Scalable Collaborative Human–Robot Learning (SCHooL)

Ken Goldberg, Pieter Abbeel, Anca Dragan, Stuart Russell, UC Berkeley

NSF NRI 2.0: Award 1734633:
Sept 2017 - Aug 2021

Scalable **C**ollaborative **H**uman–R**o**b**o**t **L**earning

(SCHooL)

Sept 2017 - Aug 2021


Ken Goldberg


Pieter Abbeel


Anca Dragan


Stuart Russell

Berkeley
UNIVERSITY OF CALIFORNIA

BAIR
BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

AUTOLAB

schoolproject.berkeley.edu

Learning from Demonstrations

- Imitation Learning
- Inverse Reinforcement Learning (IRL)
  - Pioneered by co-PI Russel (1998)
  - Apprenticeship Learning (Abbeel & Ng, 2004)
- HRI Optimizing Legibility (Dragan, 2013)
- Cloud Robotics (Kuffner 2015, Kehoe, Abbeel, Goldberg, 2017)

Cooperative IRL

- A CIRL is a 2-player cooperative Markov game $\langle \mathcal{S}, (\mathcal{A}^H, \mathcal{A}^R), P, R_\theta \rangle$
- Human and robot take simultaneous actions, get same reward parametrized by θ
- Human preference R_θ initially unknown to the robot
- This incentivizes the human to teach and the robot to learn this preference
- Both agents can reason about the robot's belief state, making it a sufficient state representation (together with the environment state)
- Example: human signals which objects should not be decluttered by replacing them in the environment when the robot removes them

4 Research Objectives

$$\tau^{\mathbf{H}} \leftarrow \underset{\tau}{\operatorname{argmax}}\, \phi(\tau)^\top \theta - \eta ||\phi_\theta - \phi(\tau)||^2$$

1. Extend CIRL Formal Framework:
2. Distributed Sensing, Reward Models using Deep Learning
3. Learning Hierarchical Task and Reward Structure
4. Bidirectional / Active Human–Robot Communication

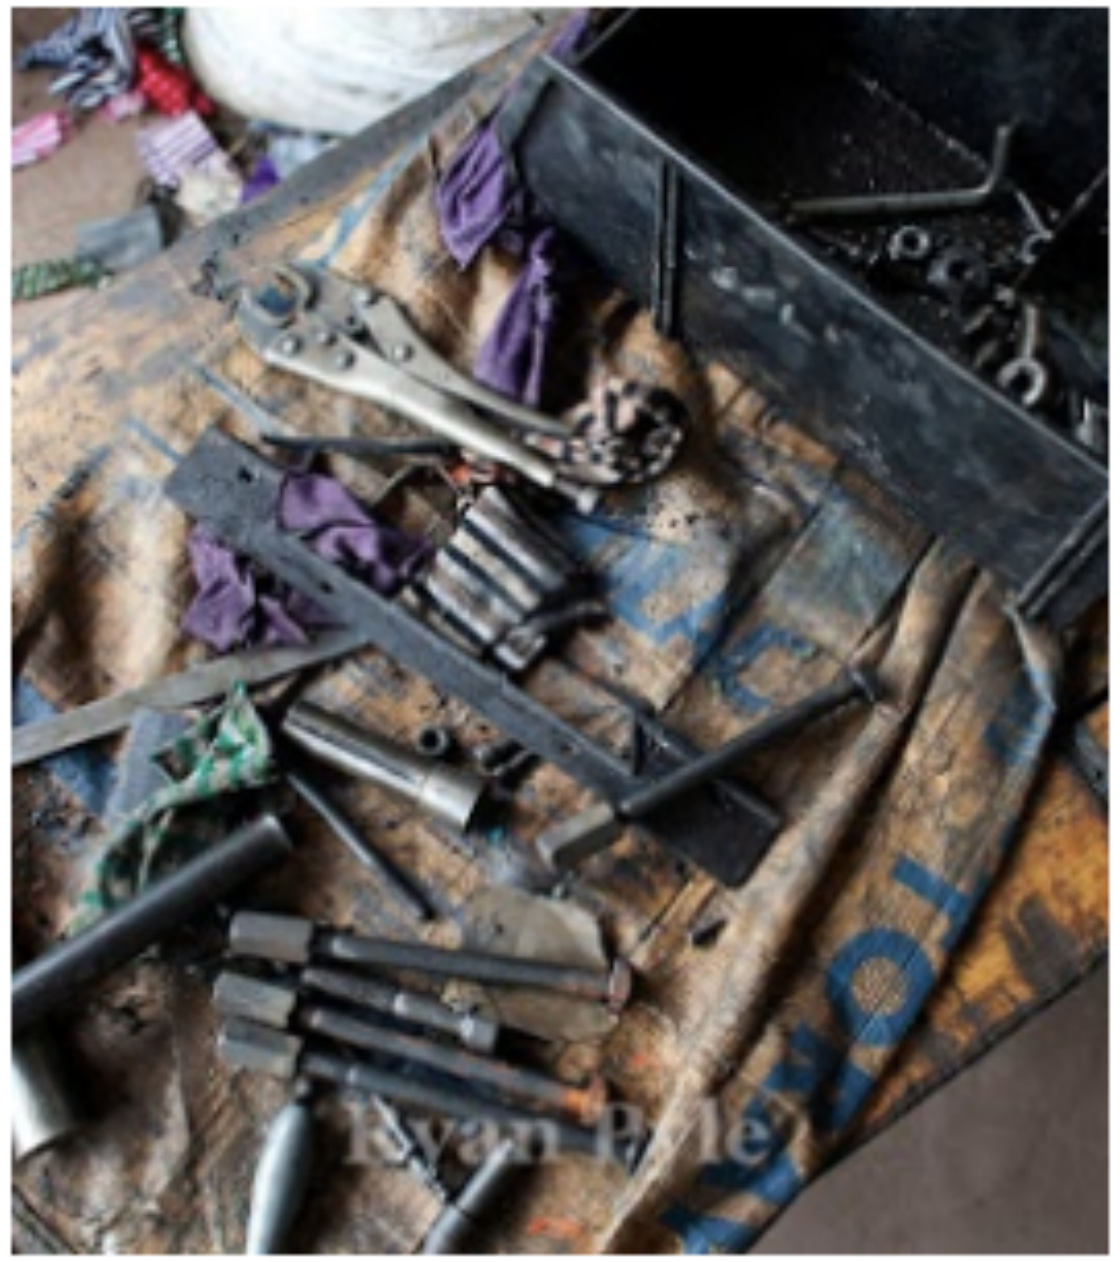

Integrative Application:

Surface Decluttering

To increase productivity and safety in homes, machine shops, warehouses, offices, and retail stores.

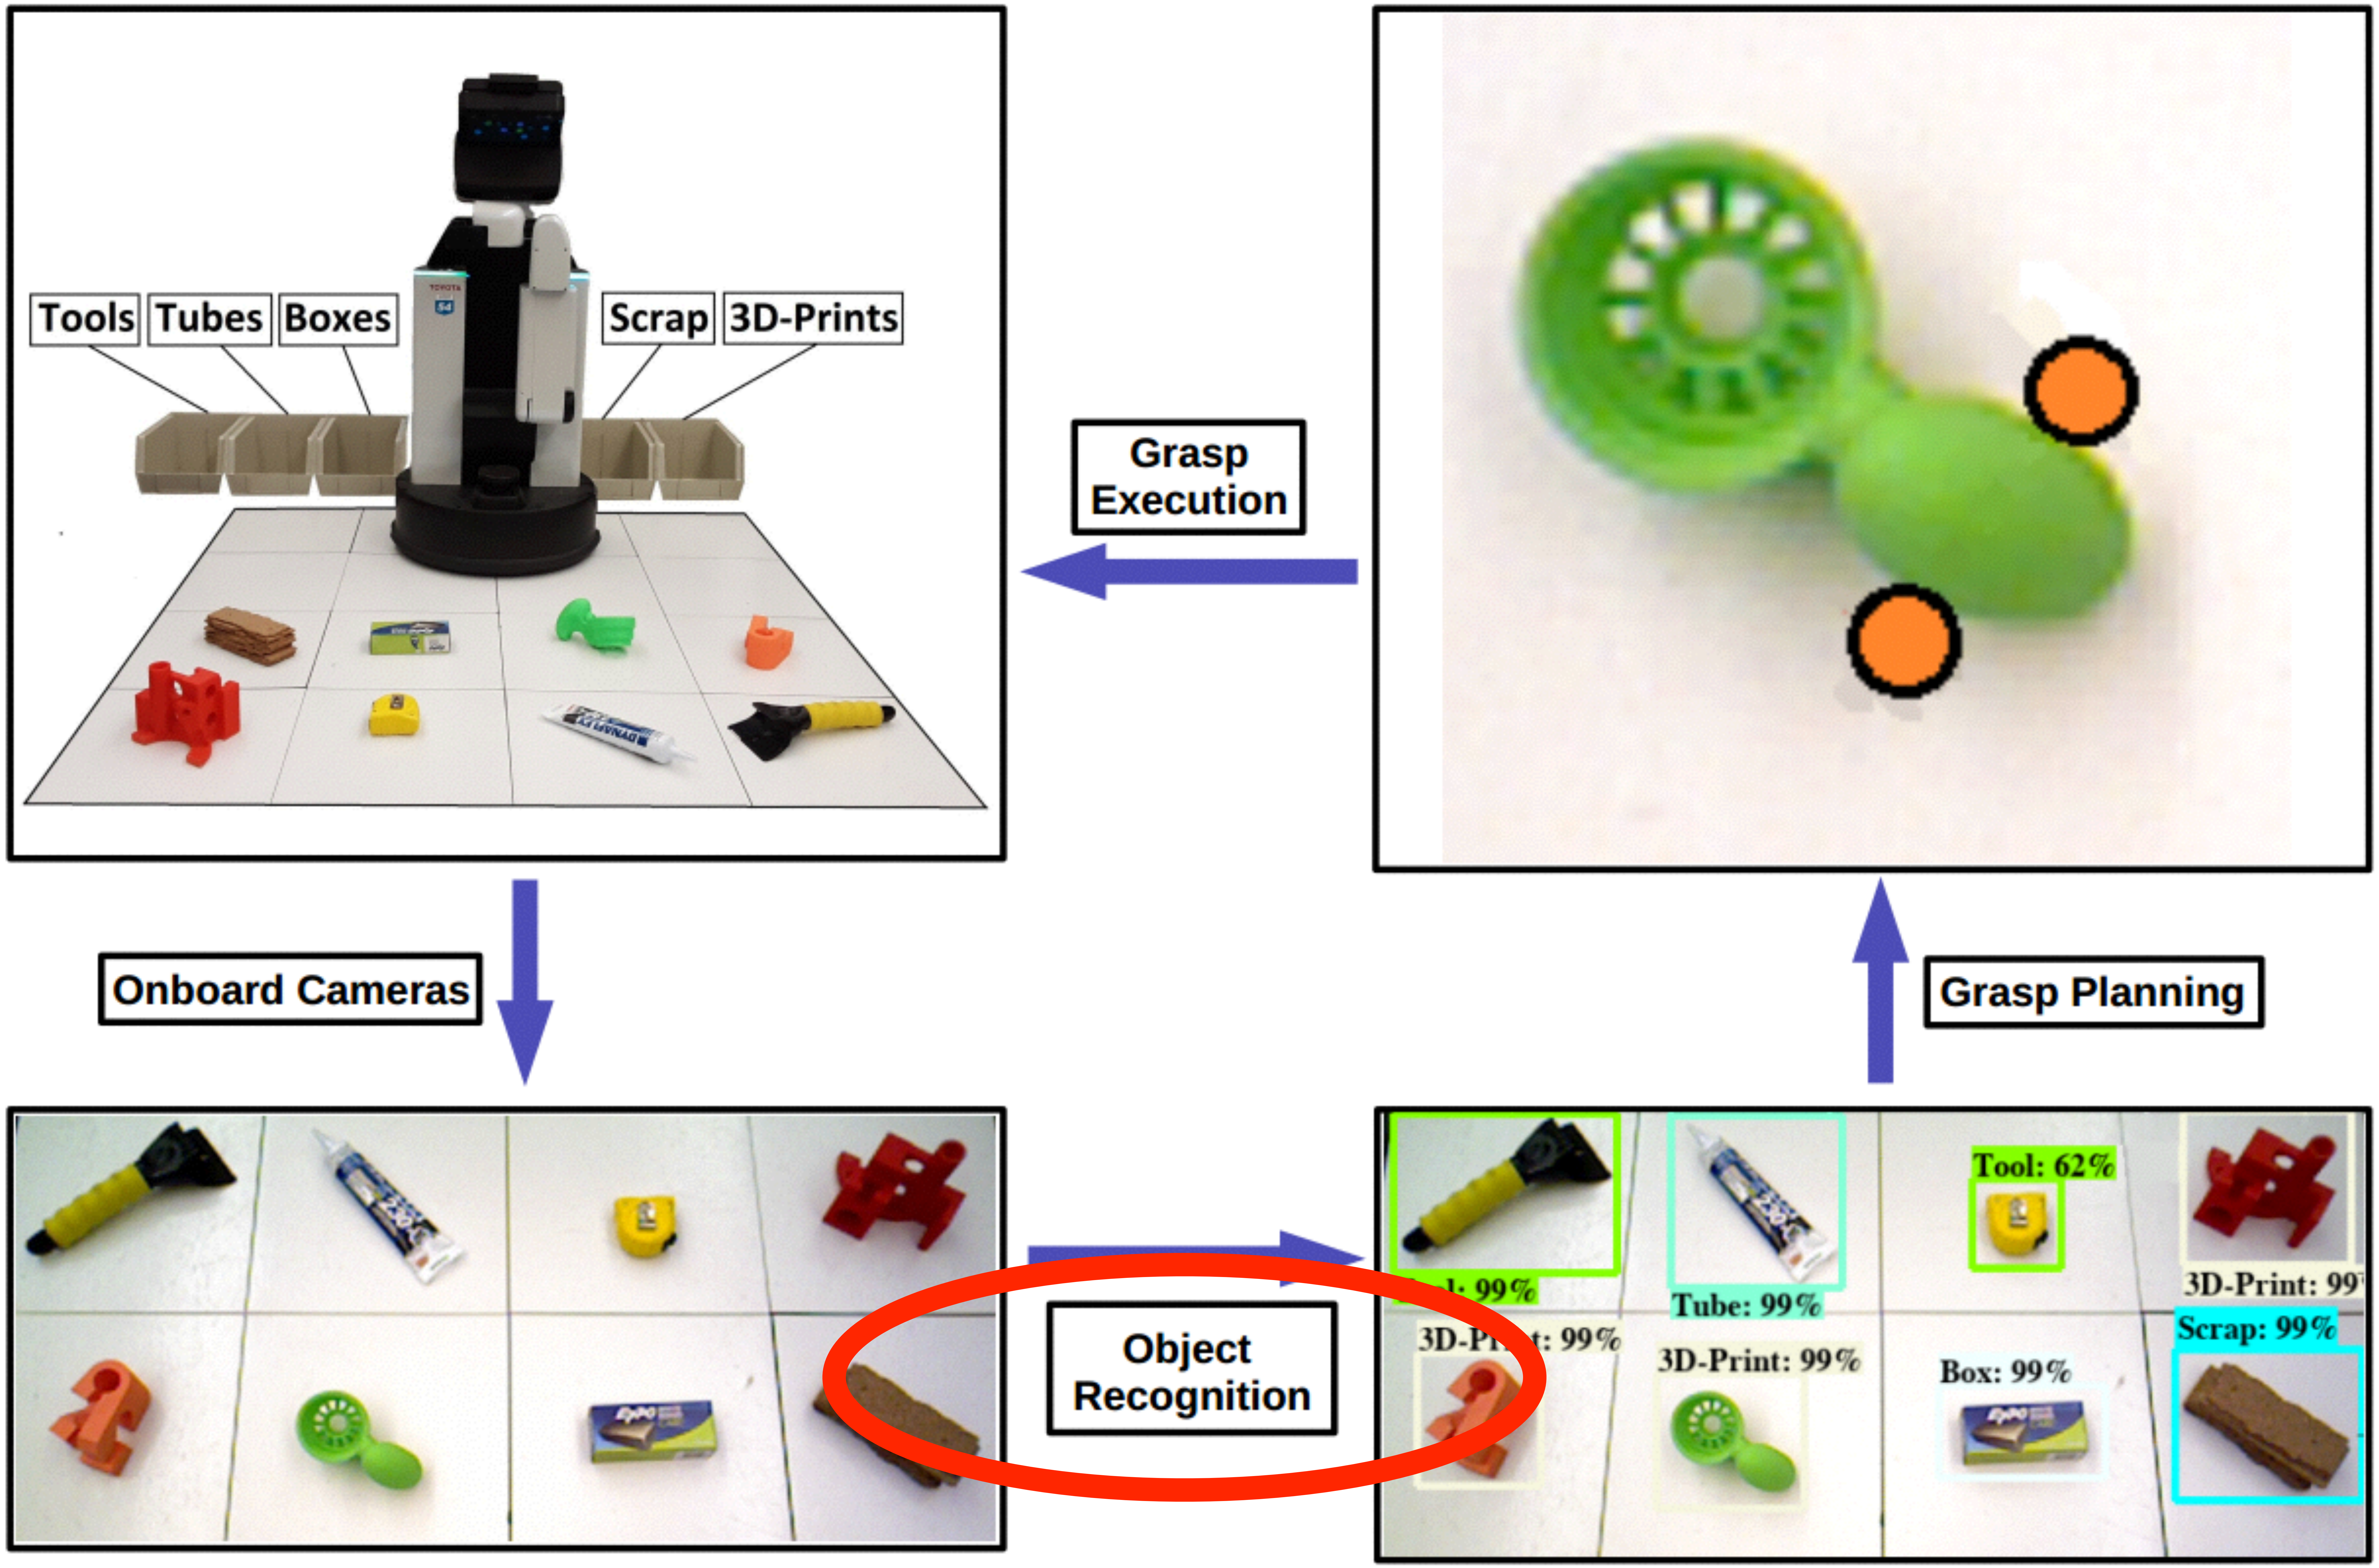

Tools
Tubes
Boxes
Scrap
3D-Prints
Grasp Execution
Onboard Cameras
Grasp Planning
Object Recognition
Tool: 62%
Tube: 99%
3D-Print: 99%
Box: 99%
Scrap: 99%

Cloud

Local Severs (eg. 5G):

Edge Robots:

Robust Task-Based Grasping as a Service

Jingyi Song[1], Ajay Tanwani[1], Jeffrey Ichnowski[1], Michael Danielczuk[1], Kate Sanders[1], Jackson Chui[1], Juan A. Ojea[2], Ken Goldberg[1]

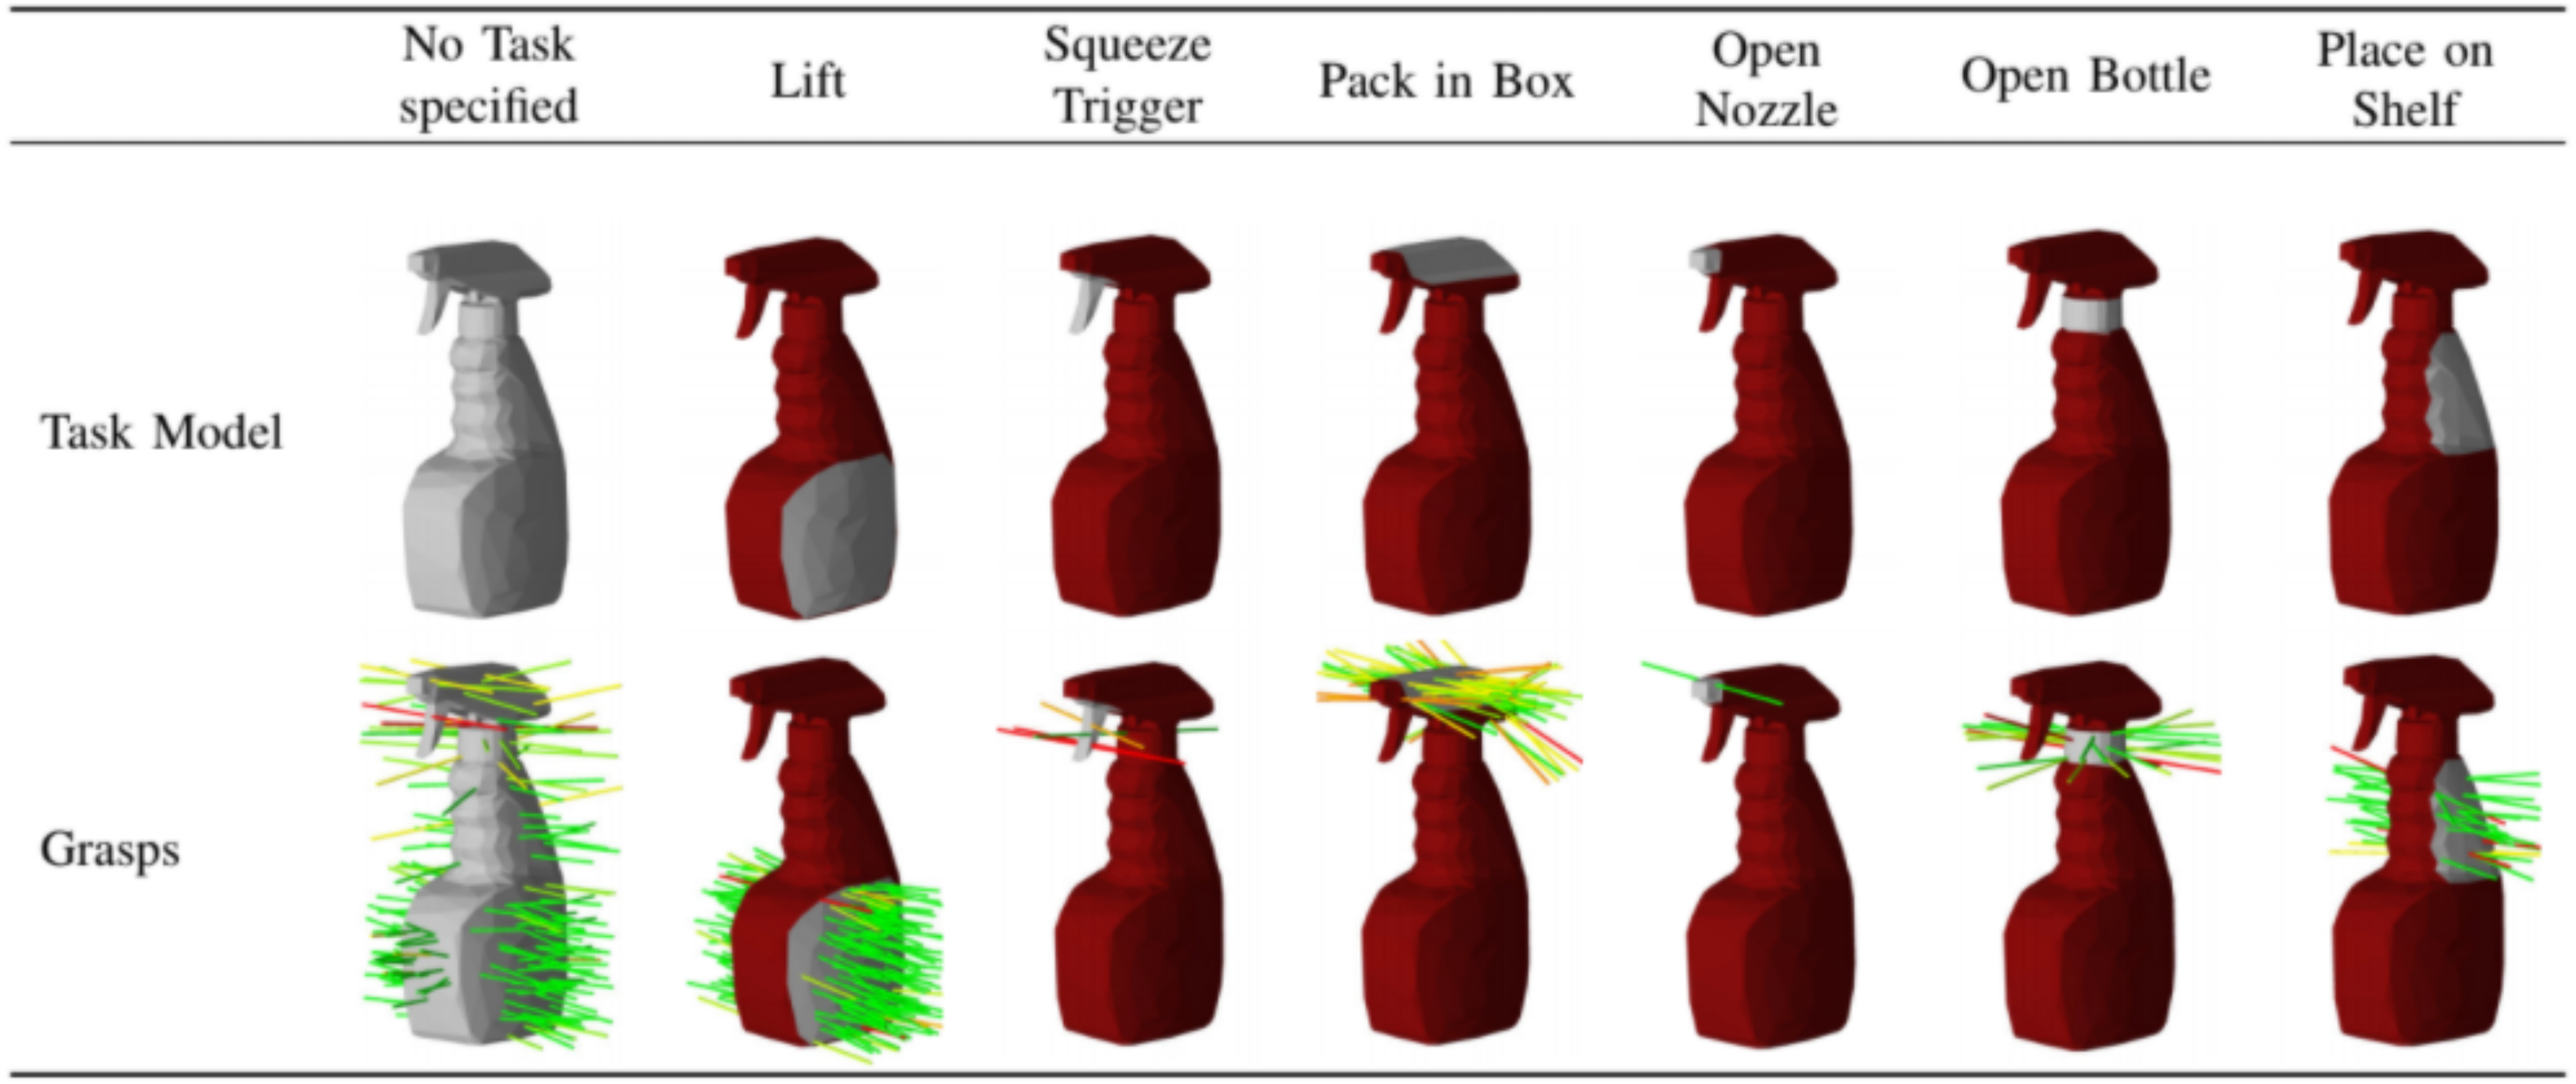

Fig. 3: **Task-directed grasping for spray bottle.** The stay-out zones of the spray bottle object displayed on the second row model the tasks described on the first row, resulting filtered grasps showed on the third row.

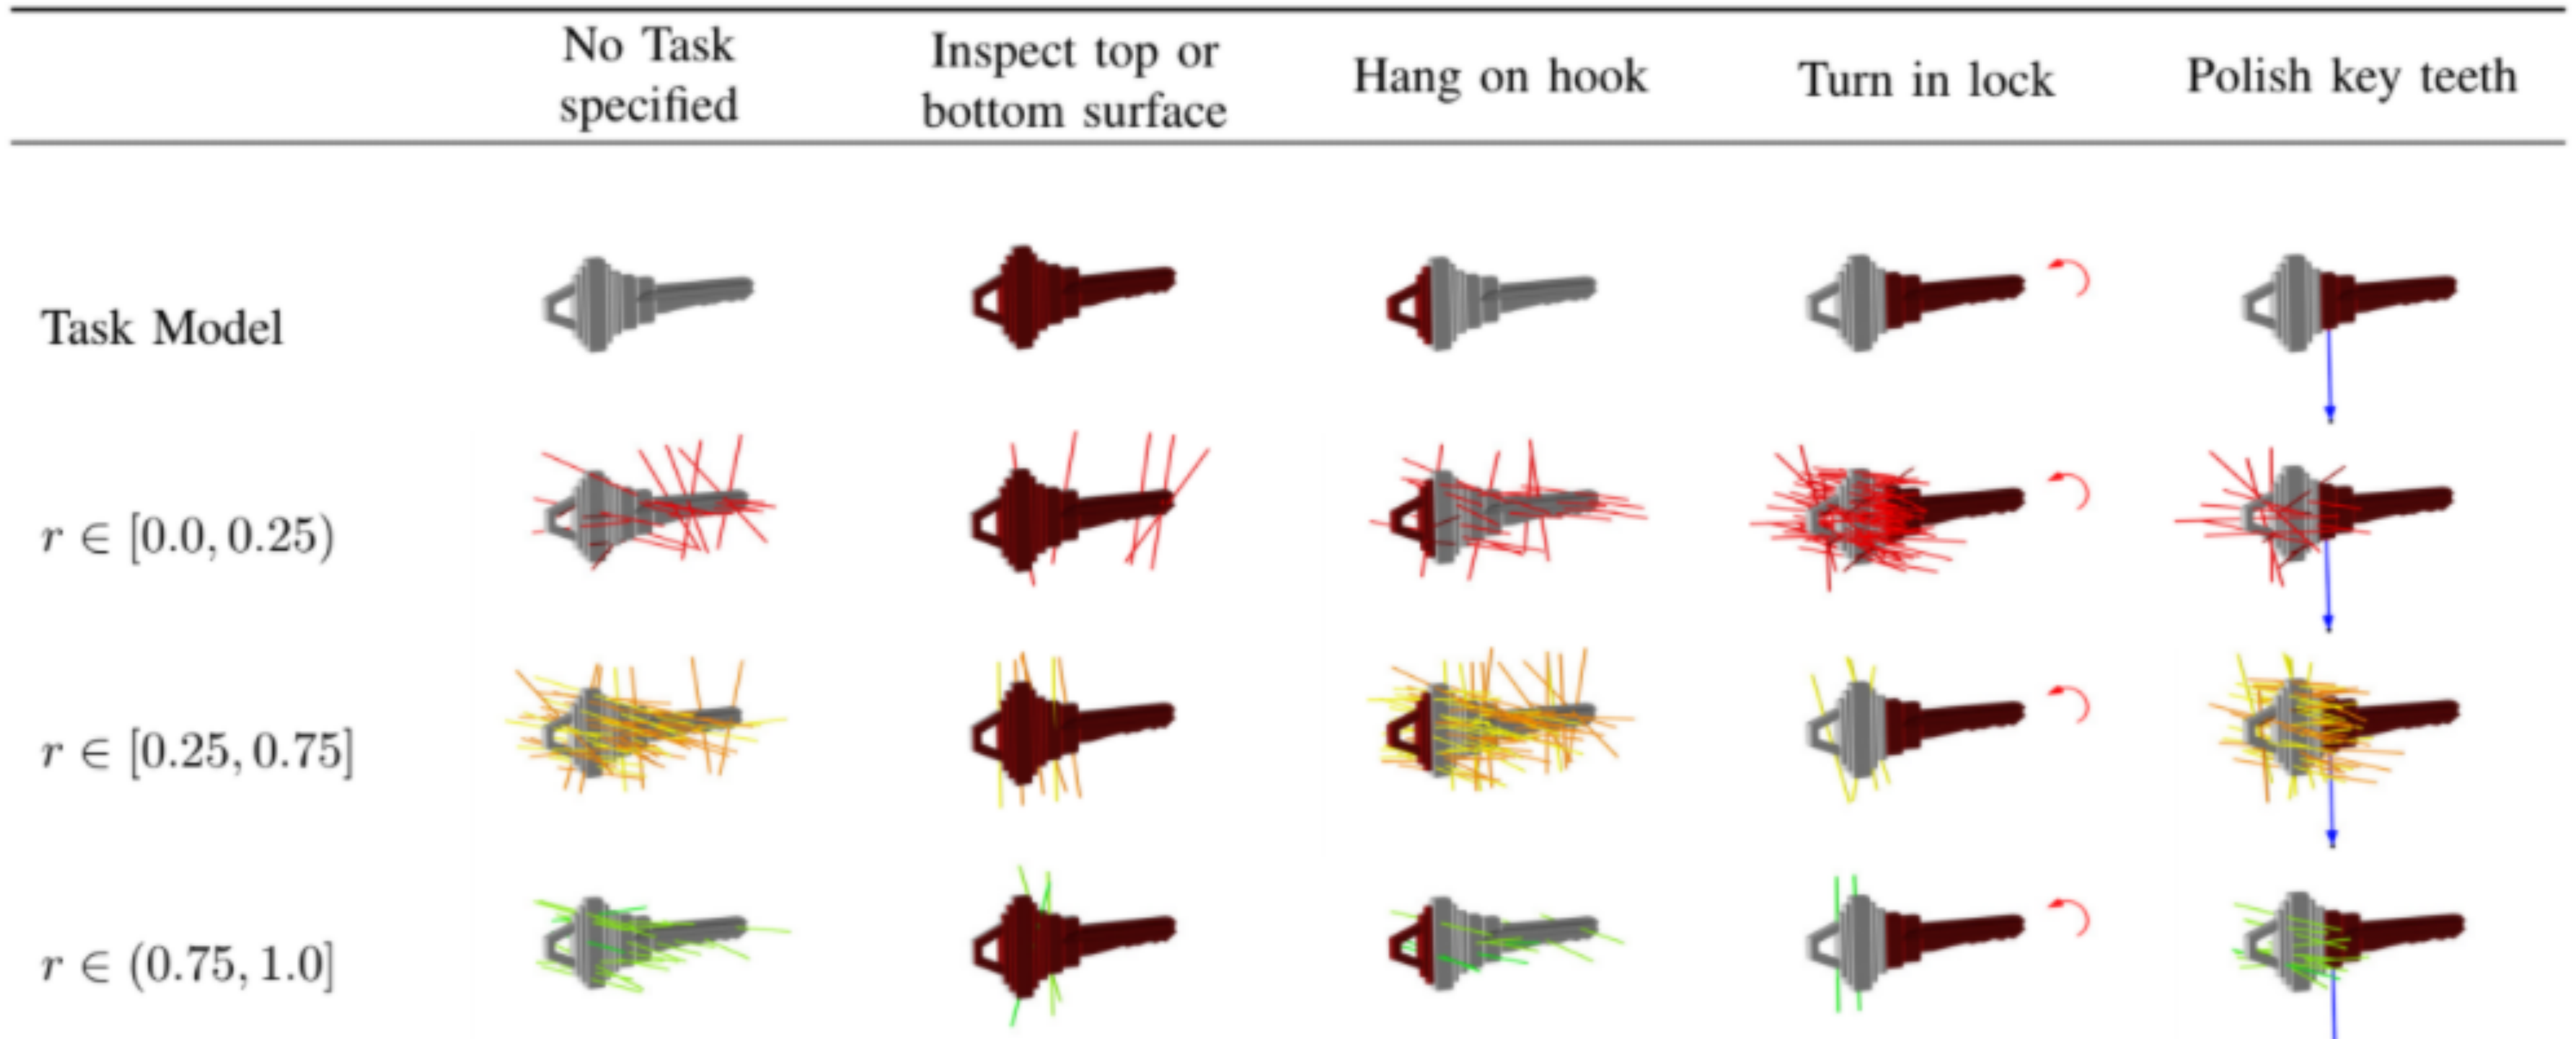

Preference learning along multiple criteria: A game-theoretic perspective

Kush Bhatia
EECS, UC Berkeley
kush@cs.berkeley.edu

Ashwin Pananjady
Simons Institute, UC Berkeley
ashwinpm@berkeley.edu

Peter L. Bartlett
EECS and Statistics, UC Berkeley
peter@berkeley.edu

Anca D. Dragan
EECS, UC Berkeley
anca@berkeley.edu

Martin J. Wainwright
EECS and Statistics, UC Berkeley
wainwrig@berkeley.edu

NeurIPS 2020

Multi-Principal Assistance Games: Definition and Collegial Mechanisms

Arnaud Fickinger
Department of EECS
University of California, Berkeley
arnaud.fickinger@berkeley.edu

Simon Zhuang
Department of EECS
University of California, Berkeley
simonzhuang@berkeley.edu

Andrew Critch
Department of EECS
University of California, Berkeley
critch@berkeley.edu

Dylan Hadfield-Menell
Department of EECS
University of California, Berkeley
dhm@berkeley.edu

Stuart Russell
Department of EECS
University of California, Berkeley
russell@berkeley.edu

NeurIPS 2020

‹ BACK TO EVENTS

PRIVATE EVENT

EECS Special Seminar: Dylan Hadfield-Menell "The Value Alignment Problem in Artificial Intelligence"

APRIL 06 2020

4:00P - 5:00P

LOCATION
online

SPEAKER
Dylan Hadfield-Menell
UC Berkeley

HOST
David Sontag
CSAIL

LazyDAgger: Reducing Context Switching in Interactive Robot Imitation Learning

Ryan Hoque[1], Ashwin Balakrishna[1], Brijen Thananjeyan[1], Carl Putterman[1], Michae
Daniel Seita[1], Daniel S. Brown[1], Ken Goldberg[1]

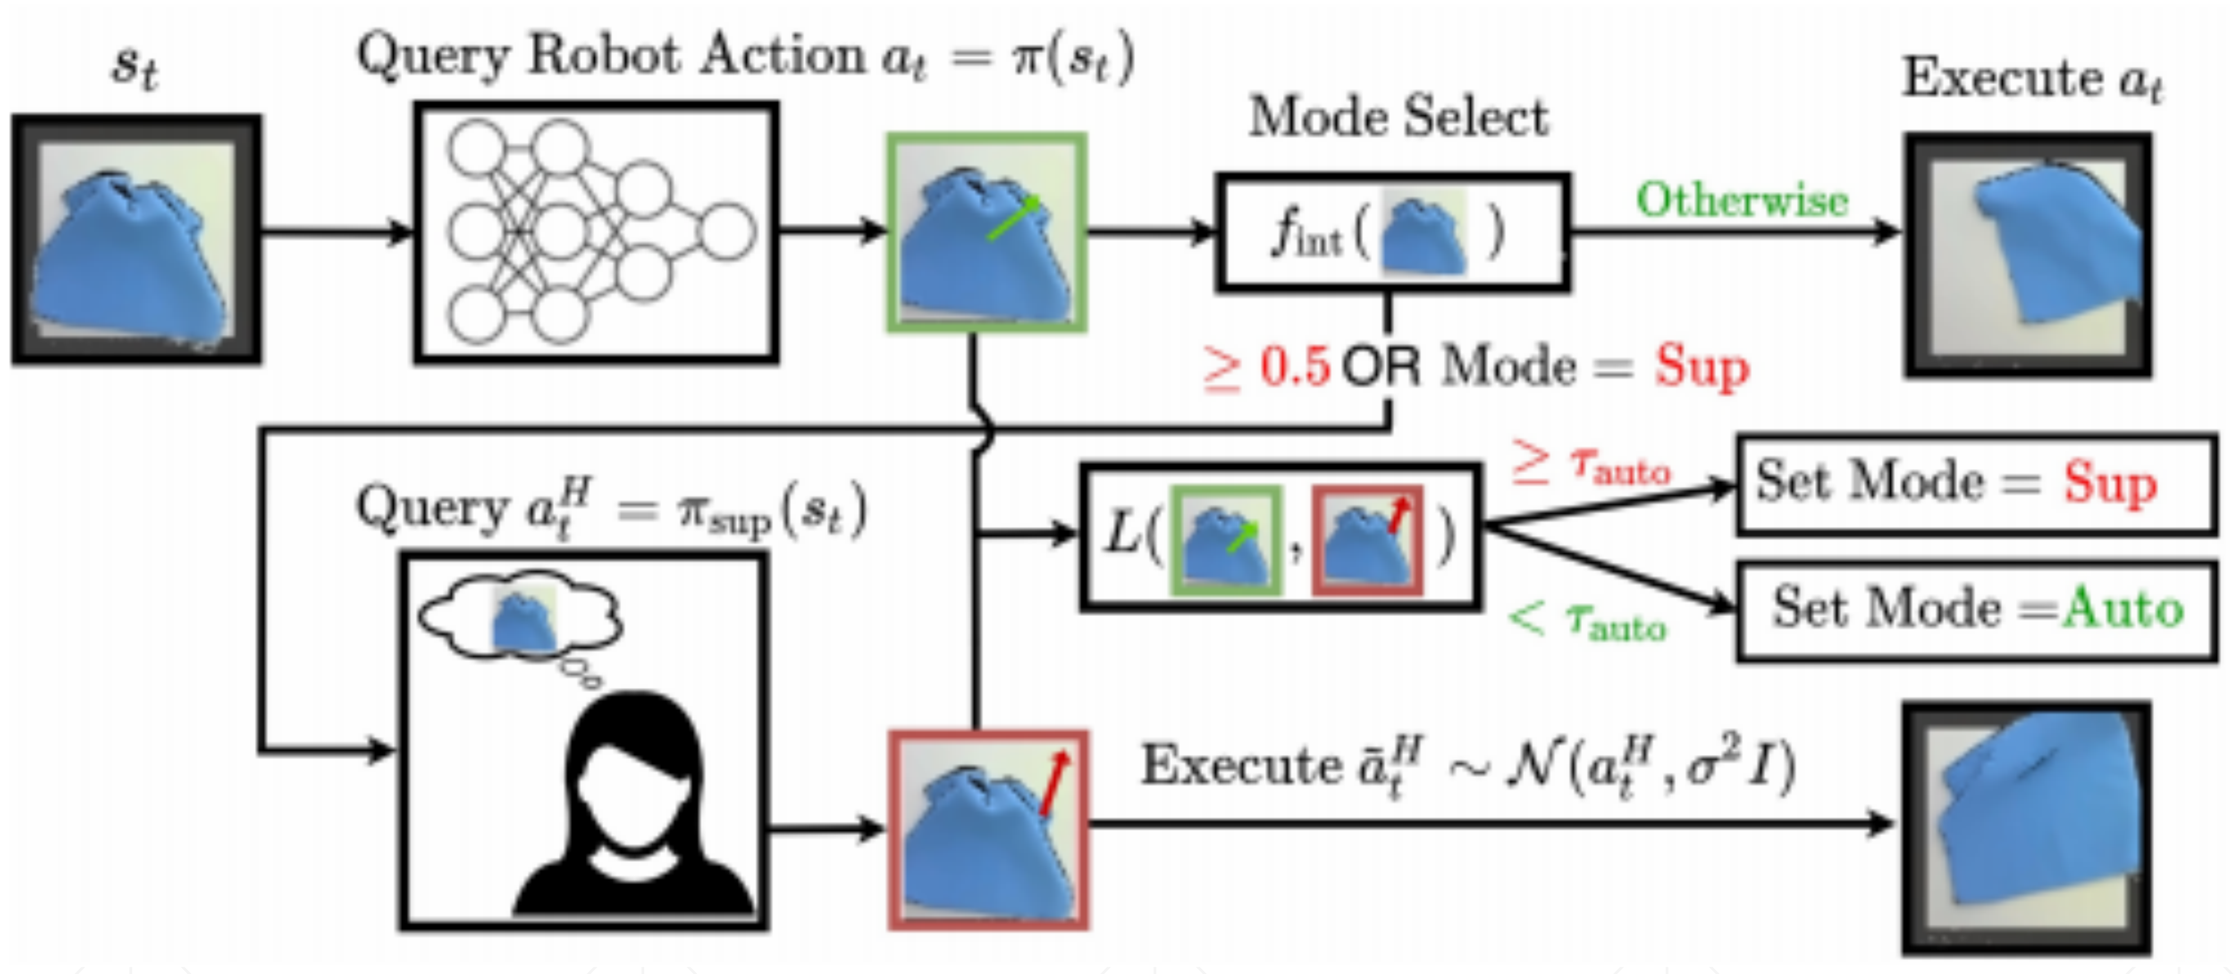

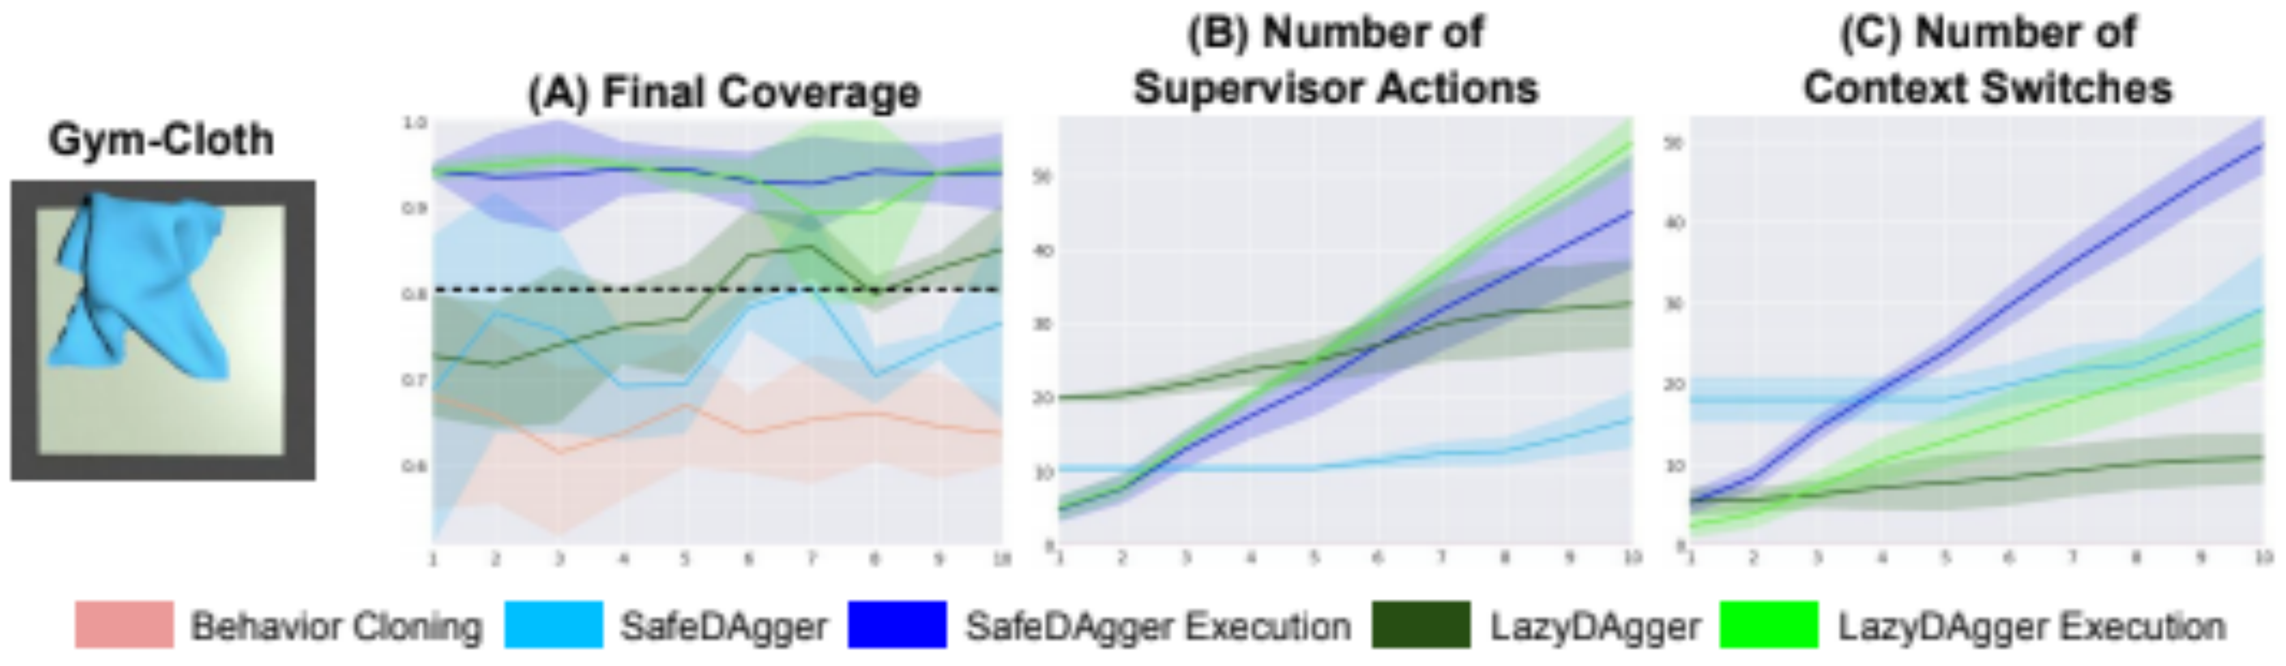

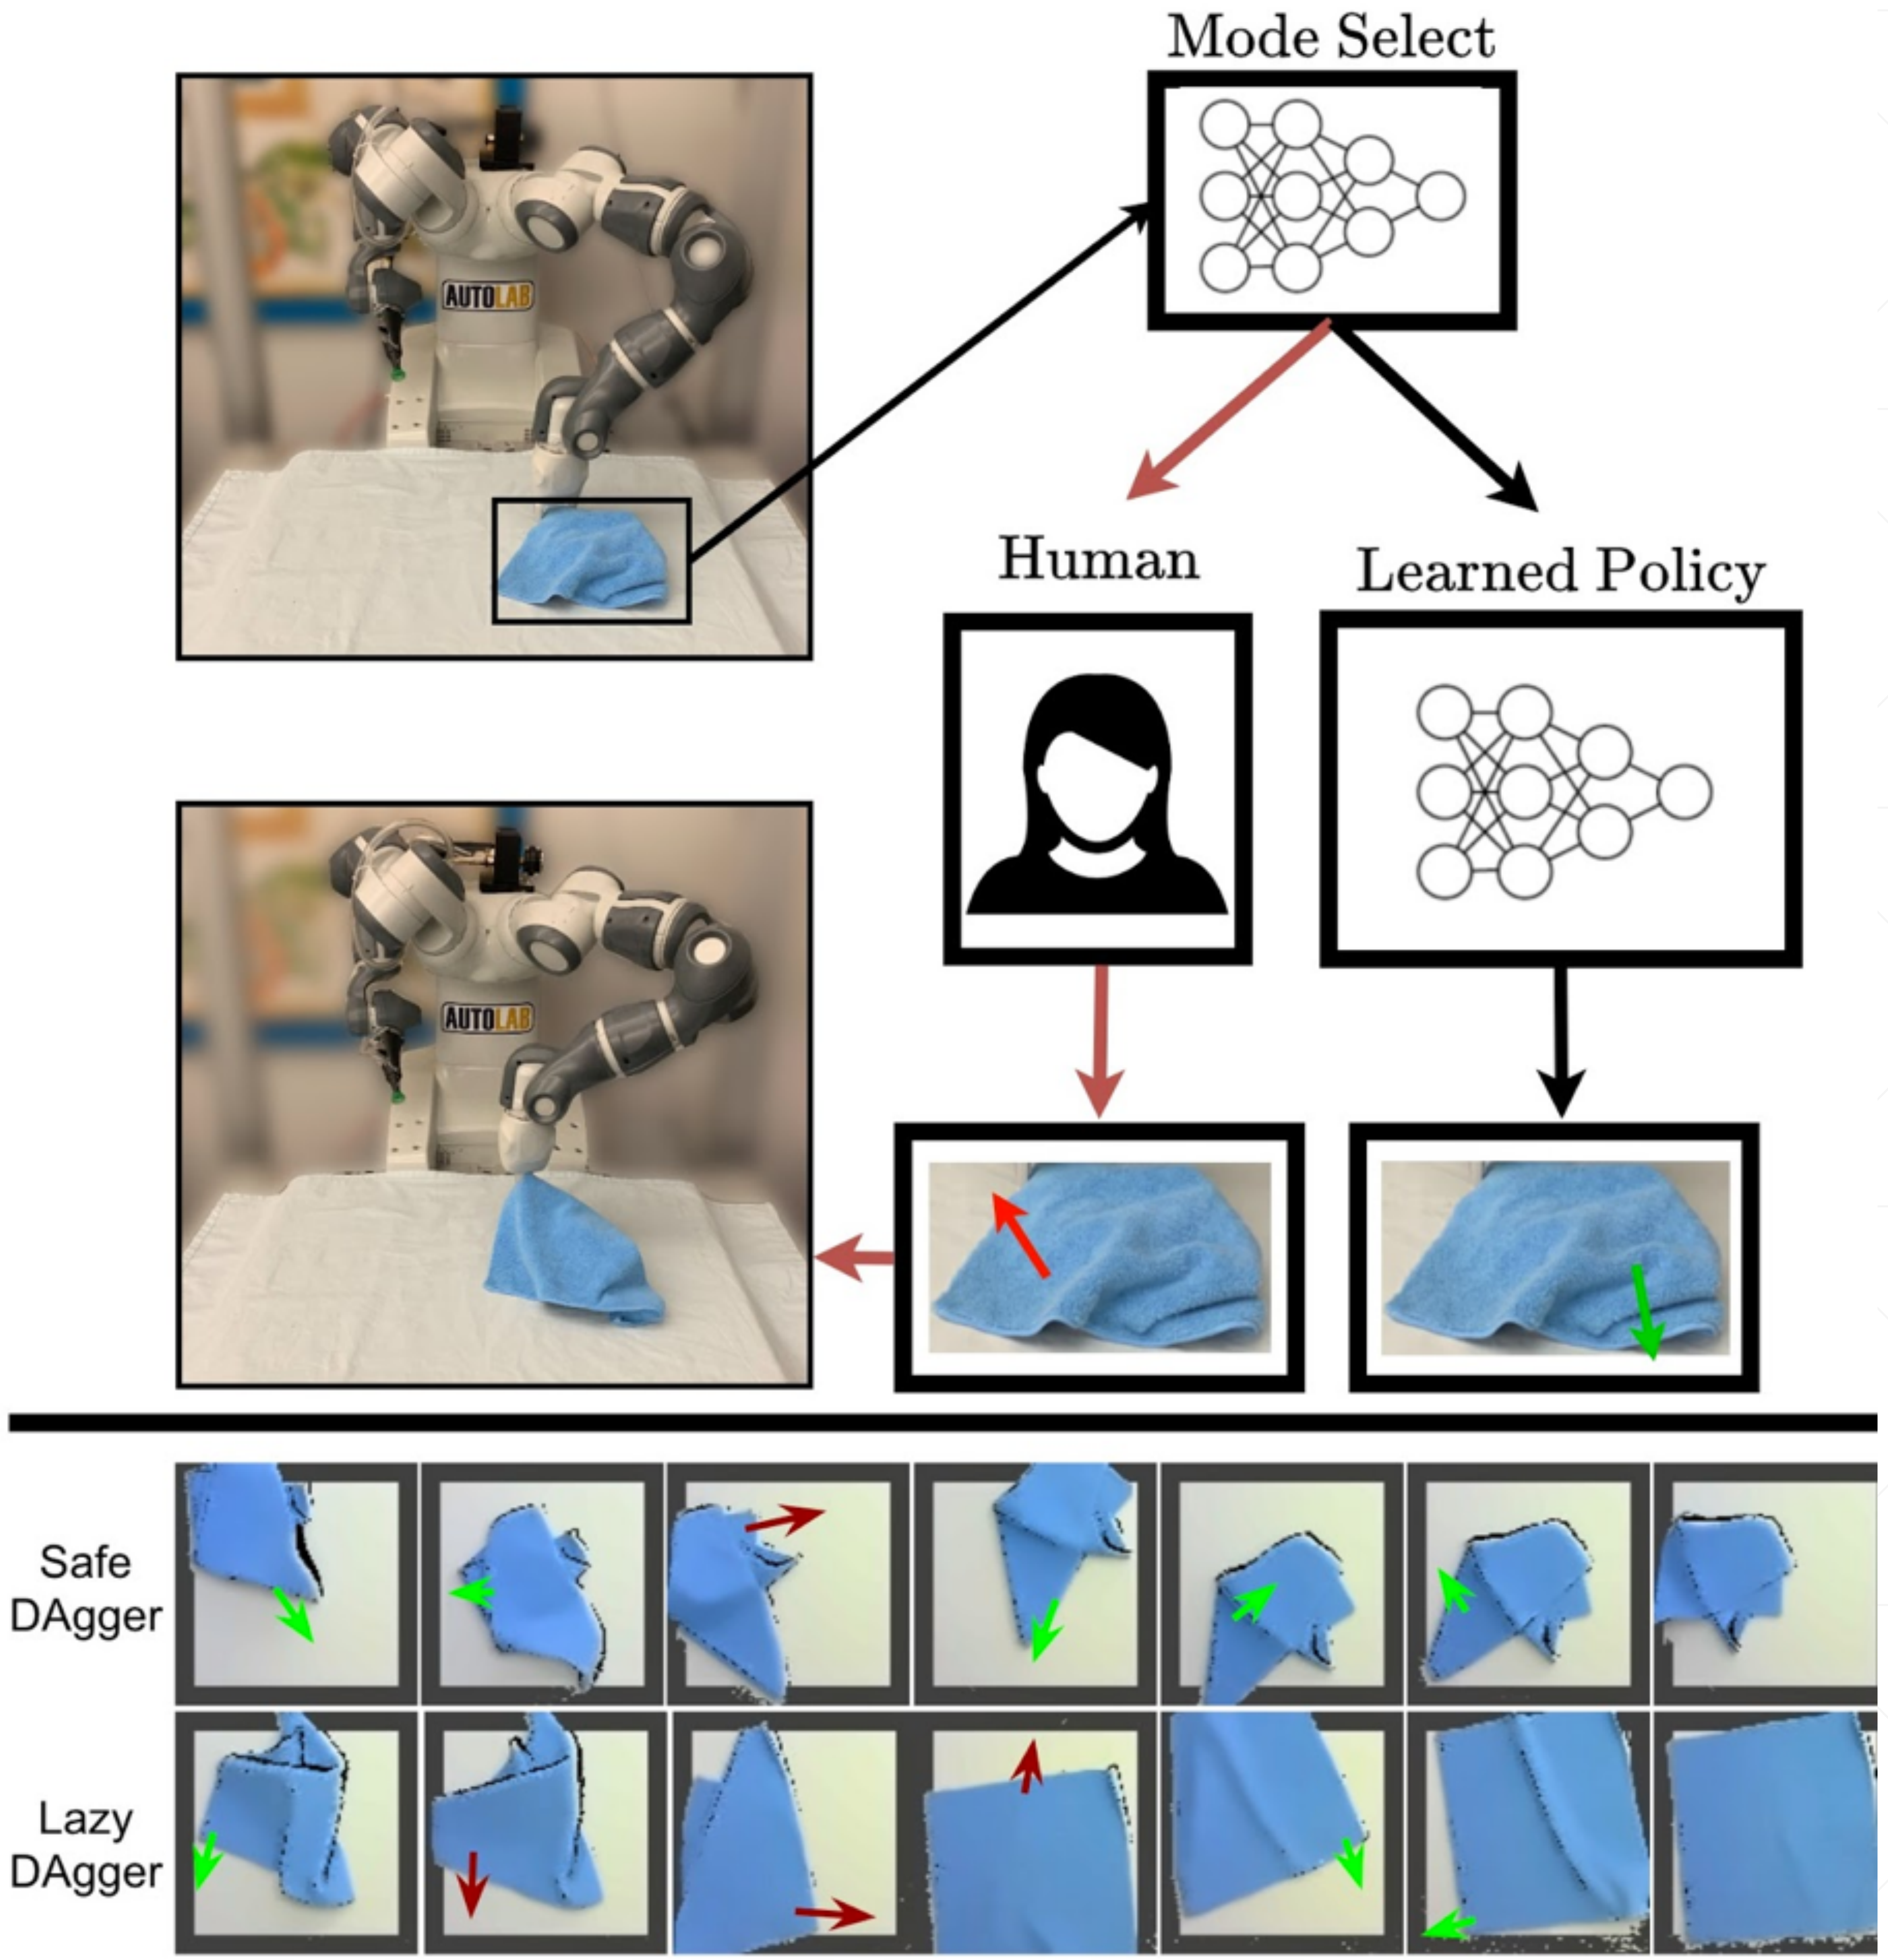

Broader Impacts

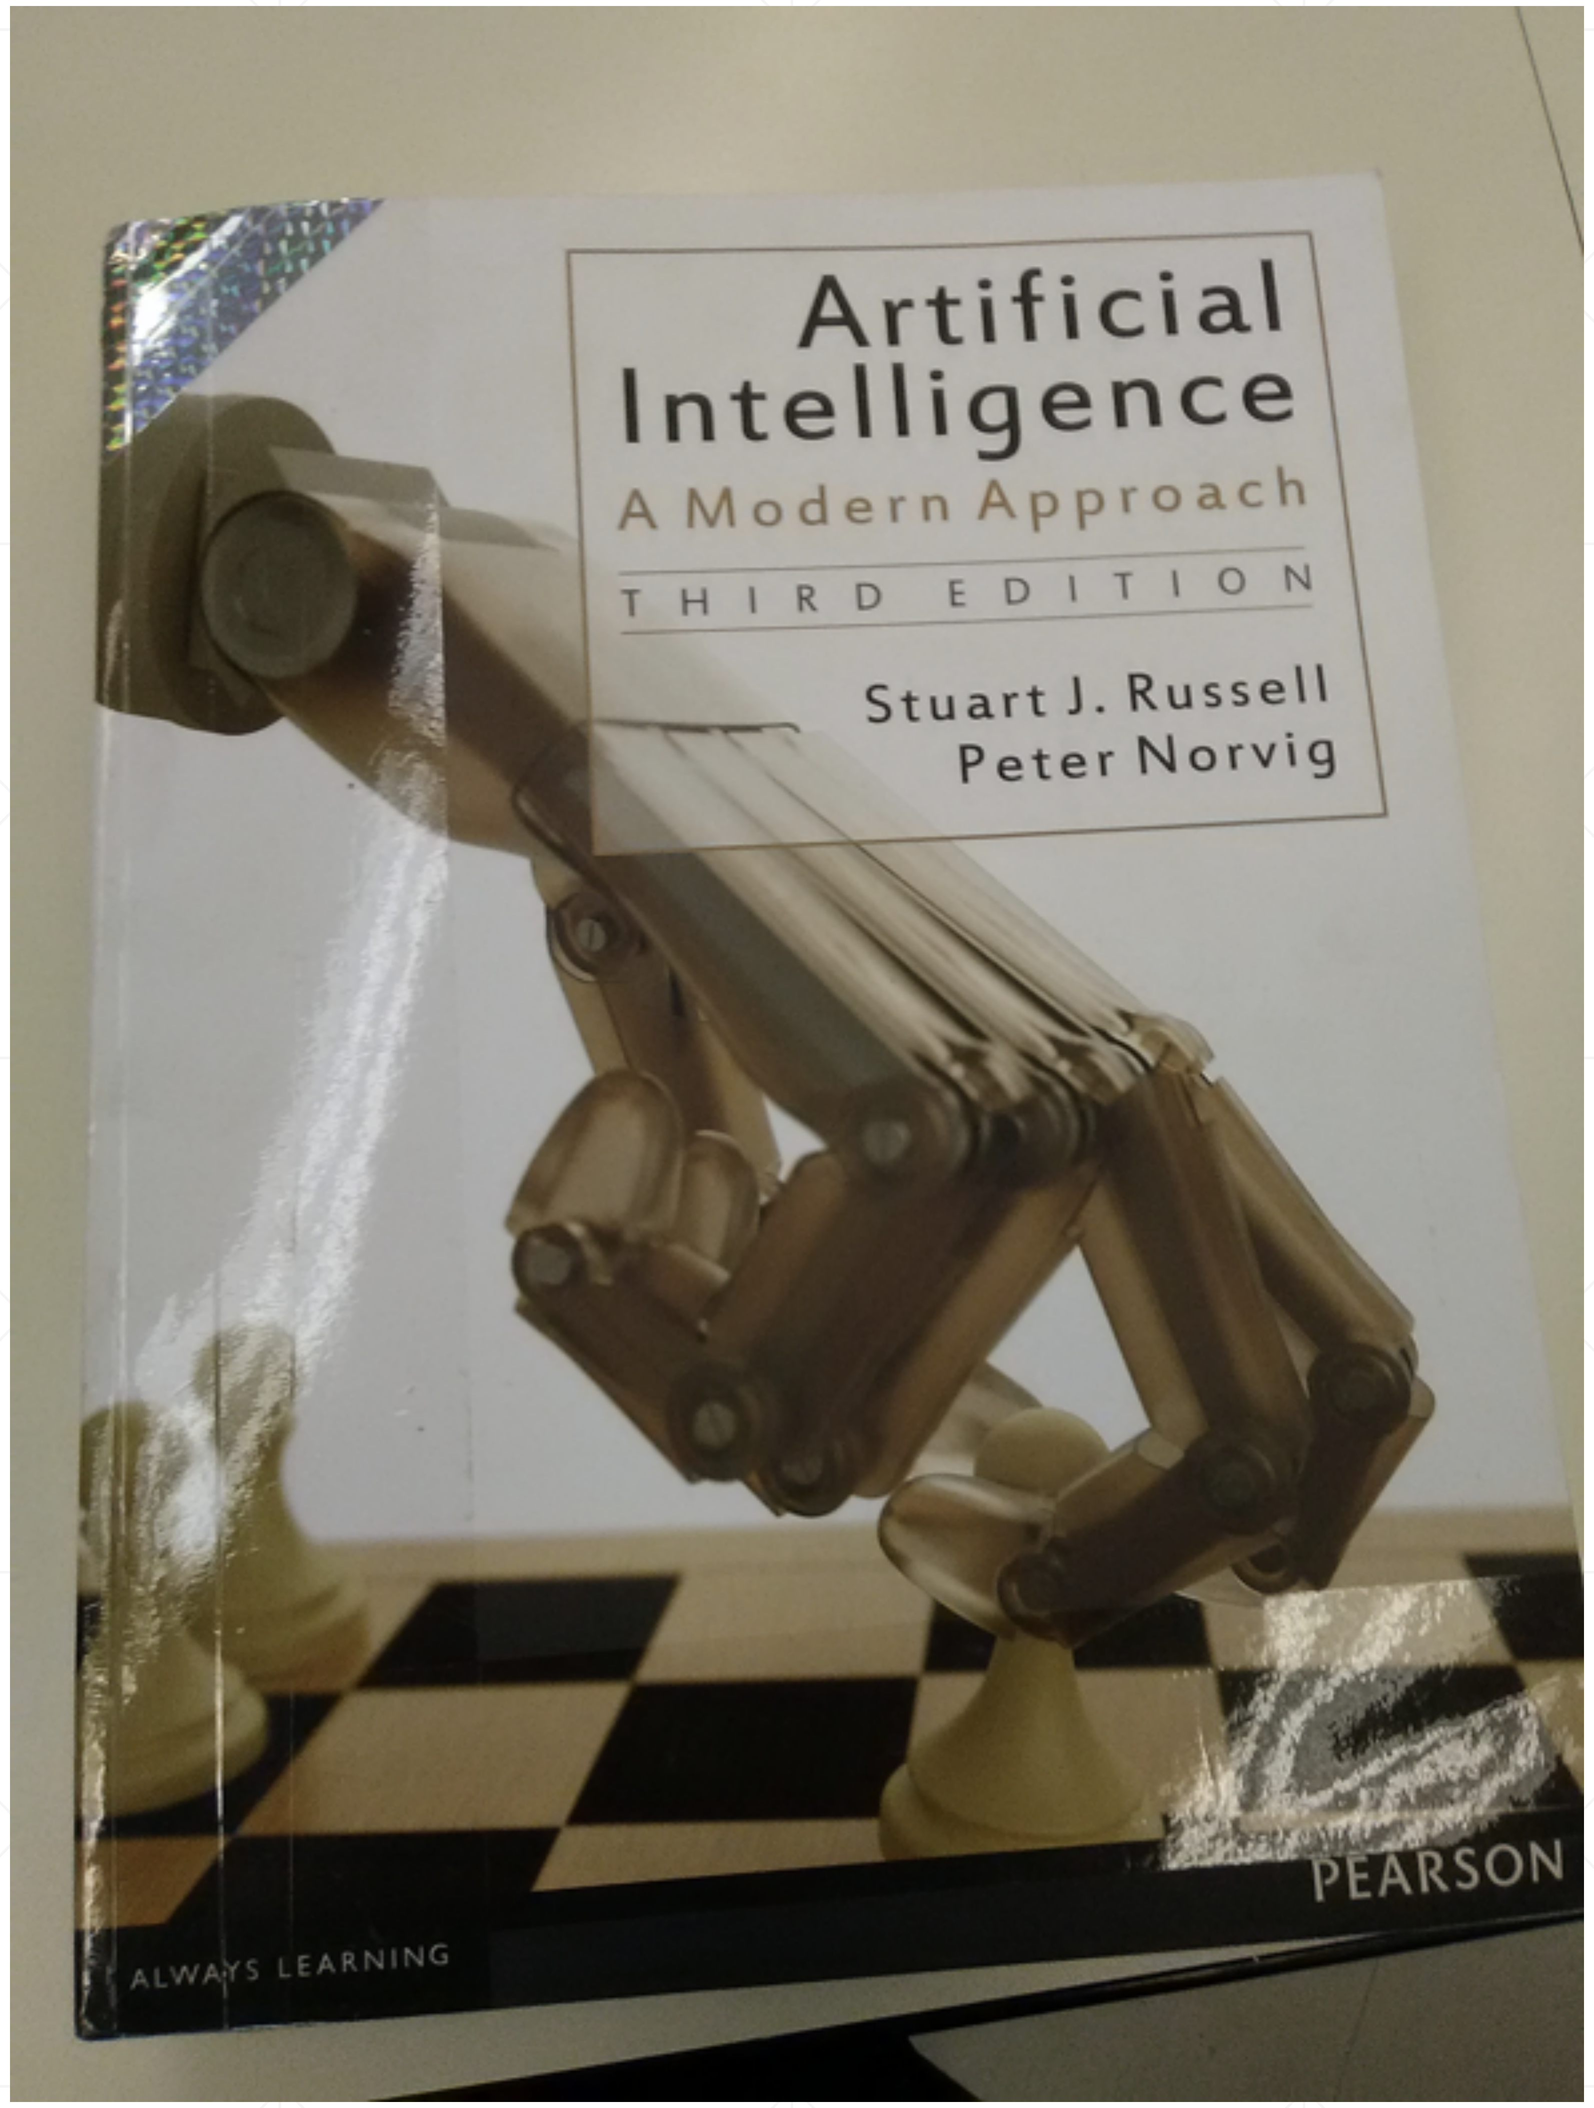

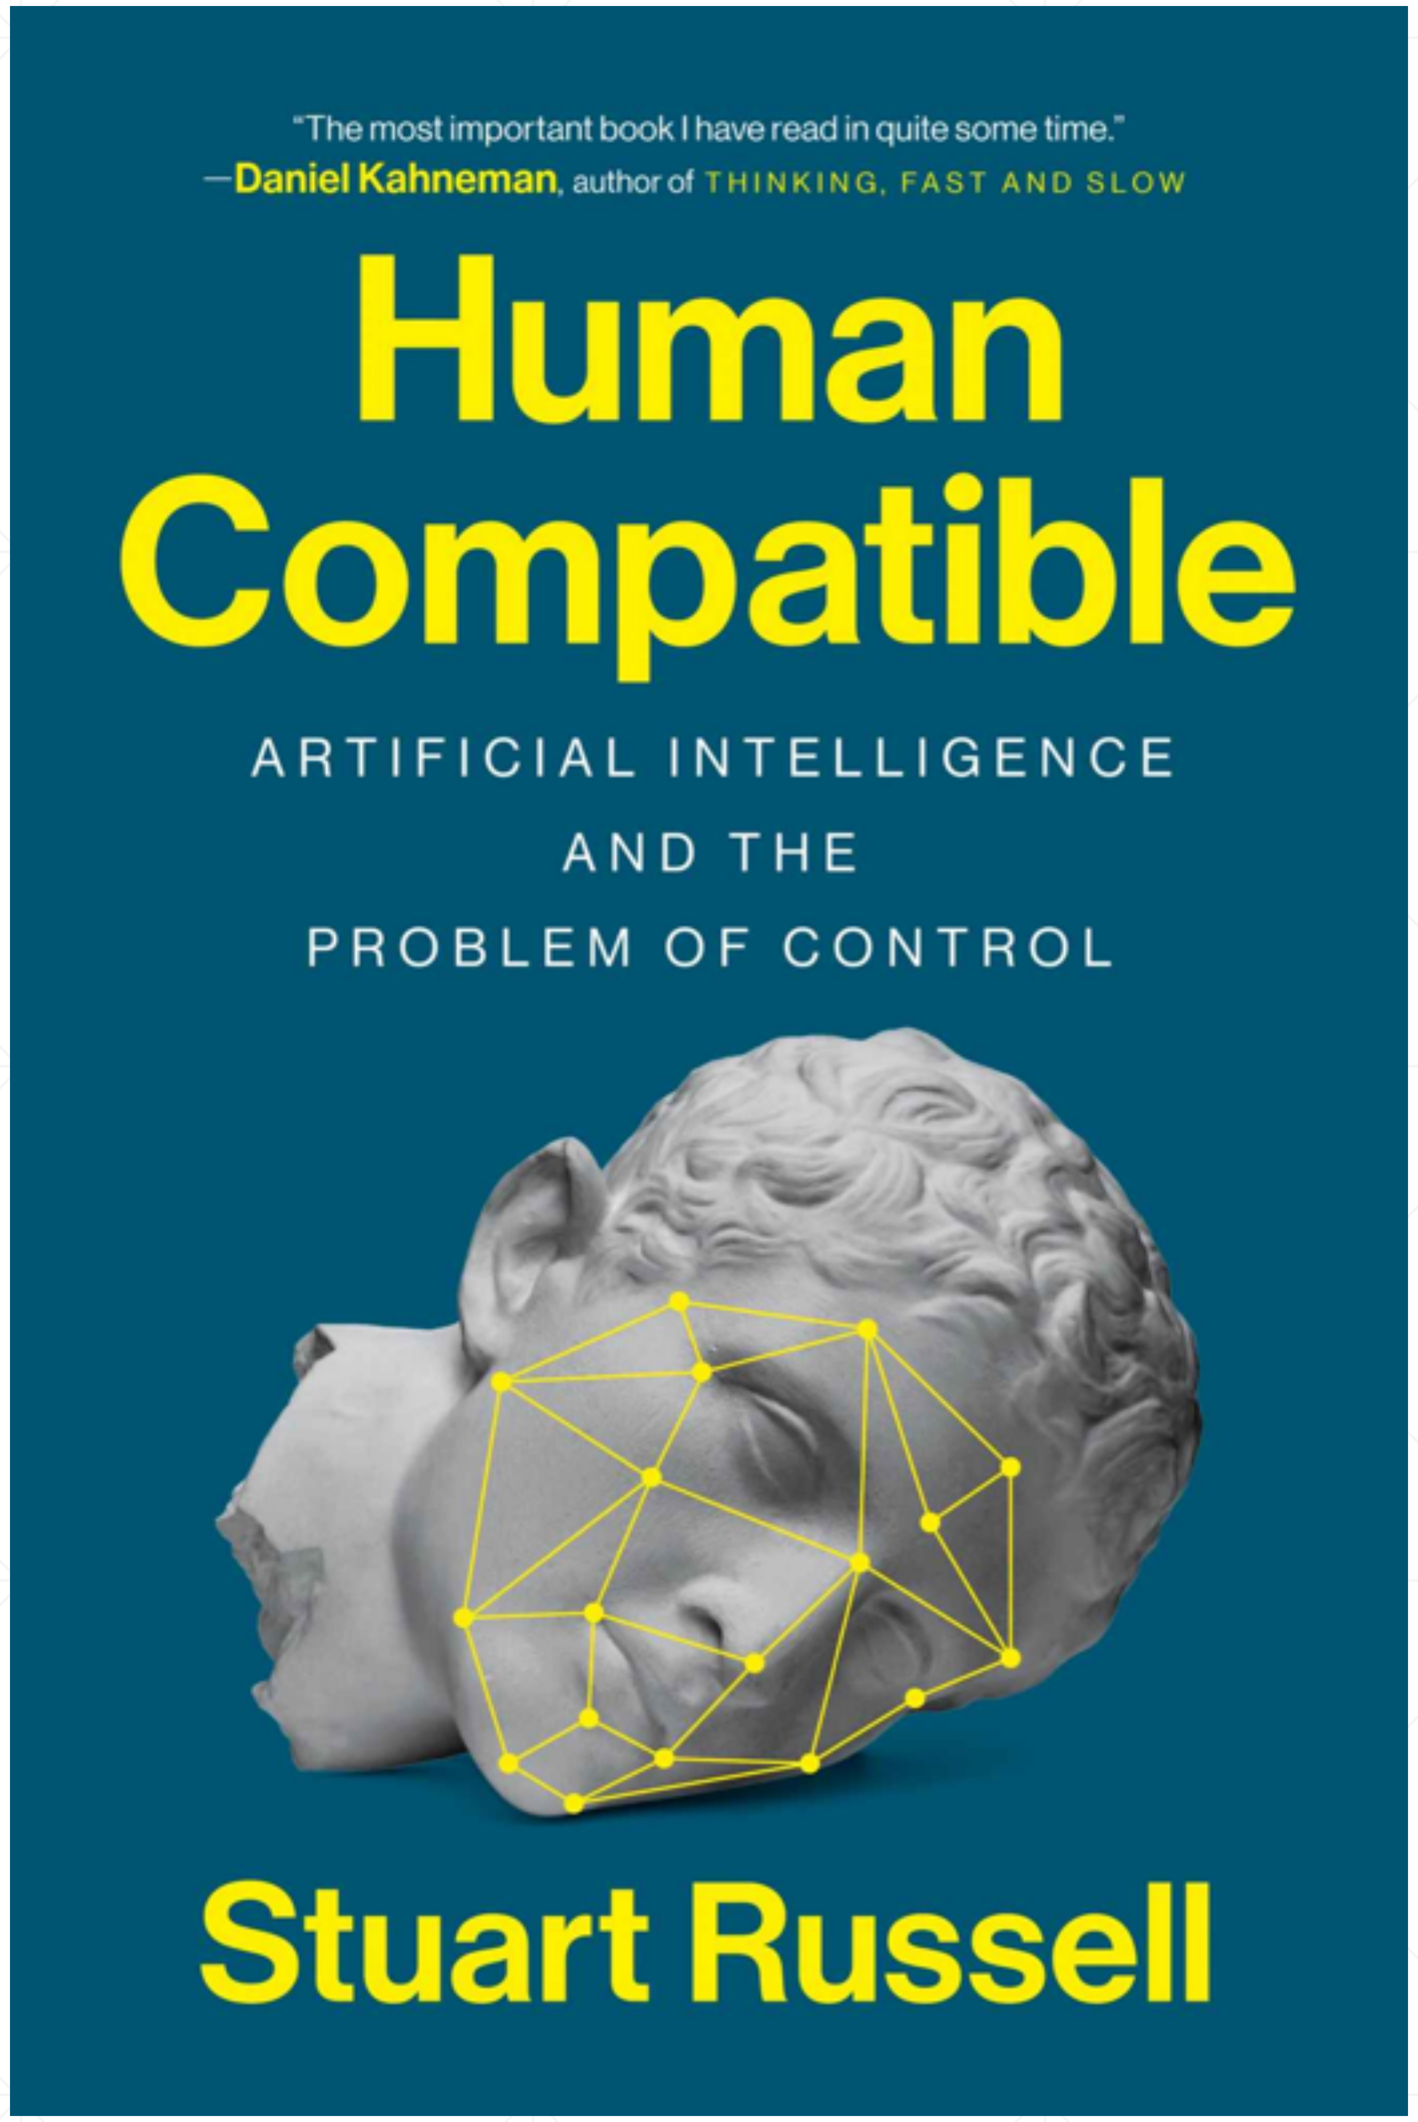

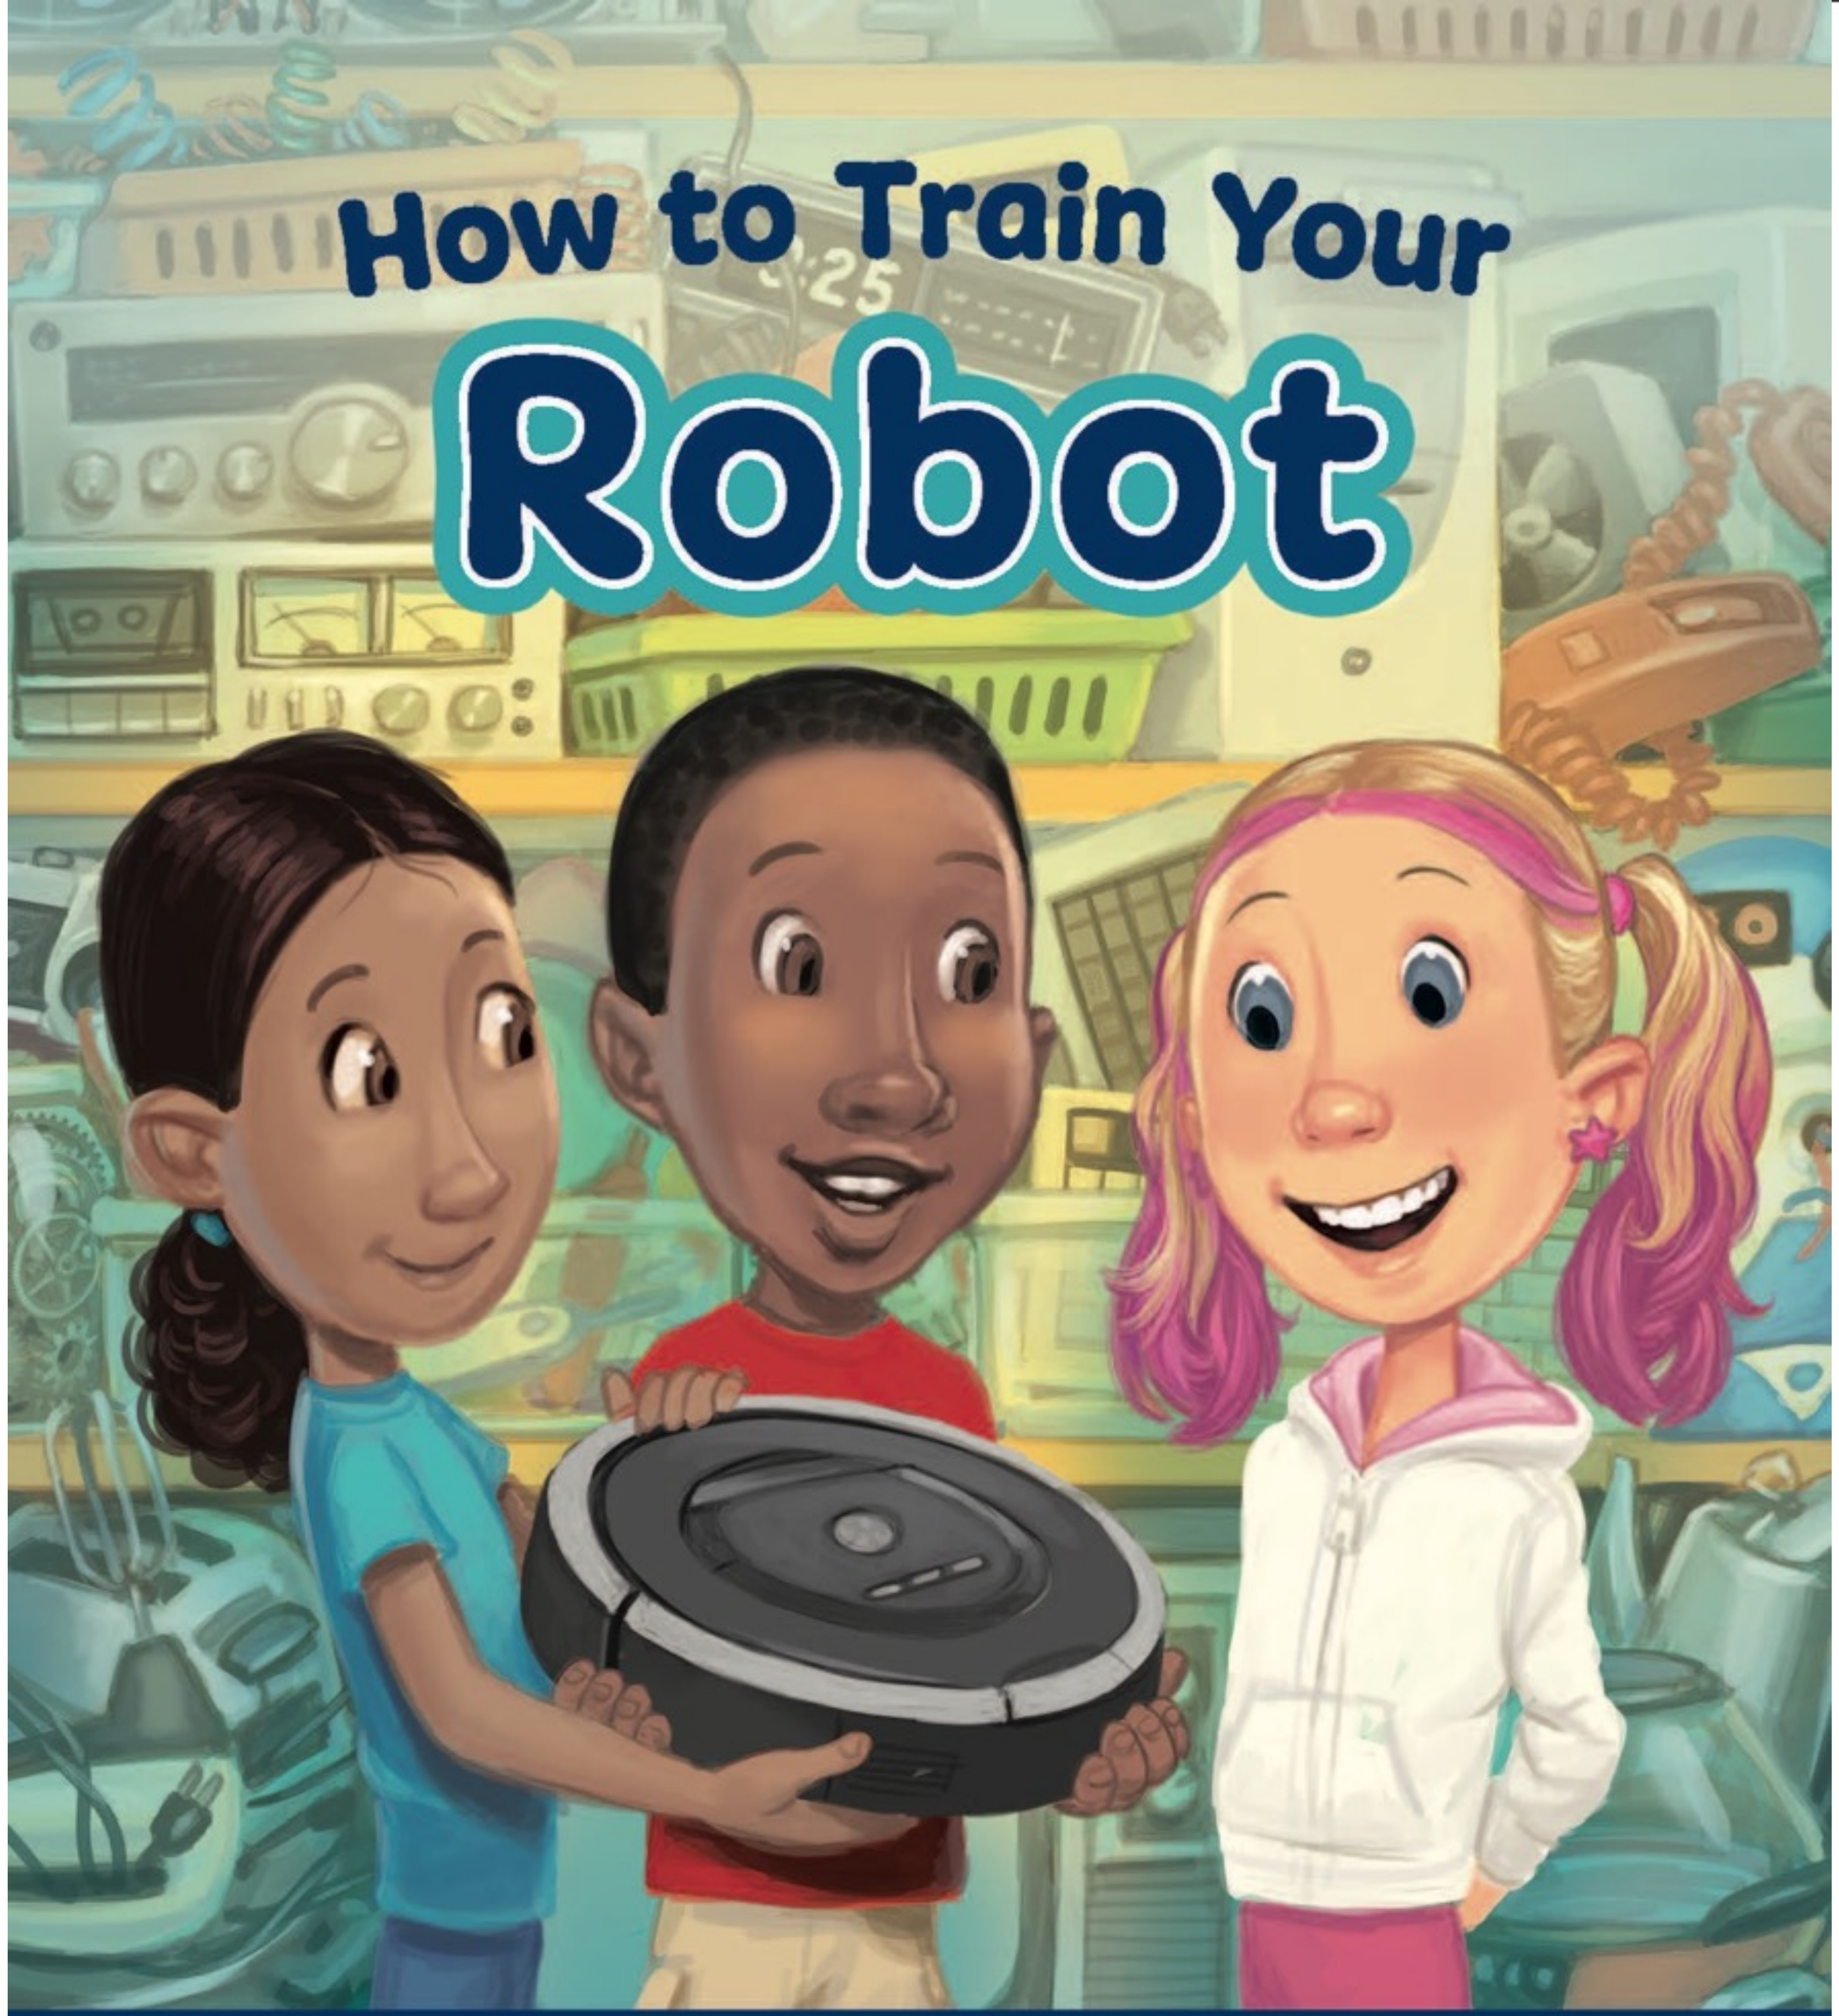

by Blooma Goldberg, Ken Goldberg, and Ashley Chase
illustrated by Dave Clegg

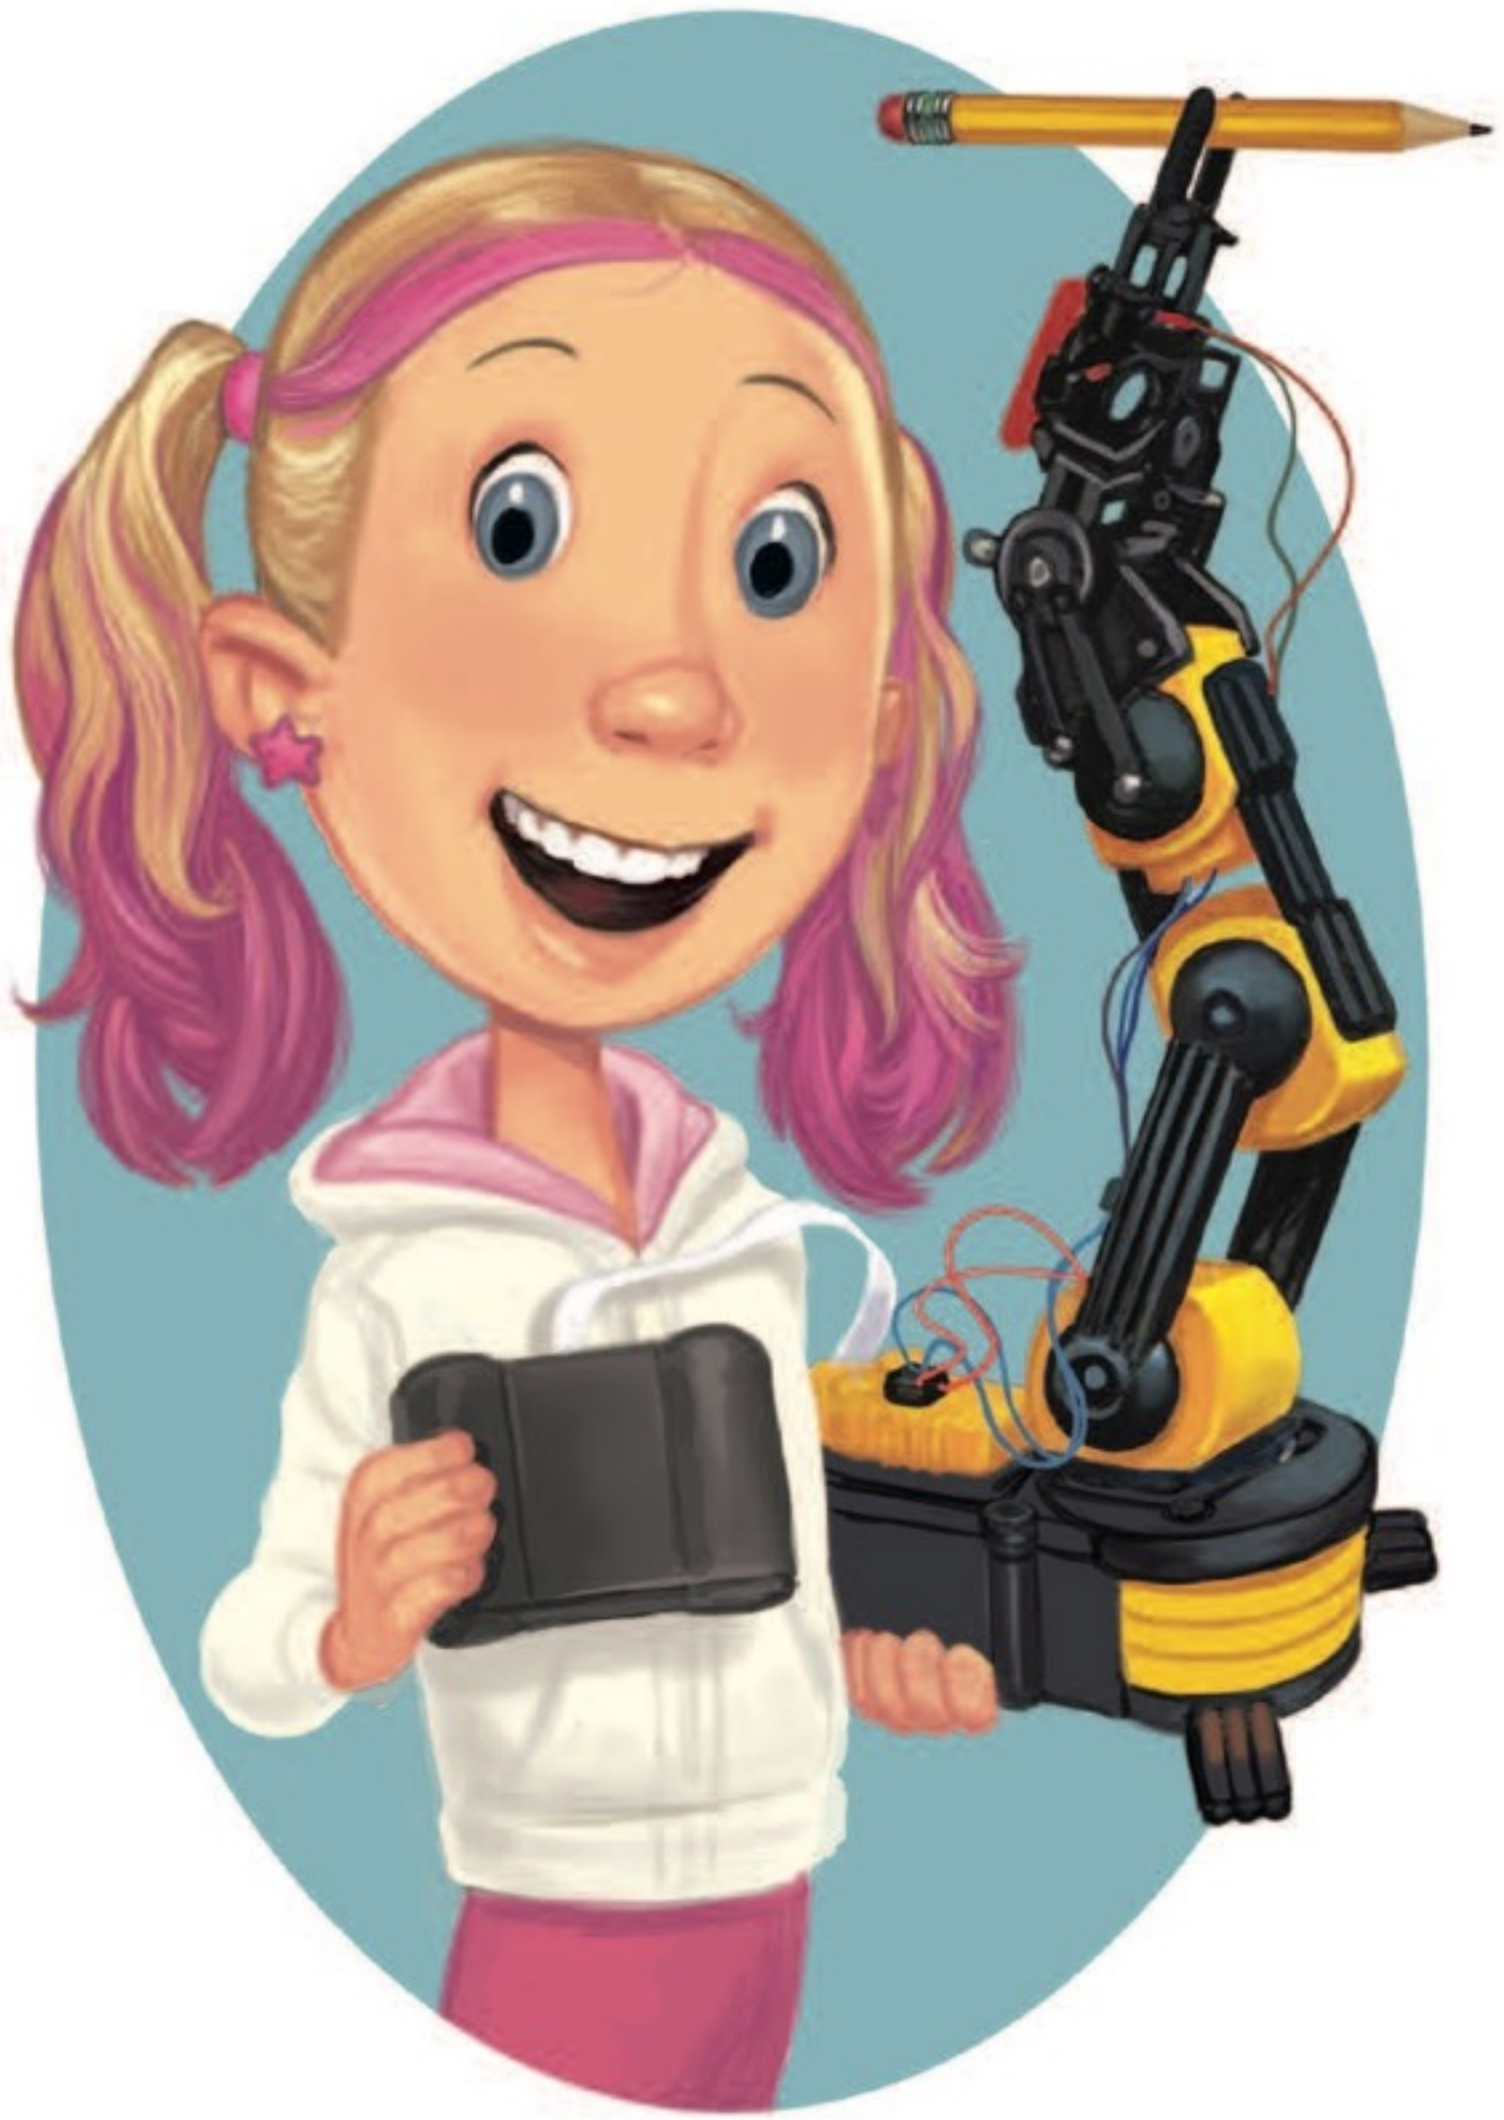

I'm Blooma. I love art, science, math, dancing, basketball, and roller skating, but my favorite activity is inventing new things... especially robots! And I'm super excited to tell you how I figured out the best way to train a robot.

How to Train Your Robot

How to Train Your Robot

video based on the book: How To Train Your Robot
by B. Goldberg, K. Goldberg, and A. Chase, illustrated by D. Clegg
(with support from NSF and UC Berkeley's Lawrence Hall of Science)

Info and Resource Page: (Google Doc)
https://bit.ly/How-To-Train-Your-Robot-Info

To view subtitled Spanish, Japanese, Hindi, and Chinese (Simplified), just click the small blue CC (closed caption) button on the right.

15 min Video
Nov 2020
Subtitled in Spanish,
Japanese, Hindi, and simplified Chinese

https://vimeopro.com/citrisproductions/how-to-train-your-robot