Learning objects by autonomous exploration and human interaction

Serena Ivaldi

Team LARSEN, INRIA
serena.ivaldi@inria.fr

WS "Open-ended learning for object perception & grasping" @ IROS - 8/11/2019

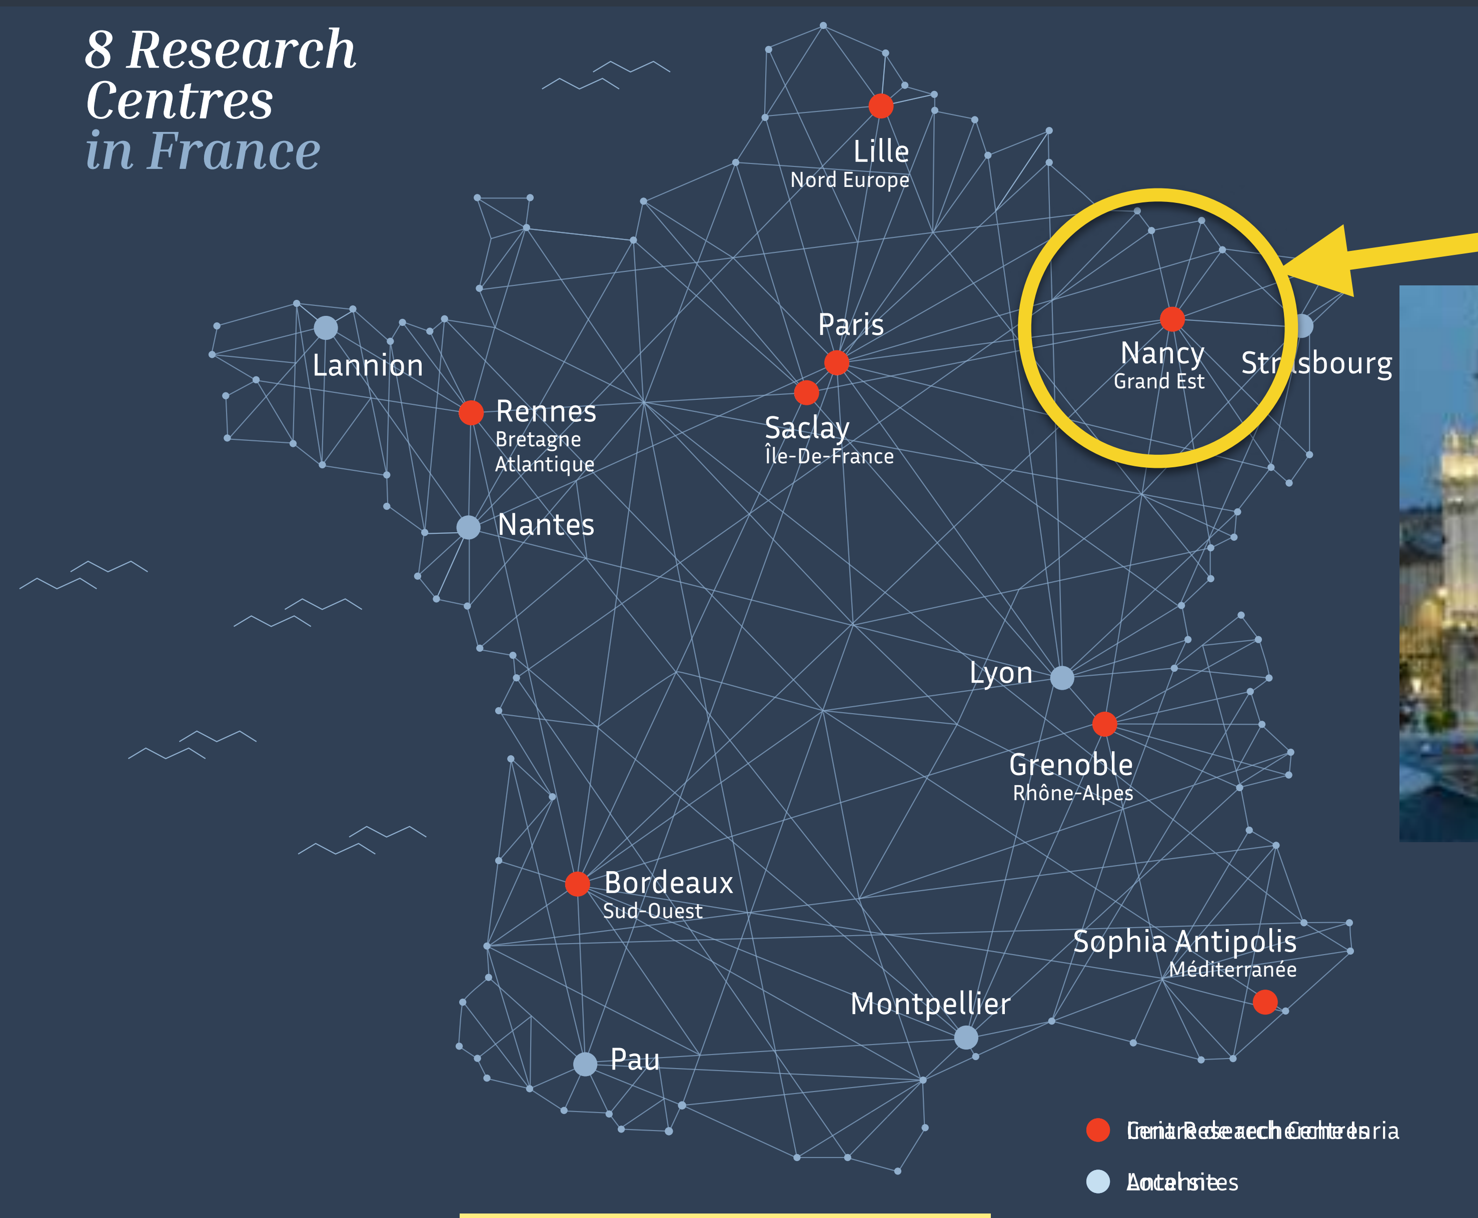

Inria Nancy / Loria
Team Larsen

European Commission | Horizon 2020 European Union funding for Research & Innovation

erc European Research Council

smart apartment

arena

Outline

Learning the visual appearance of objects from scratch - with curiosity-based autonomous exploration and social guidance

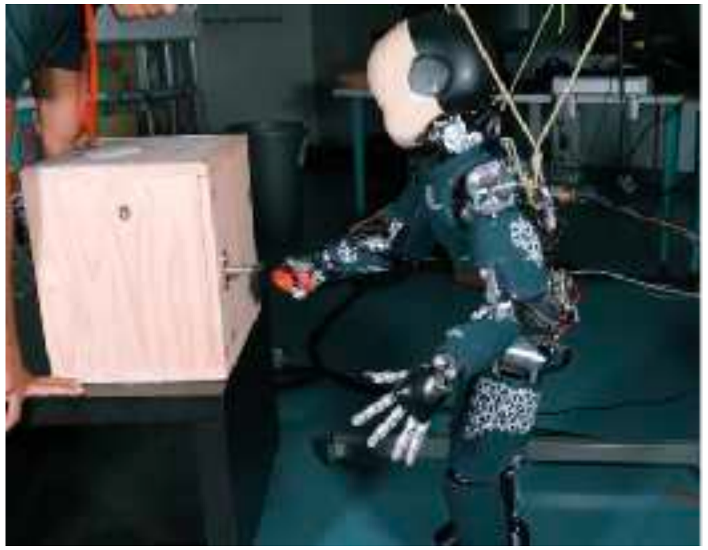

Learning by demonstration whole-body manipulation on humanoids

HEAP project - Incorporating human preferences to improve grasping of irregular objects in a heap

Outline

Learning the visual appearance of objects from scratch - with curiosity-based autonomous exploration and social guidance

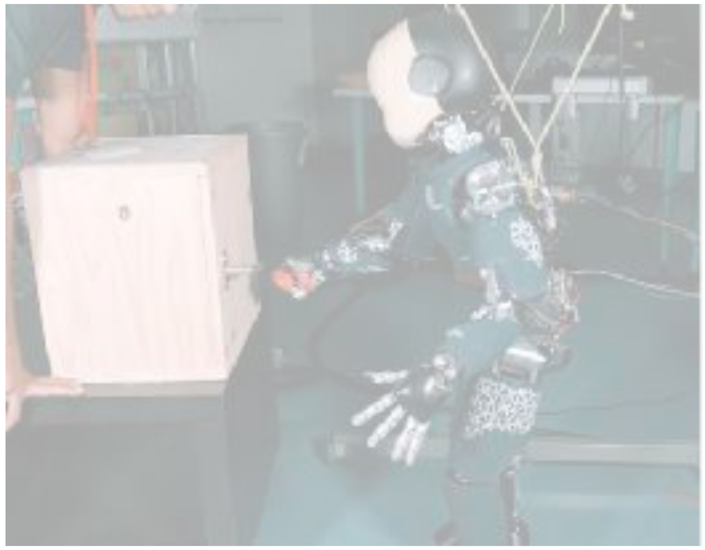

Learning by demonstration whole-body manipulation on humanoids

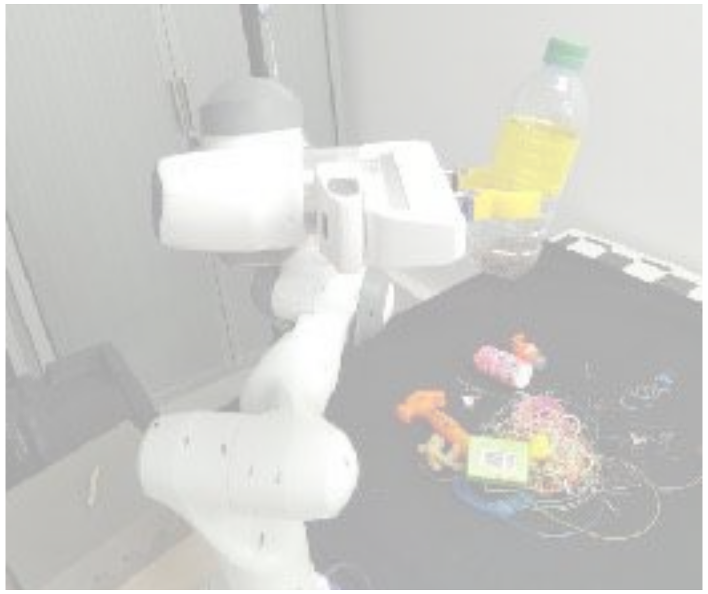

HEAP project - Incorporating human preferences to improve grasping of irregular objects in a heap

Learning like a child

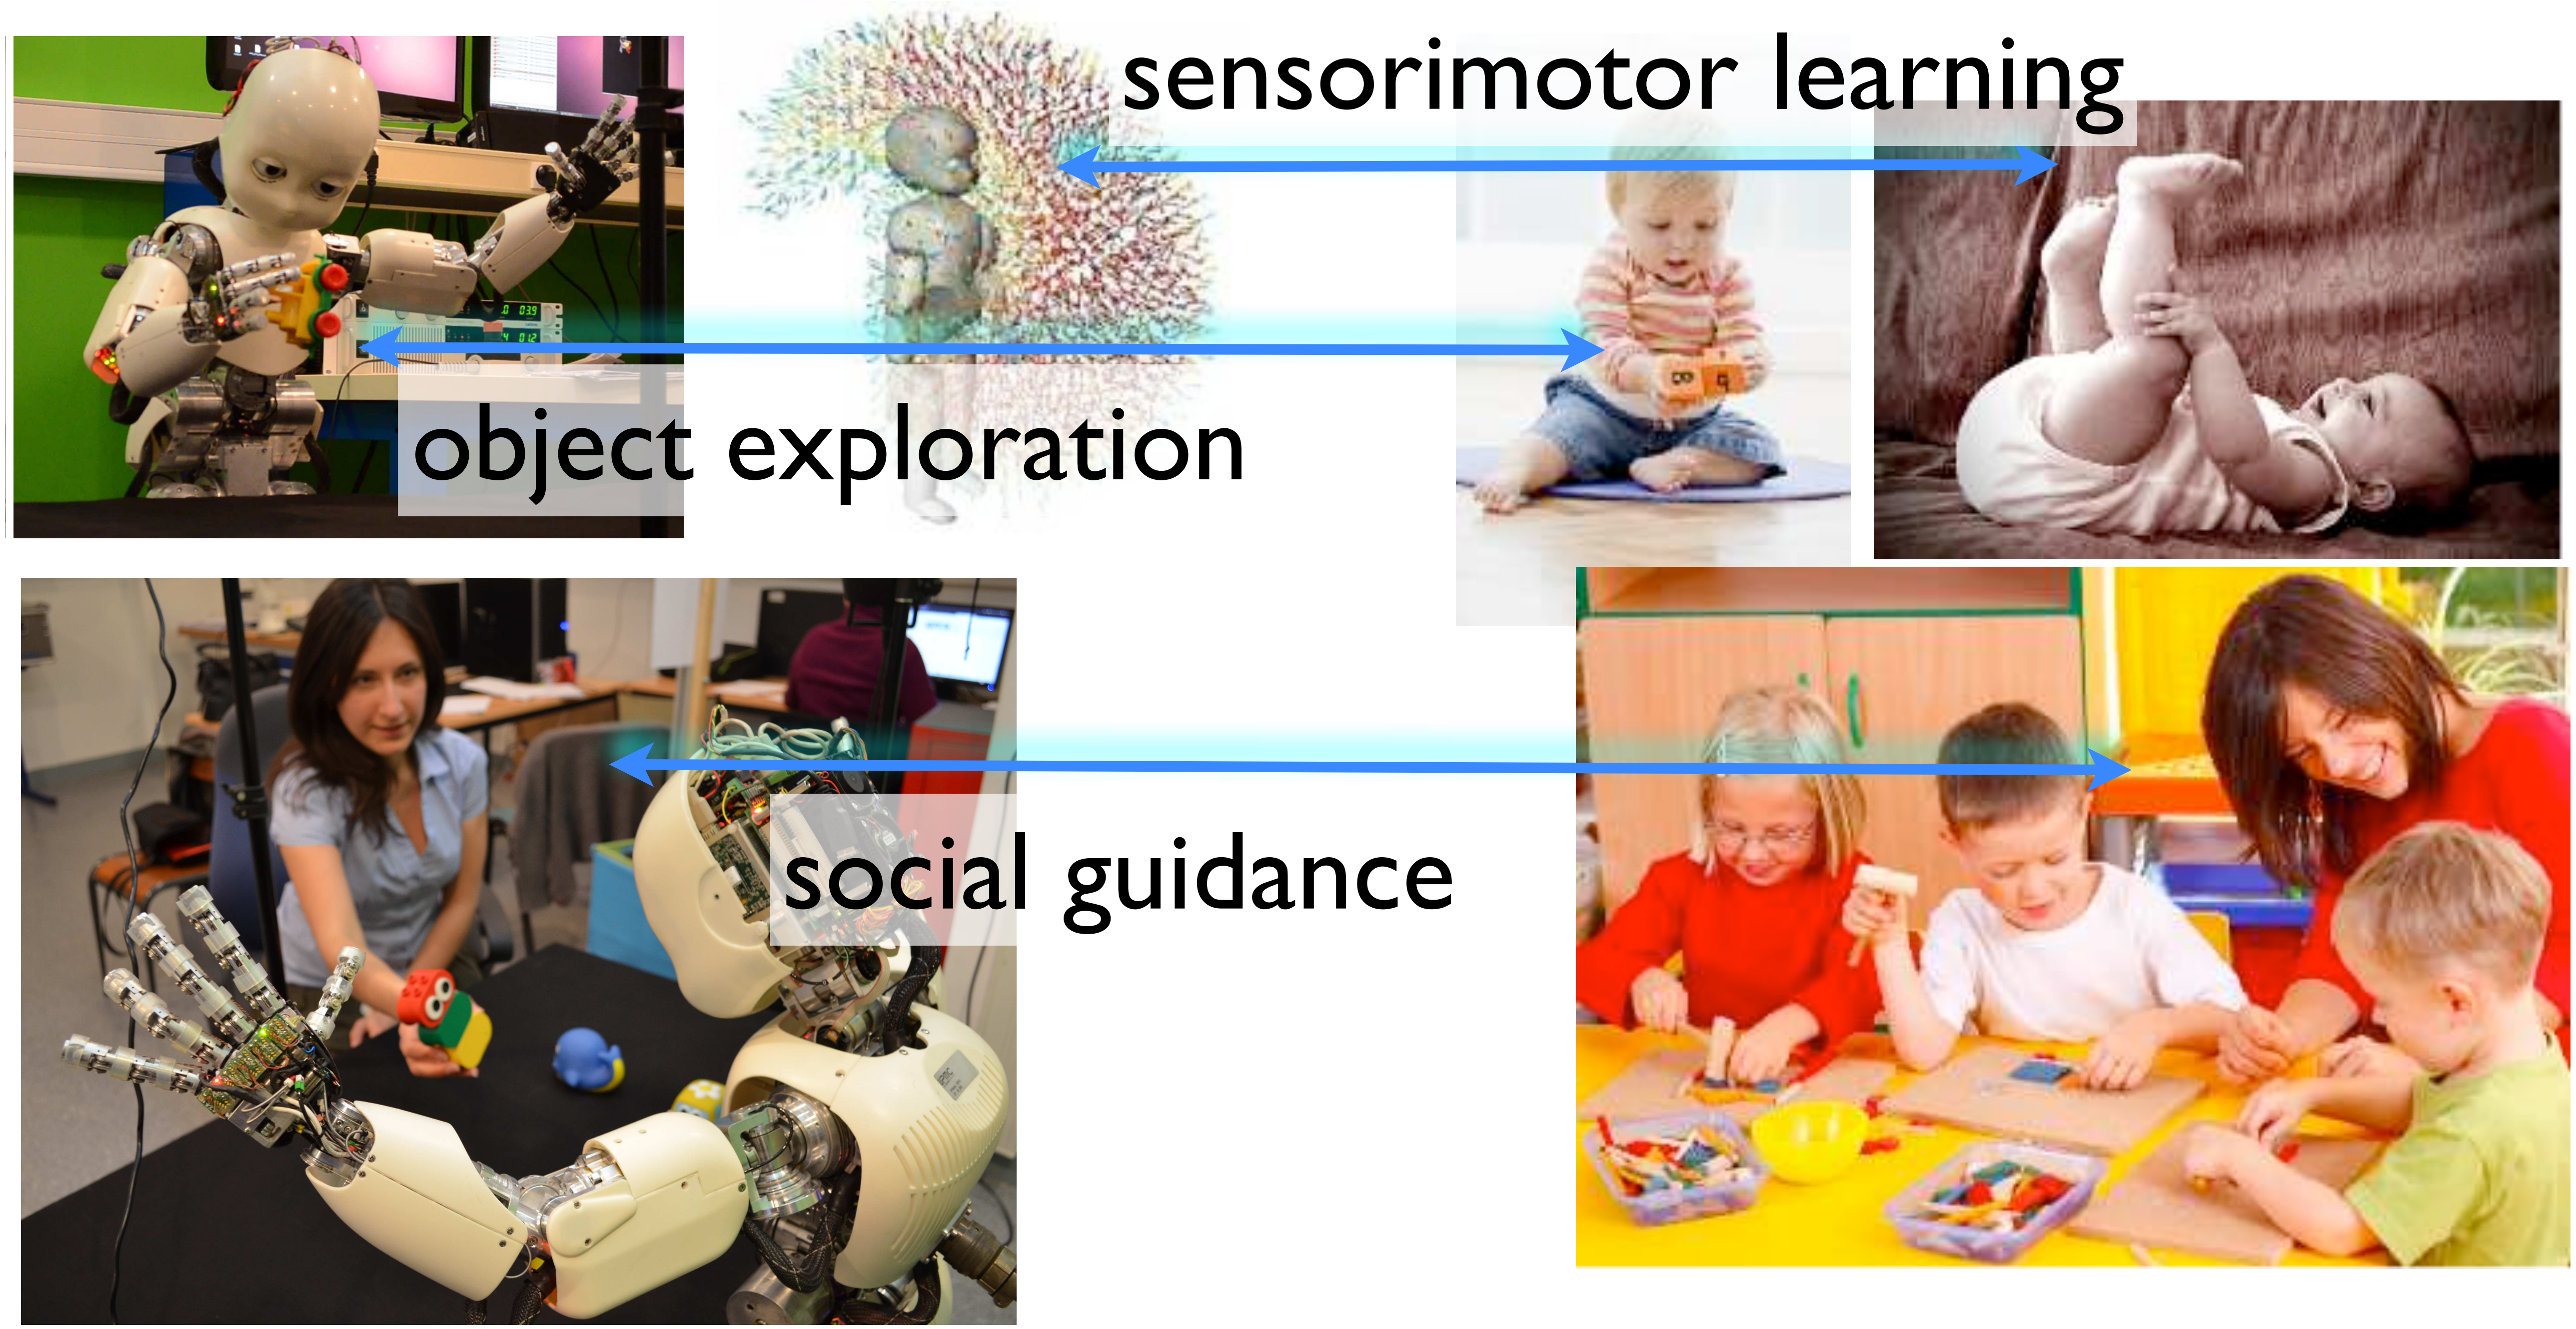

Object learning through interaction

autonomous learning of visuomotor models

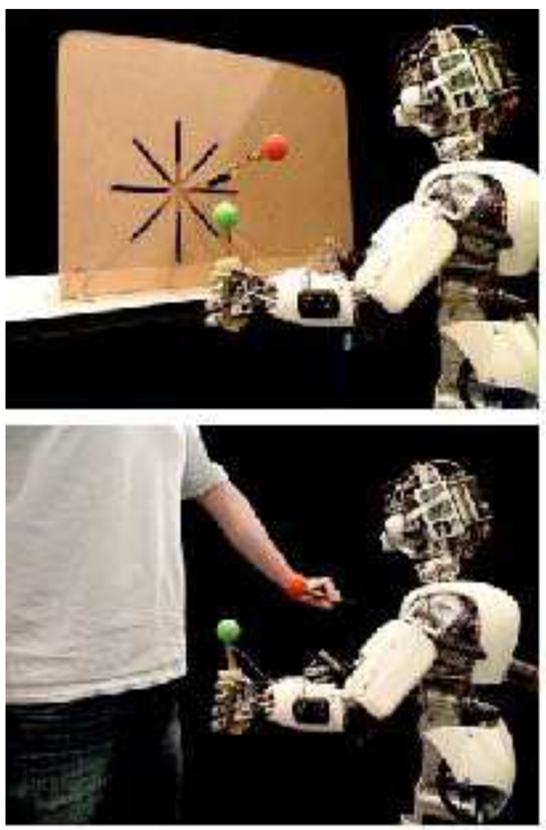

powerful learning tools

IROS 2012

object learning through active exploration

intrinsic motivation + aut. exploration + social guidance

IEEE TAMD 2014

object learning through joint attention

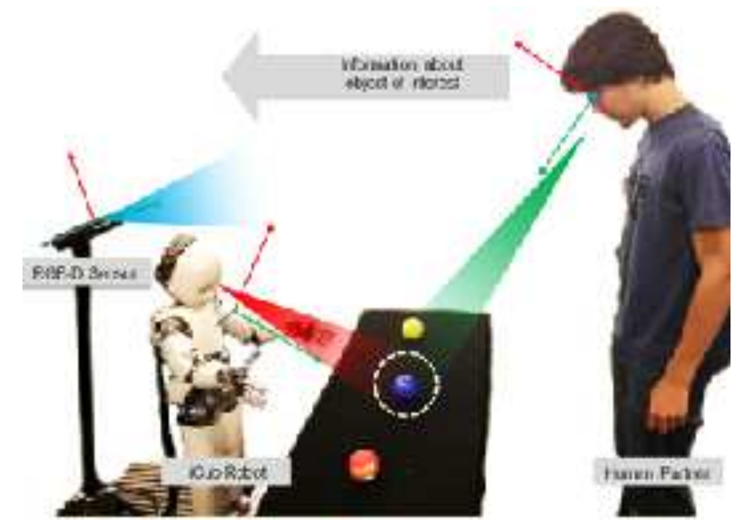

social guidance

Front. Neurorob. 2013

Multimodality for object learning

exploration and interaction
(better models with categories)

active exploration
(manipulation, better models)

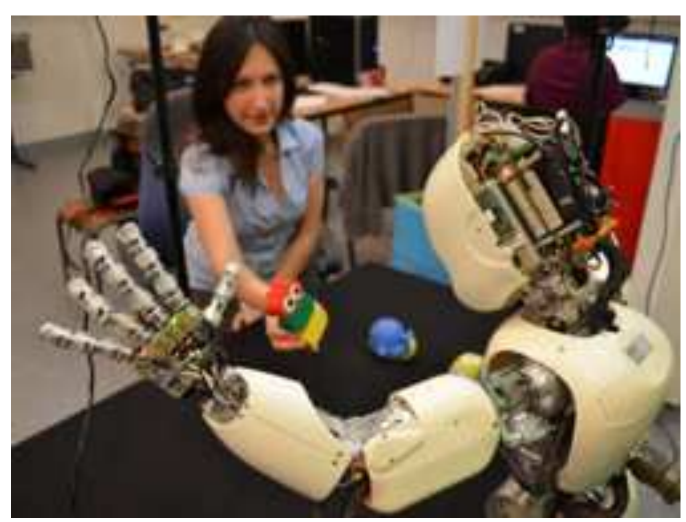

observation
(pure vision: models and entities)

Multimodality for object learning

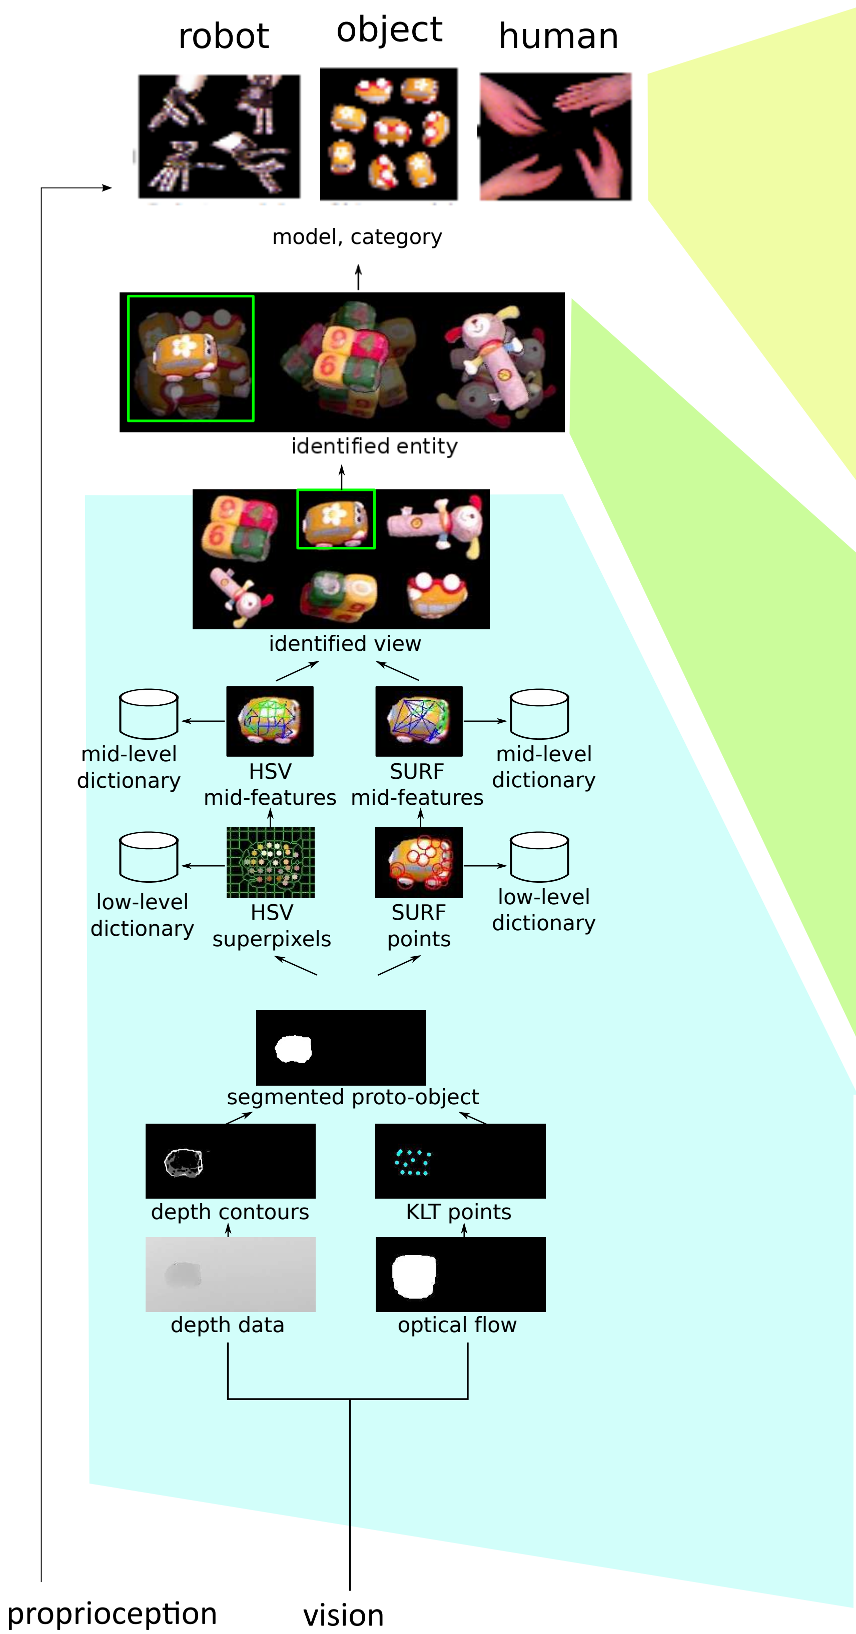

exploration and interaction
(better models with categories)

active exploration
(manipulation, better models)

observation
(pure vision: models and entities)

Multimodality for object learning

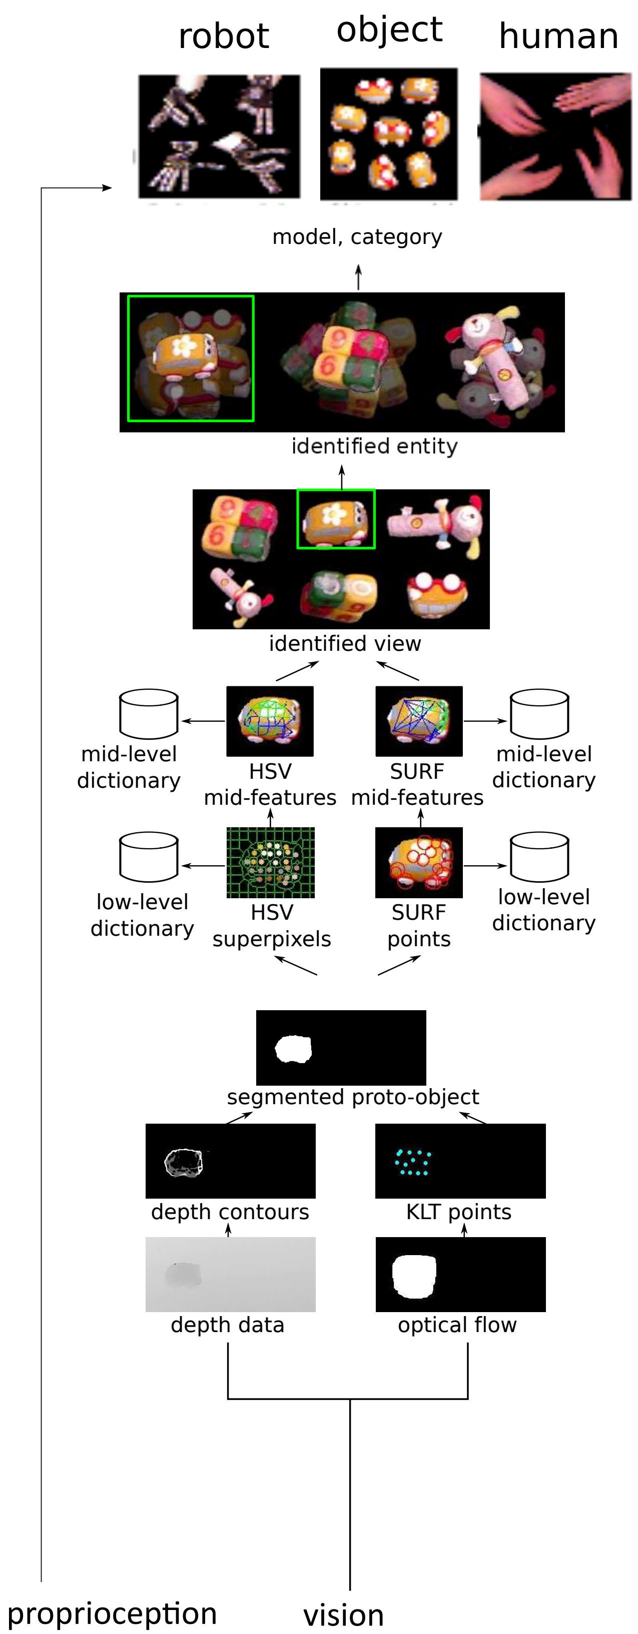

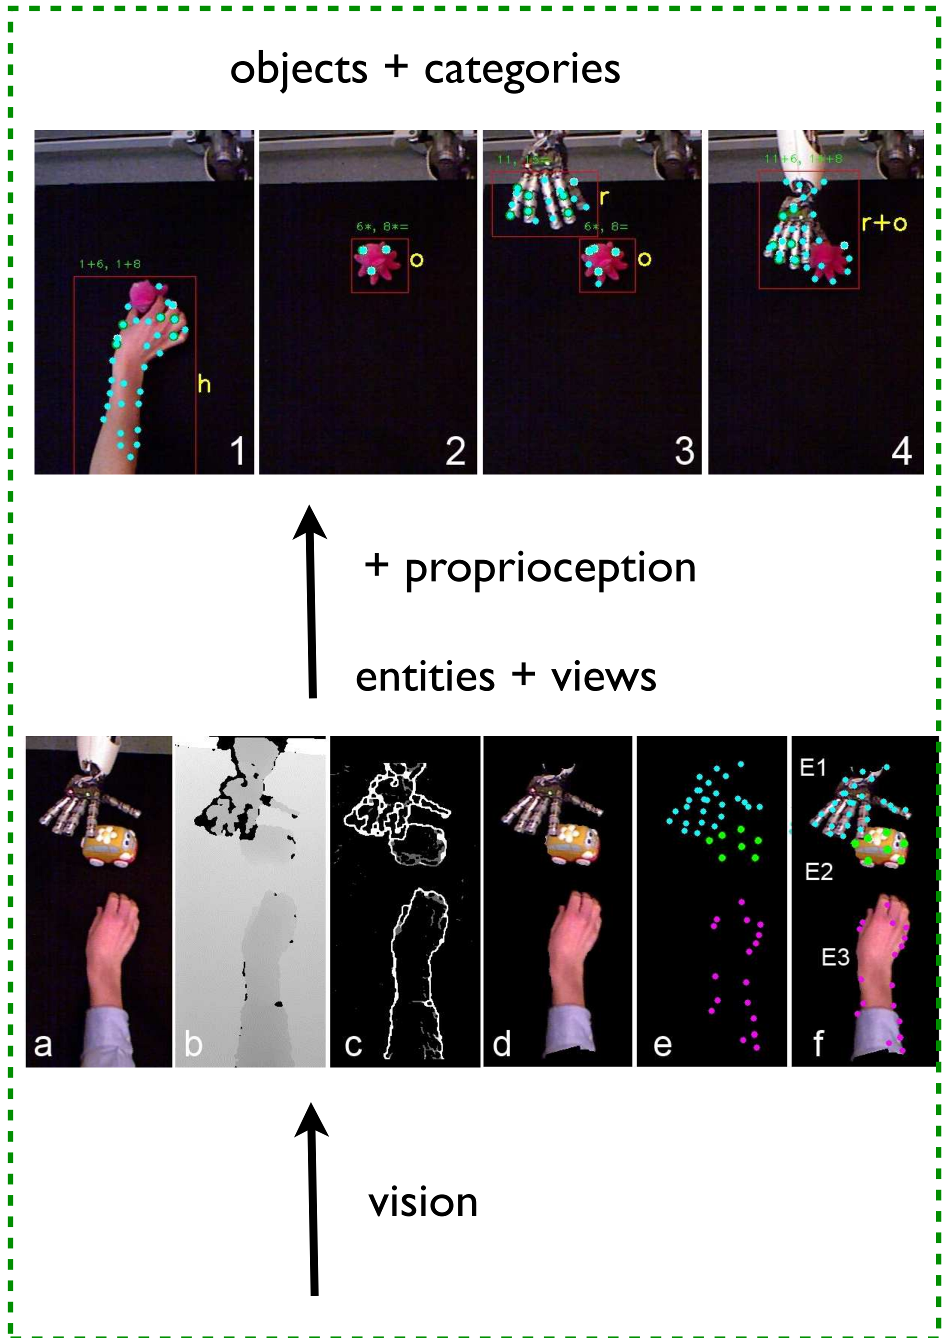

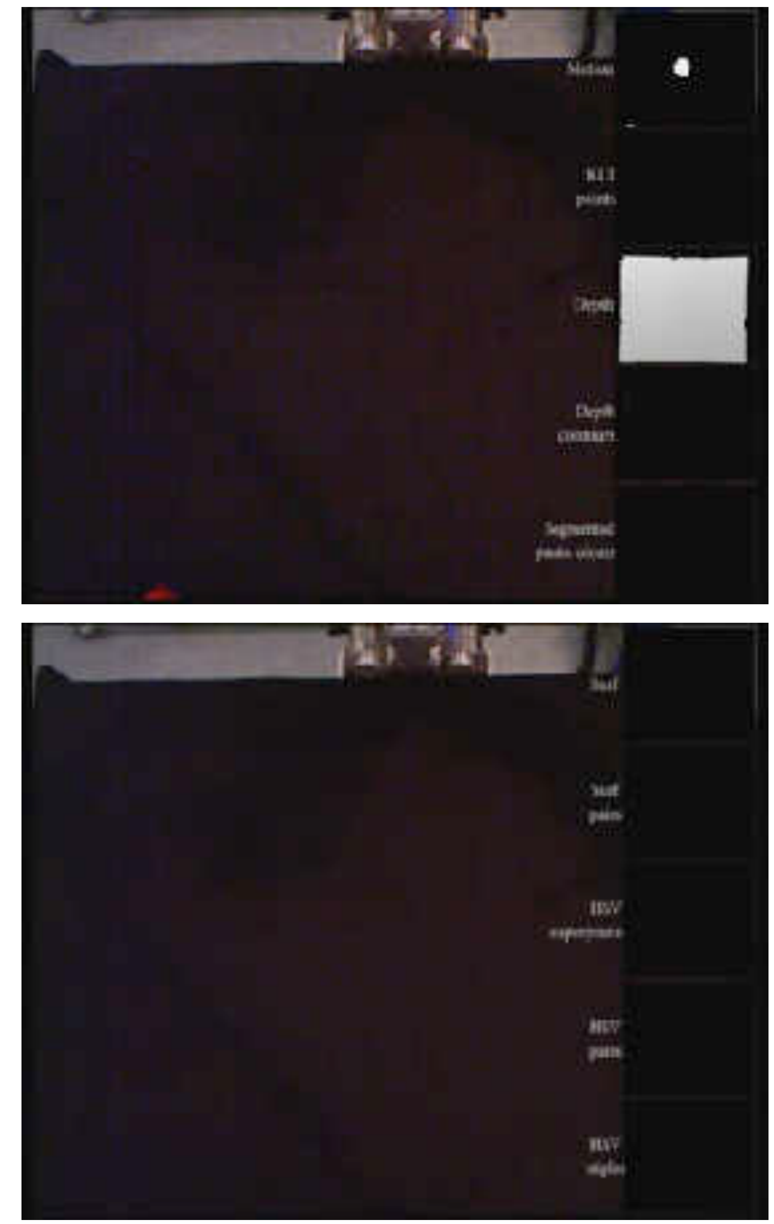

Observation alone is not enough

vision

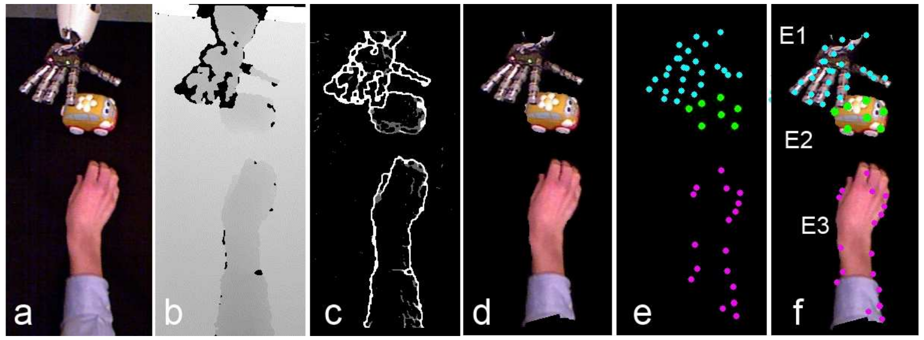

entities + collected views

model

The robot learns the objects demonstrated by the human.

The robot has not yet learnt to identify its body, hence all entities are labeled by an "unknown" category.

Lyubova, Ivaldi, Filliat (2016) From passive to interactive object learning and recognition through self-identification on a humanoid robot. Autonomous Robots, 40(1):33-57.

Acquire better models through action

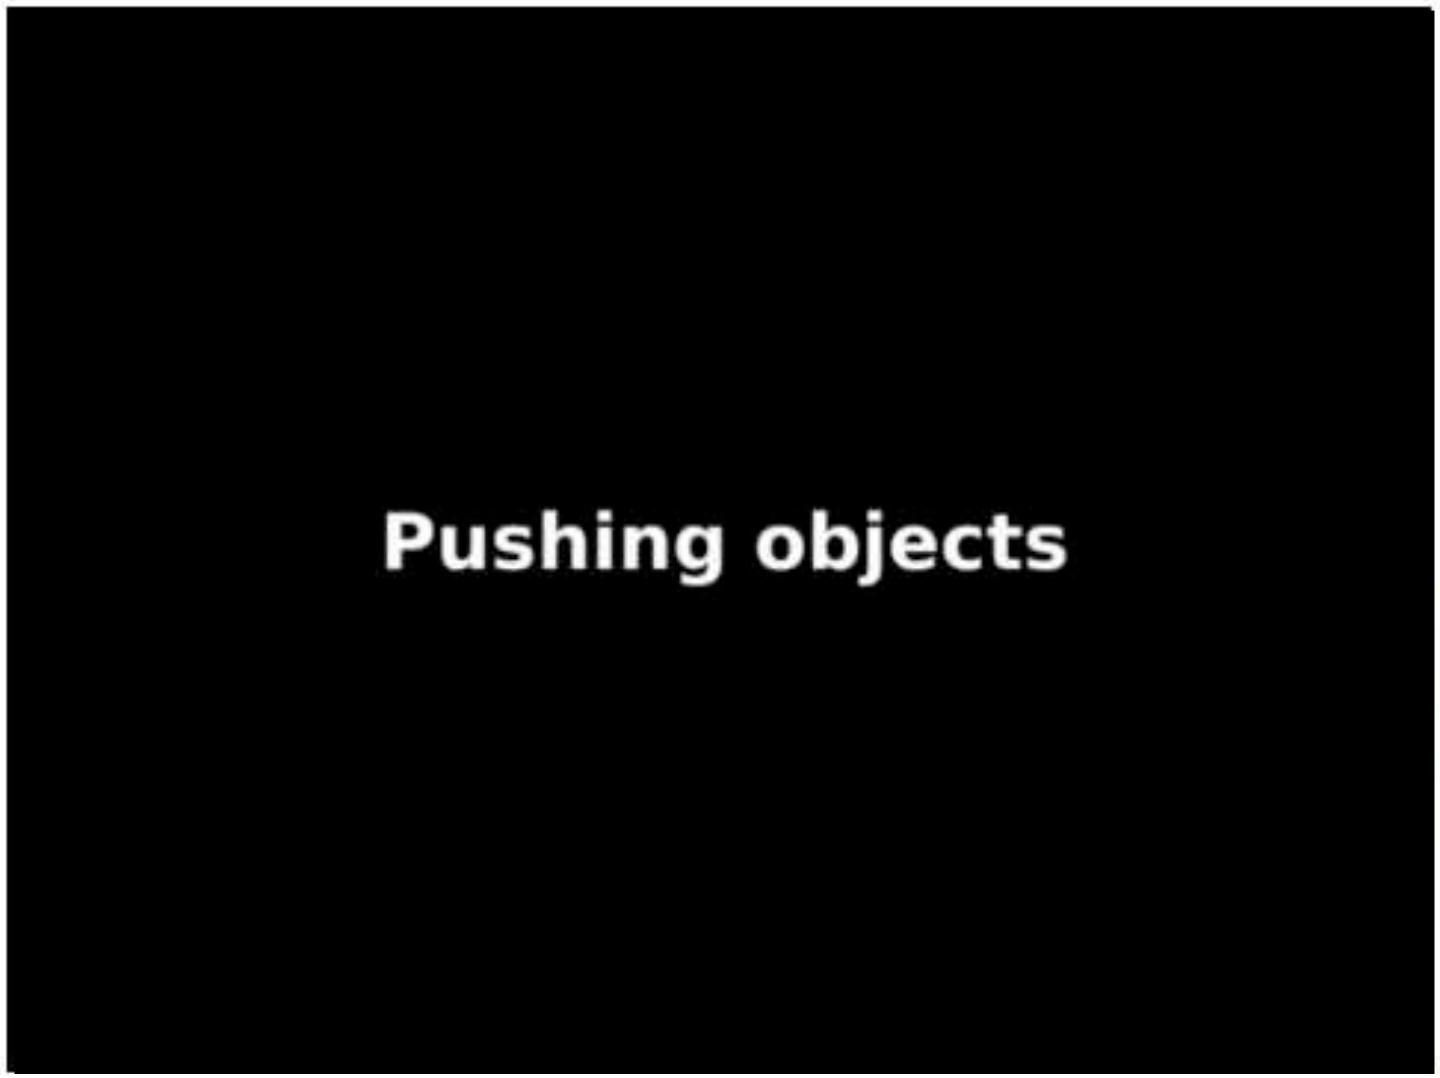

side push, pick & place

grasp
lift
throw

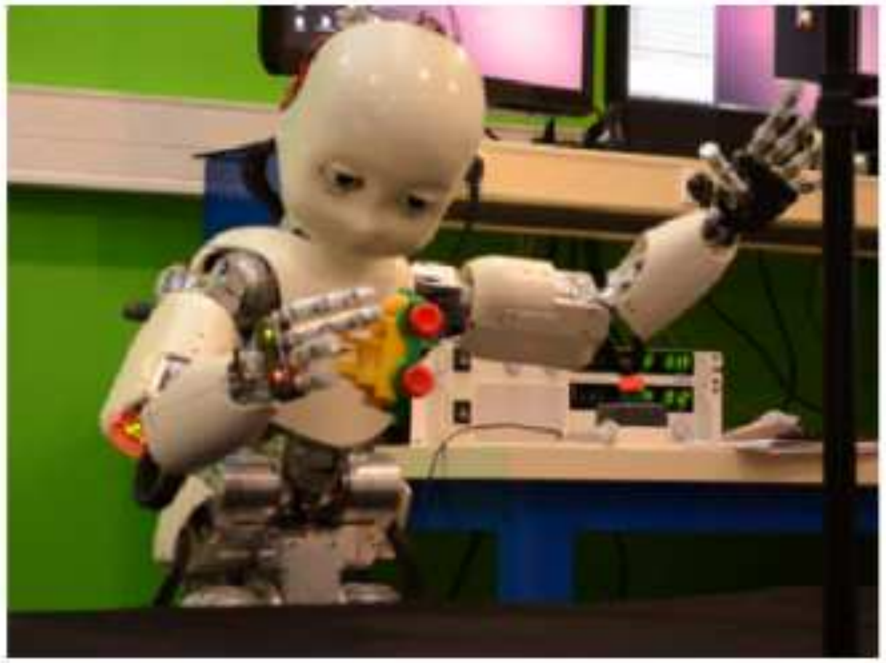

grasp
lift
rotate
put

Ivaldi, Nguyen, Lyubova, Droniou, Padois, Filliat, Oudeyer, Sigaud (2014) Object learning through active exploration. IEEE Transactions on Autonomous Mental Development.

Active exploration of objects

action does not change the view

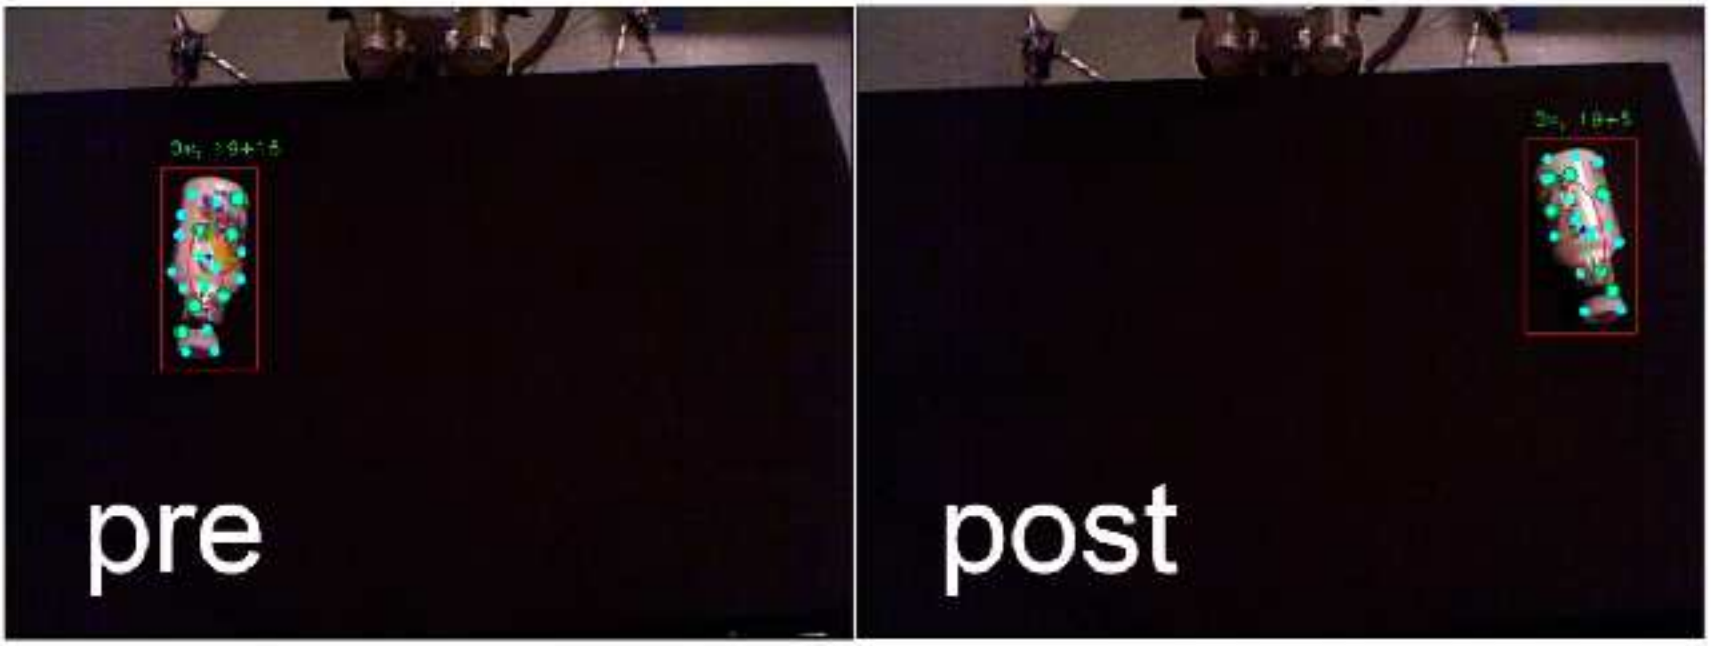

action provokes a new view

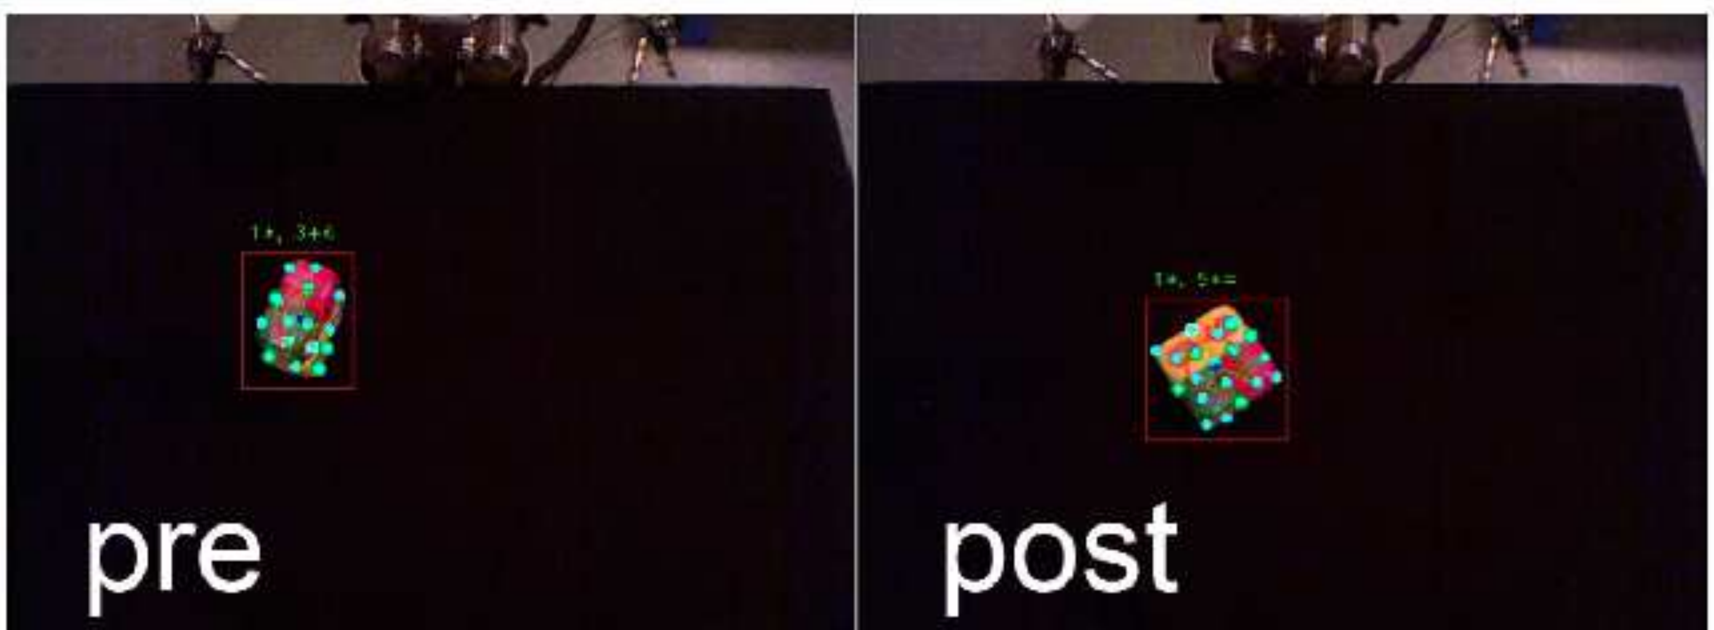

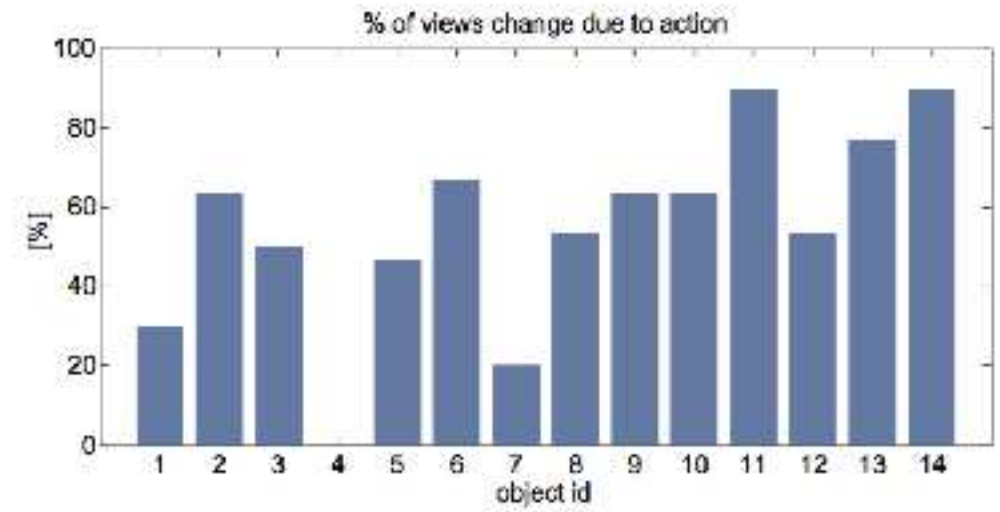

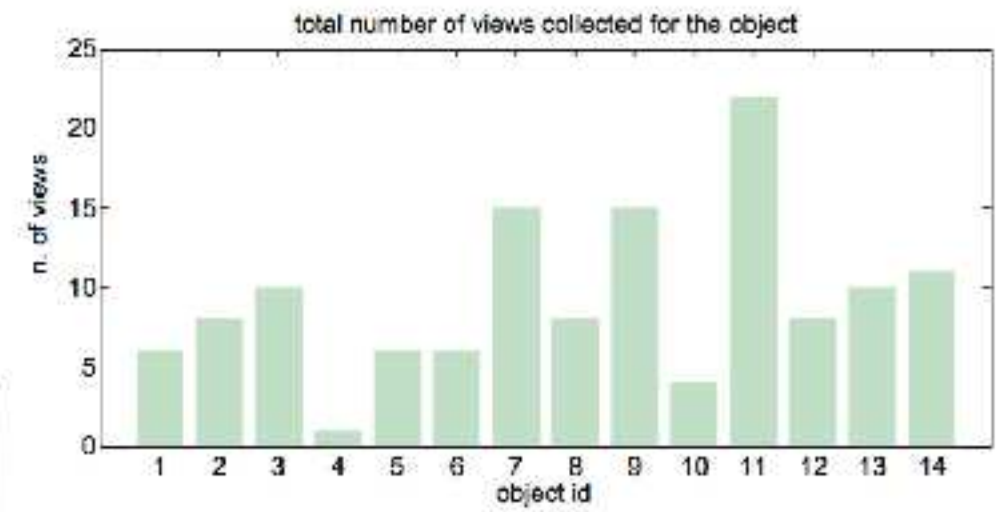

Active exploration & social guidance

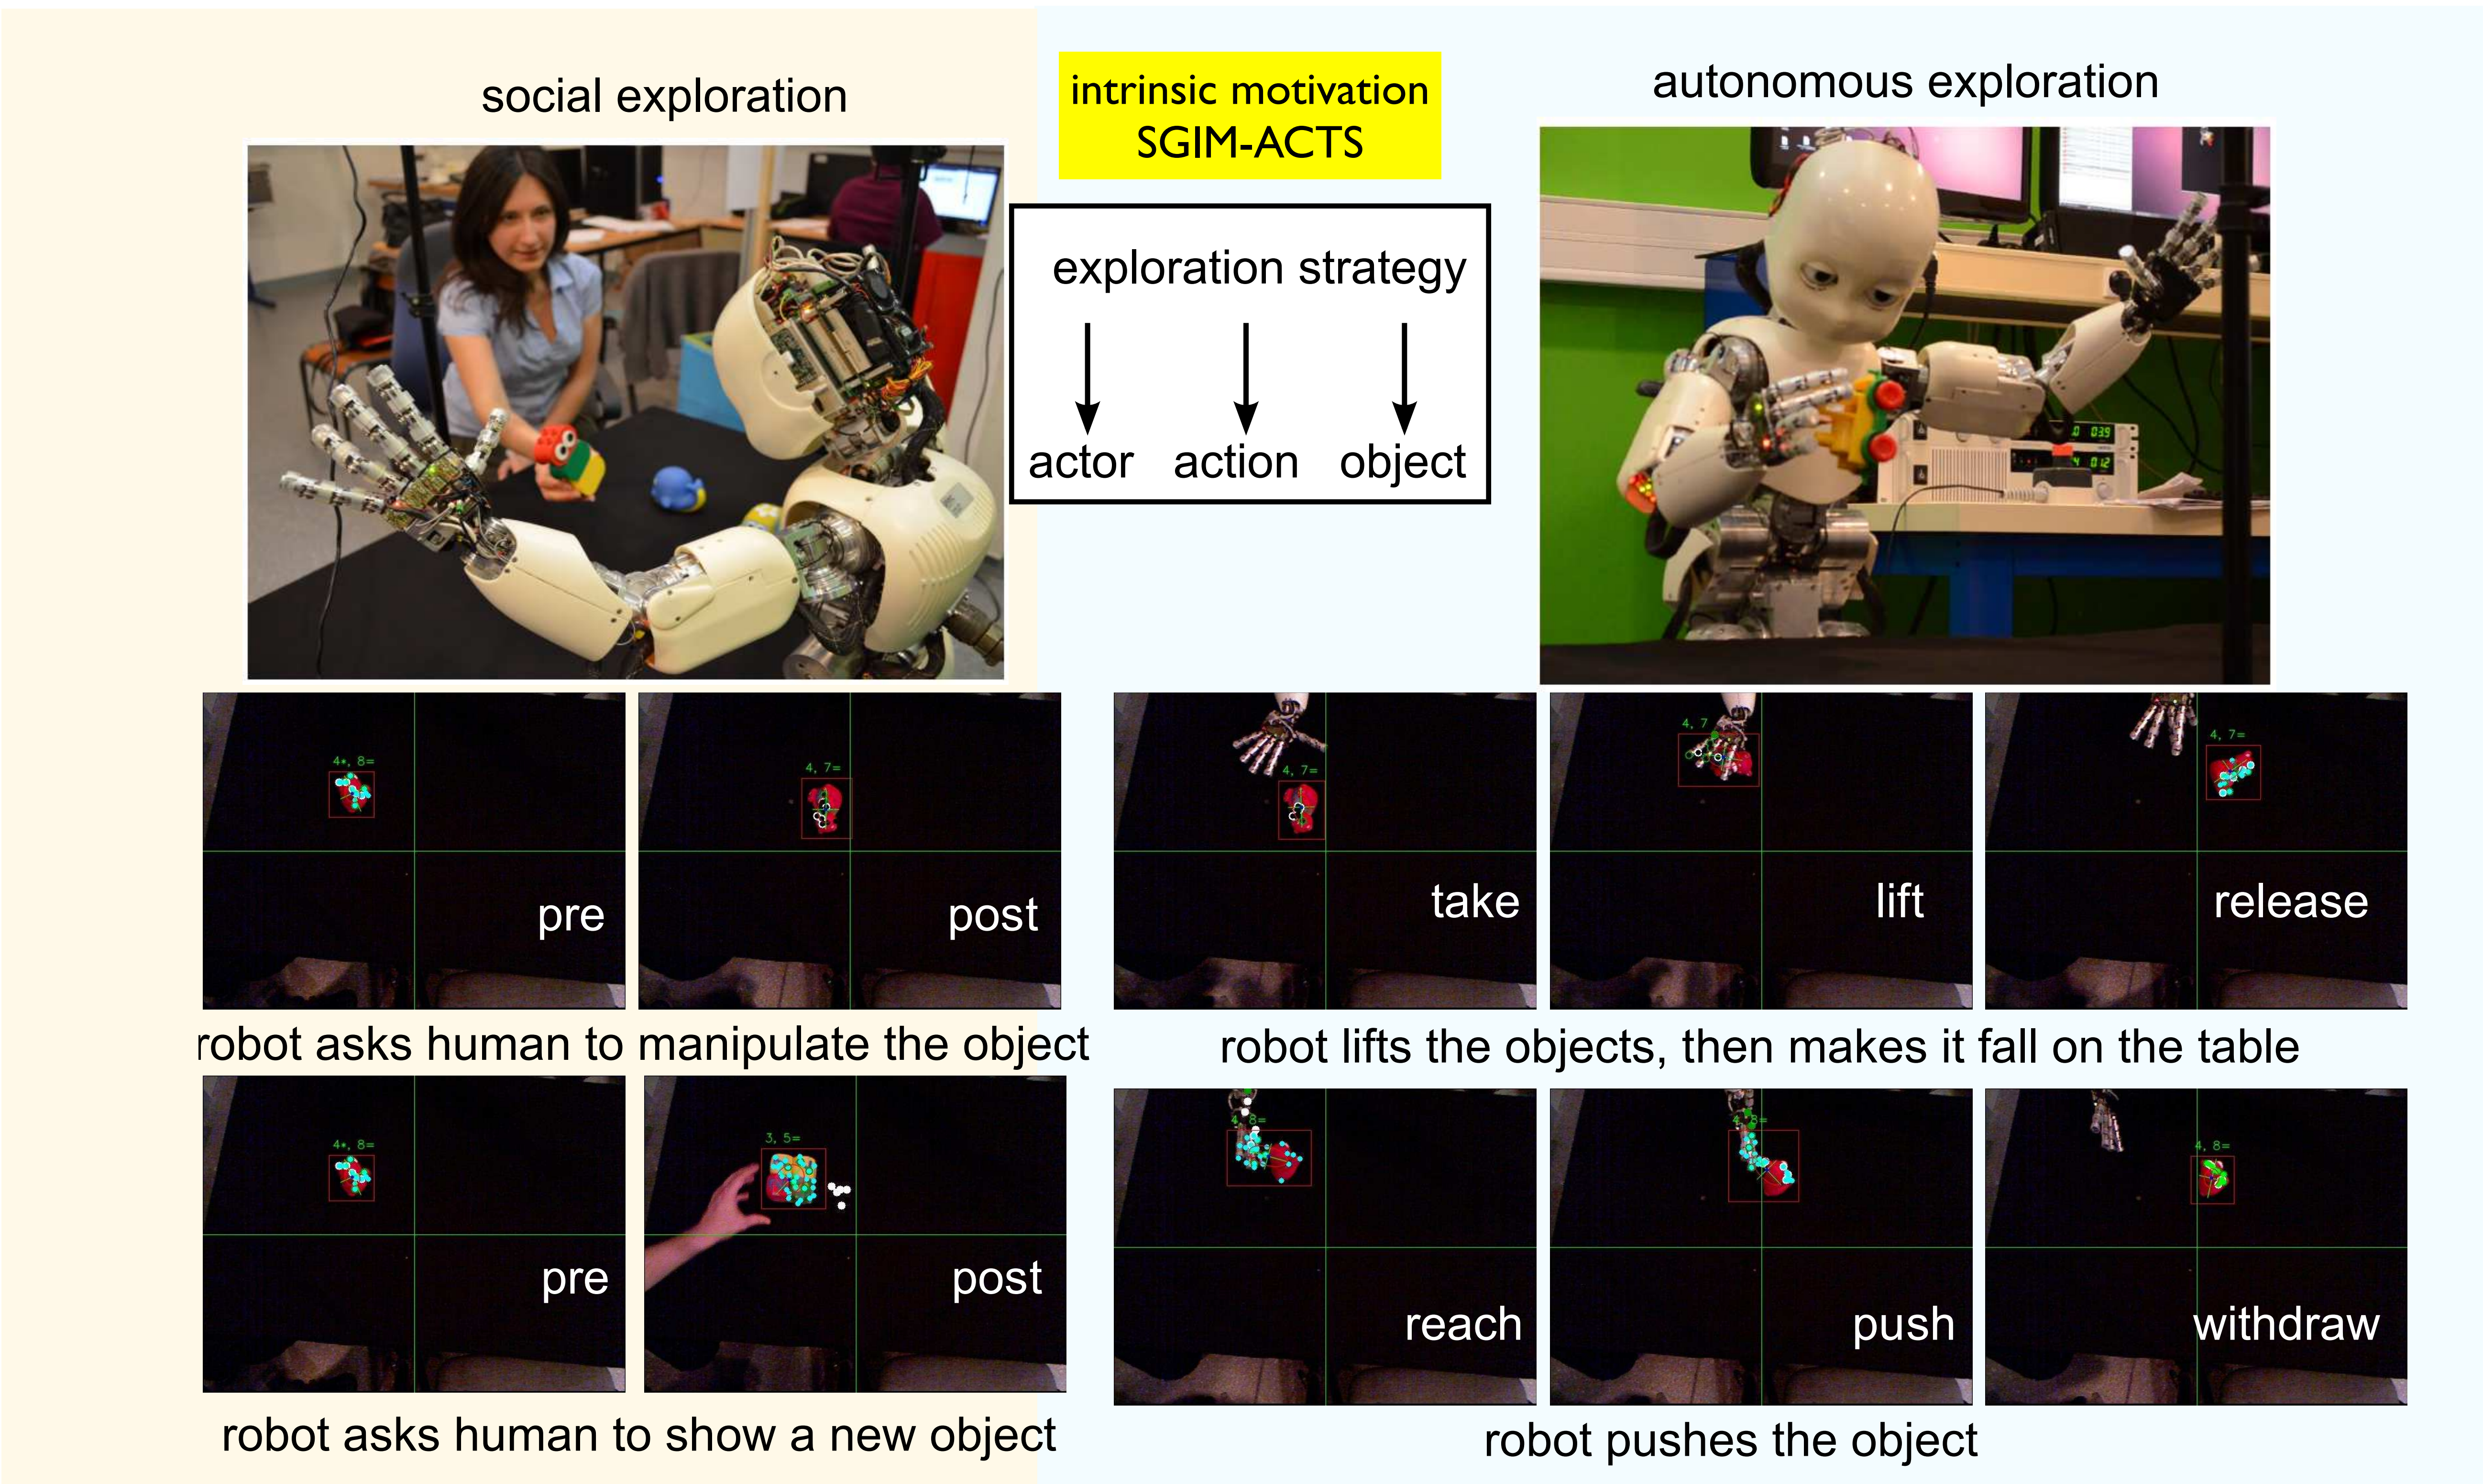

Active exploration & social guidance

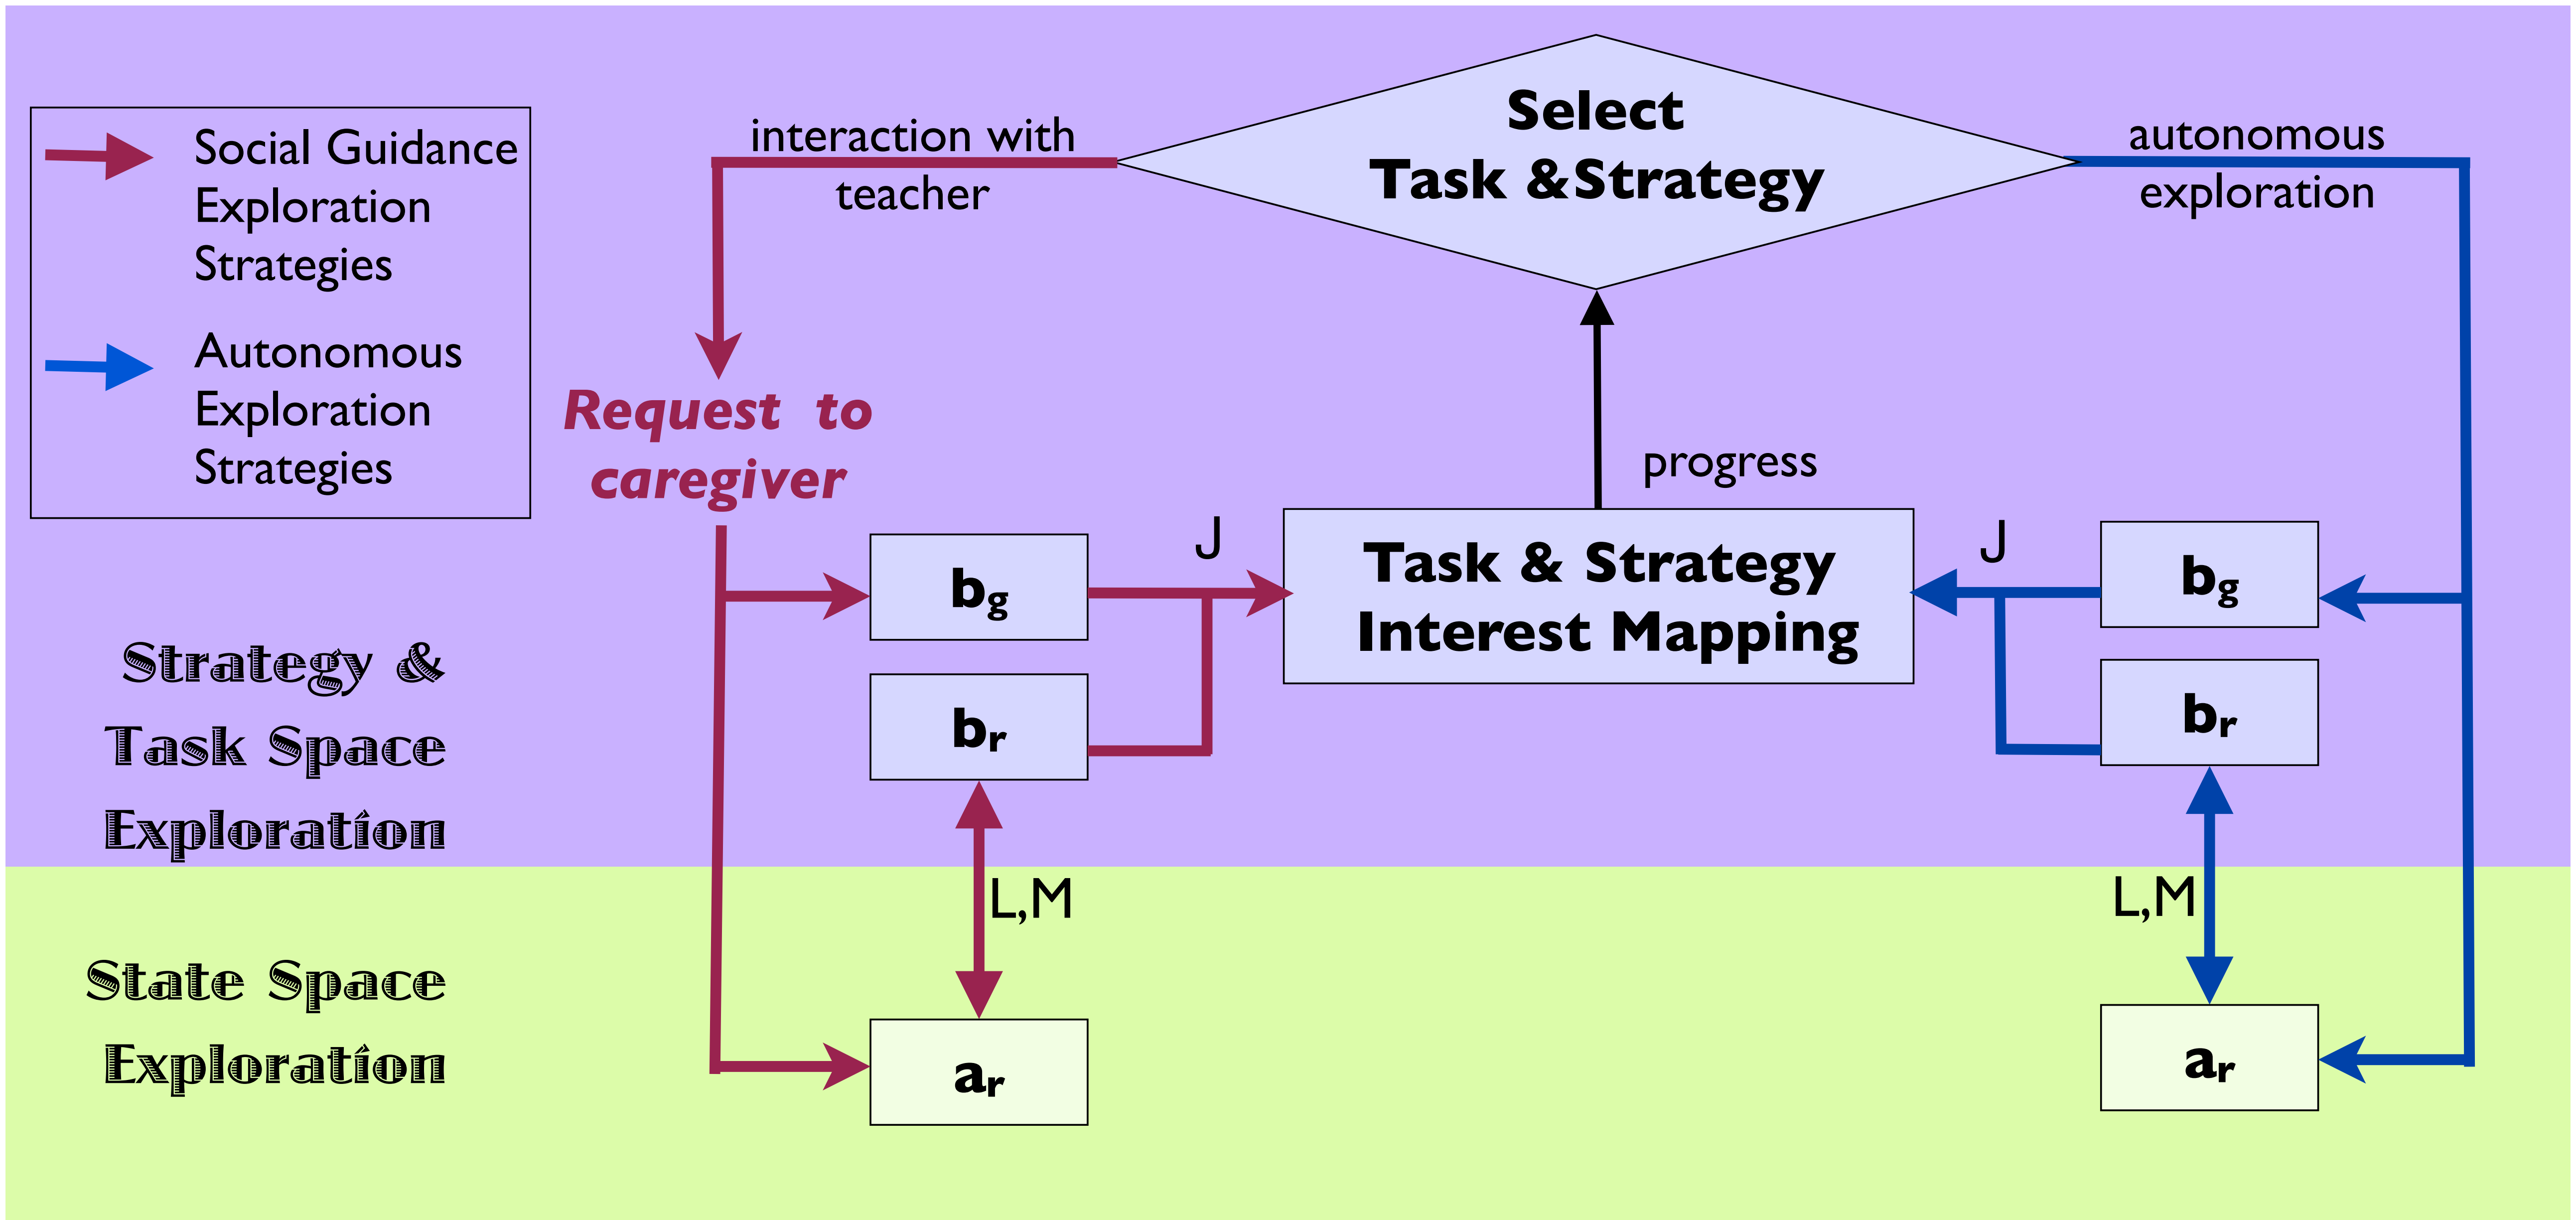

Ivaldi, Nguyen, Lyubova, Droniou, Padois, Filliat, Oudeyer, Sigaud (2014) Object learning through active exploration. IEEE Transactions on Autonomous Mental Development.

Curiosity-driven exploration of objects

- Focusing on the objects that are not yet learned
- choosing the appropriate action for each object

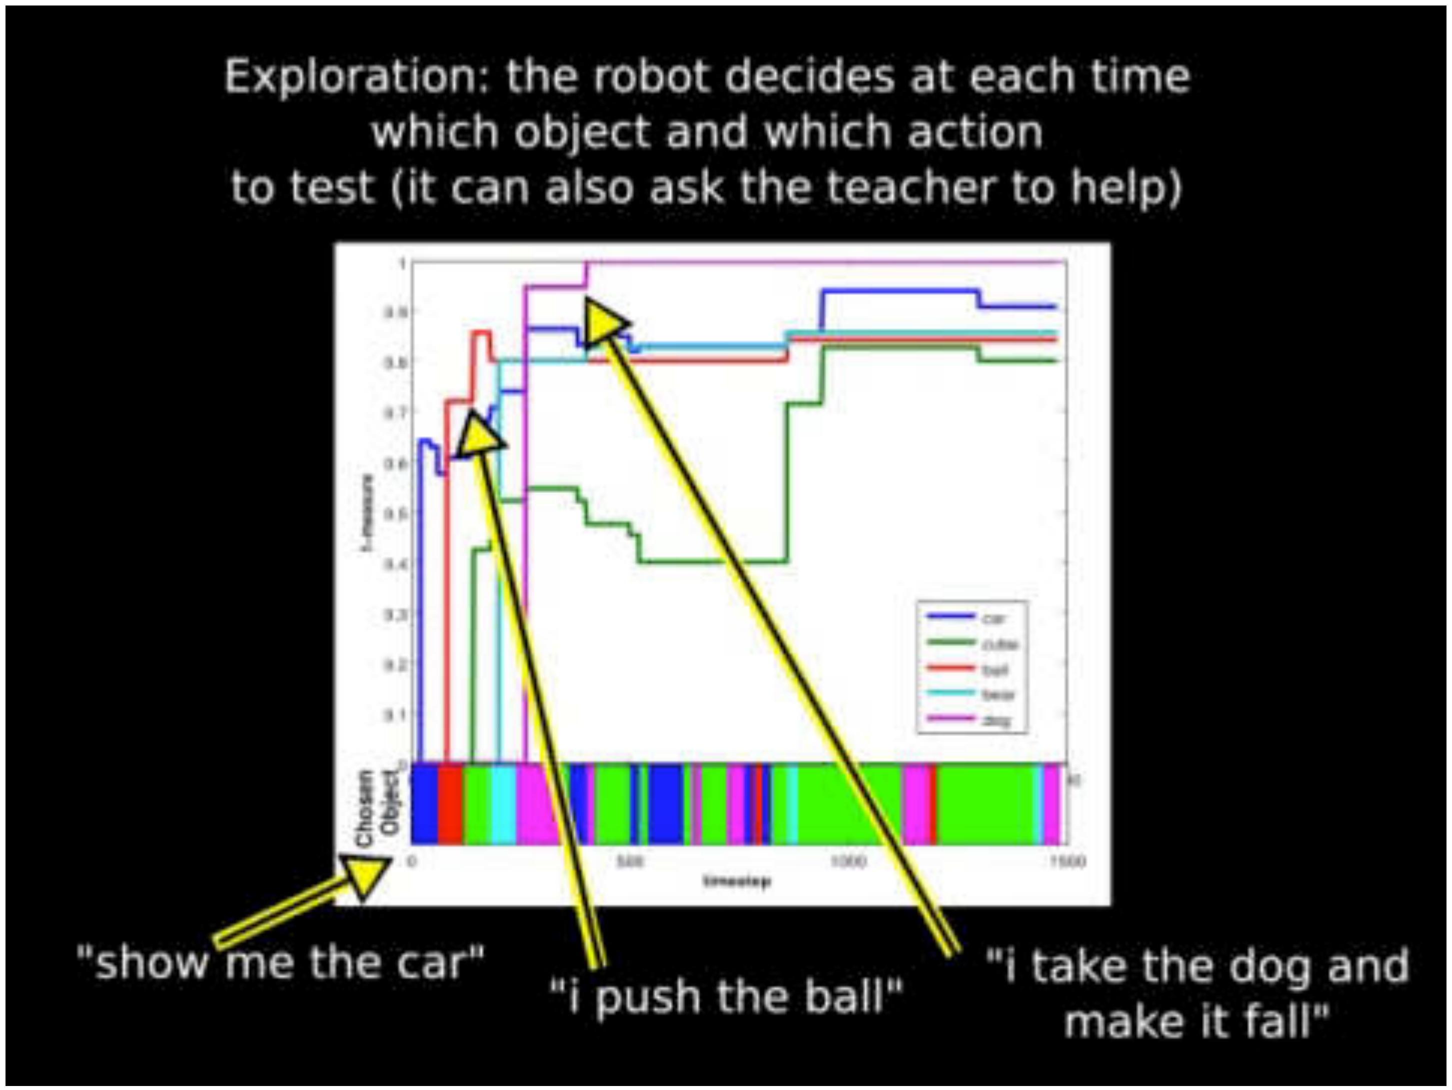

ball → yellow car → red bear → …

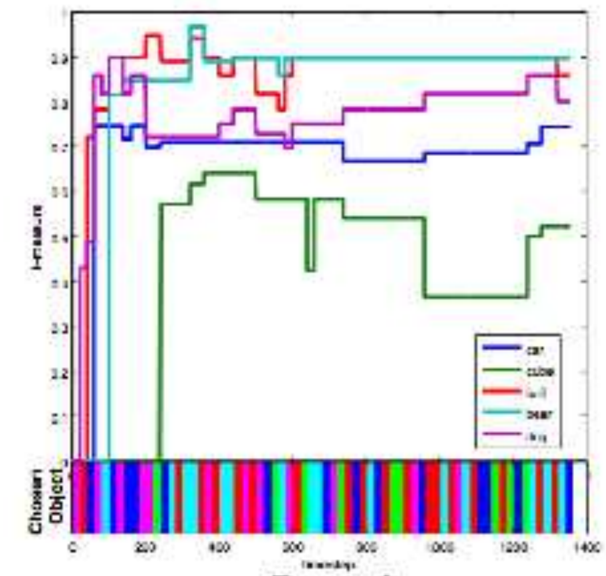

Random

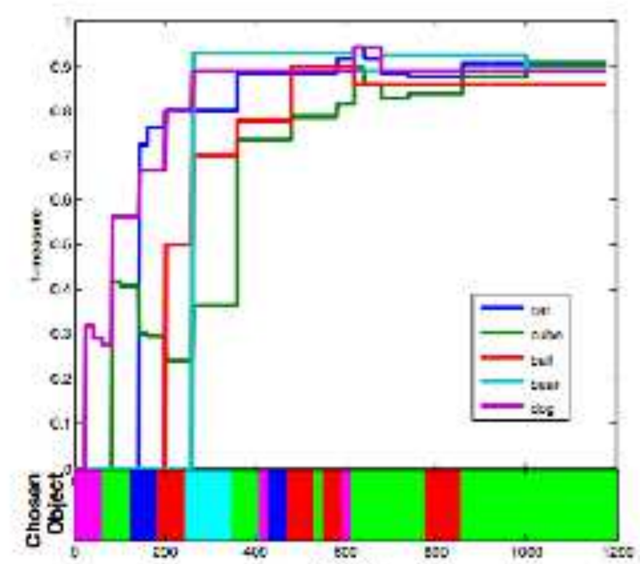

Intrinsic motivation

Ivaldi, Nguyen, Lyubova, Droniou, Padois, Filliat, Oudeyer, Sigaud (2014) Object learning through active exploration. IEEE Transactions on Autonomous Mental Development.

Active exploration & social guidance

- with the "bad"teacher the robot takes and throws more the objects (41% vs 24%)

- the "good" teacher has a catalysing effect: learning process is 25% faster

- with intrinsic motivation, the robot spends most of its time in learning the cube (54% and 51% with bad and good teacher respectively)

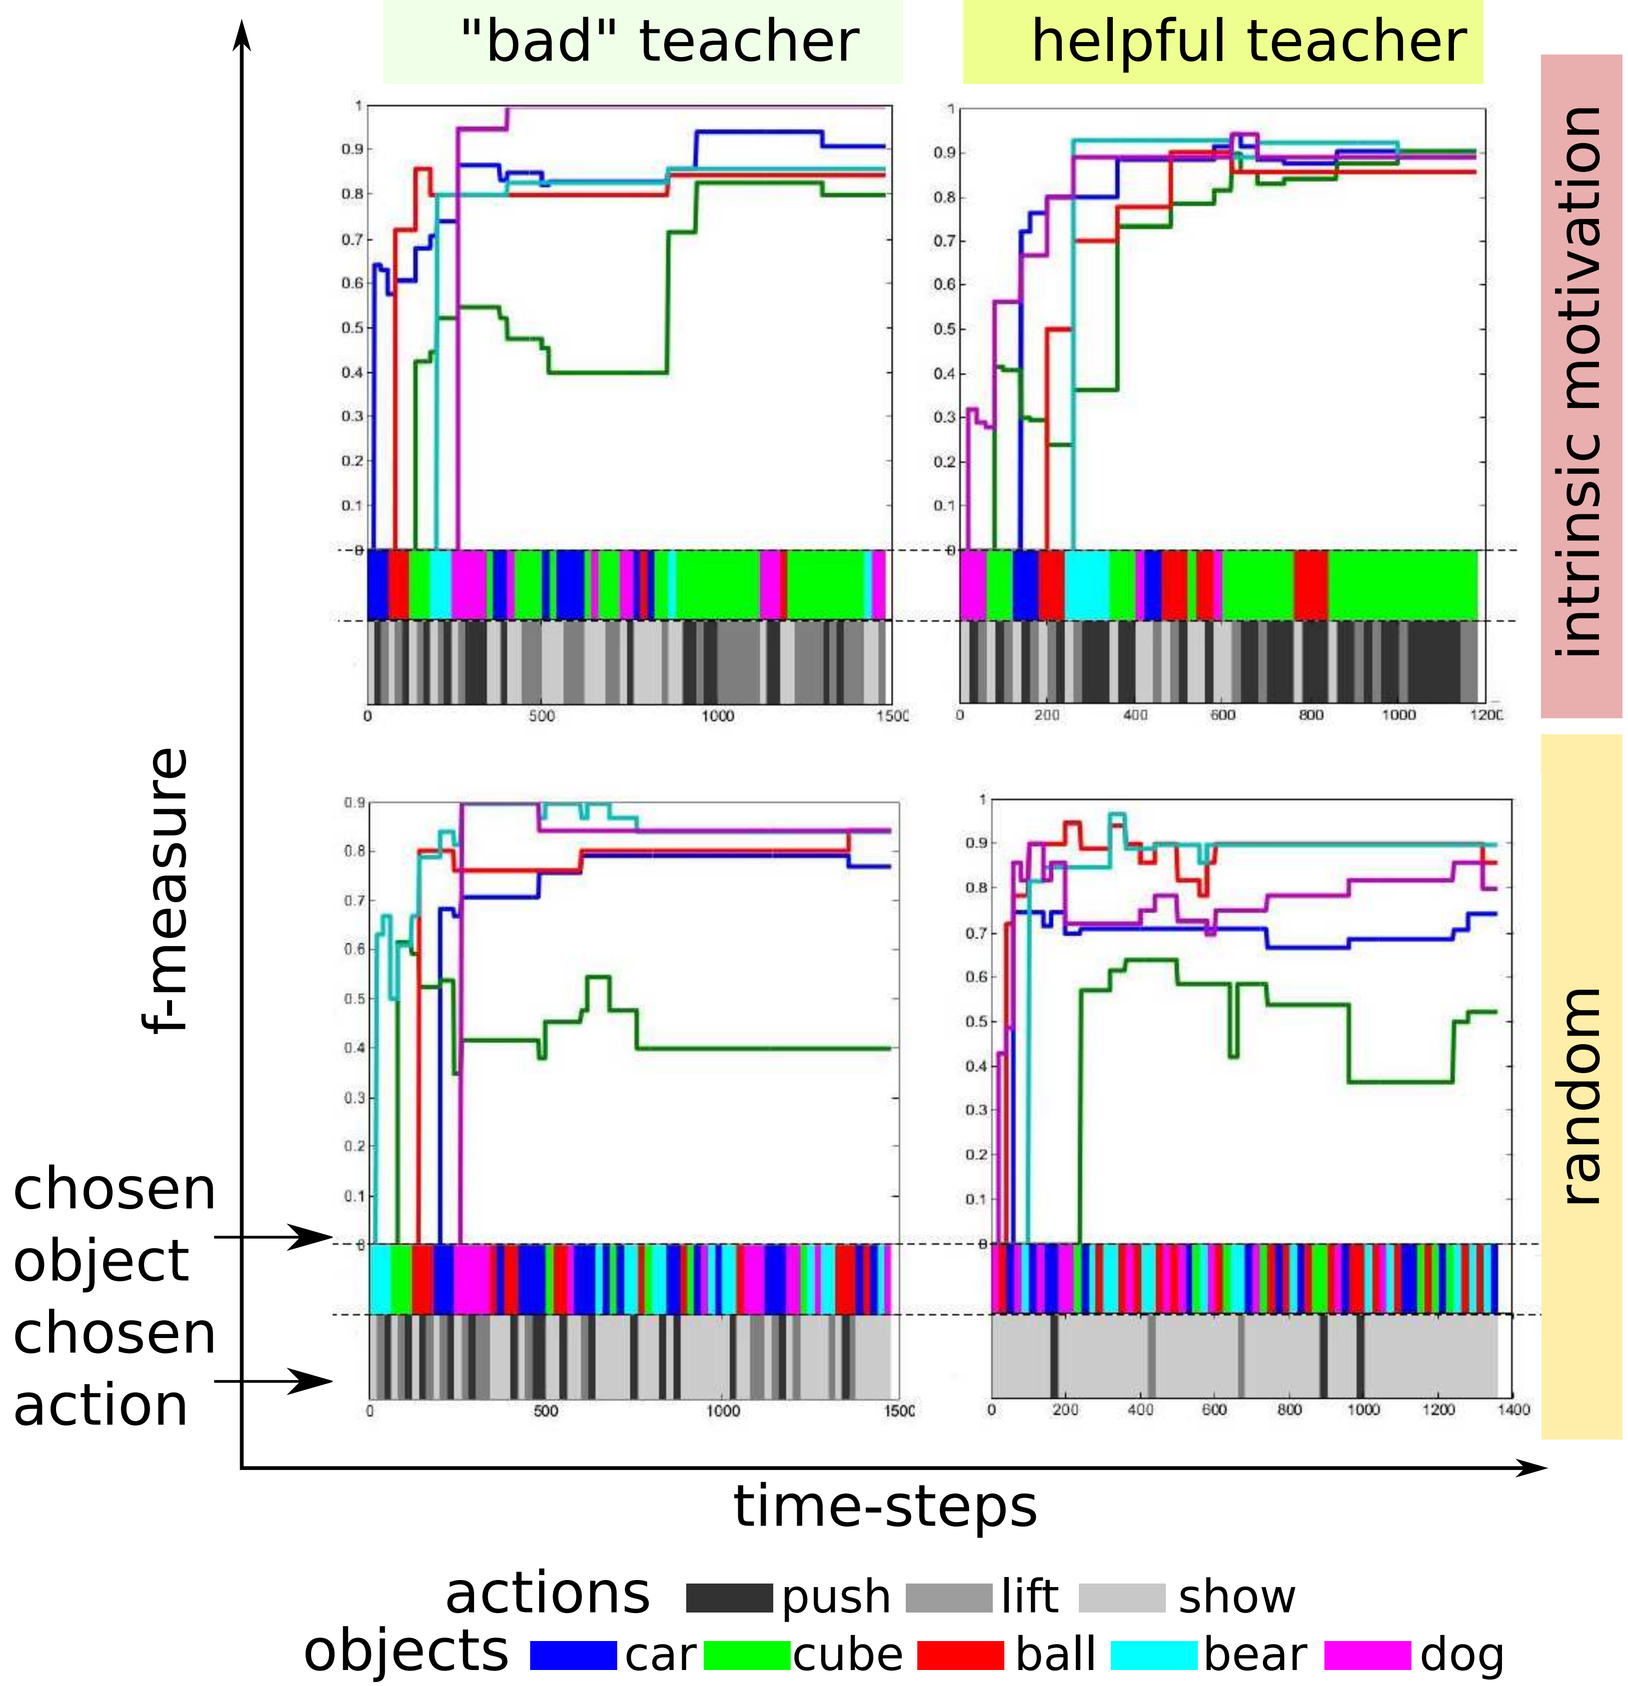

Better learning with action and interaction

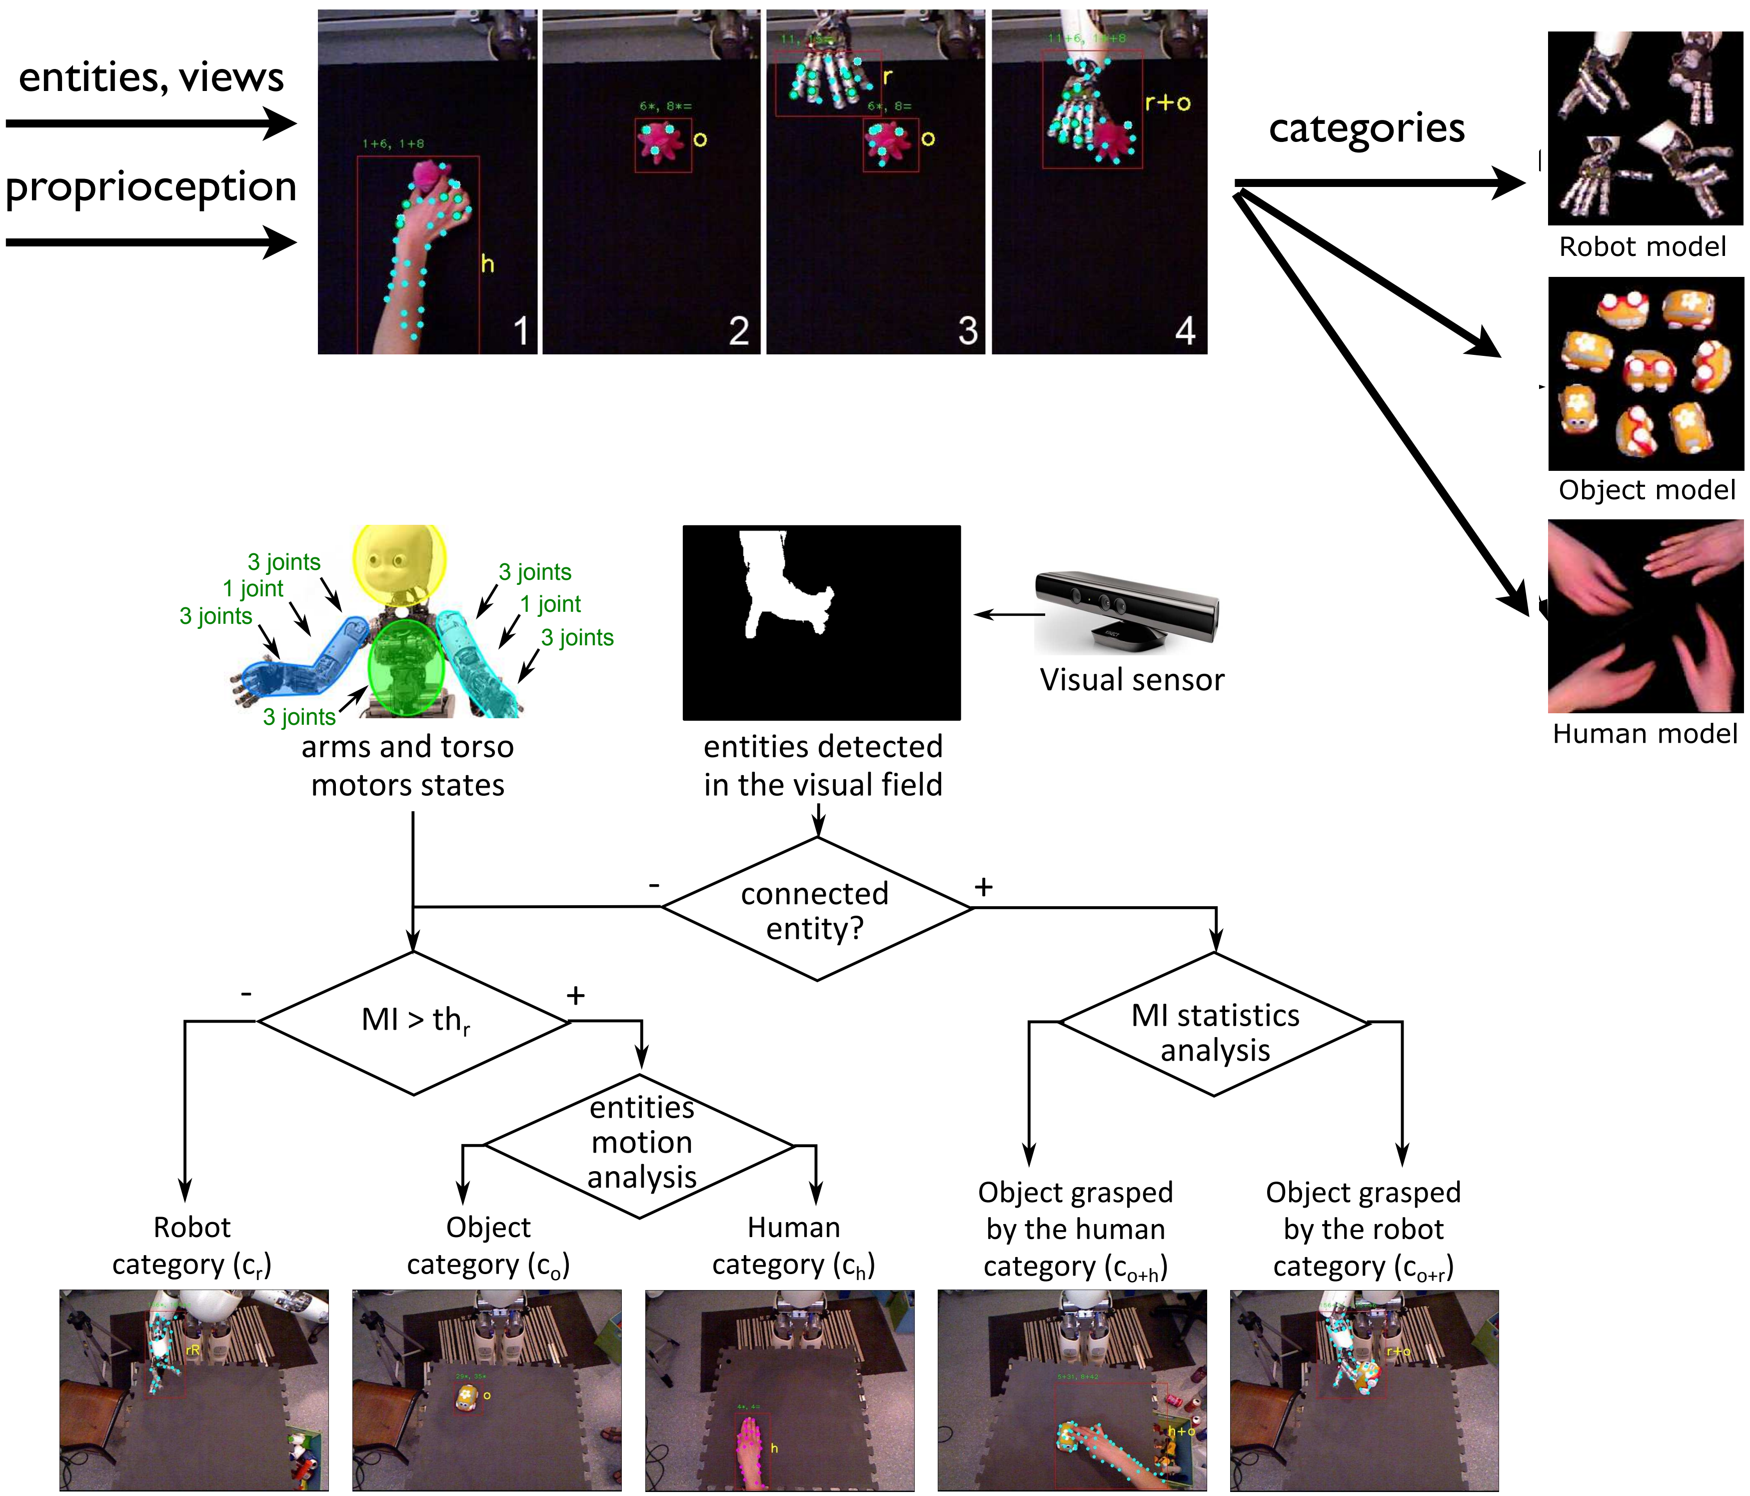

The robot learns the objects through manipulation.

The robot learns to identify its body, hence entities can be categorized as "robot hand", "human hand" and "object".

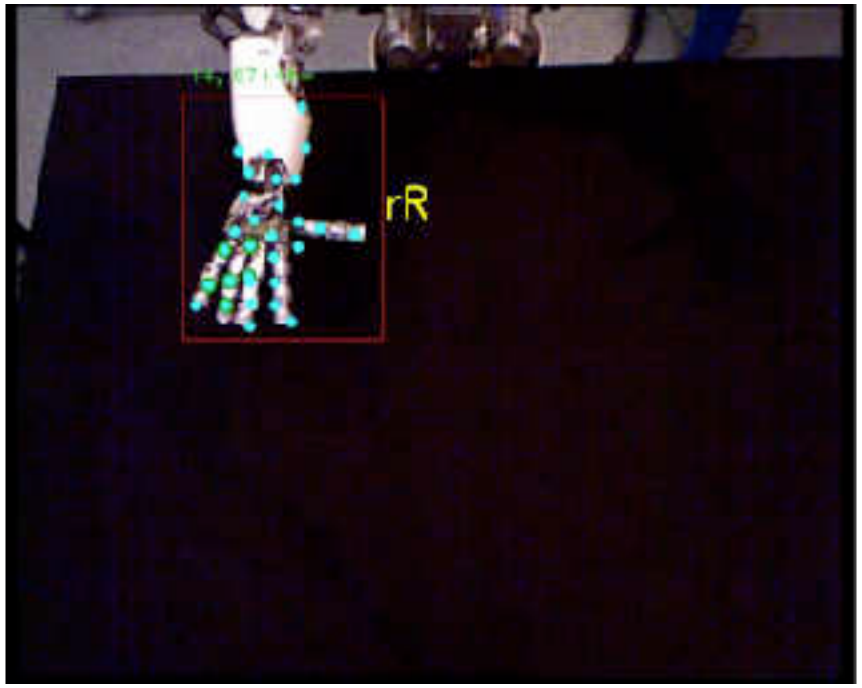

Lyubova, Ivaldi, Filliat (2016) From passive to interactive object learning and recognition through self-identification on a humanoid robot. Autonomous Robots, 40(1):33-57.

Better object recognition

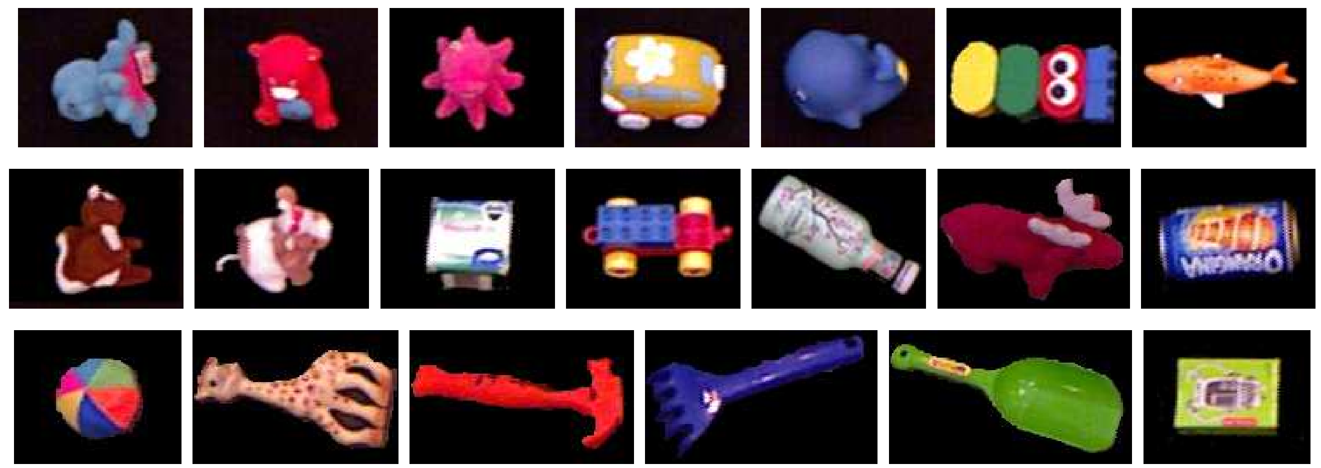

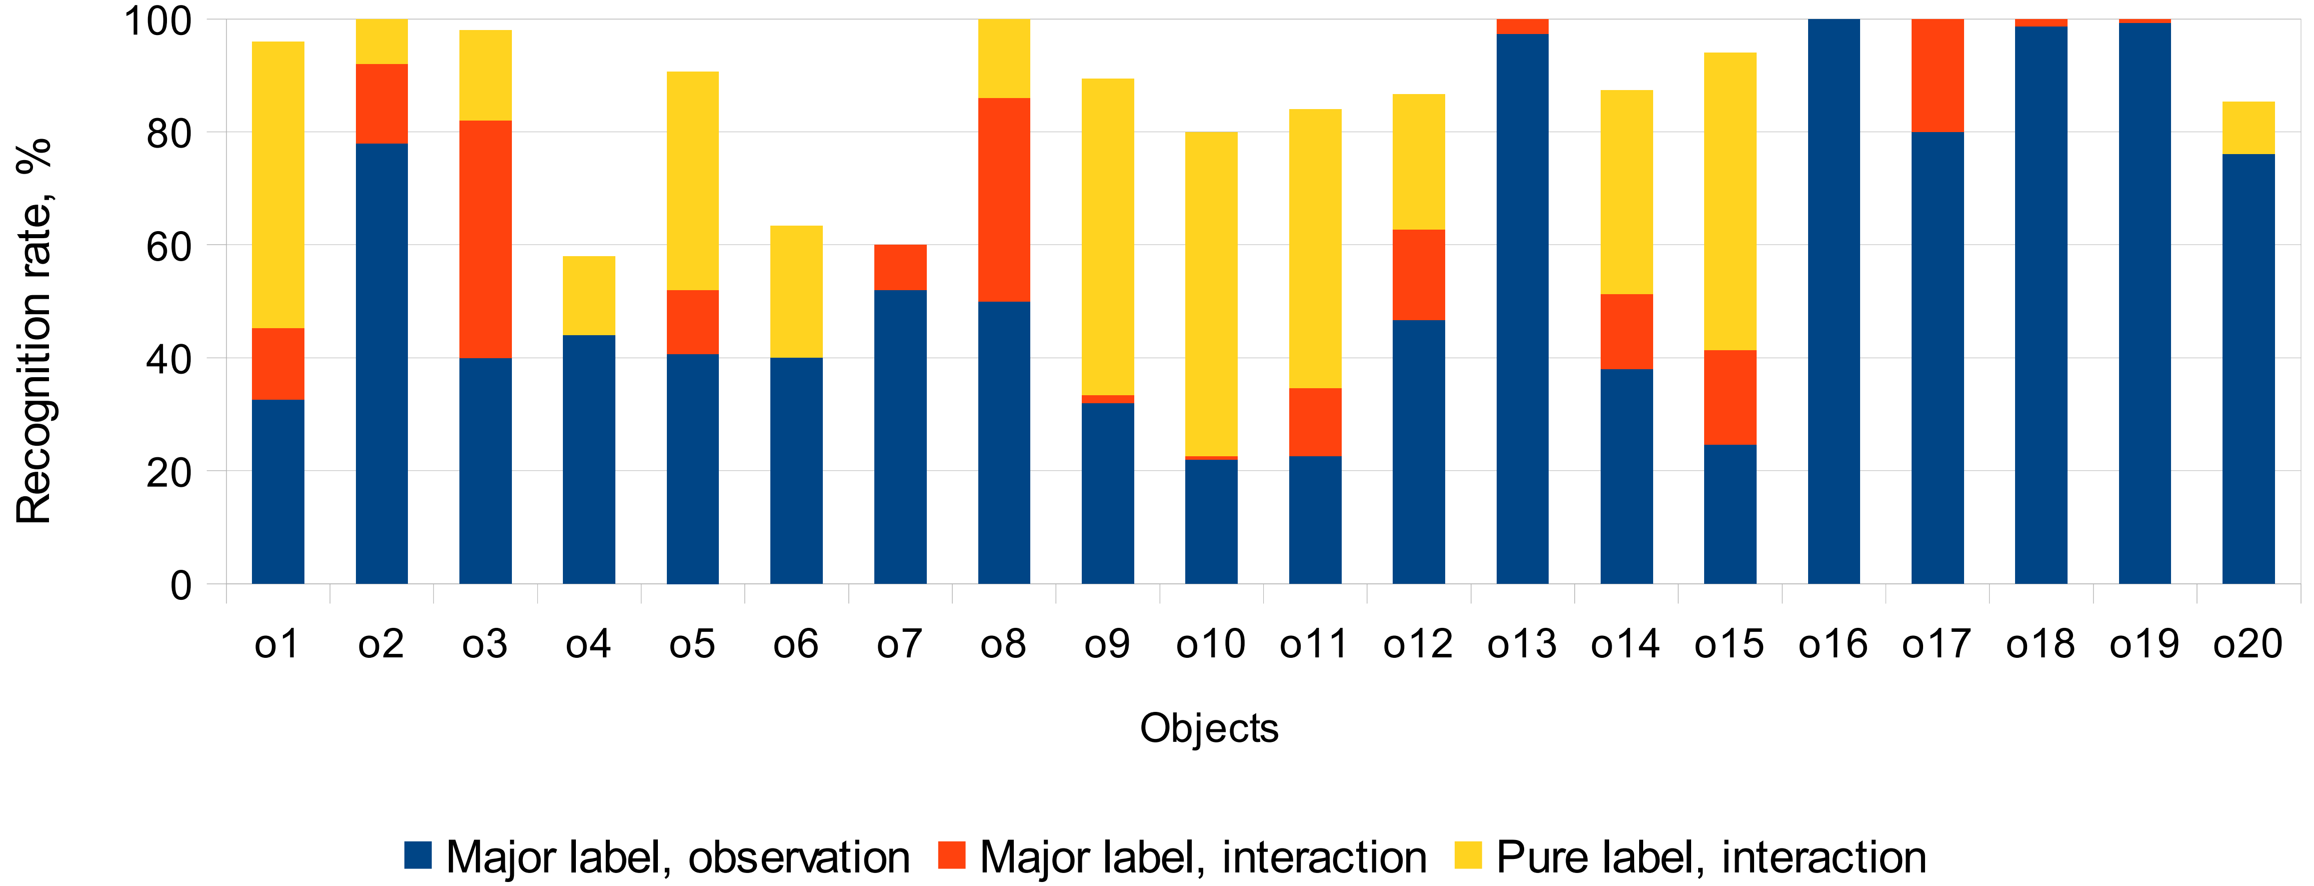

Lyubova, Ivaldi, Filliat (2016) From passive to interactive object learning and recognition through self-identification on a humanoid robot. Autonomous Robots, 40(1):33-57.

Better object recognition

This is open-ended!!
No limit to the number of objects, they are learned incrementally.

The limit is how much you are patient and interact with the robot :)

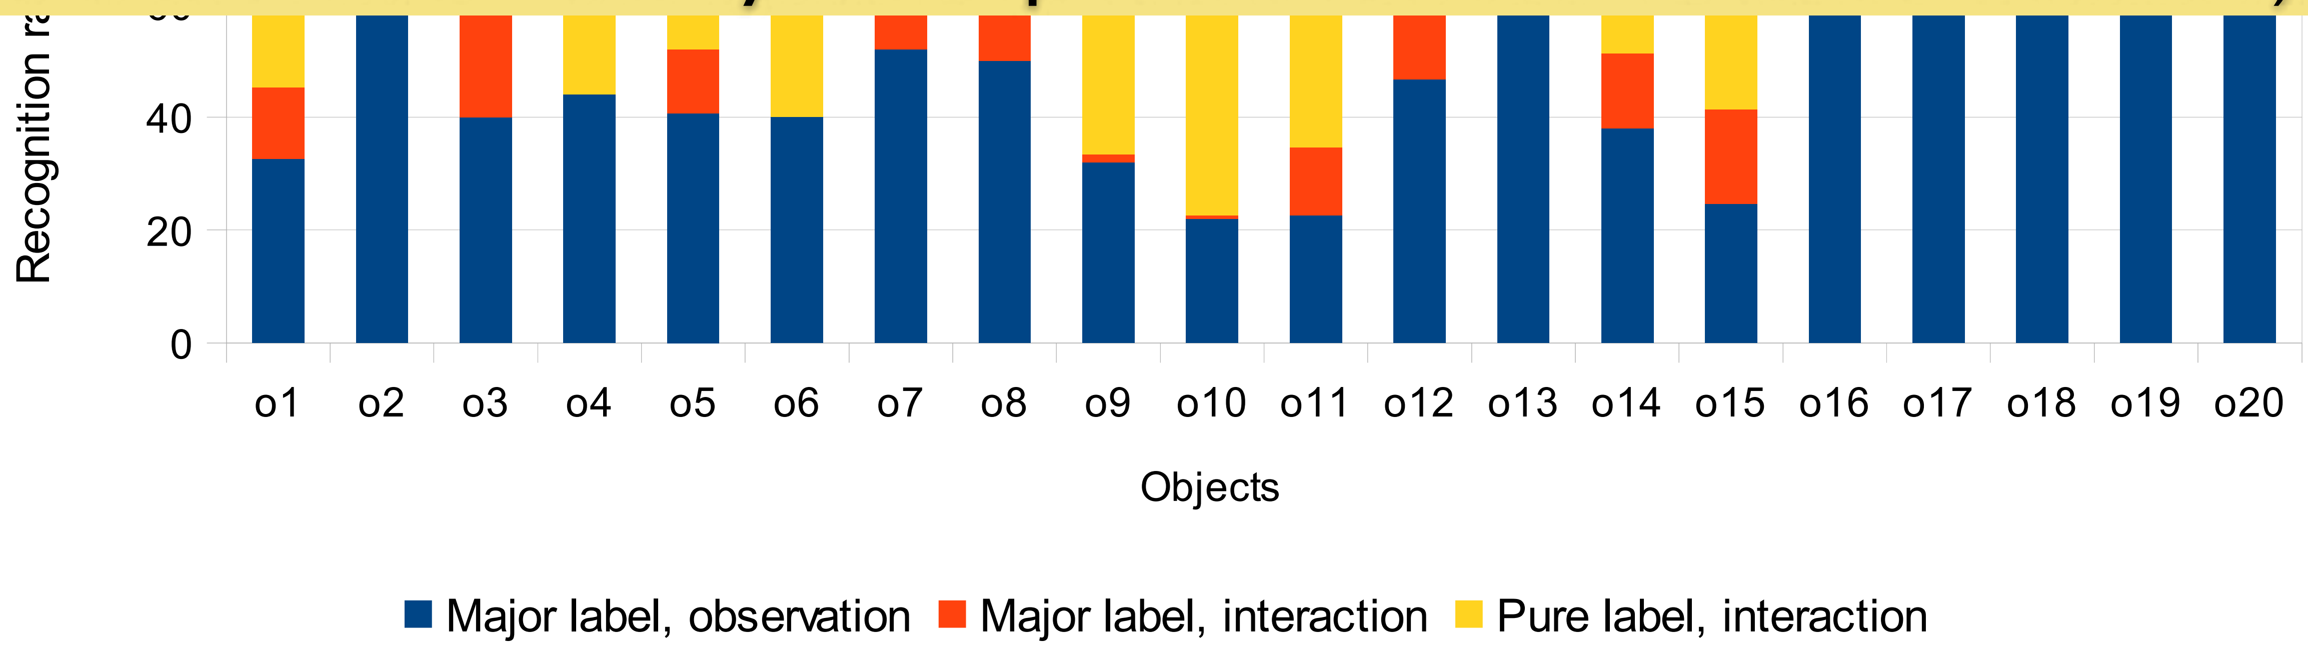

Lyubova, Ivaldi, Filliat (2016) From passive to interactive object learning and recognition through self-identification on a humanoid robot. Autonomous Robots, 40(1):33-57.

Outline

Learning the visual appearance of objects from scratch - with curiosity-based autonomous exploration and social guidance

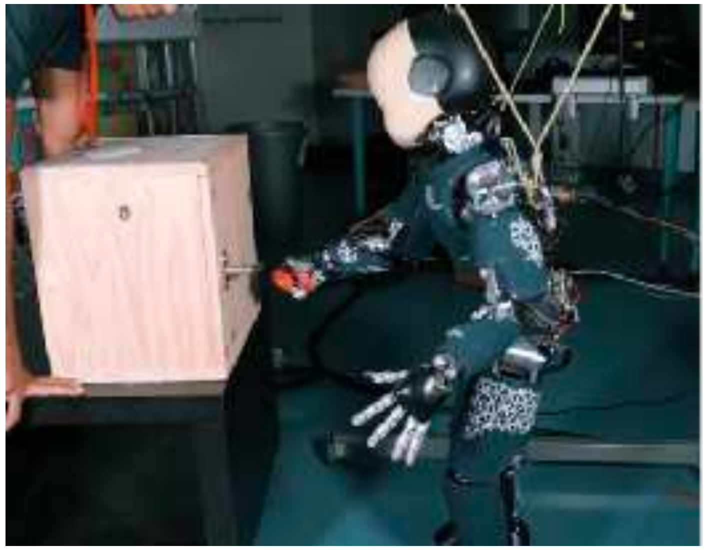

Learning by demonstration whole-body manipulation on humanoids

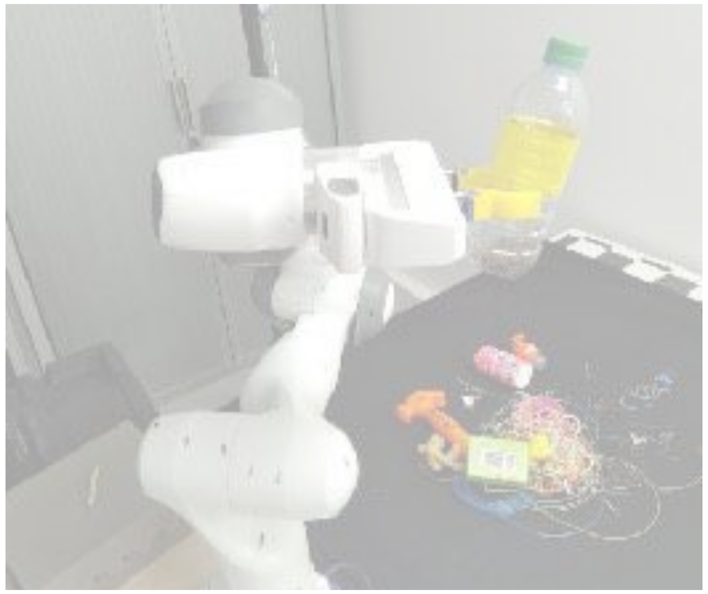

HEAP project - Incorporating human preferences to improve grasping of irregular objects in a heap

Towards whole-body manipulation

Grasping an object is a particular task adding to the balancing and locomotion tasks of the whole-body robot controller.

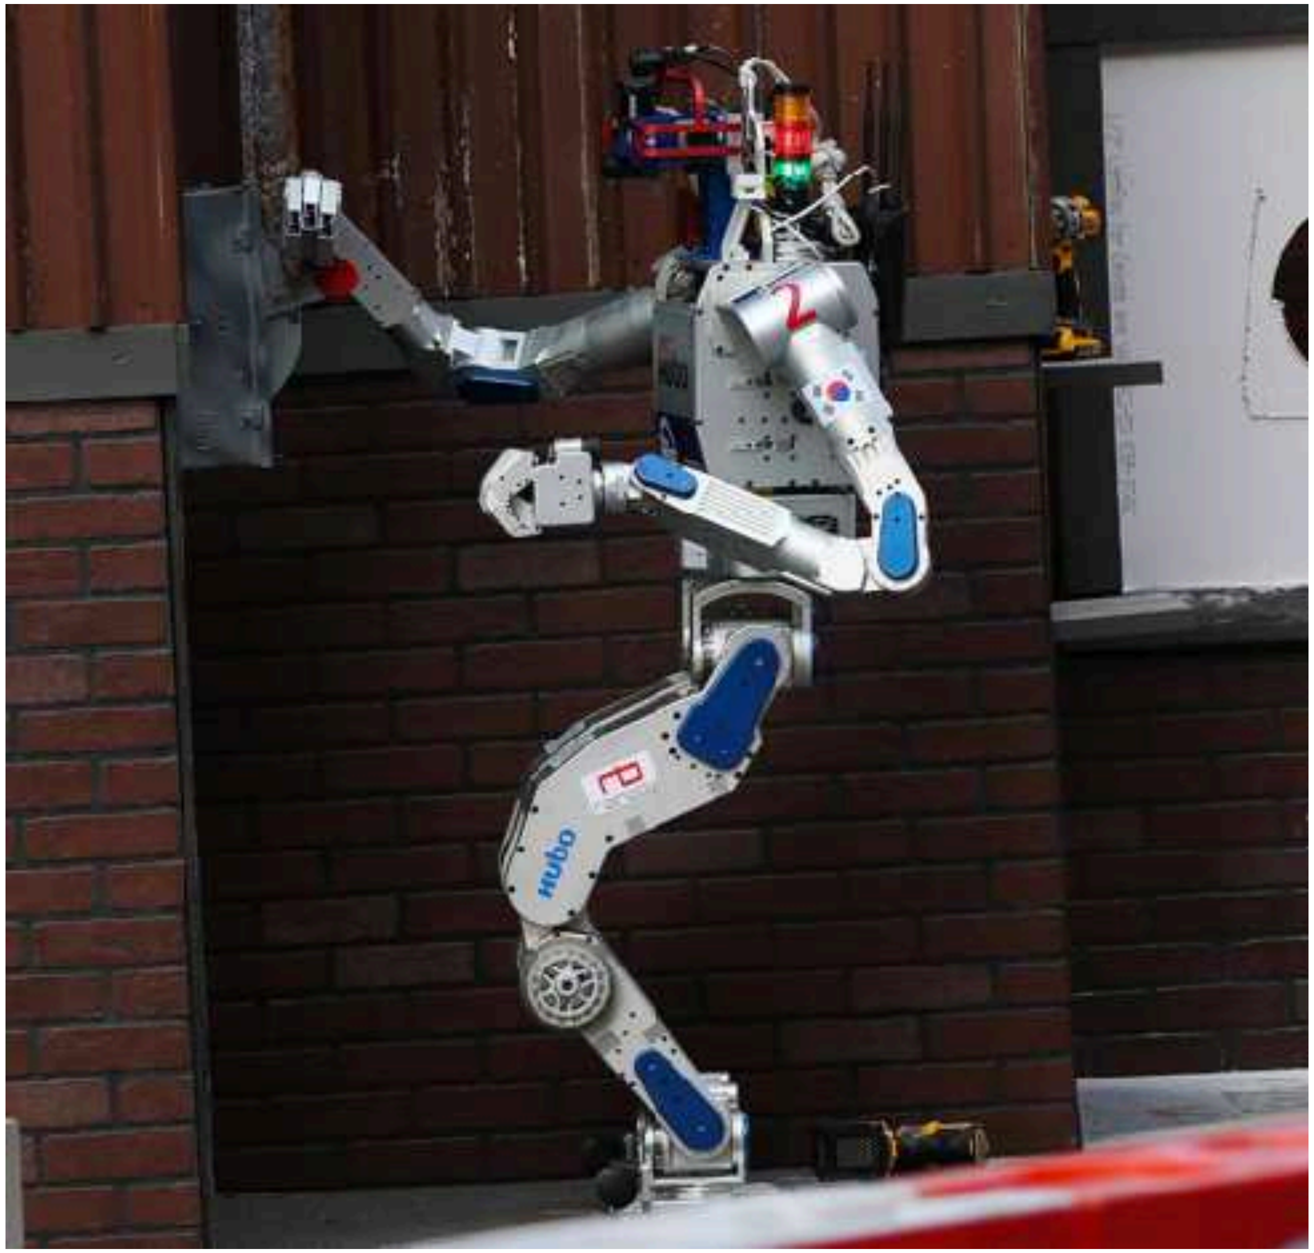

Towards whole-body grasping & manipulation

Tele-operation/retargeting is the whole-body kinesthetic teaching

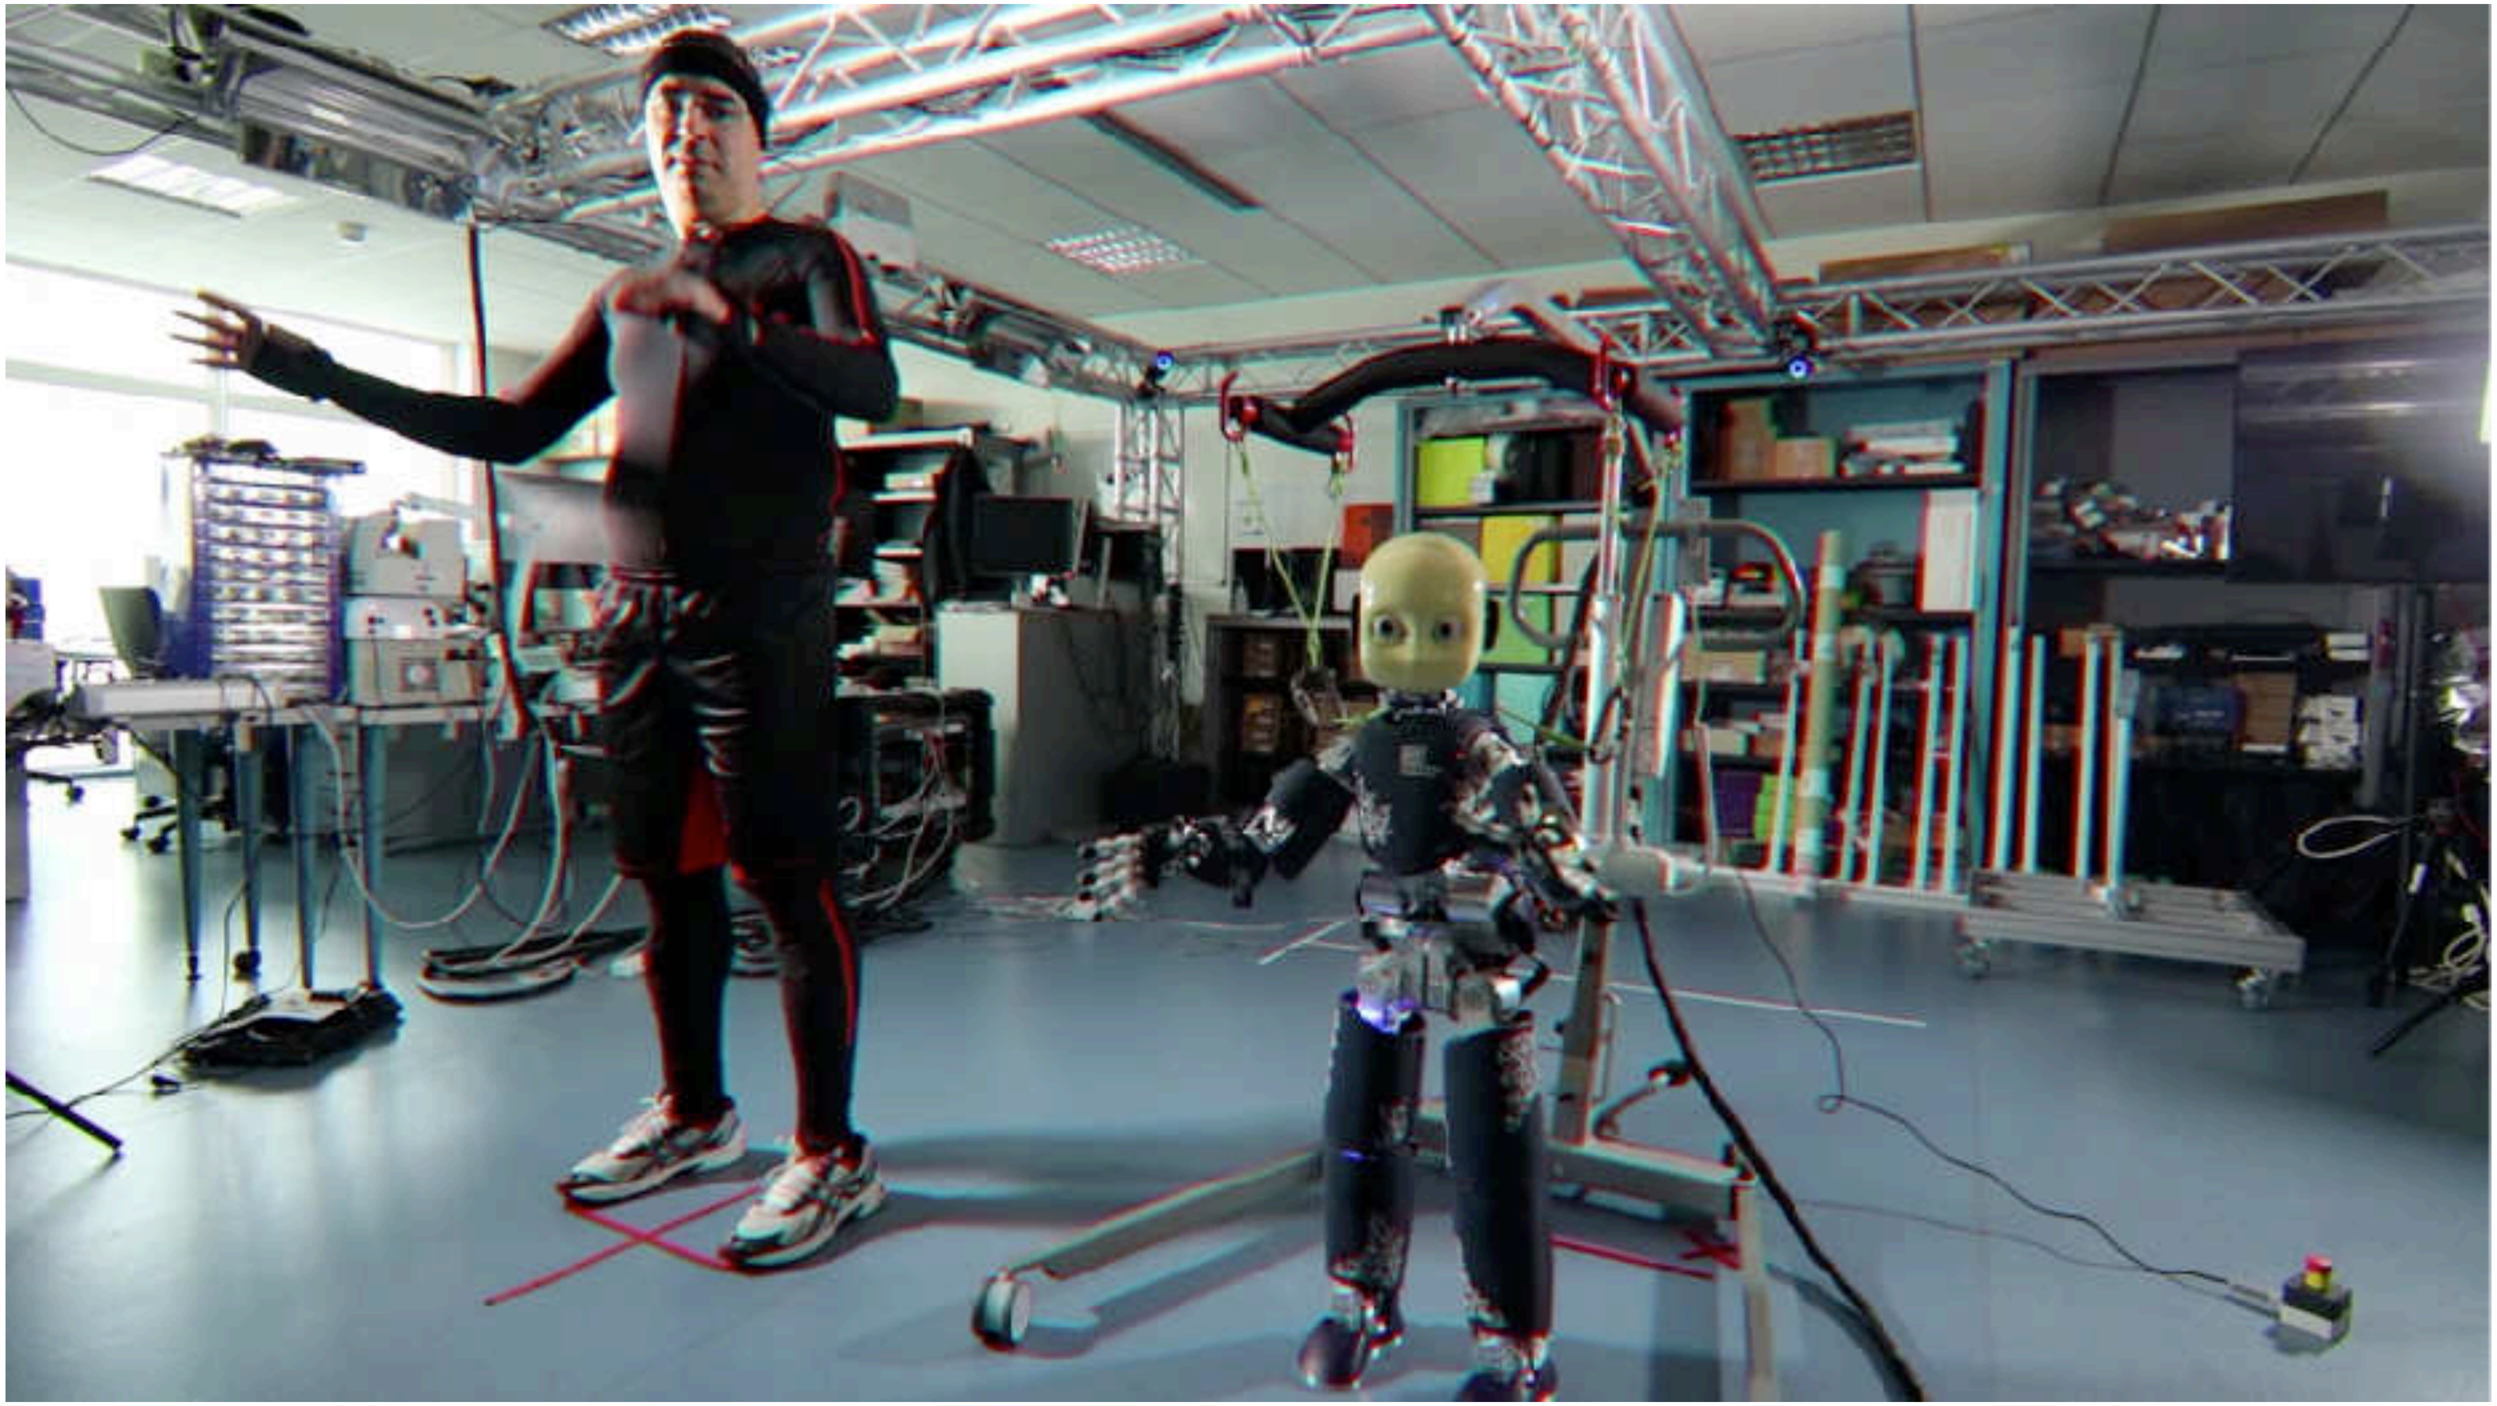

Penco et al (2018) **Robust real-time whole-body motion retargeting from human to humanoid.**
Proc. IEEE/RAS International Conf. on Humanoid Robots (HUMANOIDS).

Transfer paradigm: from humans to robots

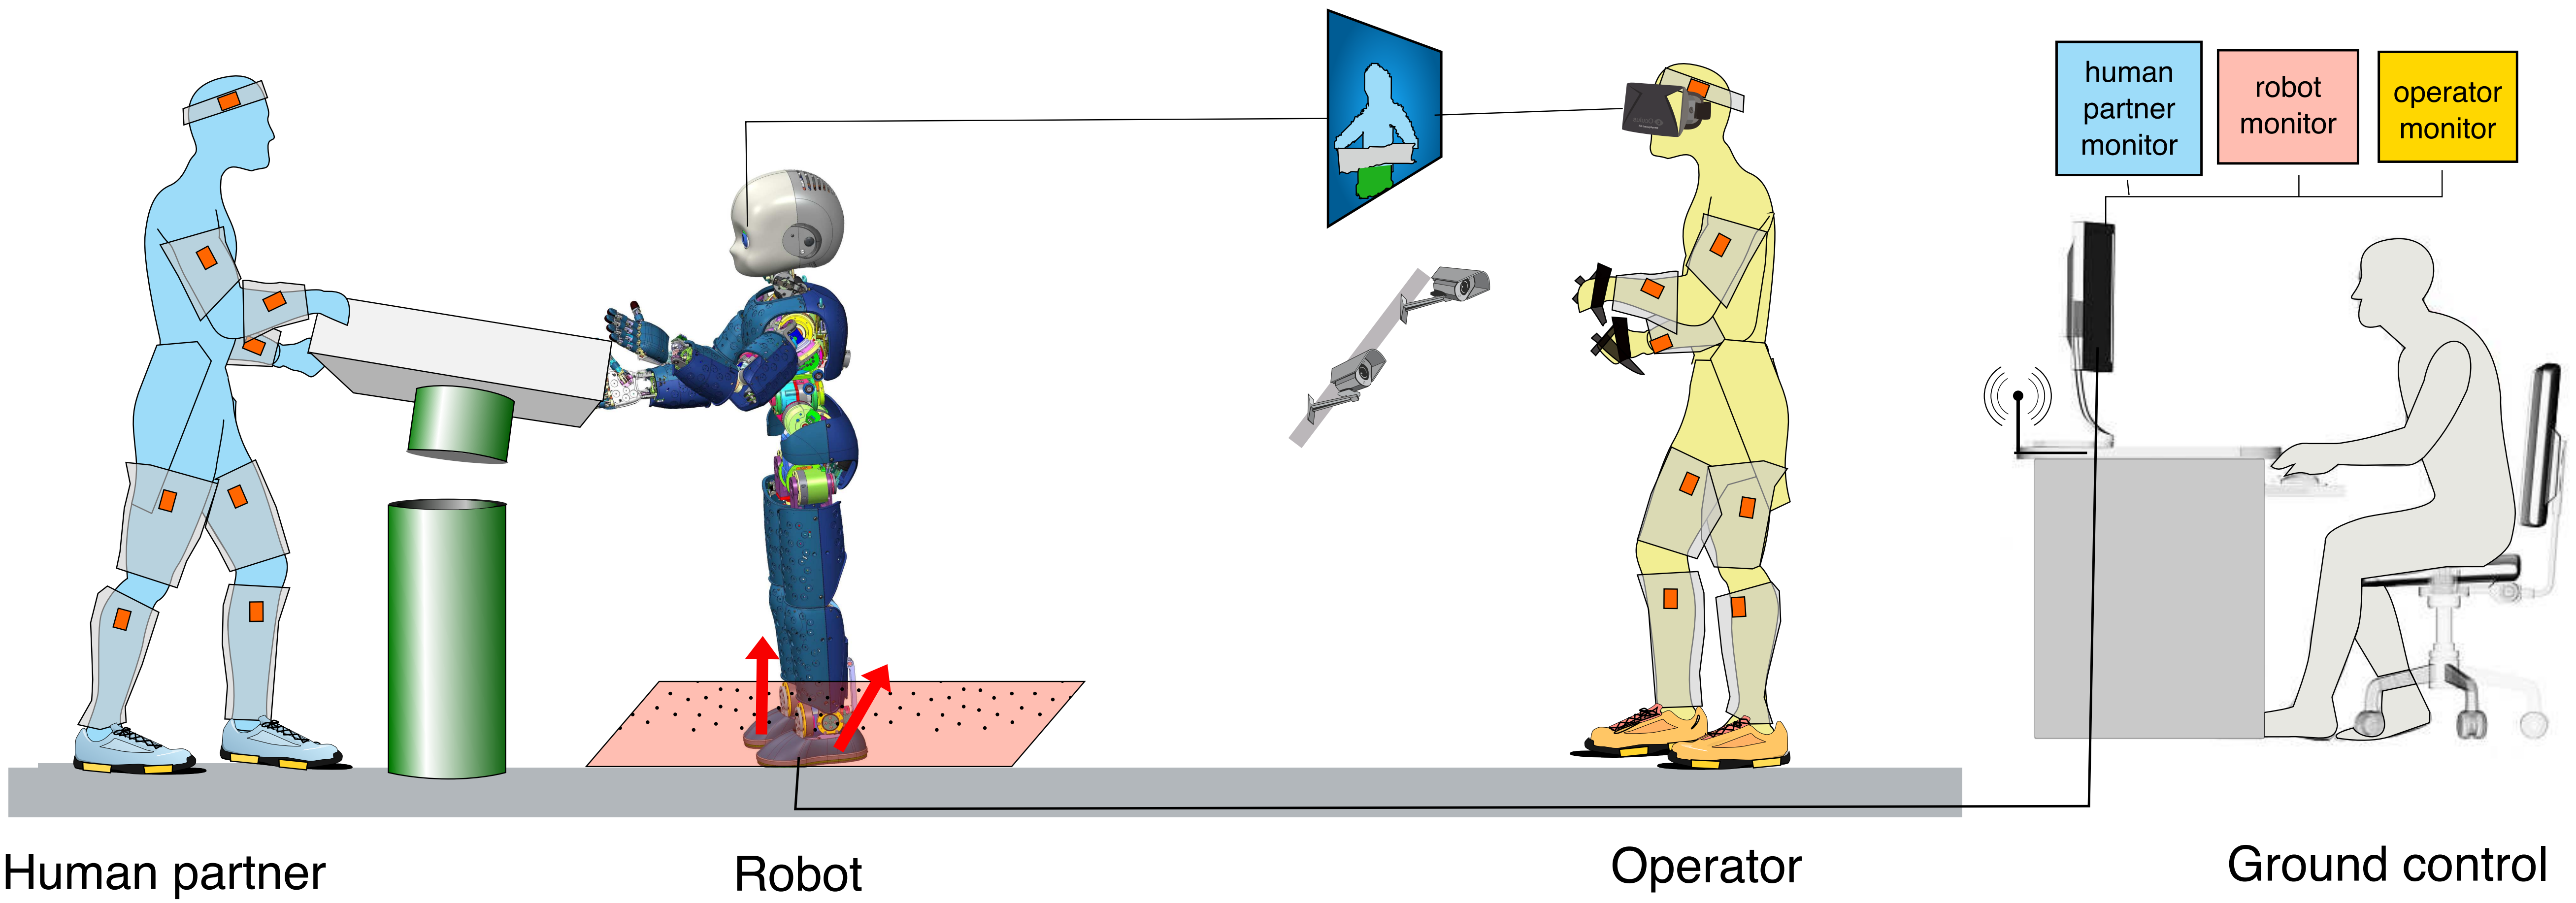

Towards whole-body grasping & manipulation

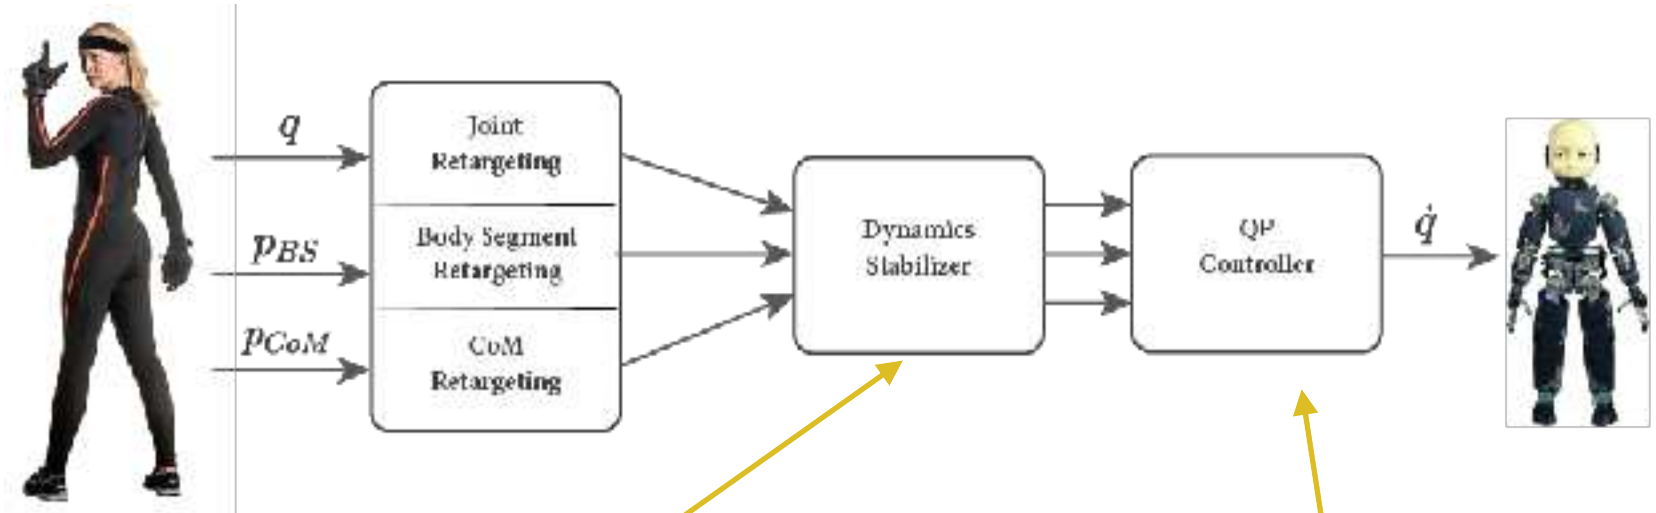

$$
\begin{aligned}
\min_{p_{ZMP}} \quad & (\dot{p}^{des}_{CoM} - \dot{p}_{CoM})^T R \, (\dot{p}^{des}_{CoM} - \dot{p}_{CoM}) \\
\text{s.t.} \quad & \dot{p}_{CoM} = \dot{p}^{t-1}_{CoM} + \frac{Tg}{h^{t-1}}(p_{CoM} - p_{ZMP}) \\
& lb_{SP} < p_{ZMP} < ub_{SP}
\end{aligned}
$$

ZMP correction

$$\min_{x} \; || \, Ax - b \, ||^2_W \quad s.t. \quad l \leq Cx \leq u$$

$$stack = (lFoot + rFoot + head_{sub})/$$
$$(com_{sub} + base_{sub} + torso_{sub} + lArm_{sub} + rArm_{sub});$$

Multi-task QP controller with strict priorities (OpenSoT)

Manipulation-related tasks must be put here!

Penco et al (2018) **Robust real-time whole-body motion retargeting from human to humanoid.** Proc. IEEE/RAS International Conf. on Humanoid Robots (HUMANOIDS).

Auto-tuning the controller

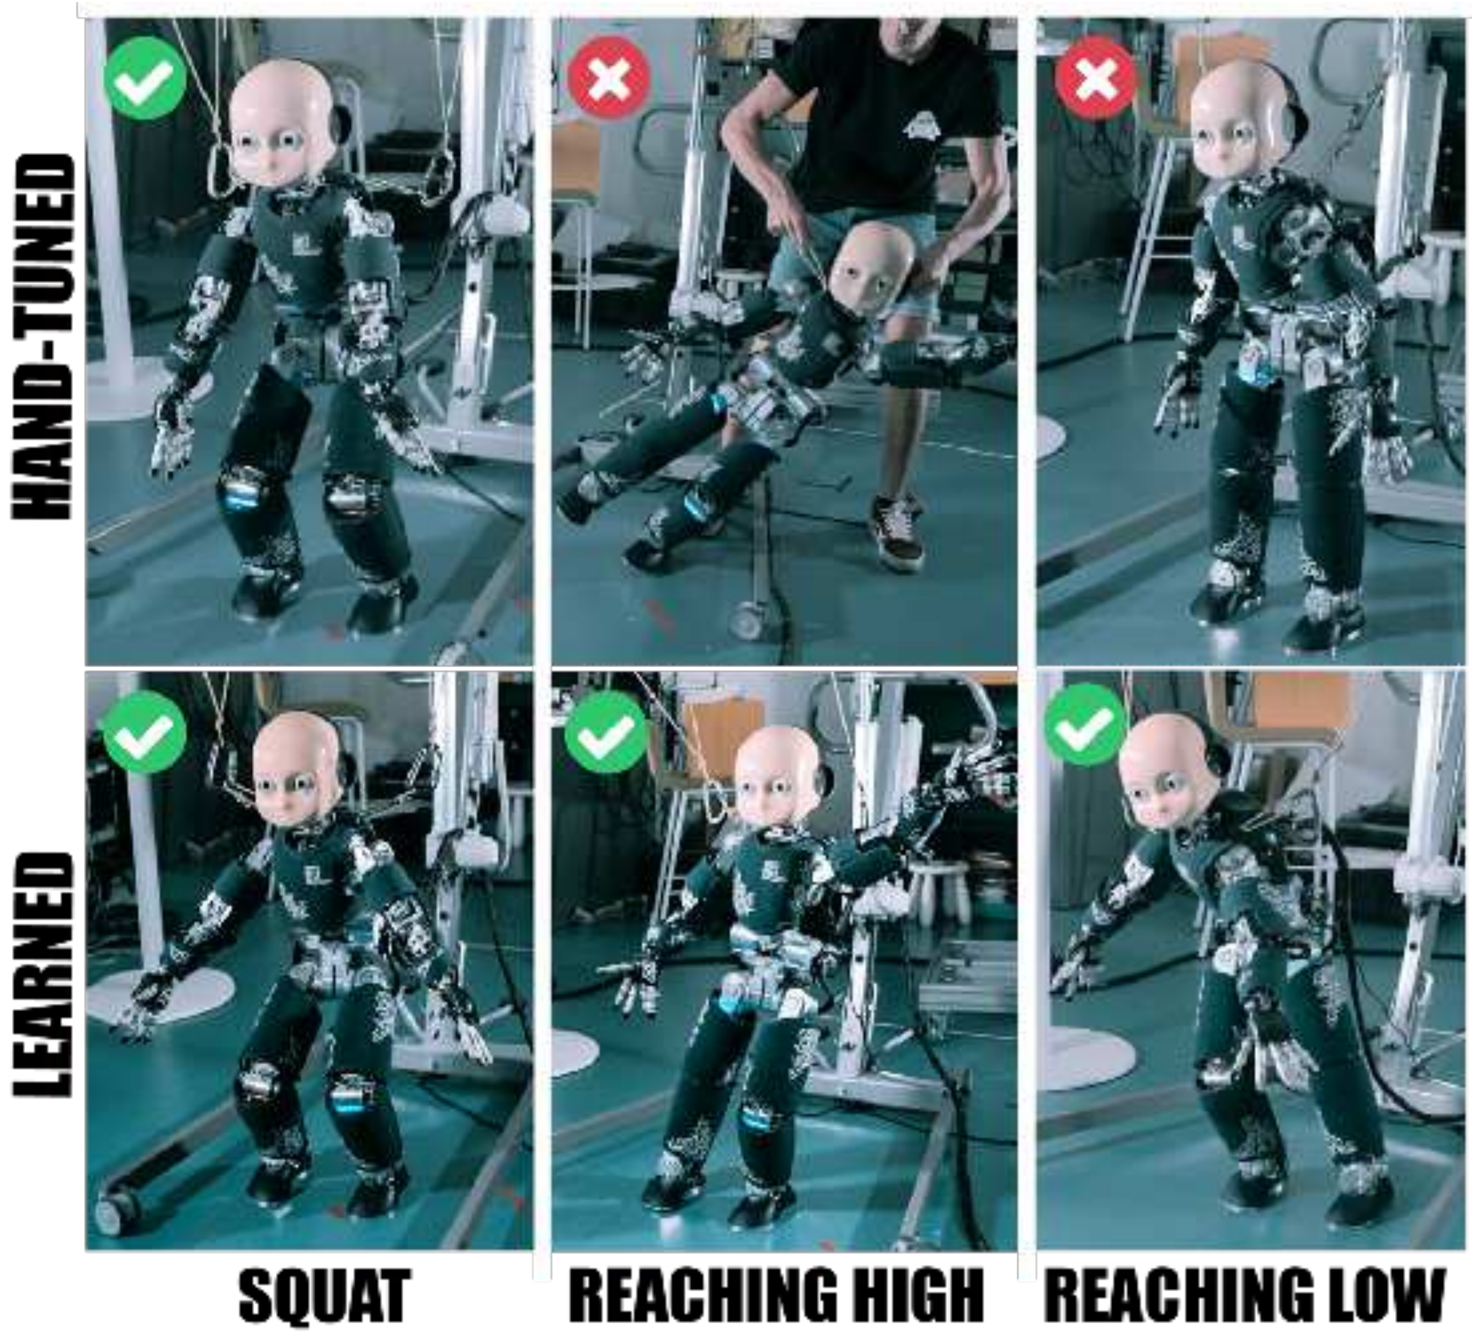

Can we make sure that the robot controller can execute in principle any retargeted human motion in real-time?

Auto-tuning the controller for teleoperation

Learning the control structure and the parameters that enable the robot to perform a variety of motions

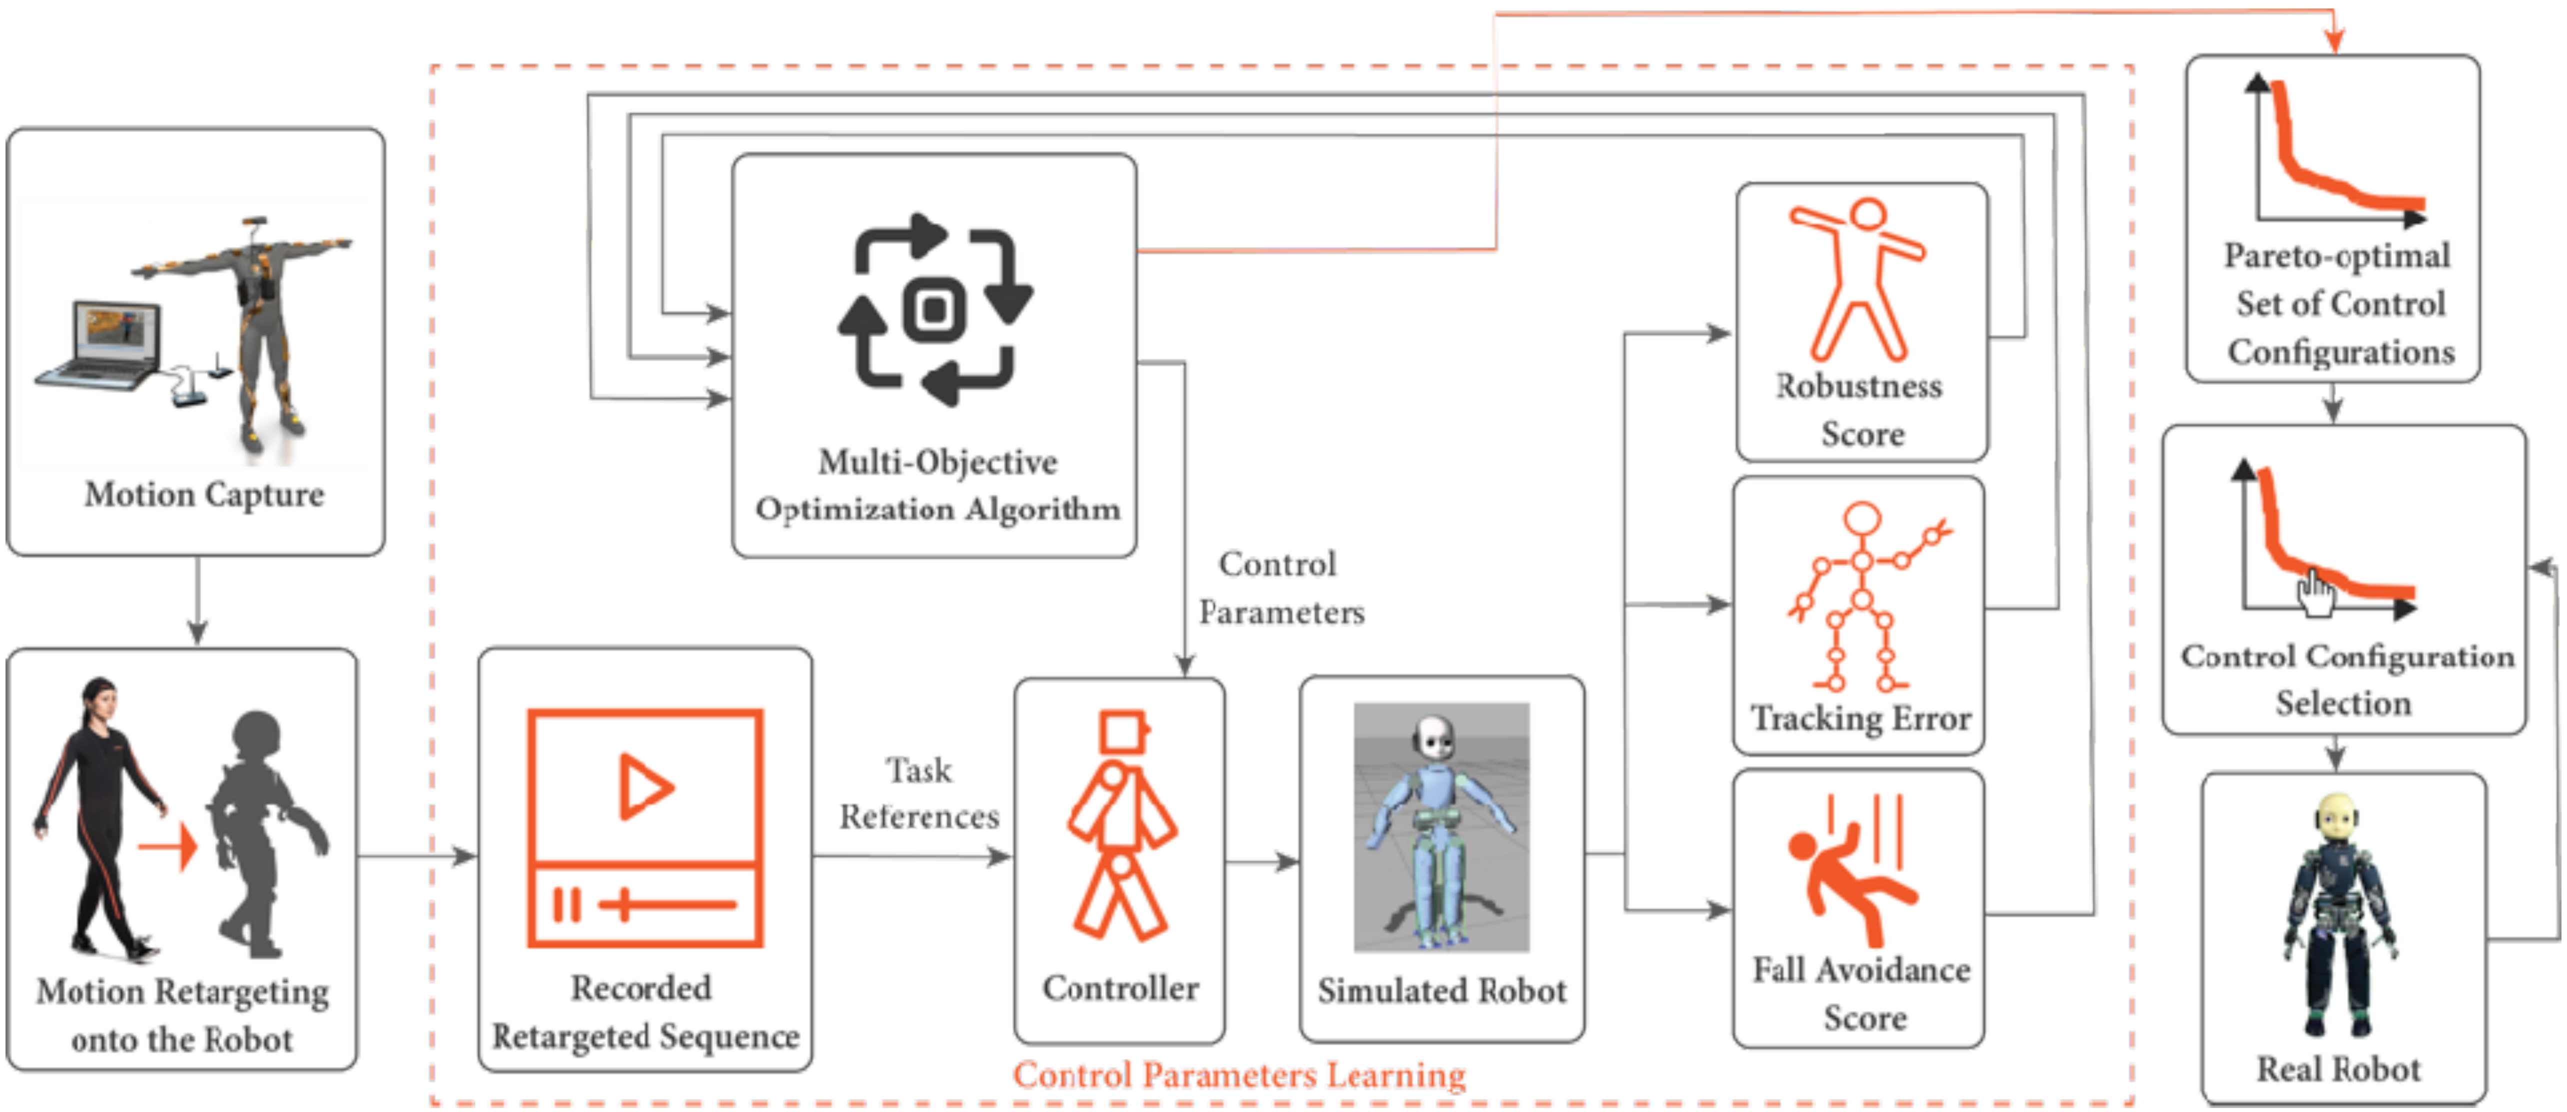

Penco et al. - under review

Auto-tuning the controller for teleoperation

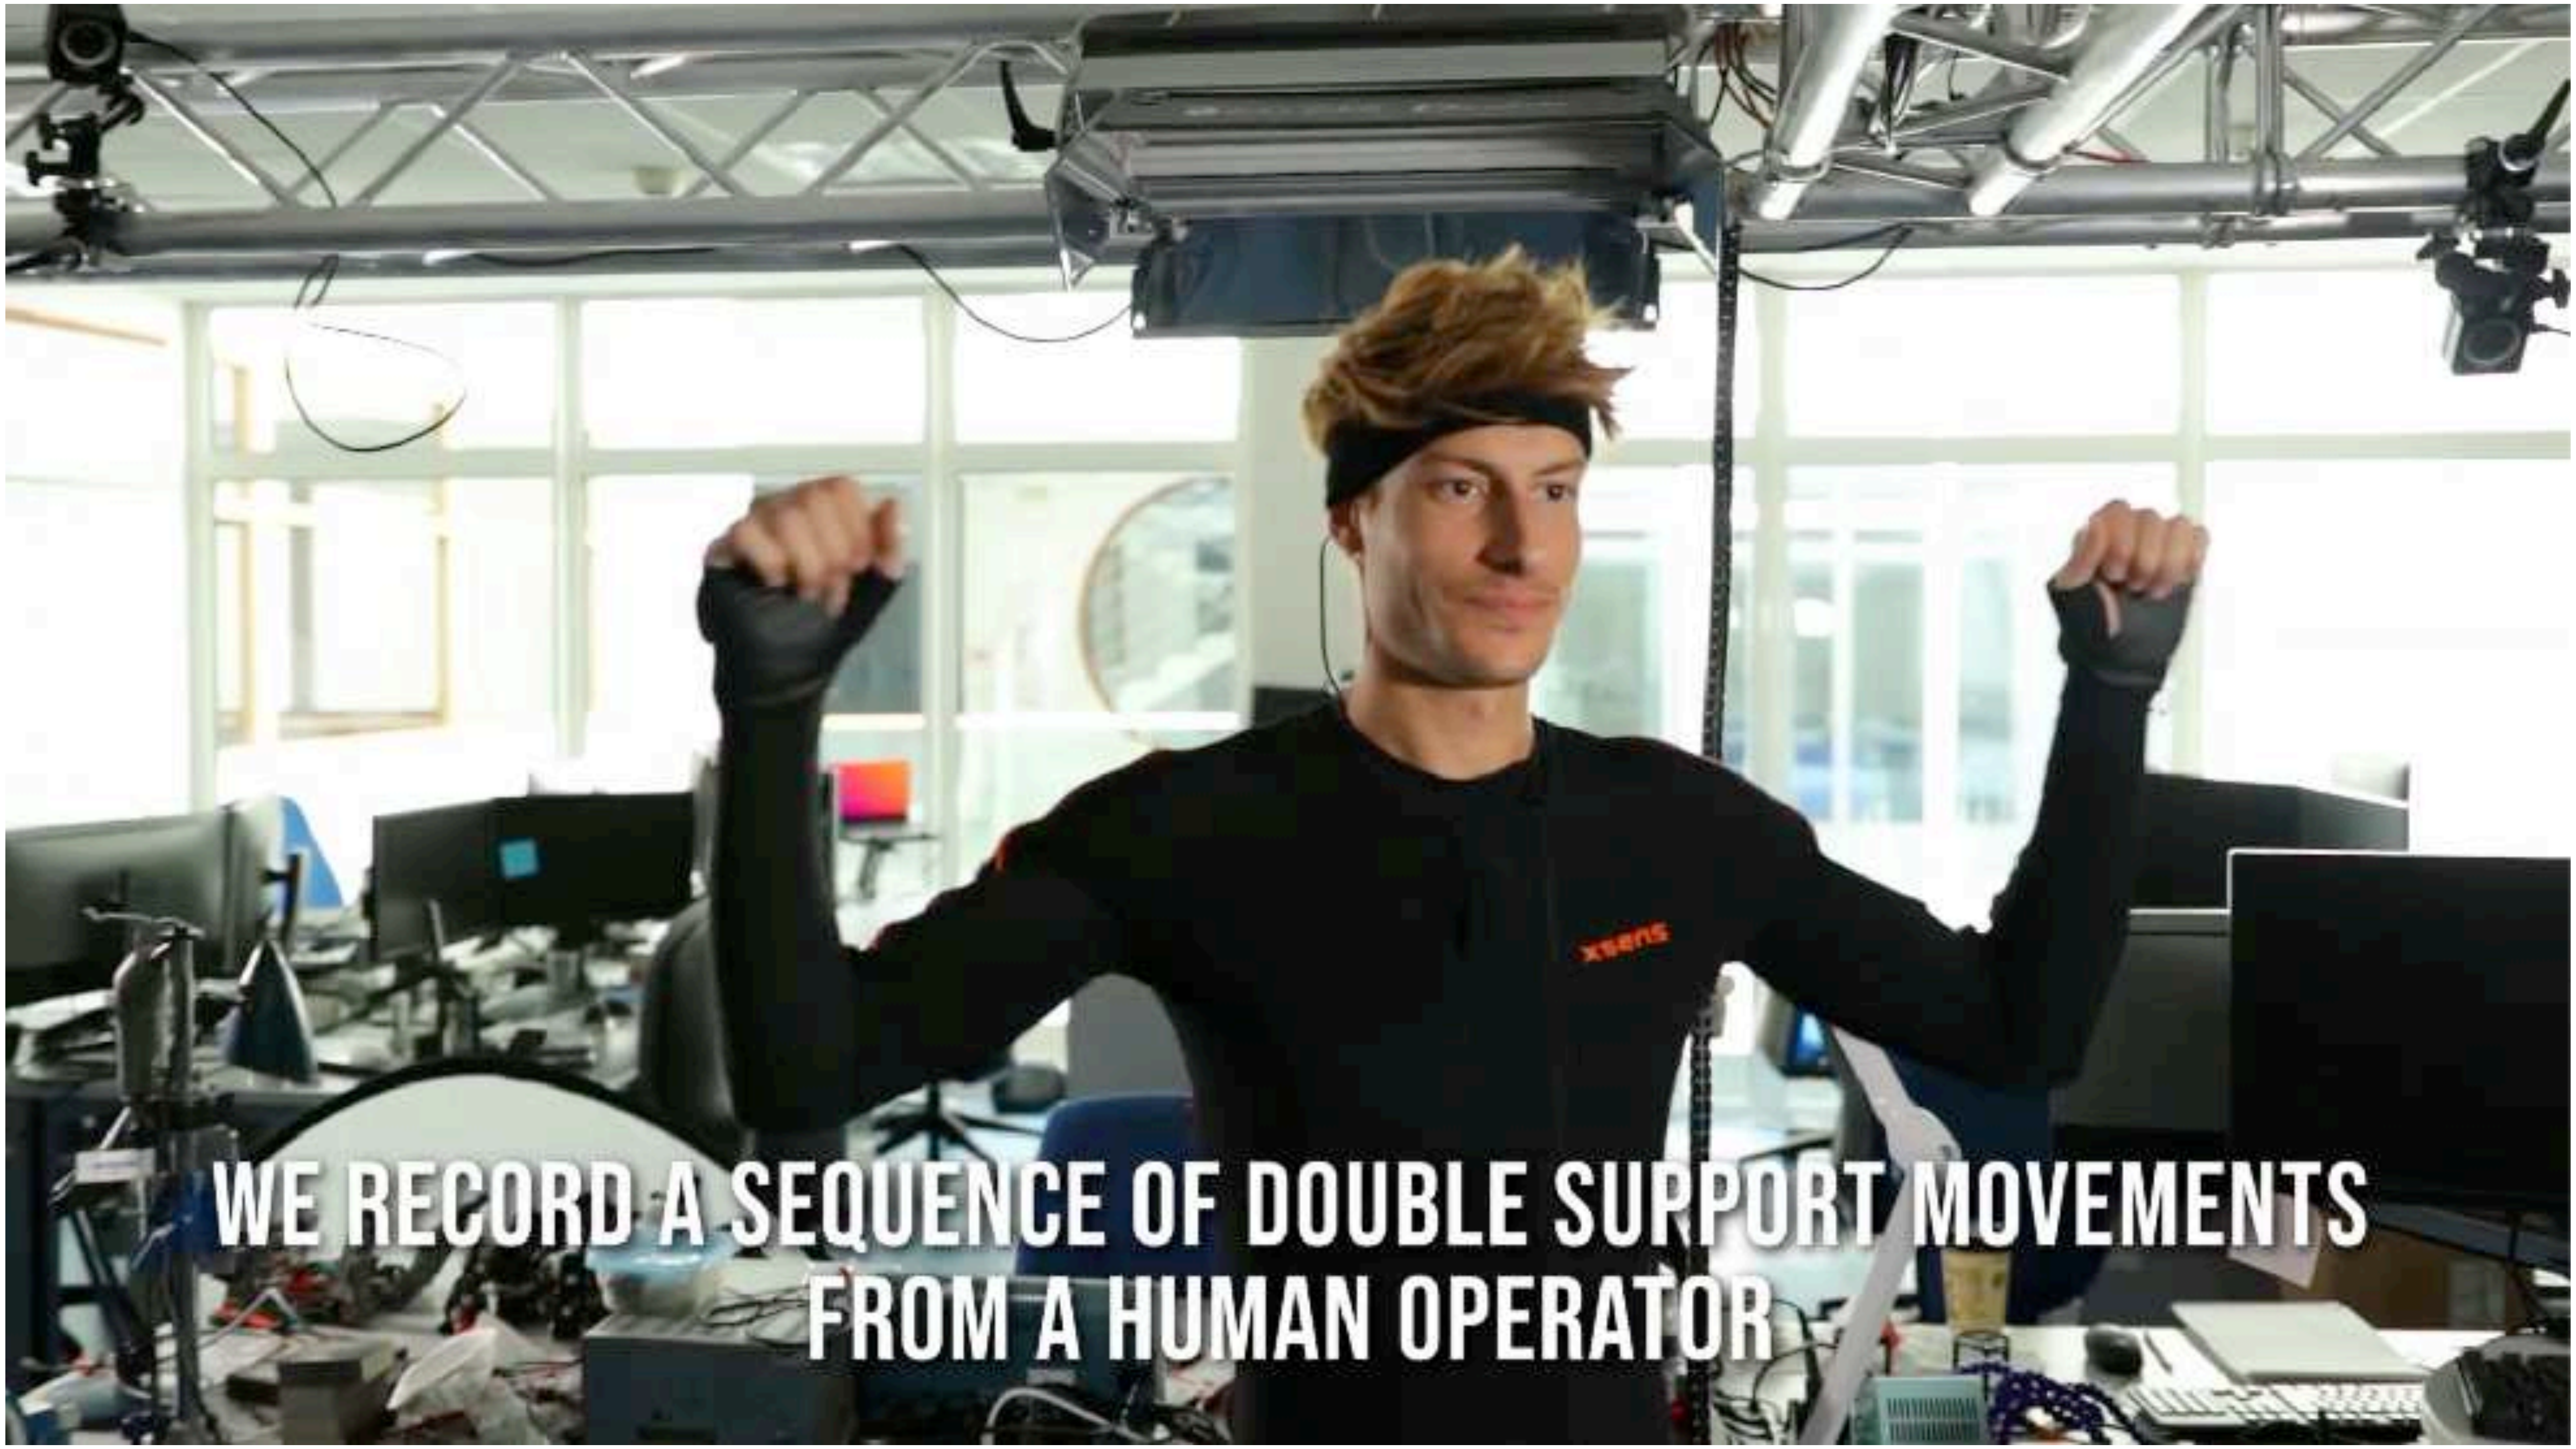

Penco et al. - under review

Auto-tuning the controller for teleoperation

Learning the control structure and the parameters that enable the robot to perform a variety of motions

Penco et al. - under review

Demonstrating whole-body manipulation

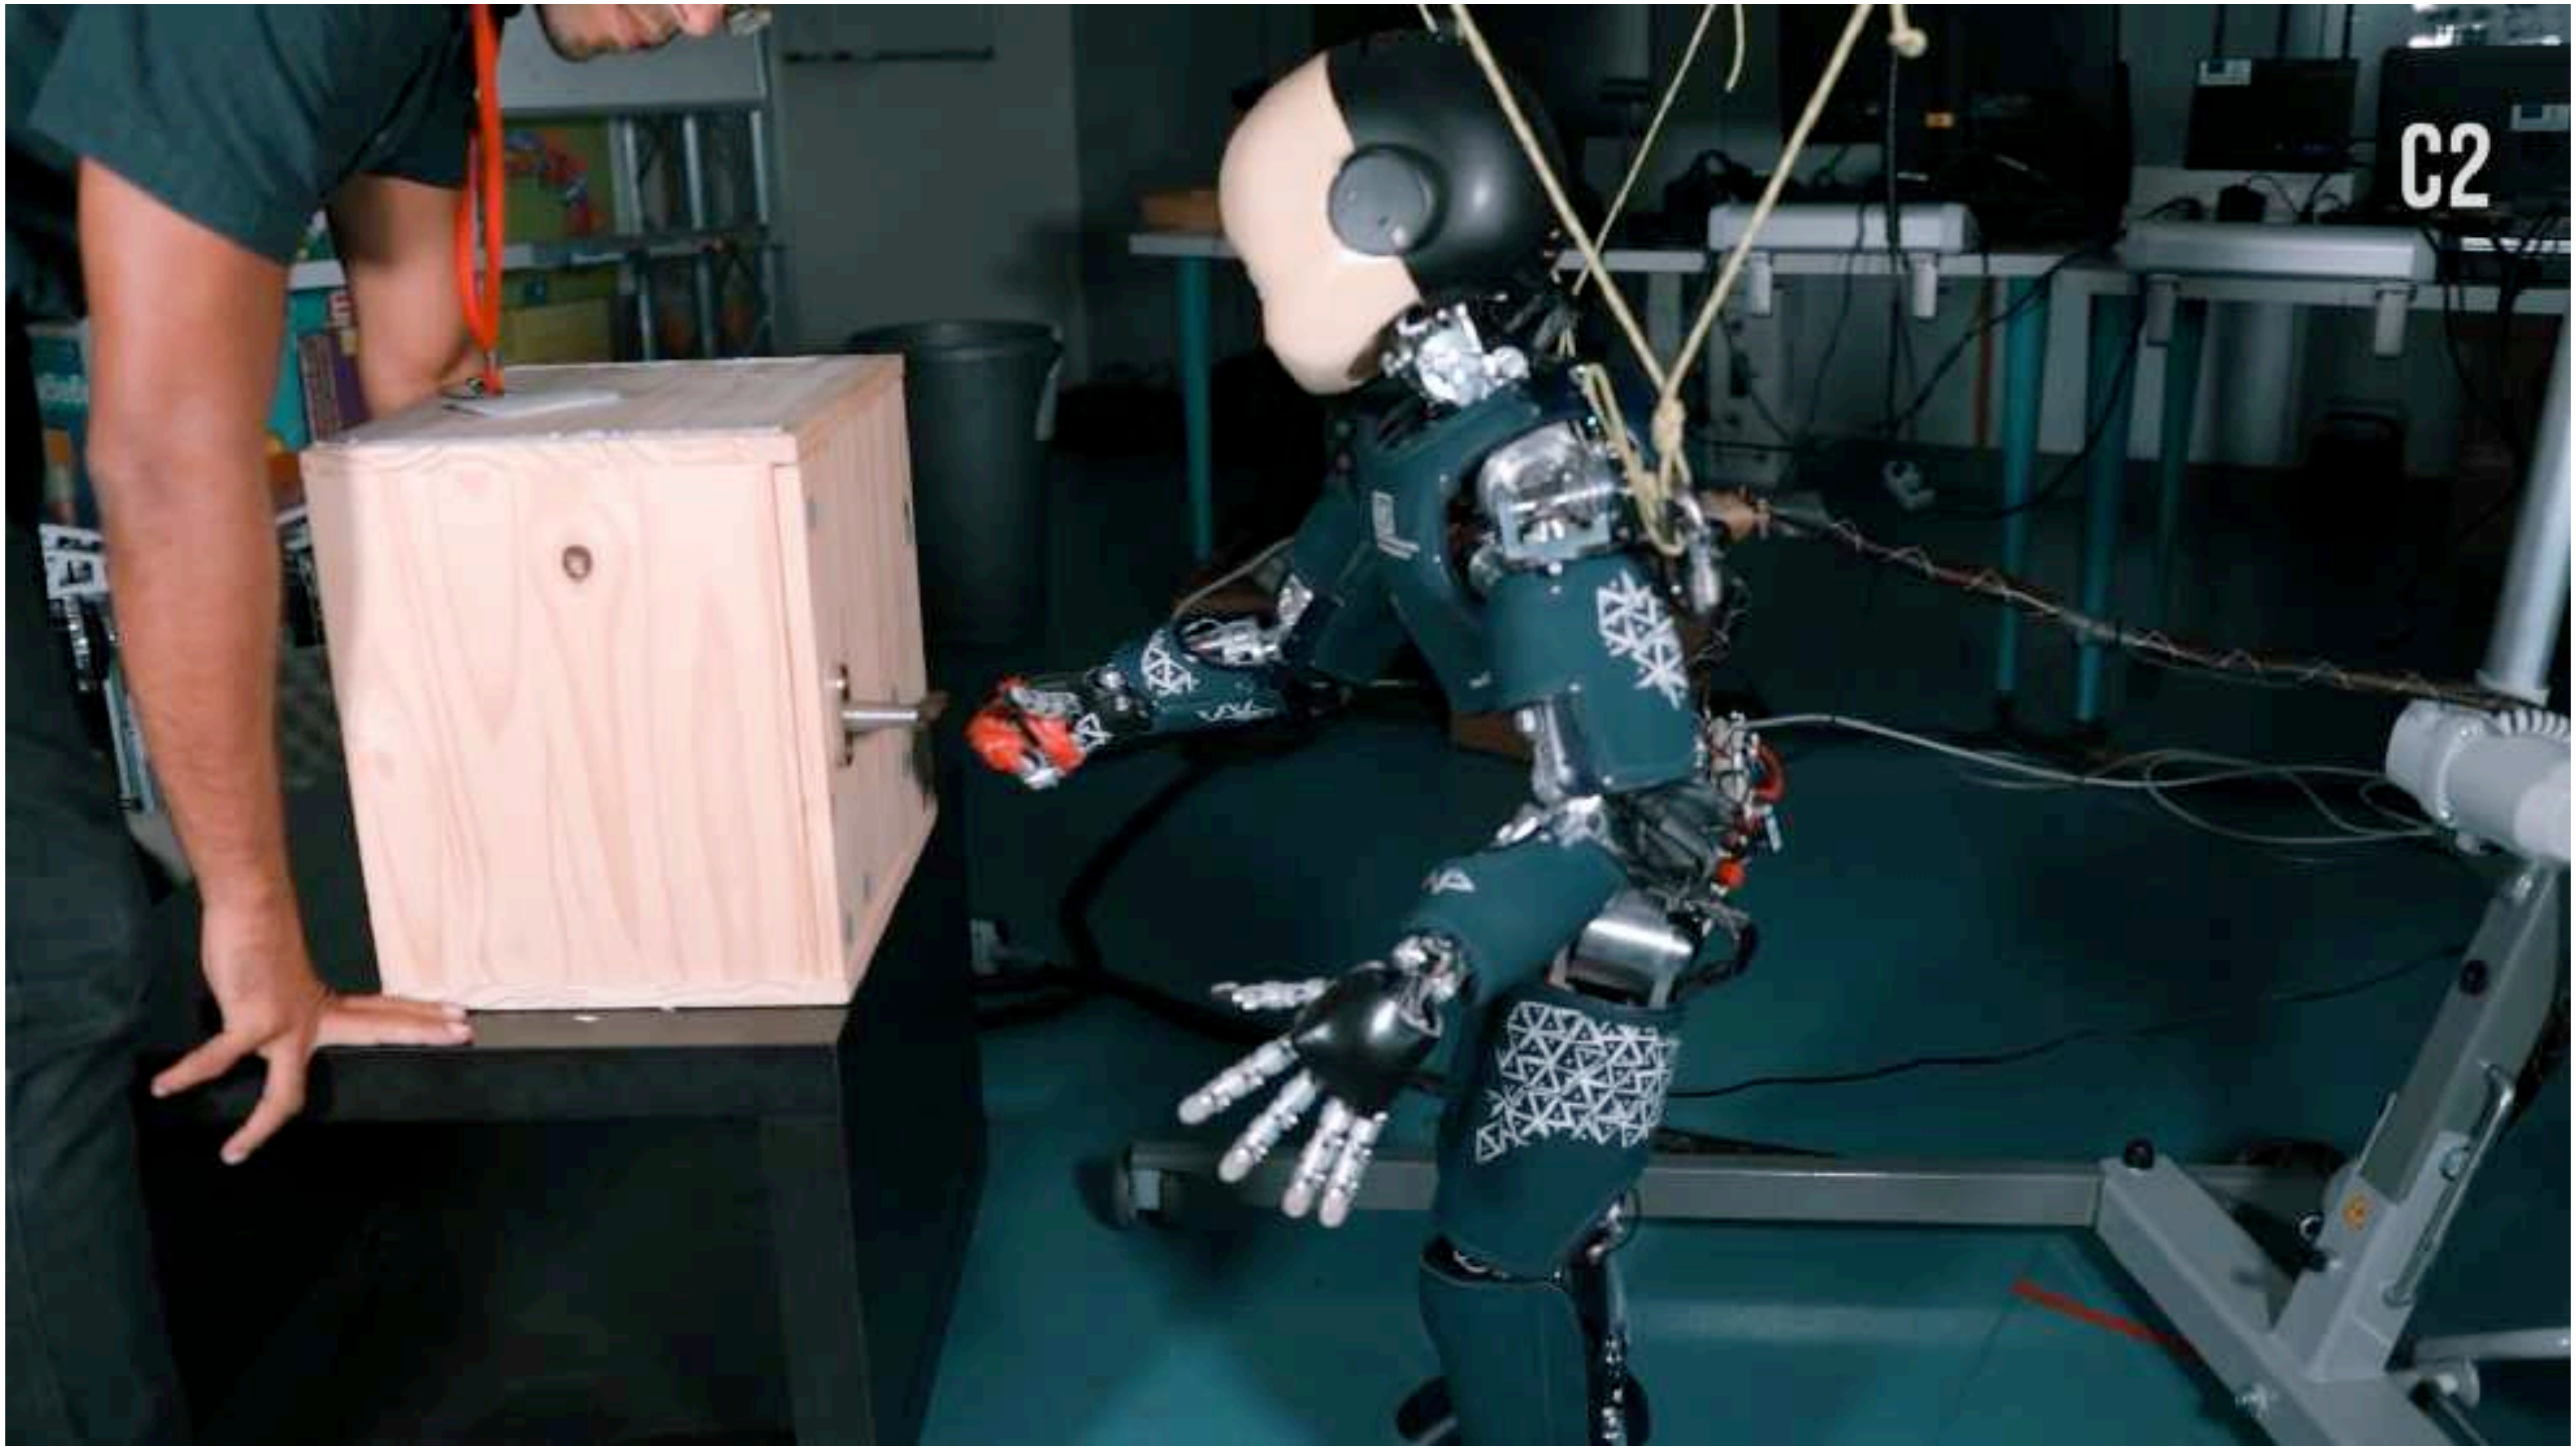

Penco et al. - under review

Whole-body co-manipulation with a human

- Take into account the entire human dynamics in a multi-task QP controller for collaborative manipulations
- Joint level controller for the robot, but capable of reacting to the human

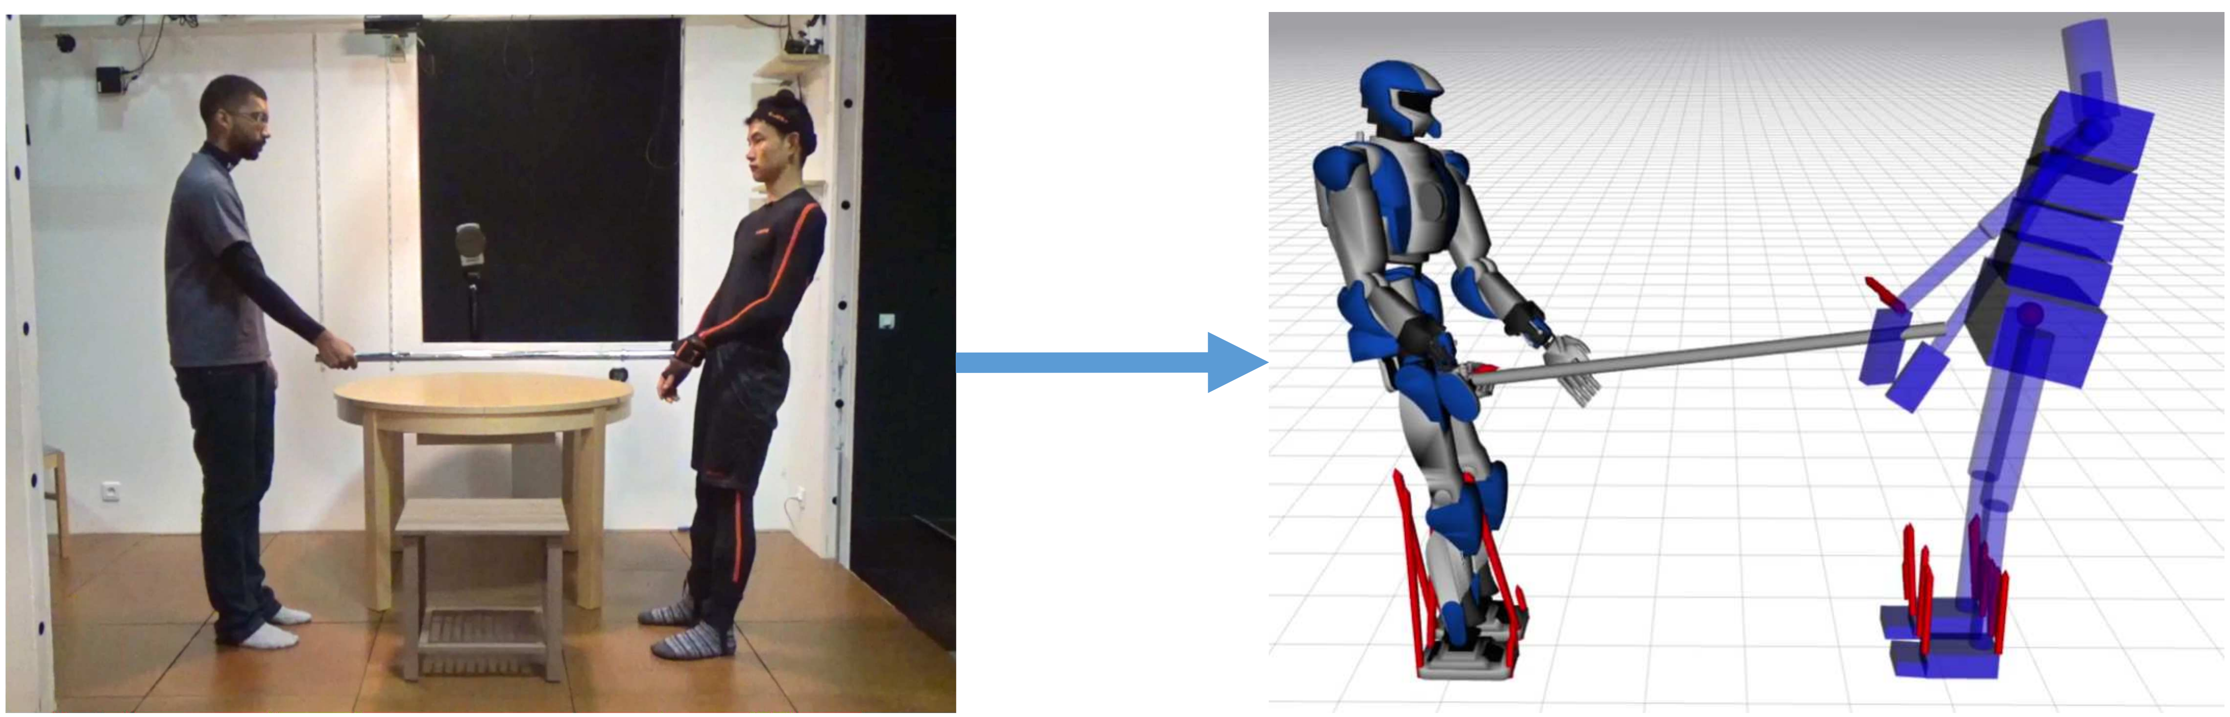

Human collaborator (to be replaced by robot) *Recorded human* *Robot collaborator* *Simulated human*

K. Otani, K. Bouyarmane, S. Ivaldi (2018) ICRA

Whole-body co-manipulation with a human

Proposed approach:

- Model the human as a robot ➞ multi-robot QP controller
- Reason in terms of balance of the couple human+robot, not robot only

$$\underset{\ddot{q},\tau,f}{\text{minimize}} \quad \sum_k w_k ||\ddot{g}_k - \ddot{g}_k^{des}||^2$$

Individual or combined tasks (e.g. combined CoM for balance)

$$\text{subject to} \quad M\ddot{q} + N = J_0^T F^0 + (J_1 - \Psi^T J_2)F^- + S\tau$$

$$J_0 \dot{q} = 0$$

Equal and opposite contact forces between robots

$$(J_1 - \Psi^T J_2)\dot{q} = 0$$

Non-slipping contacts between robots

$$f \in \mathcal{C}$$

torque limits, joint limits, collision avoidance

K. Otani, K. Bouyarmane, S. Ivaldi (2018) ICRA

Whole-body co-manipulation with a human

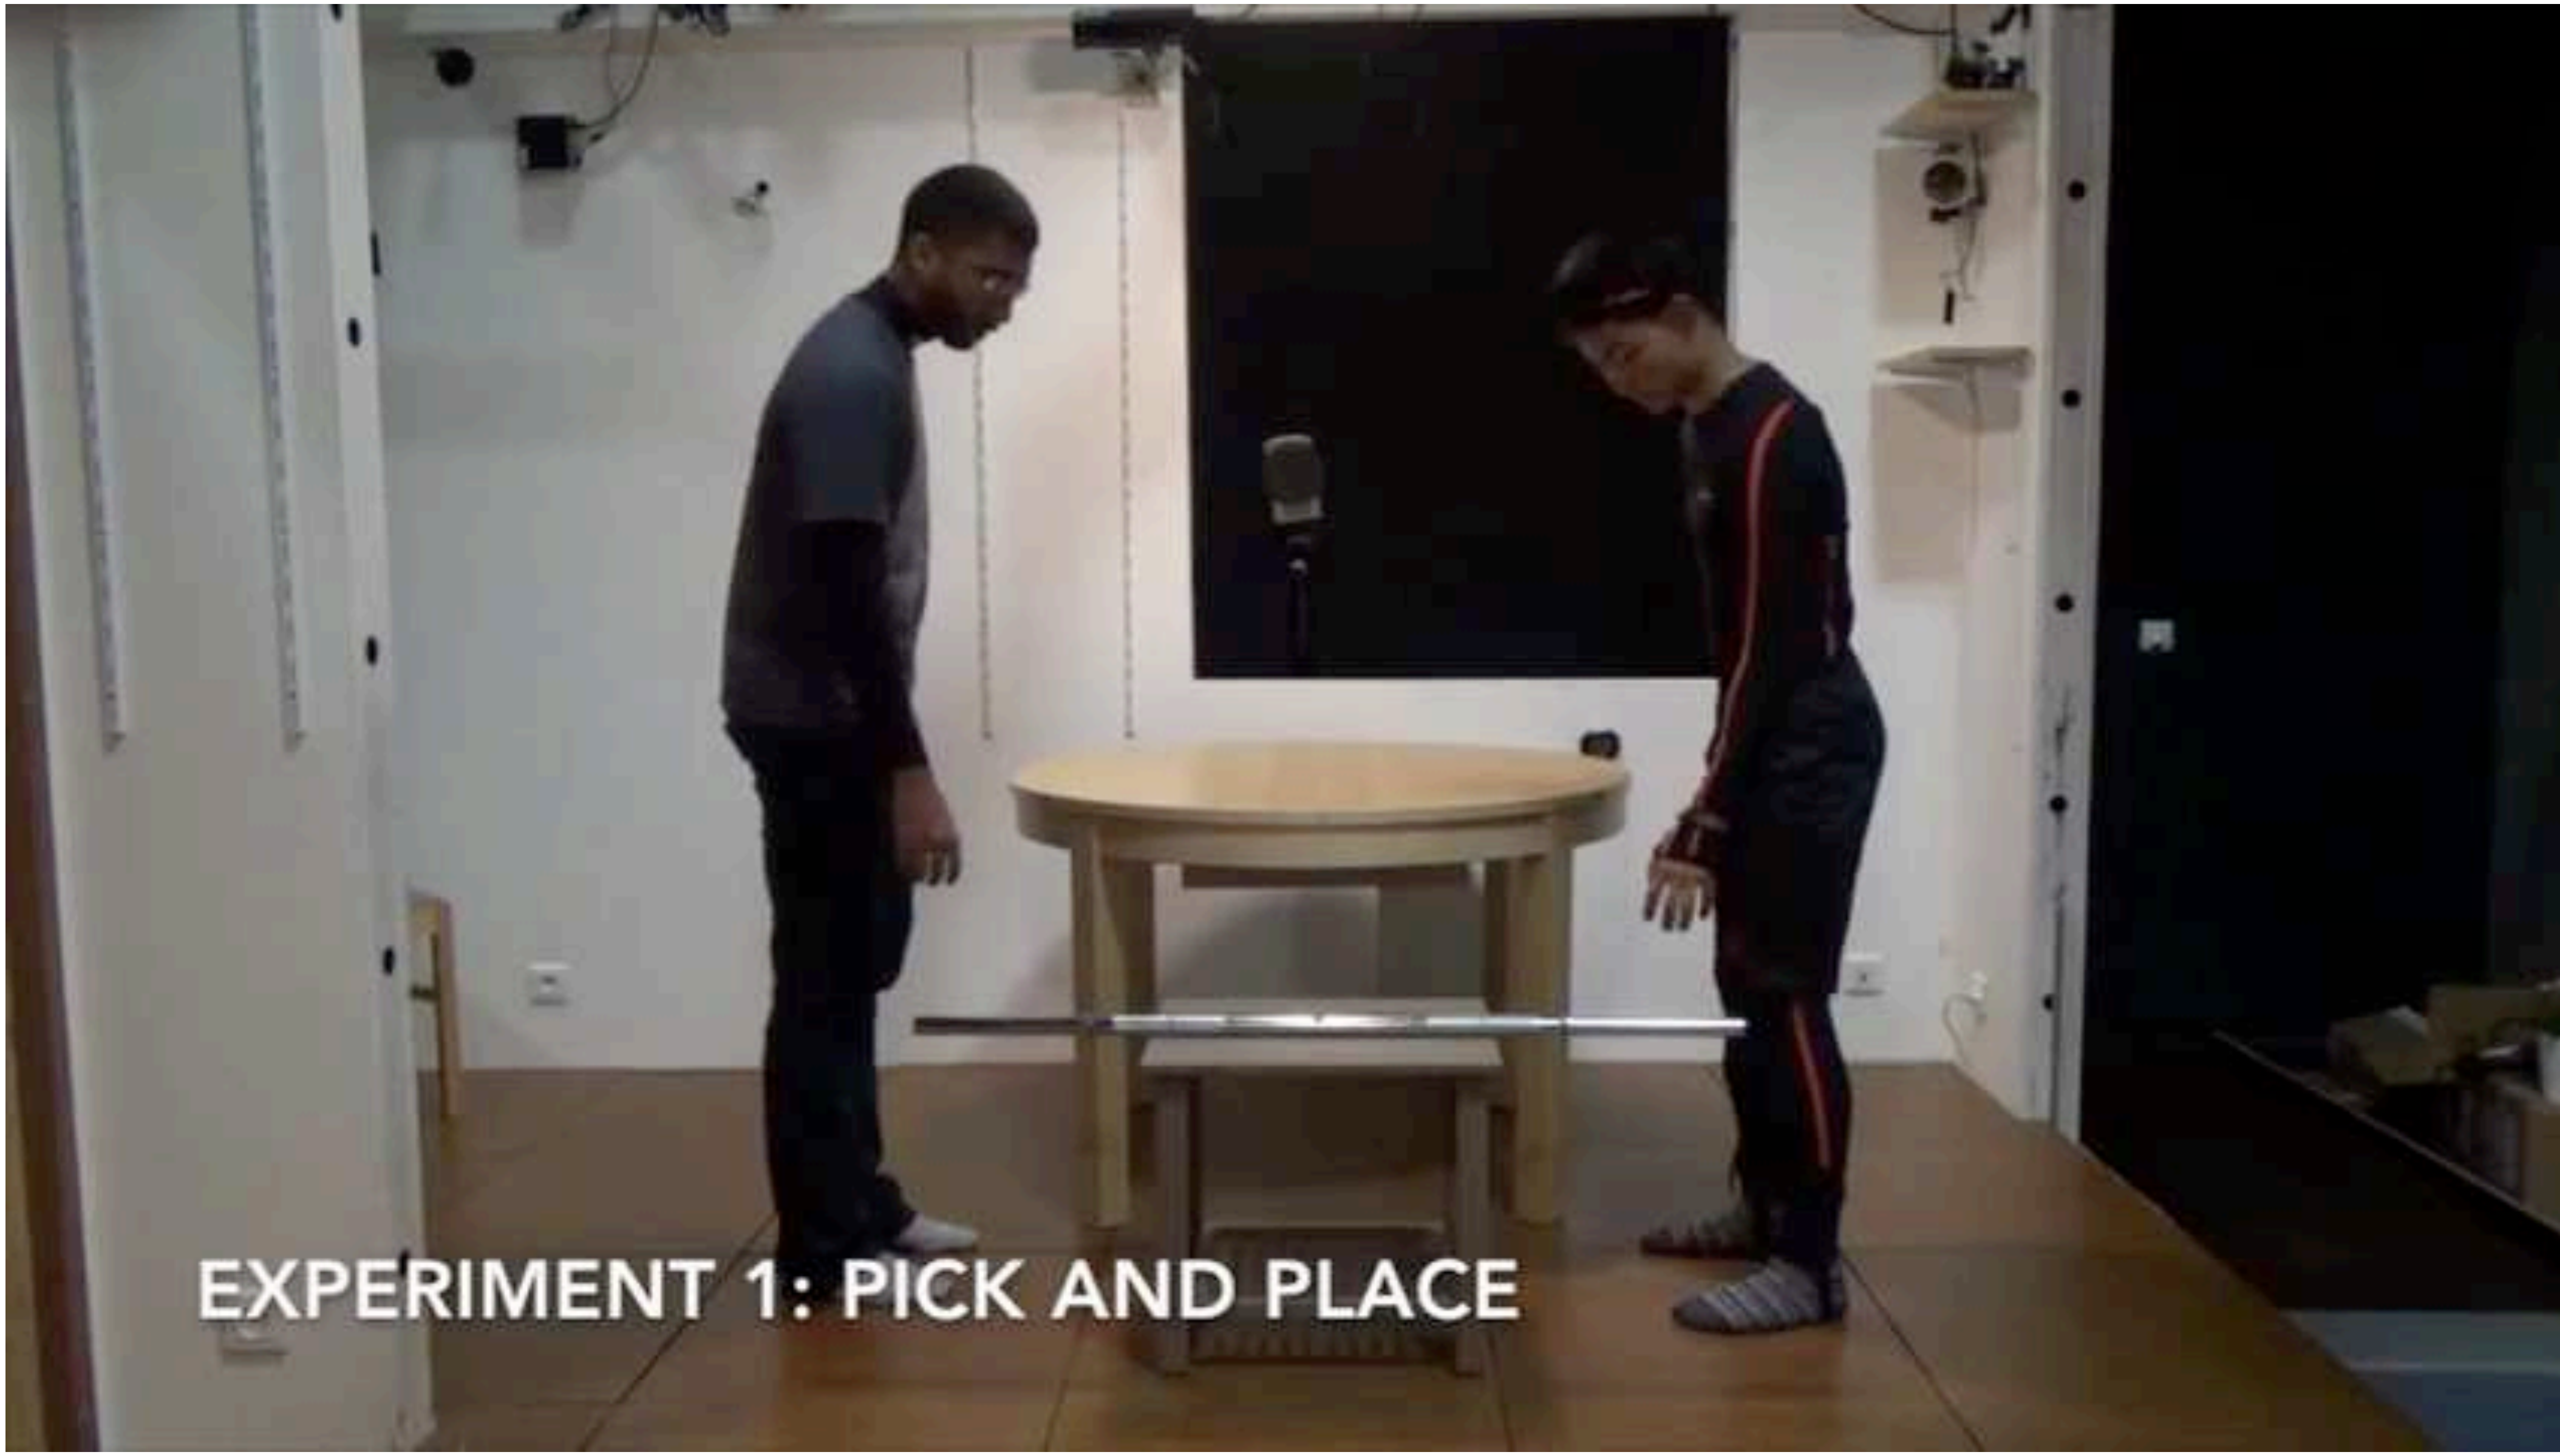

K. Otani, K. Bouyarmane, S. Ivaldi (2018) ICRA

Outline

Learning the visual appearance of objects from scratch - with curiosity-based autonomous exploration and social guidance

Learning by demonstration whole-body manipulation on humanoids

HEAP project - Incorporating human preferences to improve grasping of irregular objects in a heap

EU Project HEAP

chist-era

Human guided learning of robotic heap sorting




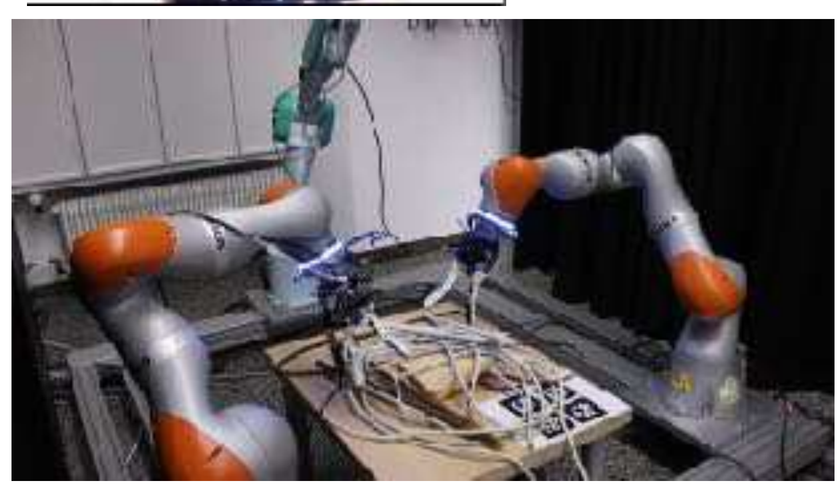


Why do we need human guidance?

HEAP setup & context

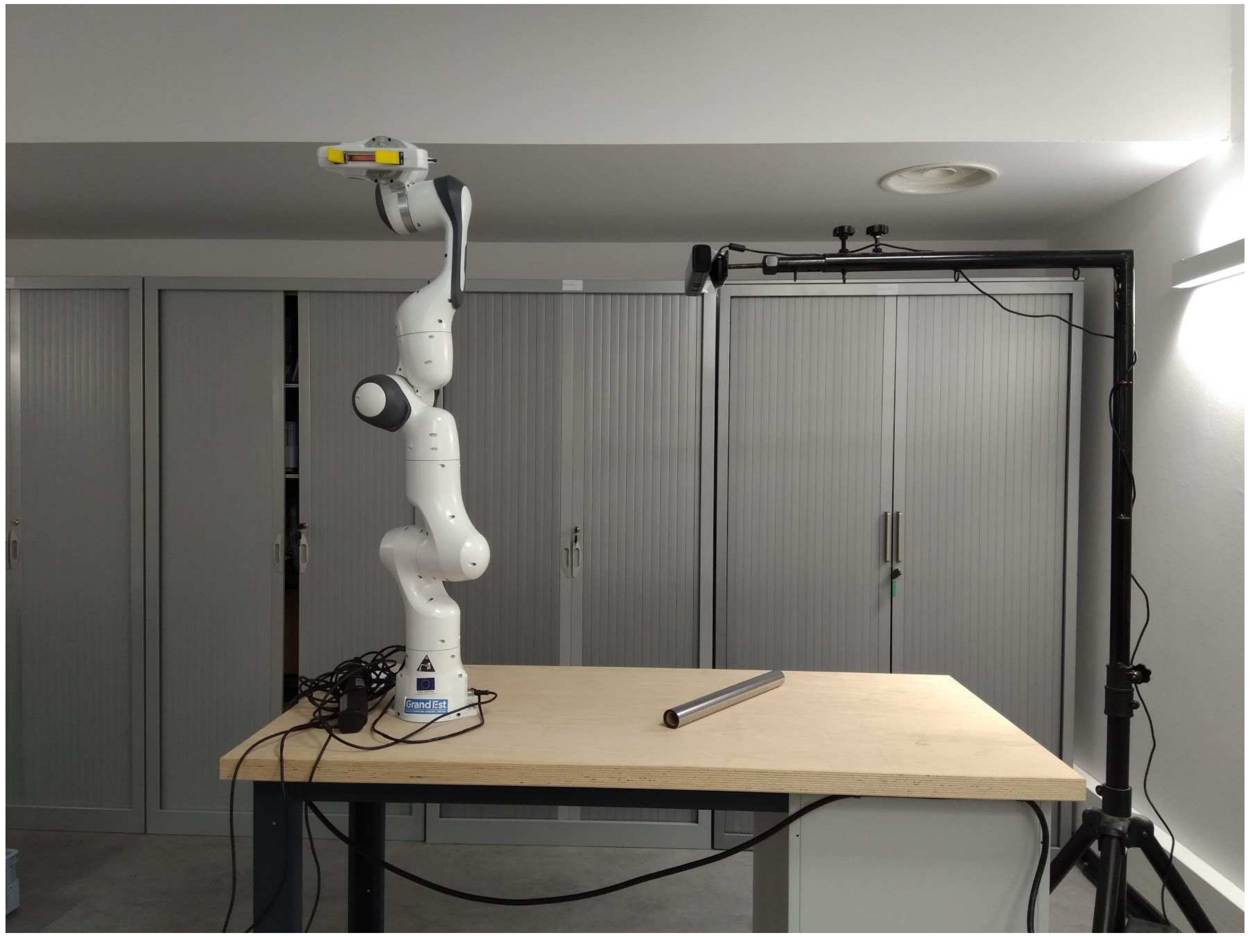

Human guidance / expertise / preference → Grasping algorithms

Why do we need human guidance

We already have plenty of grasping algorithms that we can use to find the best candidate grasp for the objects in the scene…

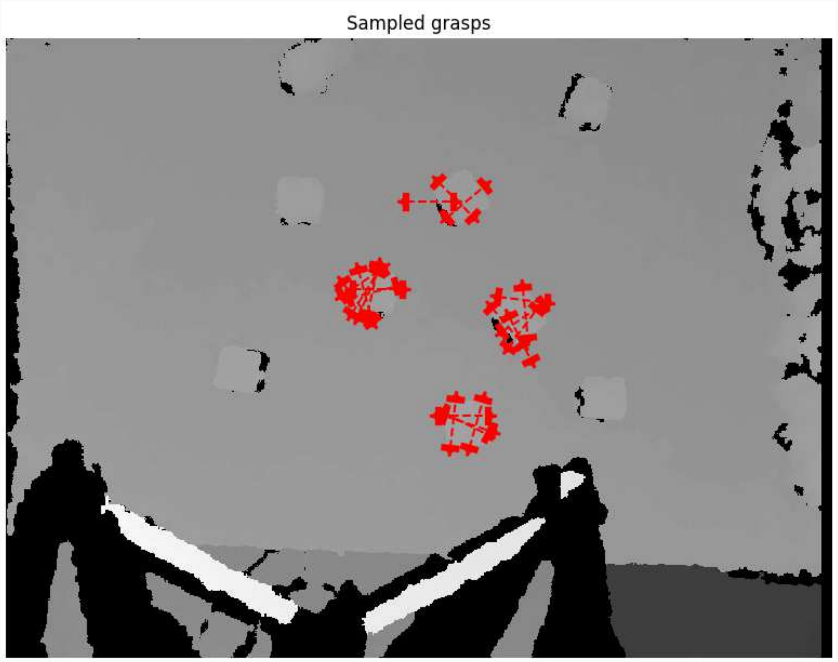

=> Dexnet 2.0
Malher et al., RSS 2017

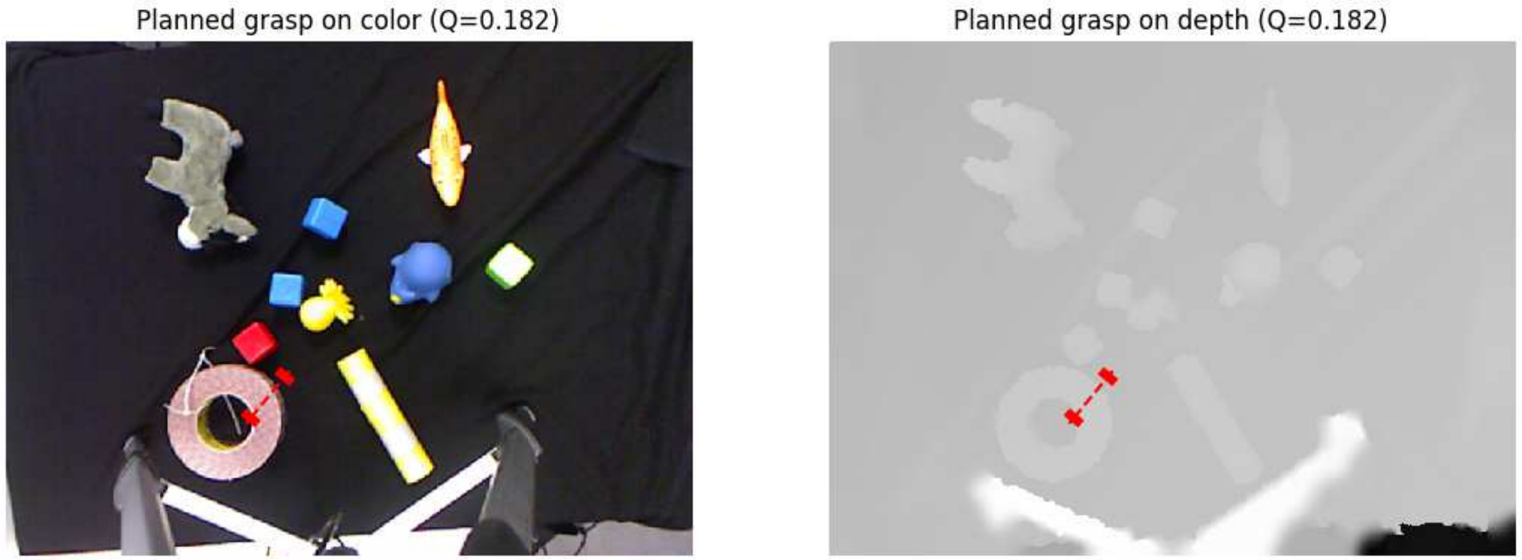

https://github.com/BerkeleyAutomation/gqcnn

Dexnet


Dex-Net 4.0:
Learning Ambidextrous Robot Grasping Policies
AUTOLAB
Science Robotics Journal 2019
berkeleyautomation.github.io/dex-net

Dexnet

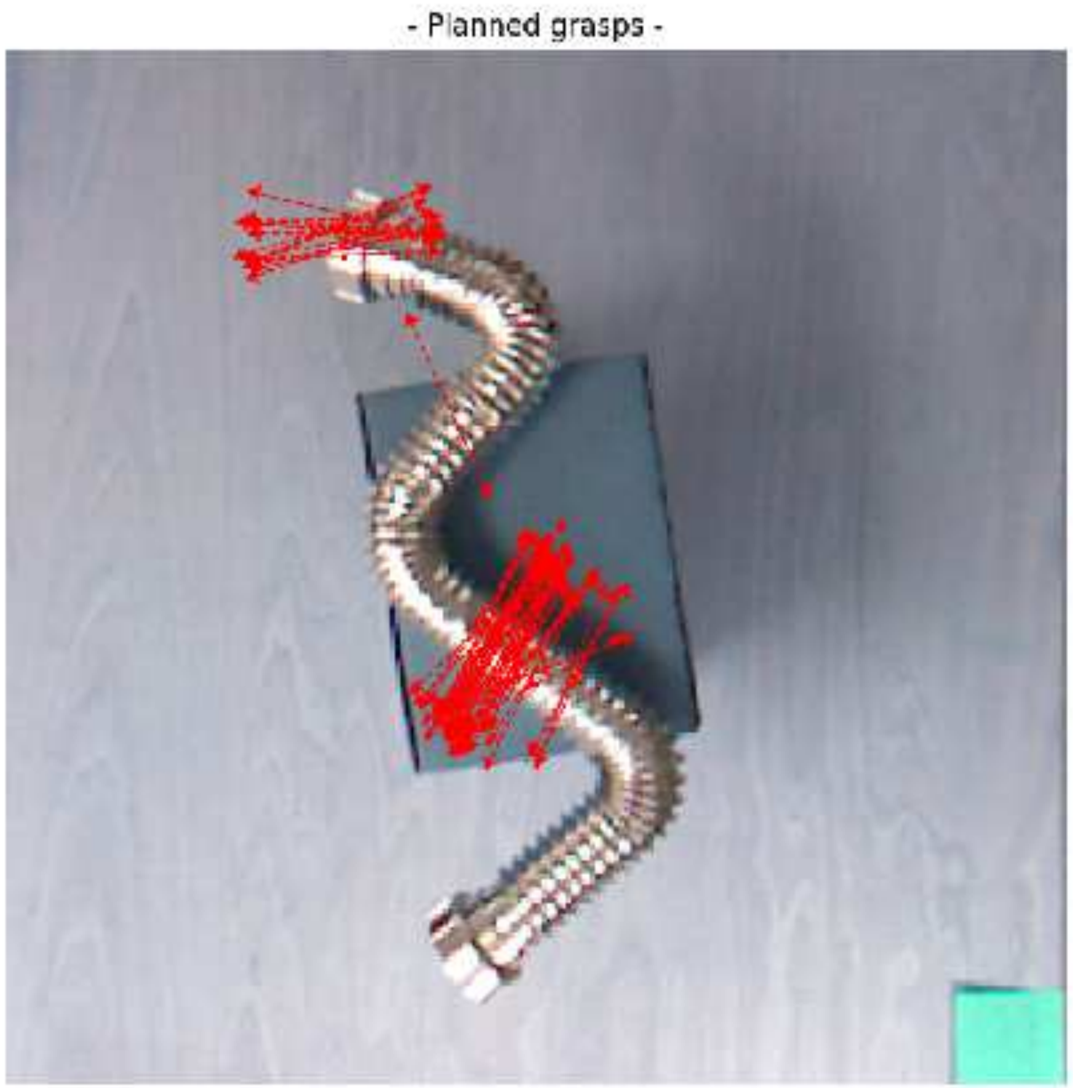

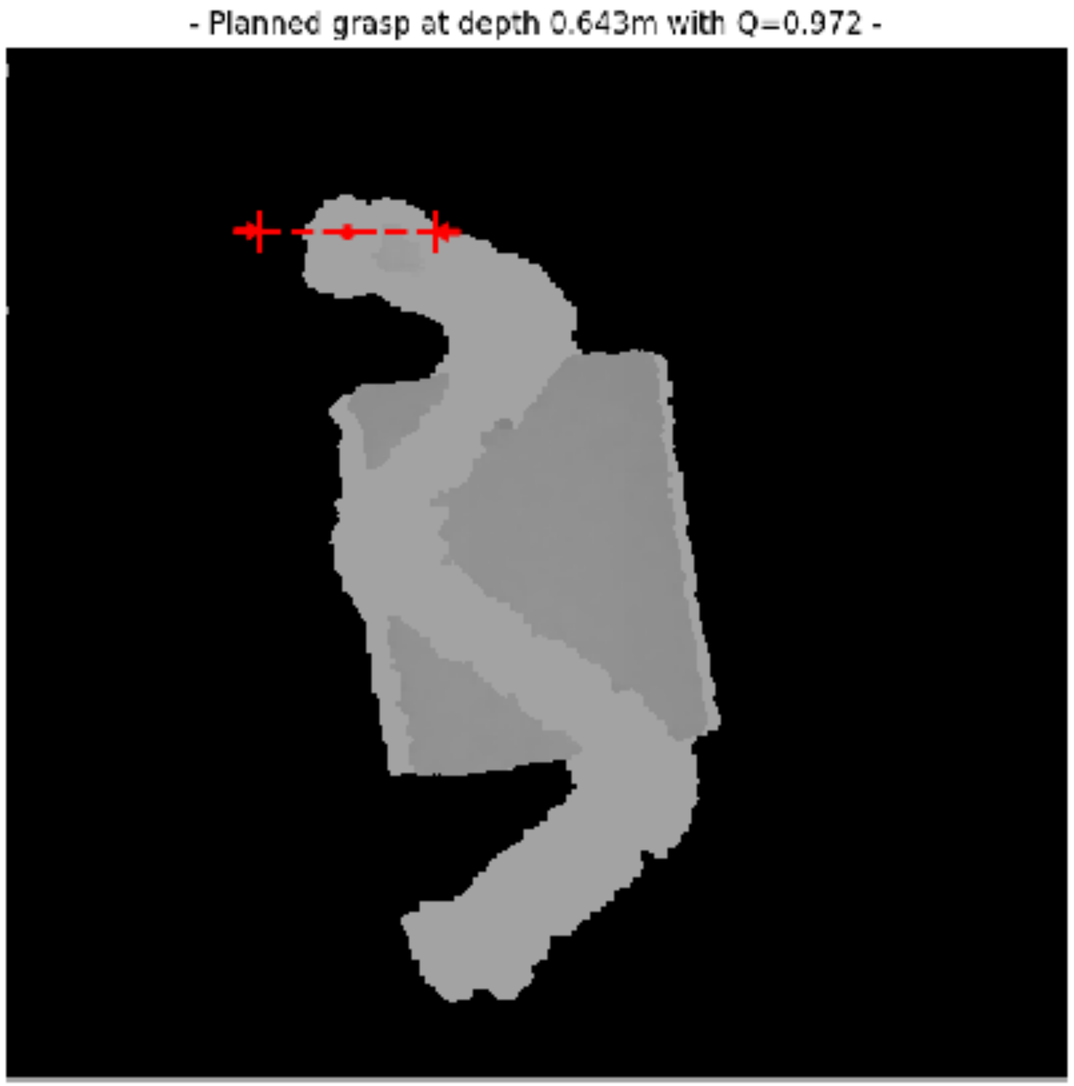

=> Dexnet 2.0, Malher et al., RSS 2017

https://github.com/BerkeleyAutomation/gqcnn

Why do we need human guidance

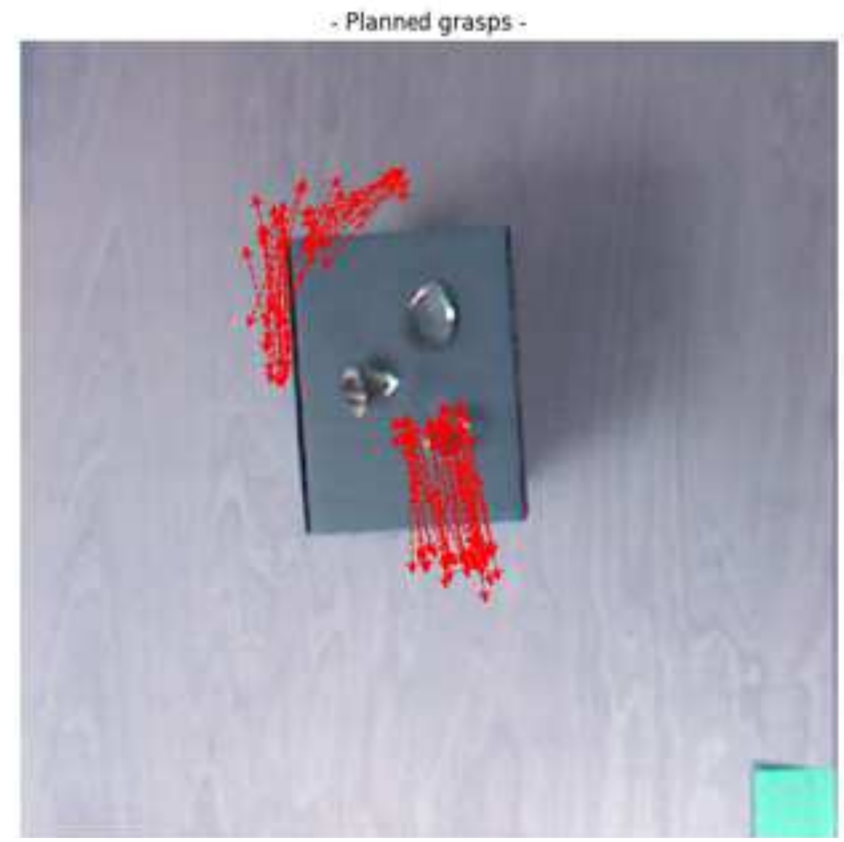

1) Because some objects challenge our cameras

➡ Crystals

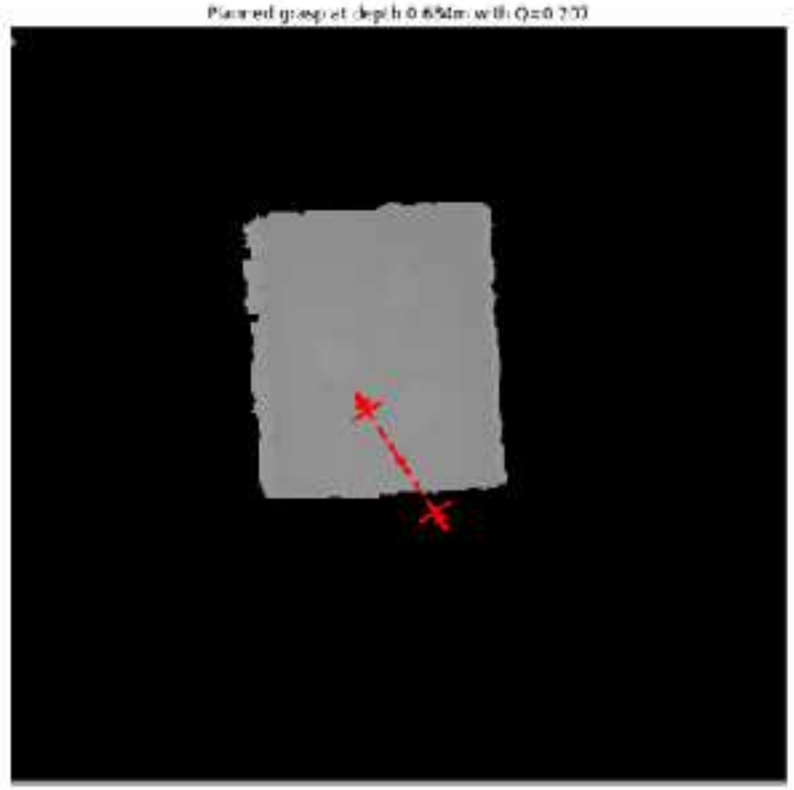

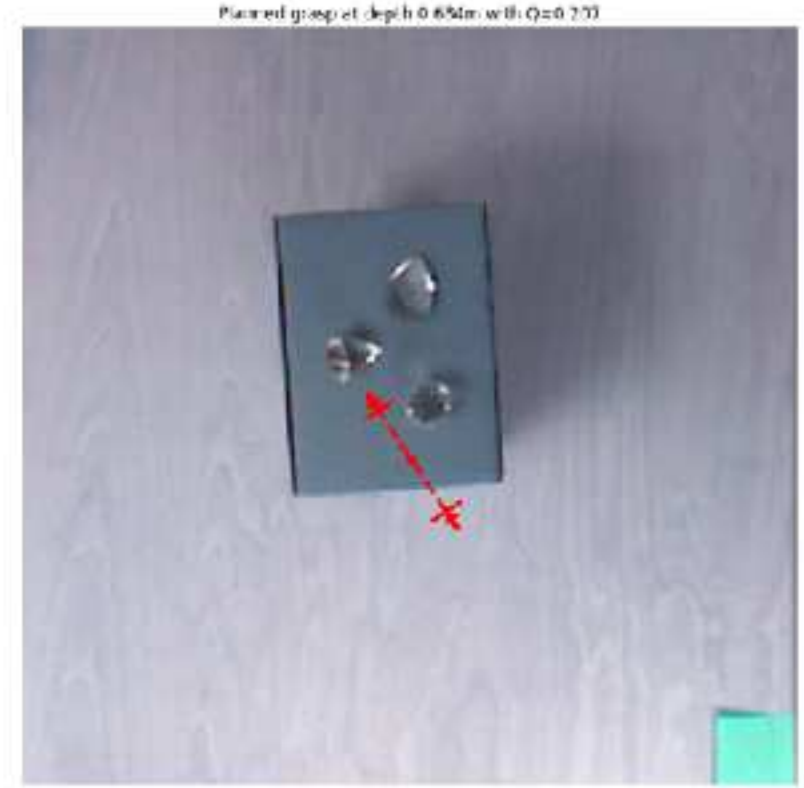

Why do we need human guidance

1) Because some objects challenge our cameras

➡Mirror

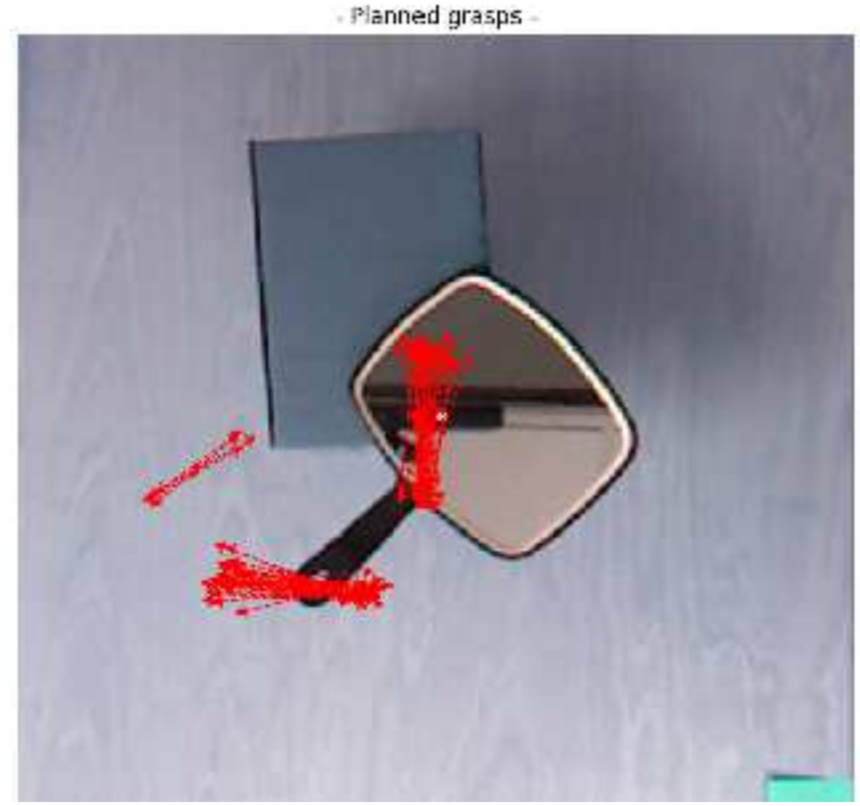

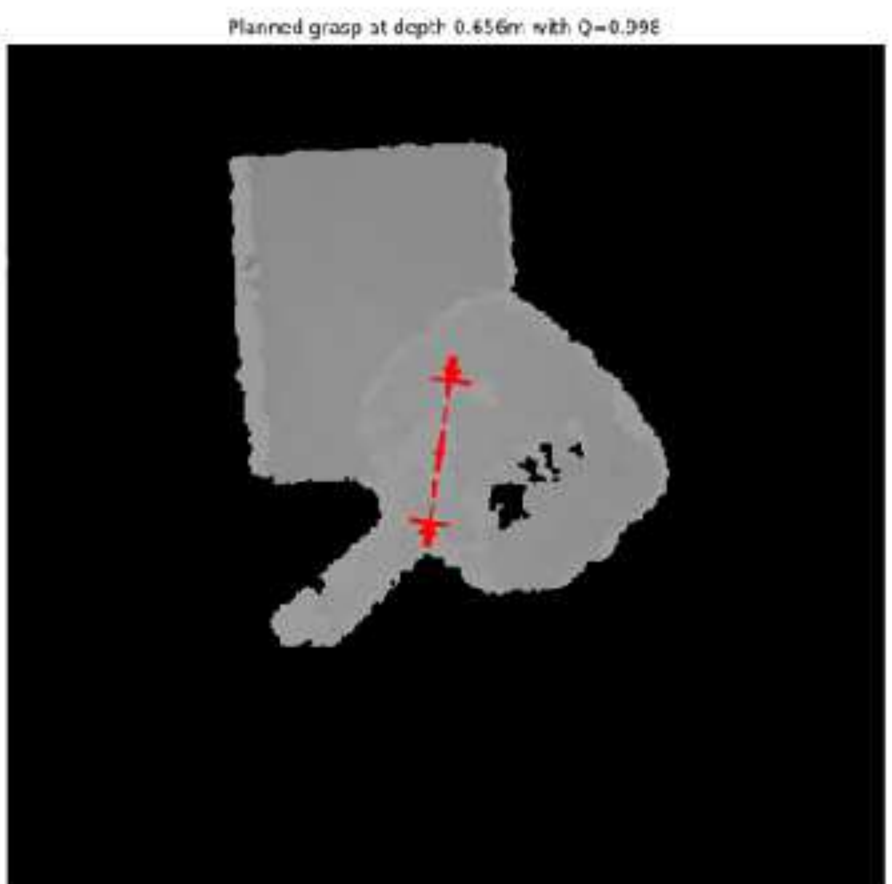

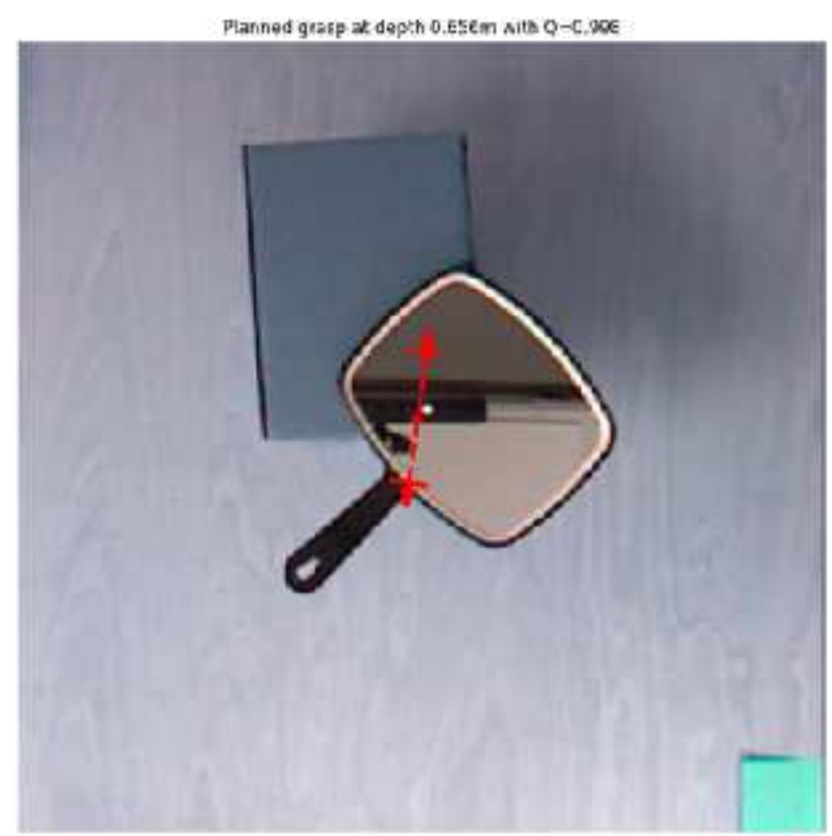

Why do we need human guidance

2) Because the best grasp candidate automatically computed by a grasping algorithm (here: Dexnet 2.0) is not the best choice according to the human, it is not what the human would do

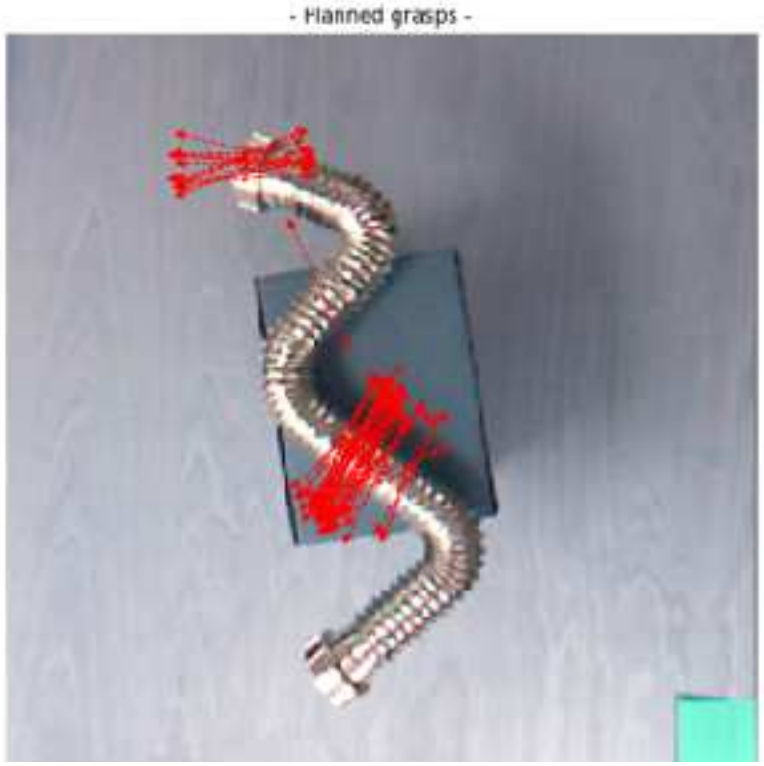

➡Pipe

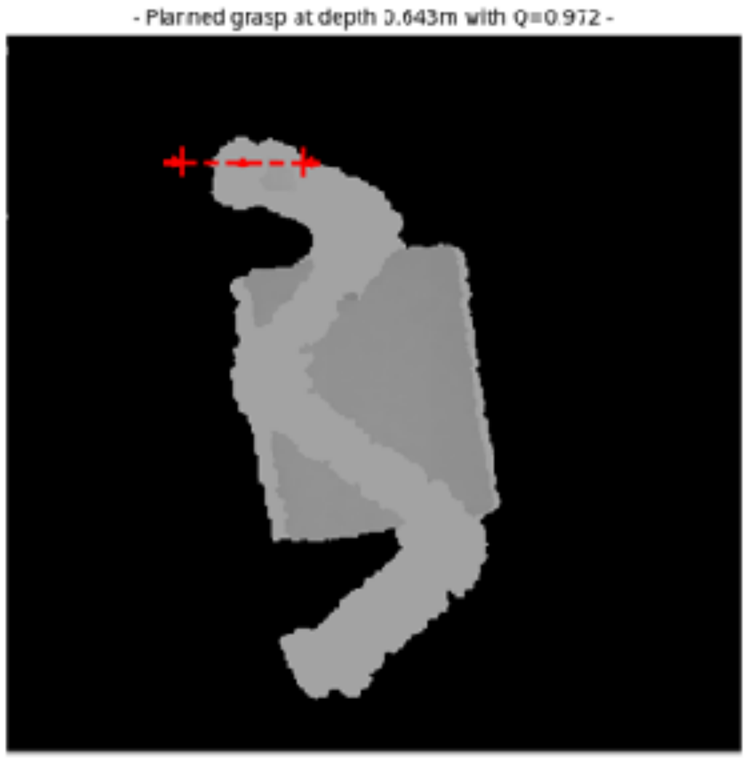

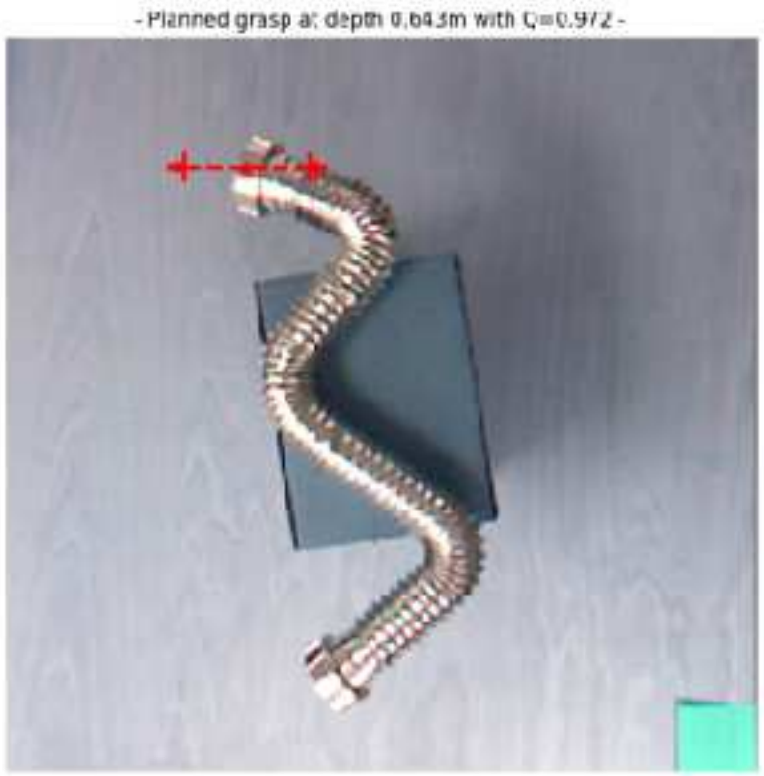

Why do we need human guidance

2) Because the best grasp candidate automatically computed by a grasping algorithm (here: Dexnet 2.0) is not the best choice according to the human, it is not what the human would do

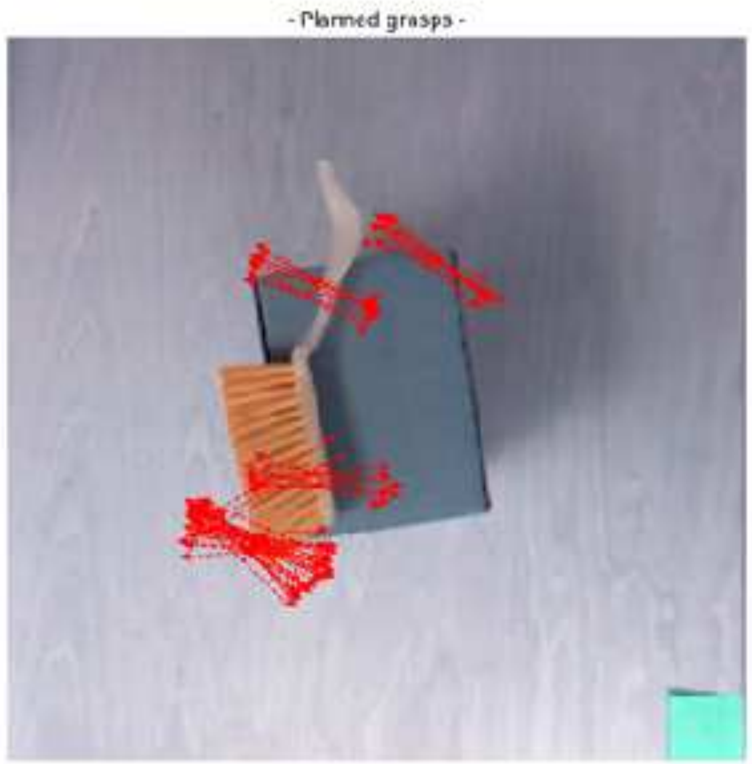

➡Broom

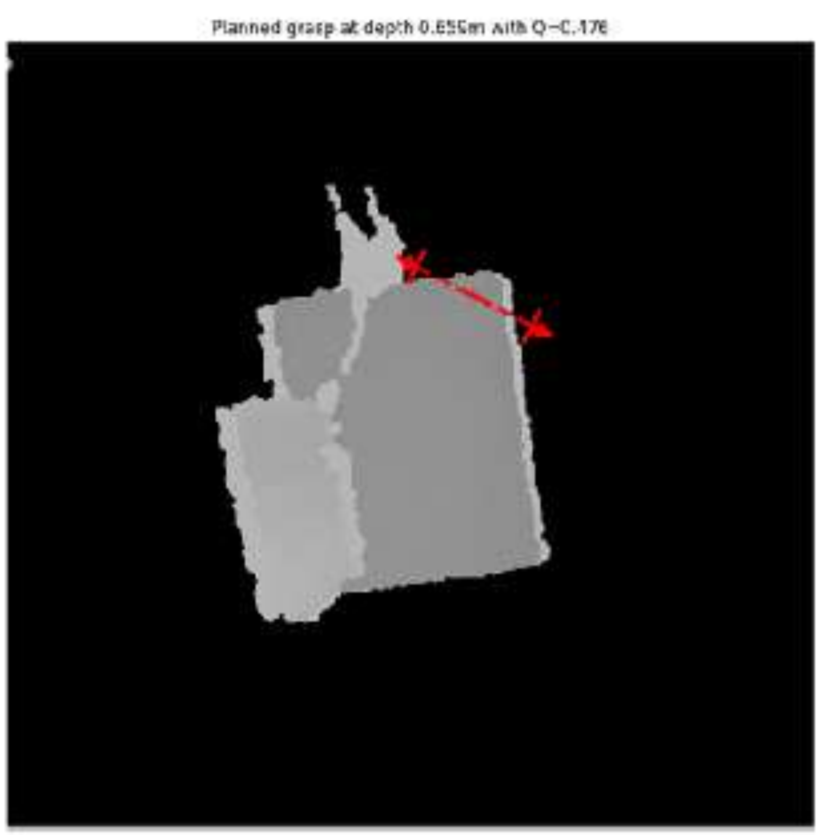

Why do we need human guidance

3) Because grasping algorithms do not reason about the objects fragility

➡ the Duplo tower will break during grasping and transportation

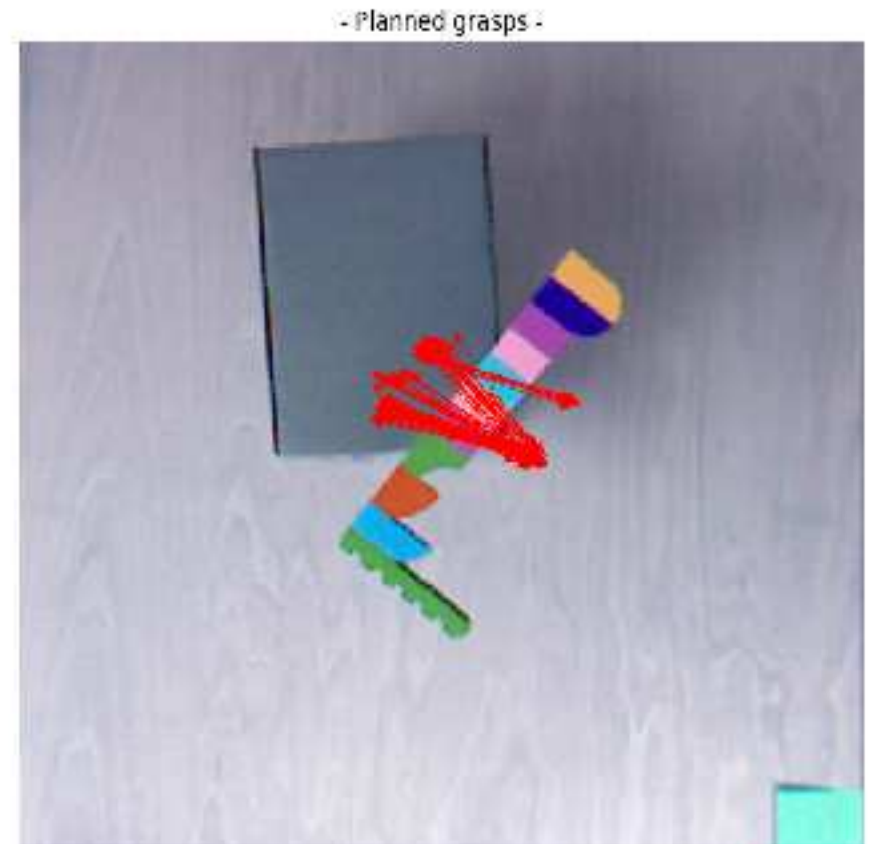

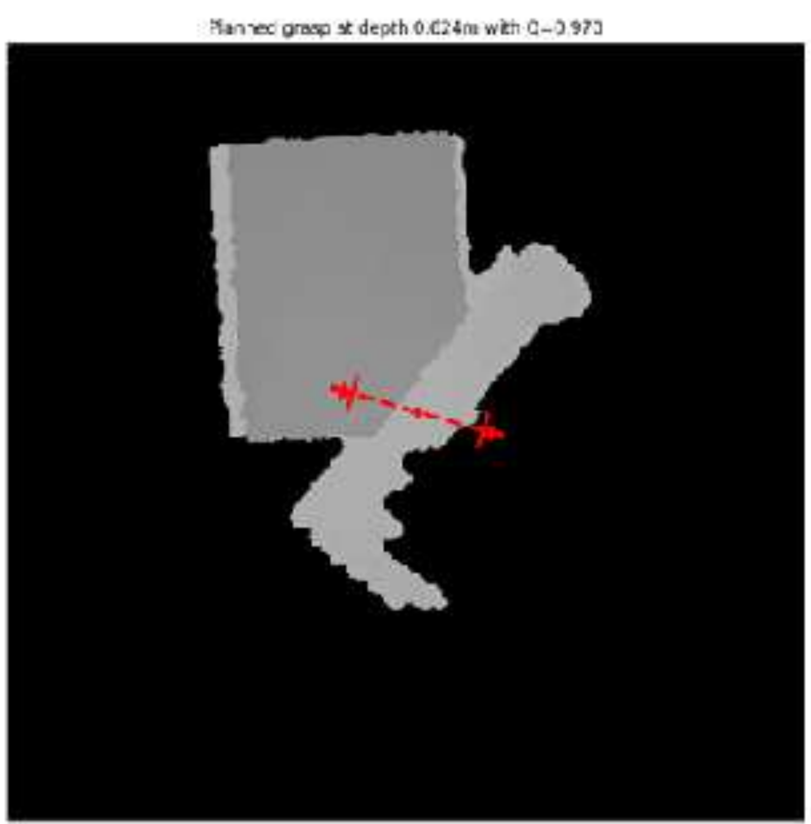

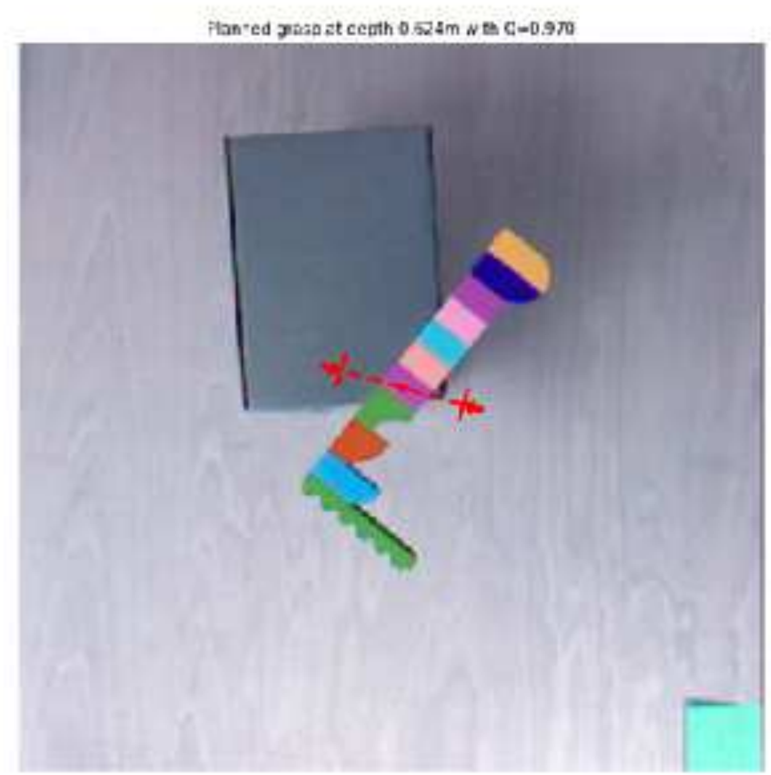

Integrating human preferences

GRASP PREDICTIONS | BAD DEXNET GRASP SELECTION | BETTER DEXNET GRASP SELECTION

Good demonstrations

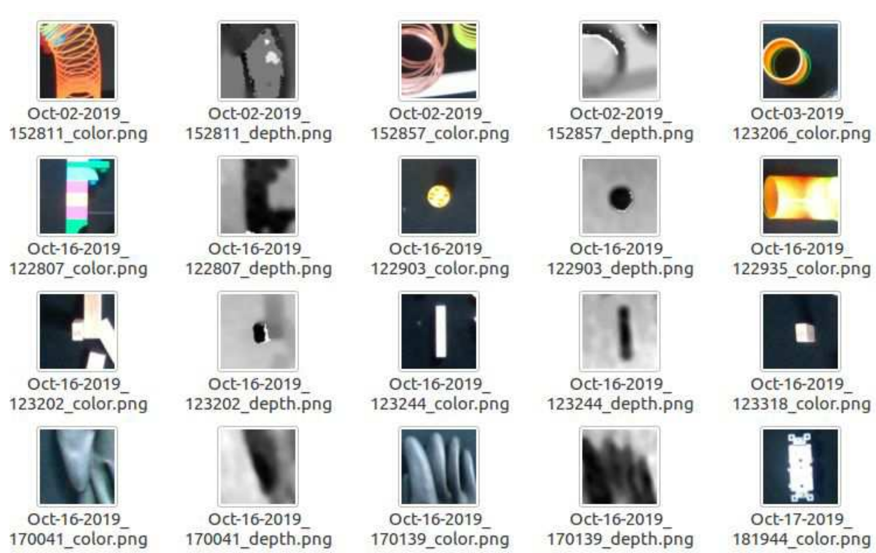

Bad demonstrations

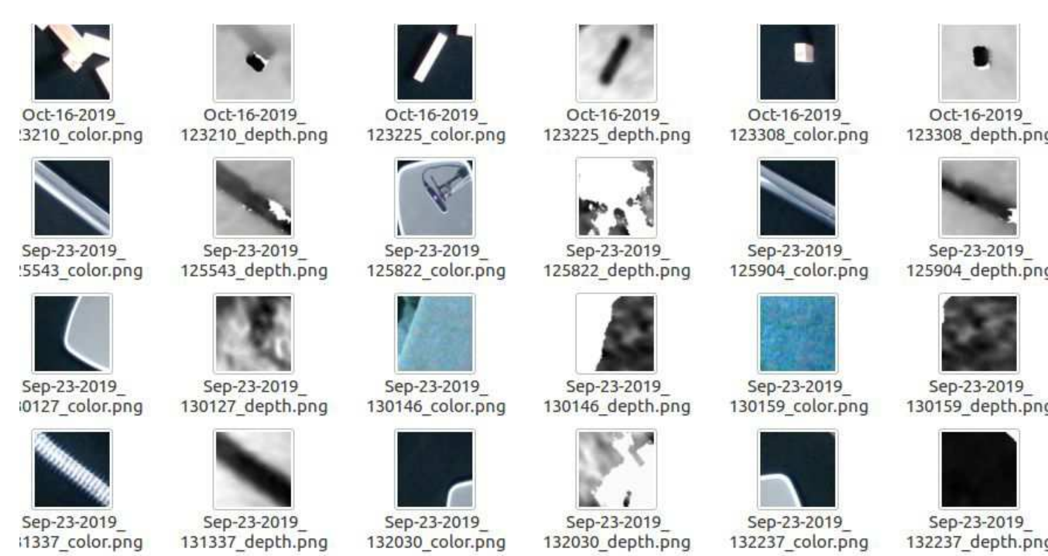

Integrating human preferences

Learning grasp preference models for classes of objects

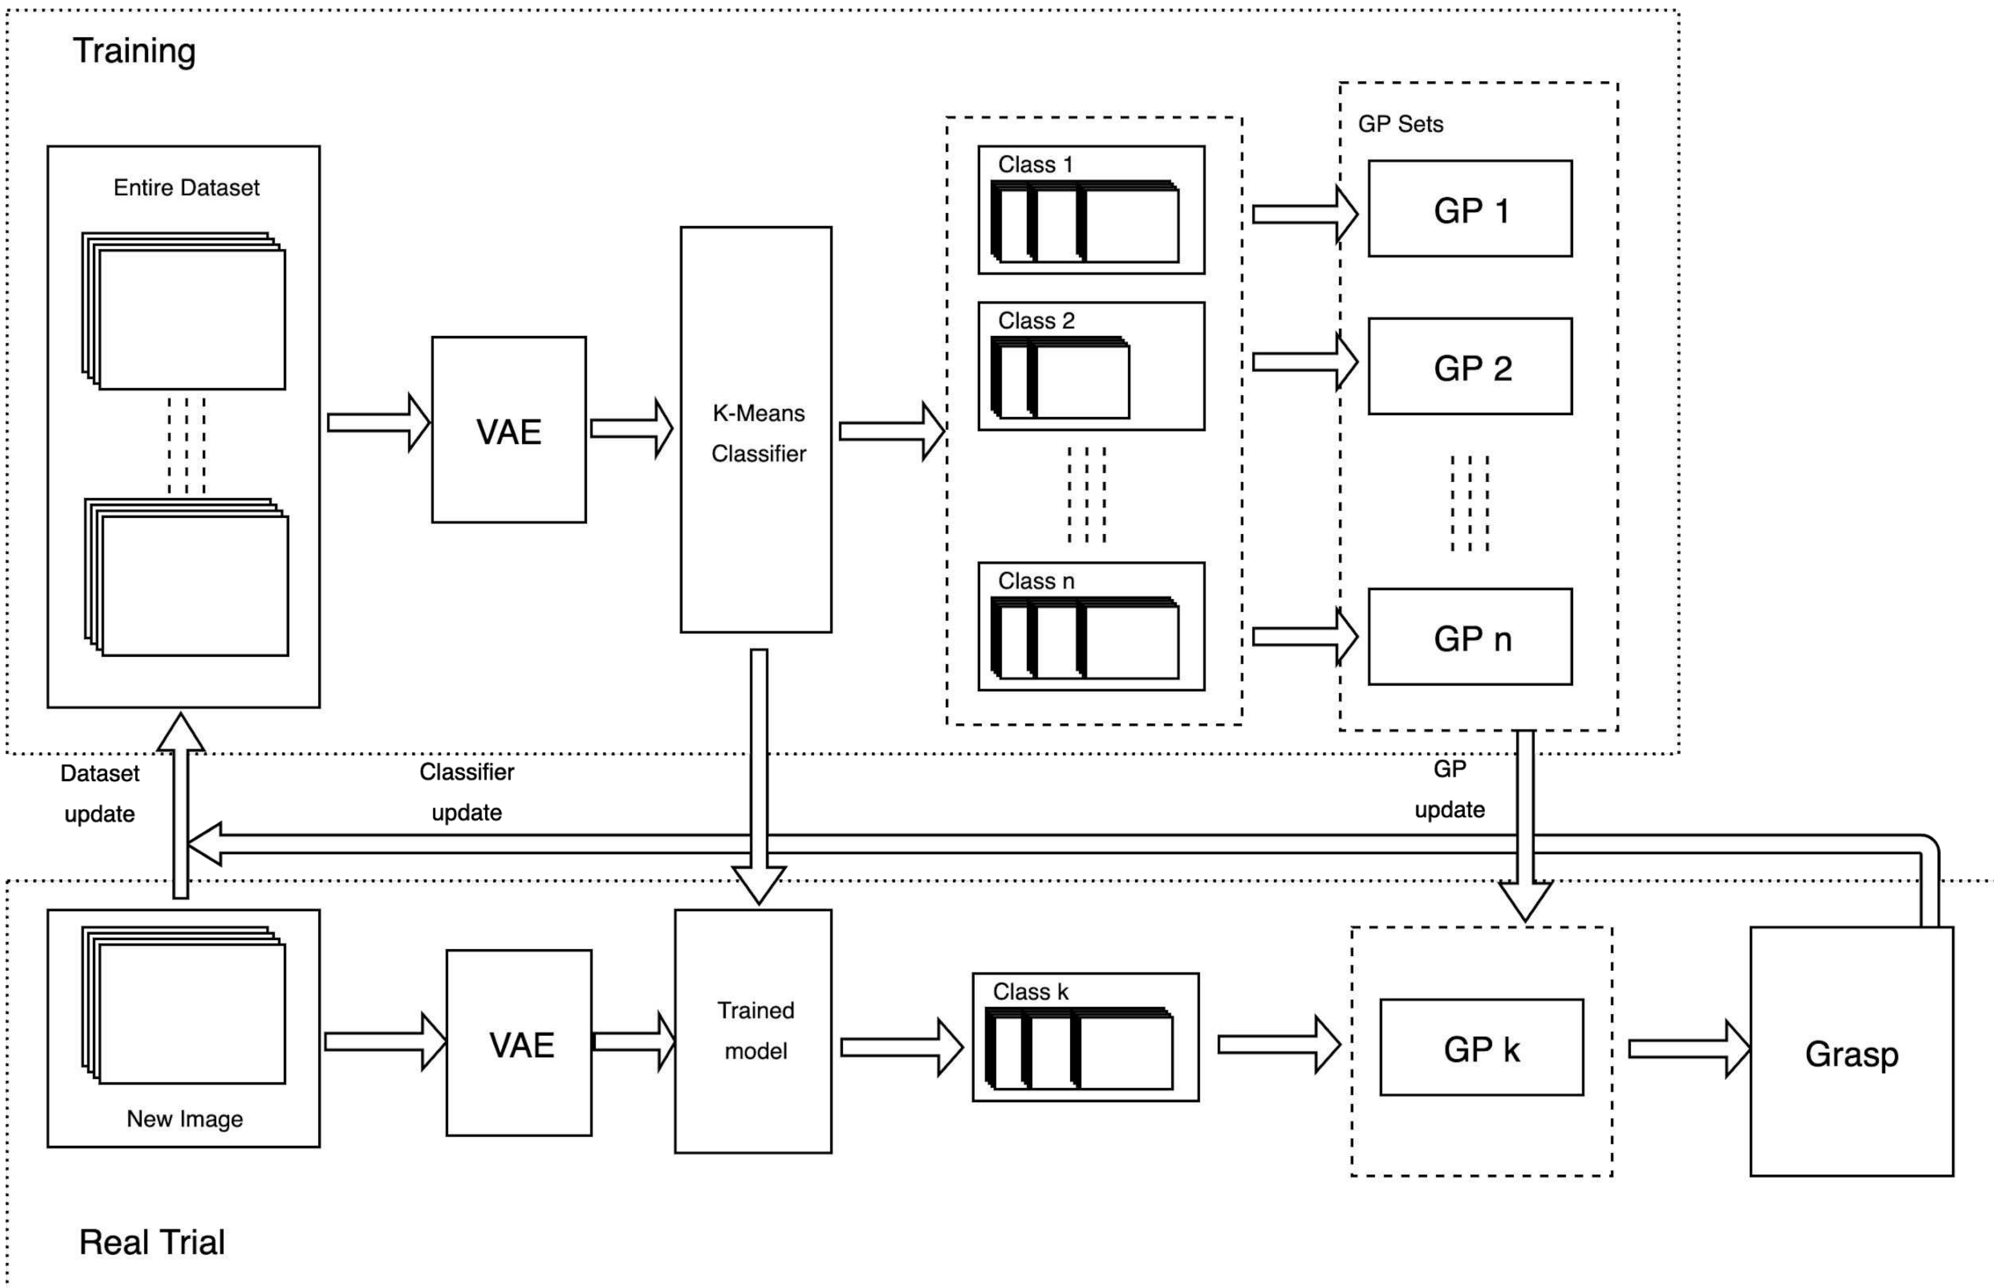

Human guidance & preferences

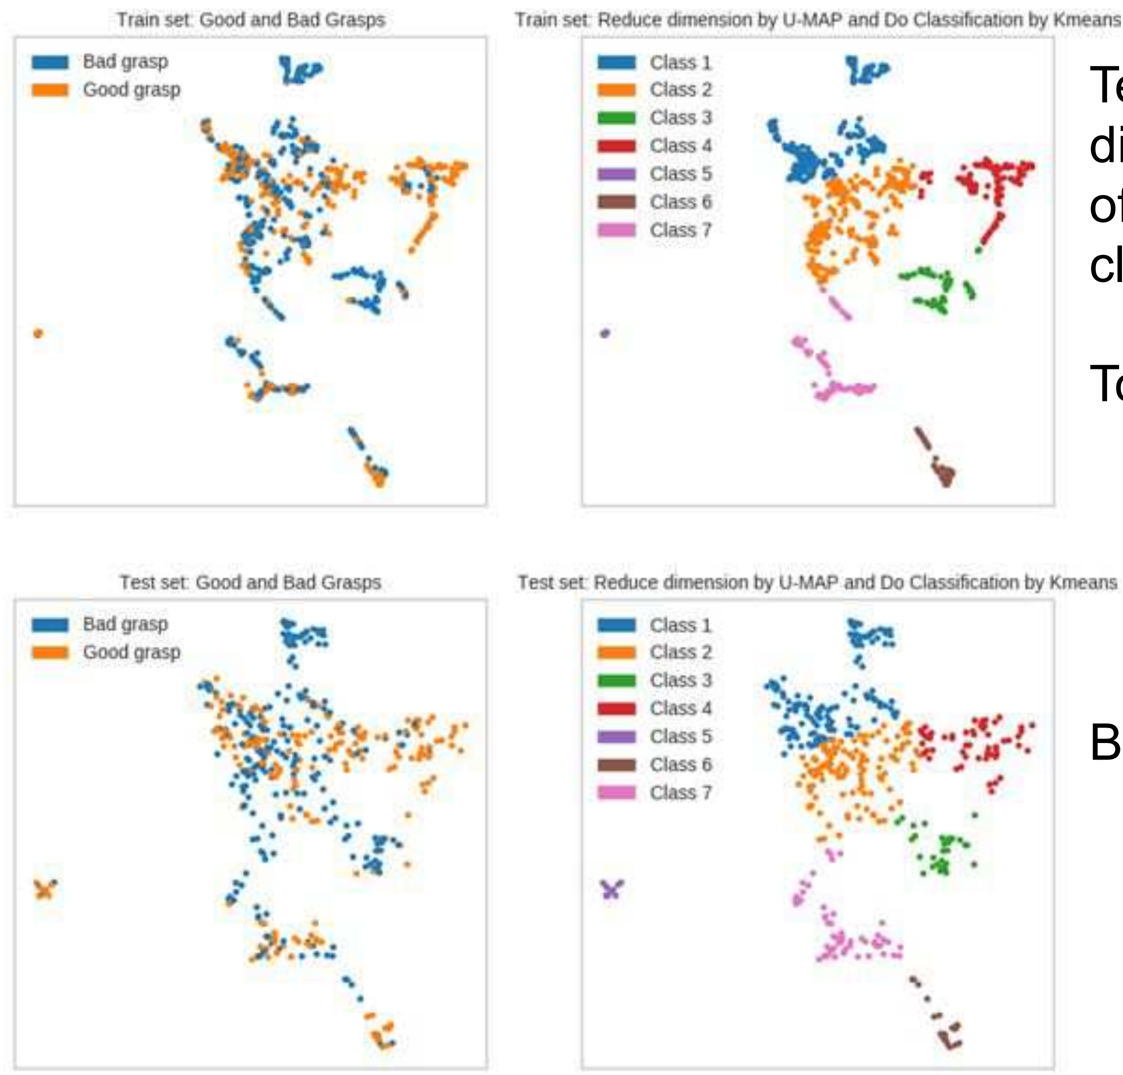

Tests with different number of targeted classes

Top Training

Bottom validation

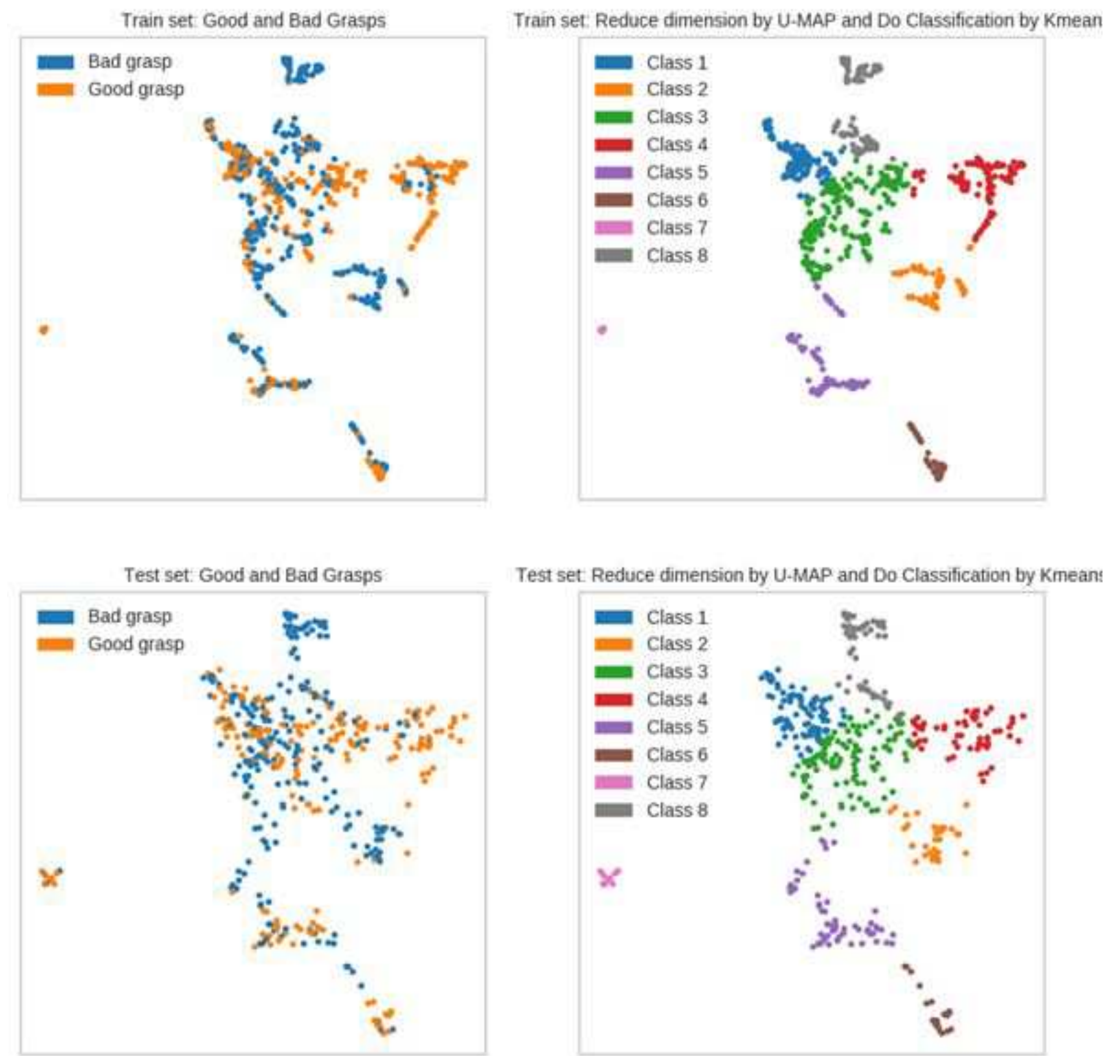

(unsupervised - k-means & u-maps)

Human guidance & preferences

Classes of objects that were automatically created

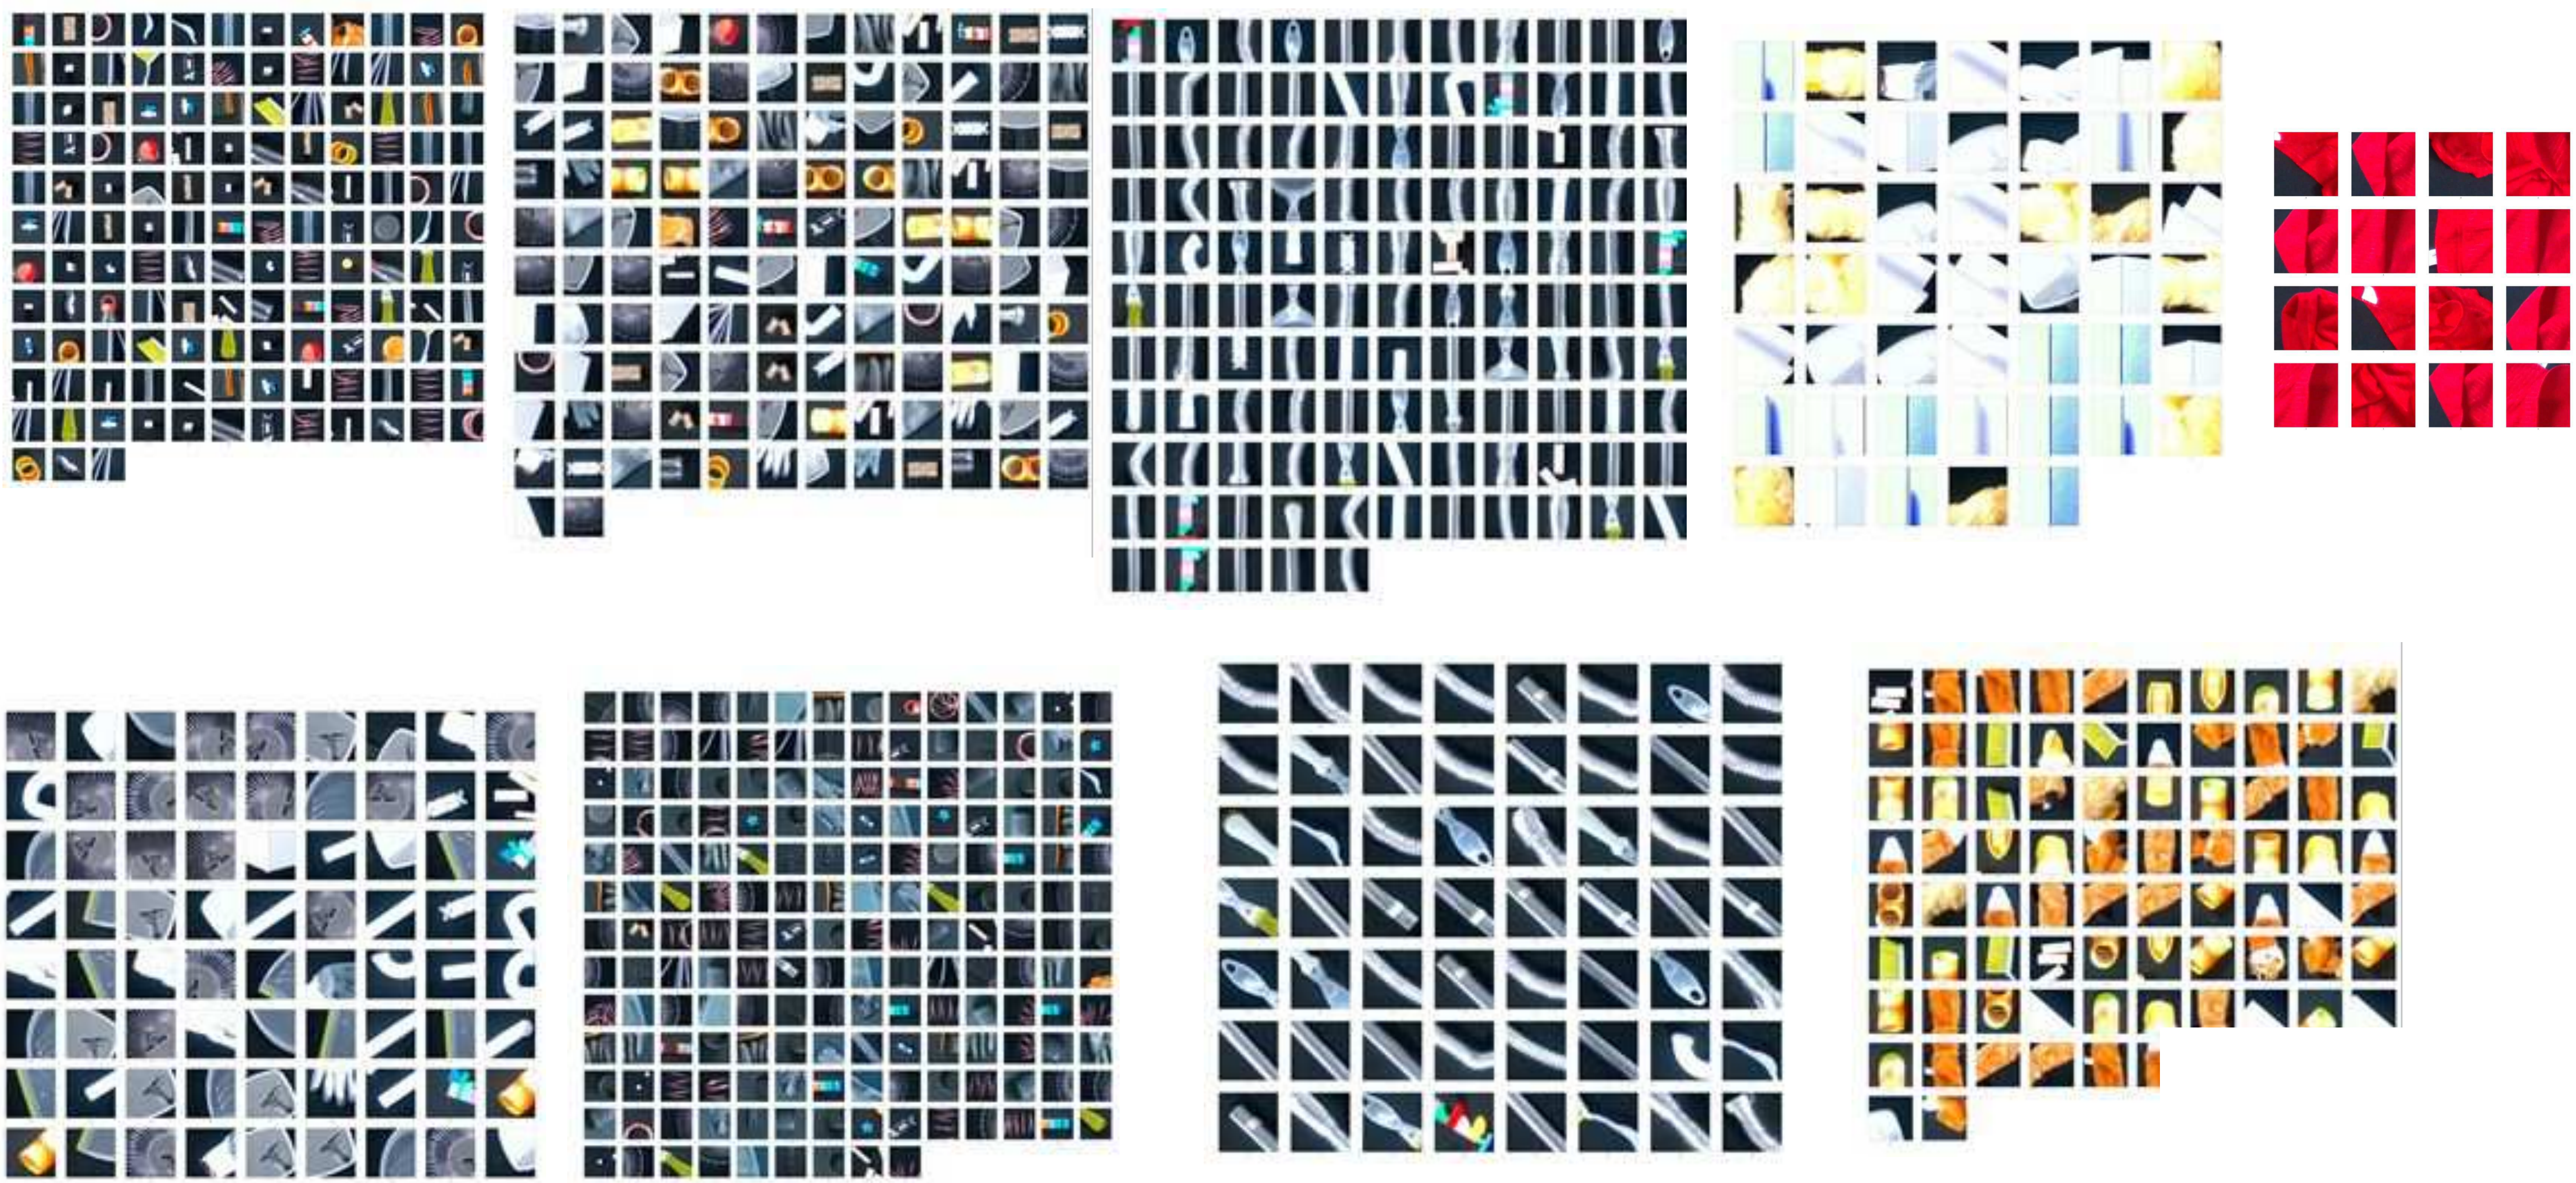

Still work in progress…

Summary

Autonomous curiosity based exploration
Human guidance in exploring objects

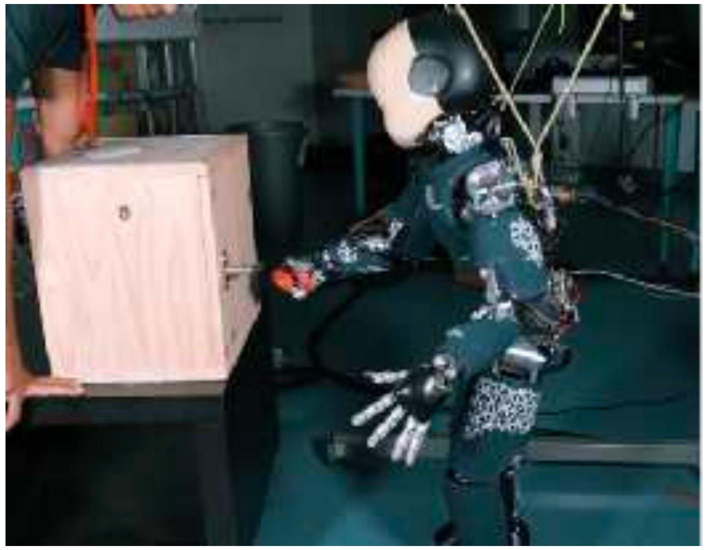

Methods for demonstrating whole-body manipulation (open-ended, without constraints)

Incorporating human preferences and learning class-specific grasp preferences (incrementally)

Acknowledgments

N.Lyubova, D. Filliat,
S.M. Nguyen, A. Droniou,
O. Sigaud, P.-Y. Oudeyer

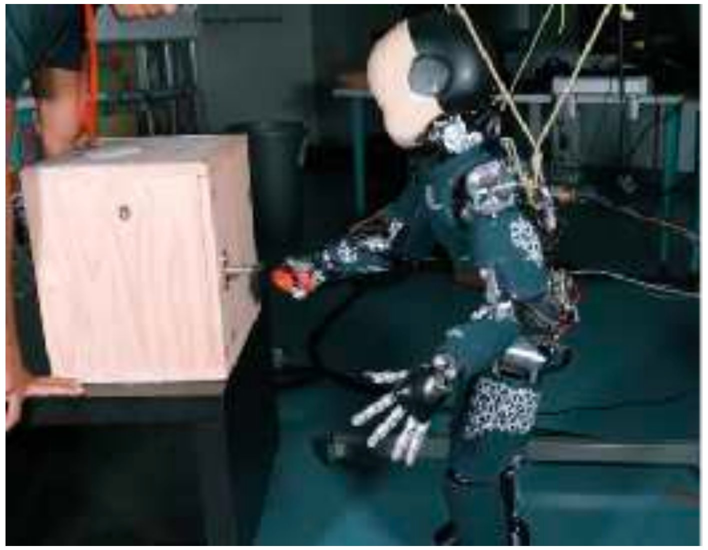

L. Penco, J.B. Mouret,
V. Modugno, W. Gomes

L. Vianello, Y. Fleytoux,
A. Ma, J.B. Mouret

Thank you! Questions?

N.Lyubova, D. Filliat,
S.M. Nguyen, A. Droniou,
O. Sigaud, P.-Y. Oudeyer

AGENCE NATIONALE DE LA RECHERCHE ANR

MACSi

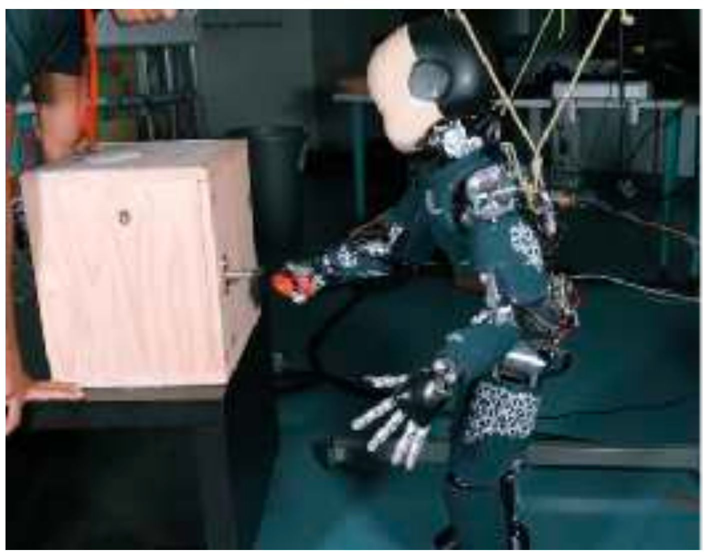

L. Penco, J.B. Mouret,
V. Modugno, W. Gomes

AnDy

L. Vianello, Y. Fleytoux,
A. Ma, J.B. Mouret

HEAP