Deep Multi-Task and Meta-Learning

CS 330

Course Logistics

Information & Resources

Chelsea Finn
Instructor

Suraj Nair
TA

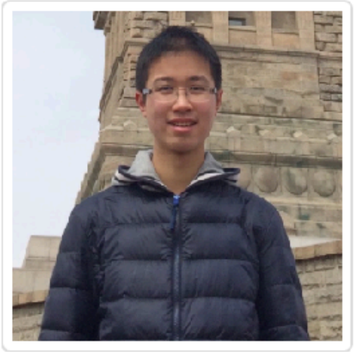

Tianhe (Kevin) Yu
TA

Abhishek Sinha
TA

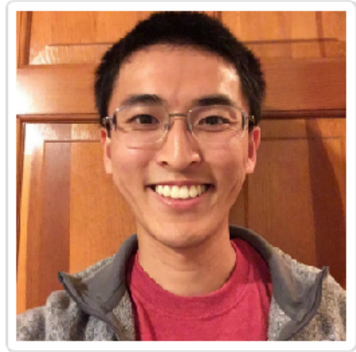

Tim Liu
TA

Course website: http://web.stanford.edu/class/cs330/

Piazza: Stanford, CS330

Staff mailing list: cs330-aut1920-staff@lists.stanford.edu

Office hours: Check course website. (Mine are Weds after class)

Pre-Requisites and Enrollment

Pre-requisites: CS229 or equivalent, previous RL experience highly recommended

If you are not enrolled: fill out enrollment form on course website.

- We will enroll subject to availability
- Fill out the form as soon as possible!

Lectures are recorded, will be internally released on Canvas, will be publicly released after the course.

SCPD: There are ~20 remote students from SCPD as part of the course.

Assignment Infrastructure

Assignments will require training networks in **TensorFlow (TF)**.

TF review section:

- Suraj Nair will hold a TF review session on Thursday, September 26.
- You should be able to understand the overview here: https://www.tensorflow.org/guide/low_level_intro
- If you don't, go to the review session & ask questions!

Topics

1. Problem definitions
2. Multi-task learning basics
3. Meta-learning algorithms: black-box approaches, optimization-based meta-learning, metric learning
4. Hierarchical Bayesian models & meta-learning
5. Multi-task RL, goal-conditioned RL, hierarchical RL
6. Meta-reinforcement learning
7. Open problems, invited lectures, research talks

Emphasis on **deep learning**, **reinforcement learning**

Topics We Won't Cover

Won't cover **AutoML** topics:

- architecture search
- hyperparameter optimization
- learning optimizers

Emphasis will be on:
deep learning approaches

Course Format

Three types of course sessions:

- **lecture** (9)
- **student reading**: presentations & discussions (7)
- **guest lecture** (3)

All students responsible for **one group paper presentation**.

[Instructions posted on Piazza.]

Participation in **discussions** is highly encouraged.

[This will change in future offerings.]

Assignments & Final Project

Homework 1: Multi-task data processing, black-box meta-learning

Homework 2: Gradient-based meta-learning & metric learning

Homework 3: Multi-task RL, goal relabeling

Final project: Research-level project of your choice

Form groups of 1-3 students, you're welcome to start early!

Grading: 20% paper presentation, 30% homework (10% each), 50% project

5 late days total across: homeworks, project paper submission

Homework Today

1. Sign up for Piazza
2. Fill out paper presentation preferences **(by Thursday!)**
3. Start forming final project groups if you want to work in a group
4. Review this: https://www.tensorflow.org/guide/low_level_intro

Two more things

Ask questions!

Because it is new, this course will be rough around the edges.

Some of My Research

(and why I care about multi-task learning and meta-learning)

How can we enable agents to learn skills in the real world?

Robots.

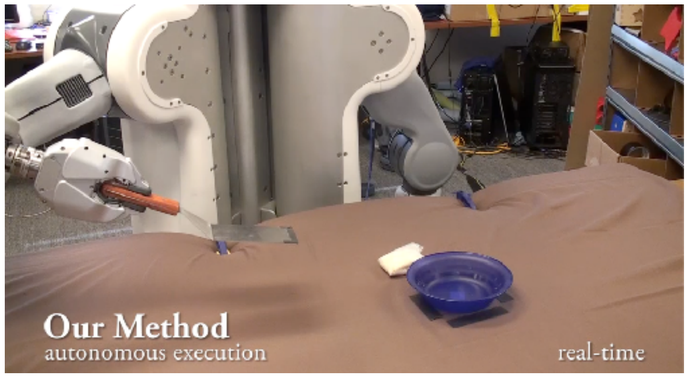

Finn, Tan, Duan, Darrell, Levine, Abbeel.
ICRA '16

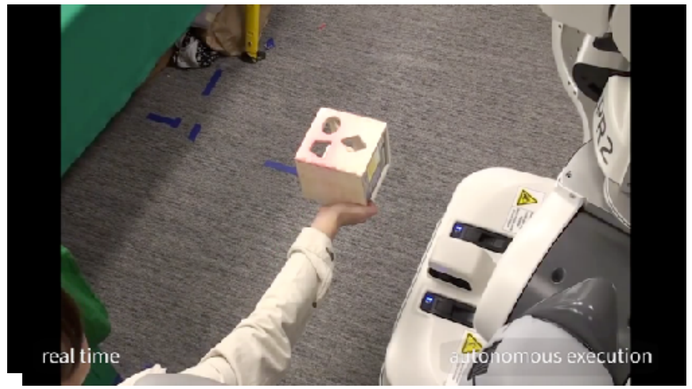

Levine*, Finn*, Darrell, Abbeel.
JMLR '16

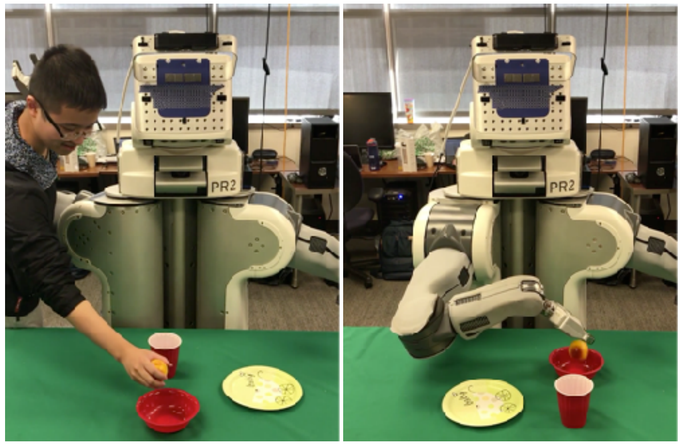

Yu*, Finn*, Xie, Dasari, Zhang,
Abbeel, Levine, RSS '18

Why robots? Robots can teach us things about intelligence.

faced with the **real world**

must **generalize** across tasks, objects, environments, etc

need some **common sense understanding** to do well

supervision can't be taken for granted

Beginning of my PhD

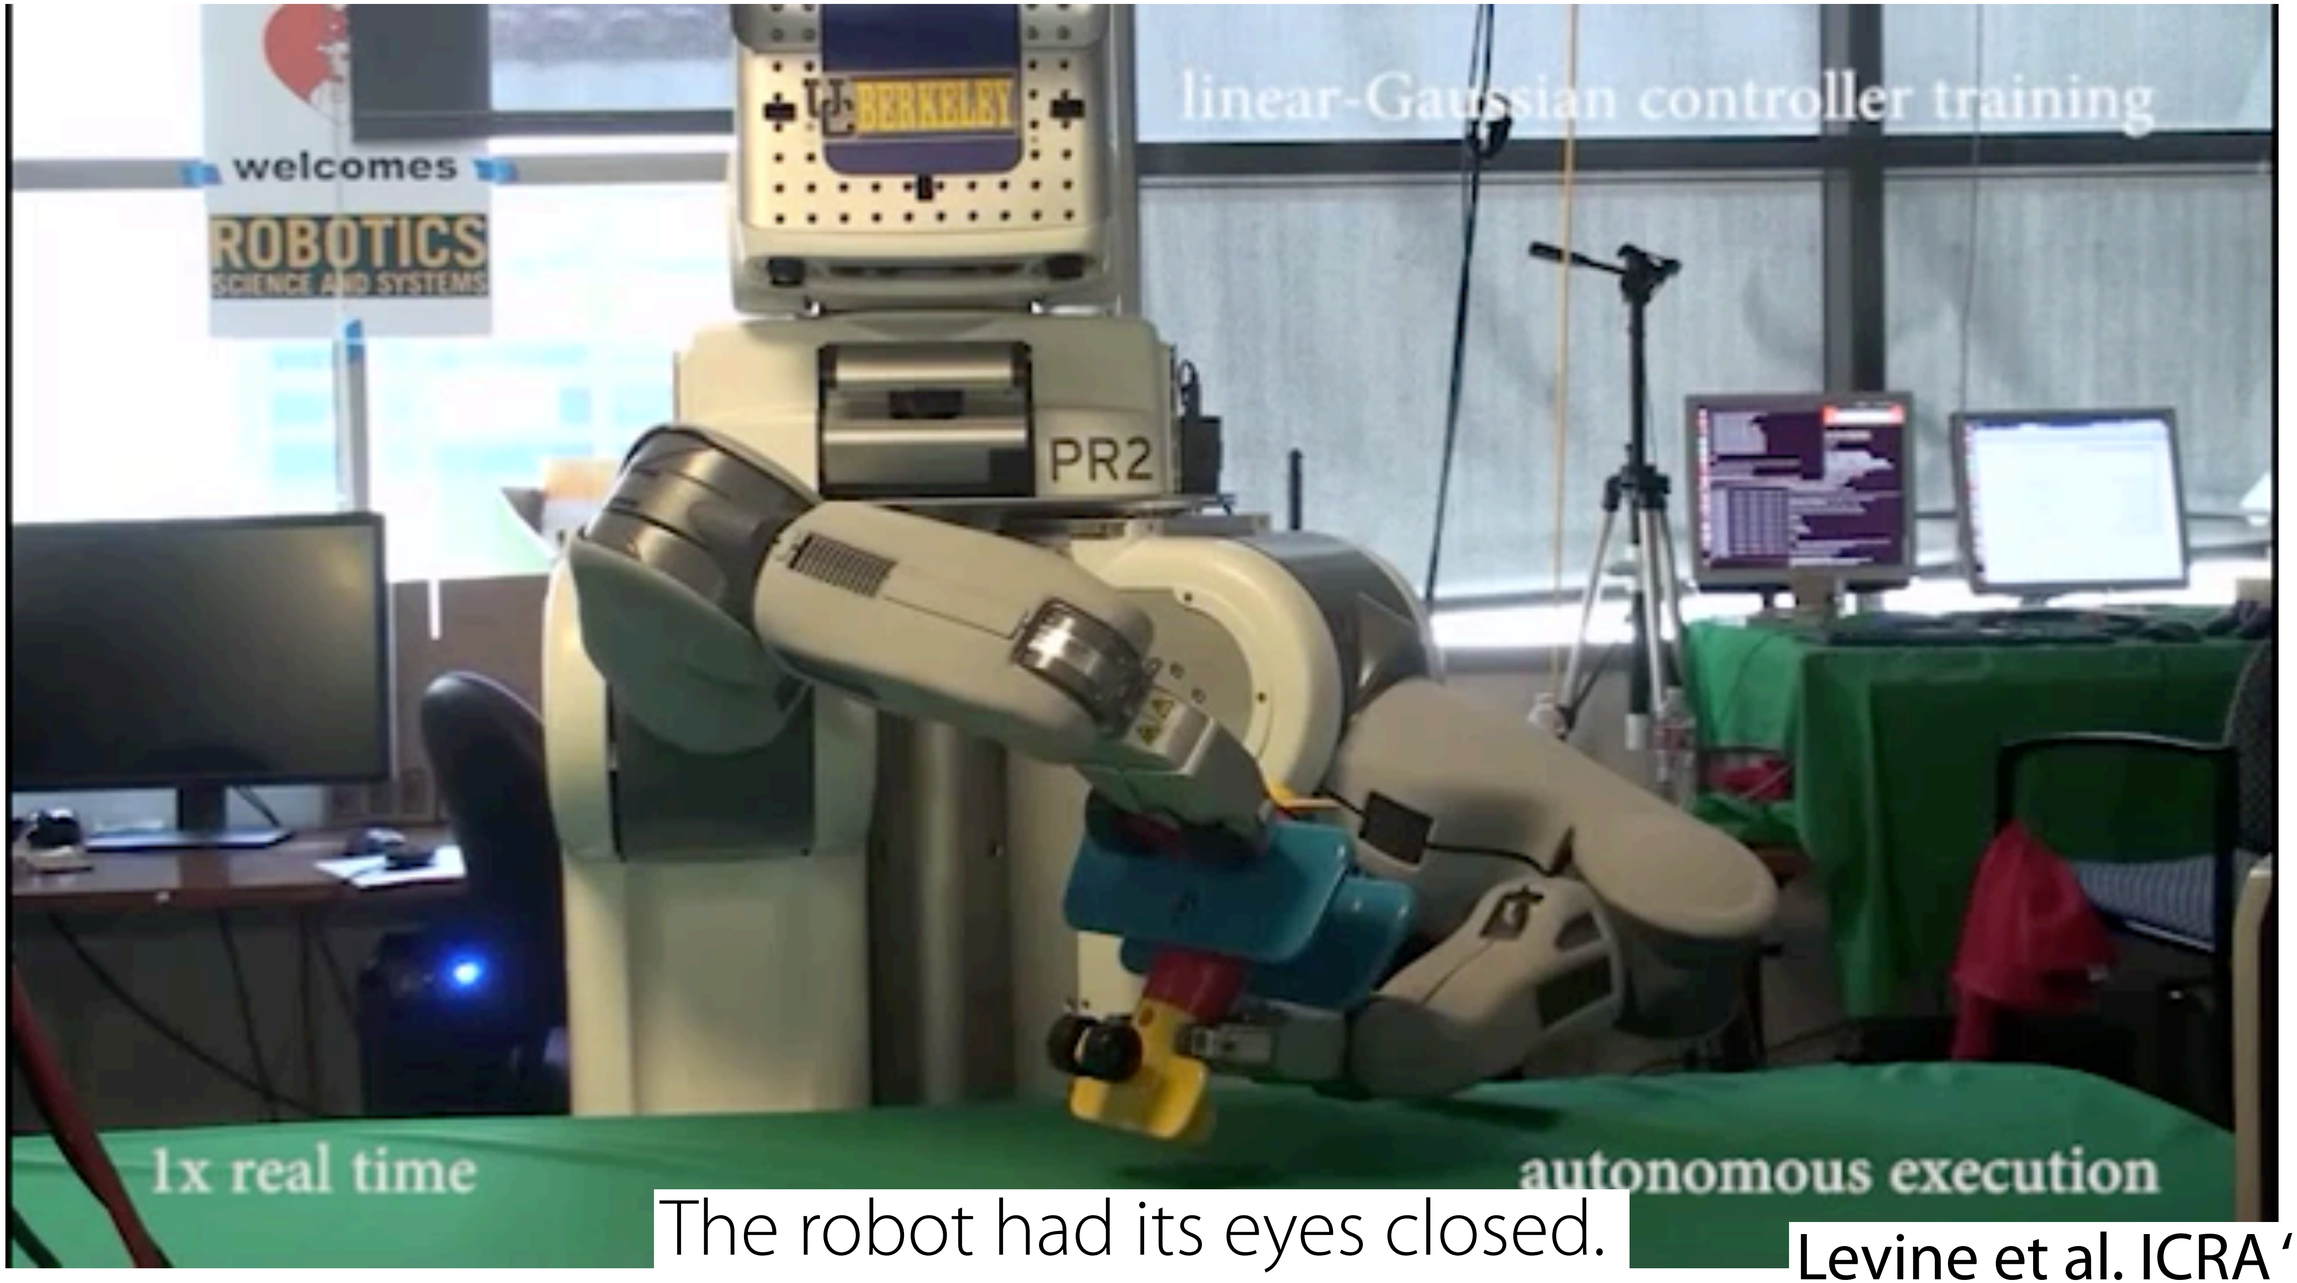

The robot had its eyes closed.

Levine et al. ICRA '15

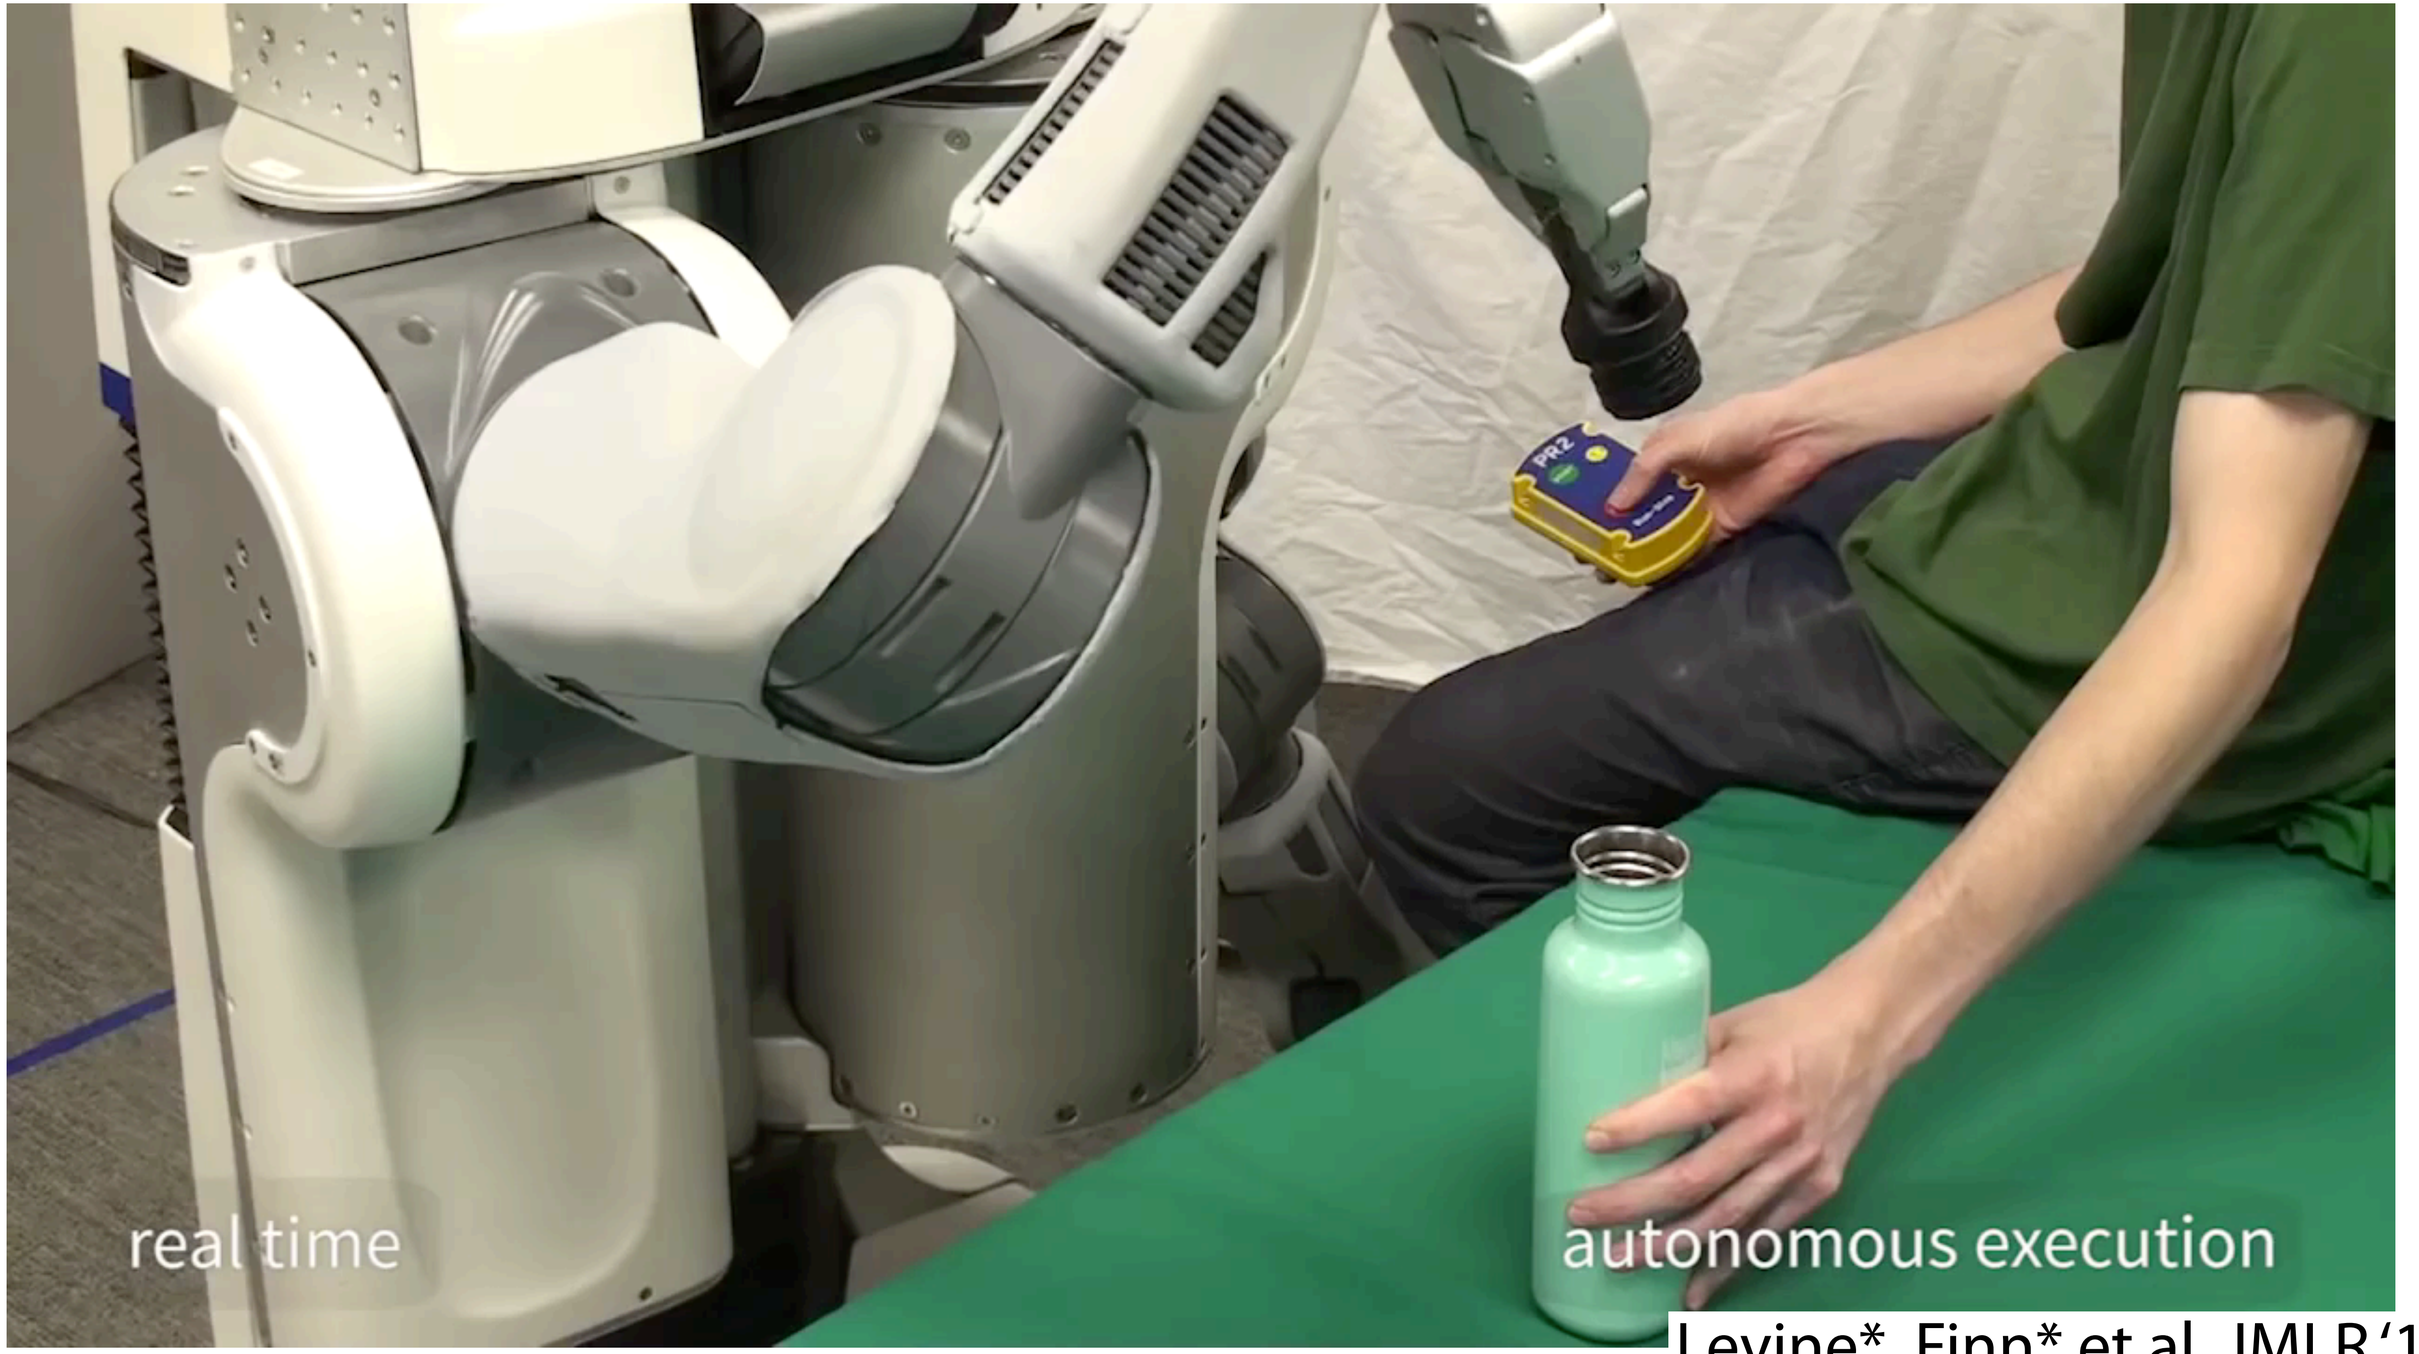

Levine*, Finn* et al. JMLR '16

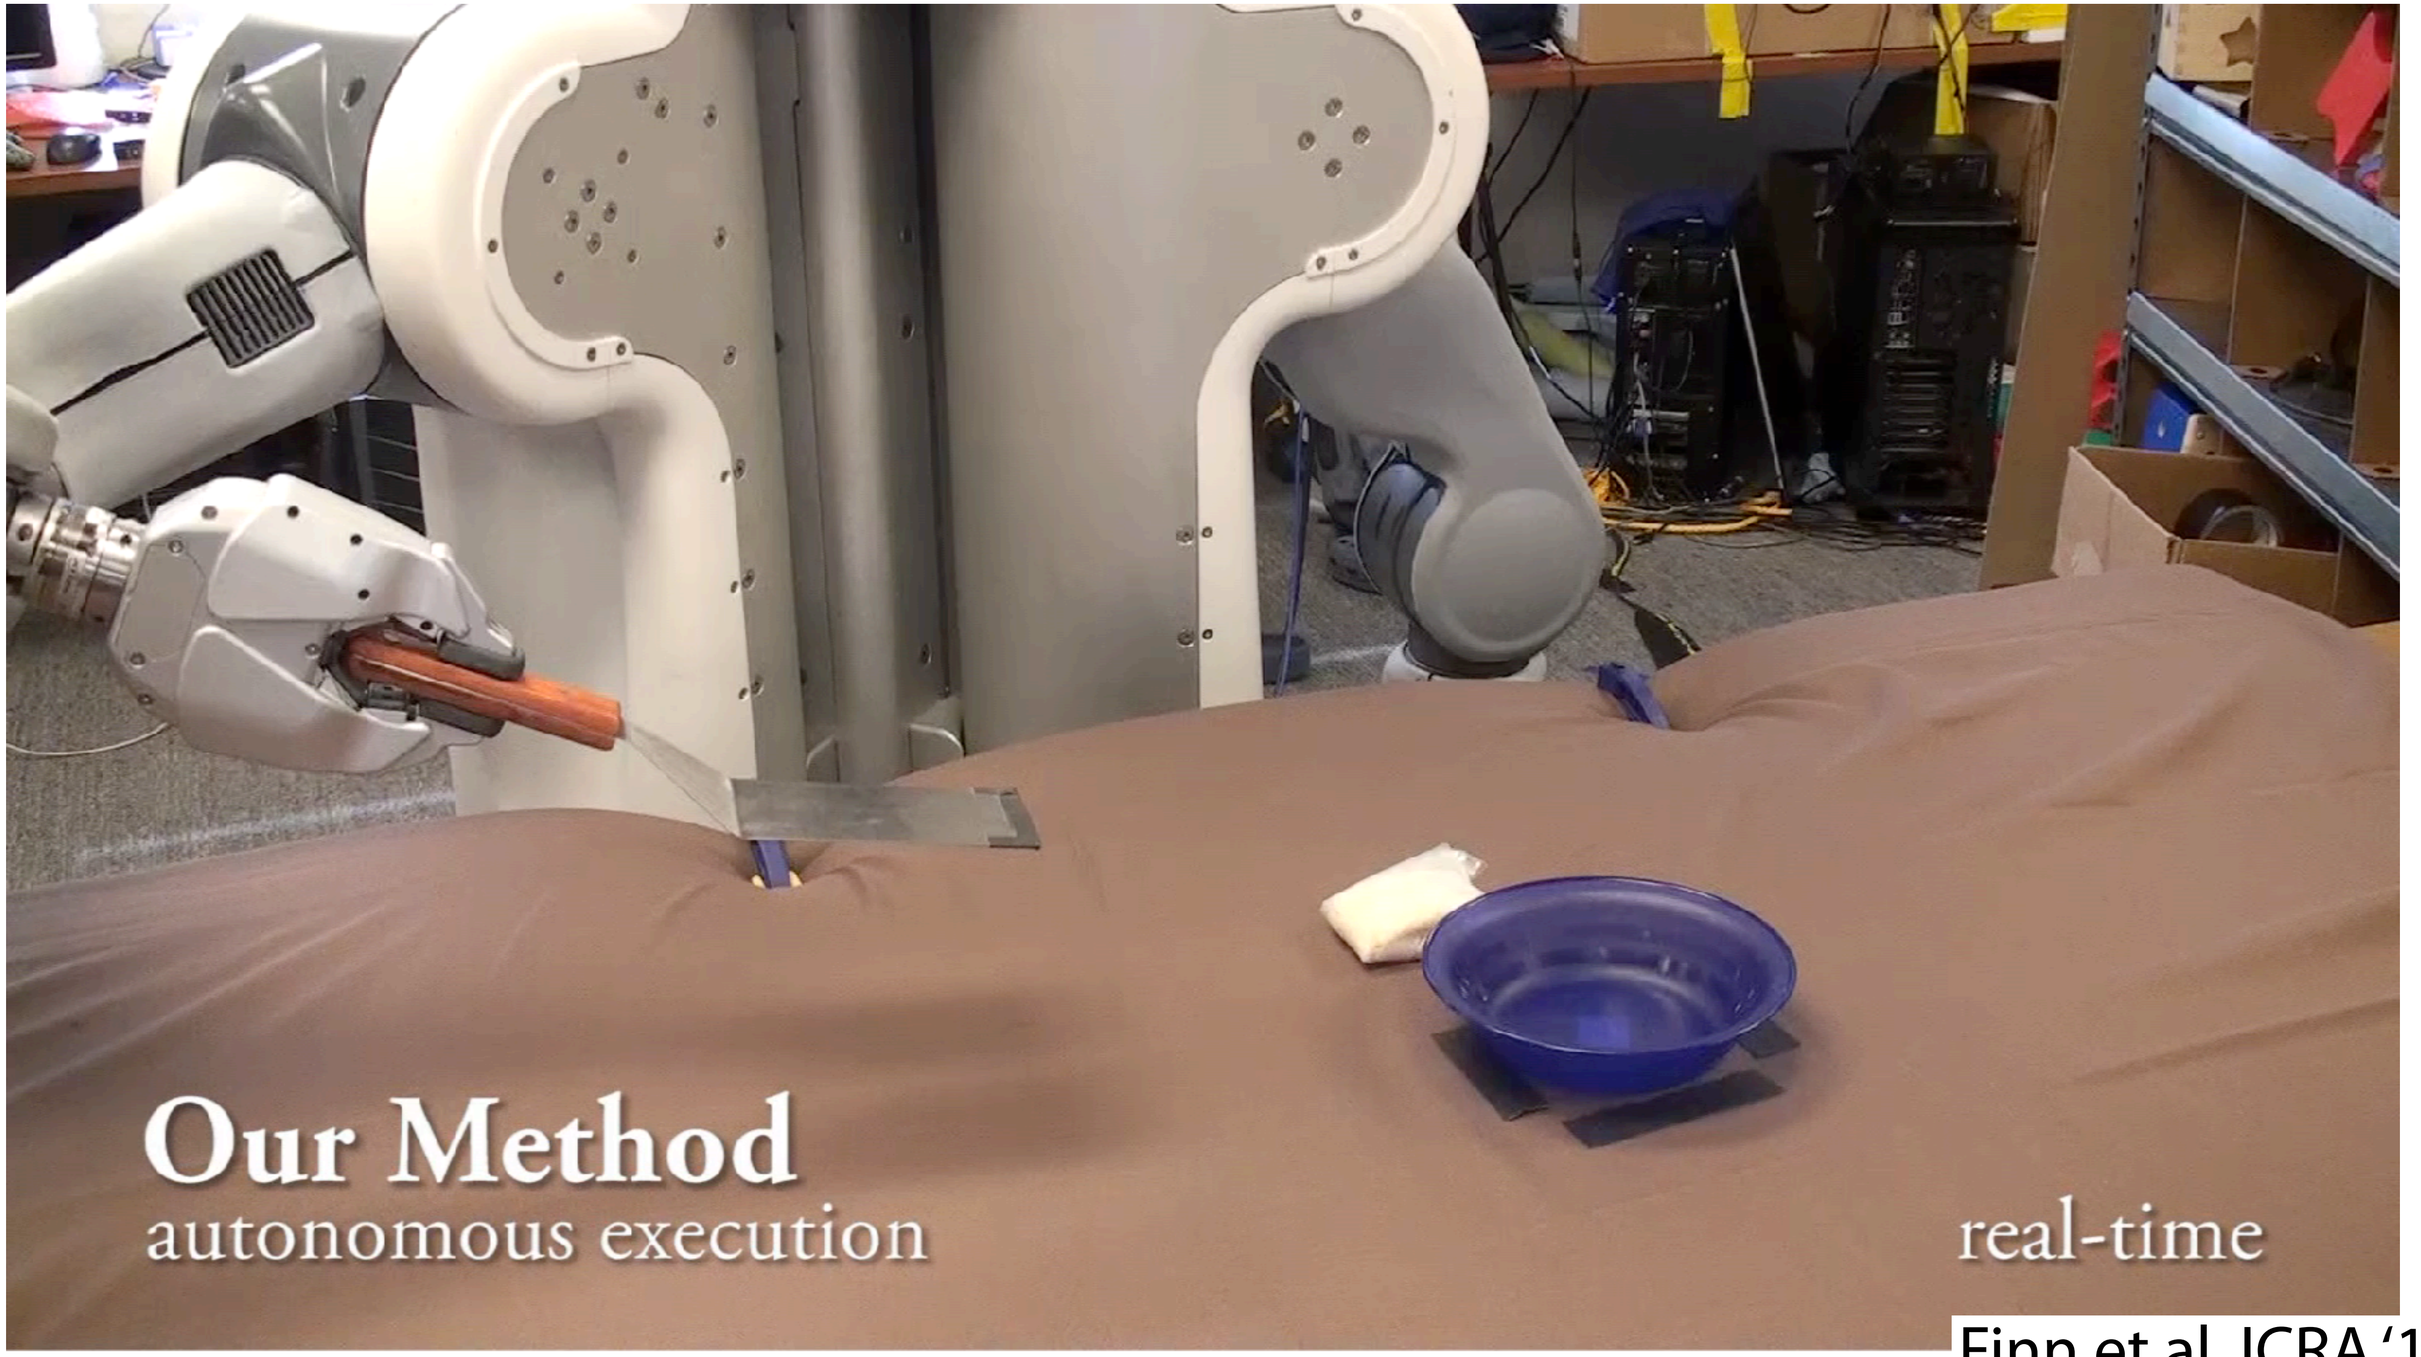

Finn et al. ICRA '16

Robot reinforcement learning

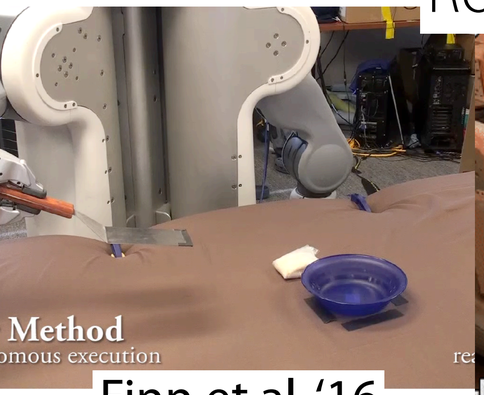

Finn et al. '16

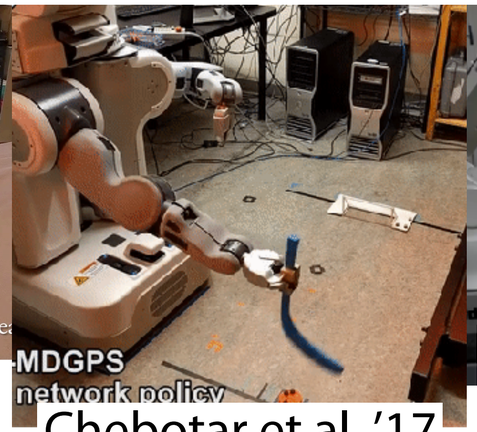

Chebotar et al. '17

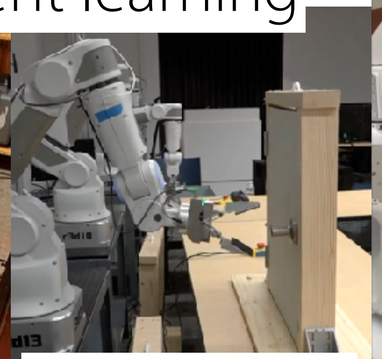

Yahya et al. '17

Ghadirzadeh et al. '17

Reinforcement learning

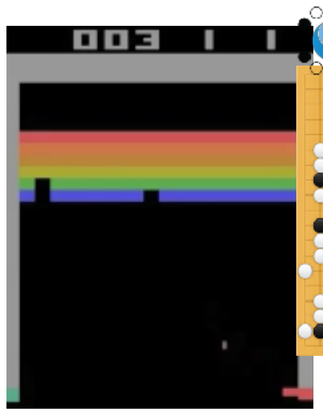

Atari

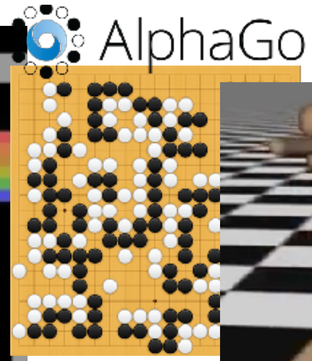

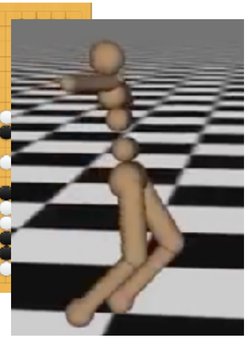

locomotion

Learn **one task** in **one environment**, **starting from scratch**

Behind the scenes…

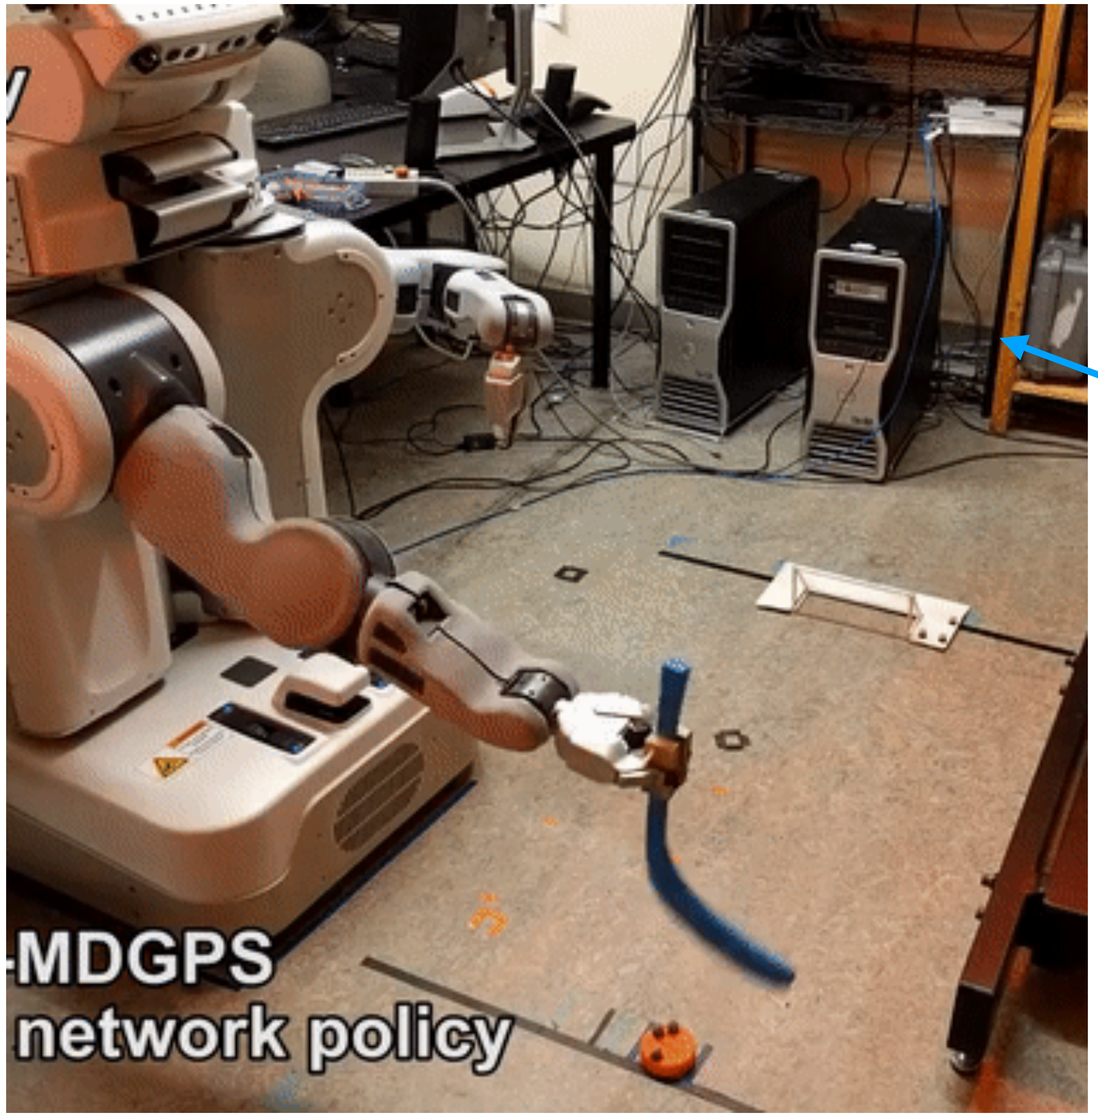

Yevgen

Yevgen is doing more work than the robot!

It's not practical to collect a lot of data this way.

Robot reinforcement learning

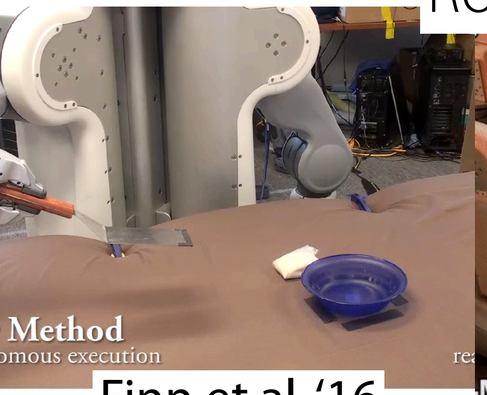

Finn et al. '16

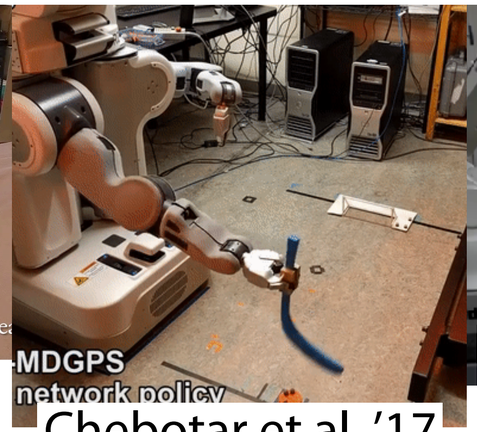

Chebotar et al. '17

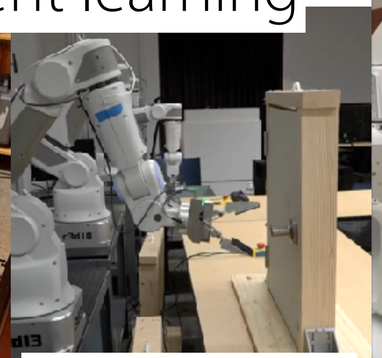

Yahya et al. '17

Ghadirzadeh et al. '17

Reinforcement learning

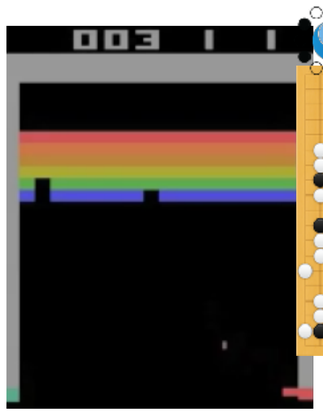

Atari

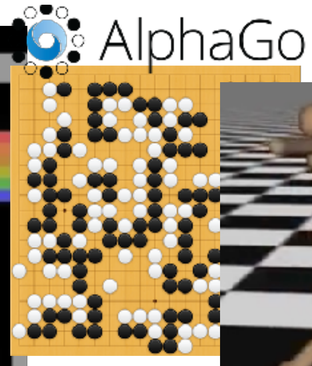

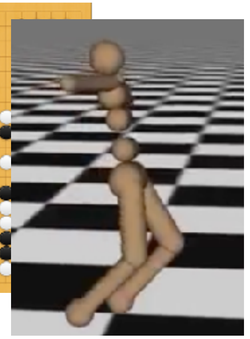

locomotion

Learn **one task** in **one environment**, **starting from scratch**
rely on **detailed supervision and guidance**.

Not just a problem with reinforcement learning & robotics.

specialists

[single task]

machine translation

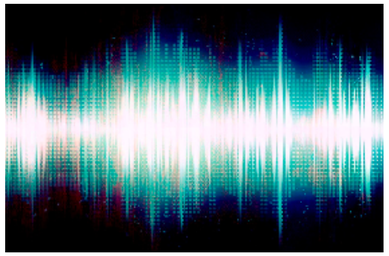

speech recognition

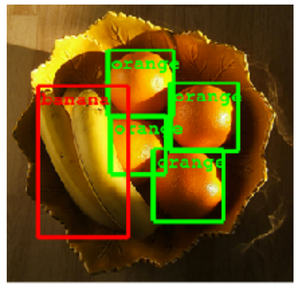

object detection

More diverse, yet still **one task**, **from scratch**, with **detailed supervision**

Source: https://youtu.be/8vNxjwt2AqY

vs.

Source: https://i.imgur.com/hJlVfZ5.jpg

Why should we care about multi-task & meta-learning?

…beyond the robots and general-purpose ML systems

deep
V
Why should we care about multi-task & meta-learning?

…beyond the robots and general-purpose ML systems

Standard computer vision:
hand-designed features

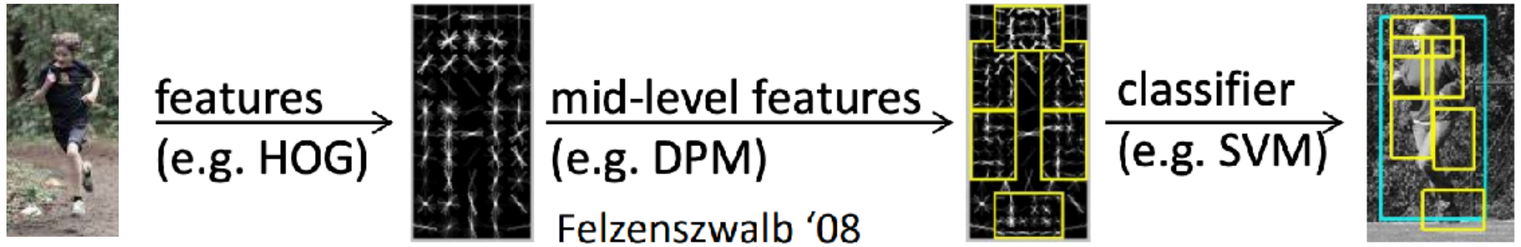

Modern computer vision:
end-to-end training

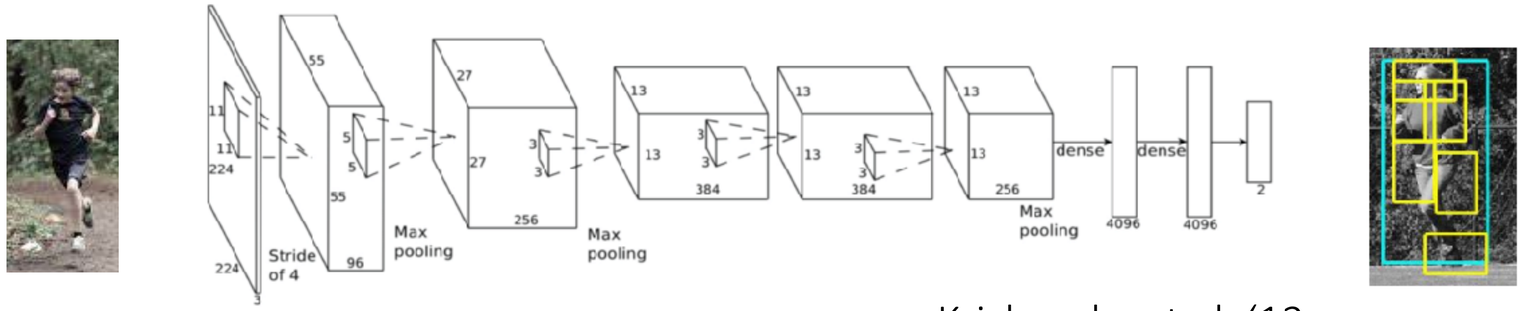

Krizhevsky et al. '12

Deep learning allows us to handle *unstructured inputs* (pixels, language, sensor readings, etc.) without hand-engineering features, with less domain knowledge

Slide adapted from Sergey Levine

Deep learning for object classification

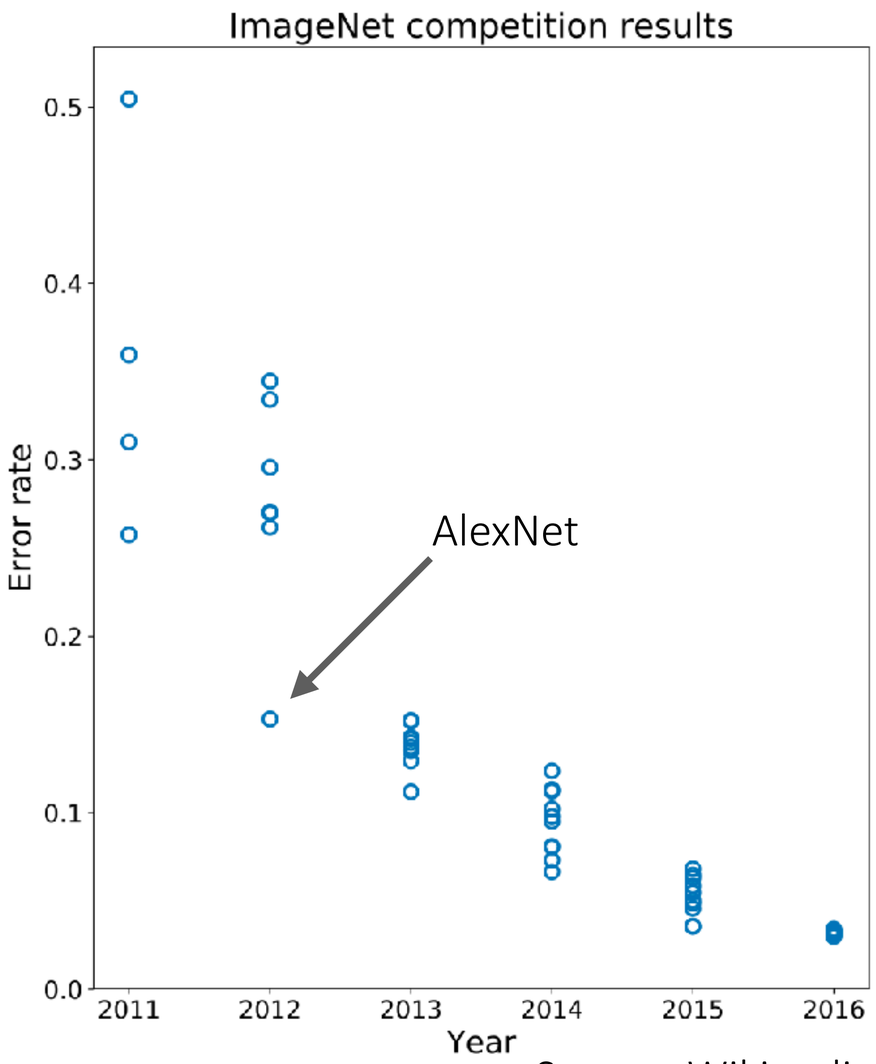

Source: Wikipedia

Deep learning for machine translation

Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi
yonghui,schuster,zhifengc,qvl,mnorouzi@google.com

Table 10: Mean of side-by-side scores on production data

	PBMT	GNMT	Human	Relative Improvement
English → Spanish	4.885	5.428	5.504	87%
English → French	4.932	5.295	5.496	64%
English → Chinese	4.035	4.594	4.987	58%
Spanish → English	4.872	5.187	5.372	63%
French → English	5.046	5.343	5.404	83%
Chinese → English	3.694	4.263	4.636	60%

Human evaluation scores on scale of 0 to 6

PBMT: Phrase-based machine translation

GNMT: Google's neural machine translation (in 2016)

Why deep **multi-task** and **meta-learning**?

Large, diverse data (+ large models) $\xrightarrow{\text{deep learning}}$ Broad generalization

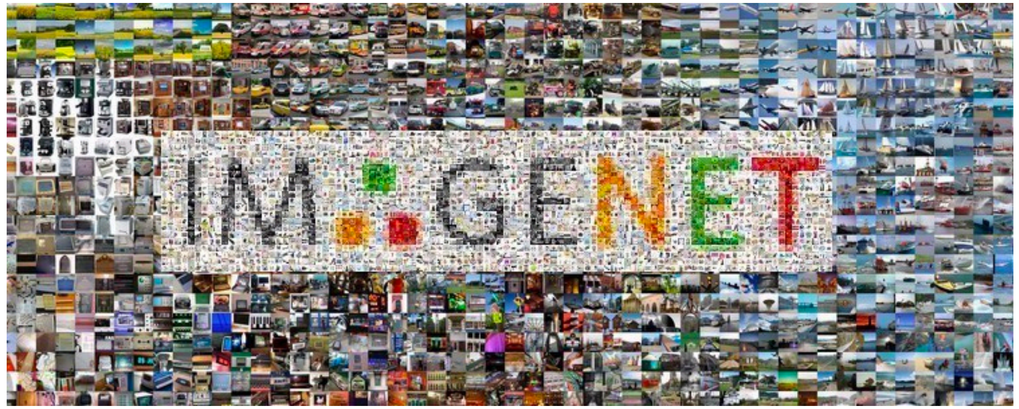

Russakovsky et al. ‘14

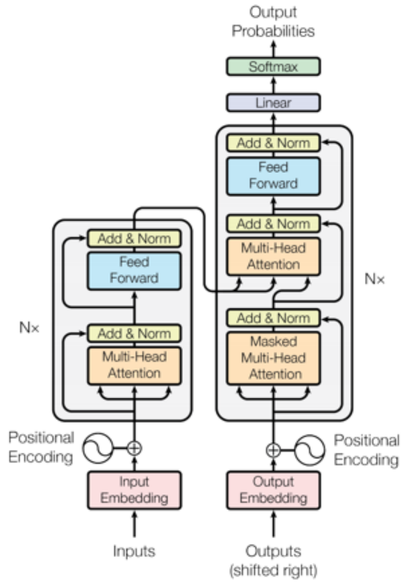

Figure 1: The Transformer - model architecture.

Vaswani et al. ‘18

GPT-2

Radford et al. ‘19

What if you don’t have a large dataset?

medical imaging robotics personalized education,

tr ations

Impractical to learn from scratch for each disease, each robot, each person, each language, each **task**

What if your data has a long tail?

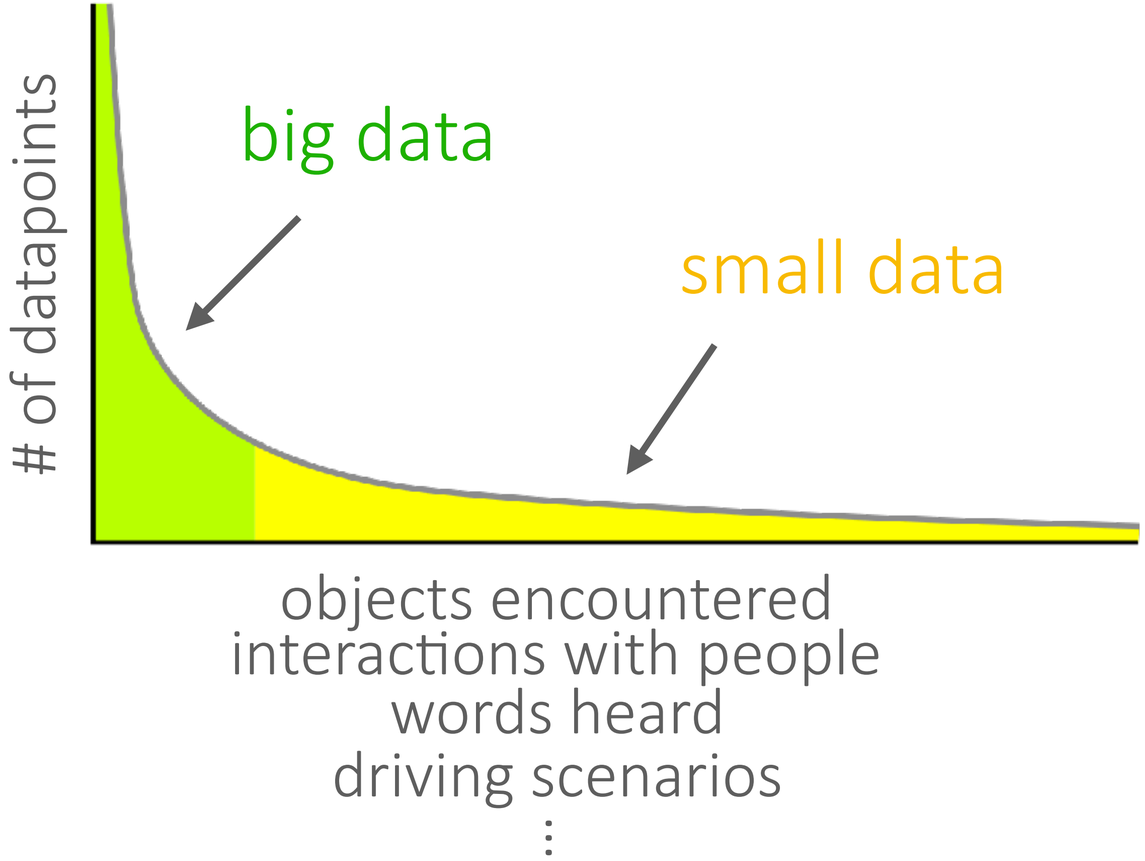

This setting breaks standard machine learning paradigms.

What if you need to quickly learn something new?

about a new person, for a new task, about a new environment, etc.

training data

Braque

Cezanne

test datapoint

By Braque or Cezanne?

What if you need to quickly learn something new?

about a new person, for a new task, about a new environment, etc.

“few-shot learning”

How did you accomplish this?

by leveraging prior experience!

What if you want a more general-purpose AI system?

Learning each task from scratch won't cut it.

What if you don't have a large dataset?

medical imaging

robotics

personalized education, medicine, recommendations

translation for rare languages

What if your data has a long tail?

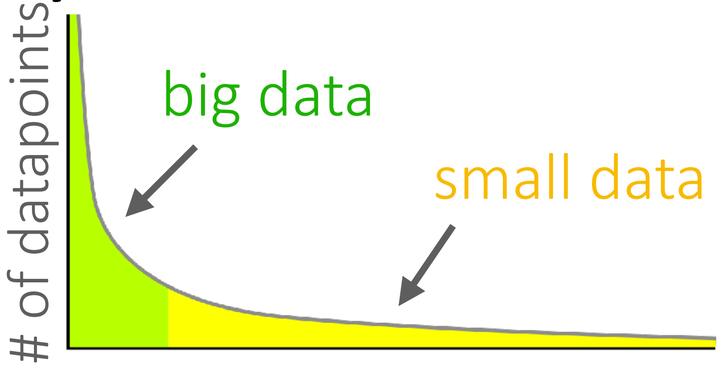

What if you need to quickly learn something new?

about a new person, for a new task, about a new environment, etc.

This is where elements of multi-task learning can come into play.

What is a task?

What is a task?

For now: dataset $\mathcal{D}$ loss function $\mathcal{L}$ $\longrightarrow$ model f_θ

Different tasks can vary based on:

- different objects
- different people
- different objectives
- different lighting conditions
- different words
- different languages
- …

Not *just* different "tasks"

Critical Assumption

The bad news: Different tasks need to share some structure.

If this doesn't hold, you are better off using single-task learning.

The good news: There are many tasks with shared structure!



Even if the tasks are seemingly unrelated:

- The **laws of physics** underly real data.
- People are **all organisms with intentions**.
- The **rules of English** underly English language data.
- Languages all develop for **similar purposes**.

This leads to far greater structure than random tasks.

Informal Problem Definitions

We'll define these more formally next time.

The multi-task learning problem: Learn all of the tasks more quickly or more proficiently than learning them independently.

The meta-learning problem: Given data/experience on previous tasks, learn a new task more quickly and/or more proficiently.

This course: anything that solves these problem statements.

Doesn't multi-task learning reduce to single-task learning?

$$\mathcal{D} = \bigcup \mathcal{D}_i \qquad \mathcal{L} = \sum \mathcal{L}_i$$

Are we done with the course?

Doesn’t multi-task learning reduce to single-task learning?

Yes, it can!

Aggregating the data across tasks & learning a single model is one approach to multi-task learning.

But, we can often do better!

Exploit the fact that we know that data is coming from different tasks.

Why now?

Why should we study deep multi-task & meta-learning now?

Multitask Learning*

RICH CARUANA

Multitask Learning (MTL) is an inductive transfer mechanism whose principle goal is to improve generalization performance. MTL improves generalization by leveraging the domain-specific information contained in the training signals of *related* tasks. It does this by training tasks in parallel while using a shared representation. In effect, the training signals for the extra tasks serve as an inductive bias. Section 1.2 argues that inductive transfer is important if we wish to scale tabula rasa learning to complex, real-world tasks. Section 1.3 presents the simplest method we know for doing multitask inductive transfer, adding extra tasks (i.e., extra outputs) to a backpropagation net. Because the MTL net uses a shared hidden layer trained in parallel on all the tasks, what is learned for each task can help other tasks be learned better. Section 1.4 argues that it is reasonable to view training signals as an inductive bias when they are used this way.

Caruana, 1997

Is Learning The n-th Thing Any Easier Than Learning The First?

Sebastian Thrun[1]

They are often able to generalize correctly even from a single training example [2, 10]. One of the key aspects of the learning problem faced by humans, which differs from the vast majority of problems studied in the field of neural network learning, is the fact that humans encounter a whole stream of learning problems over their entire lifetime. When faced with a new thing to learn, humans can usually exploit an enormous amount of training data and experiences that stem from other, related learning tasks. For example, when learning to drive a car, years of learning experience with basic motor skills, typical traffic patterns, logical reasoning, language and much more precede and influence this learning task. The transfer of knowledge across learning tasks seems to play an essential role for generalizing accurately, particularly when training data is scarce.

Thrun, 1998

On the Optimization of a Synaptic Learning Rule

Samy Bengio Yoshua Bengio Jocelyn Cloutier Jan Gecsei

Université de Montréal, Département IRO

This paper presents a new approach to neural modeling based on the idea of using an automated method to optimize the parameters of a synaptic learning rule. The synaptic modification rule is considered as a parametric function. This function has *local* inputs and is the same in many neurons. We can use standard optimization methods to select appropriate parameters for a given type of task. We also present a theoretical analysis permitting to study the *generalization* property of such parametric learning rules. By generalization, we mean the possibility for the learning rule to learn to solve *new* tasks. Experiments were performed on three types of problems: a

Bengio et al. 1992

These algorithms are continuing to play a fundamental role in machine learning research.

Multilingual machine translation

Massively Multilingual Neural Machine Translation

Roee Aharoni*
Bar Ilan University
Ramat-Gan
Israel
roee.aharoni@gmail.com

Melvin Johnson and **Orhan Firat**
Google AI
Mountain View
California
melvinp,orhanf@google.com

while supporting up to 59 languages. Our experiments on a large-scale dataset with 102 languages to and from English and up to one million examples per direction also show promising results, surpassing strong bilingual baselines and encouraging future work on massively multilingual NMT.

2019

One-shot imitation learning from humans

DAML Yu et al. RSS 2018

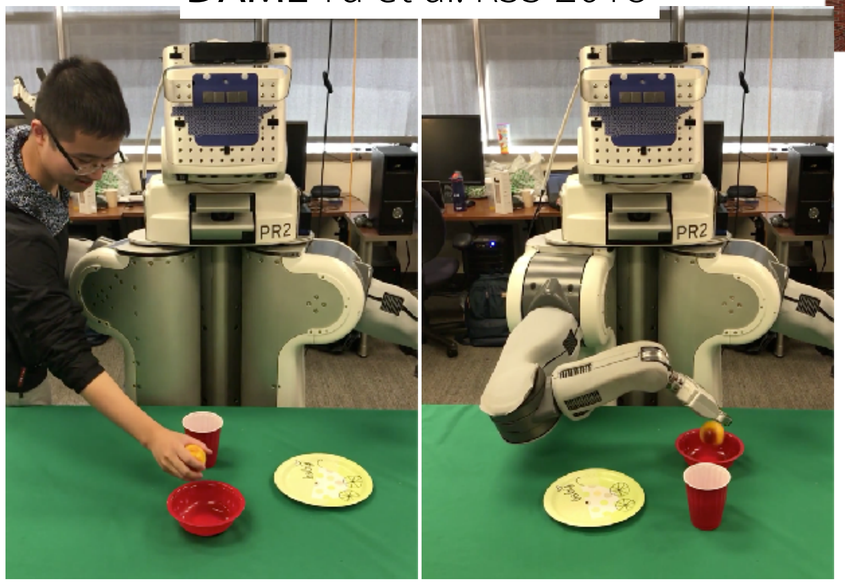

Multi-domain learning for sim2real transfer

CAD²RL Sadeghi & Levine, 2016

YouTube recommendations

Recommending What Video to Watch Next: A Multitask Ranking System

Zhe Zhao, Lichan Hong, Li Wei, Jilin Chen, Aniruddh Nath, Shawn Andrews, Aditee Kumthekar, Maheswaran Sathiamoorthy, Xinyang Yi, Ed Chi
Google, Inc.
{zhezhao,lichan,liwei,jilinc,aniruddhnath,shawnandrews,aditeek,nlogn,xinyang,edchi}@google.com

In this paper, we introduce a large scale multi-objective ranking system for recommending what video to watch next on an industrial video sharing platform. The system faces many real-world challenges, including the presence of multiple competing ranking objectives, as well as implicit selection biases in user feedback. To

2019

These algorithms are playing a fundamental, and **increasing** role in machine learning research.

Interest level via search queries

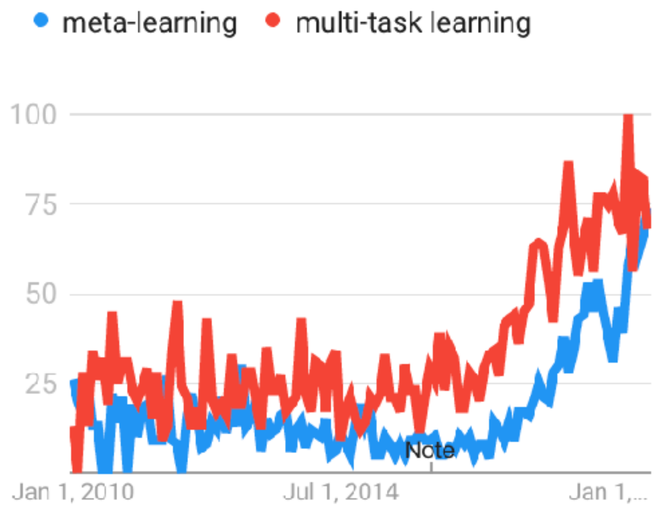

How transferable are features in a deep neural network?
Yosinski et al. '15

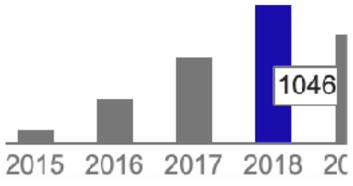

Learning to learn by gradient descent by gradient descent
Andrychowicz et al. '15

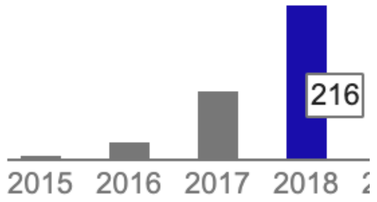

Model-agnostic meta-learning for fast adaptation of deep networks
Finn et al. '17

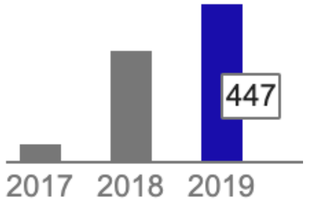

An overview of multi-task learning in neural networks
Ruder '17

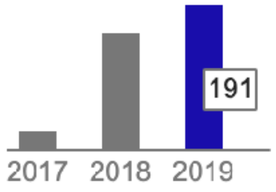

Graph sources: Google scholar, Google trends

Its success will be critical for the **democratization** of deep learning.

ImageNet	1.2 million images and labels
WMT ’14 English - French	40.8 million paired sentences
Switchboard Speech Dataset	300 hours of labeled data

Kaggle’s Diabetic Retinopathy Detection dataset	35K labeled images
Adaptive epilepsy treatment with RL Guez et al. ‘08	< 1 hour of data
Learning for robotic manipulation Finn et al. ‘16	< 15 min of data

But, we still have many open questions and challenges!

Reminder: Homework Today

1. Sign up for Piazza
2. Fill out paper presentation preferences **(by Thursday!)**
3. Start forming final project groups if you want to work in a group
4. Review this: https://www.tensorflow.org/guide/low_level_intro

Next time (Weds): Multi-Task and Meta-Learning Basics