Robotics and AI

Brian Bingham

Mobile Robots:
Flying, driving, walking, swimming, grasping

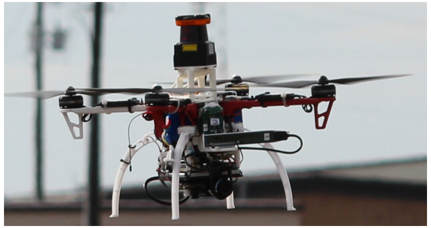

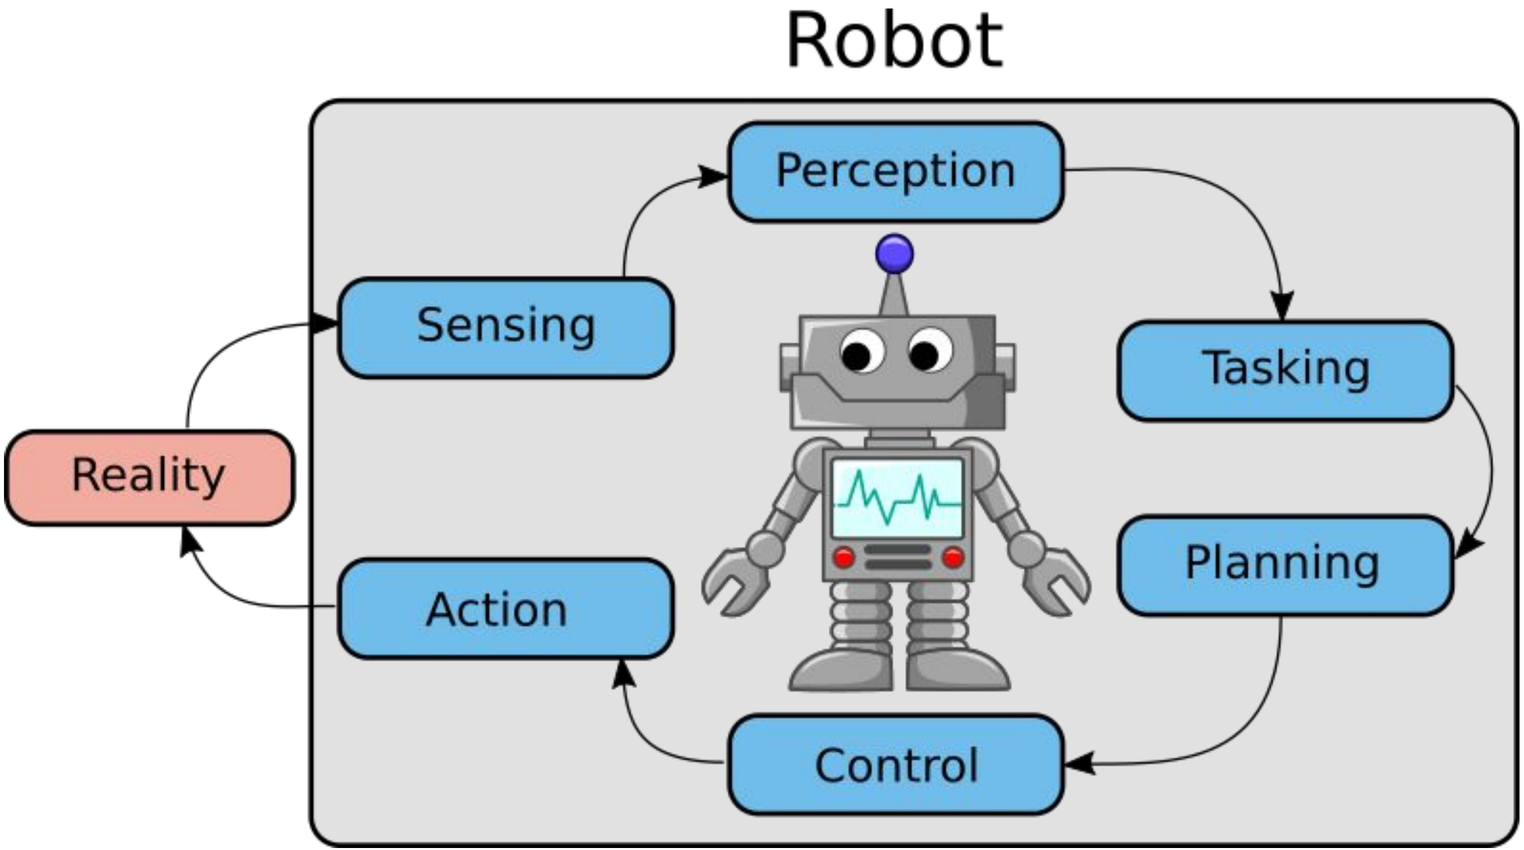

Robotics: “The intelligent connection of perception to action”,
J. M. Bradley, MIT AI Lab, 1986


SLOCUM

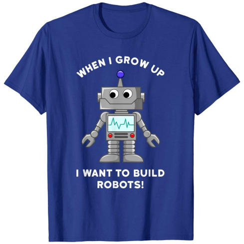
WHEN I GROW UP
I WANT TO BUILD
ROBOTS!

Challenges to engineer a general purpose robot

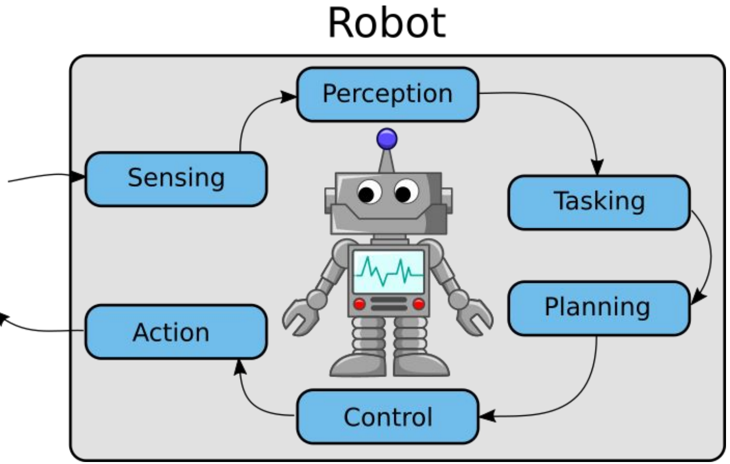

Deep Learning Successes

Computer Vision

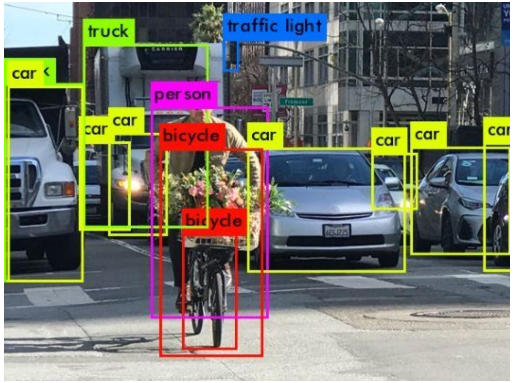

Speech Recognition

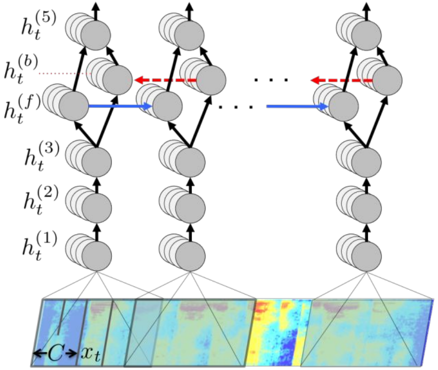

Machine Translation

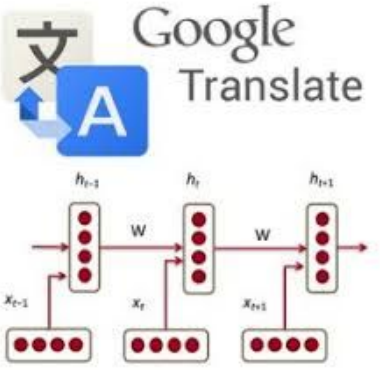

Robotics

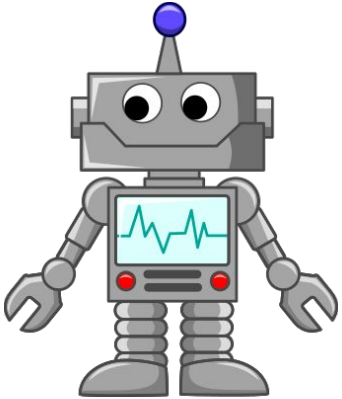

14 M Label Images

Millions of phrases

40 B sentence pairs

Training data?
Labels?

Deep Learning Successes

Mobile Learning Machines (MLMs)

Typically we have **engineered solutions** to connect sensing to actuation,

Can we design systems to **learn solutions**?

Learn: acquire a new capacity for action

- Machine A is more powerful than B when

  A can learn functions B cannot

Web of Science: Robotics and Learning Citations

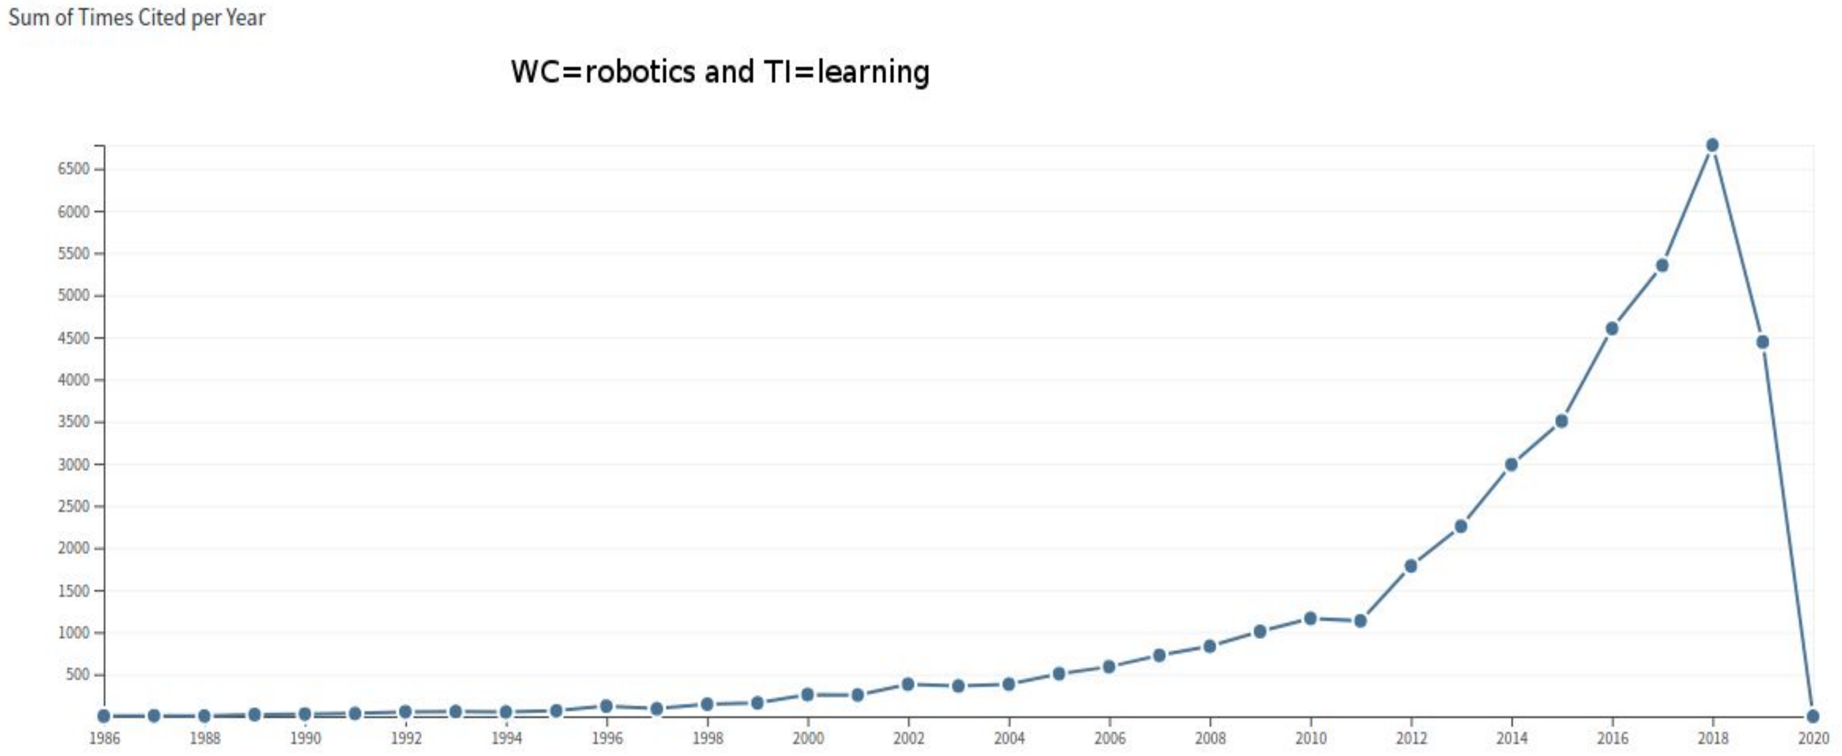

The Ice Cream Model

Scoops

Swirls

Shakes

Scoops: Perception, Payloads

Leveraging advances in AI/ML as **independent** components of a robotics system.

The AI/ML doesn't know its on a robot.

Learning

Robot

Automatic Target Recognition (ATR)

Teach a mobile robot to find an object by leveraging ML-based perception

- Learned perception: Mathworks AlexNet CNN (ImageNet)

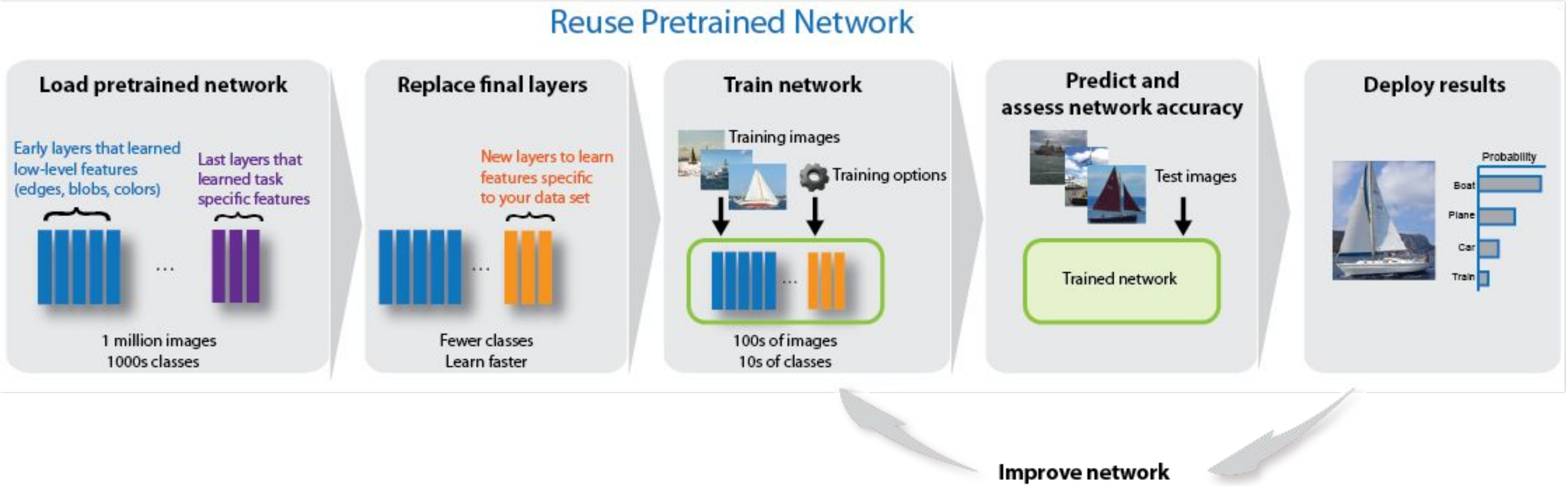

- Engineered tasking, planning, control and action
    - Wander, explore, avoid obstacles, etc.

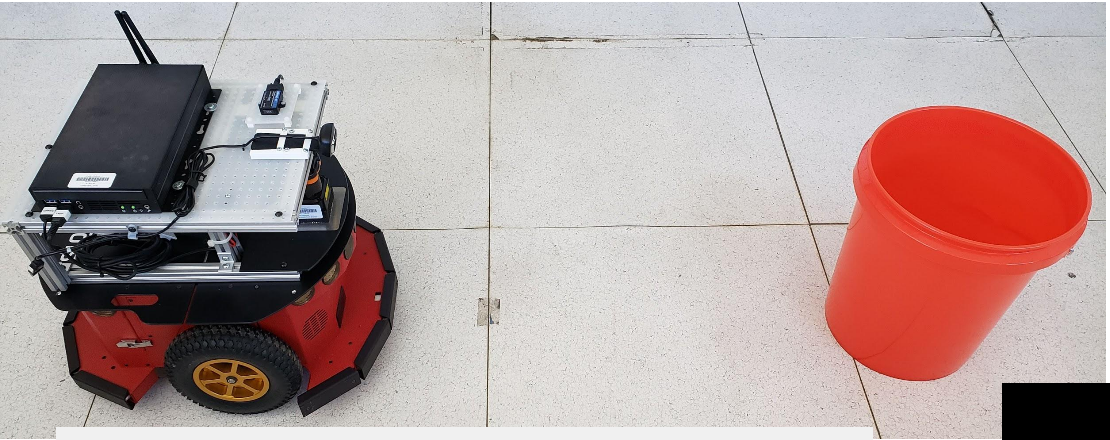

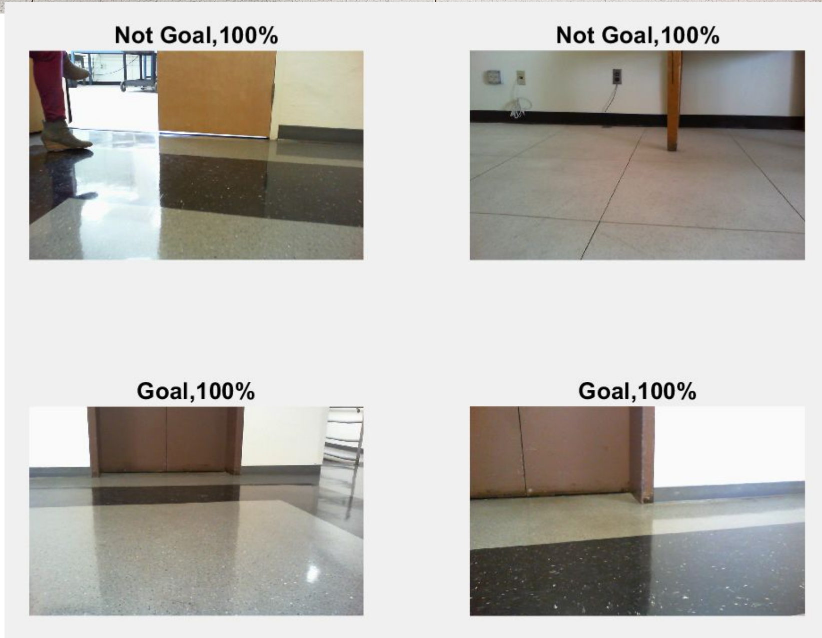

Not Goal,100%
Not Goal,100%
Goal,100%
Goal,100%

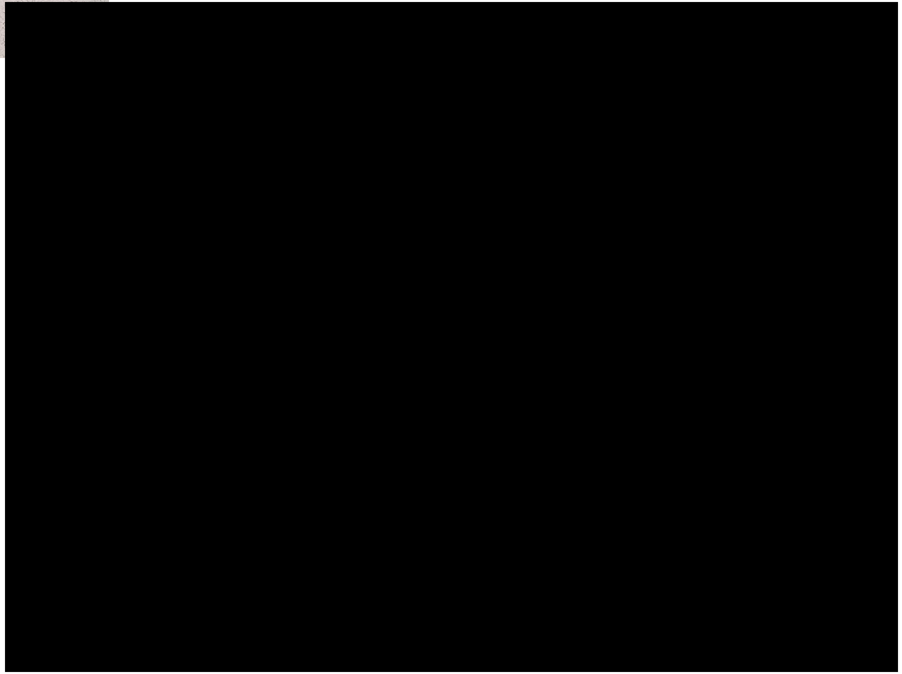

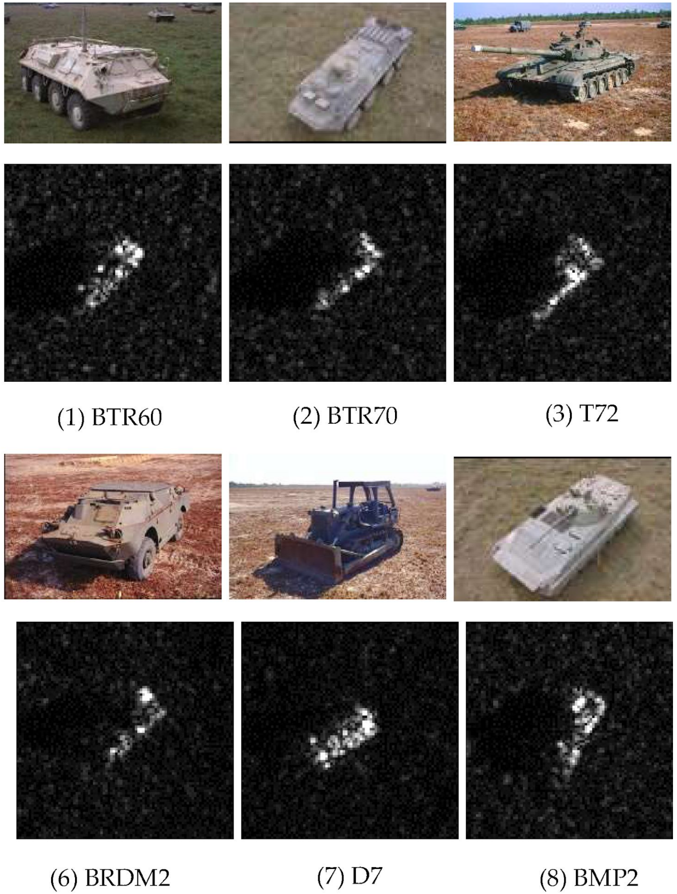
(1) BTR60
(2) BTR70
(3) T72
(6) BRDM2
(7) D7
(8) BMP2

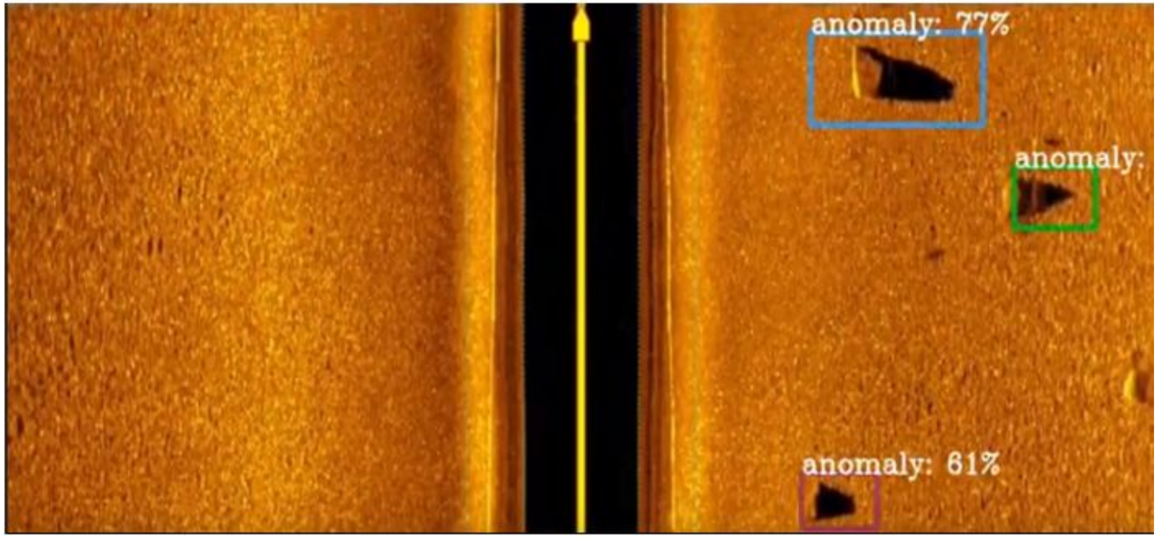
anomaly: 77%
anomaly:
anomaly: 61%

Swirls: Coupling Motion and Learning

Customizing AI/ML for robotic deployment.

New tools for robotics systems

Adaptive sampling

System identification

Shakes: Inseparable Integration

Robots that learn by doing.
Robot learns optimal behavior through trial-and-error interactions with the environment.

Where learning replaces engineering at the system level.

Tools for the design of hard-to-engineer capabilities that generalize.

Reinforcement Learning

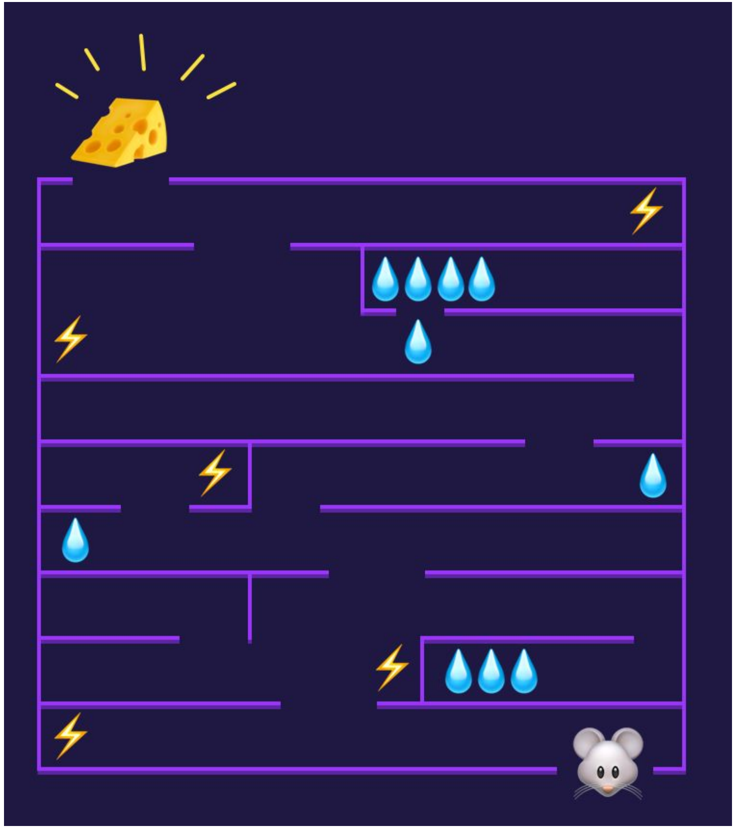

Automated sequential decision making.

Learn a *policy* to map *states* of the system to *actions*.

$$\mathbf{a_t} = \pi(\mathbf{s_t})$$

Maximize rewards

$$\max_{\pi} \sum_{t=0}^{\infty} \gamma^t r(\mathbf{s_t}, \mathbf{a_t})$$

Balance *exploration* and *exploitation*

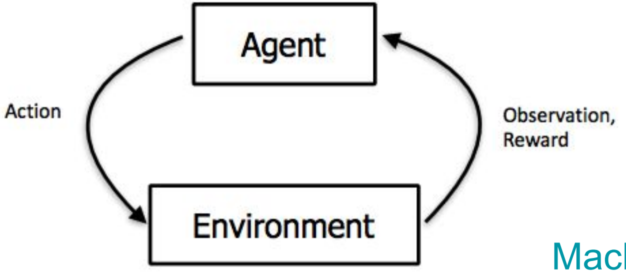

[Machine Learning for Humans]

Jens Kober and Jan Peters, Technische Universität Darmstadt and Max Planck Institute for Intelligent Systems.

Learning Hand-Eye Coordination for Robotic Grasping

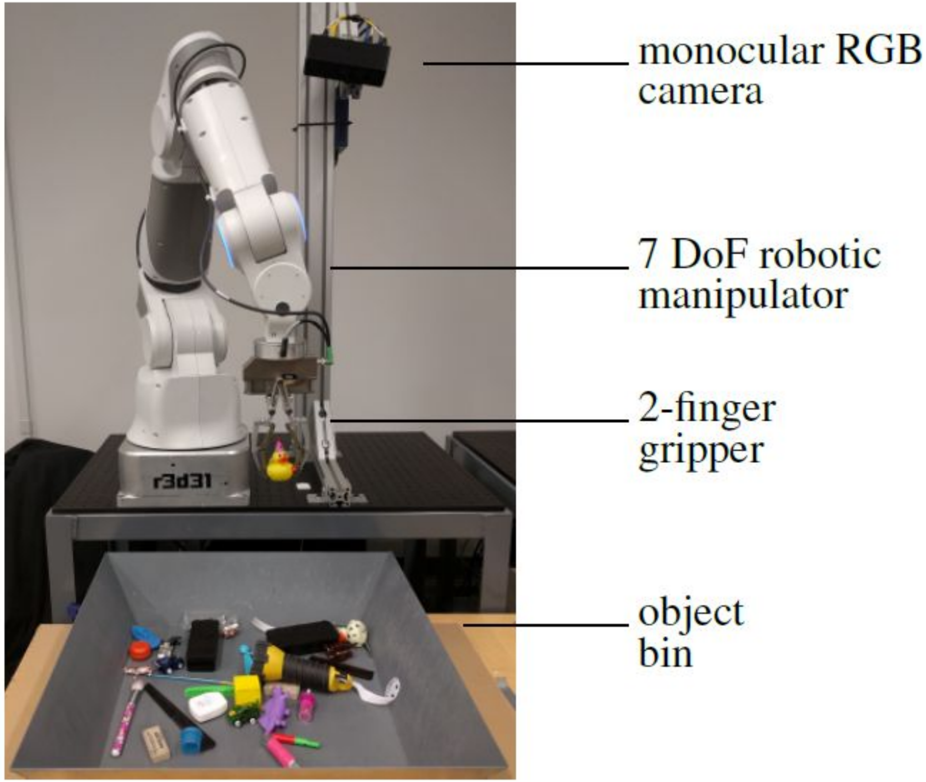

“Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection”
Sergey Levine et al., UC Berkeley and Google Brain

Arm Farm Approach

6-14 robotic arms

800,000 attempts over 2 months

80% success

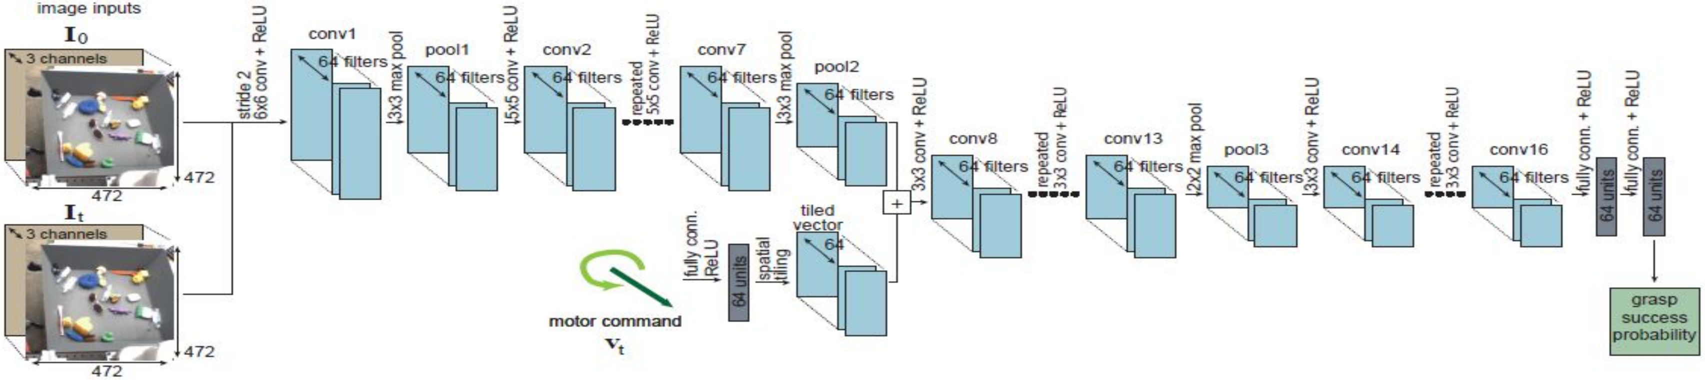

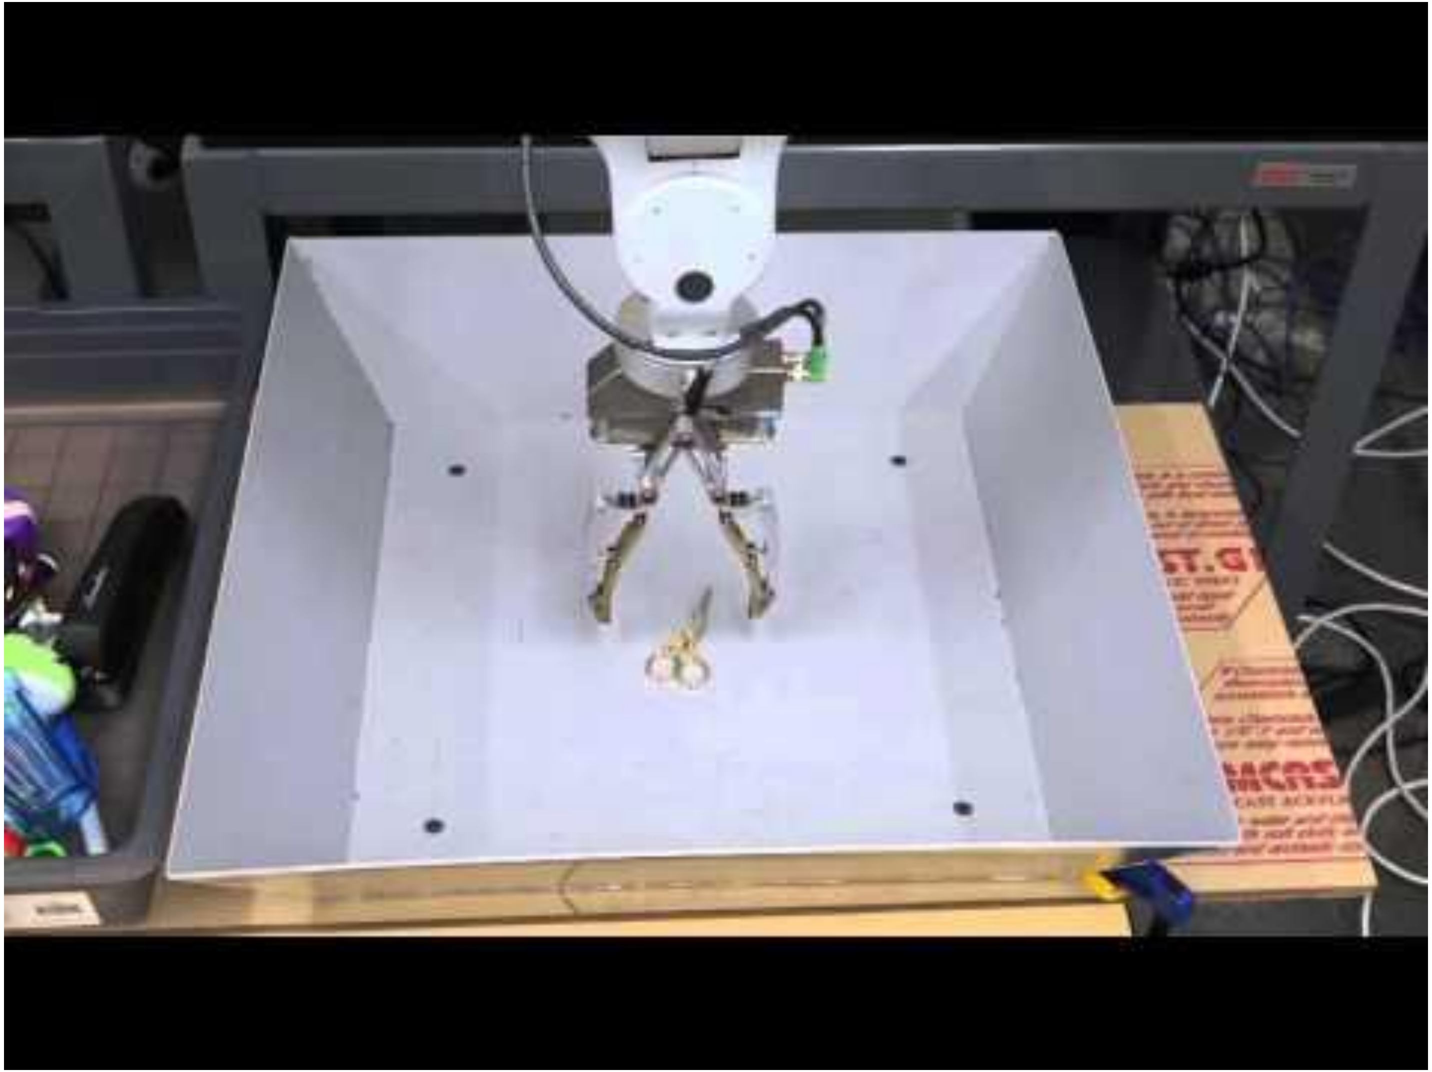

Challenges for Learning and Robotics

Experience/exploration is expensive and hard to reproduce,
and simulations are rarely sufficient (*under-modeling)*.

Training data sets don't readily generalize and labeling isn't straightforward.

Actions influence the data.

Reward-shaping may be harder than engineering a solution.
"Do what I do, not what I say"

Curse of dimensionality: Continuous, high-dimensional states and actions.

Partial observability: Uncertainty in state.

Summary

Learning can complement, but doesn’t replace engineering, in robotic systems.

Low-hanging fruit: Integrating ML-based perception (Scoops).

Unique challenges of applying ML to robotics systems (Shakes).

How much of ML successes will translate to robotics?