OSARO | AI FOR INDUSTRIAL ROBOTS

NEXT GENERATION AI-ENABLED ROBOTICS

NVIDIA GTC

Bastiane Huang_Product @ Osaro

ROBOT EVOLUTION: Anthropomorphic?

HOW ARE AI-ENABLED ROBOTS DIFFERENT?

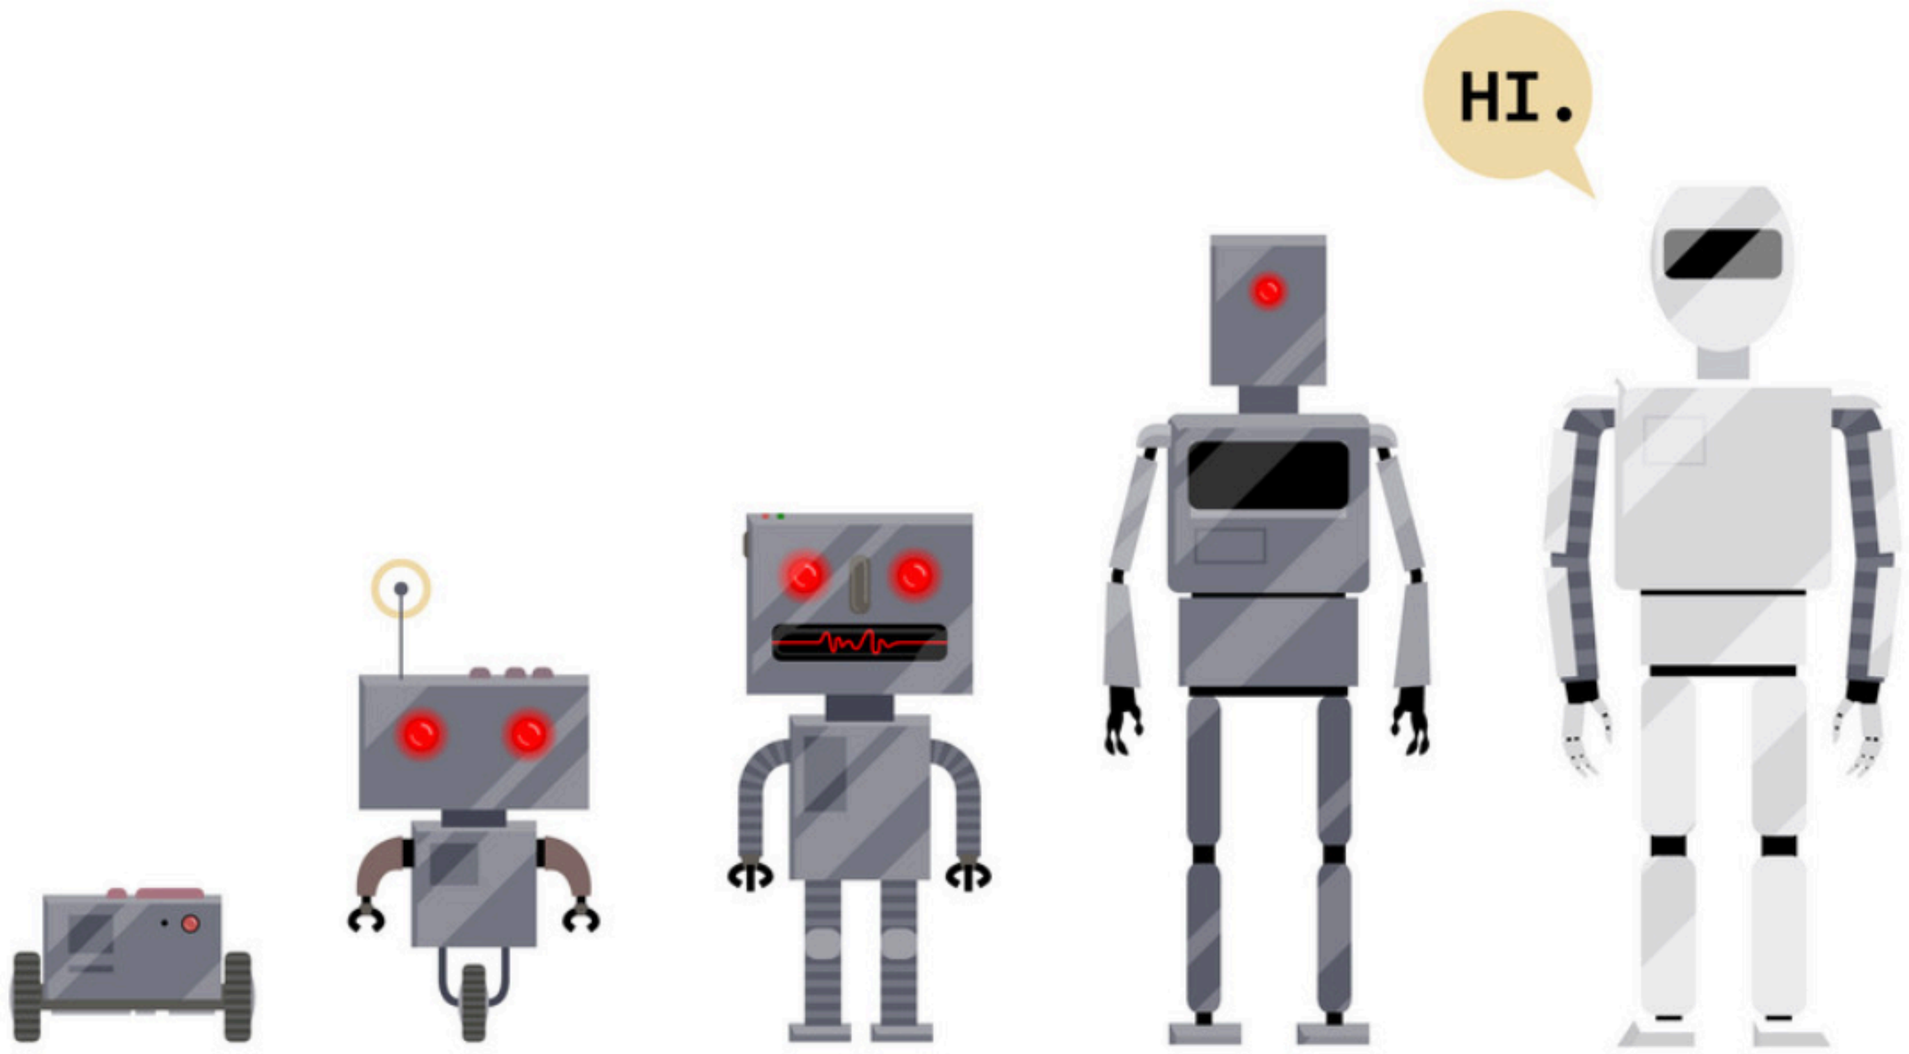

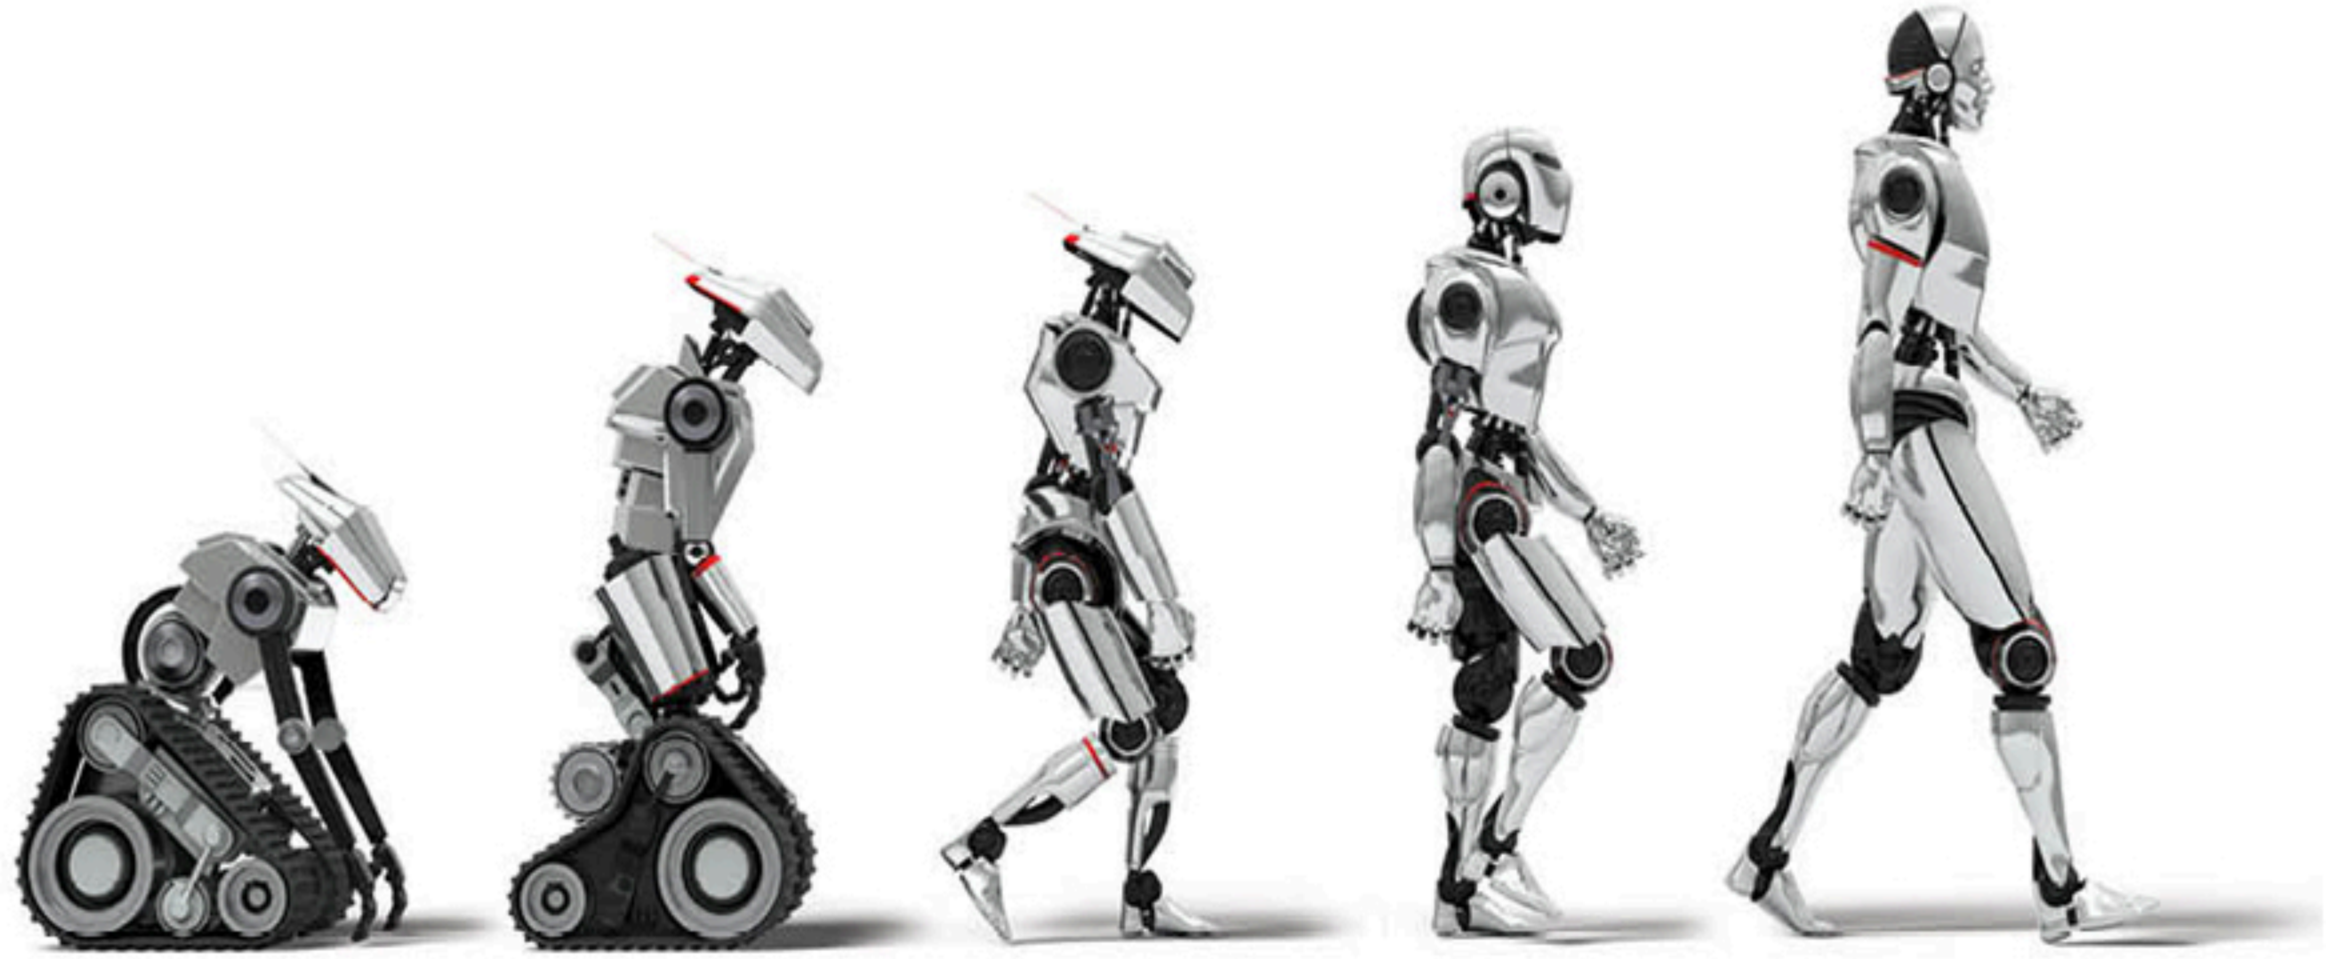

ROBOT 2.0:
NEXT GENERATION AI-ENABLED ROBOTICS

- Osaro Introduction

- AI-enabled Robotics
  - What, How, Why is it important, Why now?
  - Machine learning research areas
    - DL, motion planning, simulation
    - RL, IL/Behavioral cloning, Meta learning,
    - Other challenges

- Challenges and Opportunities

- Osaro's Approach

- Thoughts on the future

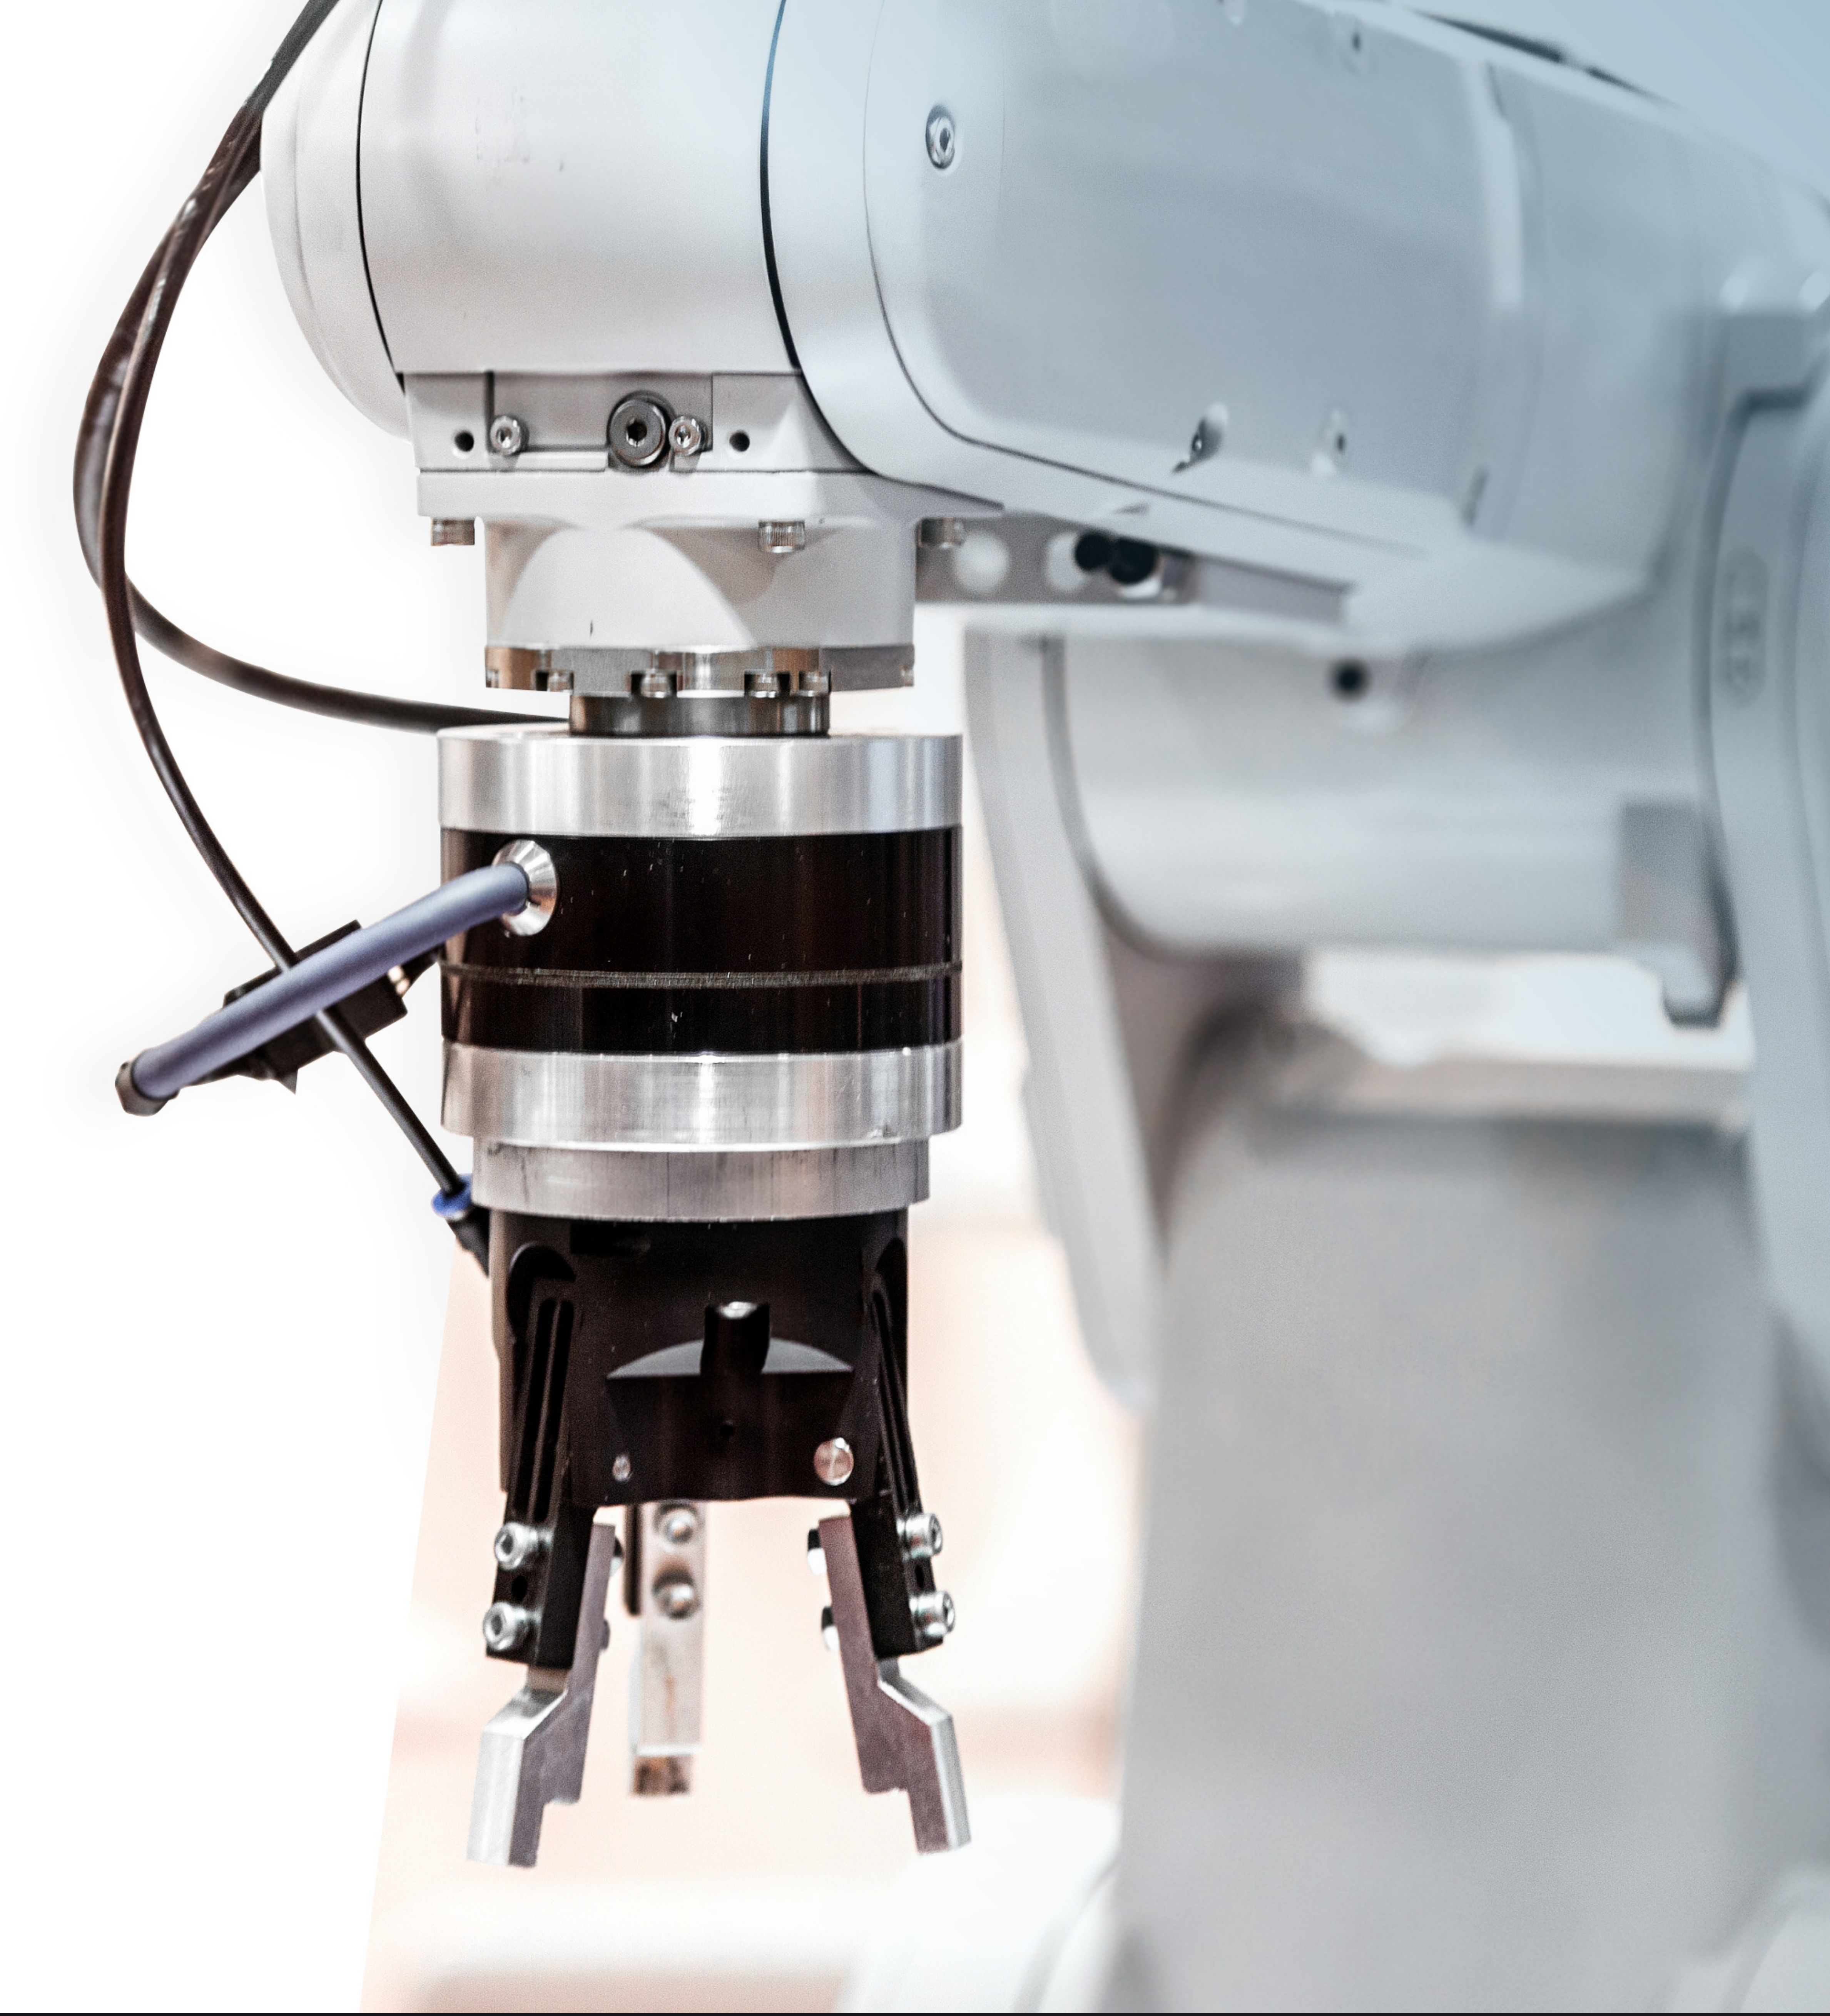

OSARO MISSION

Build machine learning software that enables robots to learn

advanced manufacturing

ecommerce fulfillment

food assembly

OSARO OVERVIEW

- Osaro builds brains for robots
- Based in San Francisco, founded in 2015
- 24 engineers, 6 business people
- Focus: Vision and control software for large scale robotic deployments (manufacturing, ASRS, food prep…)

Derik Pridmore
CEO

Led DeepMind Series A

Investor - Founders Fund

Founder - Arda Capital

Physics & CS - MIT

Michael Kahane
CTO

Serial Entrepreneur

Senior Engineer - Samsung

Elite R&D Unit - IDF

CS, EE - Ben Gurion

OSARO OVERVIEW

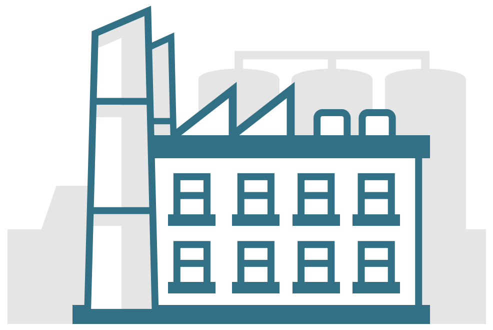

Product
Edge-deployed perception and control software.

First use case is automating piece picking in warehouses at scale.

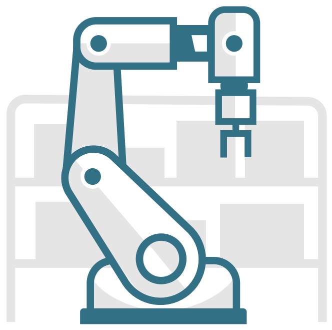

Business Model
We sell software and have a SaaS model. We have per robot recurring revenue.

Distribution Strategy
We distribute our software via system integrators.

Traction
We are working with 45% of our target market (materials handling integrators) and 6 top robotics companies.

ROBOT EVOLUTION: from automation (hard-programmed) to true autonomy (self-directed)

FROM AUTOMATION TO AUTONOMY

Level 0	Level 1	Level 2	Level 3	Level 4	Level 5
No Automation	Driver Assistance (single automated operation)	Partial Automation	Conditional Autonomy	High Autonomy	Full Autonomy

DEEP NETWORKS
LEARN REPRESENTATIONS

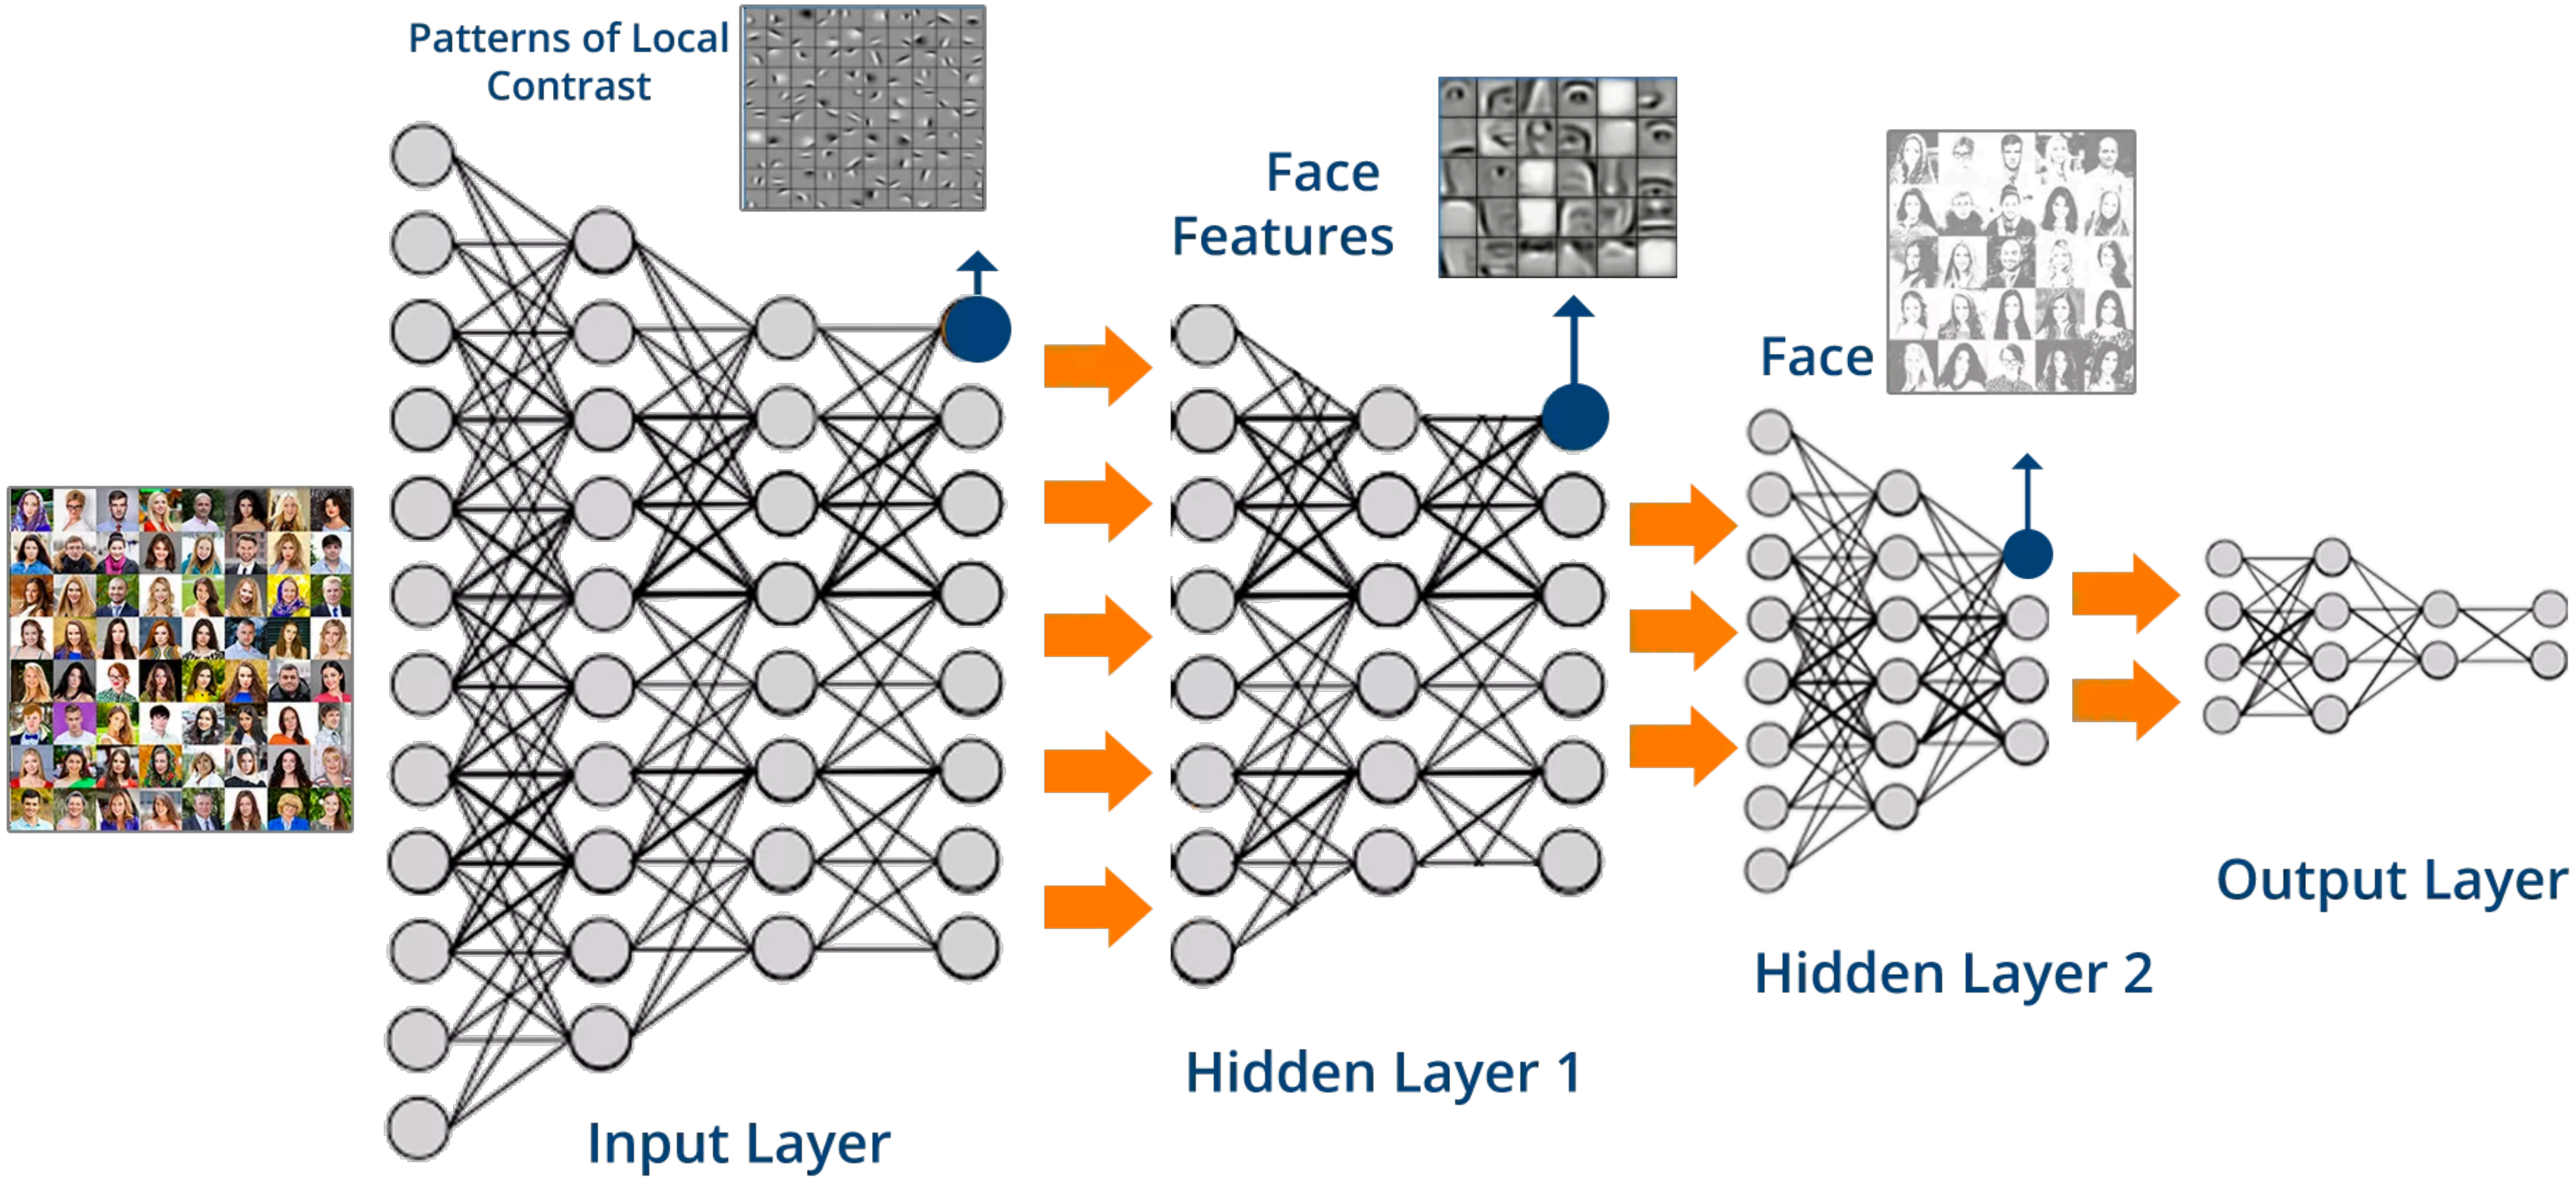

(Image from: https://cdn.edureka.co/blog/wp-content/uploads/2017/05/Deep-Neural-Network-What-is-Deep-Learning-Edureka.png)

REINFORCEMENT LEARNING: maximizing the long term return

LEARNS CONTROL POLICIES

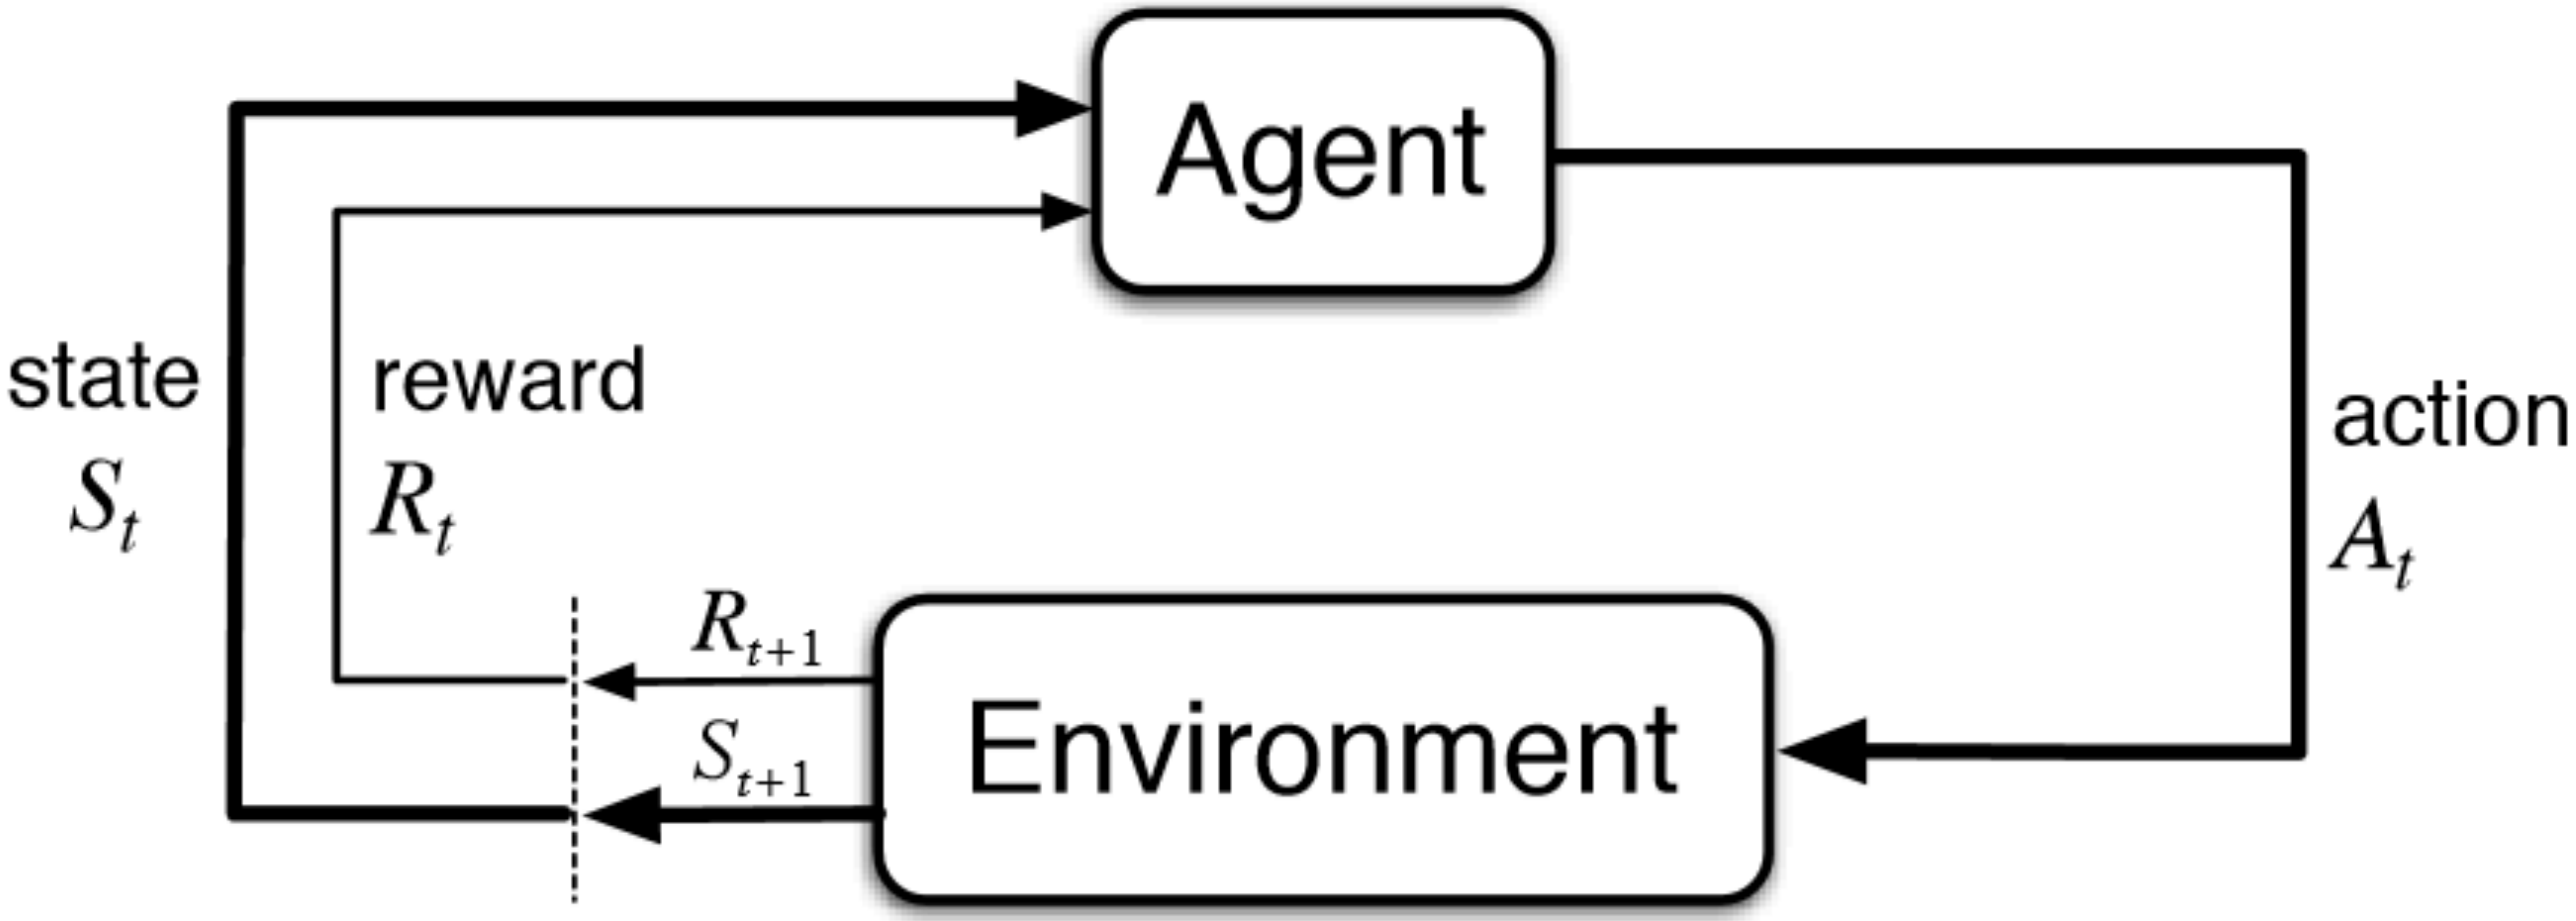

(Image from: Towards Data Science)

DEEP REINFORCEMENT LEARNING
COMBINES THESE TWO APPROACHES

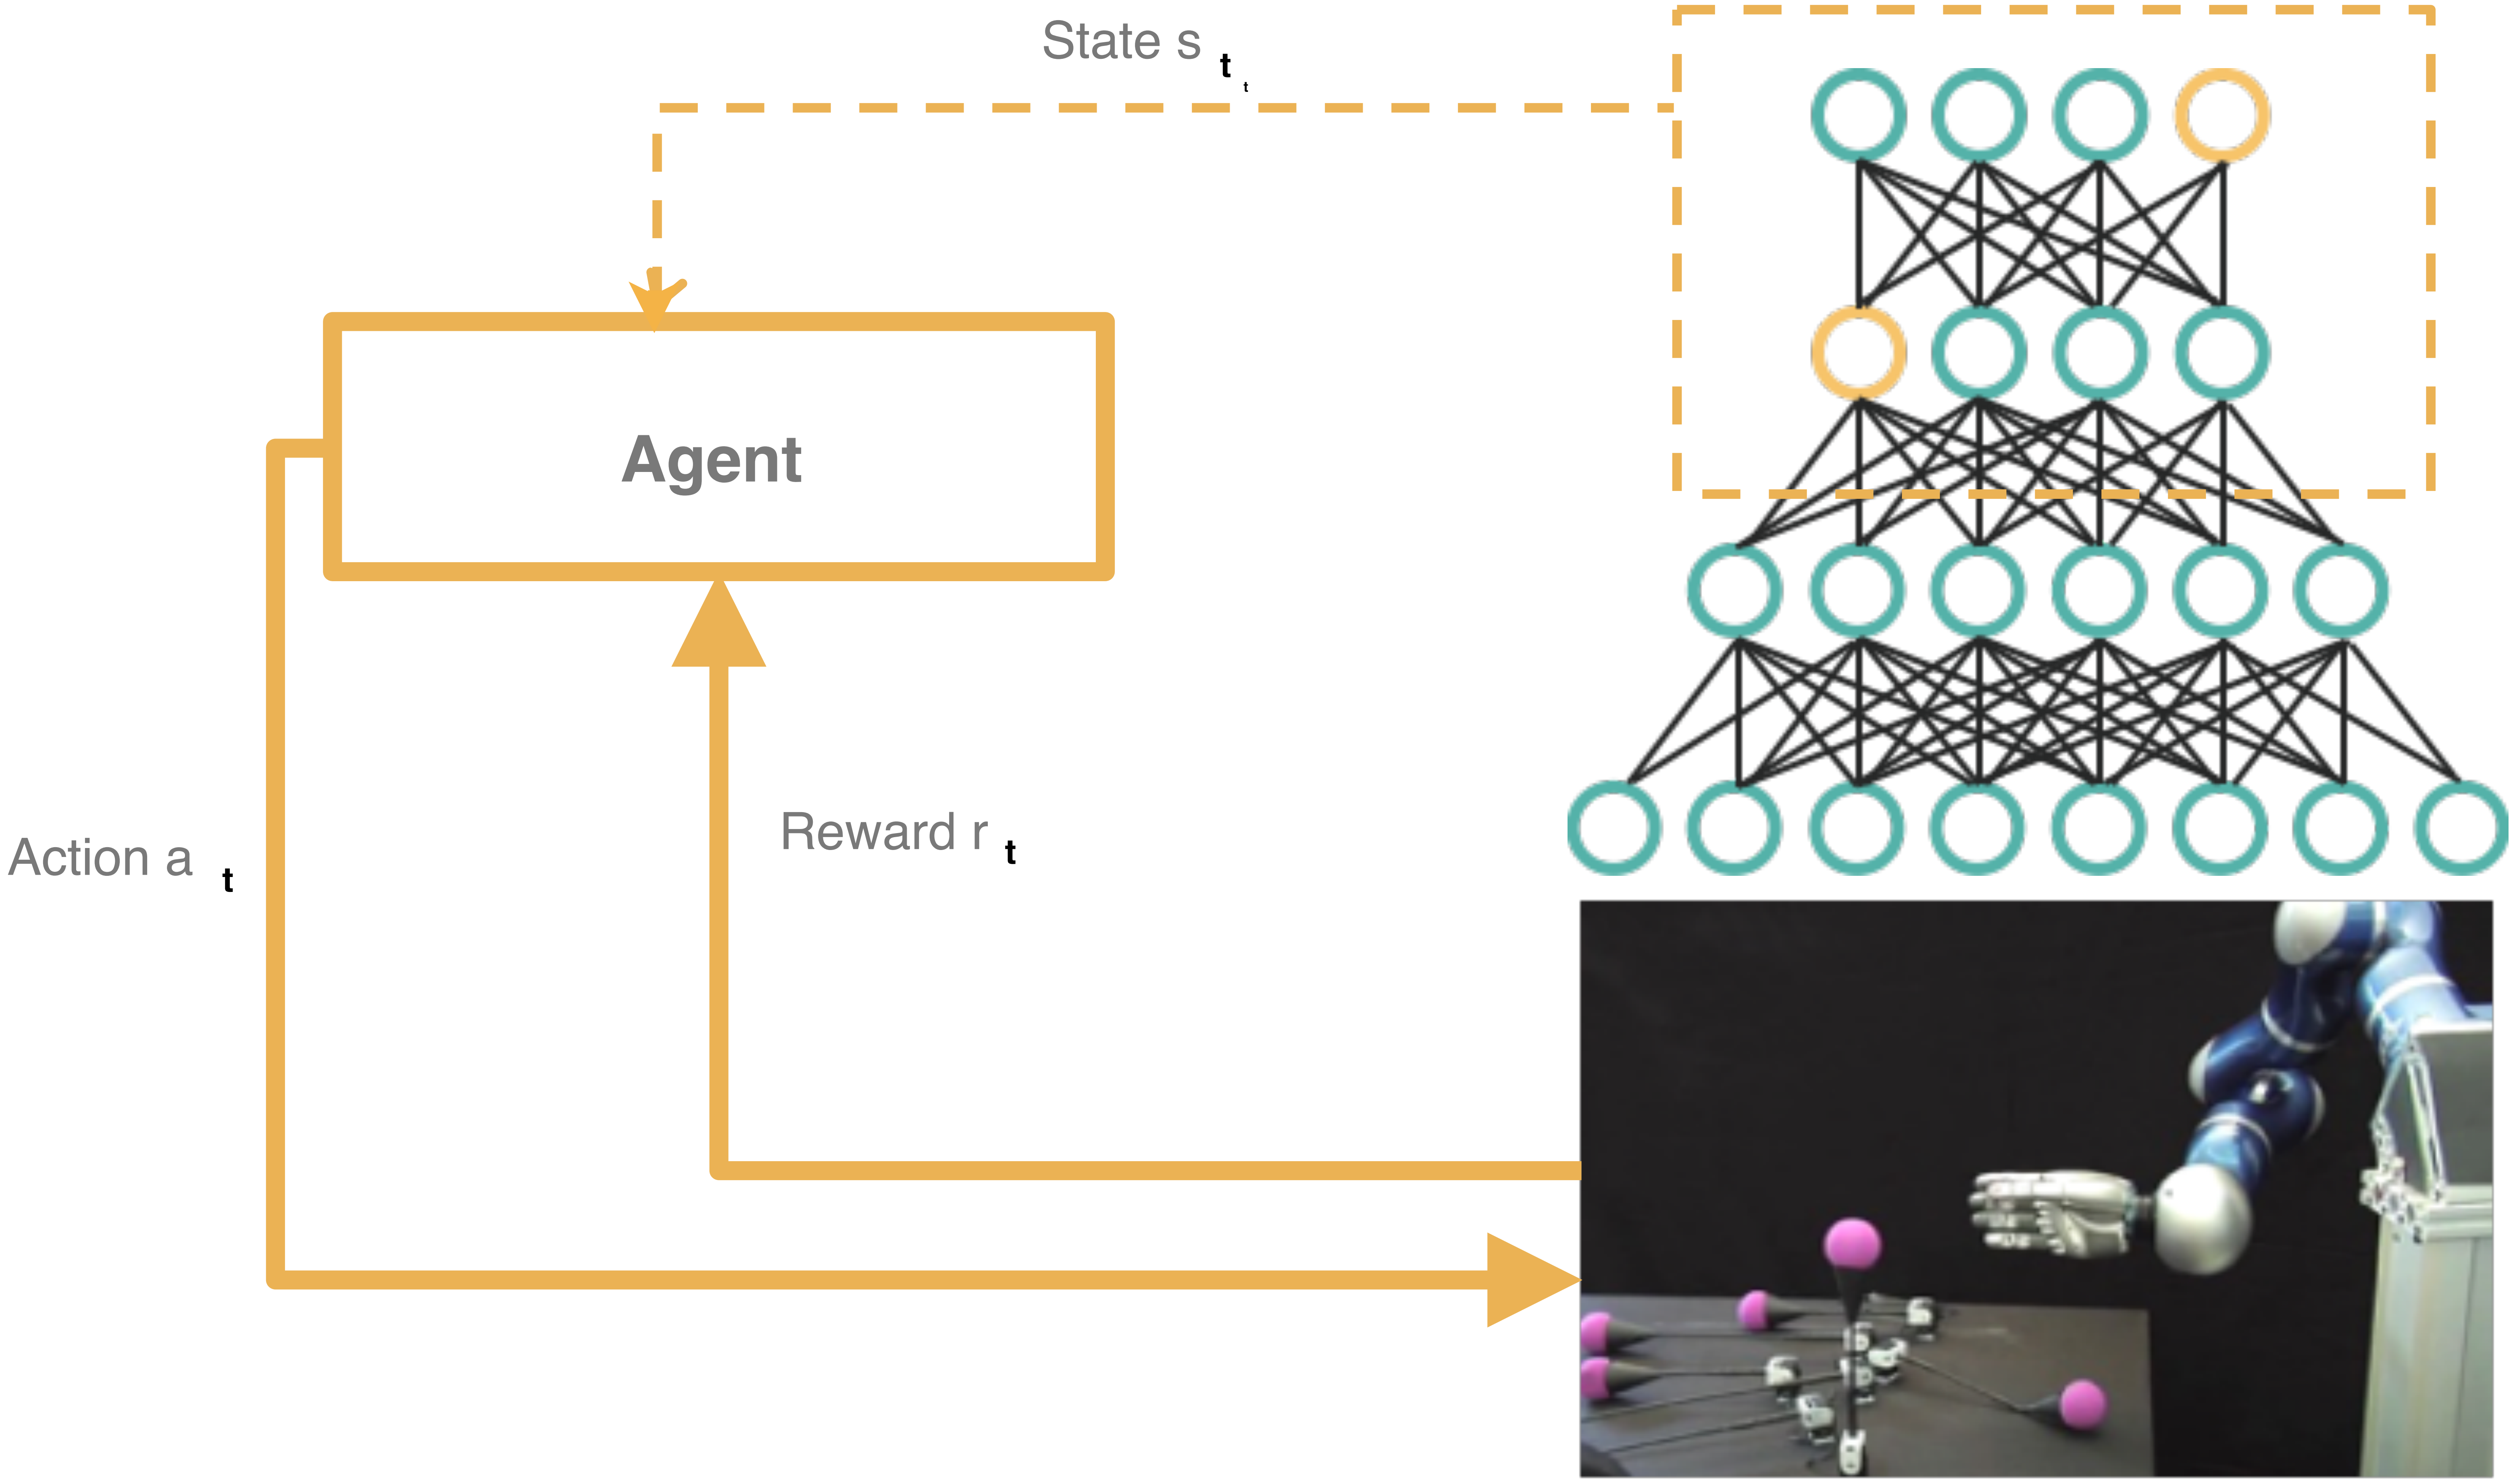

WHY IS IT IMPORTANT?

- Moving away from programming robots
- Addressing questions that are “hard to describe to computers”
- Handling variabilities and react to changes (dexterity, flexibility)
- Learning at scale (connected robots/cloud)

PICKING -

IN THE REAL WORLD

Why is it necessary?

Today’s problems have too much variability

Millions of products (Clear, Reflective, Flexible, Varying Surfaces)

Examples from e-commerce:

MACHINE LEARNING -

HELPS RECOGNIZE AND HANDLE CHALLENGING ITEMS

Why is it necessary?

Suction brush, not thin handle

Don't get fooled by **reflections**

Pick **clear box**, not what's inside

Avoid pump, **aim for body**

Target label, not gloves

Avoid contour, pick flat parts

WHY NOW?
ROBOT SALES SKYROCKETING

2017: robot sales increase 30%, record growth of industrial robots

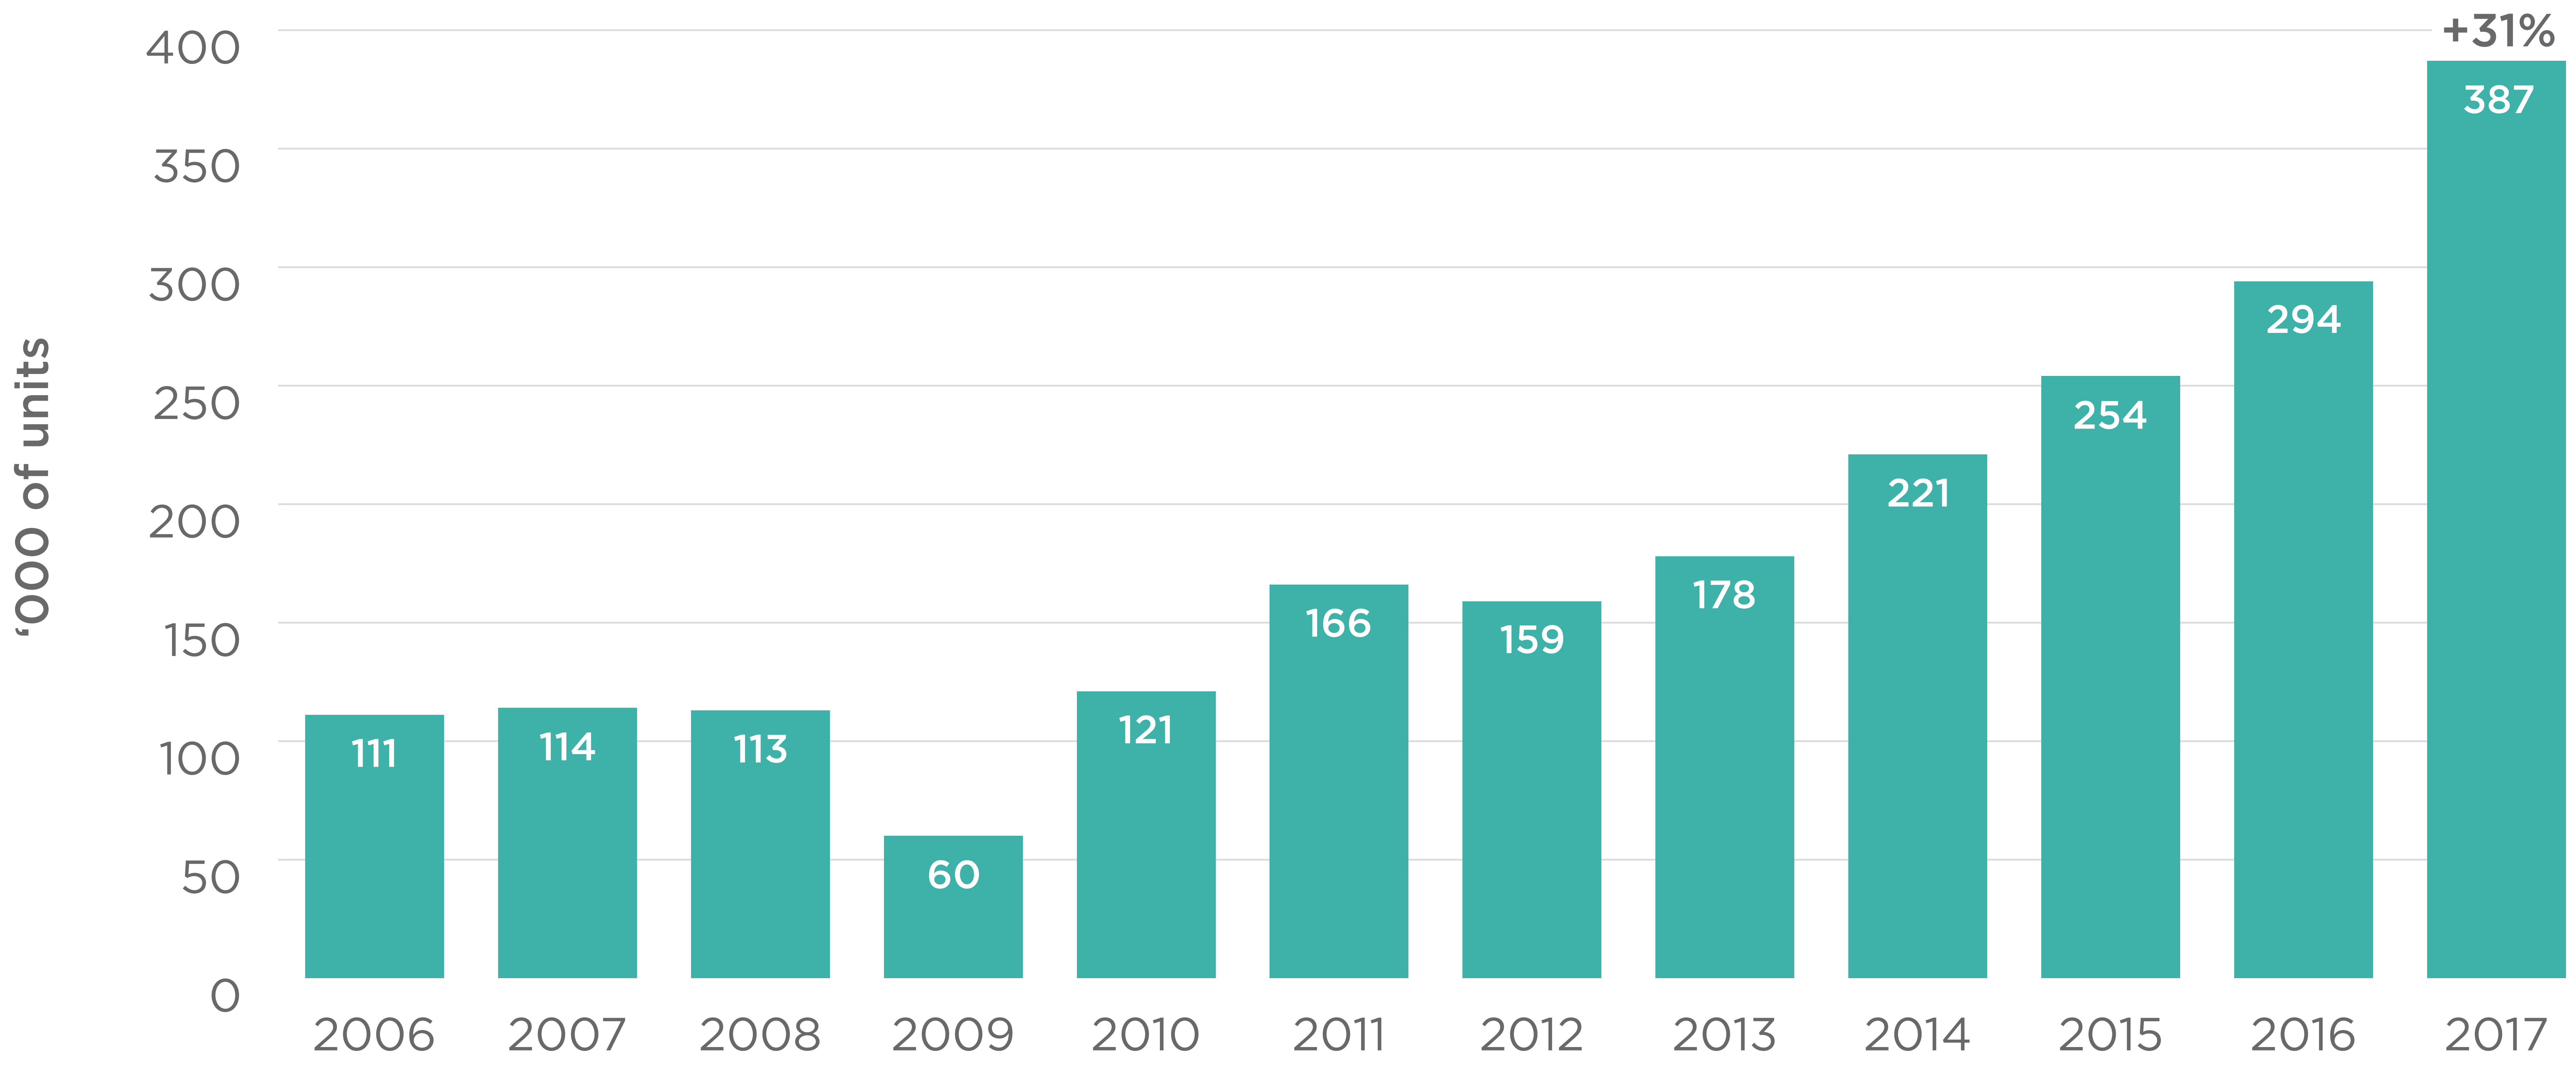

IFR International Federation of Robotics

Billions of robots in the future means 30% growth for 80 years

WHY NOW?

FALLING ROBOT PRICES

Robot production increased; costs go down

Robot price fell by half over the past 30 years

Demand from emerging economies

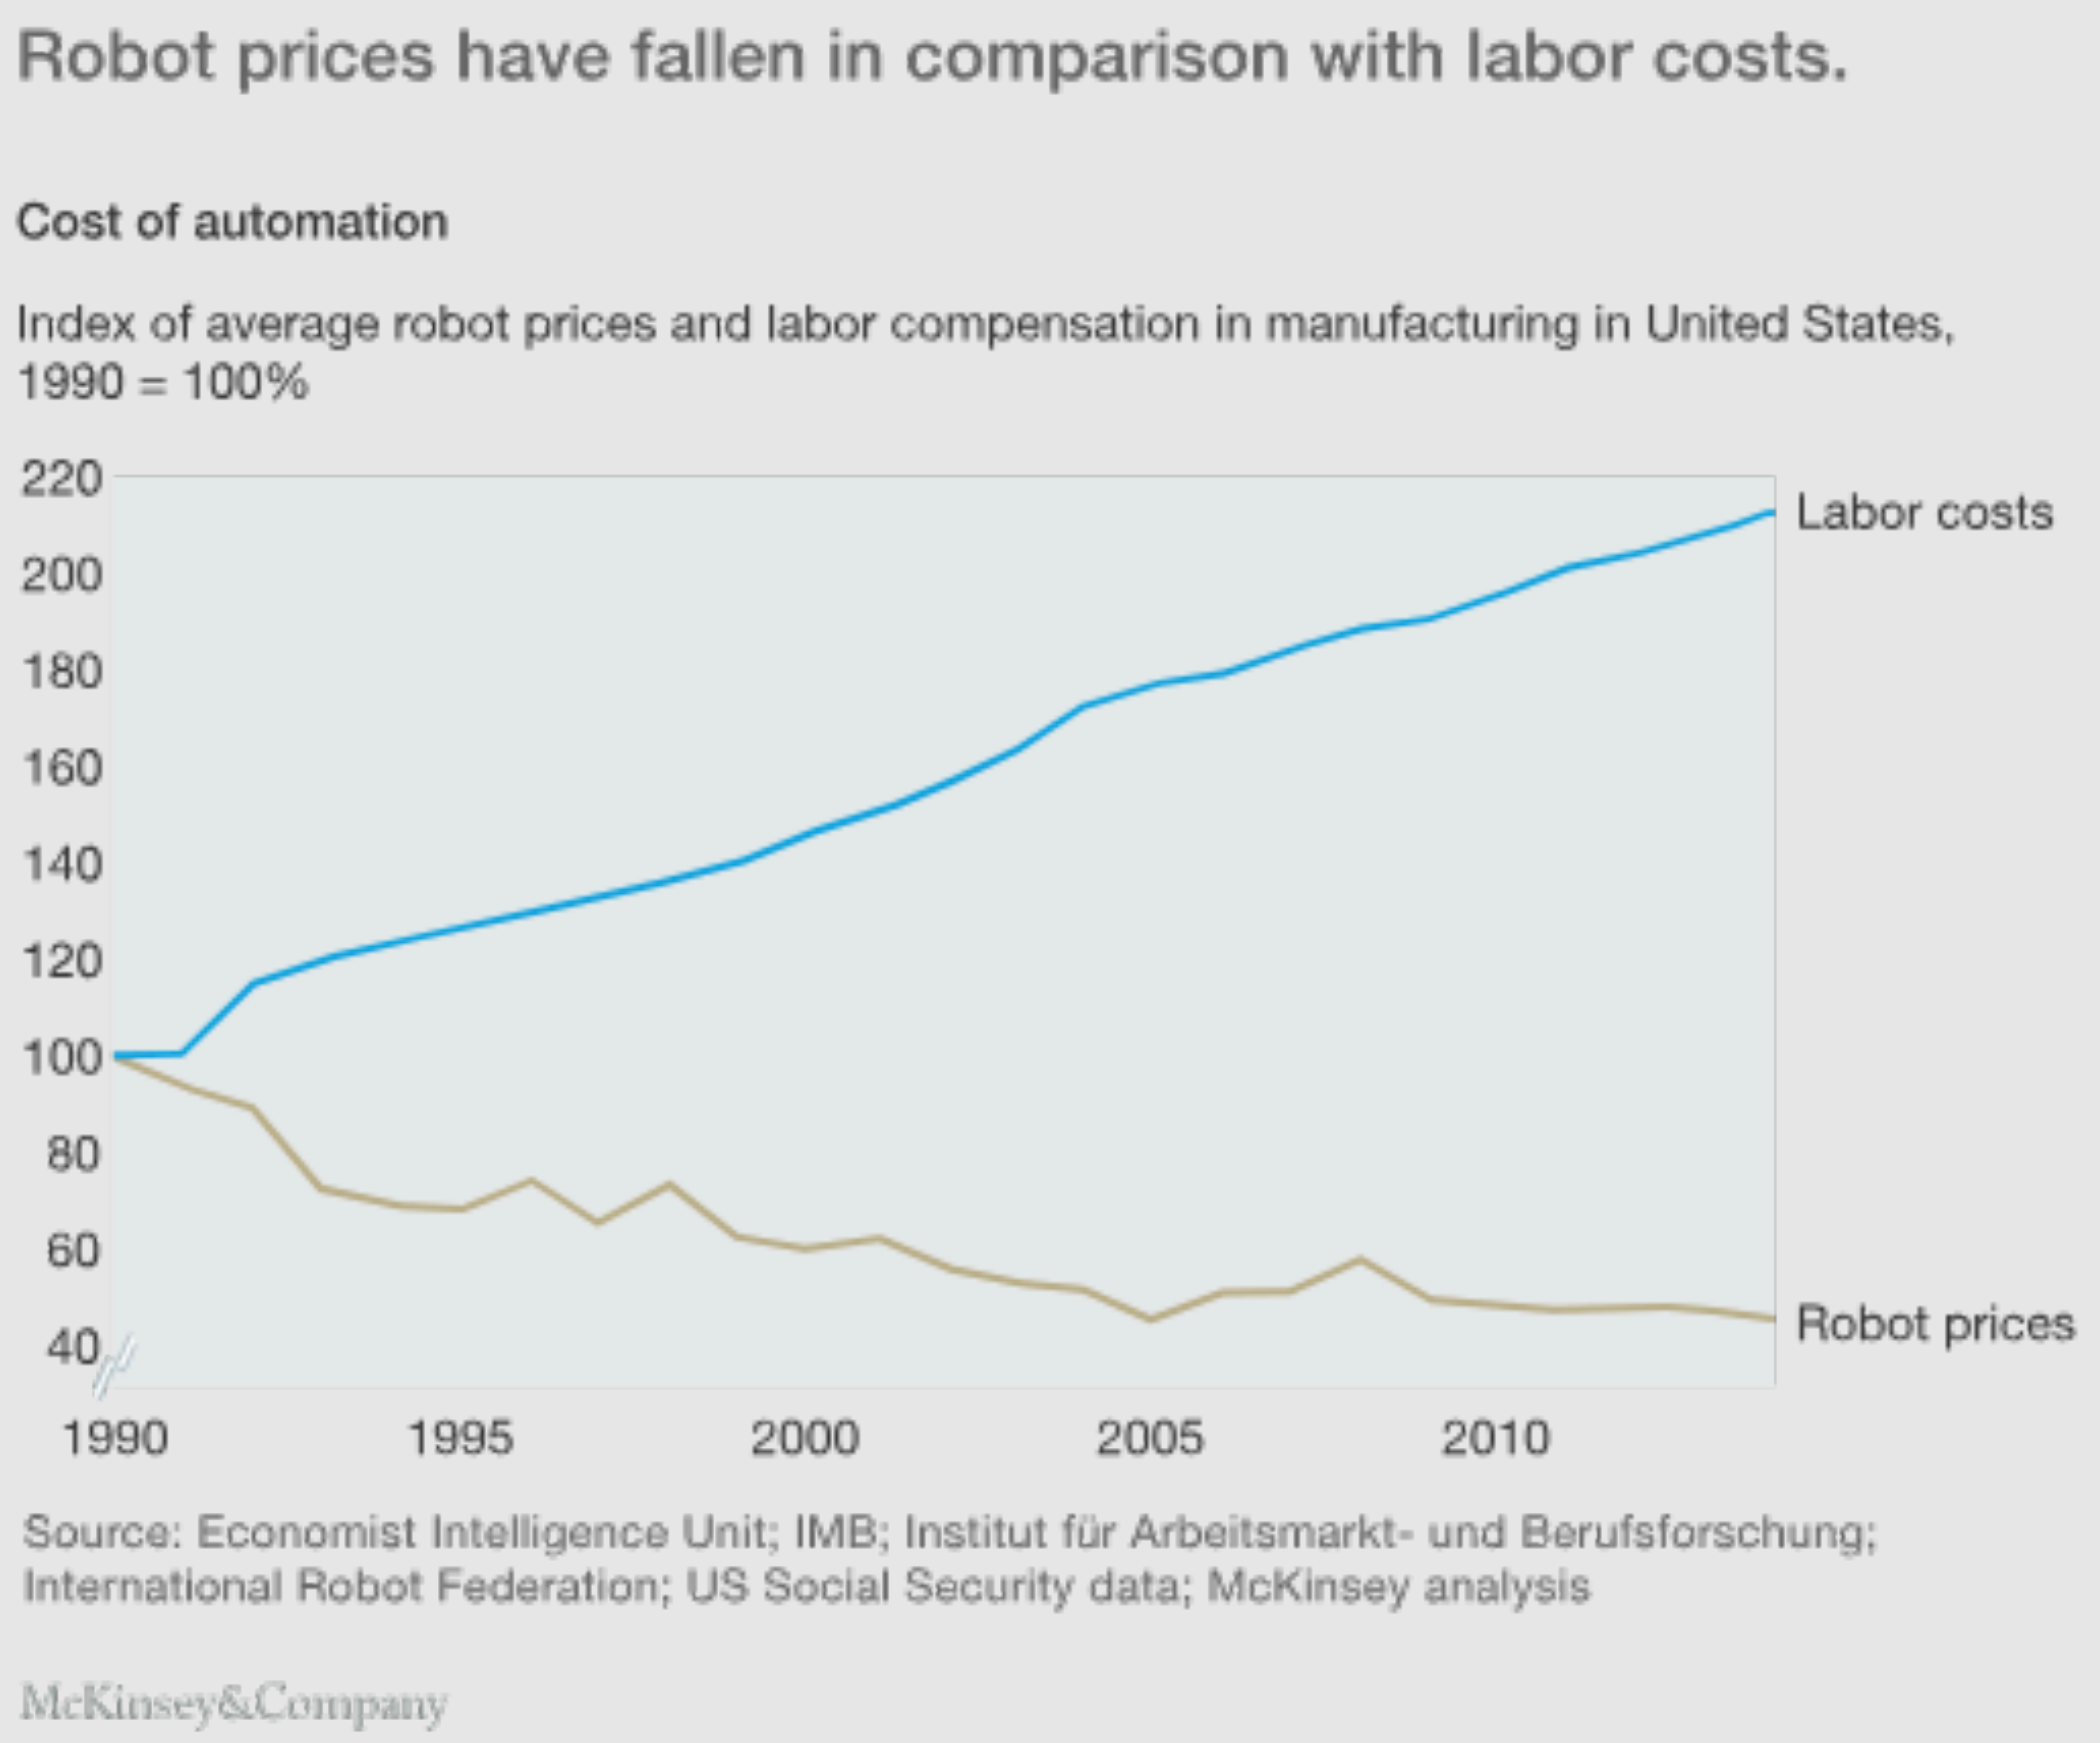

WHY NOW? GLOBAL PROBLEM:

SHRINKING LABOR FORCES NEED AUTOMATION

Global labor shortages

Acute in extremely monotonous jobs

Variability in tasks, difficulty of items

Osaro targets countries with:

→ Declining Labor Force

→ High Labor Costs

→ Ex: Japan, Germany, Australia, etc.

Japan's working-age population is set to decline at an even faster pace than the overall population and more swiftly than that of other advanced economies.

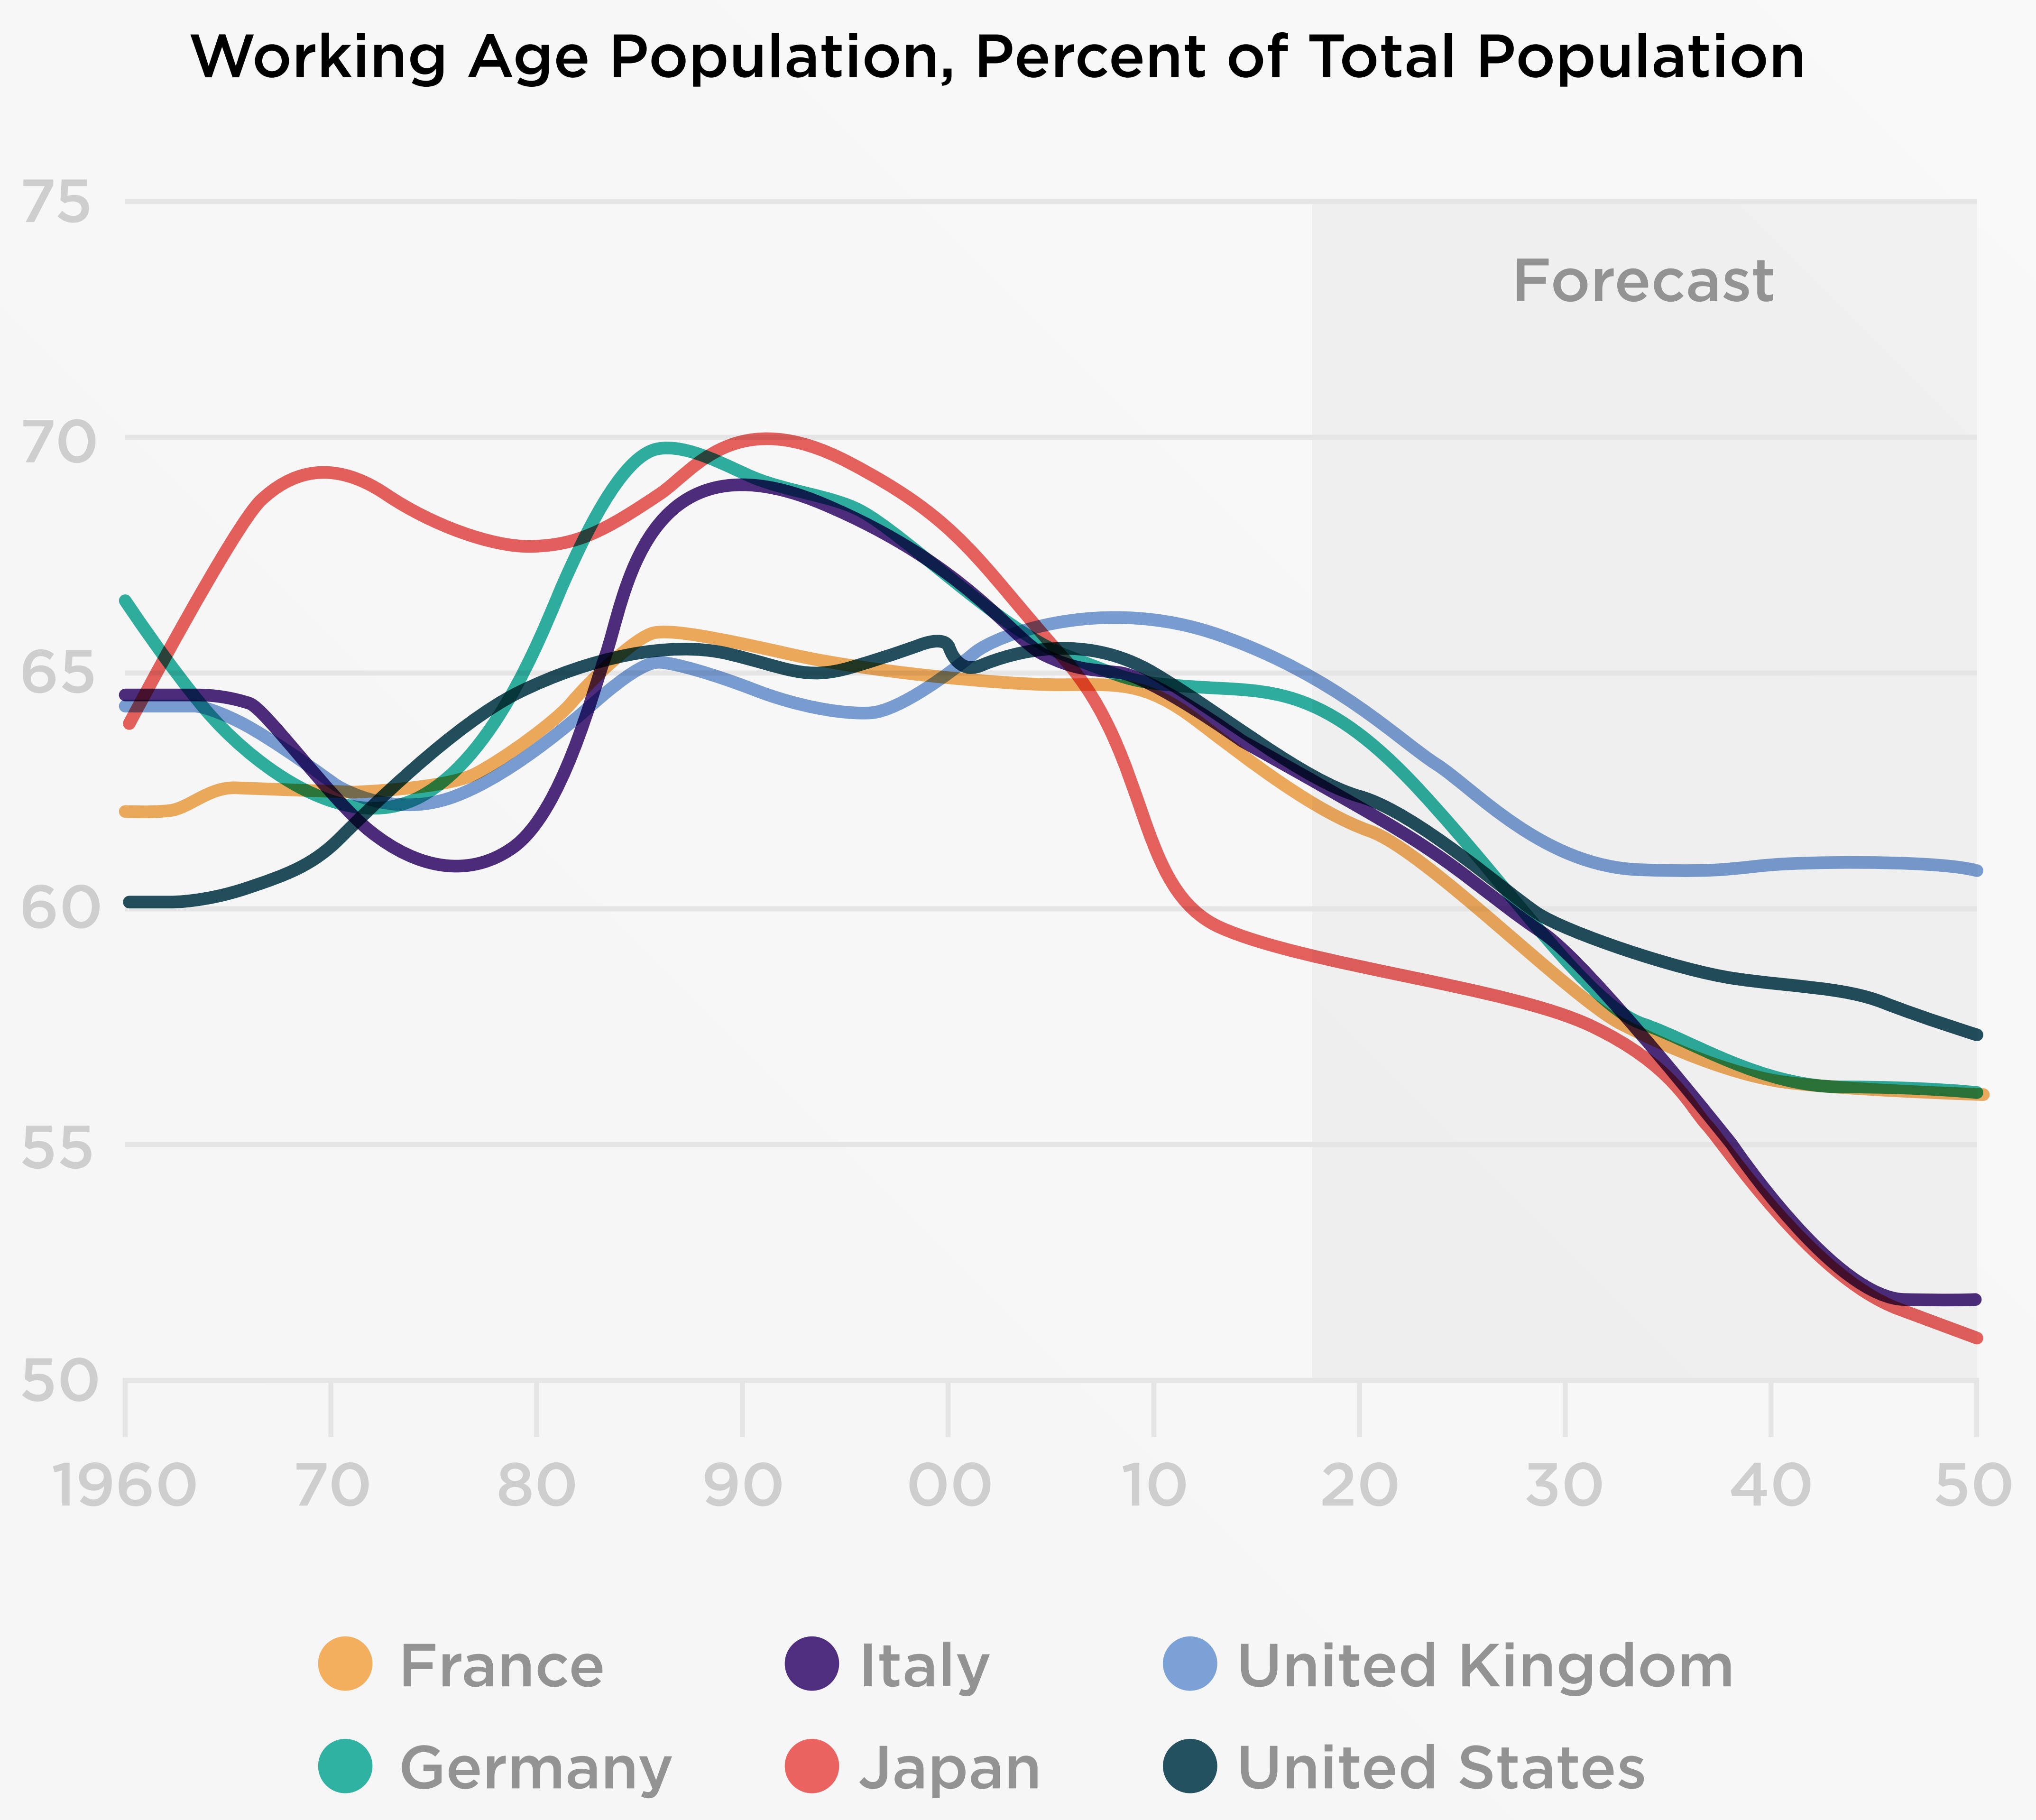

Source: United Nations, World Population Prospects, 2017 revision.

WHY NOW? FOCUSED PROBLEM:

E-COMMERCE FULFILLMENT

Shortages are Acute in E-commerce

Warehouse picking jobs:

→ Costly

→ Hard to Fill and Retain

→ Highly Repetitive

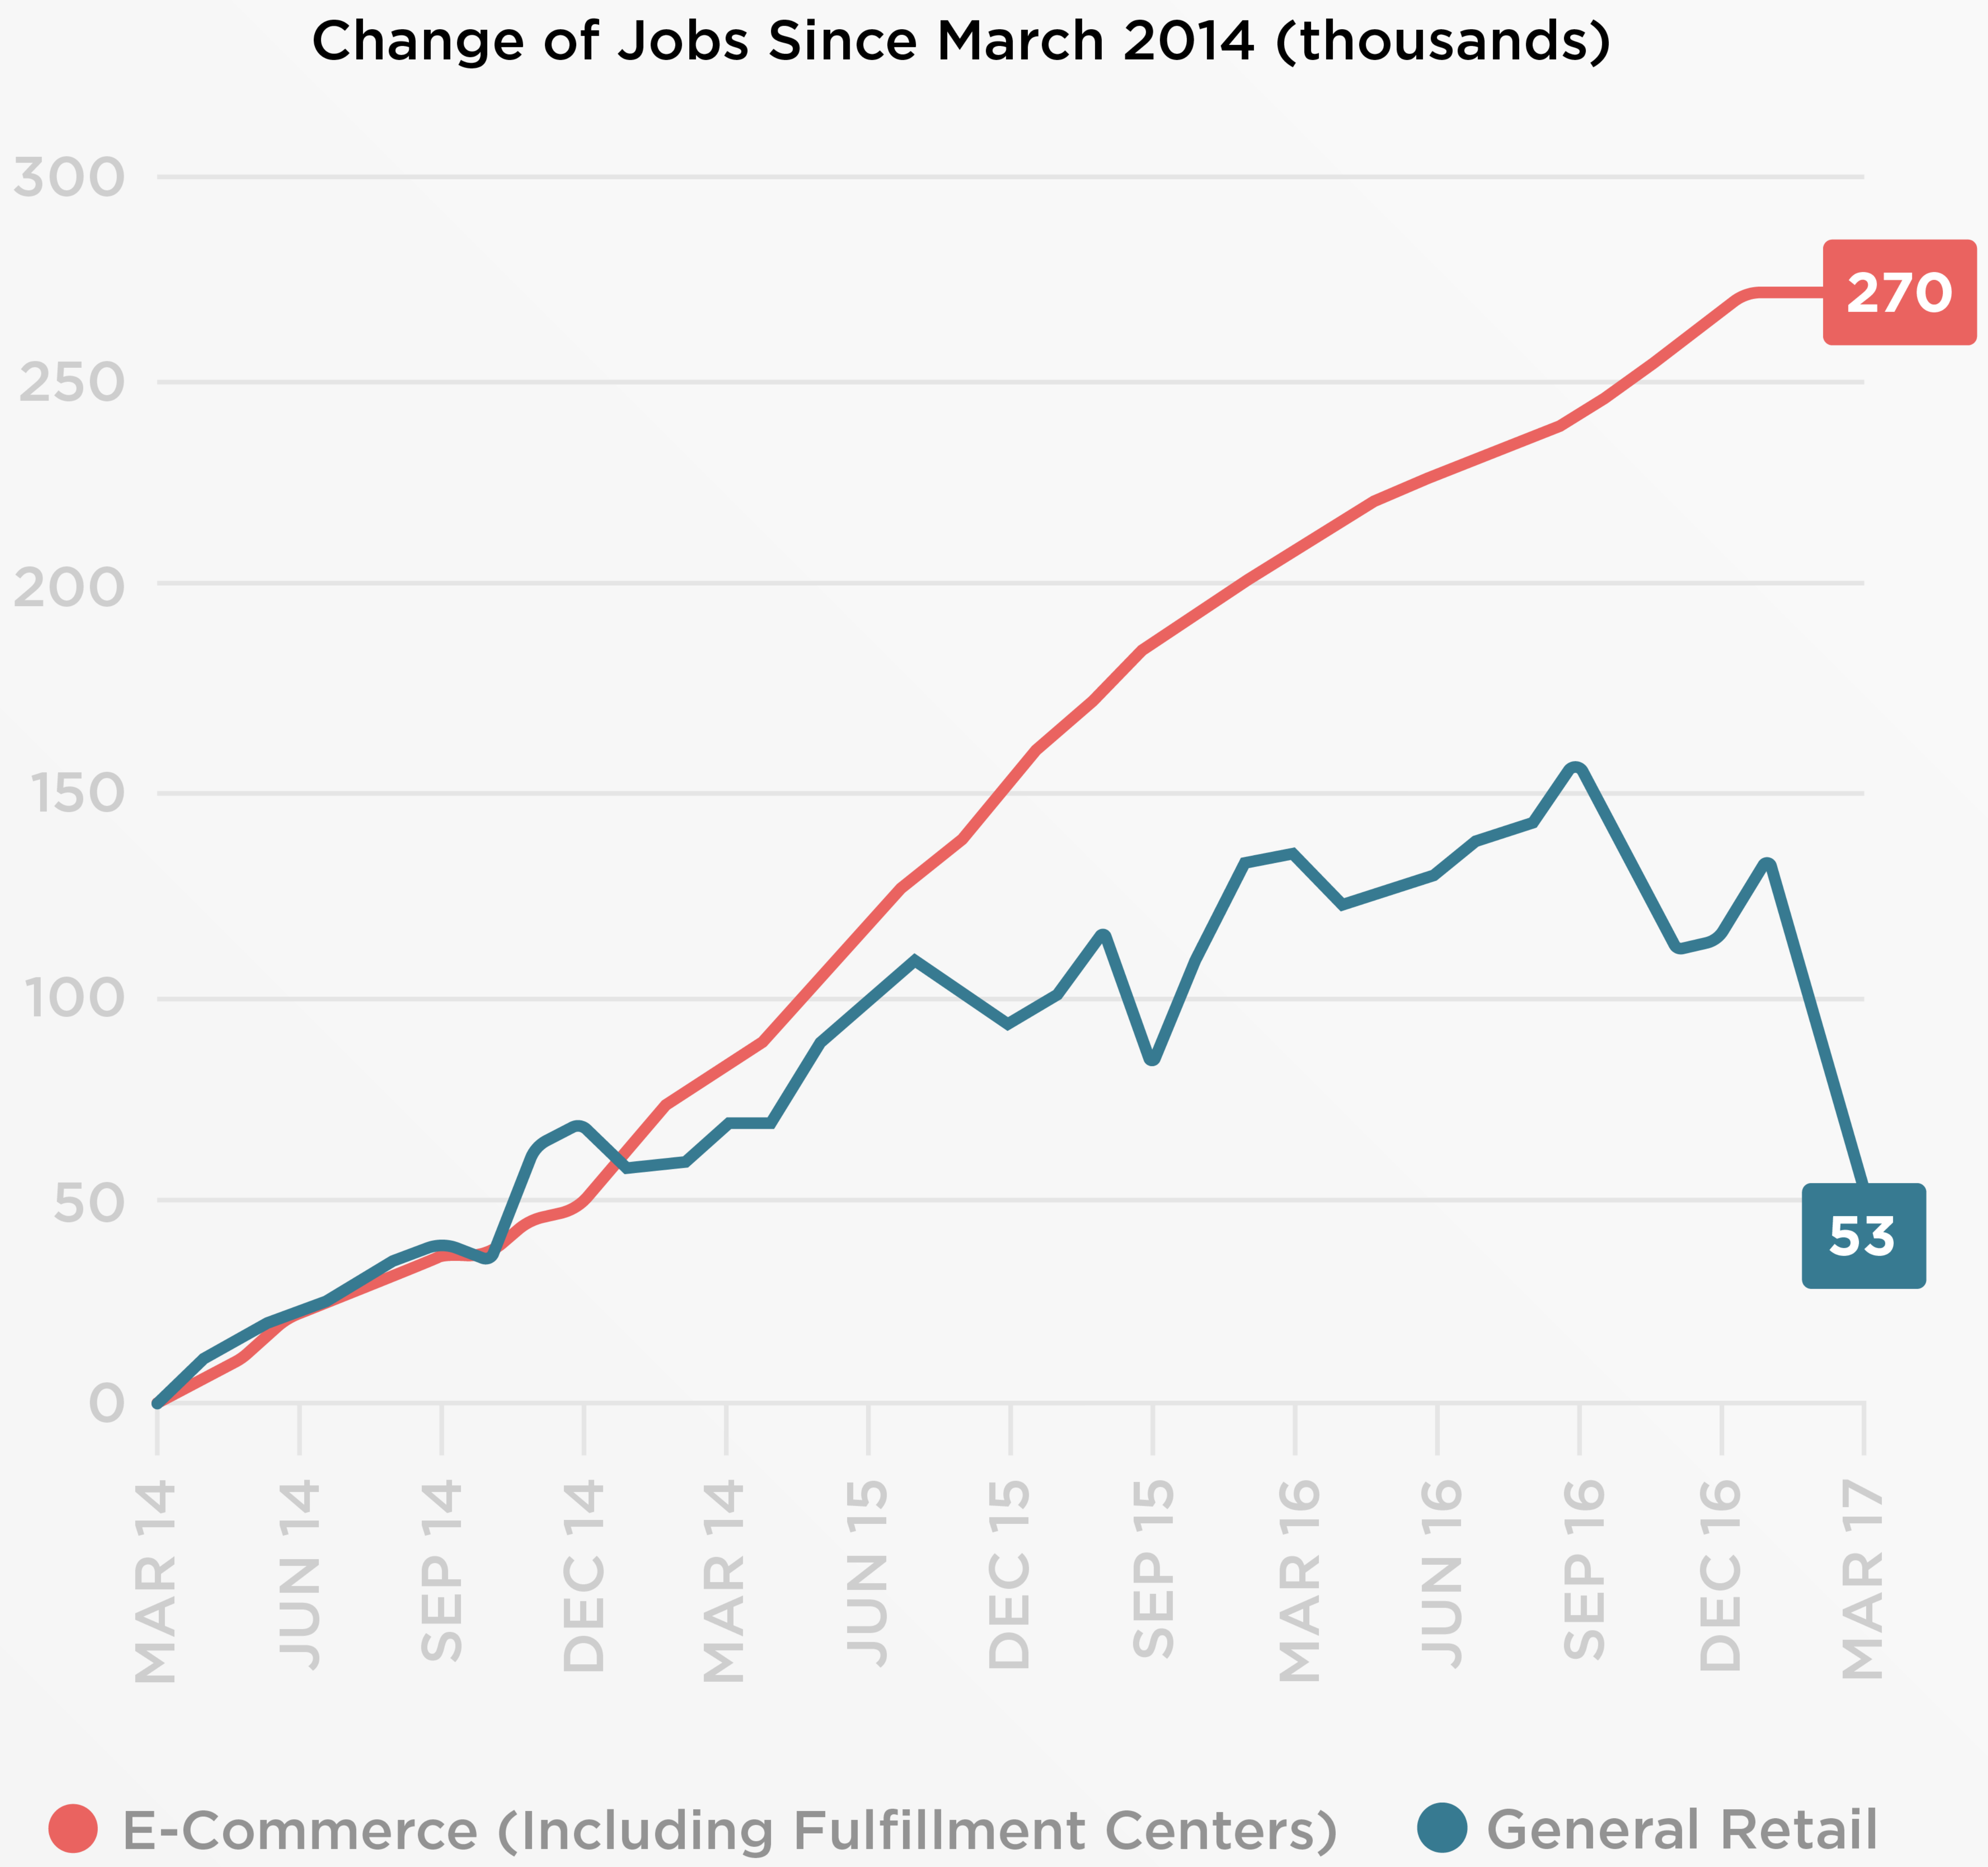

Data: BLS, Center for Emerging Employment (PPI)

CASE STUDY
WAREHOUSE PIECE PICKING

Picking: a mind is a terrible thing to waste

“Employee turnover rates currently stand at 13.7% for distribution, warehouse and manufacturing functions.”

“U.S. businesses lose $11 billion annually due to employee turnover.”

“As many as half of all hourly workers leave new jobs within the first 120 days.”

“Direct costs to replace an employee can reach as high as 50% to 60% of an employee’s annual salary.

WAREHOUSE PIECE PICKING (LEVEL 4 AUTONOMY)

https://www.youtube.com/watch?v=XbxiyYVSJ2M

OSARO AUTOMATES PICKING AND PLACING TASKS (ASRS)

OsaroPick **automates stationary picking stations** in "goods to person" distribution centers.

This is done by integrating with automated storage and retrieval systems (ASRS) for Level-4 automation (**no human in the loop**).

OSARO:
SCALABLE SOFTWARE-BASED PICKING SOLUTION

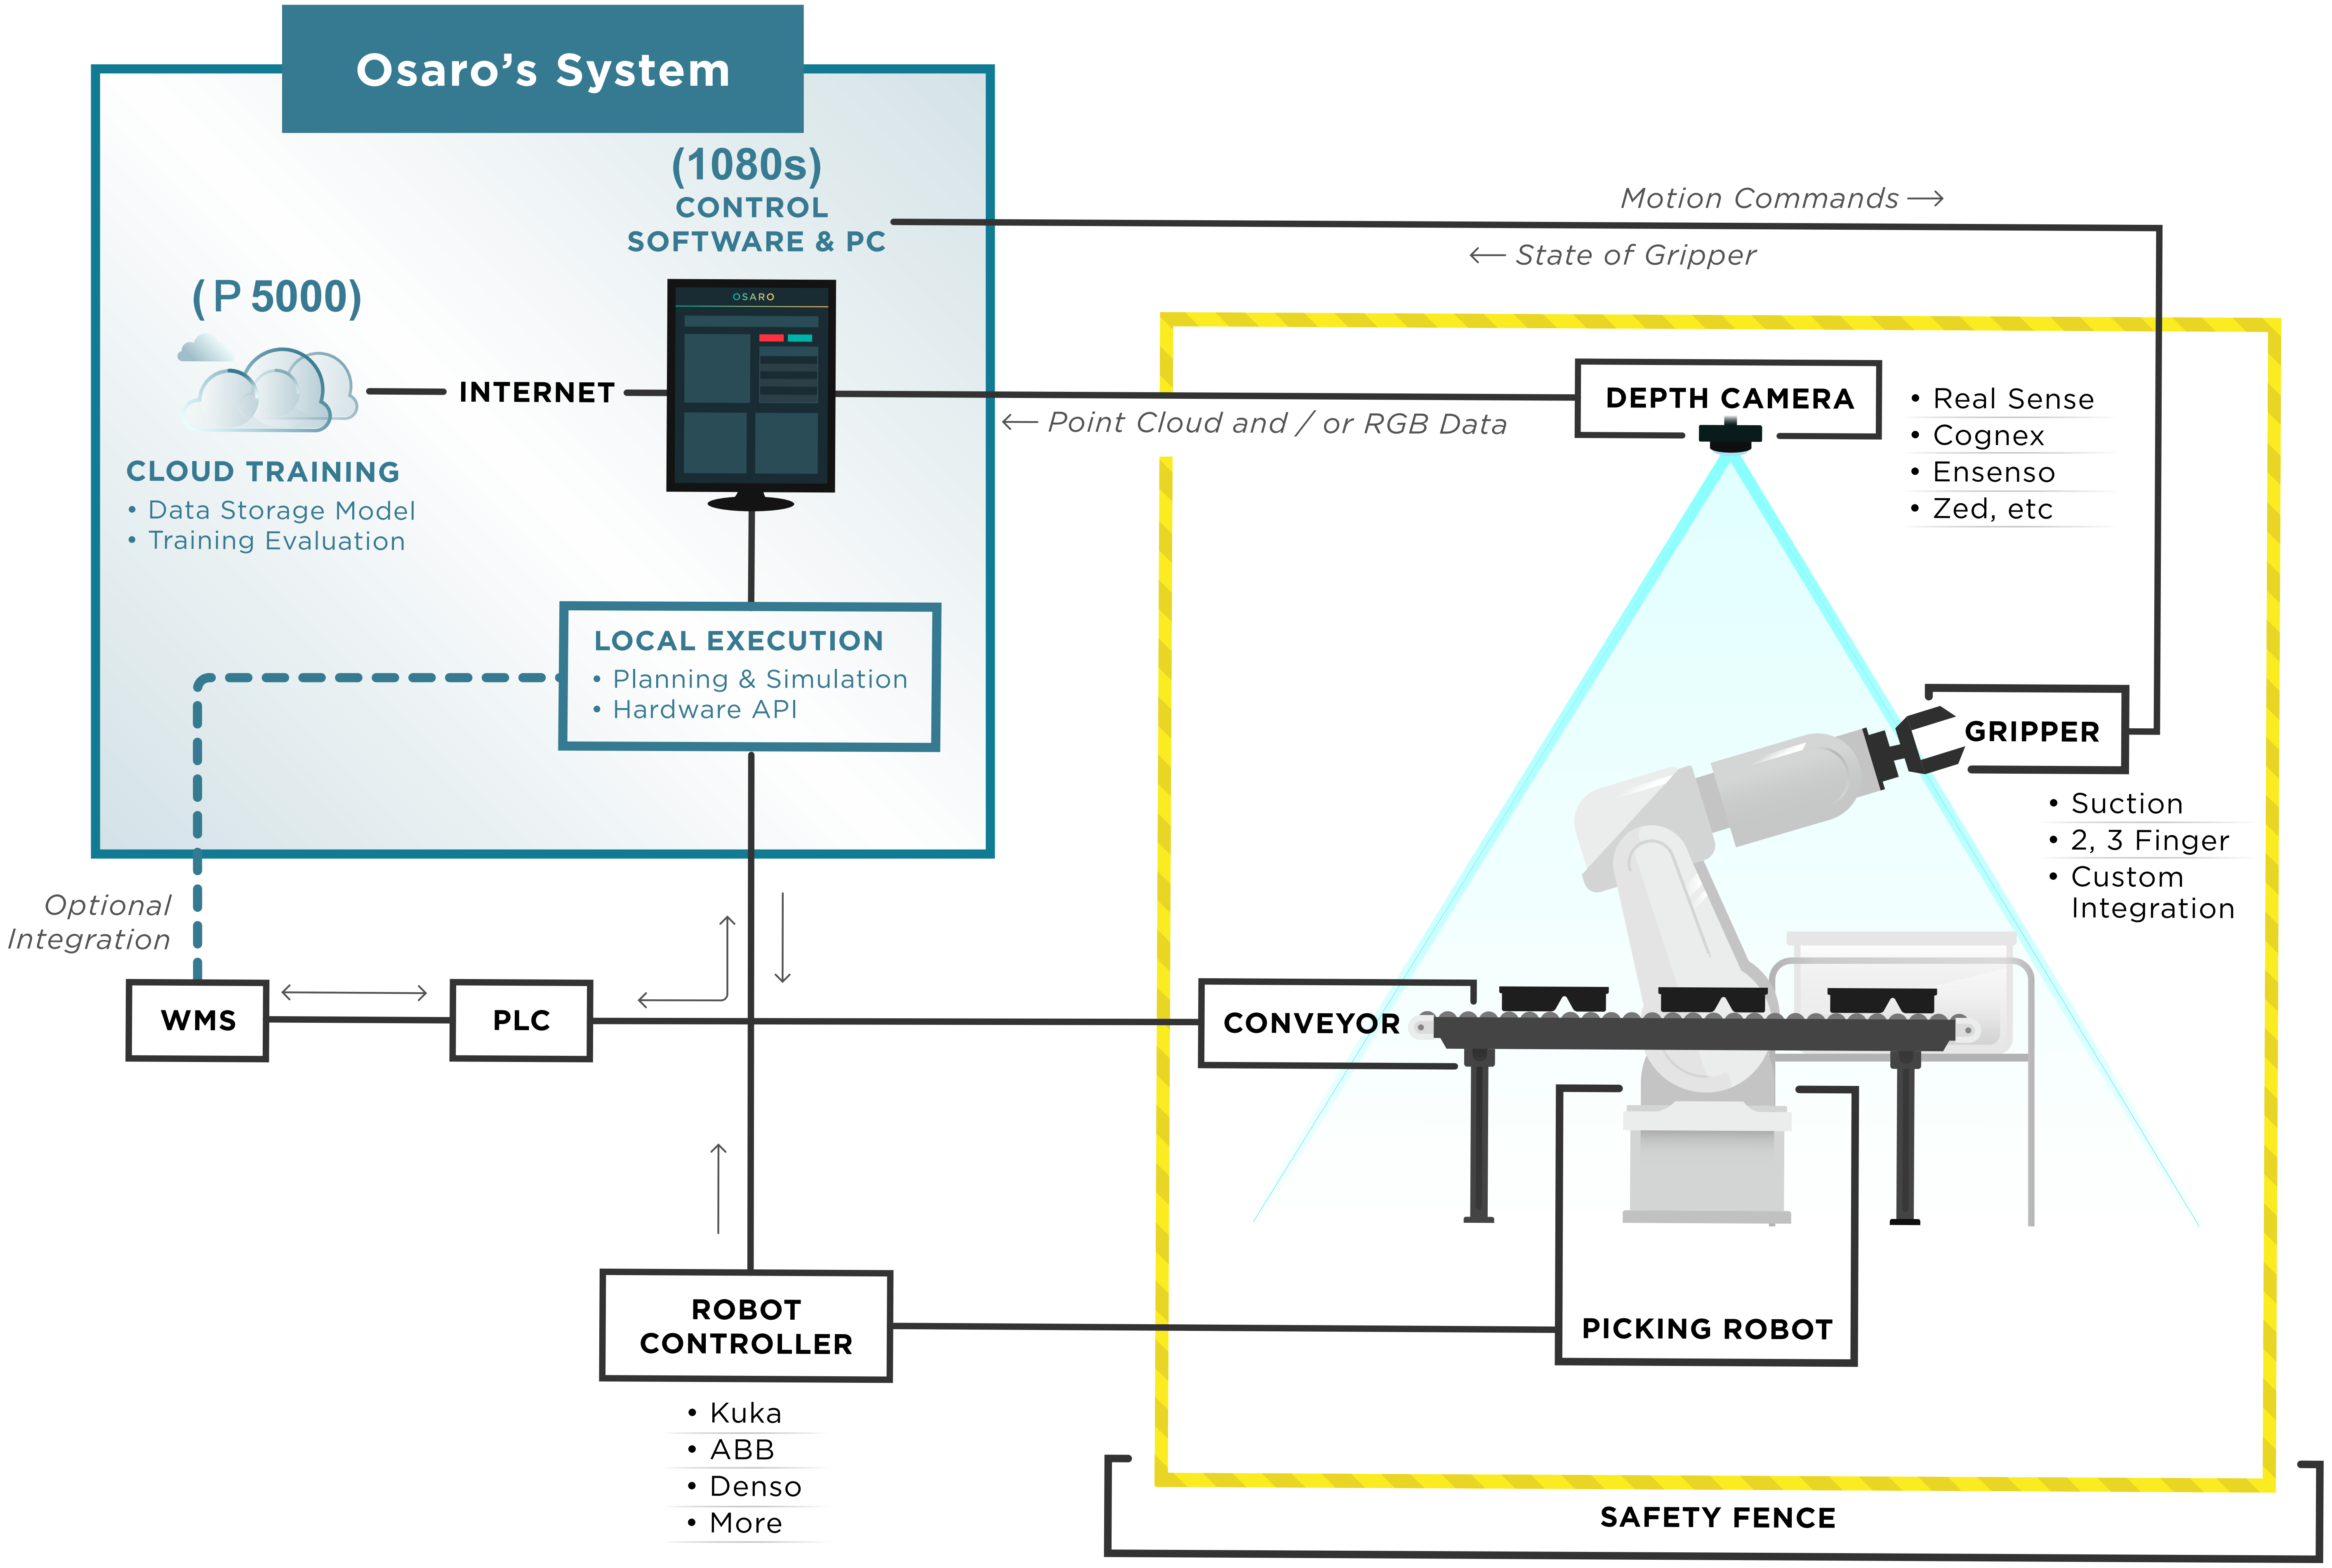

MACHINE LEARNING

WHAT, WHY, AND HOW

How is it applied? Idealized process:

- **Collect Data**
- **“Train a model”**
- **Test its accuracy, generalization**
- **Repeat**

MACHINE LEARNING RESEARCH AREAS

DEEP LEARNING, MOTION PLANNING, SIMULATION

	What is it?	Why use it?	Issues
Deep Learning	Deep learning methods exploit the unknown structure in the data distribution to discover rich representations using deep neural networks.	- Powerful in supervised learning problems - Success across a wide range of problems without domain understanding for feature introspection	- Large amount of labled data - Current success mostly in supervised learning problems - Interpretability - High end infrastructure
Motion Planning	A technique to determine a path between two given points, provided the kinematic and dynamic limits of the robot and the kinematic constraints of the environment. Sampling, Gradient/Optimization based; Learning from experience; imitation learning - motion planner learn initial policy	- Operate the robot optimally in terms of speed, accuracy and smoothness while avoiding obstacles in the environment. - Safe/collision free robot path planning	- Optimization can be time consuming (tradeoff pick time) - Optimal planning is NP-hard - Curse of dimensionality - Sensitive to fast environment changes
Simulation	Sim to Real. Data augmentation, Synthetic Data - Use realistic simulated data Domain Randomization (learn tasks and model arbitrary changes)	- Data efficient - Less expensive/dangerous - Faster/more scalable - Easier to label	- Requires an accurate model of the problem - Simulation to real is not trivial

REINFORCEMENT LEARNING

What is it?	Why use it?	Issues
Class of algorithms that aim to solve the sequential decision making problems.	- Minimal supervision - Autonomously learns by interacting with the environment	- Very large sample complexity (expensive and not feasible in real world domains) - Current success limited mostly to simulated environments such as games - Stability issues when used simultaneously with representation learning.

OSARO EARLY WORK - MARCH 2016

IMITATION LEARNING

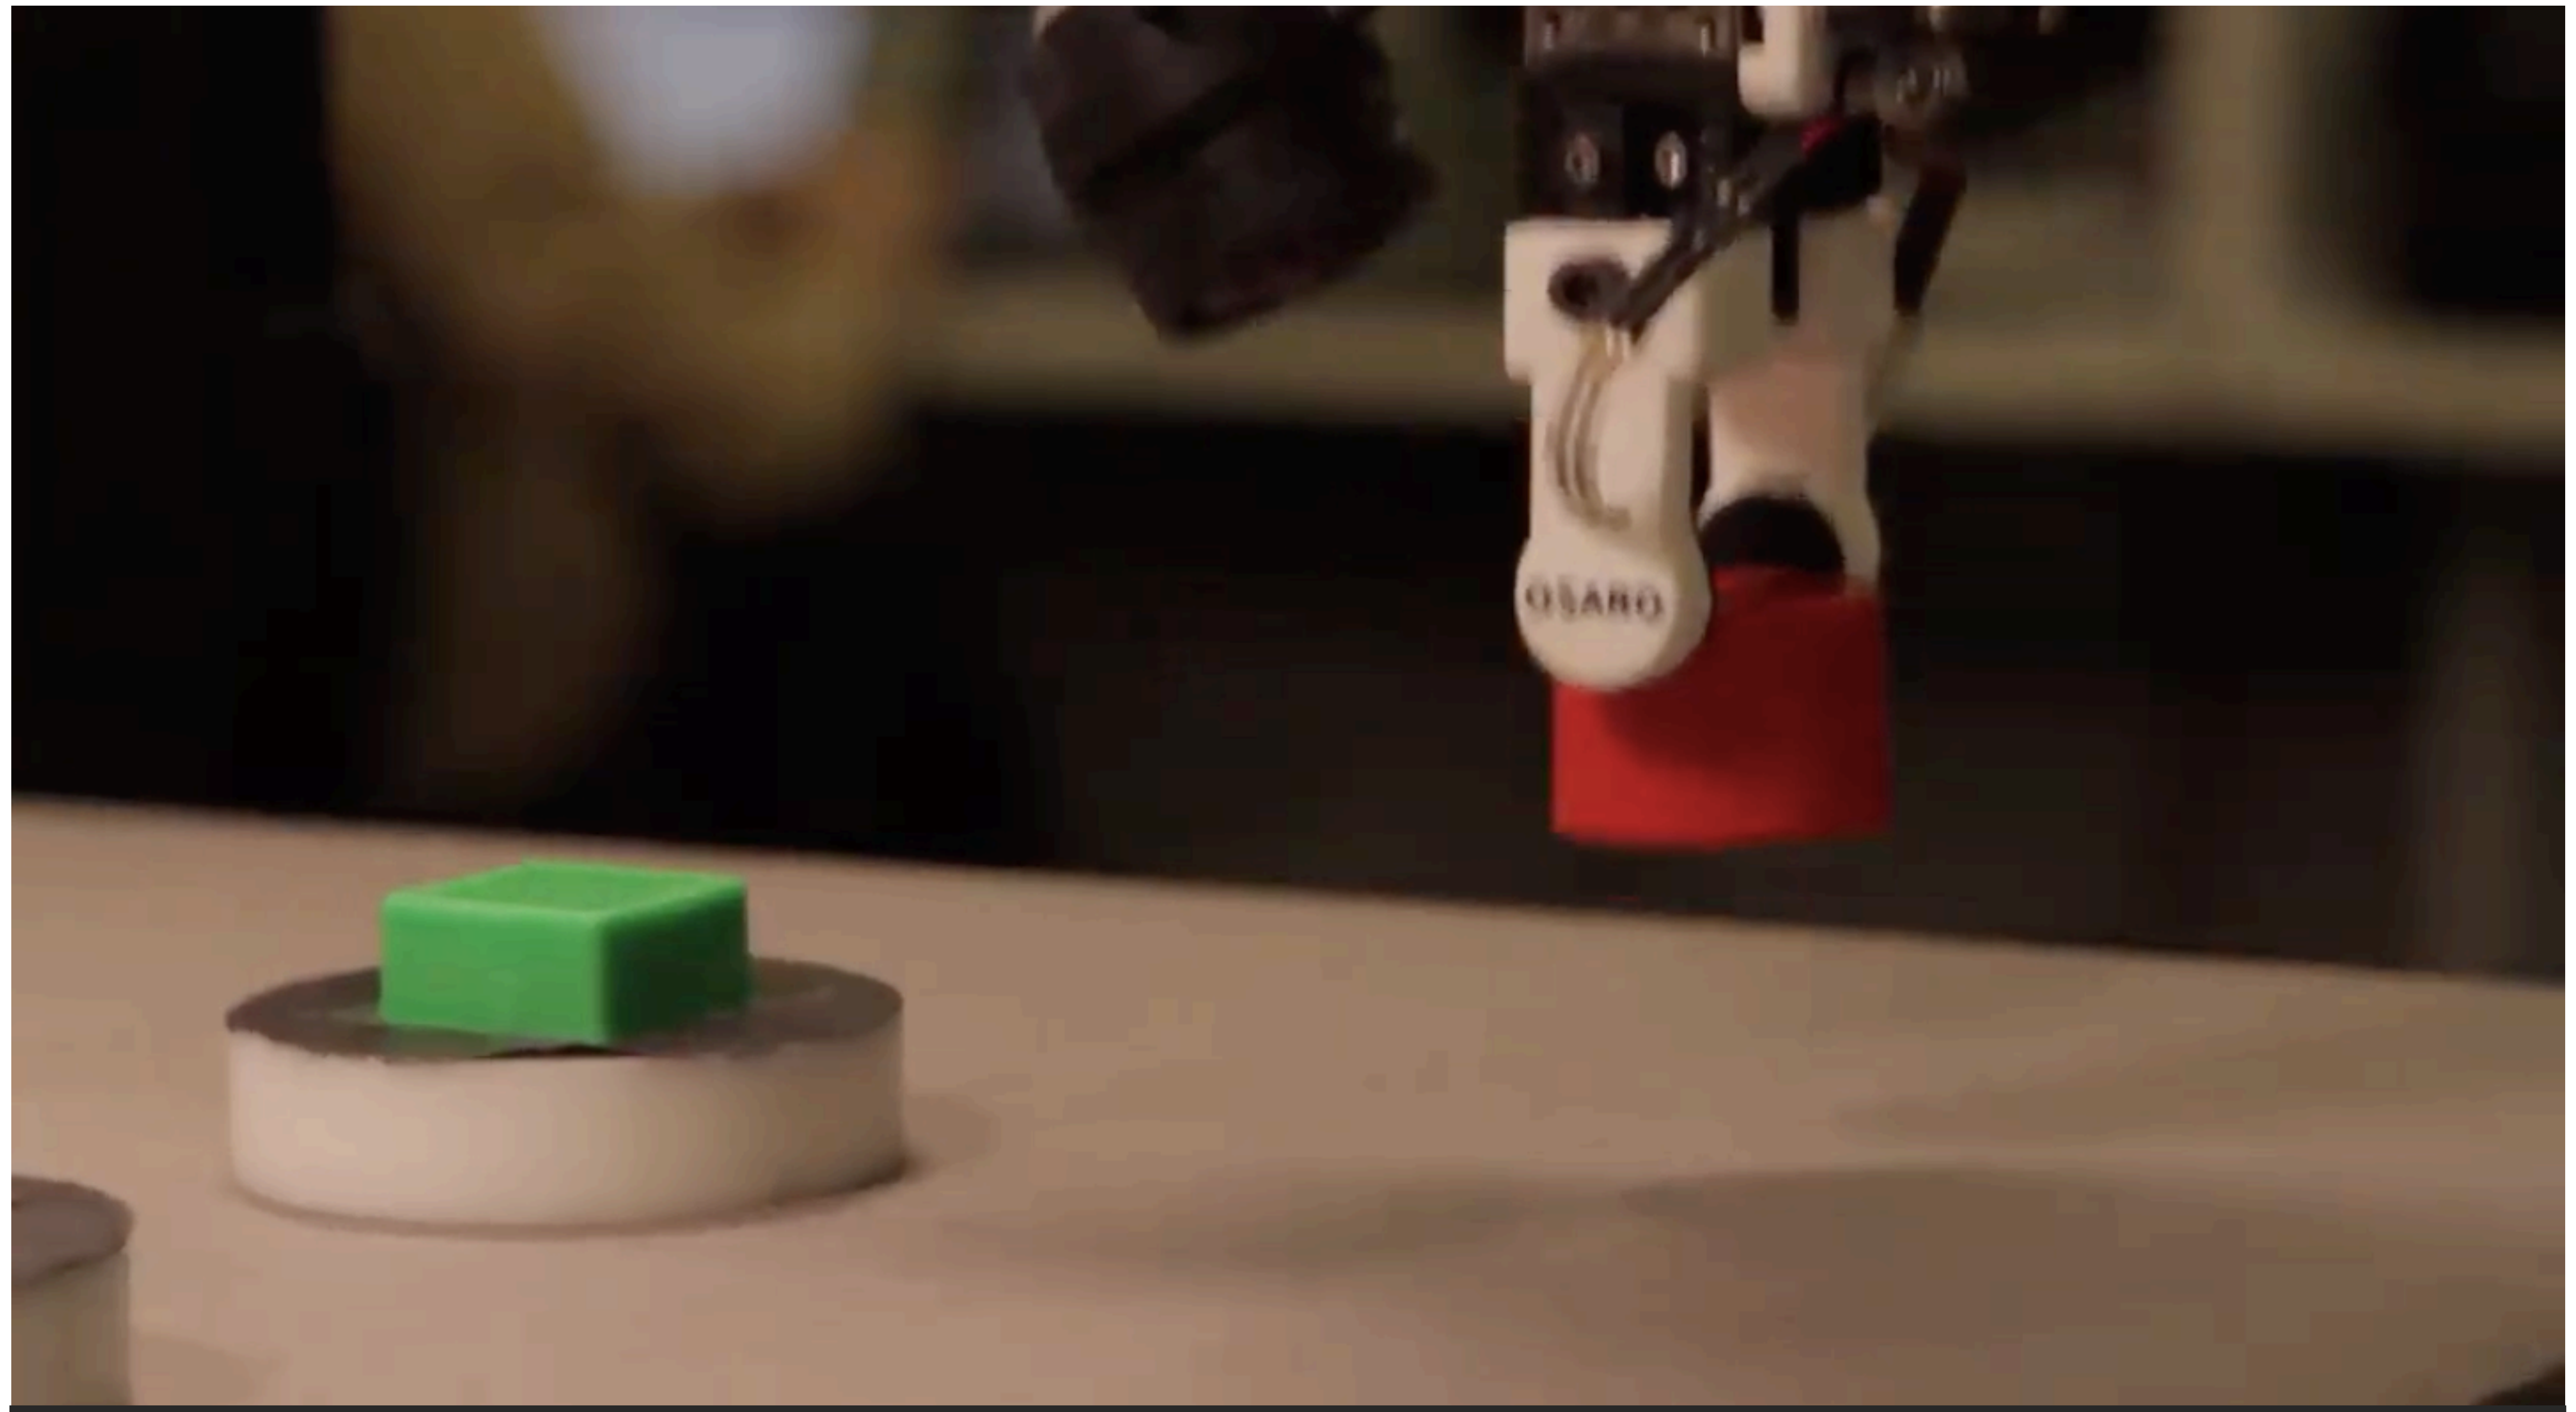

https://youtu.be/d6cIVueQAxE

IMITATION LEARNING

What is it?	Why use it?	Issues
Aims at learning control using expert demonstrations, focusing on imitating expert demonstrations.	- Address sample inefficiency and computational feasibility by giving an agent prior information about the world through mimicking human behavior. - When reward function is sparse, difficult to specify, or optimize directly	- Compounding error problem - Requires expert demonstrations (expensive and difficult to collect) - Copies the behavior, does not learn the problem objective necessarily - Trained policy only as good as demonstration

MACHINE LEARNING

META LEARNING / TRANSFER LEARNING?

	What is it?	Why use it?	Issues
Meta Learning	Models that can learn new skills or adapt to new environments rapidly with a few training examples. Early approaches: Automatically fine tune hyper parameters (learning algorithm) Now: Learning to learn Meta learner (agent) and Learner (model): common feature representations of tasks	- The ability to quickly adapt to new unseen tasks - Scalability - Sample efficiency	- Complex algorithms - Training time - Careful design of training procedure
Transfer Learning	Broadly defined as using experiences and knowledge from solving one task to another related problem Examples: Pre-trained models, sim-to-real	Many tasks share a similar underlying structure/ representation	- Theoretically less understood - Negative transfer issue when source and target domains do not overlap much - Measure transferability

MACHINE LEARNING AT SCALE
REAL WORLD SENSORS ARE NOISY

- Sensor Considerations
  - Sensors are noisy and fail
  - Real world products confound sensors

  - Training deep neural architectures
  - Semantic segmentation; background

Suction brush, not thin handle

Don't get fooled by **reflections**

Pick **clear box**, not what's inside

Avoid pump, **aim for body**

Target label, not gloves

Avoid contour, pick flat parts

Case2(hard contact)/Person_A/Success

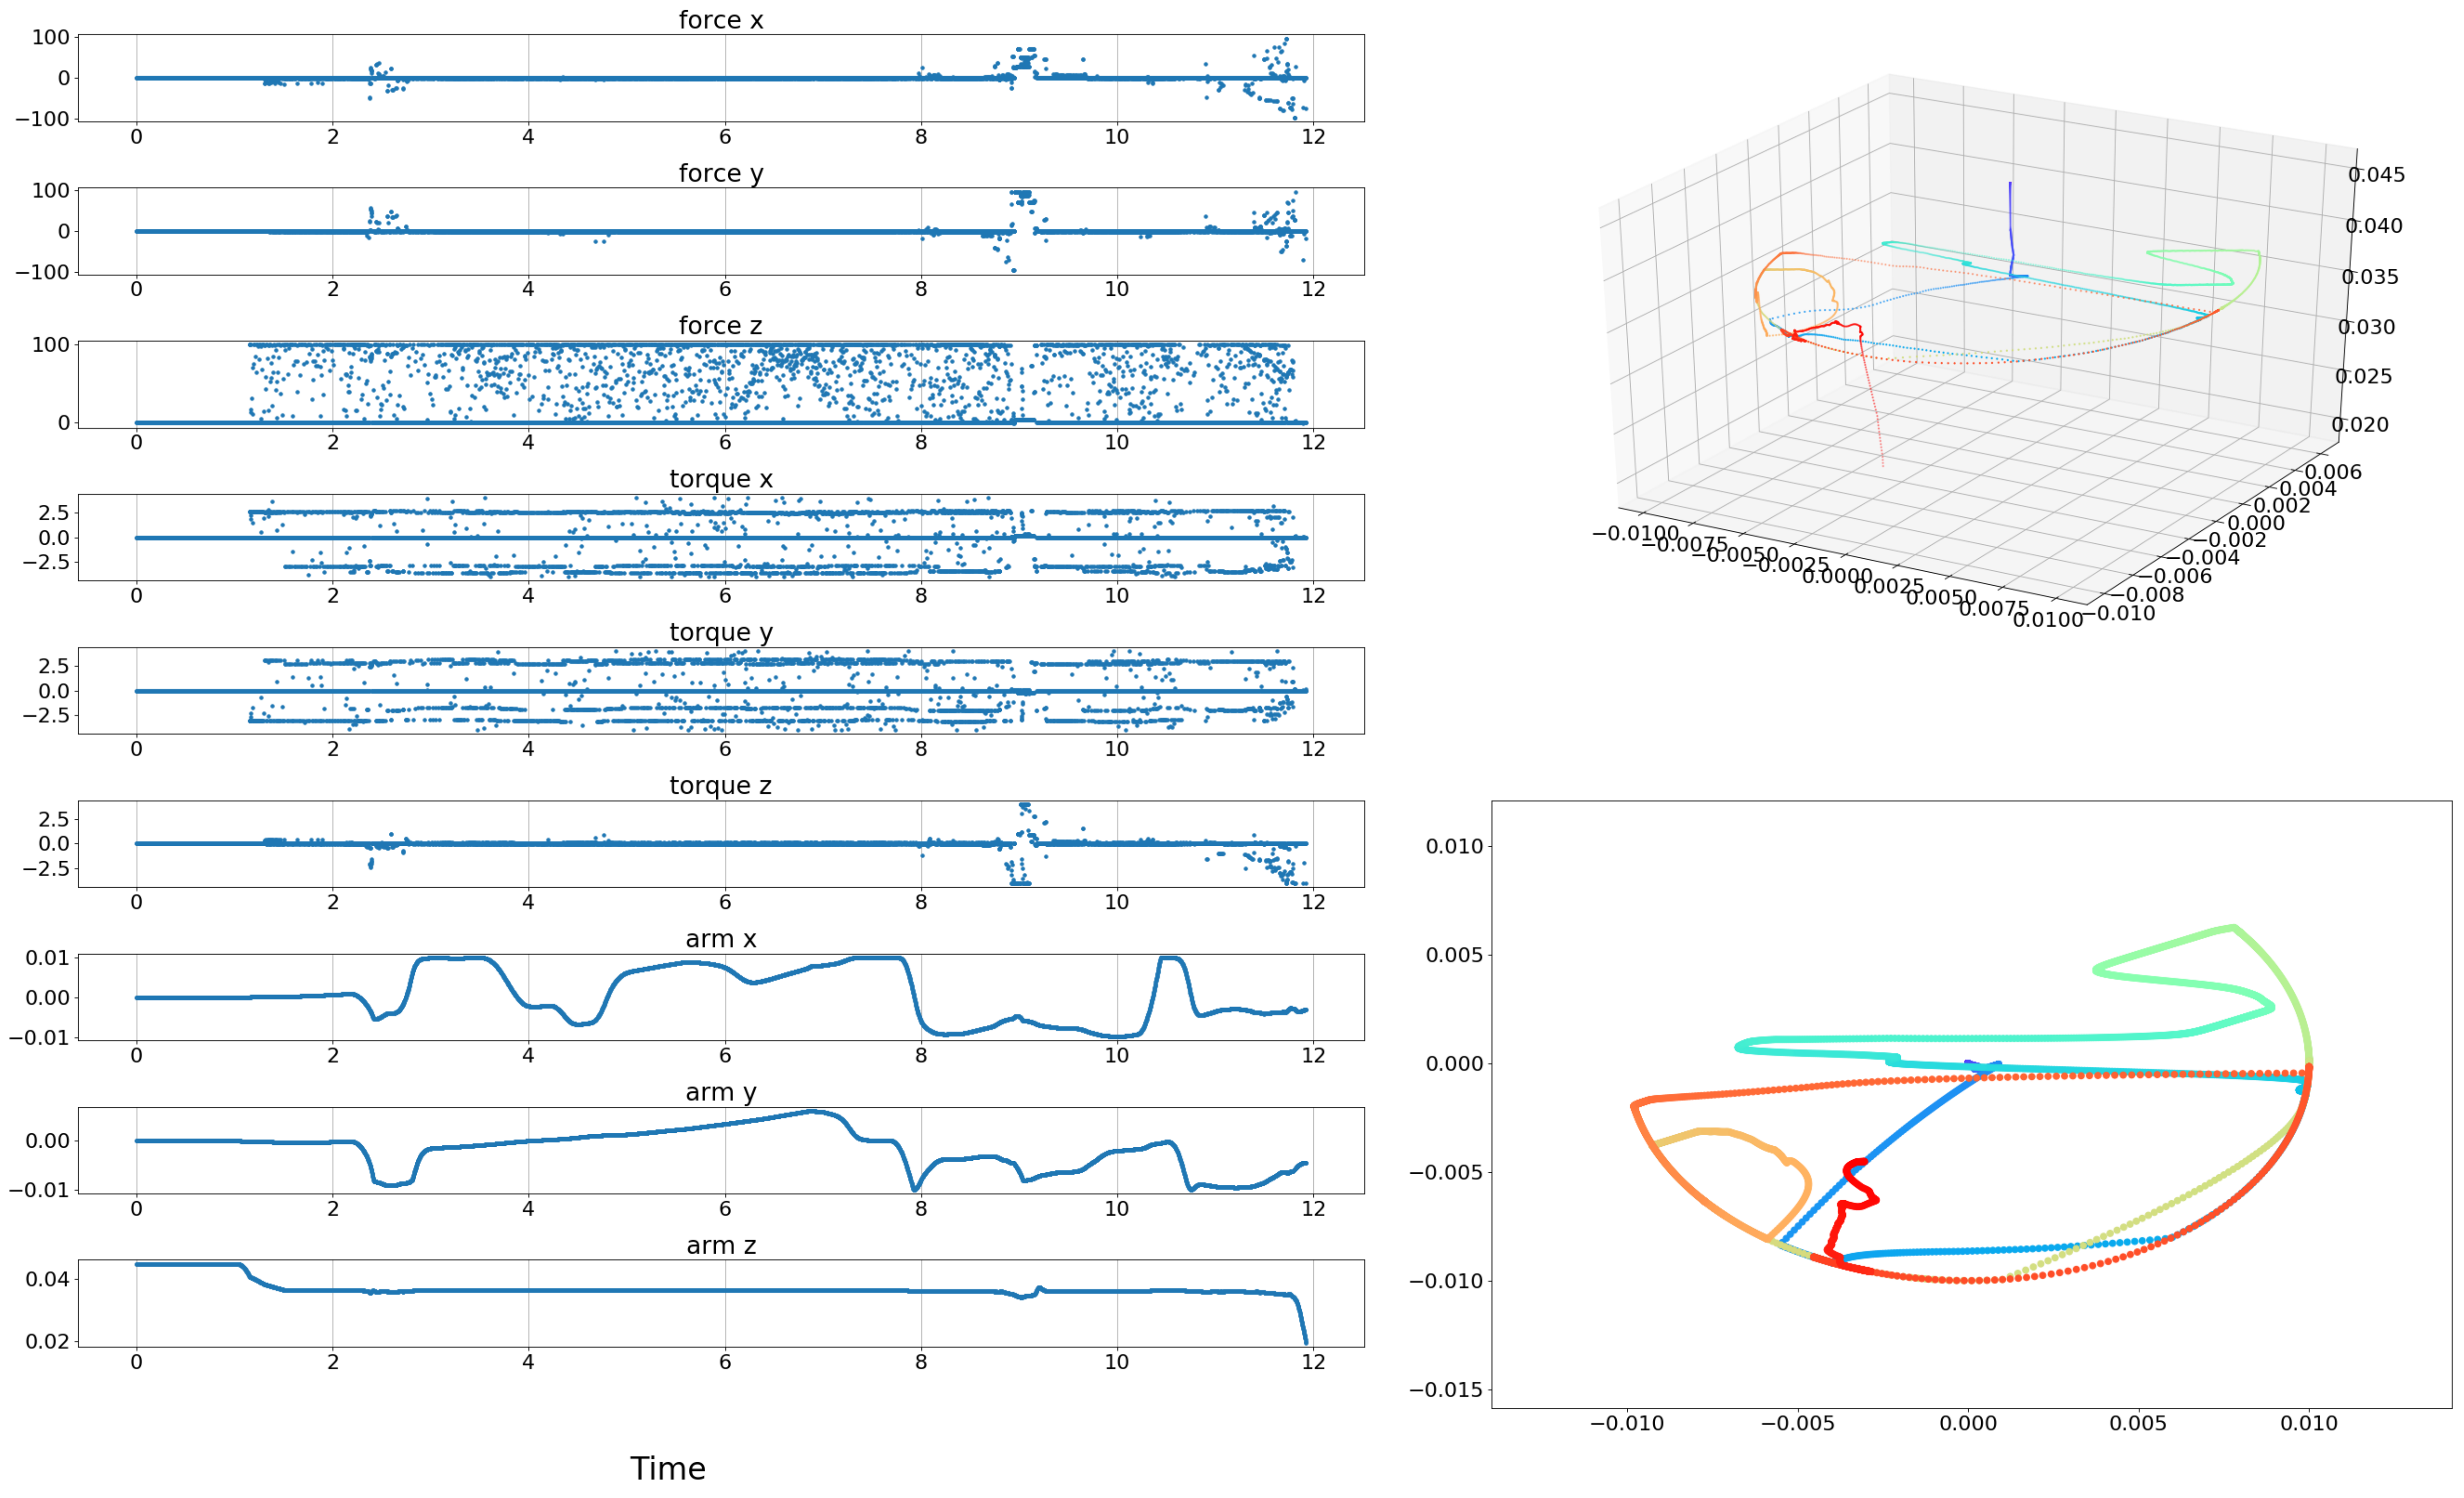

MACHINE LEARNING AT SCALE

INTEGRATION AND CONNECTIVITY

- Integrations and connectivity
  - Bandwidth constraints => data sparsity
  - Connectivity constraints => asynchronous
  - Must integrate with surrounding systems
  - Must integrate with commodity hardware

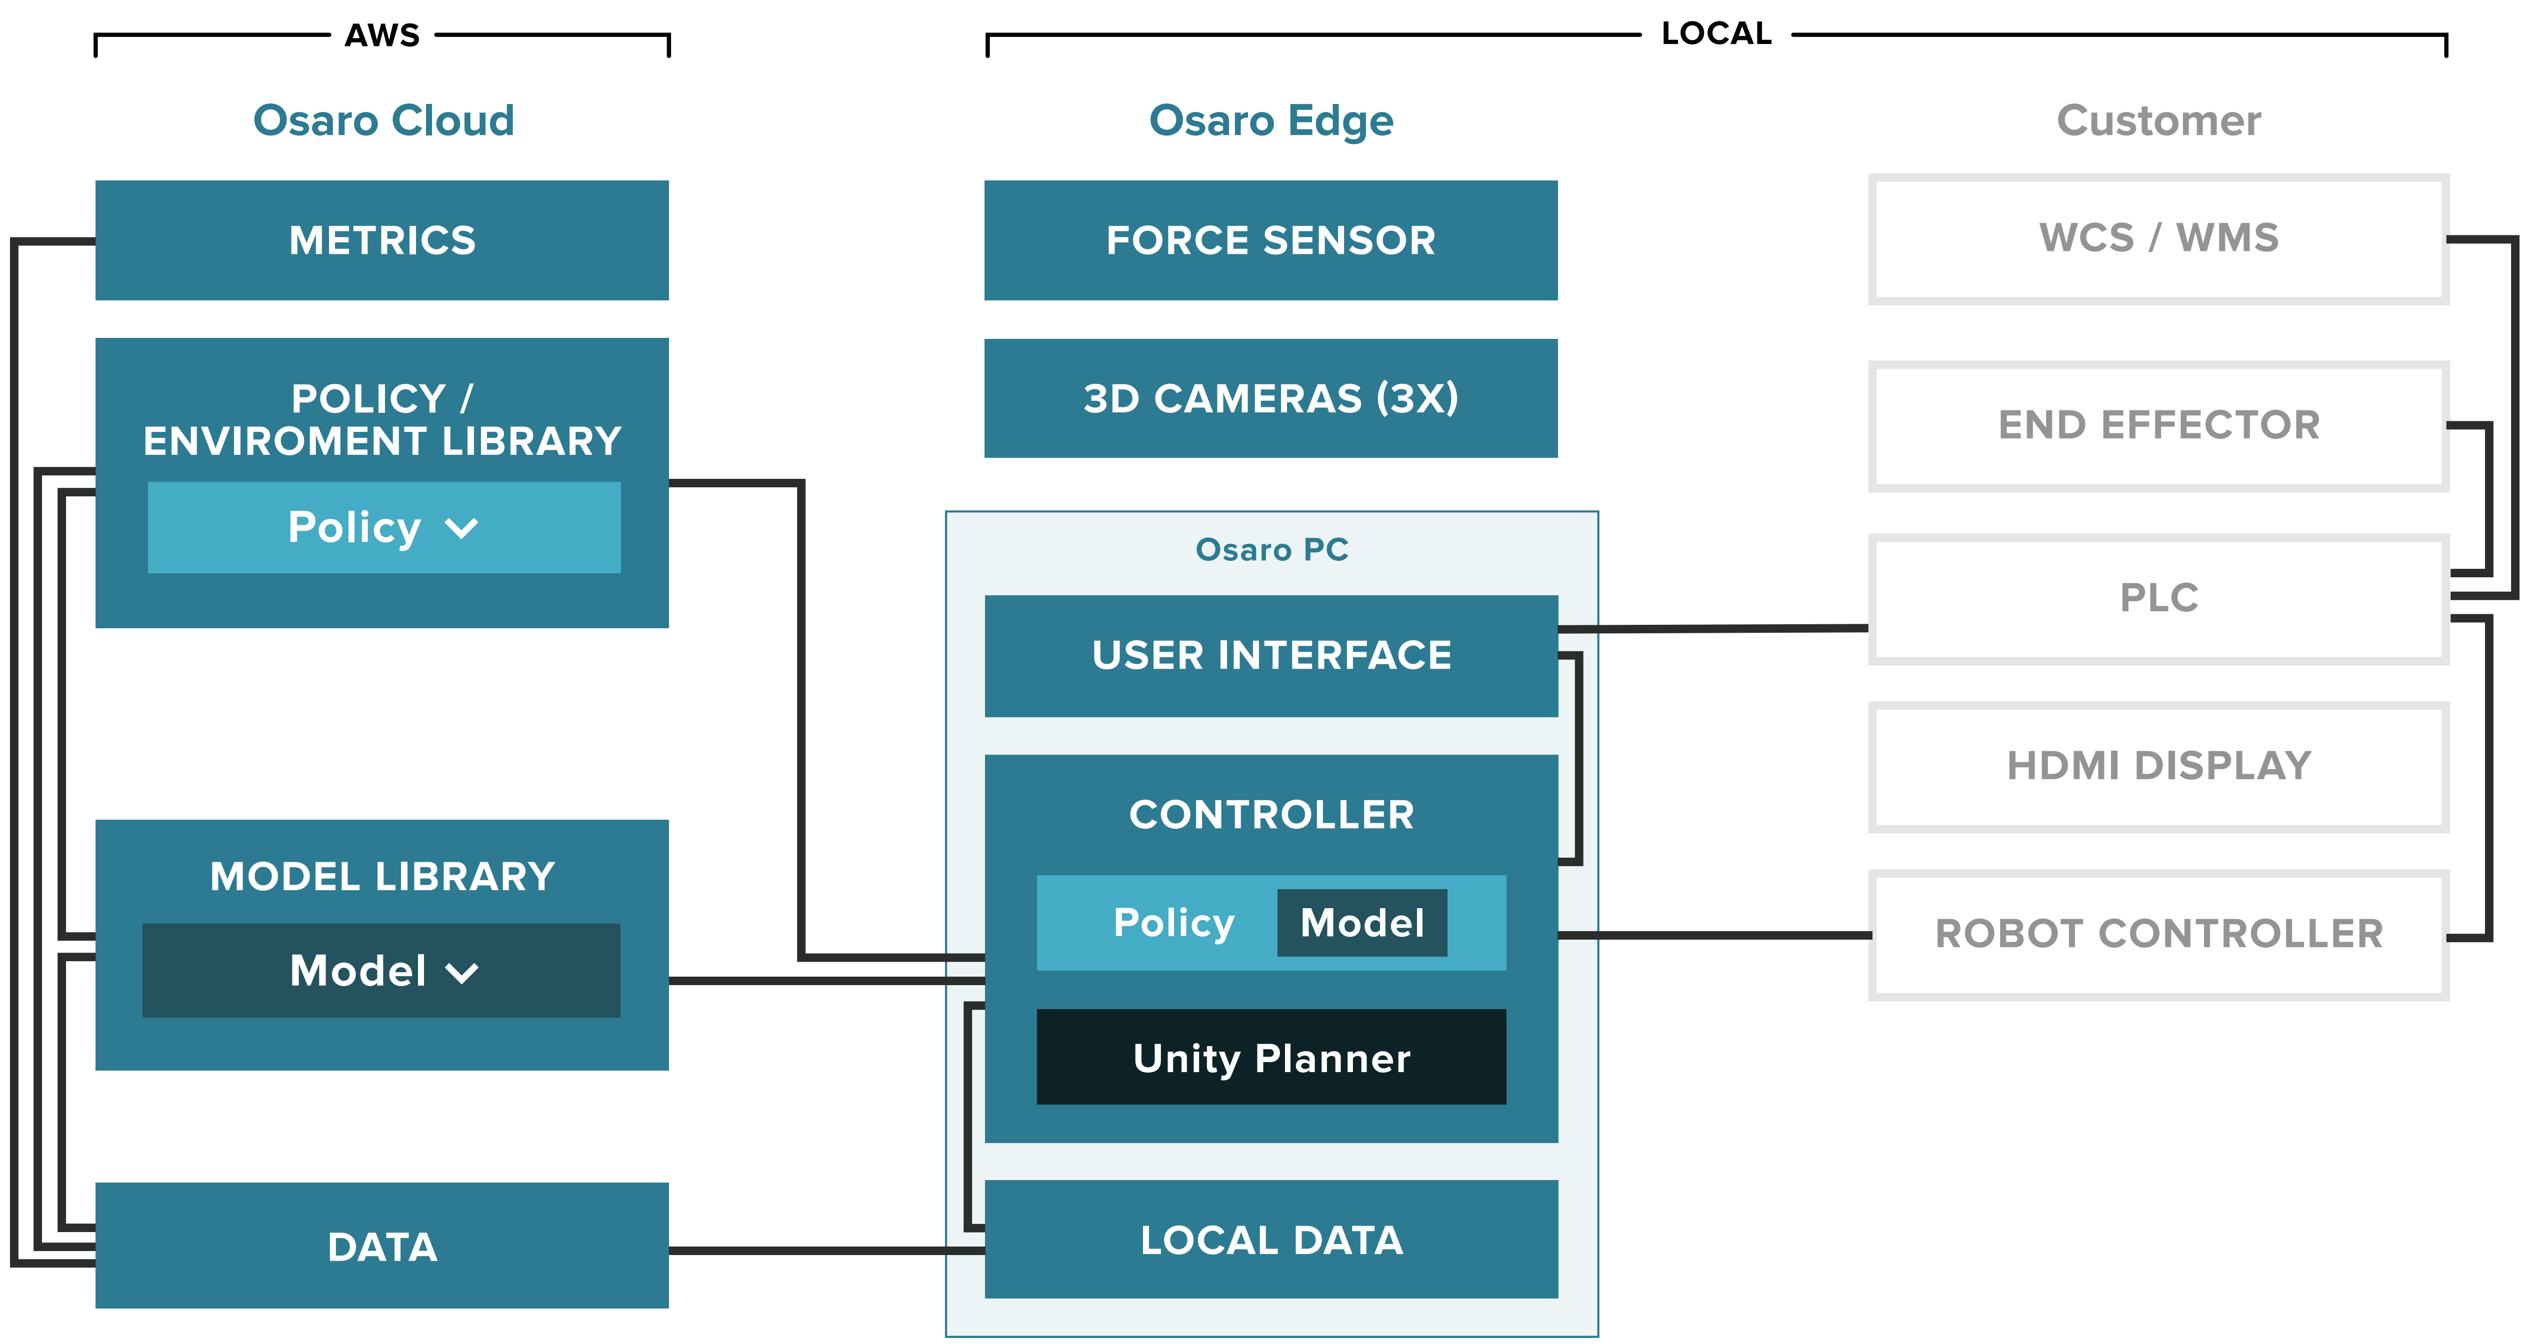

CUSTOMERS WANT FLEXIBILITY

Robots: Kuka, ABB, FANUC, Denso, UR, etc

End effectors: Schmaltz, Schunk, Nitta, etc.

Sensors: Realsense, Pointgrey, etc.

WMS Integrations

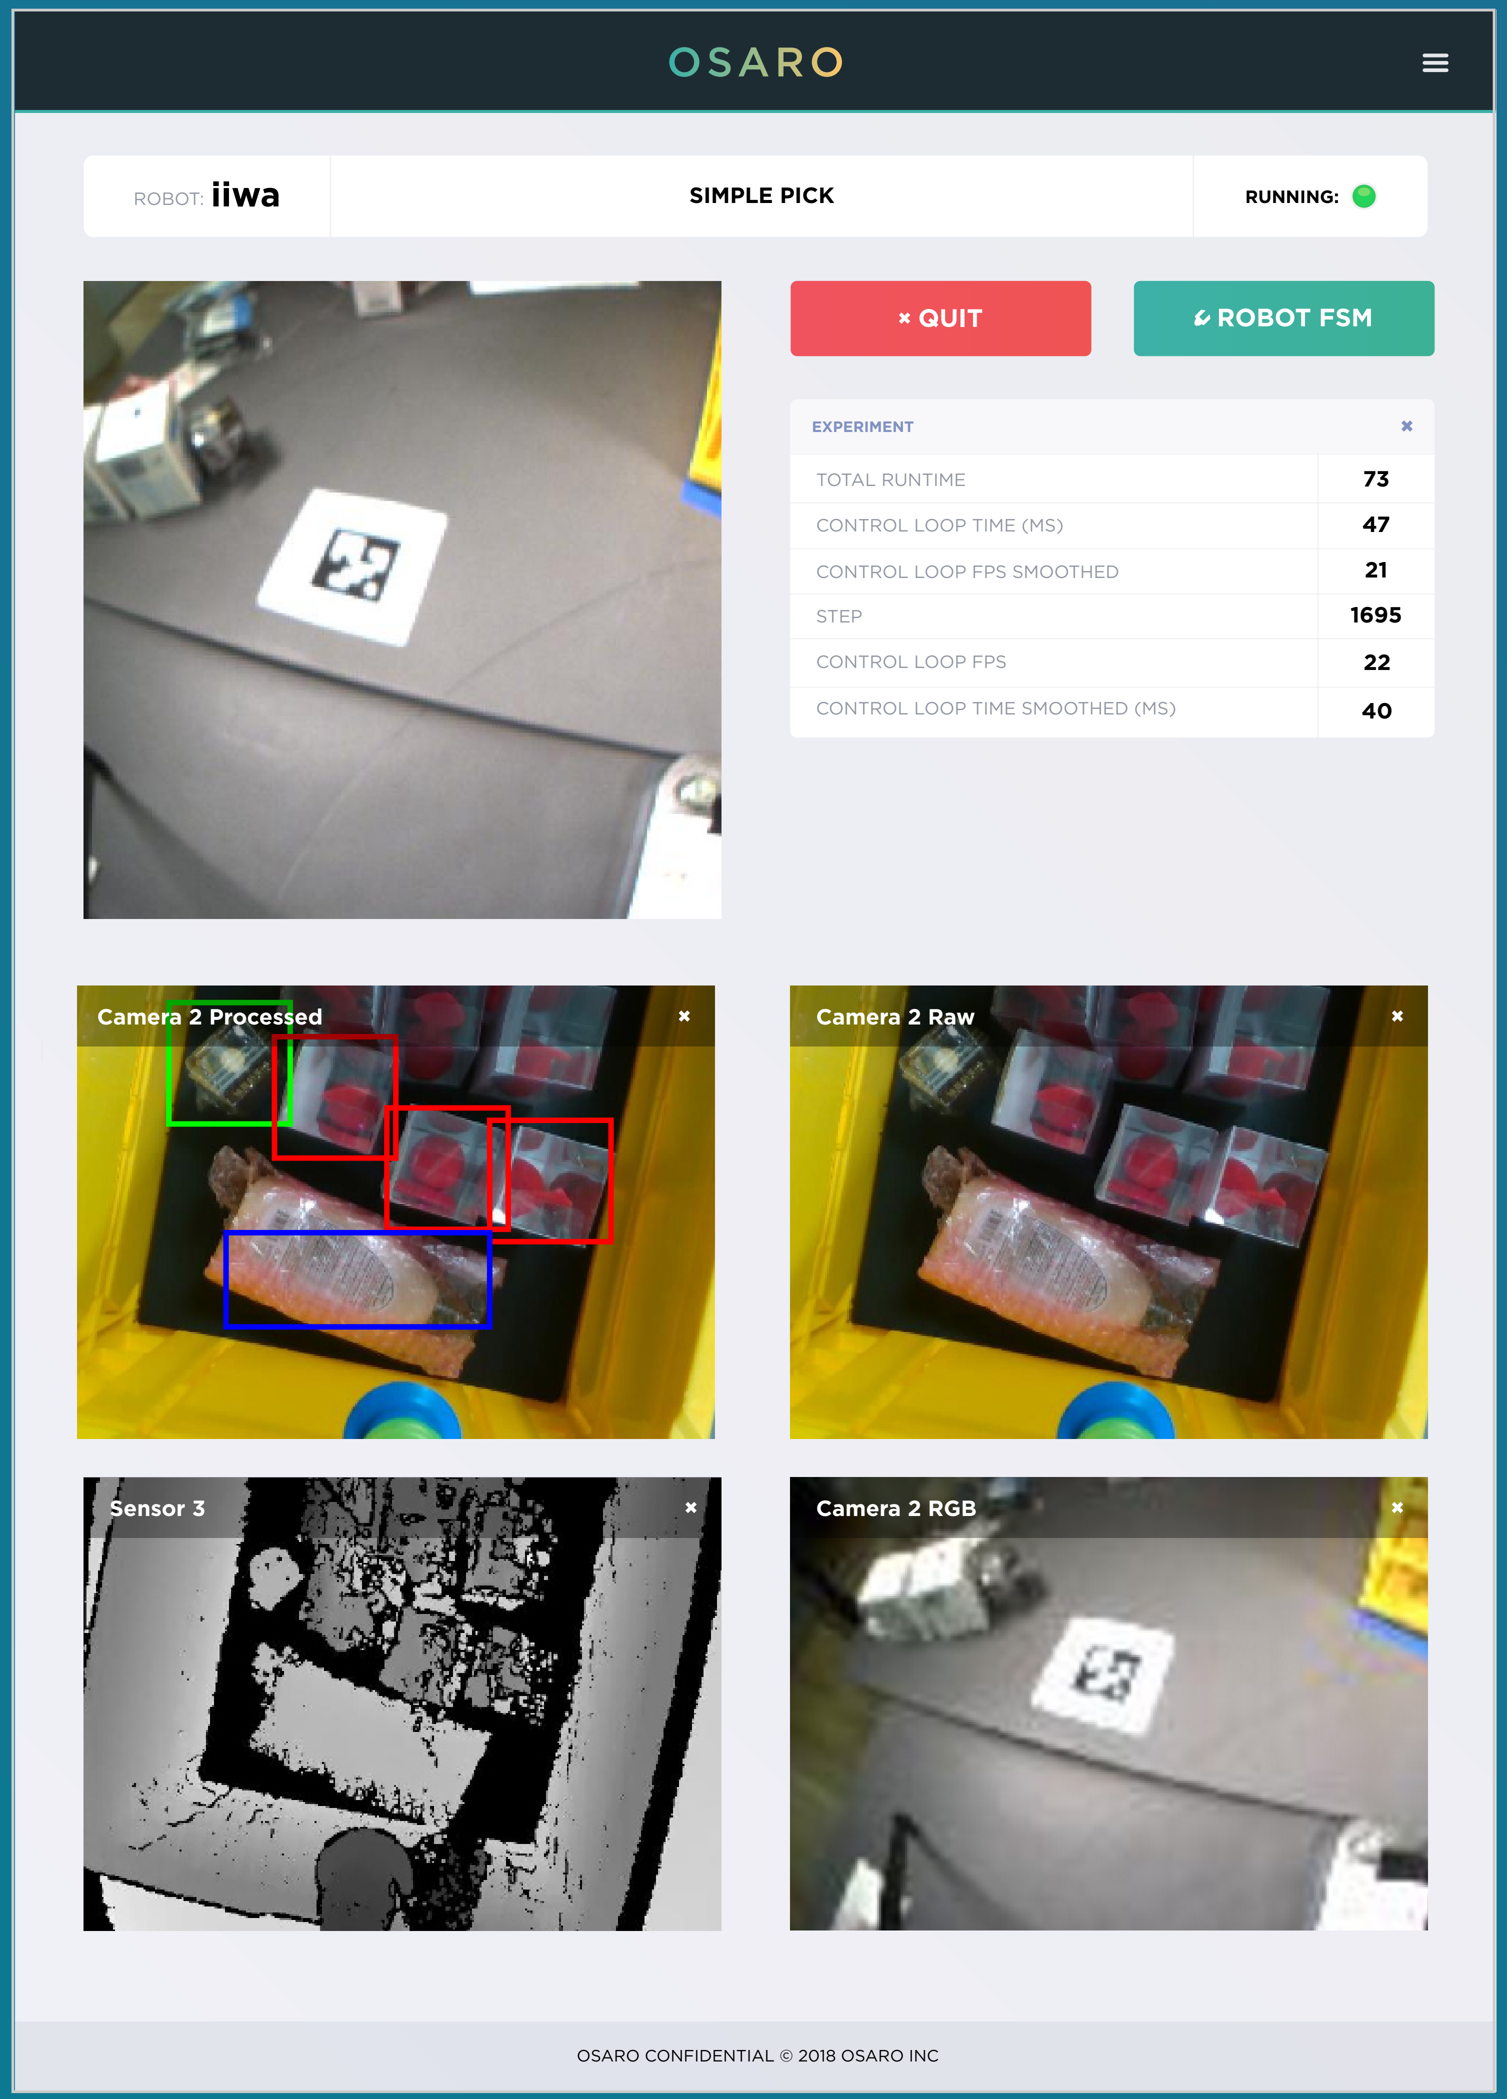

MACHINE LEARNING AT SCALE
MODEL TRAINING AND EVALUATION

- Model training and evaluation
  - Models must “just work”
  - Training must be quick, continuous
  - Inference must be measurable, interpretable
  - Must be robust to changes

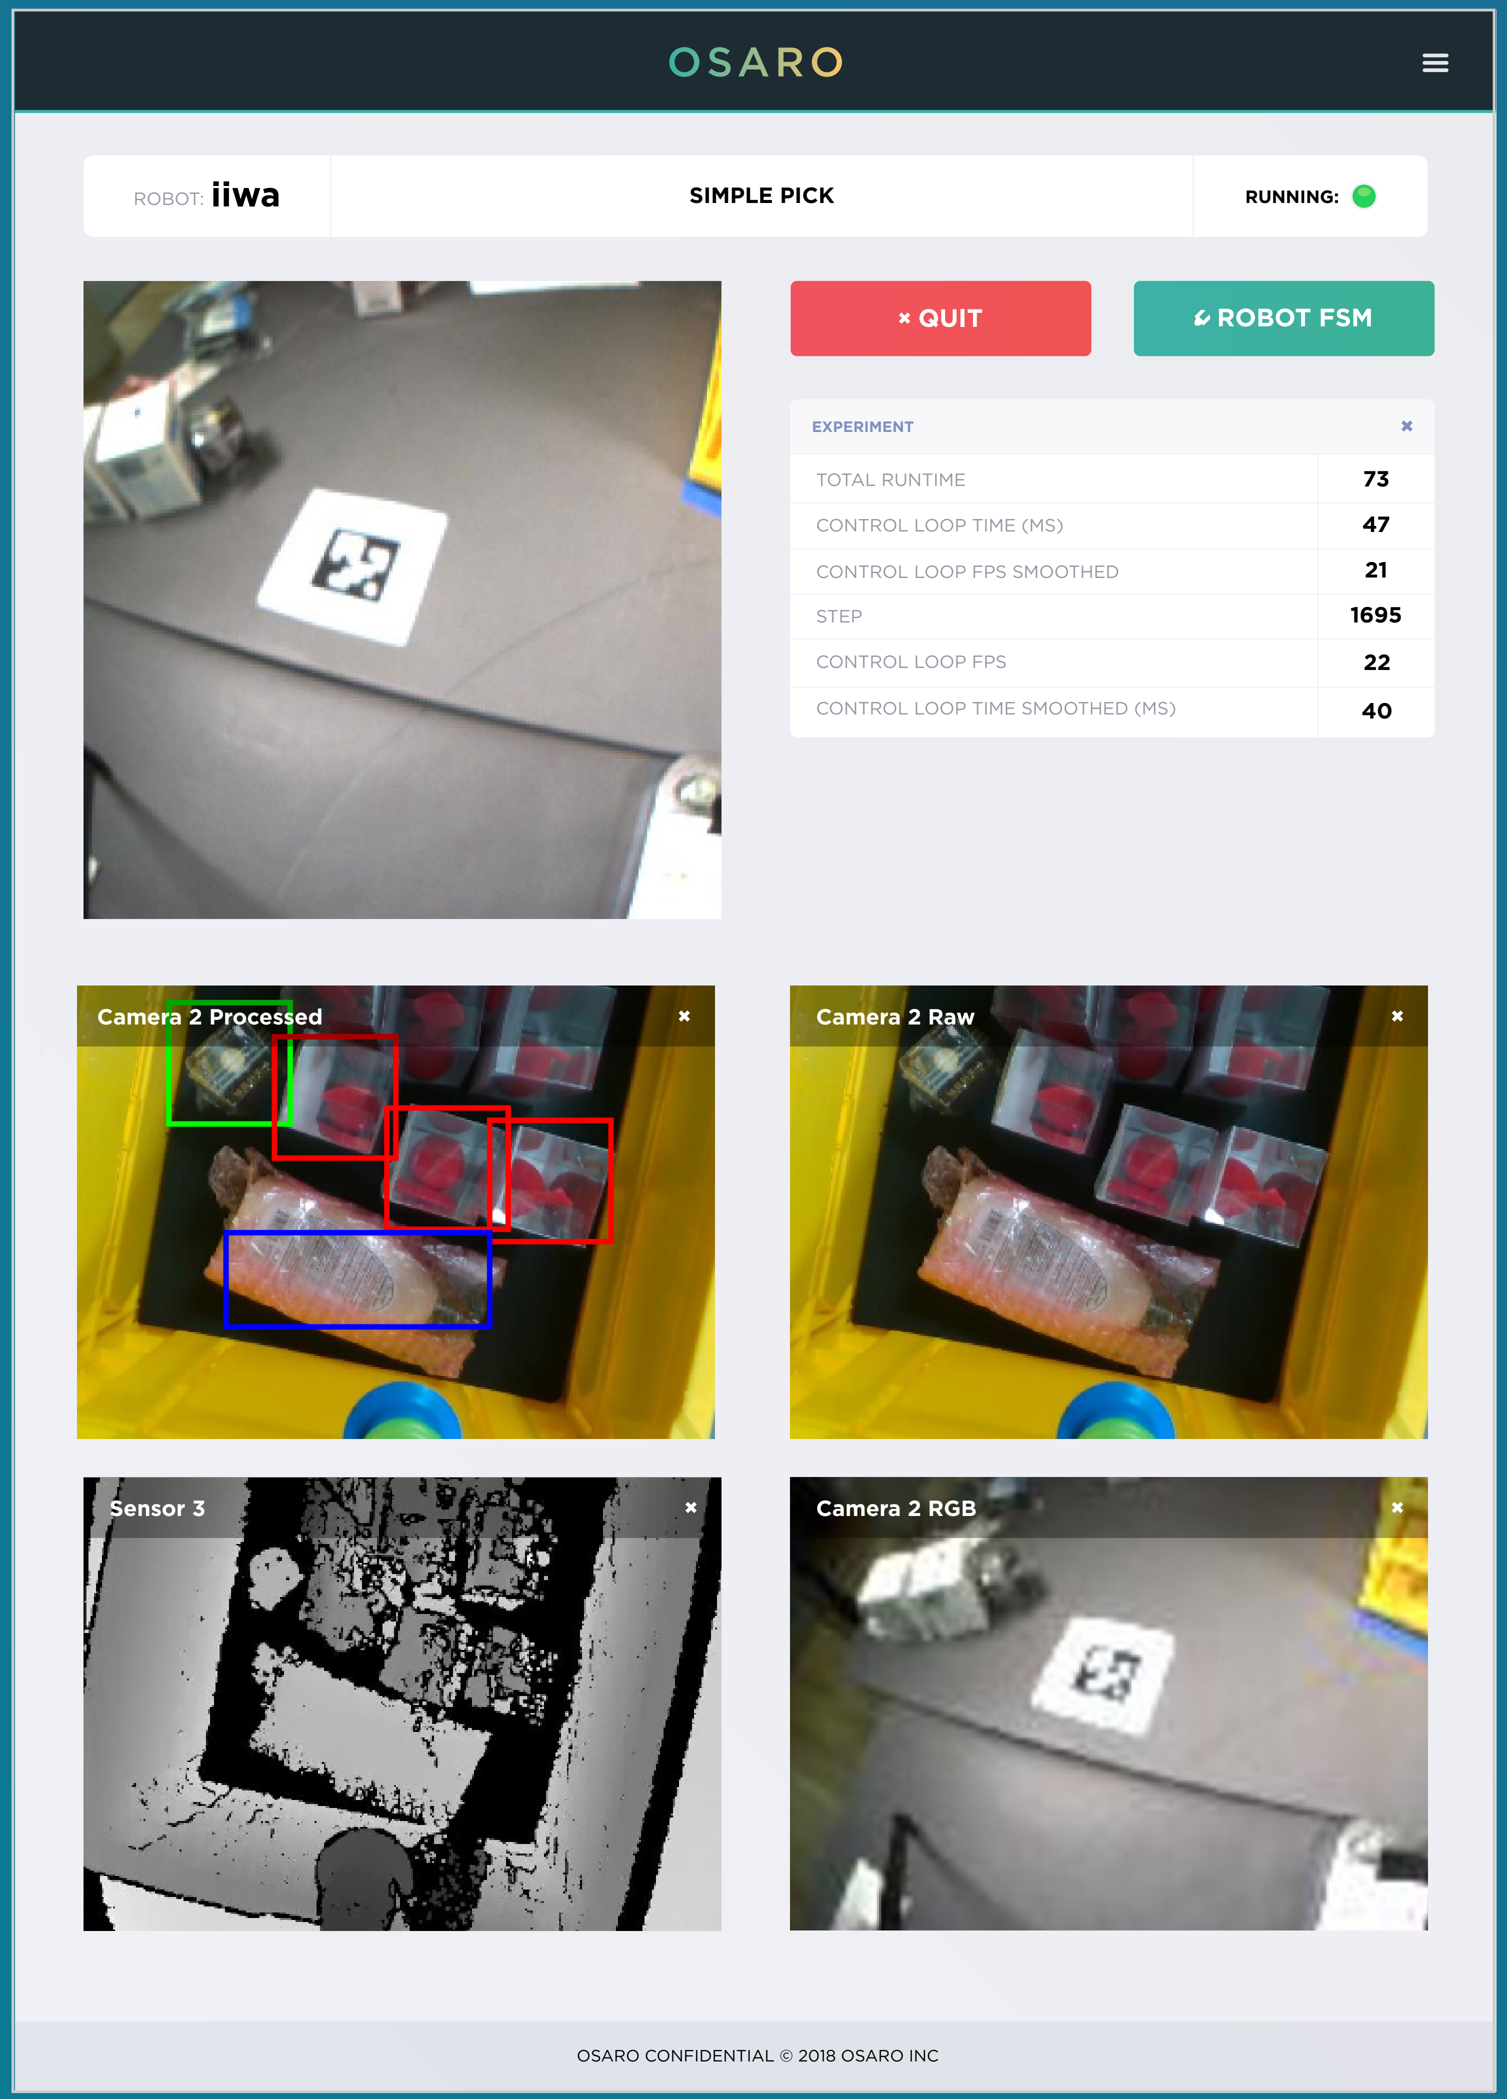

MACHINE LEARNING AT SCALE

METRICS: pick time, accuracy, error, system performance

OSARO PICKING EVALUATION

Evaluation Profile
Robot: Iiwa
Item Type: ointment products
# of Items: 11

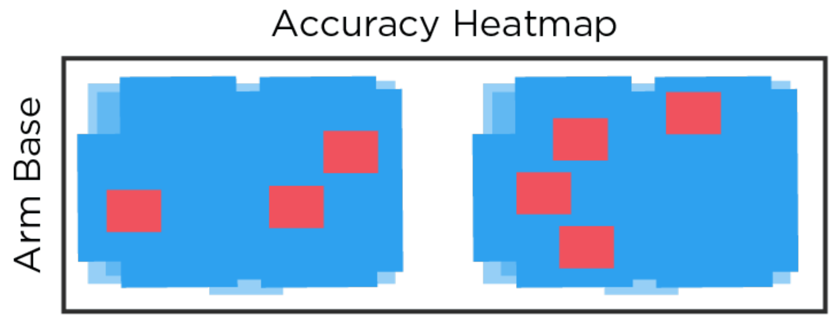

Accuracy

Bins Cleared:
- 111 Attempts
- 105 Successes (94.6%)

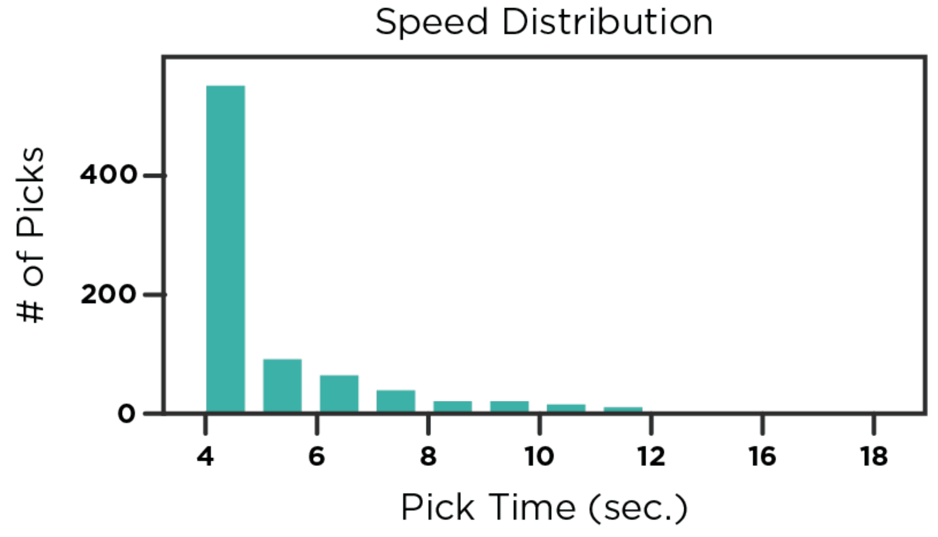

Pick Breakdown:
- Successful First Tries: 1102 (90.8%)
- Successful Fallbacks: 94 (7.7%)
- Retries: 11 (0.9%)
- Errors: 7 (0.6%)
- Total Pick Episodes: 1214

Evaluation Summary:
- Successful Picks: 1196 (98.5%)
- Retries and Errors: 18 (1.5%)
- Total Pick Episodes (1214)

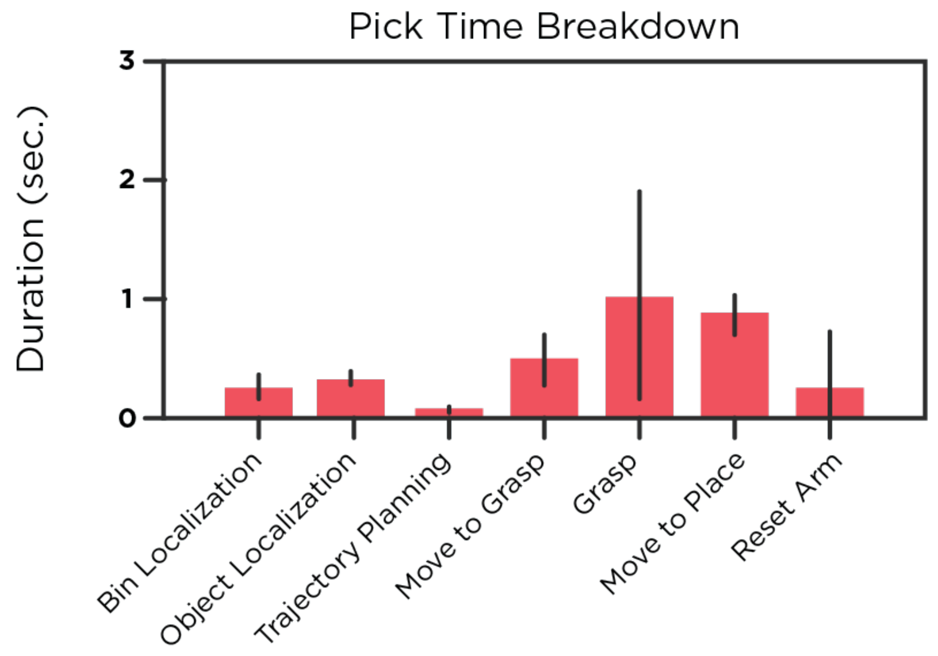

Efficiency
Average Takt Time: 5.1s
Average Success Rate: 11.8 items/min (706.1 items/hr)
Total Evaluation Time: 103.2mm

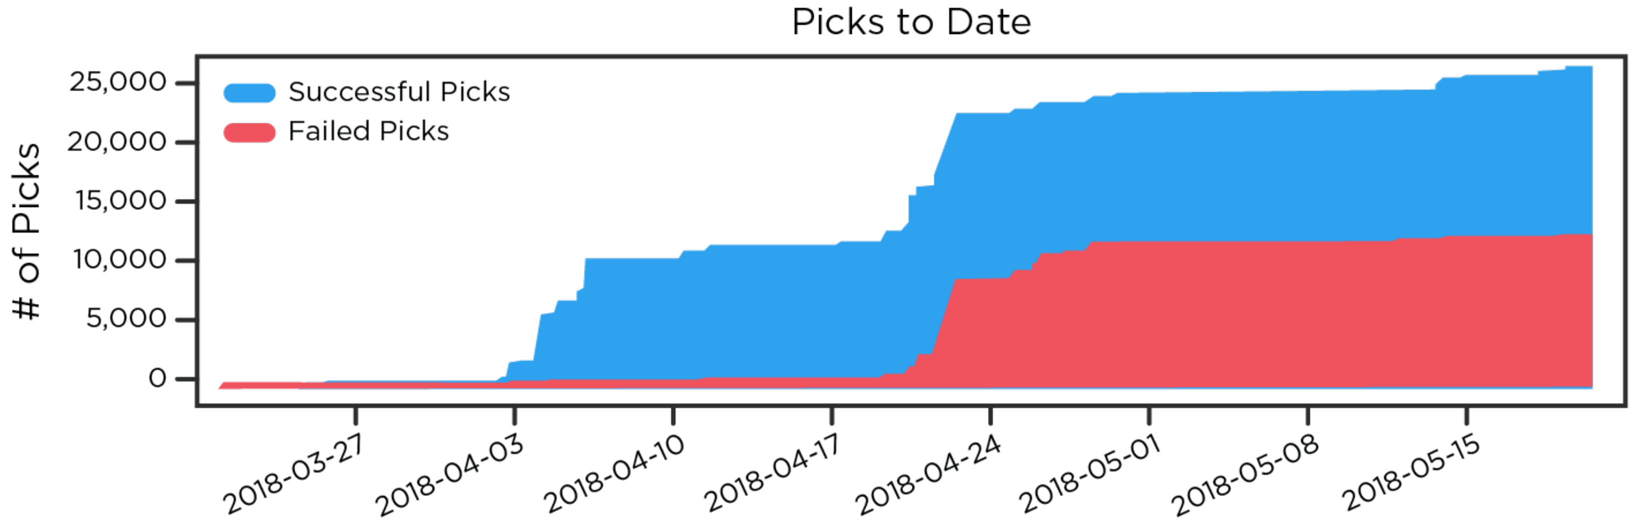

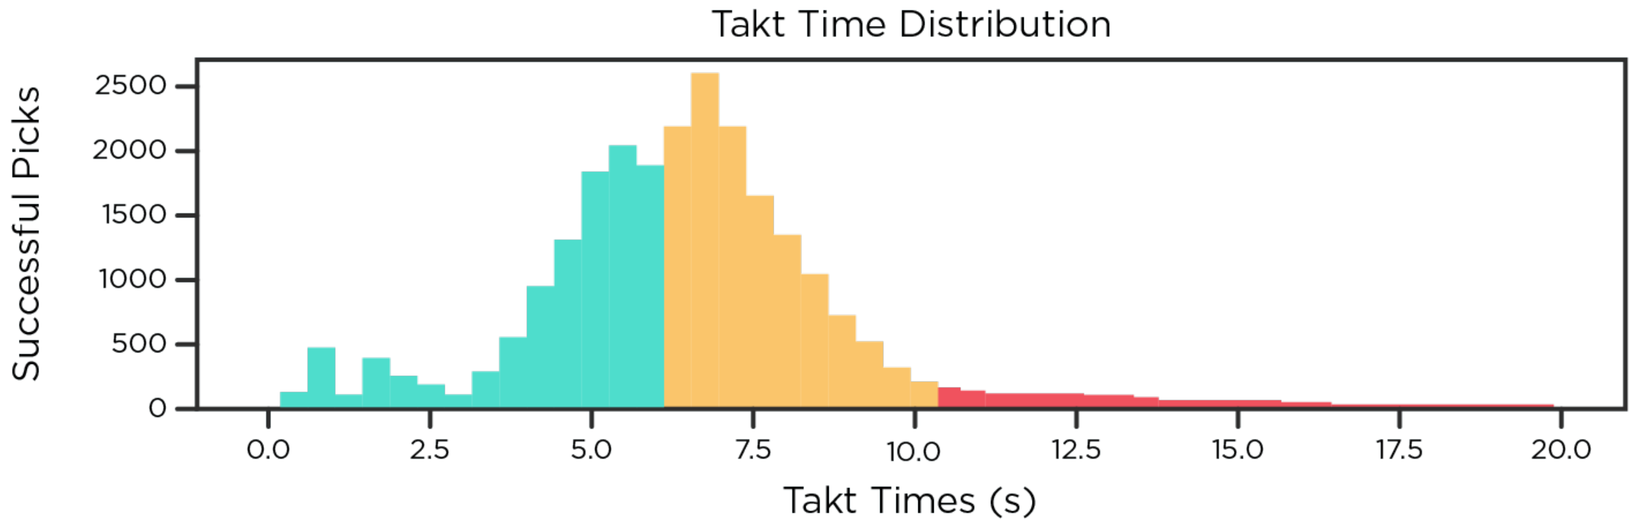

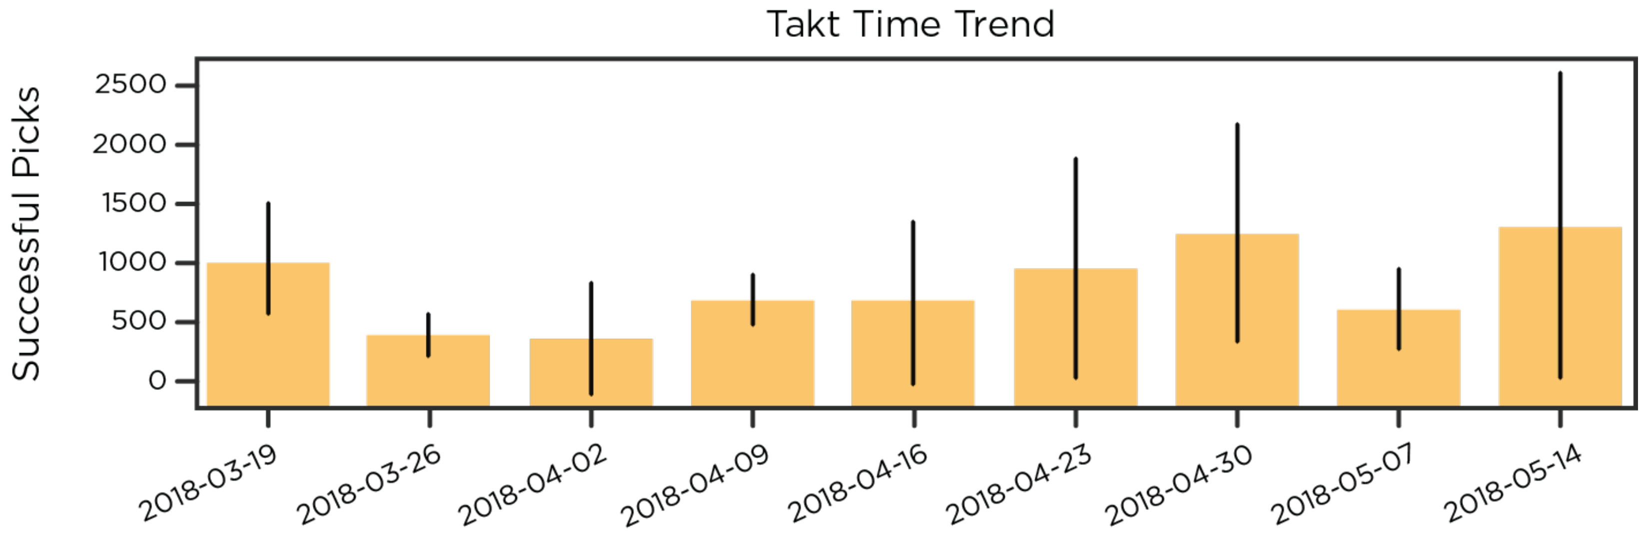

MACHINE LEARNING AT SCALE

UI MATTERS: simple, intuitive vs. functionality, flexibility

MACHINE LEARNING AT SCALE

TRAINING DATA

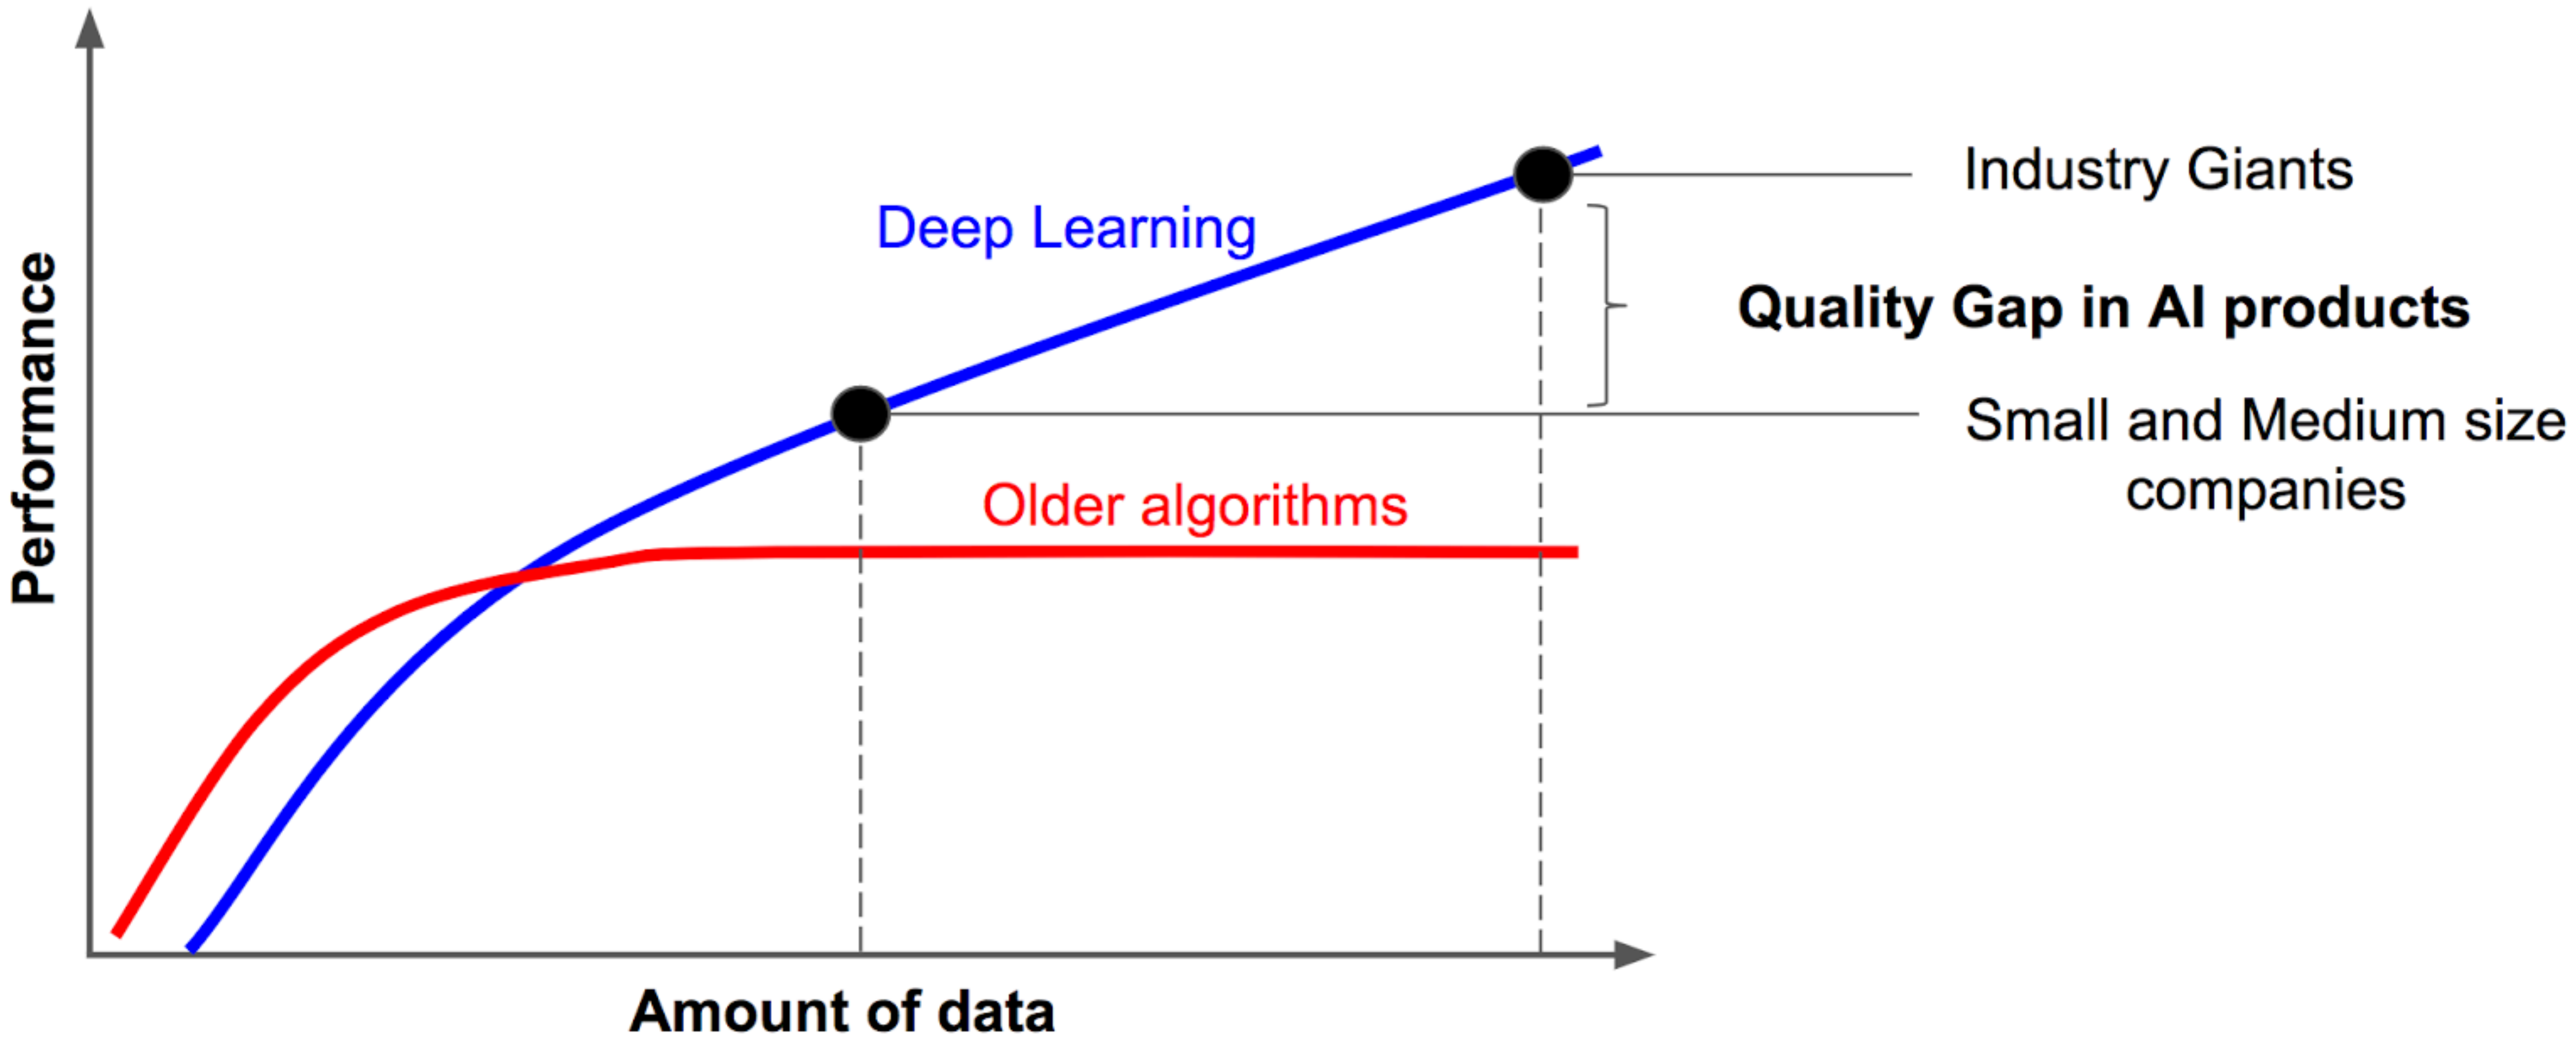

THE ULTIMATE DREAM: "FOLD MY LAUNDRY. THEN FIX MY CAR."

SCALABILITY & GENERALITY

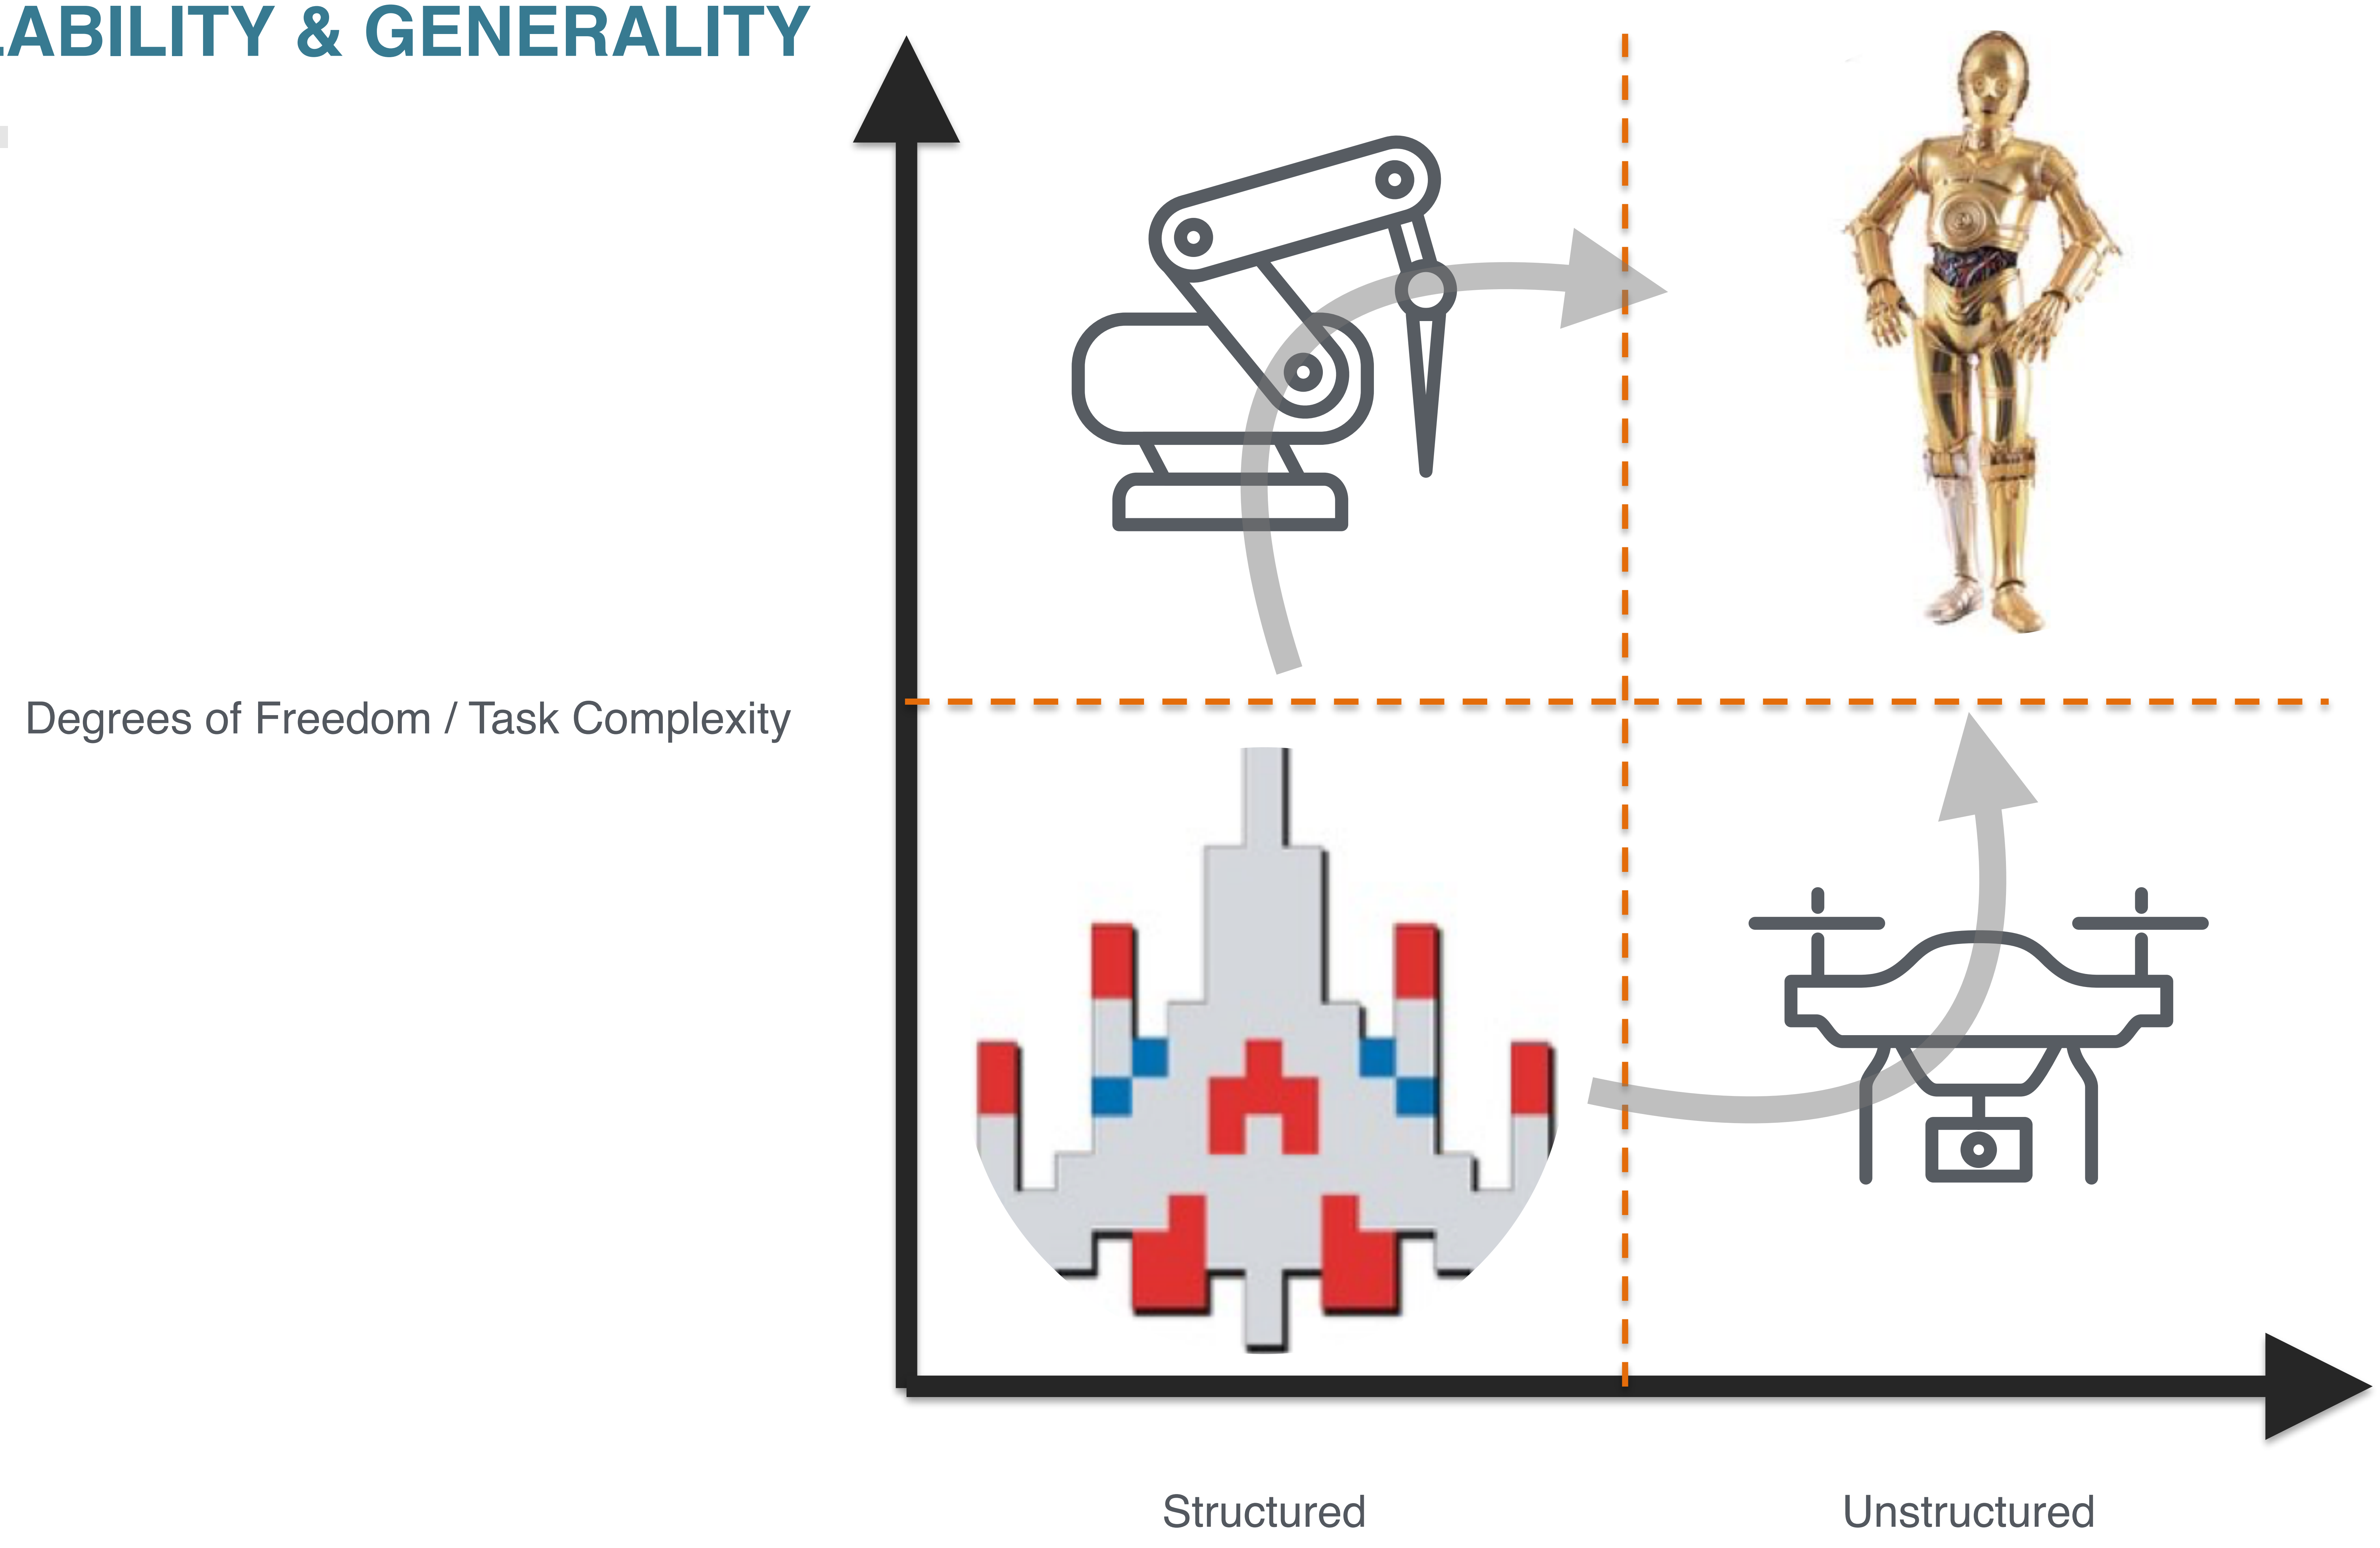

OSARO'S
APPROACH: Interdisciplinary

- Designed for tight integration
- Supports commodity hardware and sensors
- Verifiable, provides analytics
- Pipelined approach
- Pragmatic - we choose the right tool for the job:

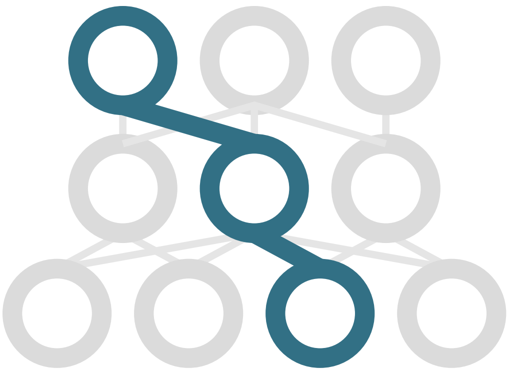

Deep Learning

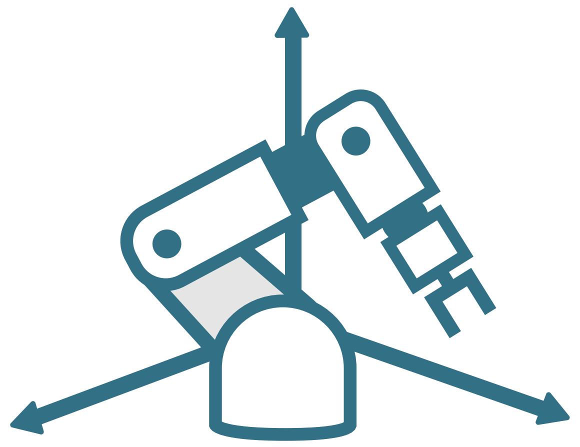

Simulation and
Data Augmentation

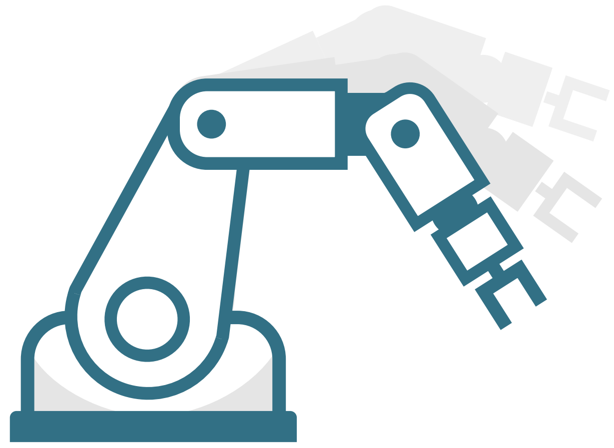

Classical Robotics,
Planning and Control

Reinforcement
Learning

OSARO'S
APPROACH

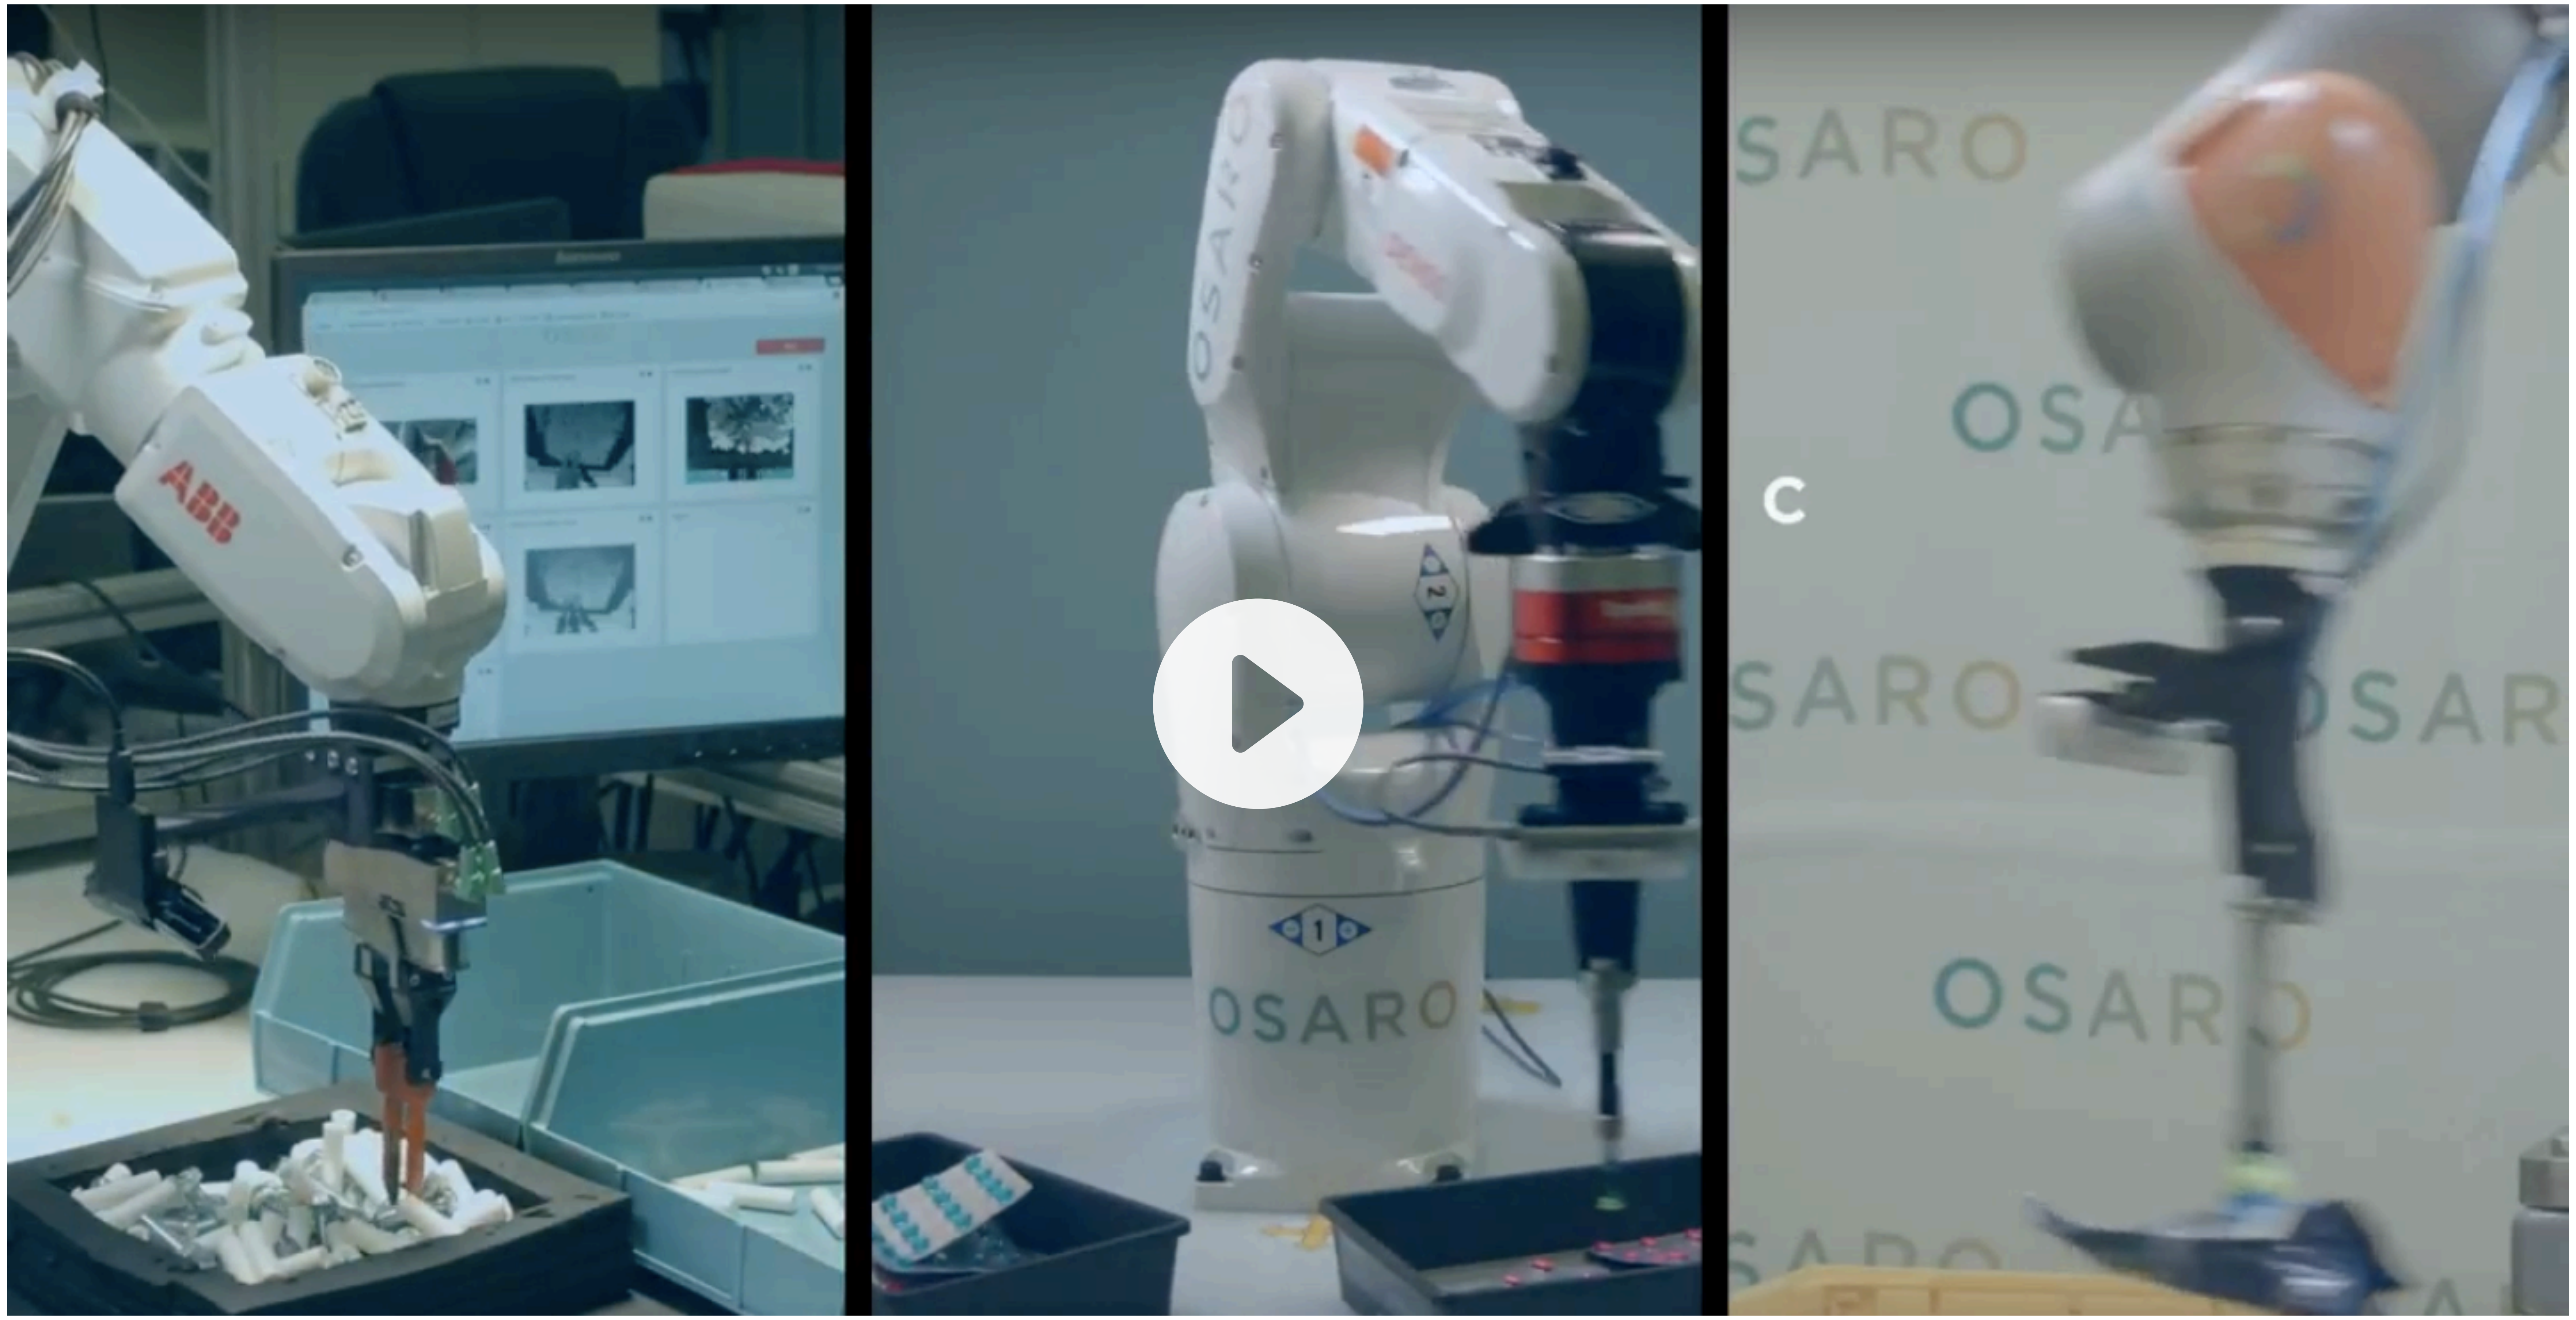

https://www.osaro.com/video

OSARO'S APPROACH

Support a variety of DL based vision models

- Pipelined approach
- Edge optimized
- Data cloud

Models Include:

→ Proprietary MASK Algorithm

→ Grasp Affordance Learning

→ Grasp Quality Network

→ Sim to Real Techniques

→ Few-Shot Meta Learner

→ And More…

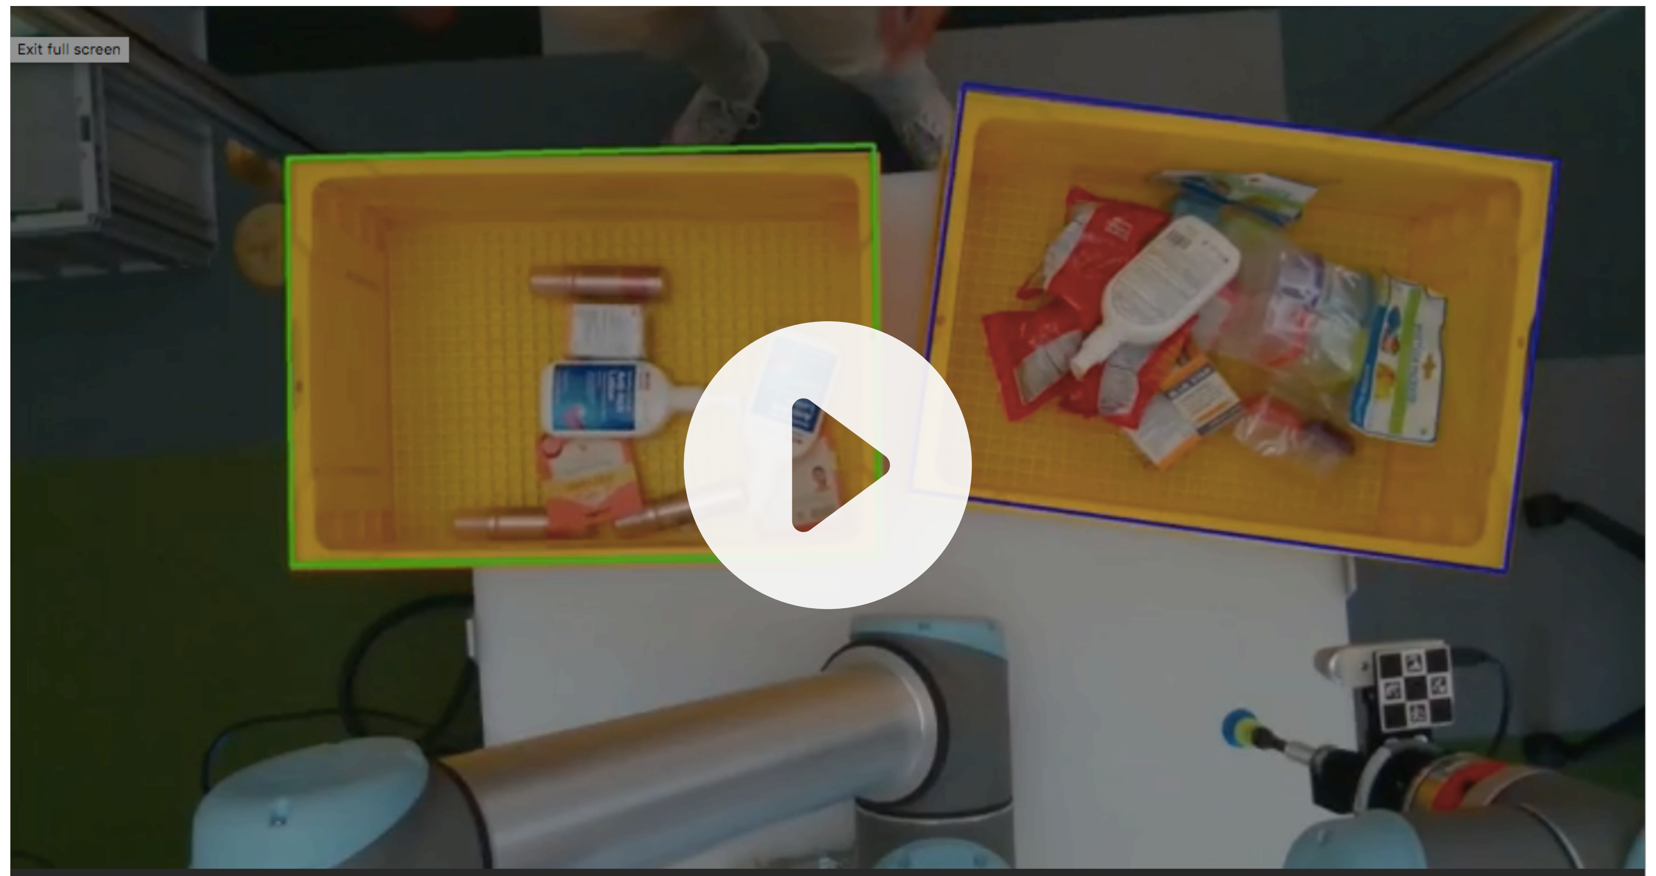

https://youtu.be/Z7bAD-OnU3k

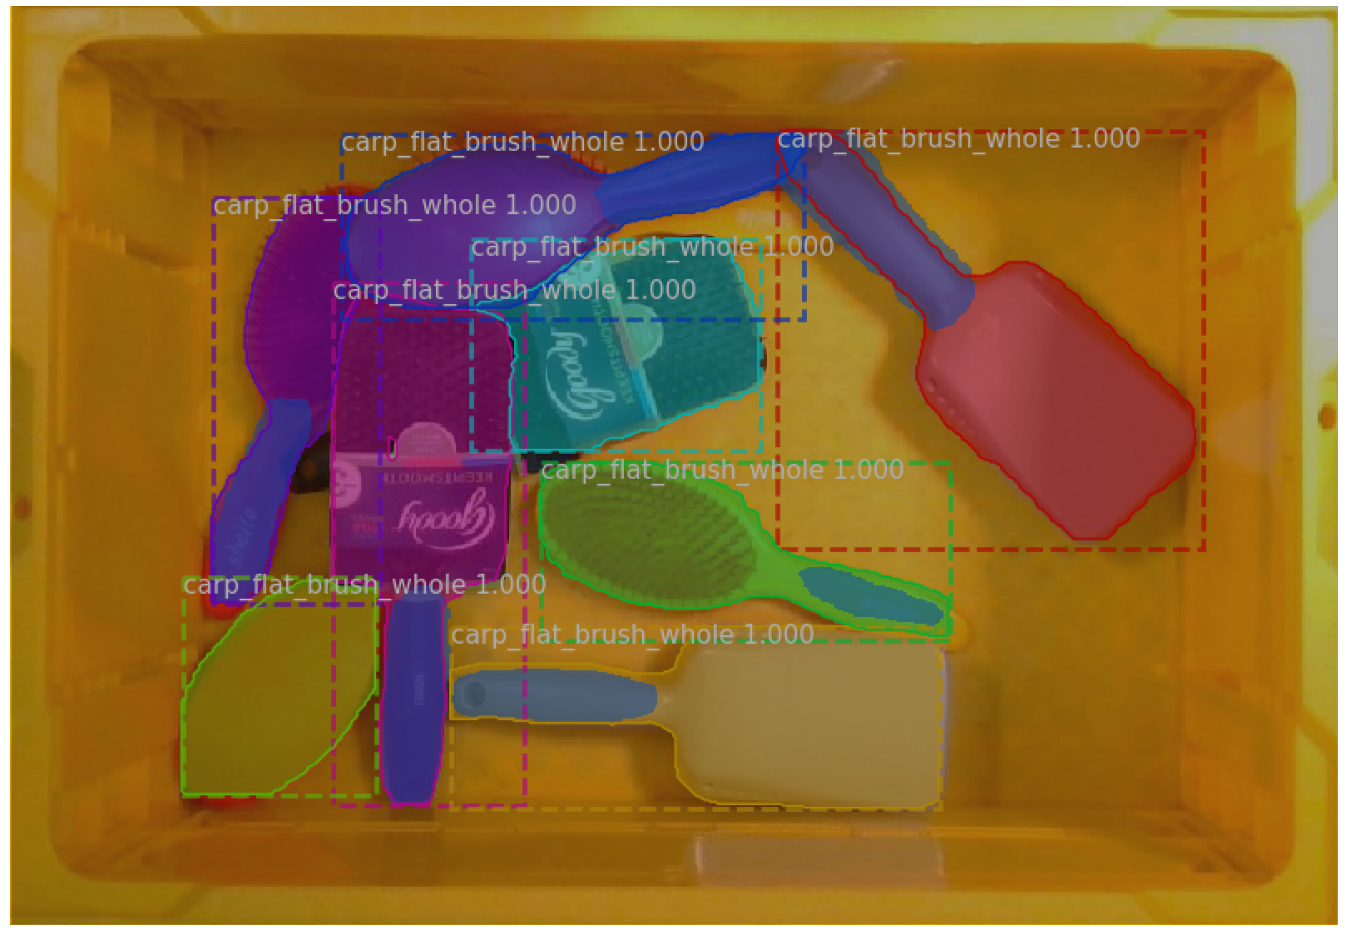

https://youtu.be/G-Cbp8Py9Bs

THE FUTURE OF ROBOTIC ML AT SCALE: ADJACENT MARKETS

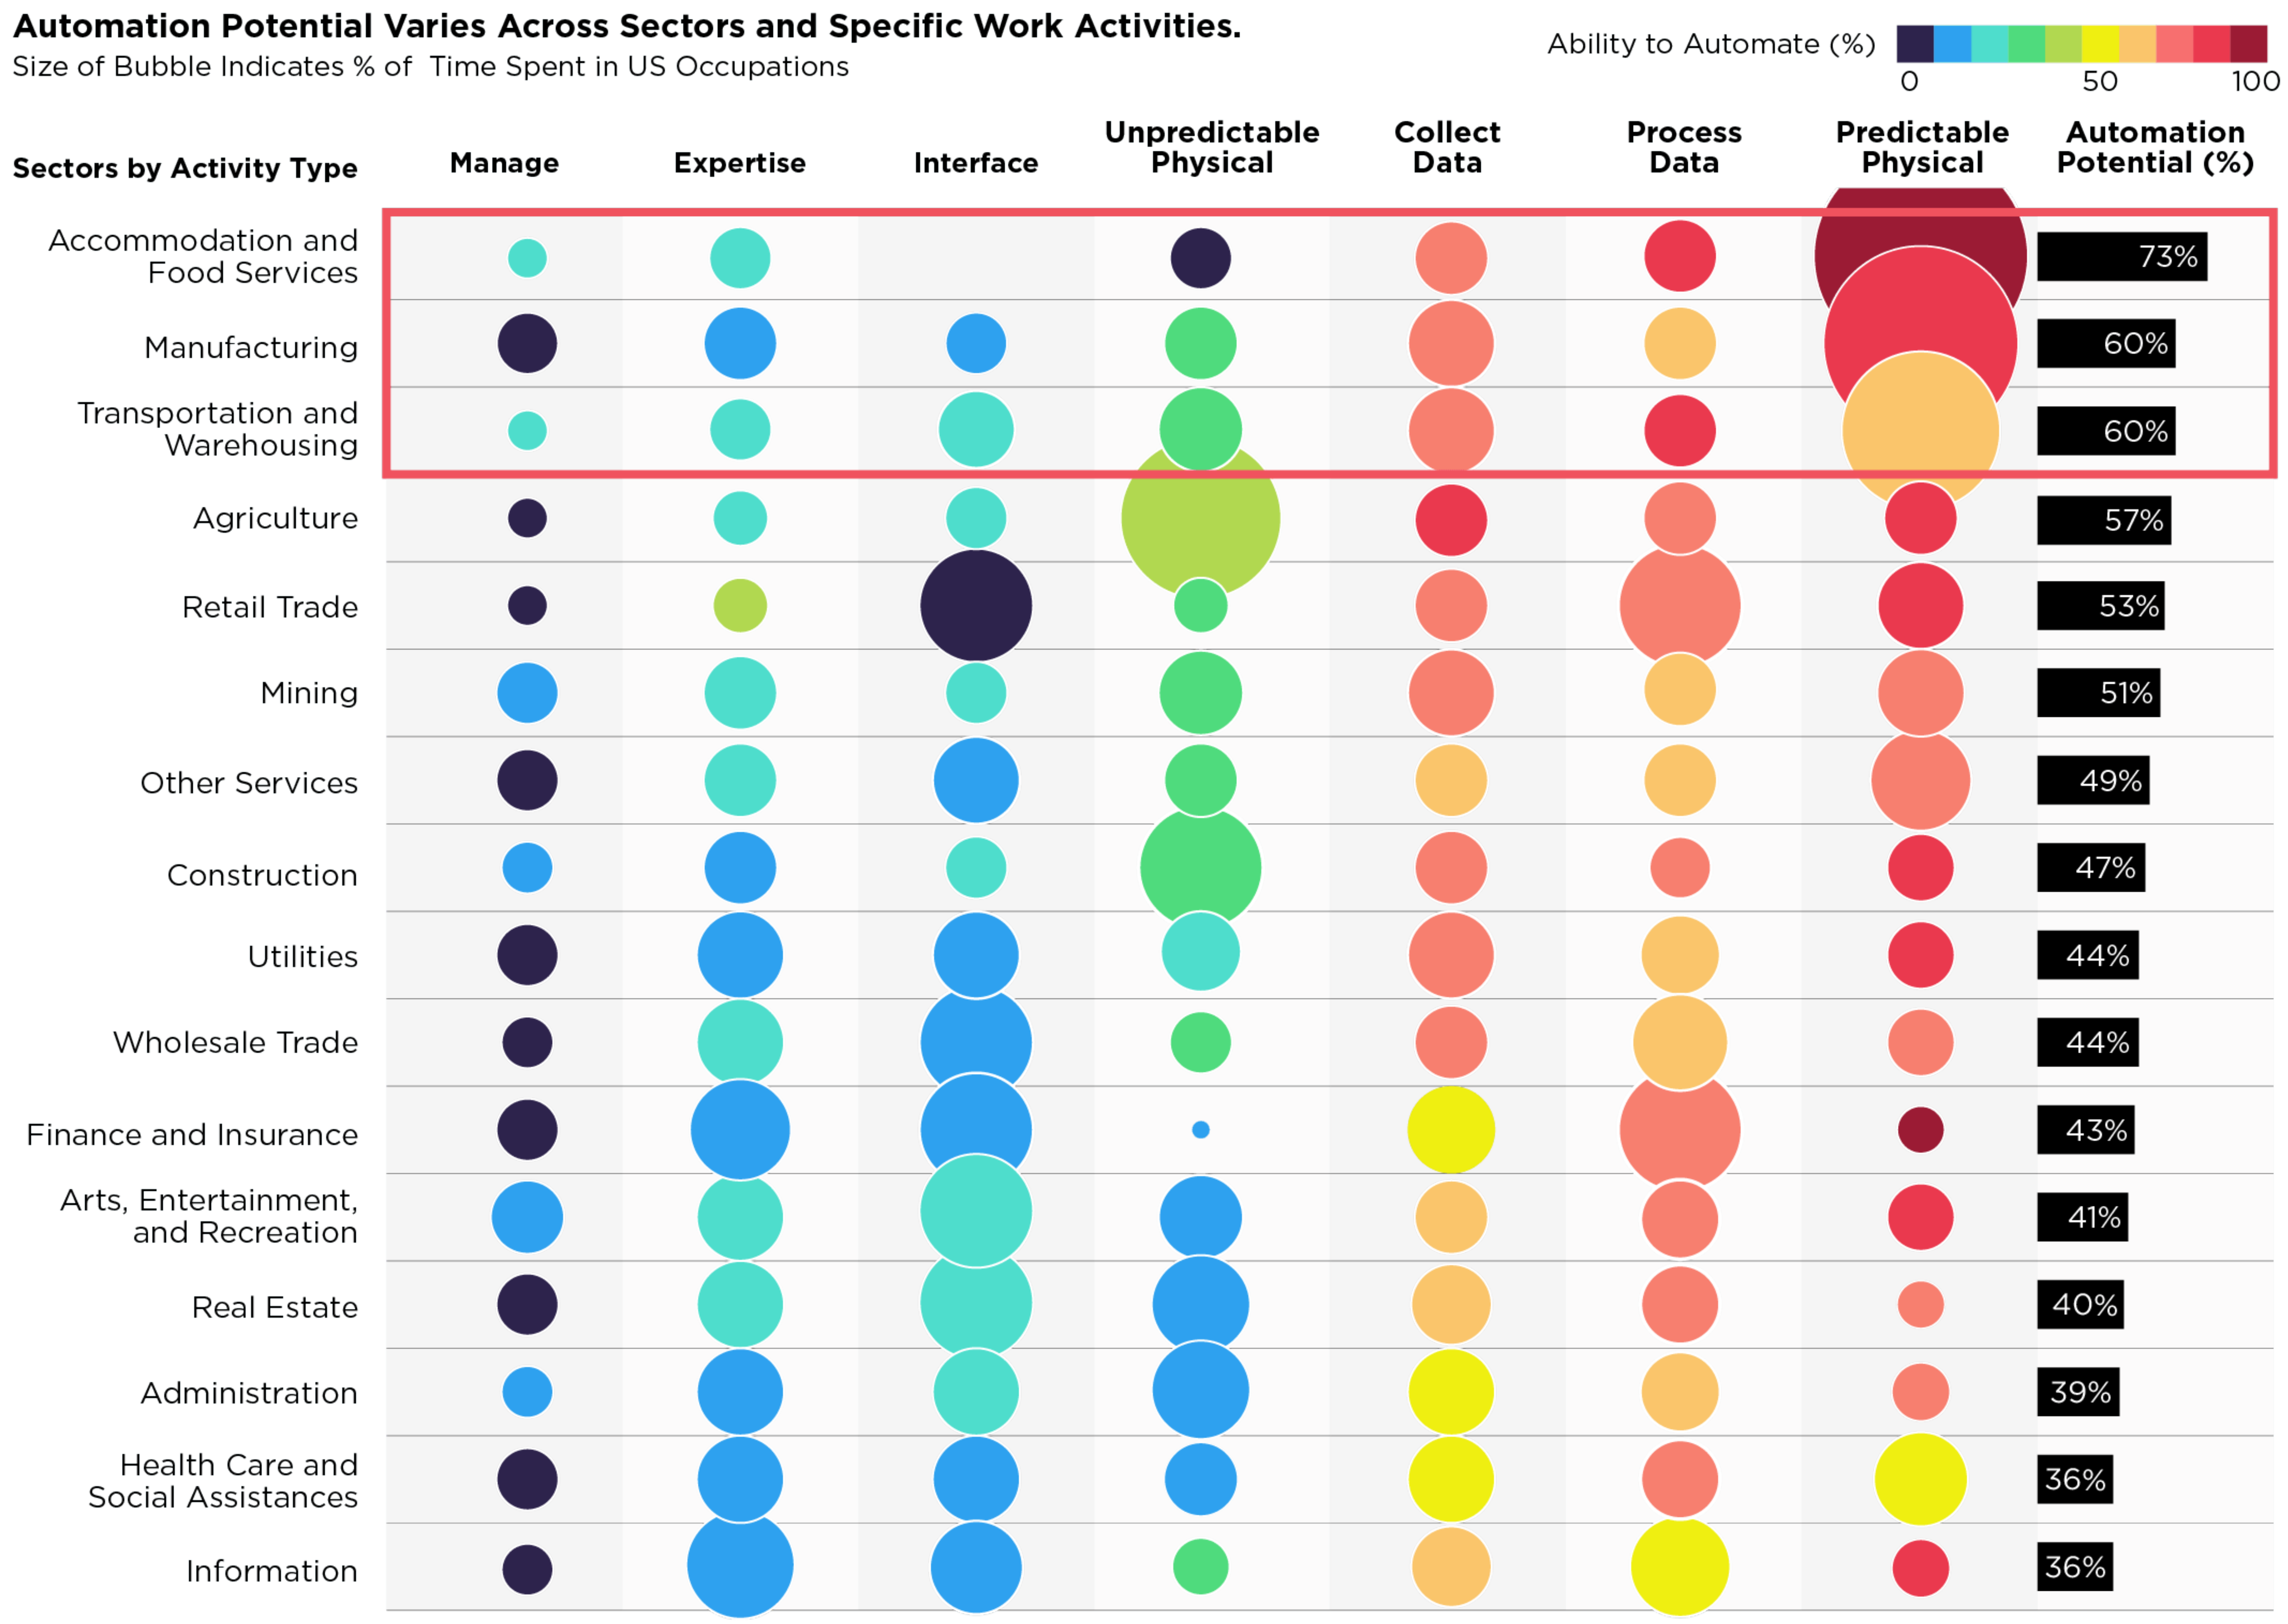

THE FUTURE OF ROBOTIC ML AT SCALE

- Automation will move from simple to complex

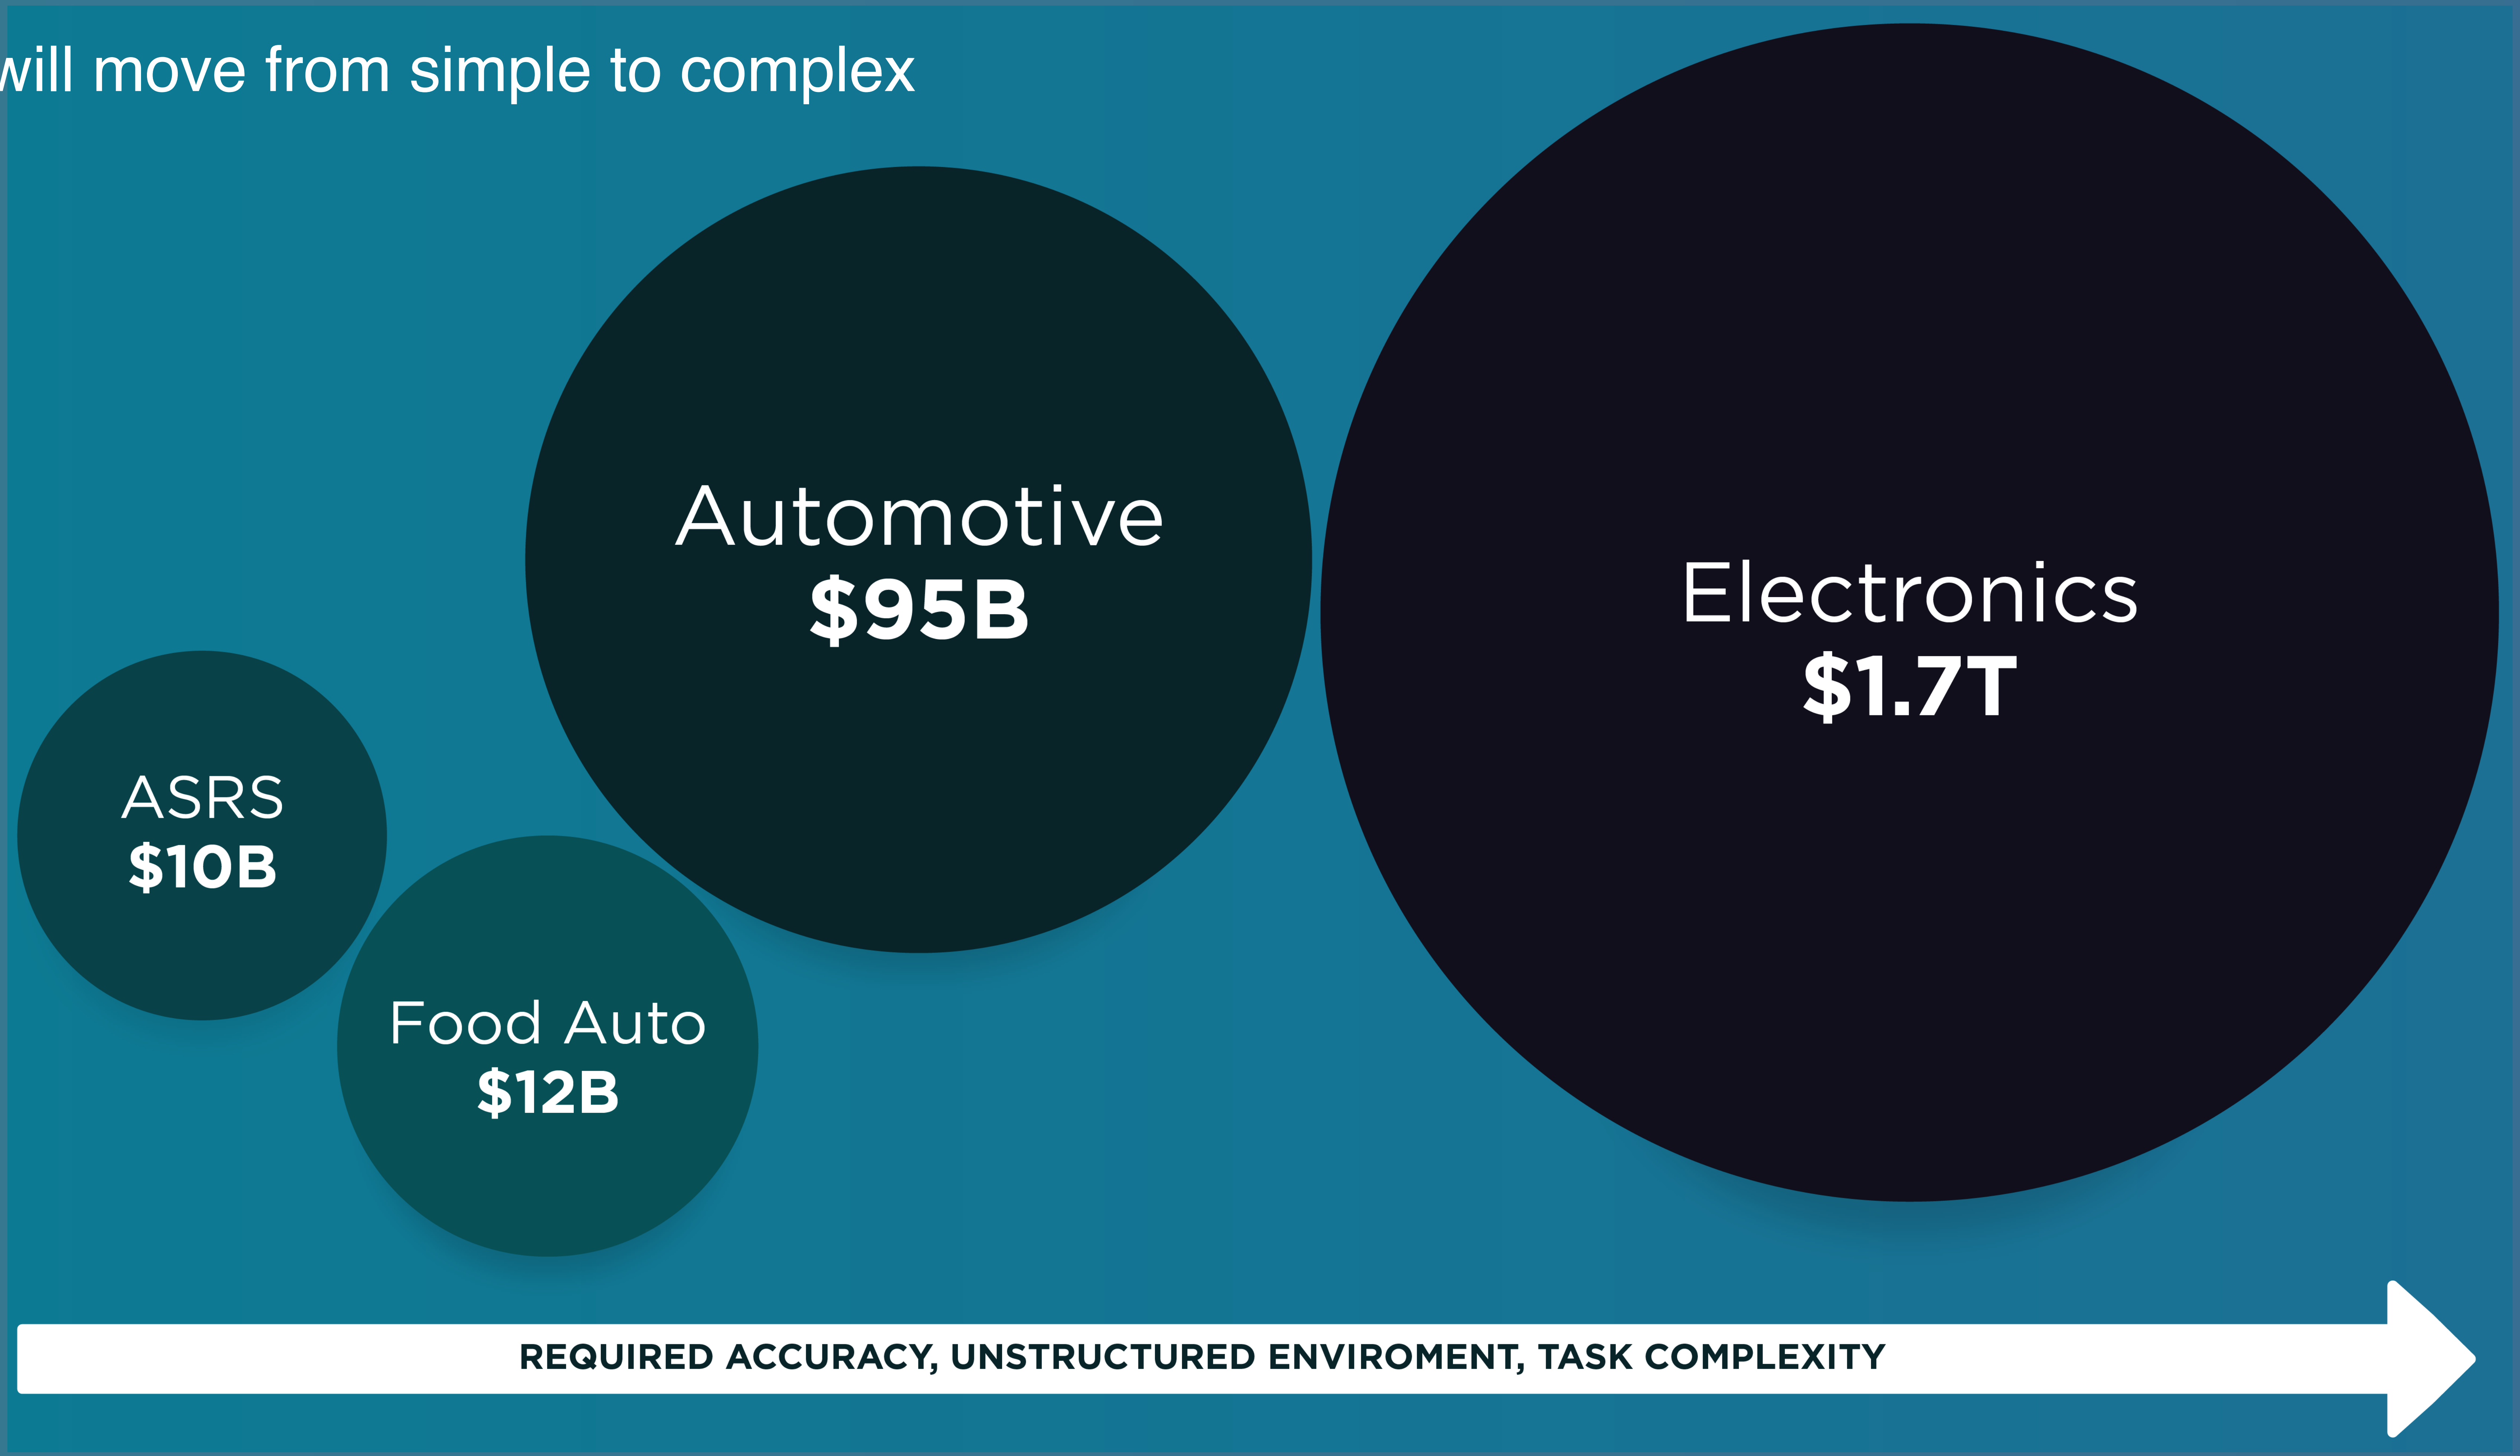

THE FUTURE OF ROBOTIC ML AT SCALE: Force Assembly

Models Include:

→ Tactile Perception (no vision)

→ Imitation Learning

→ Fit pegs/holes robustly

→ Works with off the shelf sensors

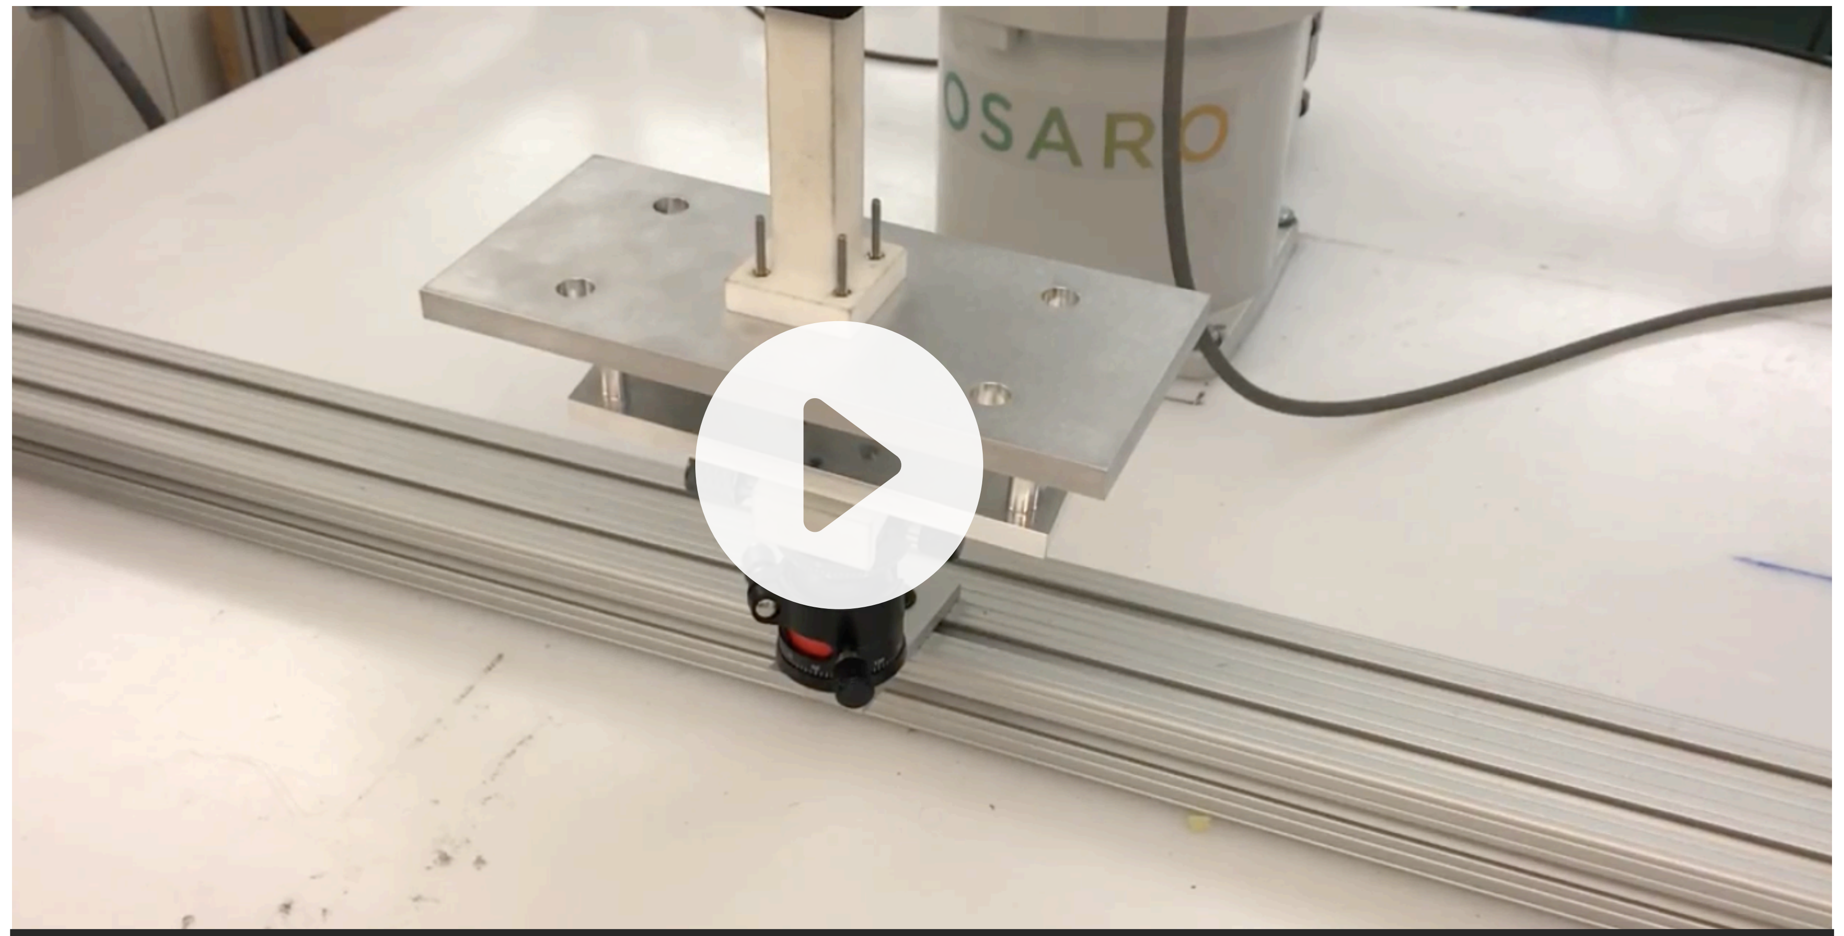

https://youtu.be/SSRNgaBGvSA

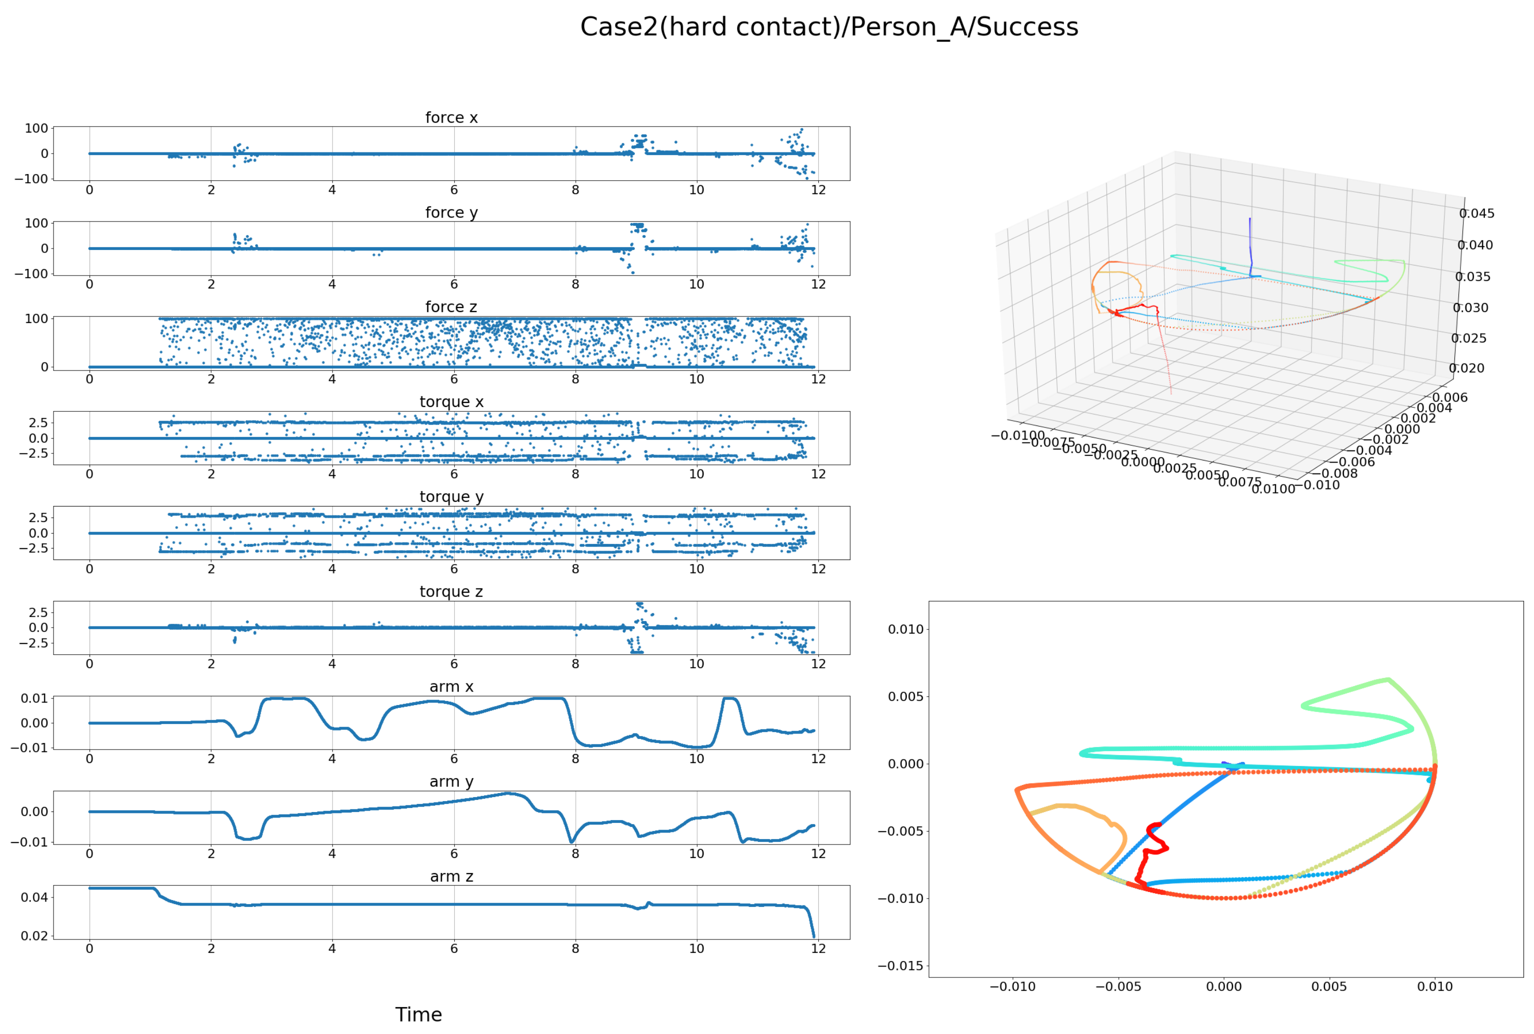

LEVEL 4 AUTONOMY:

FOOD ASSEMBLY AND PREPARATION

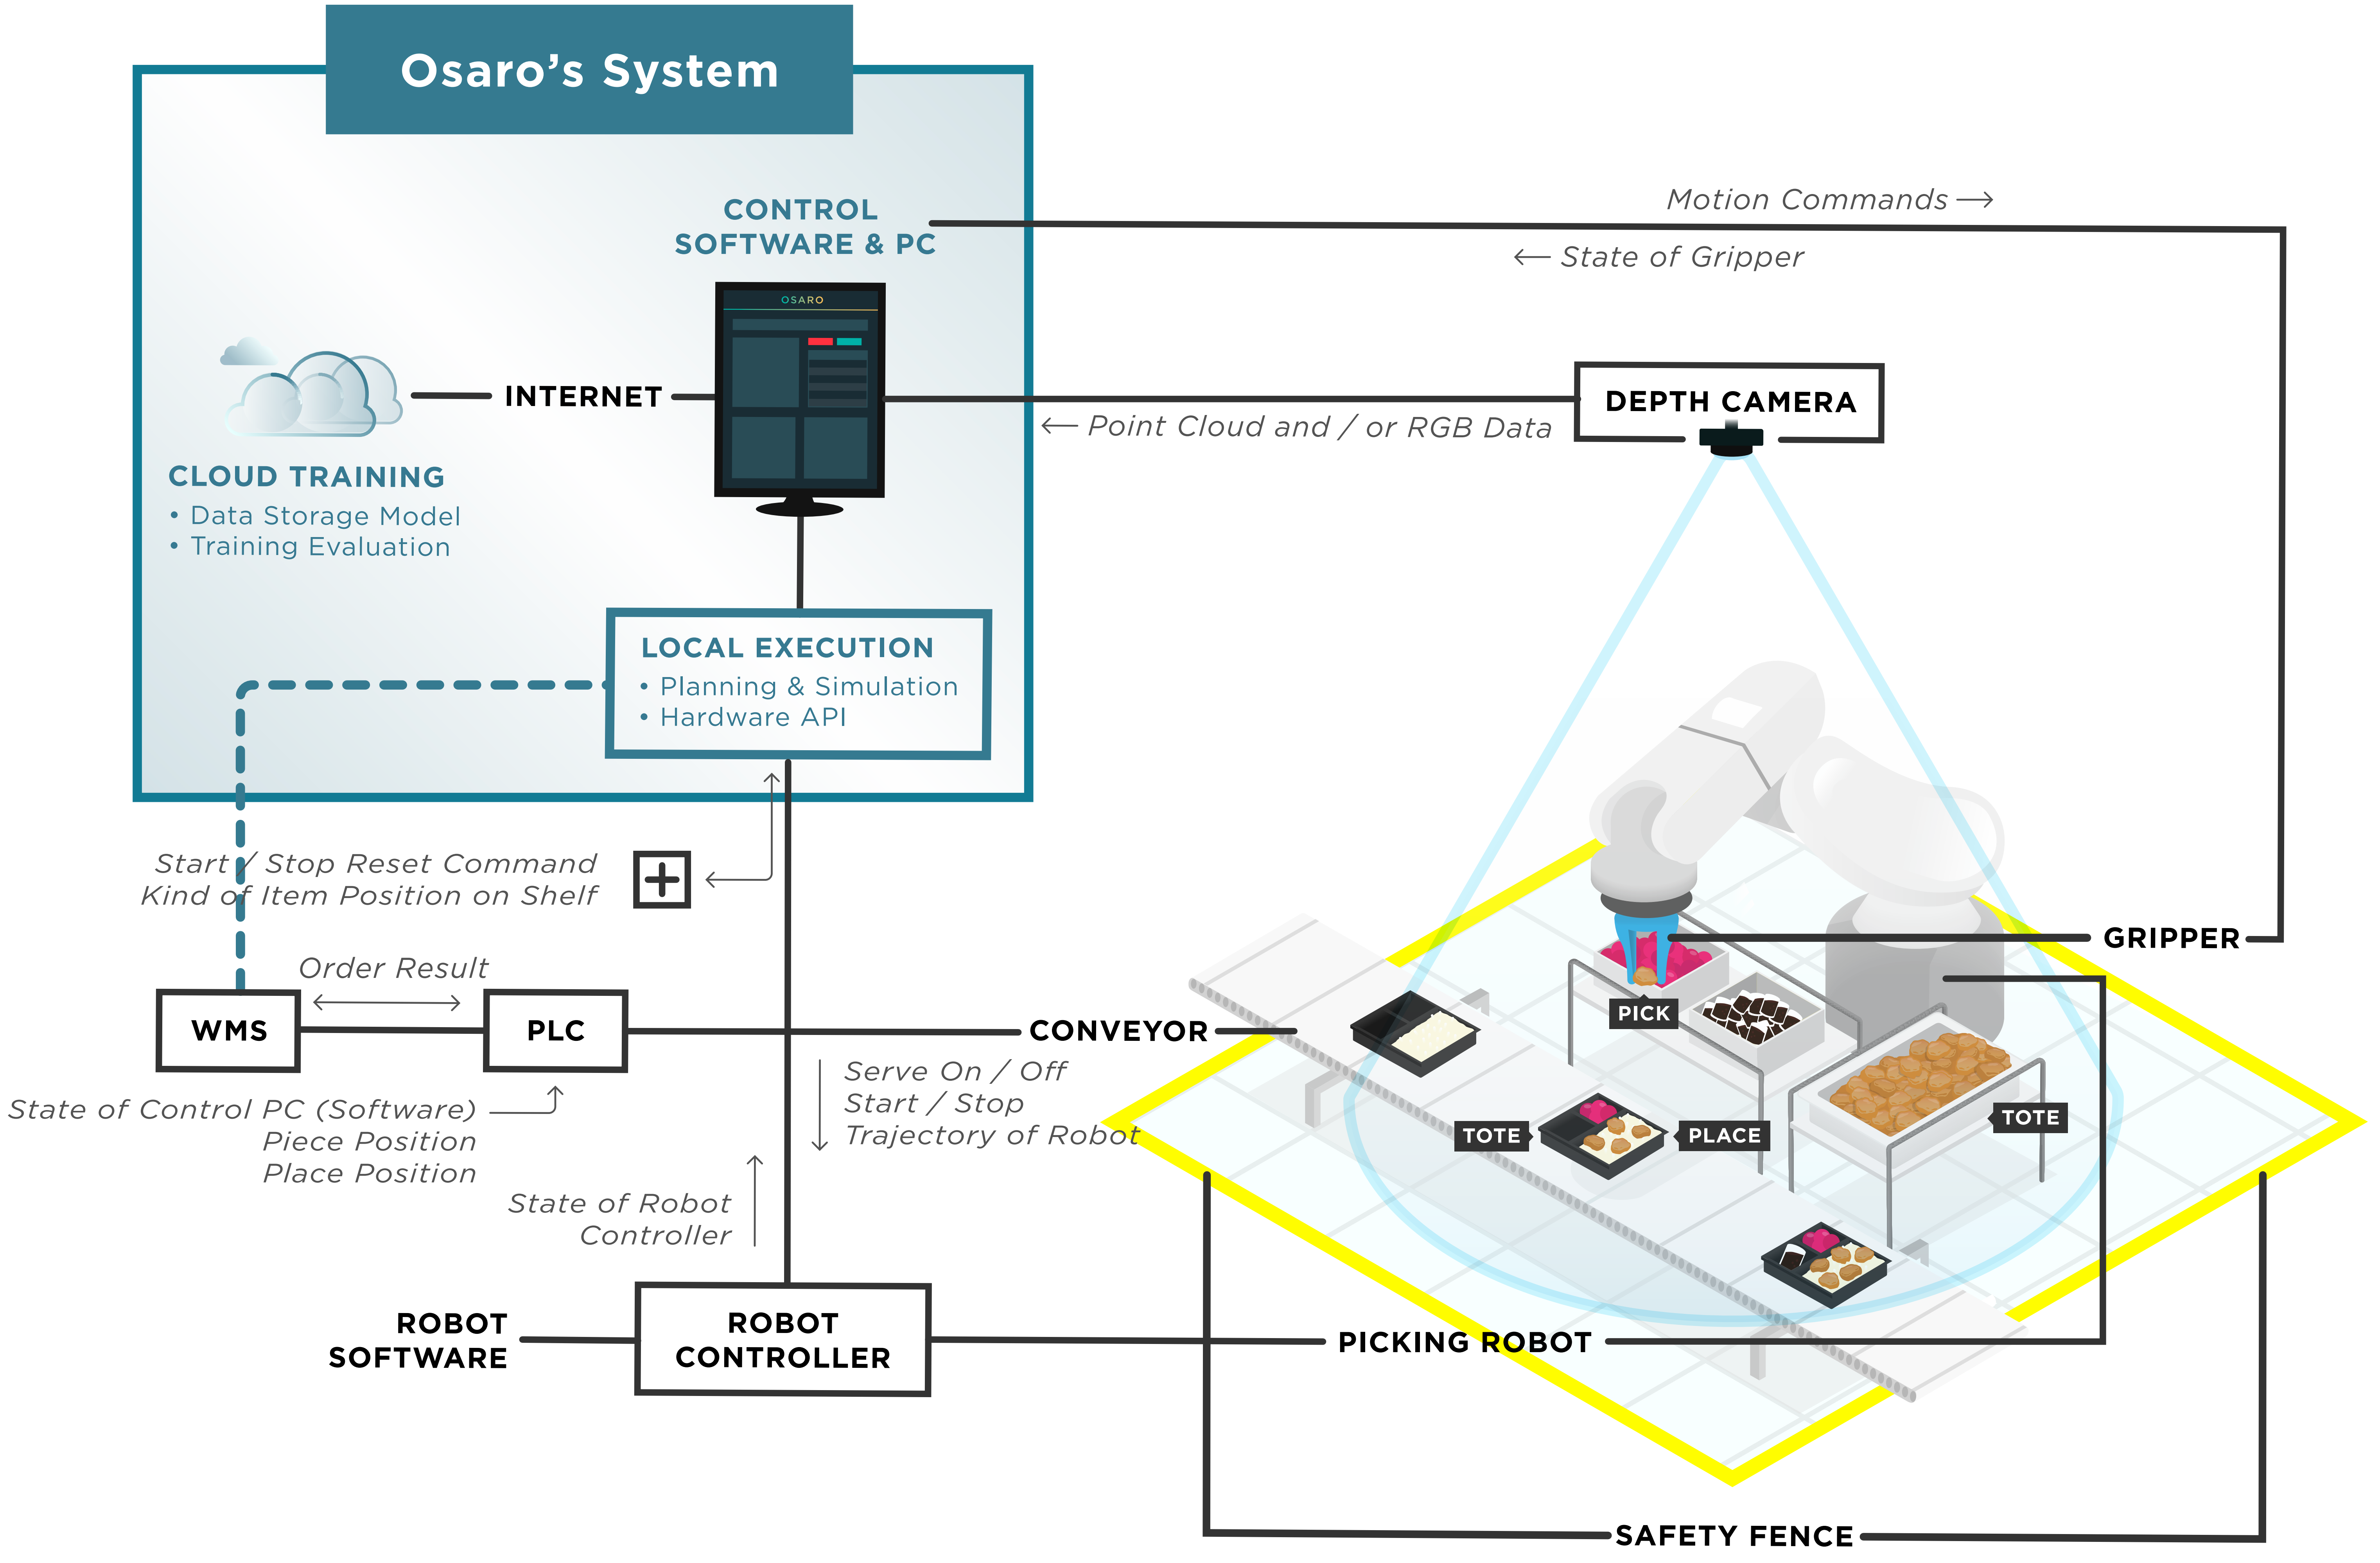

THE FUTURE OF ROBOTIC ML - CASE STUDY:

BENTO BOX ASSEMBLY WITH DENSO AT FOOMA 2018

Osaro invited to exhibit in Denso's booth for Japan's largest food automation show.

10 day integration

Drove record leads for Denso

2 weeks from invite to deployment

https://youtu.be/YbpEe-xDE1E?t=41s

Thank You

Bastiane Huang

OSARO I bastiane@osaro.com

https://medium.com/@Bastiane

We're Hiring!

https://www.osaro.com/careers

US MARKET:

$766 BILLION TOTAL WAGES FOR PREDICTABLE PHYSICAL WORK

Technical automation potential is concentrated in countries with the largest populations and/or high wages

Potential impact due to automation, adapting currently demonstrated technology (46 countries)

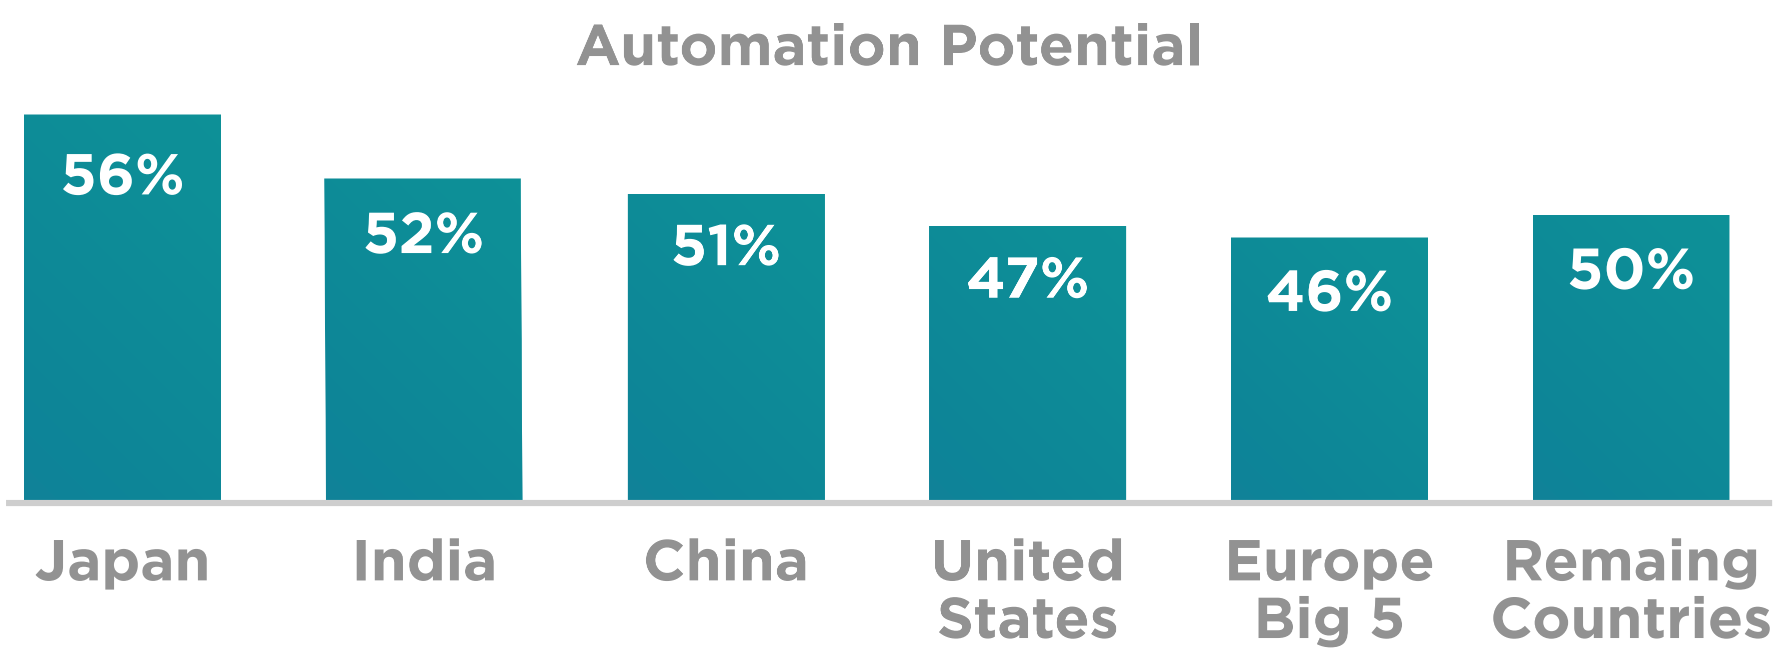

Wages Associated with Technically Automatable Activities
($ Trillion)

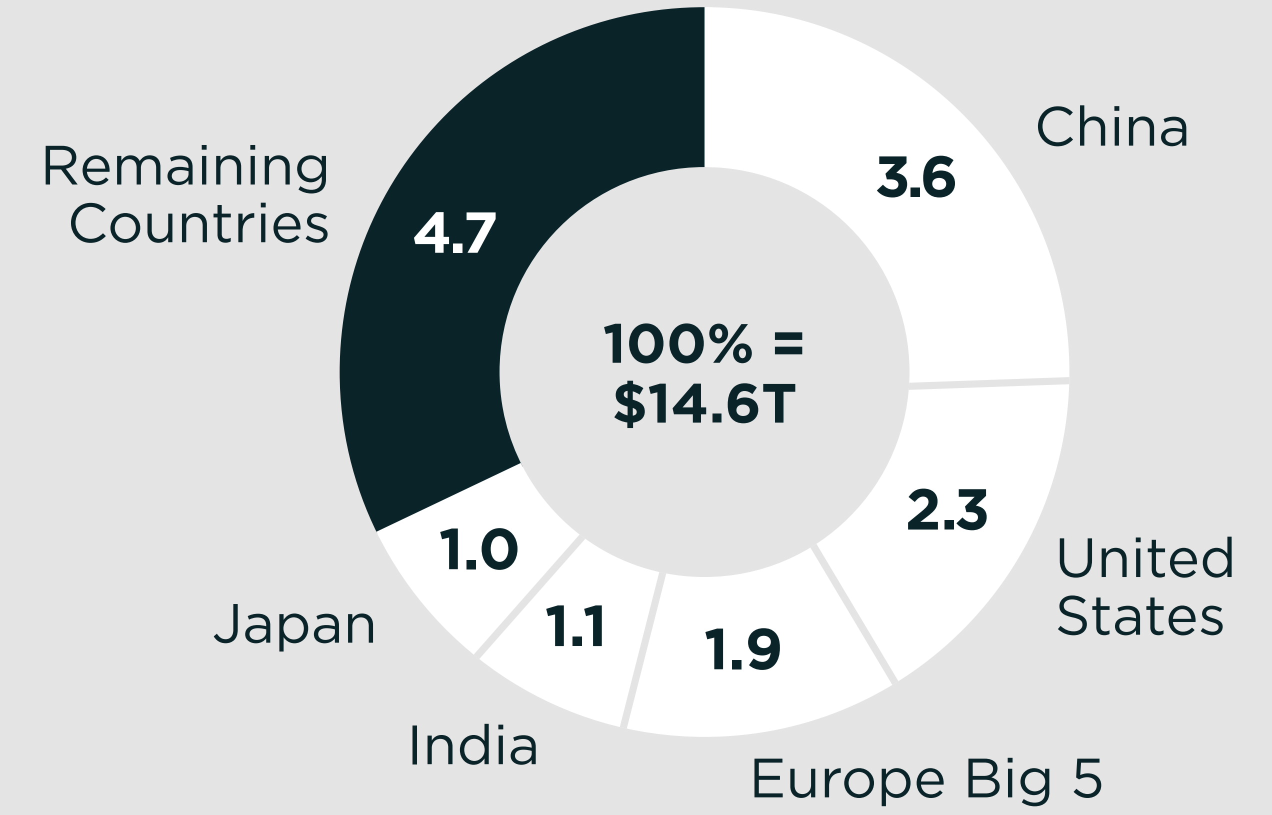

Labor Associated with Technically Automatable Activities
(Million FTE)

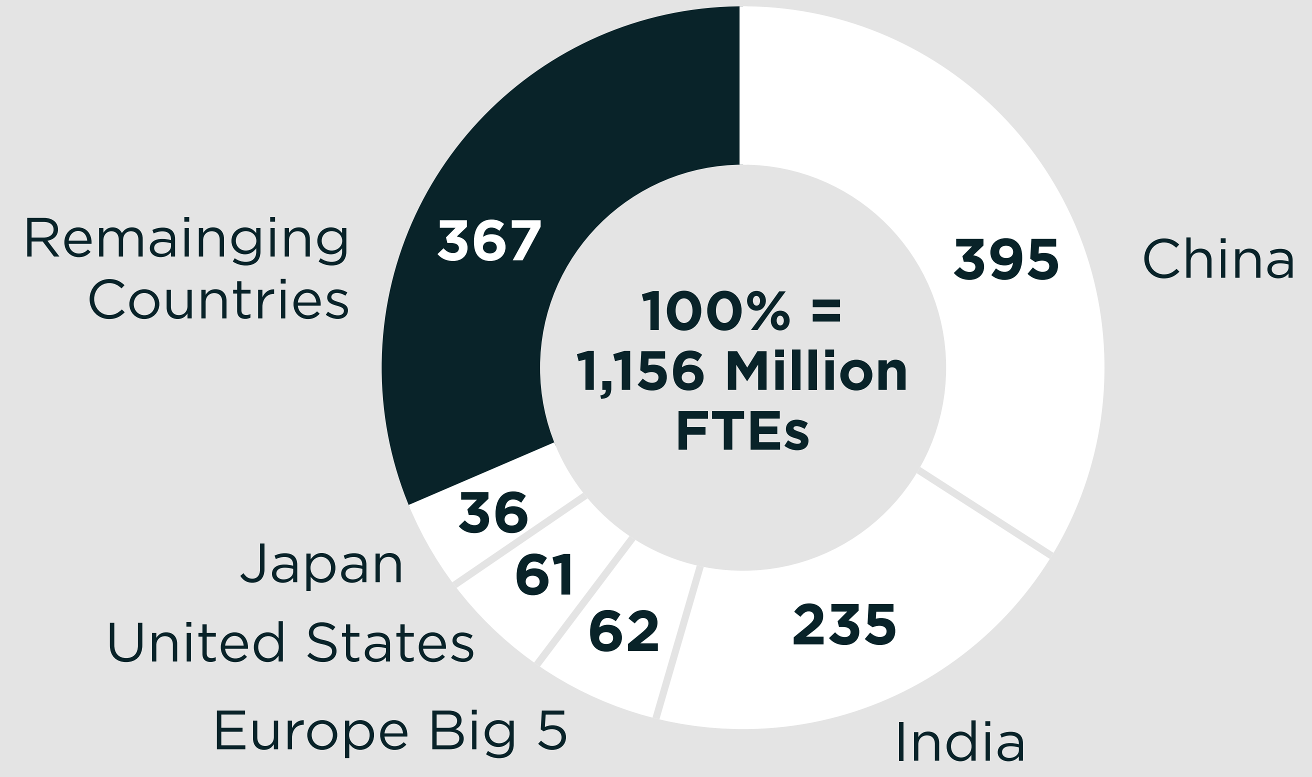

US MARKET:

$766 BILLION TOTAL WAGES FOR PREDICTABLE PHYSICAL WORK

Sectors[3]	Managing others	Applying expertise[1]	Stakeholder interactions	Unpredictable physical work[2]	Data collection	Data processing	Predictable physical work[2]
Overall average	9%	18%	20%	25%	64%	69%	78%
Accommodations and food service	2	5	22	5	8	10	48
Manufacturing	5	13	8	8	22	11	33
Agriculture	3	5	7	51	11	9	13
Transportation and warehousing	4	8	14	14	22	14	24
Retail trade	3	6	26	5	15	28	17
Mining	7	11	8	24	21	12	17
Other services	7	12	17	13	15	11	25
Construction	5	10	8	41	15	11	10
Utilities	7	14	13	19	23	13	12
Wholesale trade	5	12	24	11	17	19	12
Finance and insurance	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]

MORE AUTOMATABLE INDUSTRIES

(Source: McKinsey)

Technical feasibility: % of time spent on activities that can be automated by adapting currently demonstrated technology

0 50 100

MACHINE LEARNING

Machine learning (ML) is the scientific study of algorithms and statistical models that computer systems use to effectively perform a specific task without using explicit instructions, relying on patterns and inference instead. It is seen as a subset of artificial intelligence. Machine learning algorithms build a mathematical model of sample data, known as "training data", in order to make predictions or decisions without being explicitly programmed to perform the task.

Machine learning at its most basic is the practice of using algorithms to parse data, learn from it, and then make a determination or prediction about something in the world. So rather than hand-coding software routines with a specific set of instructions to accomplish a particular task, the machine is “trained” using large amounts of data and algorithms that give it the ability to learn how to perform the task.

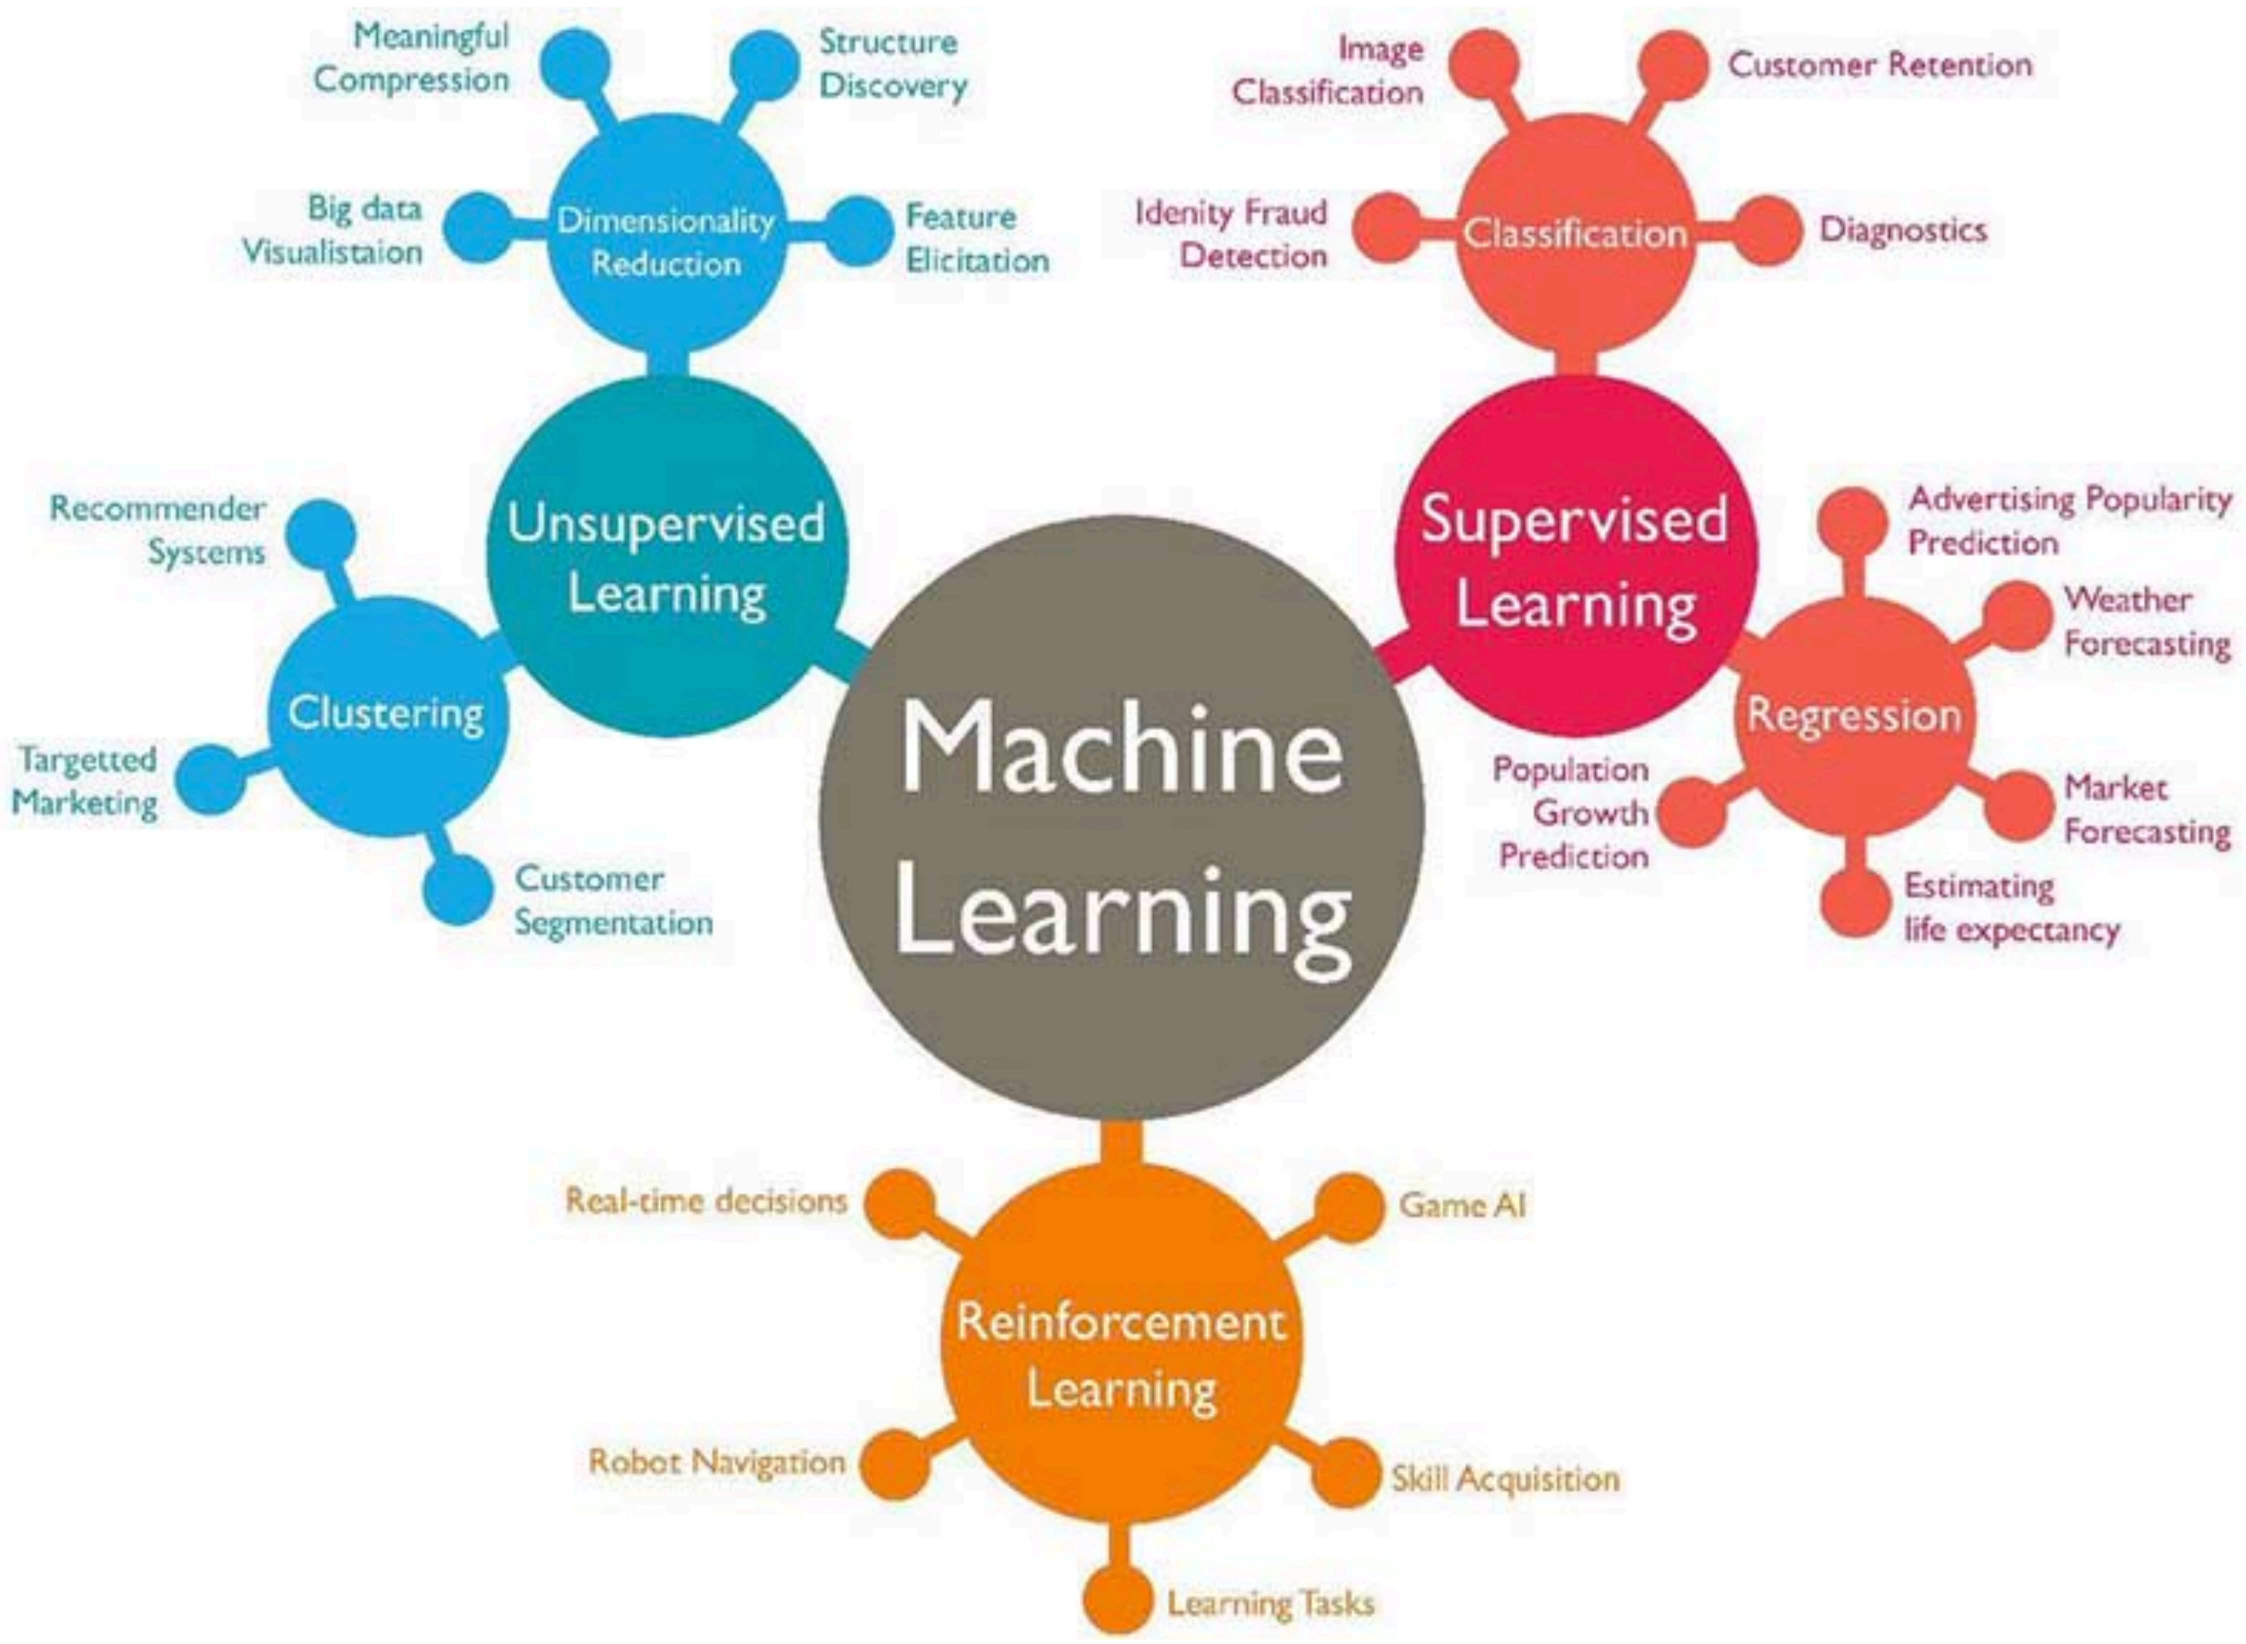

Levels of Autonomy

Level 2: Partial automation. Robots can only operate by themselves at certain times under certain conditions.

Level 3: Conditional Automation. The robot takes over actively monitoring the environment when the system is engaged. However, human must be prepared to respond to a "request to intervene"

Level 4: High automation. The robot will be able to handle most "dynamic tasks" but will still require human intervention from time to time in unusual scenarios.

Robot

A robot is a machine—especially one programmable by a computer— capable of carrying out a complex series of actions automatically.[2] Robots can be guided by an external control device or the control may be embedded within. Robots may be constructed to take on human form but most robots are machines designed to perform a task with no regard to how they look.

GOOGLE RESEARCH:
ARM FARM

- Result: 14 arms, 800k grasps improves random policy from 30% to 70% in 3 months.
- Importance: early demonstration of the parallelism of optimization across robots.
- Problems: super slow, not practical in real-world use cases

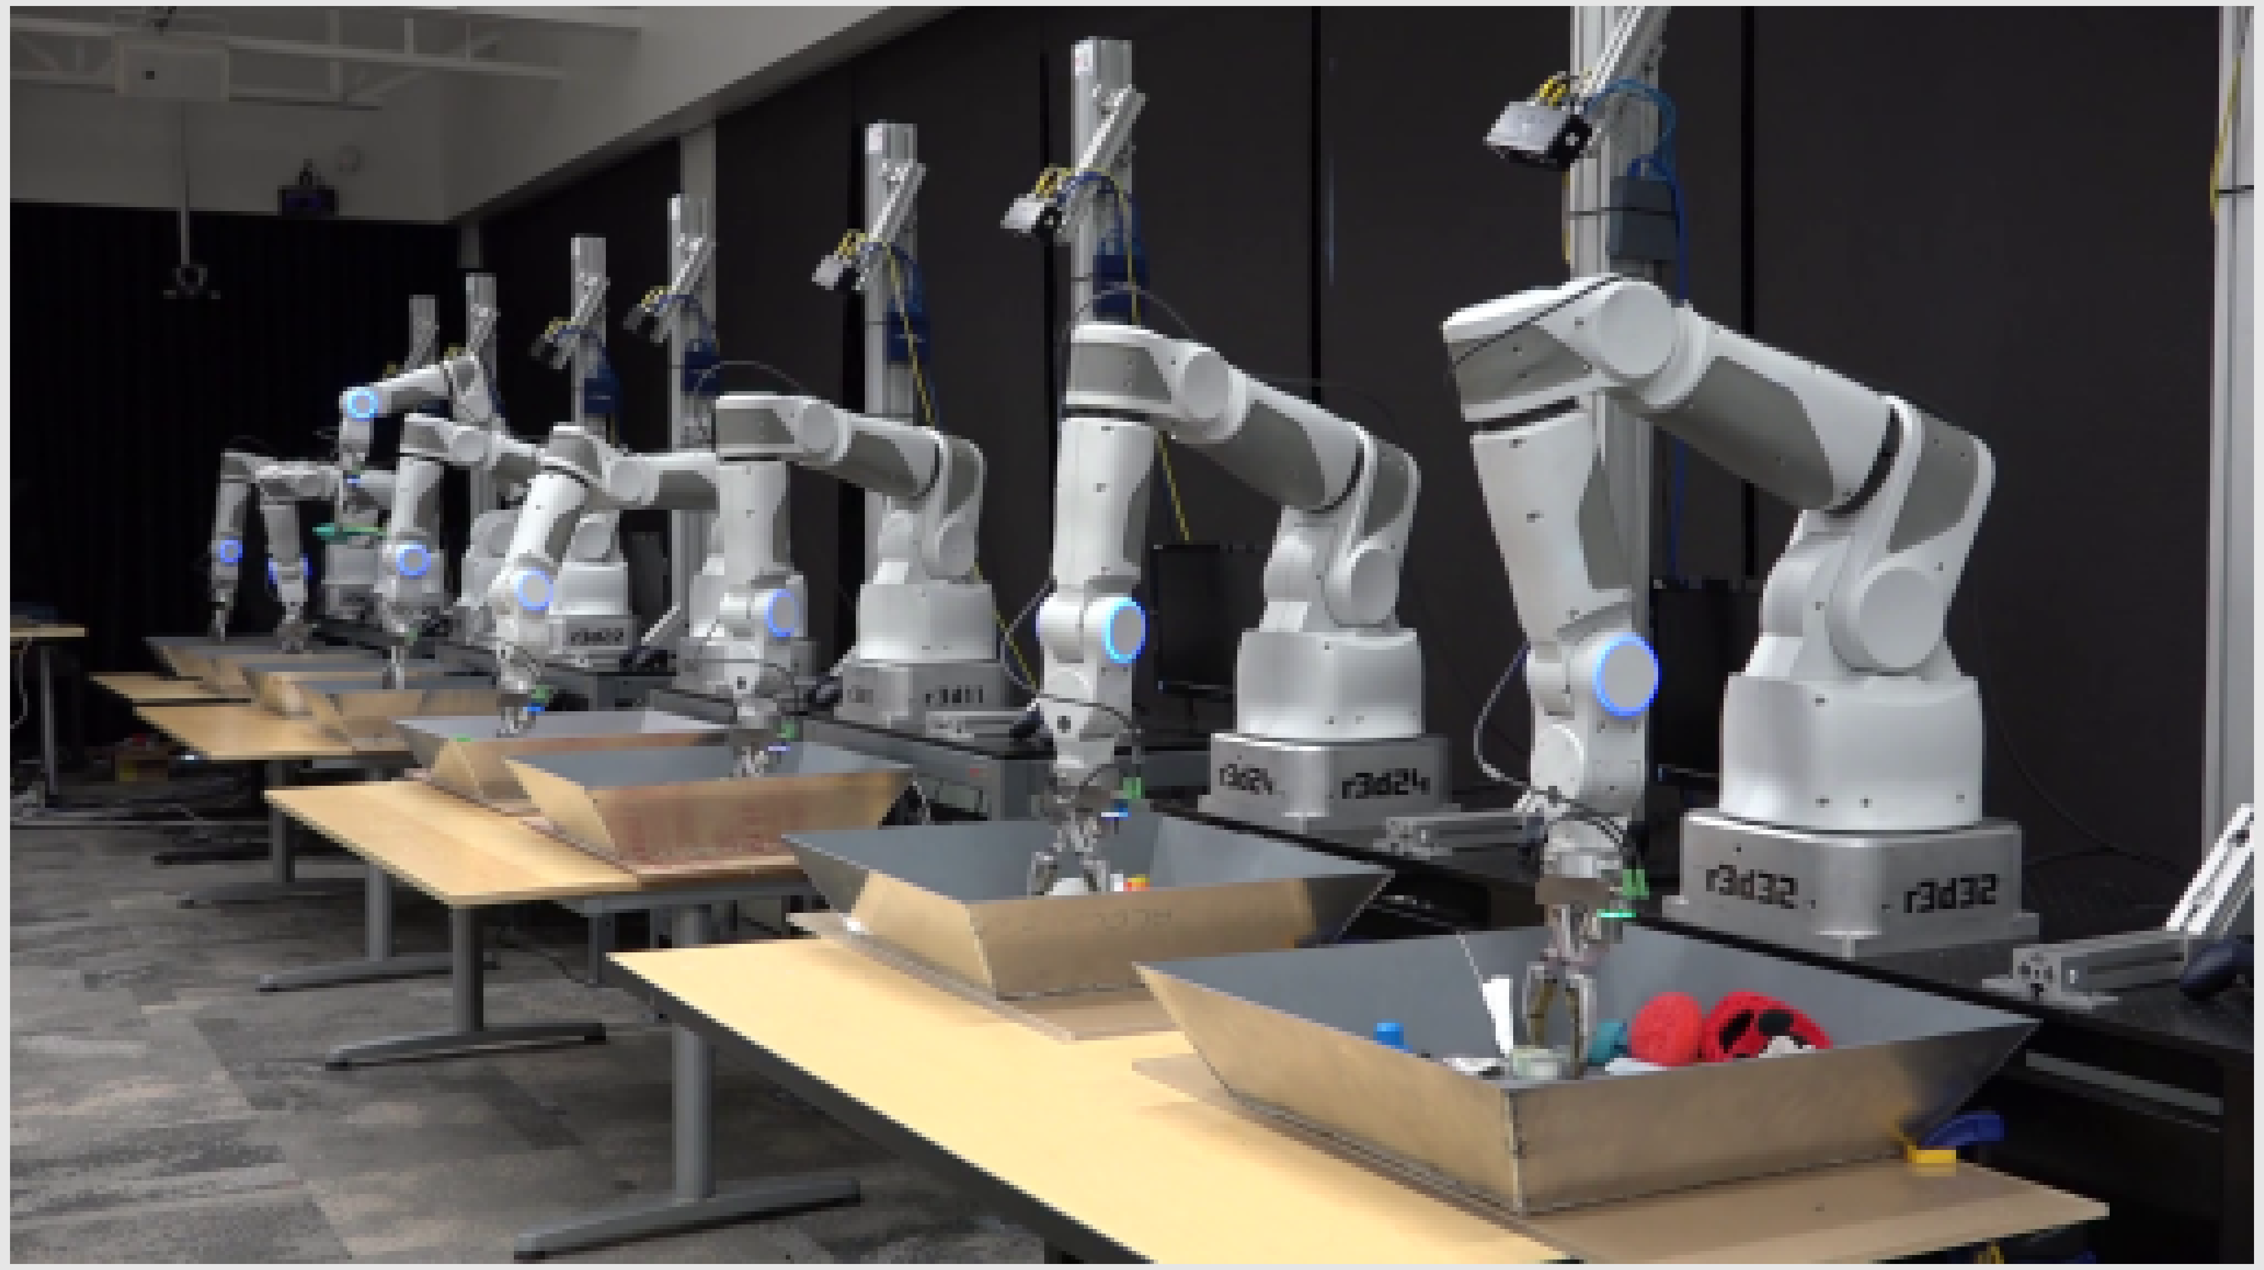

- "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection" Levine et al. March 2016.

EARLY WORK
AT OSARO (2015)

- Result: 100-1000x speedup.

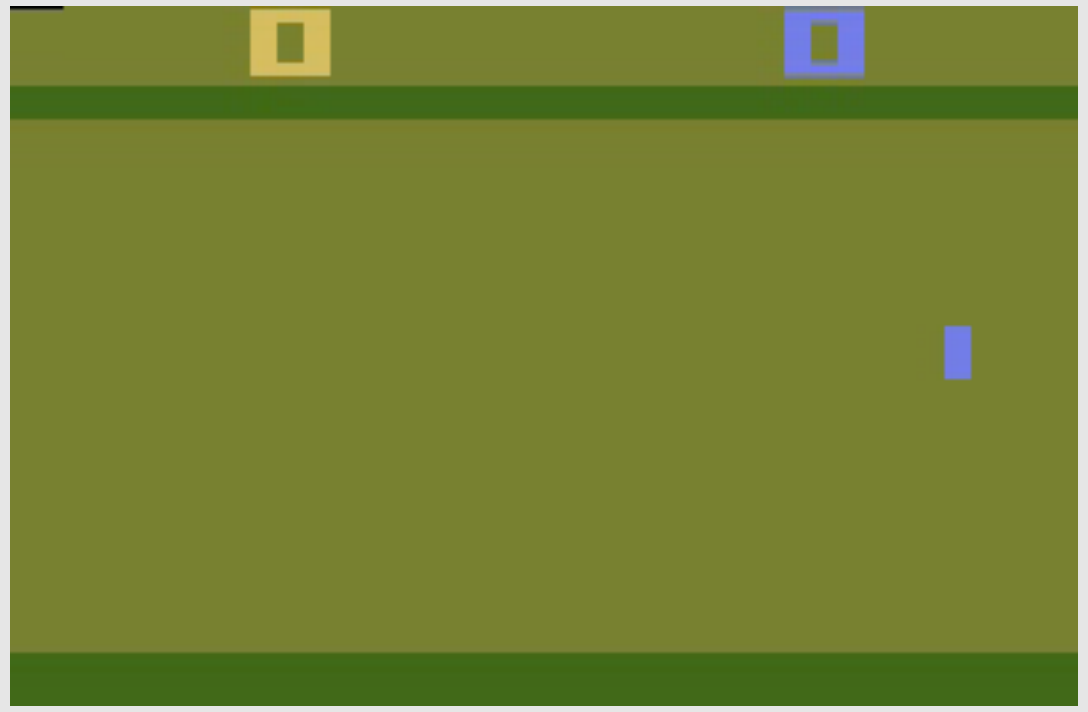

NAIVE
DRL

https://youtu.be/9X25FpI16kA

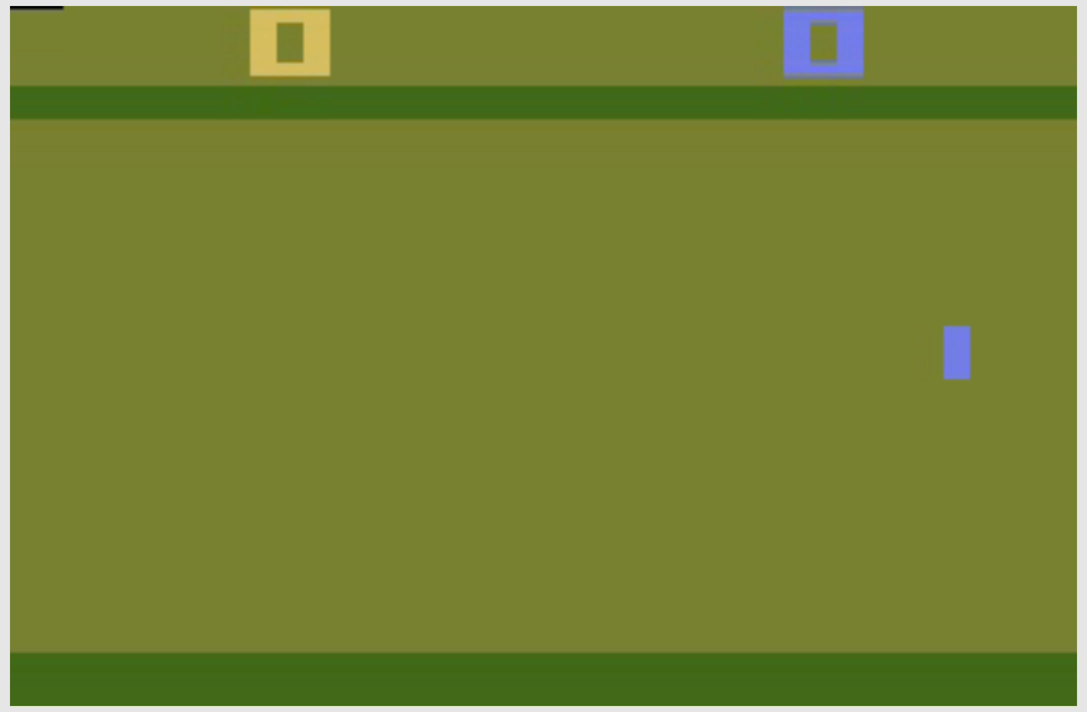

MENTOR
DRL

https://youtu.be/dvYKpMURNVE