Part 5

Two main meanings for SSL

- Systems (typically robots) that collect their own training data but then solve a standard supervised learning task
- Systems that learn to extract meaningful representations from the data itself

Learning to see by listening

A robot that acquires its own training data...

… and learns useful representations through an audio pretext task

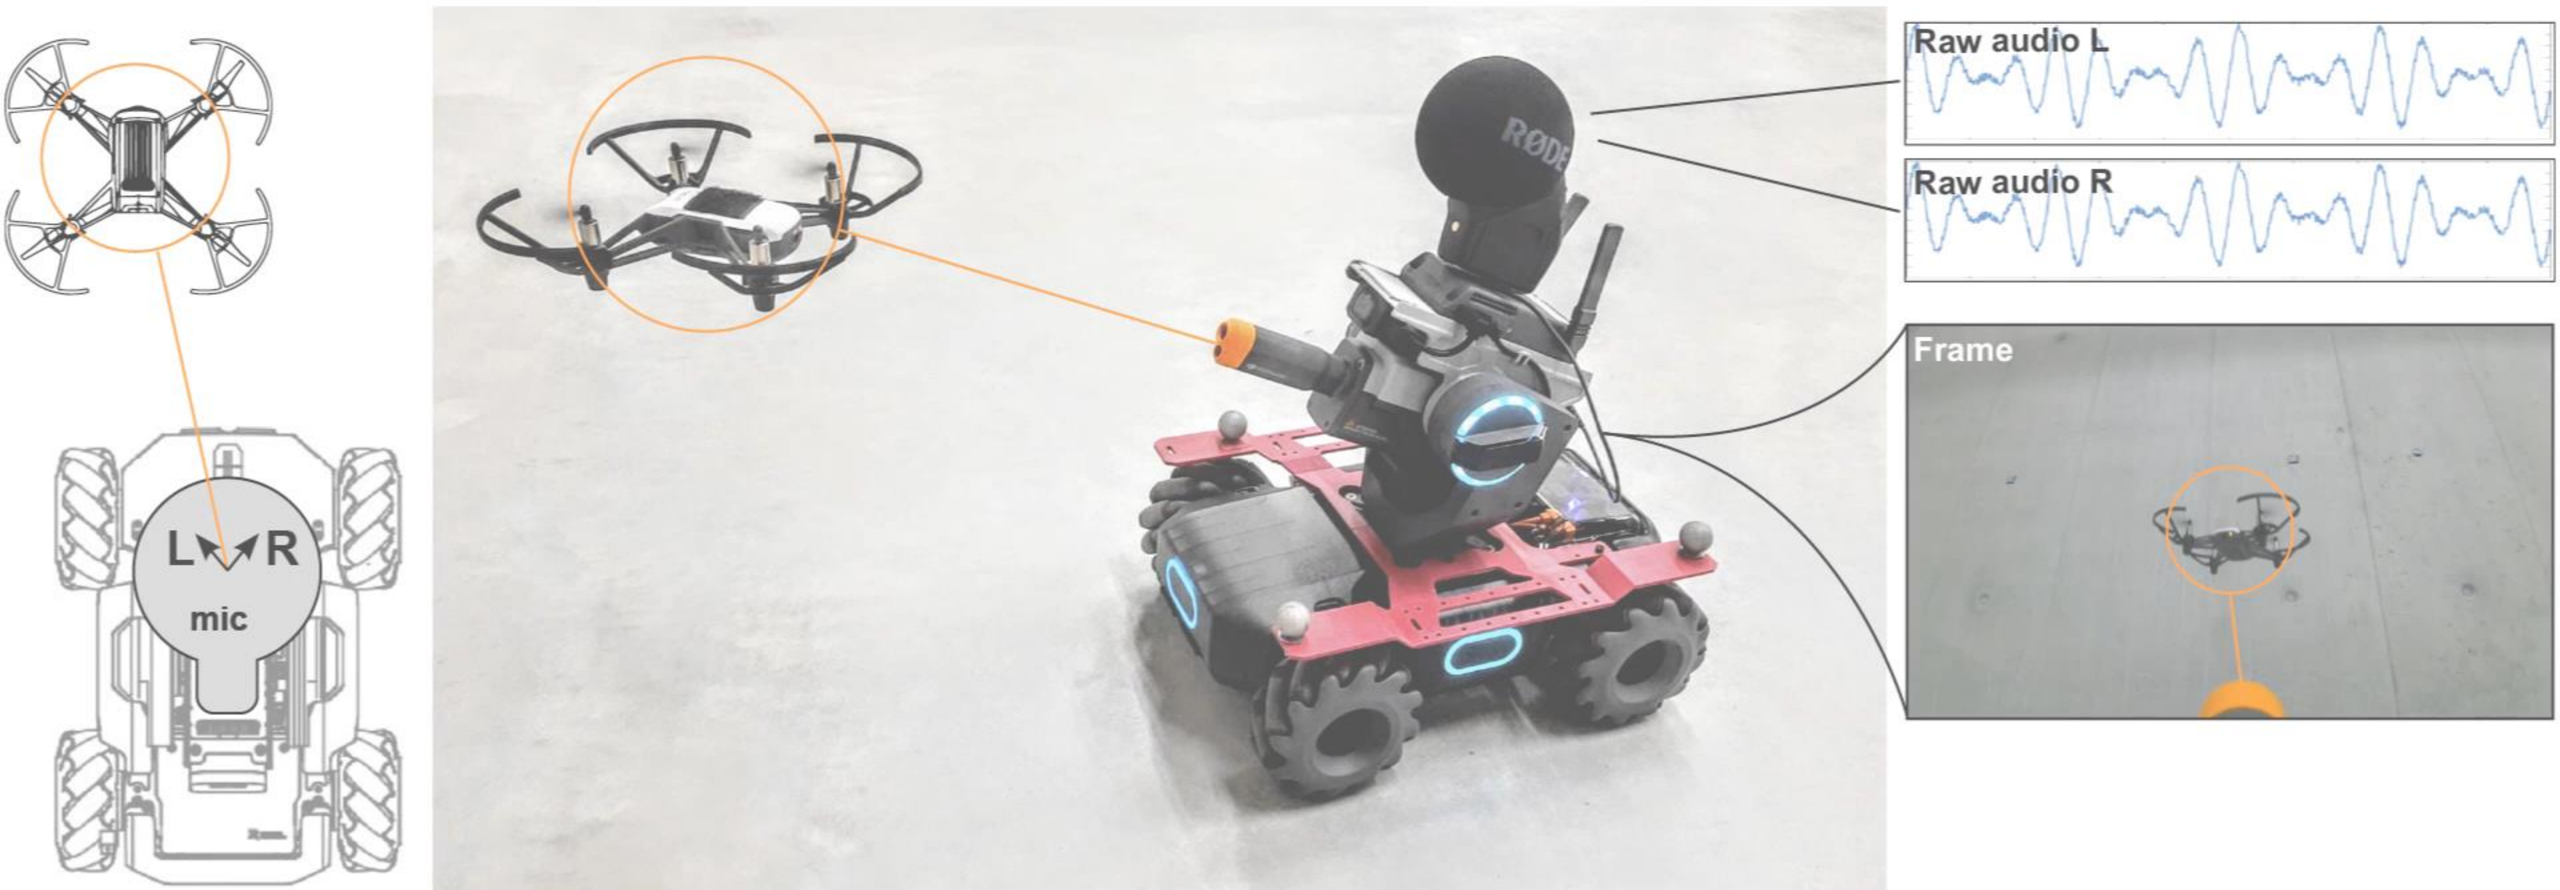
L
R
mic
RØDE
Raw audio L
Raw audio R
Frame

An experiment in self-supervised learning for Robots

A robot that acquires its own training data...
but then solves a standard Supervised Learning problem

Mighty Thymio

- 5 front-facing infra-red sensors
- 720p camera
- ODRIOD C1
- Wi-Fi connectivity

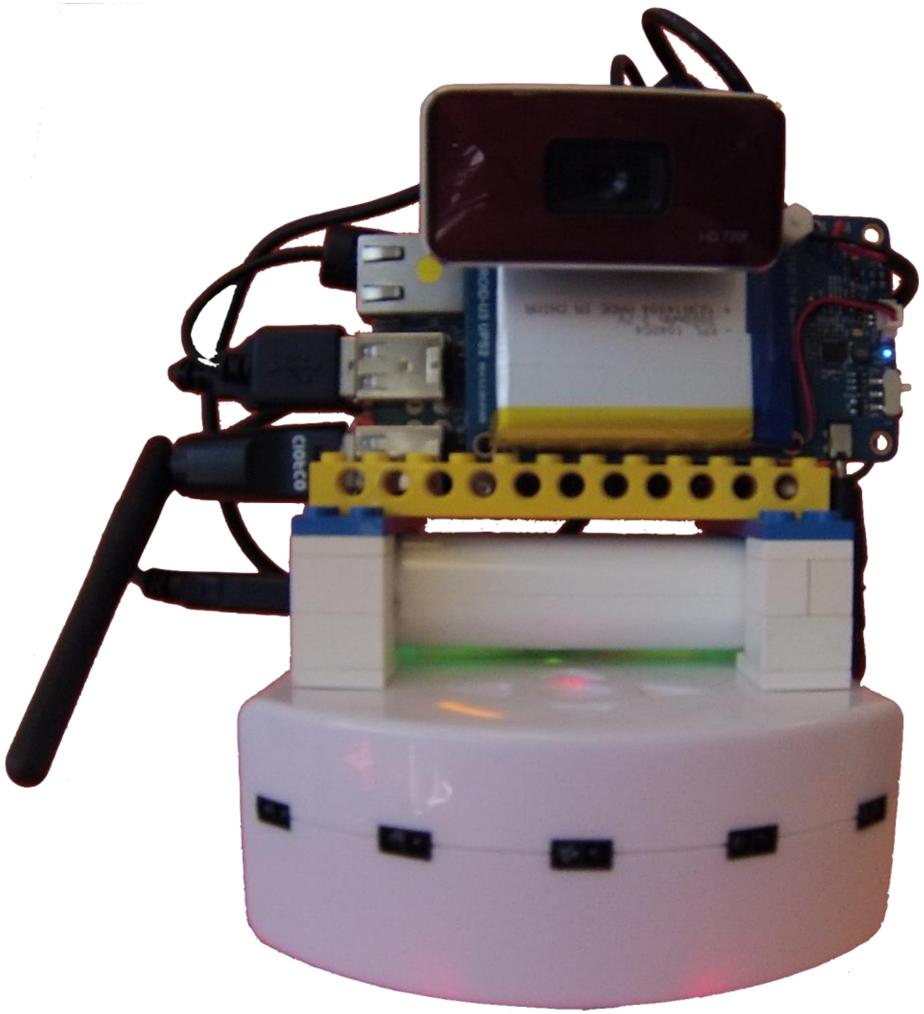

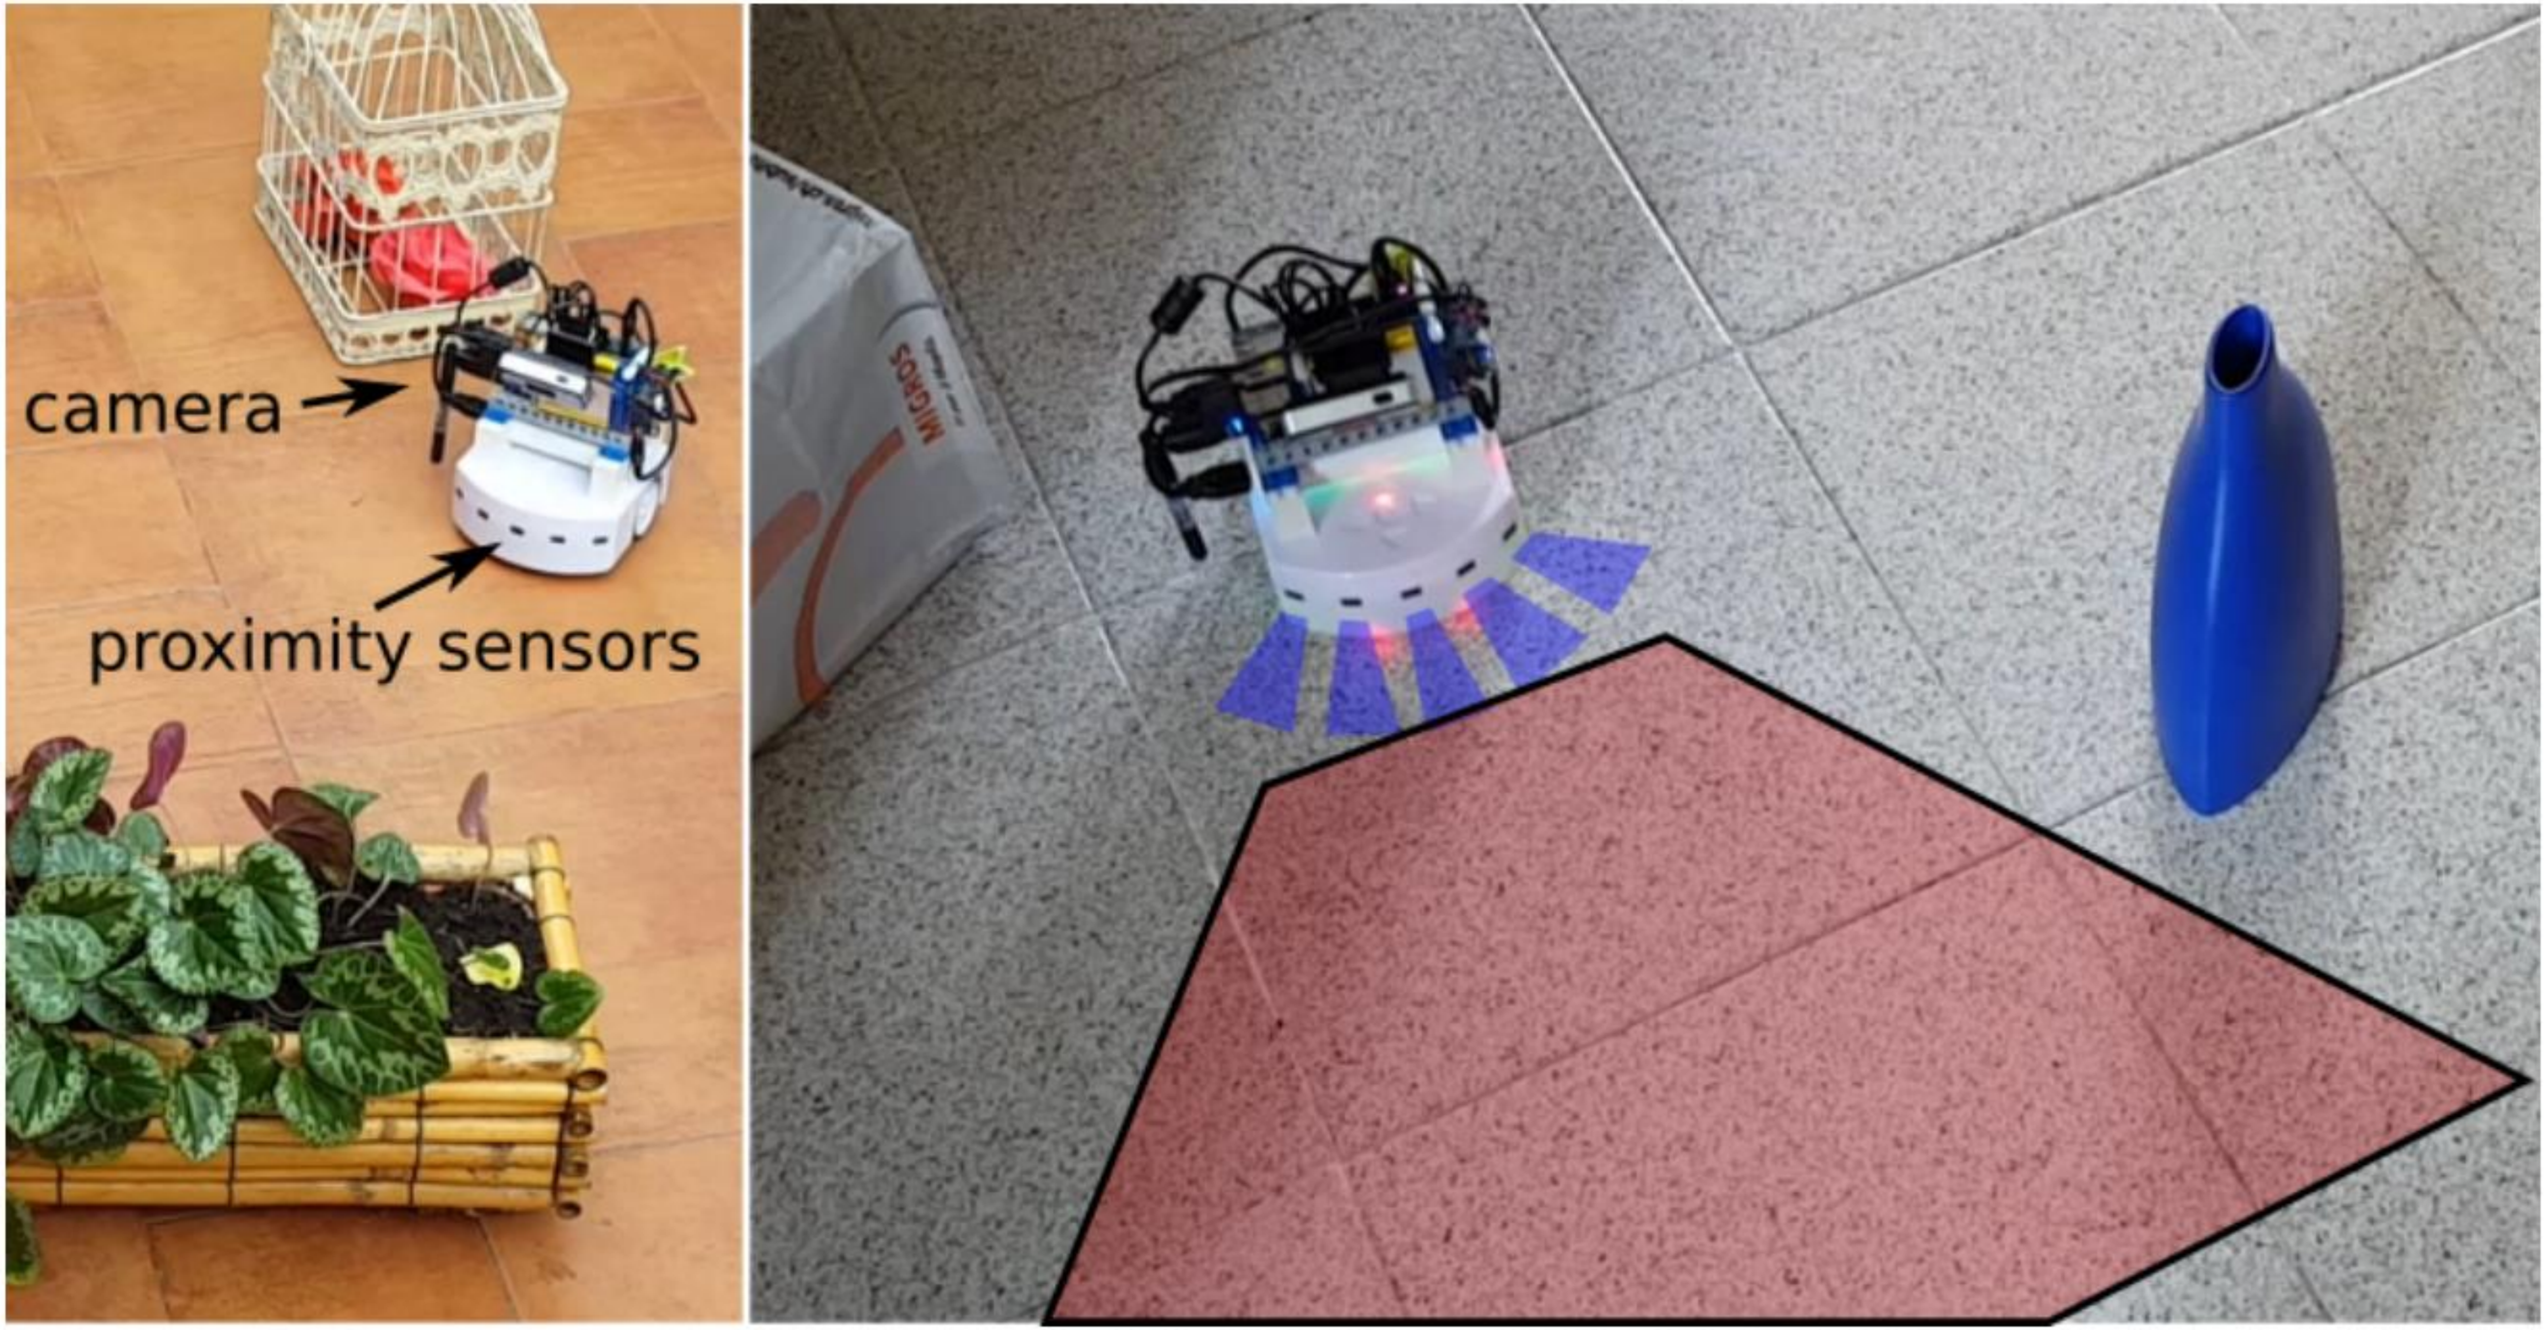
camera
proximity sensors
MIGROS

Cross-sensor prediction (image -> proximity)

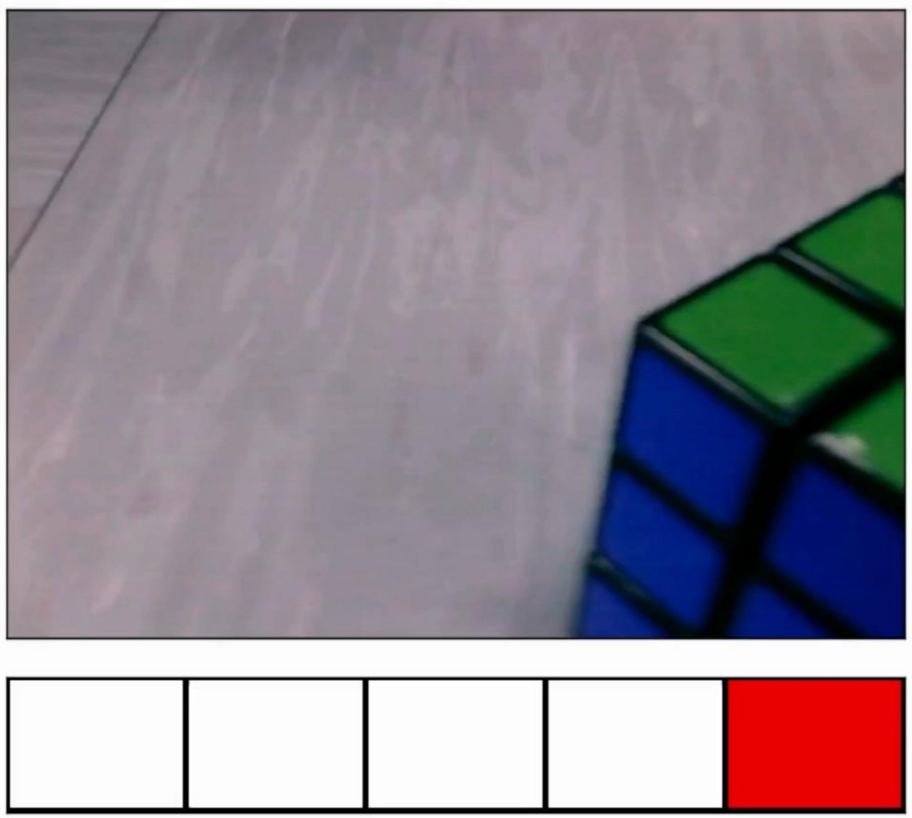

Problem definition example

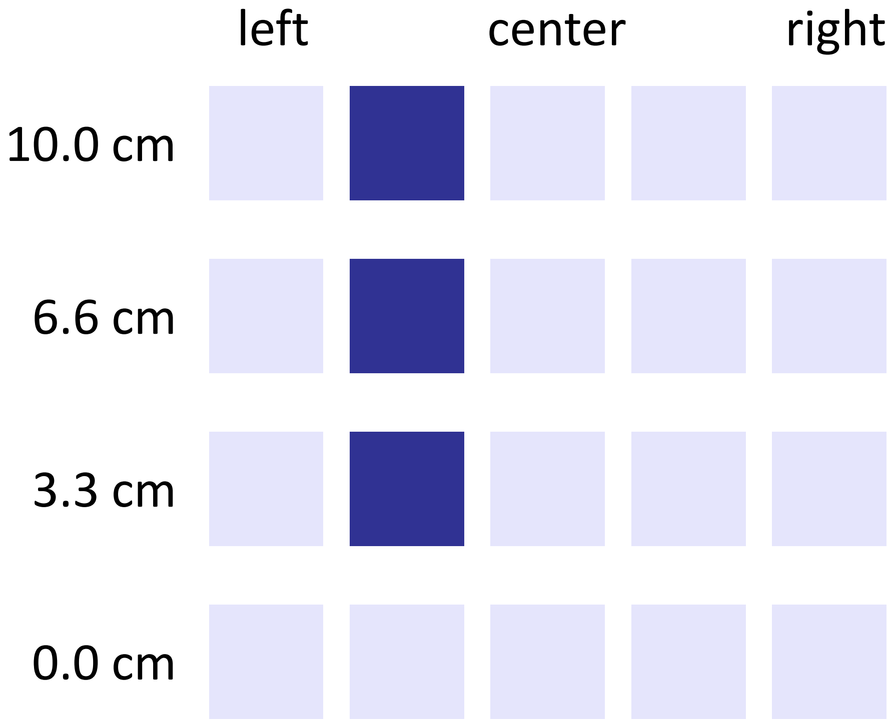

Problem definition example

	left		center		right
10.0 cm					
6.6 cm					
3.3 cm					
0.0 cm					


Obstacle
Sensor's range
Thymio

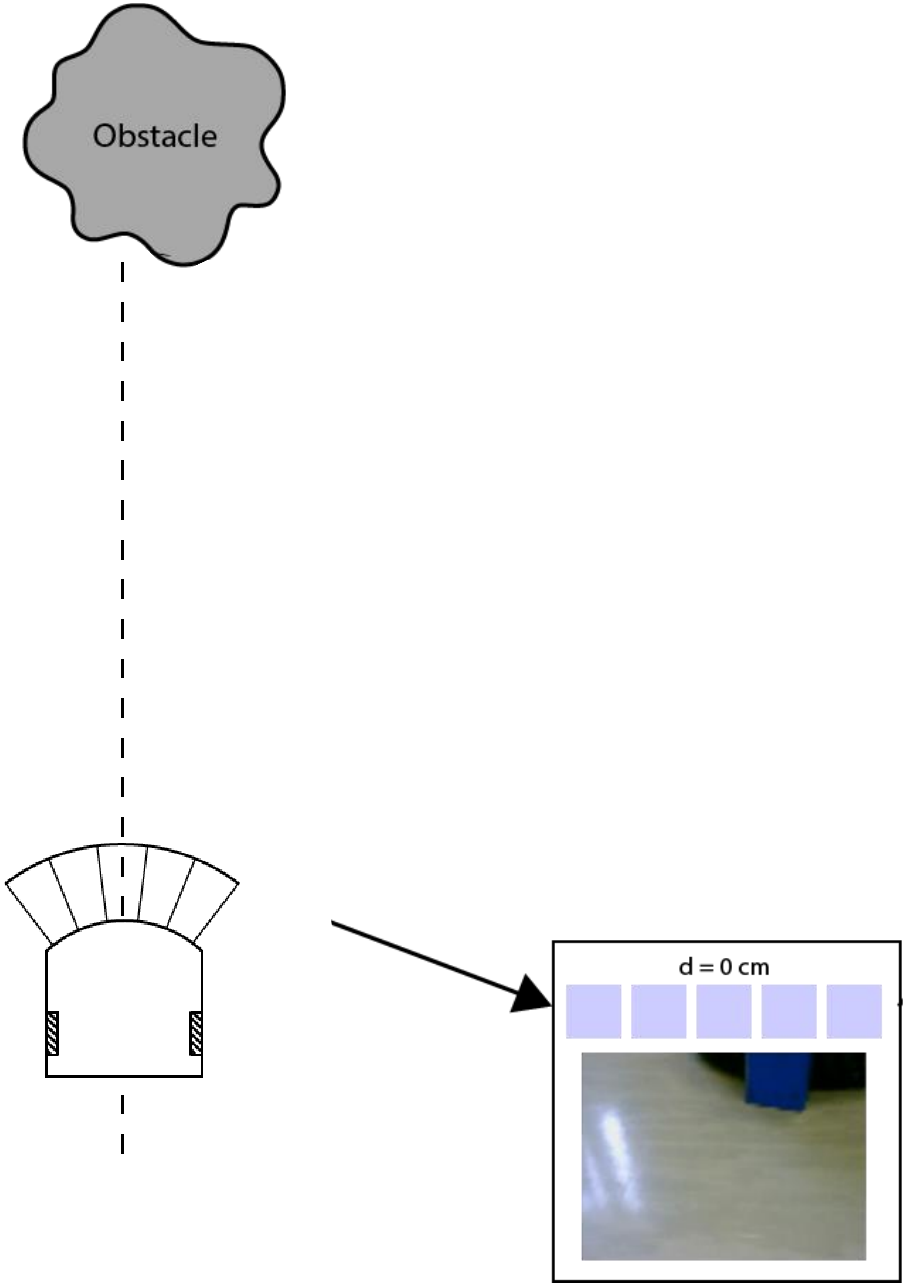
Obstacle
d = 0 cm

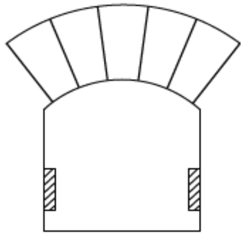
Sensor’s range
Thymio

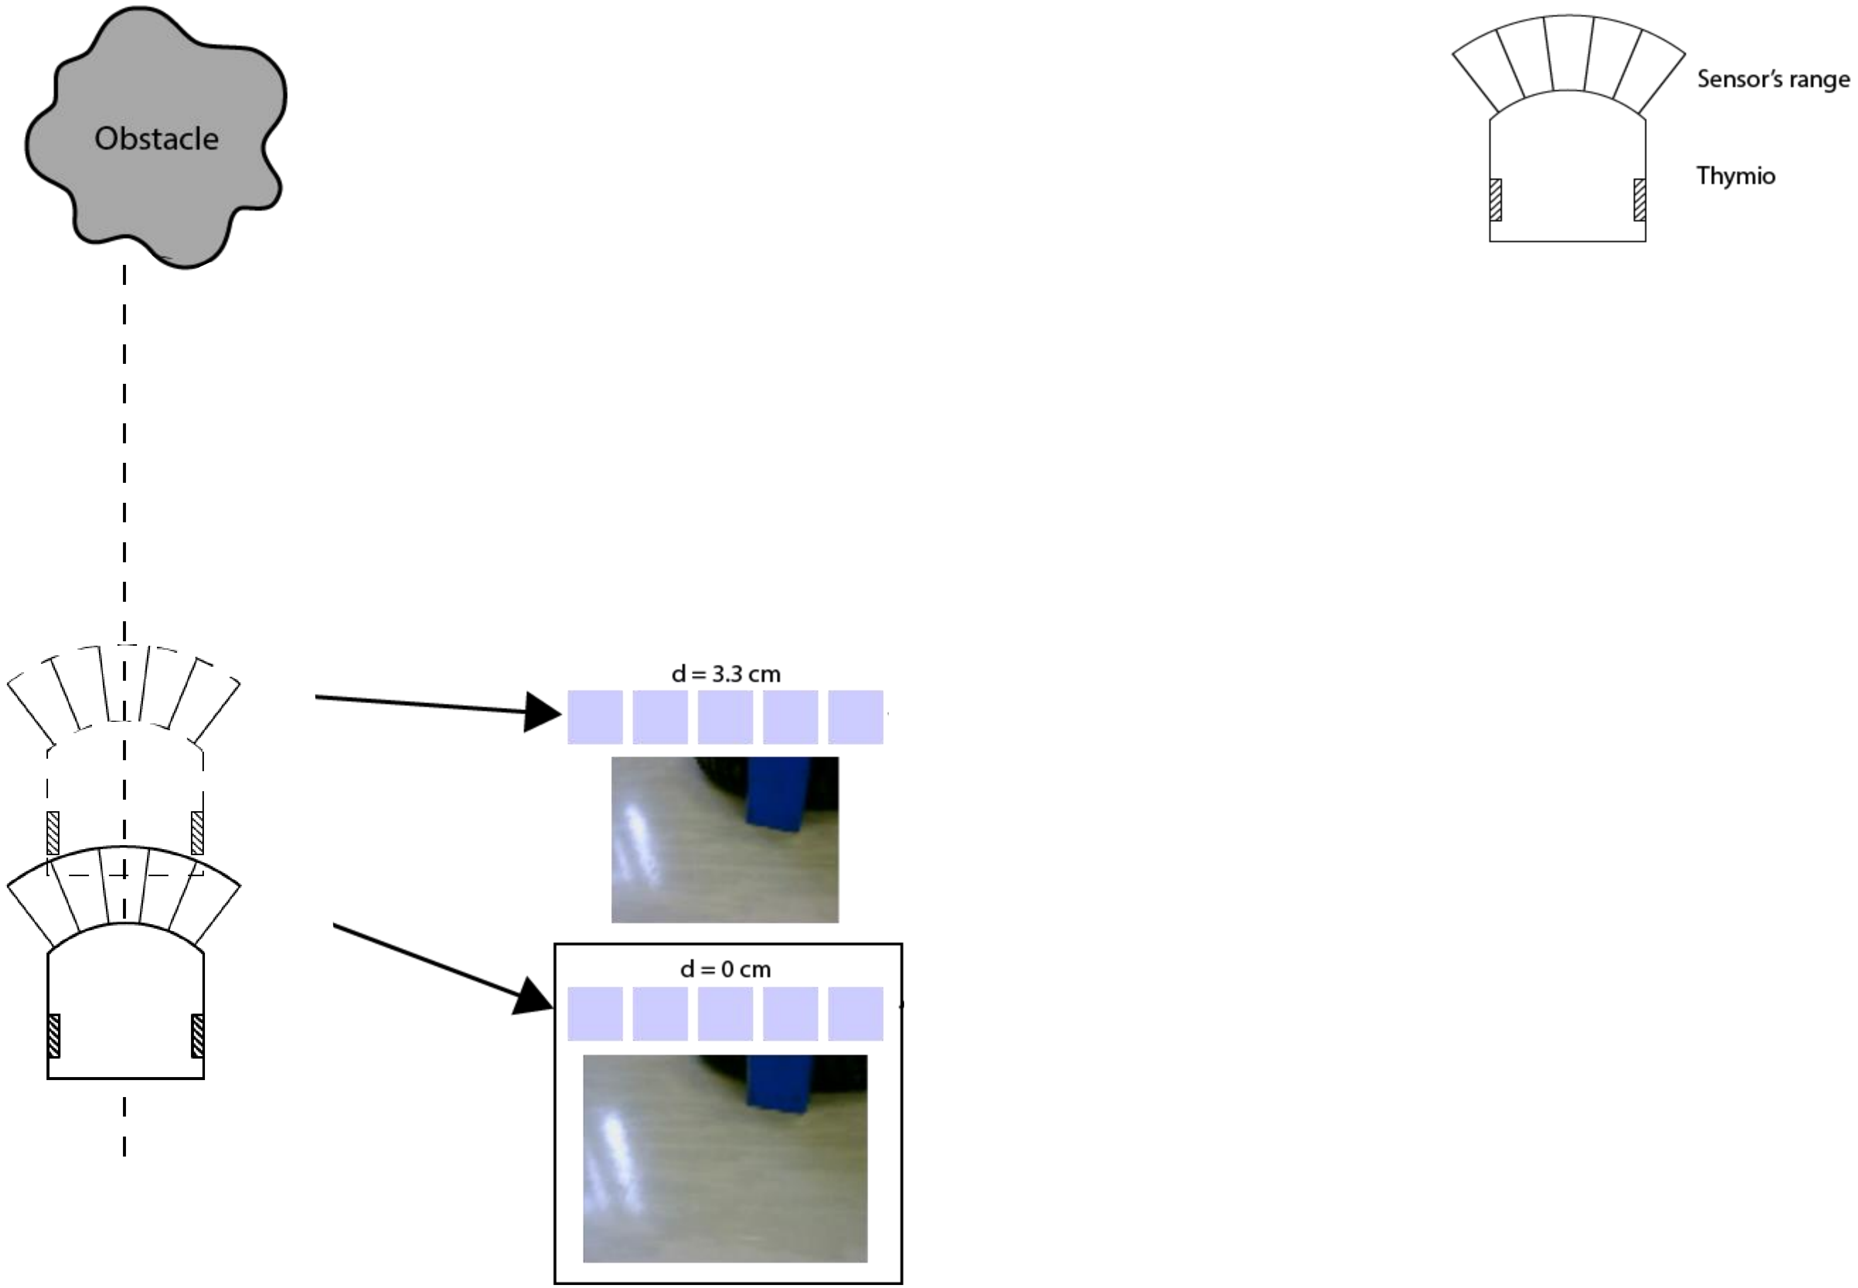
Obstacle
d = 3.3 cm
d = 0 cm
Sensor's range
Thymio

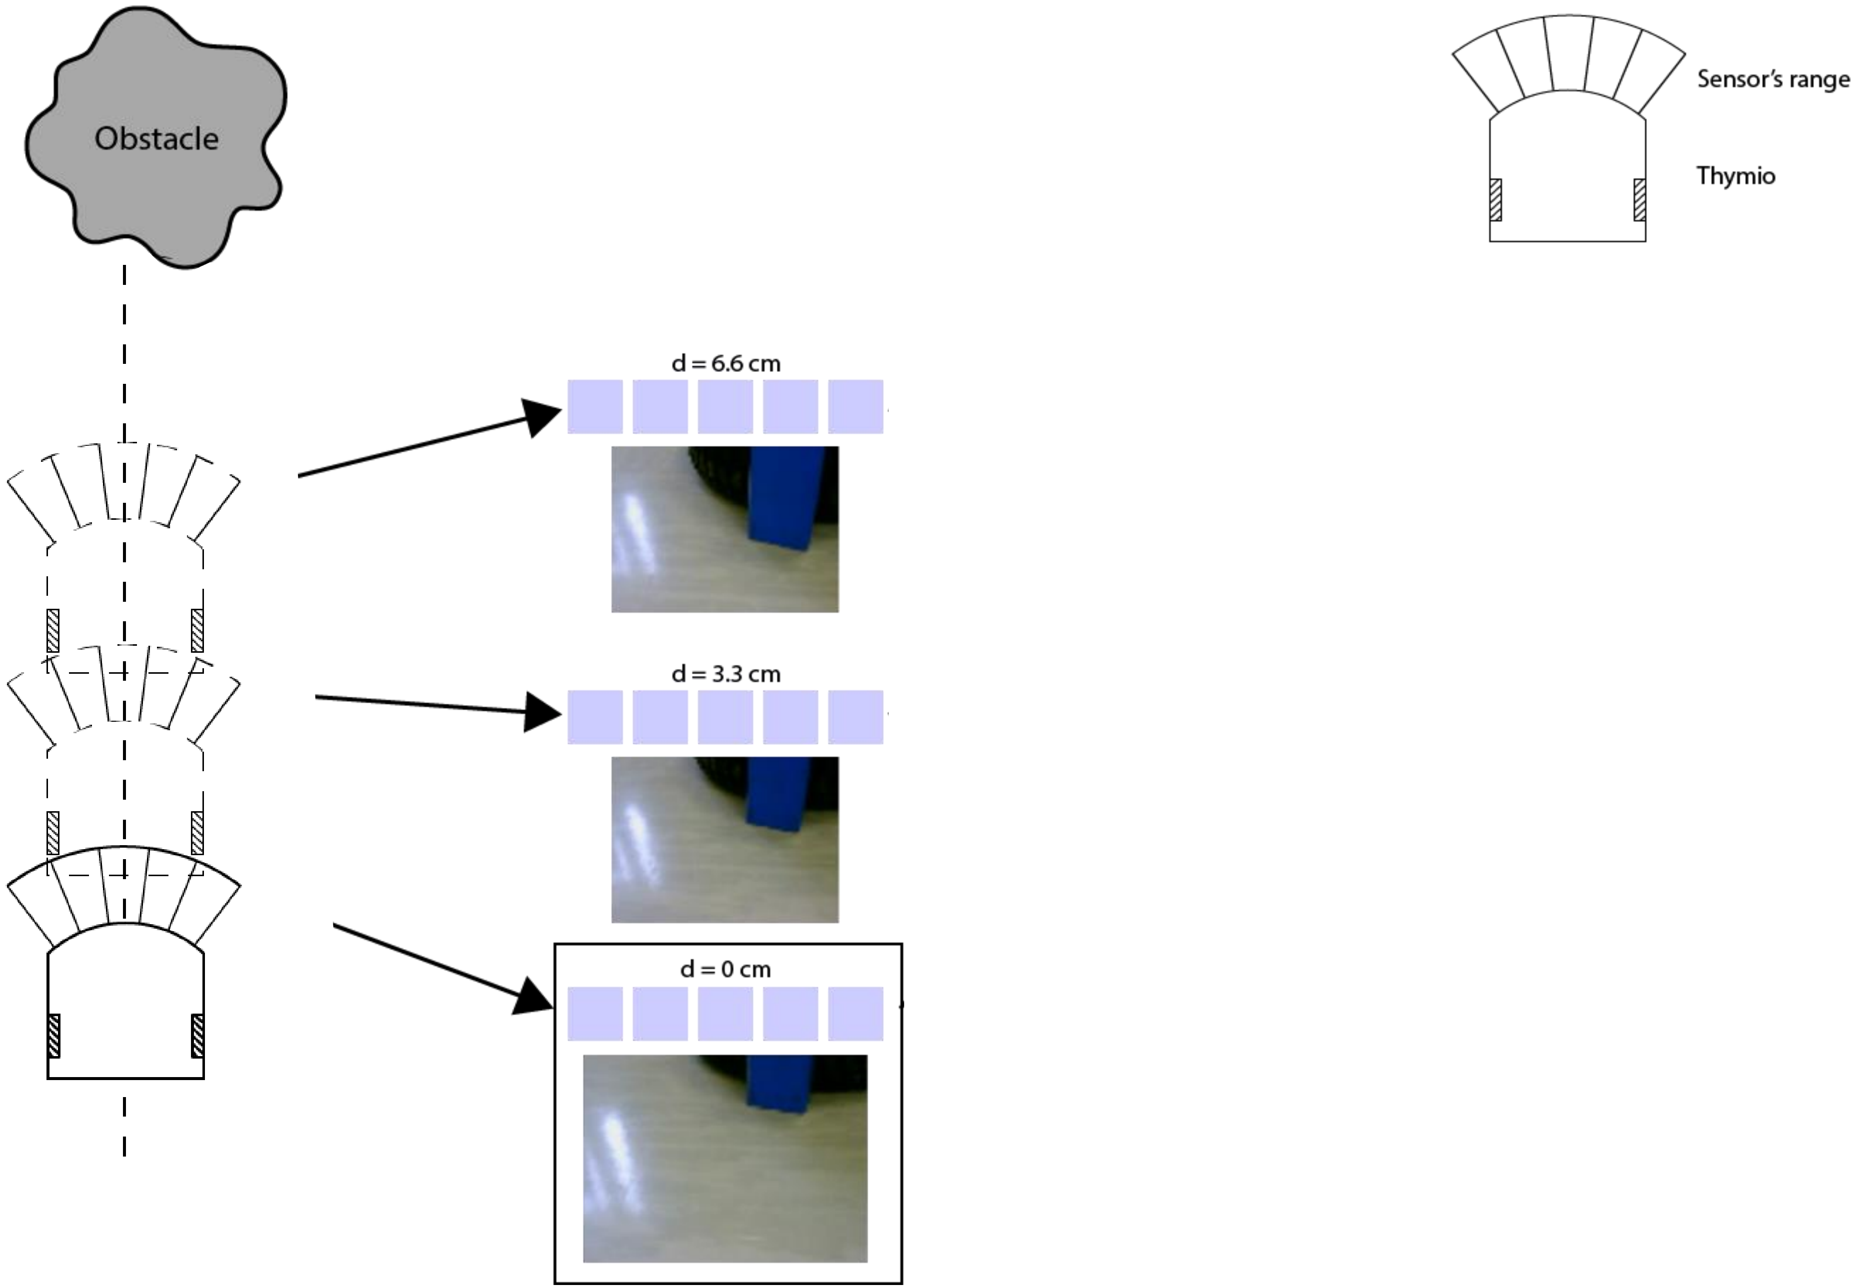
Obstacle
d = 6.6 cm
d = 3.3 cm
d = 0 cm
Sensor's range
Thymio

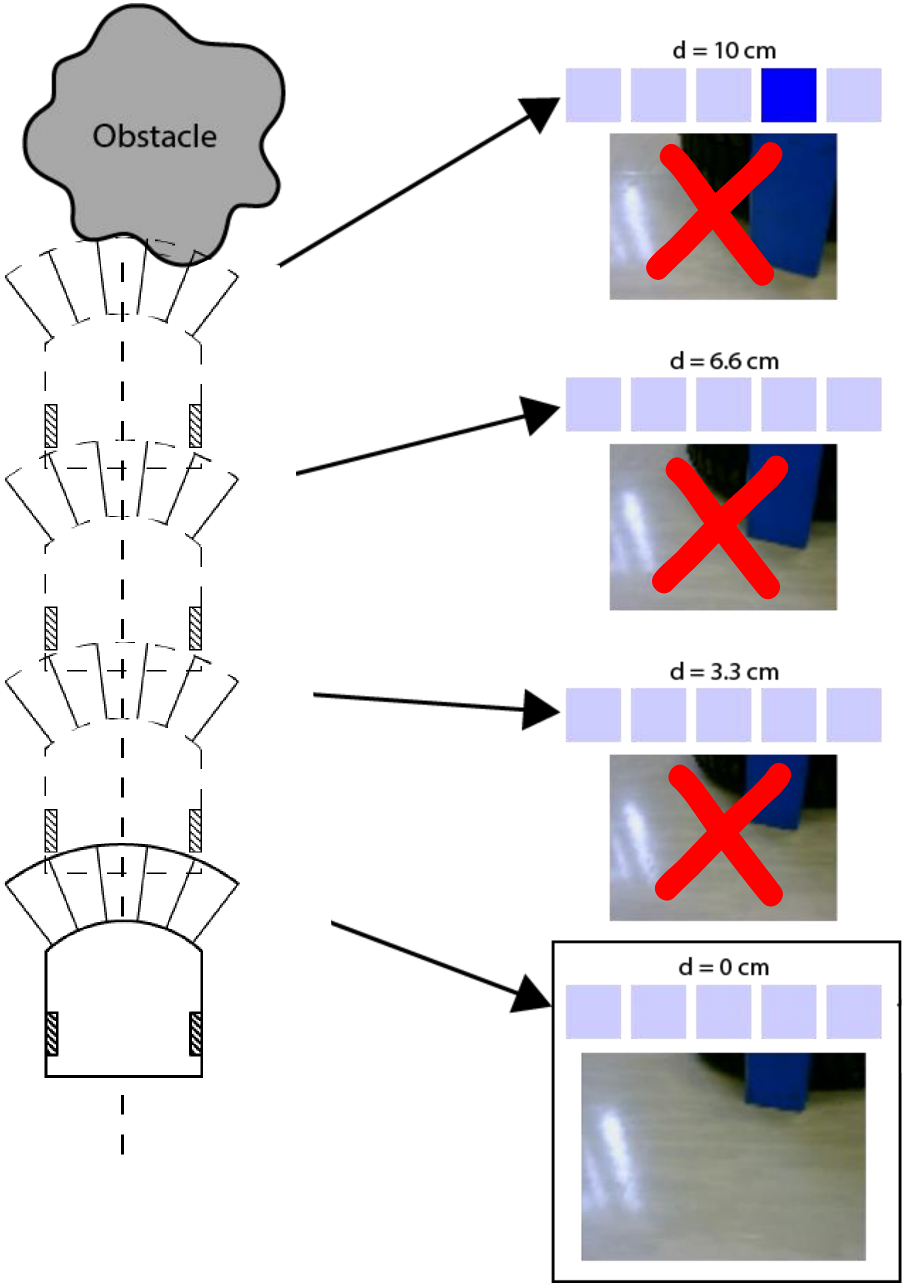
Obstacle
d = 10 cm
d = 6.6 cm
d = 3.3 cm
d = 0 cm

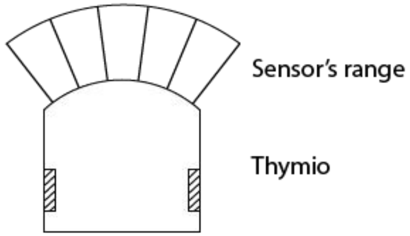
Sensor's range
Thymio

Data gathering

- Various examples:
  - Different distances and directions
  - Floors with different textures
  - Obstacles with different shapes, materials and colors
- 8 recording sessions
- 36k training examples

An (optional) controller for efficient data gathering

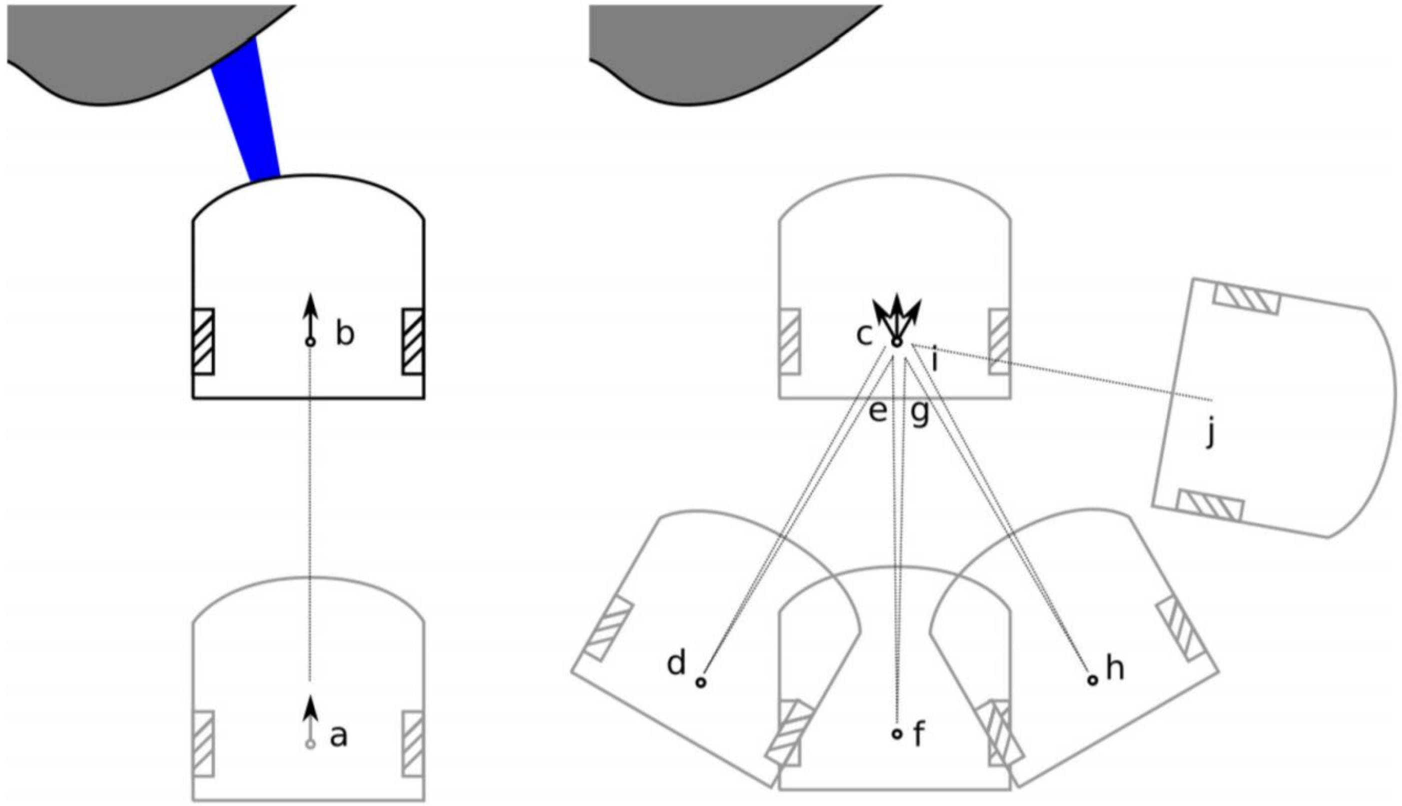

Fig. 4. Example trajectory generated by the data acquisition controller.

Quantitative evaluation

Area Under the Receiver Operating Characteristic Curve

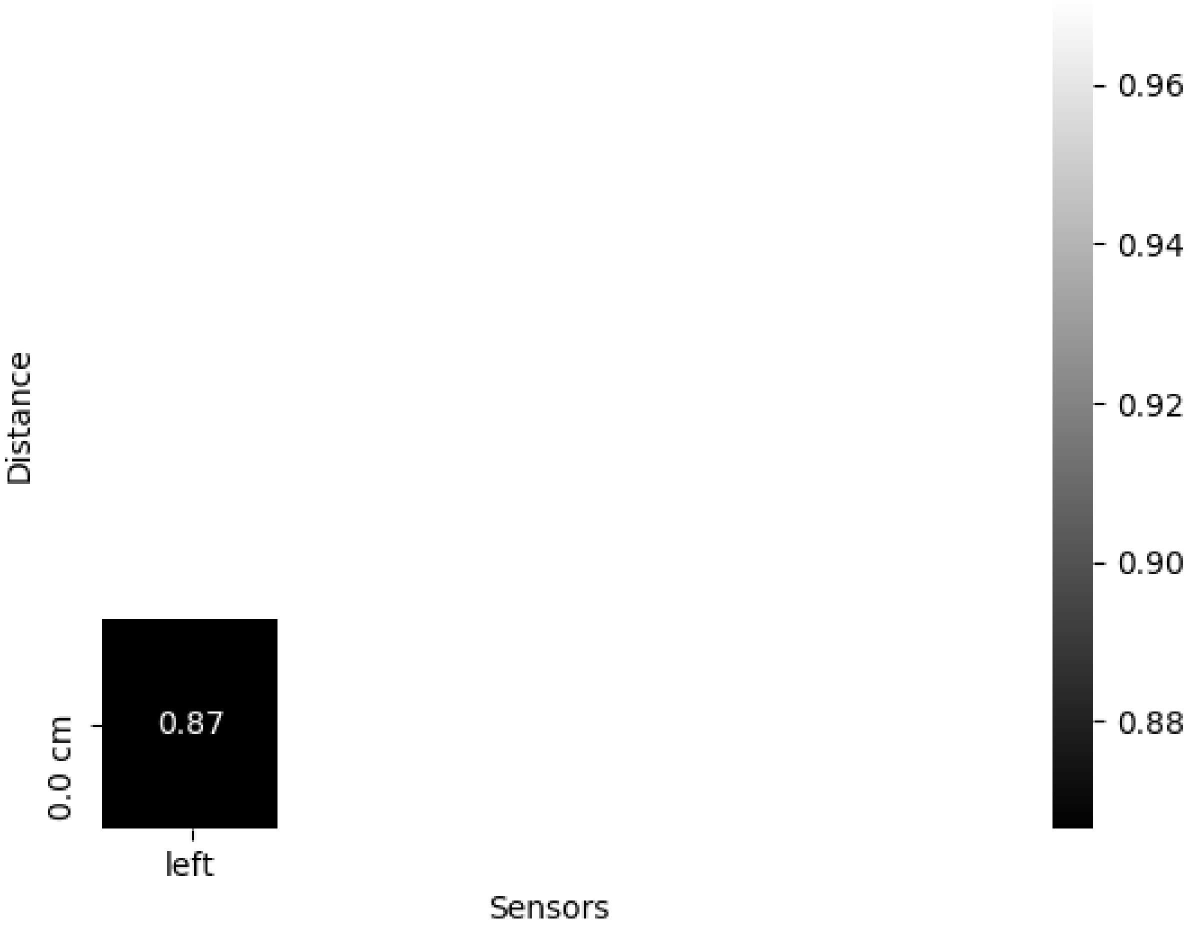

- Symmetric
- Decreases on sides
- Decreases with distance
- Distance = 0 cm is the hardest

Why 0cm is so hard? the camera blind spot!

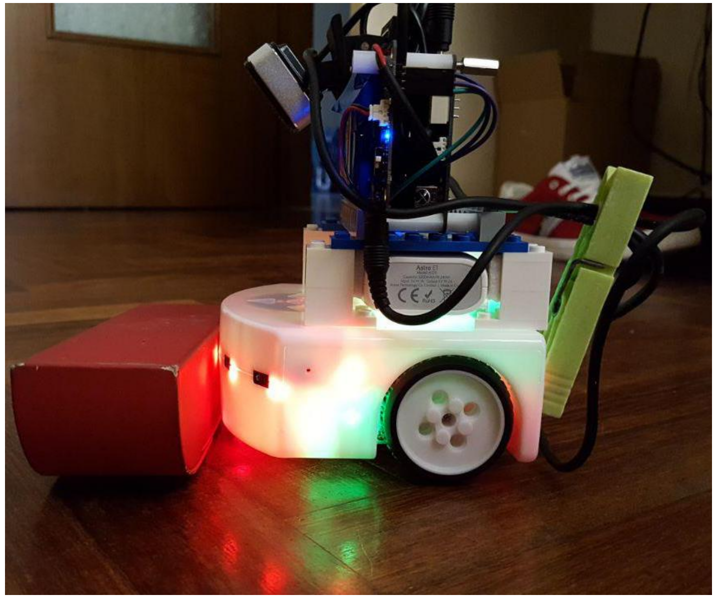

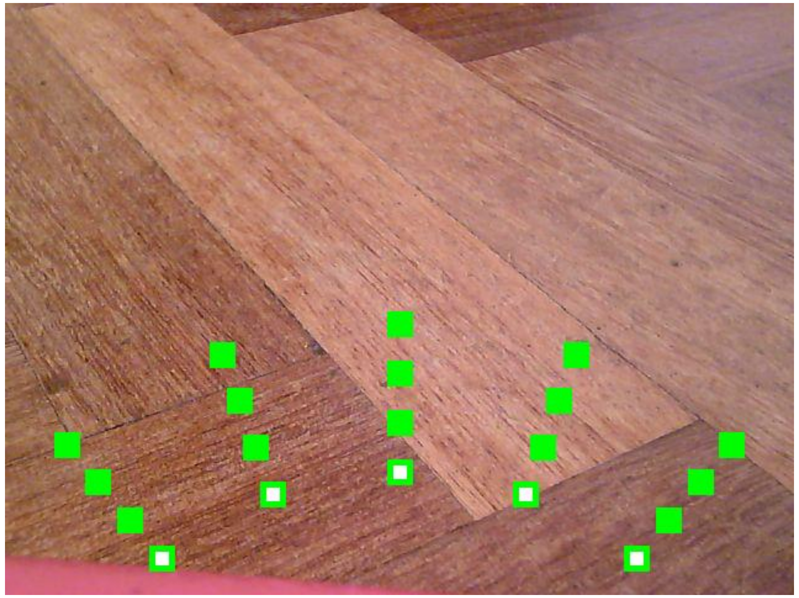

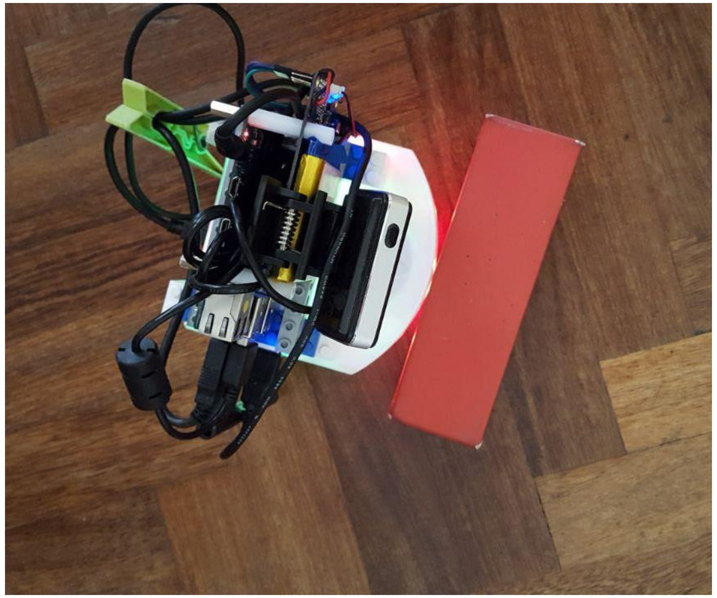

Prediction Target

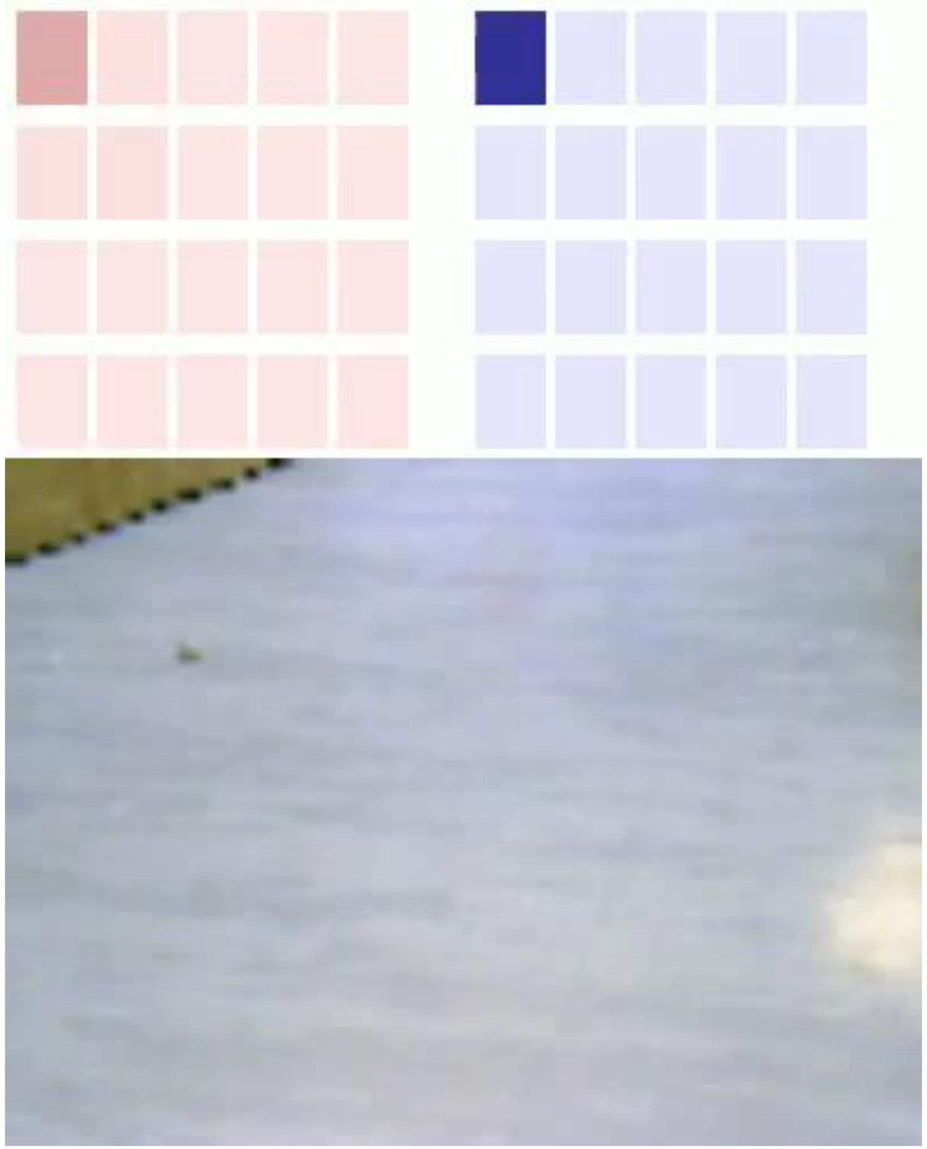

Prediction Target

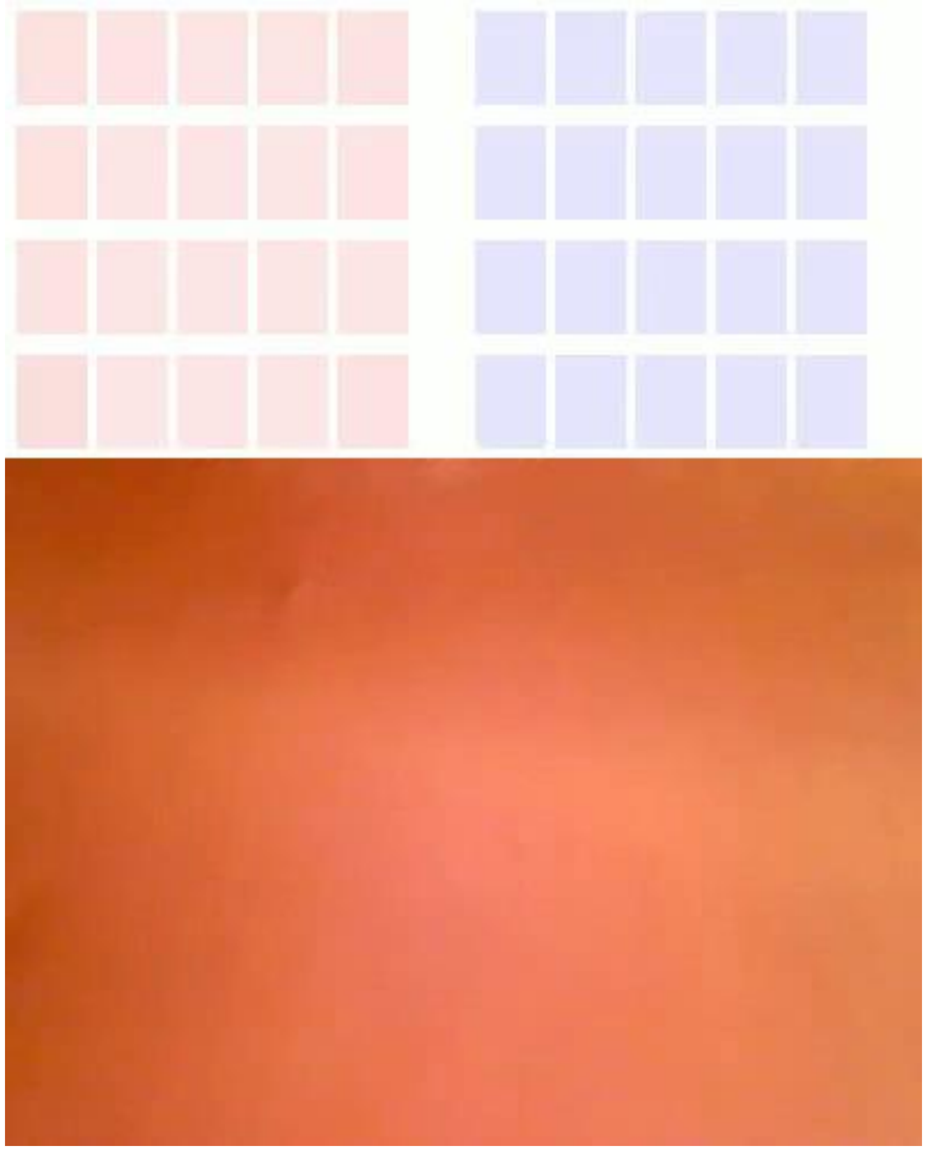

It works!

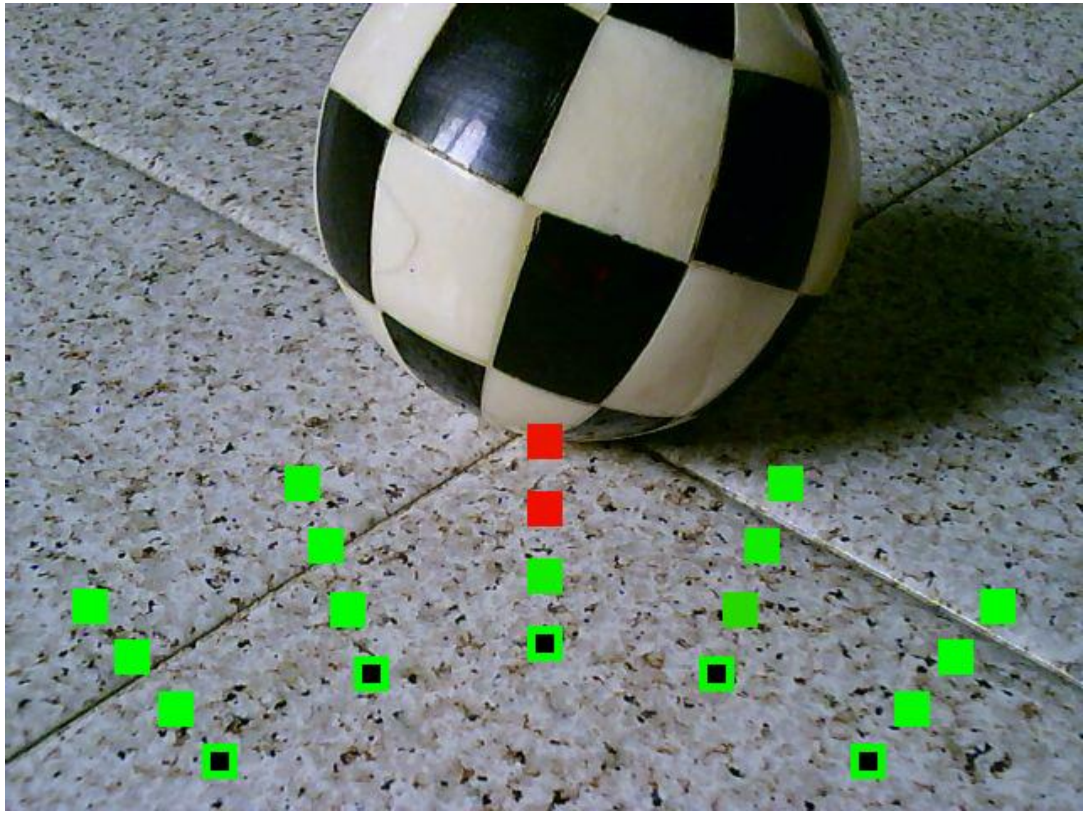

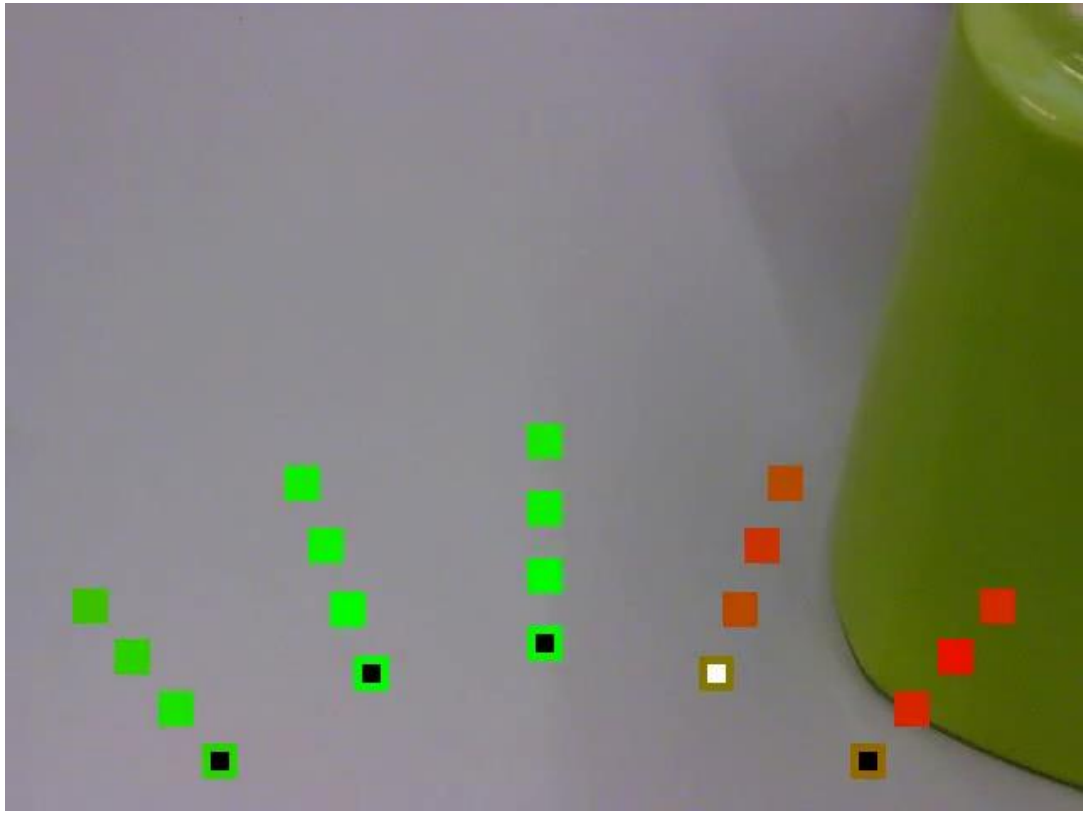

Video

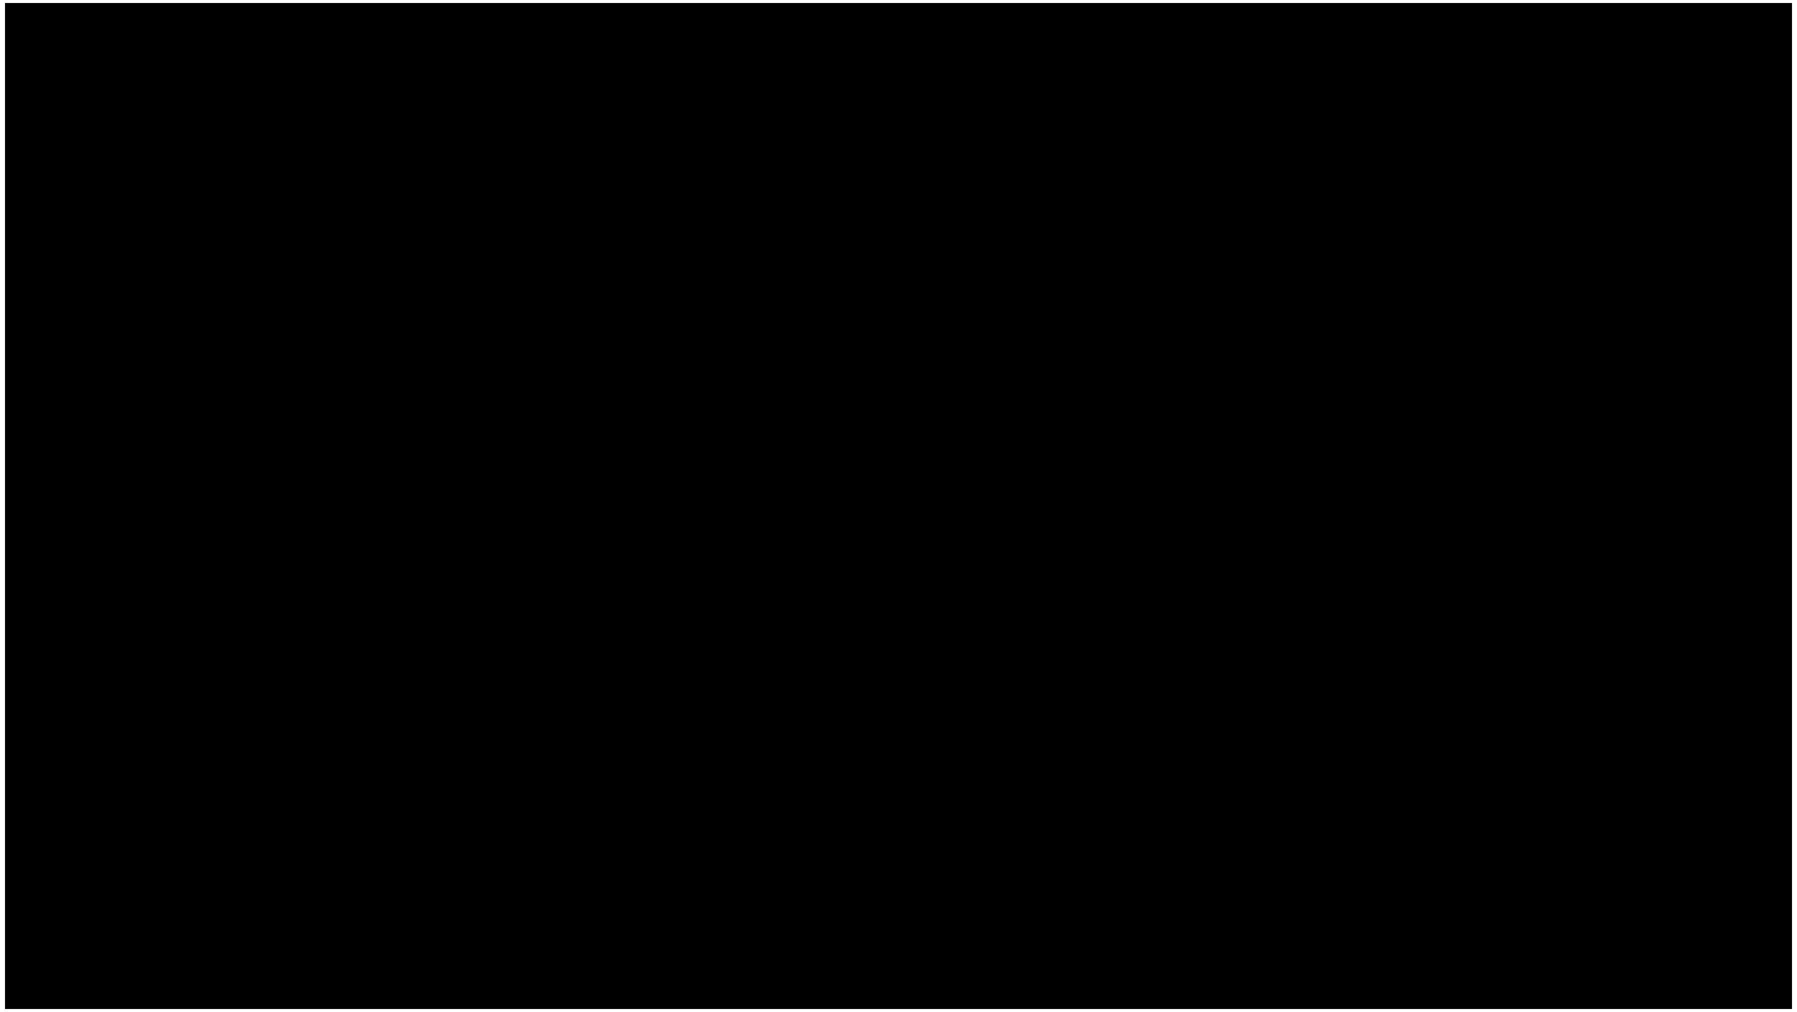

Video

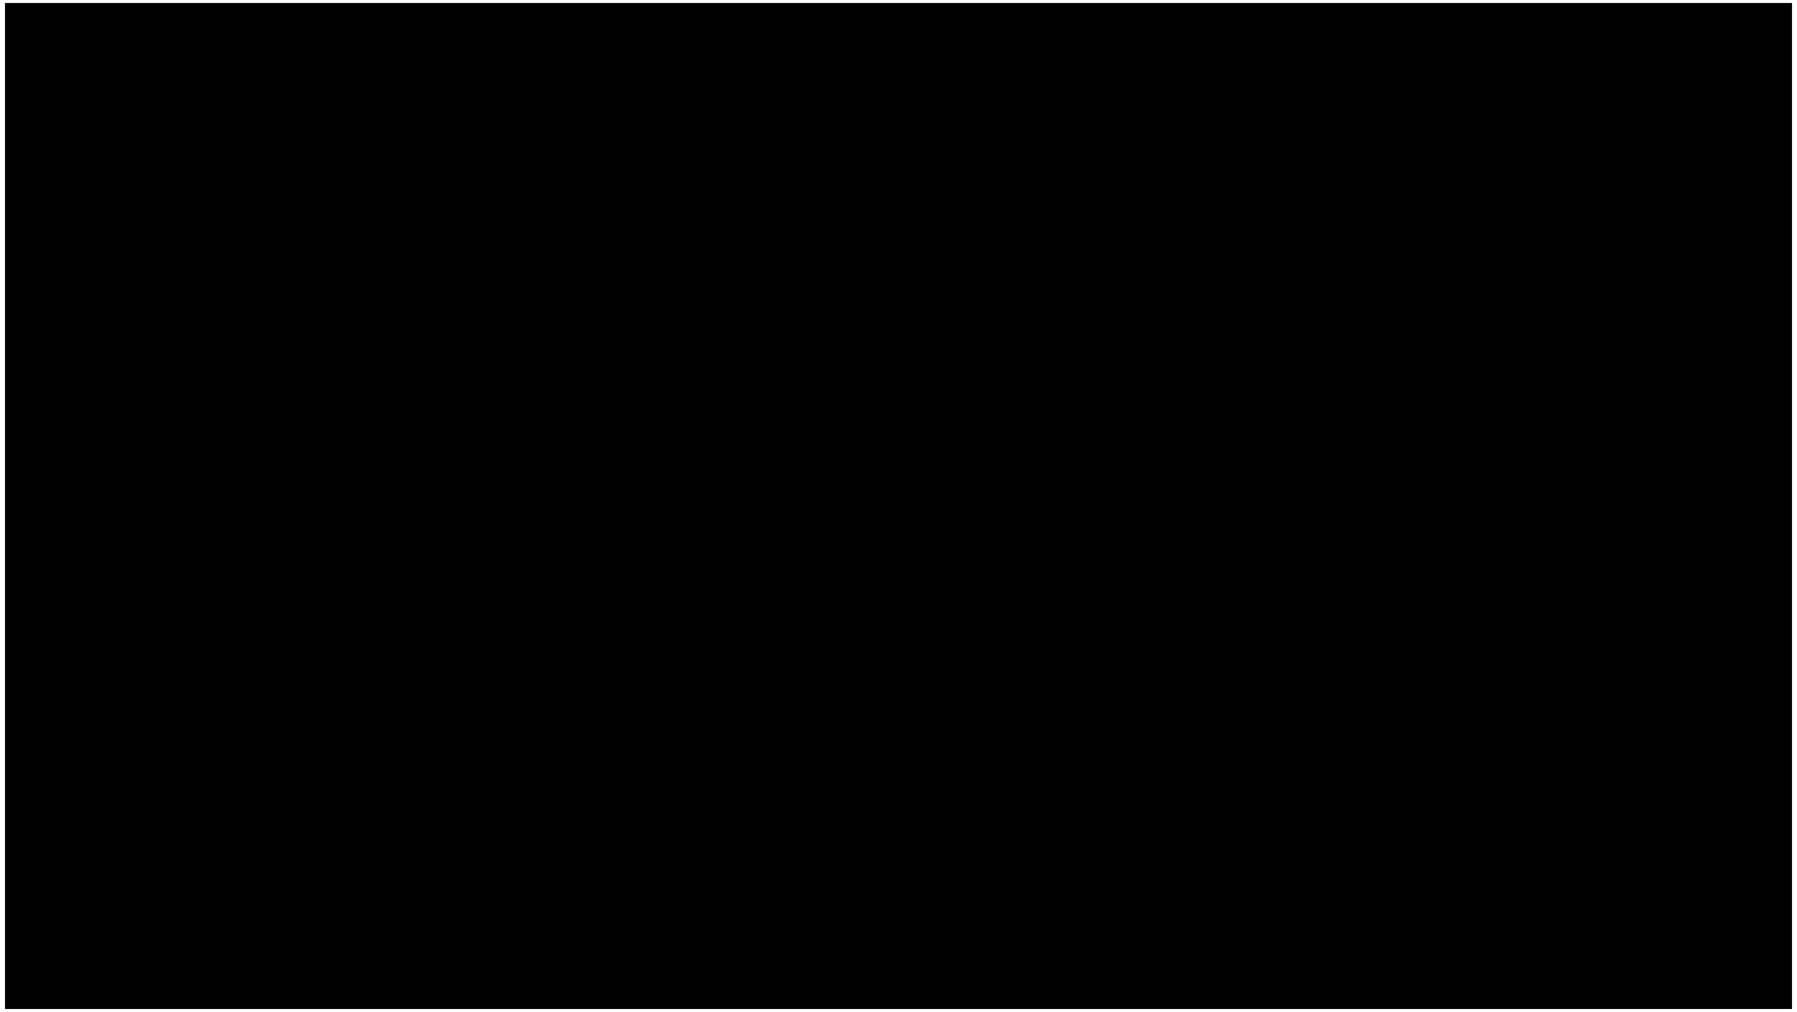

Generalizing…

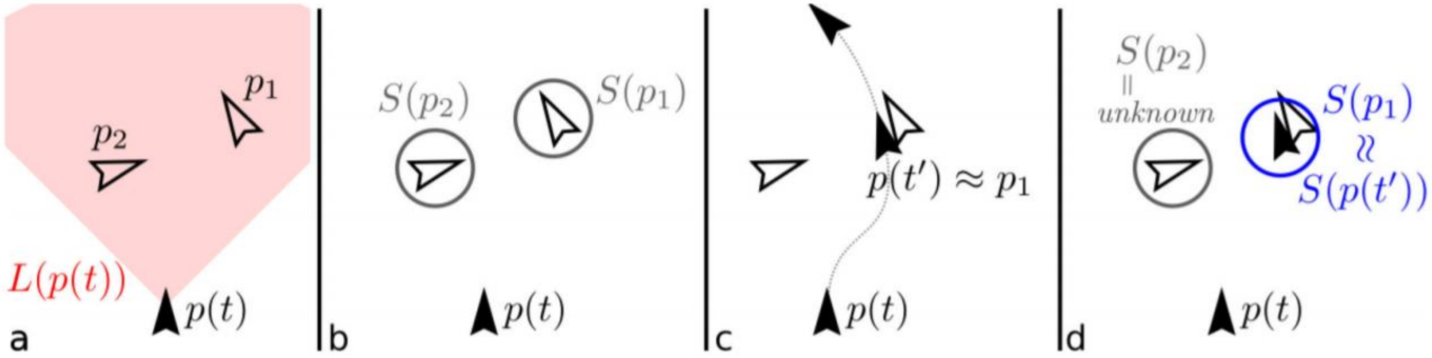

Fig. 2. (a) A mobile robot at pose $p(t)$ has a long-range sensor L (red) and (b) a short-range sensor S. Our objective is to predict the value of S at n target poses $p_1, p_2, \ldots p_n$ from the value of $L(p(t))$. (c, d) For a given instance, we generate ground truth for a subset of labels by searching the robot's future trajectory for poses close to the target poses.

IEEE ROBOTICS AND AUTOMATION LETTERS, VOL. 4, NO. 2, APRIL 2019

1279

Learning Long-Range Perception Using Self-Supervision From Short-Range Sensors and Odometry

Mirko Nava, Jérôme Guzzi, R. Omar Chavez-Garcia, Luca M. Gambardella, and Alessandro Giusti

Abstract—We introduce a general self-supervised approach to predict the future outputs of a short-range sensor (such as a proximity sensor) given the current outputs of a long-range sensor (such as a camera). We assume that the former is directly related to some piece of information to be perceived (such as the presence of an obstacle in a given position), whereas the latter is information rich but hard to interpret directly. We instantiate and implement the approach on a small mobile robot to detect obstacles at various distances using the video stream of the robot's forward-pointing camera, by training a convolutional neural network on automatically-acquired datasets. We quantitatively evaluate the quality of the predictions on unseen scenarios, qualitatively evaluate robustness to different operating conditions, and demonstrate usage as the sole input of an obstacle-avoidance controller. We additionally instantiate the approach on a different simulated scenario with complementary characteristics, to exemplify the generality of our contribution.

Index Terms—Range sensing, computer vision for other robotic applications, deep learning in robotics and automation.

VIDEOS, DATASETS, AND CODE

Videos, datasets, and code to reproduce our results are available at: https://github.com/idsia-robotics/learning-long-range-perception/

I. INTRODUCTION

WE CONSIDER a mobile robot capable of odometry and equipped with at least two sensors: a long-range one, such as a camera or laser scanner; and a short-range sensor such as a proximity sensor or a contact sensor (bumper). We then consider a specific perception task, such as detecting obstacles while roaming the environment. Regardless on the specific choice of the task and sensors, it is often the case that the long-range sensors produce a large amount of data, whose interpretation for the task at hand is complex; conversely, the short-range sensor readings directly solve the task, but with limited range. For example, detecting obstacles in the video stream of a forward-pointing camera is difficult but potentially allows us to detect them while they are still far; solving the same task with a proximity sensor or bumper is straightforward as the sensor directly reports the presence of an obstacle, but only works at very close range.

In this letter we propose a novel technique for solving a perception task by learning to interpret the long-range sensor data; in particular, we adopt a self-supervised learning approach in which future outputs from the short-range sensor are used as a supervisory signal. We develop the complete pipeline for an obstacle-detection task using camera frames as the long-range sensor and proximity sensor readings as the short-range sensor (see Figure 1). In this context, the camera frame acquired at time t (input) is associated to proximity sensor readings obtained at a different time $t' \neq t$ (labels); for example, if the robot's odometry detects it has advanced straight for 10 cm between t and t', the proximity sensor outputs at t' correspond to the presence of obstacles 10 cm in front of the pose of the robot at t. These outputs at time t' can be associated to the camera frame acquired at time t as a label expressing the presence of an obstacle 10cm ahead. The same reasoning can be applied to other distances, so that we define a multi-label classification problem with a single camera frame as input, and multiple binary labels expressing the presence of obstacles at different distances.

The approach is *self-supervised* because it does not require any explicit effort for dataset acquisition or labeling: the robot acquires labeled datasets unattended and can gather additional

camera
proximity sensors

Fig. 1. The Mighty Thymio robot in two environments; five proximity sensors can easily detect obstacles at very close range (blue areas), whereas the camera has a much longer range (red area) but its outputs are hard to interpret.

Manuscript received September 10, 2018; accepted January 15, 2019. Date of publication January 23, 2019; date of current version February 15, 2019. This letter was recommended for publication by Associate Editor L. Paull and Editor C. Stachniss upon evaluation of the reviewers' comments. This work was supported by the Swiss National Science Foundation through the NCCR Robotics. (*Corresponding author: Mirko Nava.*)

The authors are with the Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI, Lugano 6928, Switzerland (e-mail: mirko@idsia.ch; jerome@idsia.ch; omar@idsia.ch; luca@idsia.ch; alessandrog@idsia.ch).

Digital Object Identifier 10.1109/LRA.2019.2894849

https://github.com/idsia-robotics/learning-long-range-perception

A seminal paper from 2006

A robot that acquires its own training data...
but then solves a standard Supervised Learning problem

Self-supervised online learning for big-ass Robots

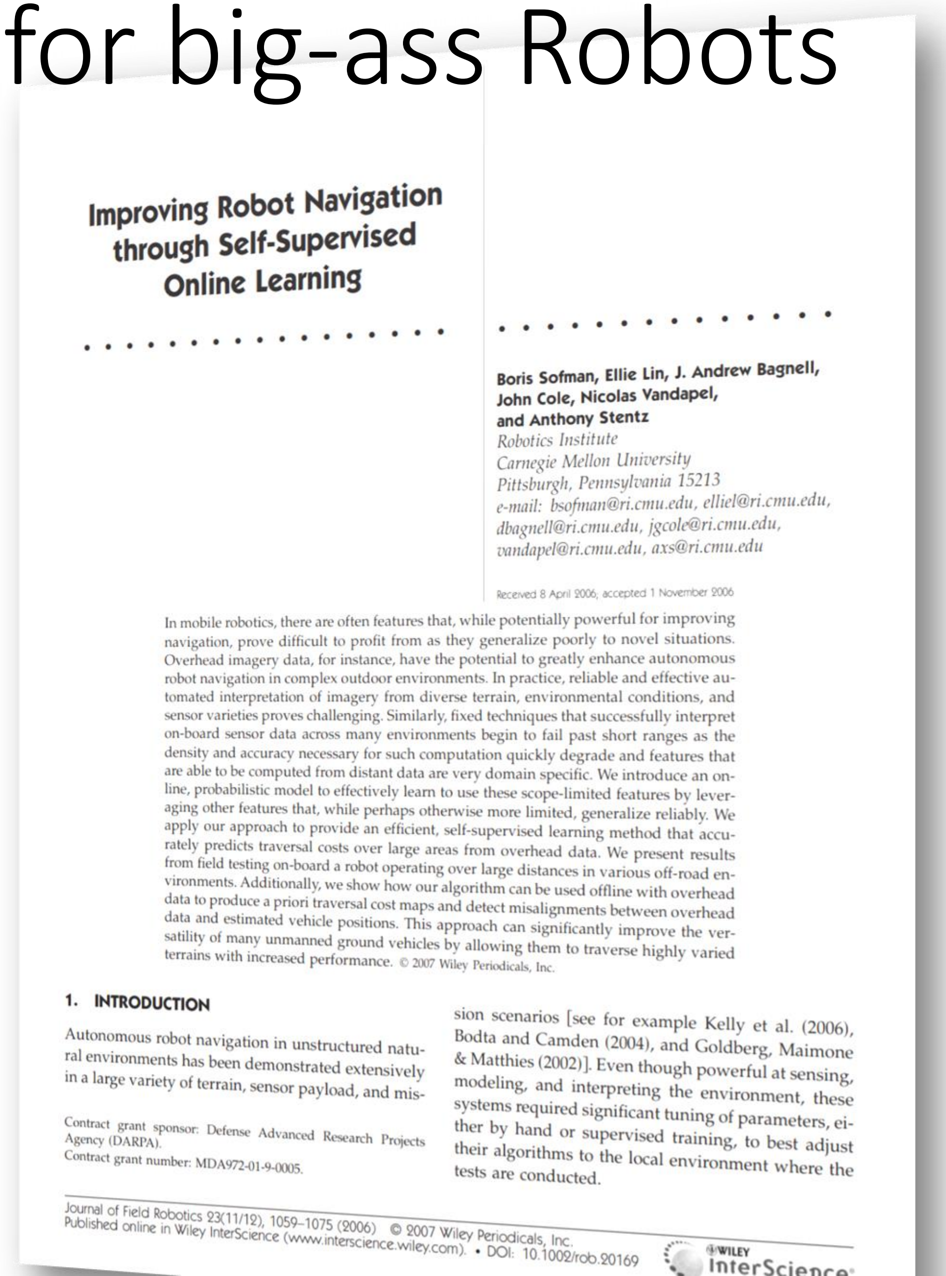

Improving Robot Navigation through Self-Supervised Online Learning

Boris Sofman, Ellie Lin, J. Andrew Bagnell, John Cole, Nicolas Vandapel, and Anthony Stentz
Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
e-mail: bsofman@ri.cmu.edu, elliel@ri.cmu.edu, dbagnell@ri.cmu.edu, jgcole@ri.cmu.edu, vandapel@ri.cmu.edu, axs@ri.cmu.edu

Received 8 April 2006; accepted 1 November 2006

In mobile robotics, there are often features that, while potentially powerful for improving navigation, prove difficult to profit from as they generalize poorly to novel situations. Overhead imagery data, for instance, have the potential to greatly enhance autonomous robot navigation in complex outdoor environments. In practice, reliable and effective automated interpretation of imagery from diverse terrain, environmental conditions, and sensor varieties proves challenging. Similarly, fixed techniques that successfully interpret on-board sensor data across many environments begin to fail past short ranges as the density and accuracy necessary for such computation quickly degrade and features that are able to be computed from distant data are very domain specific. We introduce an online, probabilistic model to effectively learn to use these scope-limited features by leveraging other features that, while perhaps otherwise more limited, generalize reliably. We apply our approach to provide an efficient, self-supervised learning method that accurately predicts traversal costs over large areas from overhead data. We present results from field testing on-board a robot operating over large distances in various off-road environments. Additionally, we show how our algorithm can be used offline with overhead data to produce a priori traversal cost maps and detect misalignments between overhead data and estimated vehicle positions. This approach can significantly improve the versatility of many unmanned ground vehicles by allowing them to traverse highly varied terrains with increased performance. © 2007 Wiley Periodicals, Inc.

1. INTRODUCTION

Autonomous robot navigation in unstructured natural environments has been demonstrated extensively in a large variety of terrain, sensor payload, and mission scenarios [see for example Kelly et al. (2006), Bodta and Camden (2004), and Goldberg, Maimone & Matthies (2002)]. Even though powerful at sensing, modeling, and interpreting the environment, these systems required significant tuning of parameters, either by hand or supervised training, to best adjust their algorithms to the local environment where the tests are conducted.

Contract grant sponsor: Defense Advanced Research Projects Agency (DARPA).
Contract grant number: MDA972-01-9-0005.

The task

Predict the traversal cost of terrain given overhead data

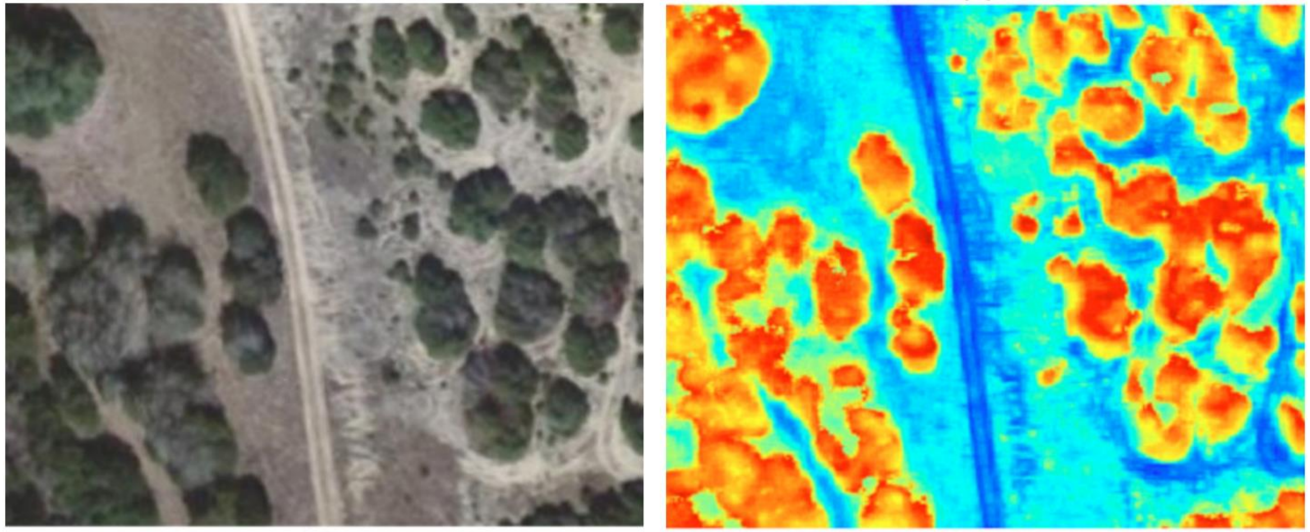

(a) (b)

Figure 2. Sample results of terrain traversal cost predictions. (a) 0.35 m resolution color overhead imagery used by our online learning algorithm and (b) corresponding predictions of terrain traversal costs. Traversal costs are color-scaled for improved visibility. Blue and red correspond to lowest and highest traversal cost estimates, respectively.

Supervision

- Short range ladar
- Robot assigns traversal costs to areas in front of itself from features computed by interpreting the position, density, and point cloud distributions of sensed obstacles

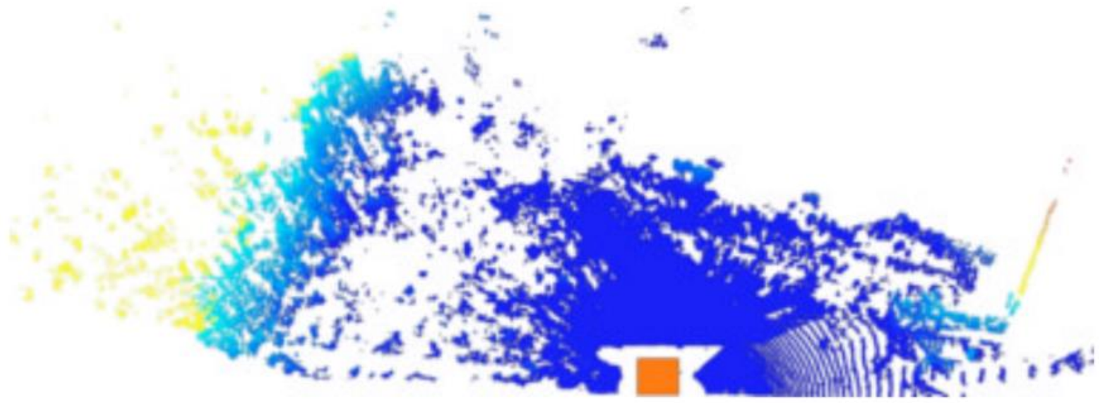

Figure 1. Typical ladar response from vehicle's perception system. Ladar points are color coded by elevation with lowest points appearing in blue and highest points appearing in yellow. Vehicle position is shown by the orange square. Notice the large drop in ladar response density (especially on the ground) as distance from the vehicle increases. Large objects such as the trees on the left generate ladar responses even at far ranges but are difficult to interpret through fixed techniques across different environments.

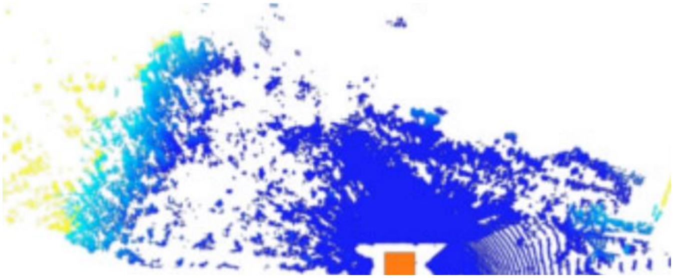

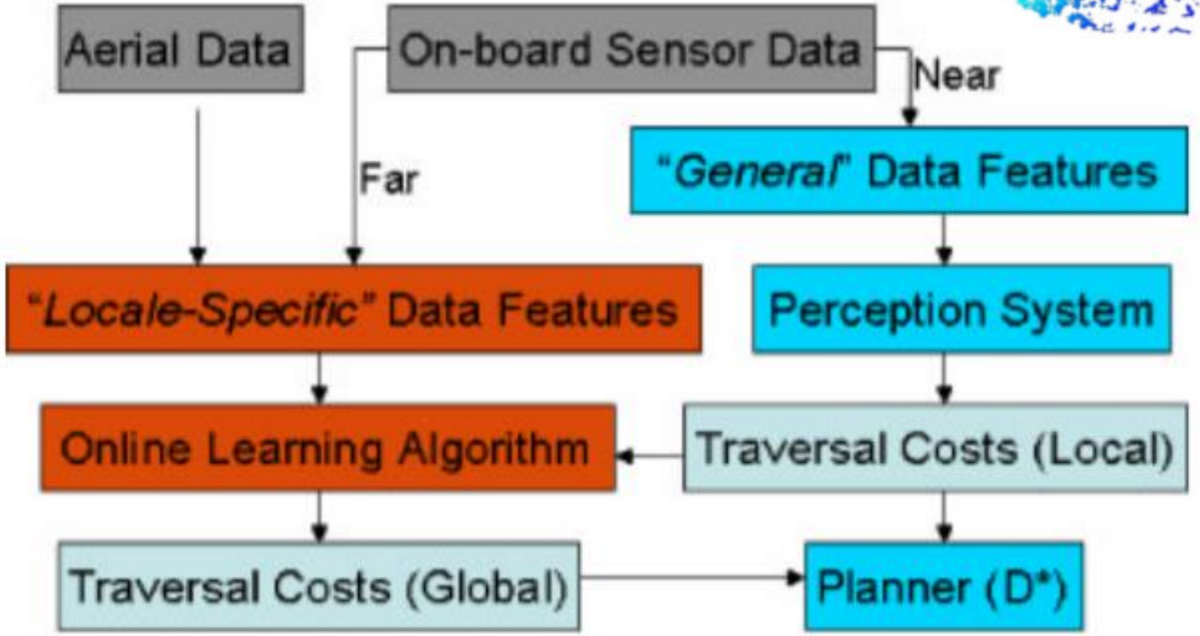

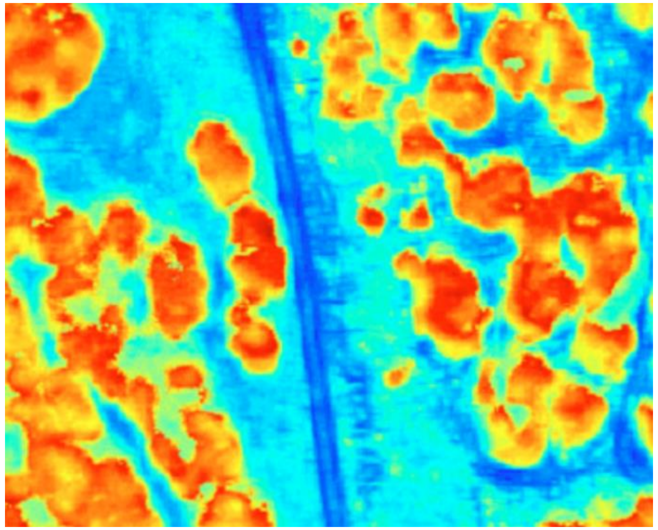

This!

Quiz:
What would you call "supervisory signal" here?

A big advantage: online learning

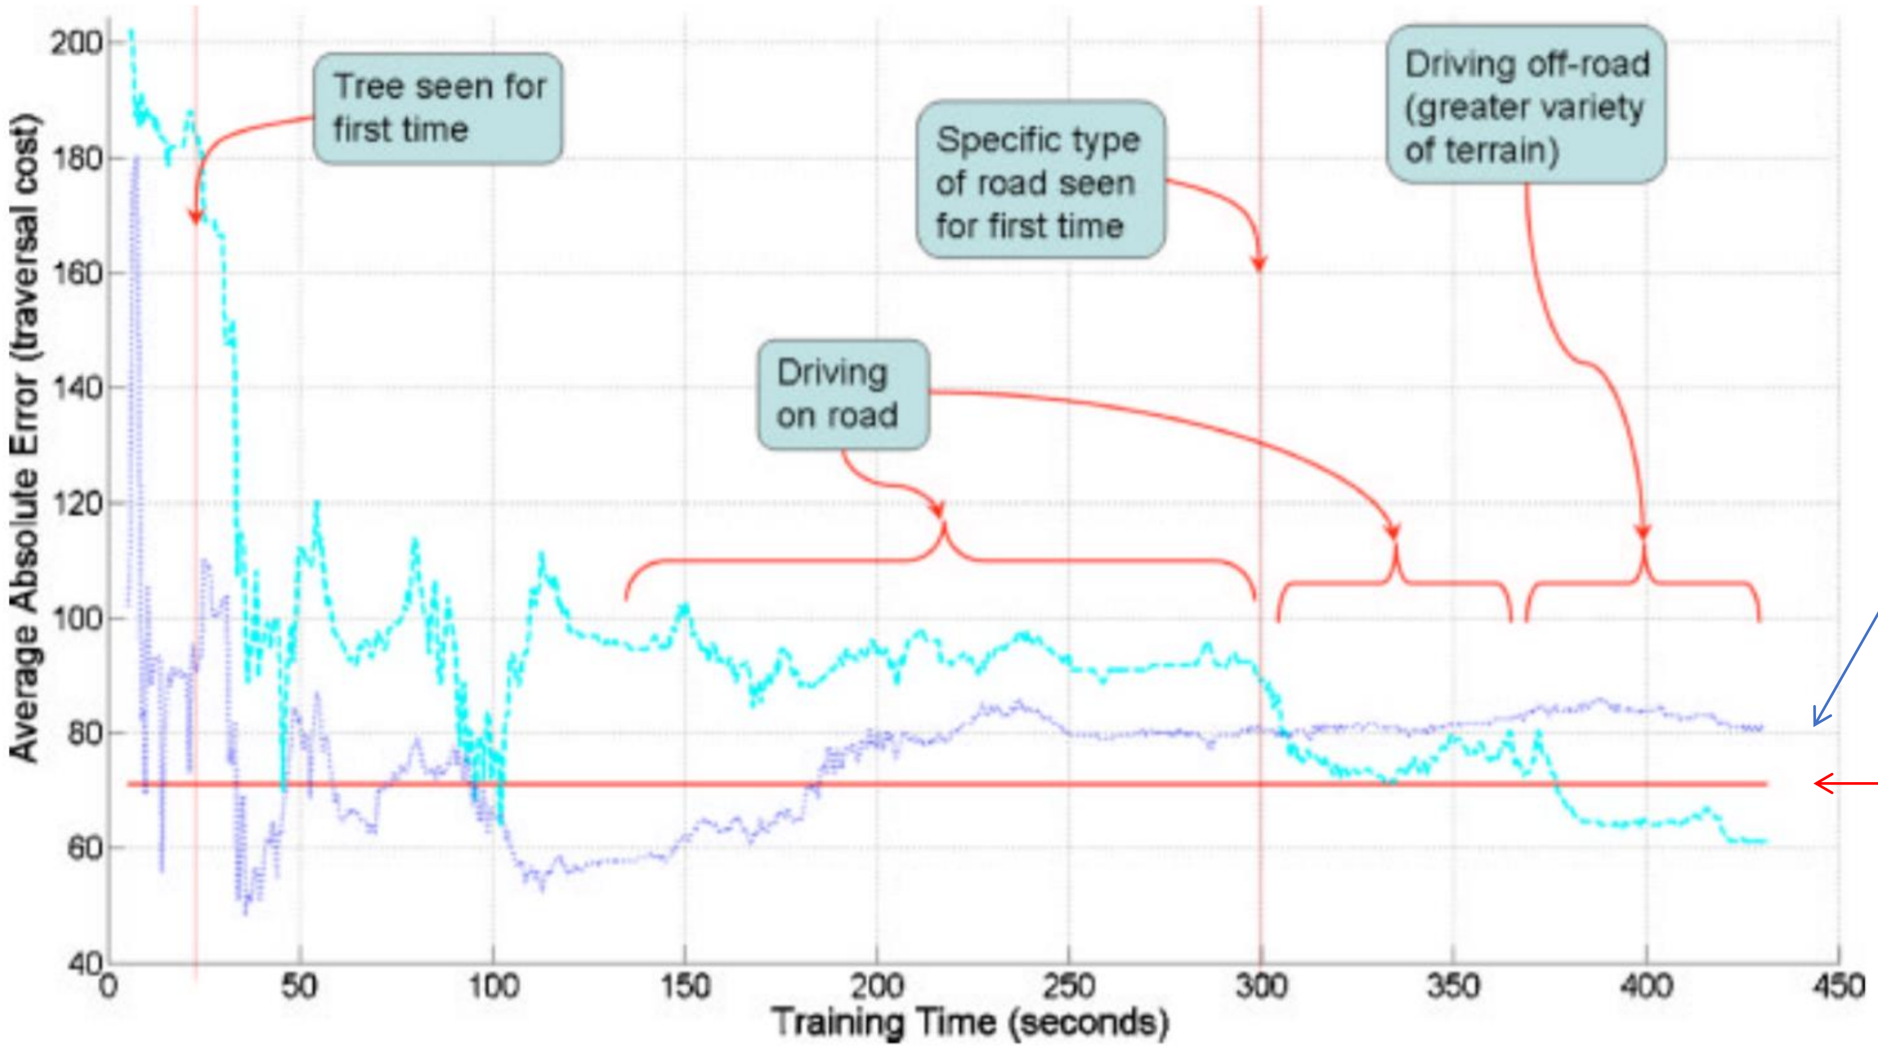

Performance of self-supervised models

Performance of a supervised system trained from human collected data

How do you evaluate something like this?

Figure 7. Comparison of paths executed by our robot for shown course when using only on-board perception (in solid red) and with OOLL (in dashed blue) and FROLL (in dotted cyan) used in real-time on-board the robot. Course started at the top right and ended at the bottom left.

Table I. Statistics for course traversals with and without online learning algorithm

	Without algorithm	With OOLL
Total Traversal time (s)	1369.86	1000.82
Total distance traveled (m)	1815.71	1681.73
Average speed (m/s)	1.33	1.68
No. of interventions	1	0

How do you evaluate something like this?

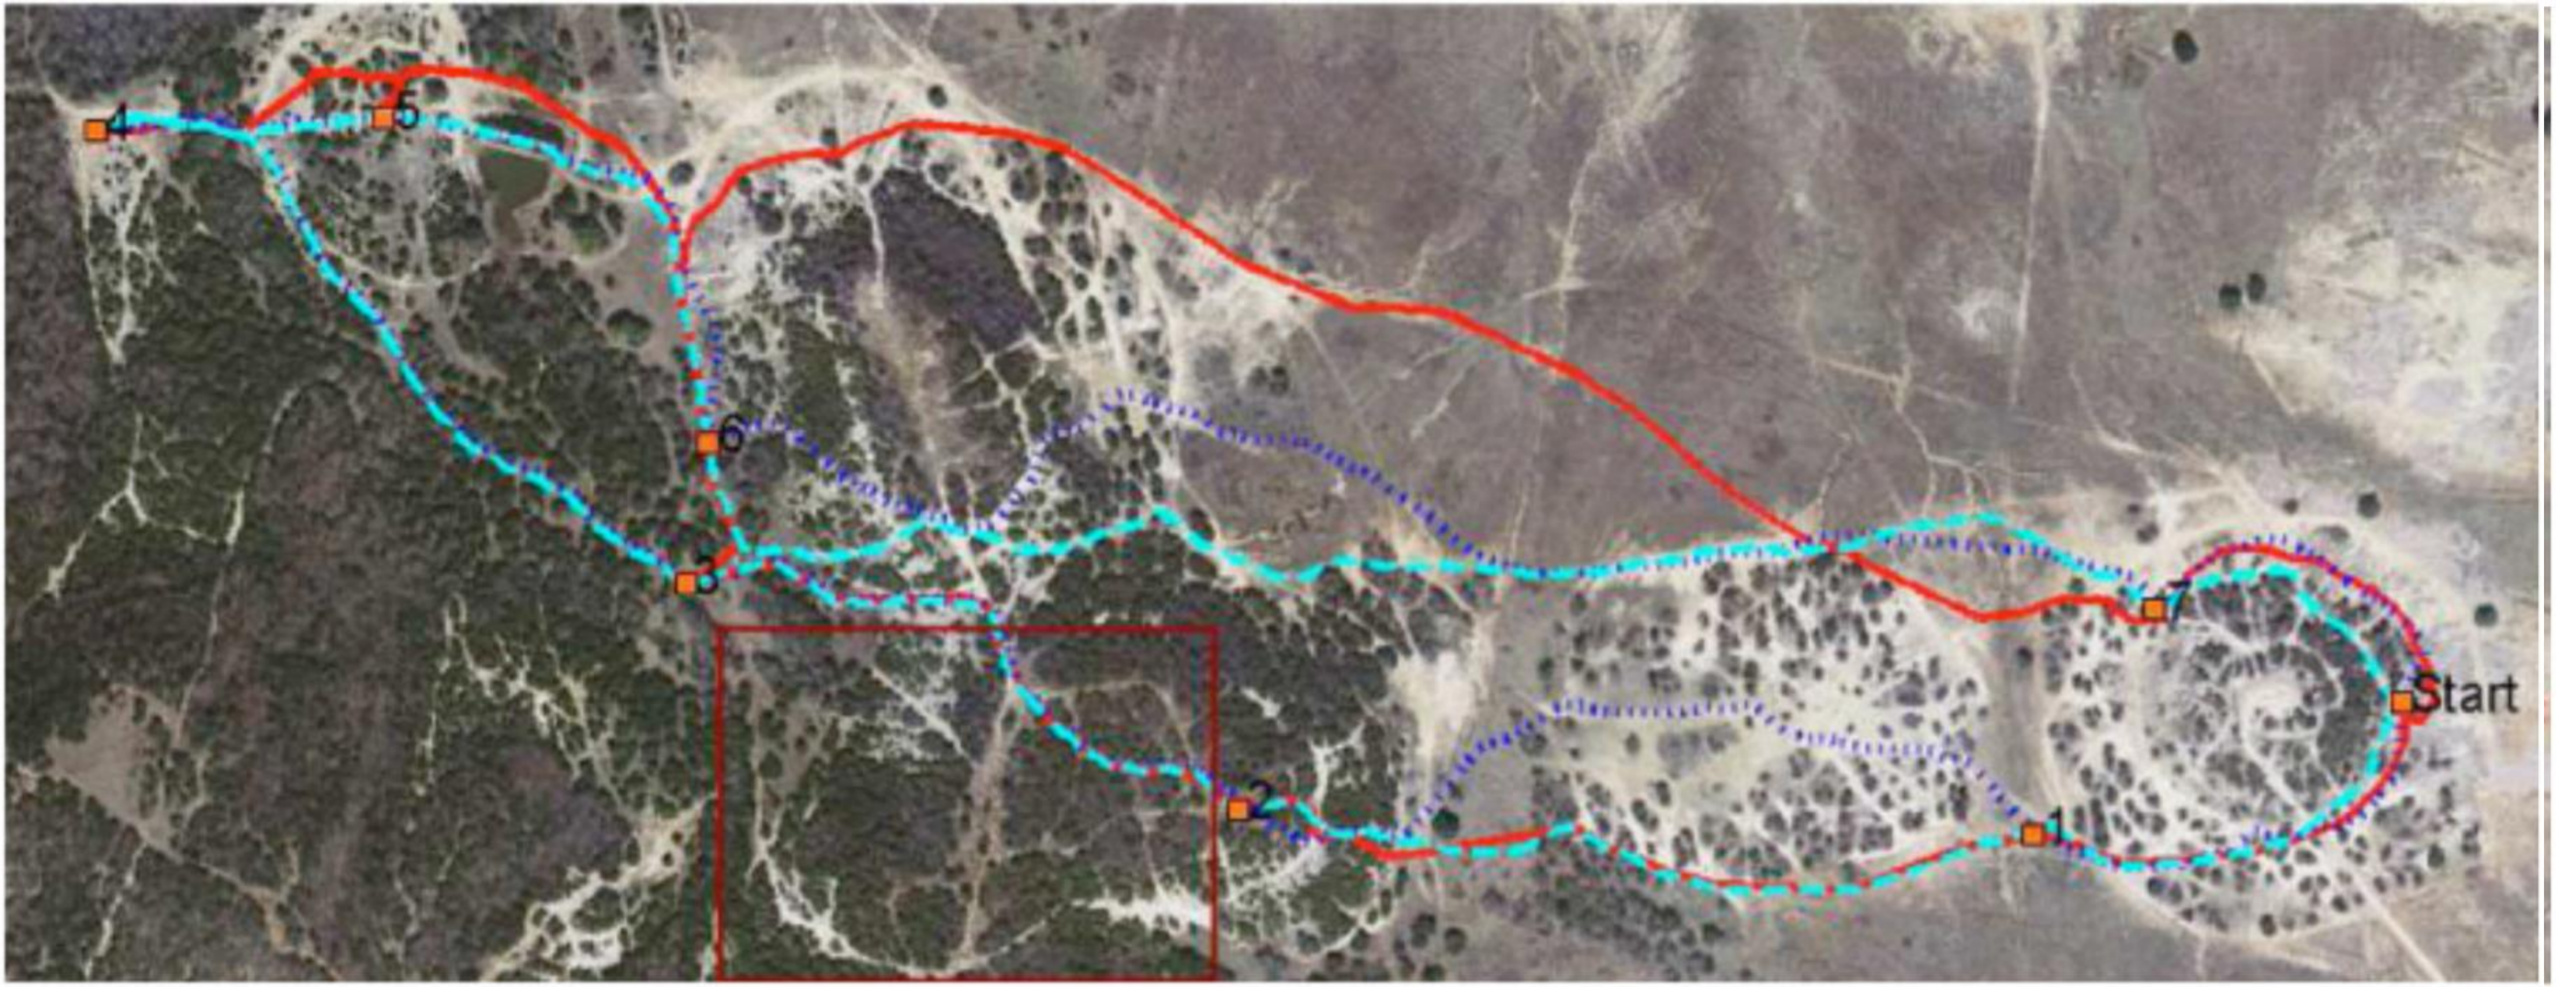

Conclusions for the whole lecture

- In deep learning:
  - Labeled training data is precious
  - Unlabled training data is often abundant
  - Self-supervised methods are used to learn useful representations from unlabeled data, using pretext tasks
- In robotics:
  - Labeled training data is precious
  - Robots can collect large amounts of labeled data cheaply
- Learning from limited supervision is possible!