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

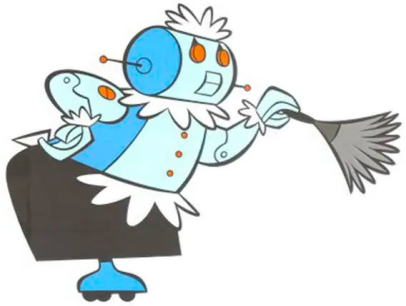

How the future is going to be.

Moravec's Paradox

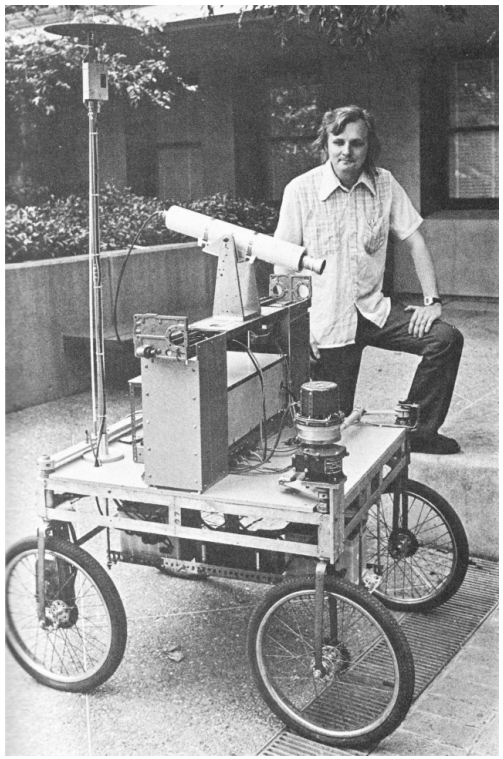

It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility.

Mind Children, H. Moravec, 1988

Hans Moravec's PhD Thesis

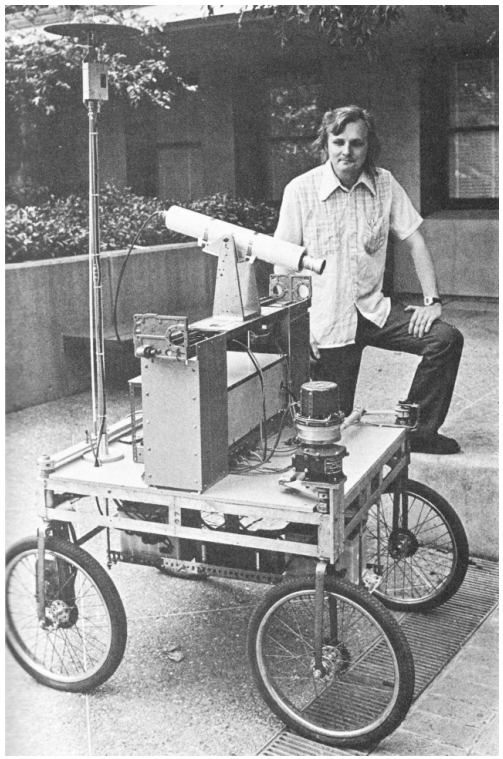

During this period a number of incoming students were assigned to the "cart project'. Each correctly perceived the situation within a year, and went on to something else. The cart's reputation as a serious piece of research apparatus, never too high, sank to new depths.

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

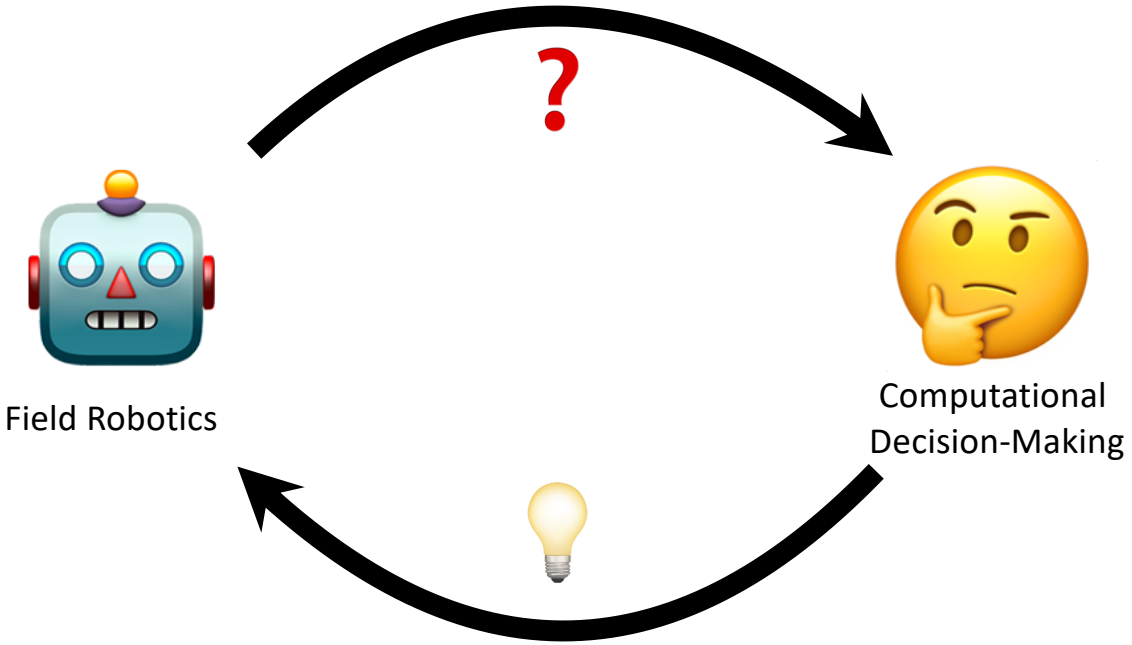

Why care about the **Real World**?

"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 make 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


CGTN

What's different about this class?

- We will see foundational concepts together with fancy systems
  - Study theoretical concepts, then come back to the latest and greatest
- 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

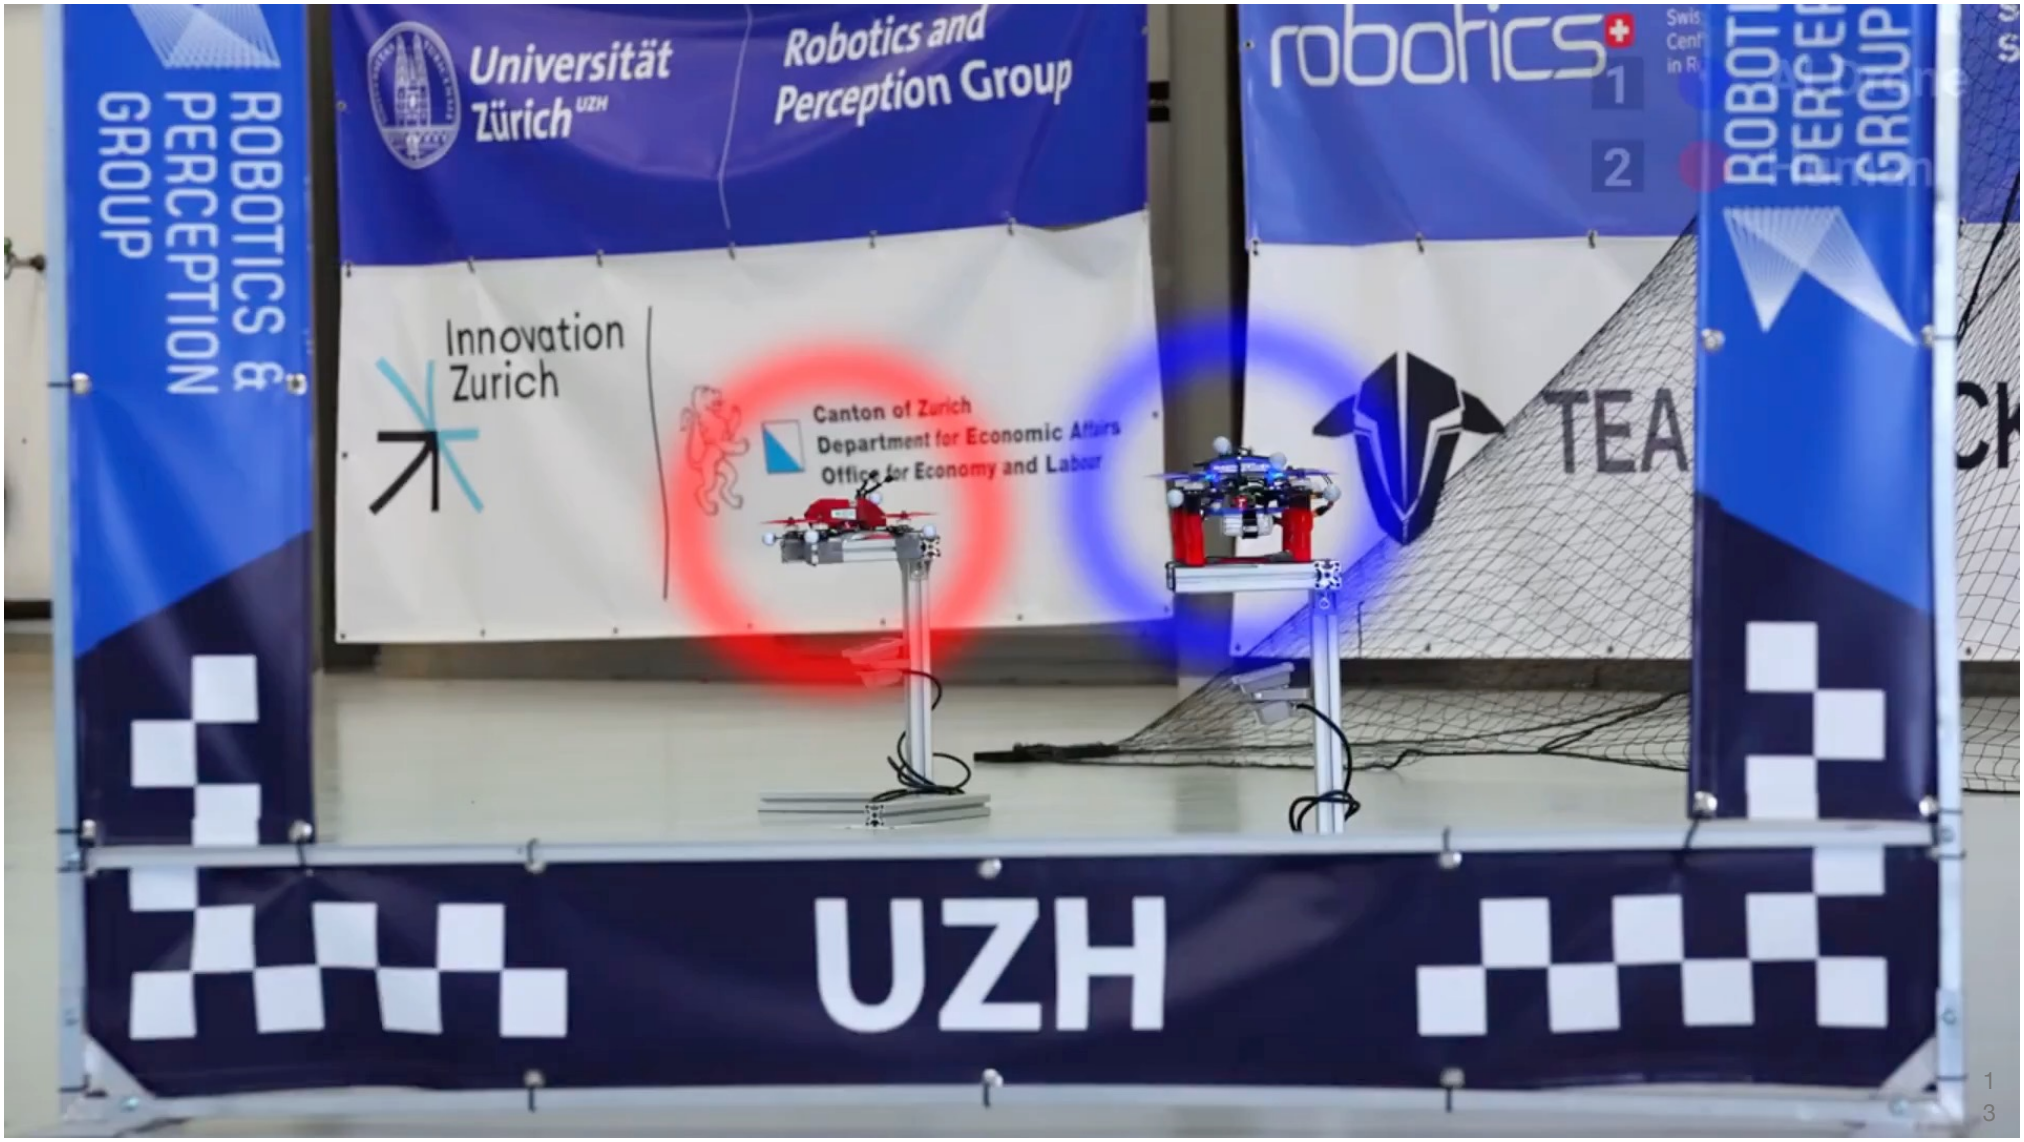
ROBOTICS & PERCEPTION GROUP
Universität Zürich UZH
Robotics and Perception Group
Innovation Zurich
Canton of Zurich
Department for Economic Affairs
Office for Economy and Labour
robotics
1
2
UZH

Easy to Miss the Complexity of the System

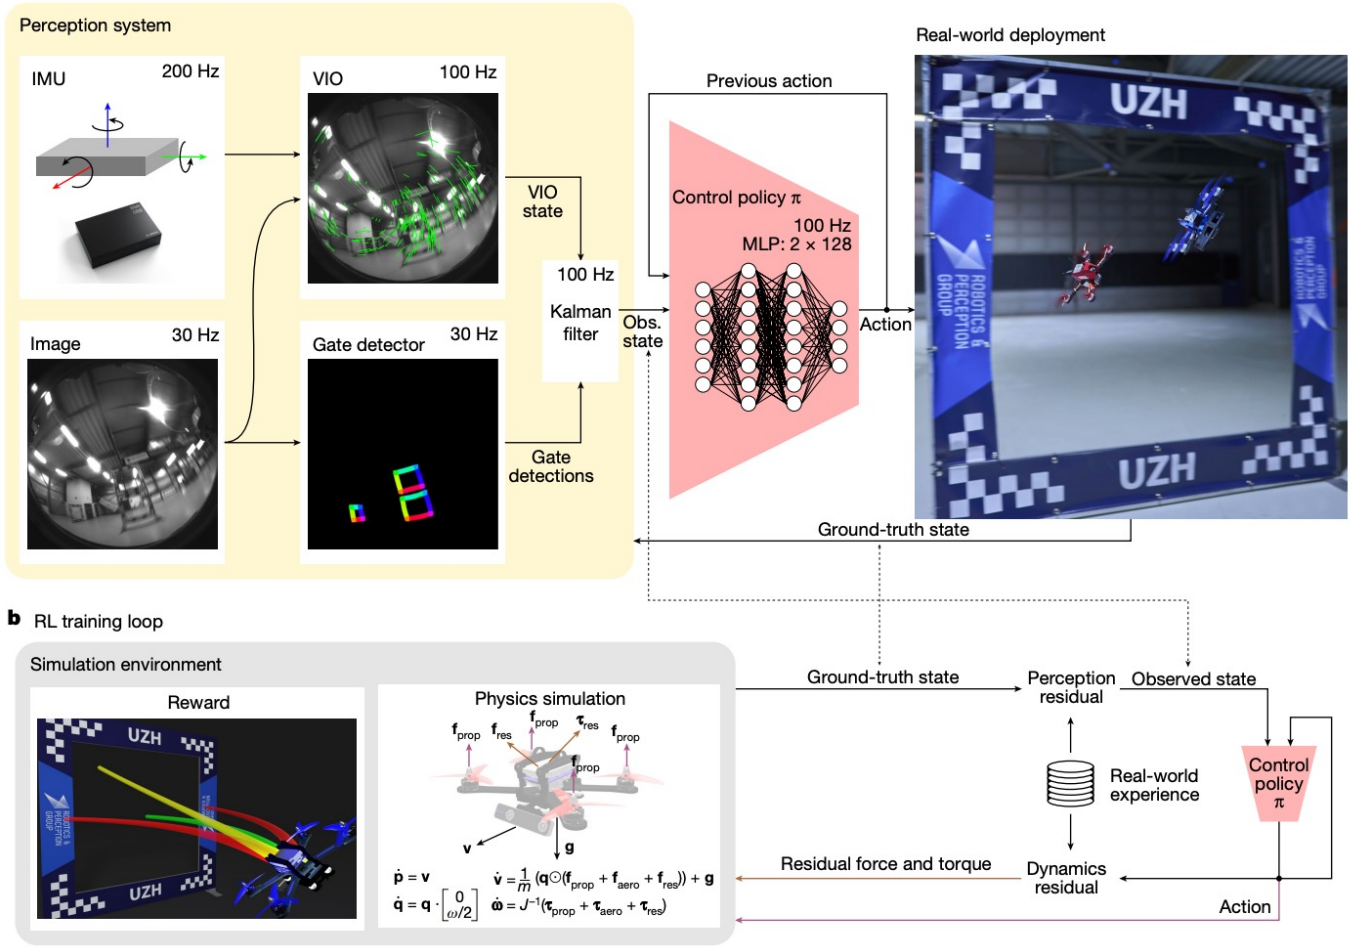

Easy to Miss the Complexity of the System

- The “Infrastructure” (built over three years of work)

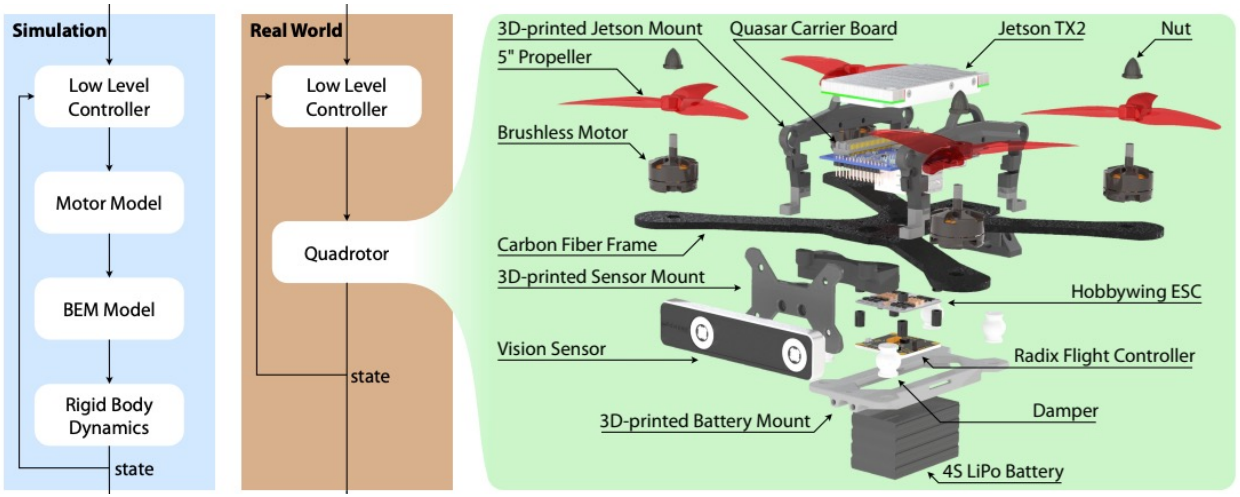

- Latency Control
- Gaussian Mixture Models for Noise Identification
- Custom code to run vision models in embedded hardware

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 Organization

Grading:

1. Class Participation (5%)
   - **Attend class.** Ask questions, be involved!
2. Midterm (30%)
3. Final Exam (50%)
   - Probably oral (depending on class size)
4. Drone Race (15%)
   - More details on this later
   - Can be done solo or in groups of 2
   - First 3 can skip the final exam

Hands-on Tutorials (3X)

- We will do exercises/coding together in class
- There will be some quizzes during the tutorial
  - Best place to get your participation points!
- 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. Unseen tracks.
- All groups must participate
- The top 3 teams move to the next phase.
  - Team Members can skip the exam!
- Three grading tears
  1. Top 10 (one bonus point)
  2. Not completing the track (grade decided depending on what was implemented)
  3. Everybody else (full score)
- Need to submit a short (1-2pp) report on what was implemented

The Drone Racing Competition: Phase II

The Drone Racing Competition: Phase II

- Drones are flying at the same time. Unseen tracks.

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

- Two grading tears
  1. First position (guaranteed top grade)
  2. Everybody else (grade decided depending on what was implemented)

- Need to do a 5 min presentation on the method after the race

The secret to being successful at the race

… and in Robotics in general

The *Quality* vs *Quantity* Groups

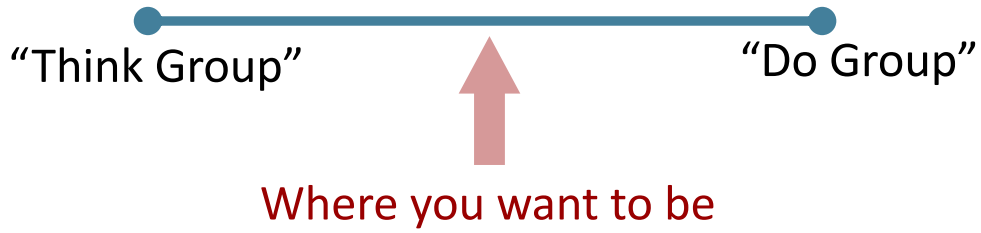

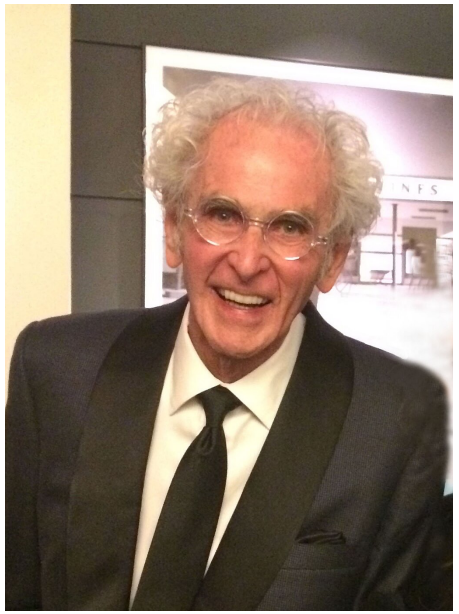

Jerry Uelsmann
Photographer &
Teaching Professor at UoF

Who this class is for

- This is not an intro level class. If you don't have previous experience with Robotics/CV/ML this will be challenging.

- We will cover the foundations of imitation and reinforcement learning, but we will have a robotics-first perspective.

- Everybody is welcome to audit

Overlap with RWRL Seminar

- Topics will be similar, but we will go into much more detail about **hardware** and **algorithms**

- Less focus on future directions

- Many more opportunities to gain hands-on experience

Some Polls!

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

If you're not yet admitted to the class, please indicate in waitlisty if you have access to compute resources (check the minimum requirement for training policies in Isaac Lab).

TAs

Chunwei Xing
(Class and Exercises)

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


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

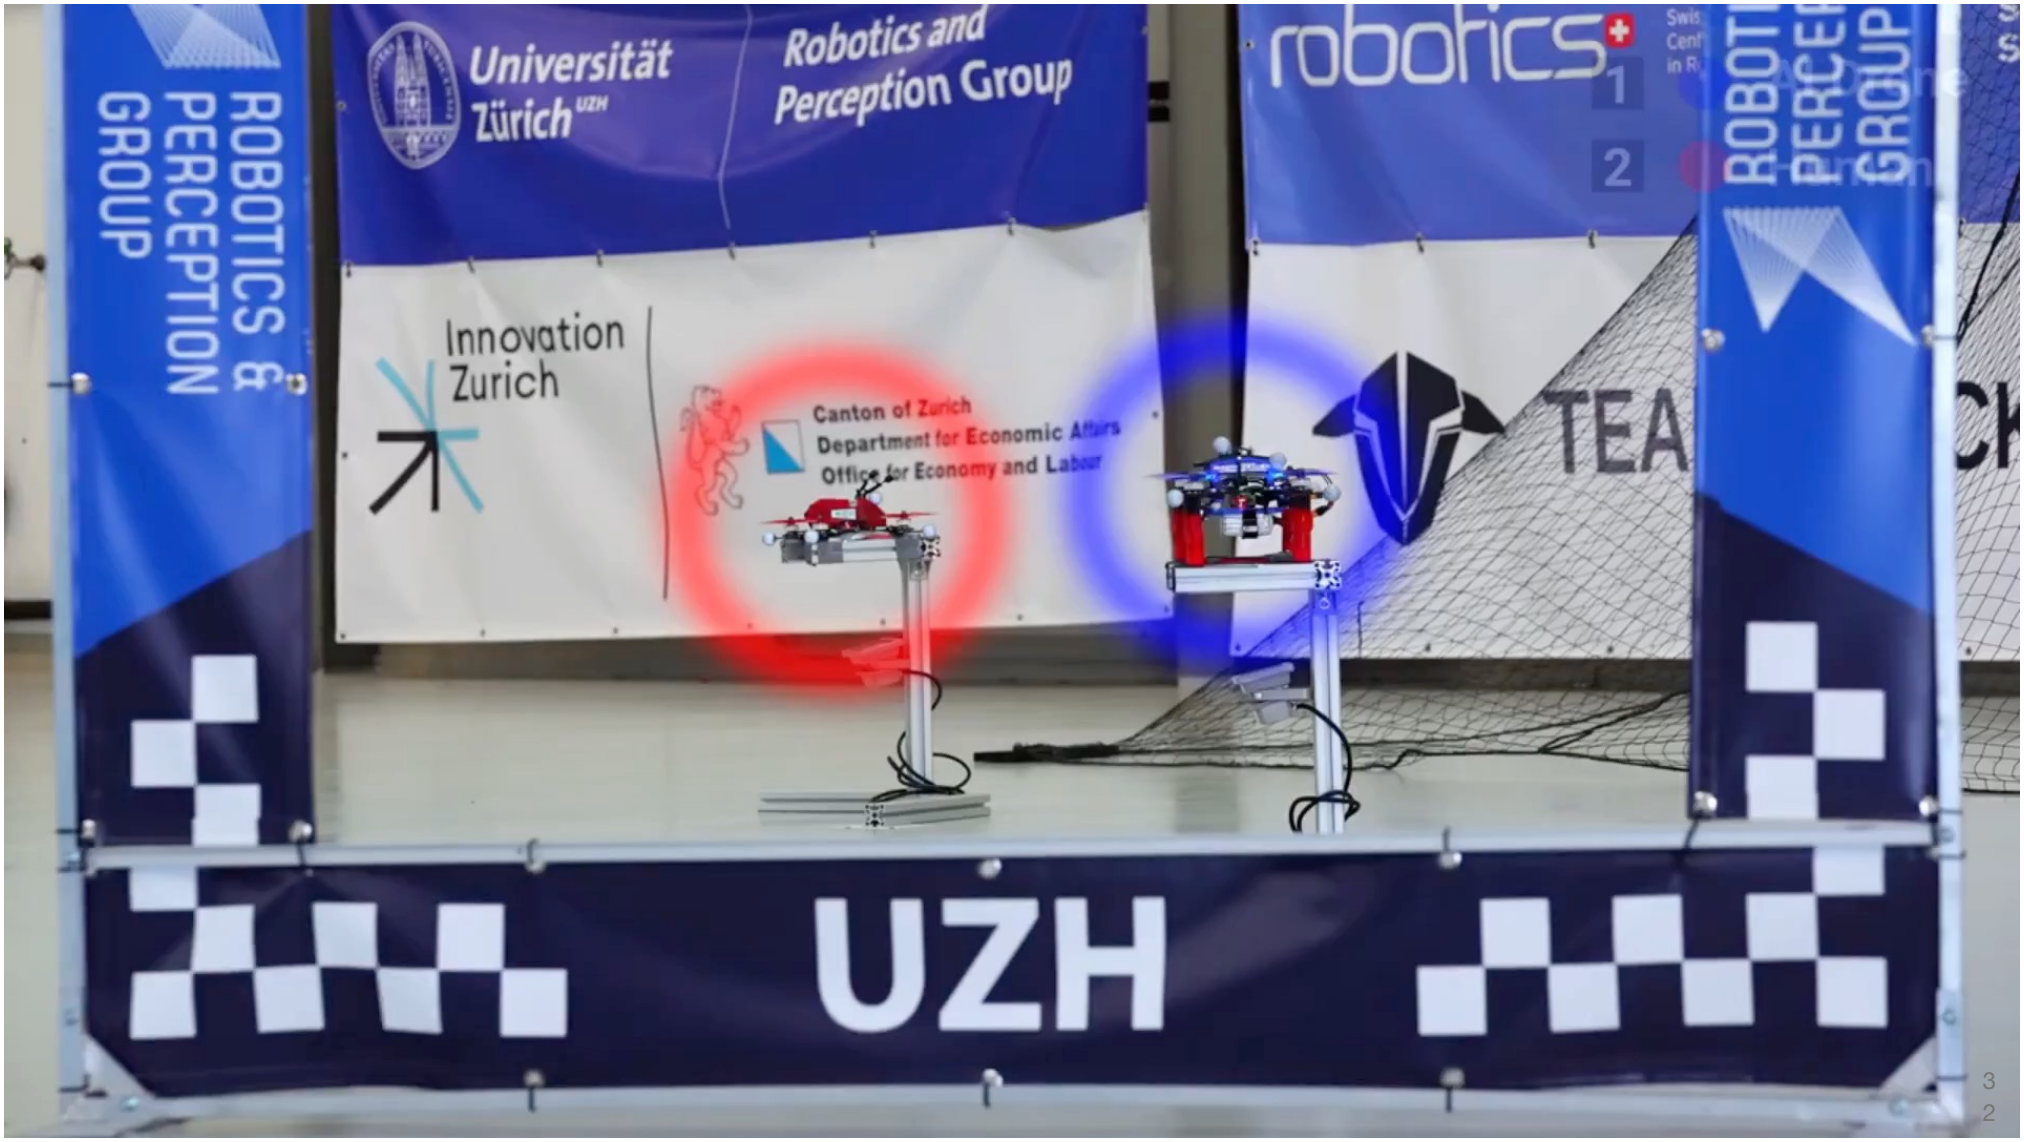
ROBOTICS & PERCEPTION GROUP
Universität Zürich UZH
Robotics and Perception Group
Innovation Zurich
Canton of Zurich
Department for Economic Affairs
Office for Economy and Labour
robotics
1
2
UZH

The Human Champions

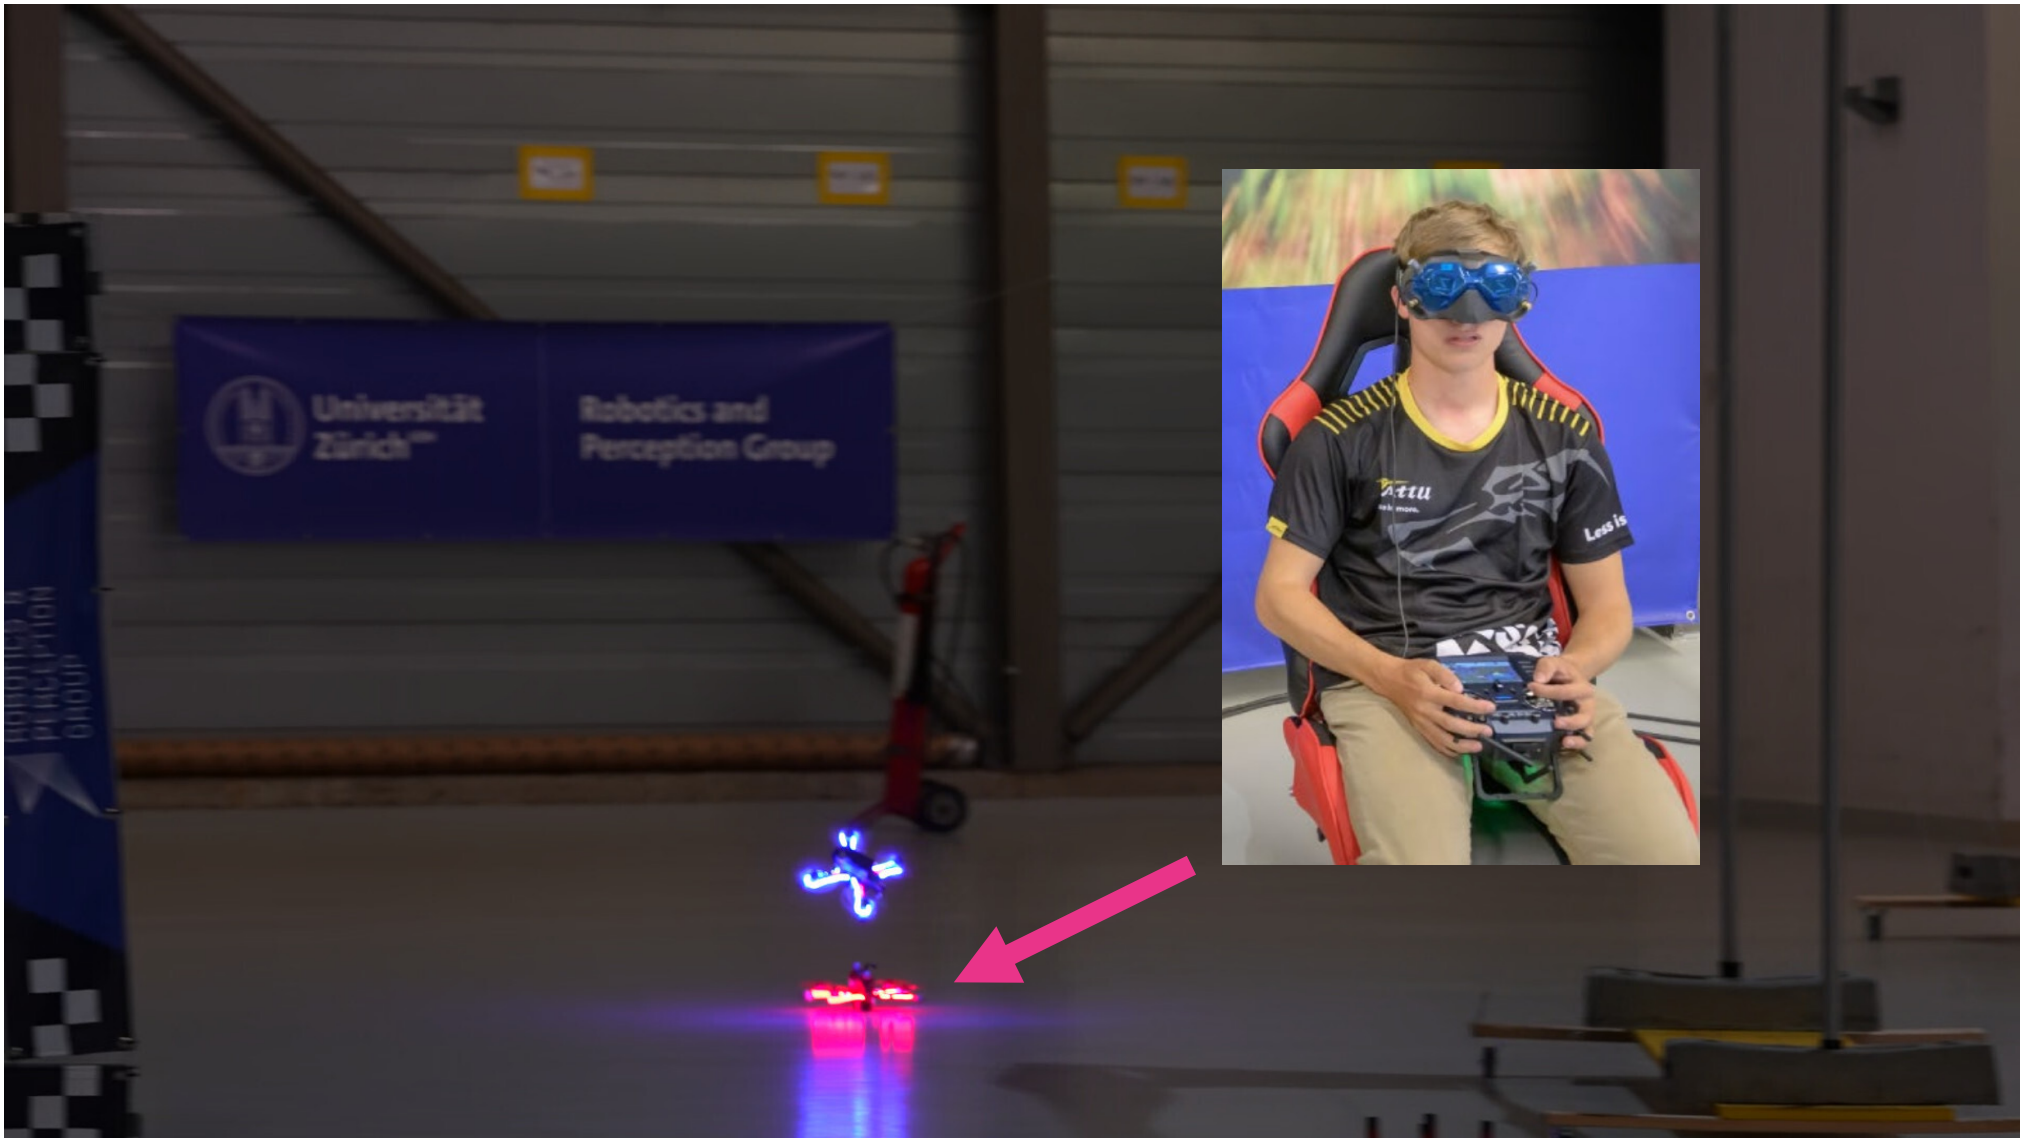
Universität Zürich
Robotics and Perception Group

Co-Adaptation

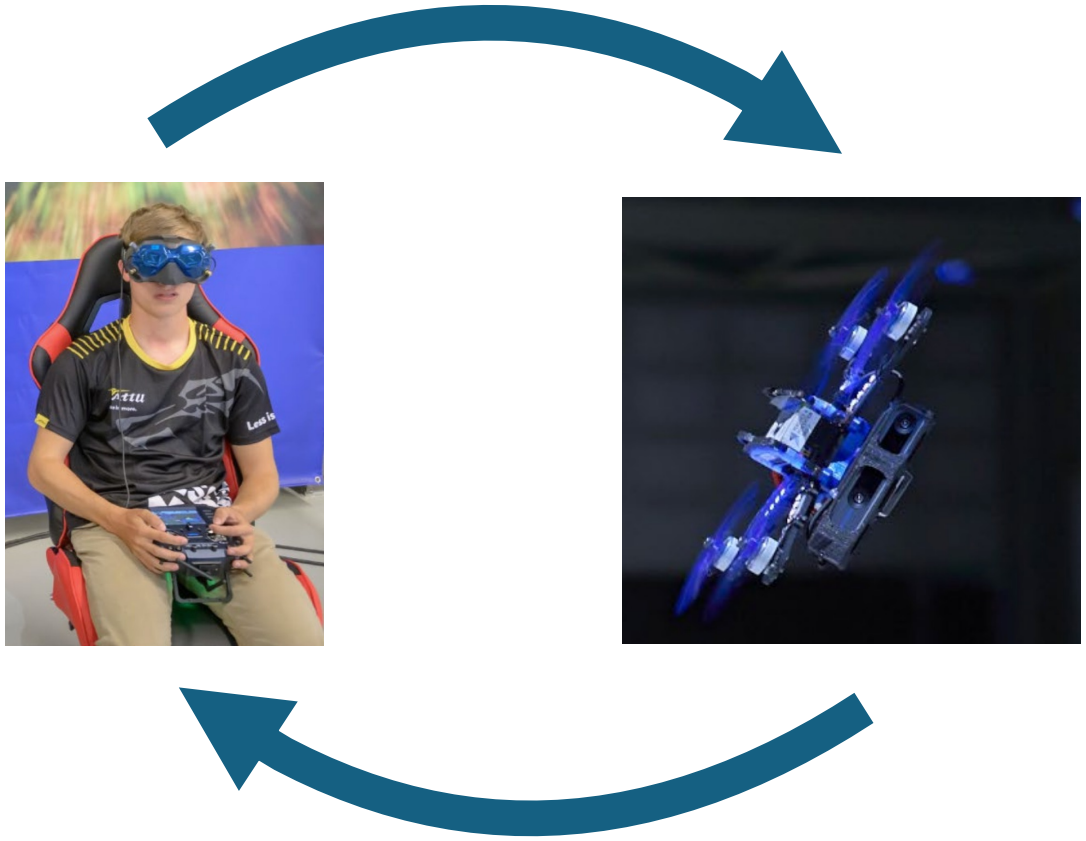

Learning Visual Locomotion with Cross-Modal Supervision

Antonio Loquercio

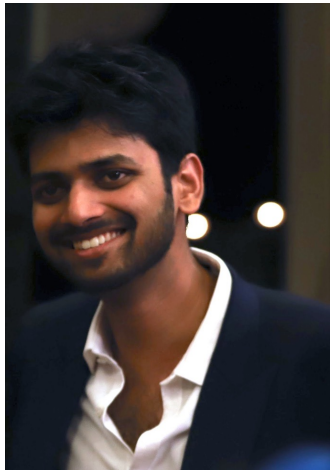

Ashish Kumar

Jitendra Malik


Day 1 (2X)
Loquercio et. al,
ICRA, 2023

Loquercio et. al, ICRA, 2023


Loquercio et. al,
ICRA, 2023

Visual Plasticity

Before Adaptation

After 1min of data

Learning Vision-Based Pursuit-Evasion Policies



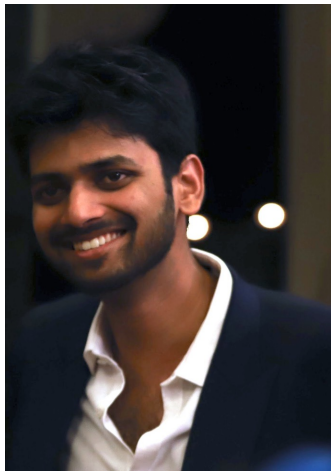


Andrea Bajcsy* Antonio Loquercio* Ashish Kumar Jitendra Malik

Vision-based Pursuit-Evasion Game

Learning from Videos

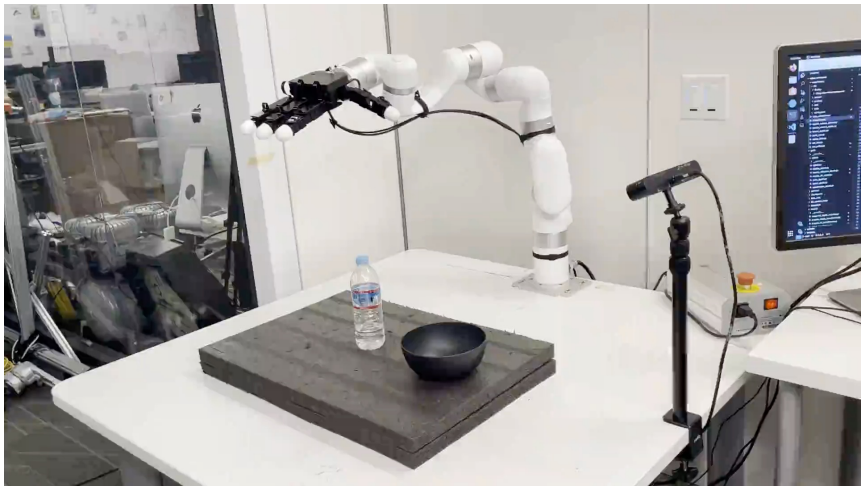

Human videos for scaling up robot data

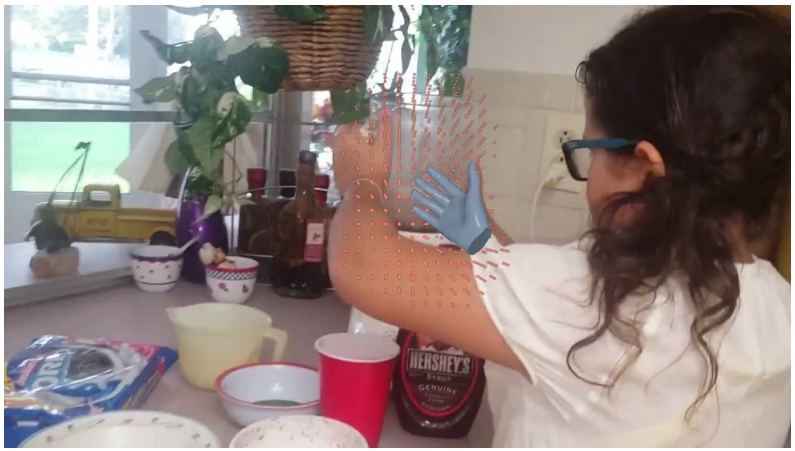

1. Intuitive physics
2. Contact poses
3. Pre/Post-contact trajectories
4. Human preferences
5. ….

Sample-efficient BC Finetuning