Perspectives on Dexterous Manipulation

16-848 January 24, 2022

Nancy Pollard

Robotics: Science and Systems 2021
Held Virtually, July 12–16, 2021

http://www.roboticsproceedings.org/rss17/p089.pdf

Surprisingly Robust In-Hand Manipulation: An Empirical Study

Aditya Bhatt* Adrian Sieler* Steffen Puhlmann Oliver Brock

Abstract—We present in-hand manipulation skills on a dexterous, compliant, anthropomorphic hand. Even though these skills were derived in a simplistic manner, they exhibit surprising robustness to variations in shape, size, weight, and placement of the manipulated object. They are also very insensitive to variation of execution speeds, ranging from highly dynamic to quasi-static. The robustness of the skills leads to compositional properties that enable extended and robust manipulation programs. To explain the surprising robustness of the in-hand manipulation skills, we performed a detailed, empirical analysis of the skills' performance. From this analysis, we identify three principles for skill design: 1) Exploiting the hardware's innate ability to drive hard-to-model contact dynamics. 2) Taking actions to constrain these interactions, funneling the system into a narrow set of possibilities. 3) Composing such action sequences into complex manipulation programs. We believe that these principles constitute an important foundation for robust robotic in-hand manipulation, and possibly for manipulation in general.

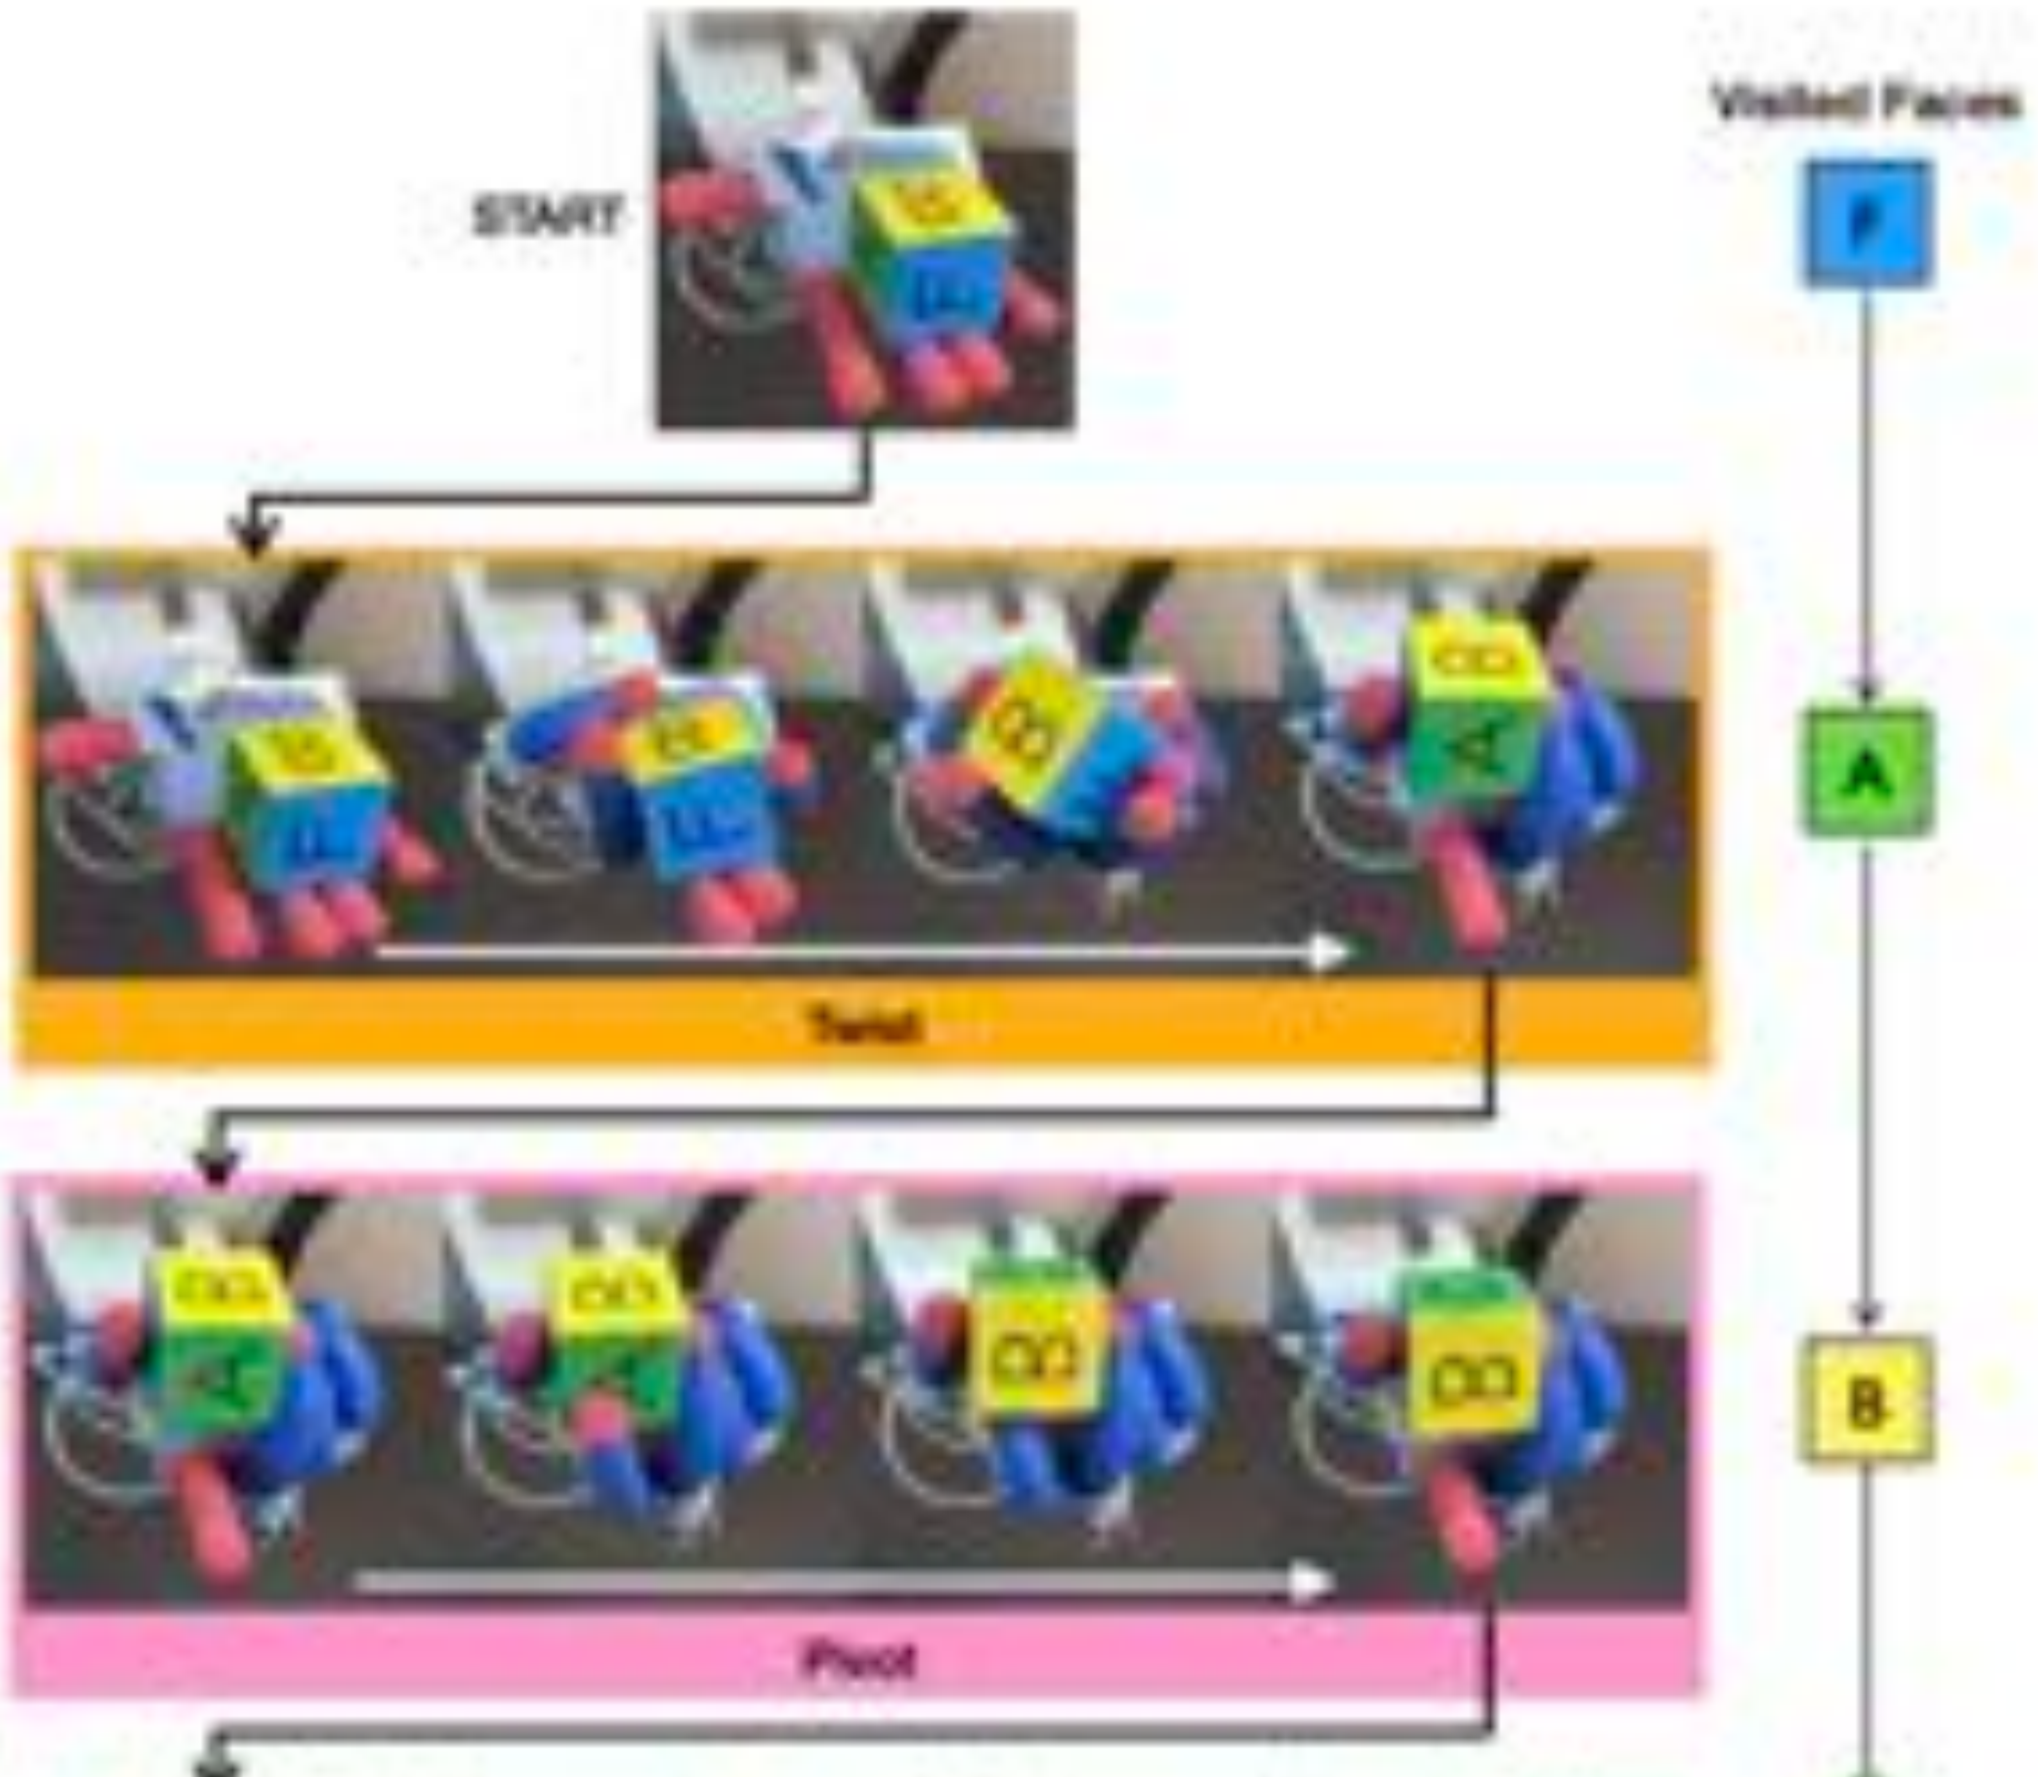


RBO
Surprisingly
Robust In-Hand
Manipulation

Oliver Brock

Robotics and Biology Laboratory, TU Berlin

Soft Hands

Soft Hands represent a departure from classical robot hand design, which often relies on exact models and precise planning of contact points. Instead, we aim to increase robustness and safety through the use of soft materials and flexible mechanics. This softness allows us to *exploit* contact with the environment and use it in robust grasping and manipulation strategies.

In our lab we develop the RBO Hand 2, research necessary Soft Robotic aspects, and formulate the concept of Morphological Computation.

RBO Hand 2

The RBO Hand 2 is a hand made from PneuFlex actuators mounted on a flexible, printed scaffold. The hand was developed to investigate the capabilities and limits of hands when relying only on soft, deformable structures. The unique deformability provides several advantageous benefits to robots trying to interact with the environment:

- very robust against blunt collisions
- very low impact energies
- passively compliant fingers and palm decouple contact from the robot arm, stabilizing force control
- mechanical adaptability to object shapes simplifies finger control
- the pneumatic actuation makes it easy to create complex hand and actuator geometries

The RBO Hand 2 ho

© RBO

The result of our research are several hand prototypes, which we refer to collectively as *Soft Hands*. RBO Hand 2 is the latest model and used in our lab for research into grasping strategies.

The RBO Hand 2 is controlled using a PneumaticBox and is relatively cheap to produce, modify and repair.

If you want to build your own, you are welcome to do so! We have published the CAD models for the PneuFlex actuators.

Contact: Raphael Deimel, Vincent Wall

Versions of the RBO Hand

Over time we have created quite a few different versions of the RBO Hand. Here are a few:

Hand prototypes for the SOMA project

© RBO

© Disney Research Zürich

© RBO

© RBO

As part of the SOMA project we develop versions of the RBO Hand, based on feedback from all partners in the consortium. These hand versions change the geometry of the fingers, palms, and wrist. Because the RBO Hand is assembled from modular parts, we can quickly switch out parts and try different ideas.

Aditya Bhatt

Adrian Sieler

Steffen Puhlmann

https://www.robotics.tu-berlin.de/menue/research/soft_hands

RBO Hand 1 https://www.youtube.com/watch?v=ziY-pHSpH5Q

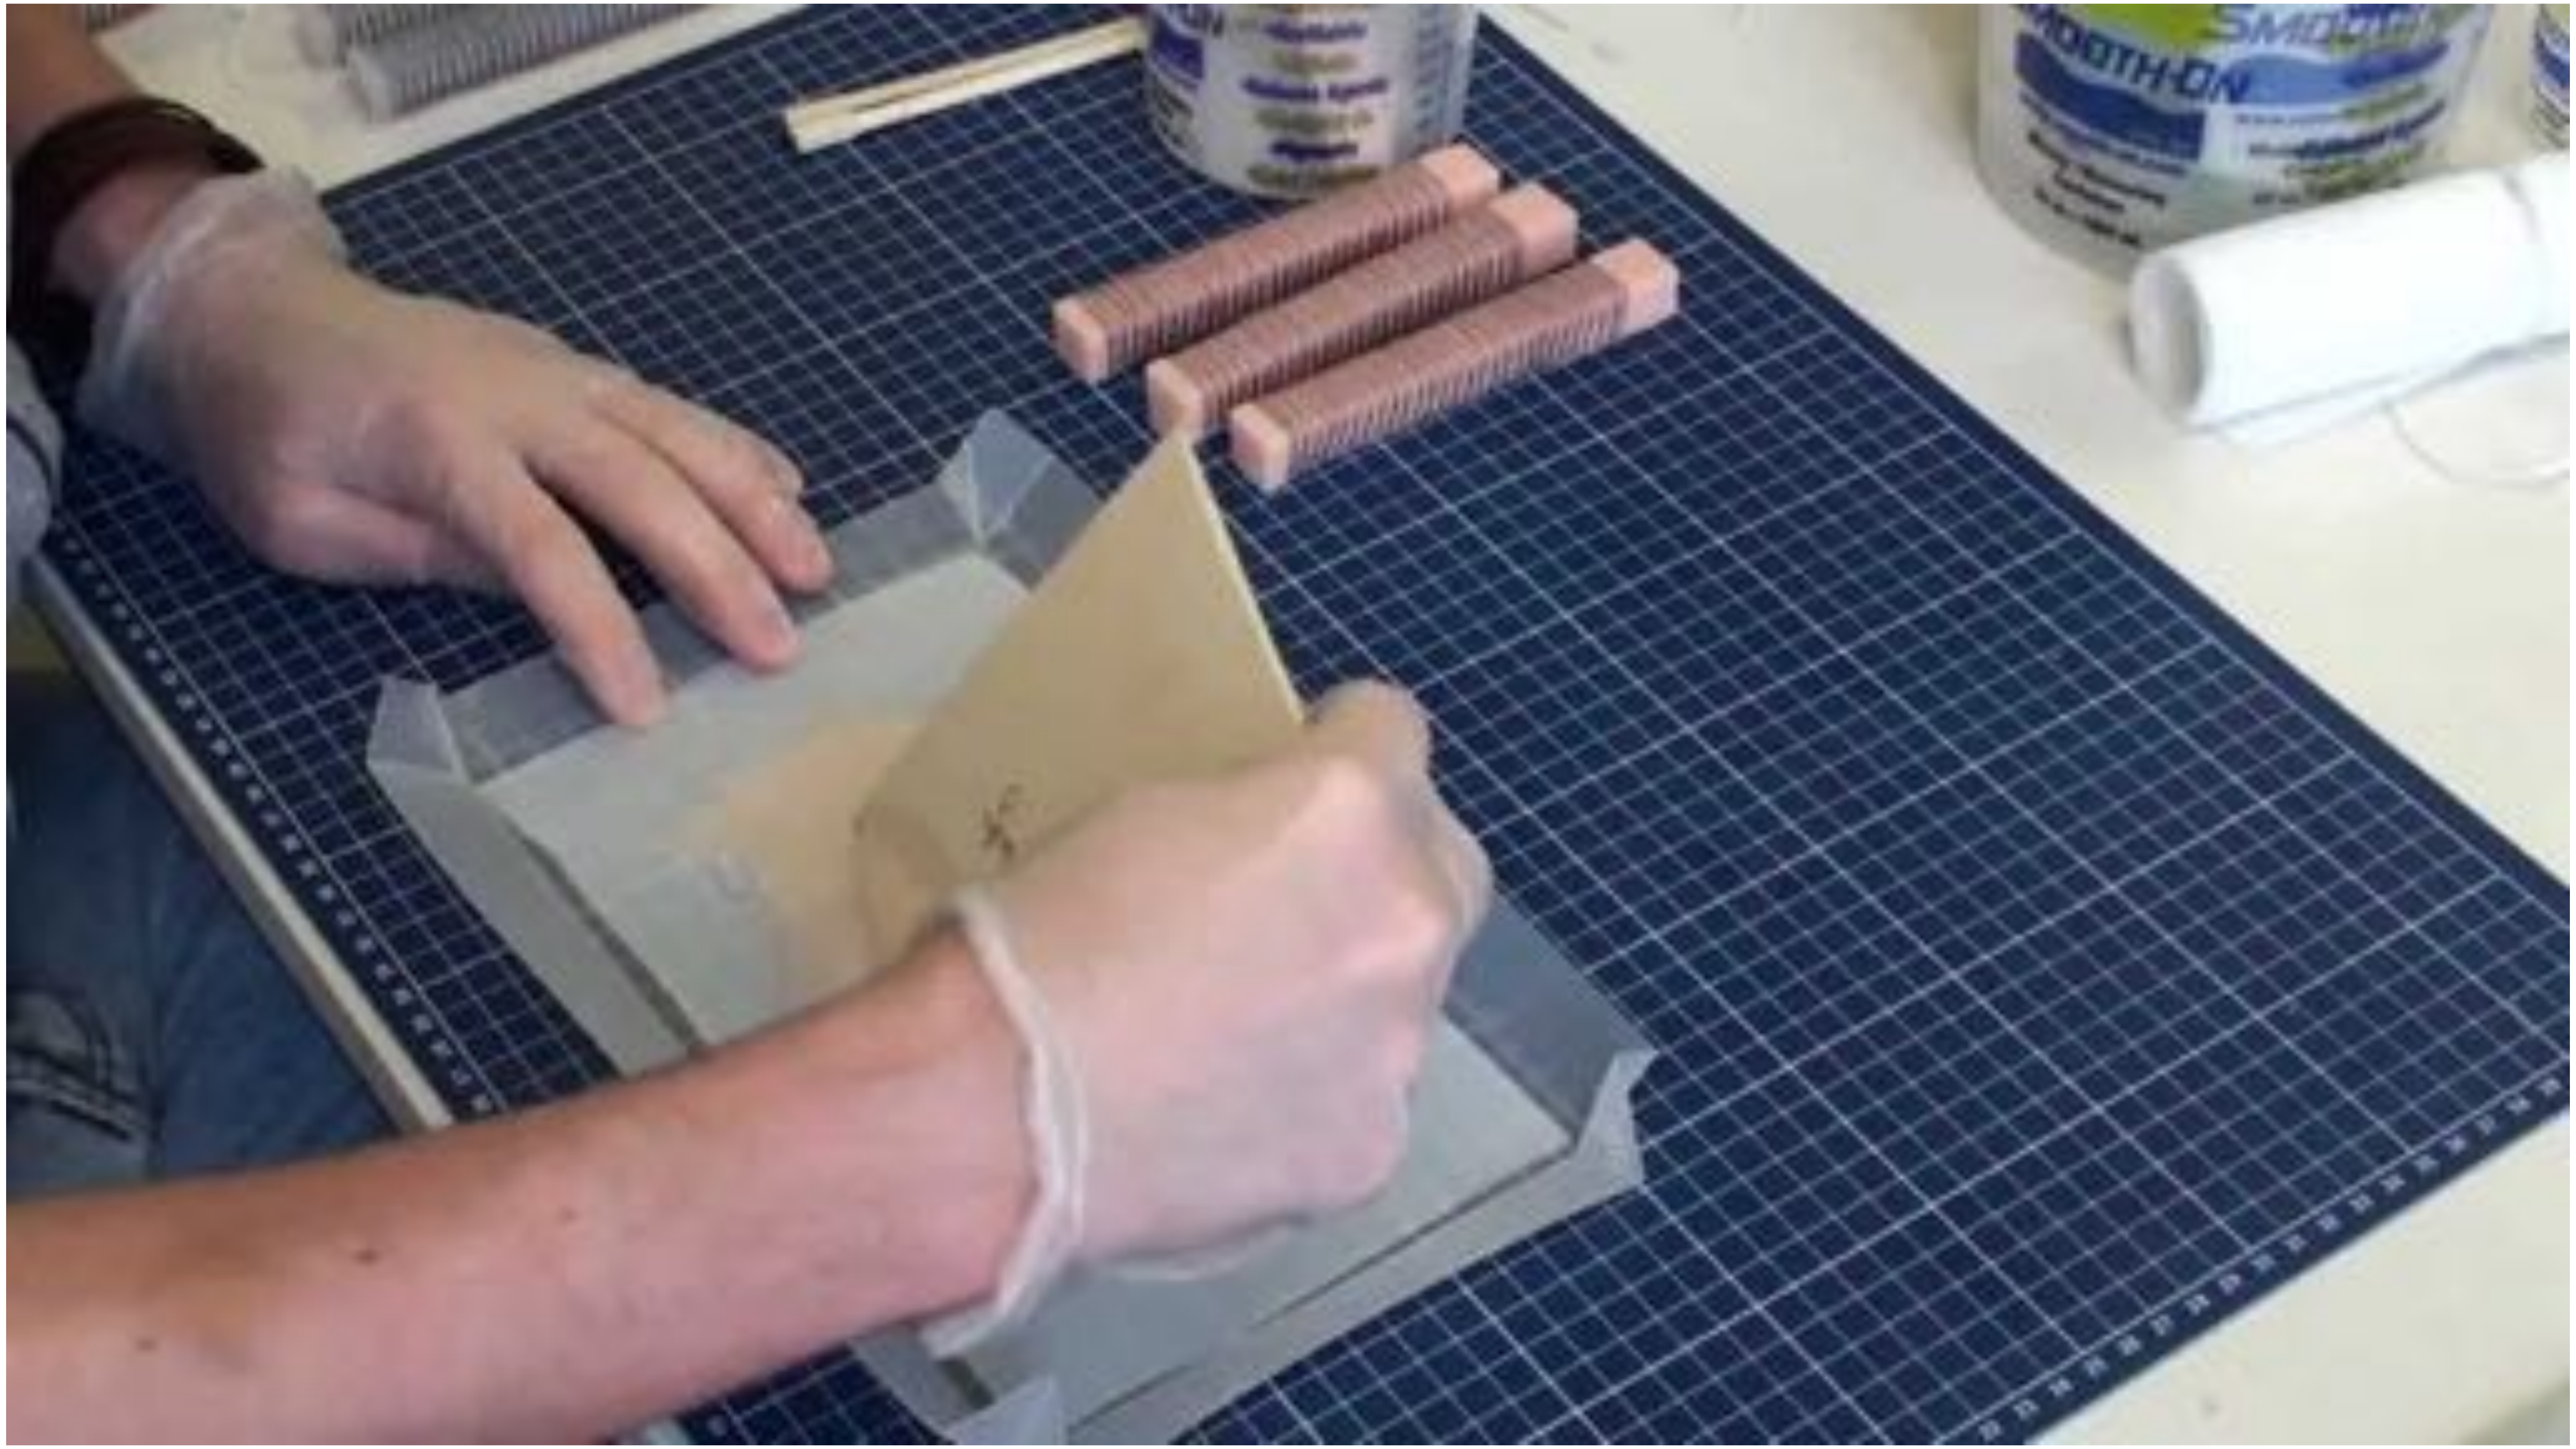

How they're made: https://www.youtube.com/watch?v=Ss-9iXRUeGc&t=814s

RBO Hand 3

- pneumatic
- 16 actuated degrees of freedom
- red and white are soft silicone paddings
- controlled by hand with mixer

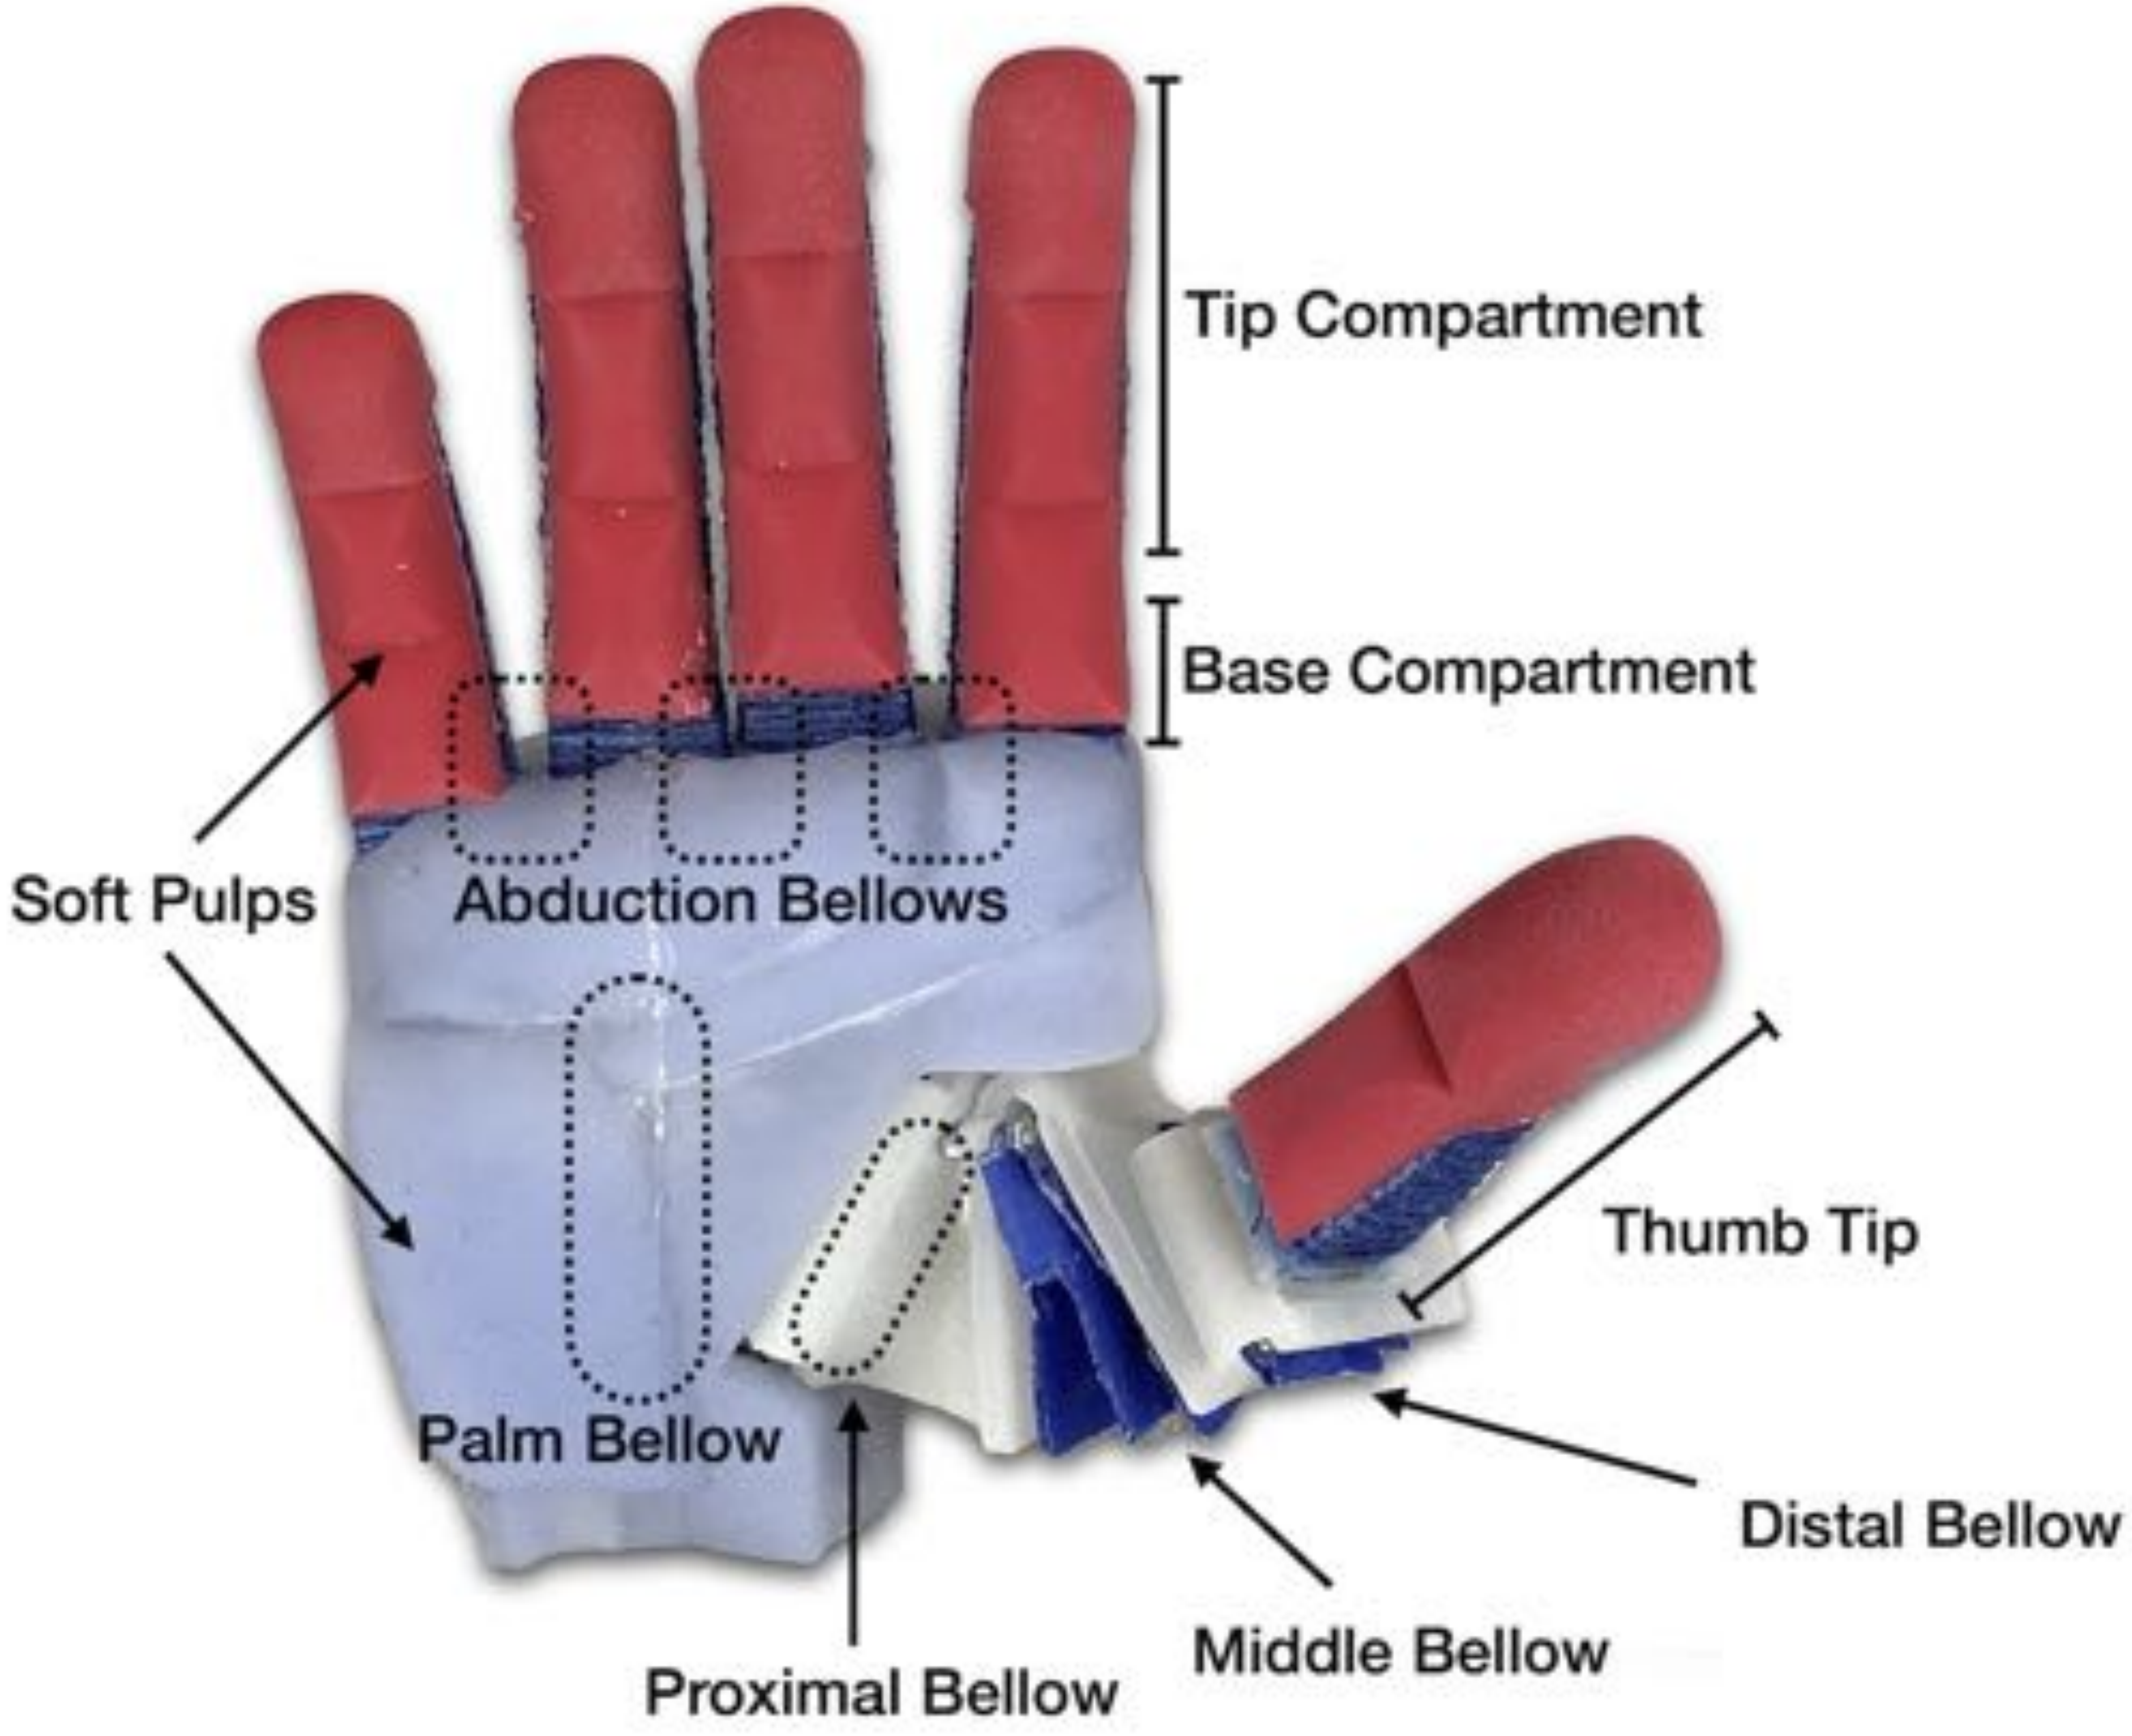

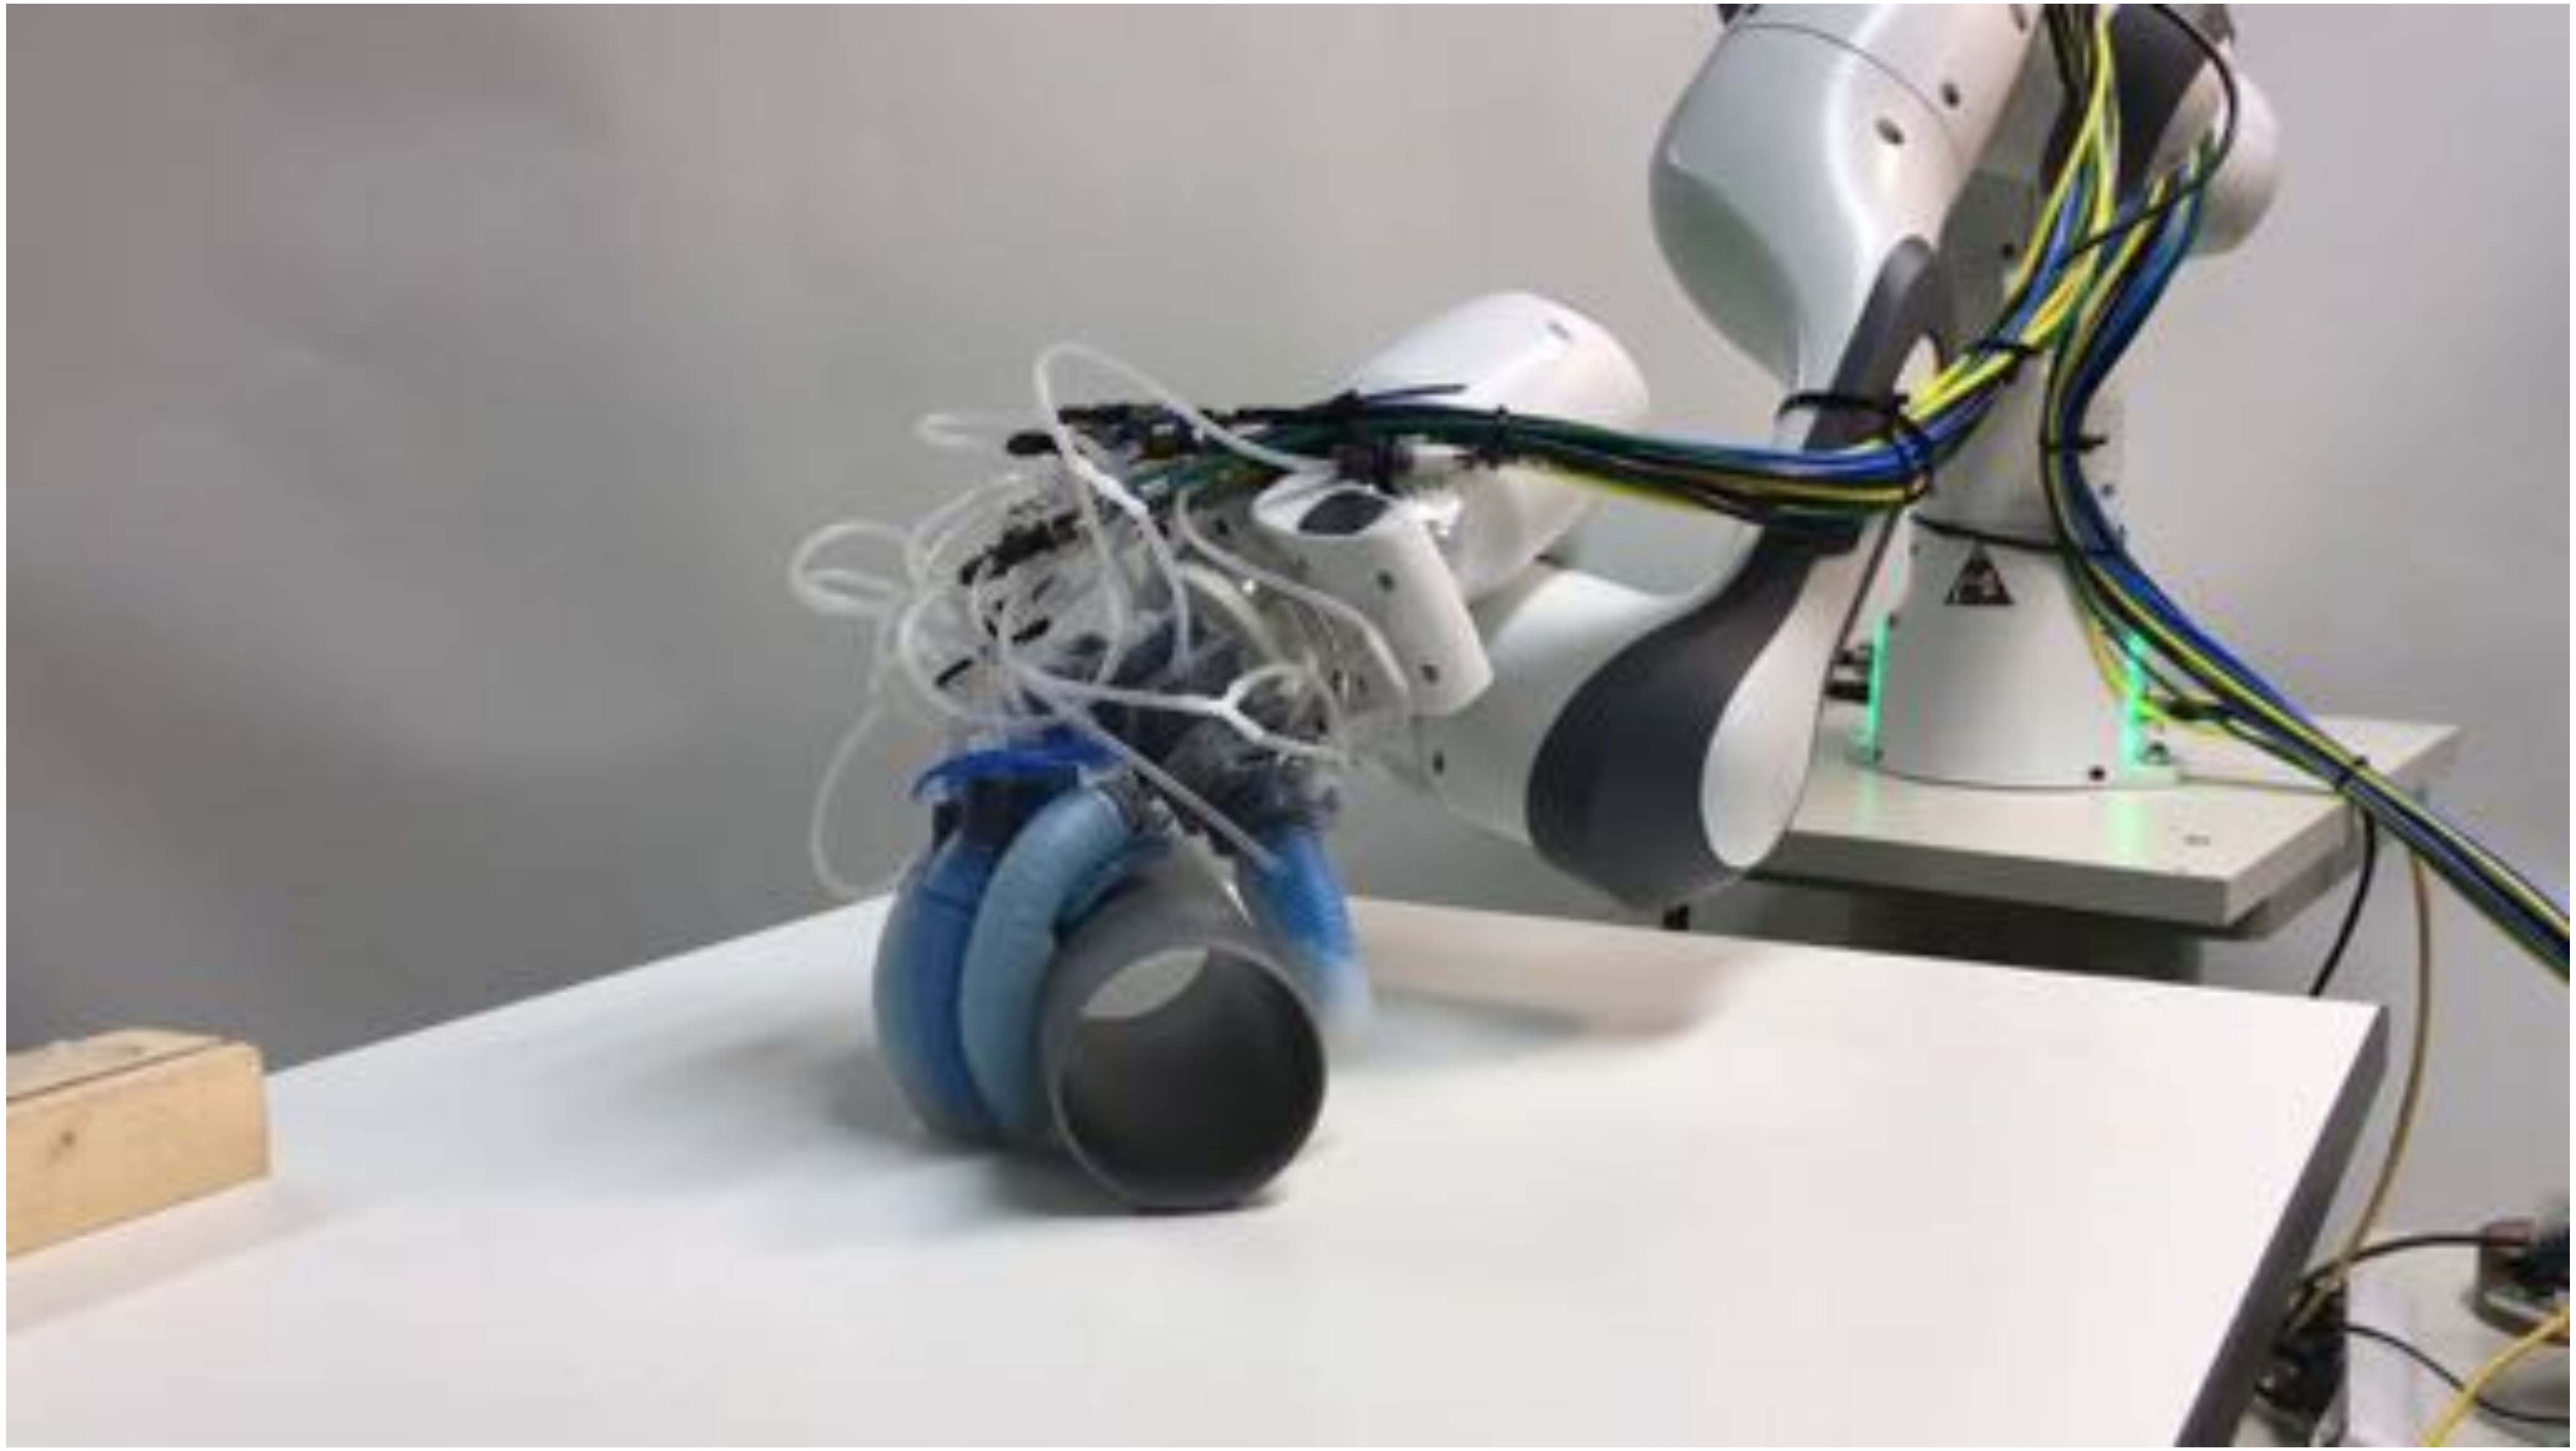

RBO Hand 3 (earlier version) https://www.youtube.com/watch?v=ENbrUOmDsSI

Context

“The state of the art, at this time, is the groundbreaking work presented by OpenAI [1], who used Deep Reinforcement Learning to produce remarkably dexterous behavior on a five- fingered robotic hand, first manipulating a cube and later even an articulated Rubik’s cube [2]. Their learned skills feature contact-rich movements like finger-gaiting, pivoting, and the exploitation of gravity.”

https://journals.sagepub.com/doi/pdf/10.1177/0278364919887447

Article

The International Journal of
Robotics Research
2020, Vol. 39(1) 3–20

DOI: 10.1177/0278364919887447
journals.sagepub.com/home/ijr
SAGE

Learning dexterous in-hand manipulation

OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Józefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng and Wojciech Zaremba

Abstract

We use reinforcement learning (RL) to learn dexterous in-hand manipulation policies that can perform vision-based object reorientation on a physical Shadow Dexterous Hand. The training is performed in a simulated environment in which we randomize many of the physical properties of the system such as friction coefficients and an object's appearance. Our policies transfer to the physical robot despite being trained entirely in simulation. Our method does not rely on any human demonstrations, but many behaviors found in human manipulation emerge naturally, including finger gaiting, multi-finger coordination, and the controlled use of gravity. Our results were obtained using the same distributed RL system that was used to train OpenAI Five. We also include a video of our results: https://youtu.be/jwSbzNHGflM.

OpenAI conducts fundamental, long-term research toward the creation of safe AGI.

- Founded in 2015, by Elon Musk, Sam Altman, …
- Performs research on a wide variety of AI problems
- Known in our community for OpenAI gym
    - Reinforcement learning environments
    - https://gym.openai.com/

Milestones

Multimodal Neurons in Artificial Neural Networks
March 4, 2021

DALL·E: Creating Images from Text
January 5, 2021

https://openai.com/research/

OpenAI Manipulation Demo: https://www.youtube.com/watch?v=jwSbzNHGflM

Shadow Dexterous Hand

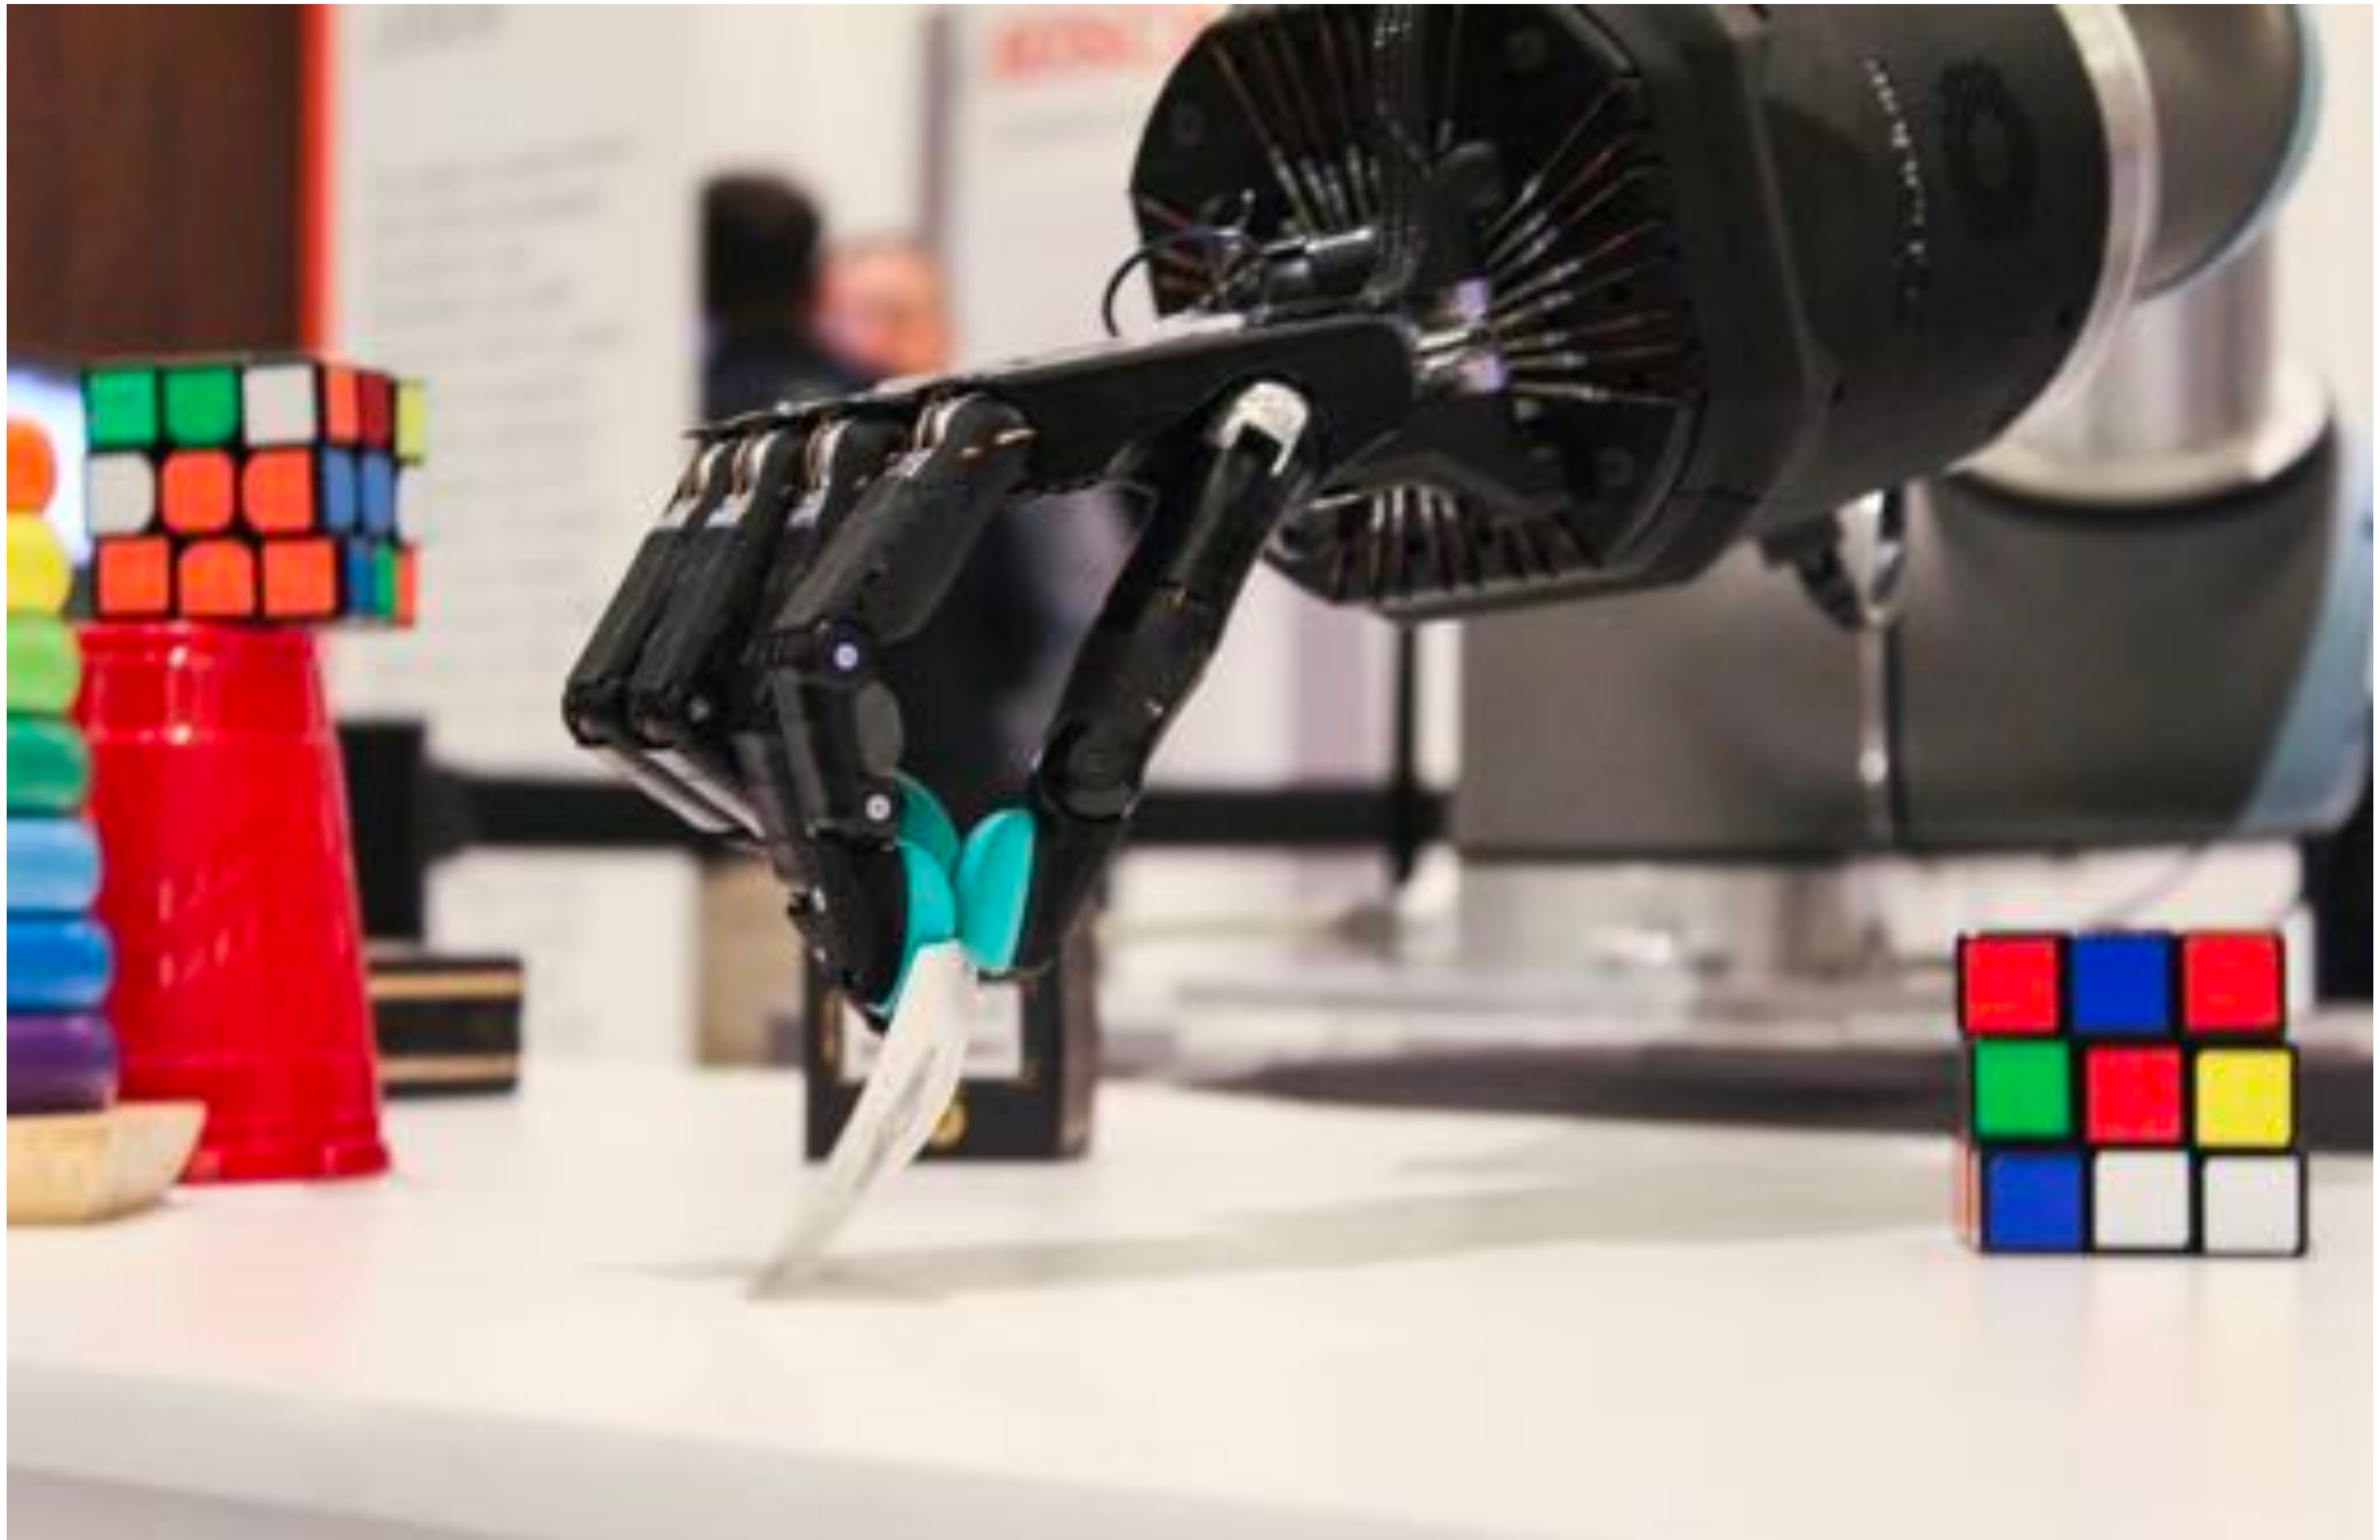

https://www.shadowrobot.com/wp-content/uploads/shadow_dexterous_hand_e_technical_specification.pdf

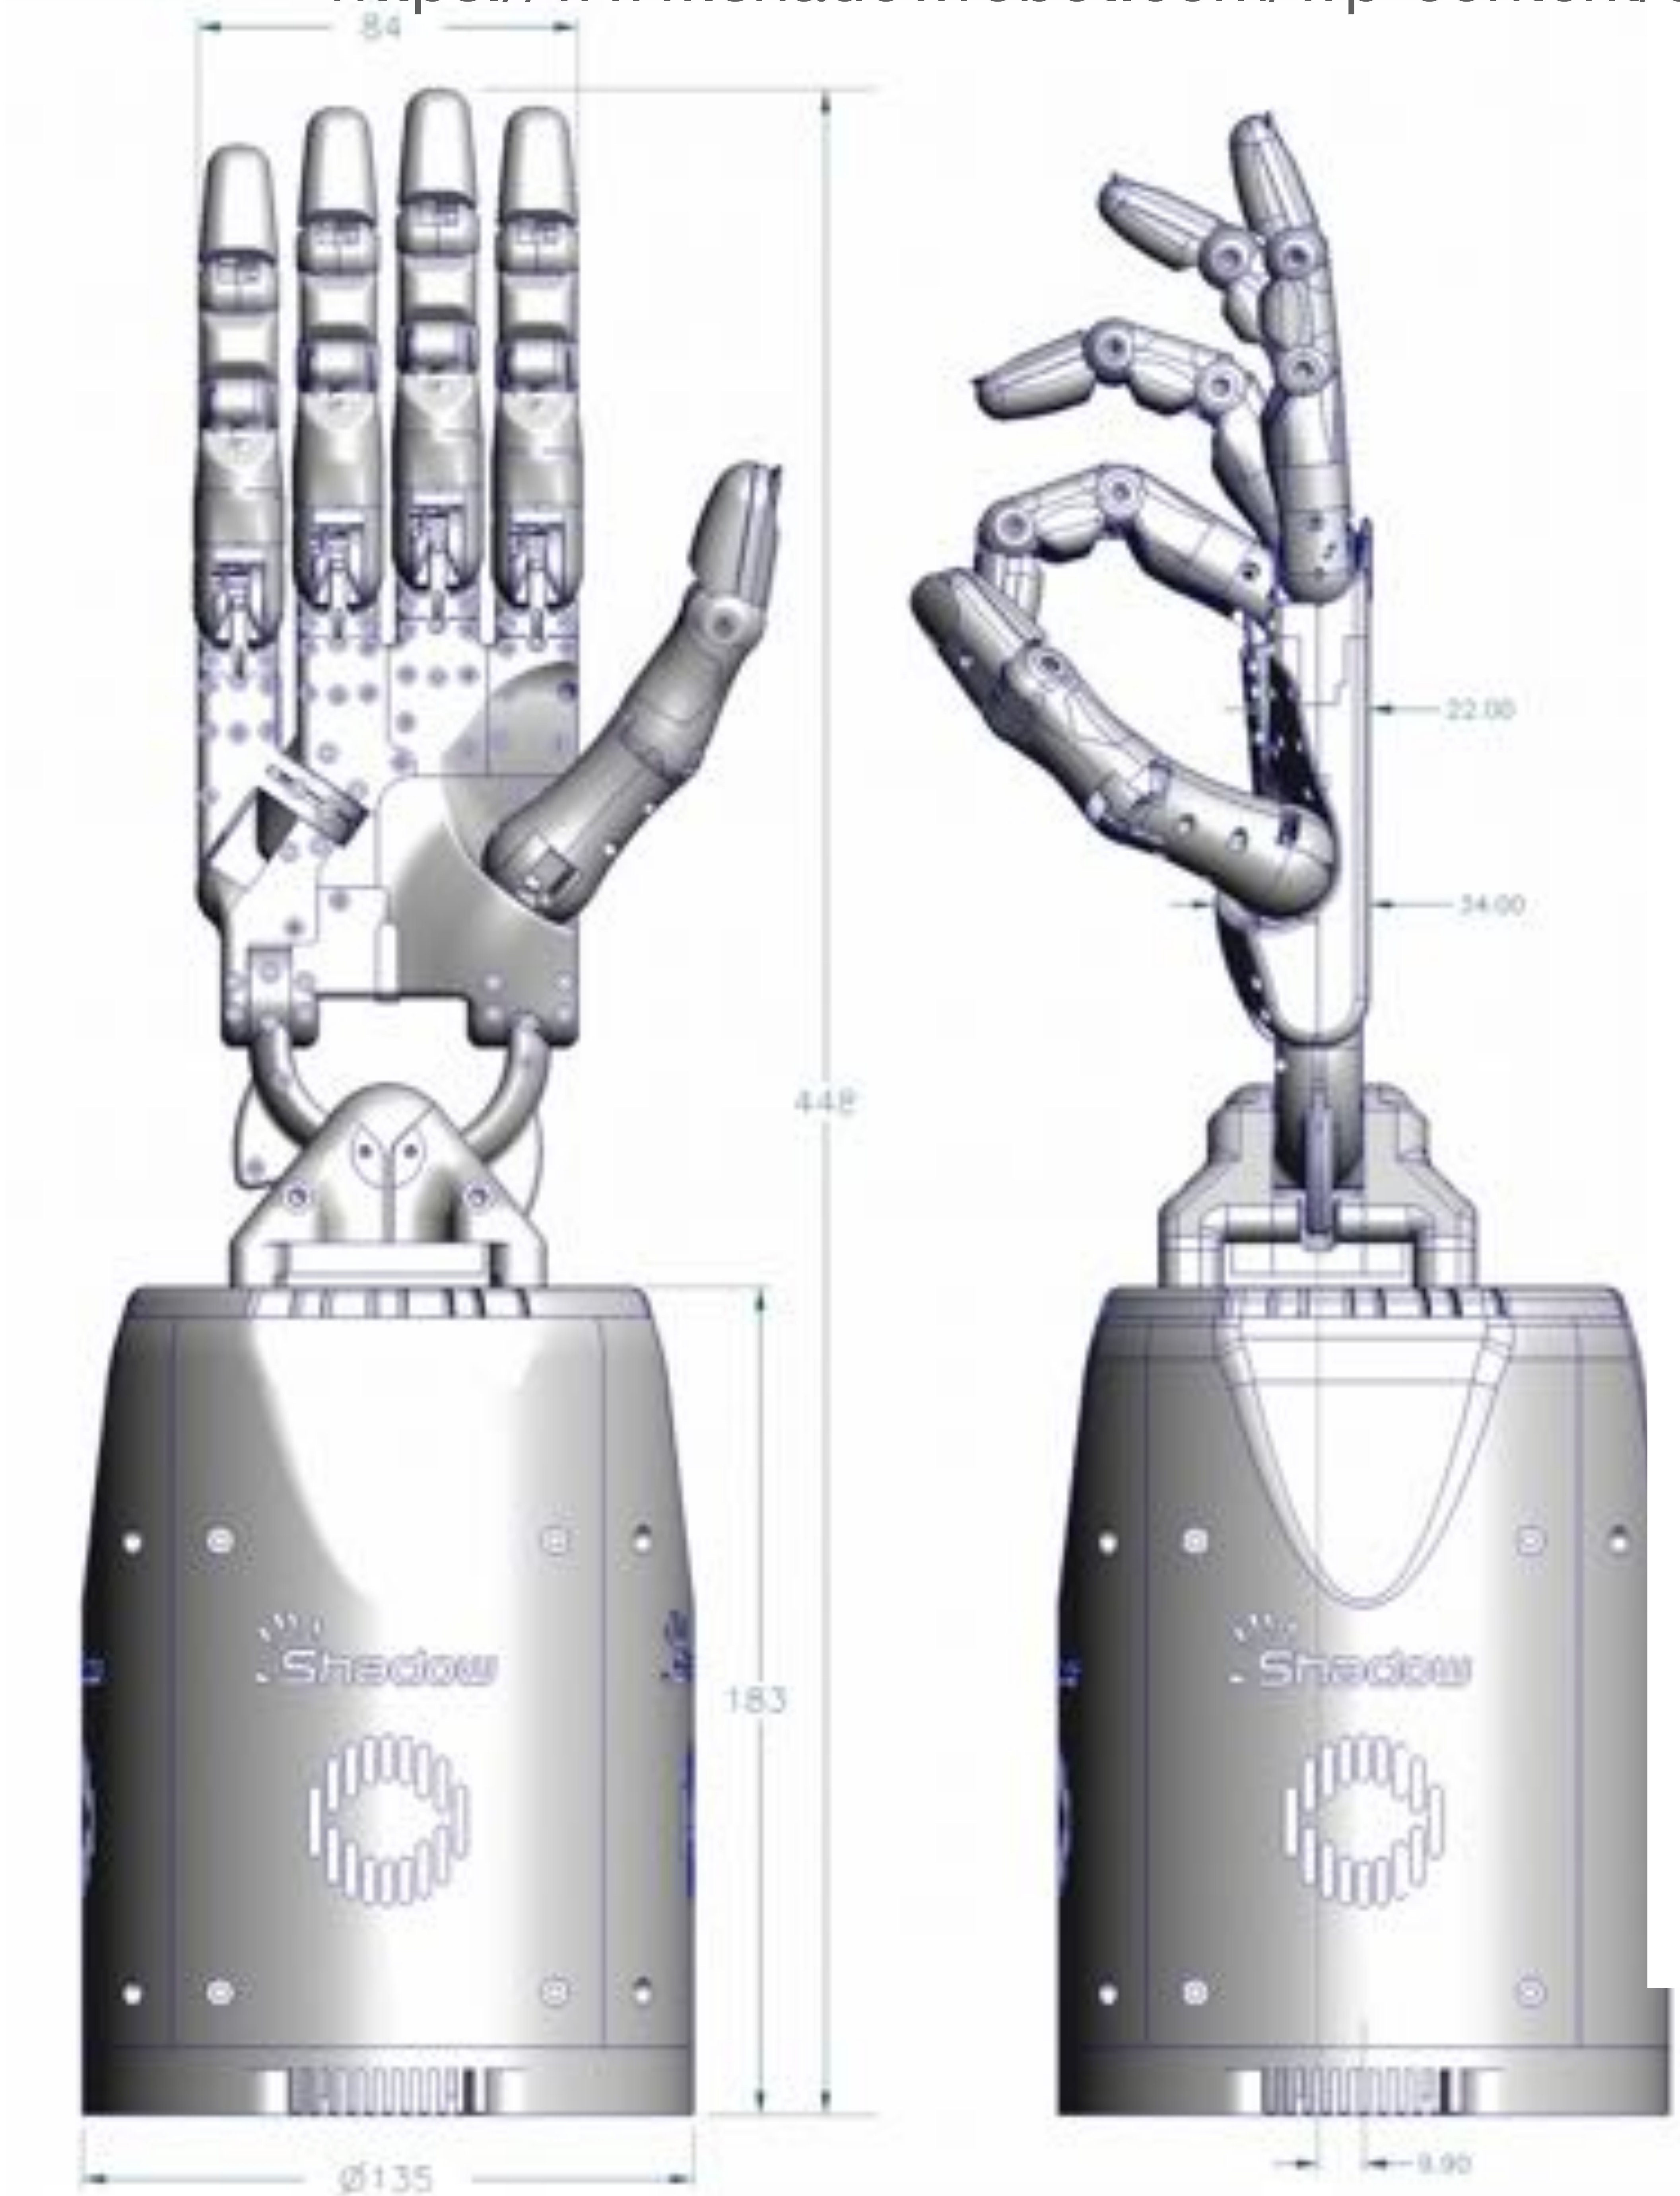

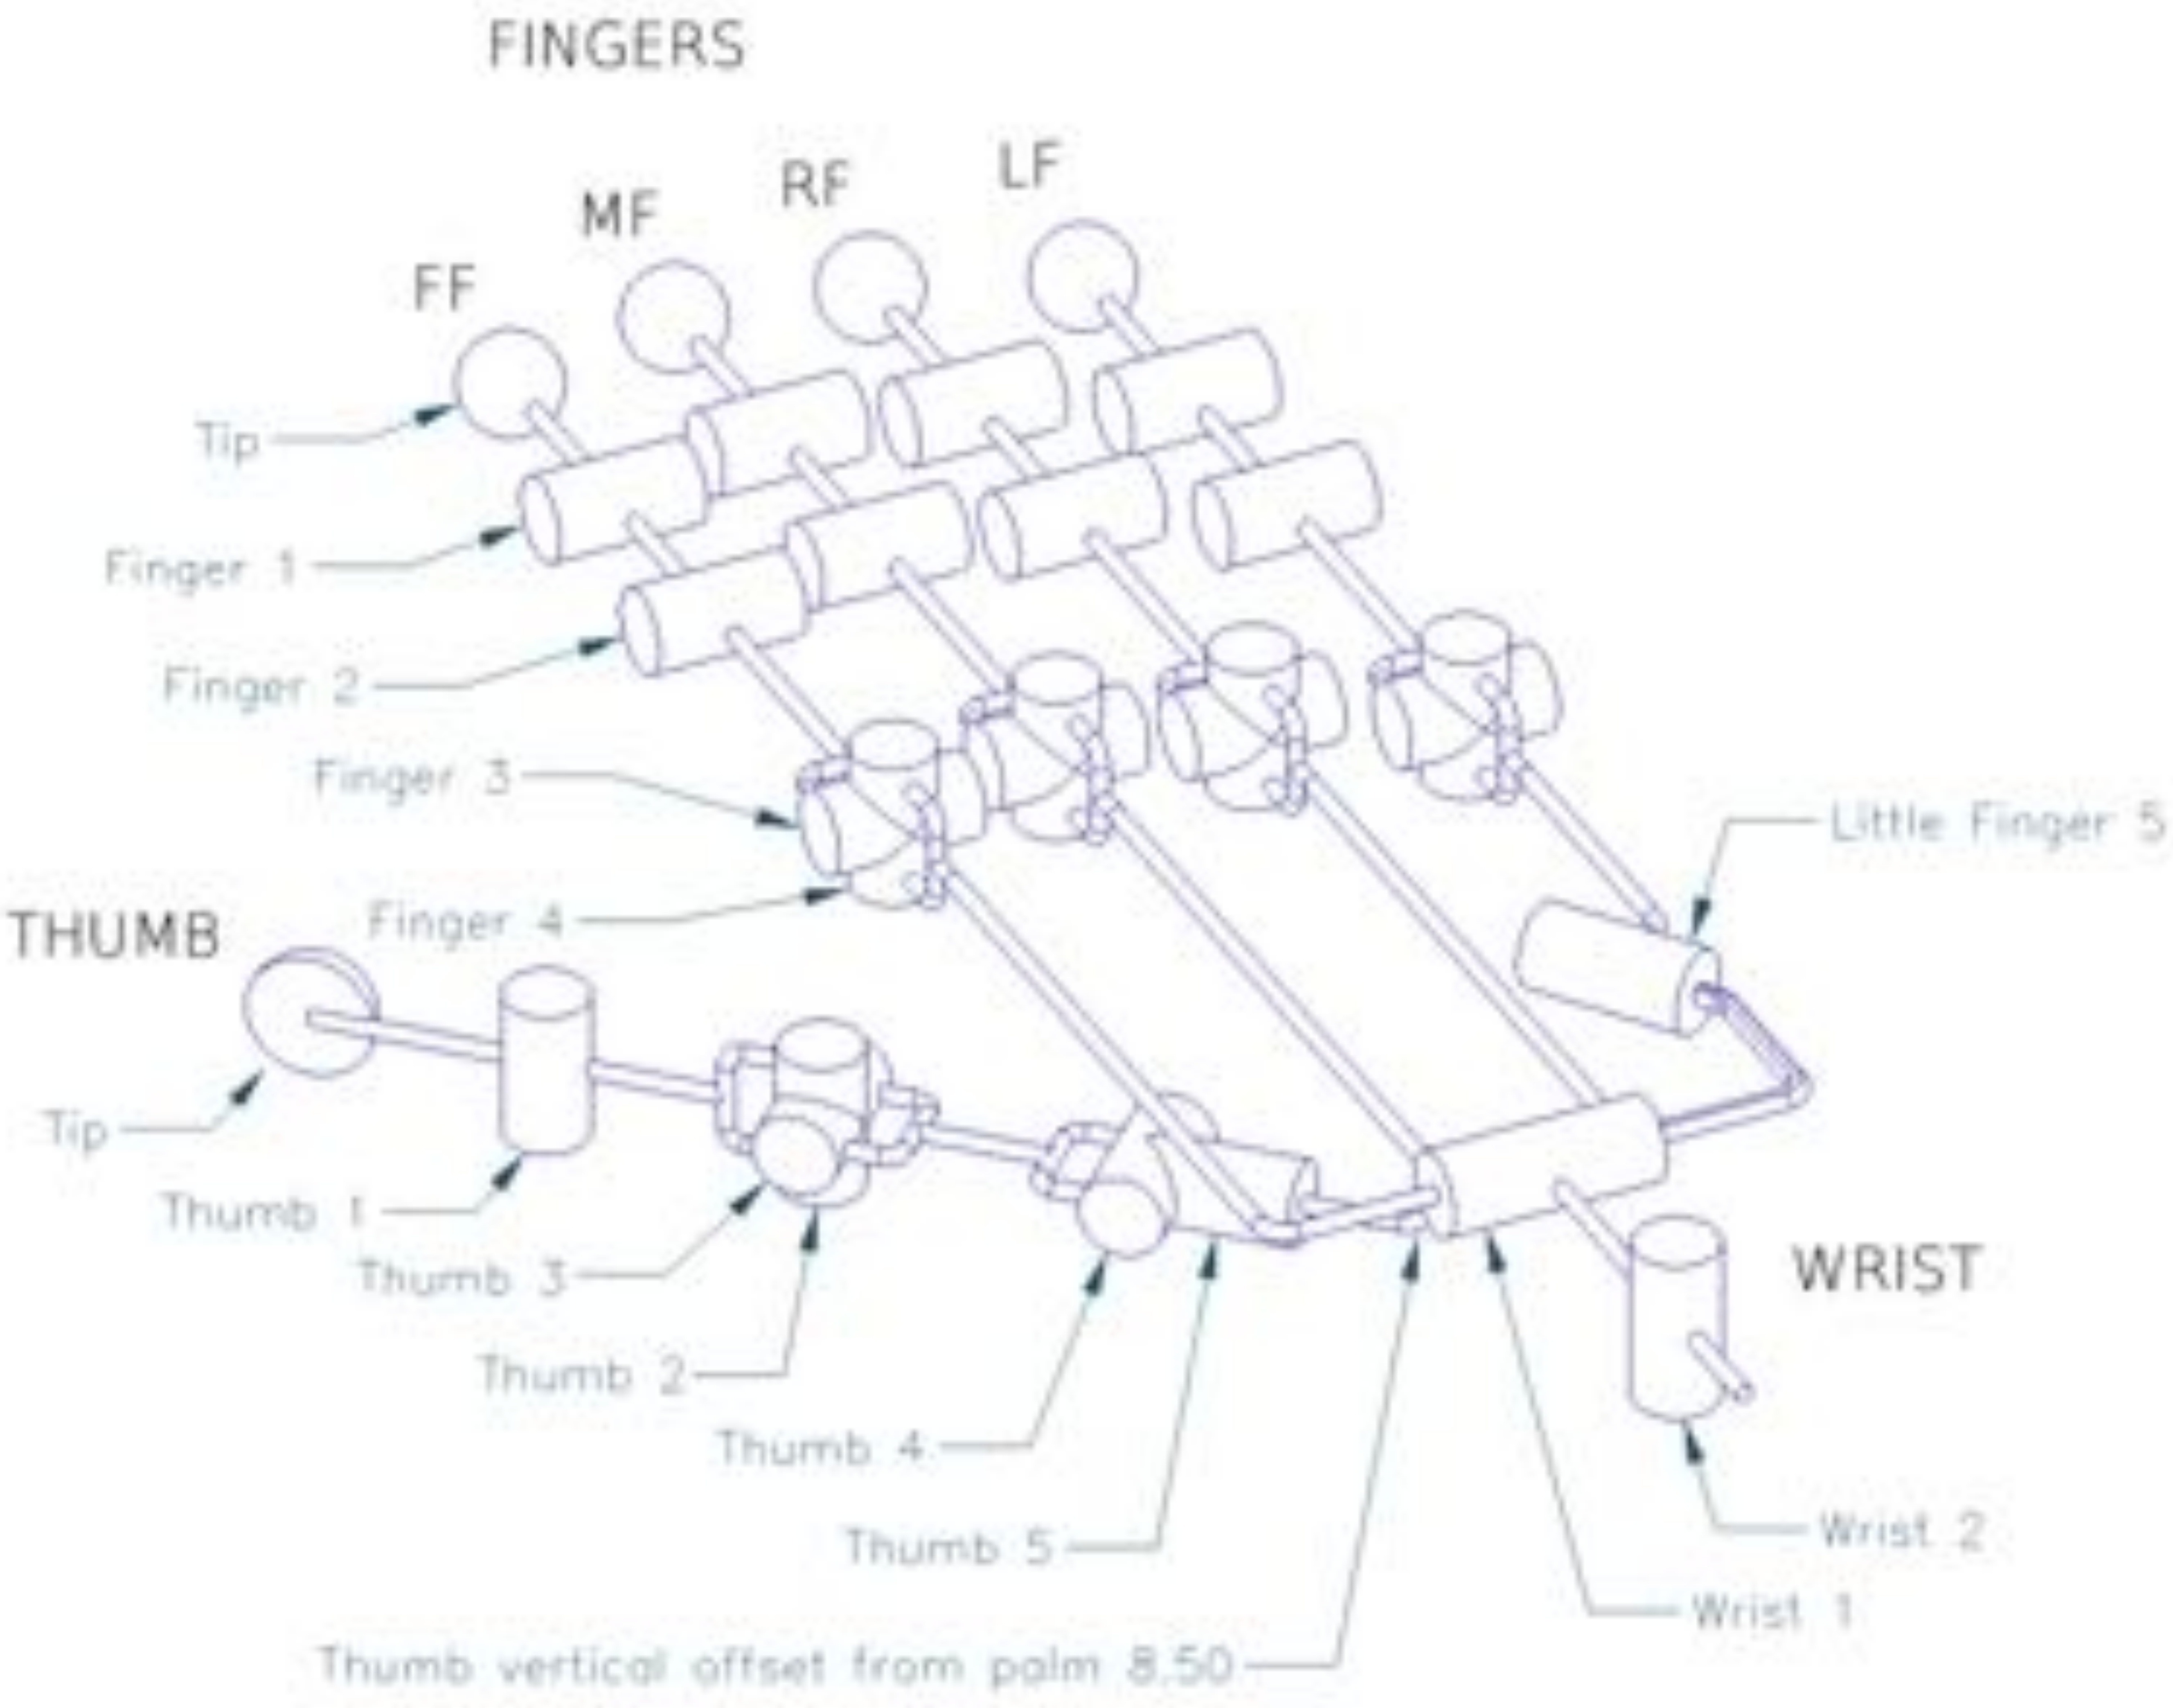

- 24 Degrees of freedom driven by 20 actuators
- Electric, tendon driven
- First commercialized in 2005

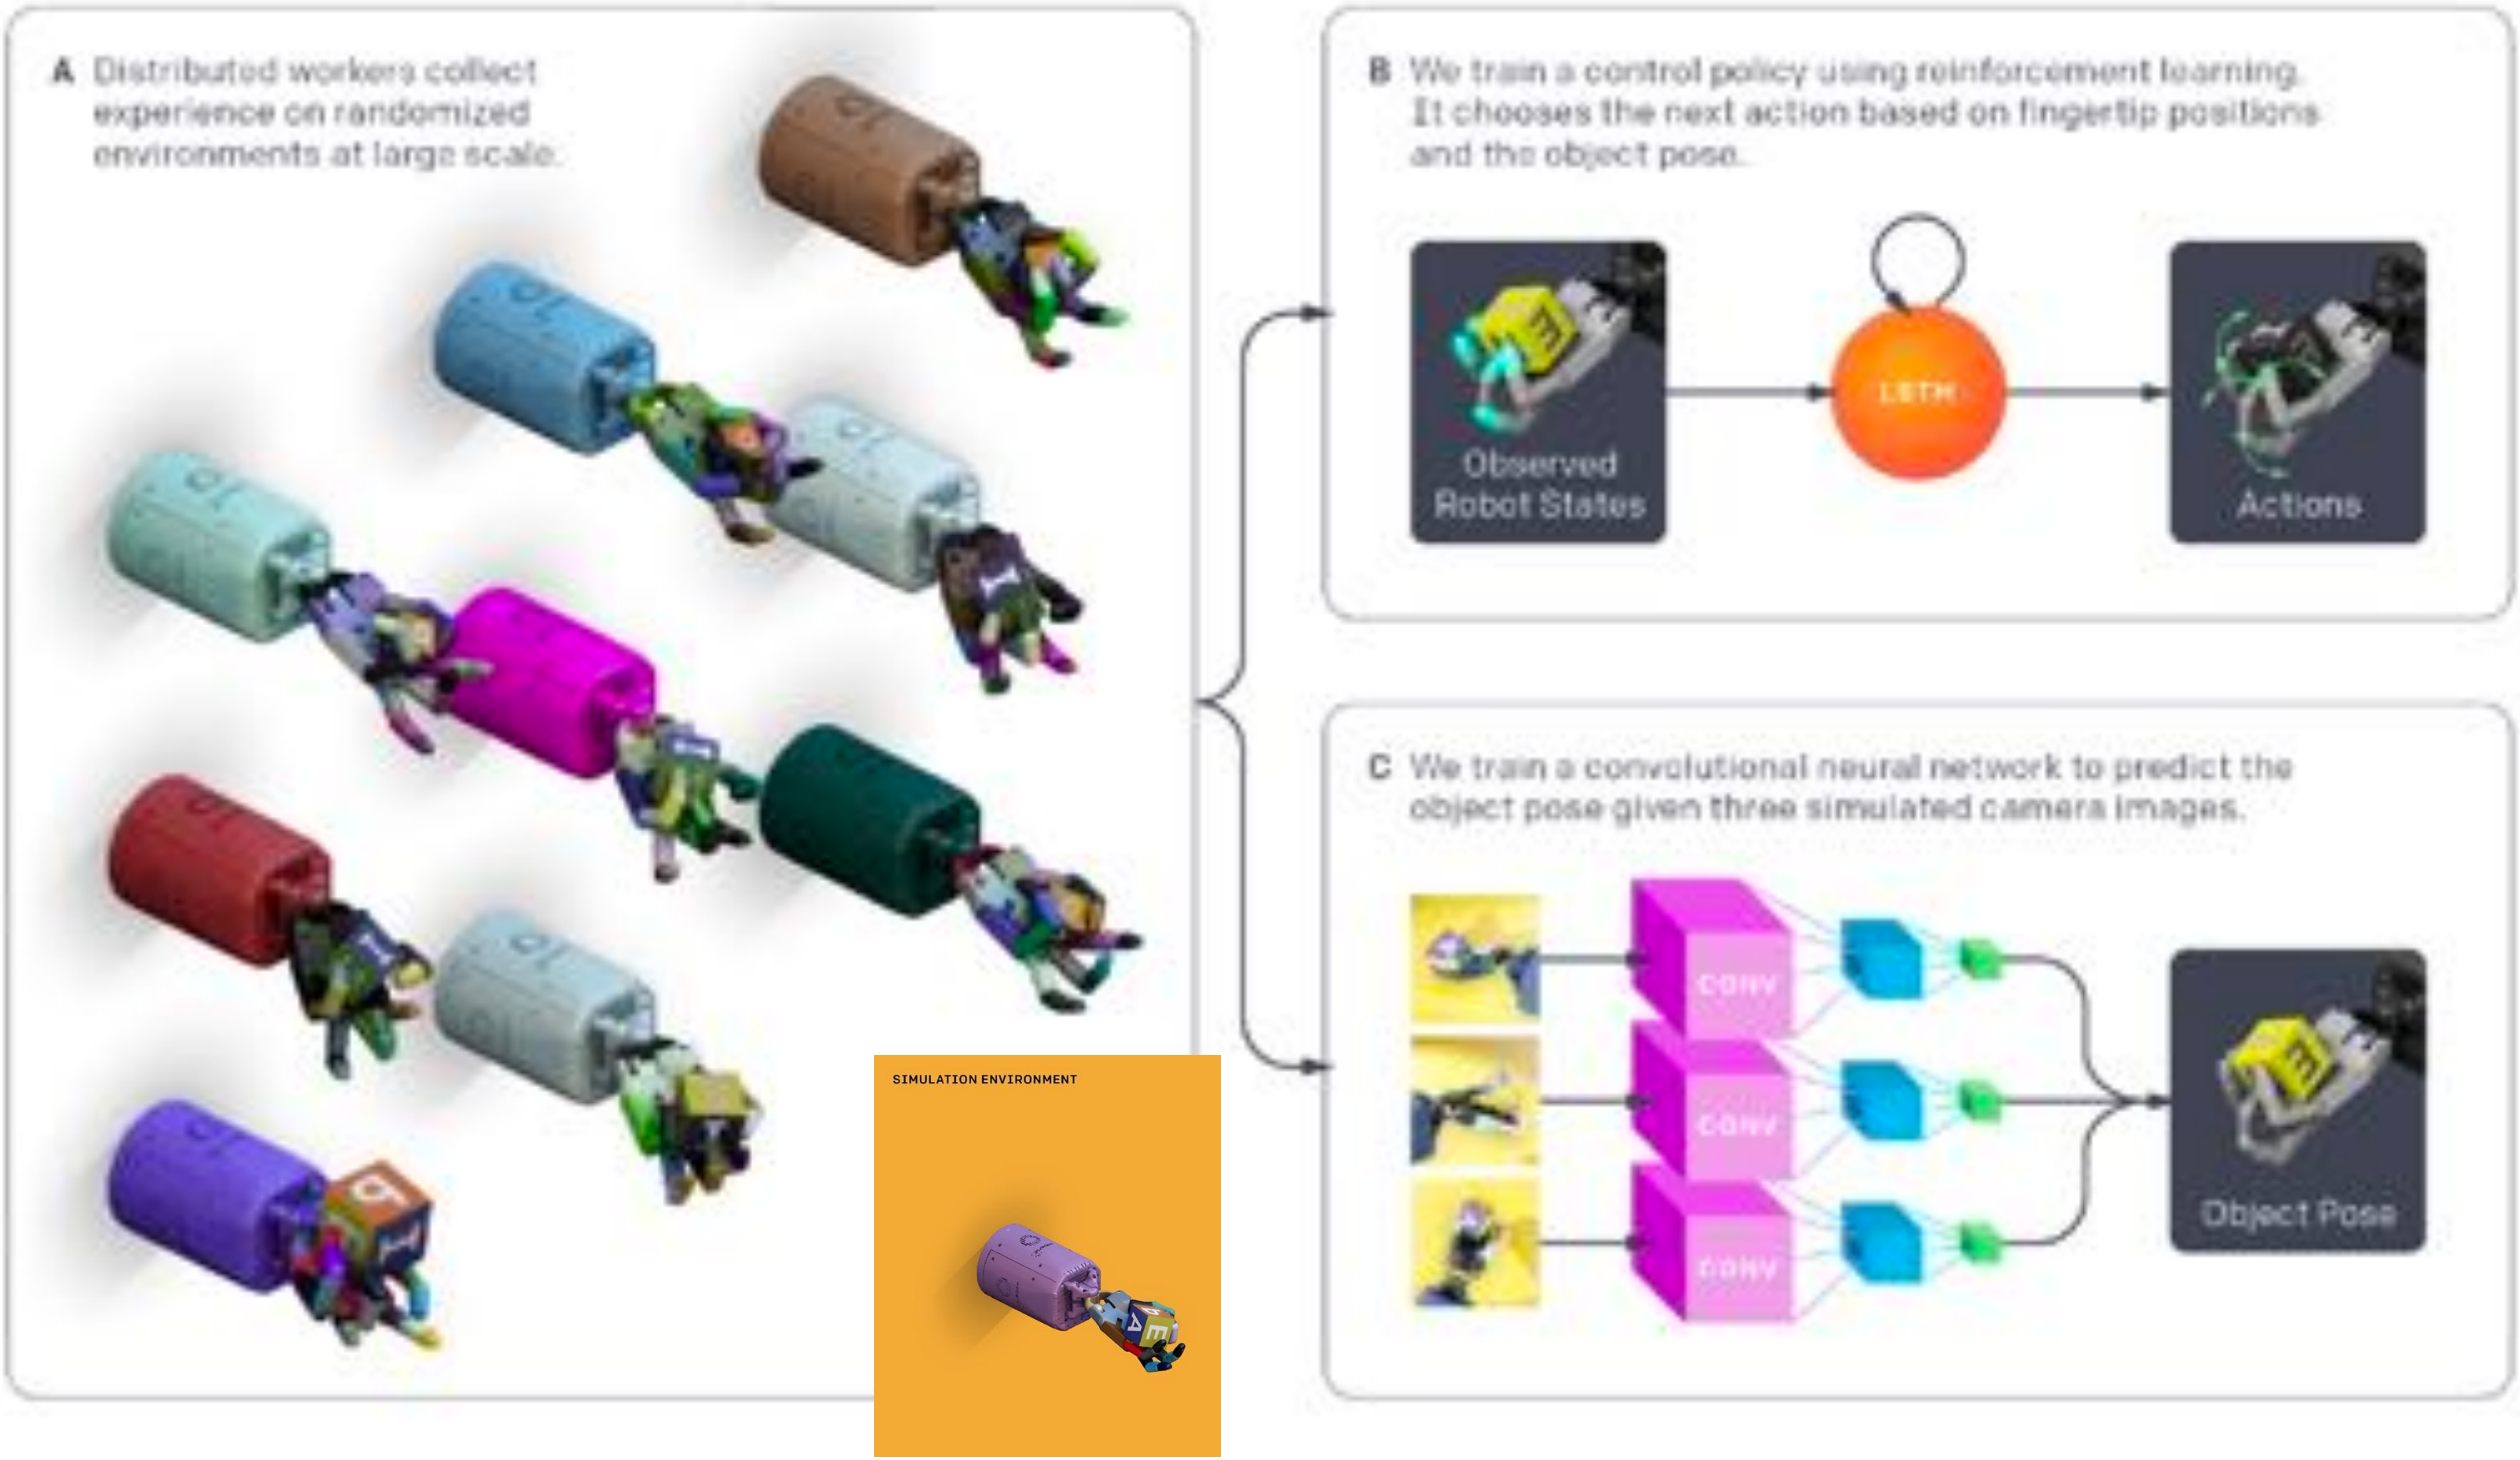
A Distributed workers collect experience on randomized environments at large scale.
B We train a control policy using reinforcement learning. It chooses the next action based on fingertip positions and the object pose.
Observed Robot States
LSTM
Actions
C We train a convolutional neural network to predict the object pose given three simulated camera images.
CONV
CONV
CONV
Object Pose
SIMULATION ENVIRONMENT


D We combine the pose estimation network
and the control policy to transfer to the real world.
CONV
CONV
CONV
Object Pose
Fingertip
Locations
LSTM
Actions

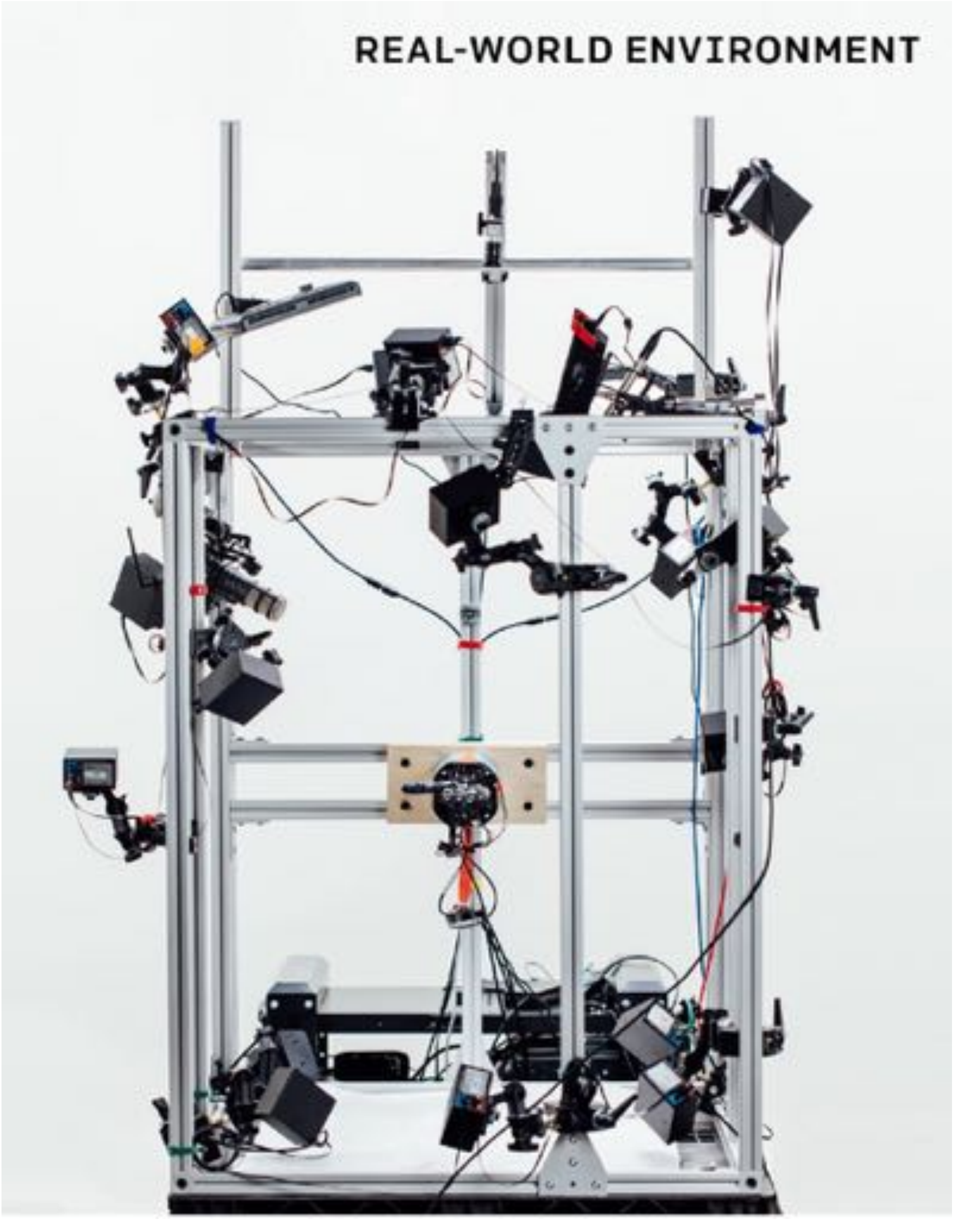
REAL-WORLD ENVIRONMENT

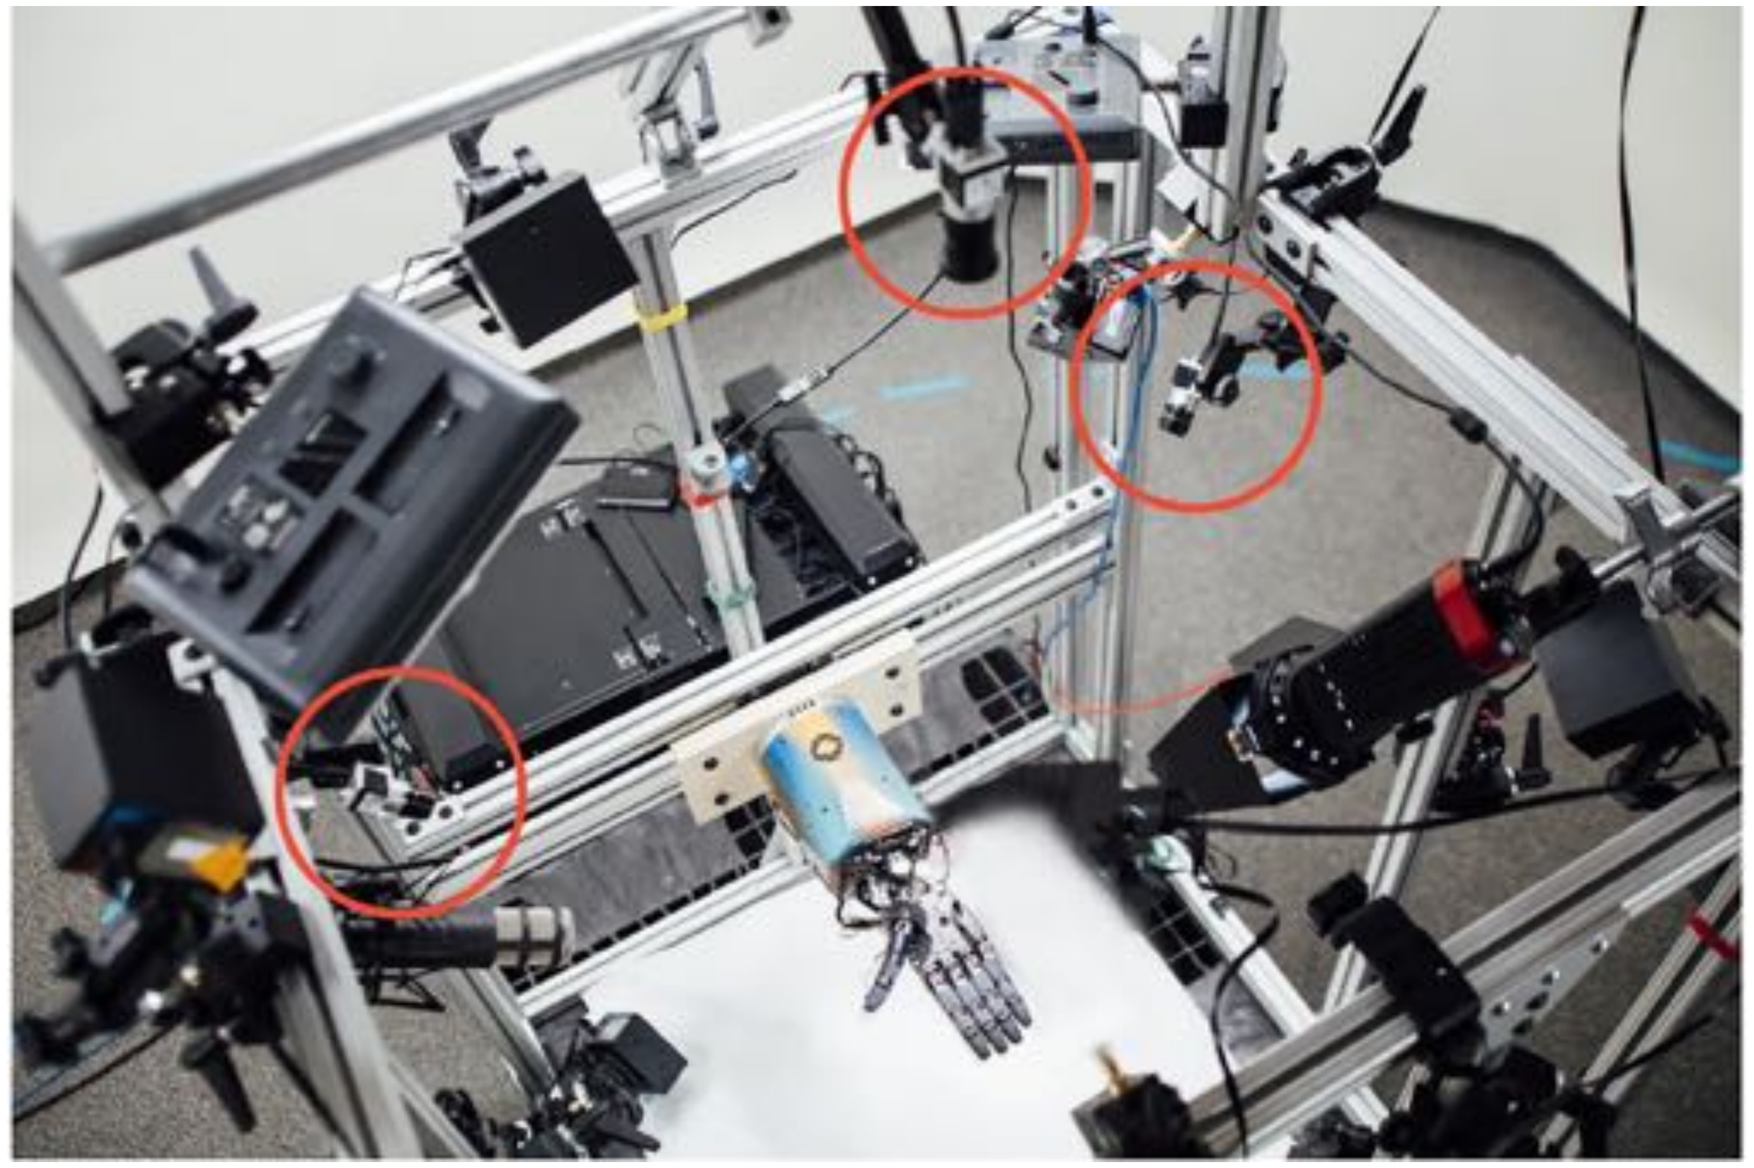

Reinforcement Learning (PPO)

- **State (60-dimensional):** angles and velocities of all robot joints as well as the position, rotation, and velocities of the object. Initial states are sampled by placing the object on the robot's palm in a random orientation and applying random actions for 100 steps

- **Goal:** desired orientation of the object

- **Action (20-dimensional):** desired angles of the hand joints. Each action dimension is discretized into 11 bins of equal size.

- **Reward:** improvement in (angular) distance to the goal orientation. An additional reward of 5 whenever a goal is achieved with some tolerance. A reward of -20 (penalty) whenever the object is dropped.

Randomizations

Table 1. Standard deviation of applied Gaussian observation noise.

Observation	Correlated	Uncorrelated
Fingertips positions	±1 mm	±2 mm
Object position	±5 mm	±1 mm
Object orientation	±0.1 rad	±0.1 rad
Fingertip marker positions	±3 mm	
Hand base marker position	±1 mm	

Table 2. Ranges of physics parameter randomizations.

Parameter	Scaling factor range
Object dimensions	uniform ([0.95, 1.05])
Object and robot link masses	uniform ([0.5, 1.5])
Surface friction coefficients	uniform ([0.7, 1.3])
Robot joint damping coefficients	loguniform ([0.3, 3.0])
Actuator force gains (P term)	loguniform ([0.75, 1.5])
Parameter	**Additive term range**
Joint limits	$\mathcal{N}(0, 0.15)$ rad
Gravity vector (per coordinate)	$\mathcal{N}(0, 0.4)$ m/s^2

Table 3. Standard deviation of action noise.

Noise type	Percentage of range
Uncorrelated additive	5%
Correlated additive	1.5%
Uncorrelated multiplicative	1.5%

Table 4. Ranges of vision randomizations.

Randomization type	Range
Number of cameras	3
Camera position	± 1.5 mm
Camera rotation	0–3° around a random axis
Camera field of view	± 1°
Robot material colors	uniform over RGB values
Robot material metallic level	5–25%[a]
Robot material glossiness level	0–100%[a]
Object material hue	± 1%
Object material saturation	± 15%
Object material value	± 15%
Object metallic level	5–15%[a]
Object glossiness level	5–15%[a]
Number of lights	4–6
Light position	uniform over upper half-sphere
Light relative intensity	1–5
Total light intensity	0–15[a]
Image contrast adjustment	50–150%
Additive per-pixel Gaussian noise	± 10%

[a]In units used by Unity. See https://unity3d.com/learn/tutorials/s/graphics.

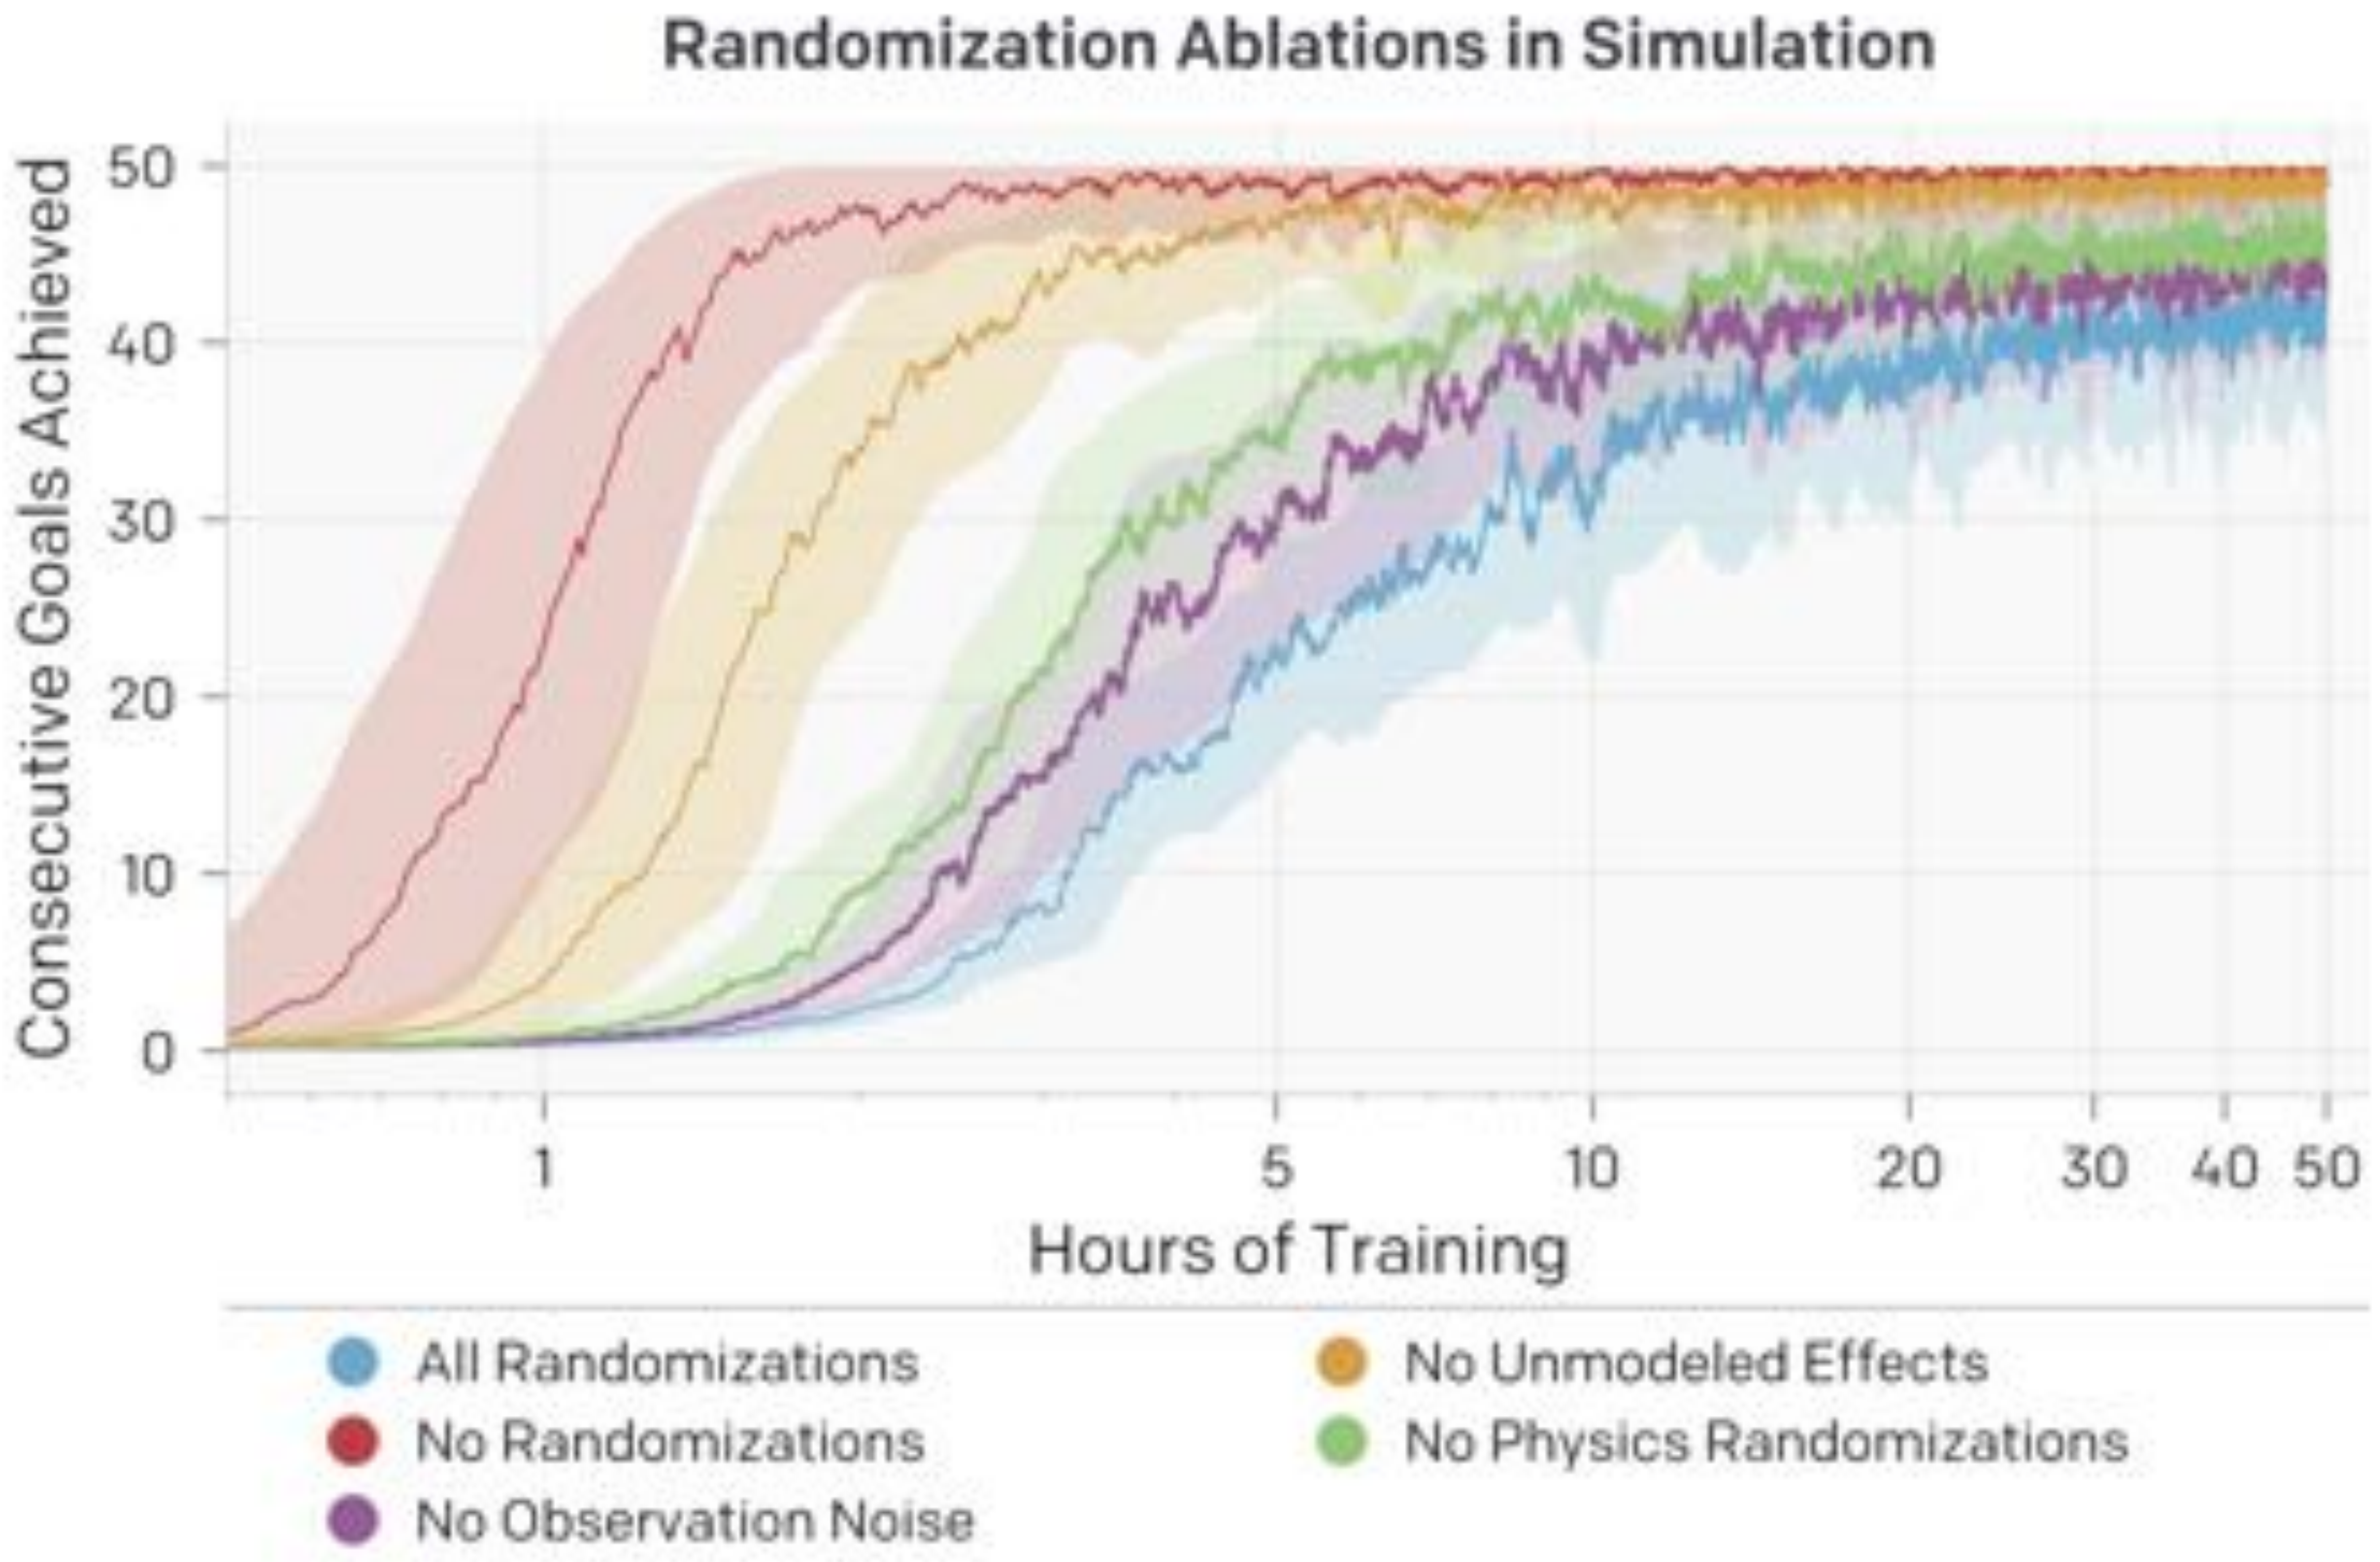

Randomization Ablations in Simulation
Consecutive Goals Achieved
0
10
20
30
40
50
1
5
10
20
30
40
50
Hours of Training
All Randomizations
No Randomizations
No Observation Noise
No Unmodeled Effects
No Physics Randomizations

Table 10. The number of successful consecutive rotations on the physical robot of five policies trained separately in environments with different randomizations held out. The first five rows use PhaseSpace for object pose estimation and were run on the same robot at the same time. Trials for each row were interleaved in case the state of the robot changed during the trials. The last two rows were measured at a different time from the first five and used the vision model to estimate the object pose.

Training environment	Mean	Median	Individual trials (sorted)									
All randomizations (state)	18.8 ± 17.1	13	50	41	29	27	14	12	6	4	4	1
No randomizations (state)	1.1 ± 1.9	0	6	2	2	1	0	0	0	0	0	0
No observation noise (state)	15.1 ± 14.5	8.5	45	35	23	11	9	8	7	6	6	1
No physics randomizations (state)	3.5 ± 2.5	2	7	7	7	3	2	2	2	2	2	1
No unmodeled effects (state)	3.5 ± 4.8	2	16	7	3	3	2	2	1	1	0	0
All randomizations (vision)	15.2 ± 14.3	11.5	46	28	26	15	13	10	8	3	2	1
No observation noise (vision)	5.9 ± 6.6	3.5	20	12	11	6	5	2	2	1	0	0

Table 11. The number of successful consecutive rotations on the physical robot of three policies with different network architectures trained on an environment with all randomizations. Results for each row were collected at different times on the physical robot.

Network architecture	Mean	Median	Individual trials (sorted)									
LSTM policy/LSTM value (state)	18.8 ± 17.1	13	50	41	29	27	14	12	6	4	4	1
FF policy/LSTM value (state)	4.7 ± 4.1	3.5	15	7	6	5	4	3	3	2	2	0
FF policy/FF value (state)	4.6 ± 4.3	3	15	8	6	5	3	3	2	2	2	0

Massive Computation

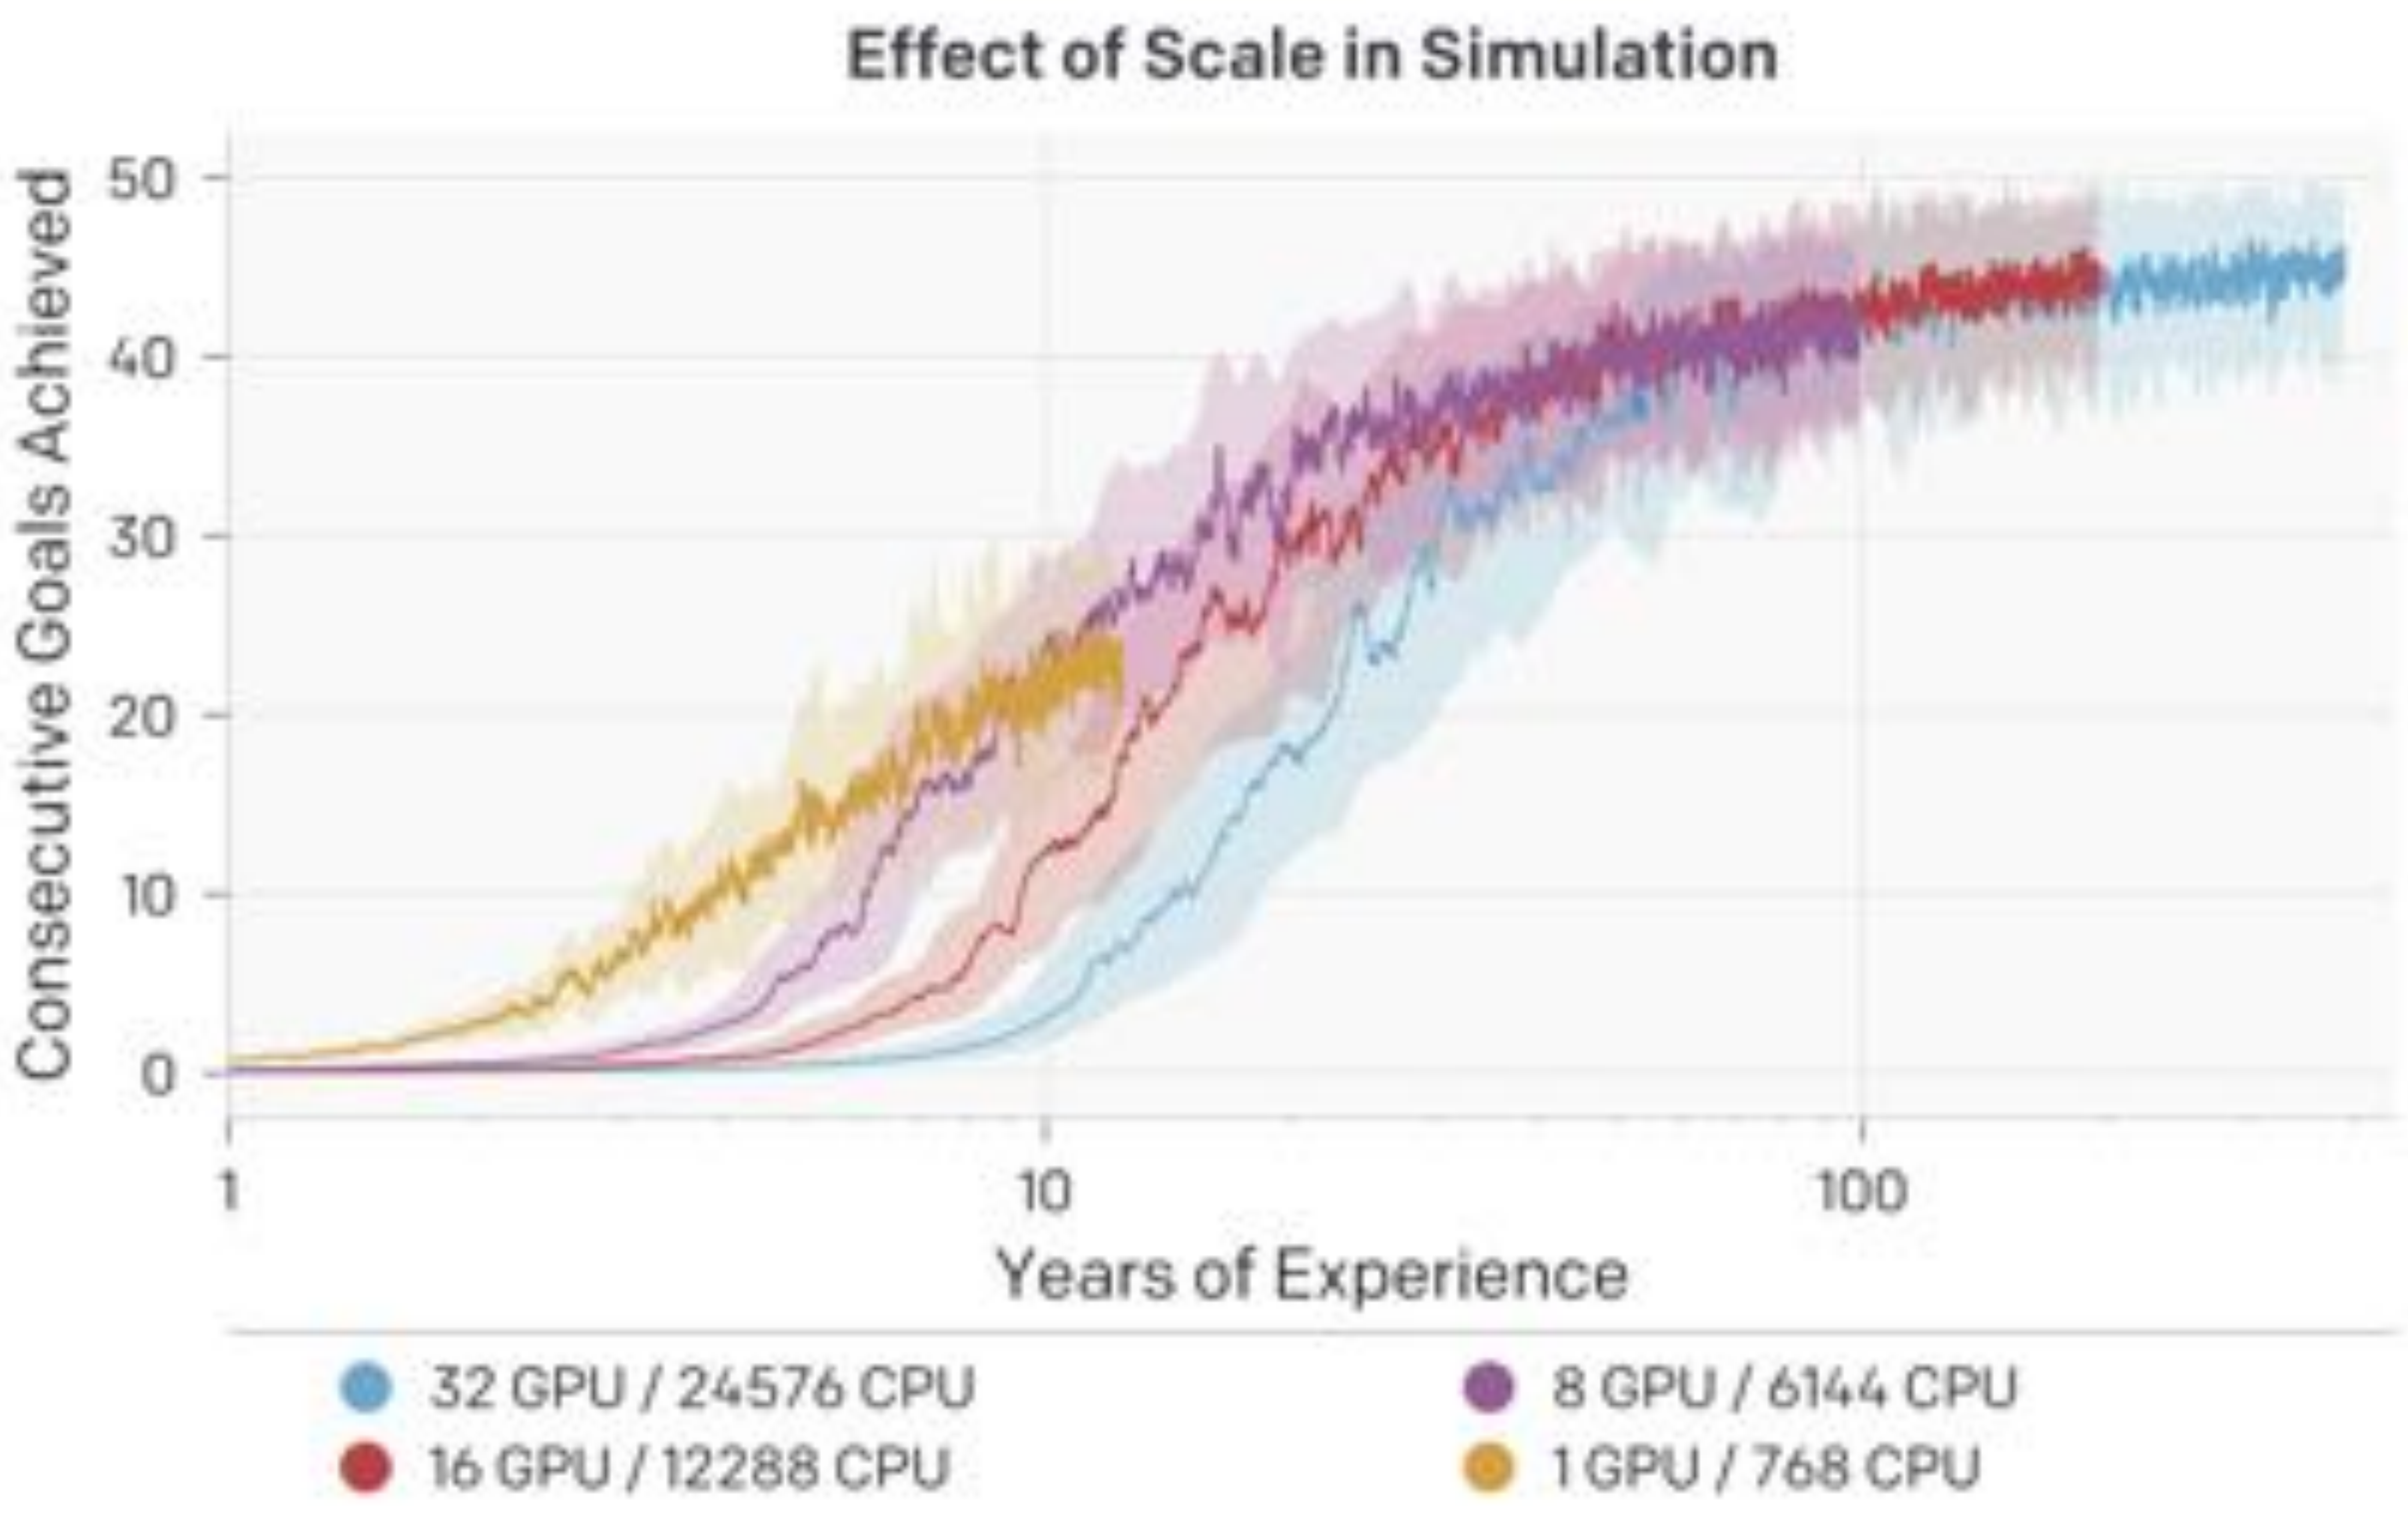

In our implementation, a pool of 384 worker machines, each with 16 CPU cores, generate experience by rolling out the current version of the policy in a sample from the previously described distribution of randomized simula- tions. … This setup allows us to generate about 2 years of simulated experience per hour.

The optimization is performed on a single machine with eight GPUs. The optimizer threads pull down generated experience … and then stage it to their respective GPU's memory for processing. After computing gradients locally, they are averaged across all threads using MPI, which we then use to update the network parameters.

https://www.youtube.com/watch?v=DKe8FumoD4E&t=198s

Results

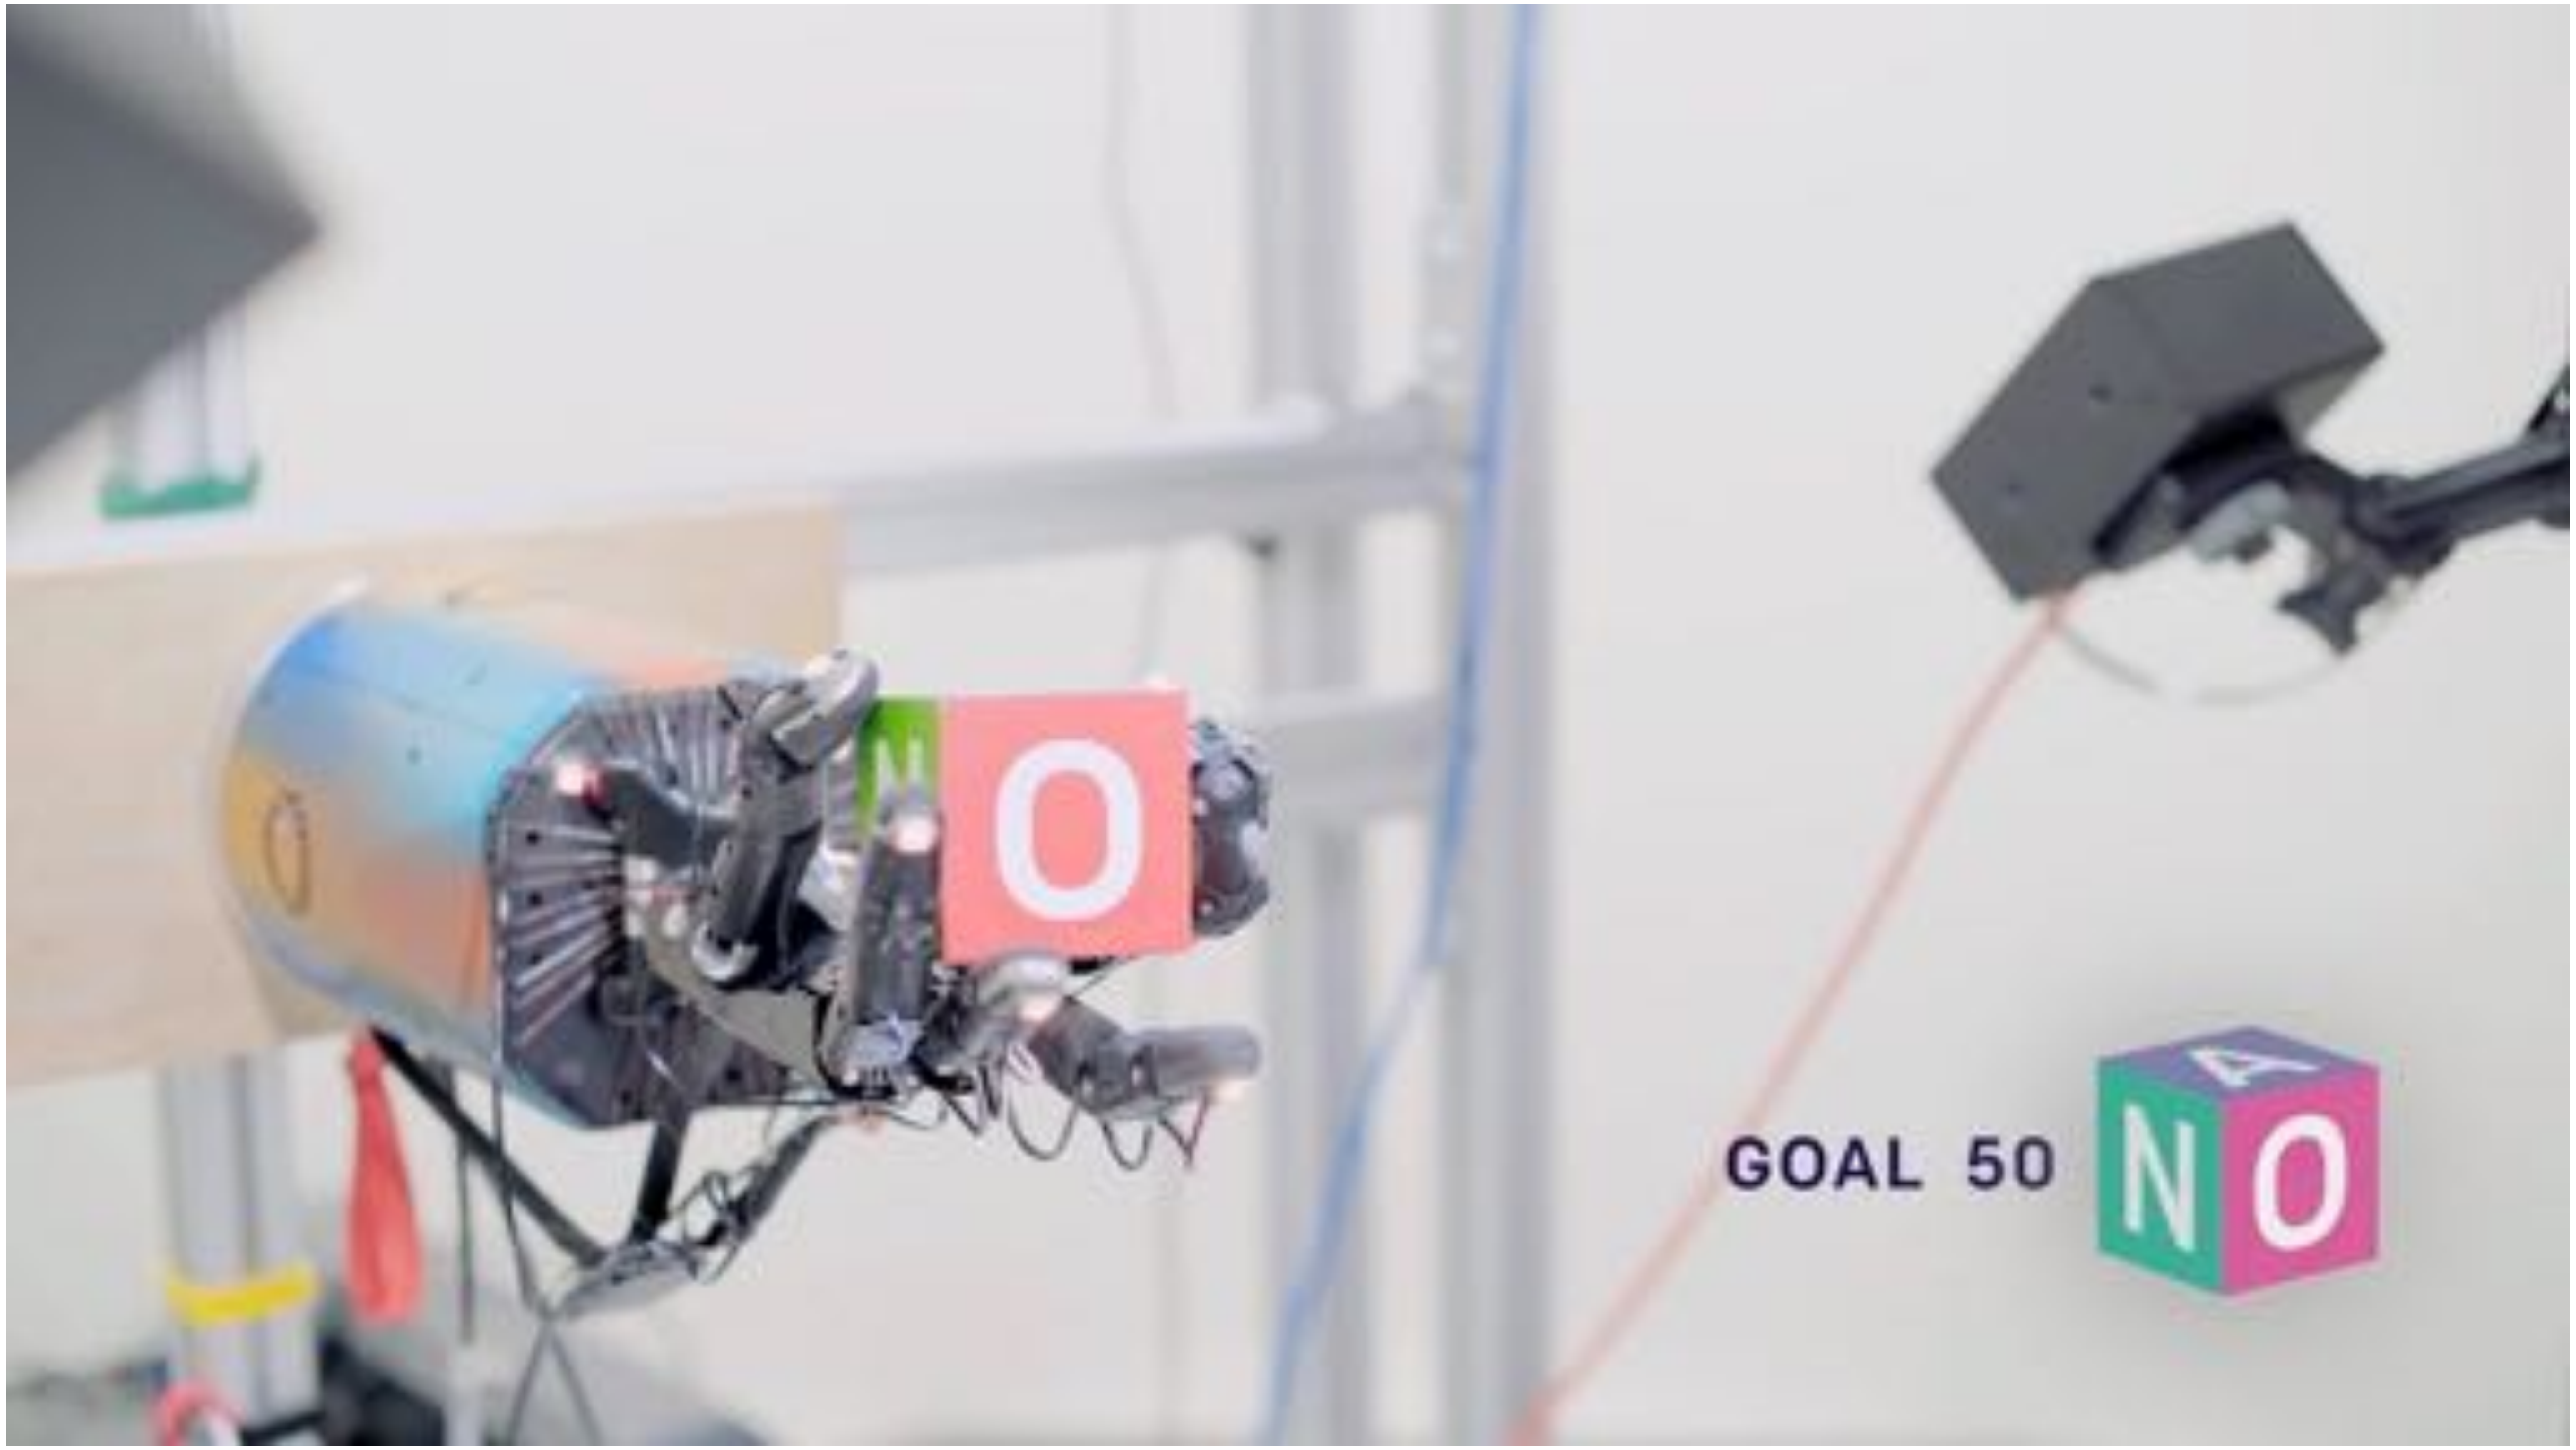

Robotics: Science and Systems 2021
Held Virtually, July 12–16, 2021

http://www.roboticsproceedings.org/rss17/p089.pdf

Surprisingly Robust In-Hand Manipulation: An Empirical Study

Aditya Bhatt* Adrian Sieler* Steffen Puhlmann Oliver Brock

Abstract—We present in-hand manipulation skills on a dexterous, compliant, anthropomorphic hand. Even though these skills were derived in a simplistic manner, they exhibit surprising robustness to variations in shape, size, weight, and placement of the manipulated object. They are also very insensitive to variation of execution speeds, ranging from highly dynamic to quasi-static. The robustness of the skills leads to compositional properties that enable extended and robust manipulation programs. To explain the surprising robustness of the in-hand manipulation skills, we performed a detailed, empirical analysis of the skills' performance. From this analysis, we identify three principles for skill design: 1) Exploiting the hardware's innate ability to drive hard-to-model contact dynamics. 2) Taking actions to constrain these interactions, funneling the system into a narrow set of possibilities. 3) Composing such action sequences into complex manipulation programs. We believe that these principles constitute an important foundation for robust robotic in-hand manipulation, and possibly for manipulation in general.

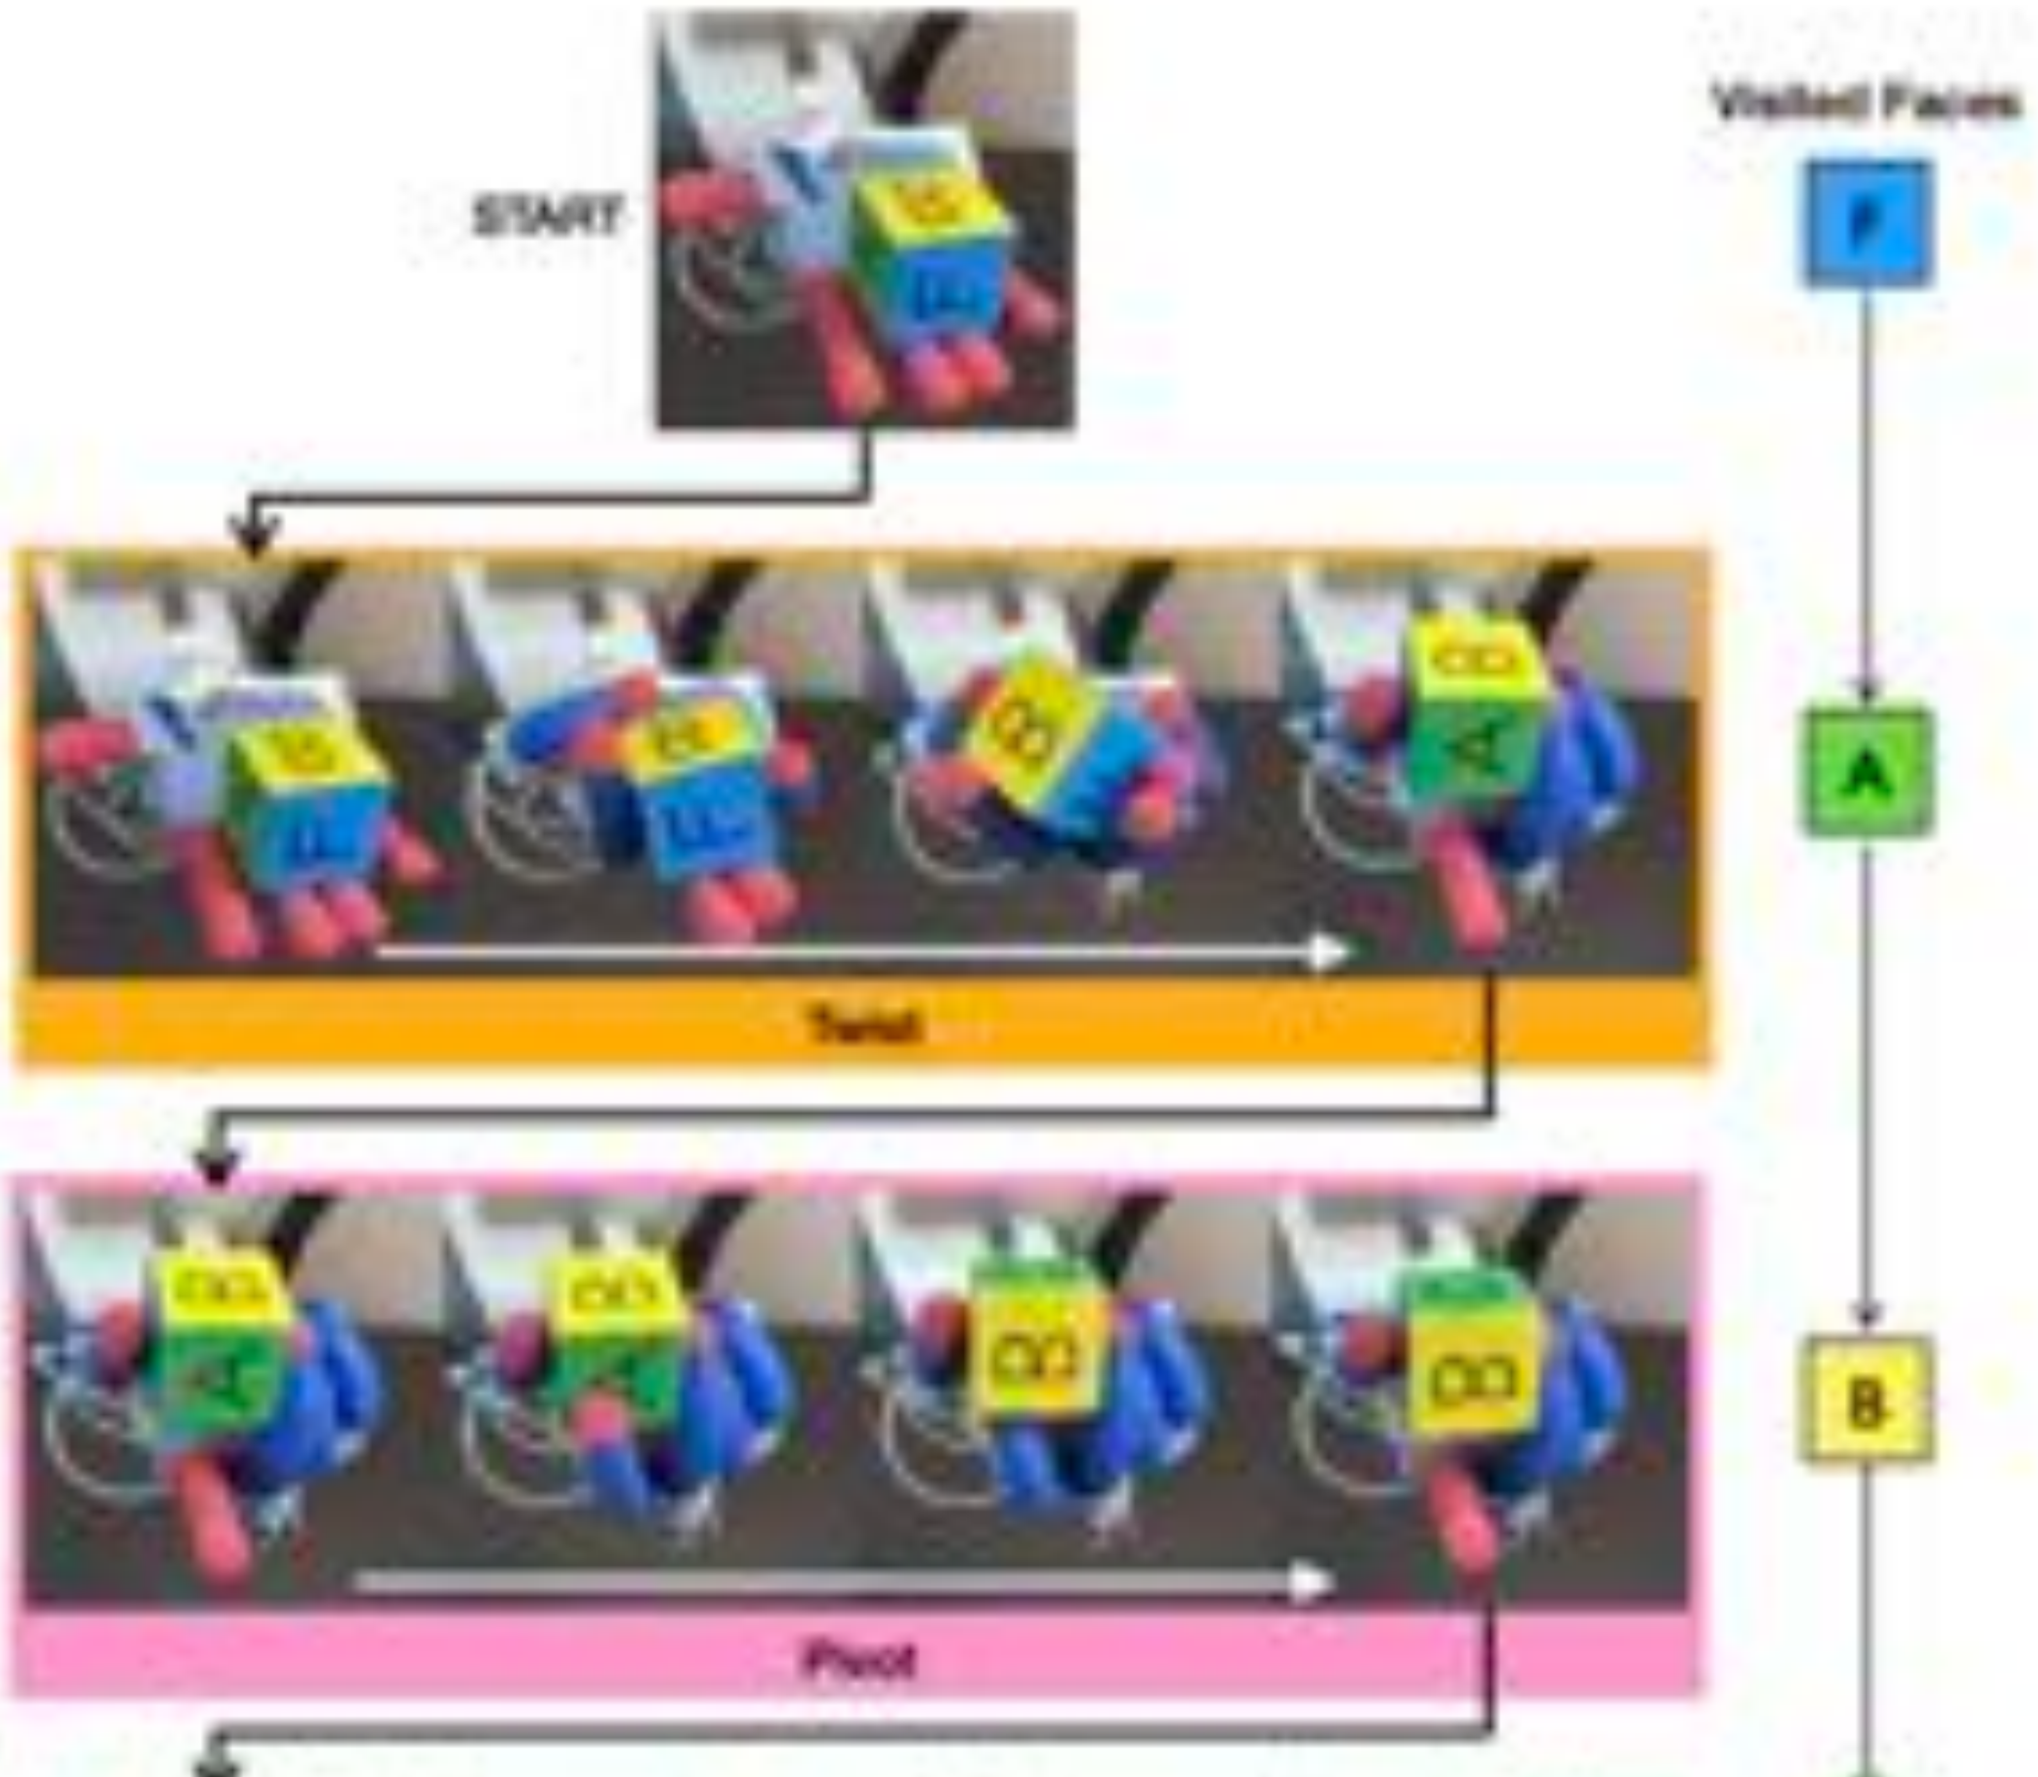

Context

"The state of the art, at this time, is the groundbreaking work presented by OpenAI [1], who used Deep Reinforcement Learning to produce remarkably dexterous behavior on a five- fingered robotic hand, first manipulating a cube and later even an articulated Rubik's cube [2]. Their learned skills feature contact-rich movements like finger-gaiting, pivoting, and the exploitation of gravity."

RBO 3 Skills

Spin and Shift

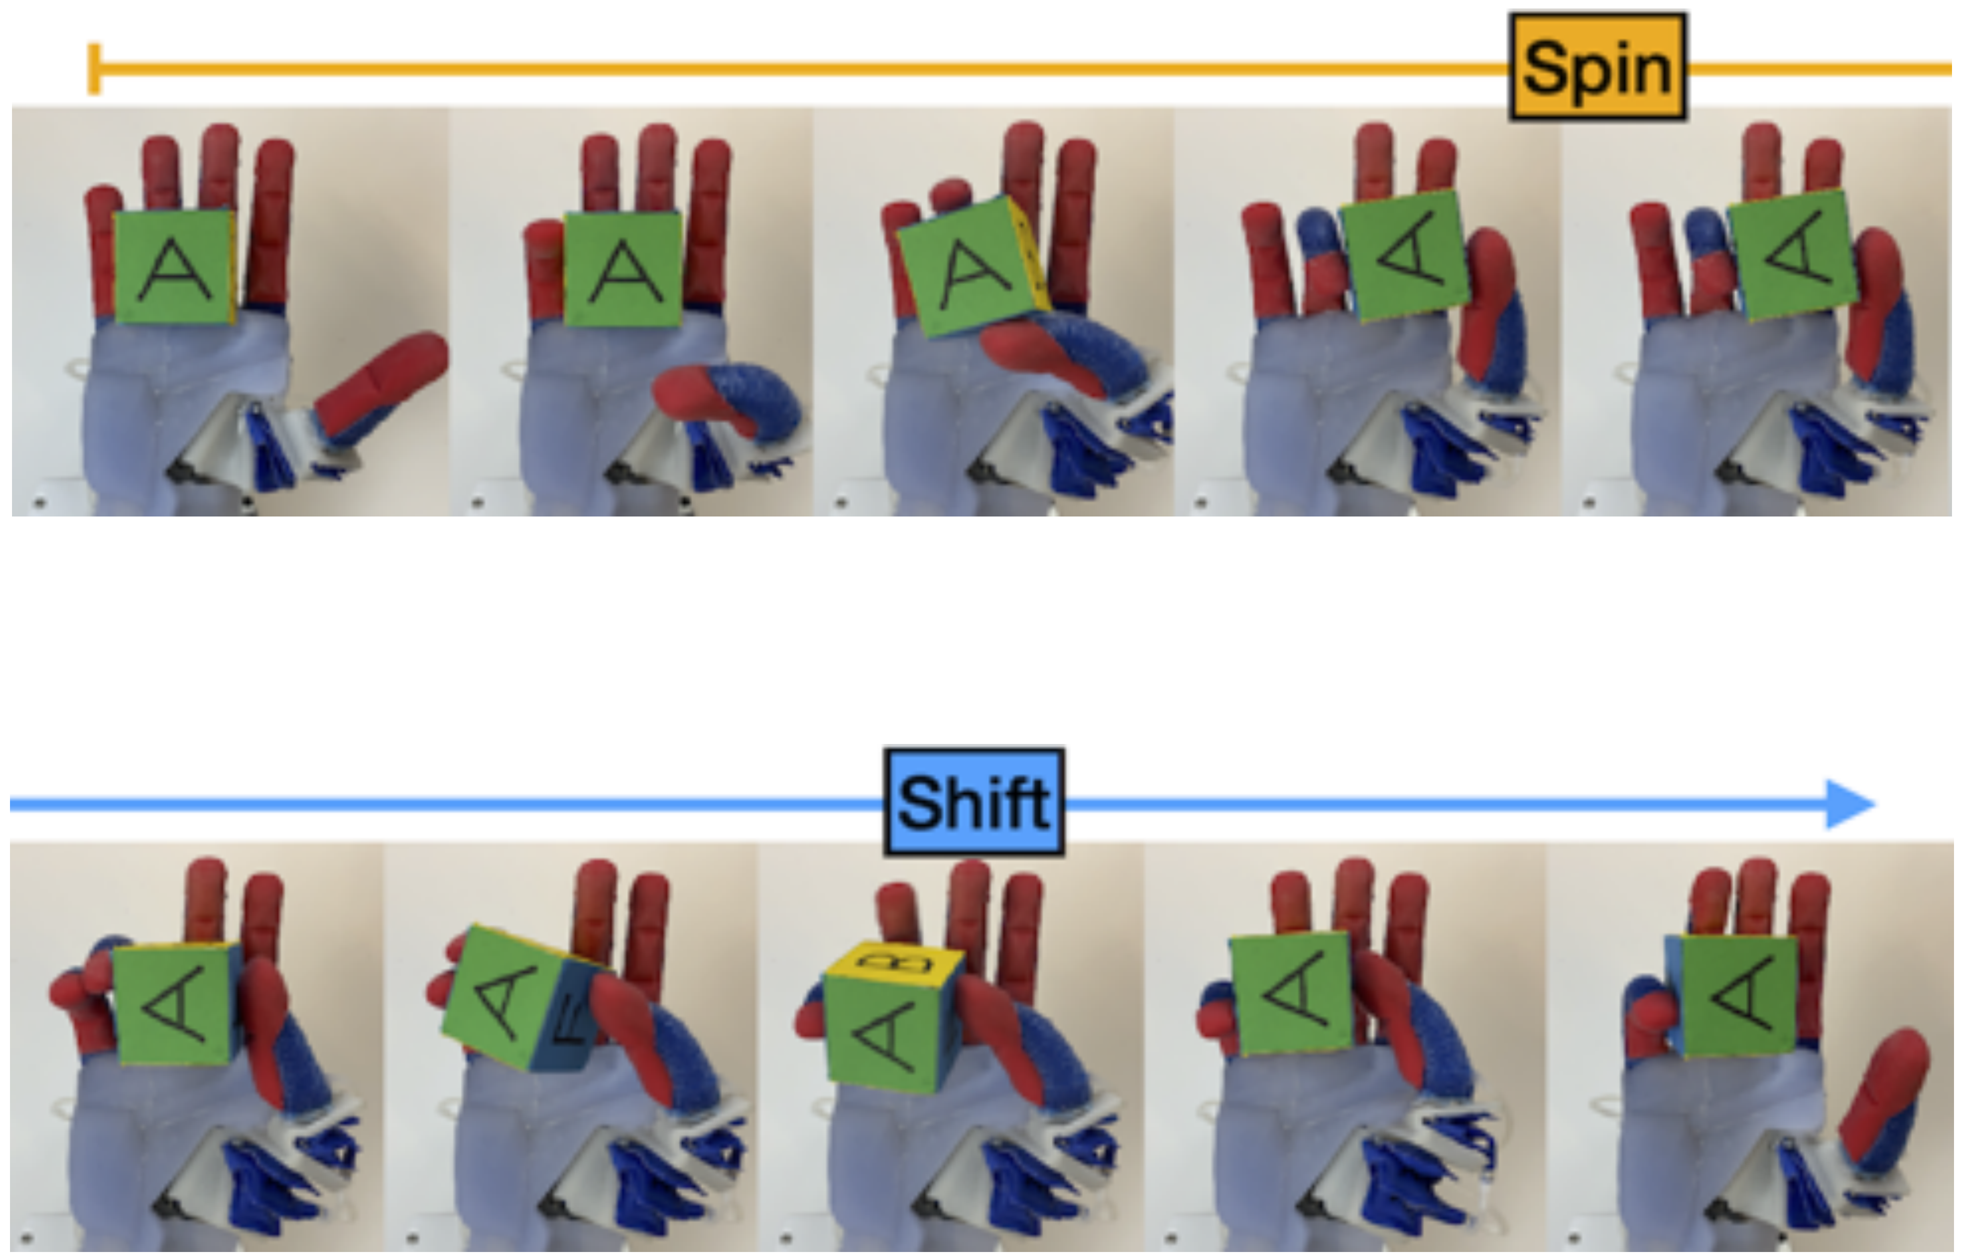

Skills are keyframed
(remember the mixer?)

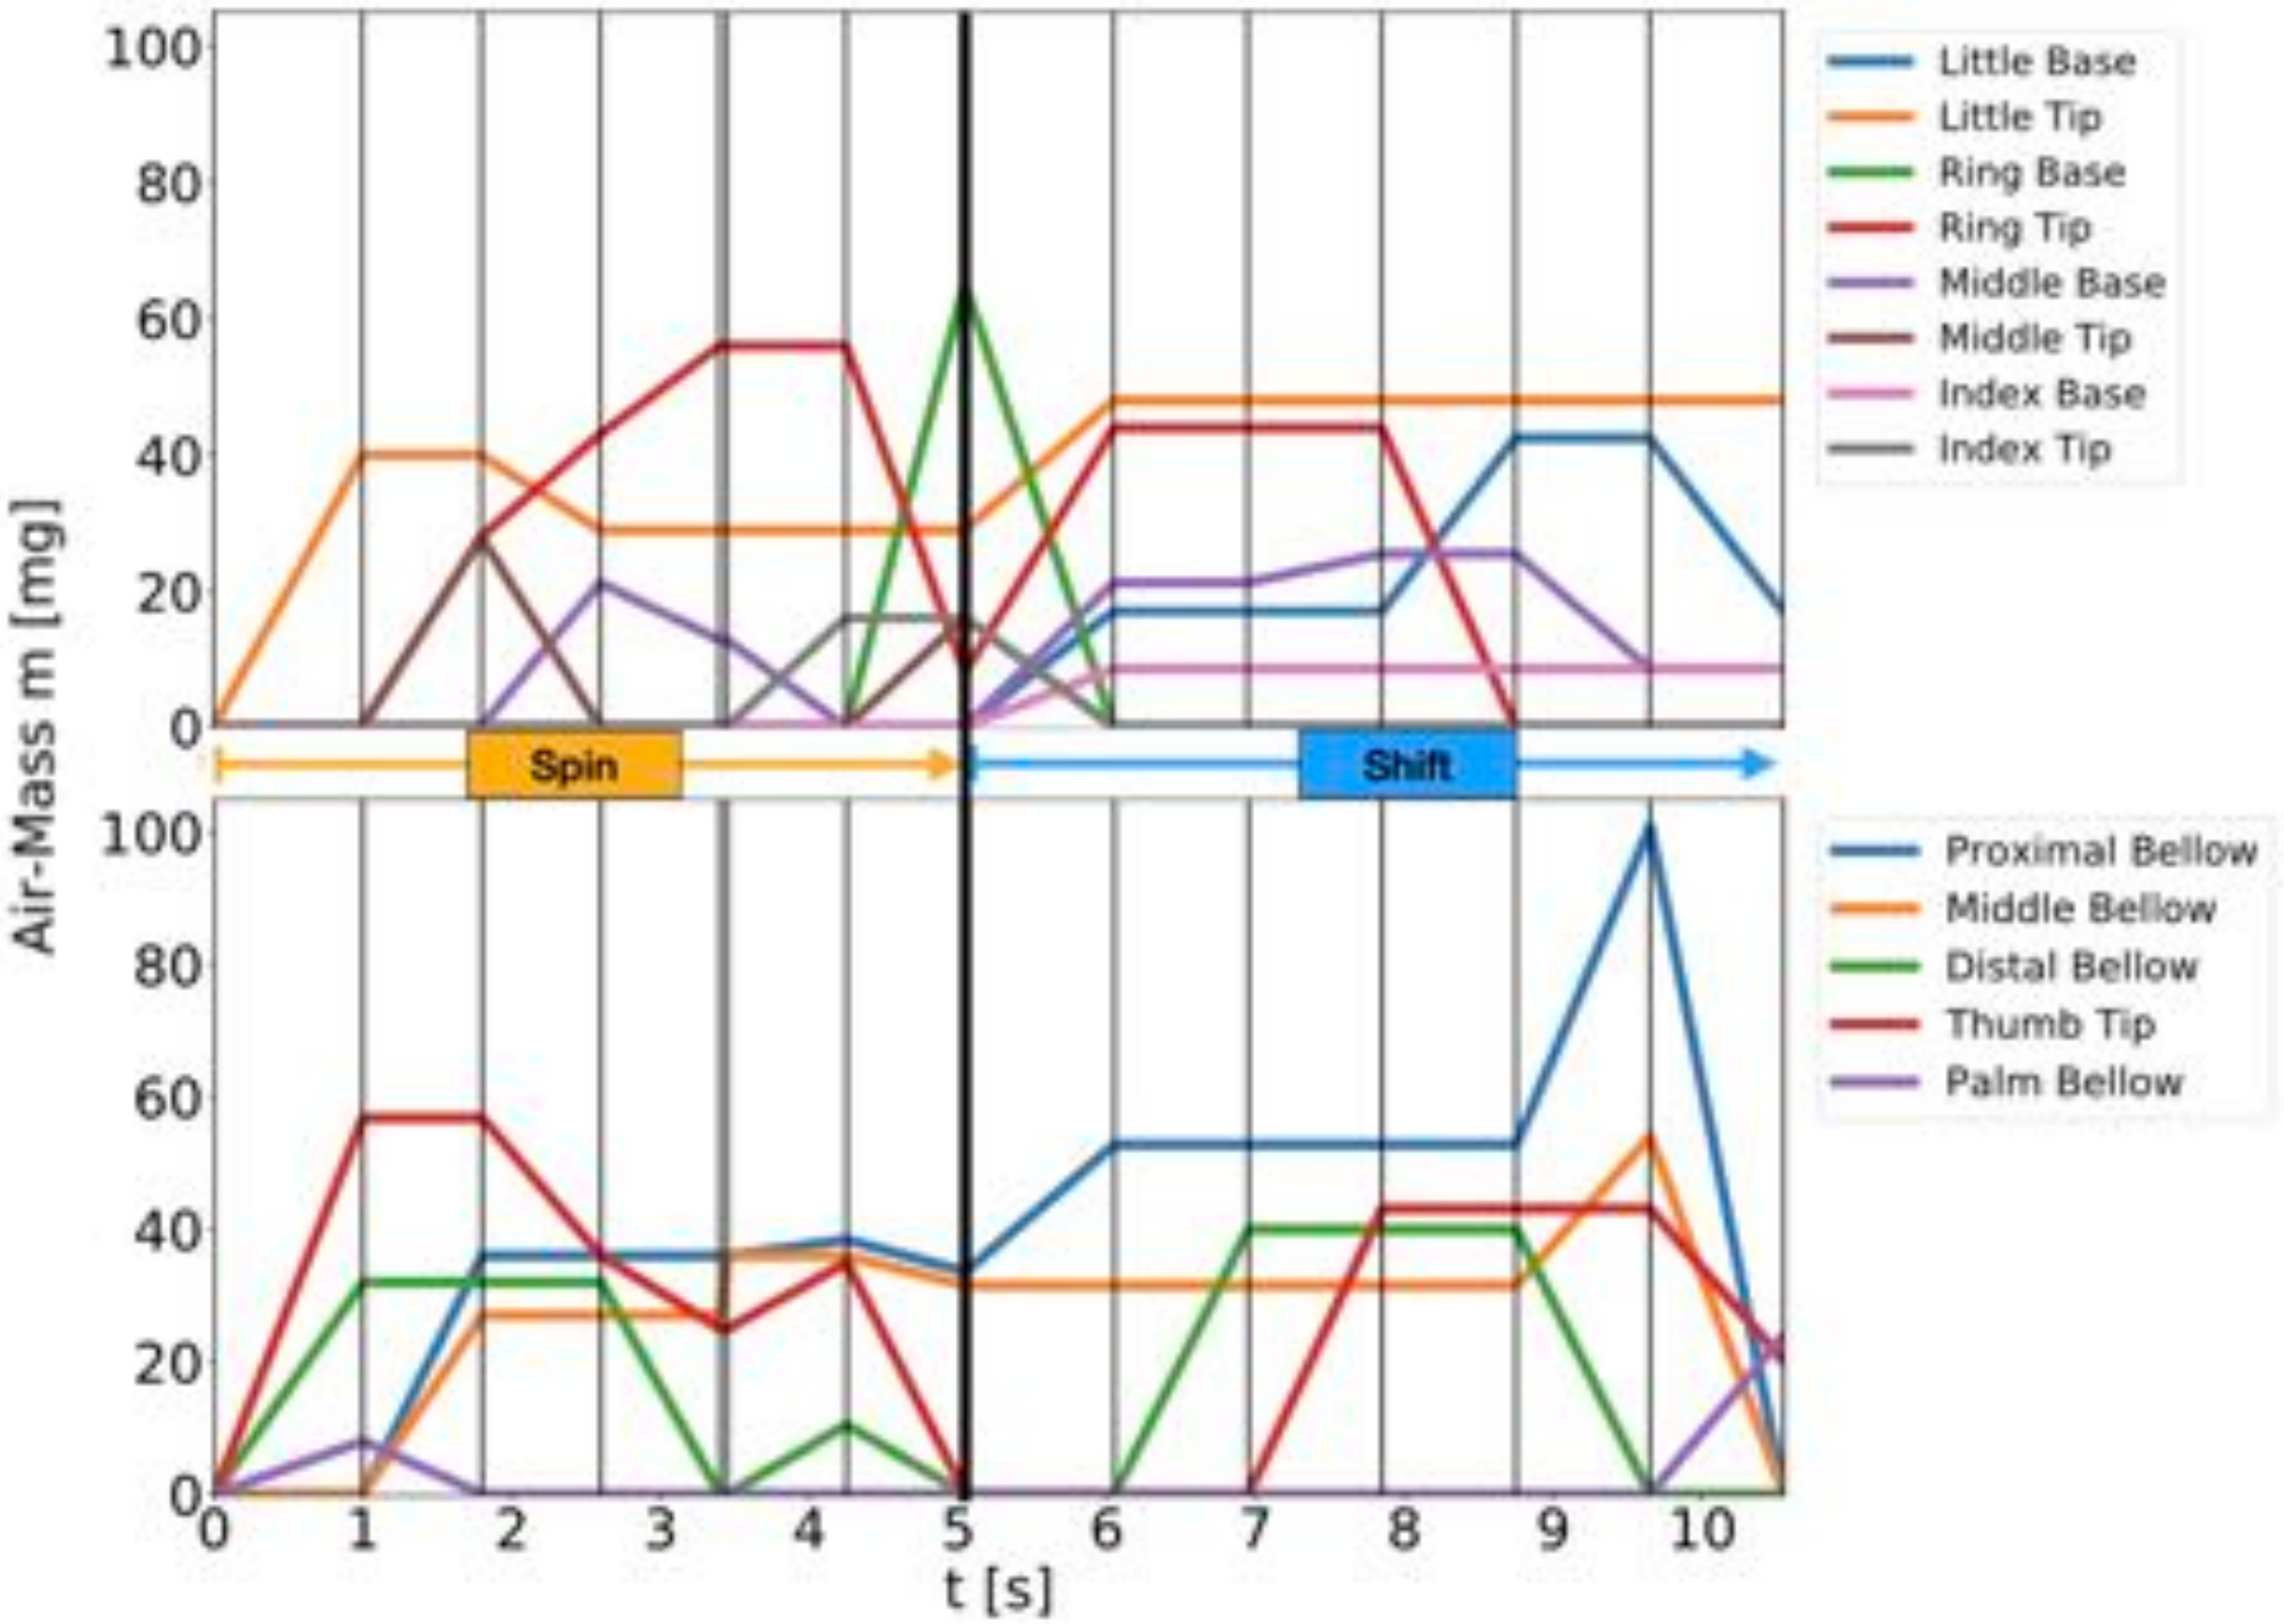

Fig. 5. **Air-mass actuation signal for the *spin* + *shift* skill.** Each vertical line corresponds to an intermediate keyframe. Appendix A contains a detailed description of each of these keyframes.

RBO 3 Skills

5 skills were keyframed

- **Spin** uses the thumb to rotate the cube ounter-clockwise by 90° around the ring-finger, and places it to rest on the index and middle finger. It does so with 7 keyframes.
- **Shift**, also called Ring→Little (RL) Finger Gait, gaits the cube from the ring-finger to the little finger, and places it to rest on the ring and middle finger (5 keyframes).
- **Twist** uses the thumb, middle, and ring fingers to lift the cube into a precision grip, in the process rotating the cube counter-clockwise by 90° (9 keyframes).
- **Pivot** maintains this posture, contacting the cube with the index finger to rotate it by 90° around the grip axis (6 keyframes).
- **Middle→Ring (MR) Finger Gait**, a variant of *shift*, gaits the cube from this posture to a grasp between the thumb and middle finger, and places it to rest on the index and ring fingers (6 keyframes).

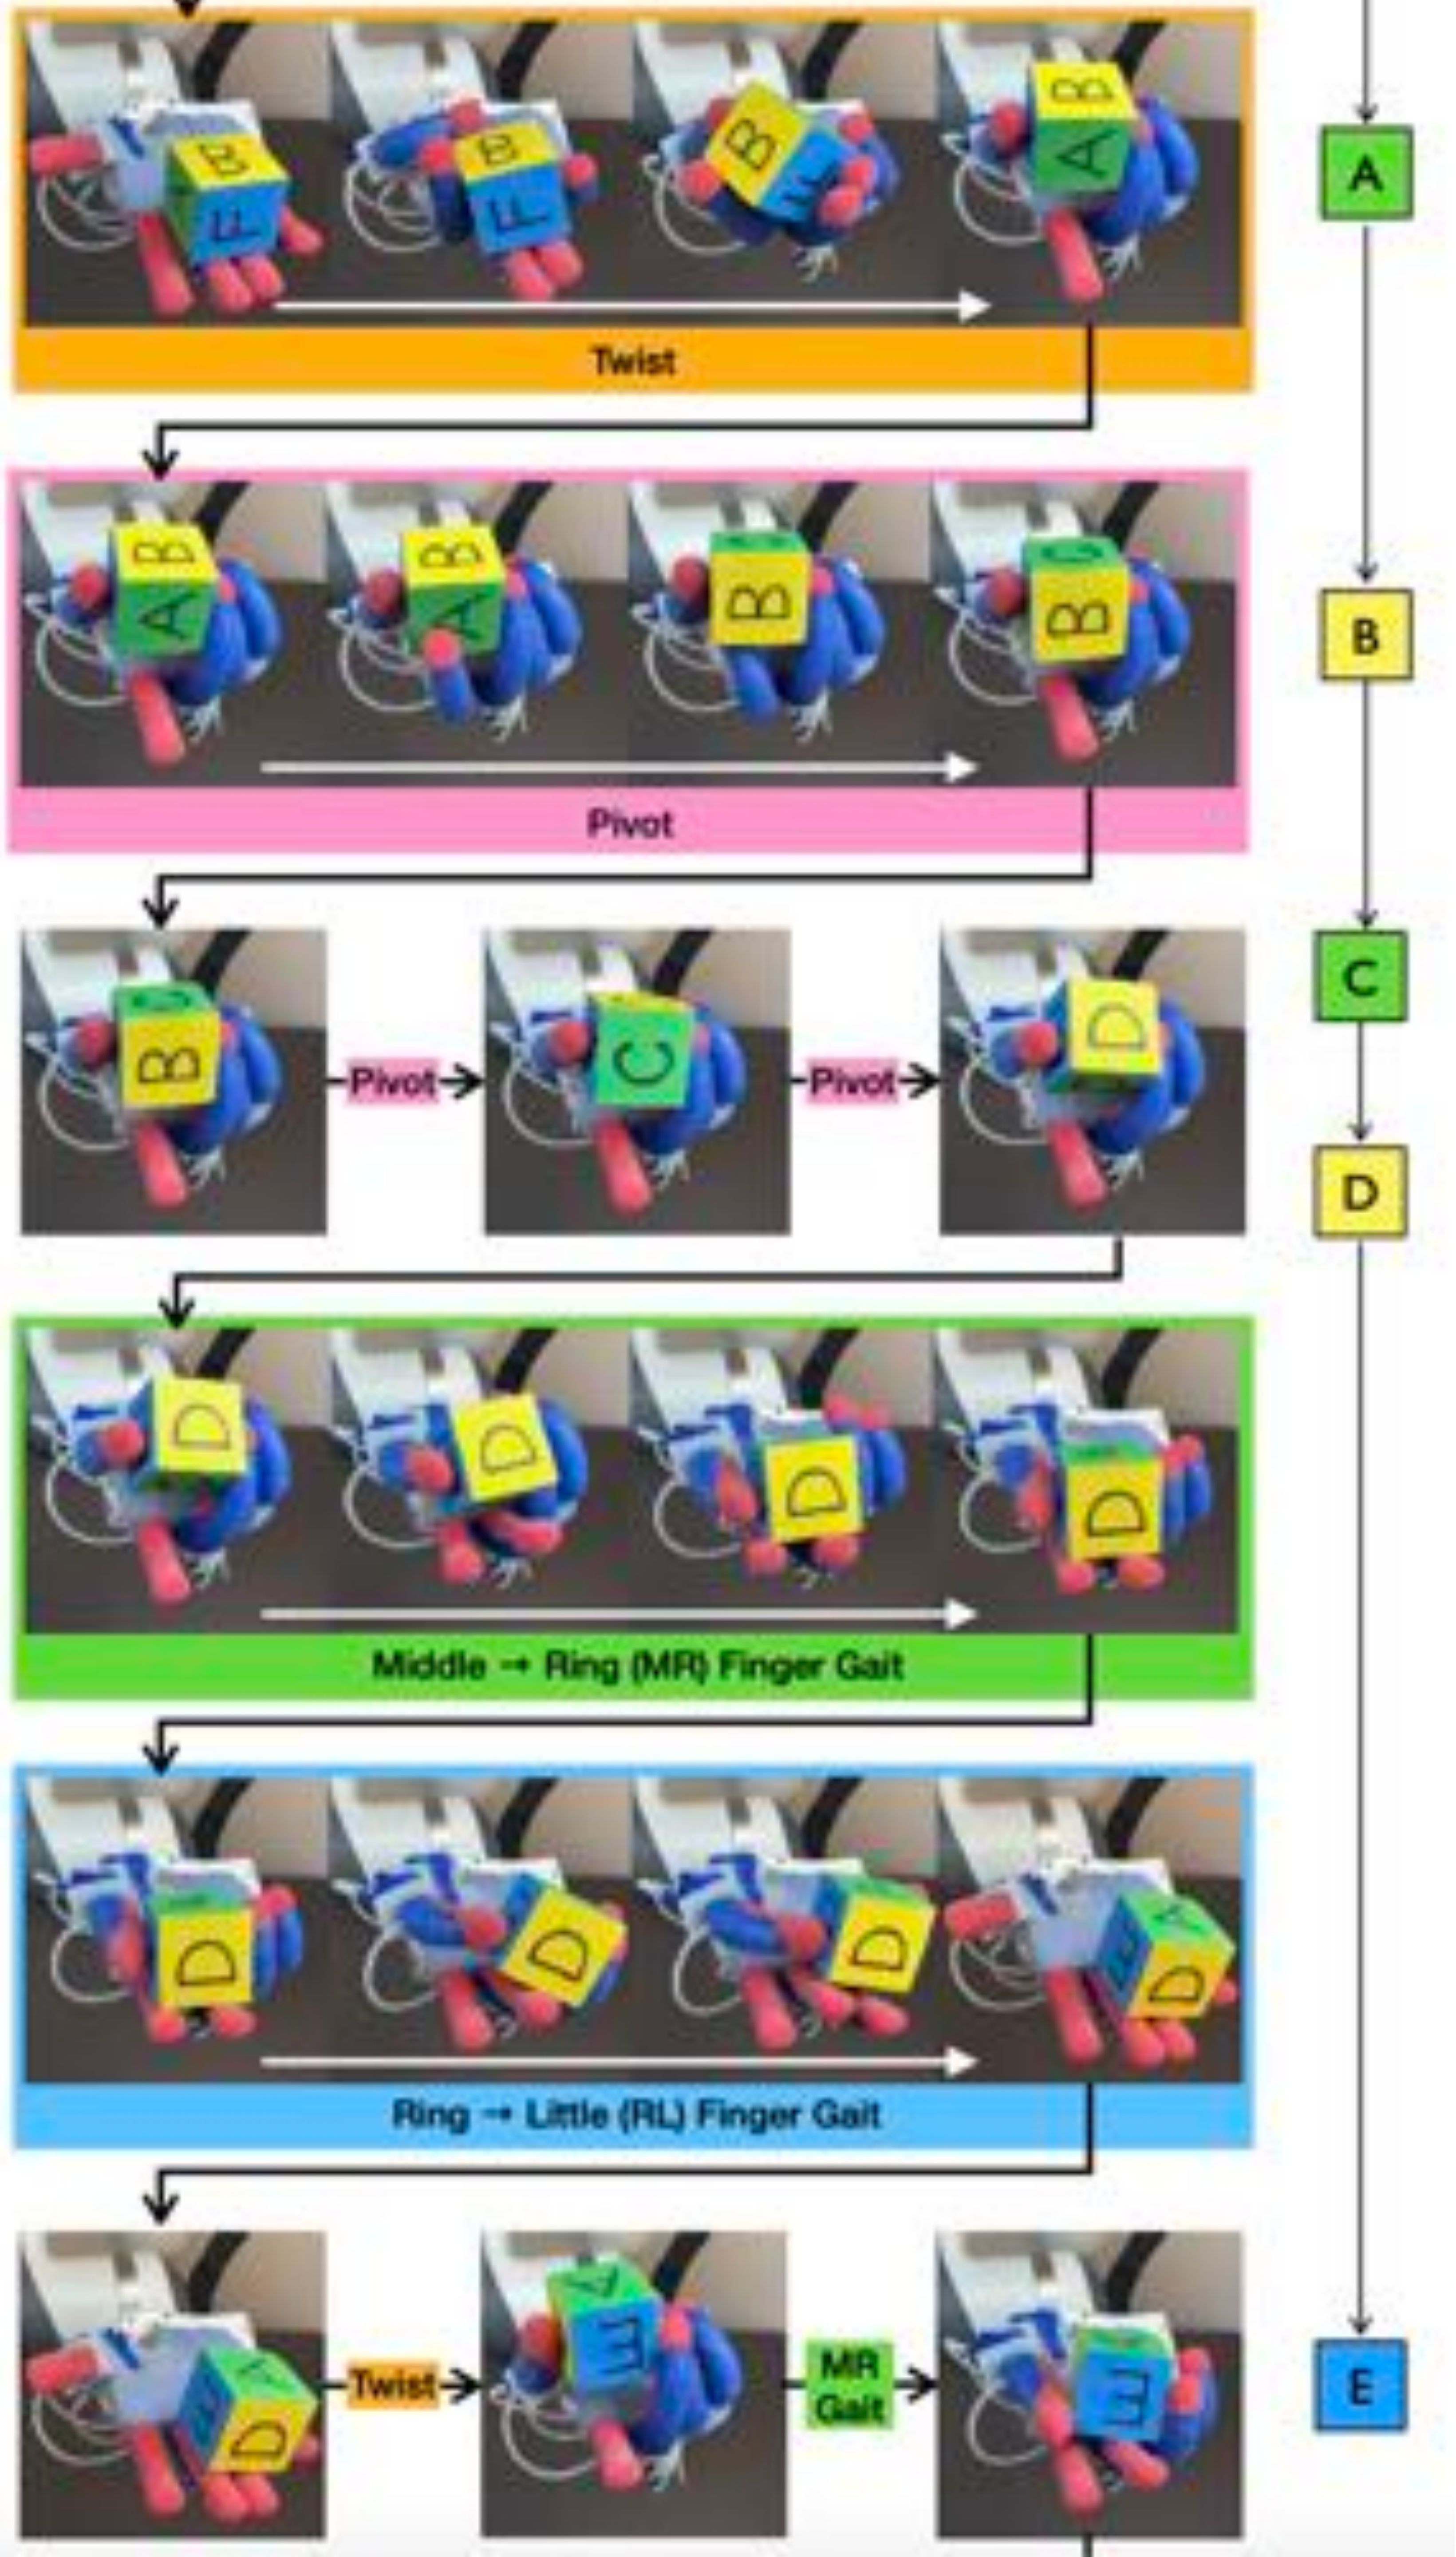

Robustness to Object Placement

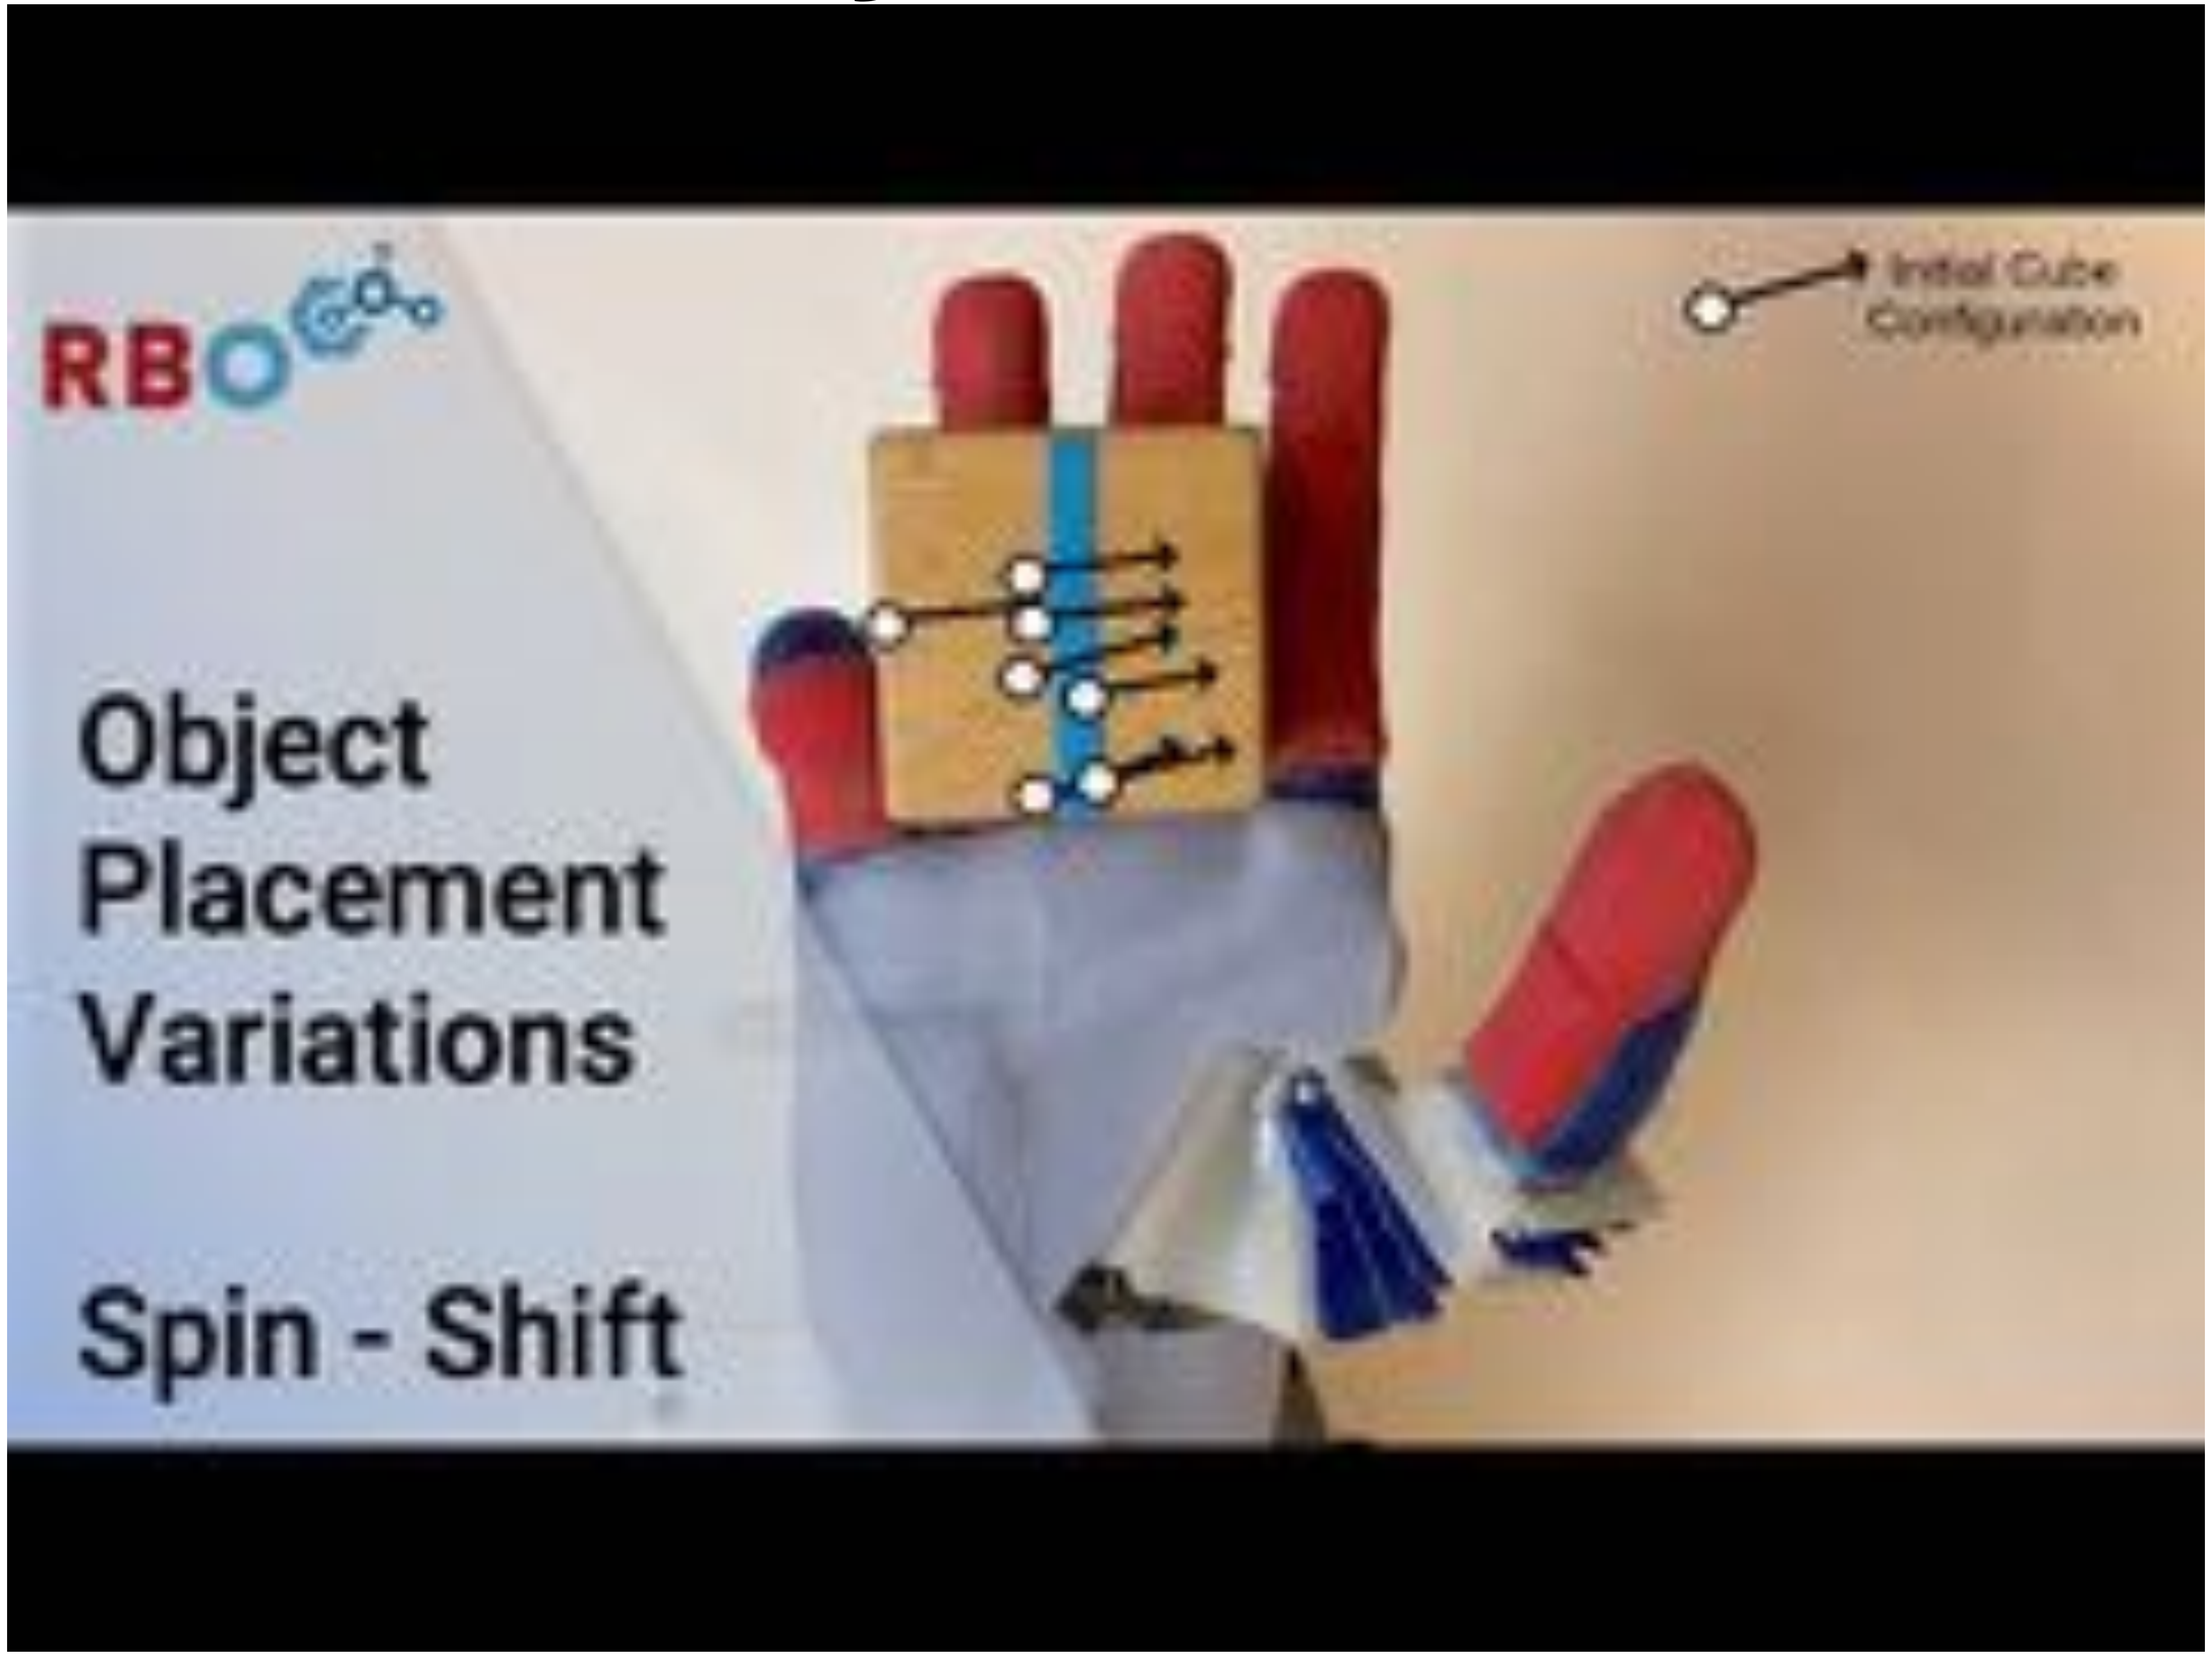

Robustness to Object Placement

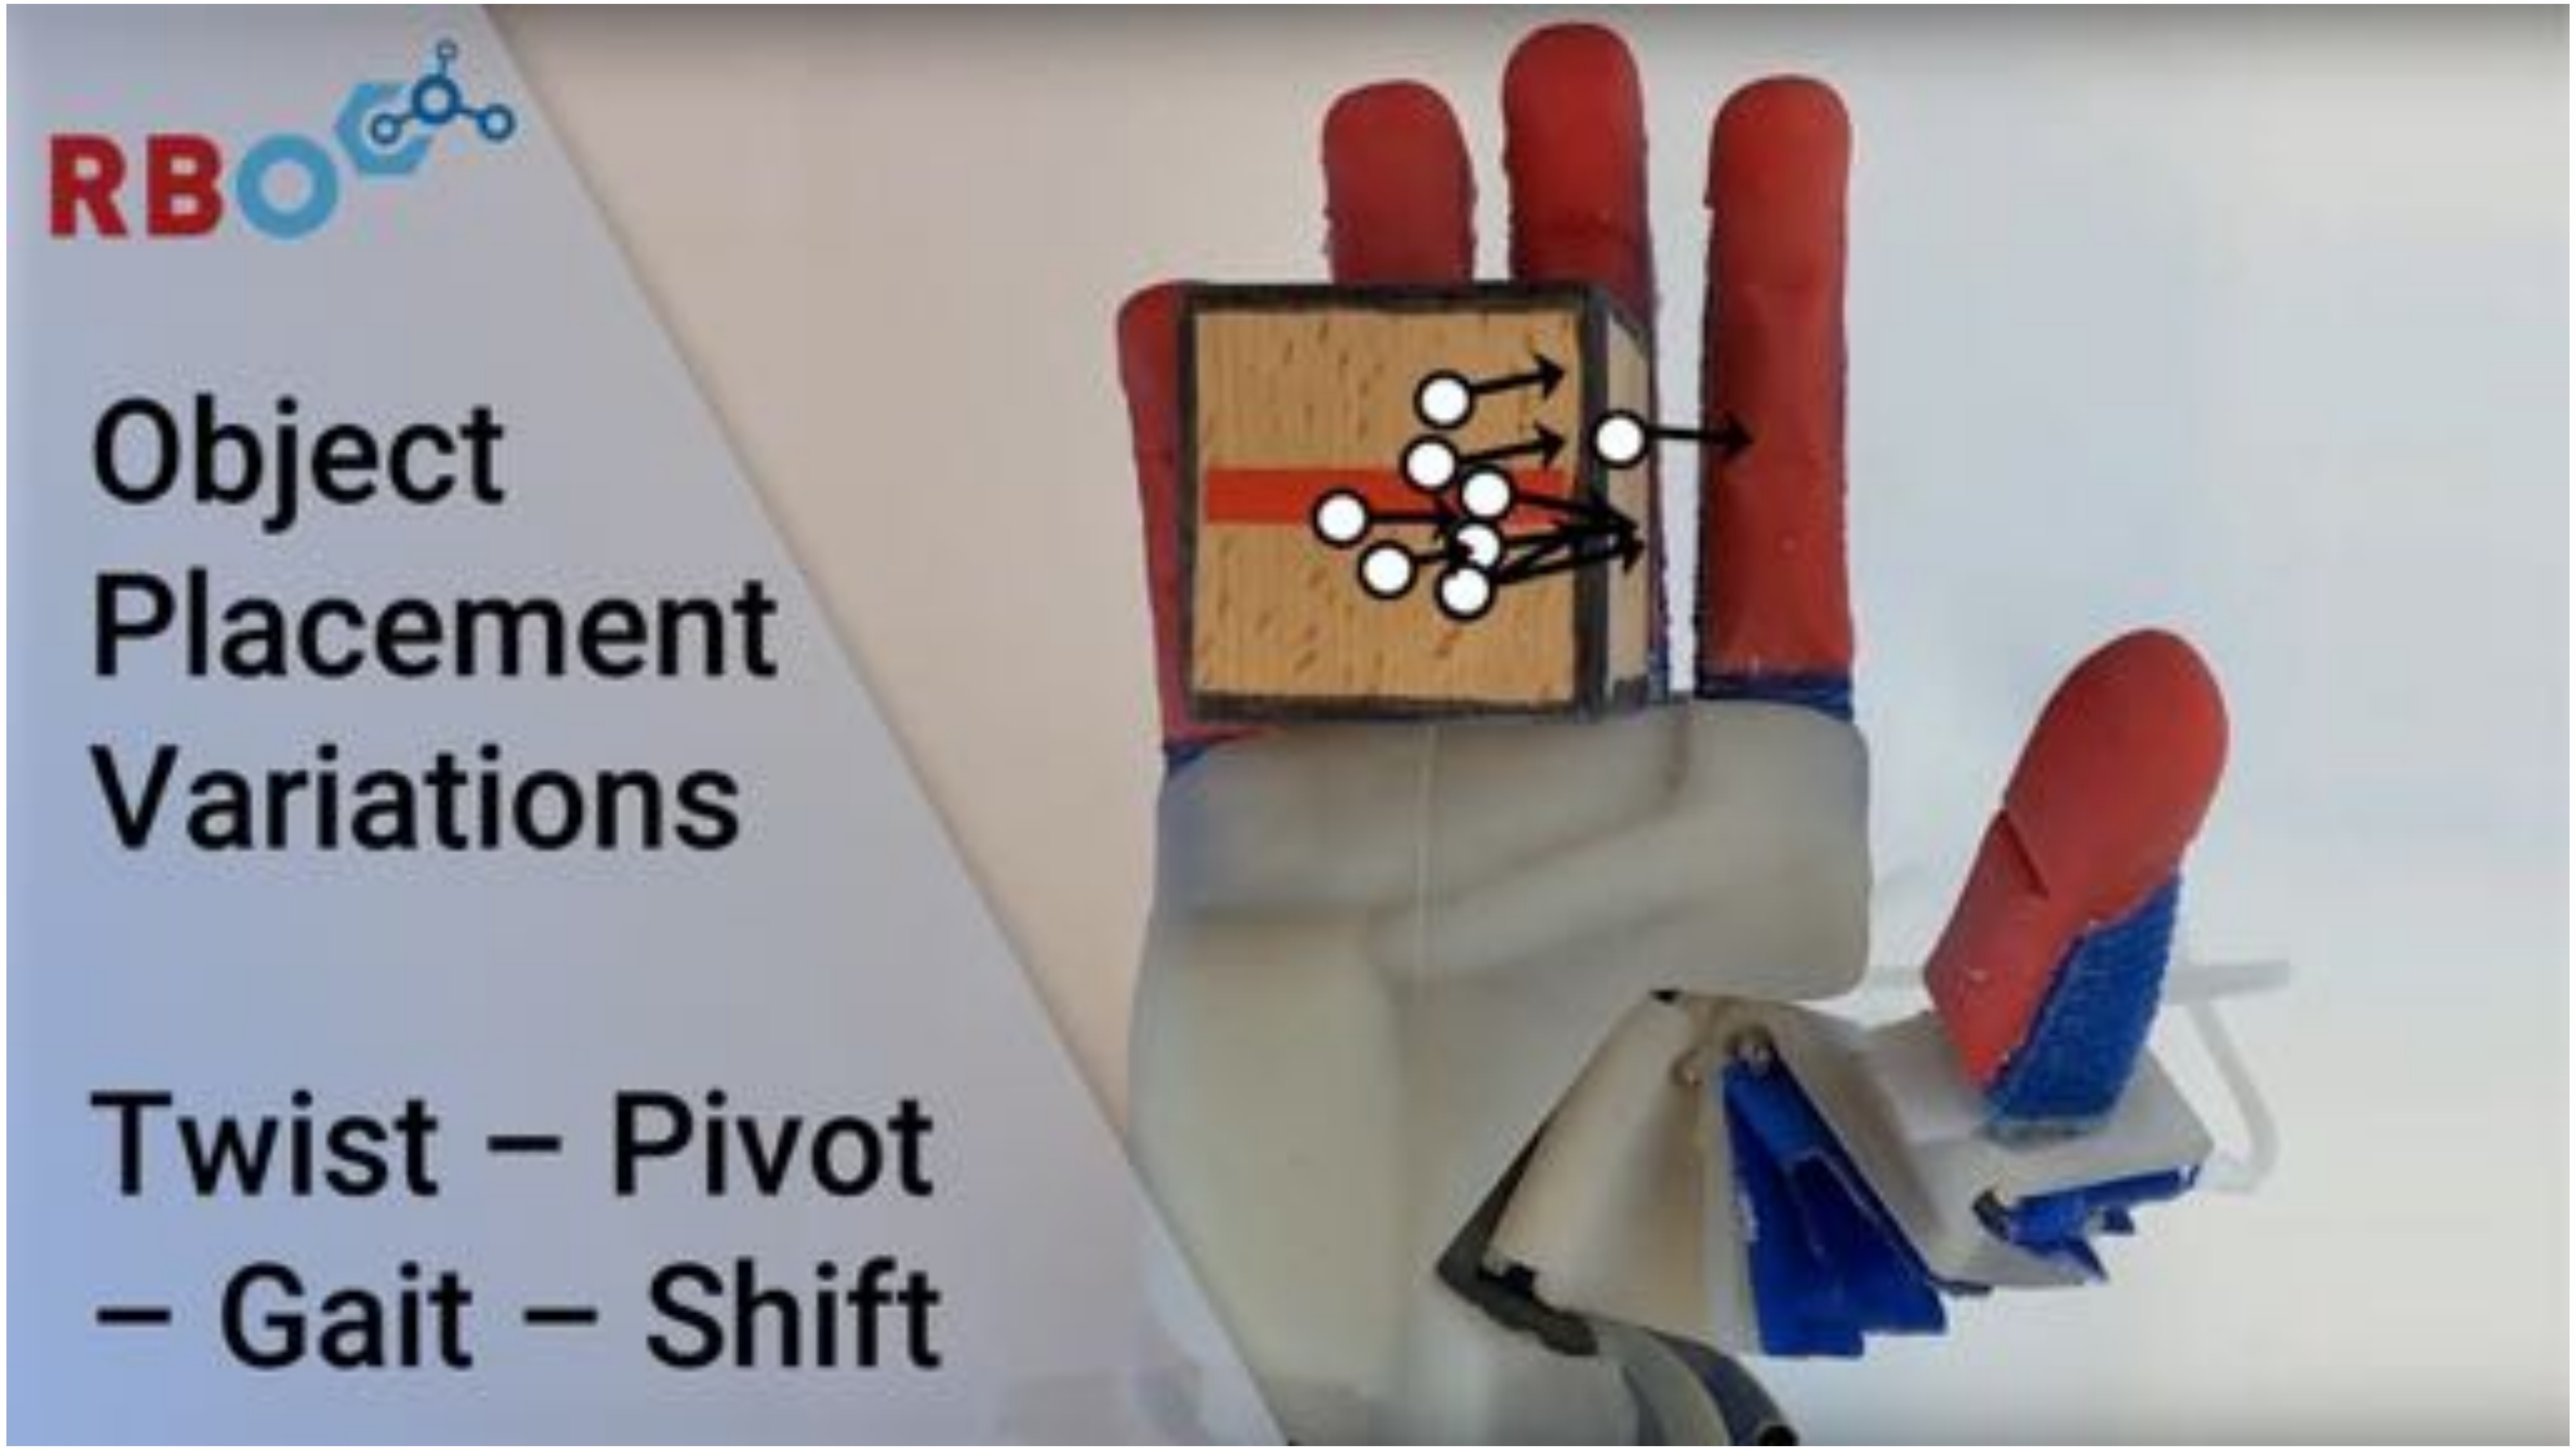

Generalizing to New Objects

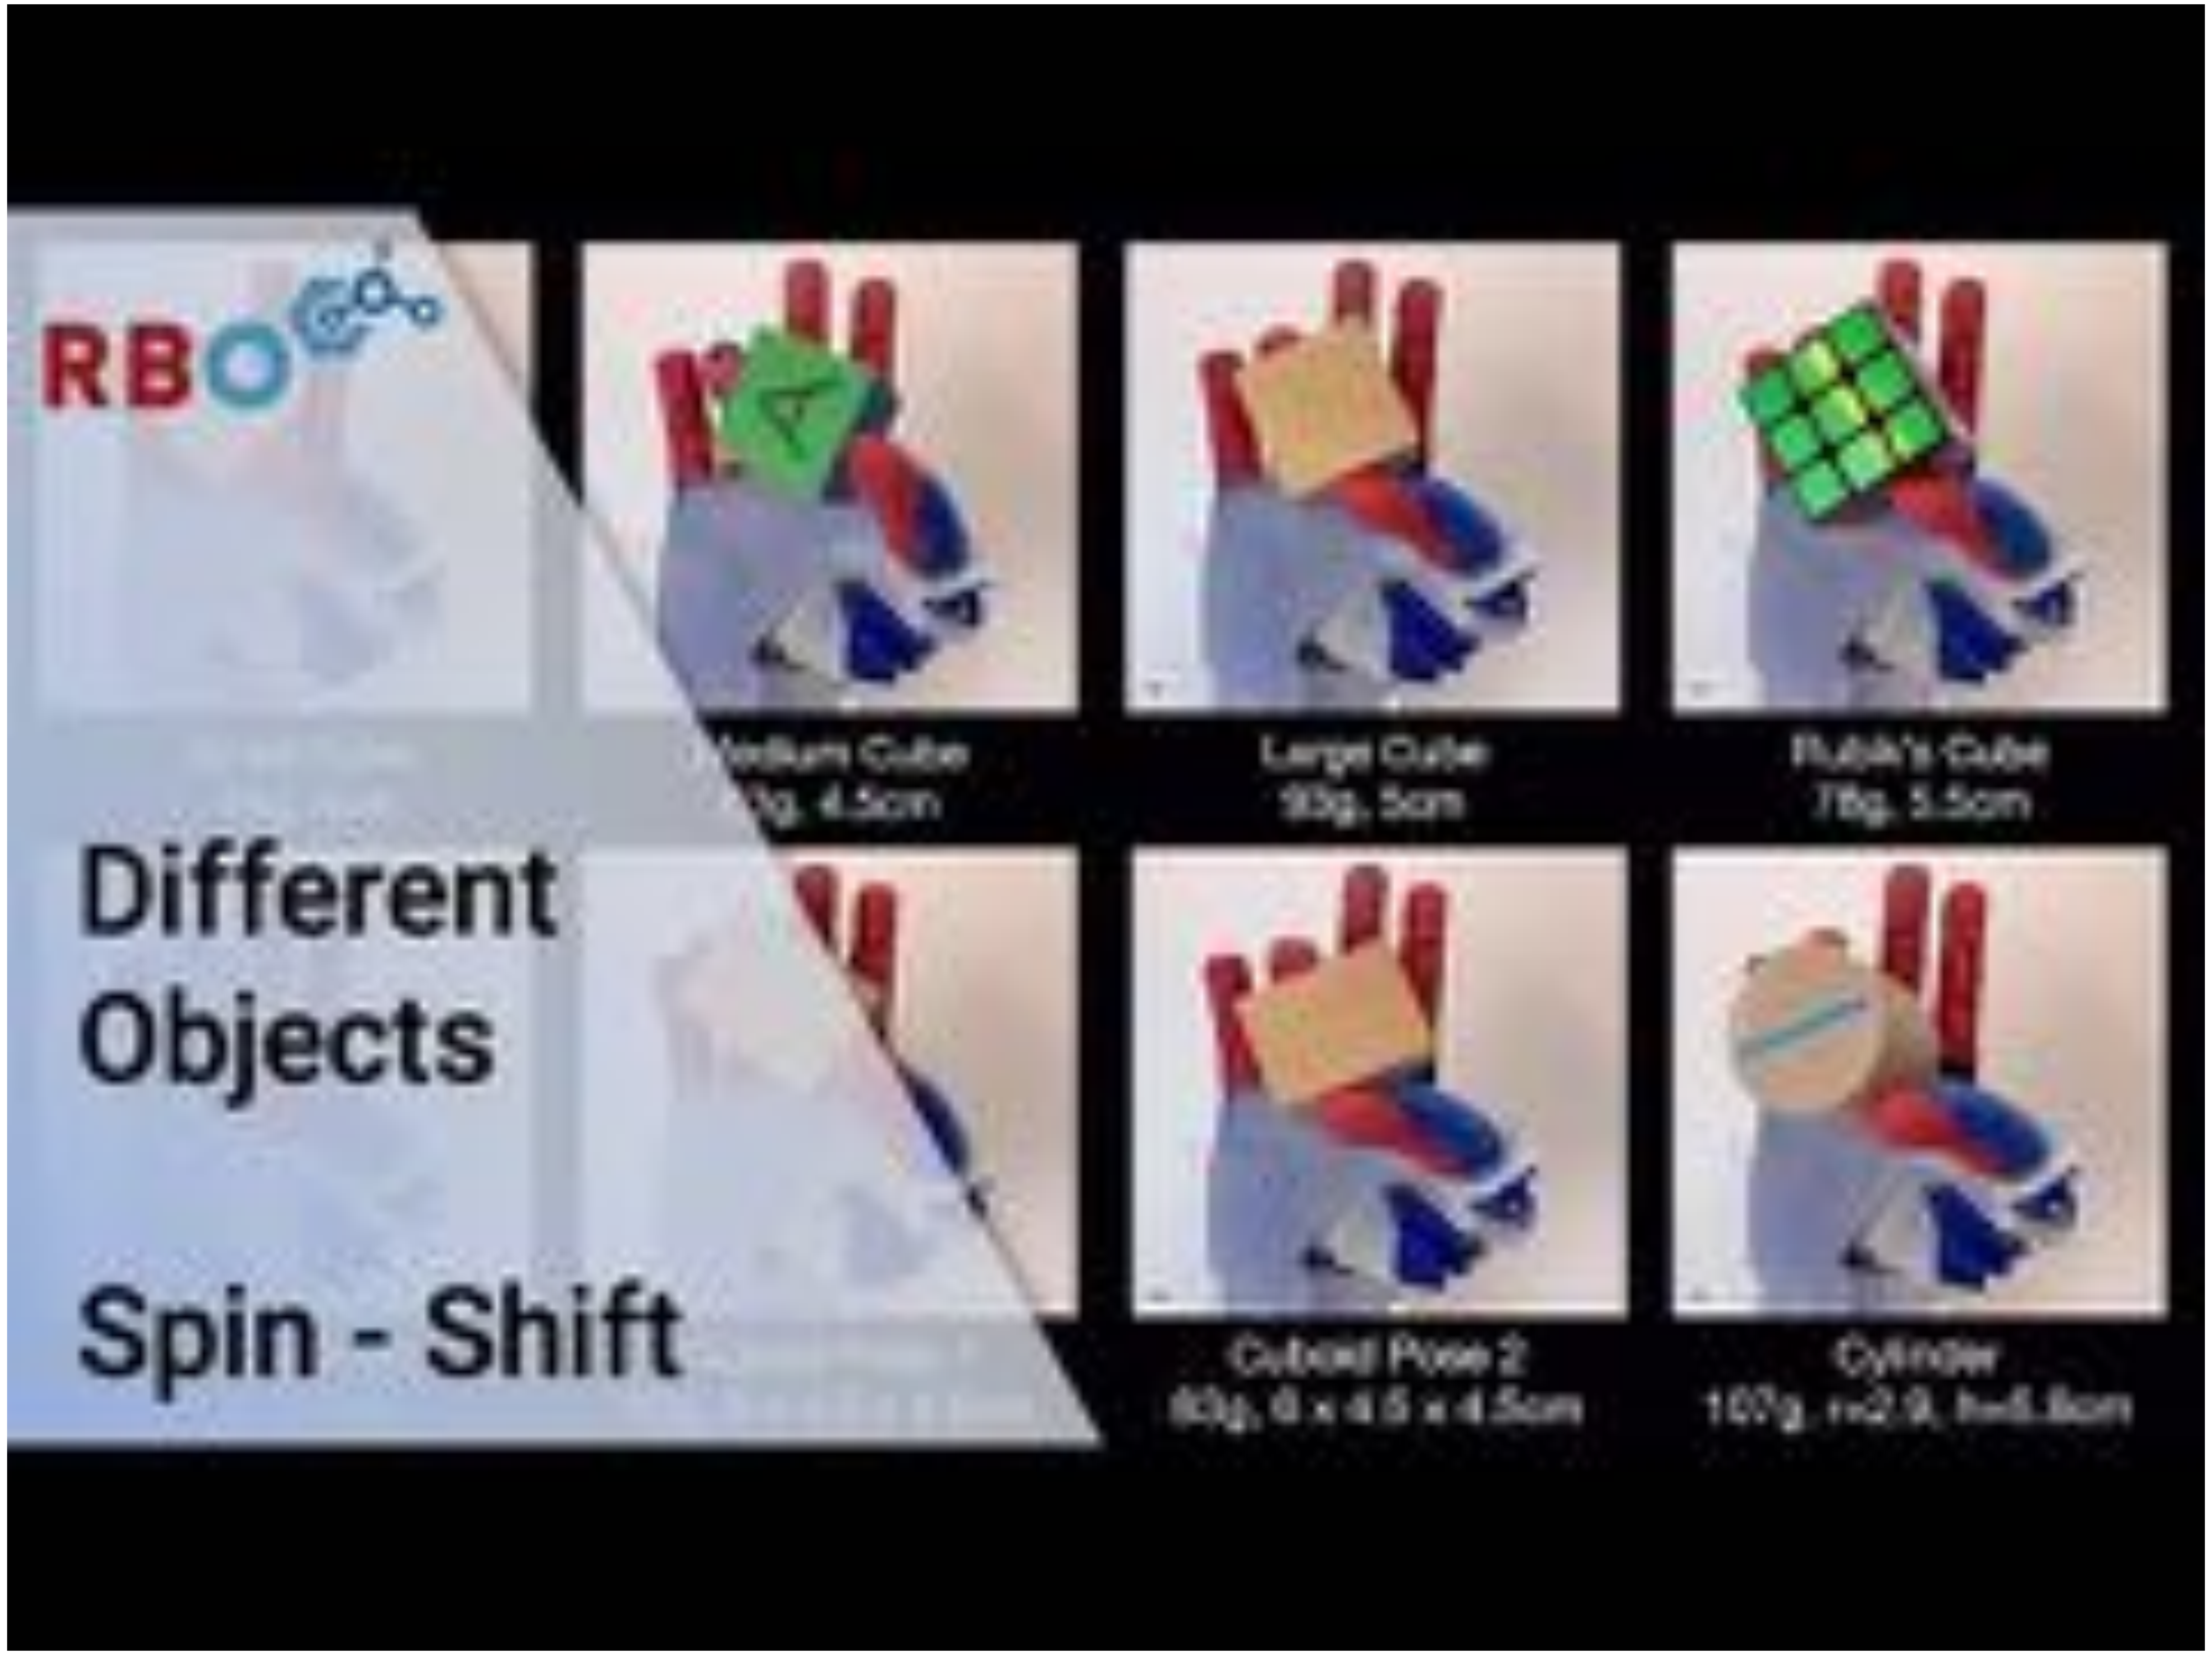

Generalizing to New Objects

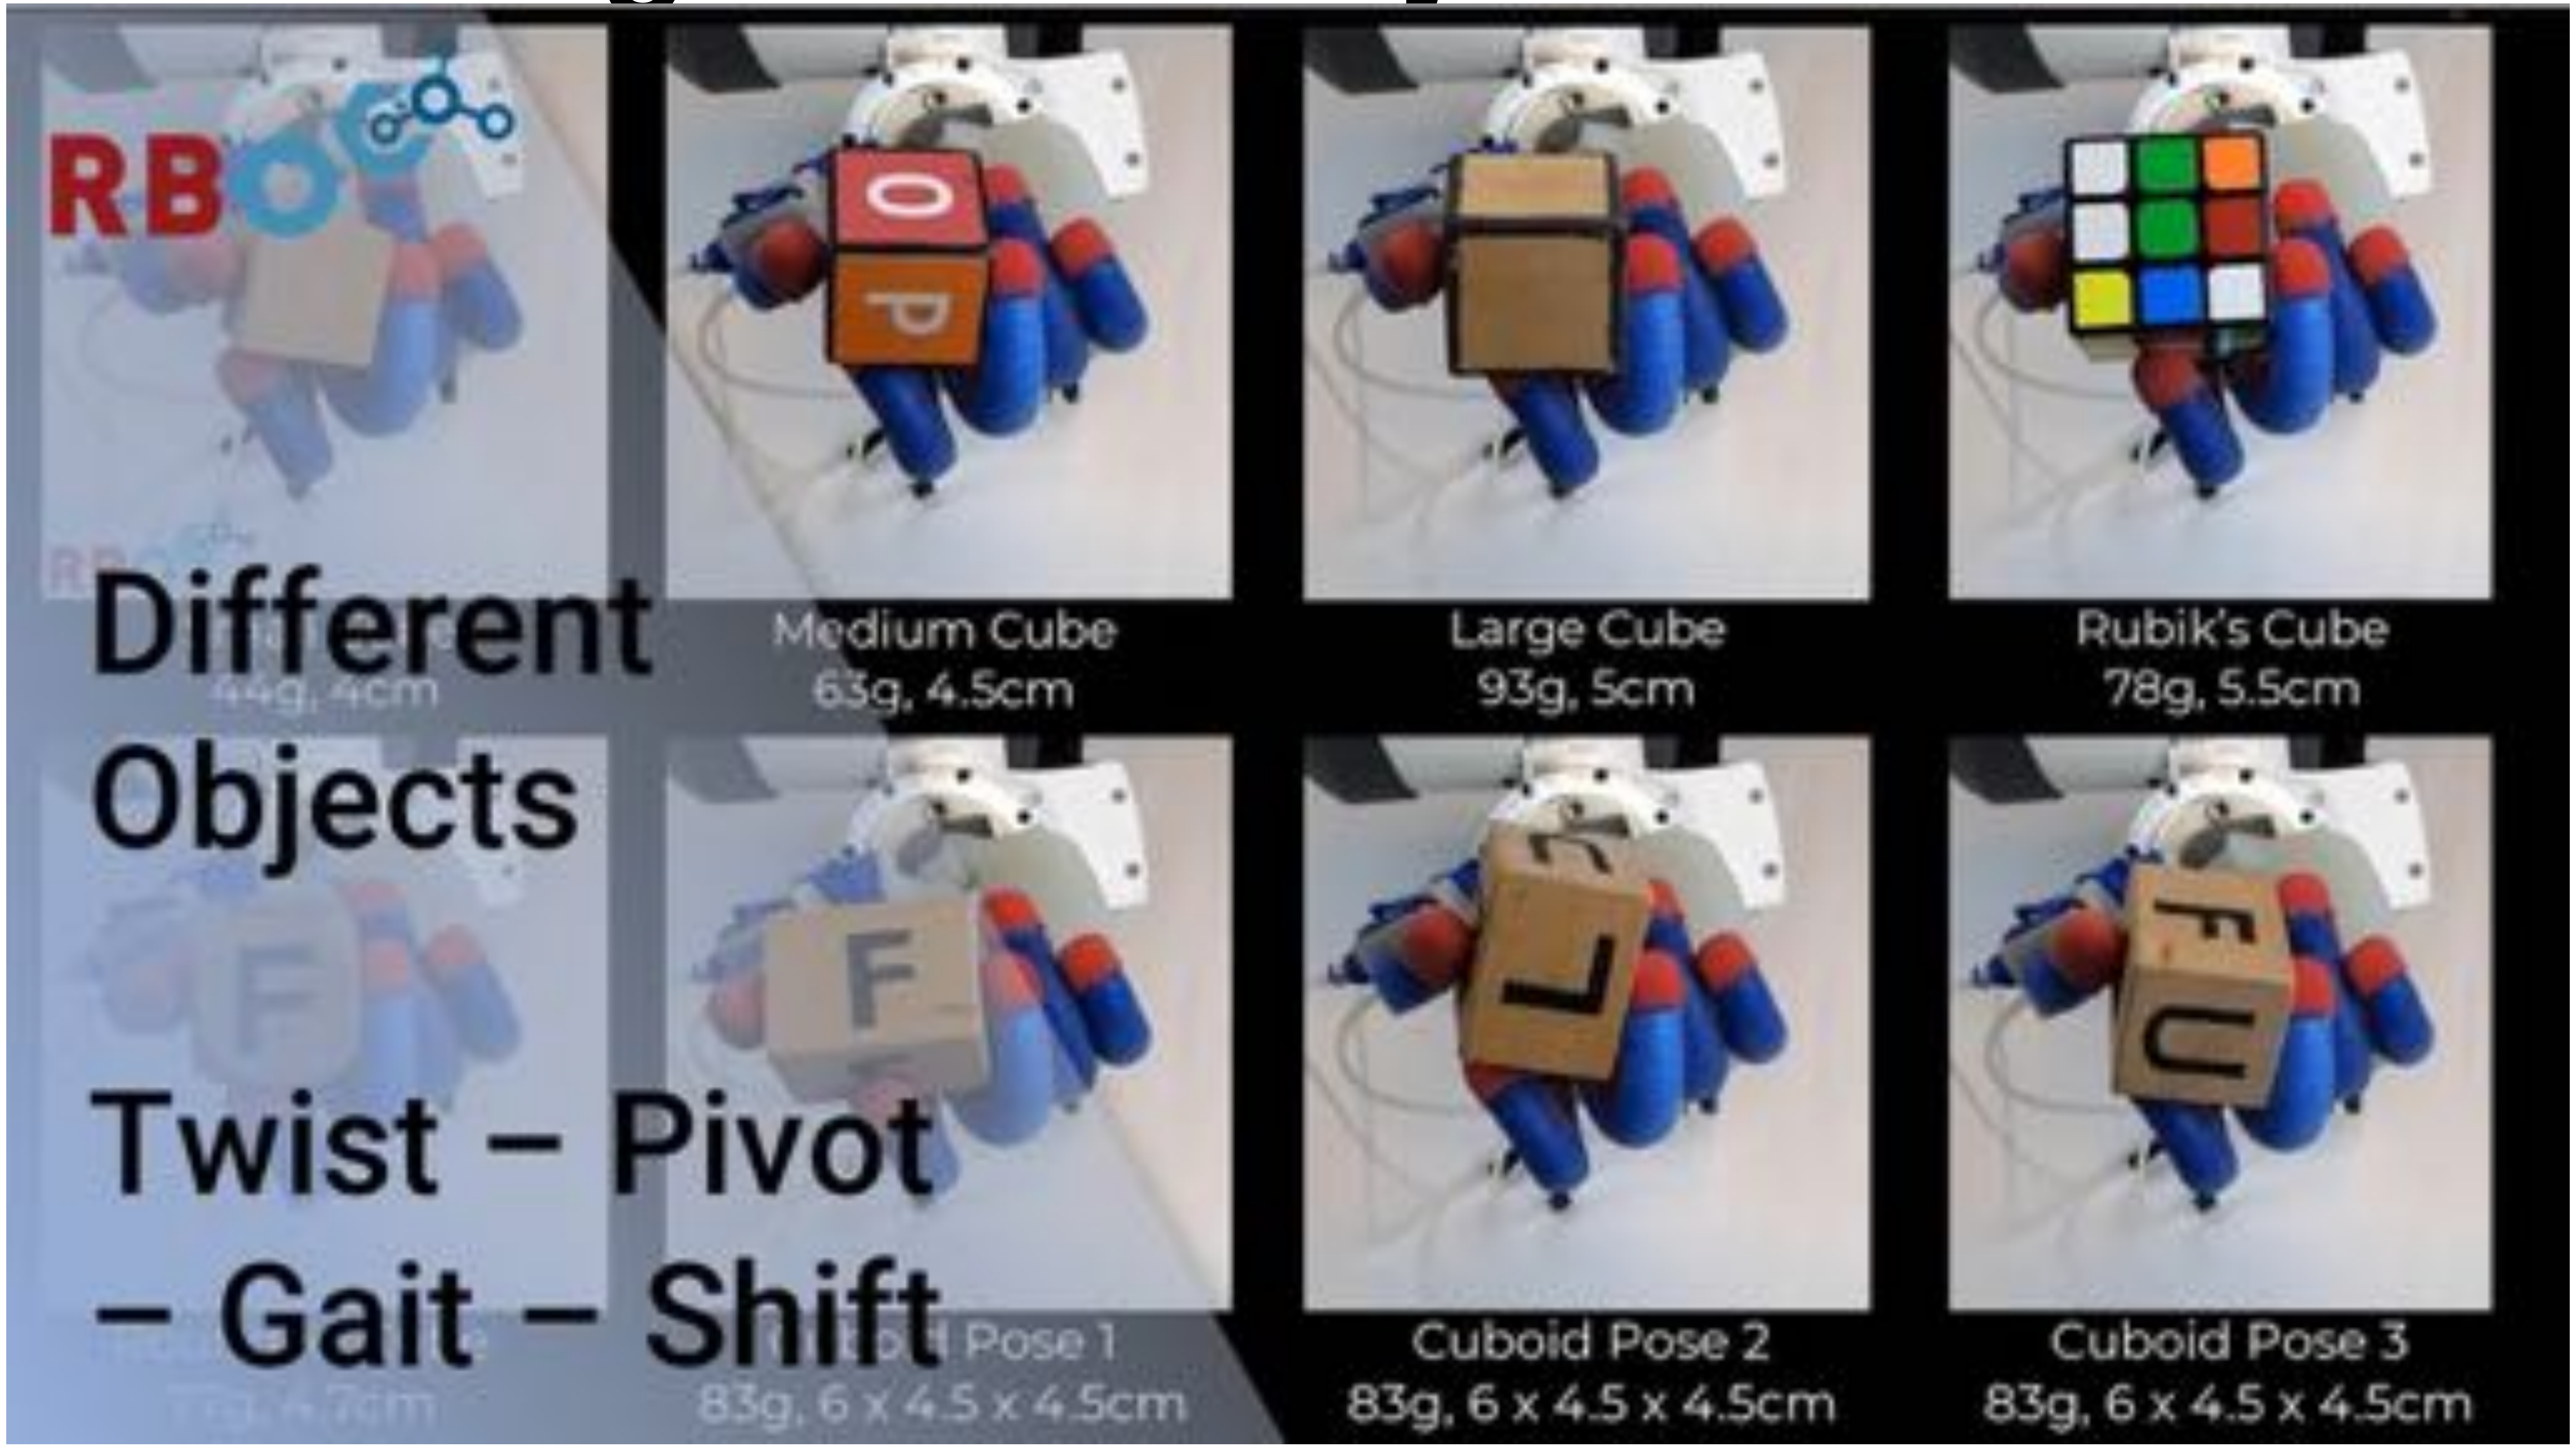

Changing Speed

Repeatability

Failures

Three Principles for Robust Manipulation

1. Morphological Computation

"Exploiting the intrinsic properties of mechanical hardware can also provide a simple, effective and reliable way of dealing with mechanical interaction."

—Neville Hogan

Hogan, N. "Impedance control-An approach to manipulation. Part I-Theory." ASME Journal of Dynamic Systems and Measurement Control (1985)

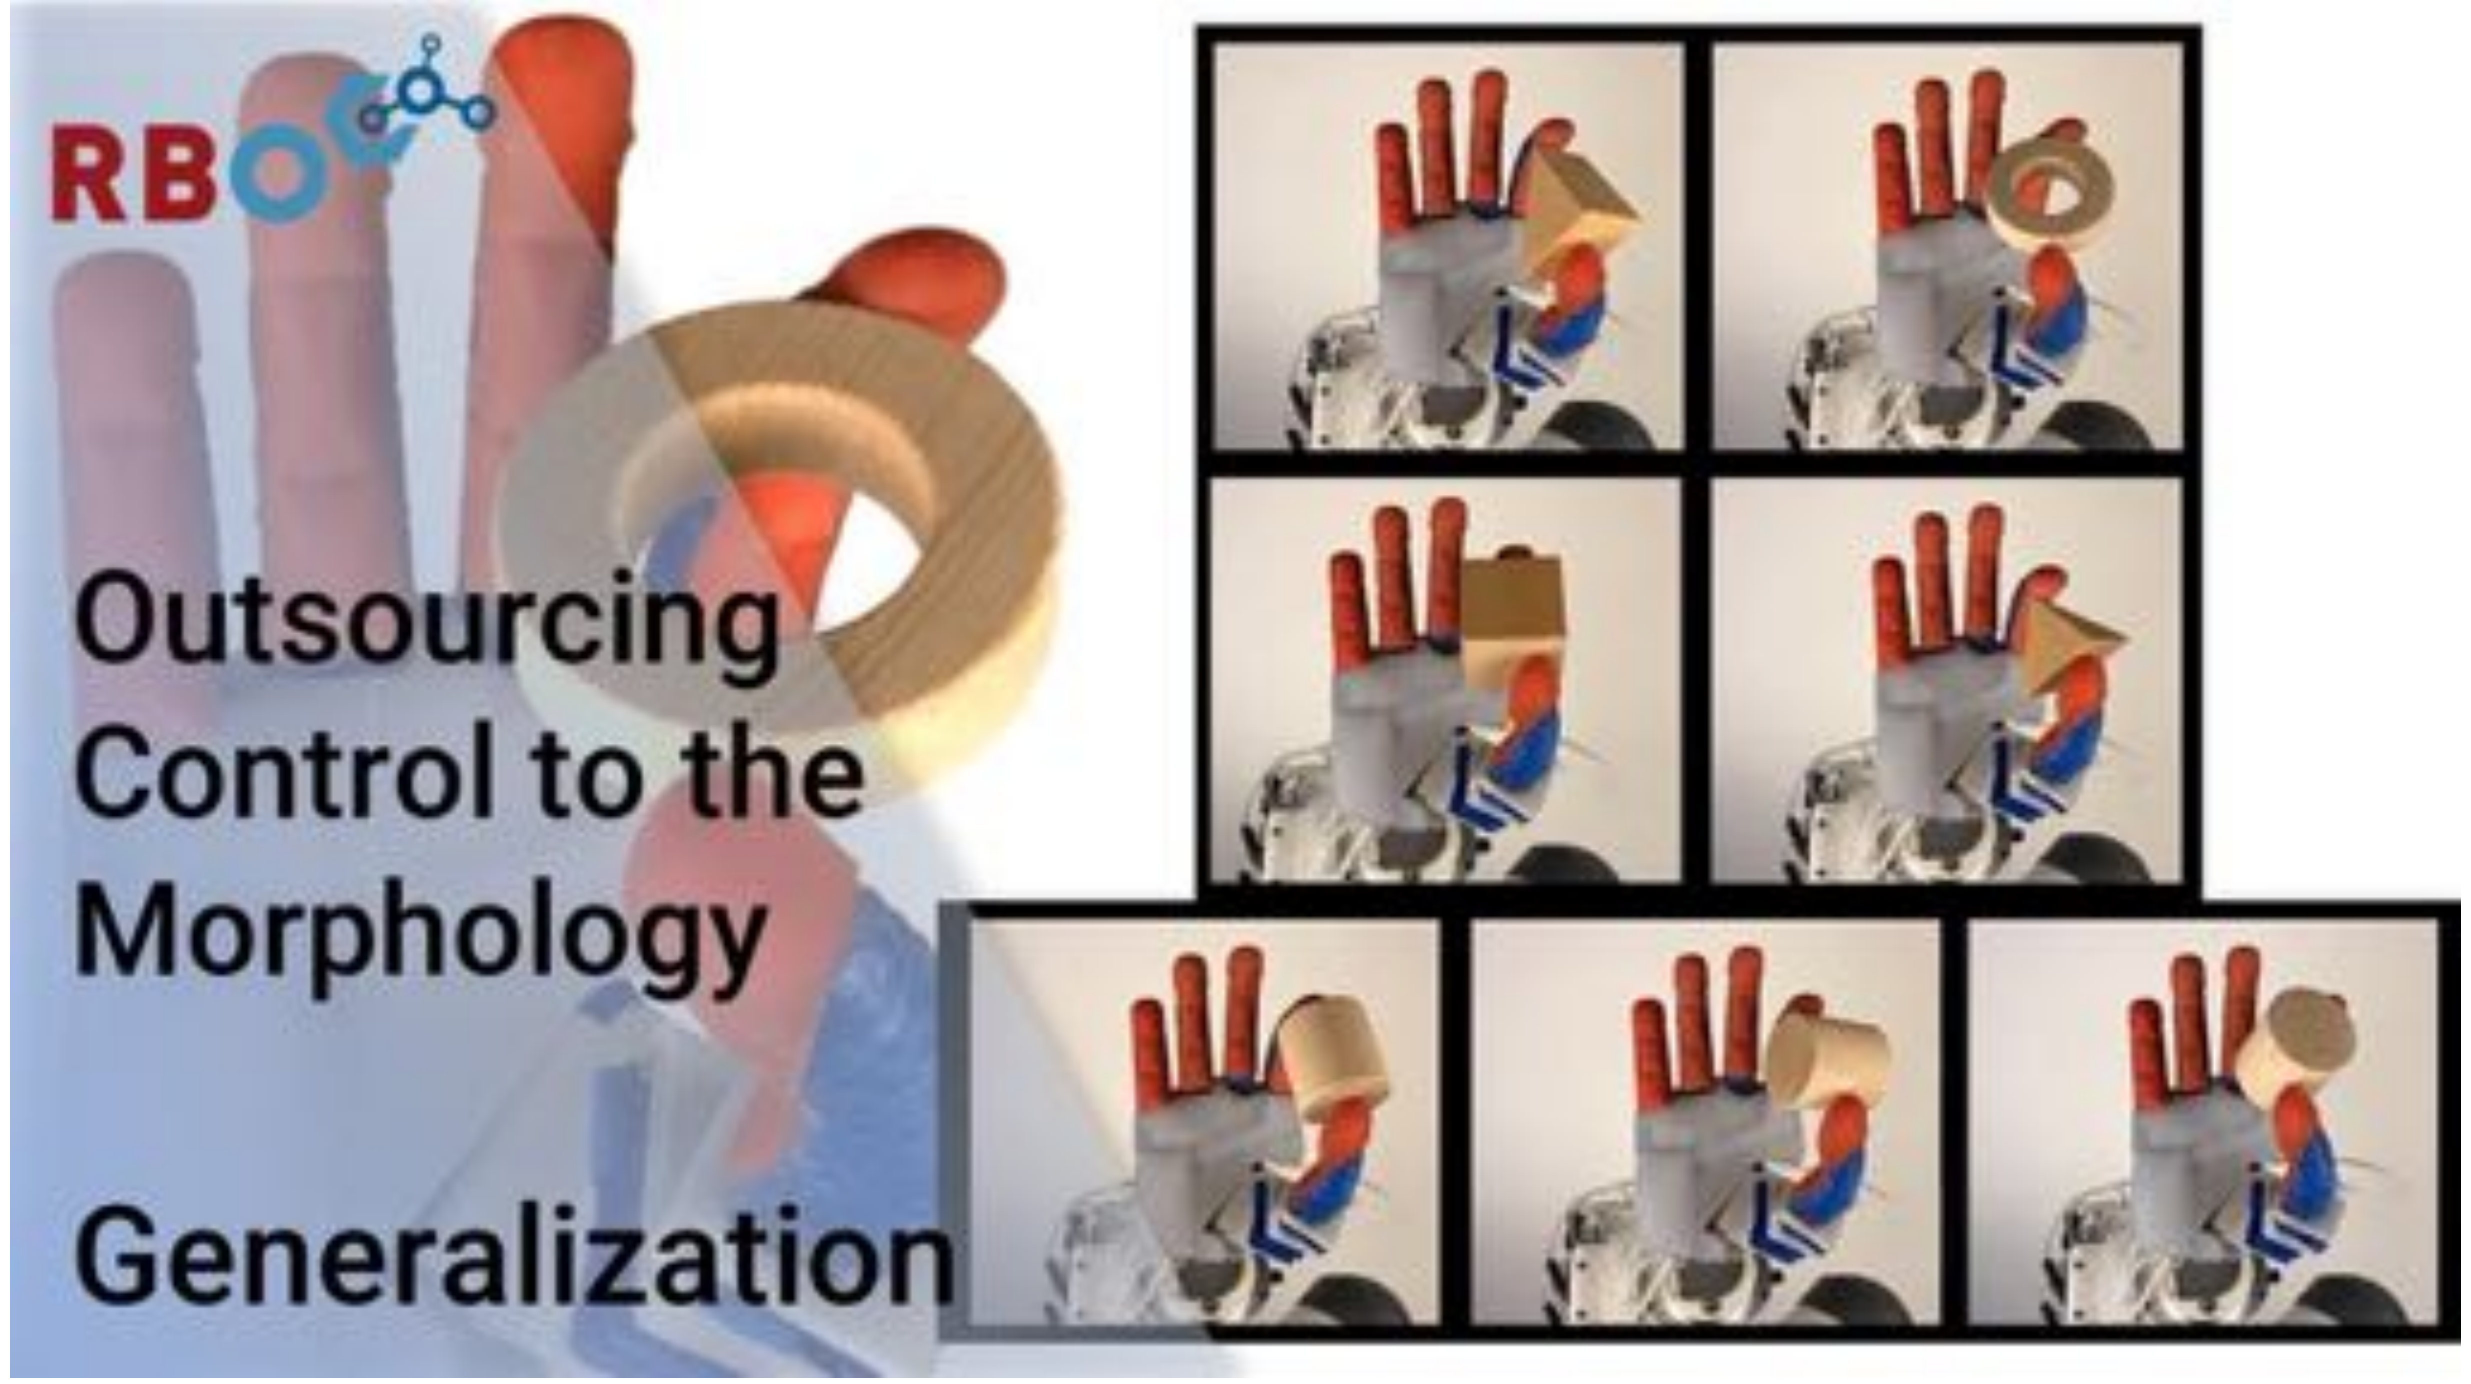

Three Principles for Robust Manipulation

2. Constraining Object Motion

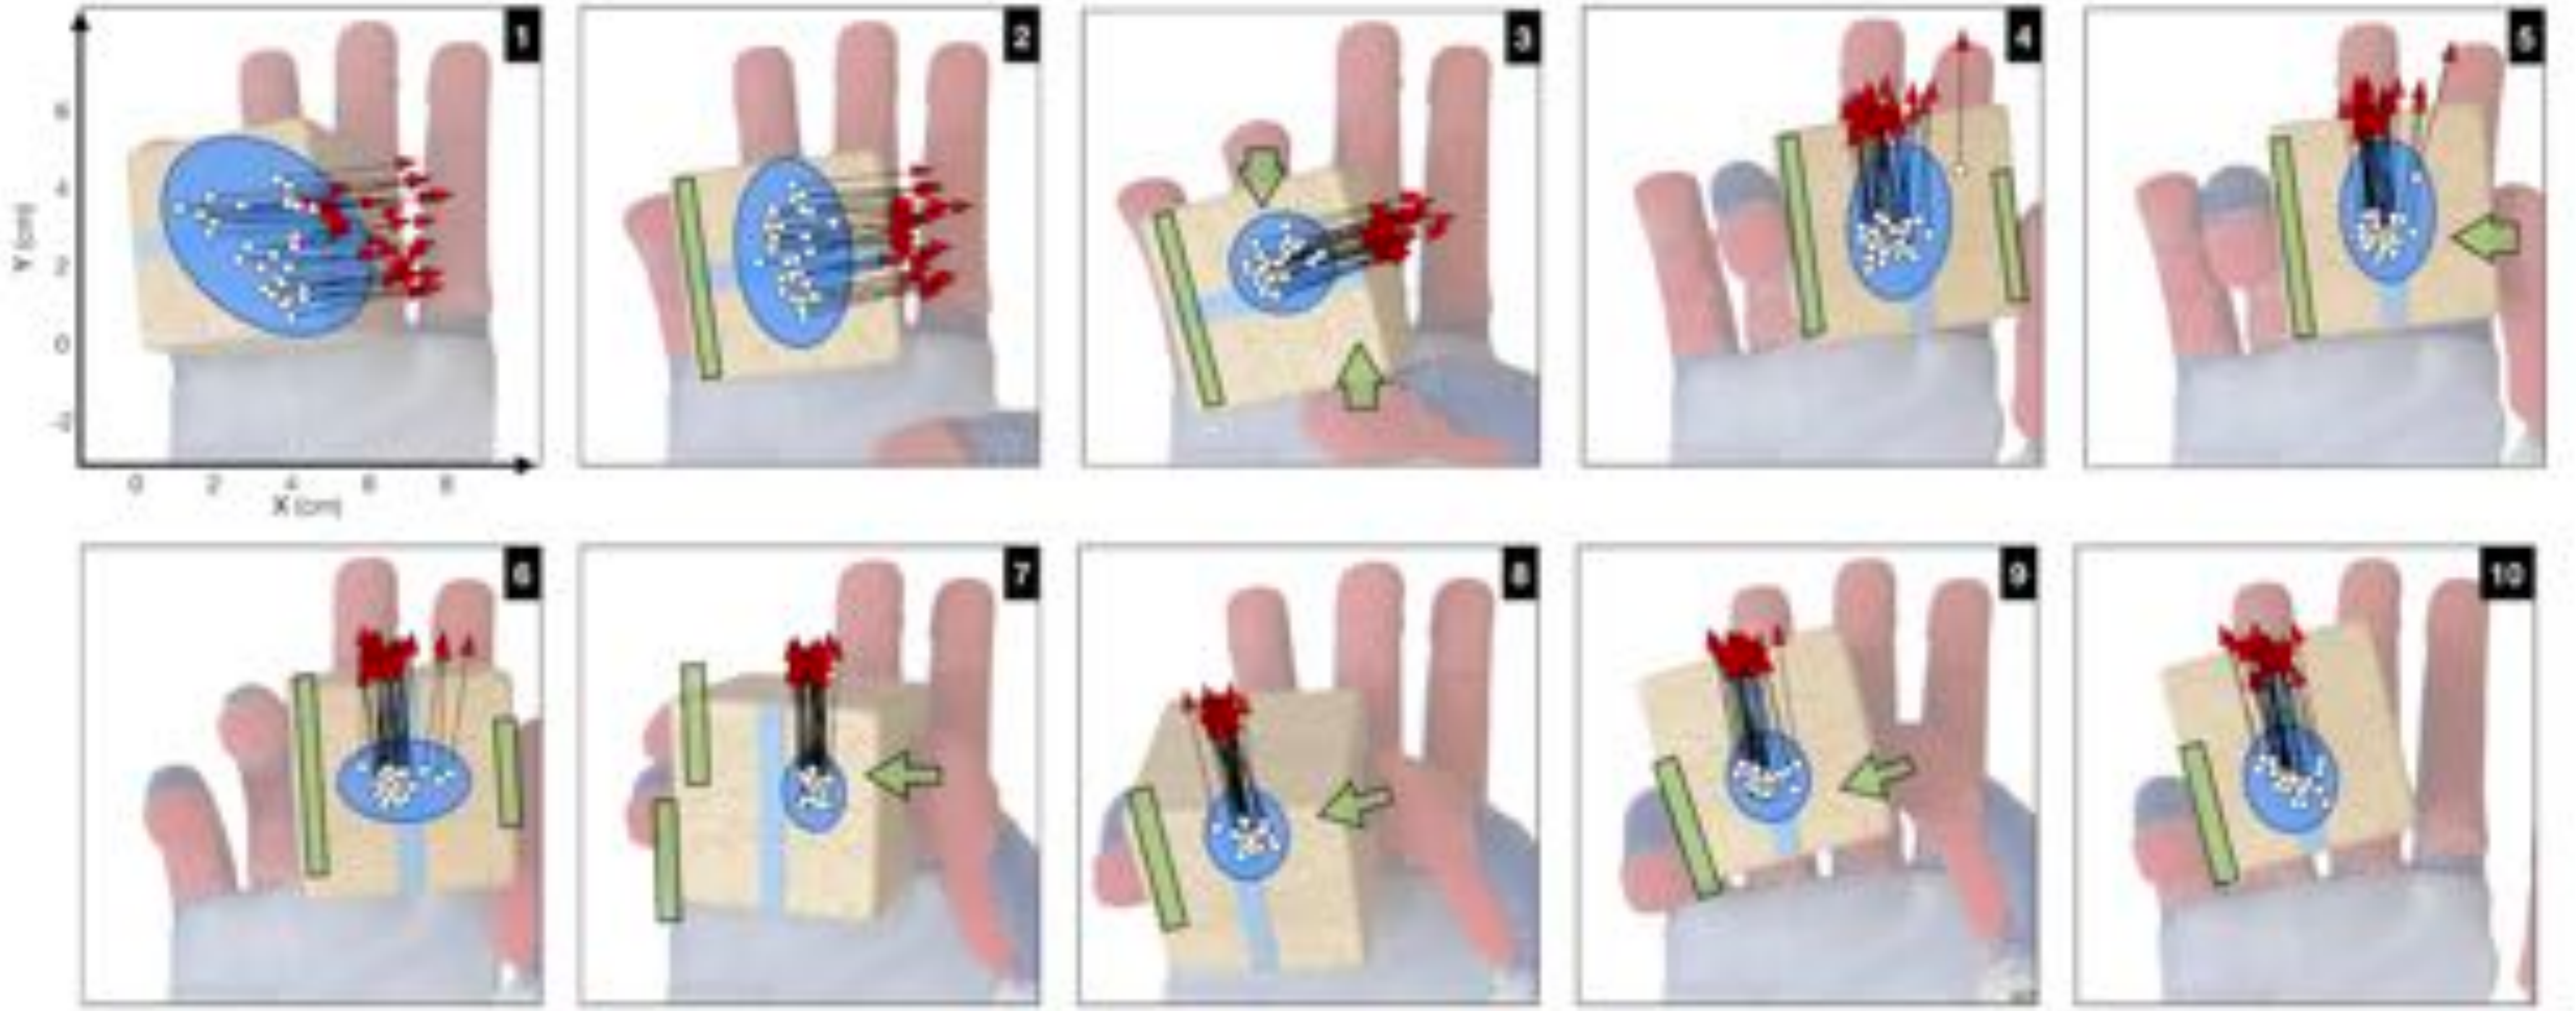

Fig. 8. **Robust manipulation as a series of uncertainty-restricting constraint exploitations:** The plots depict a keyframe-by-keyframe breakdown of cube poses (x, y, θ) gathered over 33 independent trials of *spin + shift*. The underlaid photographs are sourced from only the most salient keyframes of one illustrative execution. A white dot marks the cube position from a single execution, and the associated arrow indicates the planar rotation θ. The blue region coarsely indicates the set of observed cube positions over all trials. The green bars represent physical walls implemented by the fingers, and the green arrows represent active pushing interactions.

Three Principles for Robust Manipulation

2. Constraining Object Motion

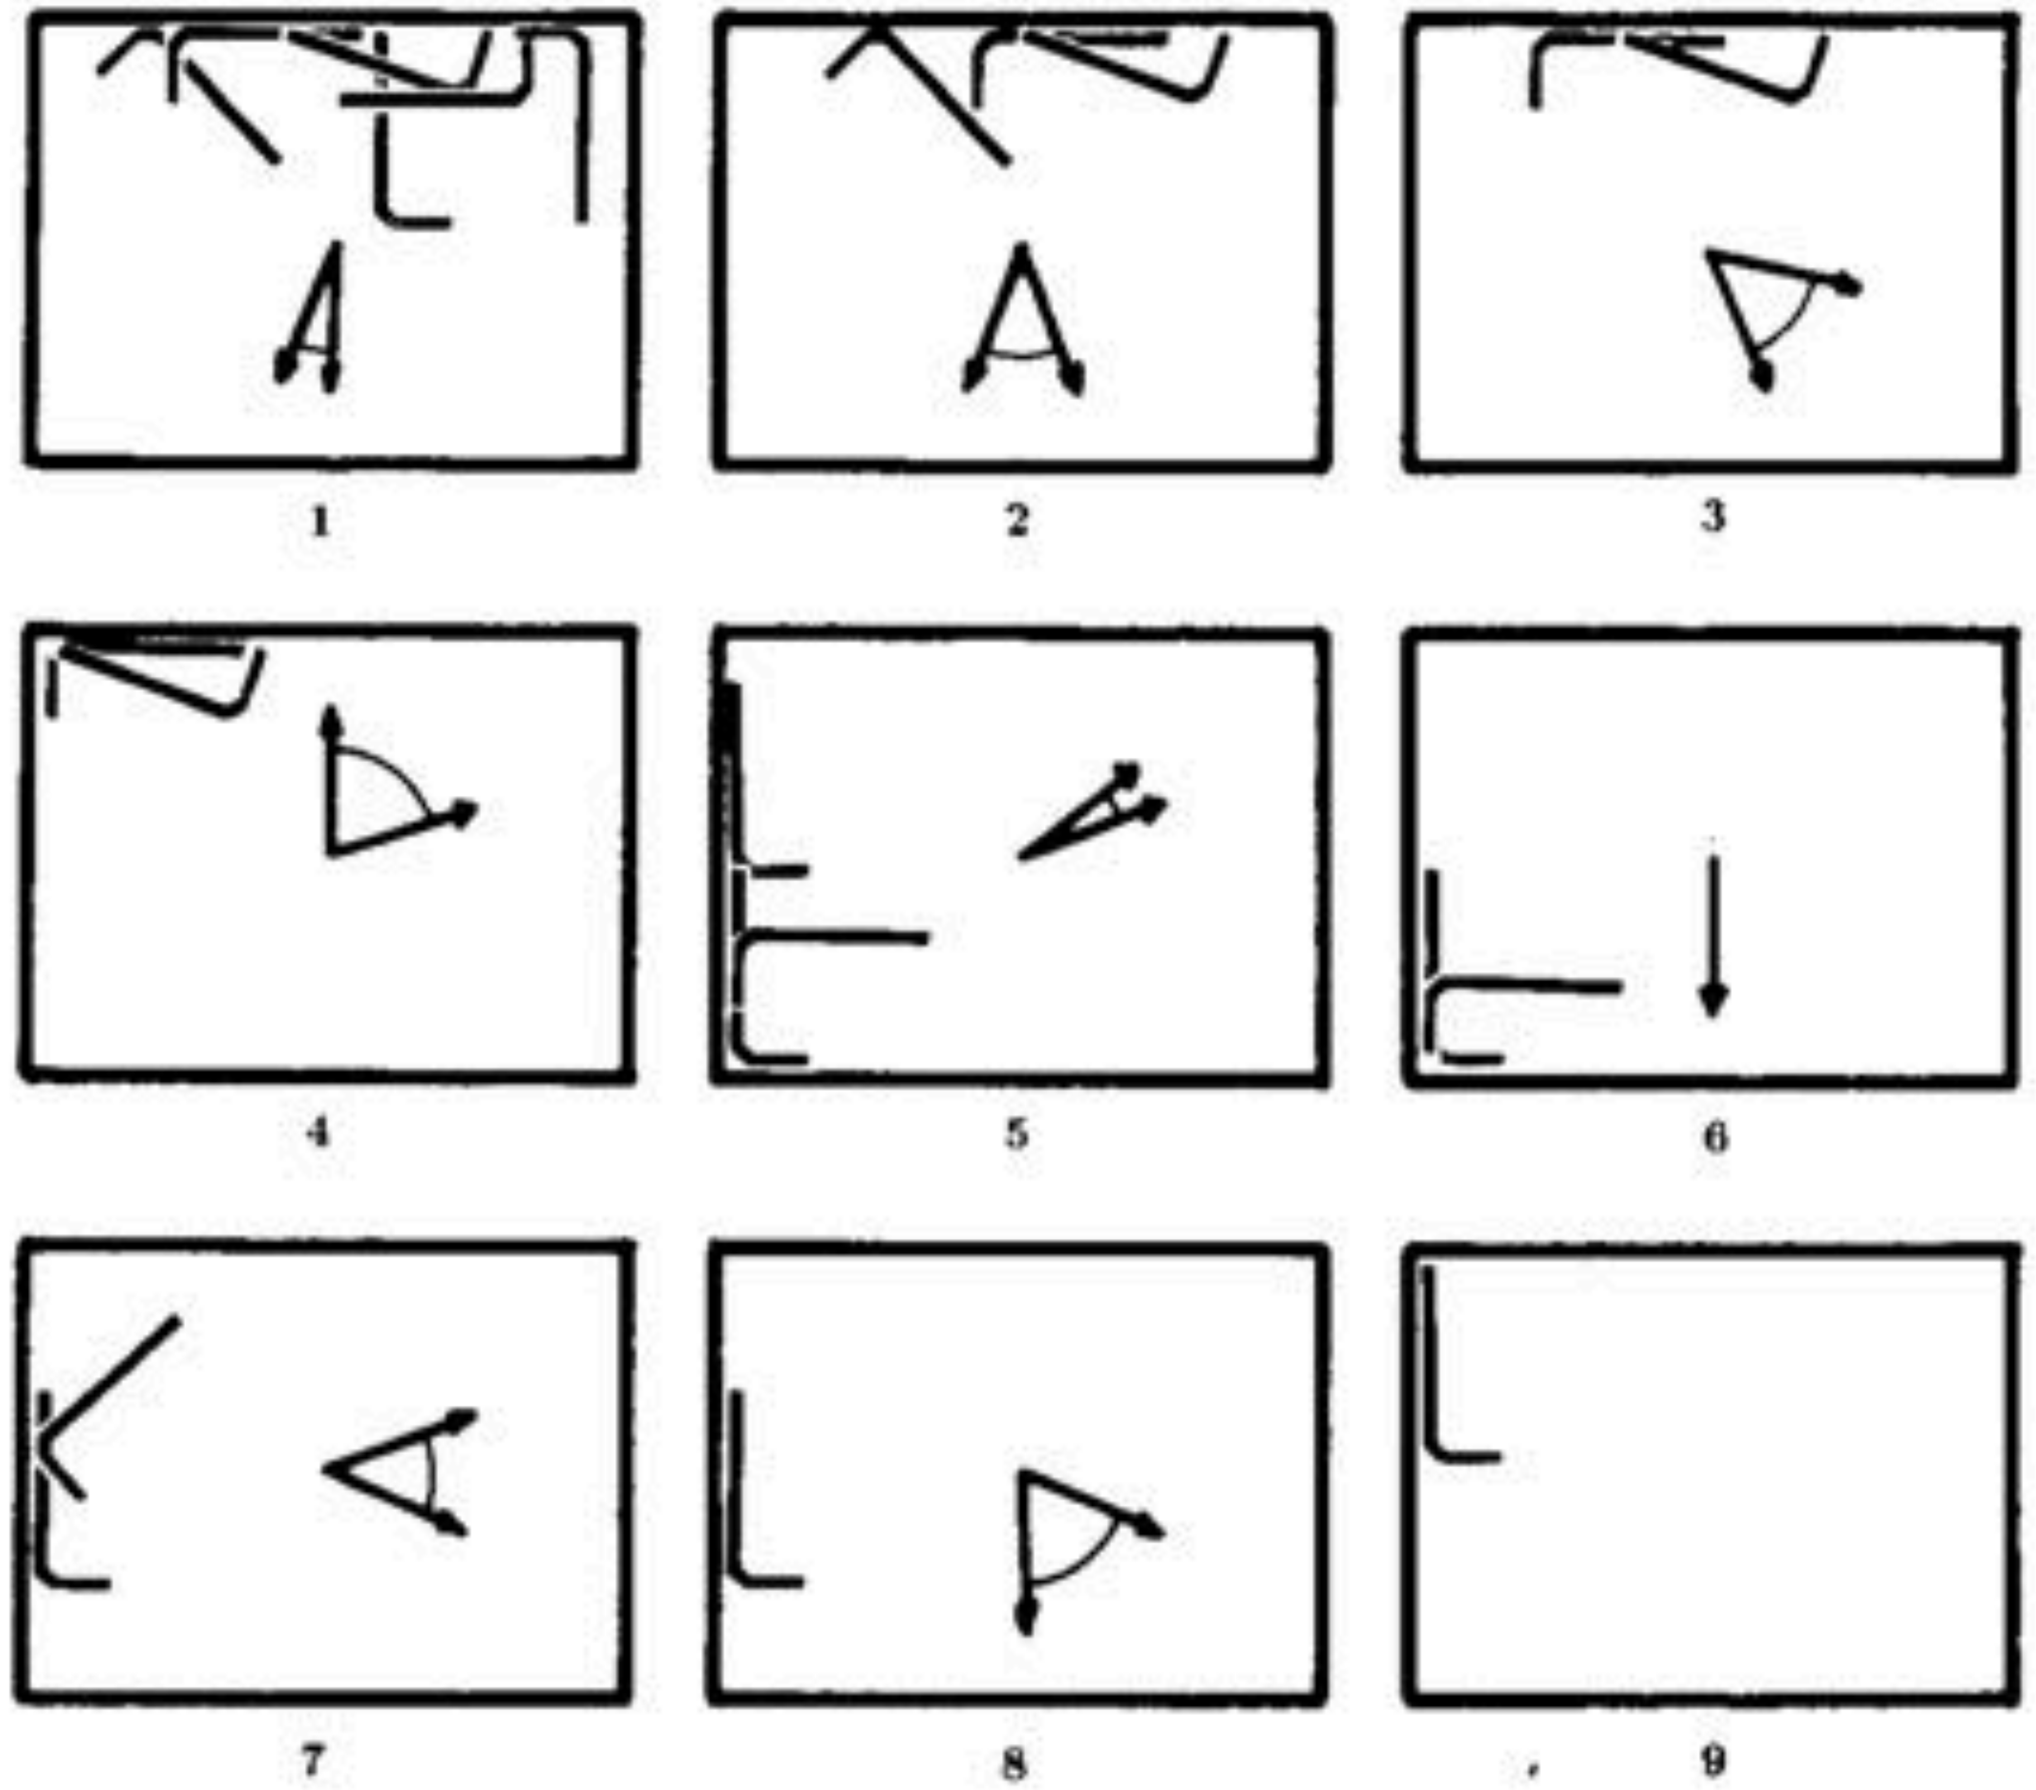

Fig. 2. Beginning at the upper left and moving from left to right, we can trace an automatically generated program that orients the wrench. Each frame shows the set of possible wrench contacts, and the operation to be applied. Each operation is represented by an interval of azimuths. The azimuth arrows indicate the tray's direction of steepest ascent; gravity acts in the opposite direction.

Erdmann, Michael A., and Matthew T. Mason. "An exploration of sensorless manipulation." *IEEE Journal on Robotics and Automation* 4, no. 4 (1988): 369-379.

Three Principles for Robust Manipulation

3. Compositing Manipulation Funnels

“Using a funnel, the goal to position an object can be accomplished despite variation in the initial locations and shapes of the objects.”

—Matt Mason

Mason, Matthew. "The mechanics of manipulation." In Proceedings. 1985 IEEE International Conference on Robotics and Automation, 1985.

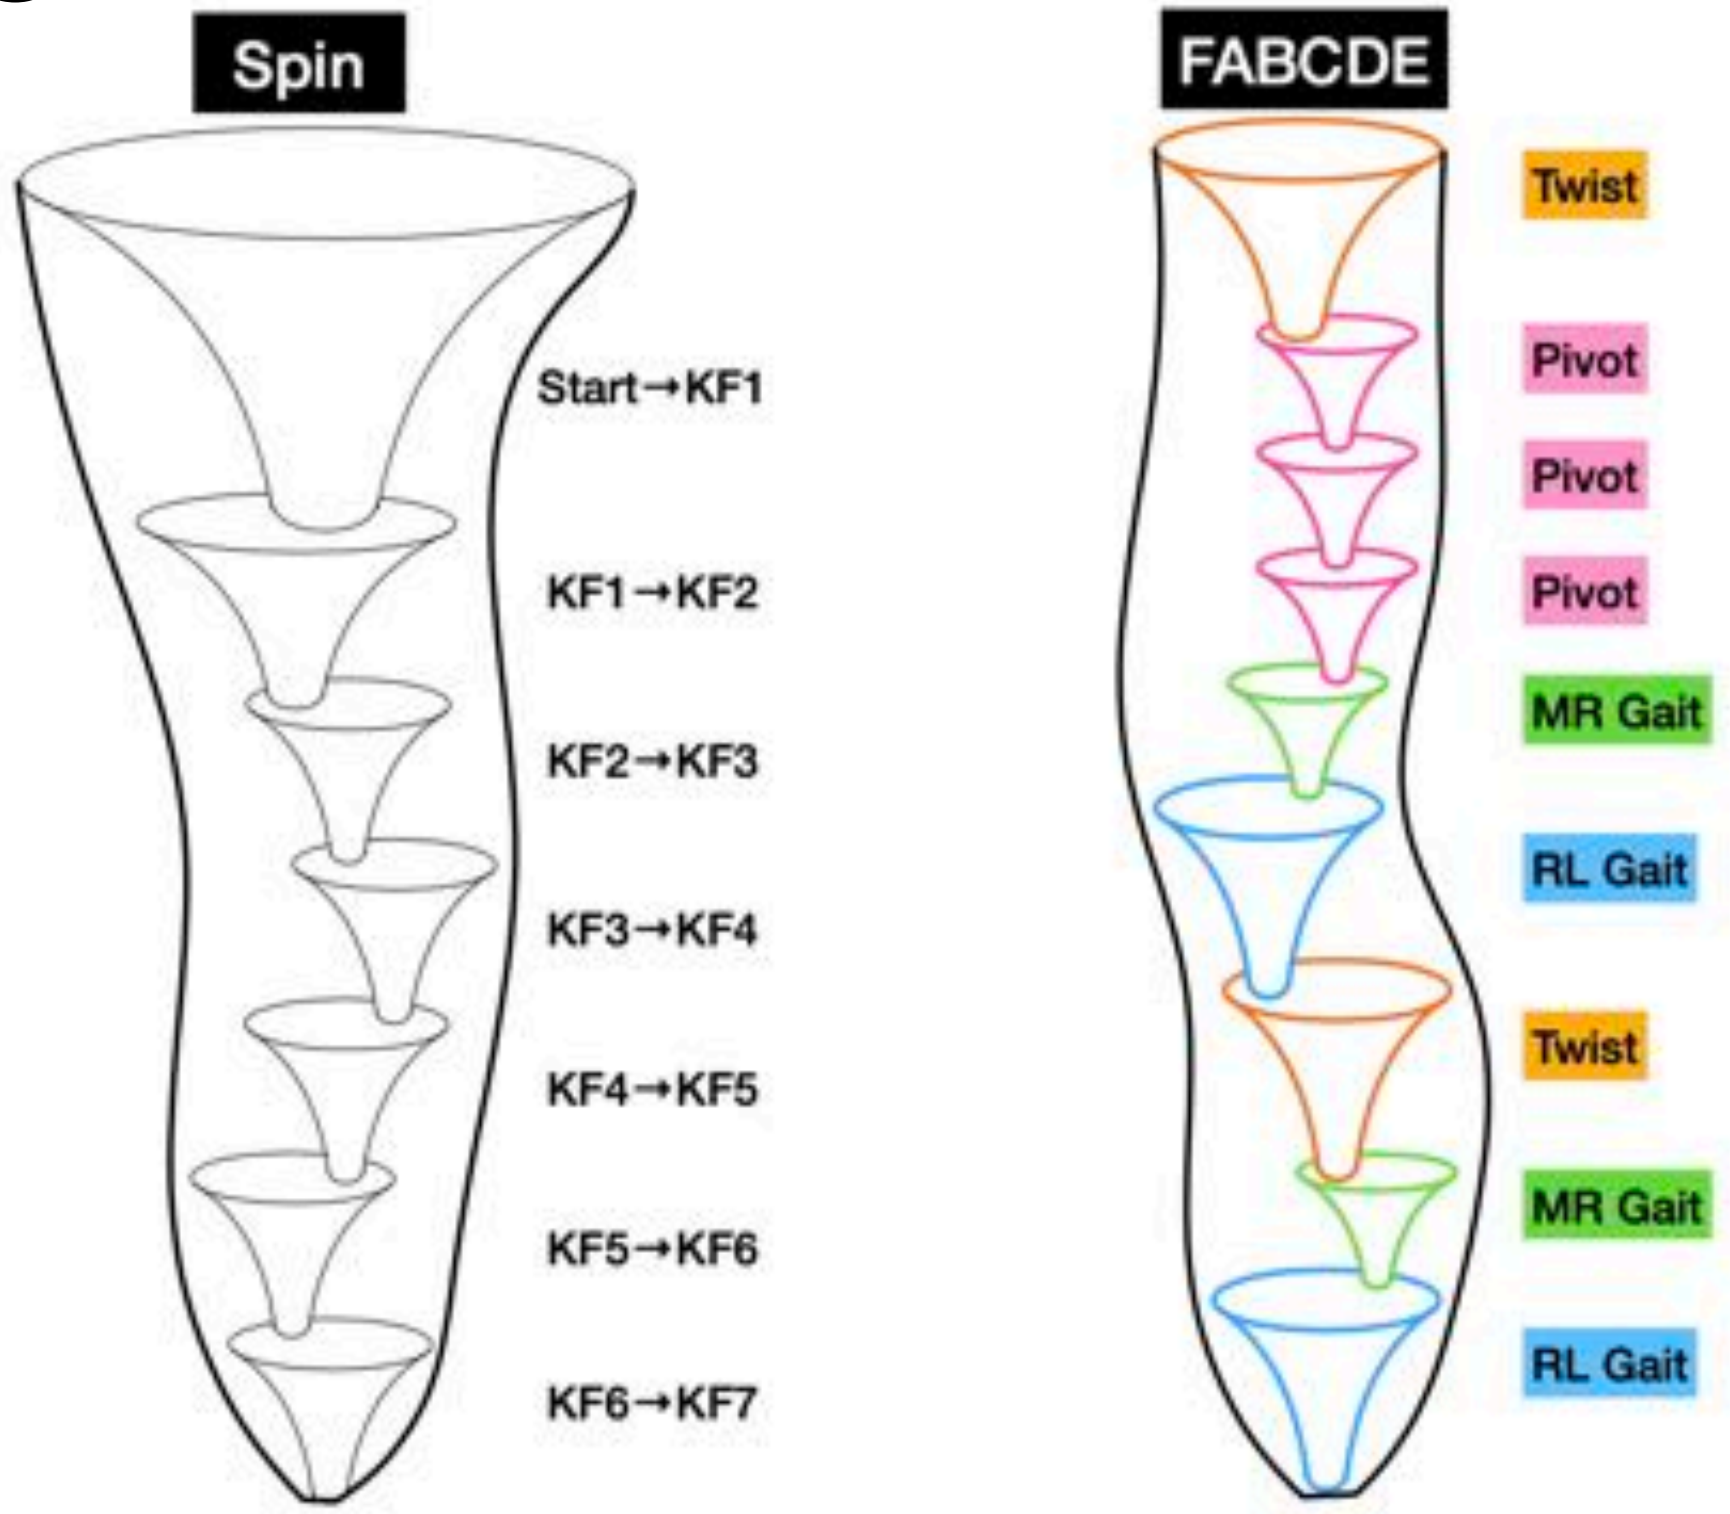

Fig. 10. **Robust manipulation through funnel composition:** Like *spin*, each skill is composed of funnel-like robust primitive manipulations, so it is a robust funnel in itself. Each funnel reliably transforms a set of hand-object configurations (its entrance) into another set (its exit). We designed our skills to funnel into each other, letting us compose longer manipulation plans like FABCDE.

Three Principles for Robust Manipulation

3. Compositing Manipulation Funnels

"Using a funnel, the goal to position an object can be accomplished despite variation in the initial locations and shapes of the objects."

—Matt Mason

Mason, Matthew. "The mechanics of manipulation." In Proceedings. 1985 IEEE International Conference on Robotics and Automation, 1985.

Other Insights

- Hand morphology and control strategies should be designed together
- Good results rely on compliance and large contact surface areas
- OpenAI may have attempted something much harder than what was required
    - rigid hand
    - physical constraints not explicitly exploited
    - compositionality using funnels not explicitly considered
- Could more effective learning techniques incorporate these ideas?