School of Computing
FACULTY OF ENGINEERING

UNIVERSITY OF LEEDS

Manipulation in cluttered environments, and interacting with robots

Anthony (Tony) G Cohn

Also: Turing Fellow at the Alan Turing Institute,
Adjunct Professor at Shandong University,
Distinguished Visiting Professor at Tongji University and
Qingdao University of Science and Technology

Today's Talk

How to facilitate robots reaching target objects in cluttered environments

- Using learned knowledge from human demonstrations
- Using human hints at run-time

Enabling humans to work with robots on collaborative tasks in a comfortable manner

Based on 3 papers published in 2020 (ICRA, RA-L, RA-L)

Some background to first part of talk

EPSRC
Engineering and Physical Sciences Research Council

E·S·R·C
ECONOMIC & SOCIAL RESEARCH COUNCIL

Human-like Computing: Call for feasibility studies

… research that could lead to the development of human-like computing systems: **machines with human-like perceptual, reasoning and learning abilities, which support collaboration and communication with human beings.**

…goes beyond designing improved AI or machine learning systems, and it is not about incorporating findings in neuroscience.

However a key component of the projects we are looking to encourage and support **is multidisciplinary research** involving cutting edge and state-of-the-art research in both **computer and cognitive science**.

Motivation of the call

- To enable better communication and collaboration between humans and machines, especially in the context of hybrid teams in the workplace.
- To support the generation by ML of explicit and debuggable hypotheses and programs which incorporate and support reasoning, which can be understood by humans.
- To improve our understanding of human cognition via combinations of psychological experiments, analysis of human-derived data, and computational modelling.
- To inspire new forms of computation based on human cognition, especially on tasks where humans currently exhibit superior abilities.

Human-like Computing (HLC)
UK EPSRC initiative

Engineering and Physical Sciences Research Council

- Much AI is not “human like”
  - Learning from millions of examples
  - e.g. Tesla Autopilot 780M miles, AlphaGo…
- Inscrutable models
- HLC aims to endow machines with human-like perceptual, reasoning and learning abilities
  - support collaboration and communication with humans
- HLC research could enrich understanding of human cognition
  - through tests of existing models
  - or development of new ones
- Call for “Feasibility Studies”: £300k each

A side note:

An output of the initiative *(funded projects, workshops, network):*

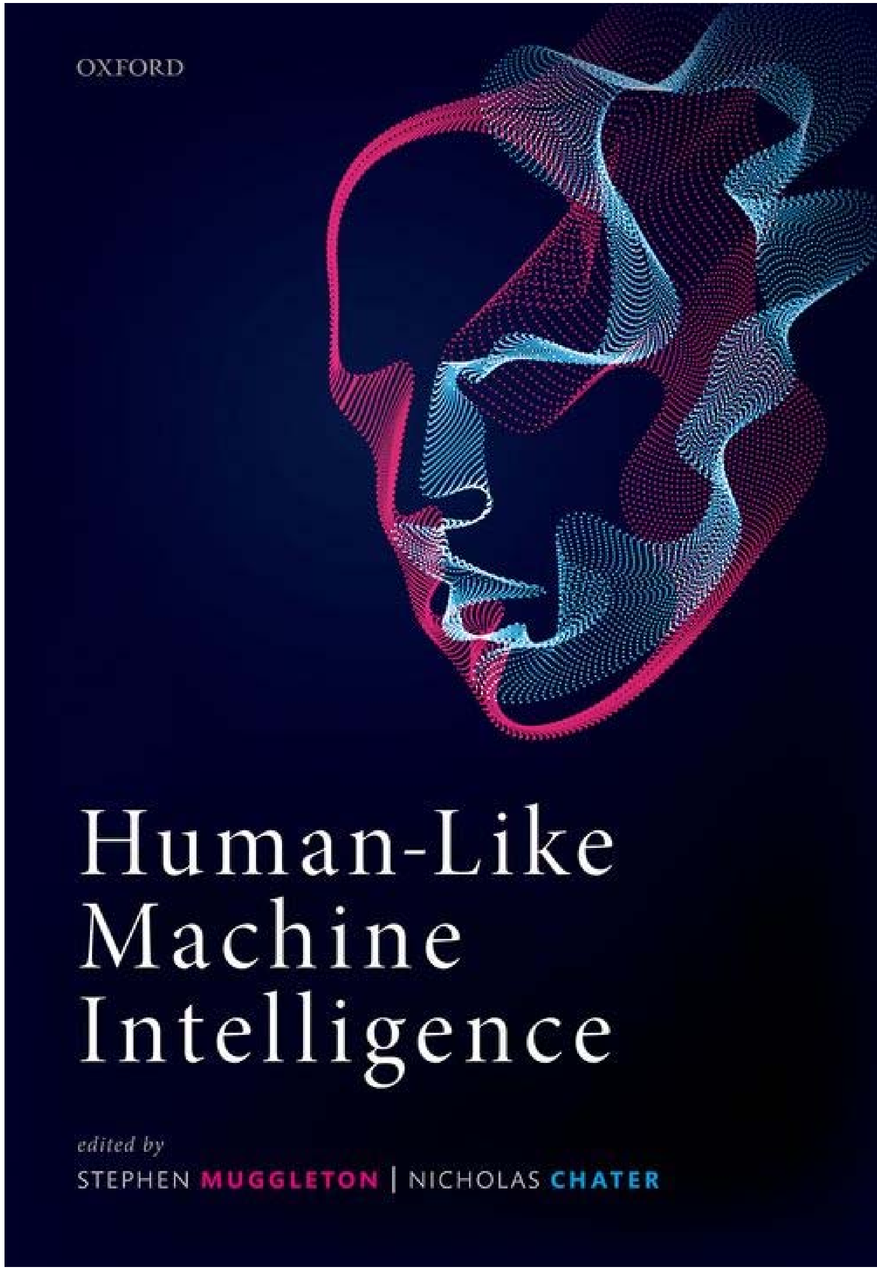

Human-like Physics Understanding for autonomous robots

State-of-the-art robot motion/manipulation planners use low-level probabilistic methods

- often based on random sampling.

Restricts robots to plan their motion at the bottom-most geometric level

- without any top-down guidance
- this results in the limited object manipulation ability displayed by today's intelligent robots.
- **Particularly in cluttered environments**

Produces randomized motion that is not legible to humans

- limits robots' collaboration capabilities with humans.

Reaching in Cluttered Environments

Grasping a target in a cluttered environment:

- Reach directly the target or firstly pushing obstacles away?

- Reaching: which path?

- Pushing: which object and where to?

Amazon Picking Challenge

Team RBO

- Open-problem.
- Hard motion-planning problem when considering physics and cluttered environments like a shelf.

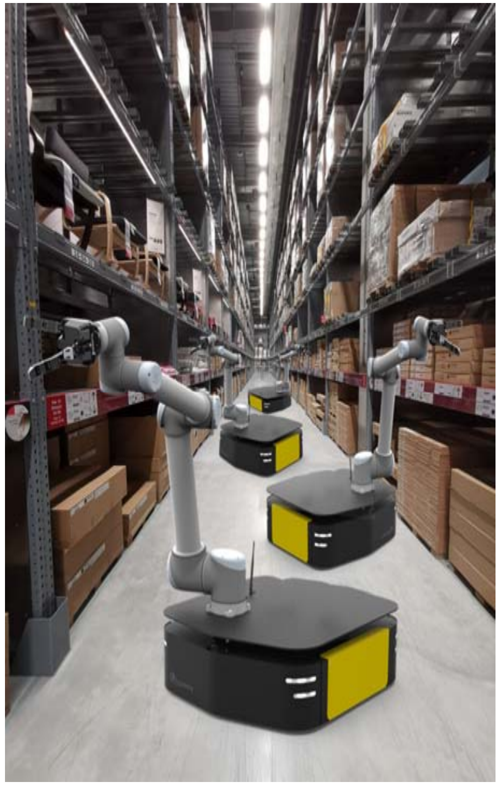

Varying levels of complexity

Existing Robot Planners

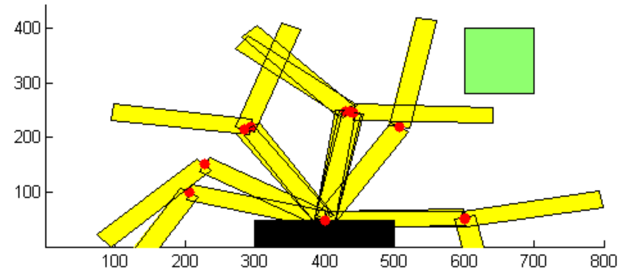

- Based on random sampling in configuration space.
- High-dimensional with large number of objects.
- Limited object manipulation.
- Long planning time.

Human-like Planning (HLP) for Reaching in Cluttered Environments

Objectives

- Learning high-level manipulation planning skills from humans.
  - "heuristics"? Cf Gerd Gigerenzer's talk from Tuesday
- Transfer these skills to robot planners.
- Plans should be "human-like"
  - "legible"
  - "explainable"

Focus on non-prehensile manipulation

Human-like planning for reaching in cluttered environments, M Hasan, M Warburton, W C Agboh, M R Dogar, M Leonetti, H Wang, F Mushtaq, M Mon-Williams, A G Cohn, ICRA-2020

Non-Prehensile Manipulation

- Manipulating objects without grasping them.
- Pushing, pulling, toppling and sweeping.
- Motivation: objects might be ungraspable, heavy or inefficient to pick & place every blocking obstacle.

Some Research Questions

- How few examples are needed?
- Which Qualitative Spatial Representations?
- Is manipulation planning more efficient?
- Can robots avoid human-like “stuck in a rut” decision making?
- Will robot actions become “human legible”?
- …

Early Experiment
(replication of known result)

1 block, 1 target.

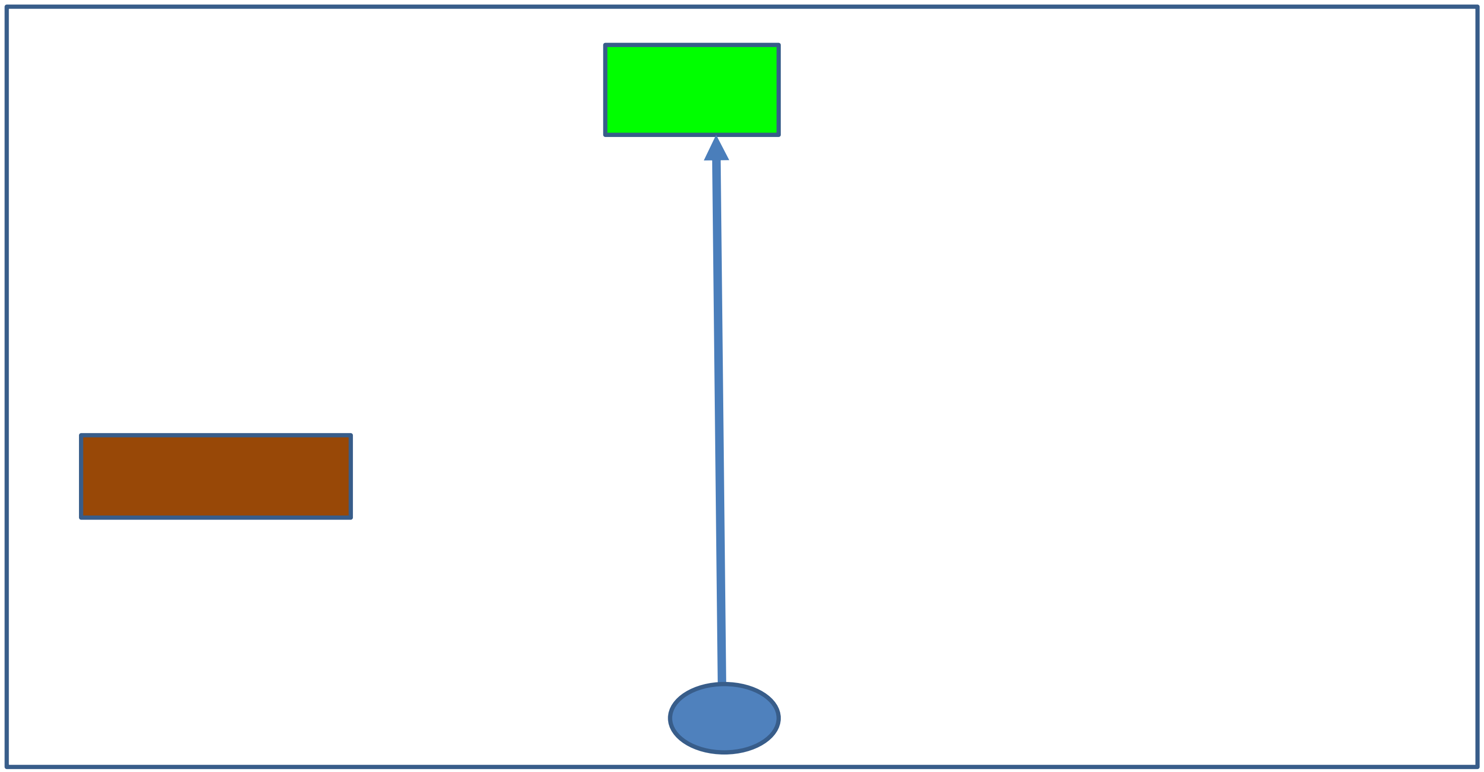

Early Experiment
(replication of known result)

1 block, 1 target.

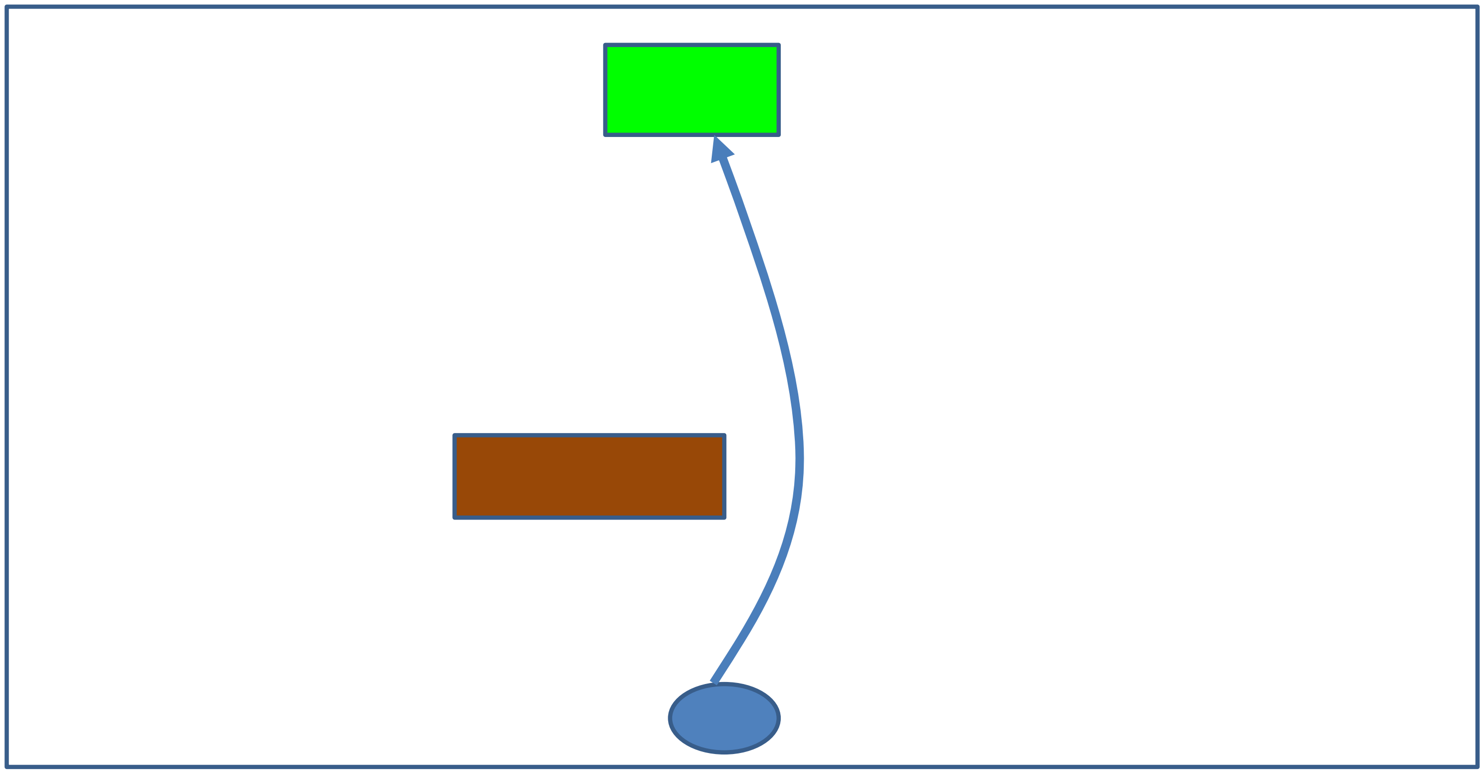

Early Experiment
(replication of known result)

1 block, 1 target.

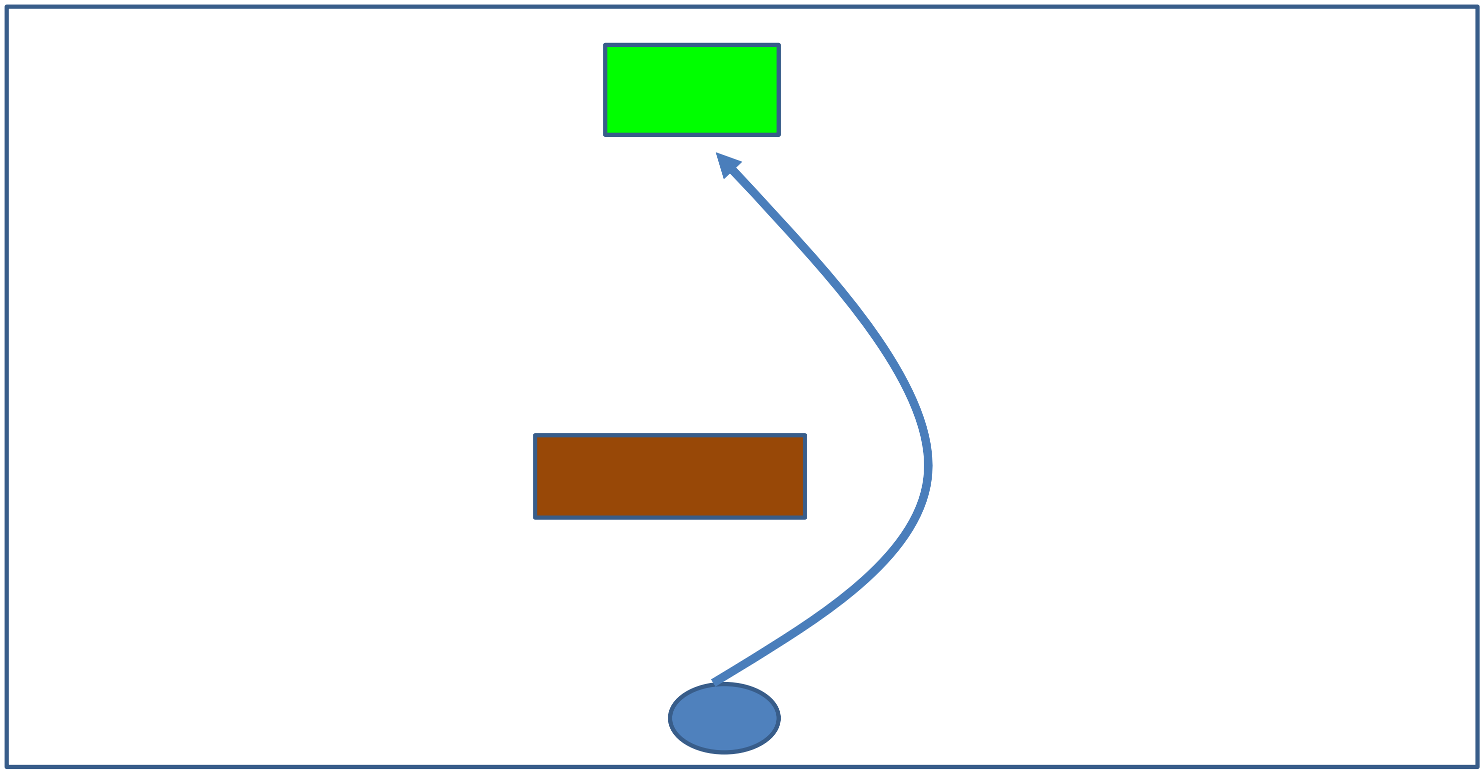

Early Experiment
(replication of known result)

1 block, 1 target.

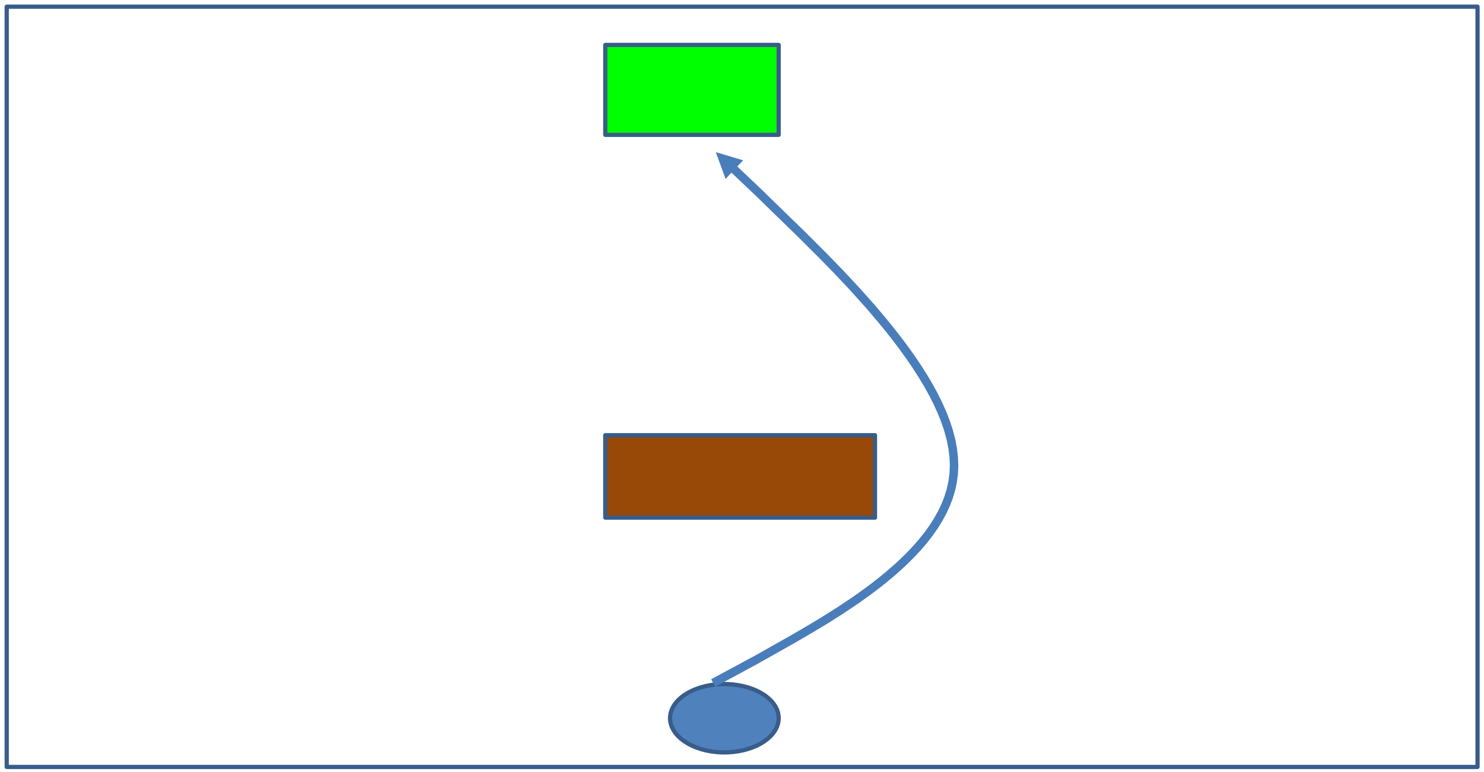

Early Experiment
(replication of known result)

1 block, 1 target.

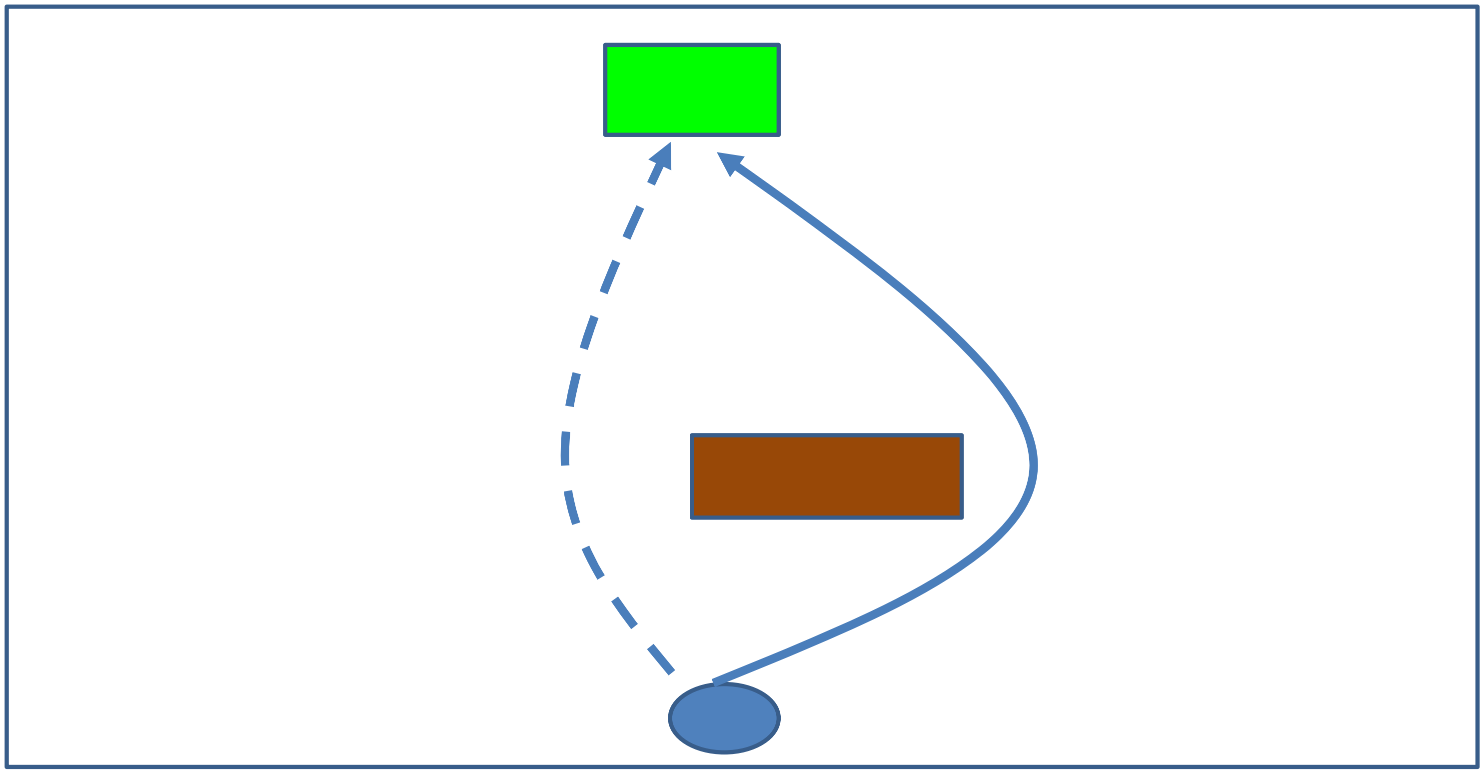

Early Experiment
(replication of known result)

1 block, 1 target.

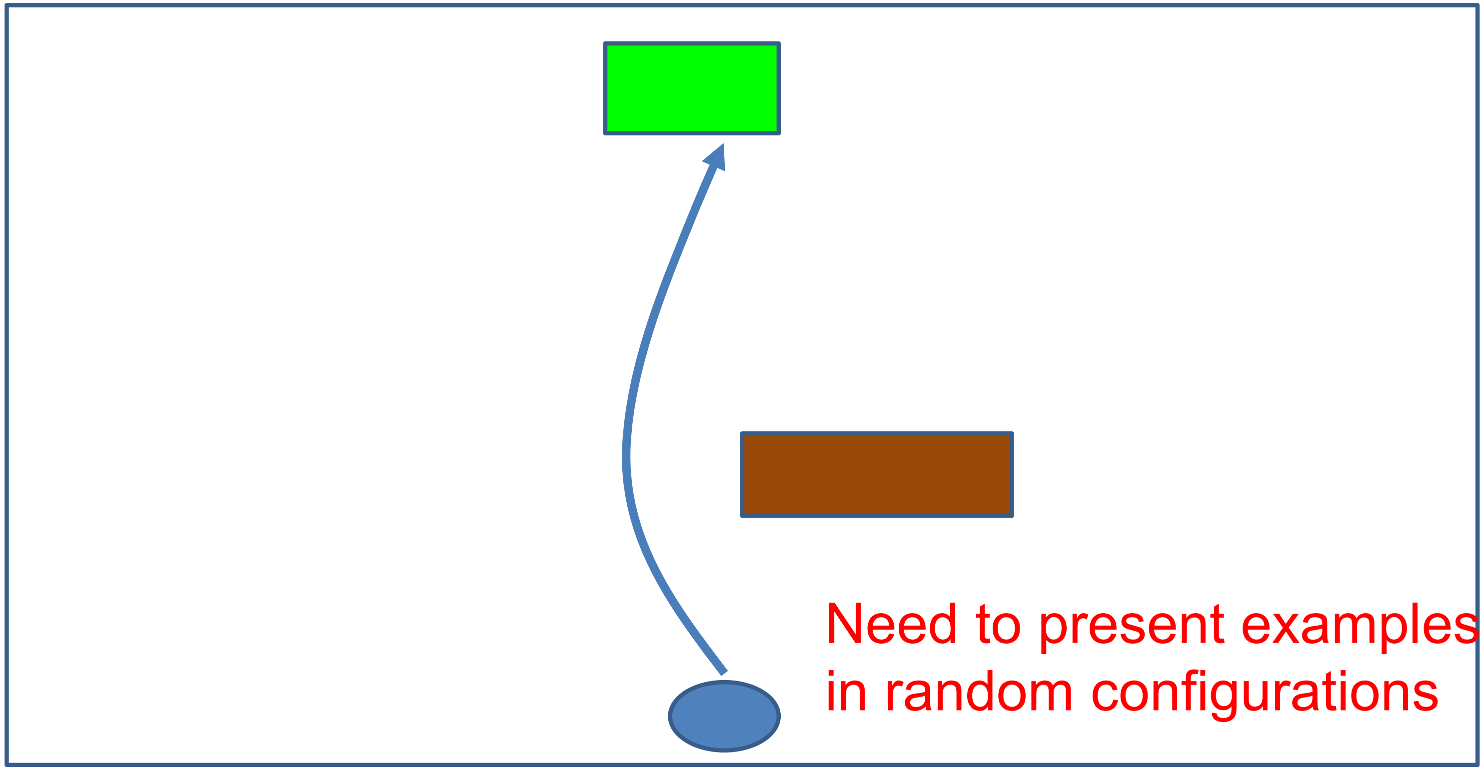

A background motivation

“Higher-order” human cognition built on sensorimotor foundations

A child’s physical interaction with the world lays foundations of their higher-order cognitive capabilities

$\Longrightarrow$ Understanding the sensorimotor world critical for high level cognitive systems

Can we capitalise on the high level understanding humans have?

HLP Overview

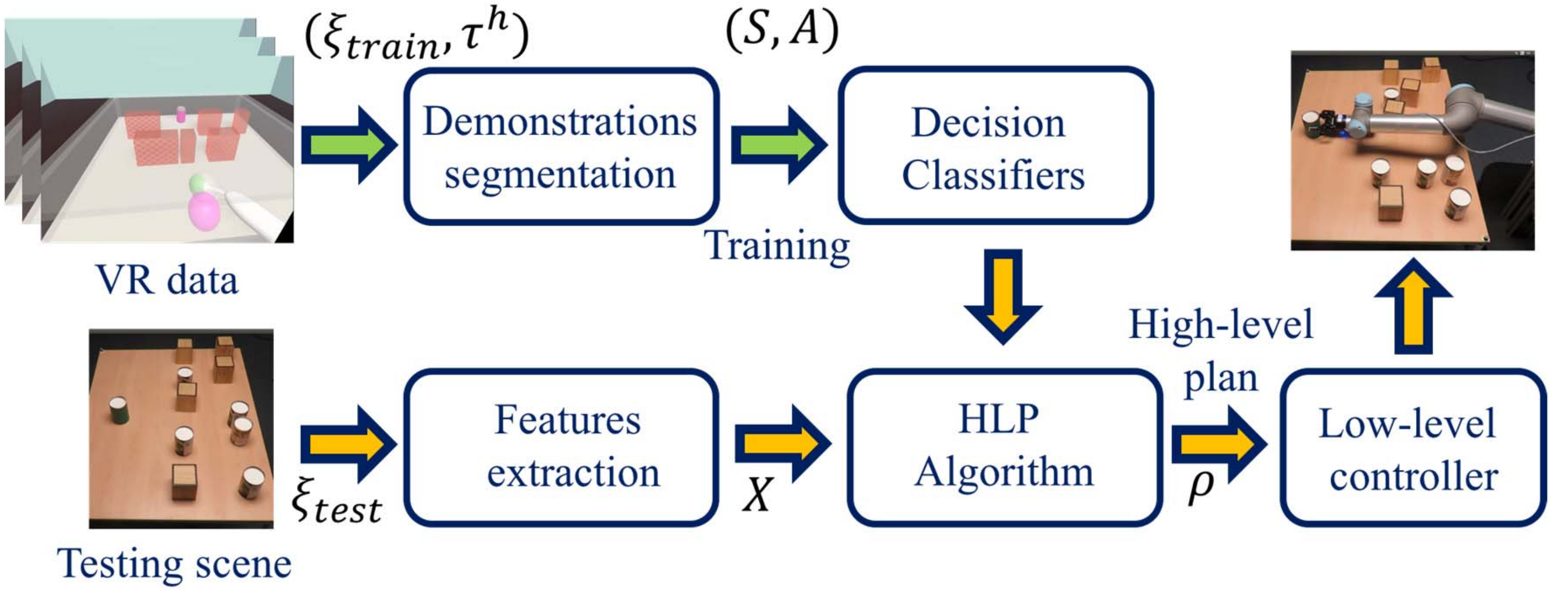

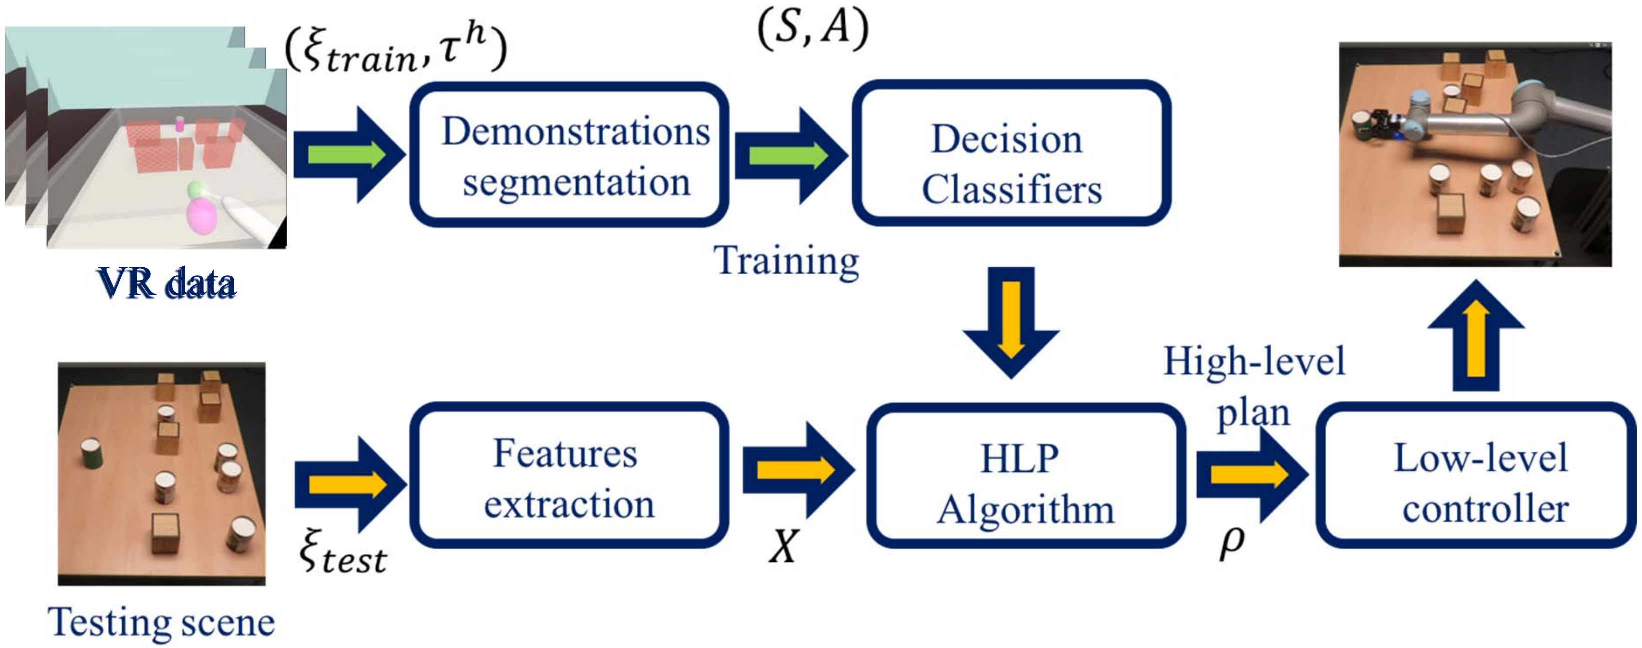

(ξ_{train}, τ^h)
Demonstrations segmentation
(S, A)
Decision Classifiers
Training
VR data
Testing scene
ξ_{test}
Features extraction
X
HLP Algorithm
High-level plan
ρ
Low-level controller

Virtual Reality Data Collection

Participant Trials

VR Environment

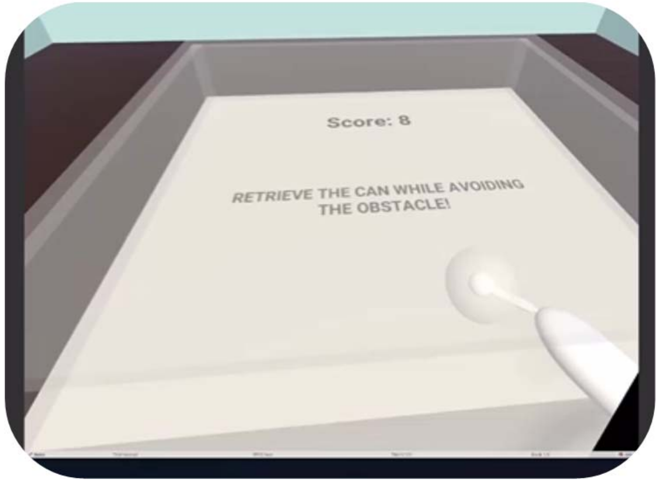

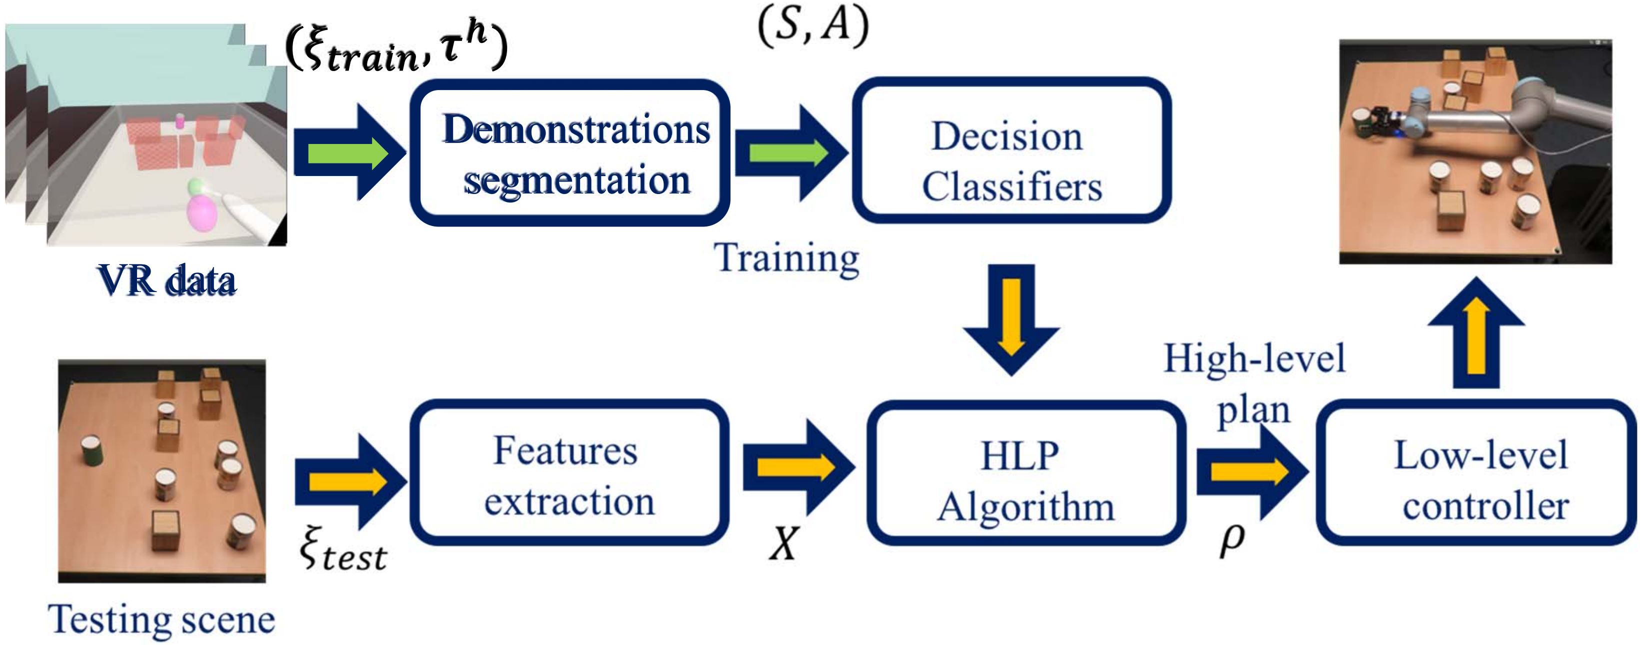

(ξ_{train}, τ^h)
(S, A)
Demonstrations segmentation
Decision Classifiers
VR data
Training
High-level plan
Features extraction
HLP Algorithm
Low-level controller
ξ_{test}
X
ρ
Testing scene

Modelling the Task Space

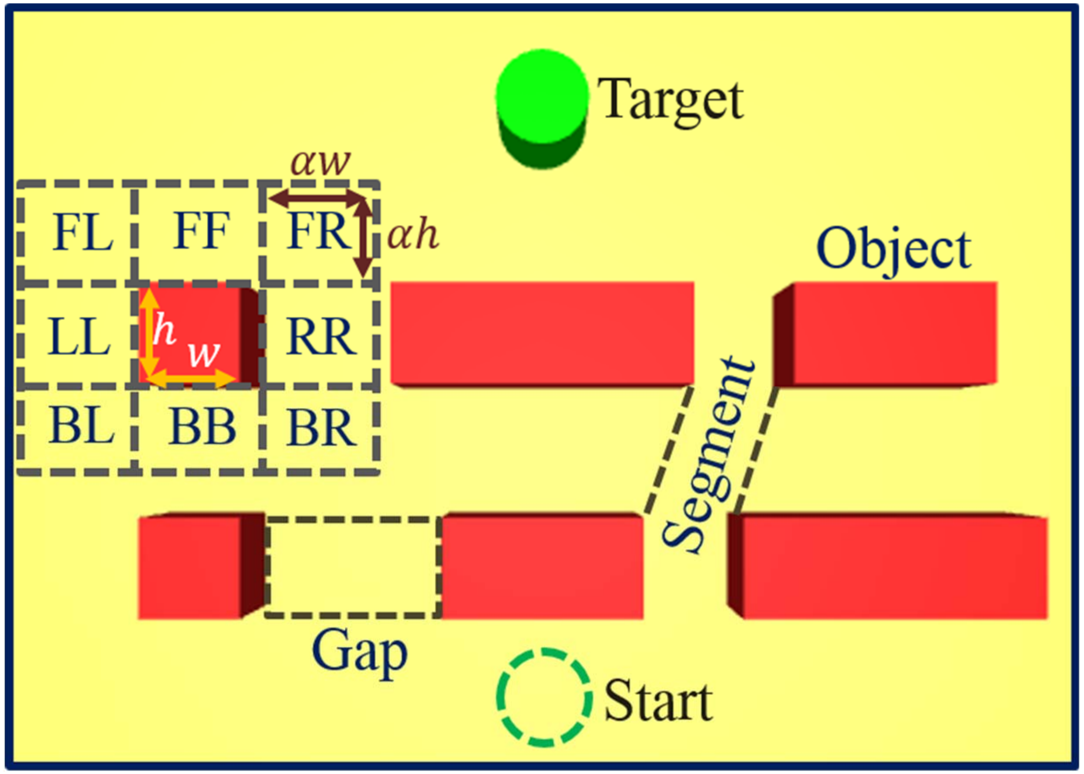

Segmenting Demonstrations

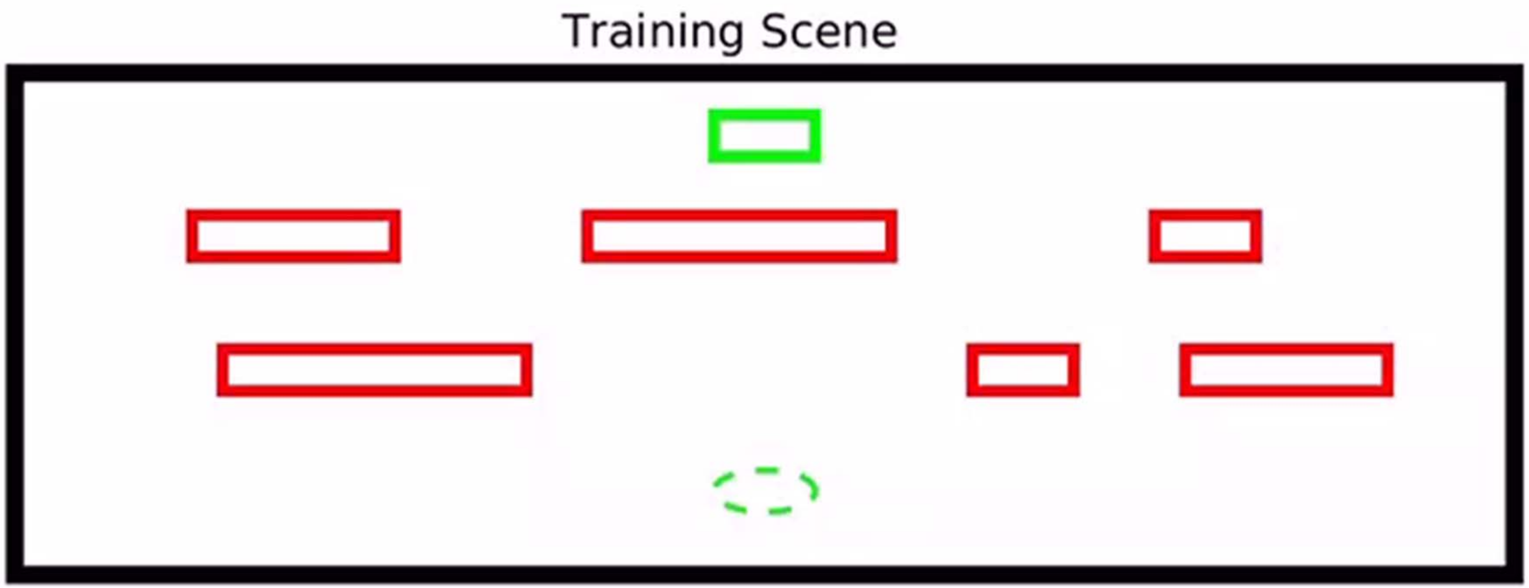

A high level plan is a sequence of keypoints in the action space connected by segments.

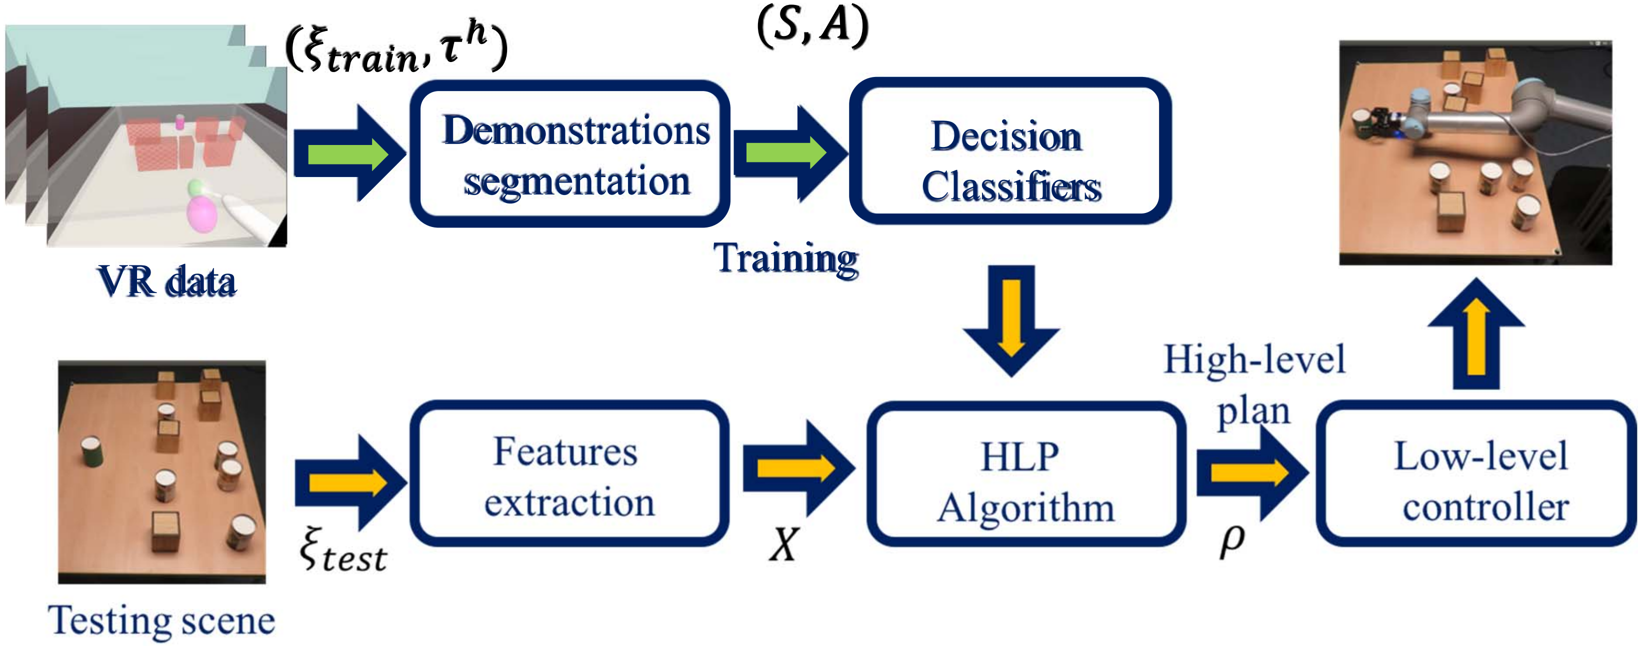
(ξ_{train}, τ^h)
(S, A)
Demonstrations segmentation
Decision Classifiers
VR data
Training
High-level plan
Features extraction
HLP Algorithm
Low-level controller
Testing scene
ξ_{test}
X
ρ

Decision Classifiers

Gaps Classifier → Which gap to go through?

Objects Classifier → Which object to move?

Object Direction Classifier → Where to move?

Segments Classifier → Which segment?

- Qualitative spatial relations used as features (distance, size & direction)
- Train binary classifiers: for each gap/object/segment learn a probabilistic classification – enables arbitrary number of objects in scene
- Object Direction Classifier is 8-way classification

Features

- **Gap classifier:** distances/orientations to target and start, gap size
- **Object classifier:** distances/orientations to target and start, object size, freespace around object, horizontal overlap with target and start.
- **Object direction classifier:** directions to target and start; free space in each of the 8 directions.
- **Segment classifier:** horizontal and vertical distances; horizontal overlap with target and start; orientation wrt start-target vector; collision measure.

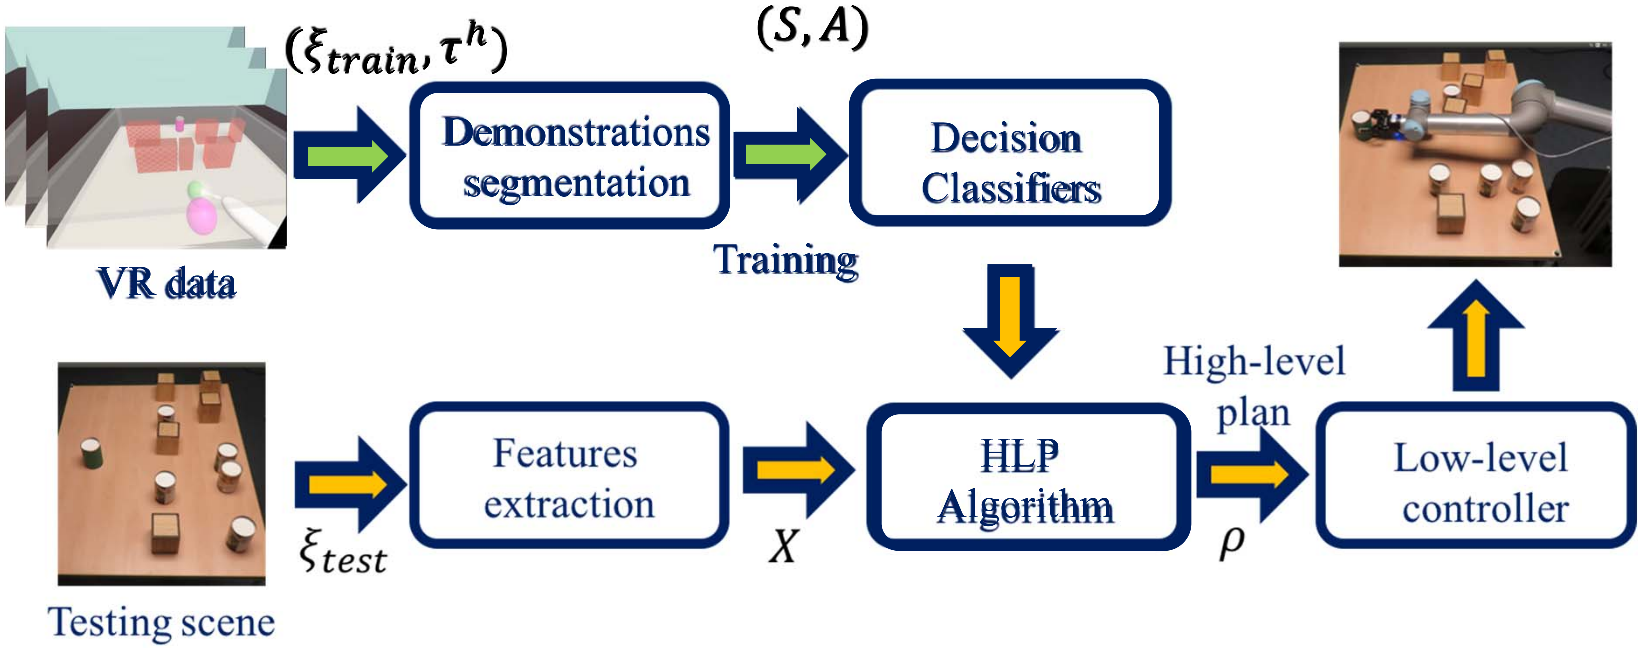

(ξ_{train}, τ^h)
Demonstrations segmentation
(S, A)
Decision Classifiers
Training
VR data
High-level plan
Testing scene
ξ_{test}
Features extraction
X
HLP Algorithm
ρ
Low-level controller

HLP Algorithm

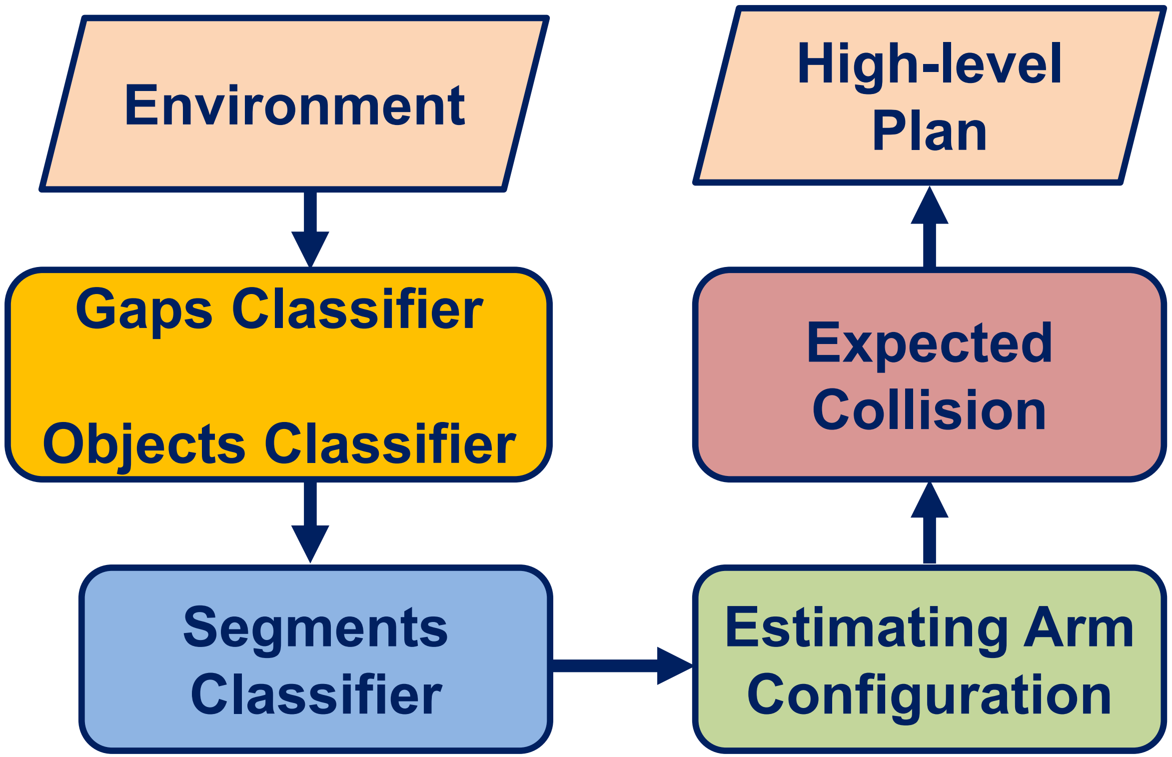

Algorithm: Human-Like Planner

The high-level plan is generated hierarchically in three levels: **path**, **segment** and **action**.

Each segment connects a pair of consecutive rows.

One action takes place at each row and applies to a gap or an object.

Human arm is modelled as a planar arm with four joints at neck, shoulder, elbow and hand.

Arm configuration is represented by two angles: ϑ_{sh} between neck-shoulder and upper arm links and ϑ_{el} between upper arm and forearm links.

Input: Environment representation $\xi = \{X^s, X^t, X_i^o\}$
Output: High-level path ρ

Locate rows R and gaps ξ^g
for all R **do**
  Compute gaps feature vector X_g
  $G_{selected} \leftarrow C_g(X_g)$
  Compute objects feature vector X_o
  $O_{selected} \leftarrow C_o(X_i^o)$
end for
for all pairs of consecutive rows **do**
  $C \leftarrow$ Segment Constructor $(G_{selected}, O_{selected})$
  Compute segments feature vector X_c
  $C_{selected} \leftarrow C_c(X_c)$
end for
for all $C_{selected}$ **do**

  if $a^o \in C_{selected}$ **then**
    Compute object-direction feature vector X_d
    Object direction = $C_d(X_d)$
    Augment $C_{selected}$ by expected object's location
  end if
  Compute arm configuration feature vector X_a
  Estimate arm configuration: $R_a(X_a)$
  Compute expected path collision ρ_ς
end for
Select the path with minimum collision score

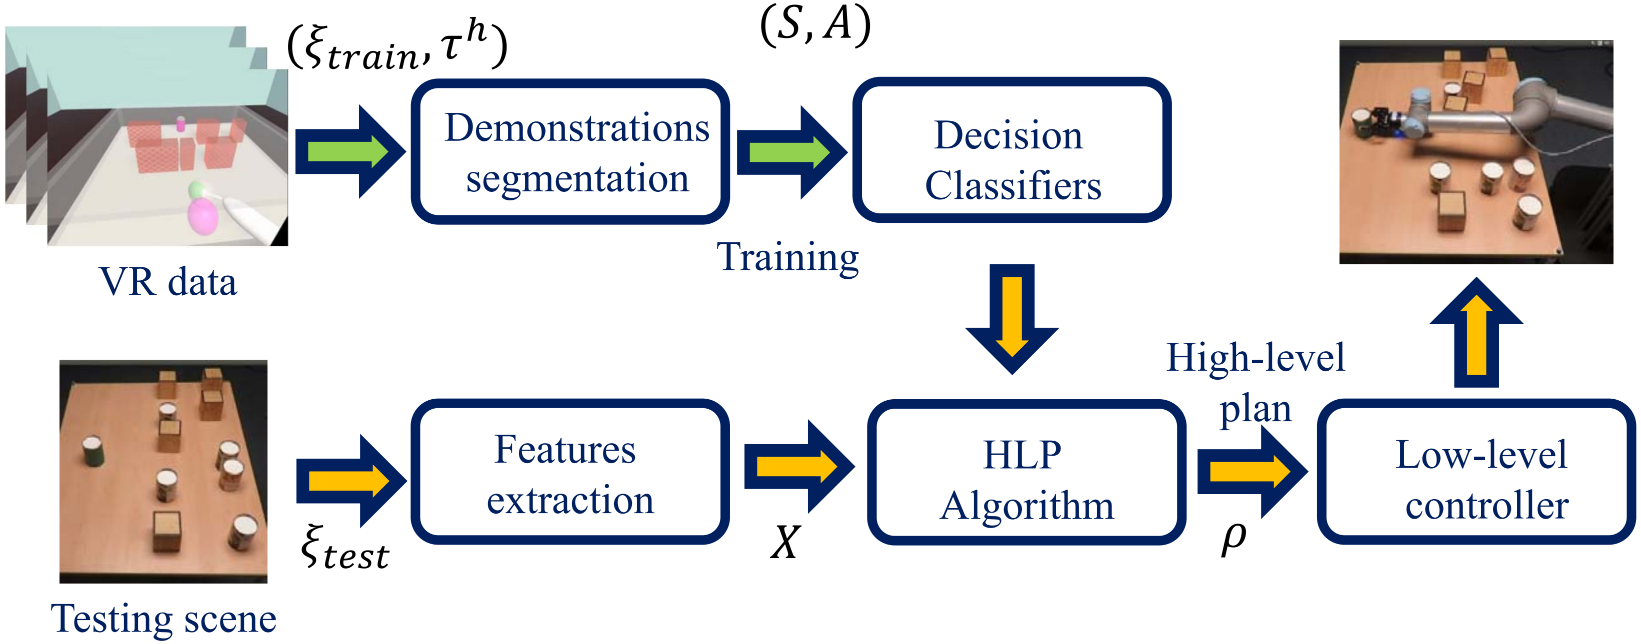

(ξ_{train}, τ^h)
Demonstrations segmentation
(S, A)
Decision Classifiers
VR data
Training
High-level plan
Testing scene
Features extraction
ξ_{test}
X
HLP Algorithm
ρ
Low-level controller

Experiments on Robot Simulation*

	HLP	STO**
Success rate (%)	**94**	84
Planning time (s)	**1.56**	17.88

* Mujoco physics engine.

** STO: Stochastic trajectory optimization.

W. C. Agboh and M. R. Dogar, “Real-time online re-planning for grasping under clutter and uncertainty,” in Humanoids, 2018.
W. C. Agboh, D. Ruprecht, and M. R. Dogar, “Combining coarse and fine physics for manipulation using parallel-in-time integration,” (ISRR), 2019.

How similar to human plans?

TABLE I: Results (mean and standard deviation) of the 5-fold *VR* test experiment.

Metric	Mean	STD
C_g accuracy	0.95	0.002
C_o accuracy	0.85	0.005
s_{HLP} (overall)	0.70	0.011
s_{HLP} $(I(D_n))$	0.79	0.016
s_{HLP} $(I(E_n))$	0.67	0.012

$$s_{HLP} = \frac{1}{2N_R} \sum_{n=1}^{N_R} I(D_n)(I(D_n) + I(E_n))$$

Comparison between HLP and Human

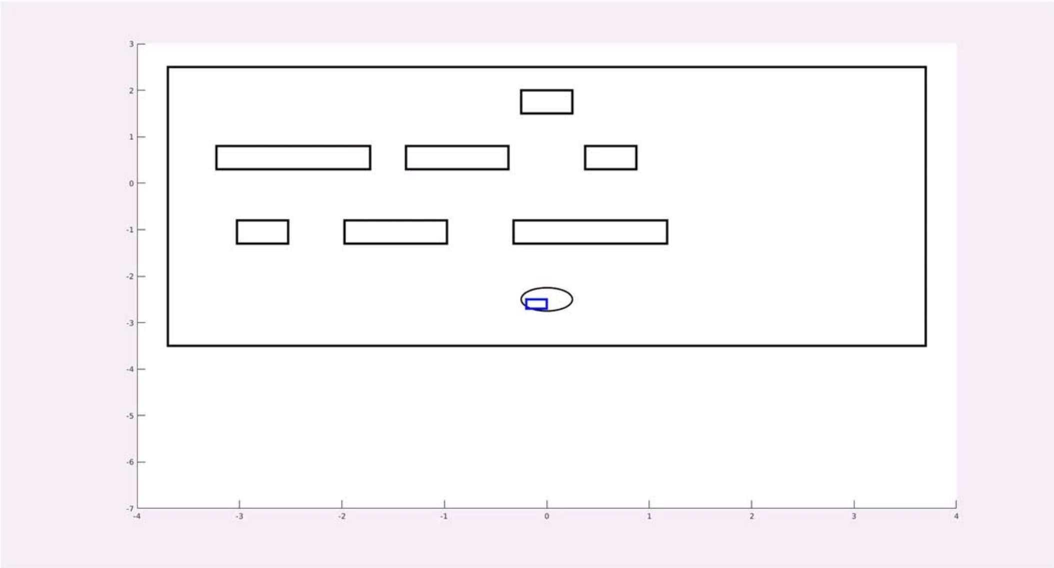

Human-**like** Plan

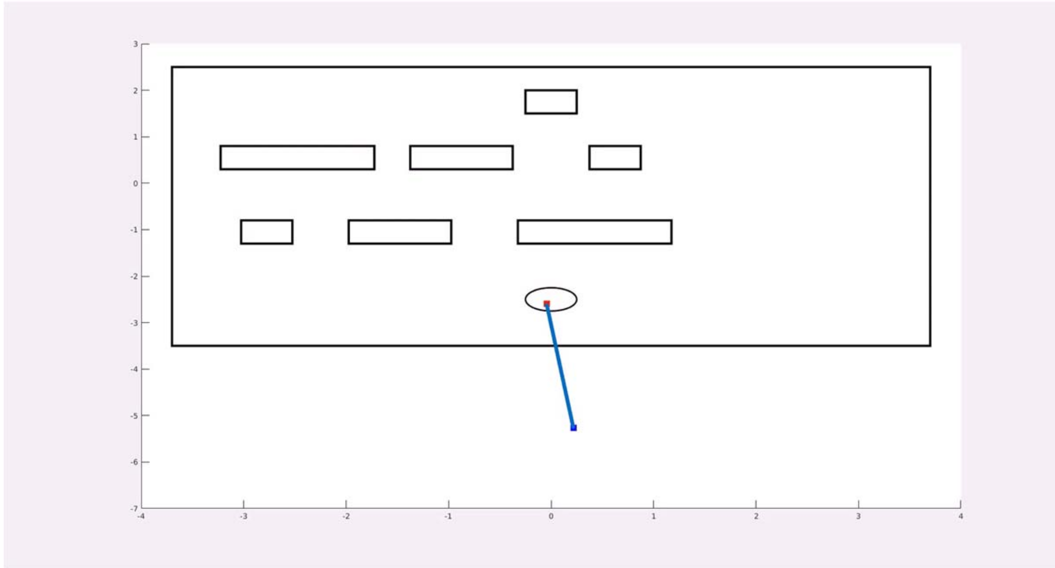

Human Plan

Real Robot Experiments

HLP

STO

Reaching in clutter (7 Objects)

Real Robot Experiments

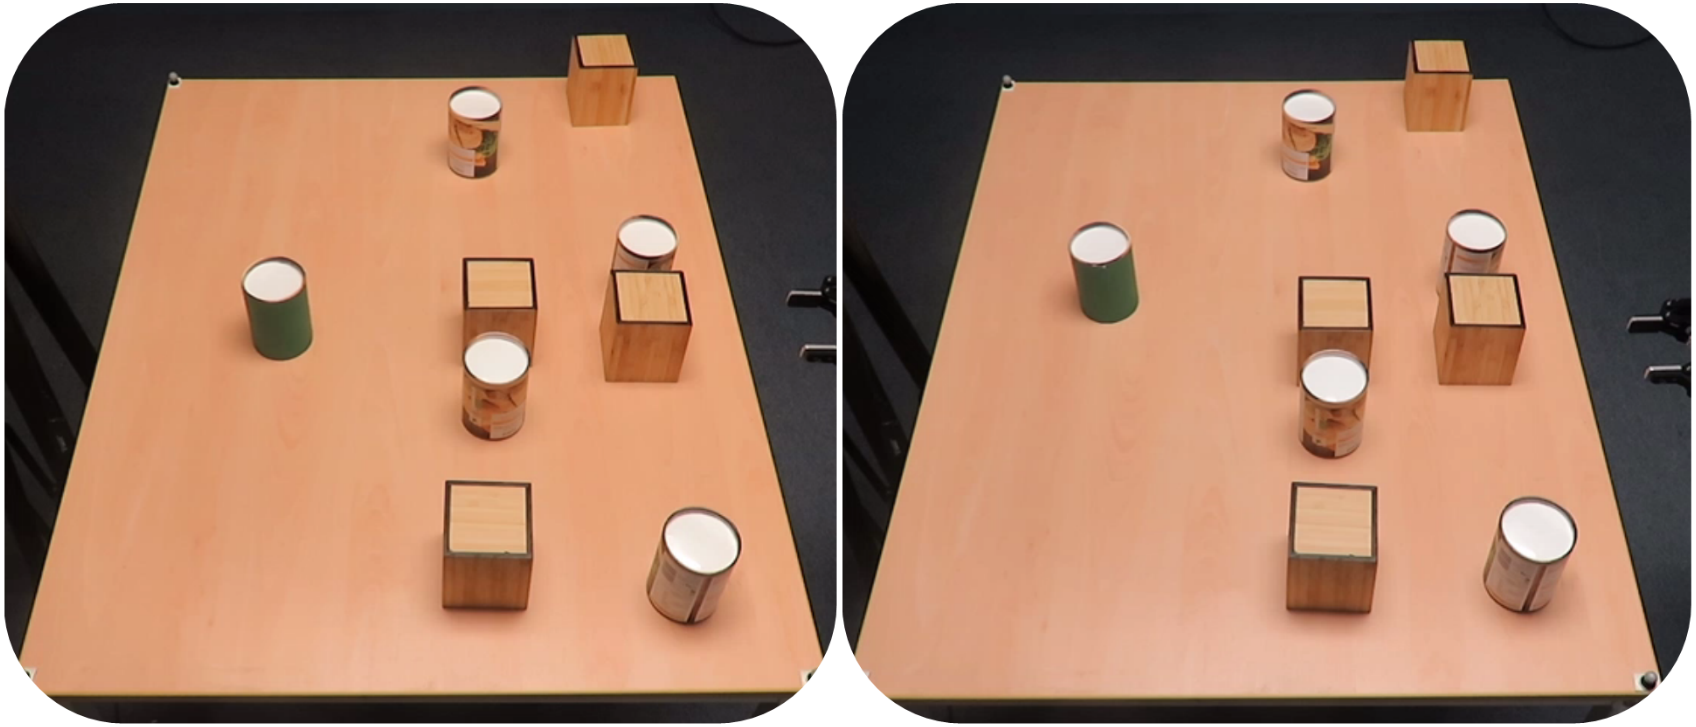

Reaching in clutter (9 Objects)

Real Robot Experiments

HLP STO

Reaching in clutter (11 Objects)

Conclusions

- Learning from humans interacting in VR.
- Qualitative representation of the task space and action space.
- High-level planning algorithm.
- Scalability.
- Working with any arbitrary robot model.

Future Work

- Non row-structured environments in training set
- Other scenarios
- Experimenting with number of training examples needed
- More powerful classifiers.
  - Already experimented successfully with ILP in an MSc project
- Closed-loop planning.

Some rules/heuristics extracted from ILP

Gaps

- Prefer
  - larger gaps
  - gaps closer to start position
  - 1st gap is in NE direction (right handed subjects?)
- Not relevant: gap direction to target

Object selection

- Large surrounding free space
- Smaller object (easier to move into free space?)
- High overlap with target

Object direction:

- Human choices seem quite random but
- Prefer FR or FL (unless blocked by moved object)

Optimization-based Motion Planning with Human-in-The-Loop for Non-Prehensile Manipulation

Rafael Papallas, Anthony G. Cohn and Mehmet R. Dogar

IEEE International Conference on Robotics and Automation (ICRA) 2020
Shared Autonomy: Learning and Control Workshop

and

IEEE Robotics and Automation Letters 2020

Human-In-The-Loop and shared-autonomy

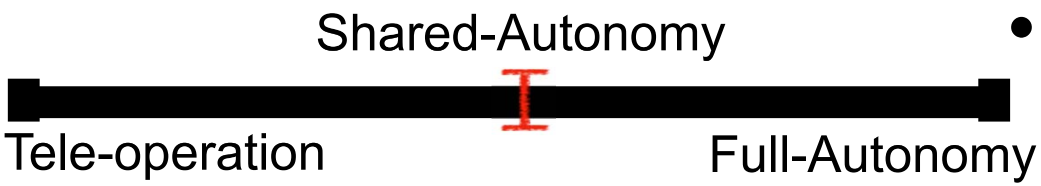

- Tele-operation: no autonomy, a human controls all DOFs of the robot.
- Full-autonomy: robot needs no input and performs everything fully autonomously.
- Shared-Autonomy/Human-in-the-Loop: Robot has some autonomy but leverage input from a human to solve the task faster and more robustly.

An alternative approach:

- Exploit human “hints” at run time
- System plans in simulation mode, optimising trajectory
- If planner fails to find a solution after fixed time then ask for human help
- Human selects an object and direction of motion for it to be moved to
- Planner incorporates the hint as an update to cost function

(Papallas et al, ICRA 2020 and RA-L 2020)

Reaching Through Clutter Problem

Source: pexels.com

Three main problems

$$Q^E = Q^r \times Q^1 \times \cdots \times Q^{|O|}$$

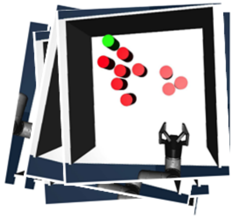

High-dimensional space

Under-actuated system

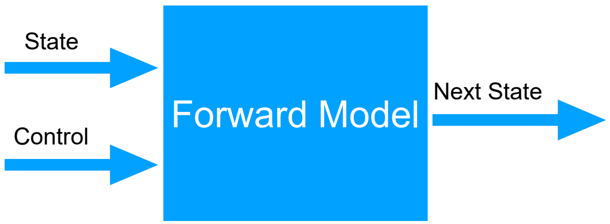

Expensive Physics Simulation

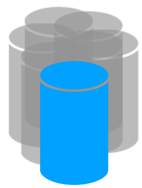

Physics Uncertainty

Online Replanning with Human-In-The-Loop Framework

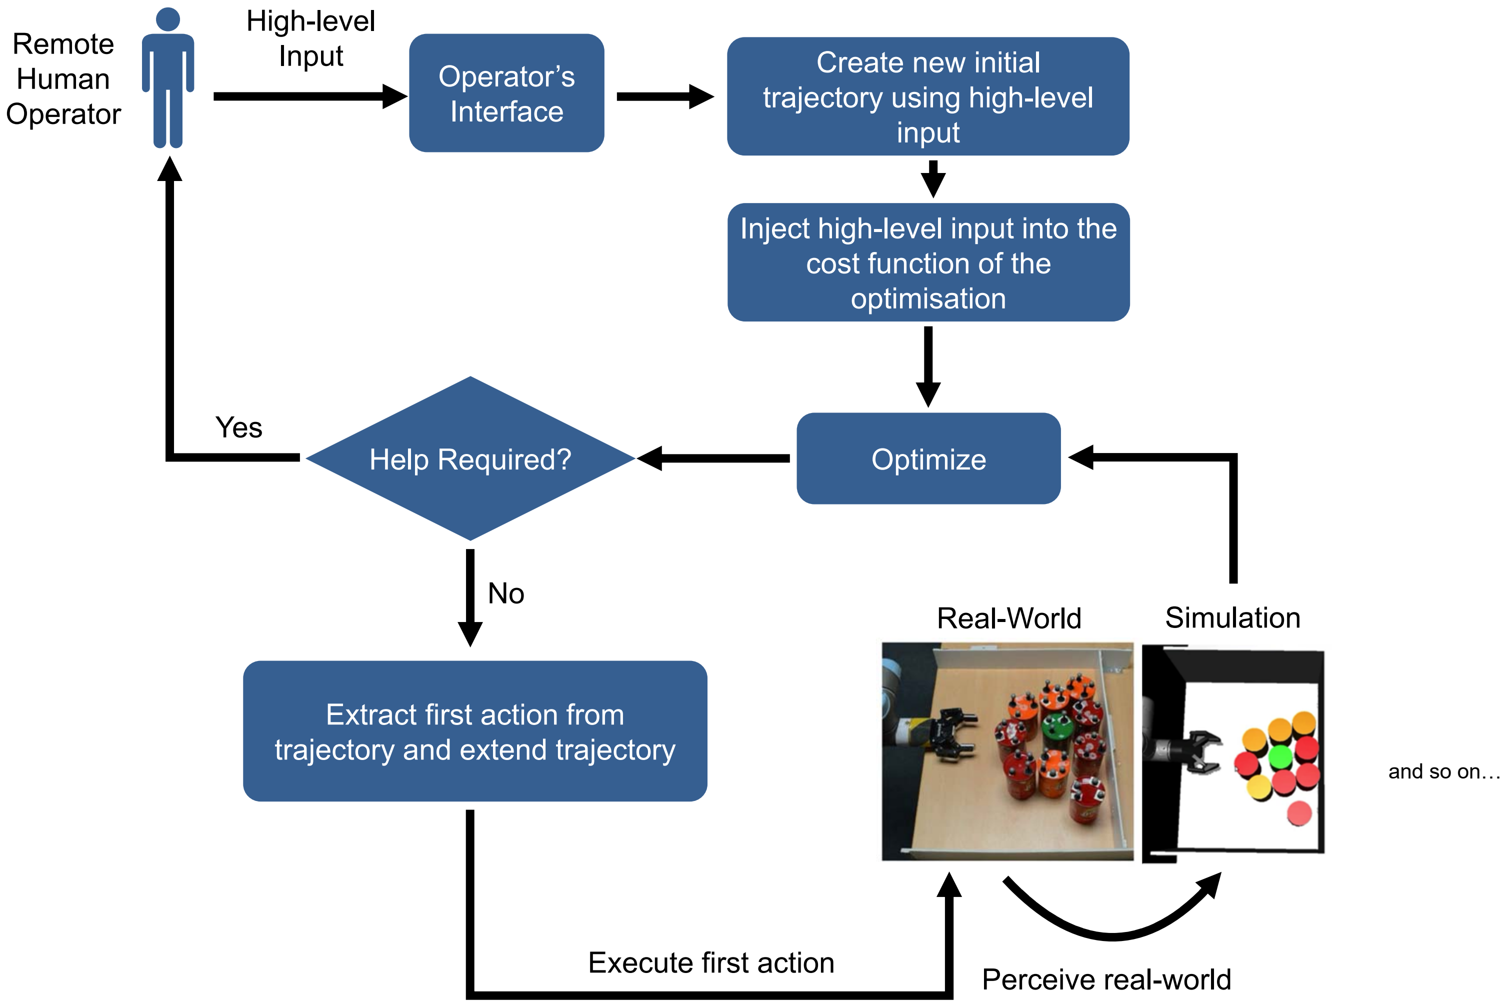

Online Replanning with Human-In-The-Loop Framework

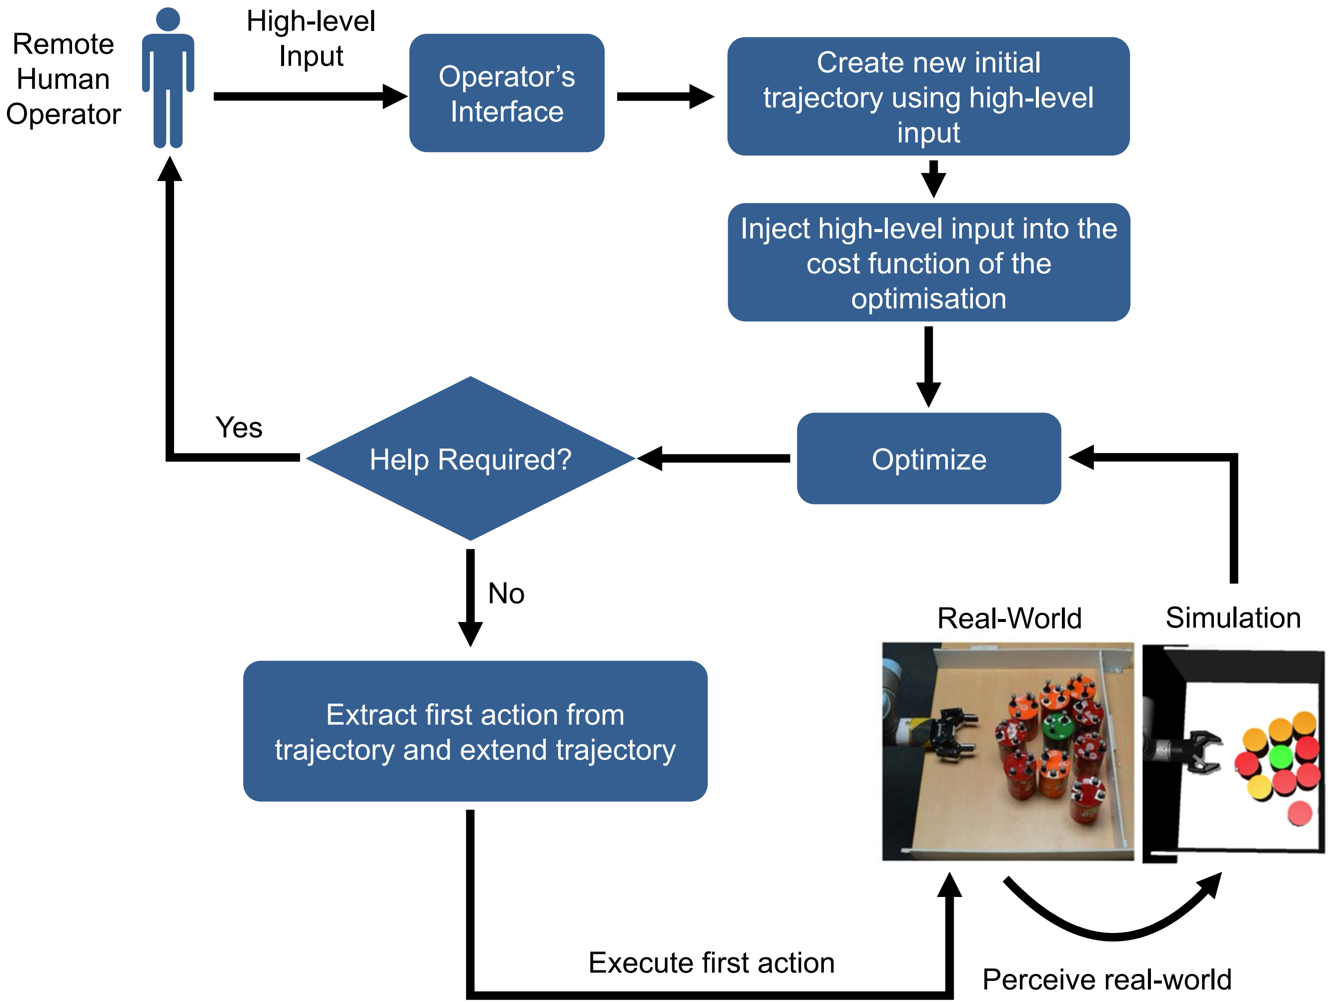

Online Replanning with Human-In-The-Loop Framework

Fixed Timeout

After X seconds of unsuccessful optimisation requests human-help

Pros/Cons

- Easy & straightforward approach.
- Hard to choose a fixed value that suits all problems.
- Could be problematic when having easy/hard problems with long/short timeout values.

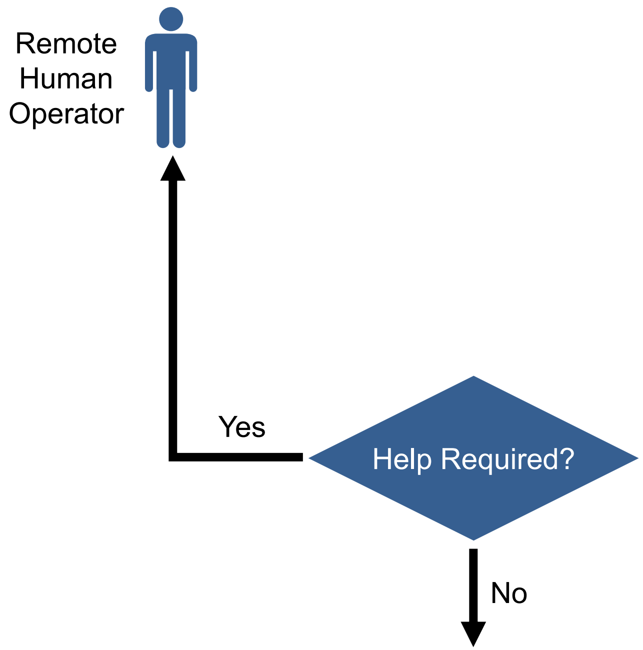

Adaptive

Ask for human-help when stuck in a local minimum.

Pros/Cons

- Adaptive, no fixed value required.
- Leverages the cost value between iterations to decide if human help is required.
- If easy/trivial problems the robot will solve them fully-autonomously.

Simulation

	Adaptive	Fixed 5	Fixed 20	Autonomous
Average Planning Time (s)	**31.0 ± 12.8**	38.1 ± 15.6	44.2 ± 13.8	79.8 ± 11.2
Success Rate	**96.6%**	90%	93.3%	74.6%
Human Time	**2.5 ± 0.9**	9.6 ± 4.1	7.0 ± 1.8	—

Takeaways

- Adaptive was more successful and yields to lower planning times on average.
- Human engagement time is low (2.5s on average).
- Fixed 5 and Fixed 20 are also better than Autonomous but more tedious to use.

Real-world

	Adaptive	OL - Adaptive
Success Rate	**86%**	53%
Planning Failure	**7%**	7%
Execution Failure	**7%**	40%

Takeaways

- Physics uncertainty in the real-world caused open-loop Adaptive to fail more frequently.
- Our online-replanning framework is more robust to physics uncertainty.

One operator can supervise multiple robots

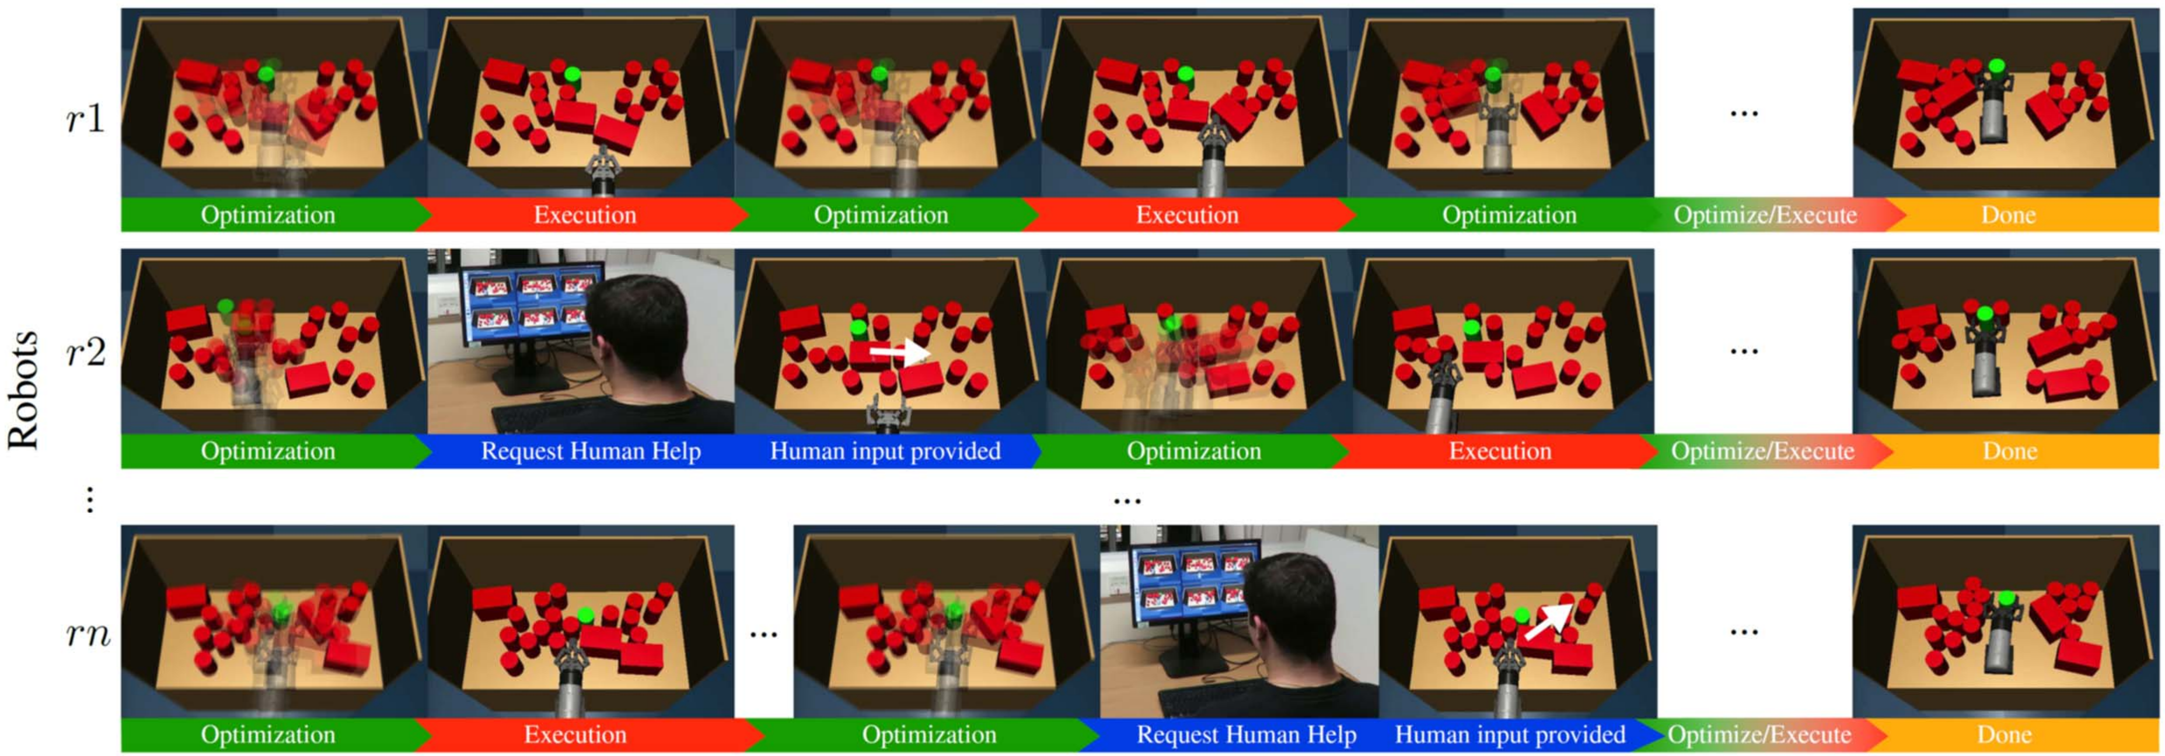

Human input indicated by white arrow

Key novel features

- Integration of human interaction into an online replanning system

  - Human help can be given dynamically, not just beforehand, as in previous work

- Trajectory optimisation using human input

  - Human input becomes part of the cost function during trajectory optimisation

- Efficient use of human time

  - Only ask for human input when likely to be beneficial
    1. If planning fails within some fixed time allocation
    2. If optimisation gets stuck at a local minimum

Higher success rate and faster

Table IV: Warehouse. Errors indicate 95% CI.

	Adaptive	Autonomous
Success	37 / 50	16 / 50
Failures	13 / 50	34 / 50
Optimization Time (s)	94.7 ± 15.1	149.7 ± 15.9
Human Time (s)	5.5 ± 1.0	-
Total Time (s)	112.2	152.7

Example of a human-in-the loop collaborative system

- Robot is autonomous where it can be
- Asks for help when stuck
- Efficient interaction with human

- Future work: learn from the hints!

Thanks:

UNIVERSITY OF LEEDS

To you for listening

For financial assistance from:

- EU project AI4EU 825619
- Marie Sklodowska-Curie projects 746143, 795714
- EPSRC projects EP/N509681/1, EP/P0195601/1, EP/R031193/1
- Alan Turing Institute

References

- Human-like planning for reaching in cluttered environments, M Hasan, M Warburton, W C Agboh, M R Dogar, M Leonetti, H Wang, F Mushtaq, M Mon-Williams, A G Cohn, ICRA-2020
- Optimization-based Motion Planning with Human-in-The-Loop for Non-Prehensile Manipulation, R Papallas, A G Cohn and M R Dogar, IEEE Robotics and Automation Letters 2020
- Human comfortability: Integrating ergonomics and muscular-informed metrics for manipulability analysis during human-robot collaboration, LFC Figueredo, RC Aguiar, L Chen, S Chakrabarty, M Dogar, A. Cohn. IEEE Robotics and Automation Letters, 2020/21

Questions?