
UZH

Easy to Miss the Complexity of the System

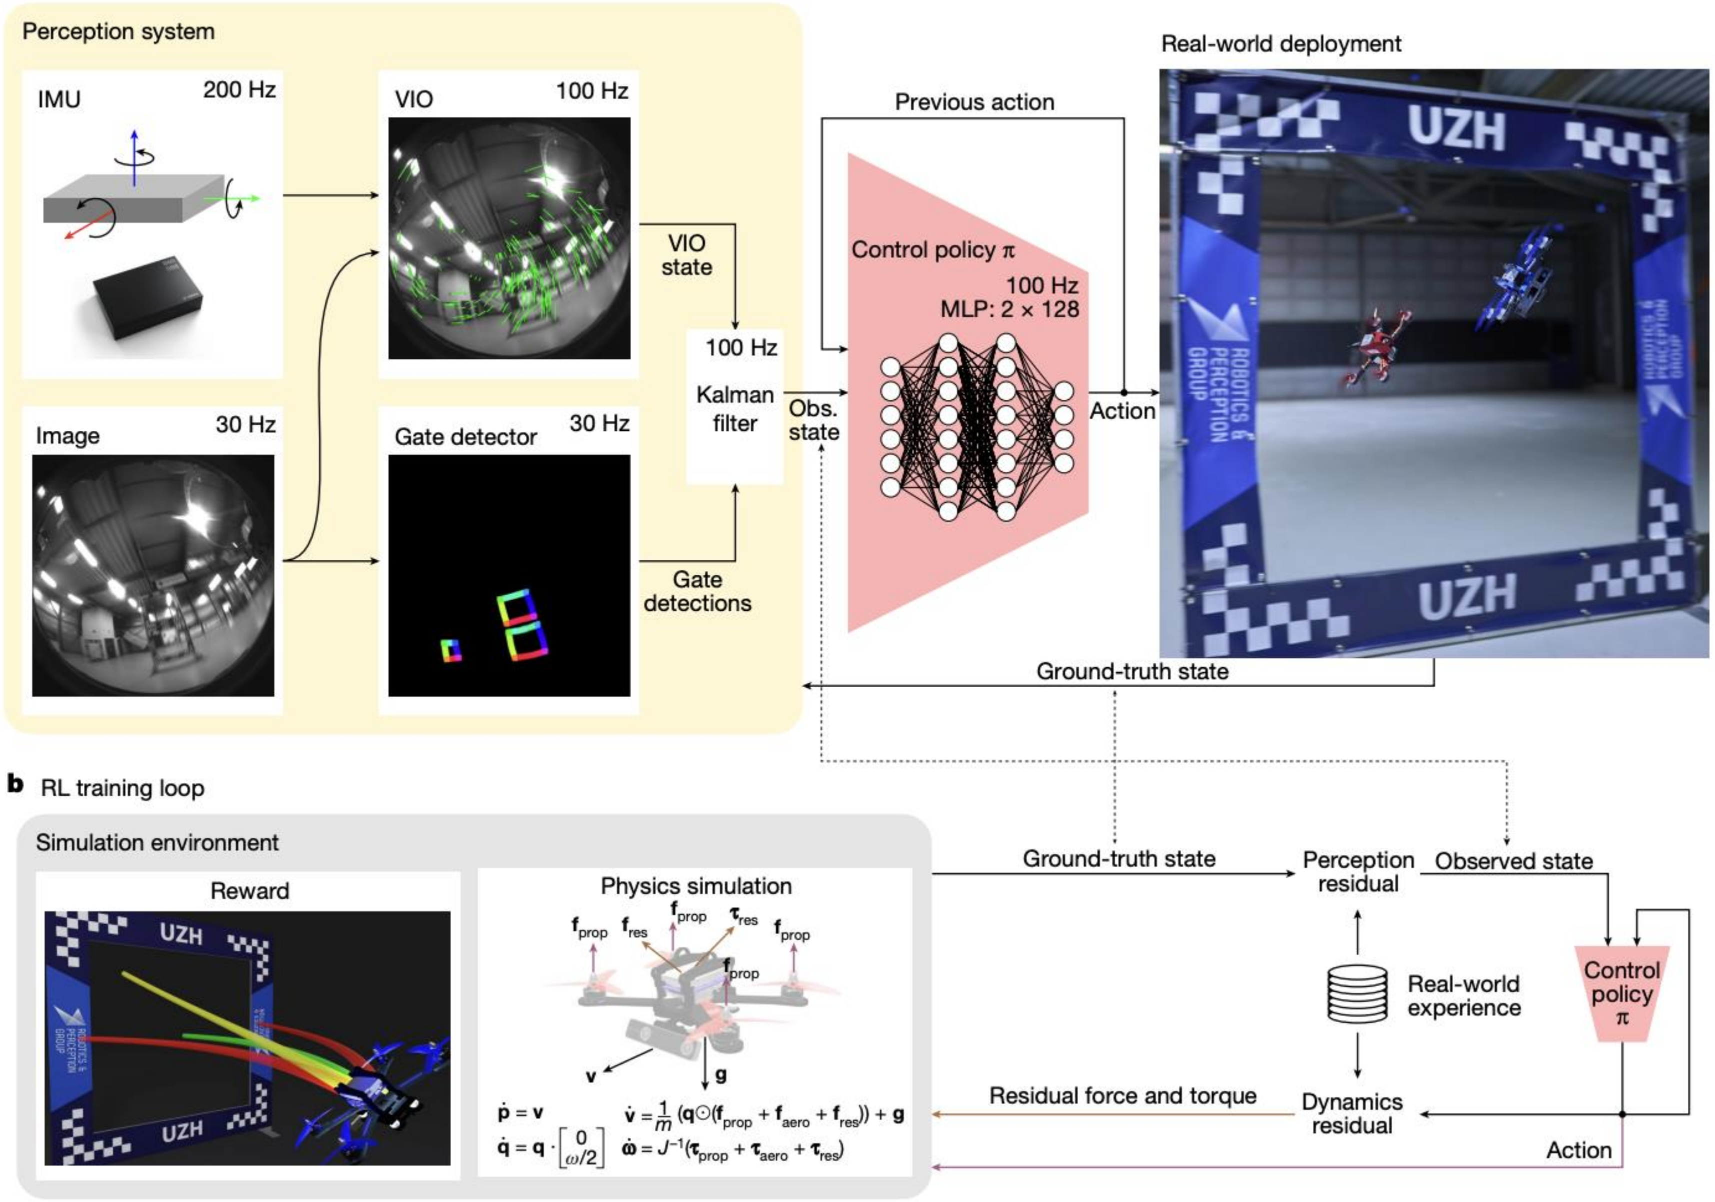

Easy to Miss the Complexity of the System

- The "Infrastructure" (built over three years of work)

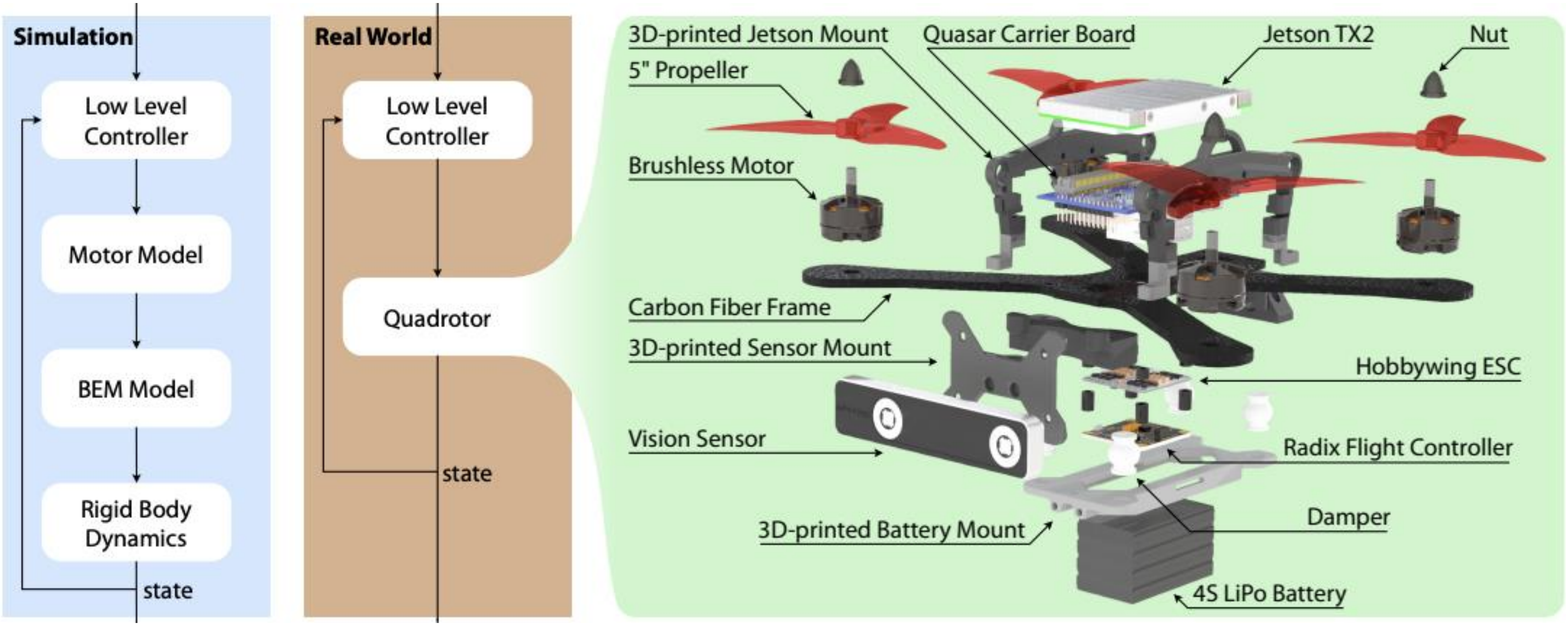

- Latency Control
- Gaussian Mixture Models for Noise Identification
- Custom code to run vision models in embedded hardware

Goals

- Be comfortable with the foundations of policy learning
- Get hands-on experience with robot learning
  - You never truly understand something until you have built it yourself.
- Use more learning-based techniques in your own work
- Learn how to think critically about robot learning papers and demos

Course Material & Tools

- All learning materials will be available on this website

- We will use lot some tools to facilitate interactions:
    - Canvas for announcements
    - Ed for discussion
    - Gradescope for coursework

Course Organization

Grading:

1. Class Participation (5%)
   - **Attend class.** Don't skip more than 2 Quizzes**.** Ask questions, be involved!

2. Midterm (35%)

3. Final Exam (40%)

4. Drone Race (20%)
   - More details on this later
   - Can be done solo or in groups of 2
   - Best teams can skip the final

Recitations (5X)

- We will do exercises/coding together in class
- There will be some quizzes during the tutorial
  - Skipping more than 2 will result in 0 points for participation.
  - Closed book (no laptop or phone)
  - Grade will be binary (1/0). You get 0 if you leave blank or write non-sense.
- You will receive additional exercises on the topic to complete at home.
- The at-home exercises will not be graded.

The Drone Racing Competition: Phase I

The Drone Racing Competition: Phase I

- Time trial, solo racing.

- All groups must participate

- Will be done in simulation.
  - More on the evaluation scheme will be announced later in the semester.

- Three grading tears
  1. The top 3 teams receive a grade bump (e.g., A -> A+).
  2. Top 10 (one bonus point)
  3. Not completing the track under a (very relaxed) maximum time (grade decided depending on what was implemented). These teams won’t be able to move to the next phase.
  4. Everybody else (full score)

- Need to submit a short (1-2pp) report on what was implemented

Phase I’s Winners Last Semester

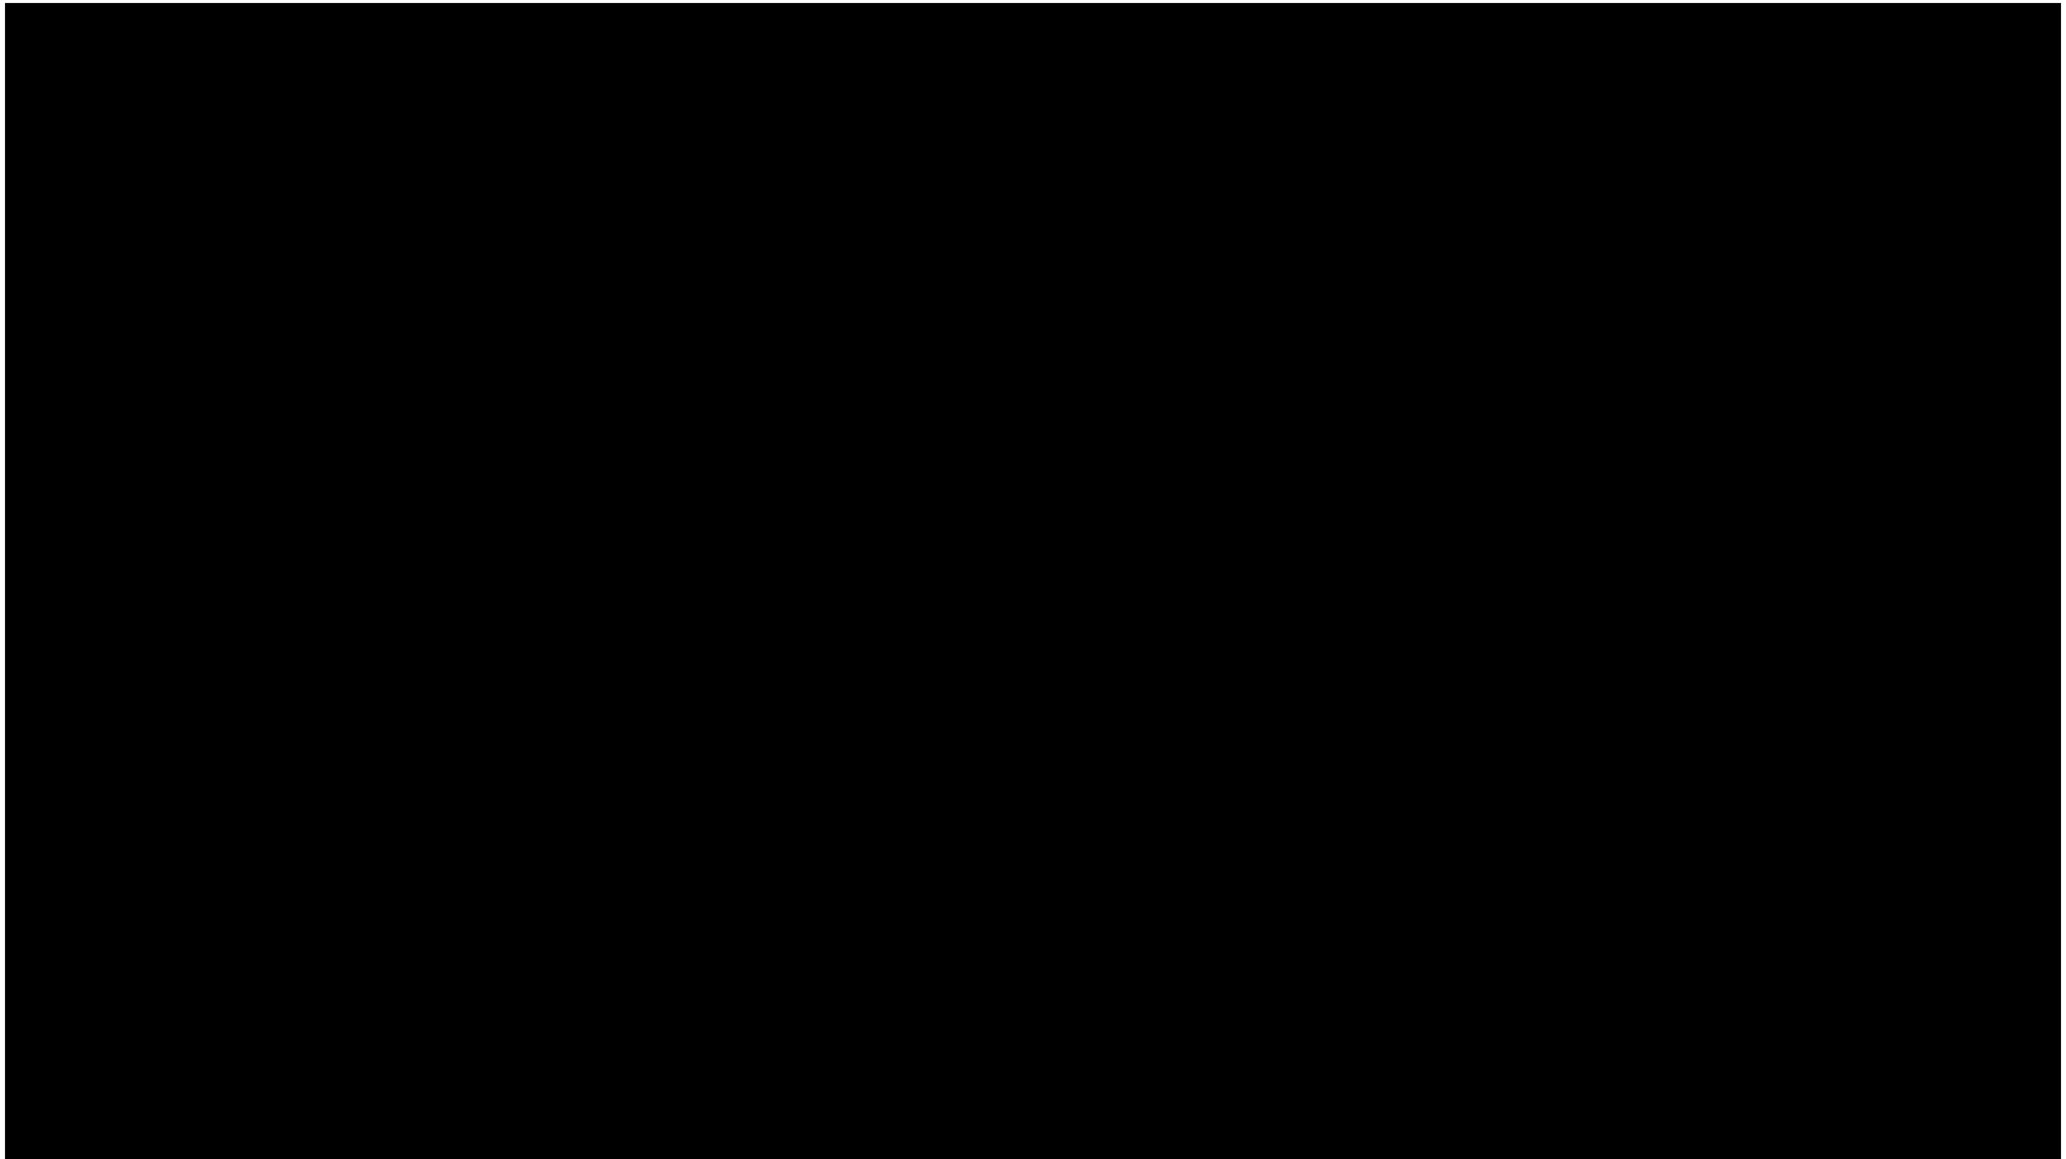

The Drone Racing Competition: Phase II

The Drone Racing Competition: Phase II

- We will race in the real world!

- Prize: dinner in my favorite restaurant in Rome (transportation not included!) or some other worthy prizes 😃

- Two grading tears
  1. The intersection of Phase I and Phase II top-3 teams can skip the final.
  2. Top 10 get one bonus point.
  3. Everybody else (grade decided depending on what was implemented)

Phase II’s Winners Last Semester

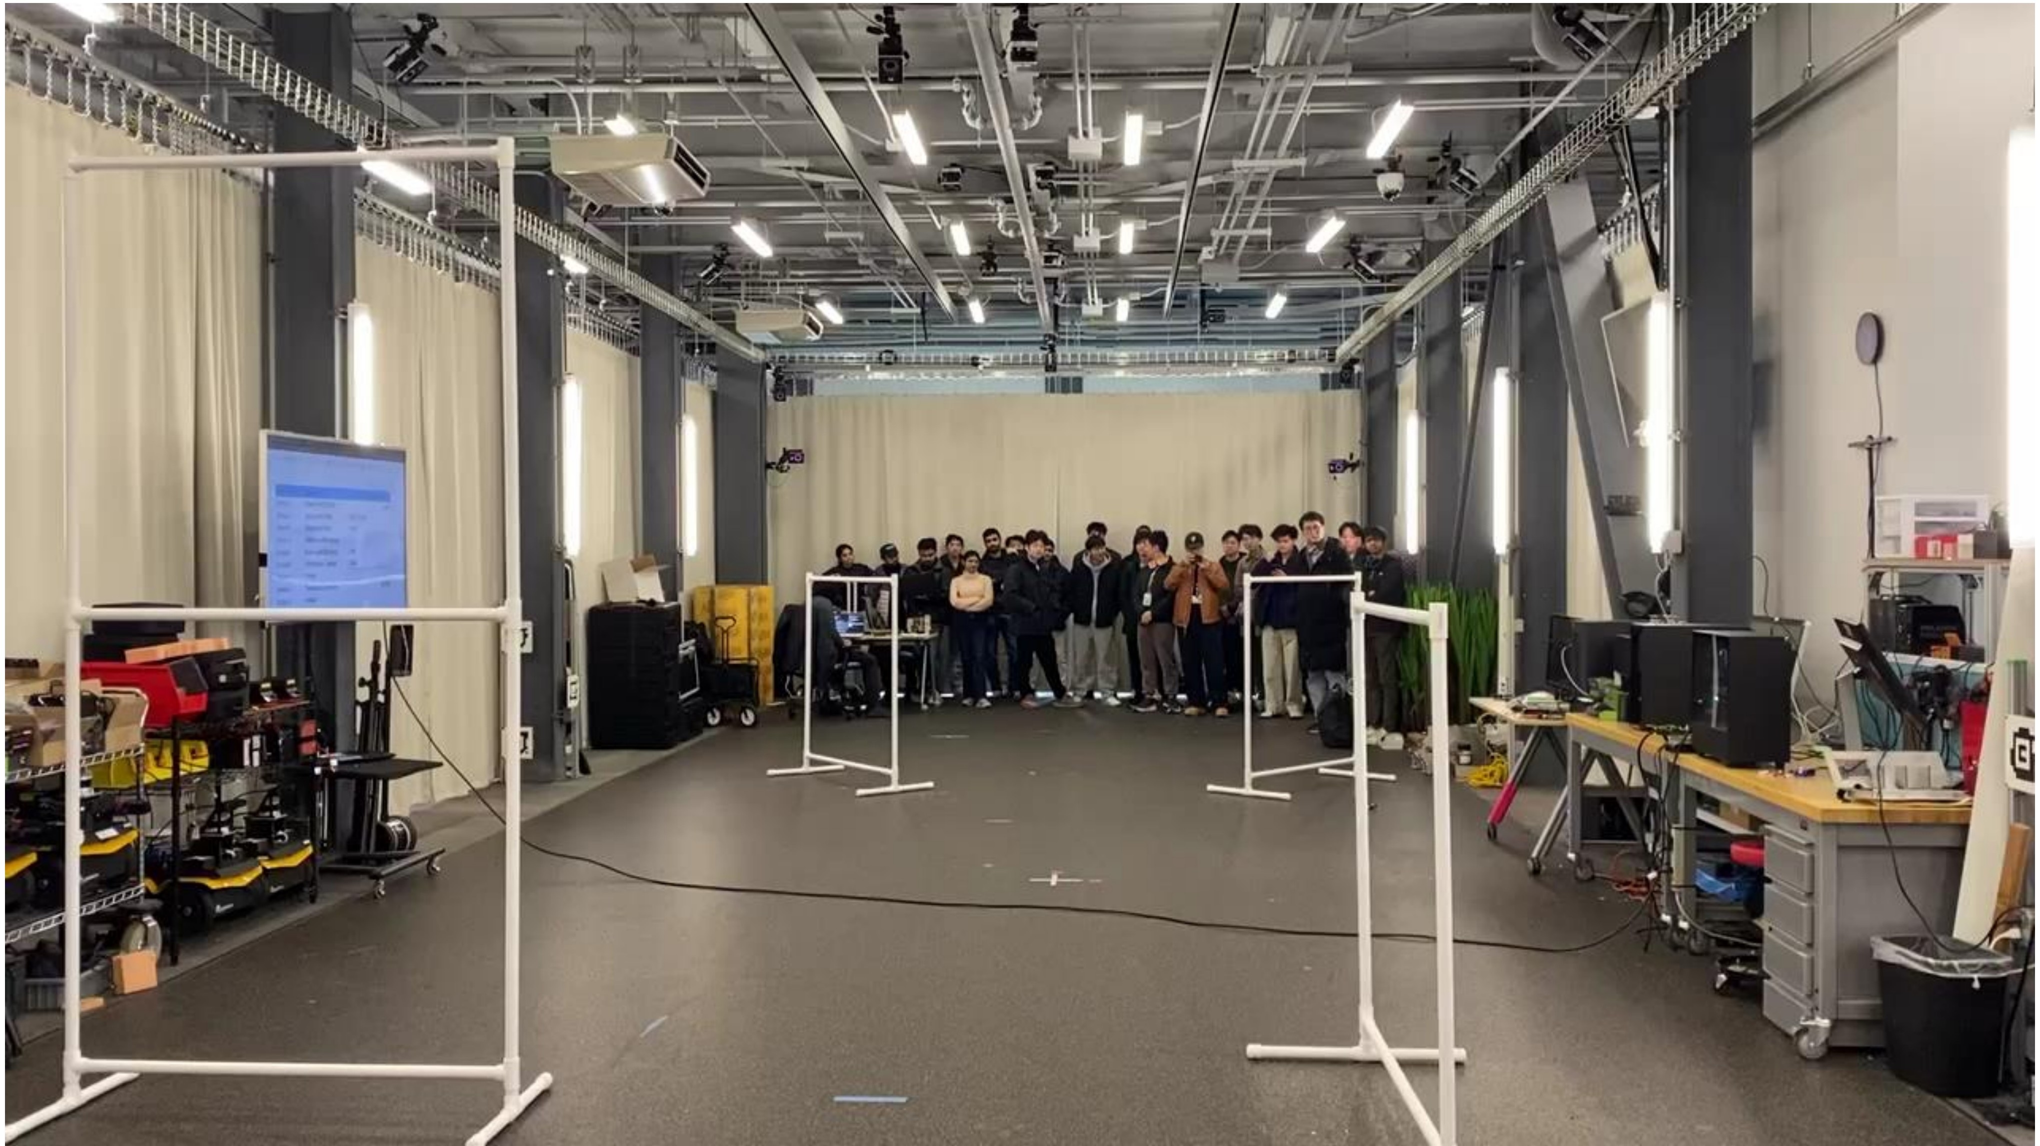

Compute Options for Policy Training

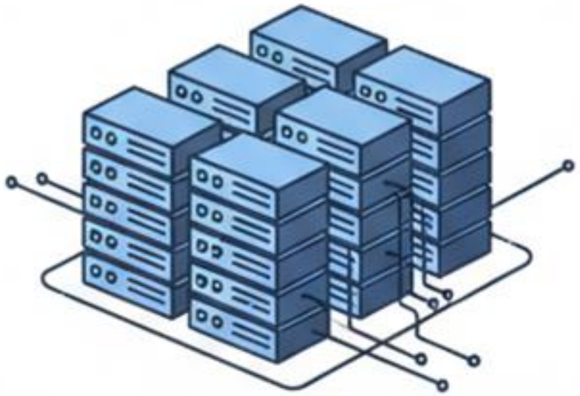

Linux Workstation (Preferred)

Personal or Lab Computer

NVIDIA GPU (RTX 3070 or better)

32 GB RAM

4 Cores

50 GB SSD

8 GB VRAM

HPC

GRASP Cluster, PRECISE, etc

For students that already have HPC access through a lab *and* your advisor approves of usage for a class.

Virtual Machine

Google Cloud, AWS

We will release formal instructions to run Isaac Lab within a VM using Google Cloud and/or AWS.

This is intended for students that do not have a Workstation or HPC access.

The secret to being successful at the race

… and in Robotics in general

The *Quality* vs *Quantity* Groups

"Think Group" "Do Group"

Where you want to be

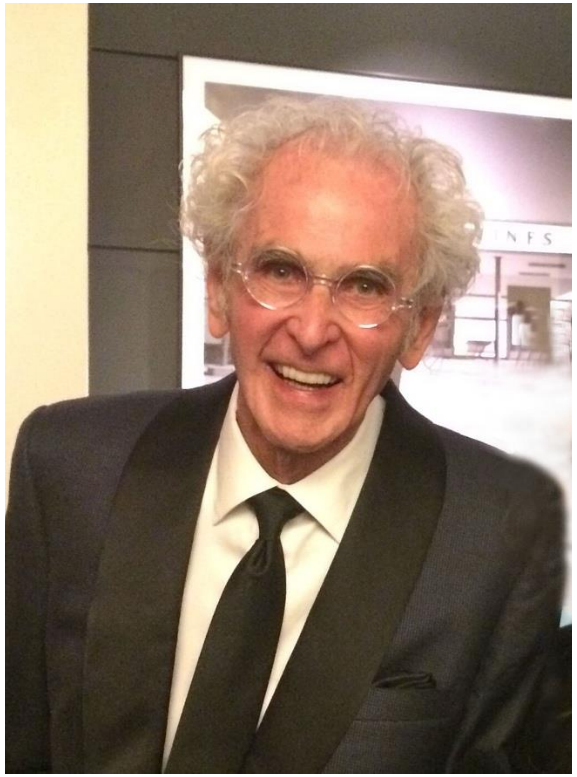

Jerry Uelsmann
Photographer &
Teaching Professor at UoF

Who this class is for

- This is not an intro-level class. If you don’t have previous experience with Robotics/CV/ML this might be challenging.
- We will cover the foundations of imitation and reinforcement learning, adopting a robotics-first perspective.
- Everybody is welcome to audit, but I don’t think you’ll get much out of it.

ESE650 vs ESE651(This class)

- ESE650 is more about “traditional” robot learning:
  - Kalman filters
  - Optimal control
  - The connections of the above to RL
- Our class will be more focused on the latest in robot learning:
  - We will look at RL/Imitation learing theory from the eye of a roboticist.
  - Sim2real, generative imitation learning, large scale pretraining, etc.
  - A bit more hands-on.
- There will be a good overlap of topics in the RL theory part.

Some Polls!

- Poll of backgrounds
  - Real robot experience?
  - Have you worked on drones before?
- Poll of programs
  - PhD
  - MS
  - Undergraduates

If you haven't been admitted to the class yet, please register soon. Make sure to fill the form when you register (otherwise I can’t admit you).

TAs

Jefferson Ng
(Class Material)

Vaibhav Thakkar
(Class Material)

Vineet Pasumarti
(Race)

Some cool things I’ve worked on

Learning High-Speed Flight in the Wild

Antonio Loquercio, Elia Kaufmann, Rene Ranflt,
Matthias Mueller, Vladlen Koltun, Davide Scaramuzza

Science Robotics, 2021

Learning to Fly From Simplified Simulators


km/h

Article | Open access | Published: 30 August 2023

Champion-level drone racing using deep reinforcement learning

Elia Kaufmann ✉, Leonard Bauersfeld, Antonio Loquercio, Matthias Müller, Vladlen Koltun & Davide Scaramuzza

Nature **620**, 982–987 (2023) | Cite this article


Universität Zürich UZH
Robotics and Perception Group
robotics
1
2
ROBOTICS & PERCEPTION GROUP
Innovation Zurich
Canton of Zurich
Department for Economic Affairs
Office for Economy and Labour
UZH

The Human Champions


Universität Zürich
Robotics and Perception Group

Co-Adaptation

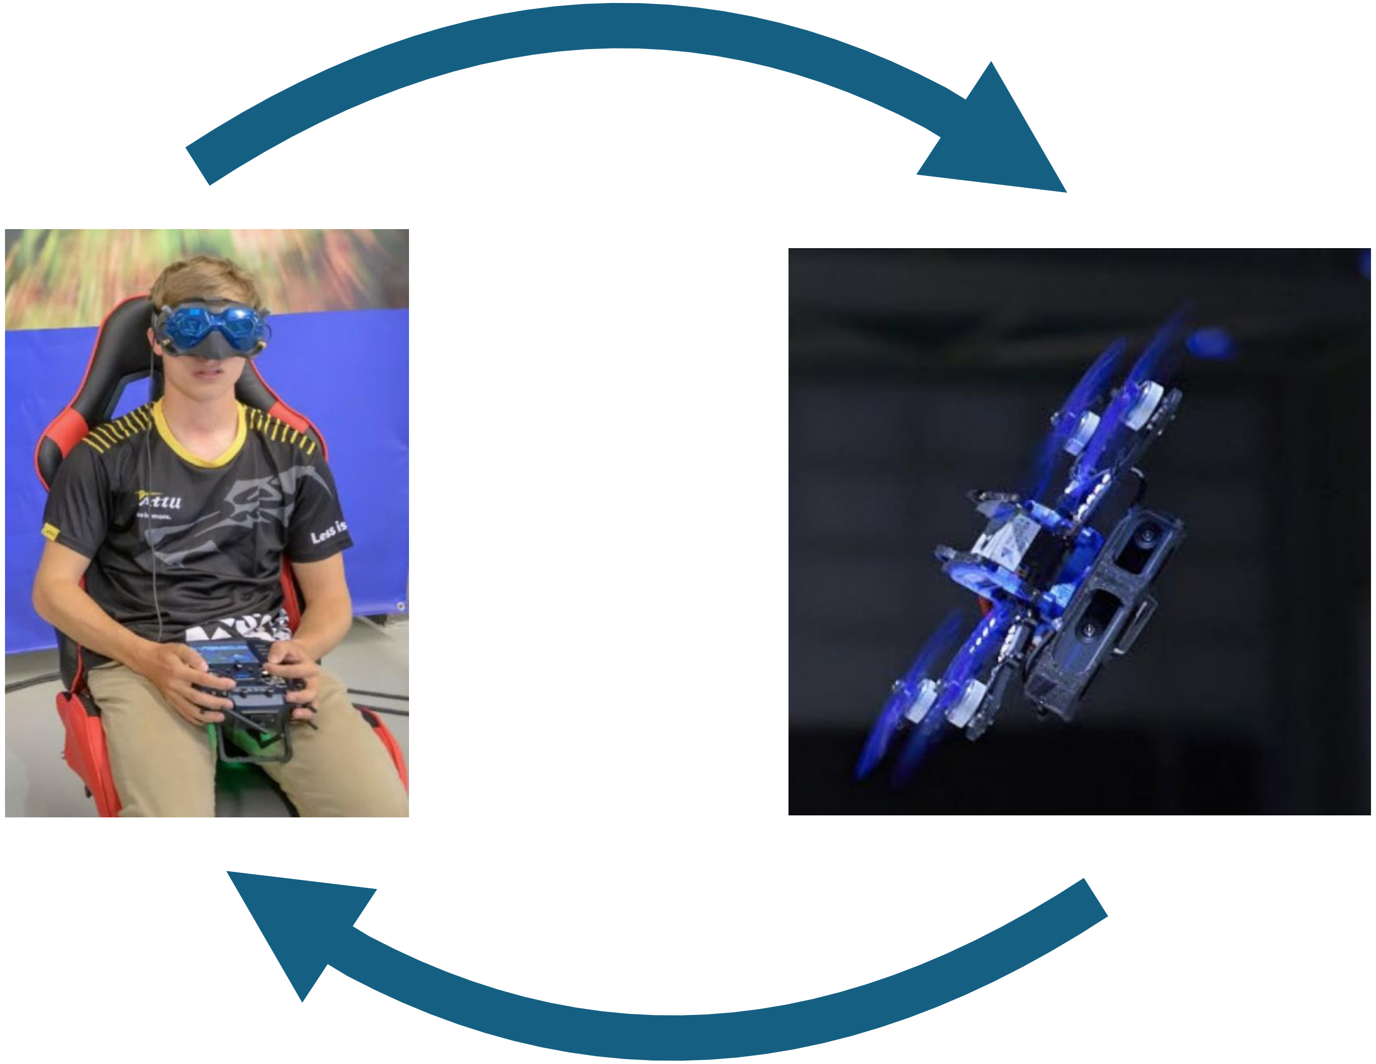

Agile Flight Emerges from Multi-Agent Competitive Racing

Vineet Pasumarti

Lorenzo Bianchi

Training Agents with Sparse Competition Rewards

Reward 1:
Pass every gate before your opponent

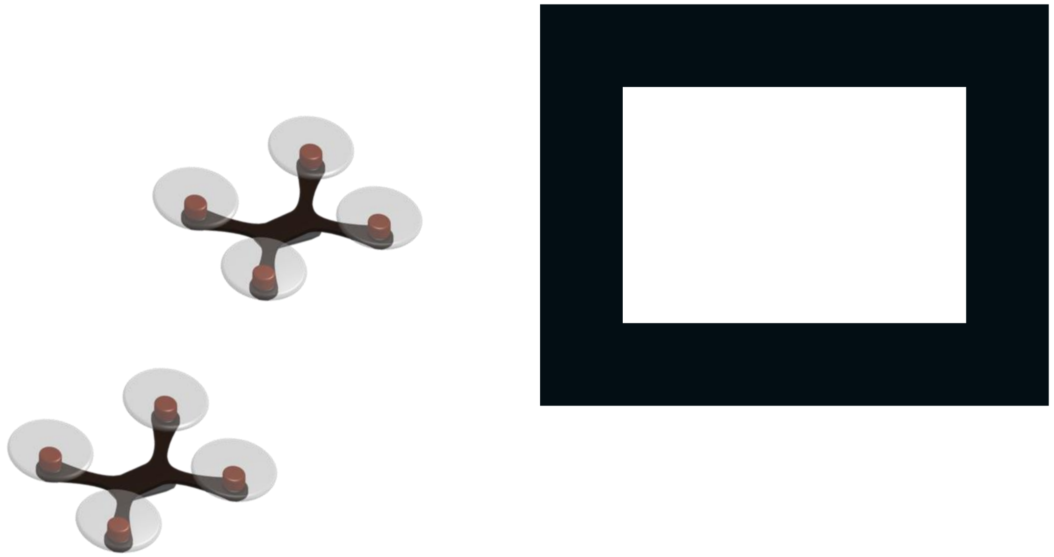

Reward 2:
Bonus if you finish a lap first

Competitive Behaviors Emerge

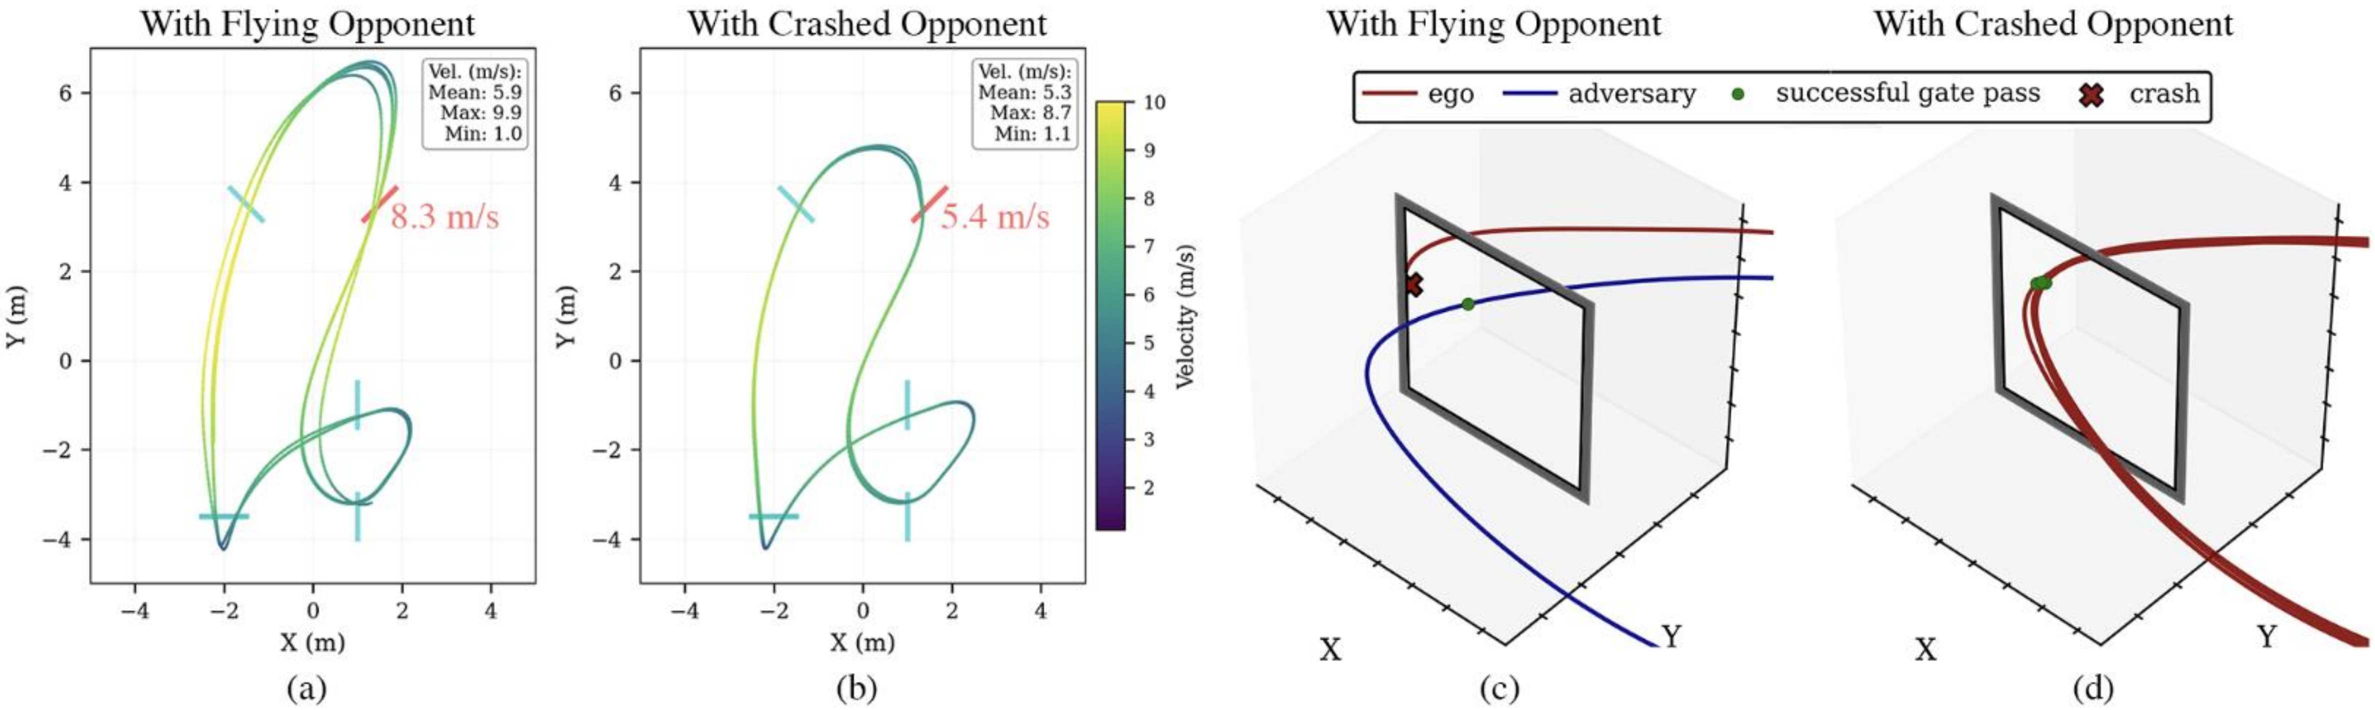

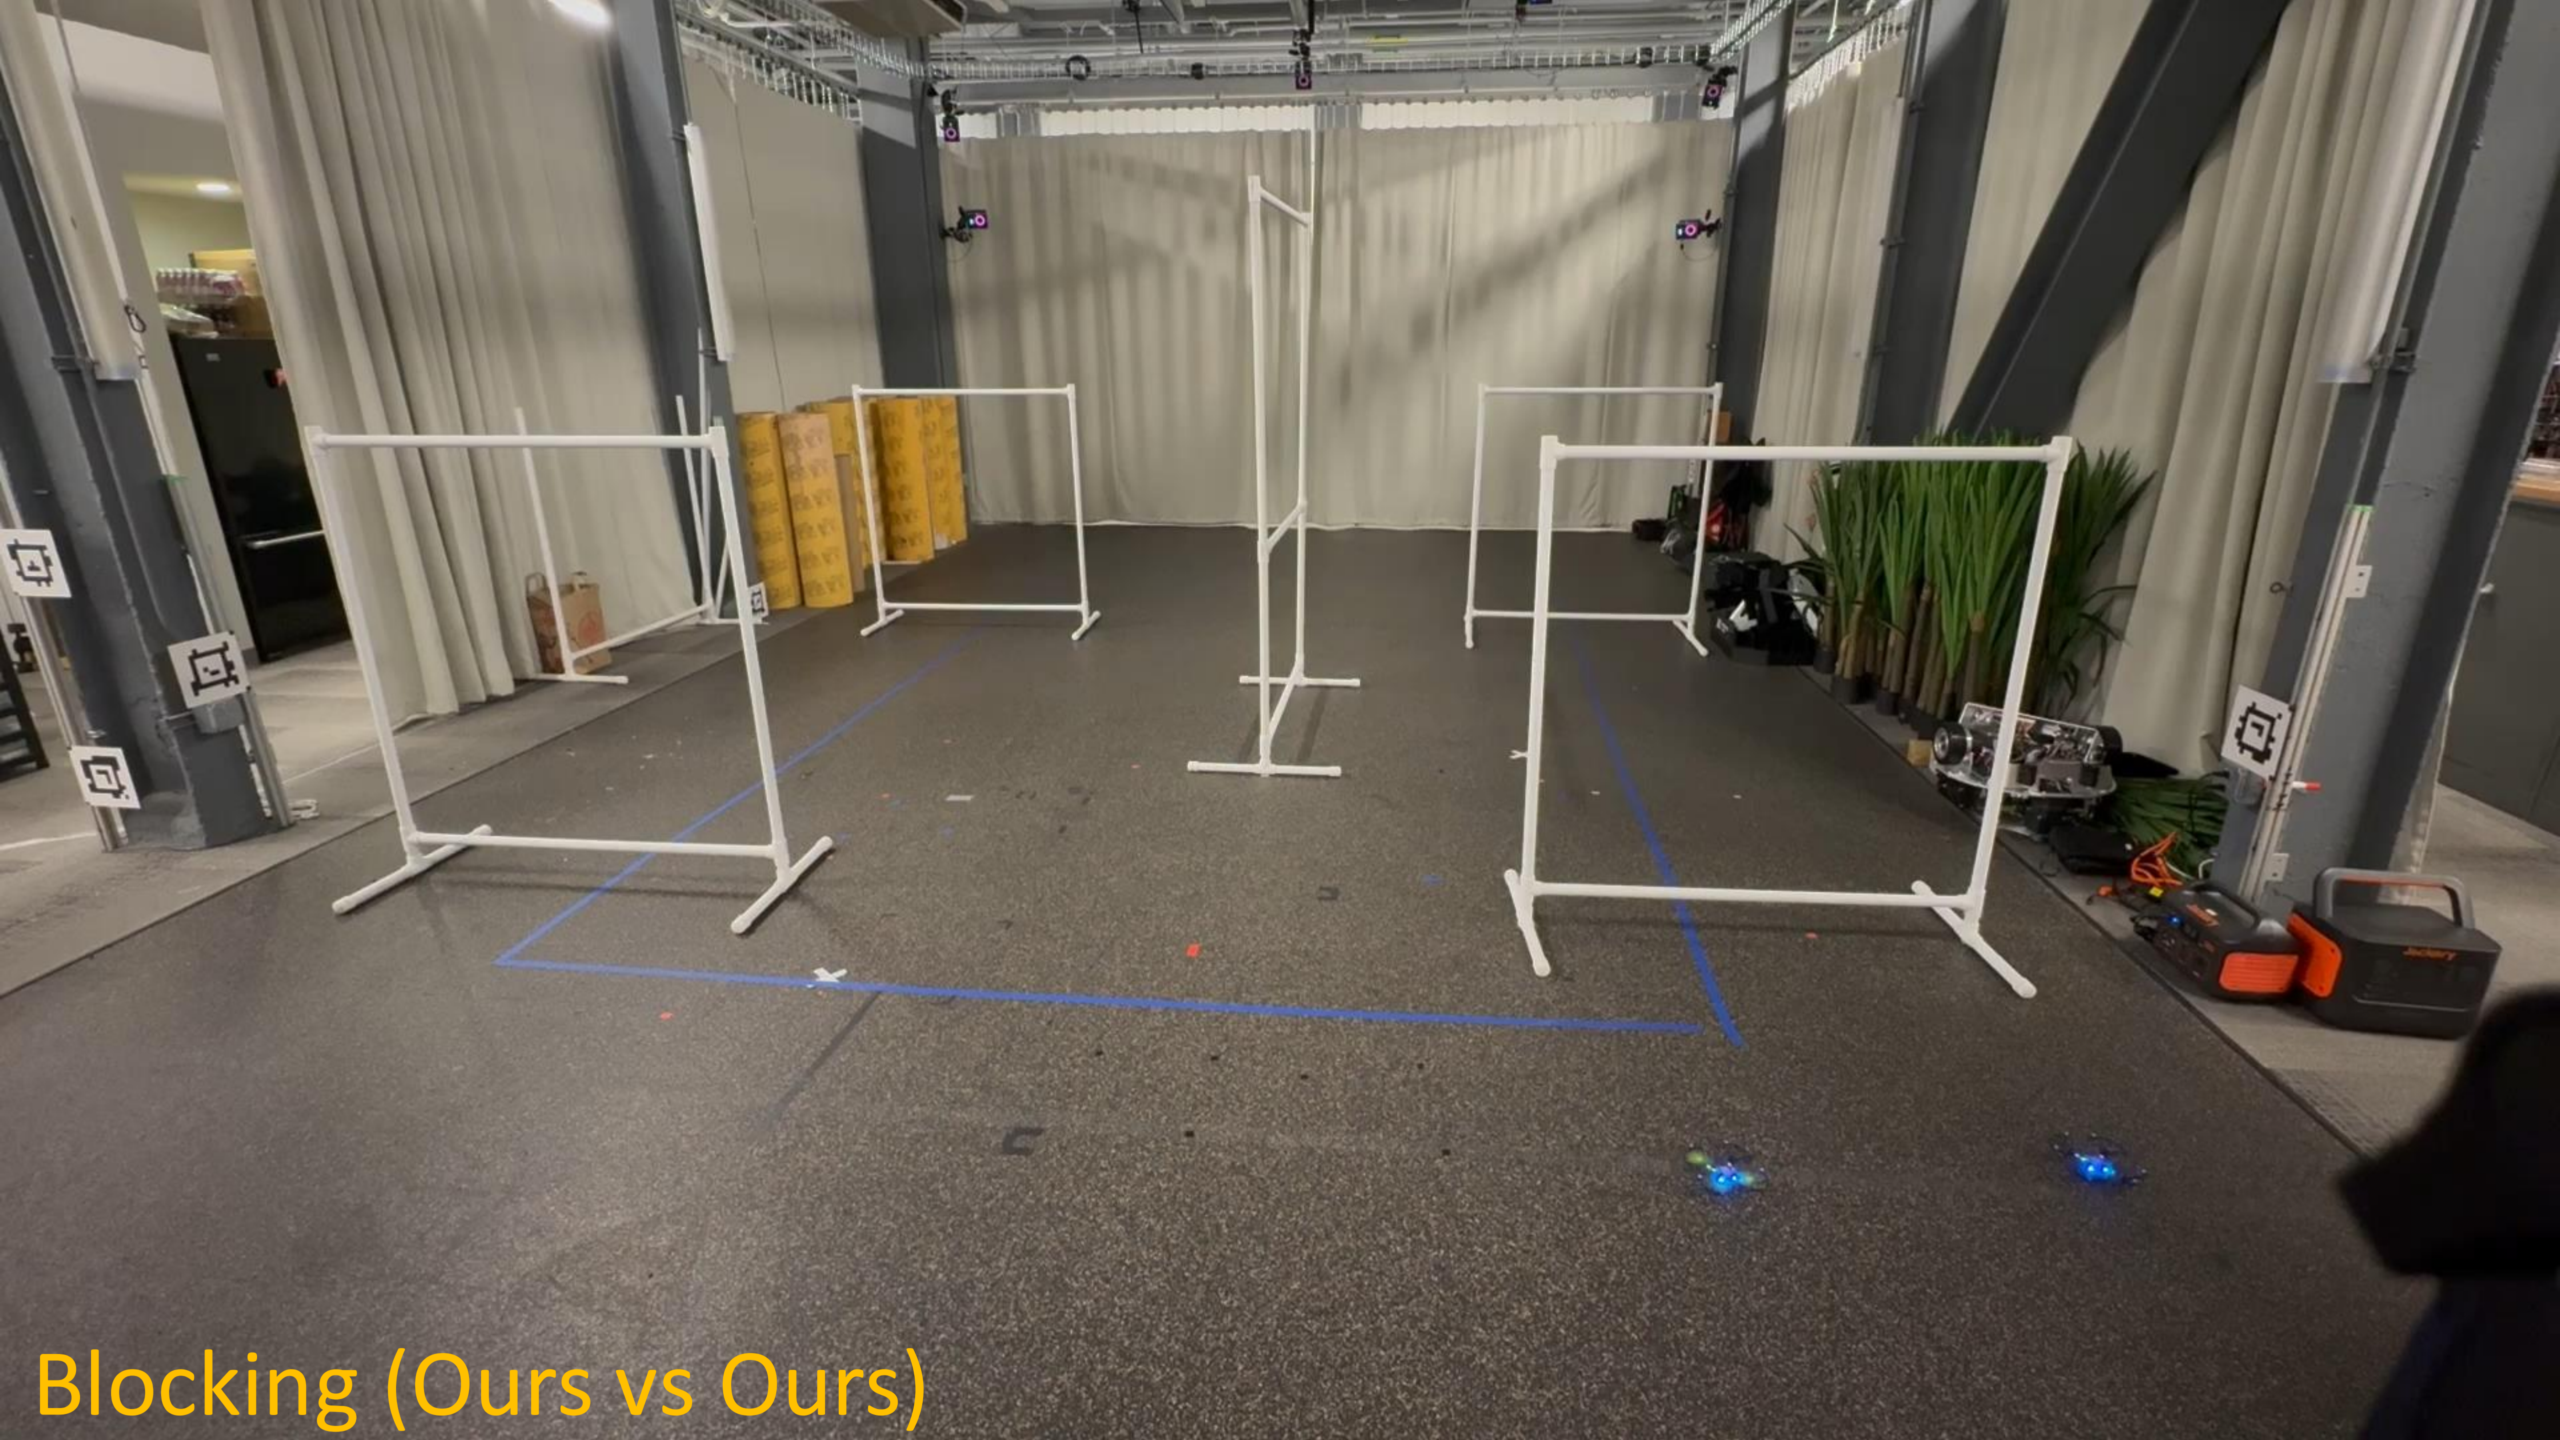

Blocking (Ours vs Ours)


Blocking (Ours vs Ours)

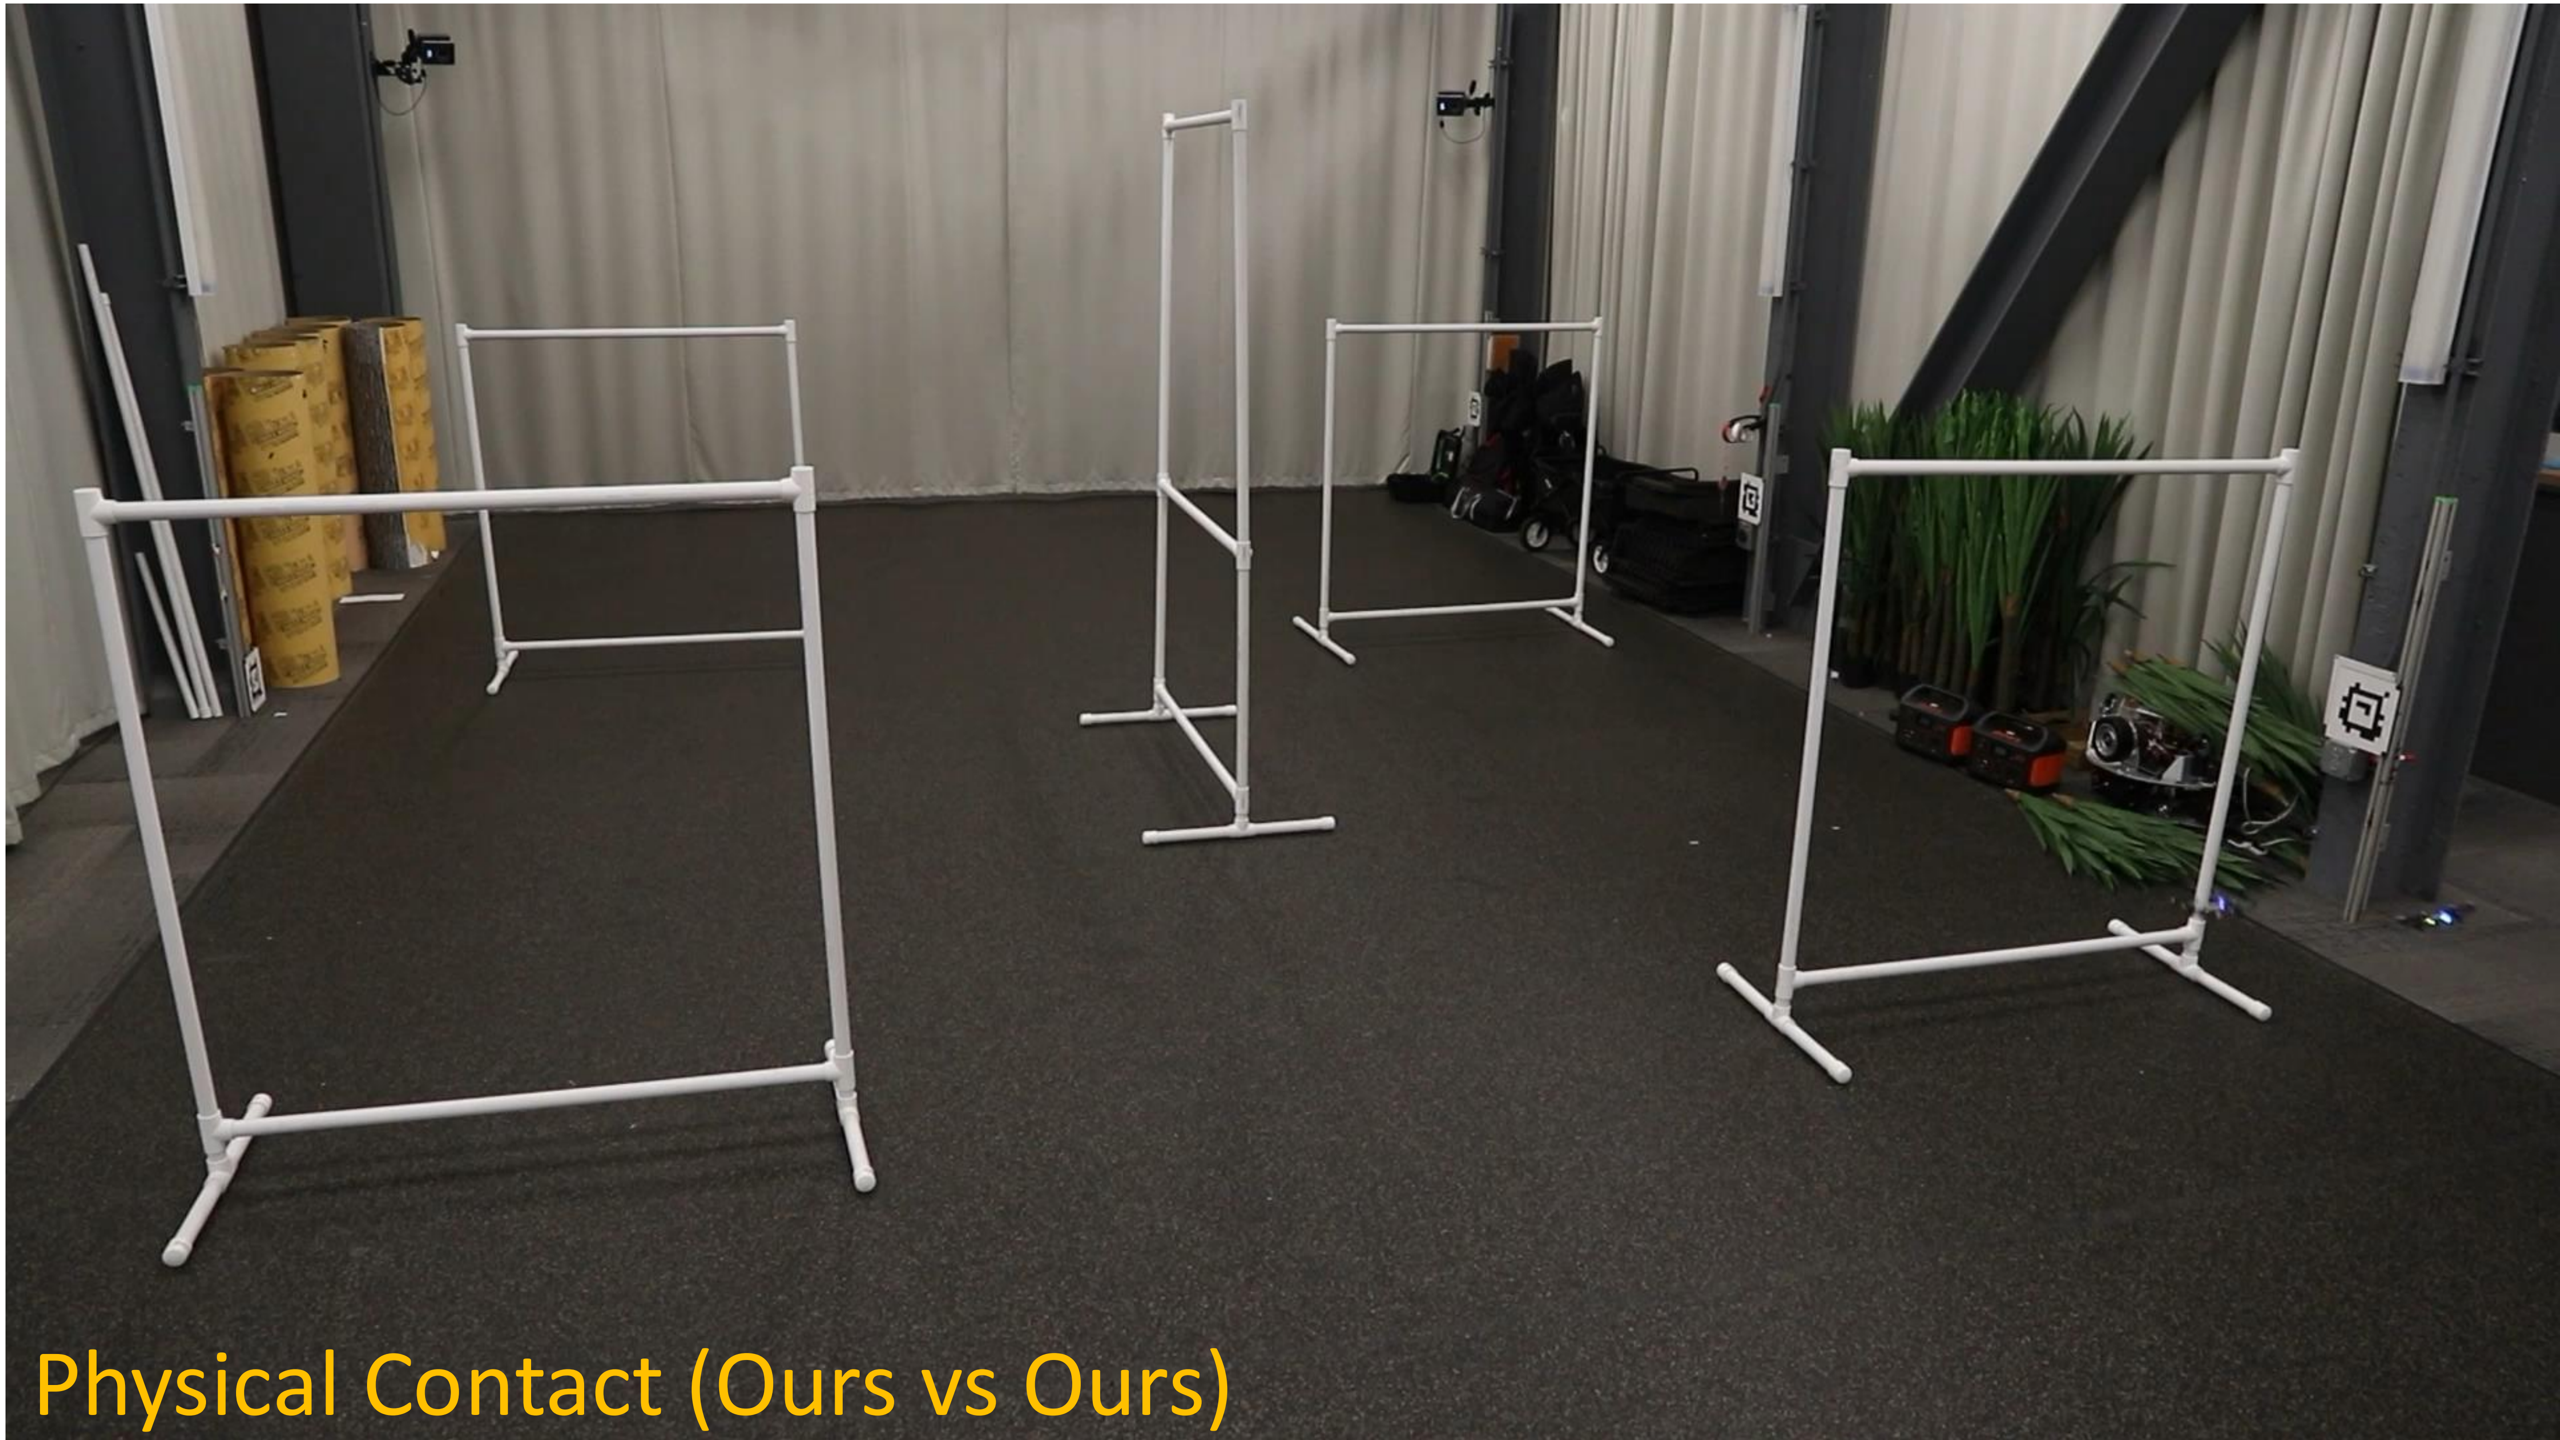

Physical Contact (Ours vs Ours)

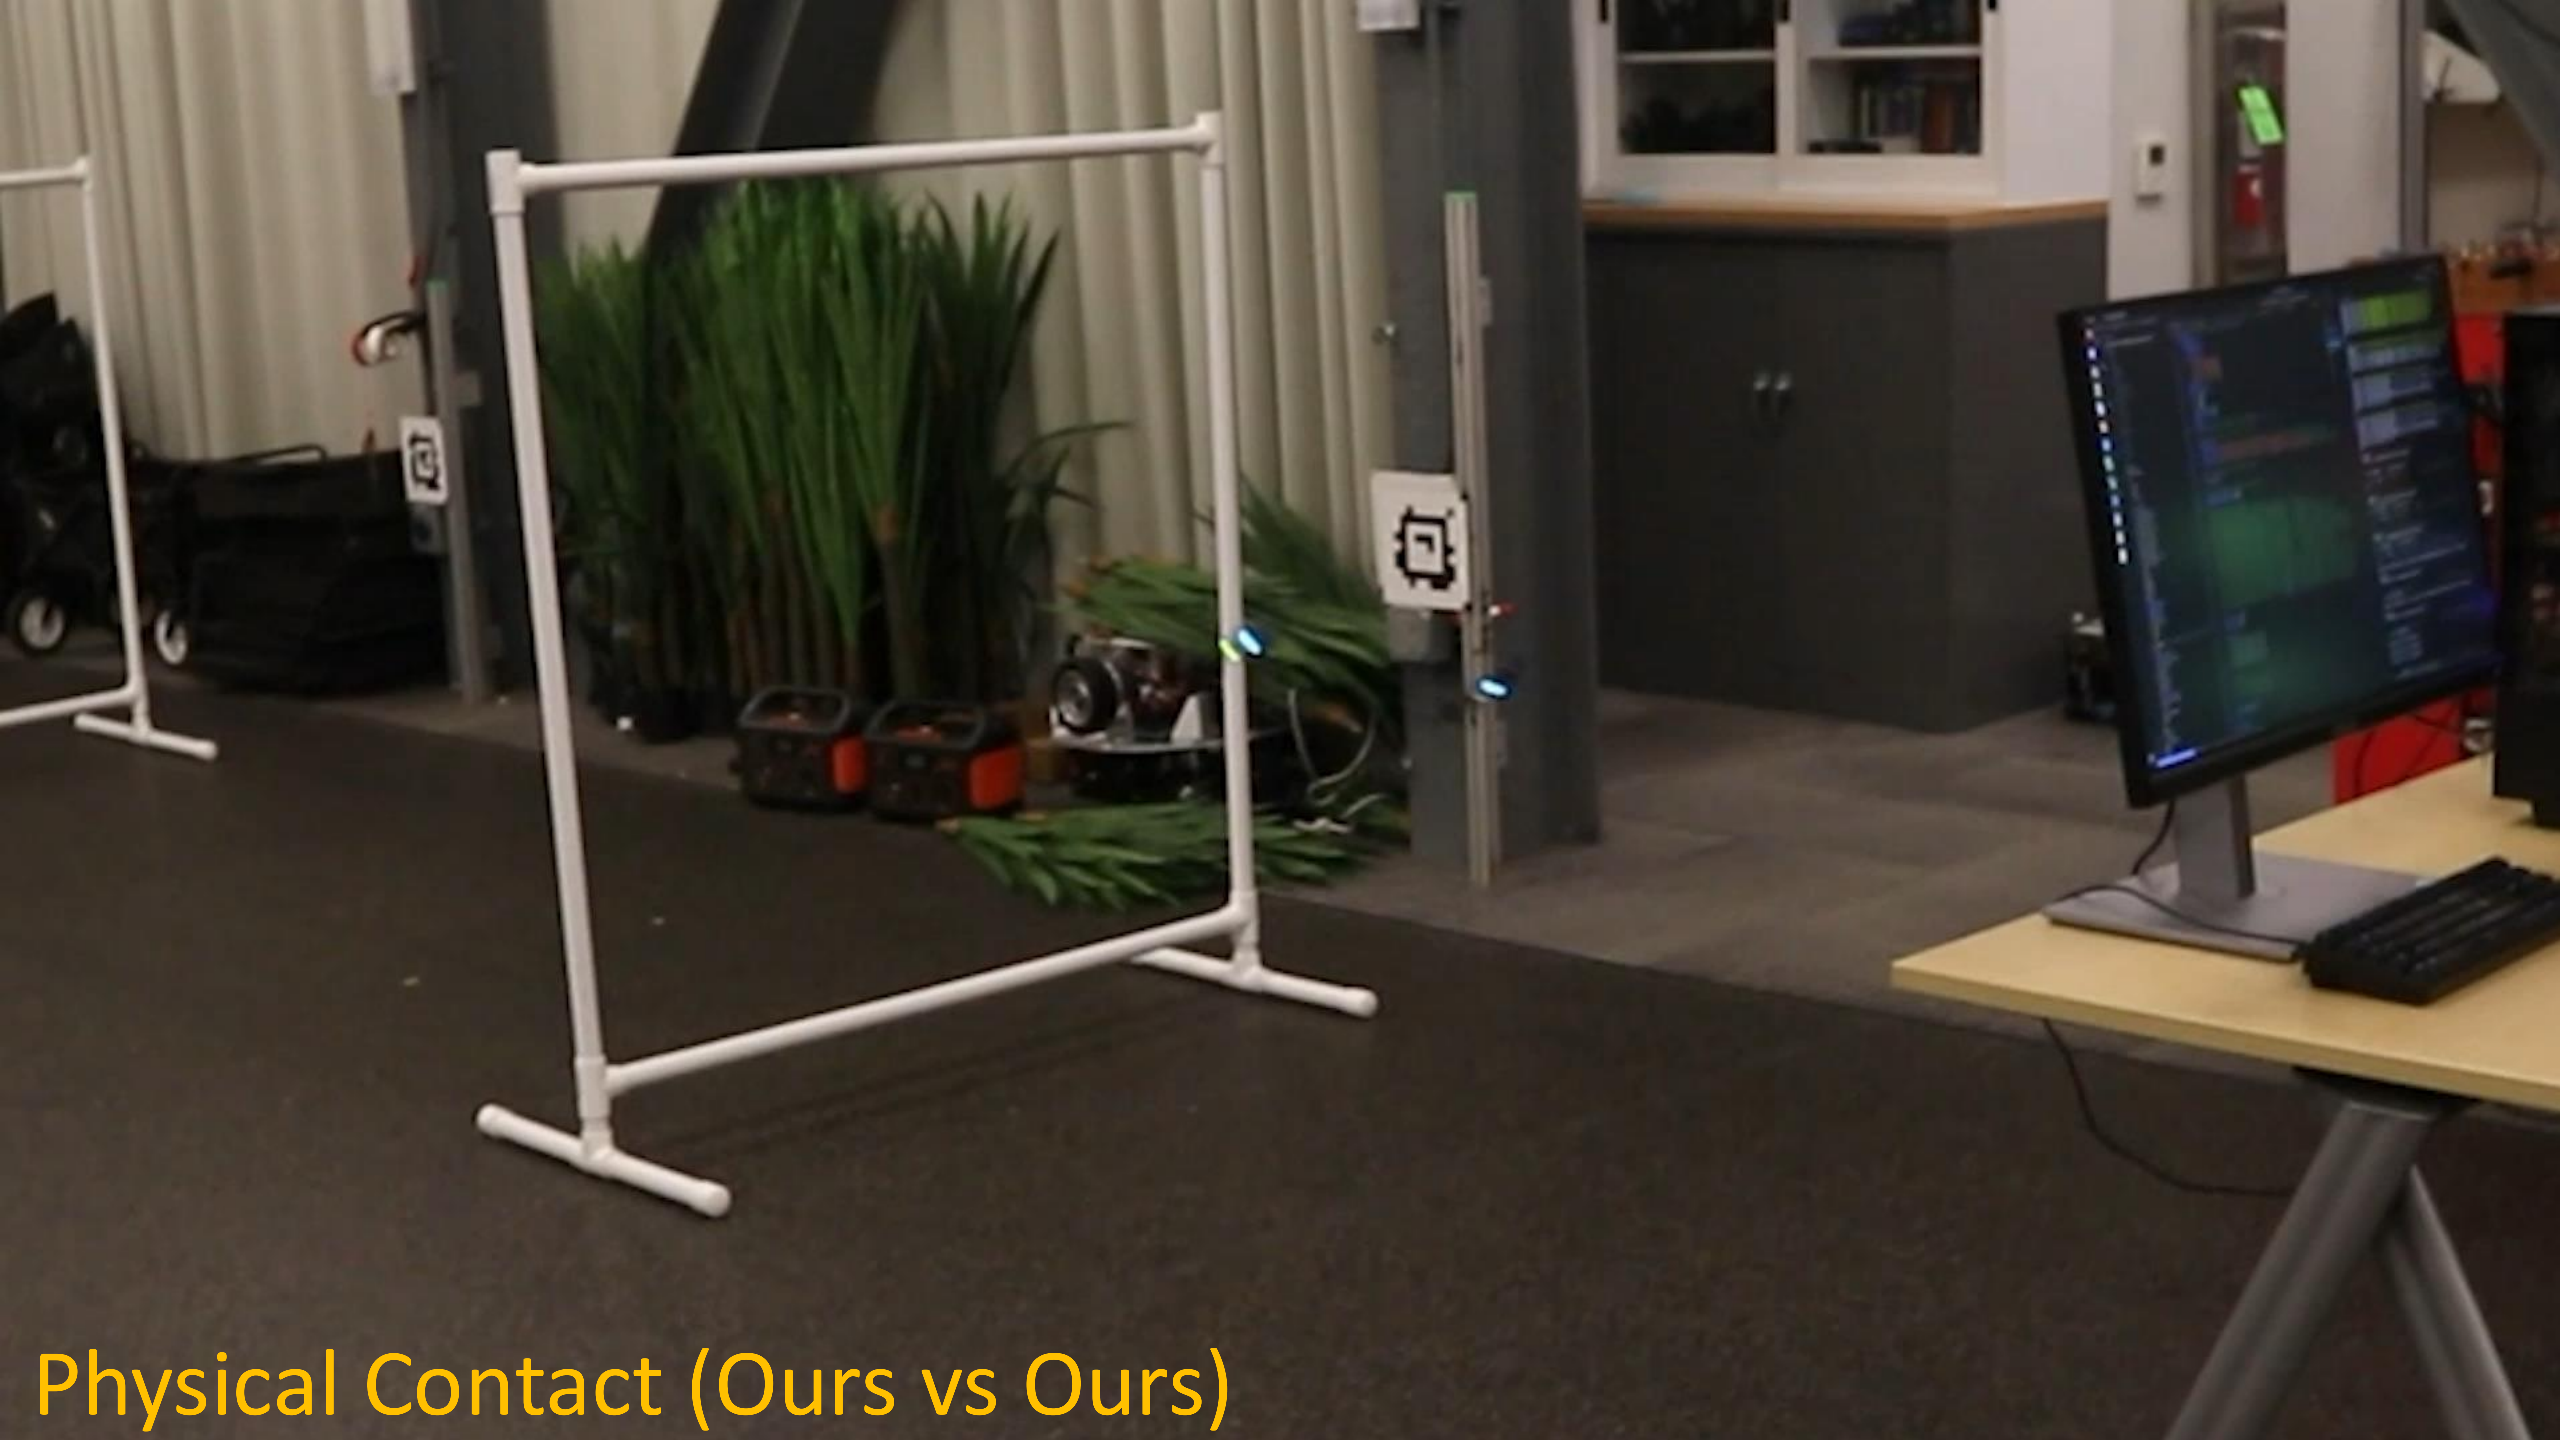
Physical Contact (Ours vs Ours)

Learning Visual Locomotion with Cross-Modal Supervision

Antonio Loquercio

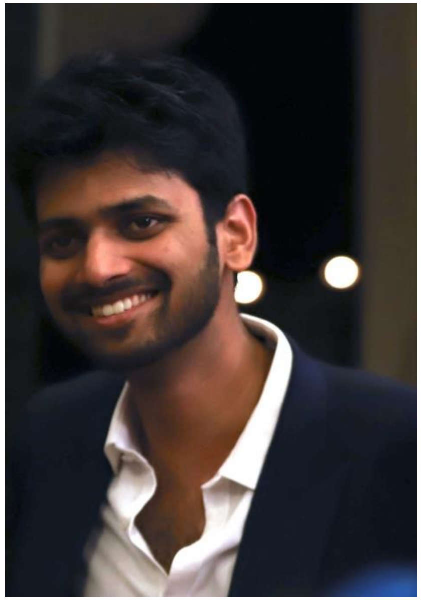

Ashish Kumar

Jitendra Malik

Loquercio et. al,
ICRA, 2023

Construction Zone

Loquercio et. al,
ICRA, 2023

Loquercio et. al,
ICRA, 2023

Visual Plasticity

Before Adaptation

After 1min of data

Learning from Videos

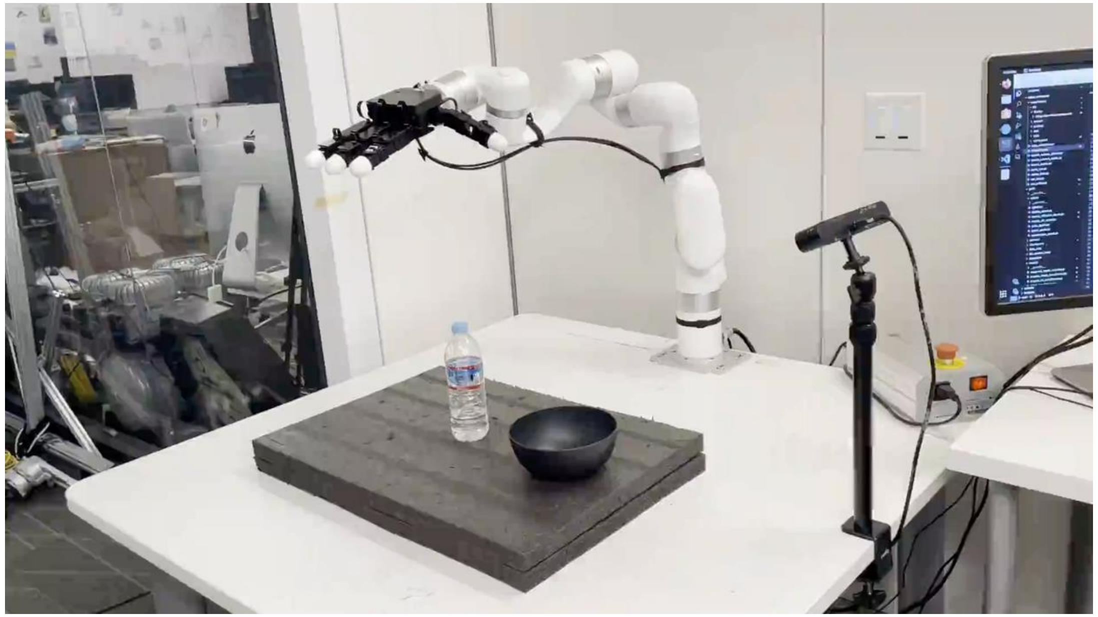

Human videos for scaling up robot data

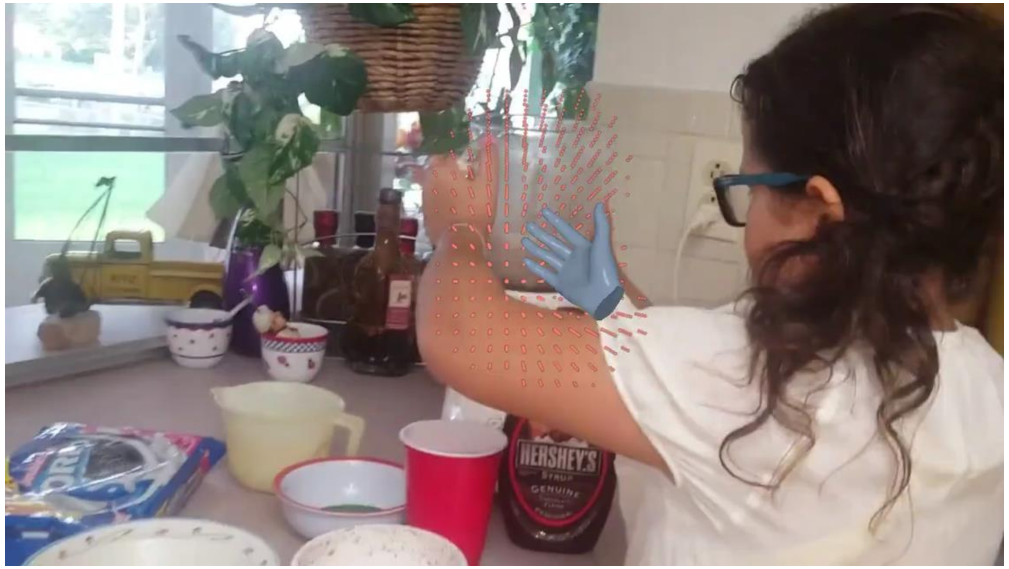

1. Intuitive physics
2. Contact poses
3. Pre/Post-contact trajectories
4. Human preferences
5. ....

Sample-efficient BC Finetuning