Learning and Autonomous Control

Jens Kober

Cognitive Robotics
Delft University of Technology
The Netherlands

j.kober@tudelft.nl
www.jenskober.de

Cognitive Robotics Department

- **Robot Dynamics** (dynamic motion control, motor control)
- **Intelligent Vehicles** (perception and modeling, dynamics, human factors)
- **Human-Robot Interaction** (physical human-robot interaction)
- **Learning and Autonomous Control** (intelligent control, cognition)

Robot Dynamics

Pick task workflow

Intelligent Vehicles

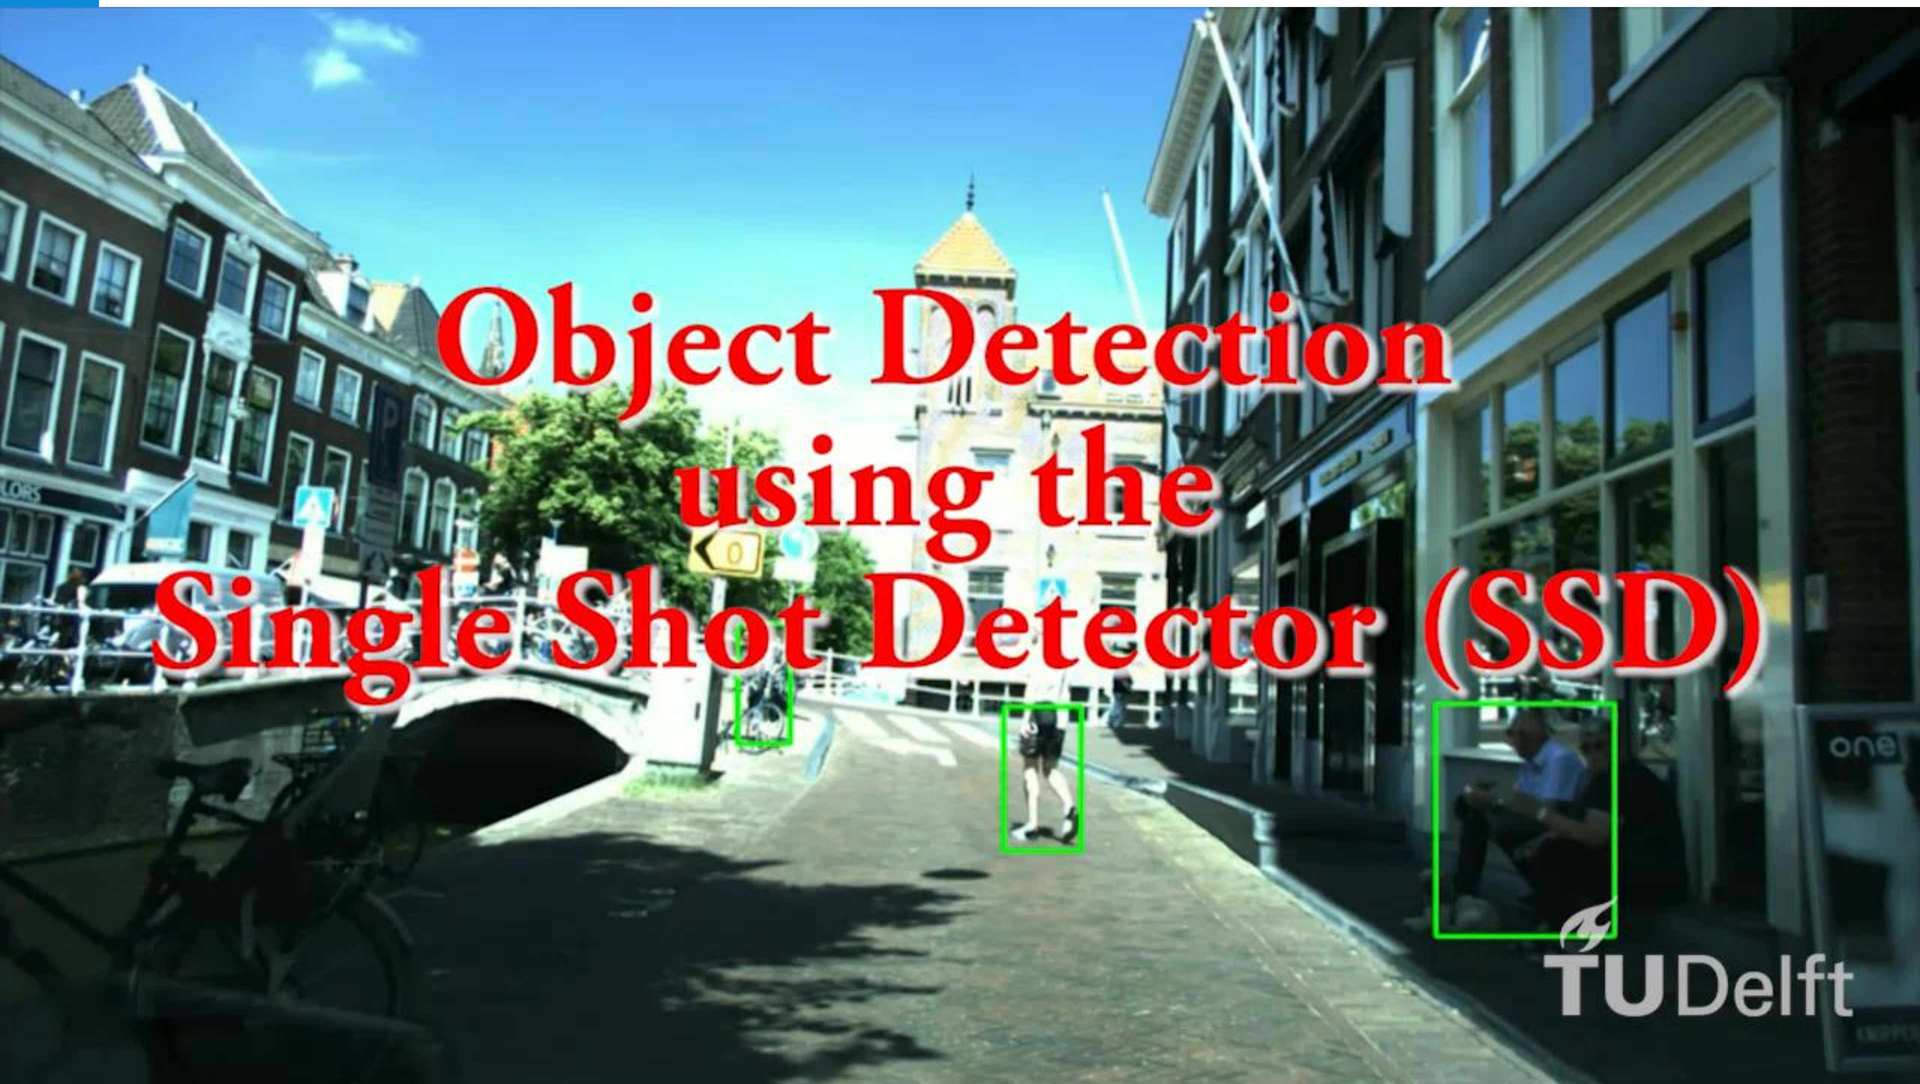

Human-Robot Interaction

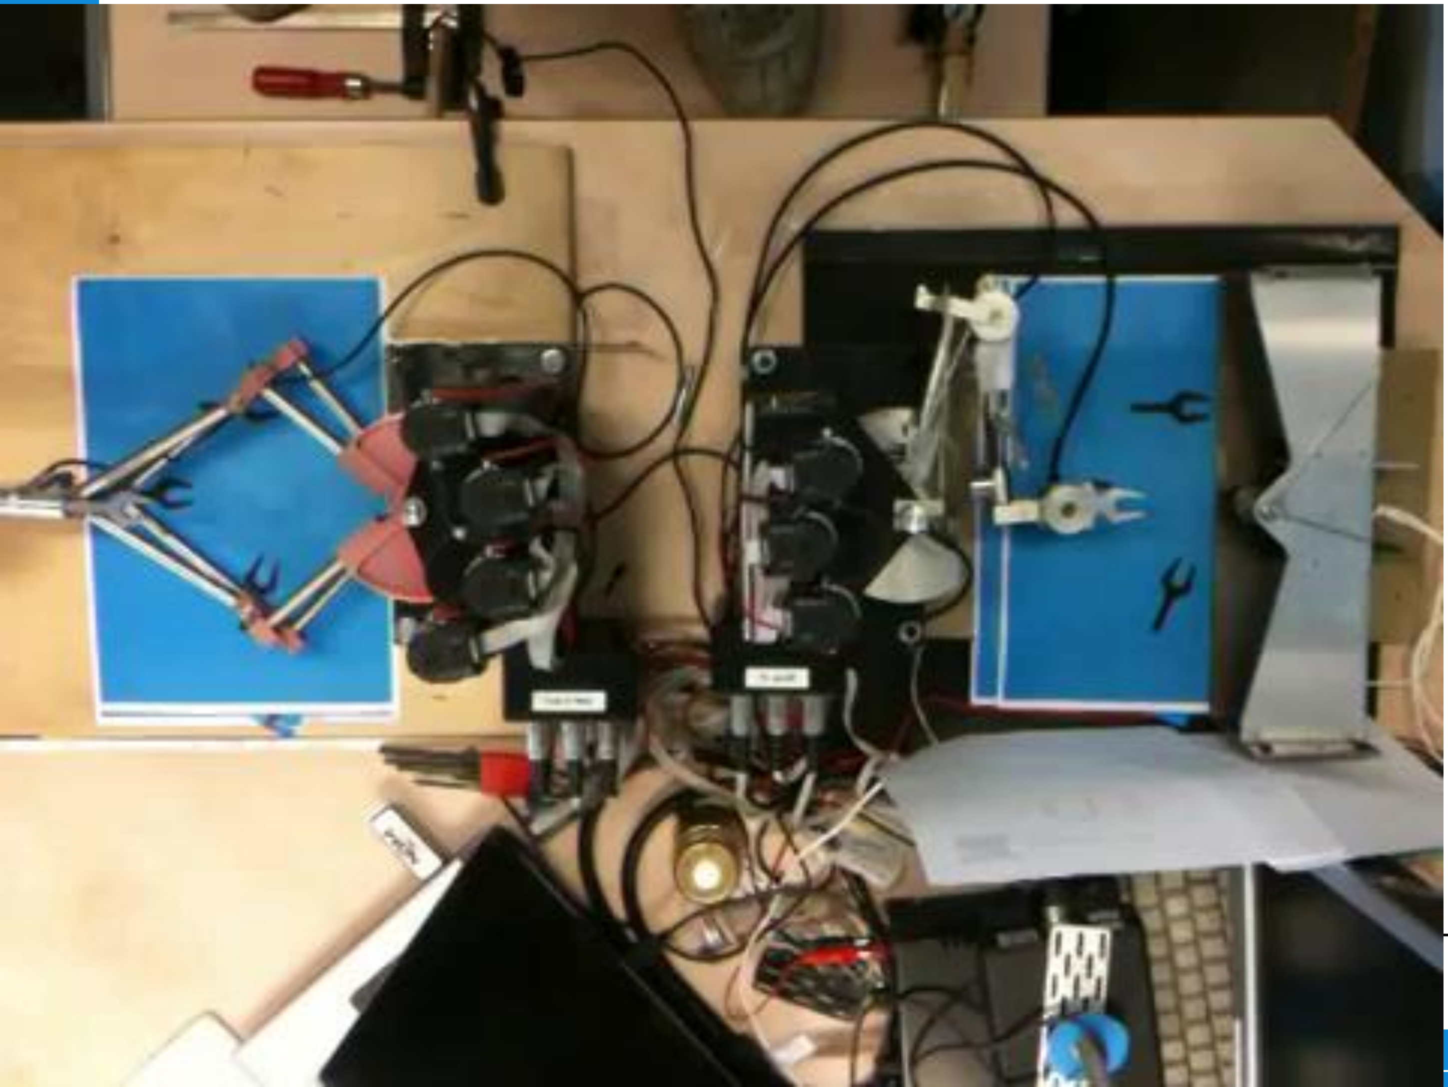

Learning and Autonomous Control

Robert Babuska
professor
learning and adaptive control

Jens Kober
assistant professor
robotics, learning motor skills

Javier Alonso Mora
assistant professor
robotics, motion planning

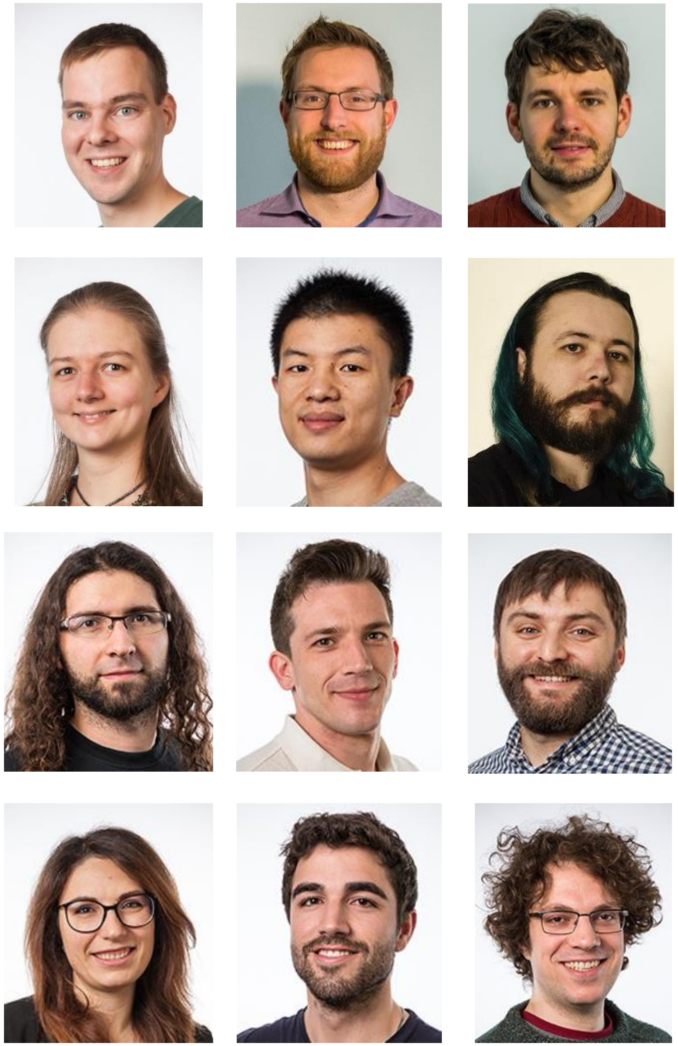

youtube/ABCUPM

Growing Interest in Robotics

A ROBOT IN EVERY HOME

The leader of the PC revolution predicts that the next hot field will be robotics

By Bill Gates

Imagine being present at the birth of a new industry. It is an industry based on groundbreaking new technologies, wherein a handful of well-established corporations sell highly specialized devices for business use and a fast-growing number of start-up companies produce innovative toys, gadgets for hobbyists and other interesting niche products. But it is also a highly fragmented industry with few common standards or platforms. Projects are complex, progress is slow, and practical applications are relatively rare. In fact, for all the excitement and promise, no one can say with any certainty when—or even if—this industry will achieve critical mass. If it does, though, it may well change the world.

Of course, the paragraph above could be a description of the computer industry during the mid-1970s, around the time that Paul Allen and I launched Microsoft. Back then, big, expensive mainframe computers ran the back-office operations for major companies, governmental departments and other institutions. Researchers at leading universities and industrial laboratories were creating the basic building blocks that would make the information age possible. Intel had just introduced the 8080 microprocessor, and Atari was selling the popular electronic game Pong. At homegrown computer clubs, enthusiasts struggled to figure out exactly what this new technology was good for.

But what I really have in mind is something much more contemporary: the emergence of the robotics industry, which is developing

AMERICAN ROBOTIC: Although a few of the domestic robots of tomorrow may resemble the anthropomorphic machines of science fiction, a greater number are likely to be mobile peripheral devices that perform specific household tasks.

AMERICAN GOTHIC, 1930, BY GRANT WOOD; ALL RIGHTS RESERVED BY THE ESTATE OF NAN WOOD GRAHAM; LICENSED BY VAGA, NEW YORK, N.Y., AND SUPERSTOCK, INC.; MODIFIED BY KENN BROWN

COPYRIGHT 2006 SCIENTIFIC AMERICAN, INC.


BEST THING
January 1, 2000
Full-scale colonization
begins


Challenges

Uncertainty due to environment

Safety

Robustness

Interaction with humans

Learning from mistakes

Compliance

Situation awareness

Autonomy

Power efficiency

Future Robots

- Need to operate under *unforeseen circumstances*
- *Interact with humans in an intelligent way*
- *Learn* from mistakes, improve over time (performance, energy)

- Impossible to preprogram all tasks or behaviors
- Economic point of view: short design time and low costs

Learning and adaptation is essential!

Main Research Topics

- Adaptive and learning control systems
  - Reinforcement learning, optimal control
  - Nonlinear adaptive control
- Deep learning for robotics
  - Data-effective learning algorithms
  - Sensor fusion
- Supervisory and distributed control design
  - Motion planning, coordination
  - Multi-agent systems

Emphasis on *novel, generic, widely applicable* techniques
Realistic / *real-world case studies* and applications

Autonomous Multi-Robots Lab

Autonomous navigation
Planning in dynamic environments

&

Multi-robot systems

AI + Control
Constrained optimization tools, combined with machine learning, formal methods and consensus

Diverse application fields
Transportation, Manipulation, Inspection, Ground & aerial vehicles

Cooperative mobile manipulation

Robot dynamics →
Environment obstacles →
Formation definition →

Constrained optimization
- **Multiple robots** — Optimize formation parameters within convex region in position-time free space
- *or*
- **Single robot** — Non-linear model predictive control

→ **Safe motion in dynamic environments**

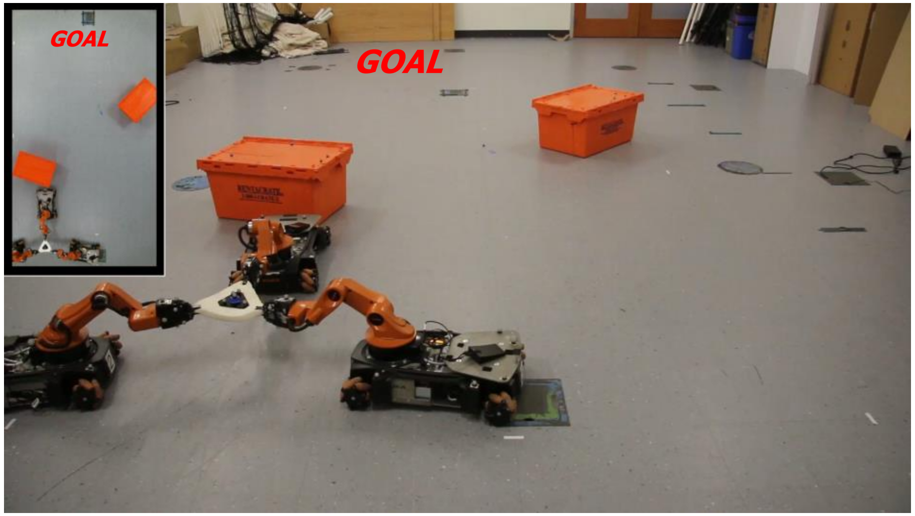

J.Alonso-Mora, et al. "Multi-robot formation control and object transport in dynamic environments via constrained optimization", IJRR, 2017

Online planning for aerial vehicles

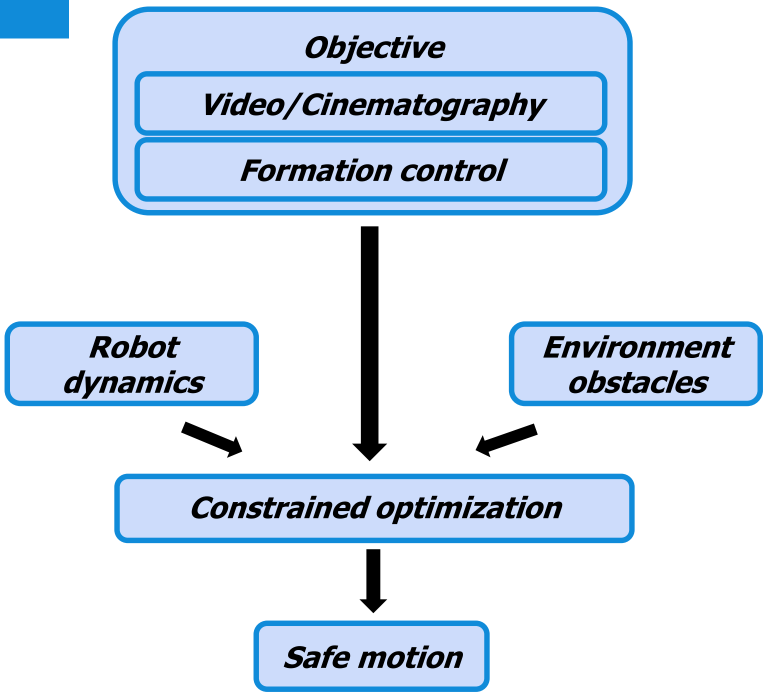

T. Naegeli, et al., "Trajectory Optimization for Multi-view Aerial Cinematography", ACM Siggraph, 2017; Videography, RA-L 2017
J. Alonso-Mora et al., "Distributed formation control in dynamic environments", ICRA 2017

Motion planning with interaction

Information from other participants

Communication
Distributed asynchronous

or

Estimation
Probabilistic models & chance constraints

Environment models

Vehicle model

Constrained optimization

Safe motion among robots & humans

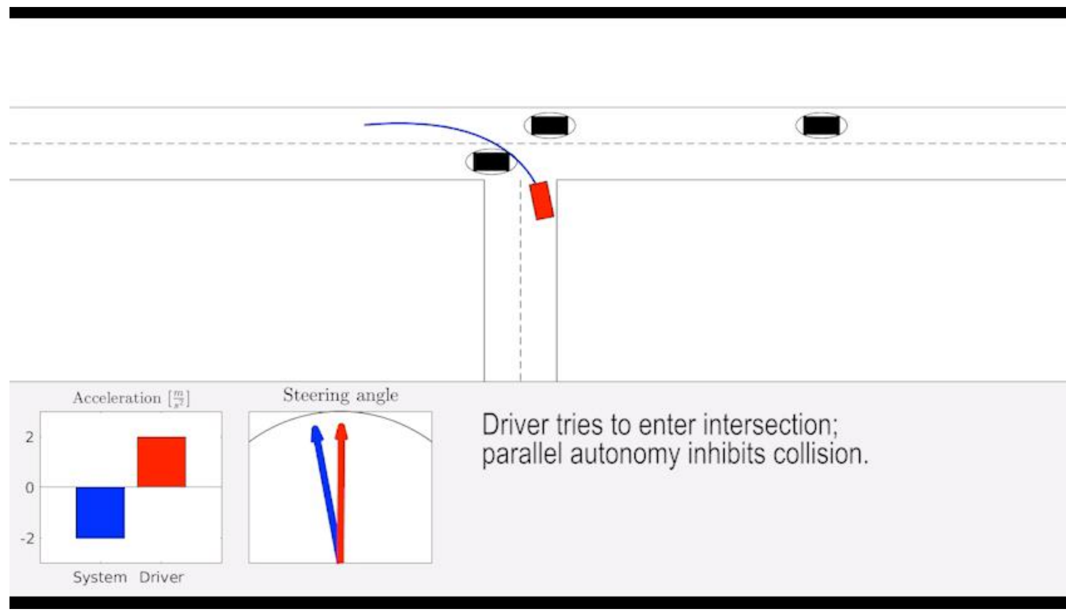

W. Schwarting et al., "Safe Nonlinear Trajectory Generation for Parallel Autonomy With a Dynamic Vehicle Model", T-ITS 2017
L. Ferranti et al., "Coordination of Multiple Vessels Via Distributed Nonlinear Model Predictive Control", ECC 2018
B. Zhou et al., " Joint Multi-Policy Behavior Estimation and Receding-Horizon Trajectory Planning for Automated Urban Driving", ICRA 2018

Intelligent Control

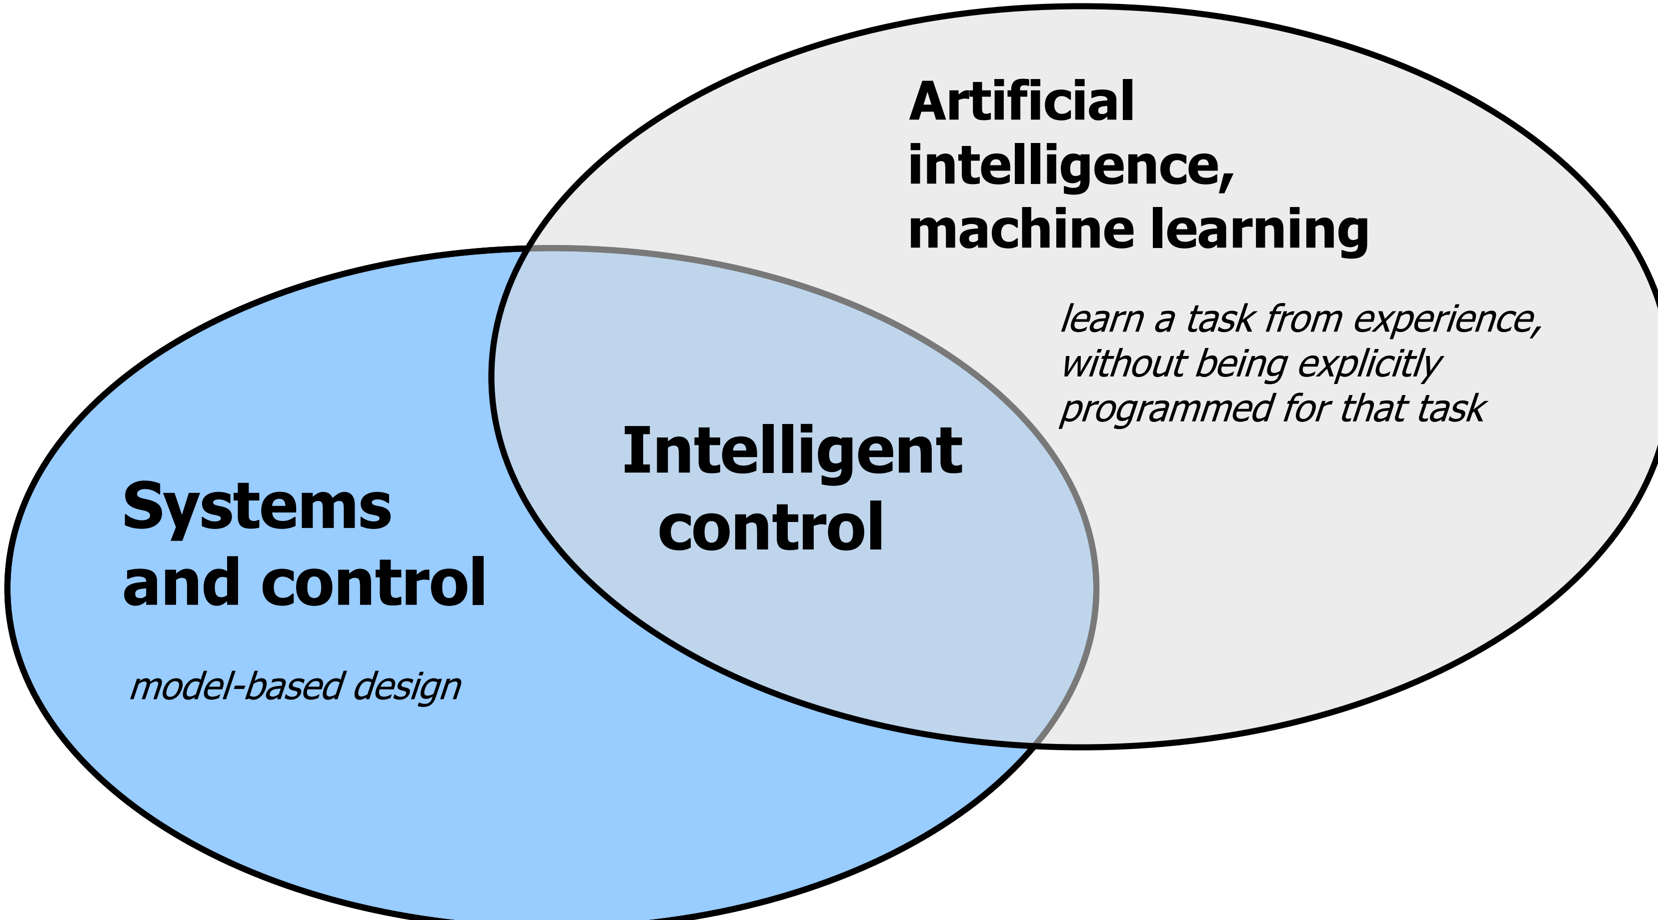

Potential of Machine Learning

- adapt to changes in environment
- find new, better solutions
- teach by demonstration or imitation

in addition:

- modeling is tedious, time consuming, expensive

Initial Optimism

- Herbert A. Simon
  Artificial Intelligence Pioneer:
  "Machines will be capable in 20 years of doing any work a man can do" (1965)

nobelprize.org

- MIT: Cog Project (1993-2003)
  Learning like human children - through constant interaction with humans and the environment

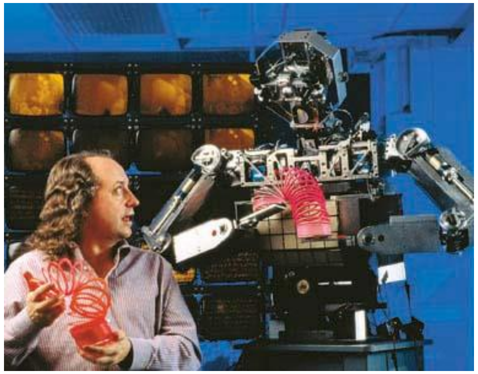

MIT Press

Spectrum of Learning Techniques

Reinforcement learning

(learn directly to control)

Supervised learning

(with 'teacher')

Unsupervised learning

(without 'teacher')

Construct a Robot Model from Data

actuate motors and observe the response

Two ways to model the system:

1. **Physical modeling**
2. **System identification**

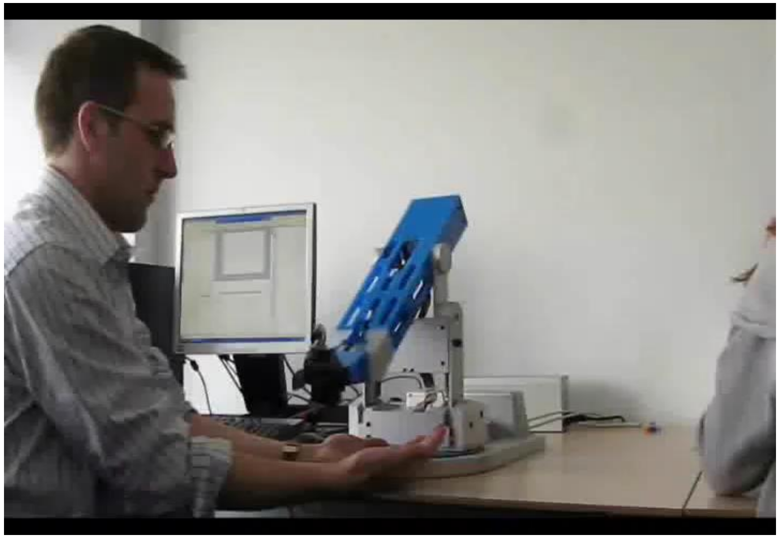

Maarten Vaandrager

Imitation Learning

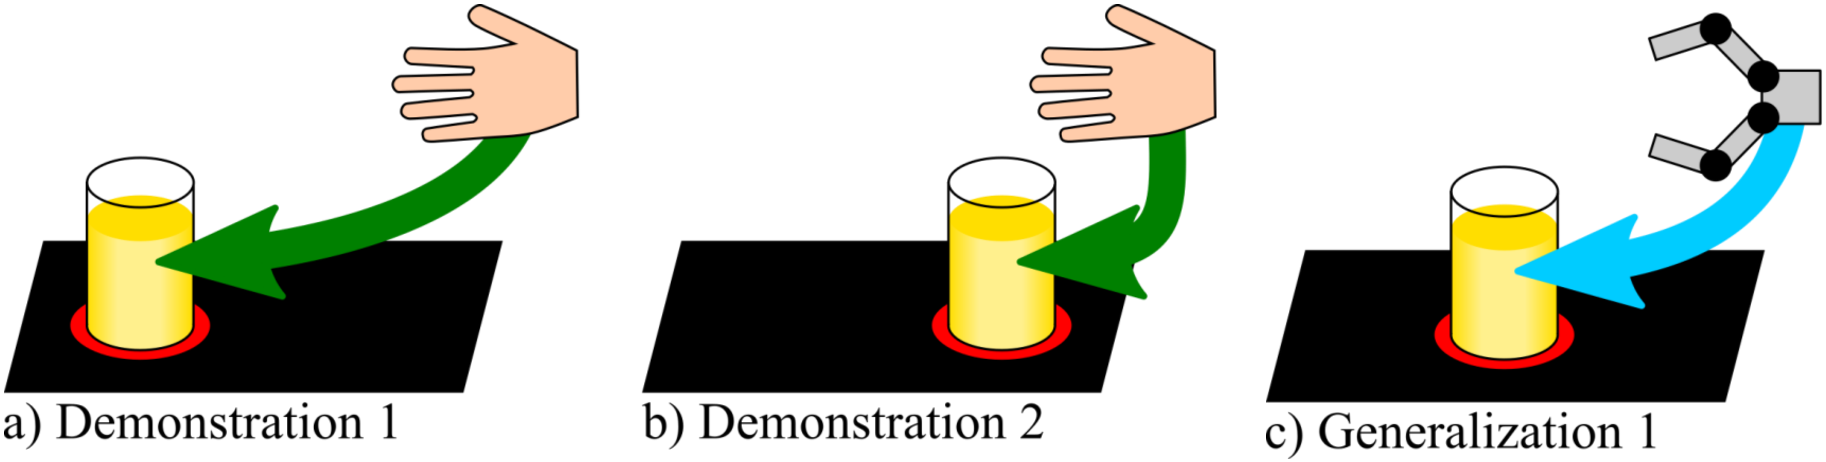

Manschitz, Kober, Gienger, & Peters, RAS 2015

Spectrum of Learning Techniques

Reinforcement learning

(learn directly to control)

Supervised learning

(with 'teacher')

Unsupervised learning

(without 'teacher')

Unsupervised Learning - Clustering

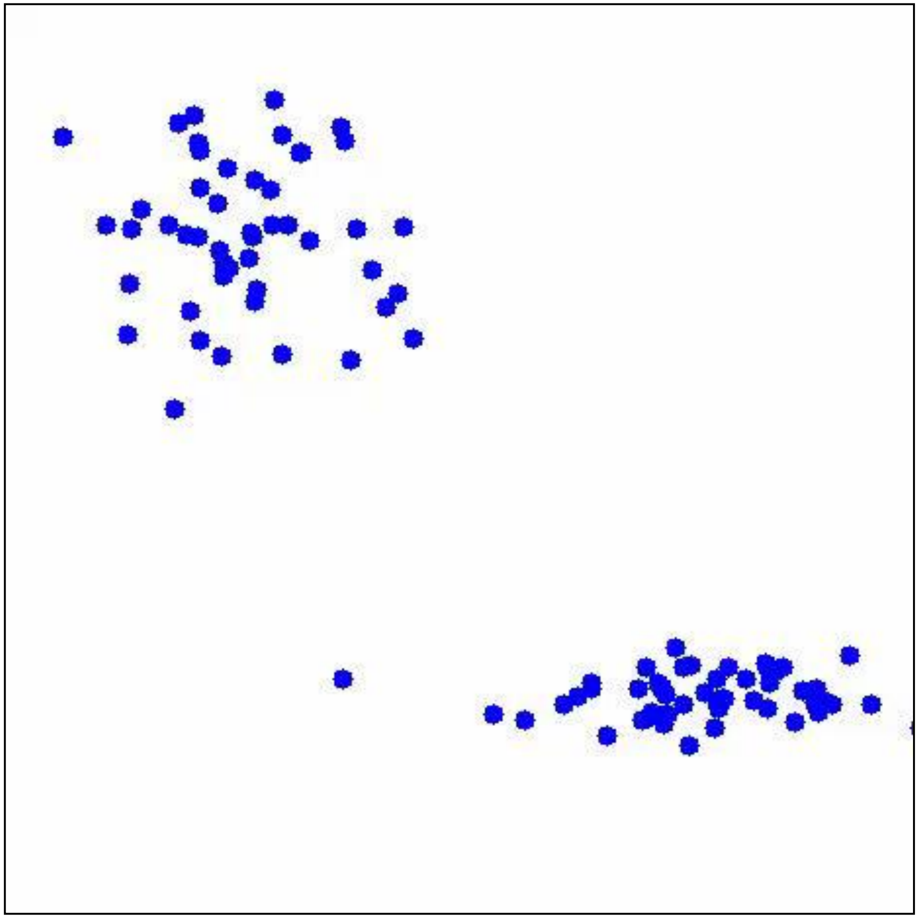

- *Discover automatically groups and structures in data*

Applications:

- Robot perception, vision
- Data-driven construction of dynamic models

Construction of Nonlinear Models

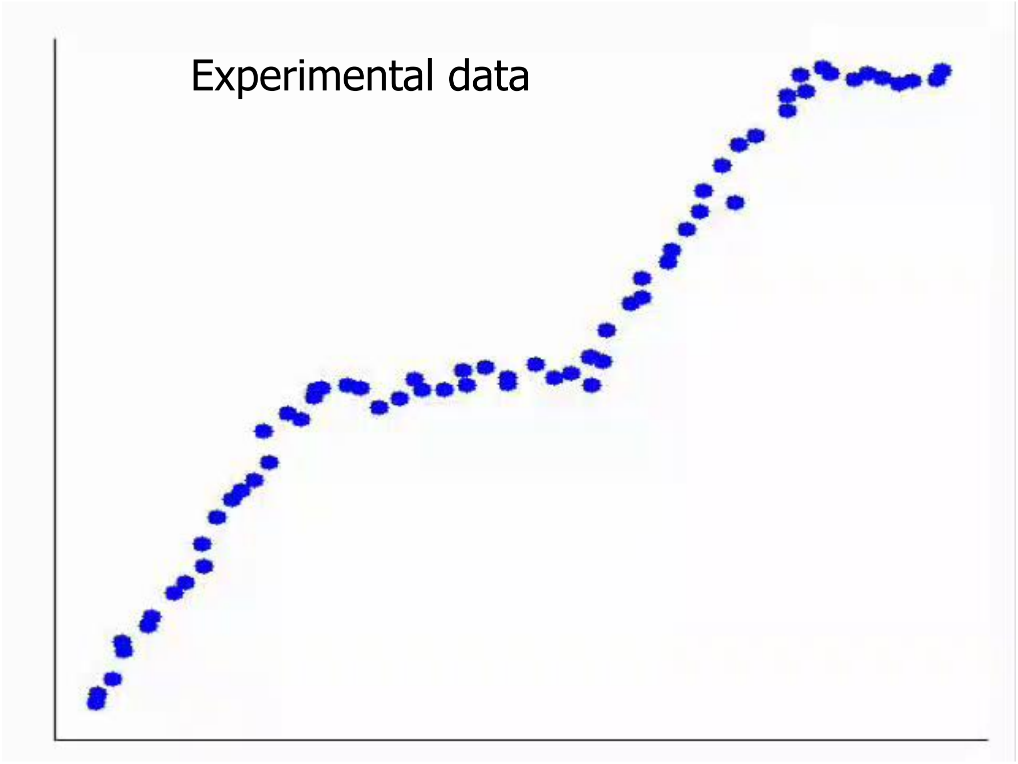

Nonlinear and uncertain behavior

Spectrum of Learning Techniques

Reinforcement learning

(learn directly to control)

Supervised learning

(with ‘teacher’)

Unsupervised learning

(without ‘teacher’)

Reinforcement Learning

Inspiration - animal learning (reward desired behavior)

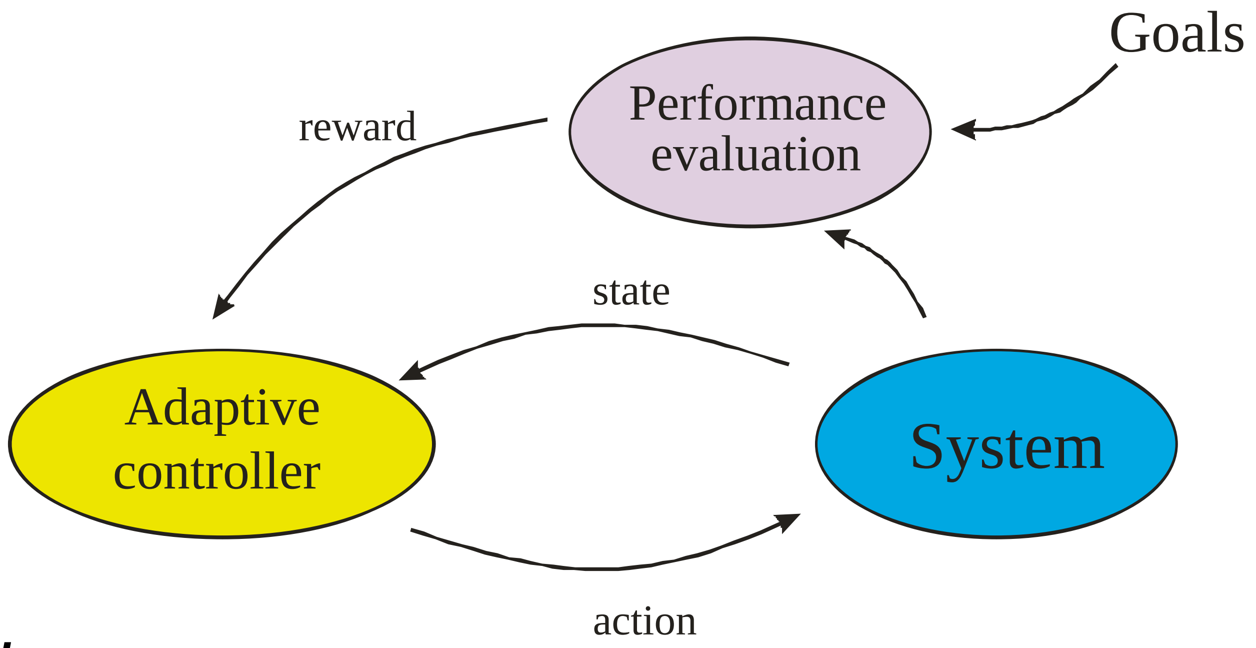

Goal:

Adapt the control strategy so that the sum of rewards over time is maximal.

Reinforcement Learning for Control

- Approximation in high-dimensional continuous spaces
- Computationally effective methods
- Constrained learning and convergence

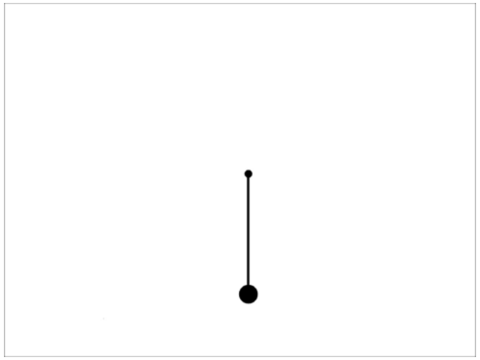

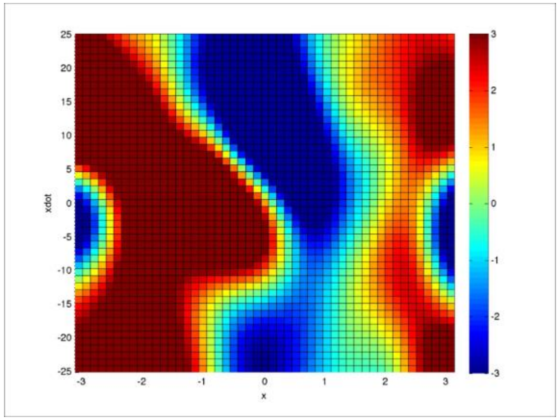

Reward-Weighted Imitation

- Maximize reward = optimize strategy
- One possibility: reward-weighted imitation

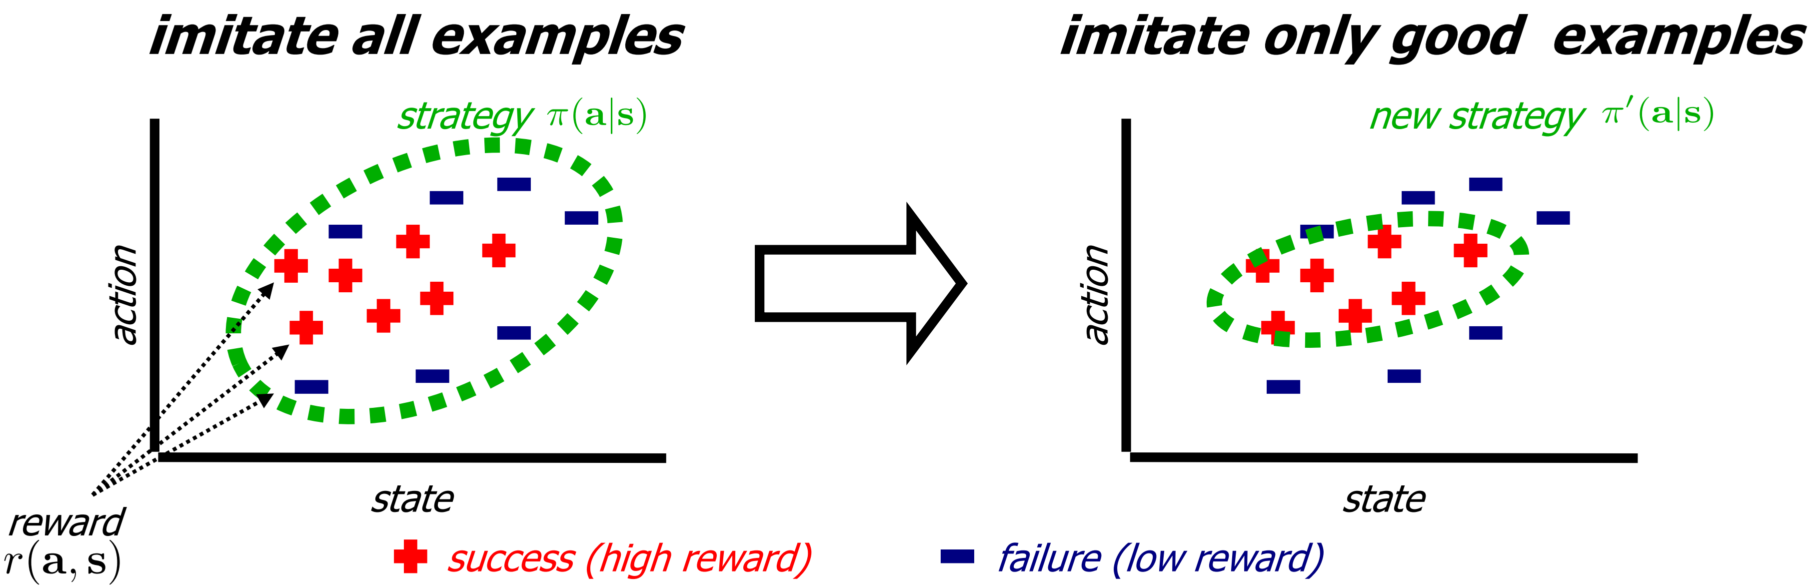

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + E\left\{\sum_{t=1}^{T} Q_t^{\mathbf{sa}} \mathbf{W}_t^{\mathbf{s}}\right\}^{-1} E\left\{\sum_{t=1}^{T} Q_t^{\mathbf{sa}} \mathbf{W}_t^{\mathbf{s}} \boldsymbol{\varepsilon}_t\right\}$$

Kober & Peters, NIPS 2008

Reward-weighted Imitation

Trial 1
Reward 0,1

Trial 2
Reward 0,3

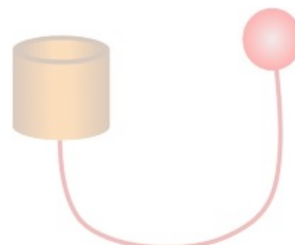

Trial 3
Reward 0,8

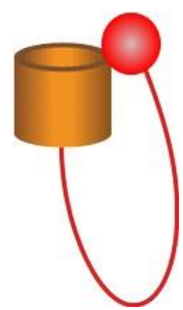

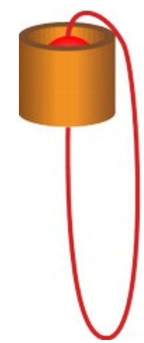

Trial 4
Reward 1,0

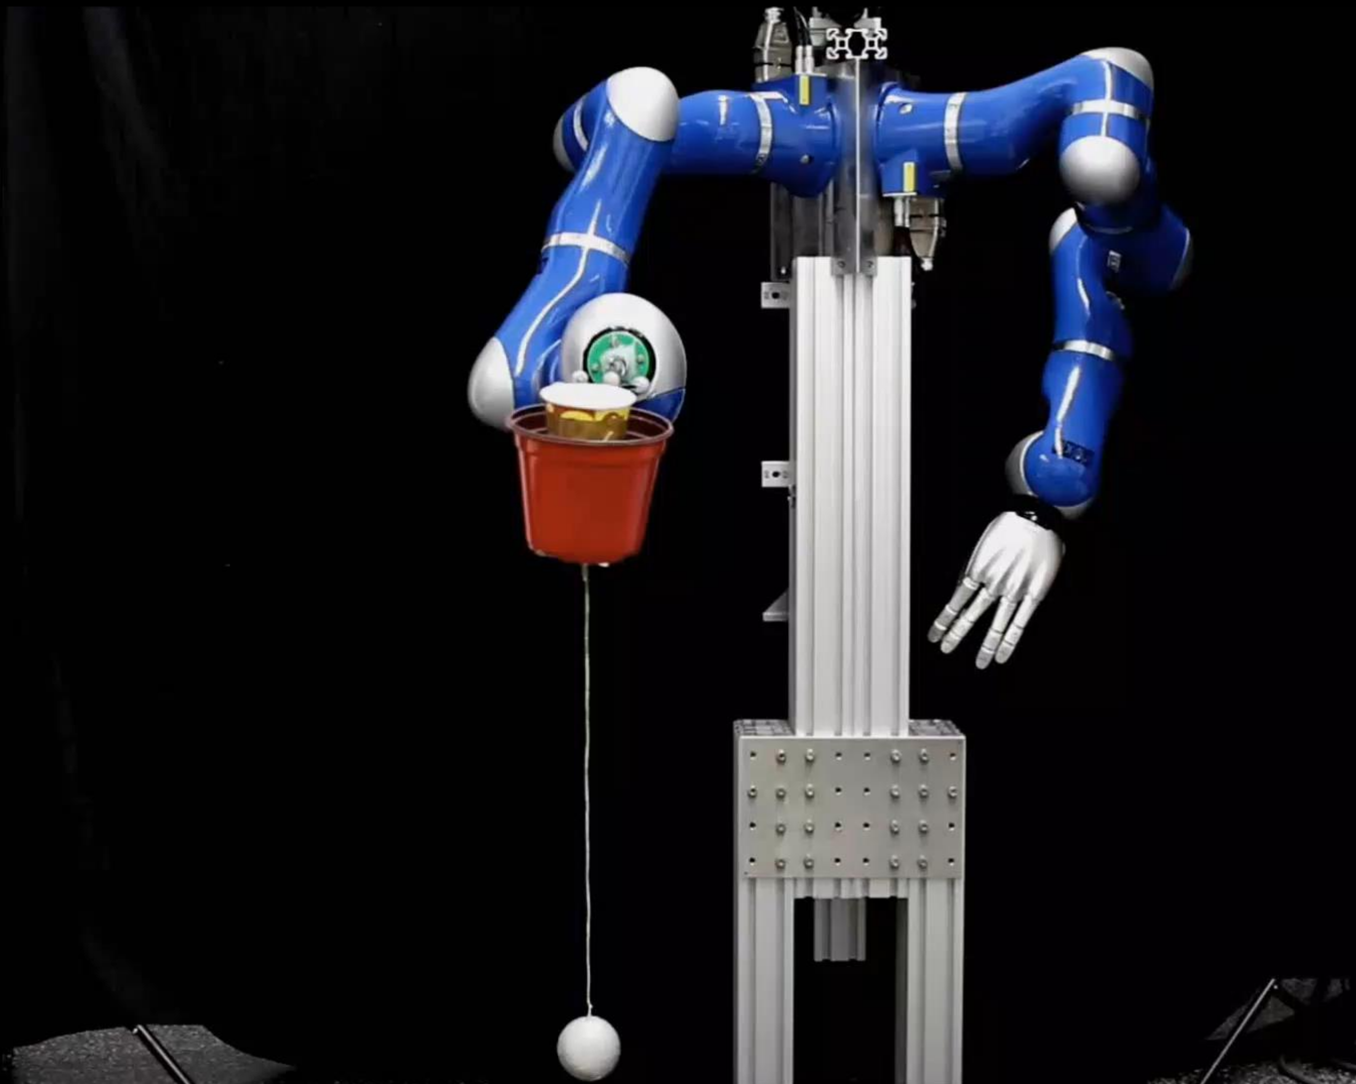

Celemin, Maeda, Ruiz-del Solar, Peters, & Kober, under review

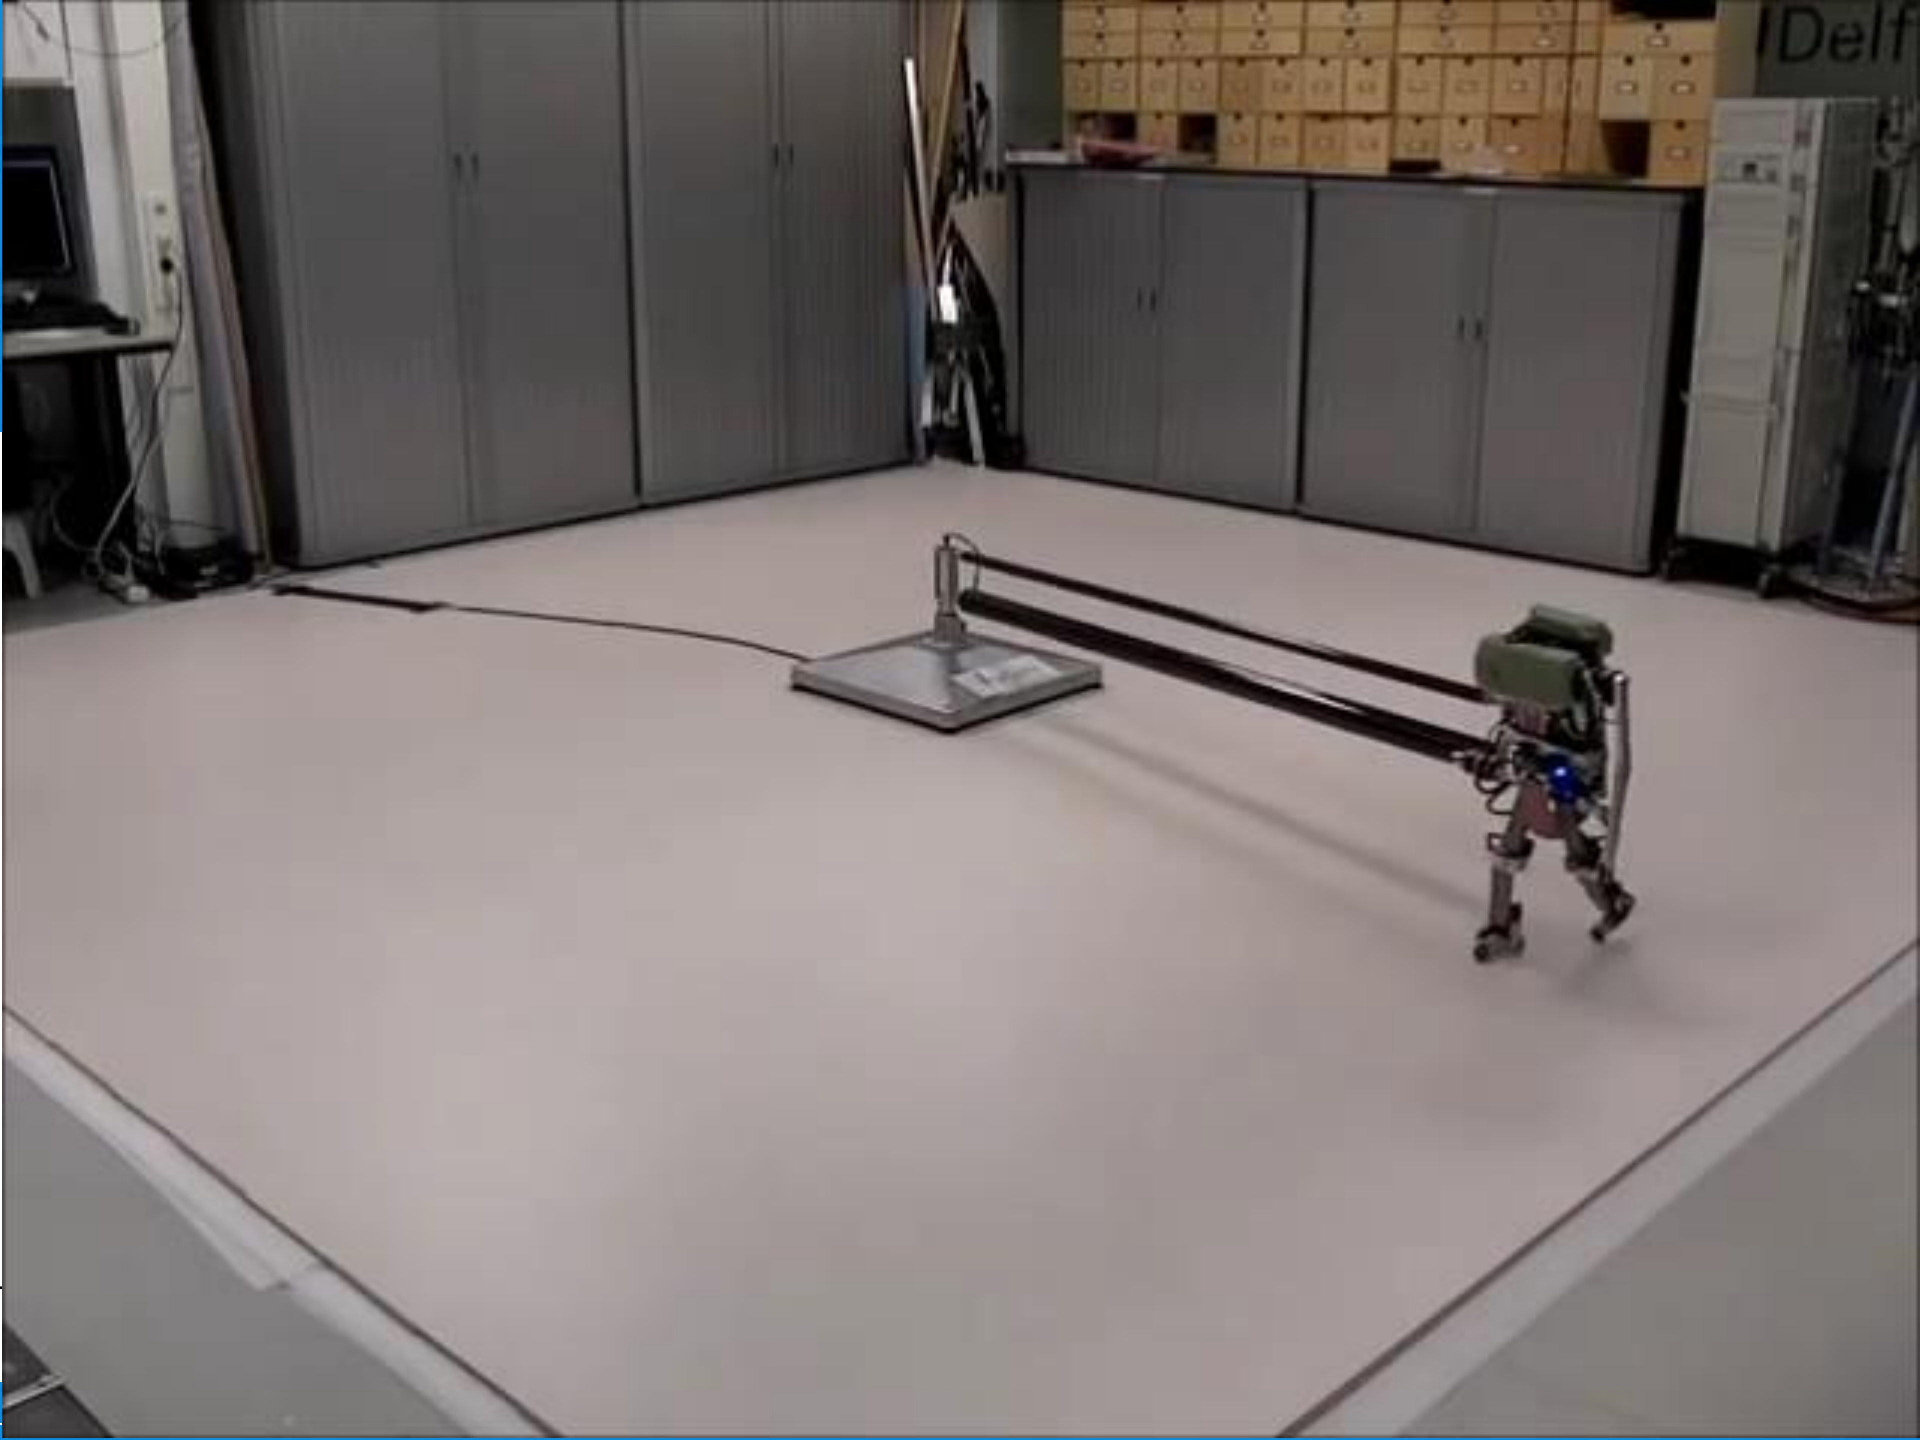

Learning a Difference Model

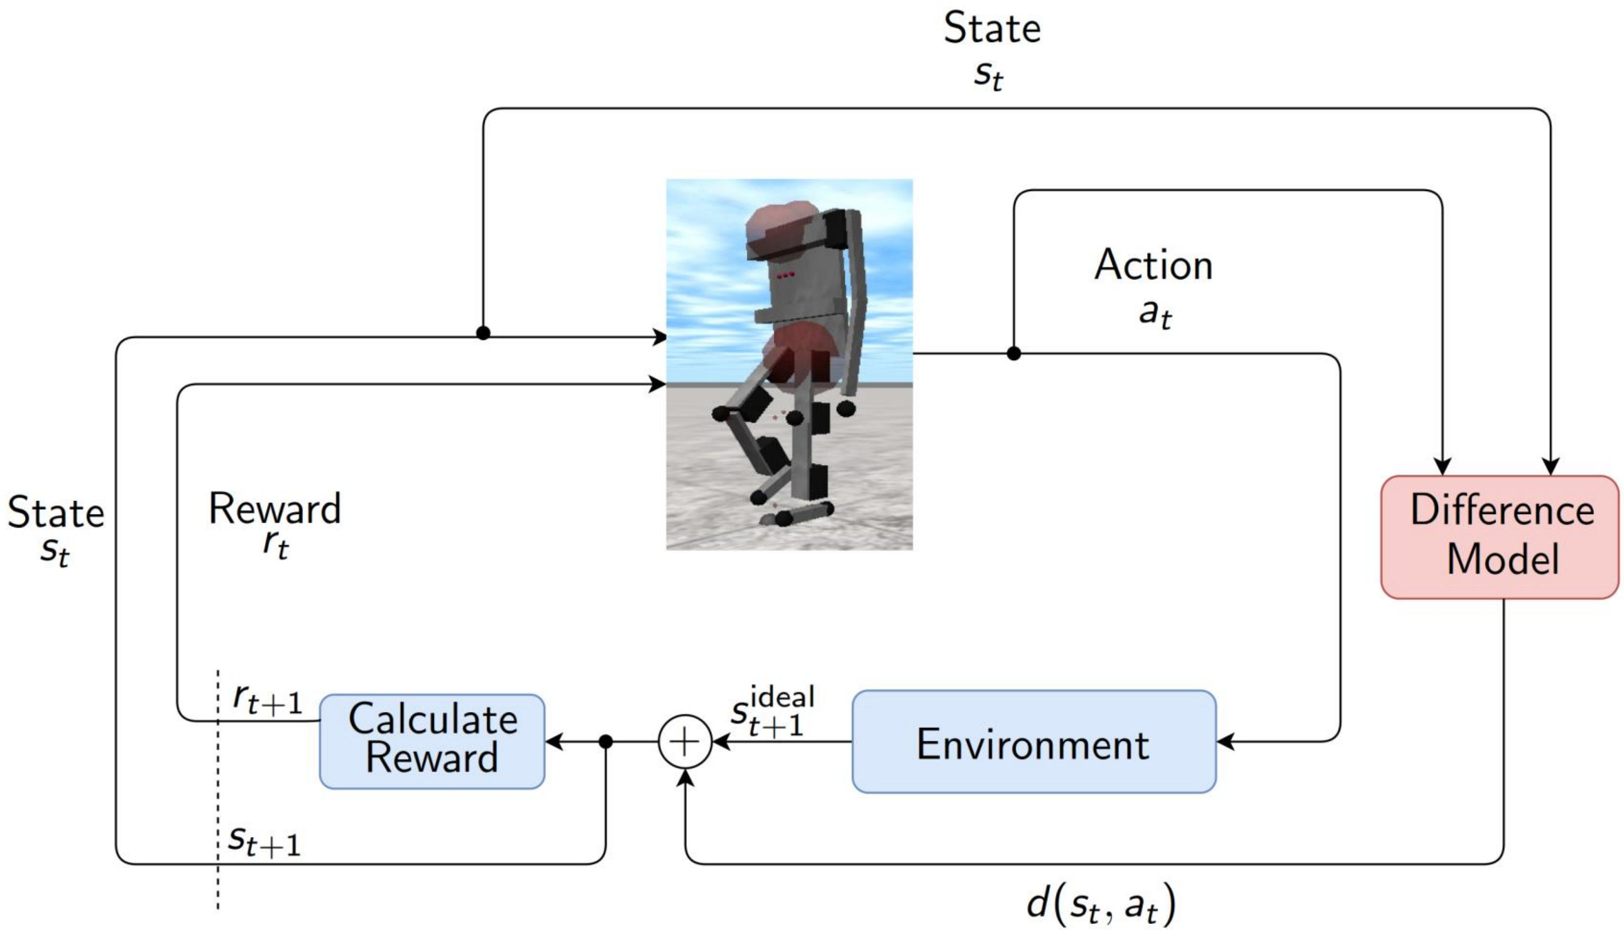

Rastogi, MSc 2017

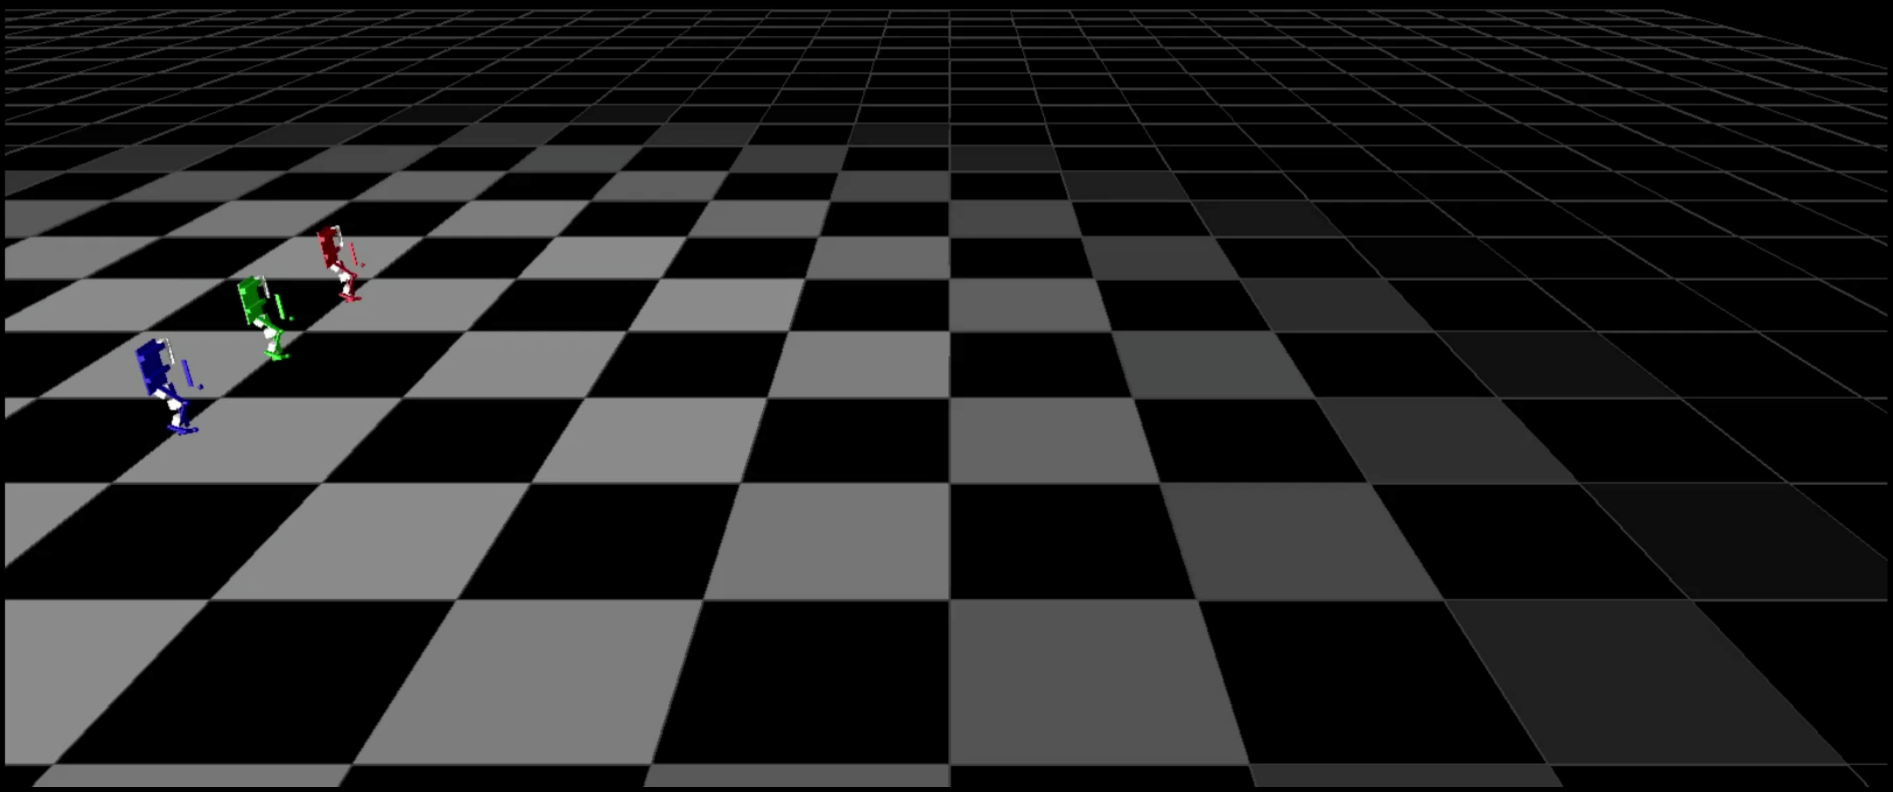

Rastogi, MSc 2017

RL-based Compensation

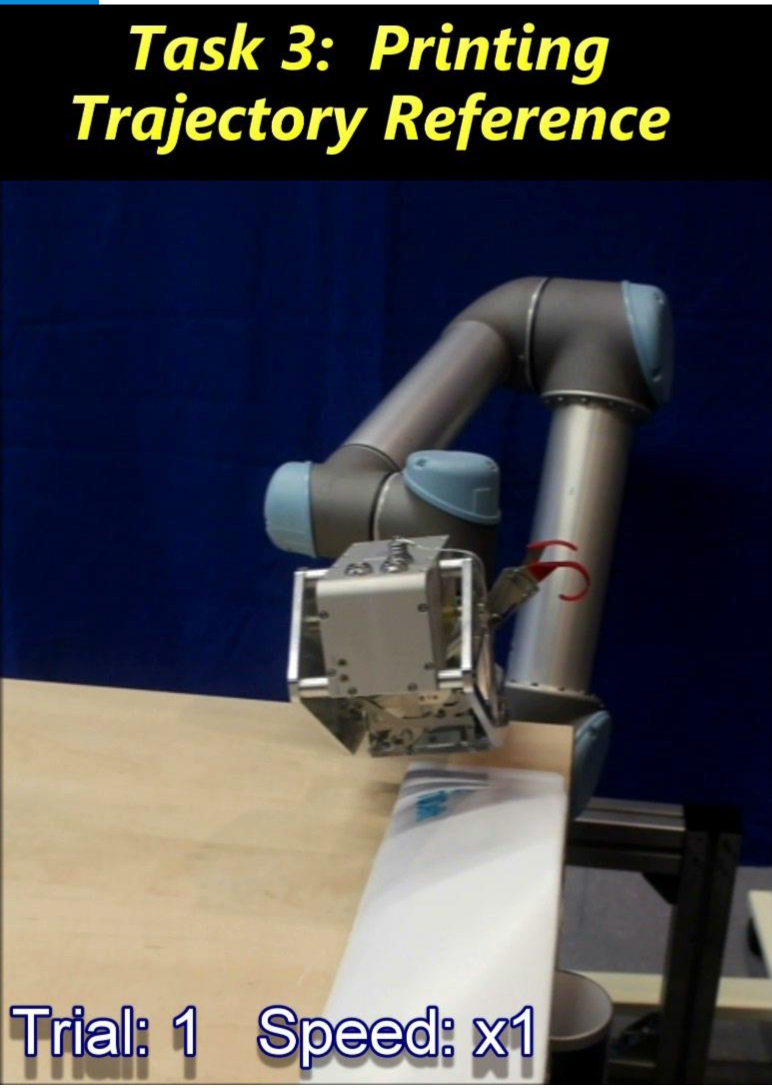

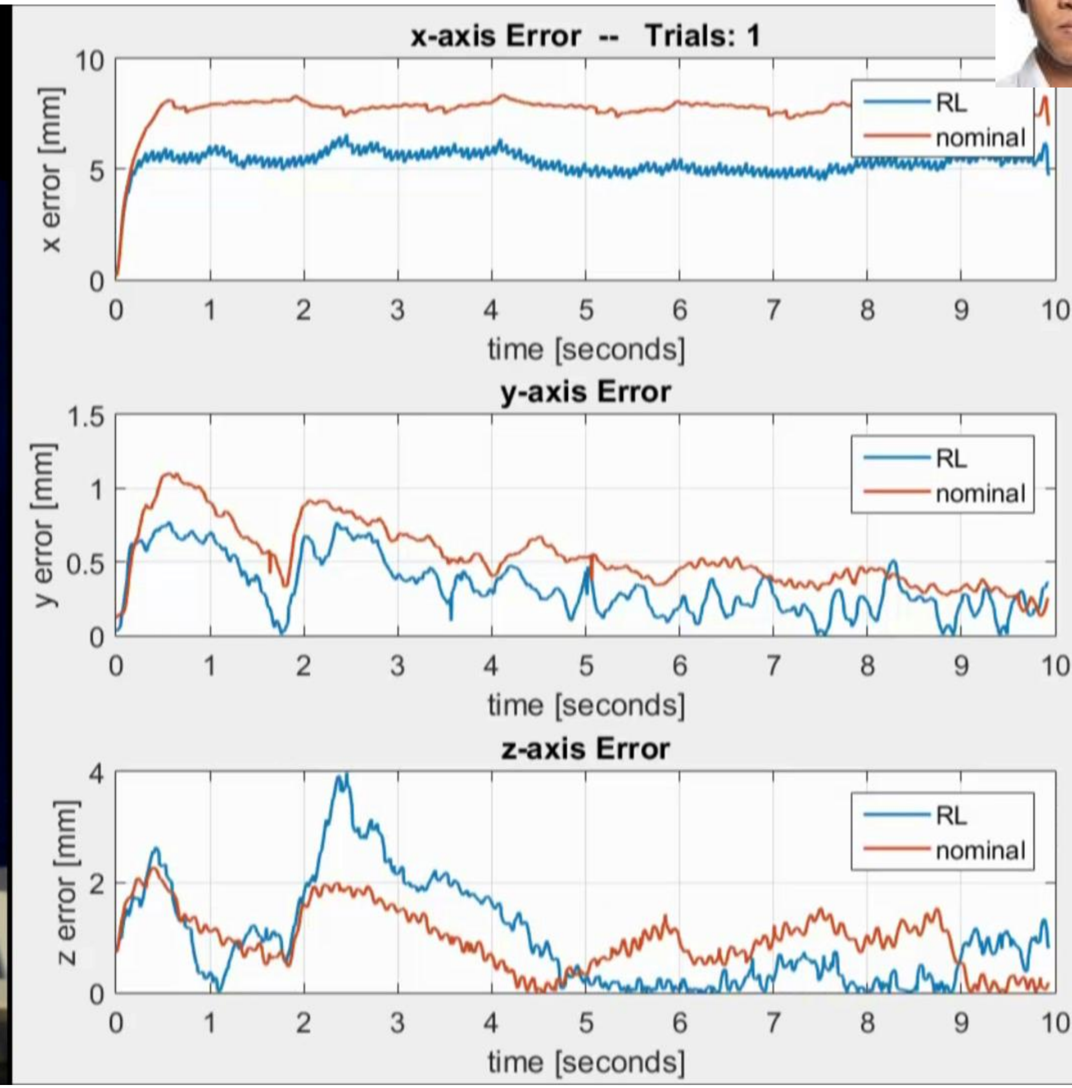

Pane, MSc 2015

*or "brilliant idiots"

Learning Abilities

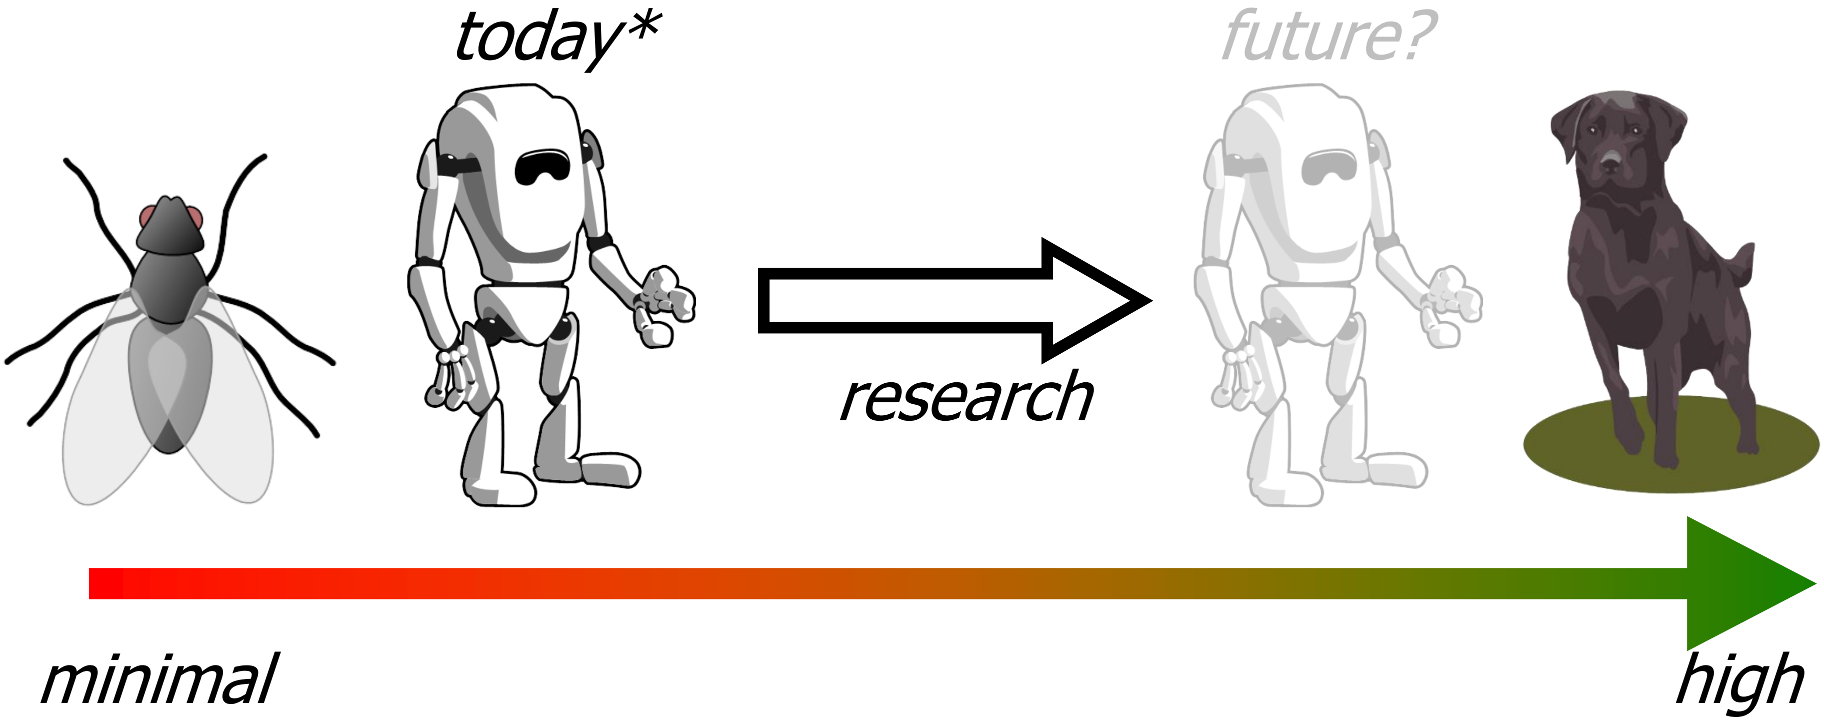

openclipart.org

Courses

- (Partially) taught by LAC
  - SC42035 Integration Project S&C
  - SC42090 Robot Motion Planning and Control
  - SC42050 Knowledge Based Control Systems
  - IN4320 Machine Learning
  - ME41025 Robotics Practicals
- Other recommended courses
  - IN4085 Pattern Recognition
  - CS4180 Deep Learning
  - IN4010 Artificial Intelligence Techniques
  - SC42100 Networked and Distributed Control Systems
  - ME41105 Intelligent Vehicles

Questions?

latd.com