


Real-World Robot Learning (ESE 6800) Vision-Based Robot Learning (CS 7000)

Instructors:

Antonio Loquercio
Dinesh Jayaraman

Spring 2025

University of Pennsylvania



Reading at a Table
Pablo Picasso

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 Real-World Robot Learning?

“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 Real-World Robot Learning?

“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!**

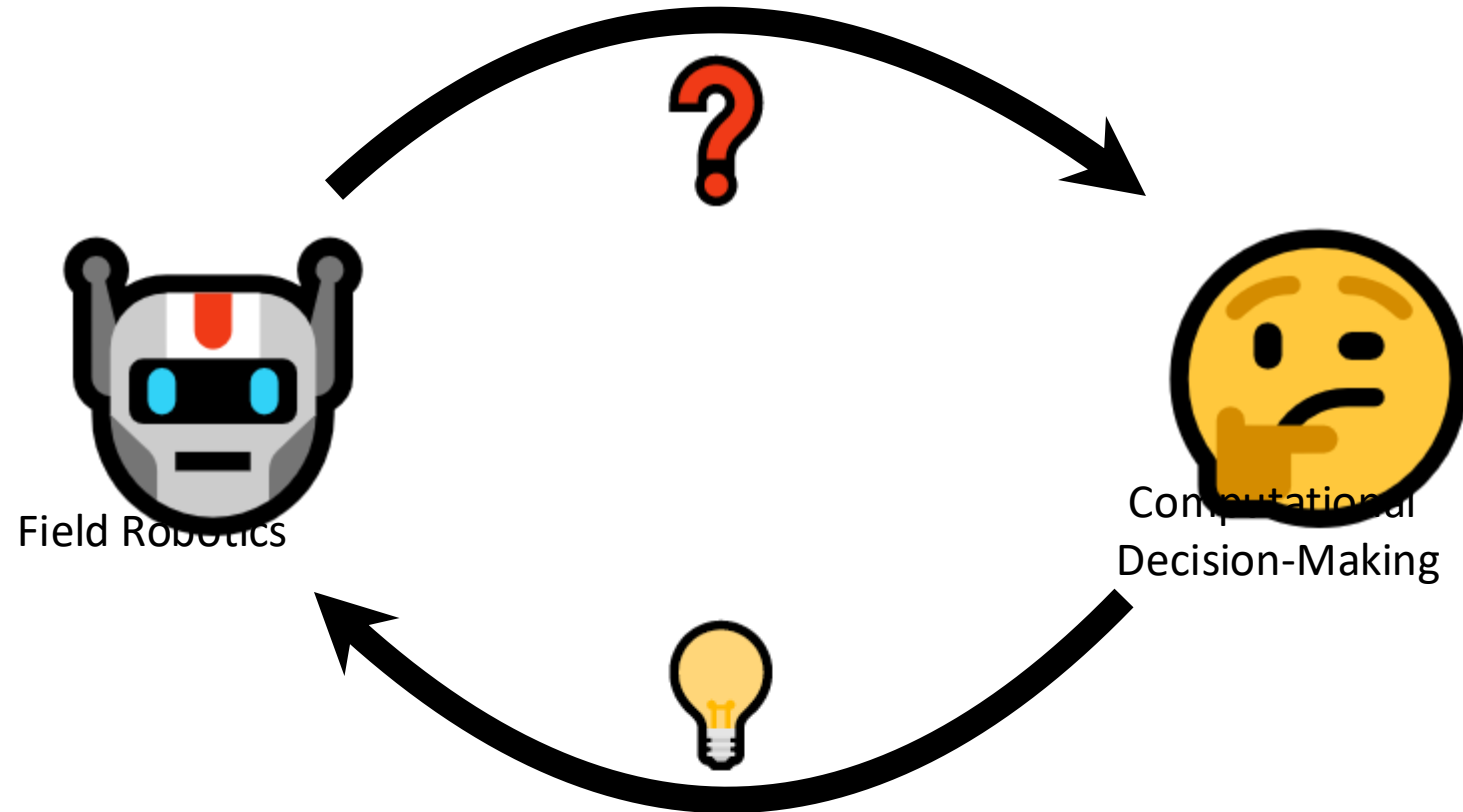
Why care about Real-World Robot Learning?

“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 \approx Reward Hacking
 - Everything works given enough constraints



Overreaching Theme



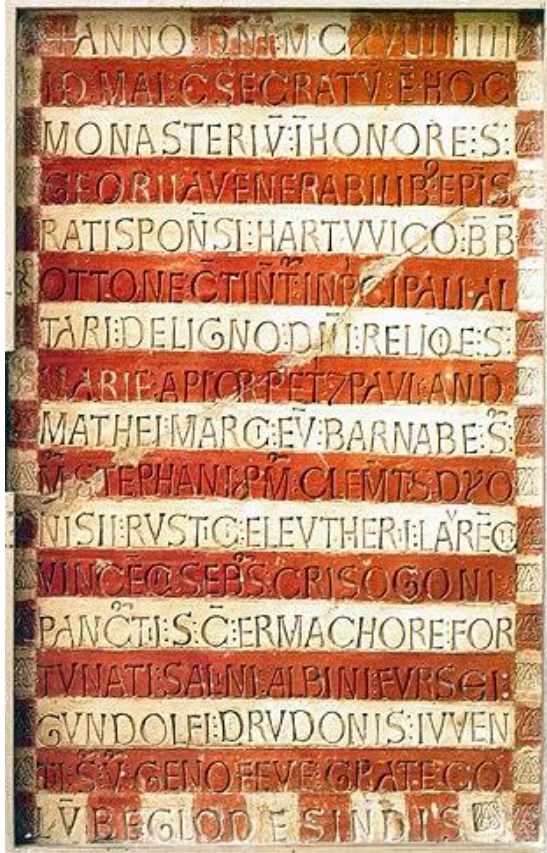
What's different about this class?

- We will try to go off the current beaten path
 - Get off the Arxiv treadmill for a bit, then come back to the latest and greatest
- Try to connect the old and the new
 - Making unexpected connections is the crux of creativity
 - Try to get inspiration from old ideas
- Understand that novelty is a nuanced concept
 - Almost all the hot things that work now are an old idea whose time finally came

Phaistos Disc (Crete, 2nd millennium BCE)



Two Thousands Years Later...



Prüfeninger Weiheinschrift
Prüfening dedicatory inscription
Germany, 1119





Movable Type Printing
Bi Sheng
China, 1041–1048

Three Hundred Years Later



Johannes Gutenberg
Printing Press
1440





Mobile ALOHA: Learning Bimanual Mobile Manipulation with Low-Cost Whole-Body Teleoperation

Zipeng Fu, Tony Z. Zhao, Chelsea Finn



BRITISH
PATHÉ

LIGHTER SIDE OF THE NEWS

Commentary by PETER ROBERTS

NEWS of the DAY

Philadelphia,
1957

Goals

- Read some interesting papers together
 - Learn something new: both you and us!
- Get up to speed on big chunk of robot learning
- Use more learning-based techniques in your own work
- Try your hand in a robot project
- Learn how to speak
- Learn how think critically about papers

Course Organization

Grading (tentative, based on class size?):

1. Class Participation (40%)

- Write and submit short summaries (1/2 page per paper) of required reading **before each class** on canvas (20%)
- Post *something** on canvas before **each class** (10%)
- **Attend class.** Ask questions / debate / fight / be involved! (10%)

2. Paper Presentations + other roles (20%)

- Give **high quality** paper presentations. More on presentations later. (15%)
- 5% for other roles (more coming up soon)
- Bonus point: run it on new data.

3. Final Project (40%)

- Can be done solo or in groups of 2
- Hopefully comes out of your paper readings / discussions
- Submit project proposal and midterm report (10%)
- Final presentations & report at end of semester (30%)

Class Participation

Keep annotated bibliography of papers you read (always a good idea!). The format is up to you. At least, it needs to have:

- Summary of key points
- A few Interesting insights, “aha moments”, keen observations, etc.
- Weaknesses of approach. Unanswered questions. Areas of further investigation, improvement.



Class Participation

Before each class:

1. Submit your summary for current paper(s) to Canvas
2. Submit a comment on Canvas:
 - **Paper summary:** if you are quick, you could be the first to create a topic for a paper and post your summary (0 extra work!)
 - **Improve posted summary:** If a summary is already posted, add missing points. Put the paper in larger context. Share your experiences with using this paper.
 - **Pose question:** if you didn't get something in this paper, chances are others might be in same boat.
 - **Highlight areas of improvement**

Impossible Mturk HITS and iterations

I have attempted to hit this mark, but I've
found it very difficult to do so, and I
cannot see how it can be done. There is a
fundamental contradiction, I believe, that the
system's design does not allow for a consistent
reading of the data. My own attempts
have been highly successful, but I've found
that it is impossible to do so. The
system is designed to be used in
a way that is not possible. I've
found that the system is not
designed to be used in a way that
is not possible. I've found that the
system is not designed to be used in a way
that is not possible. I've found that the
system is not designed to be used in a way
that is not possible.

Experiment by Greg Little

Slide from Antonio Torralba

Impossible Mturk HITS and iterations

I had intended to hit the wall, but I'm
 not a very good wall in general, and I
 would like things very much. This is a
 common mistake, I know, but it
 doesn't mean that I've been completely
 reading since it started, my own strategy
 will involve highly assuming that most of the
 reading is what it is repeated into the
 second enough that the second half is
 looking at strings, and then I'll restore
 my hand and point something away. That's
 how this goes.

- Please transcribe this text.
- Put words you are unsure about in (parentheses).
- Use (?) for words you don't know.

character count: 0/500 please type something above
 Submit

non-iterative

- ???
 ???
 ?????? ~ 1287 seconds
- - unreadable - ~ 60 seconds
- (unreadable) ~ 59 seconds
- ? ~ 74 seconds
- words are blurred ~ 17 seconds
- *Experiment stopped here, still waiting on 3 responses.*

Experiment by Greg Little

Impossible Mturk HITS and iterations

I had intended to hit the nail, but I'm not a very good aim it seems, and I ended up hitting my thumb. This is a common occurrence, I know, but it doesn't make me feel any less having done it myself. My new strategy will involve lightly tapping the nail while holding it, until it is into the (enough) that the (itself) itself is holding it straight, and then I'll remove my hand and pound (carefully) away. We'll see how (this) goes.

- Please transcribe this text.
- Put words you are unsure about in (parentheses).
- Use (?) for words you don't know.

character count: 0/500 please type something above

Submit

iterative

- I had attended to the mail, but I'm not a very good ~ 431 seconds
- I (had) (resorted) to (hit) (the) (wall), (but) I'm (not) (a) (very) (good) (?) (is) (?), (?) (I) (?) (up) (?) (my) (?). This is a (?) (?), I (know), but (it) (doesn't) (mean) (the) (?) (sure) (?) (?) (?) (it) (?). My new (strategy) (will) (involve) (lightly) (?) then (?) (while) (?) (it) (is) (?) (?) (this) (wasn't) (enough) (?) (?) (?) (?) (is) (?) (?) (straight), and then I'll (?) my hand and (?) (?) (?) (away). We'll (?) (?) (?) (?). ~ 1349 seconds
- I (had) (resorted) to (hit) (the) (wall), (but) I'm (not) (a) (very) (good) (quitter). (I) (gave) (up) (on) (my) (ability). This is a (shame) , I (know), but (it) (doesn't) (mean) (the) (end) (surely) (will) (come) (anytime) (soon). My new (strategy) (will) (involve) (lightly) (thinking) then (doing) (while) (trying) (to) (make) (it) (happen). (This) (wasn't) (supposed) (to) (happen) (this) (way). ~ 353 seconds
- I had intended to hit the nail, but I'm not a very good (aim) it seems, and I ended up hitting my thumb. This is a common occurrence, I know, but it doesn't (make) me feel any (less) (?) having done it myself. My new strategy will involve lightly tapping the nail while holding it, until it is (?) into the (?) (enough) that the (?) itself is holding it straight, and then I'll remove my hand and pound (carefully) away. We'll see how (this) goes. ~ 645 seconds
- I had intended to hit the nail, but I'm not (a) very good (aim) in it (seems), and I ended up hitting my thumb. This is a common occurrence, I know, but it doesn't make me feel (any) (less) (?) having done it myself. My new strategy will involve lightly tapping the nail (while holding it), until it is (?) into the (?) (enough) that the (?) itself is holding it straight, and then I'll remove my hand and pound (carefully) away. We'll see how (this) goes. ~ 318 seconds
- I had intended to hit the nail, but I'm not a very good (aim) it seems, and I ended up hitting my thumb. This is a common occurrence, I know, but it doesn't (make) me feel any (less) (?) having done it myself. My new strategy will involve lightly tapping the nail while holding it, until it is (?) into the (?) (enough) that the (?) itself is holding it straight, and then I'll remove my hand and pound (carefully) away. We'll see how (this) goes. ~ 103 seconds
- I had intended to hit the nail, but I'm not a very good aim it seems, and I ended up hitting my thumb. This is a common occurrence, I know, but it doesn't make me feel any less ridiculous having done it myself. My new strategy will involve lightly tapping the nail while holding it, until it is embedded into the wood enough that the wood itself is holding it straight, and then I'll remove my hand and pound carefully away. We'll see how this goes. ~ 159 seconds

Experiment by Greg Little

Impossible Mturk HITS and iterations

I had intended to hit the nail, but I'm not a very good aim it seems, and I ended up hitting my thumb. This is a common occurrence, I know, but it doesn't make me feel any less ridiculous having done it myself. My new strategy will involve lightly tapping the nail while holding it, until it is wedged into the wood enough that the wood itself is holding it straight, and then I'll remove my hand and pound carefully away. We'll see how this goes.

original:

I had intended to hit the nail, but I'm not a very good aim it seems, and I ended up hitting my thumb. This is a common occurrence, I know, but it doesn't make me feel any less ridiculous having done it myself. My new strategy will involve lightly tapping the nail while holding it, until it is **wedged** into the wood enough that the wood itself is holding it straight, and then I'll remove my hand and pound carefully away. We'll see how this goes.

Iterative:

I had intended to hit the nail, but I'm not a very good aim it seems, and I ended up hitting my thumb. This is a common occurrence, I know, but it doesn't make me feel any less ridiculous having done it myself. My new strategy will involve lightly tapping the nail while holding it, until it is **embedded** into the wood enough that the wood itself is holding it straight, and then I'll remove my hand and pound carefully away. We'll see how this goes.

Experiment by Greg Little

Slide from Antonio Torralba

Paper Presentations





1. Pick one topic from the list (we will have ~2 people per topic) by next week.
2. Each topic will have papers to be presented.
3. Prepare an **amazing** 15-20 min presentation, focusing on the picked paper but also providing context
4. We will provide a template, but the format is up to you.
5. Incorporate comments/questions from canvas.
6. **Seed** a discussion
7. Paper presentation + discussion = 35-40 min
8. Meet us the day before to show slides

The Discussion Game

We will do a role-playing paper-reading exercise.



Which roles?

-  Defender (15 mins) Reviewer / Discussion Lead (presenter)
 -  Archaeologist.
 -  Industry practitioner
 -  Private Investigator
-
- Format for presentation:
 - Presenter in Defender mode (15 mins w/ slides)
 - Archaeologist (1 min)
 - Industry Practitioner (1 min)
 - Private Investigator (1 min)
 - Presenter in Discussion lead mode. (w/ slides)

How to give a great talk

- Some good points:
 - http://www.cs.berkeley.edu/~messer/Bad_talk.html
- Great advise from David Fouhey:
 - https://web.eecs.umich.edu/~fouhey/teaching/EECS542_F21/commonissues.pdf
 - Very easy to follow. **Please go over this check-list before you come to us!**
- There are no unimportant talks.
 - There are no big or small audiences.
 - Prepare each talk with the same enthusiasm.

How to give a talk

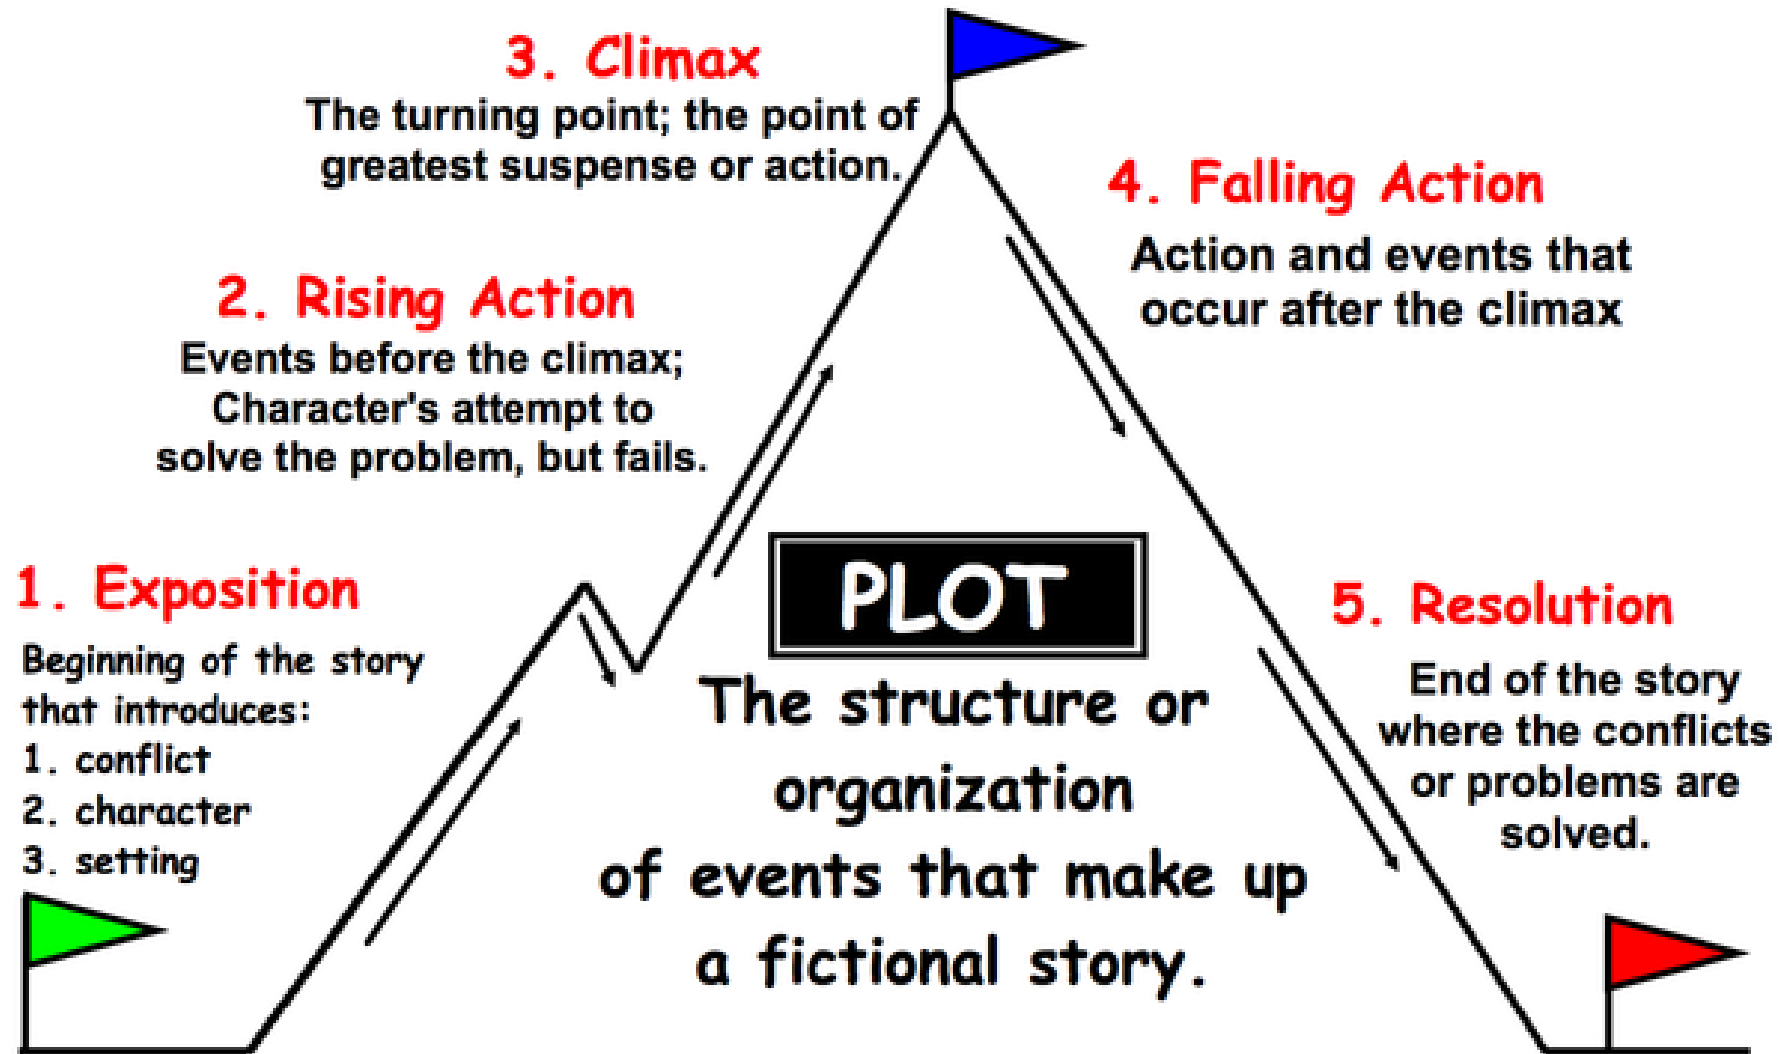
Talk organization: here there are as many theories as there are talks. Here there are some extreme advices:

1. Go into details / only big picture
 2. Go in depth on a single topic / cover as many things as you can
 3. Be serious (never make jokes, maybe only one) / be funny (it is just another form of theater)
- Corollary: ask people for advice, but at the end, it will be just you and the audience. Choose what fits best your style.
 - What everybody agrees on is that you have to practice in advance (the less your experience, the more you have to practice). Do it with an audience or without, but practice.


The best advice I got came from Alyosha Efros while preparing my job talk:

“just give a good talk”

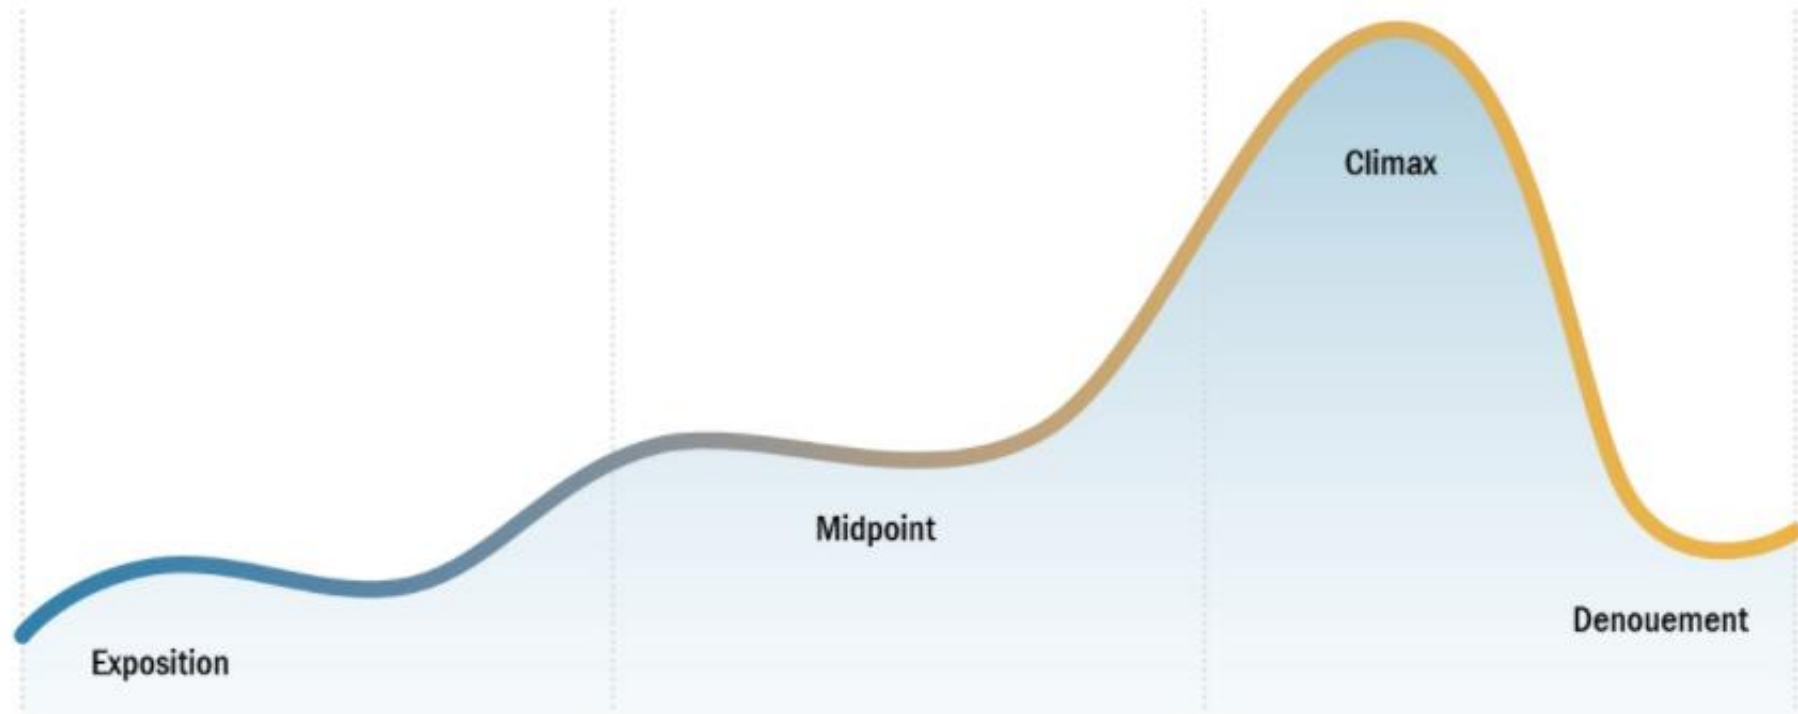
Your talk must have a **story!**



Your talk must have a story!



Your talk must have a story!



Some Bad News

“The more you work on a talk, the better it gets: if you work on it for 3 hours, the talk you give will be better than if you had only worked on it for 2 hours. If you work on it for 5 hours, it will be better still. 7 hours, better yet...”

How to give a talk

Delivering:

- Look at the audience! Try not to talk to your laptop or to the screen. Instead, look at the other humans in the room.
- You have to believe in what you present, be confident... even if it only lasts for the time of your presentation.
- Let the audience see your personality
 - They want to see you enjoy yourself.
 - They want to see what you love about the work.

Class Project

- Can grow out of your presentation, reading or your own research.
- But hopefully inspired by topics in the course.
- 1-2 people per project.
- Project proposals end of Feb
- Project midterm end of March
- Project presentations at the end of semester.
- Report presented as a CORL-style paper.
- Hopefully, a few papers may be submitted to conferences.

End of Semester Awards

- We will vote for:
 - Best Presentation
 - Best Project
- Prize: dinner in my favorite restaurant in Rome (transportation not included!) or some other worthy prizes 😊

Who this class is for

- We are hoping to make this small and cozy, facilitate fun discussions.
- This is not an intro level class. If you don't have previous experience with ML/CV this will be hard.
- You can't audit without doing paper presentations and summaries.

(Tentative) Schedule

- Intro to BC/RL (people who are already familiar can skip)
 - Poll: Pytorch or Simulators for RL?
- **Foundations**
 - Perception
 - The development perspective
- **Data**
 - Real-World
 - Others'
 - Internet
 - Synthetic
 - Lifelong
- **Challenges for the field**

Introductions!

- Each person describe (about 30 seconds):
 - Name
 - Which course, year
 - Research interests, which faculty you are working with if any.
 - Where in the world are you from
 - Favorite food / favorite travel destination

More Polls!



- Poll of backgrounds
 - Familiarity with machine learning
 - Familiarity with computer vision
 - Familiarity with robots / real robot experience?
- Poll of programs
 - PhD
 - MS
 - Undergraduates
- Poll of waitlist vs. already in class
- Poll of familiarity with the seminar format

Some cool things we have done
with robot learning techniques


Imitation Learning




Imitating Manually Operated Robots




Imitating Humans Directly?



Generalizing Across Objects, Views, Rooms, Robots!




Training Robots in Simulation




Iteration 1

Training Robots in Simulation




Iteration 2

Training Robots in Simulation




Iteration 3

Training Robots in Simulation




Iteration 4

Training Robots in Simulation



Iteration 5

Training Robots in Simulation



The international journal of science / 31 August 2023

nature

DRONE RACER


AI pilot beats human competitors in real-world championship



Vol. 620, No. 7976
© 2023 Nature.com

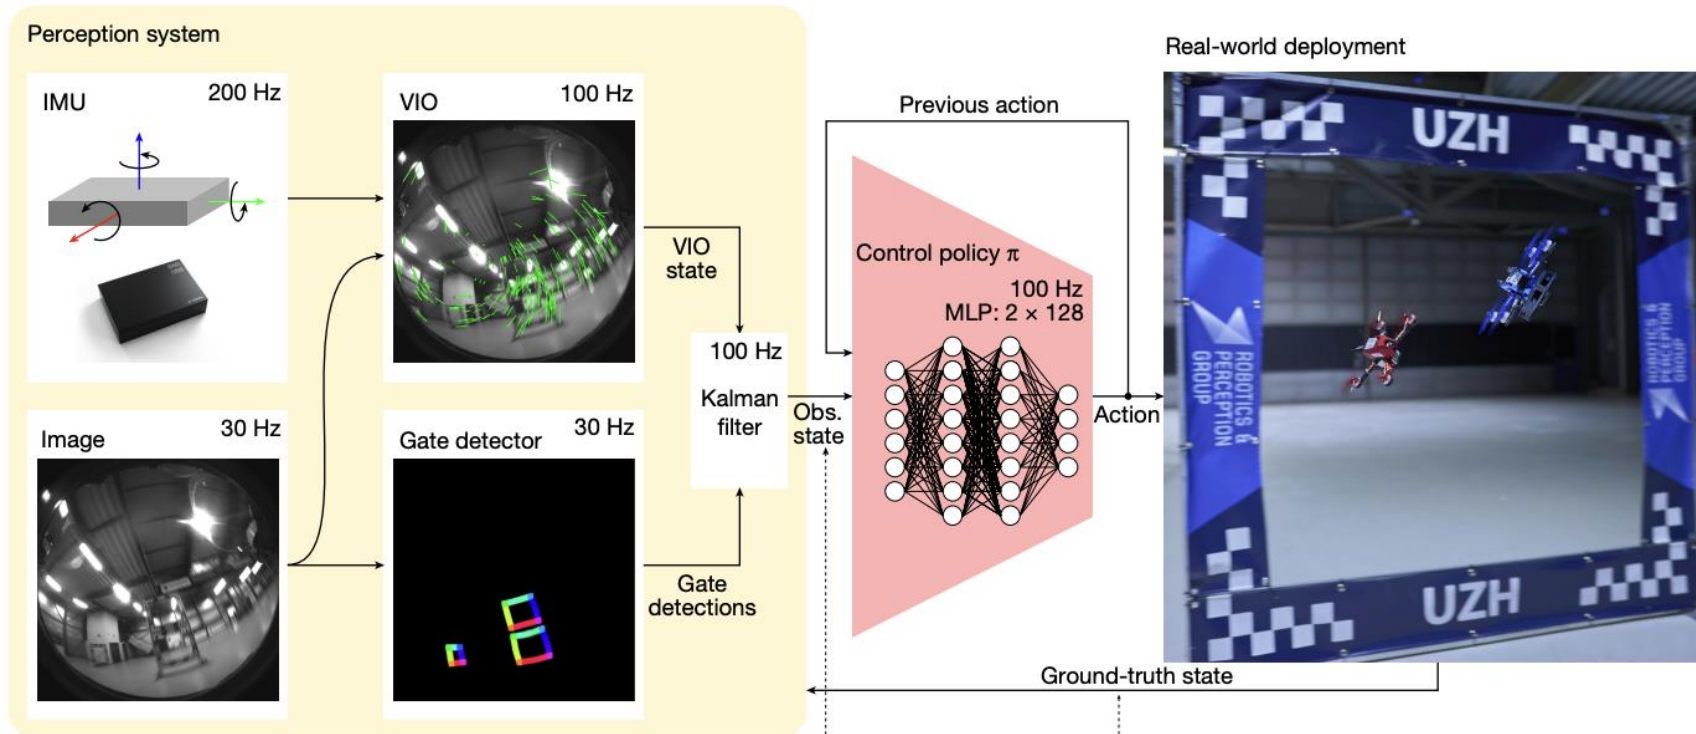
Article | [Open access](#) | Published: 30 August 2023

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](#)

Simulation to Real World Transfer



b RL training loop



ROBOTICS &
PERCEPTION
GROUP



Universität
Zürich ^{UZH}

Robotics and
Perception Group

robotics⁺

Swiss
Centr
in R

1

2

ROBOTICS
PERCEPTION
GROUP



Innovation
Zurich



Canton of Zurich
Department for Economic Affairs
Office for Economy and Labour



TEA



UZH



Learning High-Speed Flight in the Wild

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

Science Robotics, 2021

40 km/h



Construction Zone



Loquercio et. al,
ICRA, 2023



Some Cool Results of Robot Learning



OpenAI 2019



Zhao, Kumar, Levine, Finn, RSS 2023



SayCan, CORL 2022



Marco Hutter's talk at ICRA 2022, Philadelphia



Fu, Cheng, Pathak, 2022