CS391R: Robot Learning

Conclusion: Open Questions in Robot Learning

Prof. Yuke Zhu

Fall 2023

Logistics

- Spotlight presentations (November 28 and November 30)
    - 4:55 min spotlight talk (see detailed instructions on Course Project page)
    - Video submission: upload to Gradescope by November 27
    - Presentation schedule: see Ed Discussion post
    - Spotlight videos will be played on TA's laptop (while you present) for time control

Today's Agenda

- General-Purpose Robot Autonomy (GPRA)
  - Review of the key concepts
- Algorithmic Toolbox for Robot Autonomy
- Open Questions in Robot Learning
- Societal impacts of Robotics + AI

General-Purpose Robot Autonomy … in the Wild

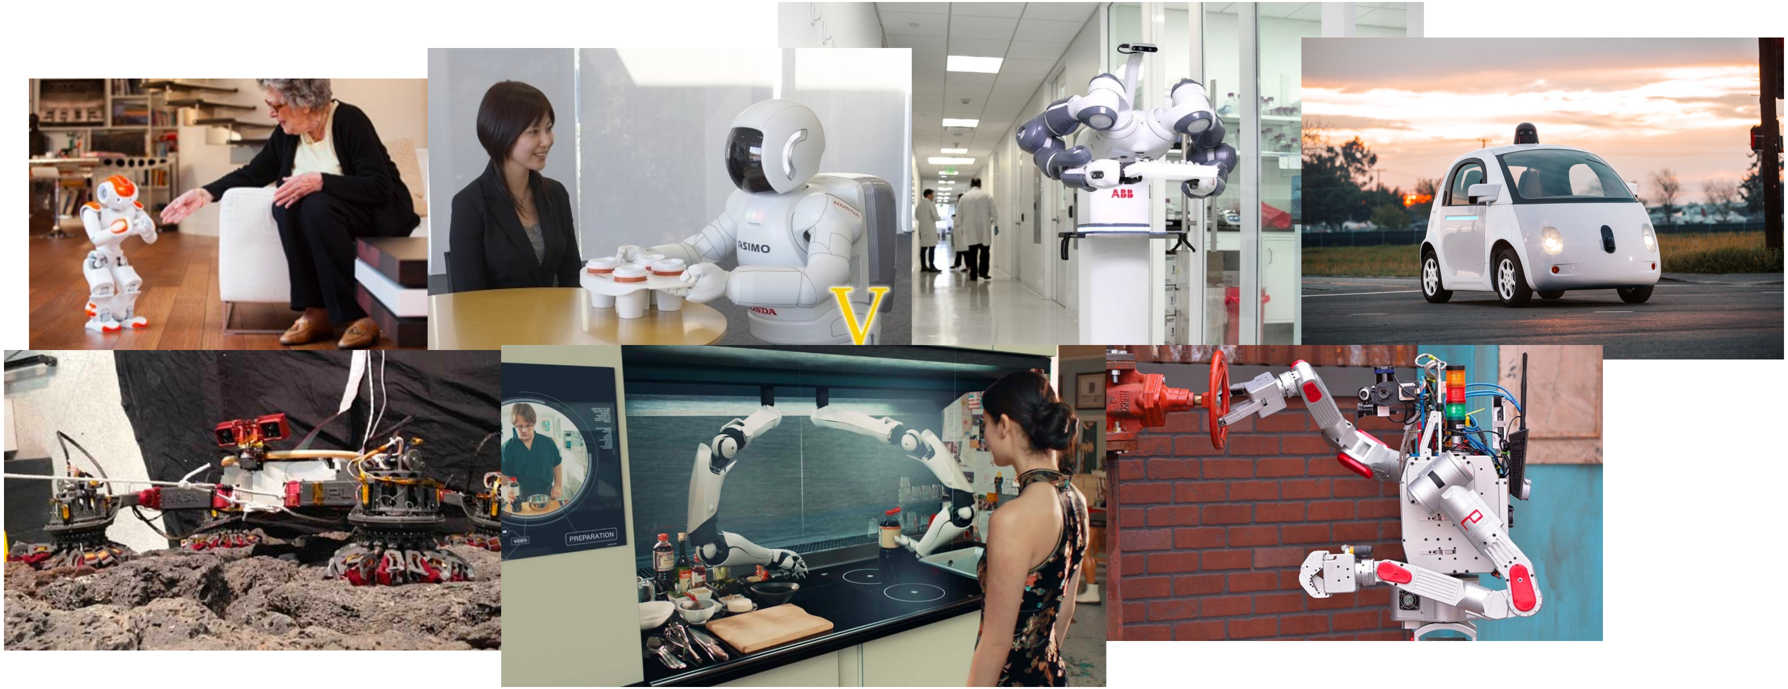

Unstructured Environments

Ever-changing Tasks

Human Involvement

Special-Purpose Robot Automation

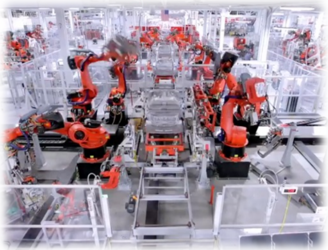

custom-built robots

human expert programming

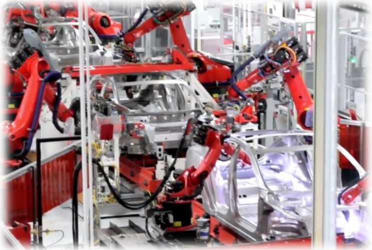

special-purpose behaviors

General-Purpose Robot Autonomy

general-purpose robots

Robot Learning

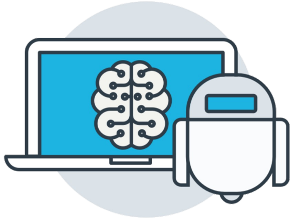

general-purpose behaviors

We have come a long way!

Our journey to the

Robot Learning wonderland

18 topics (**7** for perception, **7** for decision-making, **4** for research frontiers)

79 readings (**39** required and **40** optional)

2 guest lectures (**1** from industry and **1** from academia)

+ background lectures, tutorials, extended materials…

Robot Perception

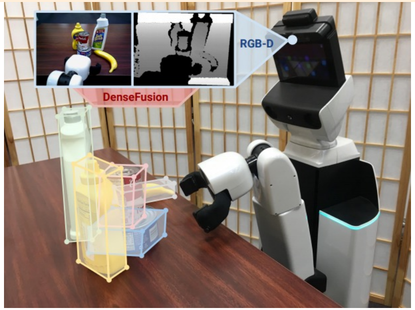

Convolutional network
PointNet / PointNet++
Vision transformer
Unsupervised learning
Self-supervised learning

Generative modeling
Neural fields
Bayes filtering
Domain randomization

Point cloud processing
Object detection
Object segmentation
Pose estimation
Visual tracking

Multimodal fusion
Visual navigation
Recursive state estimation
Visual SLAM

Decision Making

Model-free reinforcement learning
Trust-region optimization
Model-based dynamics learning
Offline (batch) reinforcement learning
Task and motion planning

Gaussian process
Behavior cloning / DAgger
Neural programming
Learning to learn
Diffusion models

Sensorimotor learning
Video prediction
Reward/utility learning
Long-horizon manipulation

Human-in-the-loop system
Learning from demonstration
Autonomous driving

Your Algorithmic Toolbox for Building Robot Autonomy

Open Questions in Robot Learning

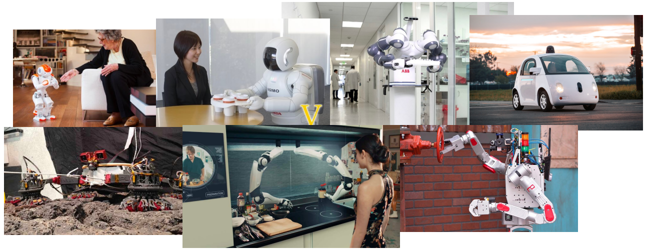

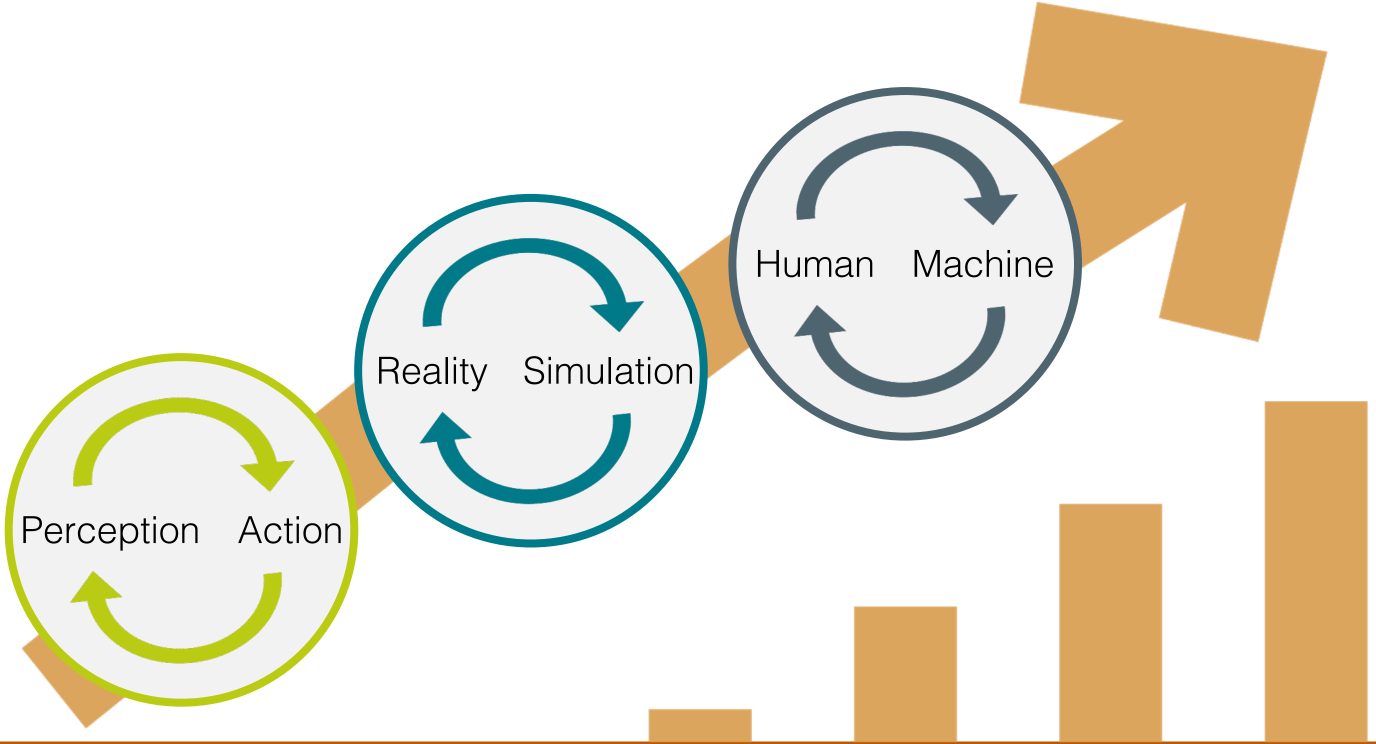

General Purpose
Robot Autonomy

Open Questions in Robot Learning

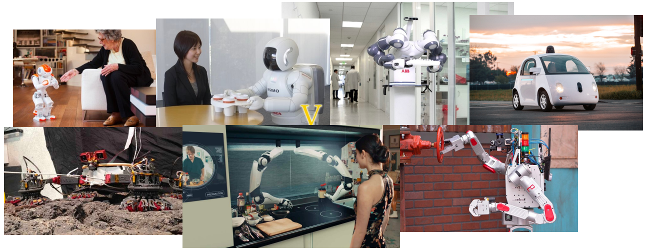

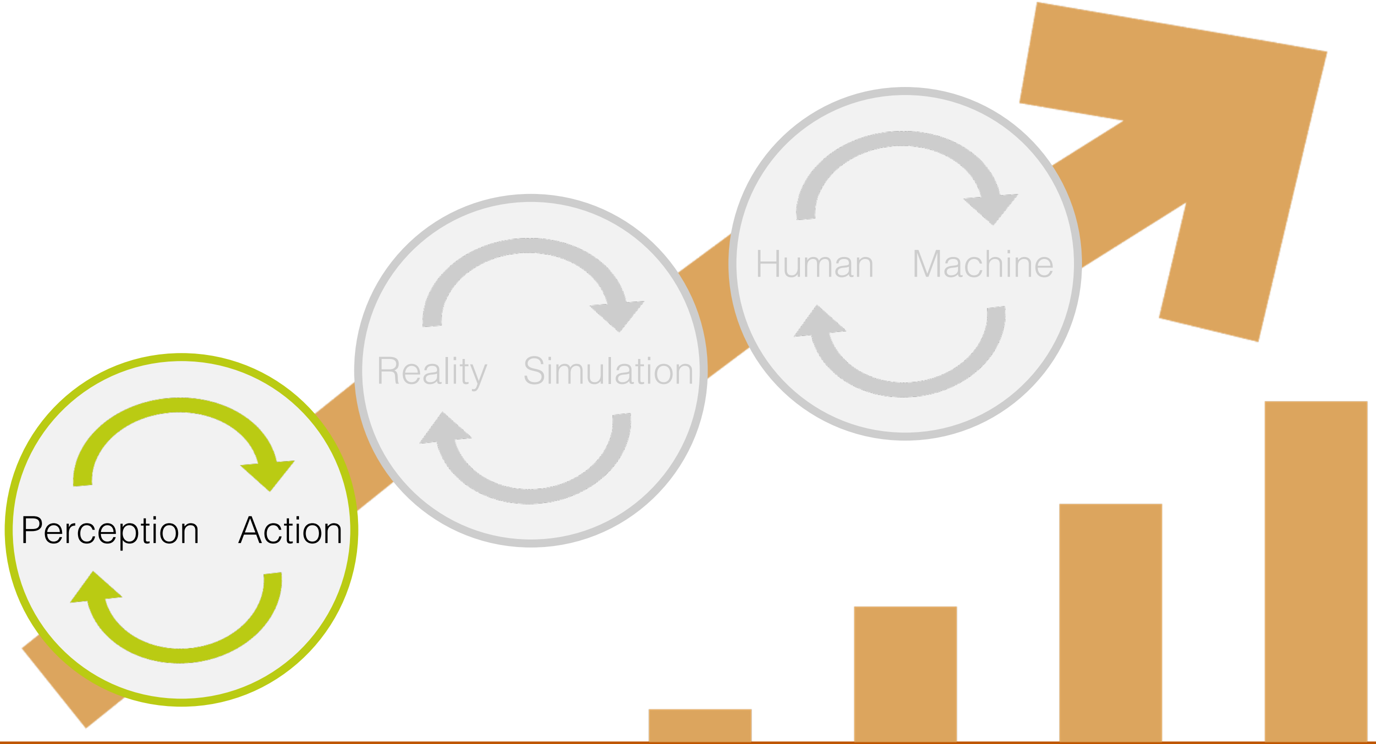

General Purpose
Robot Autonomy

Closing the **Perception-Action** Loop

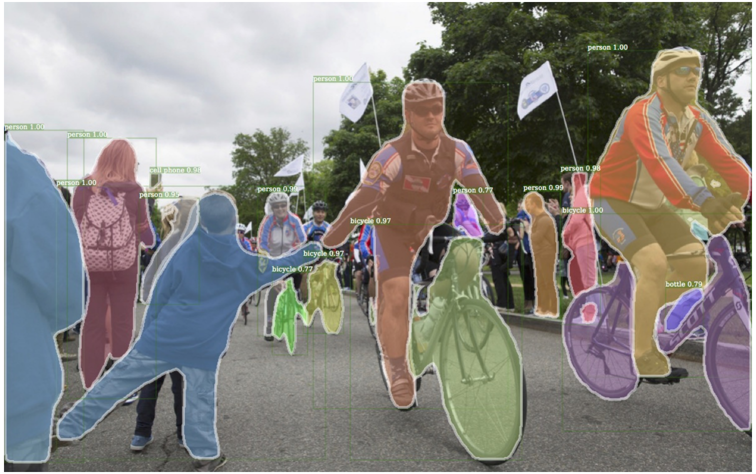

[Detectron - Facebook AI Research]

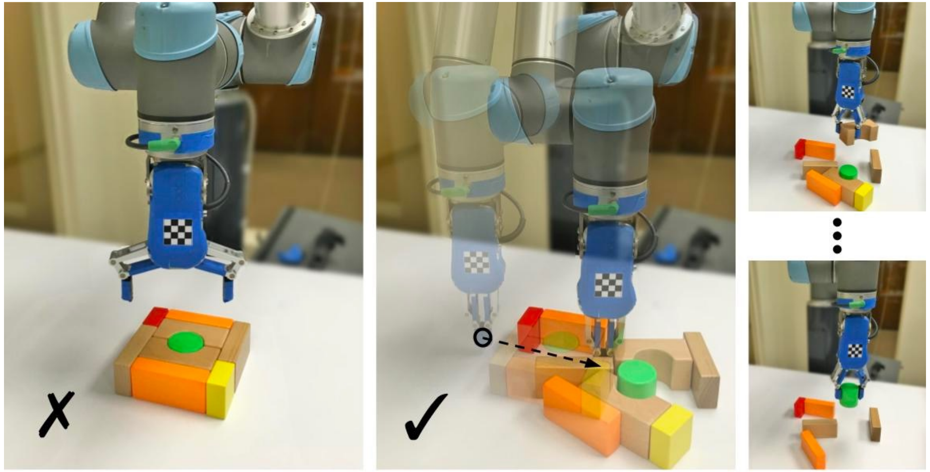

[Zeng et al., IROS 2018]

Conventional Computer Vision

Physically-Grounded Robot Perception

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

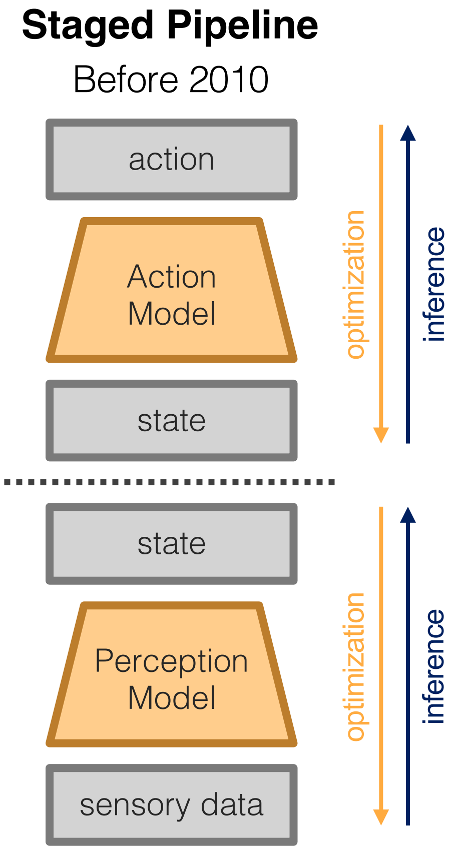

Close the **Perception-Action Loop**: A New Paradigm

Staged Pipeline

Before 2010

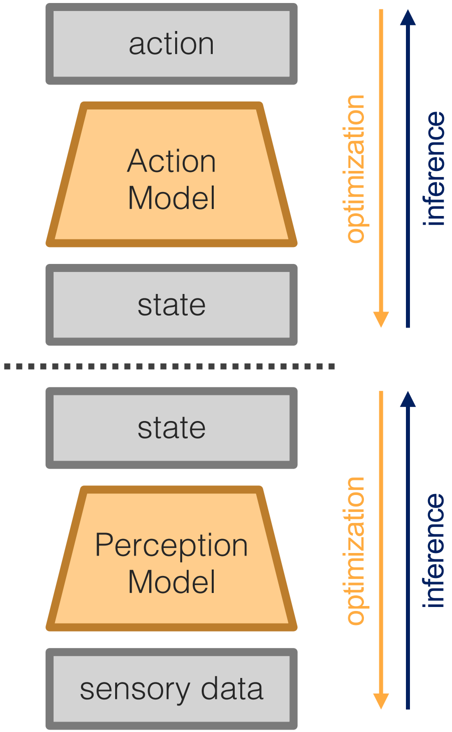

Raw Sensory Data

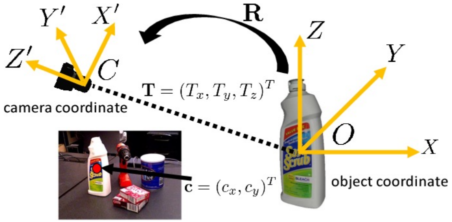

World State

Perception &
Computer Vision

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

Staged Pipeline

Before 2010

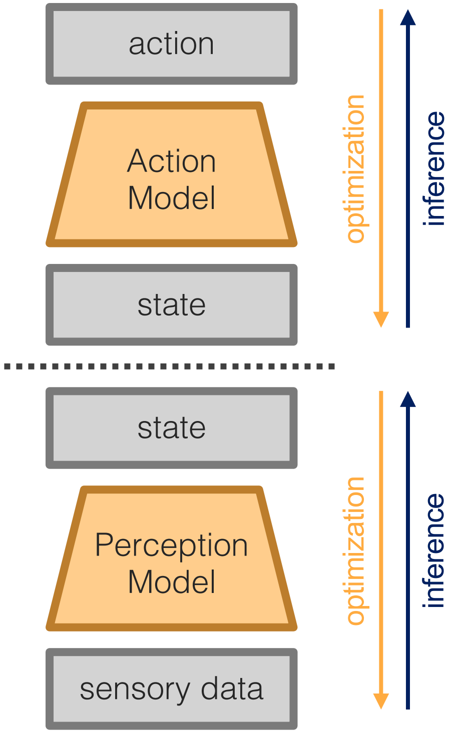

Raw Sensory Data

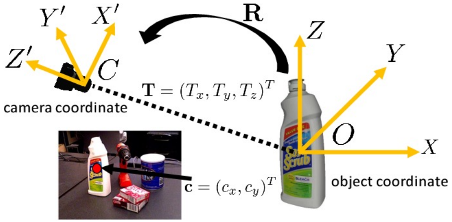

World State

Perception & Computer Vision

Robot Control & Decision Making

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

Staged Pipeline

Before 2010

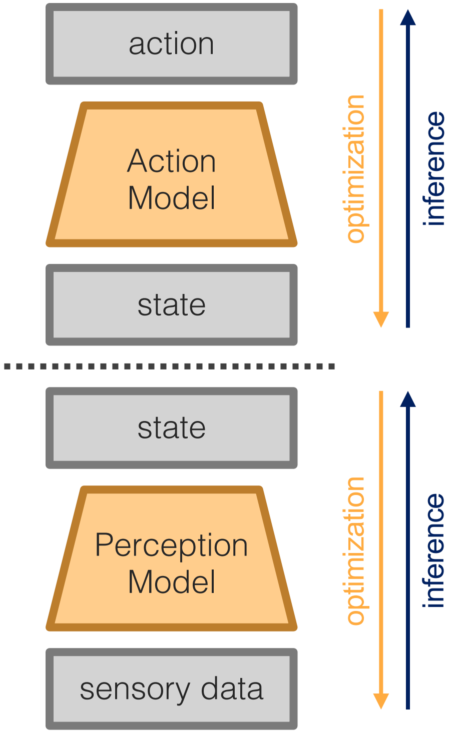

End-to-End Learning

2010 - 2020

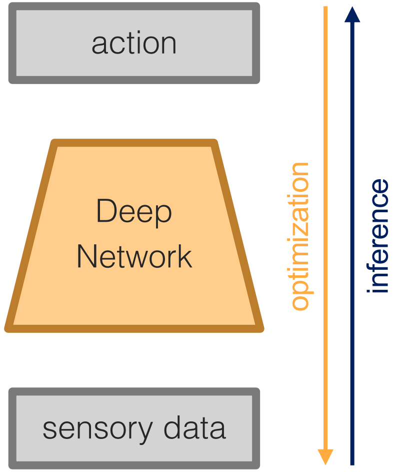

Close the **Perception-Action Loop**: A New Paradigm

Staged Pipeline

Before 2010

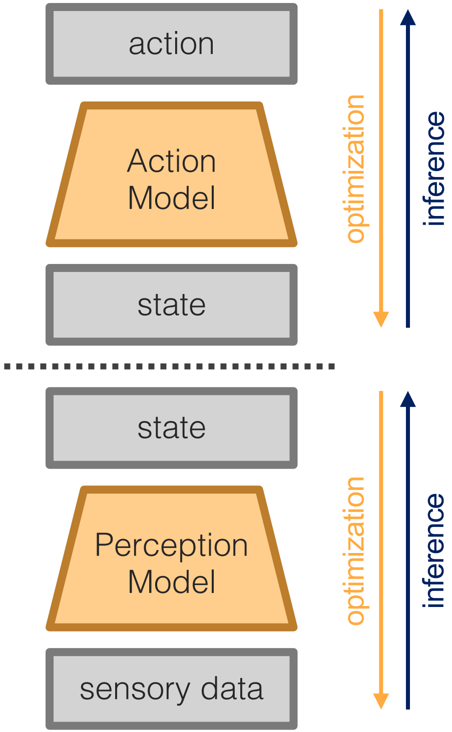

End-to-End Learning

2010 - 2020

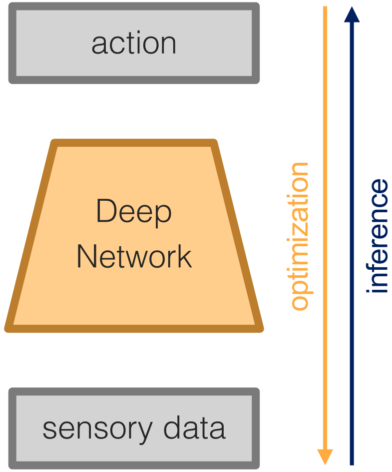

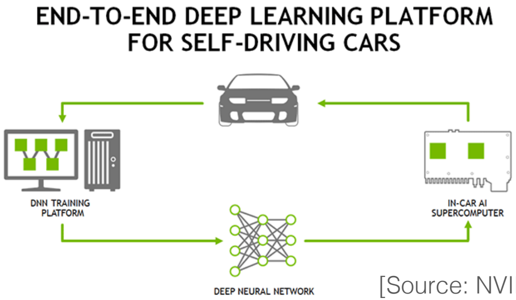

[Source: NVIDIA]

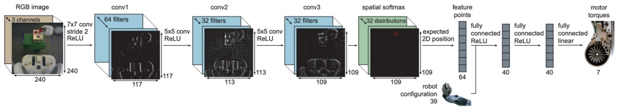

[Source: Levine et al. JMLR 2016]

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

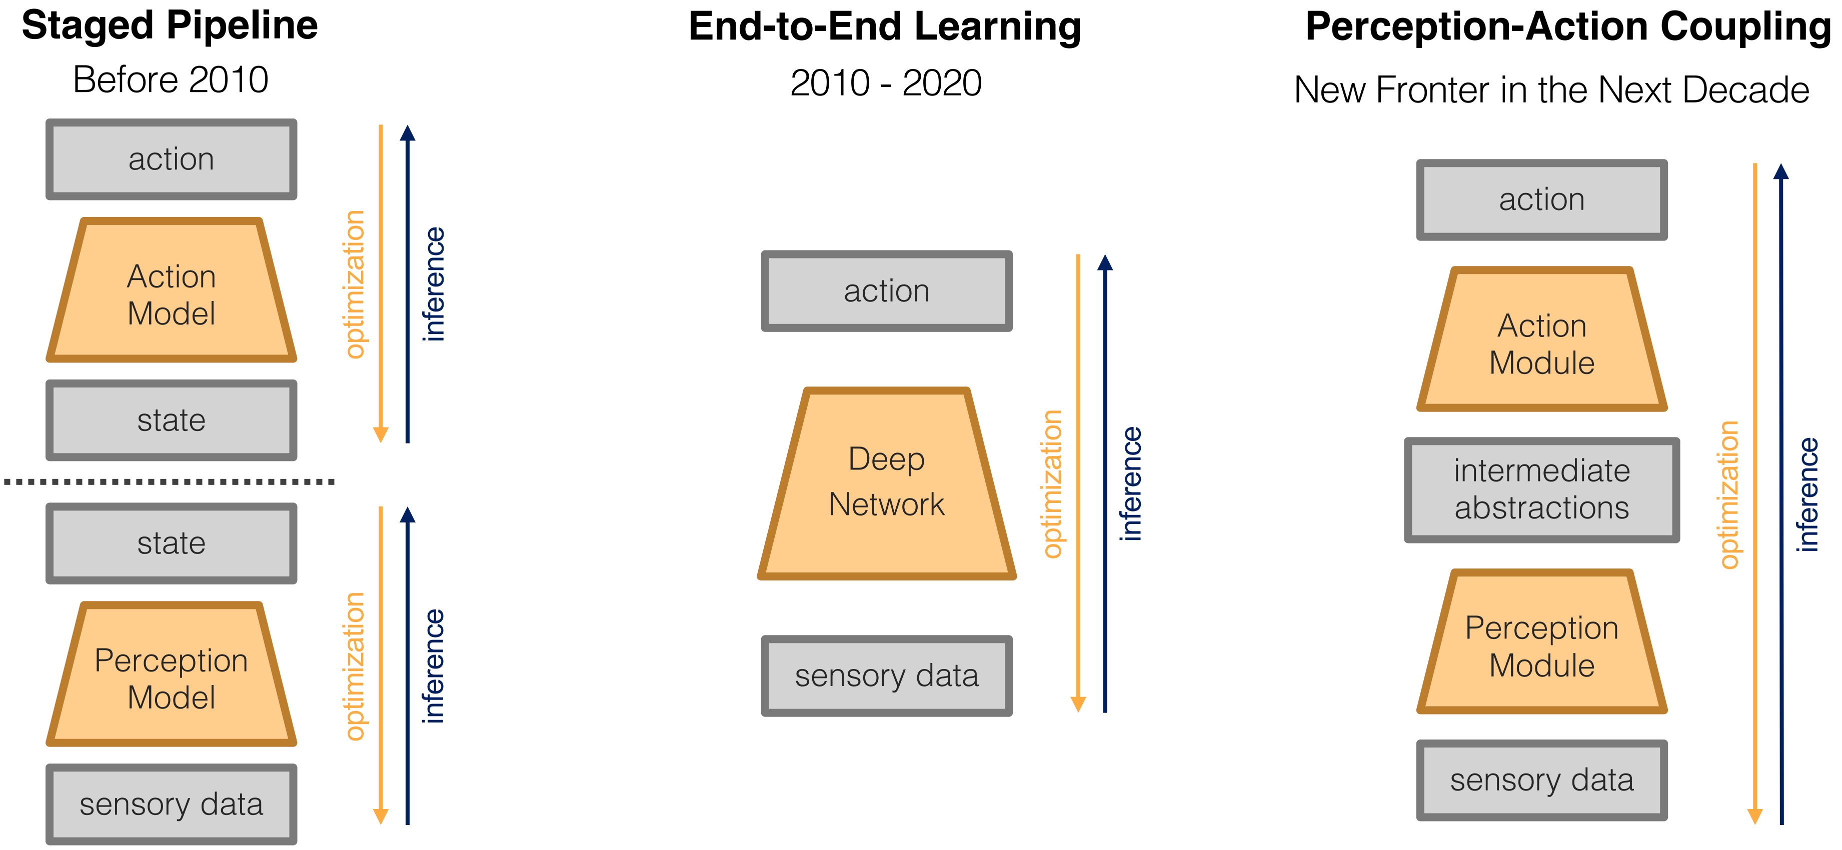

New Paradigm: Perception-Action Coupling

Rich inductive biases from model structures

Learning action-informed perceptual representation

Joint optimization of functional modules (Software 2.0)

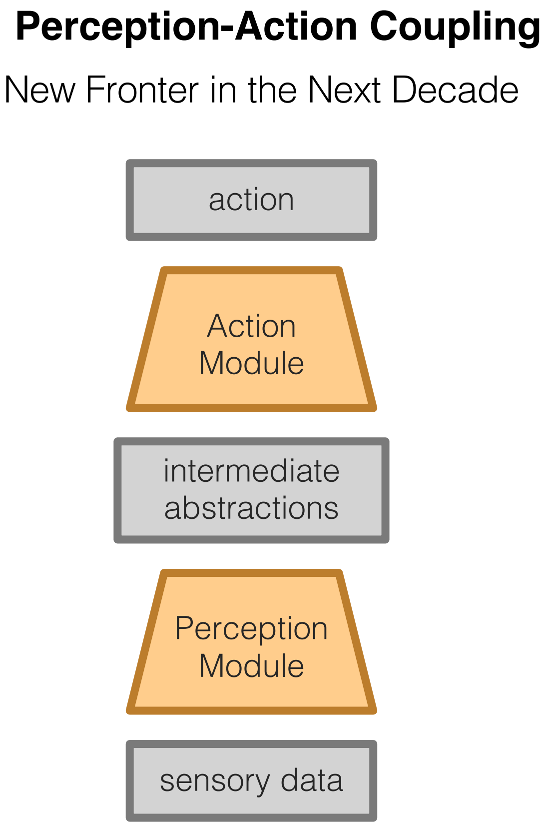

Software 2.0: https://medium.com/@karpathy/software-2-0-a64152b37c35

New Paradigm: Perception-Action Coupling

Perception-Action Coupling

New Fronter in the Next Decade

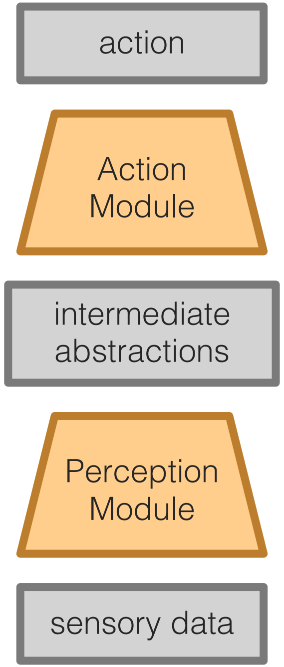

My Million-Dollar Question: “1 + 1 < 2?”

“Will a joint optimization of perception and decision-making make the computational problem of general-purpose robot autonomy easier?”


BONGU

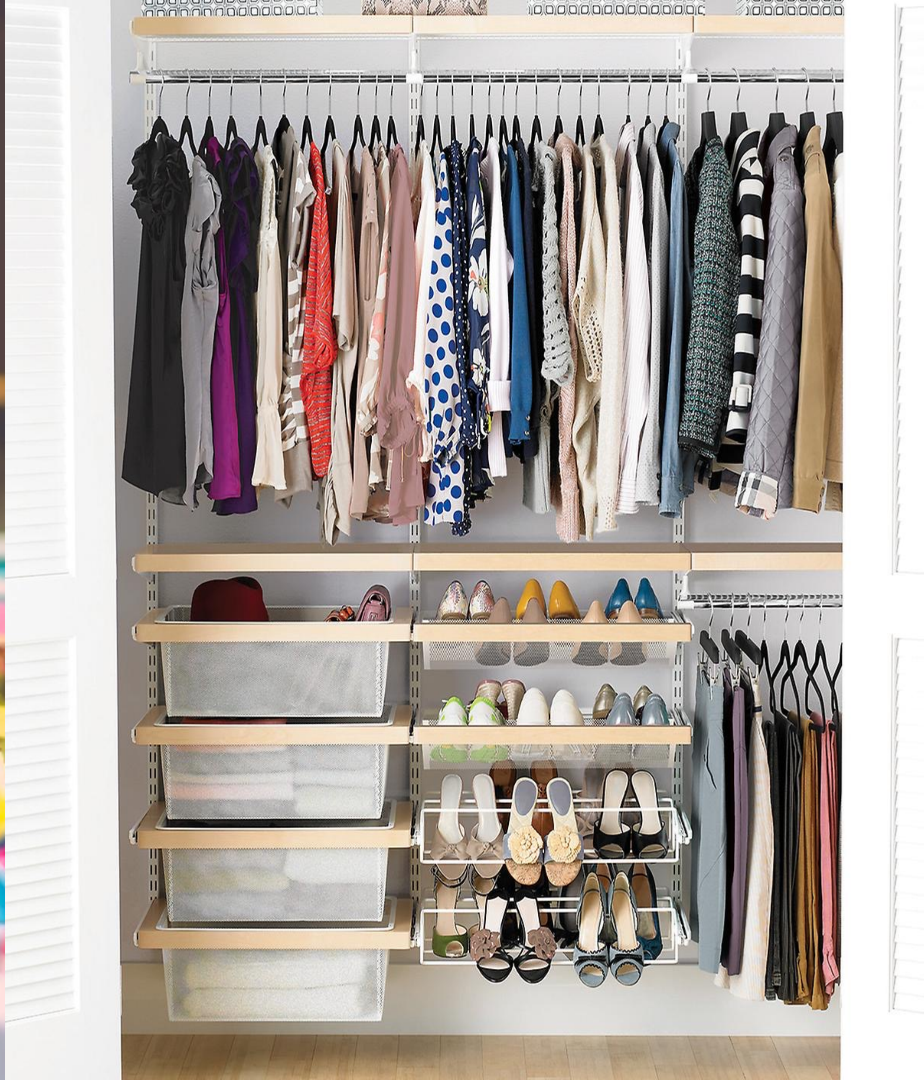

Open Questions in Robot Learning

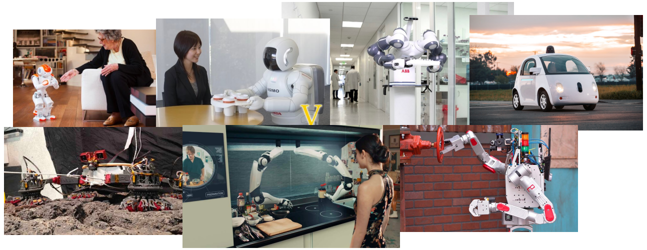

General Purpose
Robot Autonomy

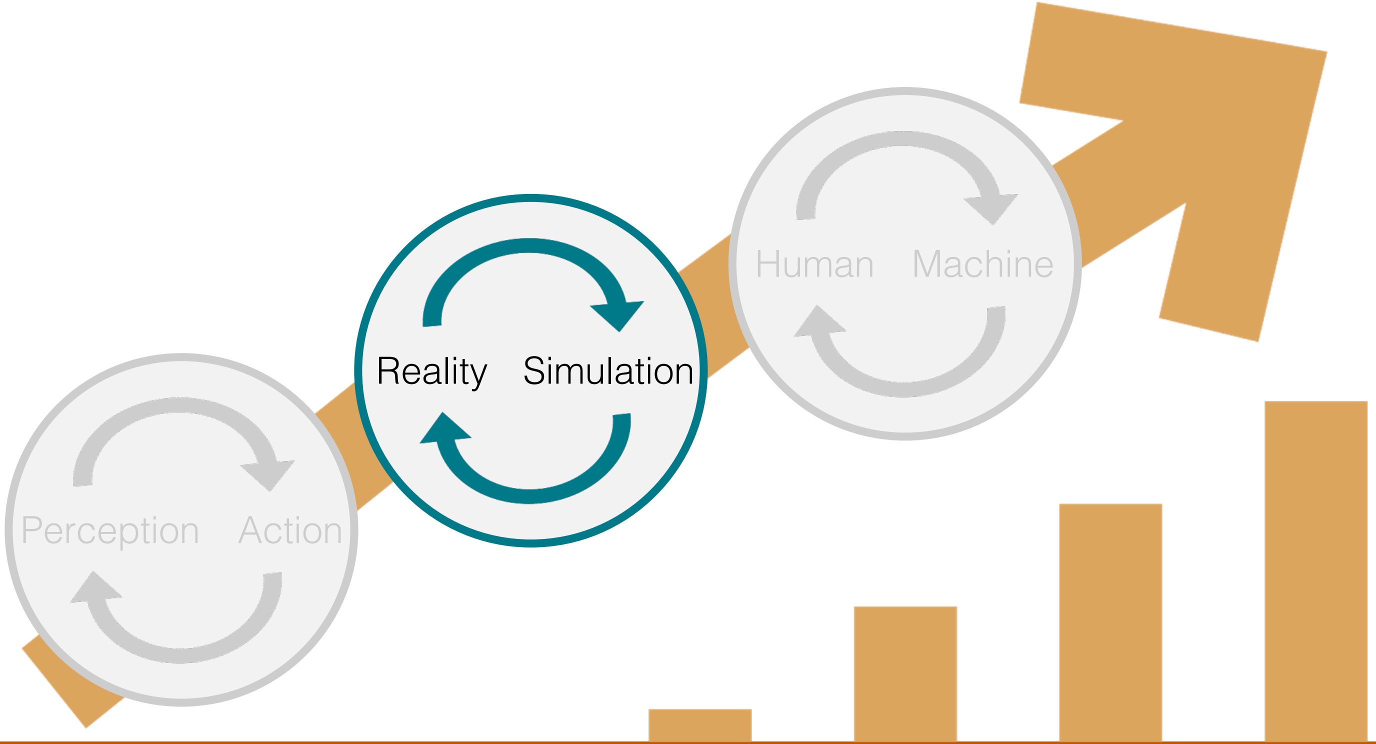

Closing the **Simulation-Reality** Loop

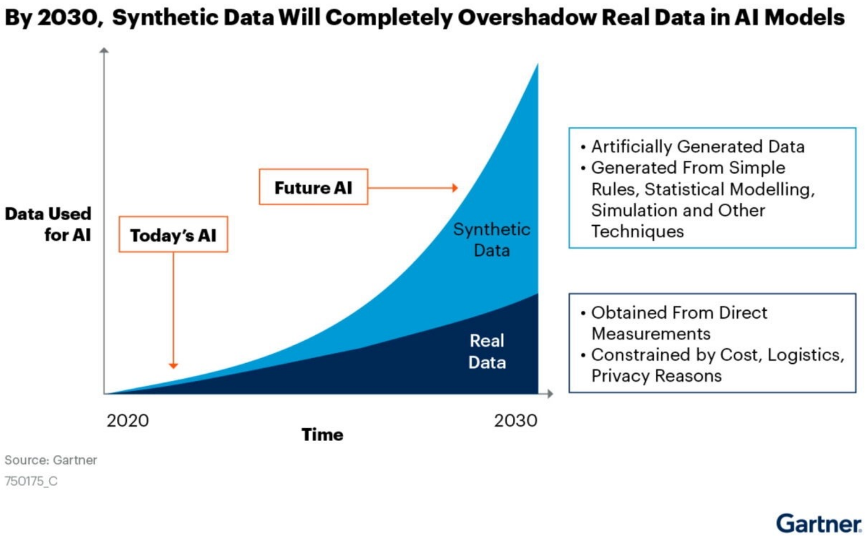

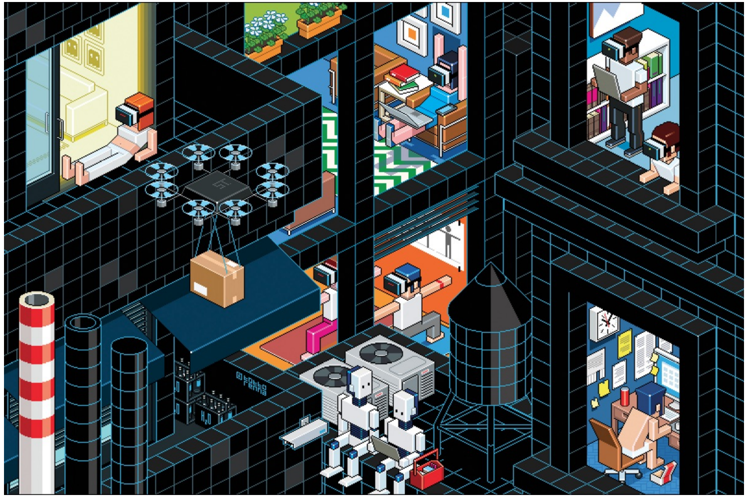

Closing the **Simulation-Reality** Loop

The reality gap is the major roadblock for learning models to benefit from synthetic data.

Photorealism (Visuals)

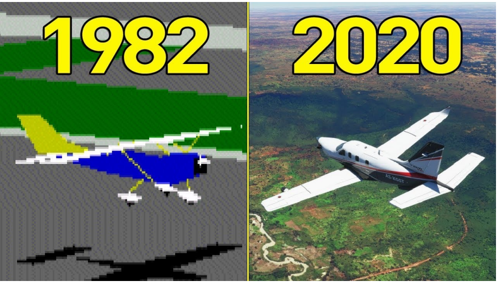

Microsoft Flight Simulator

Physical Realism (Dynamics)

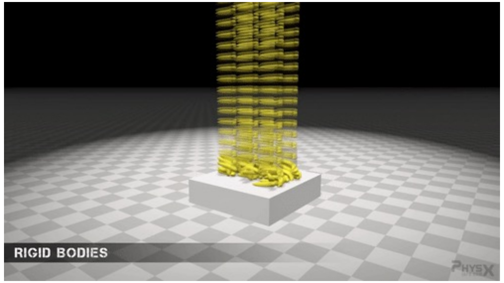

NVIDIA PhyX

Sensor Realism

BlenSor

Behavioral Realism

VirtualHome

Closing the **Simulation-Reality** Loop

The reality gap is the major roadblock for learning models to benefit from synthetic data.

real-world kitchen

photorealistic rendering of kitchen

[Roberts et al. "Hypersim" 2021]

Closing the **Simulation-Reality** Loop

Going from real world to simulation and from simulation to real world…

real-world experiences

generative modeling

realistic virtual world

Closing the **Simulation-Reality** Loop

Going from real world to simulation and from simulation to real world…

real-world experiences

generative modeling

realistic virtual world

robot learning in simulation

Closing the **Simulation-Reality** Loop

Going from real world to simulation and from simulation to real world…

real-world experiences

generative modeling

realistic virtual world

robot learning in simulation

adaptation & deployment

Closing the **Simulation-Reality** Loop

Generating objects

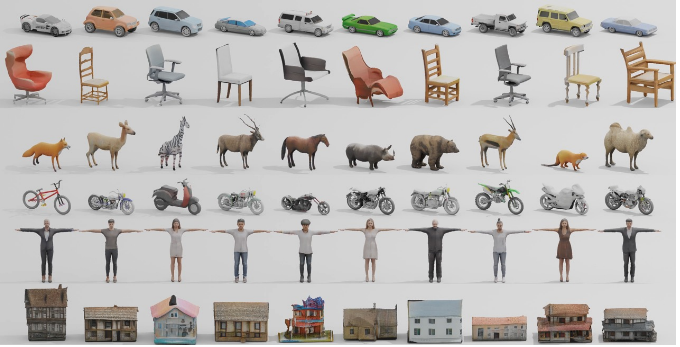

GET3D [Gao et al. NeurIPS 2022]

Generating reward functions

Eureka [Ma et al. arXiv 2023]

Generating manipulation tasks

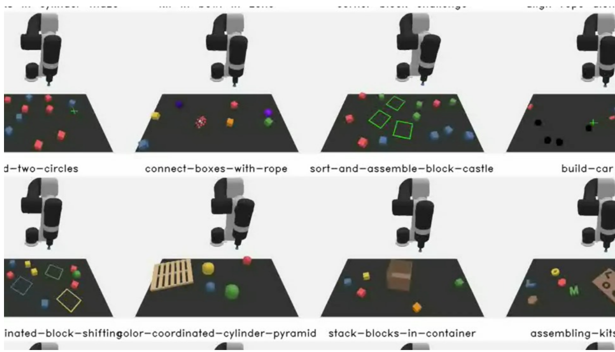

GenSim [Wang et al. arXiv 2023]

Generative AI techniques will play a pivotal role in the next generation of robotic simulation.

Open Questions in Robot Learning

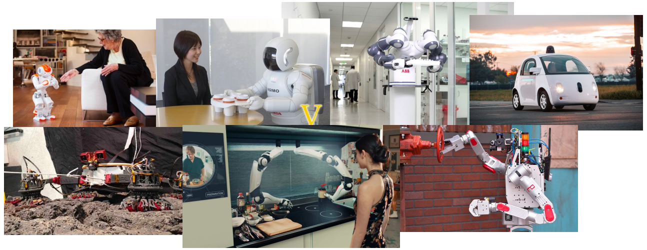

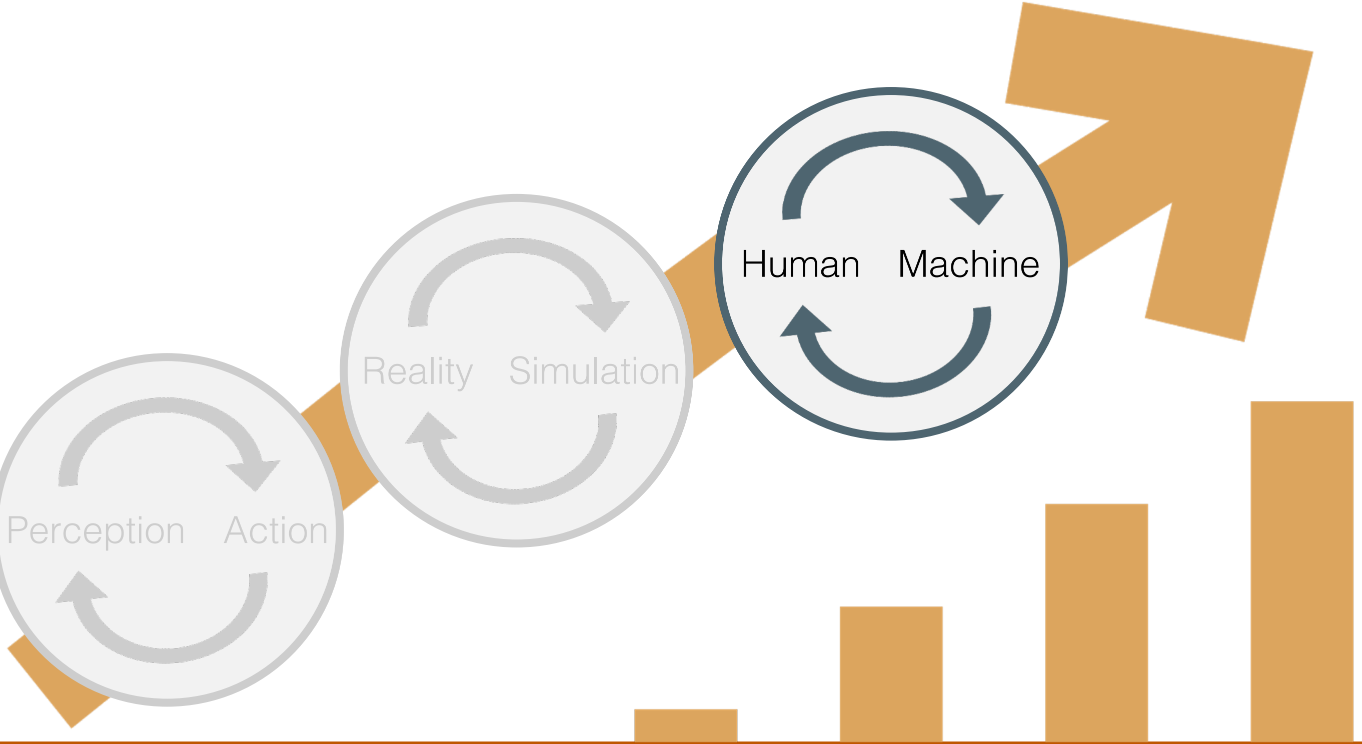

General Purpose
Robot Autonomy

Closing the **Human-Machine** Collaboration Loop

Most (if not all) deployable robot learning systems are

human-in-the-loop systems.

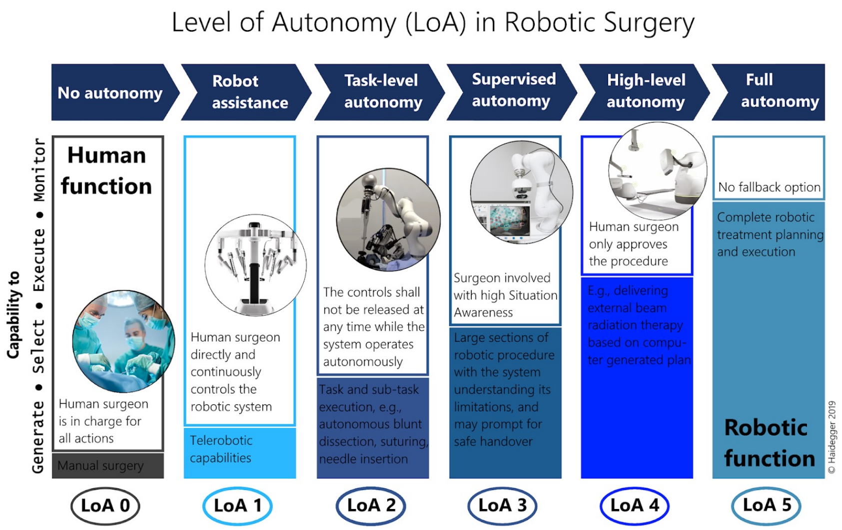

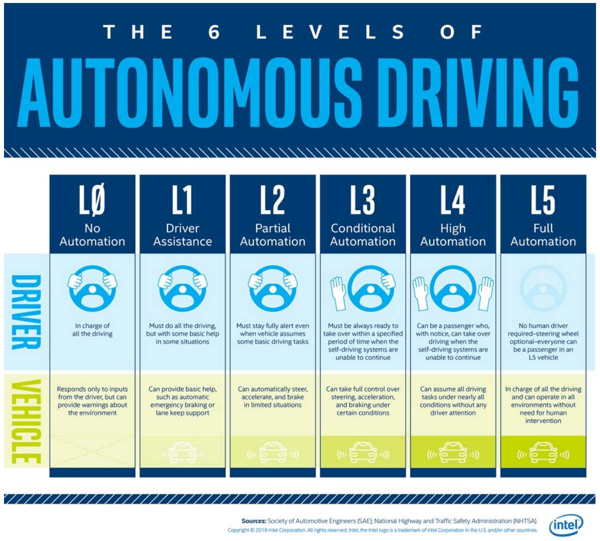

Closing the **Human-Machine** Collaboration Loop

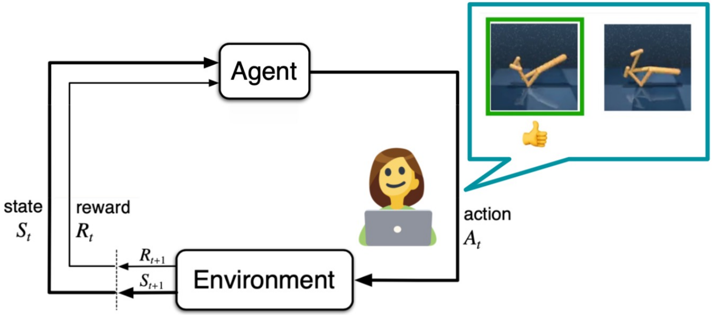

During learning

Accelerating robot learning with rich forms of human feedback

During deployment

Achieving performance guarantees through human-robot collaboration

Closing the **Human-Machine** Collaboration Loop

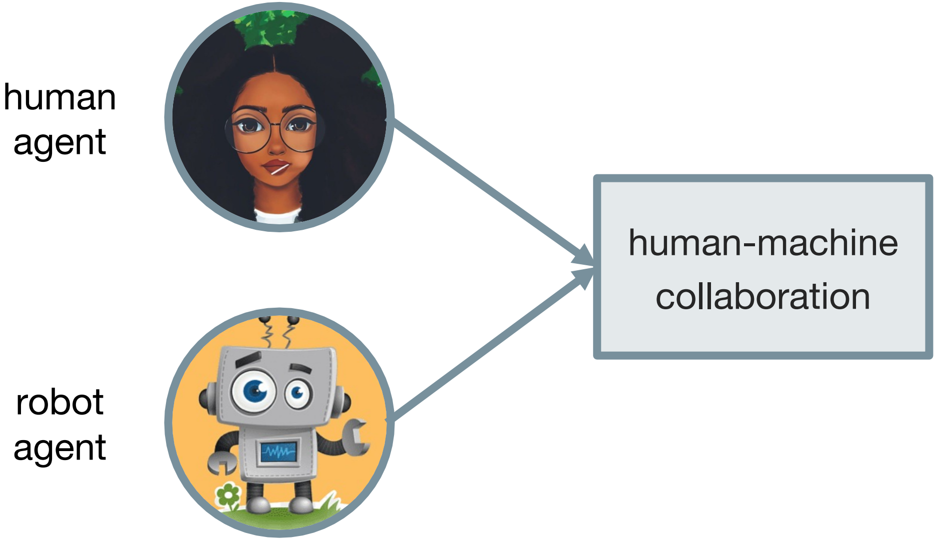

Closing the **Human-Machine** Collaboration Loop

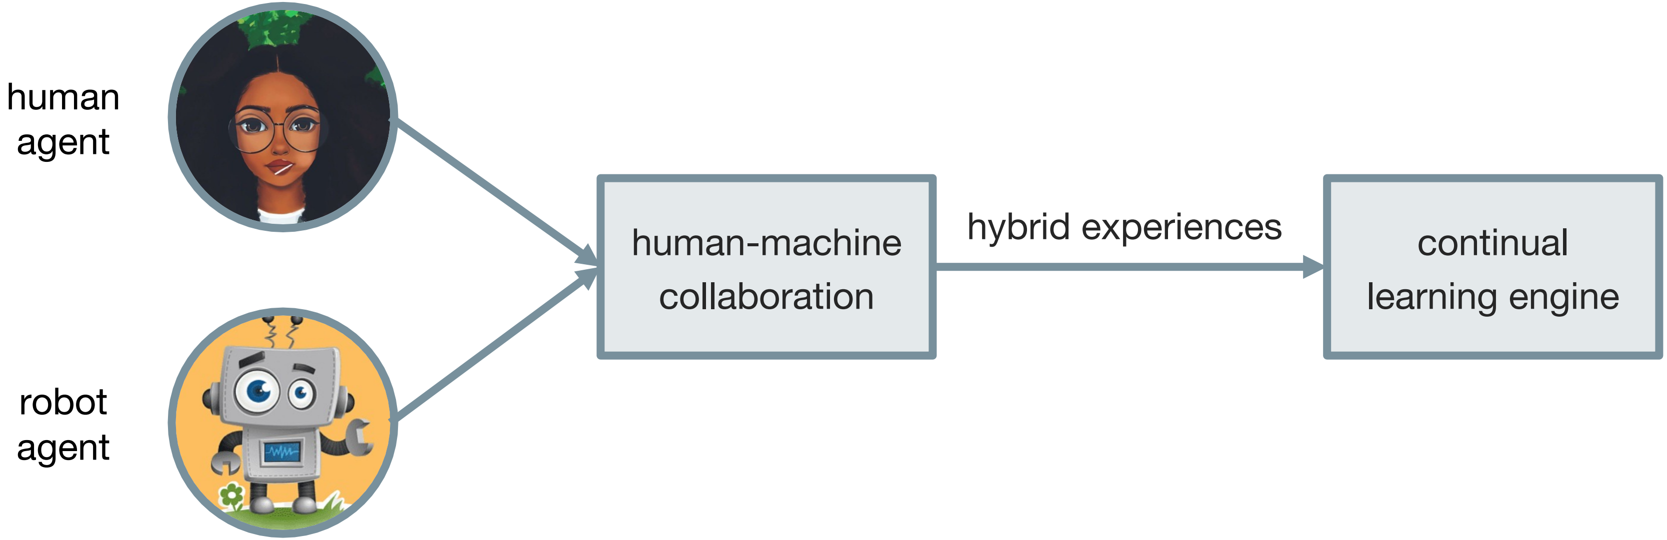

Closing the **Human-Machine** Collaboration Loop

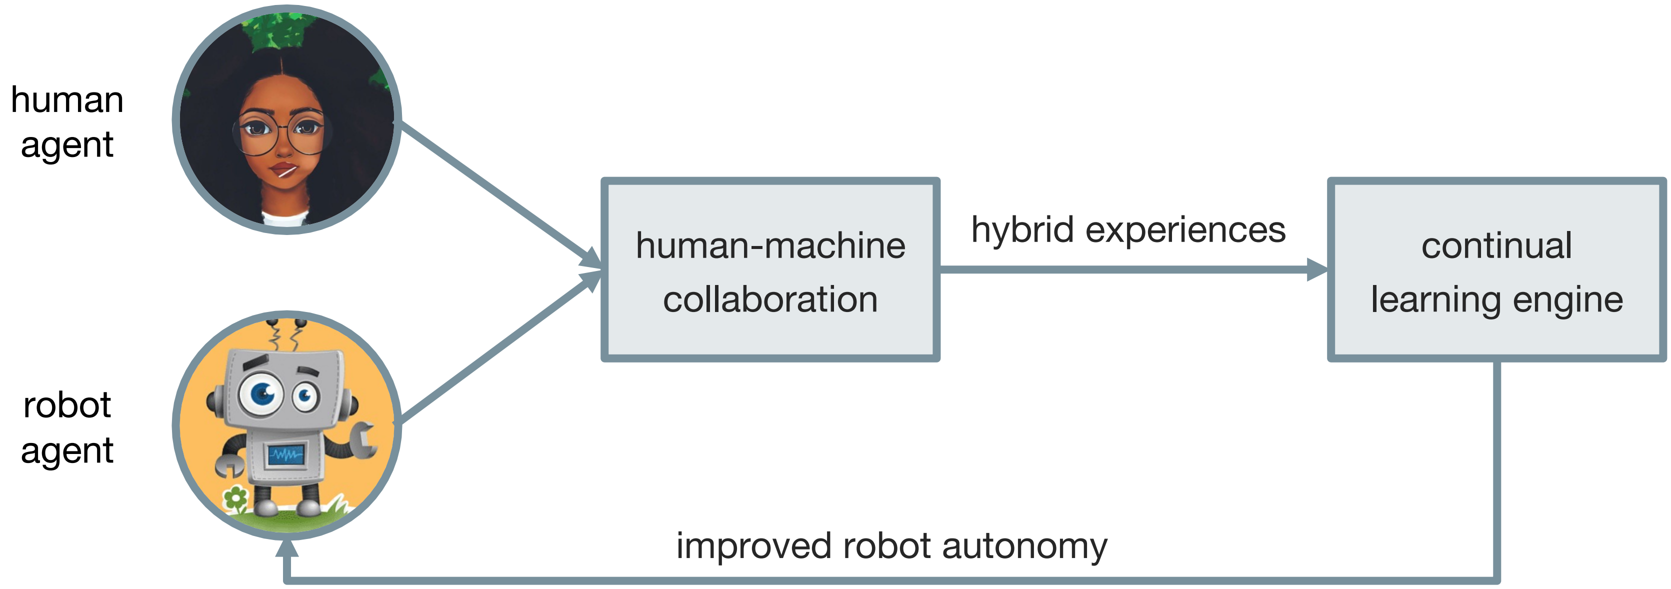

Closing the **Human-Machine** Collaboration Loop

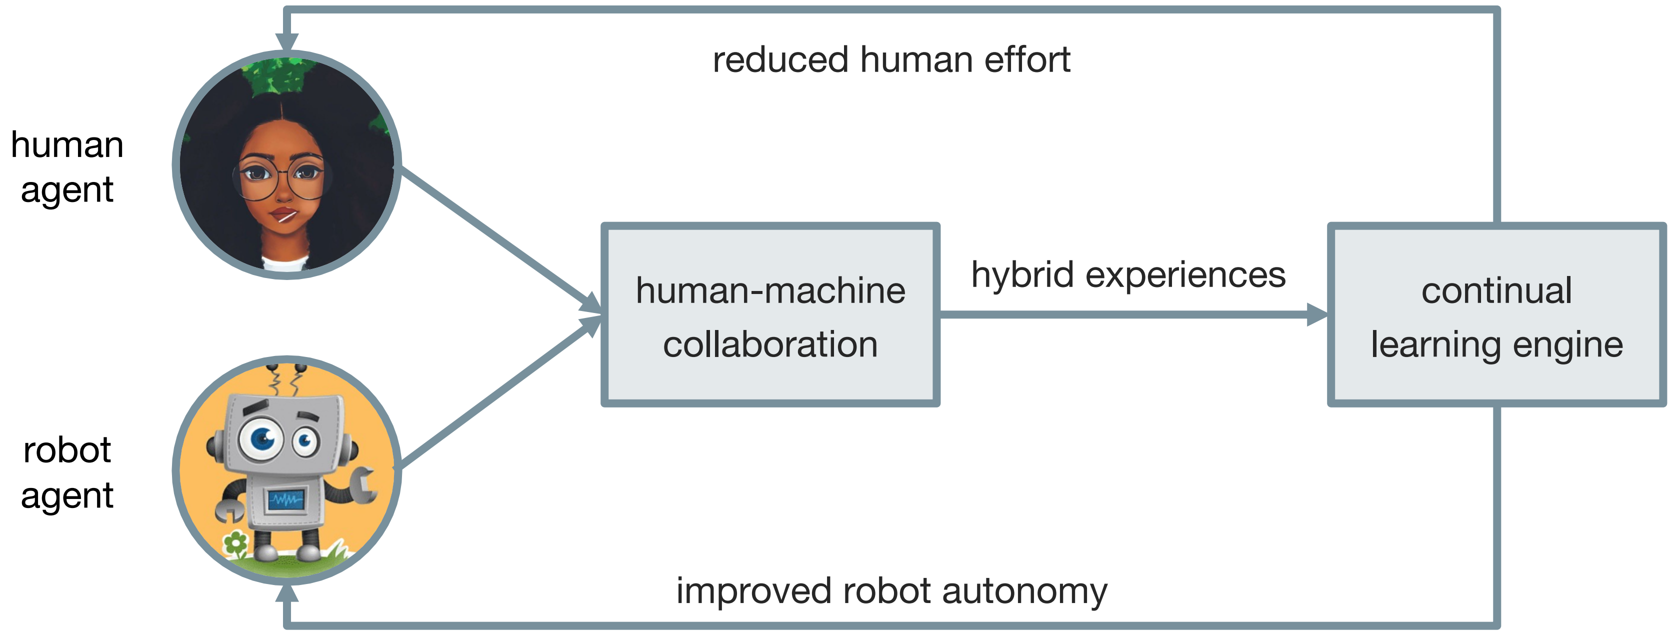

Closing the **Human-Machine** Collaboration Loop

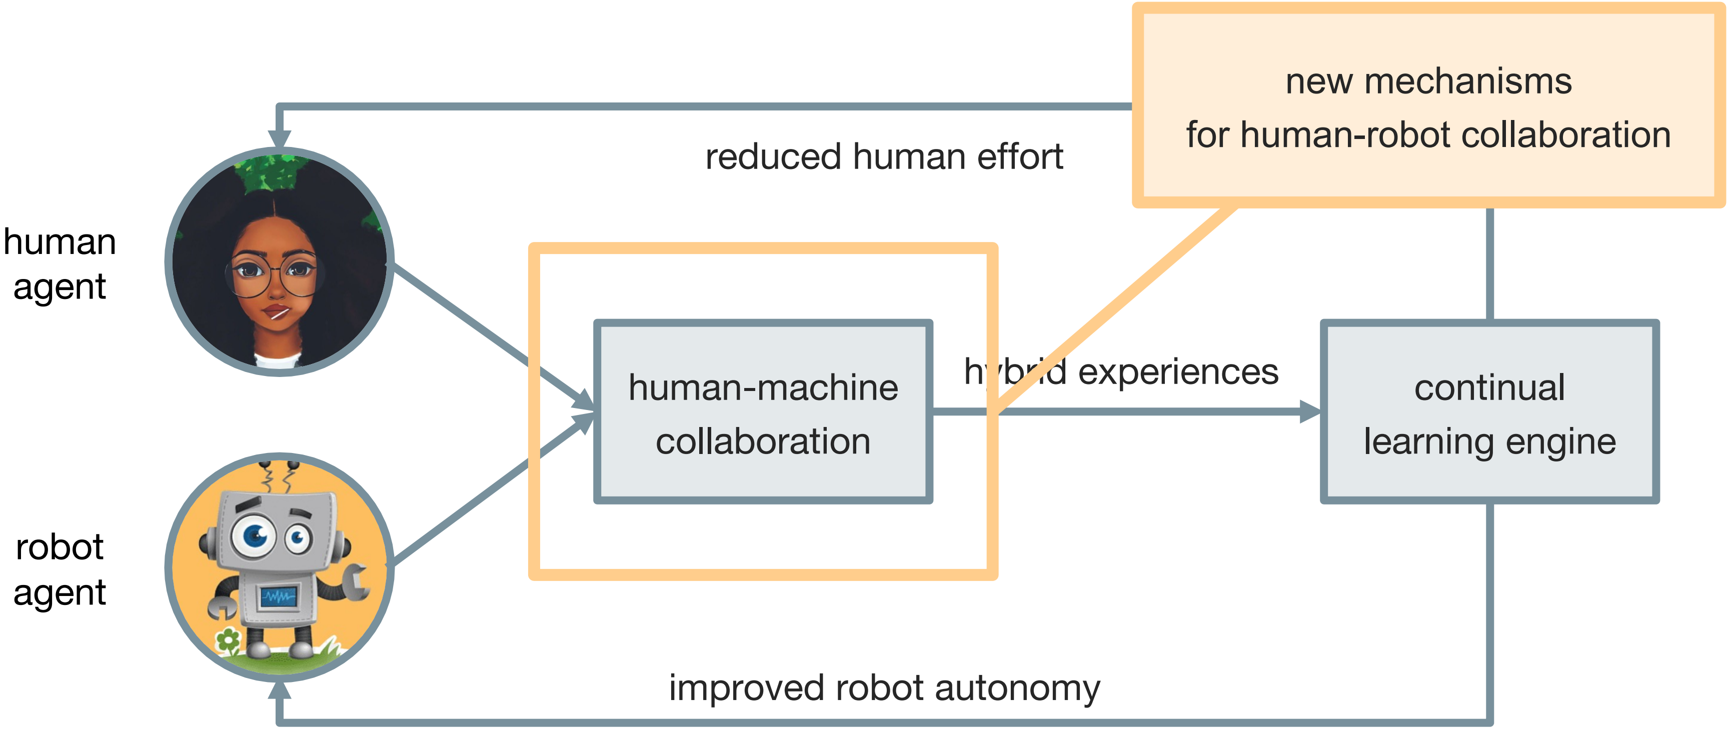

Closing the **Human-Machine** Collaboration Loop

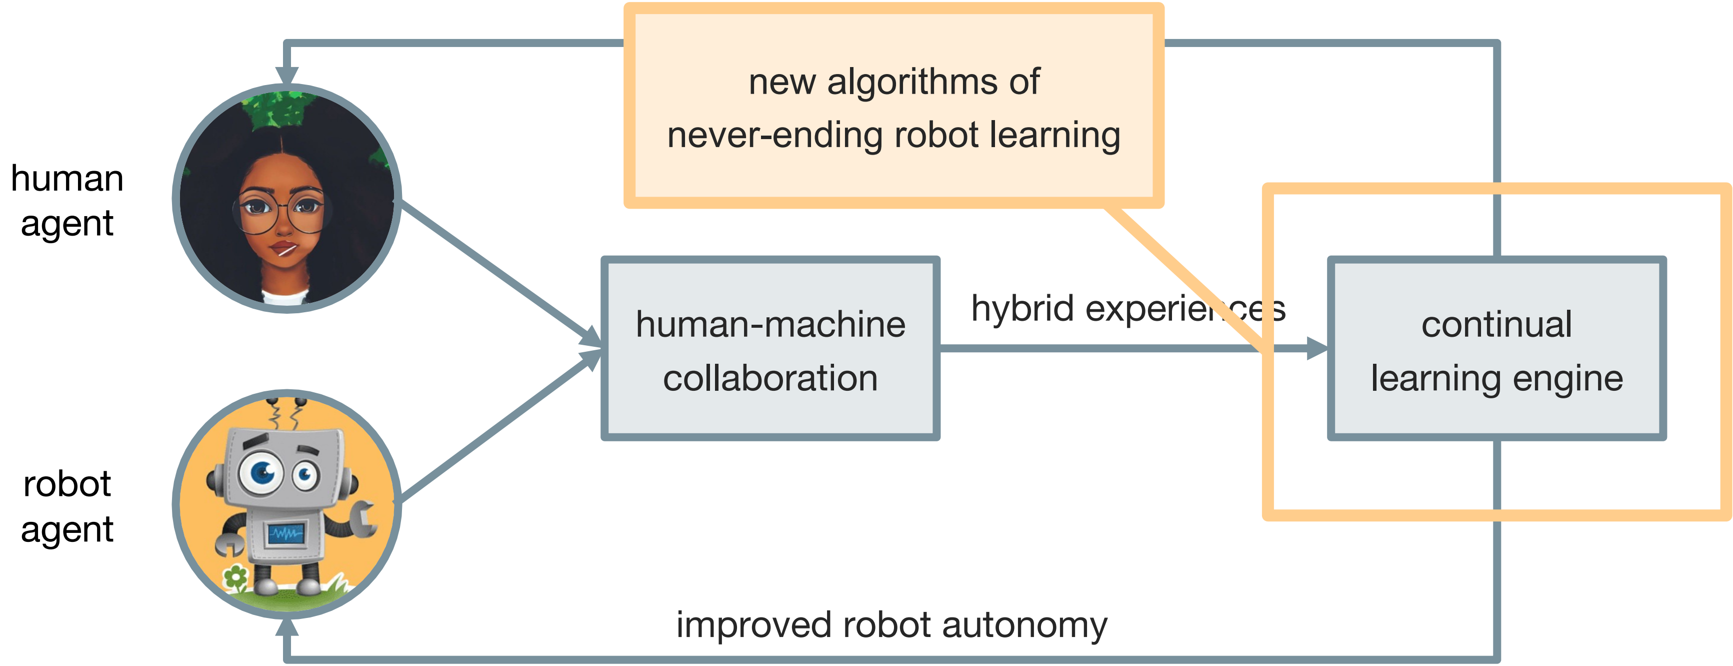

Closing the **Human-Machine** Collaboration Loop

How can we ensure trustworthy deployment?

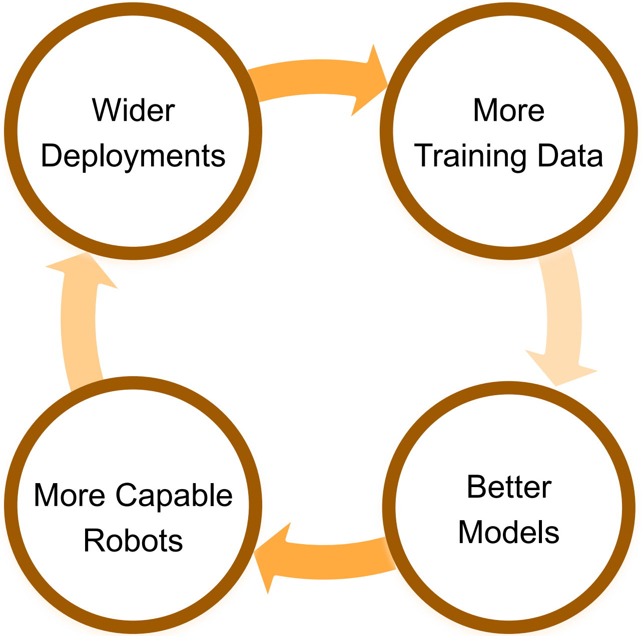

How can robots learn continually with more data?

Open Questions: **Requests for Research**

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:** harnessing the complementary strengths of the “data pyramid”, from internet data, synthetic data, to real-world data;
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, memory organization, etc.;
4. **Safety and robustness of real-world robotic systems:** simulation-to-reality gap, uncertainty quantification & safe learning, human-robot teaming, and trustworthy and verifiable AI systems.

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"

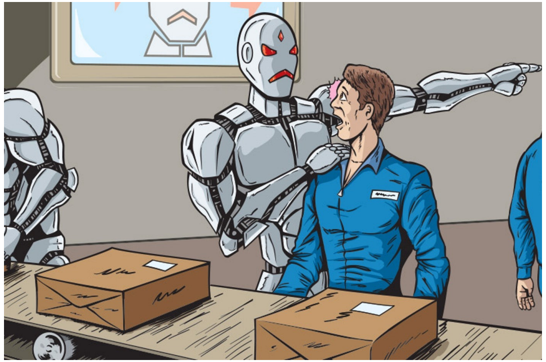

"Neo-Luddism's Tech Skepticism"

Question: What's the value of work?

Robots and Society

Personal assistive household robots in the aging society

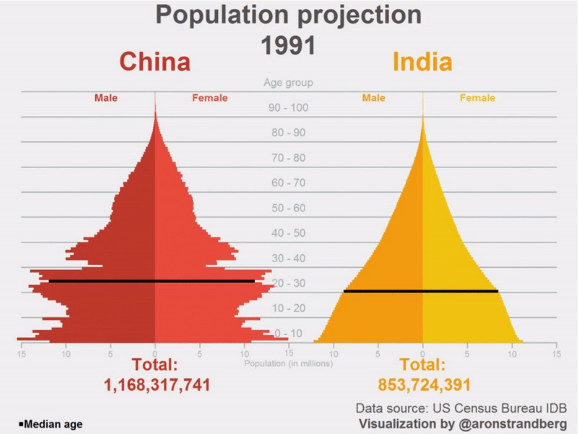

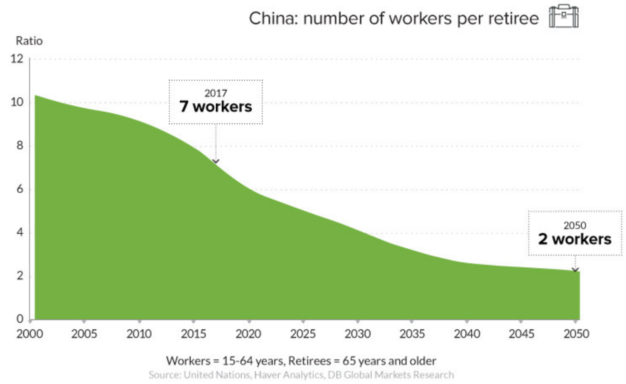

"Robot carers for the elderly in Japan"
[Source: The Times UK]

"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

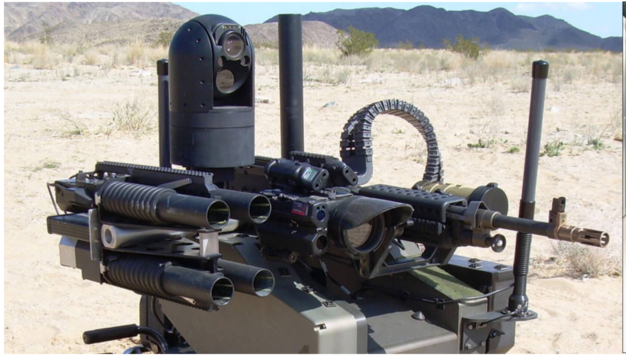

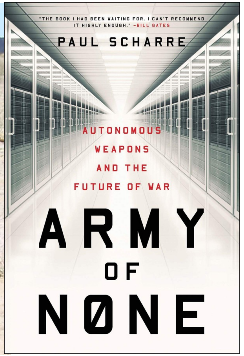

https://autonomousweapons.org/

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.

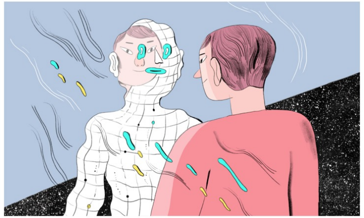

Opinion

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

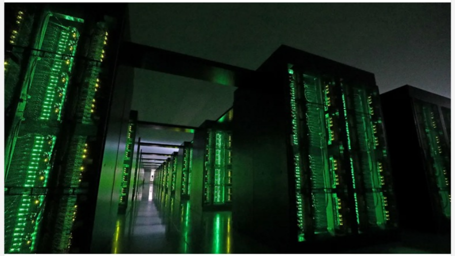

Why aligning AI to our values may be harder than we think

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.

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/