CS391R: Robot Learning

Conclusion: Open Questions in Robot Learning

Prof. Yuke Zhu

Fall 2020

Today's Agenda

- General-Purpose Robot Autonomy (GPRA)
    - Review of the key concepts
    - Computing paradigms of the perception-action loop
- Summary of knowledge: hammers (techniques) versus nails (problems)
- Open research questions
- Progressive roadmap towards GPRA
- Societal impacts of Robotics + AI

General-Purpose Robot Autonomy … in the Wild

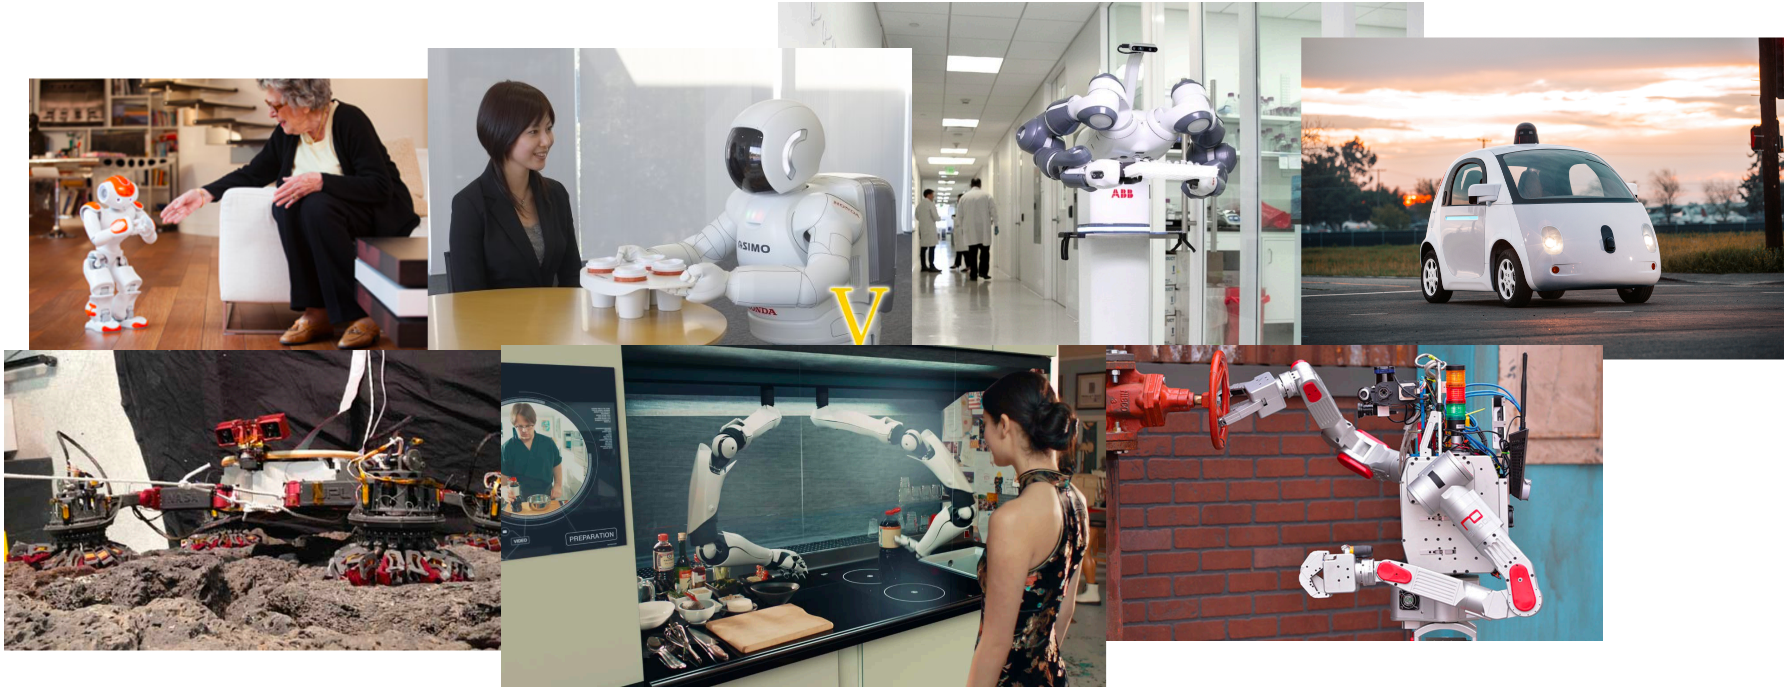

Unstructured Environments

Ever-changing Tasks

Human Involvement

Special-Purpose Robot Automation

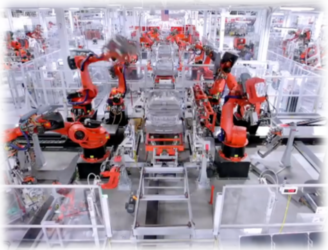

custom-built
robots

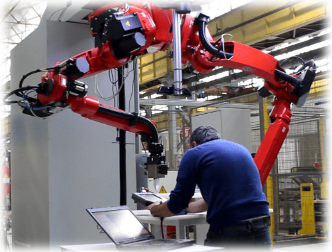

human expert
programming

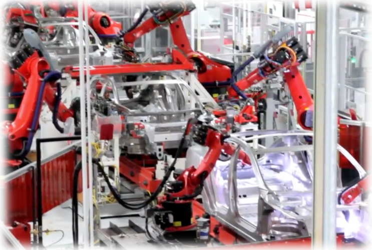

special-purpose
behaviors

General-Purpose Robot Autonomy

general-purpose
robots

Robot Learning

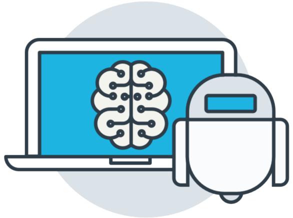

general-purpose
behaviors

A key challenge in **Robot Learning** is to close the **Perception-Action Loop**.

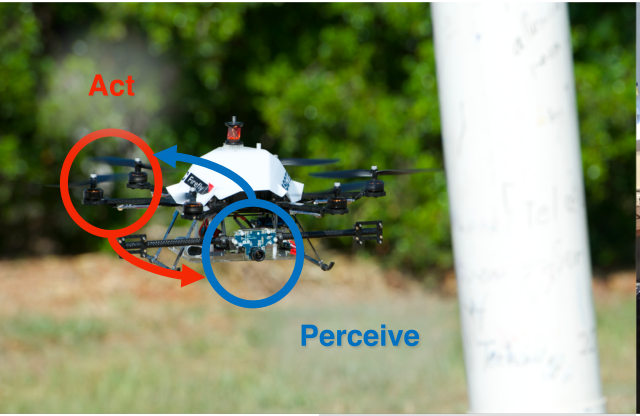

[Sa et al. IROS 2014]

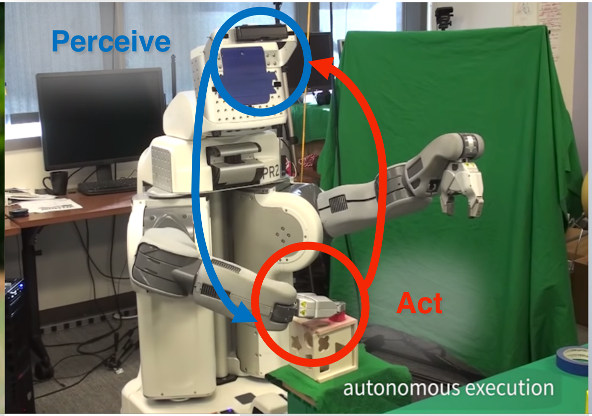

[Levine et al. JMLR 2016]

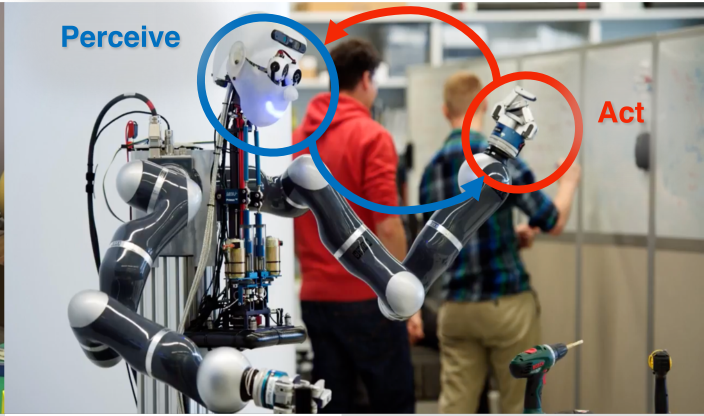

[Bohg et al. ICRA 2018]

Close the **Perception-Action Loop**: A New Paradigm

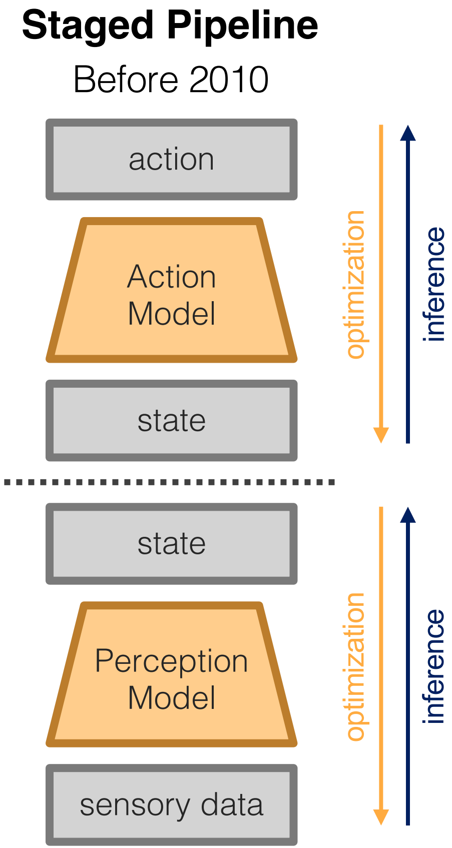

Close the **Perception-Action Loop**: A New Paradigm

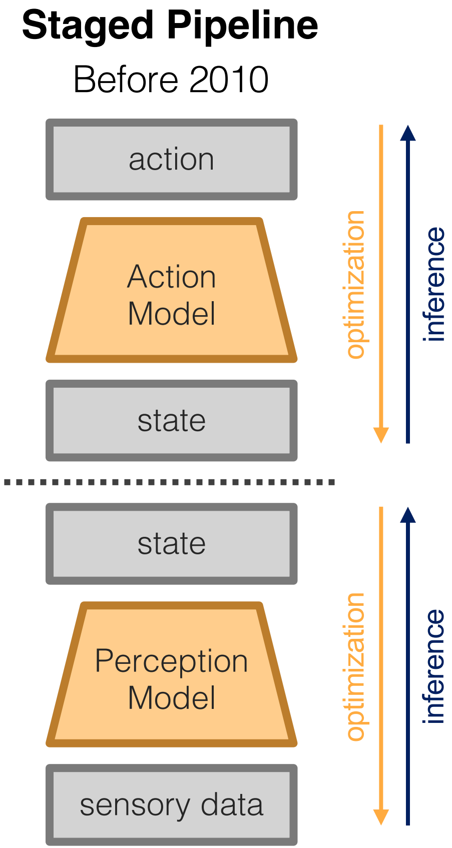

Noisy Sensory Data

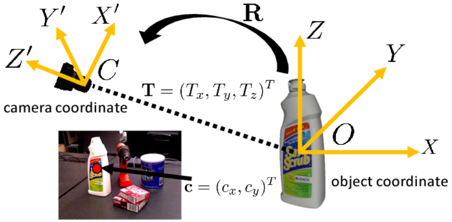

Physical State

Perception & Computer Vision

Close the **Perception-Action Loop**: A New Paradigm

Staged Pipeline

Before 2010

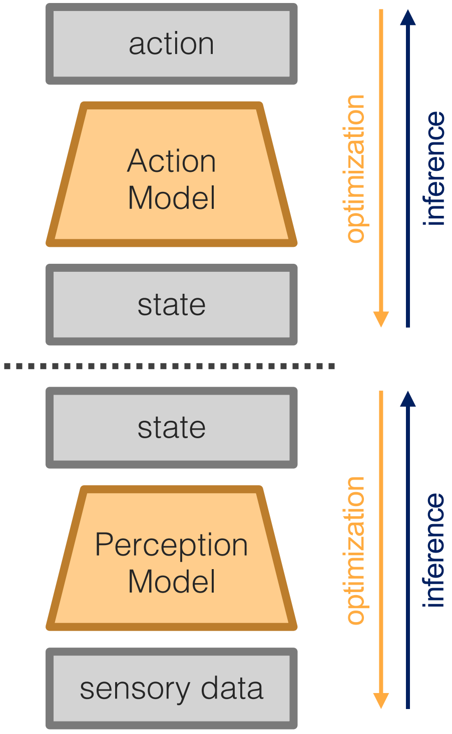

Noisy Sensory Data

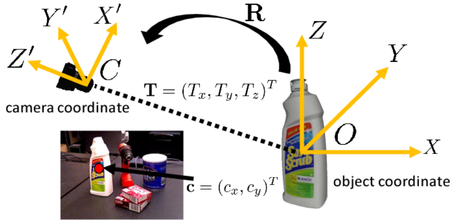

Physical State

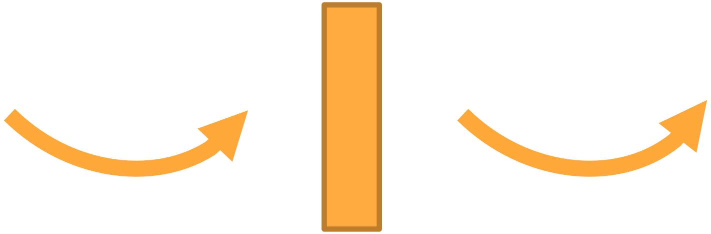

Perception & Computer Vision

Robot Control & Decision Making

Close the **Perception-Action Loop**: A New Paradigm

Staged Pipeline

Before 2010

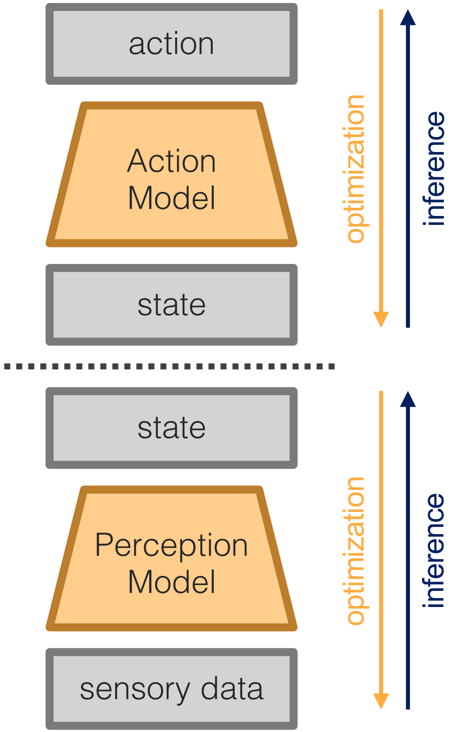

End-to-End Learning

2010 - 2020

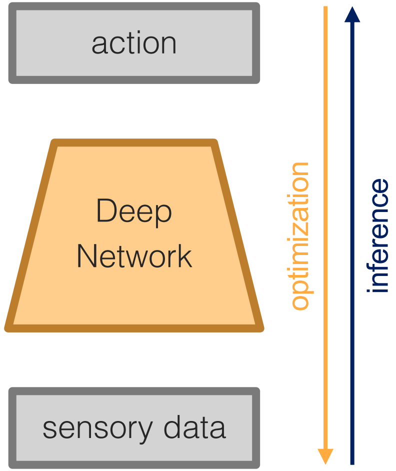

Close the **Perception-Action Loop**: A New Paradigm

Staged Pipeline

Before 2010

action

Action Model

state

---

state

Perception Model

sensory data

optimization ↓ inference ↑

optimization ↓ inference ↑

End-to-End Learning

2010 - 2020

action

Deep Network

sensory data

optimization ↓ inference ↑

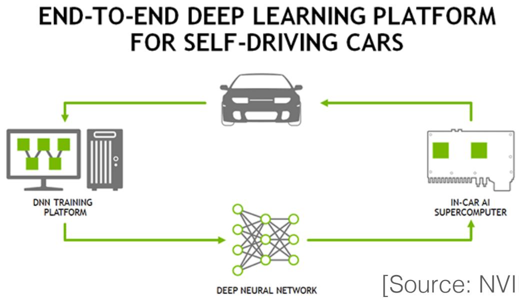

[Source: NVIDIA]

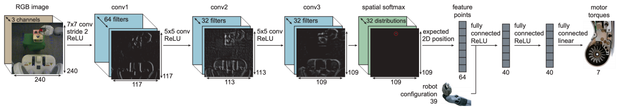

[Source: Levine et al. JMLR 2016]

Close the **Perception-Action Loop**: A New Paradigm

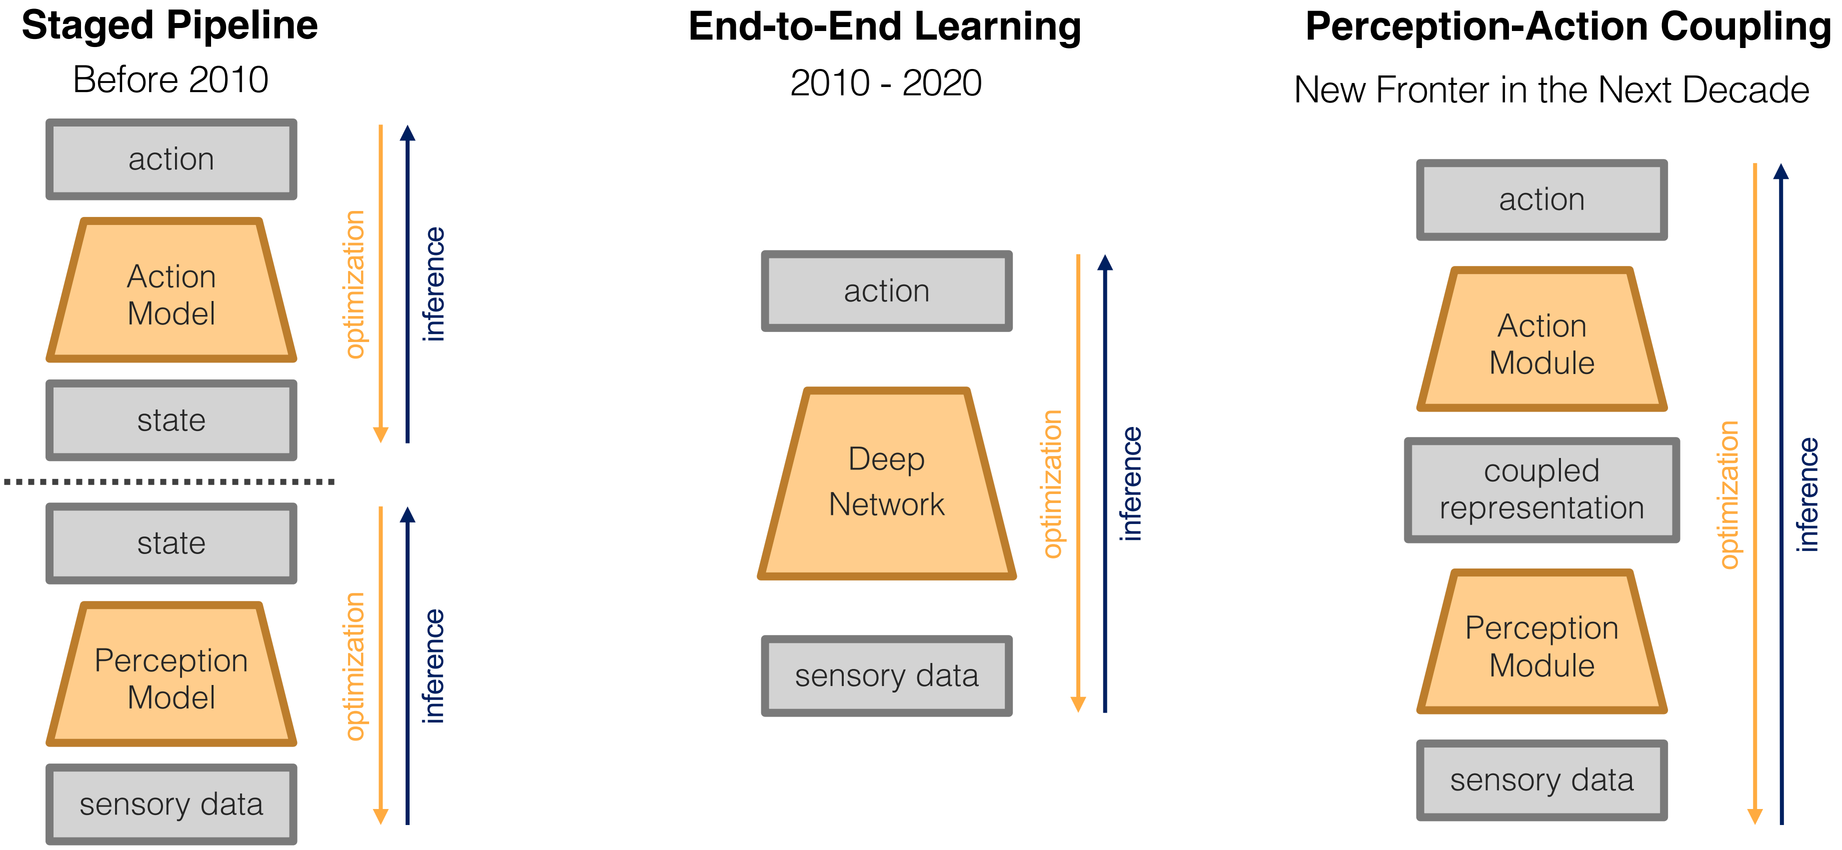

New Paradigm: **Perception-Action Coupling**

Perception-Action Coupling

New Fronter in the Next Decade

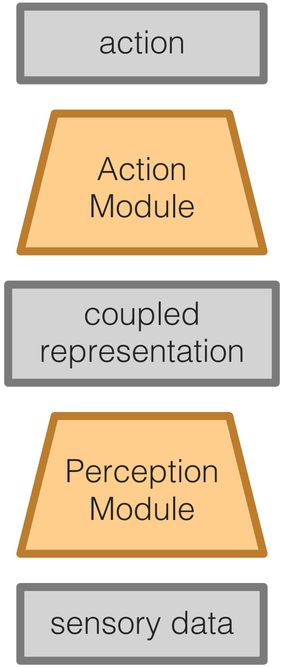

Rich inductive biases from model structures

Learning action-informed perceptual representation

Joint optimization of functional modules (Software 2.0)

Software 2.0: https://medium.com/@karpathy/software-2-0-a64152b37c35

New Paradigm: **Perception-Action Coupling**

Perception-Action Coupling

New Fronter in the Next Decade

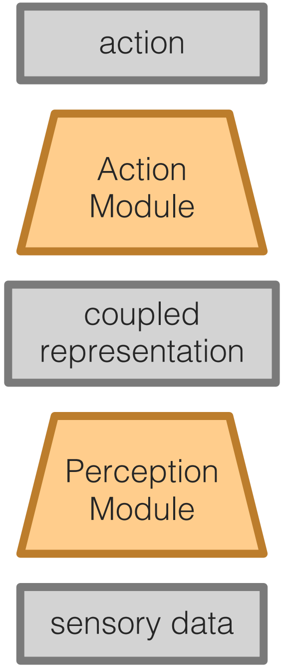

My Million-Dollar Question: “1 + 1 < 2?”

New Paradigm: **Perception-Action Coupling**

Perception-Action Coupling

New Fronter in the Next Decade

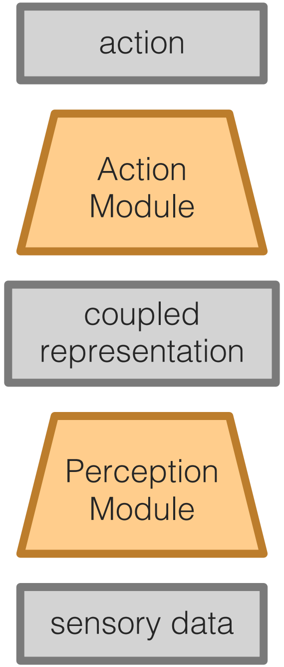

My Million-Dollar Question: “1 + 1 < 2?”

“Will joint optimization make the computational problem of perception and decision-making (towards GPRA) easier?”


BONGU
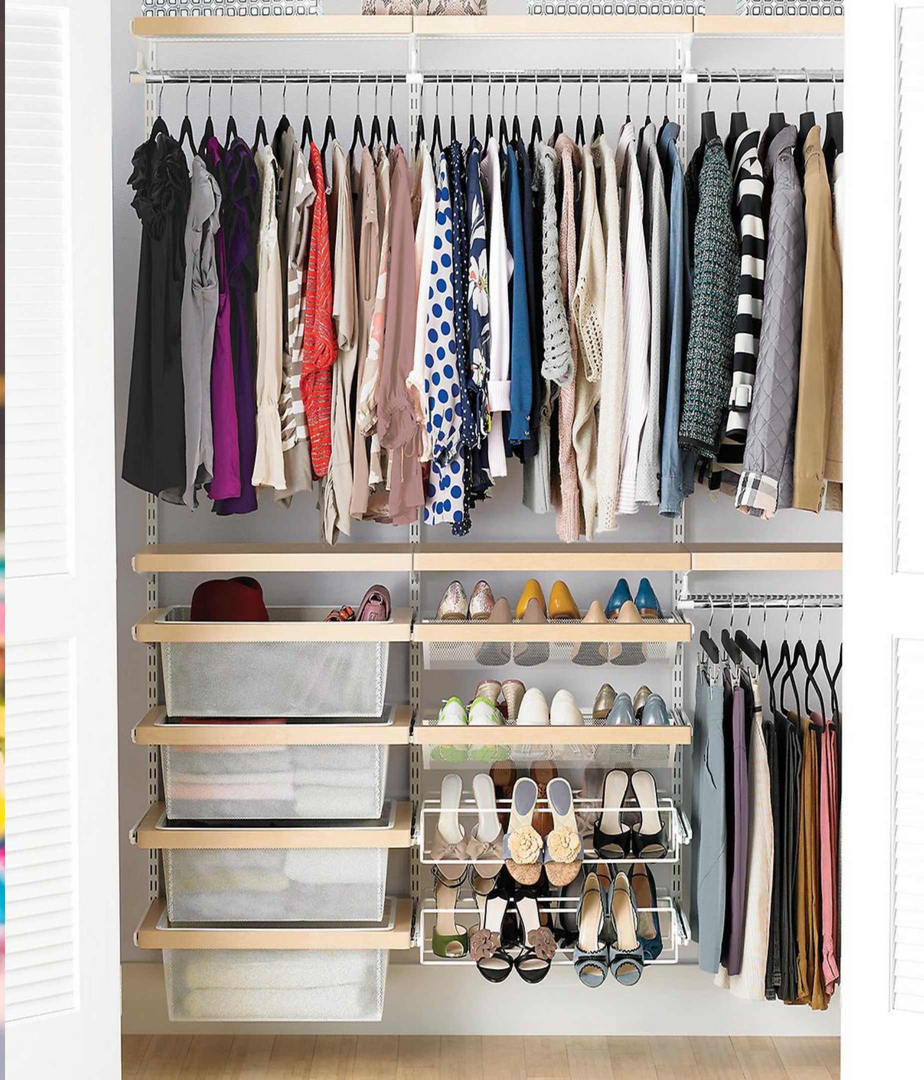

New Paradigm: **Perception-Action Coupling**

Part II: Robot Decision Making

Week 7 Tue, Oct 6 — Lecture Overview of Robot Decision Making [slides]

- Reinforcement Learning in Robotics: A Survey. Jens Kober, J. Andrew Bagnell, Jan Peters (2013)
  - Recent Advances in Robot Learning from Demonstration. Harish Ravichandar, Athanasios S. Polydoros, Sonia Chernova, Aude Billard (2020)

Part I: Robot Perception

Week 2 Tue, Sept 1 — Lecture Overview of Robot Perception [slides]

- The Limits and Potentials of Deep Learning for Robotics. Niko Sünderhauf, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, Pieter Abbeel, Wolfram Burgard, Michael Milford, Peter Corke (2018)
  - A Sensorimotor Account of Vision and Visual Consciousness. Kevin O'Regan and Alva Noë (2001)

Perception-Action Coupling

New Fronter in the Next Decade

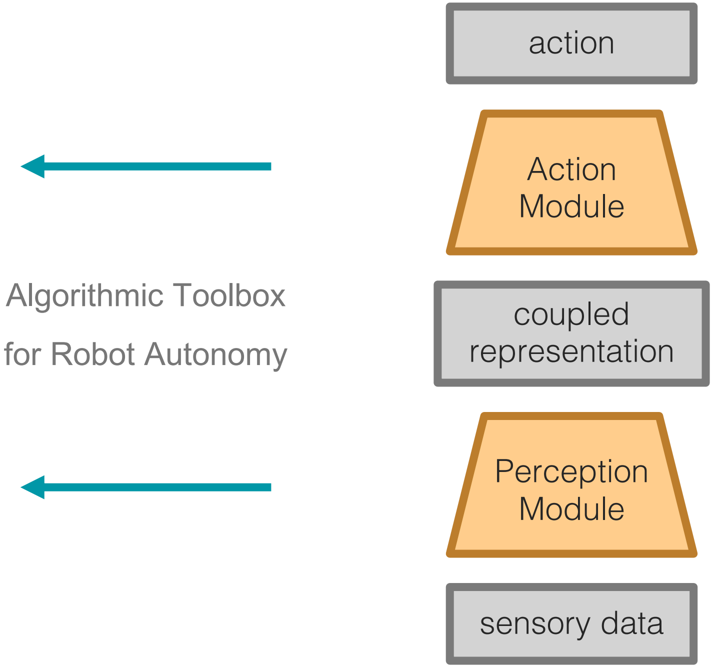

Ch 1: Perception

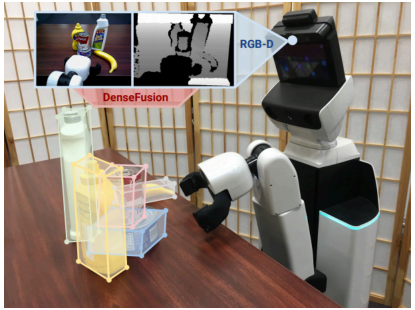

- Convolutional networks
- PointNet / PointNet++
- Contrastive learning
- Unsupervised learning
- Predictive coding
- Implicit representation
- Bayes filtering

- 3D point cloud processing
- Pose estimation
- 2D/3D visual tracking
- Multimodal understanding
- Data association
- Concept discovery
- Visual navigation
- Recursive state estimation
- Active perception

Ch 2: Decision Making

- Model-free reinforcement learning
- Trust-region optimization
- Model-based dynamics learning
- Gaussian process
- Behavior cloning / DAgger
- Inverse reinforcement learning
- Adversarial imitation learning

- Sensorimotor learning
- Video prediction
- Reward/utility learning
- Learning from demonstration
- Autonomous driving

Ch 3: Intelligence

- Meta-learning
- Neural memory networks
- Hierarchical RL
- Neural programming induction
- Task-and-motion planning (TAMP)
- Causal reasoning
- Evolutionary computation
- Knowledge ontology

- Affordance reasoning
- One-shot imitation learning
- Long-horizon manipulation
- Open-ended curriculum learning
- Knowledge transfer across tasks

Ch 4: Real-World Systems

- Bayesian inference
- Domain randomization
- Multi-armed bandits
- Cross-entropy methods
- Big data in robotics

- System identification
- Semantic segmentation
- Open/closed-loop grasping
- Vision-based manipulation
- Quadruped locomotion

The Algorithmic Toolbox for Robot Autonomy

Open Research Questions: **Perception**

1. **Making sense of the unstructured world:** unified holistic scene representations of semantics, geometry, dynamics, and agents over time;

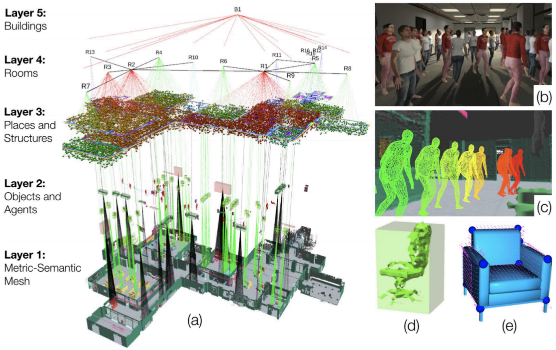

3D Dynamic Scene Graph [Rosinol et al. RSS'20]

Open Research Questions: **Decision Making**

1. **Making sense of the unstructured world:** unified holistic scene representations of semantics, geometry, dynamics, and agents over time;
2. **Learning with limited supervision and from rich data sources:** self supervision, natural language, visual demonstrations, human preferences, multimodality, web data, gaze, social interactions, etc.

Open Research Questions: General Intelligence

1. **Making sense of the unstructured world:** unified holistic scene representations of semantics, geometry, dynamics, and agents over time;
2. **Learning with limited supervision and from rich data sources:** self supervision, natural language, visual demonstrations, human preferences, multimodality, web data, gaze, social interactions, etc.
3. **Continual learning and compositional modeling of concepts:** never-ending learning of new concepts from self-directed explorations and modeling the compositionality of tasks and semantics;

Open Research Questions: **Real-World Systems**

1. **Making sense of the unstructured world:** unified holistic scene representations of semantics, geometry, dynamics, and agents over time;
2. **Learning with limited supervision and from rich data sources:** self supervision, natural language, visual demonstrations, human preferences, multimodality, web data, gaze, social interactions, etc.
3. **Continual learning and compositional modeling of concepts:** never-ending learning of new concepts from self-directed explorations and modeling the compositionality of tasks and semantics;
4. **Safety and robustness of real-world robotic systems:** simulation-to-reality gap, uncertainty quantification & safe learning, and trustworthy and verifiable AI systems.

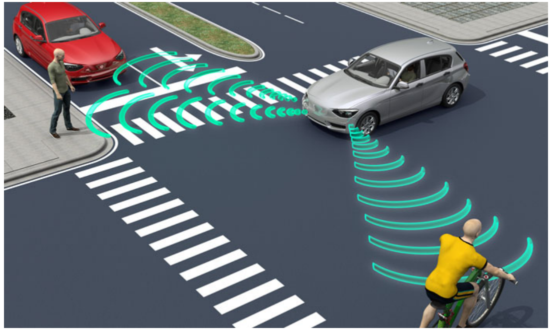

A Progressive Roadmap to General-Purpose Robot Autonomy

- learning **primitive sensorimotor repertoires** from raw perceptual input (Chapter 1, 2, 4)

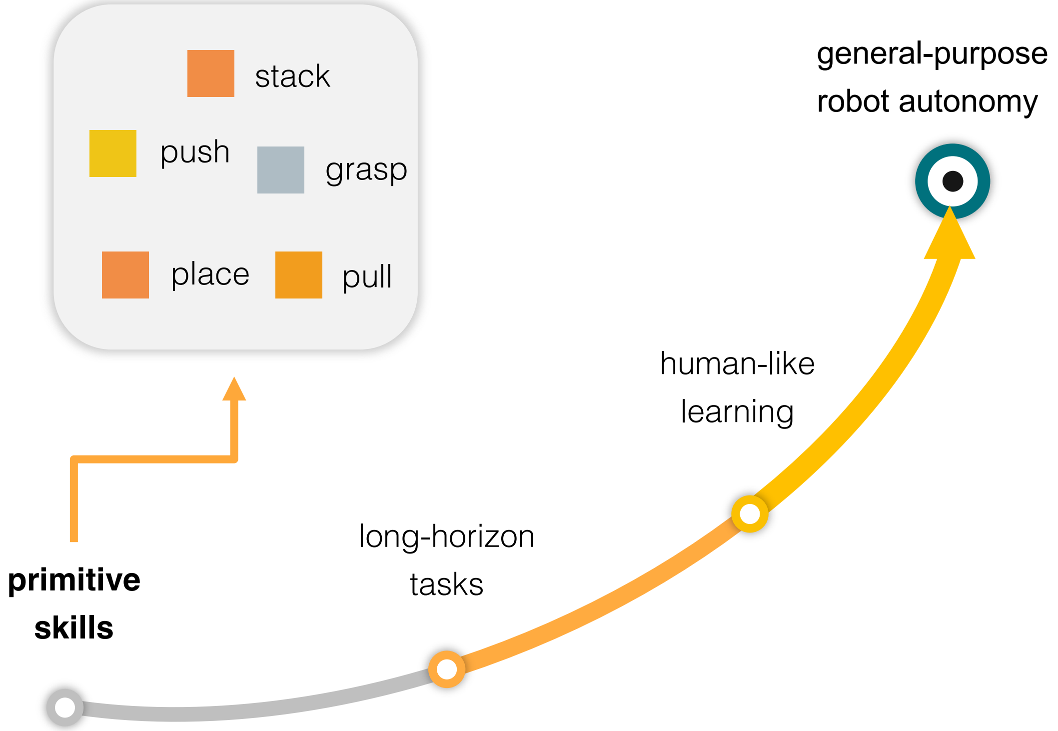

My talk "**Building General-Purpose Robot Autonomy: A Progressive Roadmap**" [video] [slides]

A Progressive Roadmap to General-Purpose Robot Autonomy

- learning **primitive sensorimotor repertoires** from raw perceptual input (Chapter 1, 2, 4)
- scaling to **long-horizon tasks** through **compositionality** and **abstraction** (Chapter 2, 3)

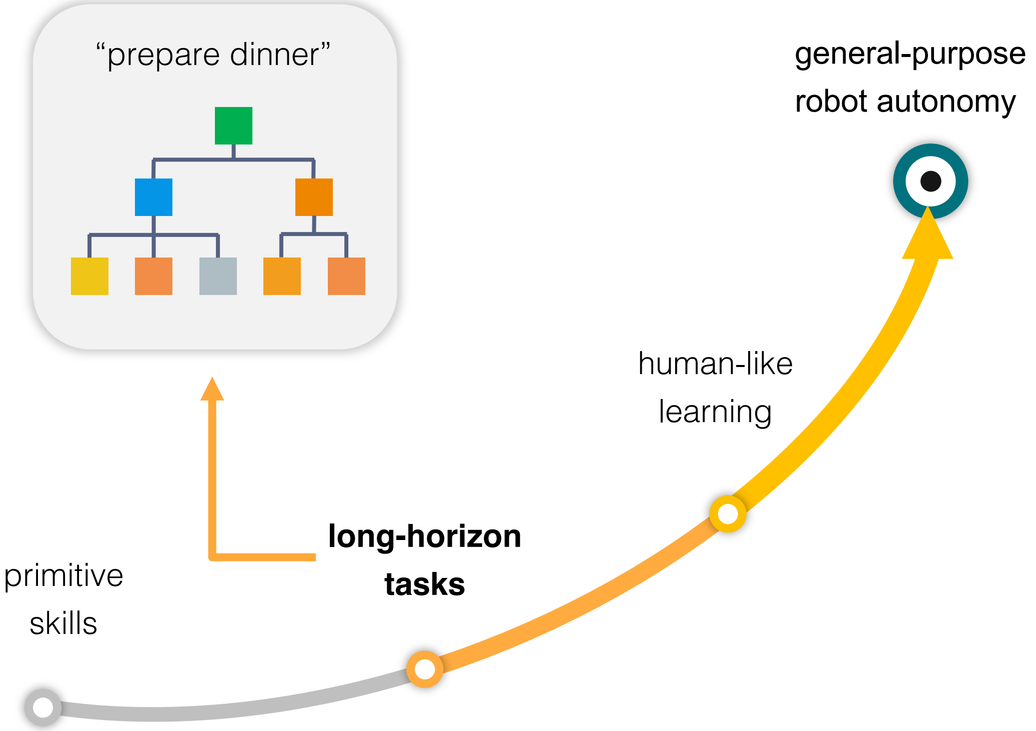

My talk "**Building General-Purpose Robot Autonomy: A Progressive Roadmap**" [video] [slides]

A Progressive Roadmap to General-Purpose Robot Autonomy

- learning **primitive sensorimotor repertoires** from raw perceptual input (Chapter 1, 2, 4)
- scaling to **long-horizon tasks** through **compositionality** and **abstraction** (Chapter 2, 3)
- **human-like learning** via **active exploration** and **model building** (Chapter 3)

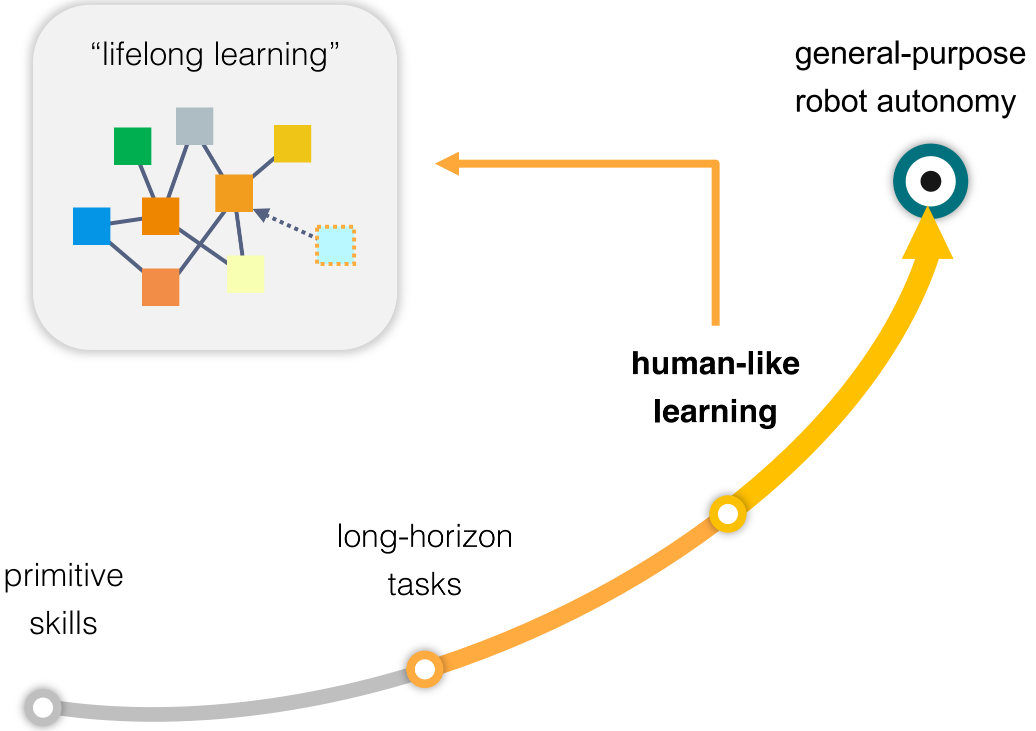

My talk "**Building General-Purpose Robot Autonomy: A Progressive Roadmap**" [video] [slides]

Robots and Society

Will intelligent robots lead to more jobs or less jobs?

More? Higher GDP per capita → More (service sector) jobs

Less? Robotics + AI is disruptive and general-purpose. “This time is different?”

“[Alaskan fishing ranked the most dangerous job in America]”

[Source: Daily Mail]

“[An early advertisement declaring the horse obsolete]”

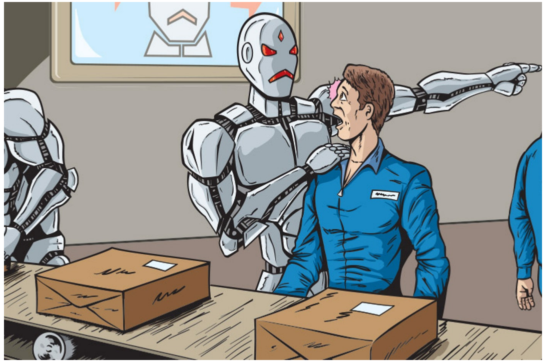

“[Neo-Luddism's Tech Skepticism]”

Question: What’s the value of work?

Robots and Society

Personal assistive household robots in the aging society

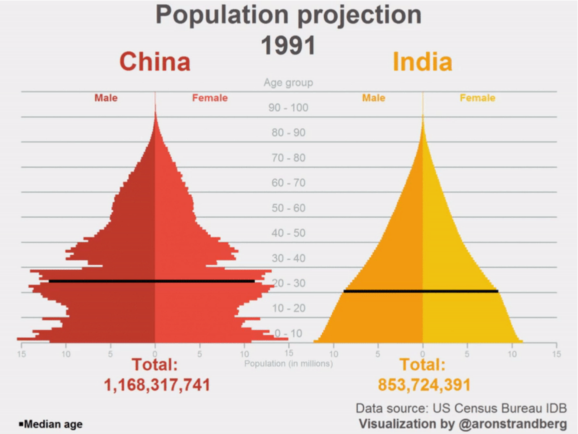

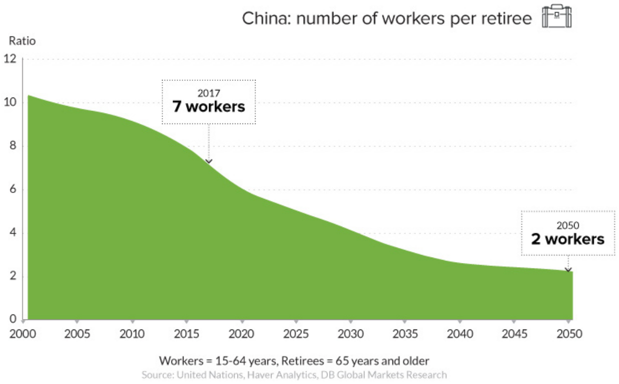

"[Robot carers for the elderly in Japan](#)"
[Source: The Times UK]

"By 2040, about one in five Americans will be age 65 or older, up from about one in eight in 2000." [source]

Robots and Society

Militarization of Robotics and AI technologies

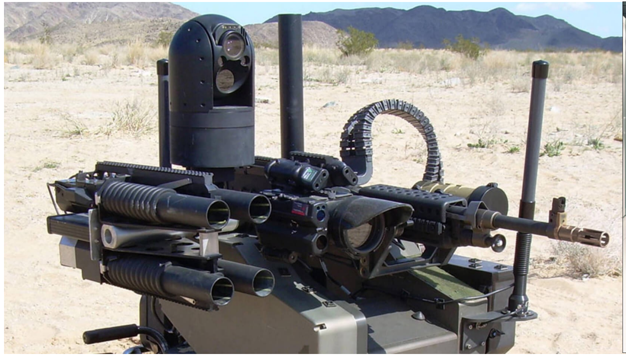

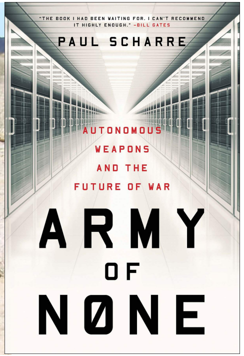

https://autonomousweapons.org/

The development of general-purpose robot autonomy calls for new approaches for ethics, philosophies, social sciences, economics, and political science.

How Can AI Systems Understand Human Values?

August 14, 2019 / by Jolene Creighton

Machine learning (ML) algorithms can already recognize patterns far better than the humans they're working for. This allows them to generate predictions and make decisions in a variety of high-stakes situations. For example, electricians use IBM Watson's predictive capabilities to anticipate clients' needs; Uber's self-driving system determines what route will get passengers to their destination the fastest; and Insilico Medicine leverages its drug discovery engine to identify avenues for new pharmaceuticals.

As data-driven learning systems continue to advance, it would be easy enough to define "success" according to technical improvements, such as increasing the amount of data algorithms can synthesize and, thereby, improving the efficacy of their pattern identifications. However, for ML systems to truly be successful, they need to understand human values. More to the point, they need to be able to weigh our competing desires and demands, understand what outcomes we value most, and act accordingly.

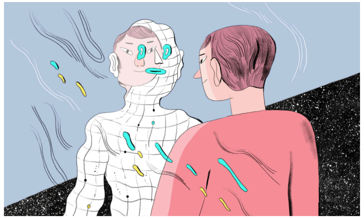

Opinion

OP-ED CONTRIBUTOR

How to Make A.I. That's Good for People

By Fei-Fei Li

March 7, 2018

Elisa Macellari

For a field that was not well known outside of academia a decade ago, artificial intelligence has grown dizzyingly fast. Tech

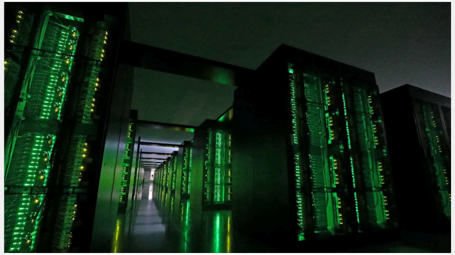

Why aligning AI to our values may be harder than we think

Can we stop a rogue AI by teaching it ethics? That might be easier said than done.

SCOTTY HENDRICKS 19 October, 2020

Eerie looking supercomputer.

Credit: STR/JIJI PRESS/AFP via Getty Images

Natalie Wolchover

Senior Writer/Editor

April 21, 2015

VIEW PDF/PRINT MODE

Artificial Intelligence
Computer Science
Deep Learning
Machine Learning
Q&A
All Topics →

Q & A

Concerns of an Artificial Intelligence Pioneer

The computer scientist Stuart Russell wants to ensure that our increasingly intelligent machines remain aligned with human values.

31

Stuart Russell, a computer scientist at the University of California, Berkeley, during a March stopover in San Antonio, Texas.

To be a **Technologist**, be a **Humanist** first.

"Artificial intelligence should treat all people fairly, empower everyone, perform reliably and safely, be understandable, be secure and respect privacy, and have algorithmic accountability. It should be aligned with existing human values, be explainable, be fair, and respect user data rights. It should be used for socially beneficial purposes, and always remain under meaningful human control."

— Tom Chatfield (2020)

[Source: There's No Such Thing As 'Ethical A.I.']

Robotics at UT-Austin

Join US!

Be part of the Robotics + AI revolution!

Robot Perception & Learning Lab

http://rpl.cs.utexas.edu/

Mission: Building General-Purpose Robot Autonomy in the Wild

TEXAS Robotics

https://robotics.utexas.edu/

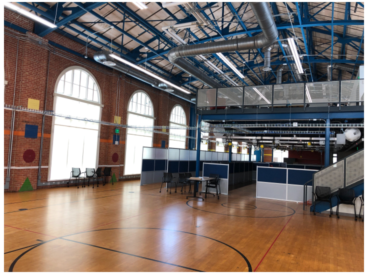

Robot Learning Reading Group

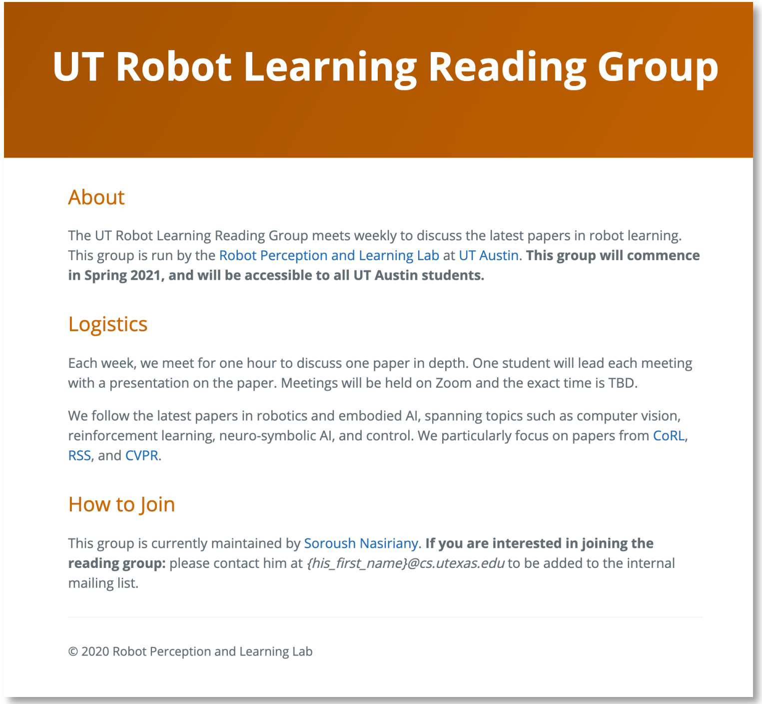

UT Robot Learning Reading Group

About

The UT Robot Learning Reading Group meets weekly to discuss the latest papers in robot learning. This group is run by the Robot Perception and Learning Lab at UT Austin. **This group will commence in Spring 2021, and will be accessible to all UT Austin students.**

Logistics

Each week, we meet for one hour to discuss one paper in depth. One student will lead each meeting with a presentation on the paper. Meetings will be held on Zoom and the exact time is TBD.

We follow the latest papers in robotics and embodied AI, spanning topics such as computer vision, reinforcement learning, neuro-symbolic AI, and control. We particularly focus on papers from CoRL, RSS, and CVPR.

How to Join

This group is currently maintained by Soroush Nasiriany. **If you are interested in joining the reading group:** please contact him at *{his_first_name}@cs.utexas.edu* to be added to the internal mailing list.

Starting in Spring 2021

http://ut-robotlearning.github.io/

Soroush Nasiriany