DeepRob

[Group 7] Lecture 7

Dual-arm Manipulation - Learning

by Ryan Roche, Matt Rajala, Adit Kadepurkar

University of Minnesota

What is Bimanual Manipulation?

What is Bimanual Manipulation?

- **The term "Bimanual Manipulation" originates from psychological studies on motor skills**
- **Refers specifically to tasks requiring the use and coordination of both hands acting on an object**
- **Used in developmental psychology studies of infants and their motor skill development**
- **Later used in robotics after robotic bimanual manipulators were developed**

"Role-differentiated bimanual manipulation (RDBM) is a complementary movement of both hands that requires differentiation between actions of the hands."

Kimmerle, Marliese et al. "Development of role-differentiated bimanual manipulation during the infant's first year." *Developmental psychobiology* vol. 52,2 (2010): 168-80. doi:10.1002/dev.20428

What is Bimanual Manipulation?

- **The term "Bimanual Manipulation" originates from psychological studies on motor skills**
- **Refers specifically to tasks requiring the use and coordination of both hands acting on an object**
- **Used in developmental psychology studies of infants and their motor skill development**
- **Later used in robotics after robotic bimanual manipulators were developed**

"Role-differentiated bimanual manipulation (RDBM) is a complementary movement of both hands that requires differentiation between actions of the hands."

Kimmerle, Marliese et al. "Development of role-differentiated bimanual manipulation during the infant's first year." *Developmental psychobiology* vol. 52,2 (2010): 168-80. doi:10.1002/dev.20428

"Behavioral studies provide evidence that bimanual tasks are more than the simple sum of unimanual tasks as they have to consider spatial and temporal coordination as well as the interactions between both hands."

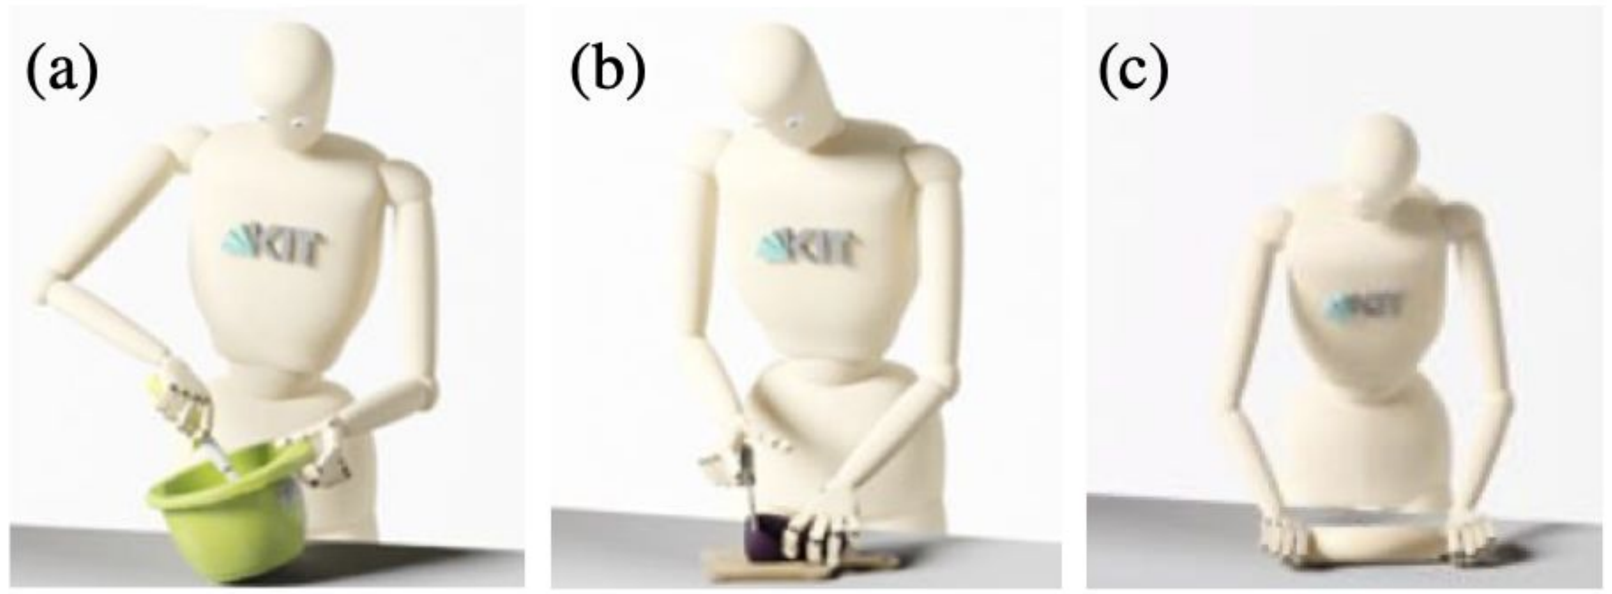

Fig. 1. Examples of bimanual actions: Asymmetric such as stir (a) and cut (b), and symmetric such as rolling (c).

Quote and figure from F. Krebs and T. Asfour, "A Bimanual Manipulation Taxonomy," in *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11031-11038, Oct. 2022, doi: 10.1109/LRA.2022.3196158

So you're talking about Humanoid Robots, then?

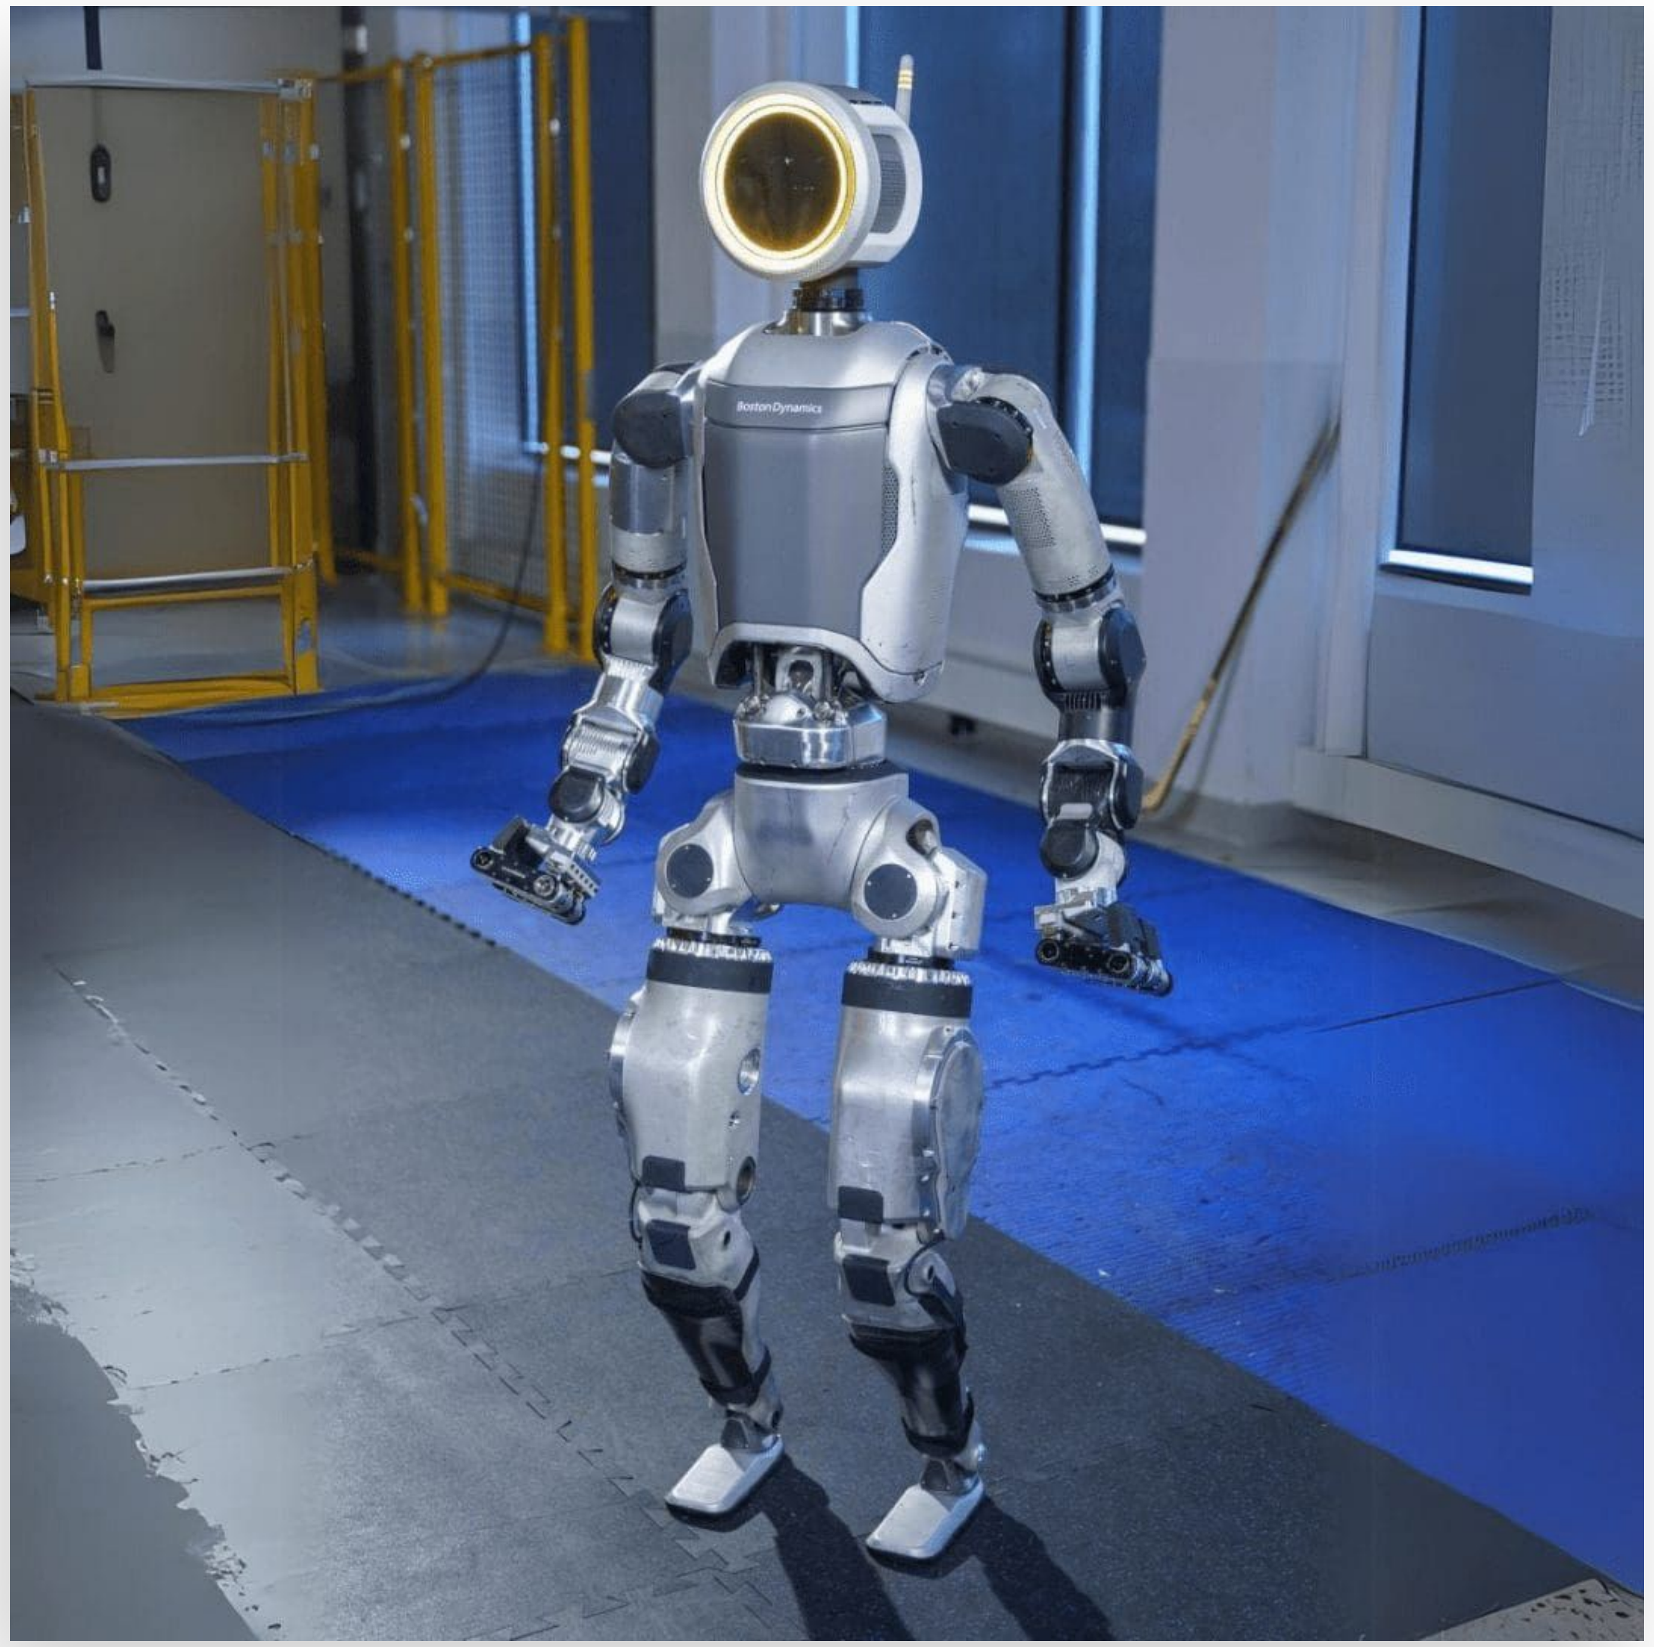

https://bostondynamics.com/blog/electric-new-era-for-atlas/

So you’re talking about Humanoid Robots, then?

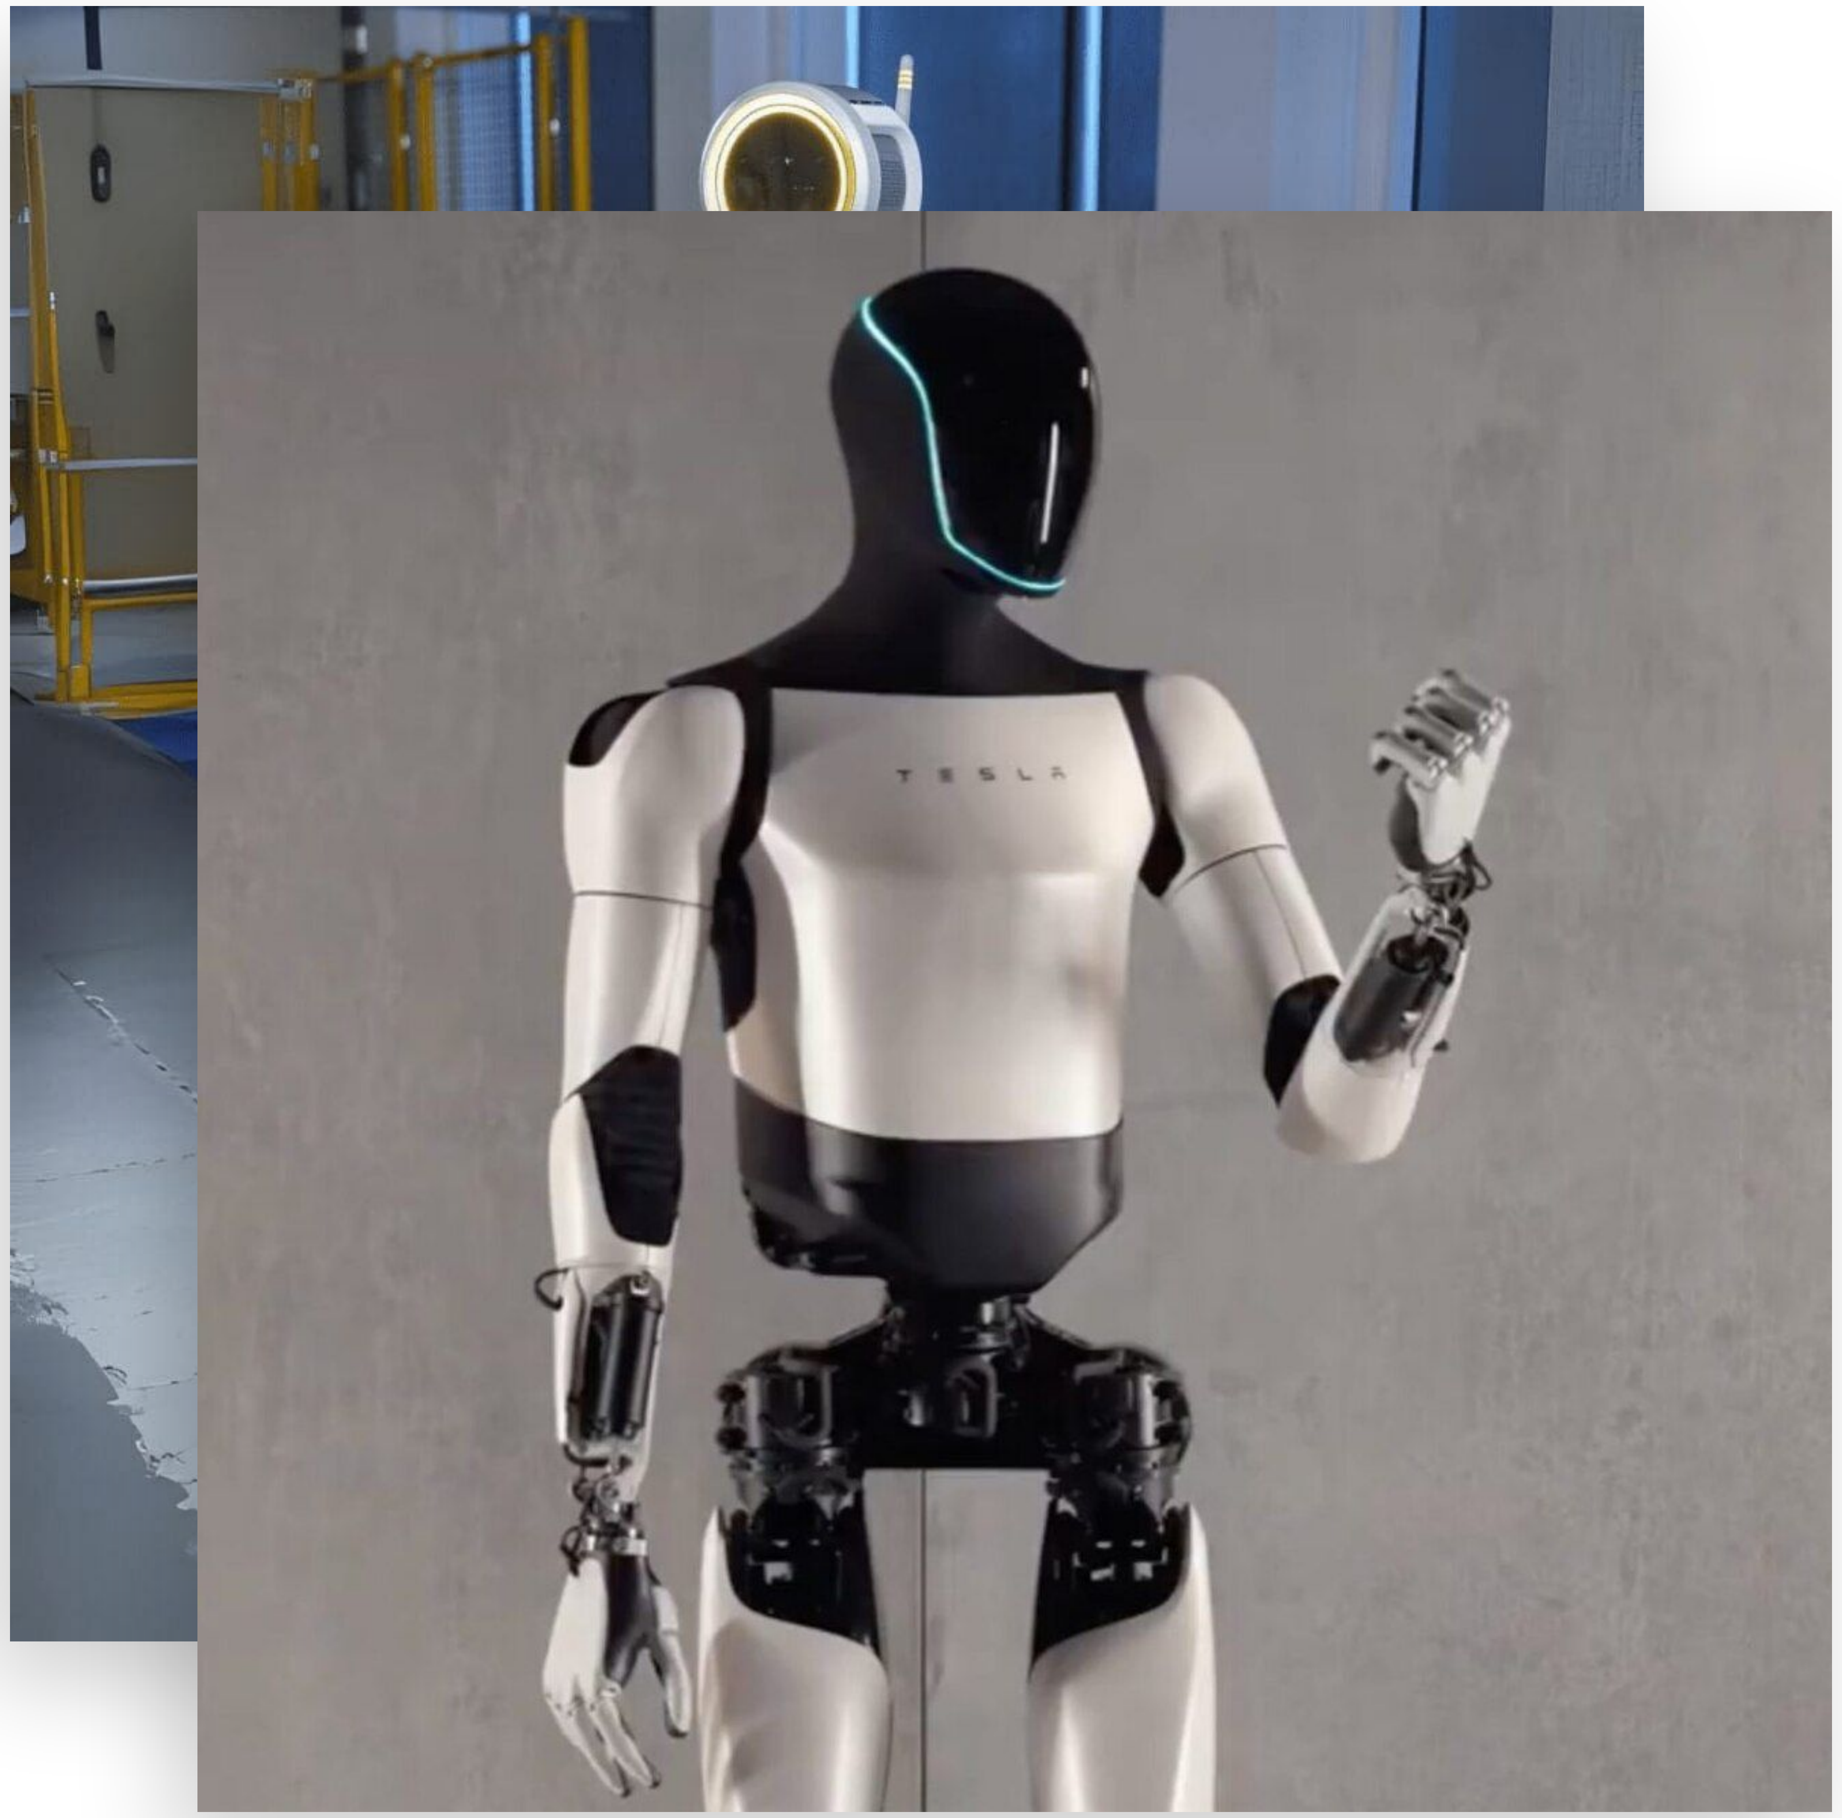

https://bostondynamics.com/blog/electric-new-era-for-atlas/

https://www.teslarati.com/tesla-shows-off-optimus-gen-2-humanoid-robot-video/

So you're talking about Humanoid Robots, then?

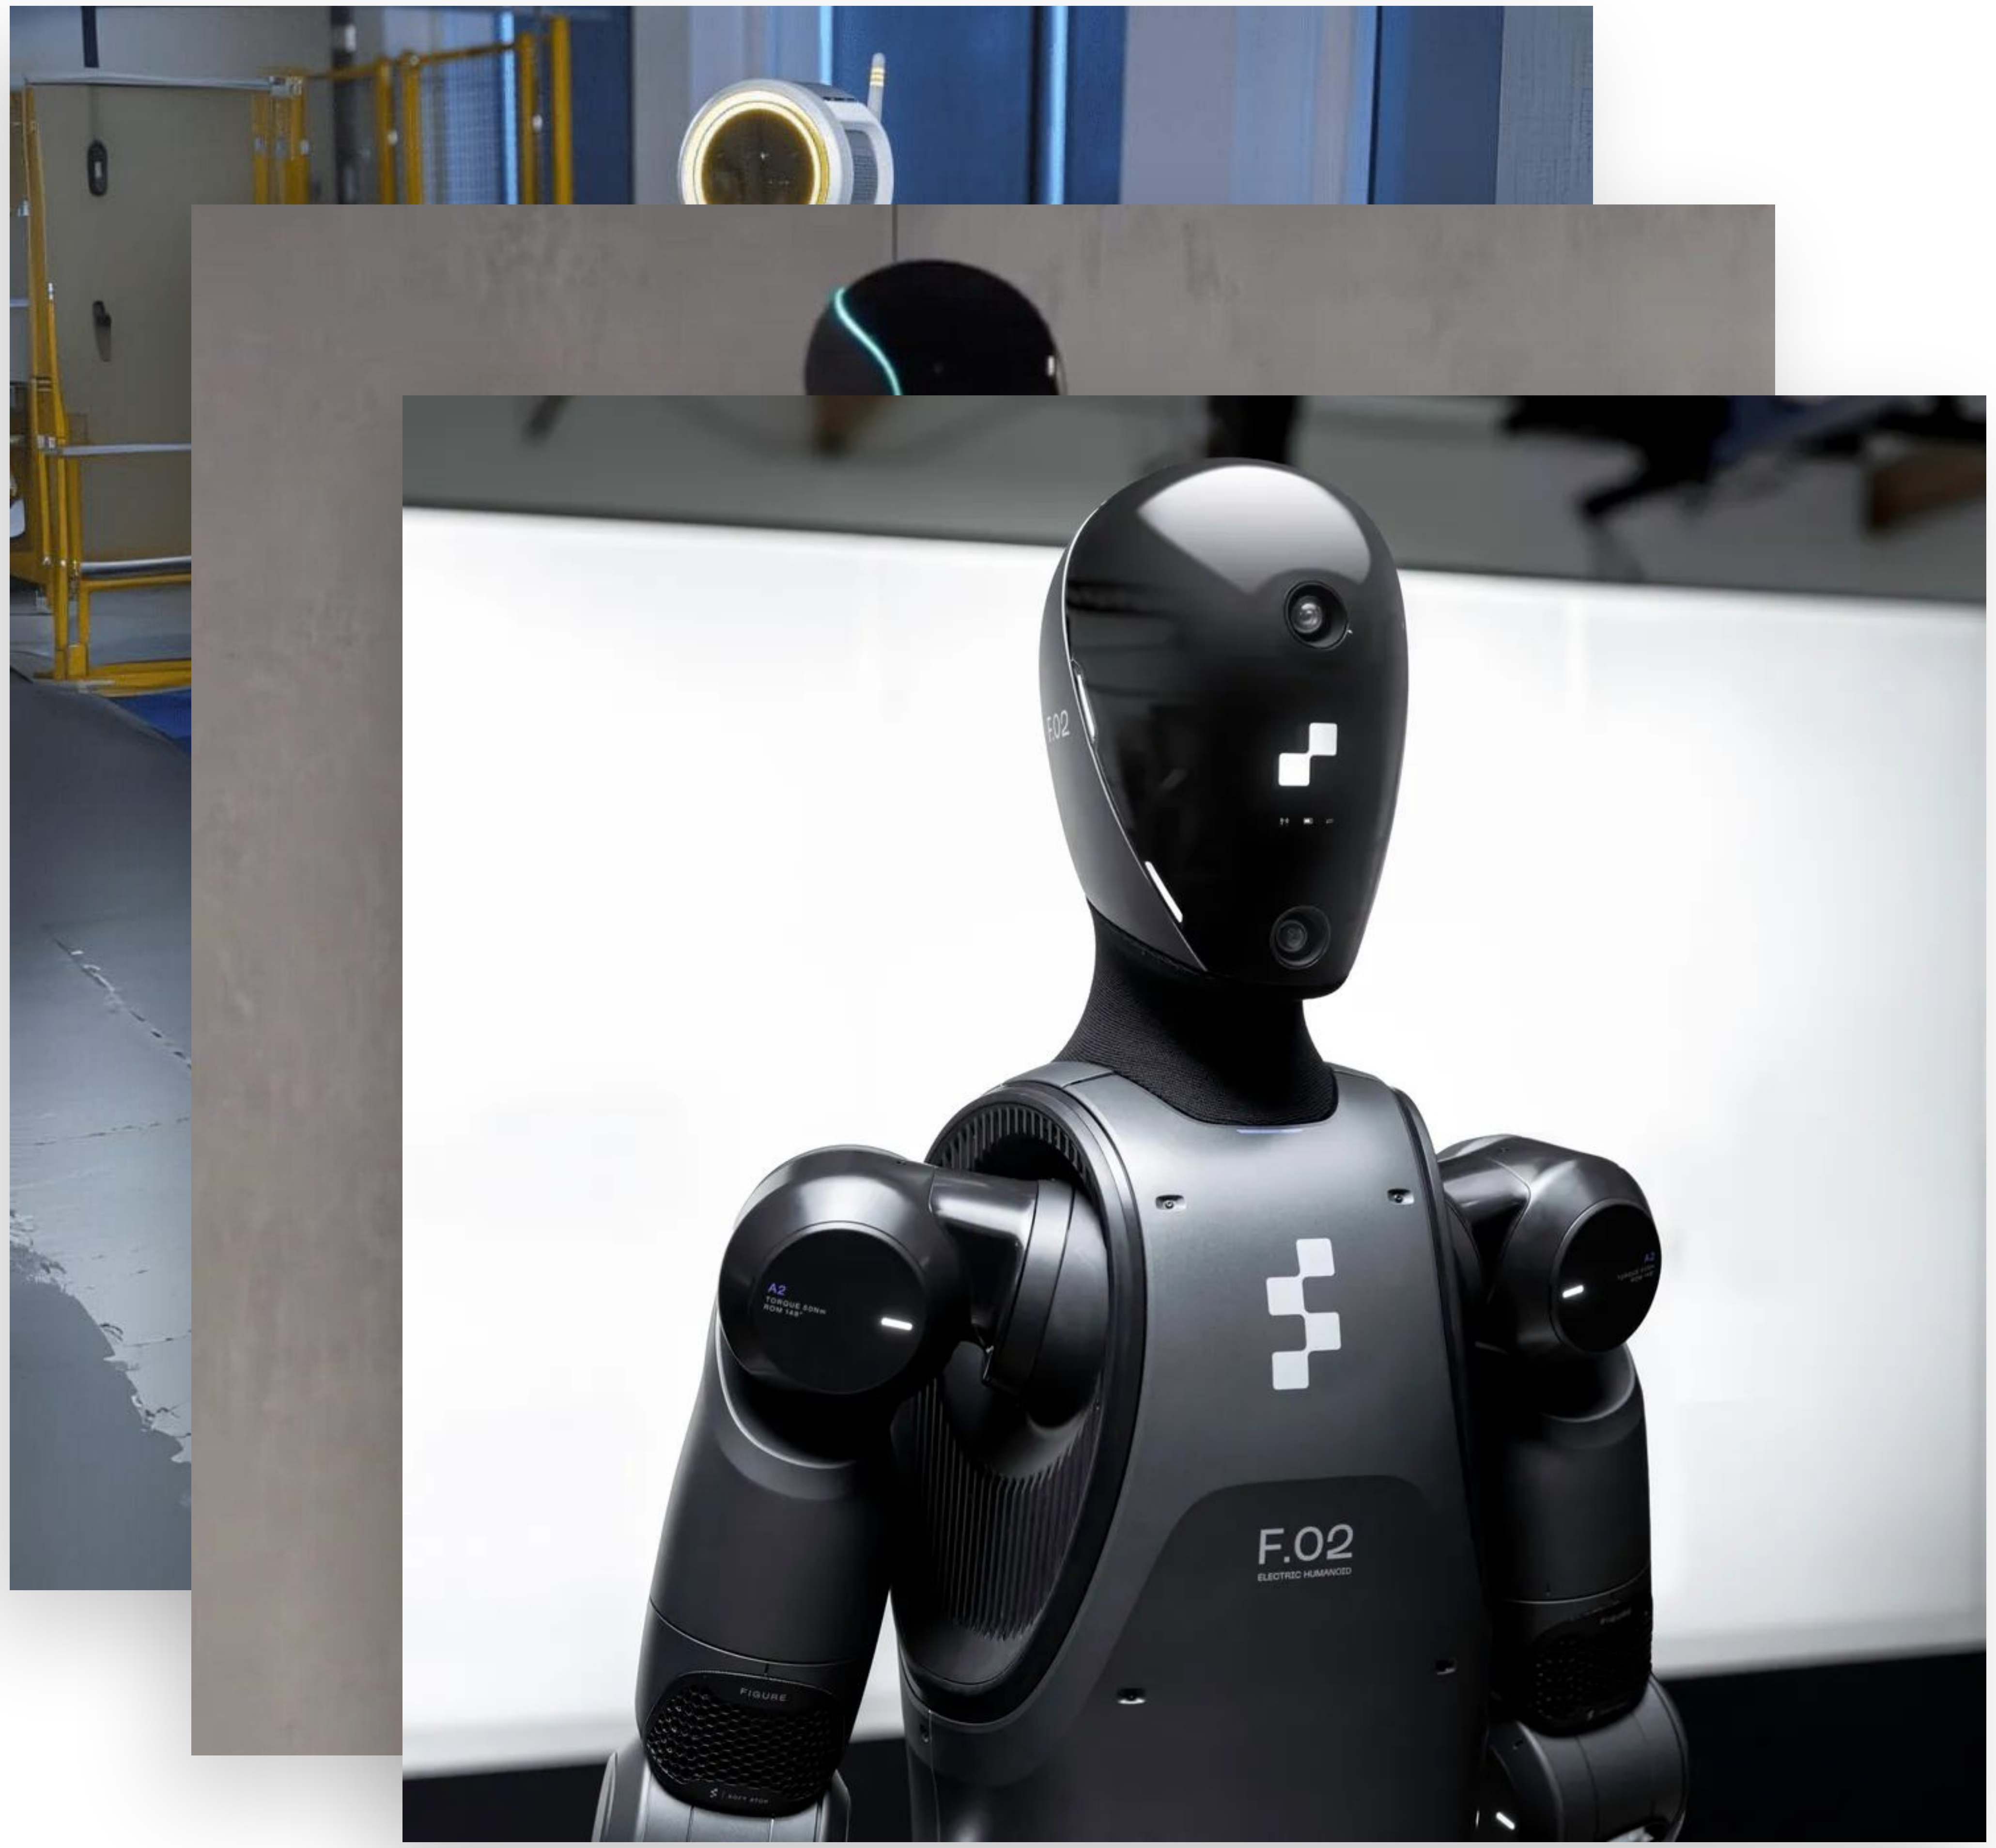

https://bostondynamics.com/blog/electric-new-era-for-atlas/

https://www.teslarati.com/tesla-shows-off-optimus-gen-2-humanoid-robot-video/

https://www.therobotreport.com/figure-02-humanoid-robot-is-ready-to-get-to-work/

So you’re talking about Humanoid Robots, then?

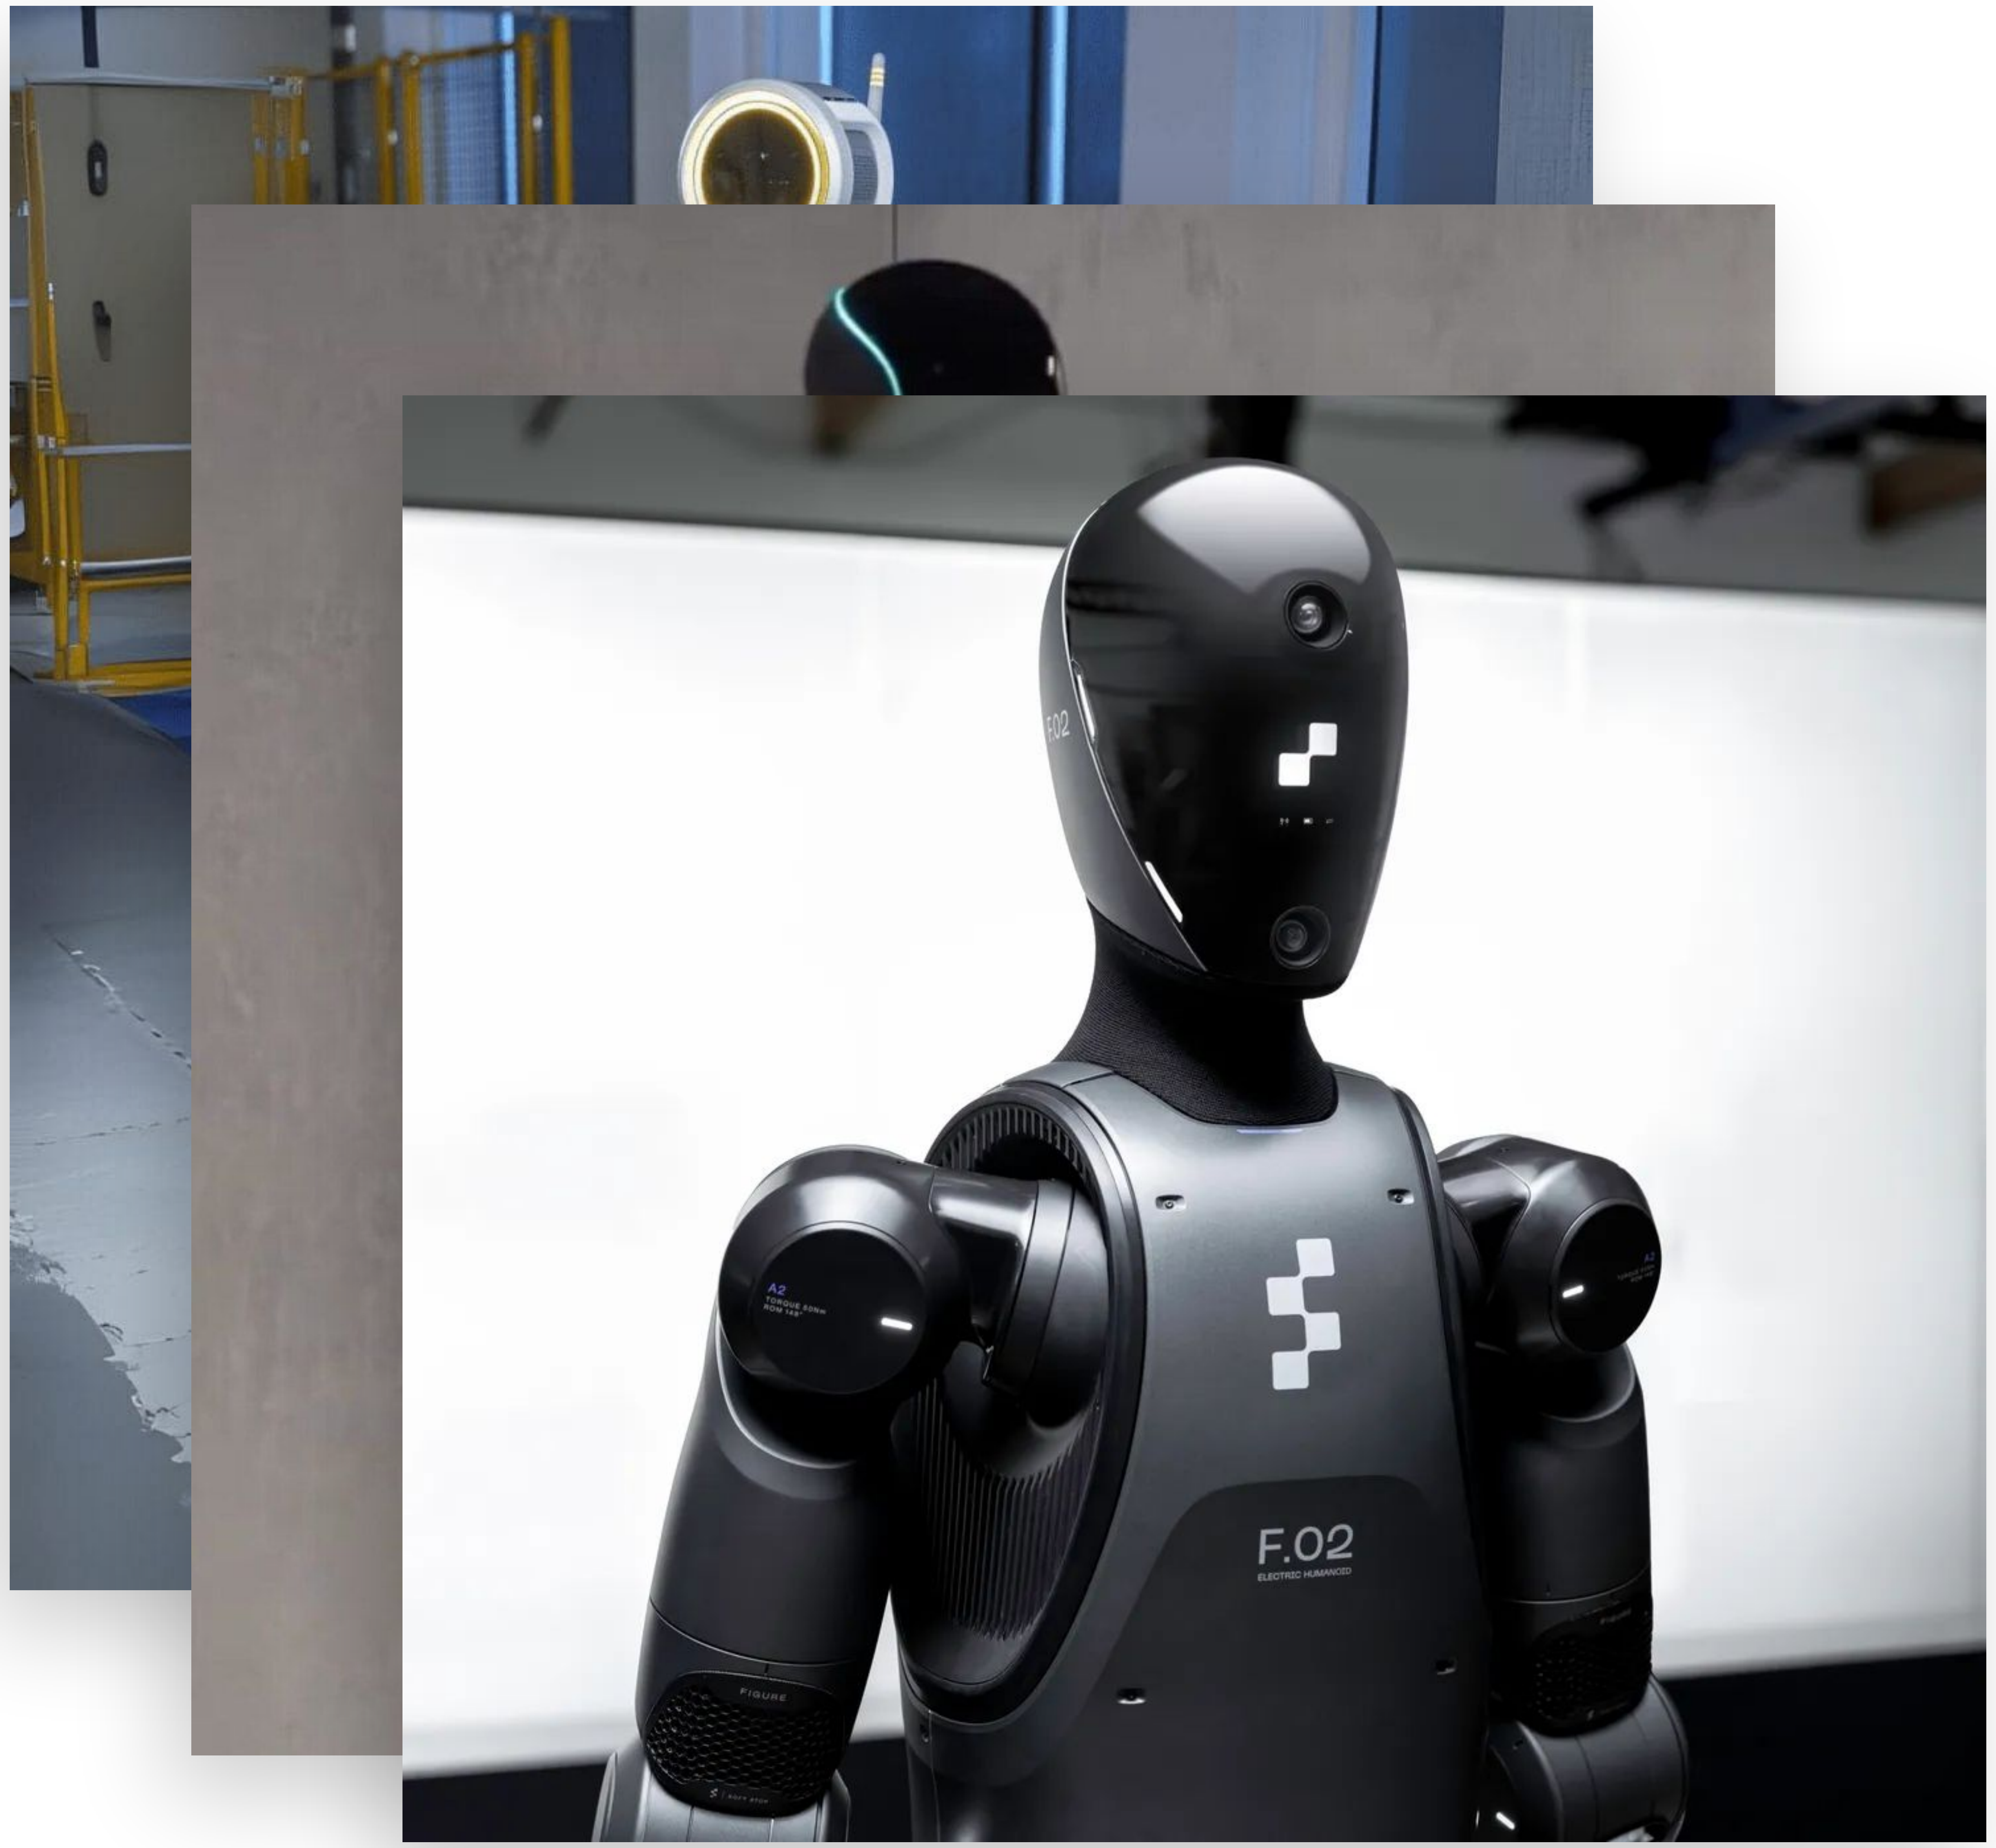

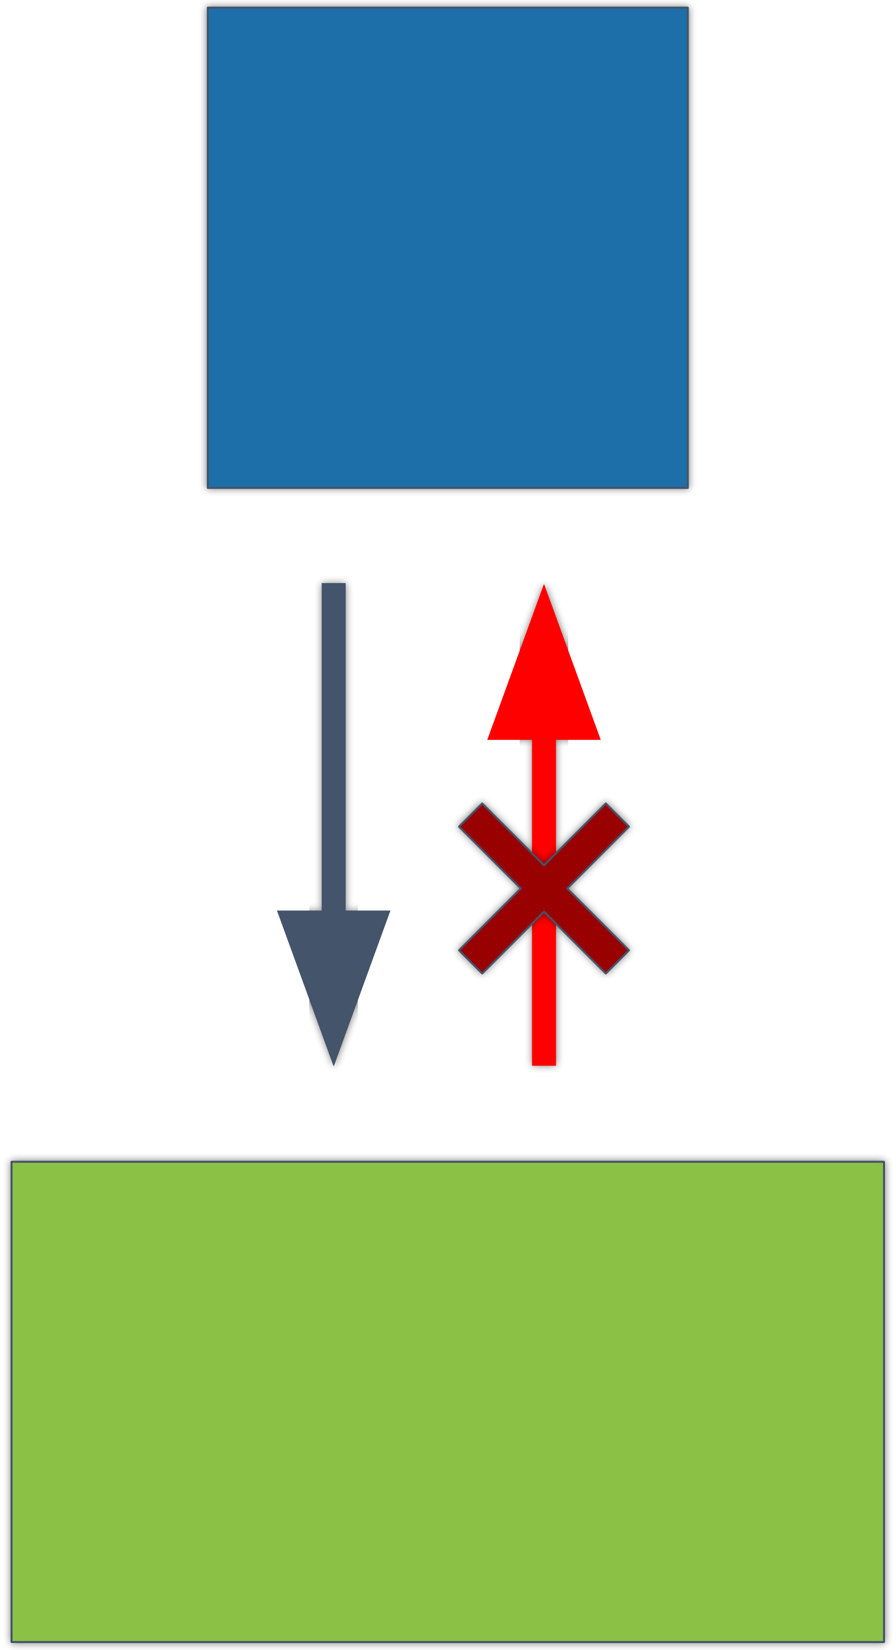

https://bostondynamics.com/blog/electric-new-era-for-atlas/

https://www.teslarati.com/tesla-shows-off-optimus-gen-2-humanoid-robot-video/

https://www.therobotreport.com/figure-02-humanoid-robot-is-ready-to-get-to-work/

So you're talking about Humanoid Robots, then?

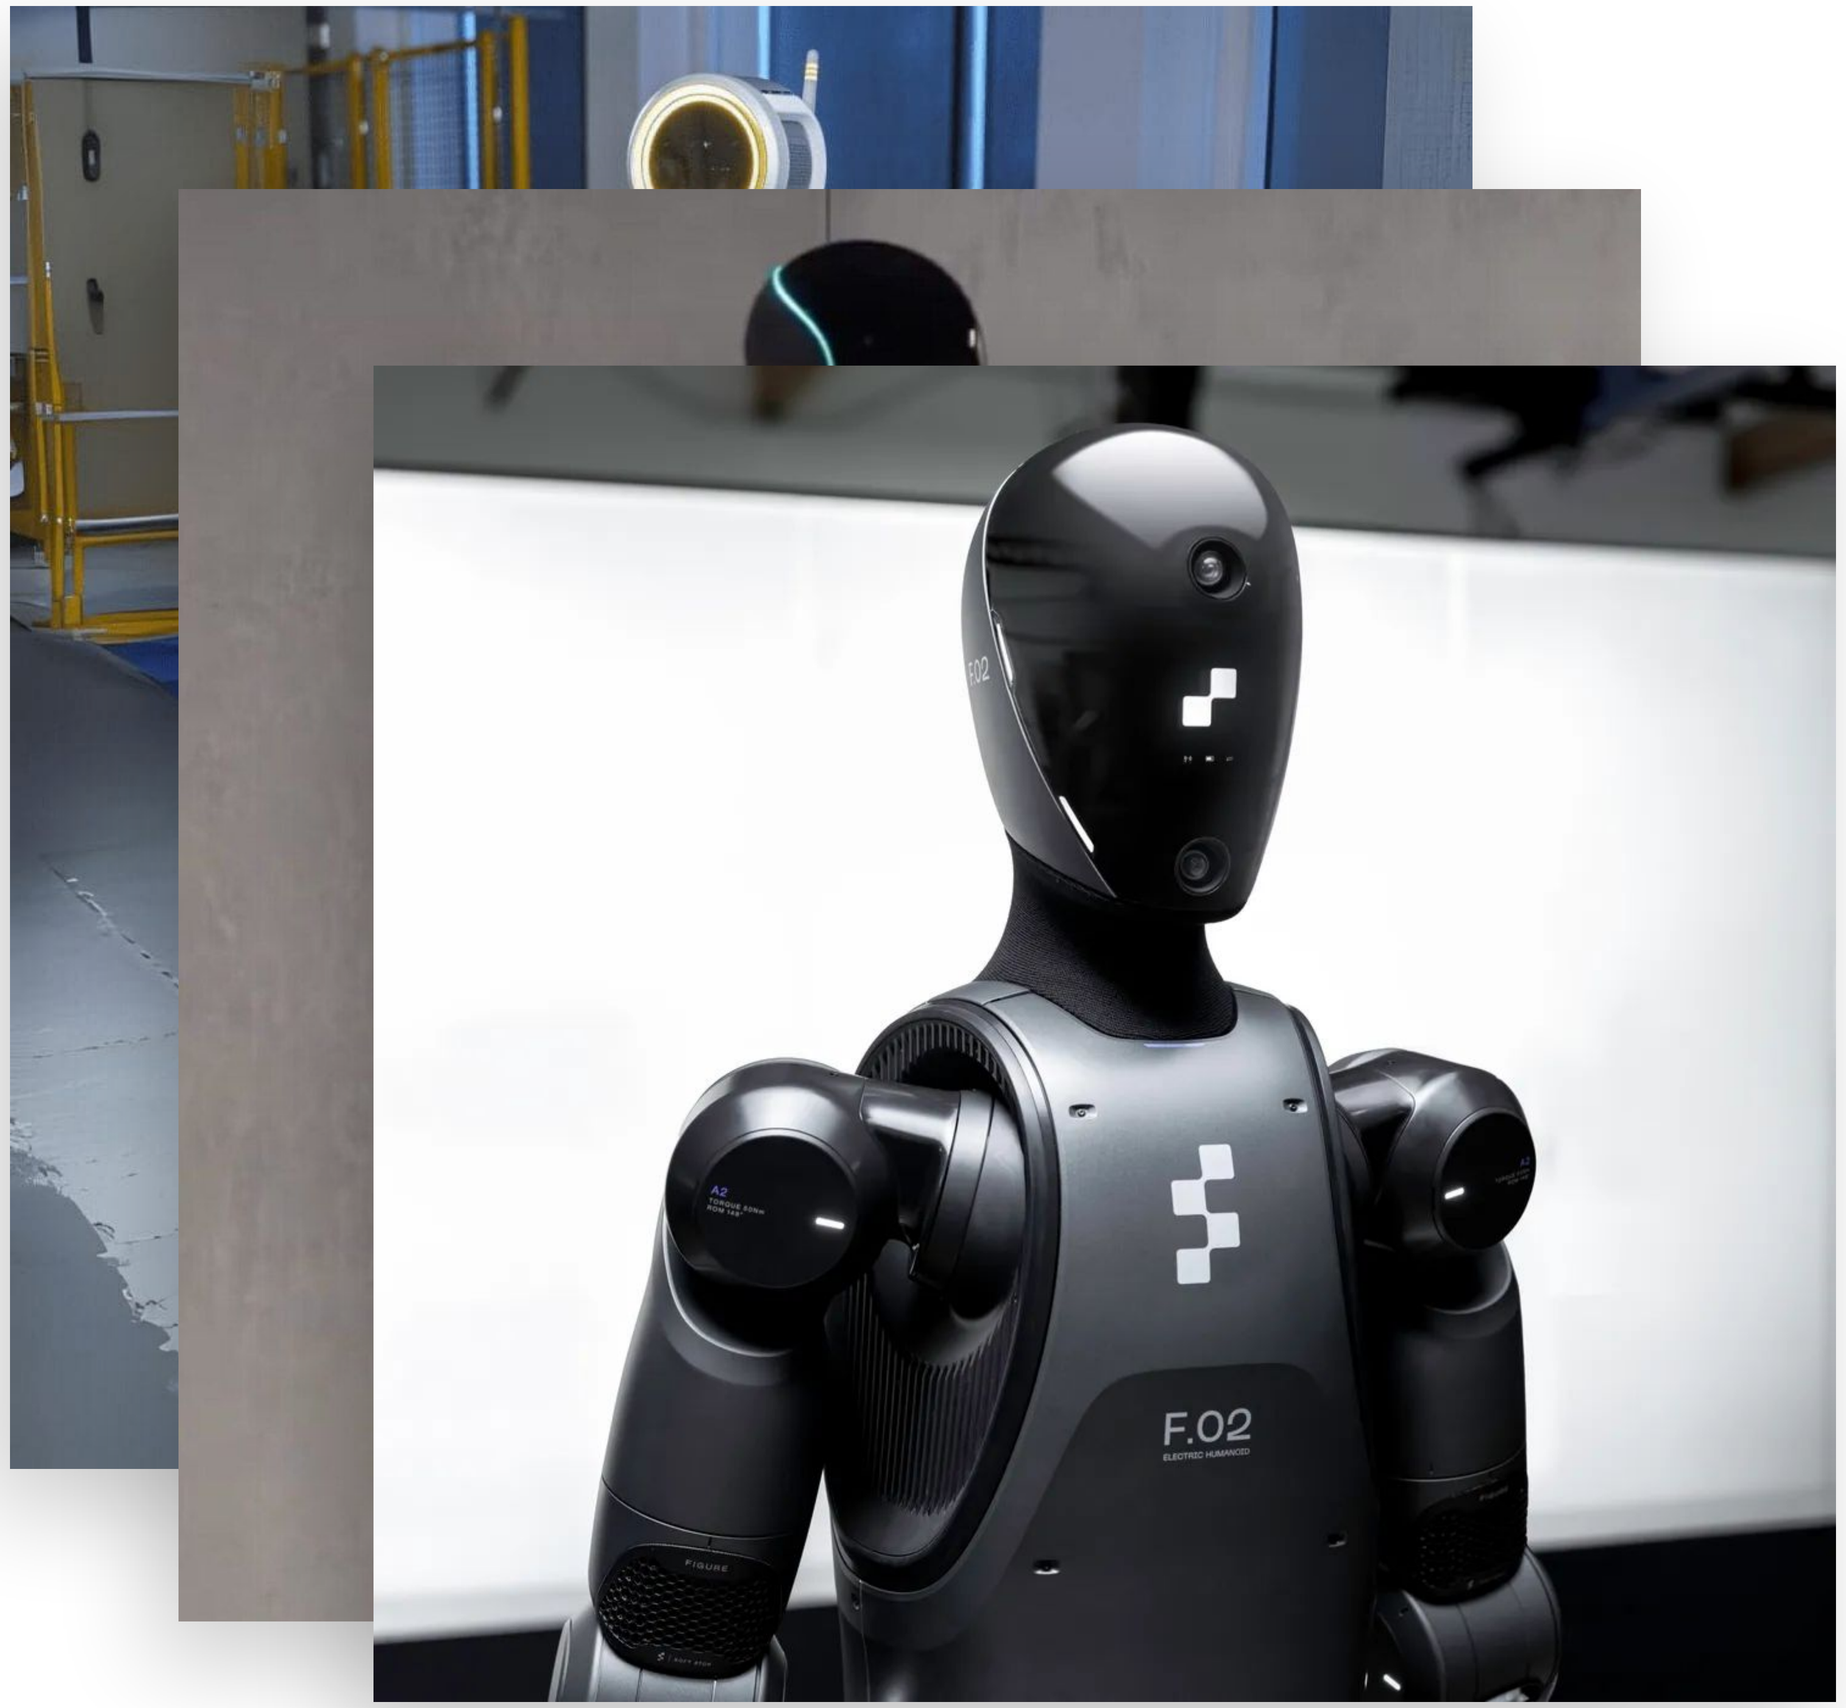

https://www.researchgate.net/figure/Robot-stir-fry-is-a-non-prehensile-manipulation-of-semi-fluid-objects-which-requires_fig1_360559814

https://bostondynamics.com/blog/electric-new-era-for-atlas/

https://www.teslarati.com/tesla-shows-off-optimus-gen-2-humanoid-robot-video/

https://www.therobotreport.com/figure-02-humanoid-robot-is-ready-to-get-to-work/

Humanoid Robots are a *subset* of **Bimanual Robots**

Why is Bimanual Manipulation important?

Why is Bimanual Manipulation important?

Objects in human environments are built for dual-arm agents

Why is Bimanual Manipulation important?

We want robots to help humans in their environments.

Objects in human environments are built for dual-arm agents

Why is Bimanual Manipulation important?

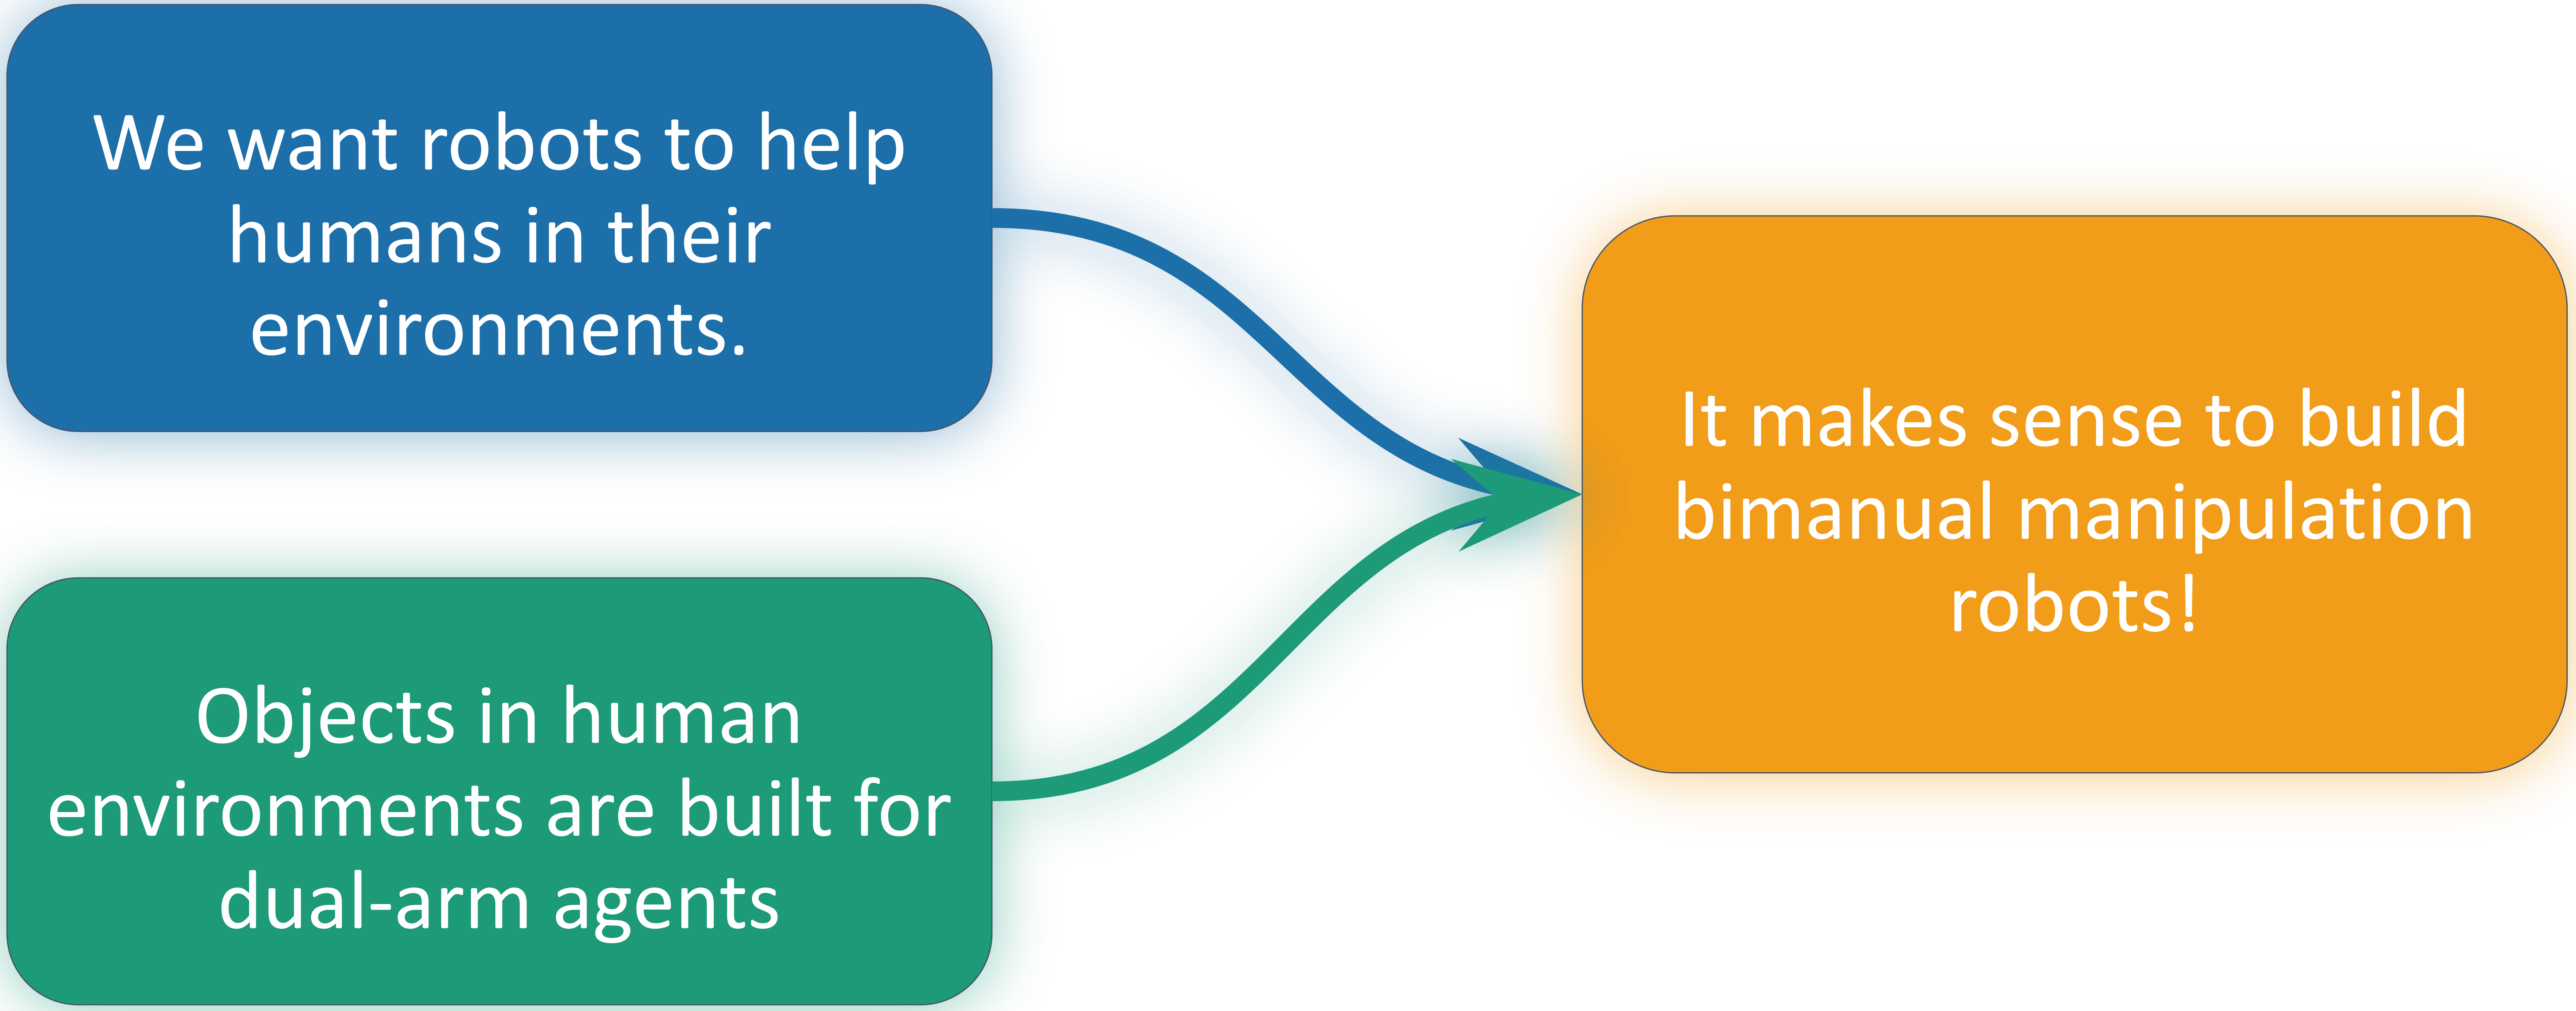

Why is Bimanual Manipulation important?

Many policy training methods require expert demonstrations

Why is Bimanual Manipulation important?

Many policy training methods require expert demonstrations

It's **much** easier to teleoperate a bimanual robot to record expert demonstrations since we already innately know how to perform these bimanual tasks ourselves!

Why is Bimanual Manipulation important?

https://www.youtube.com/watch?v=PHXQFE-Rteo

Why is Bimanual Manipulation important?

- **Similarity to operator**
  - Teleoperation becomes near-trivial as the operator’s innate bimanual manipulation skills can be applied to the robot’s operation

From *Dual arm manipulation -- A survey*
C. Smith, Y. Karayiannidis, L. Nalpantidis, X. Gratal, P. Qi, D. V. Dimarogonas, D. Kragic
KTH Royal Institute of Technology
DOI 10.1016/j.robot.2012.07.005

Why is Bimanual Manipulation important?

- Similarity to operator
  - Teleoperation becomes near-trivial as the operator’s innate bimanual manipulation skills can be applied to the robot’s operation
- **Manipulability**
  - The ability to manipulate both ends of a task (i.e. peg-in-hole or nut+bolt screw assembly) provides more avenues towards solving a task

From *Dual arm manipulation -- A survey*
C. Smith, Y. Karayiannidis, L. Nalpantidis, X. Gratal, P. Qi, D. V. Dimarogonas, D. Kragic
KTH Royal Institute of Technology
DOI 10.1016/j.robot.2012.07.005

Why is Bimanual Manipulation important?

- Similarity to operator
  - Teleoperation becomes near-trivial as the operator's innate bimanual manipulation skills can be applied to the robot's operation
- Manipulability
  - The ability to manipulate both ends of a task (i.e. peg-in-hole or nut+bolt screw assembly) provides more avenues towards solving a task
- **Cognitive Motivation**
  - Humans have an innate understanding of bimanual manipulation, so it becomes much easier to relate to and understand what a manipulator is trying to do

From *Dual arm manipulation -- A survey*
C. Smith, Y. Karayiannidis, L. Nalpantidis, X. Gratal, P. Qi, D. V. Dimarogonas, D. Kragic
KTH Royal Institute of Technology
DOI 10.1016/j.robot.2012.07.005

Why is Bimanual Manipulation important?

- Similarity to operator
  - Teleoperation becomes near-trivial as the operator's innate bimanual manipulation skills can be applied to the robot's operation
- Manipulability
  - The ability to manipulate both ends of a task (i.e. peg-in-hole or nut+bolt screw assembly) provides more avenues towards solving a task
- Cognitive Motivation
  - Humans have an innate understanding of bimanual manipulation, so it becomes much easier to relate to and understand what a manipulator is trying to do
- **Human form factor**
  - Robots are often expected to operate in environments intended for human use, thus it motivates the creation of humanlike (and thus, bimanual) robots

From *Dual arm manipulation -- A survey*
C. Smith, Y. Karayiannidis, L. Nalpantidis, X. Gratal, P. Qi, D. V. Dimarogonas, D. Kragic
KTH Royal Institute of Technology
DOI 10.1016/j.robot.2012.07.005

Timeline of Bimanual Manipulators

Select milestones from the past 30 years

Timeline of Bimanual Manipulators

Select milestones from the past 30 years

1994

Yaskawa Motoman introduces the MRC system

Allowed for synchronized control and coordination of two robotic arms by "teaching" it a sequence of movements, or programming a task on a PC

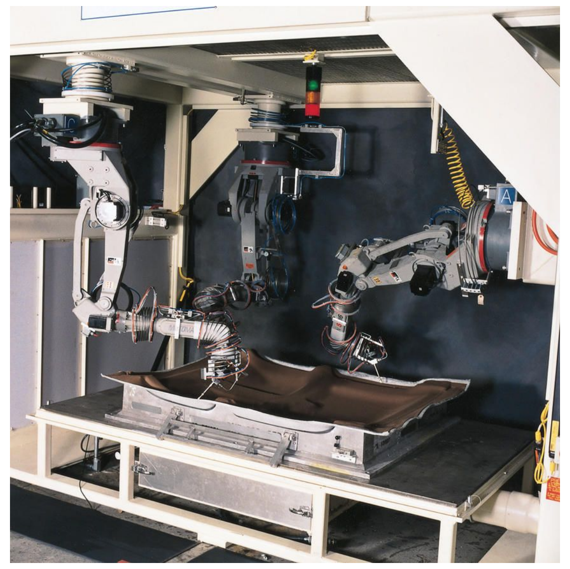

https://ifr.org/robot-history

Timeline of Bimanual Manipulators

Select milestones from the past 30 years

https://robotsguide.com/robots/davinci

Intuitive Surgical releases first Da Vinci surgical robot system

Enables less invasive surgeries through the use of smaller robotic tools with bimanual teleoperation

1999

1994

Yaskawa Motoman introduces the MRC system

Allowed for synchronized control and coordination of two robotic arms by "teaching" it a sequence of movements, or programming a task on a PC

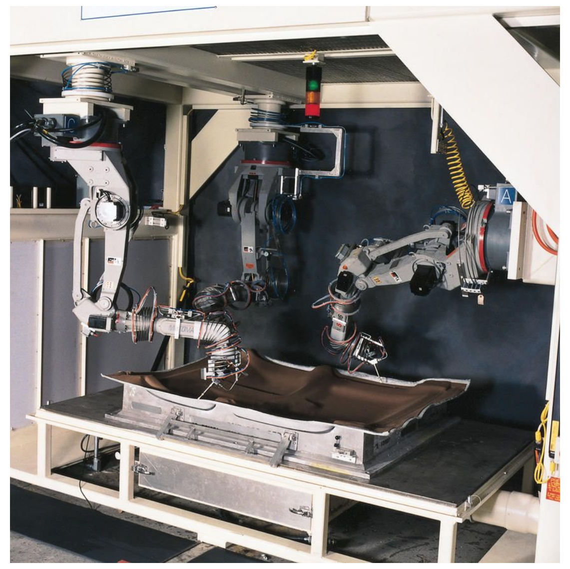

https://ifr.org/robot-history

Timeline of Bimanual Manipulators

Select milestones from the past 30 years

https://robotsguide.com/robots/davinci

Intuitive Surgical releases first Da Vinci surgical robot system

Enables less invasive surgeries through the use of smaller robotic tools with bimanual teleoperation

1999

1994

Yaskawa Motoman introduces the MRC system

Allowed for synchronized control and coordination of two robotic arms by "teaching" it a sequence of movements, or programming a task on a PC

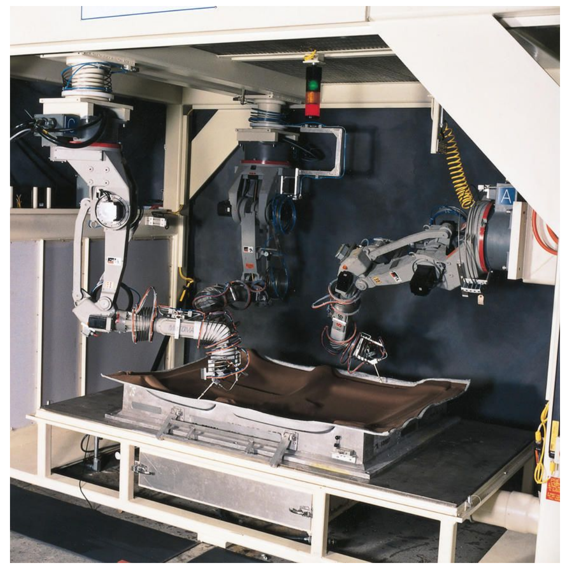

https://ifr.org/robot-history

2006

Personal Robotics Lab develops Herb (Home Exploring Robot Butler)

Bimanual robot for domestic tasks developed by the Personal Robotics Lab at CMU (now at UW Seattle)

https://robotsguide.com/robots/herb

Timeline of Bimanual Manipulators

Select milestones from the past 30 years

1994

Yaskawa Motoman introduces the MRC system

Allowed for synchronized control and coordination of two robotic arms by “teaching” it a sequence of movements, or programming a task on a PC

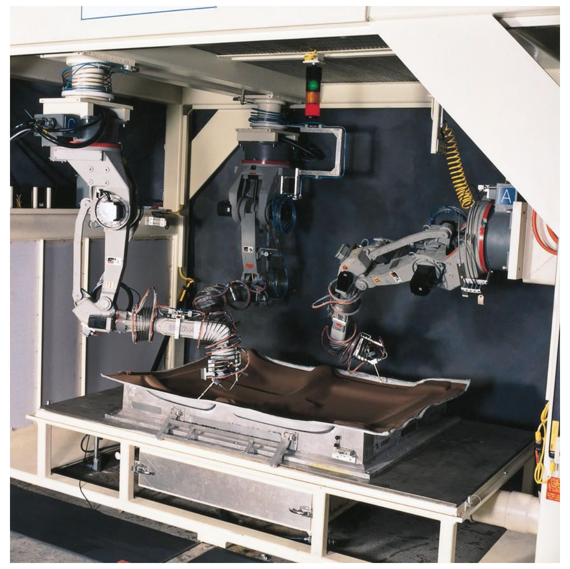

https://ifr.org/robot-history

1999

https://robotsguide.com/robots/davinci

Intuitive Surgical releases first Da Vinci surgical robot system

Enables less invasive surgeries through the use of smaller robotic tools with bimanual teleoperation

2006

Personal Robotics Lab develops Herb (Home Exploring Robot Butler)

Bimanual robot for domestic tasks developed by the Personal Robotics Lab at CMU (now at UW Seattle)

https://robotsguide.com/robots/herb

2012

https://robotsguide.com/robots/baxter

Rethink Robotics releases Baxter

Now-defunct Rethink Robotics releases Baxter, a bimanual manufacturing robot that can be programmed simply by moving its arms by hand

Timeline of Bimanual Manipulators

Select milestones from the past 30 years

https://robotsguide.com/robots/davinci

Intuitive Surgical releases first Da Vinci surgical robot system

Enables less invasive surgeries through the use of smaller robotic tools with bimanual teleoperation

https://robotsguide.com/robots/baxter

Rethink Robotics releases Baxter

Now-defunct Rethink Robotics releases Baxter, a bimanual manufacturing robot that can be programmed simply by moving its arms by hand

1999

2012

1994

Yaskawa Motoman introduces the MRC system

Allowed for synchronized control and coordination of two robotic arms by “teaching” it a sequence of movements, or programming a task on a PC

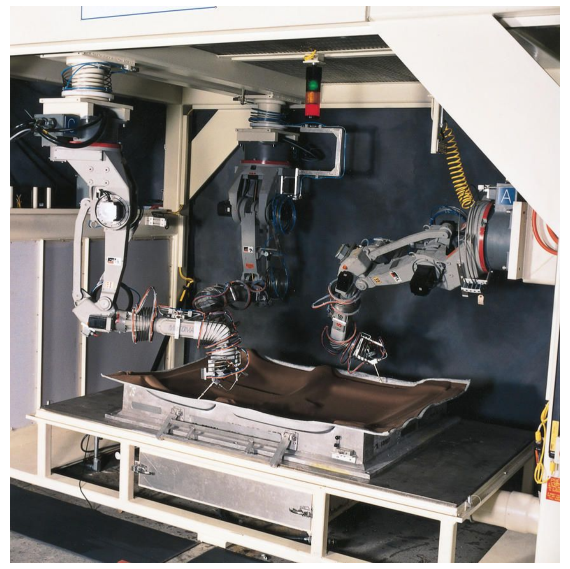

https://ifr.org/robot-history

2006

Personal Robotics Lab develops Herb (Home Exploring Robot Butler)

Bimanual robot for domestic tasks developed by the Personal Robotics Lab at CMU (now at UW Seattle)

https://robotsguide.com/robots/herb

2013

Boston Dynamics develops first iteration of Atlas Robot

Developed as a disaster-response robot for the Defence Advanced Research Projects Agency

https://robotsguide.com/robots/atlas2013

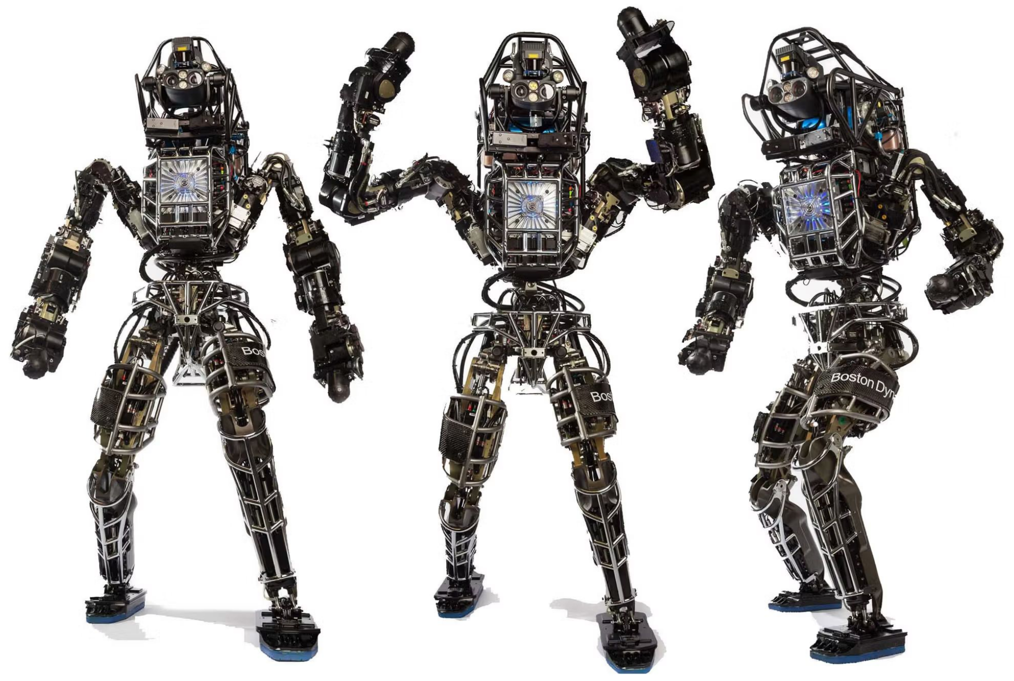

Timeline of Bimanual Manipulators

Select milestones from the past 30 years

1994

Yaskawa Motoman introduces the MRC system

Allowed for synchronized control and coordination of two robotic arms by "teaching" it a sequence of movements, or programming a task on a PC

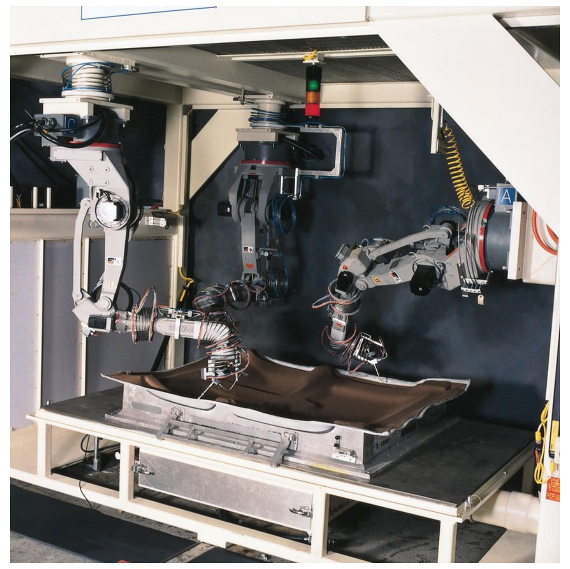

https://ifr.org/robot-history

1999

https://robotsguide.com/robots/davinci

Intuitive Surgical releases first Da Vinci surgical robot system

Enables less invasive surgeries through the use of smaller robotic tools with bimanual teleoperation

2006

Personal Robotics Lab develops Herb (Home Exploring Robot Butler)

Bimanual robot for domestic tasks developed by the Personal Robotics Lab at CMU (now at UW Seattle)

https://robotsguide.com/robots/herb

2012

https://robotsguide.com/robots/baxter

Rethink Robotics releases Baxter

Now-defunct Rethink Robotics releases Baxter, a bimanual manufacturing robot that can be programmed simply by moving its arms by hand

2013

Boston Dynamics develops first iteration of Atlas Robot

Developed as a disaster-response robot for the Defence Advanced Research Projects Agency

https://robotsguide.com/robots/atlas2013

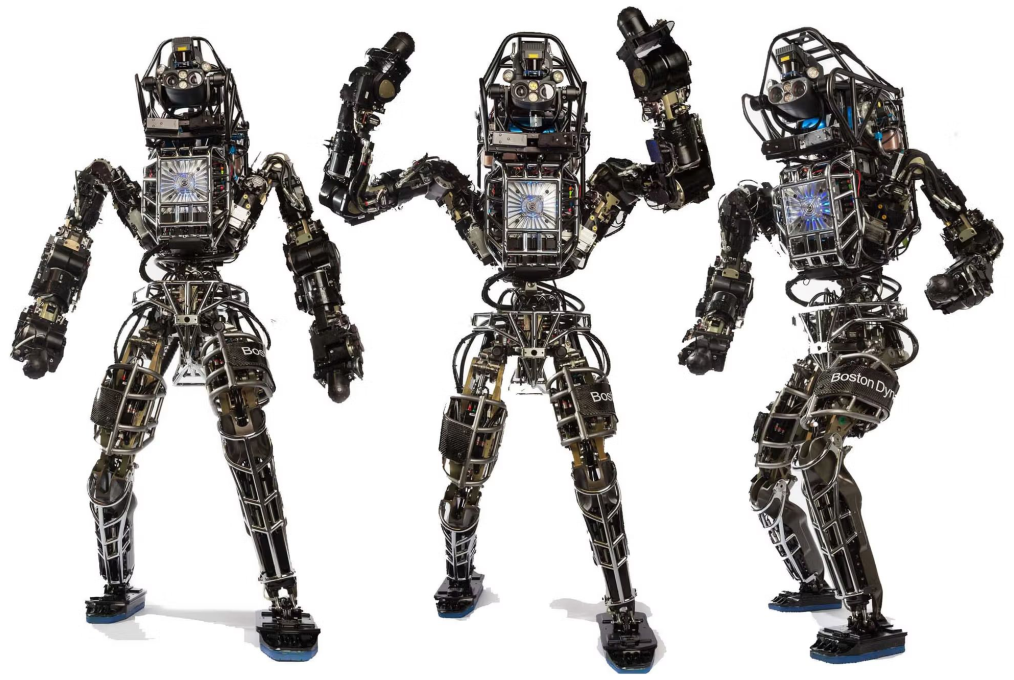

2024

https://aloha-unleashed.github.io/

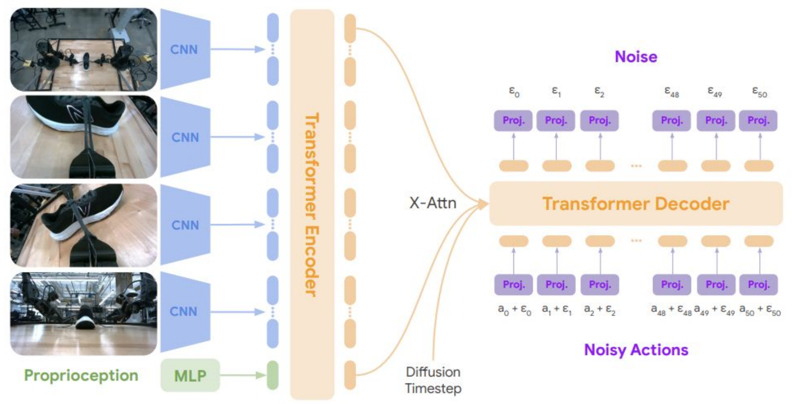

ALOHA Unleashed

Diffusion policy-backed imitation learning framework capable of learning complex bimanual tasks with deformable objects

Humanoid Robots, too!

2000

NASA Completes first iteration of Robonaut

https://robotsguide.com/robots/pr2

2010

Willow Garage releases humanoid PR2 Robot

https://robotsguide.com/robots/pr2

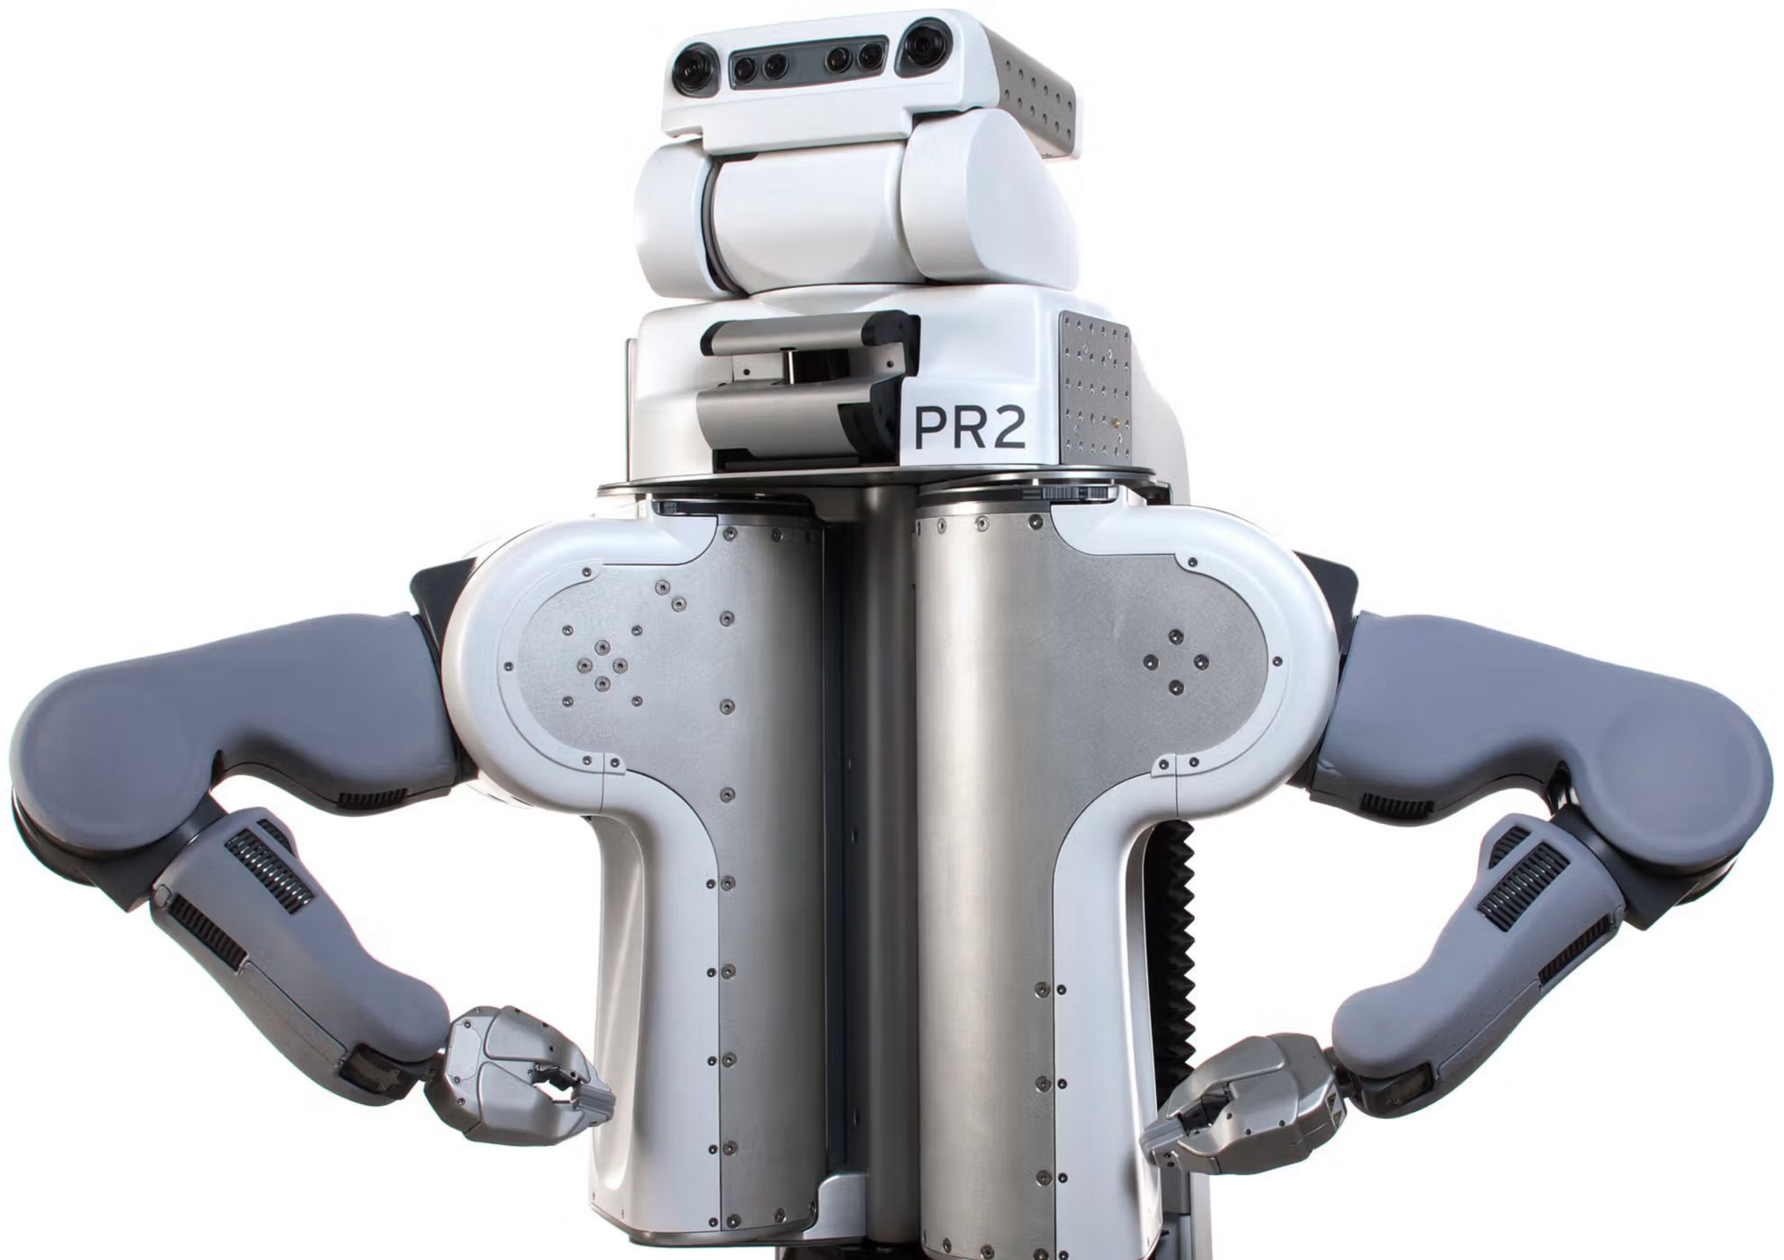

2020

1X Technologies releases EVE humanoid robot

https://robotsguide.com/robots/eve

Why stop at two manipulators?

https://www.youtube.com/watch?v=TUOmZCcRKbI

- Multi-arm robotic apple picker
- While the robot has more than two arms, it’s effectively multiple single-arm manipulation tasks in parallel
- Robots with more arms tend to be more specialized towards specific tasks
- We want a robot that can be generalized to as many domestic tasks as possible

Why stop at two manipulators?

Recall, we want to be able to perform as many tasks as possible in a domestic environment.

Why stop at two manipulators?

Challenges

Task Variety

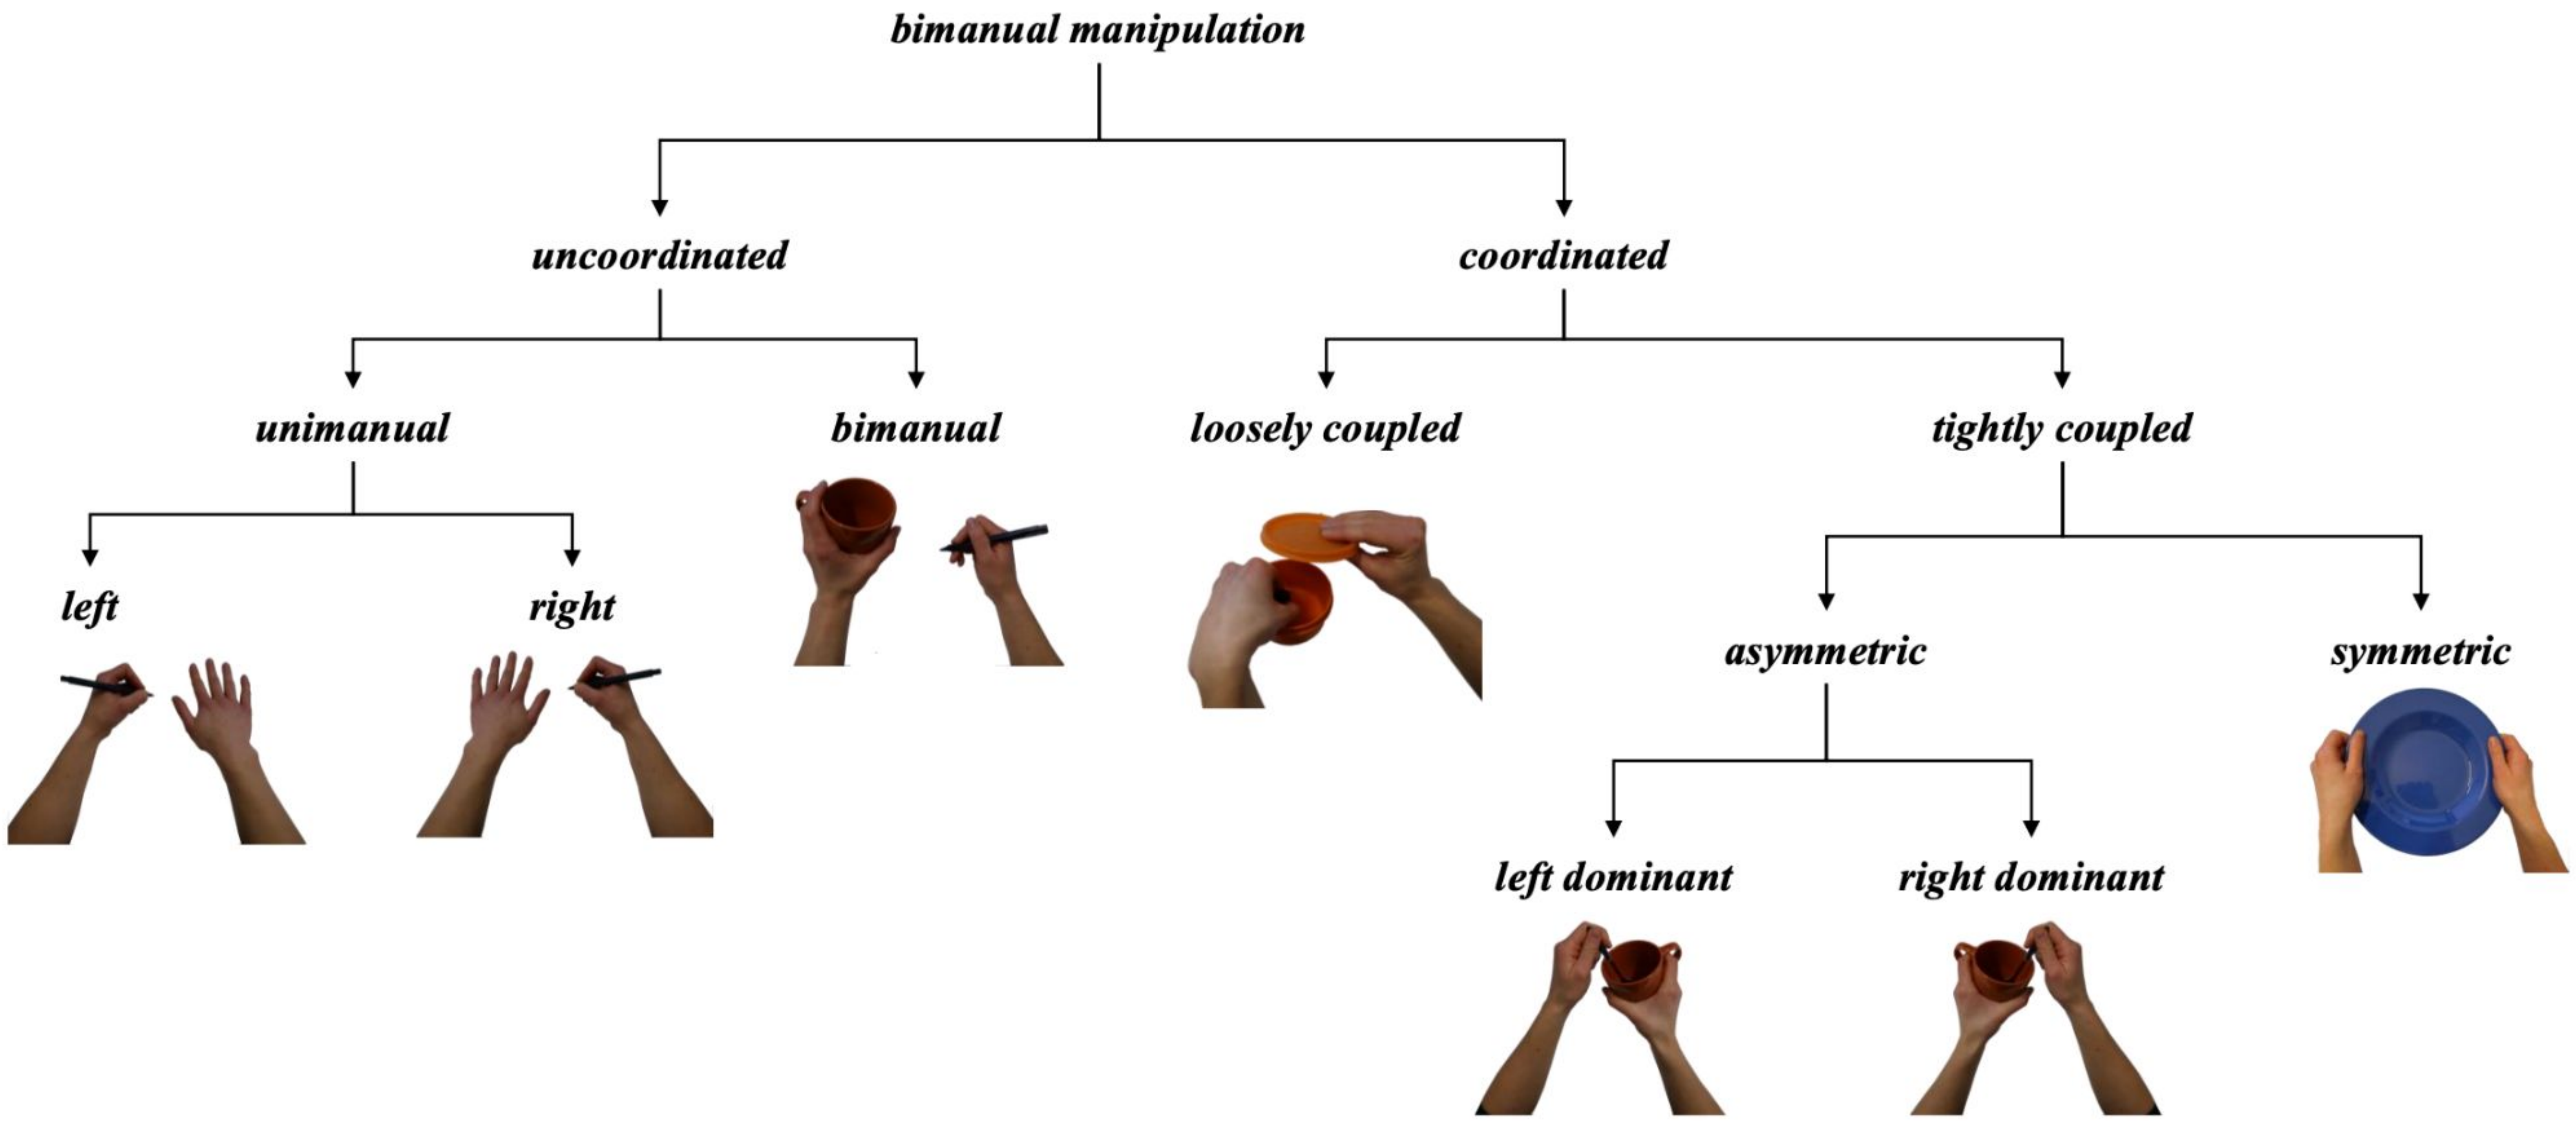

Fig. 2. Bimanual manipulation taxonomy. Tasks are classified based on the aspects *coordination*, *interaction*, *hand role* and *symmetry*.

Figure from F. Krebs and T. Asfour, "A Bimanual Manipulation Taxonomy," in *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11031-11038, Oct. 2022, doi: 10.1109/LRA.2022.3196158

Task Variety

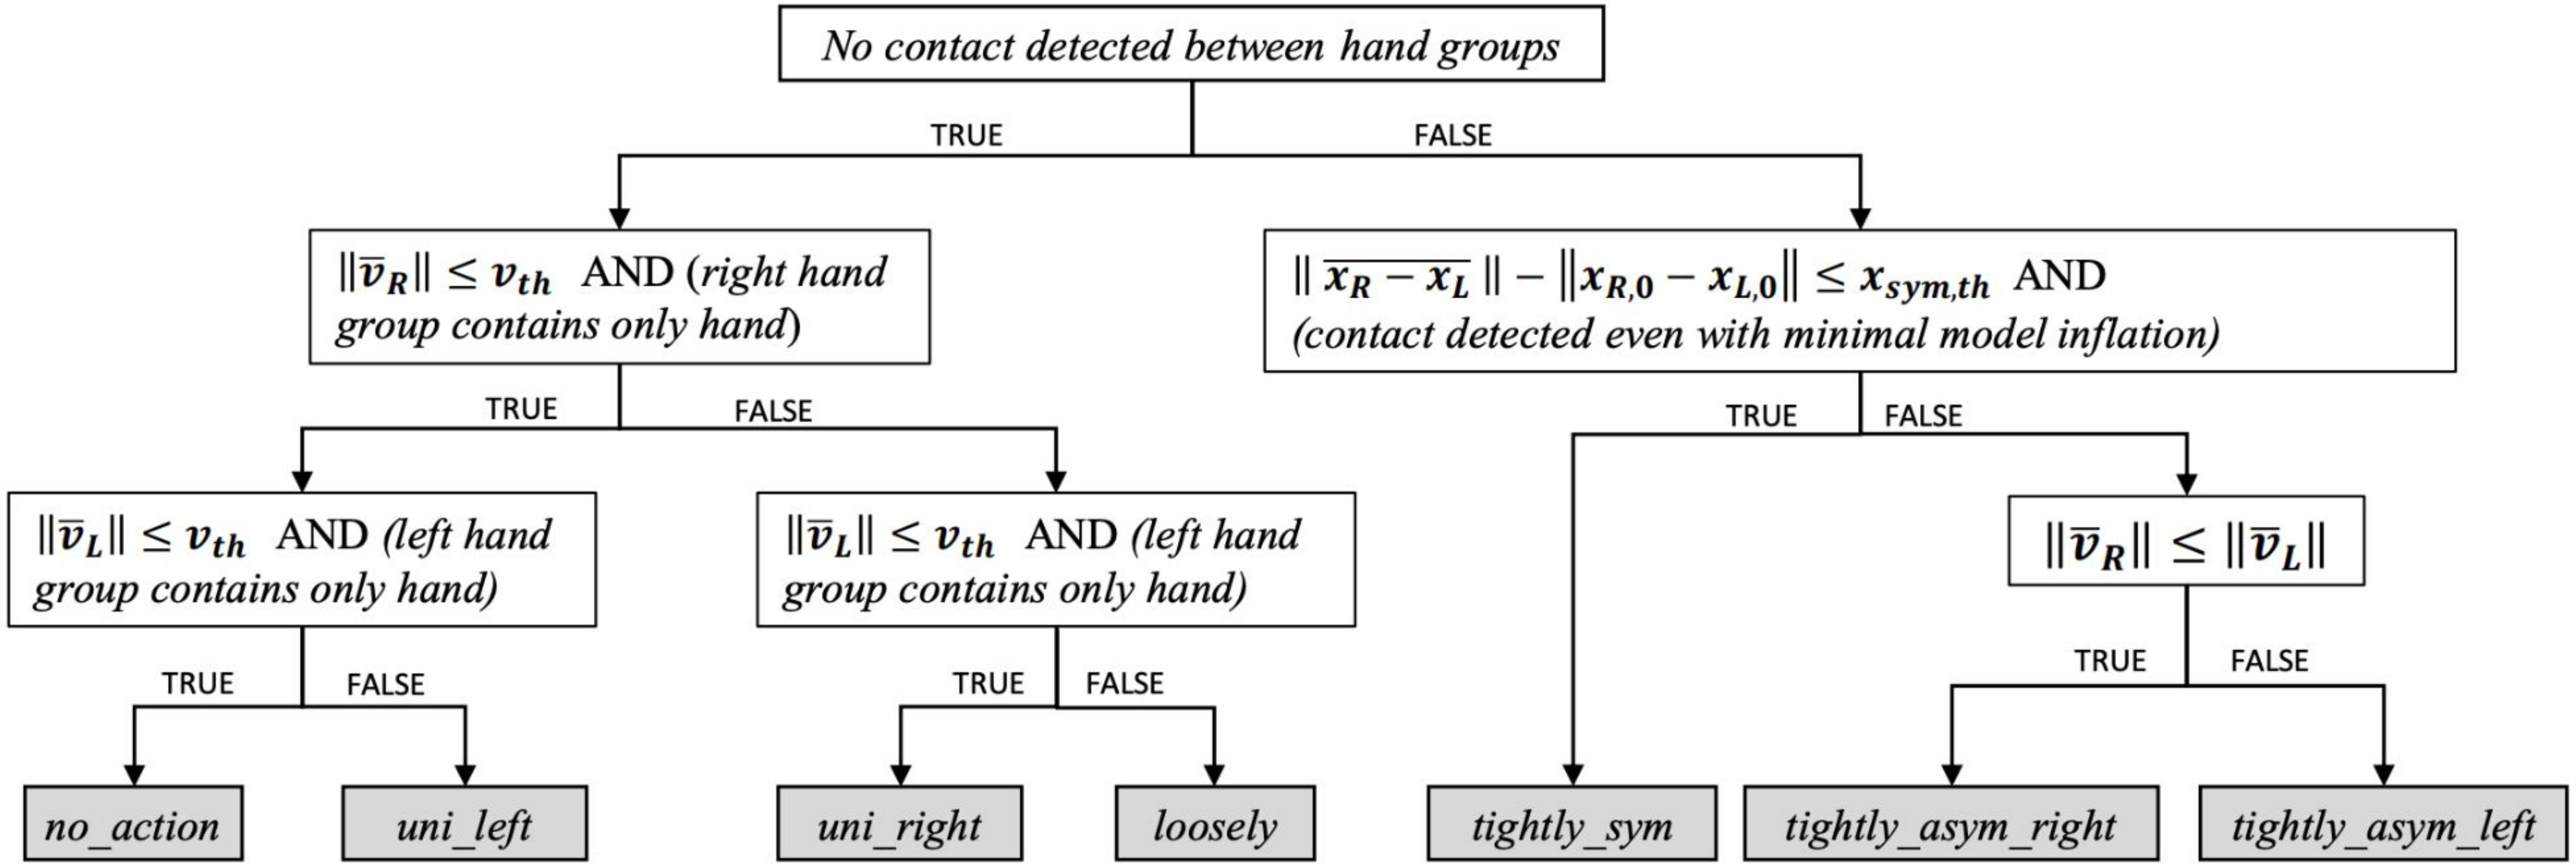

Fig. 4. Decision tree for the rule-based classification

Figure from F. Krebs and T. Asfour, "A Bimanual Manipulation Taxonomy," in *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11031-11038, Oct. 2022, doi: 10.1109/LRA.2022.3196158

Coordination

With regards to the task(s)…

- Each type of the task as defined by the decision tree shown previously warrants its own planning strategy
  - Sometimes the arms are each doing their own, uncoordinated tasks
  - Other times forces transfer between end effectors
  - Constraints for each effector can interact with each other
- The category of a task can change partway through!

Information adapted from F. Krebs and T. Asfour, "A Bimanual Manipulation Taxonomy," in *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11031-11038, Oct. 2022, doi: 10.1109/LRA.2022.3196158

Coordination

With regards to the task(s)…

- Each type of the task as defined by the decision tree shown previously warrants its own planning strategy
  - Sometimes the arms are each doing their own, uncoordinated tasks
  - Other times forces transfer between end effectors
  - Constraints for each effector can interact with each other
- The category of a task can change partway through!

With regards to the manipulators…

- The addition of a second manipulator constitutes an added, dynamic set of obstacles for each manipulator
- Imposes a whole new set of constraints upon the configuration space (more details later)
- Manipulators take on “roles” (leader + follower, fixed transformation…)

Information adapted from F. Krebs and T. Asfour, "A Bimanual Manipulation Taxonomy," in *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11031-11038, Oct. 2022, doi: 10.1109/LRA.2022.3196158

Methodologies

Methodologies

Control based method

- IK
- Control
- Manipulation

Policy learning methods

- ALOHA
- VLAs
- π_0

Kinematic-based methodologies

Forward Kinematics (FK)

- What is “forward kinematics”

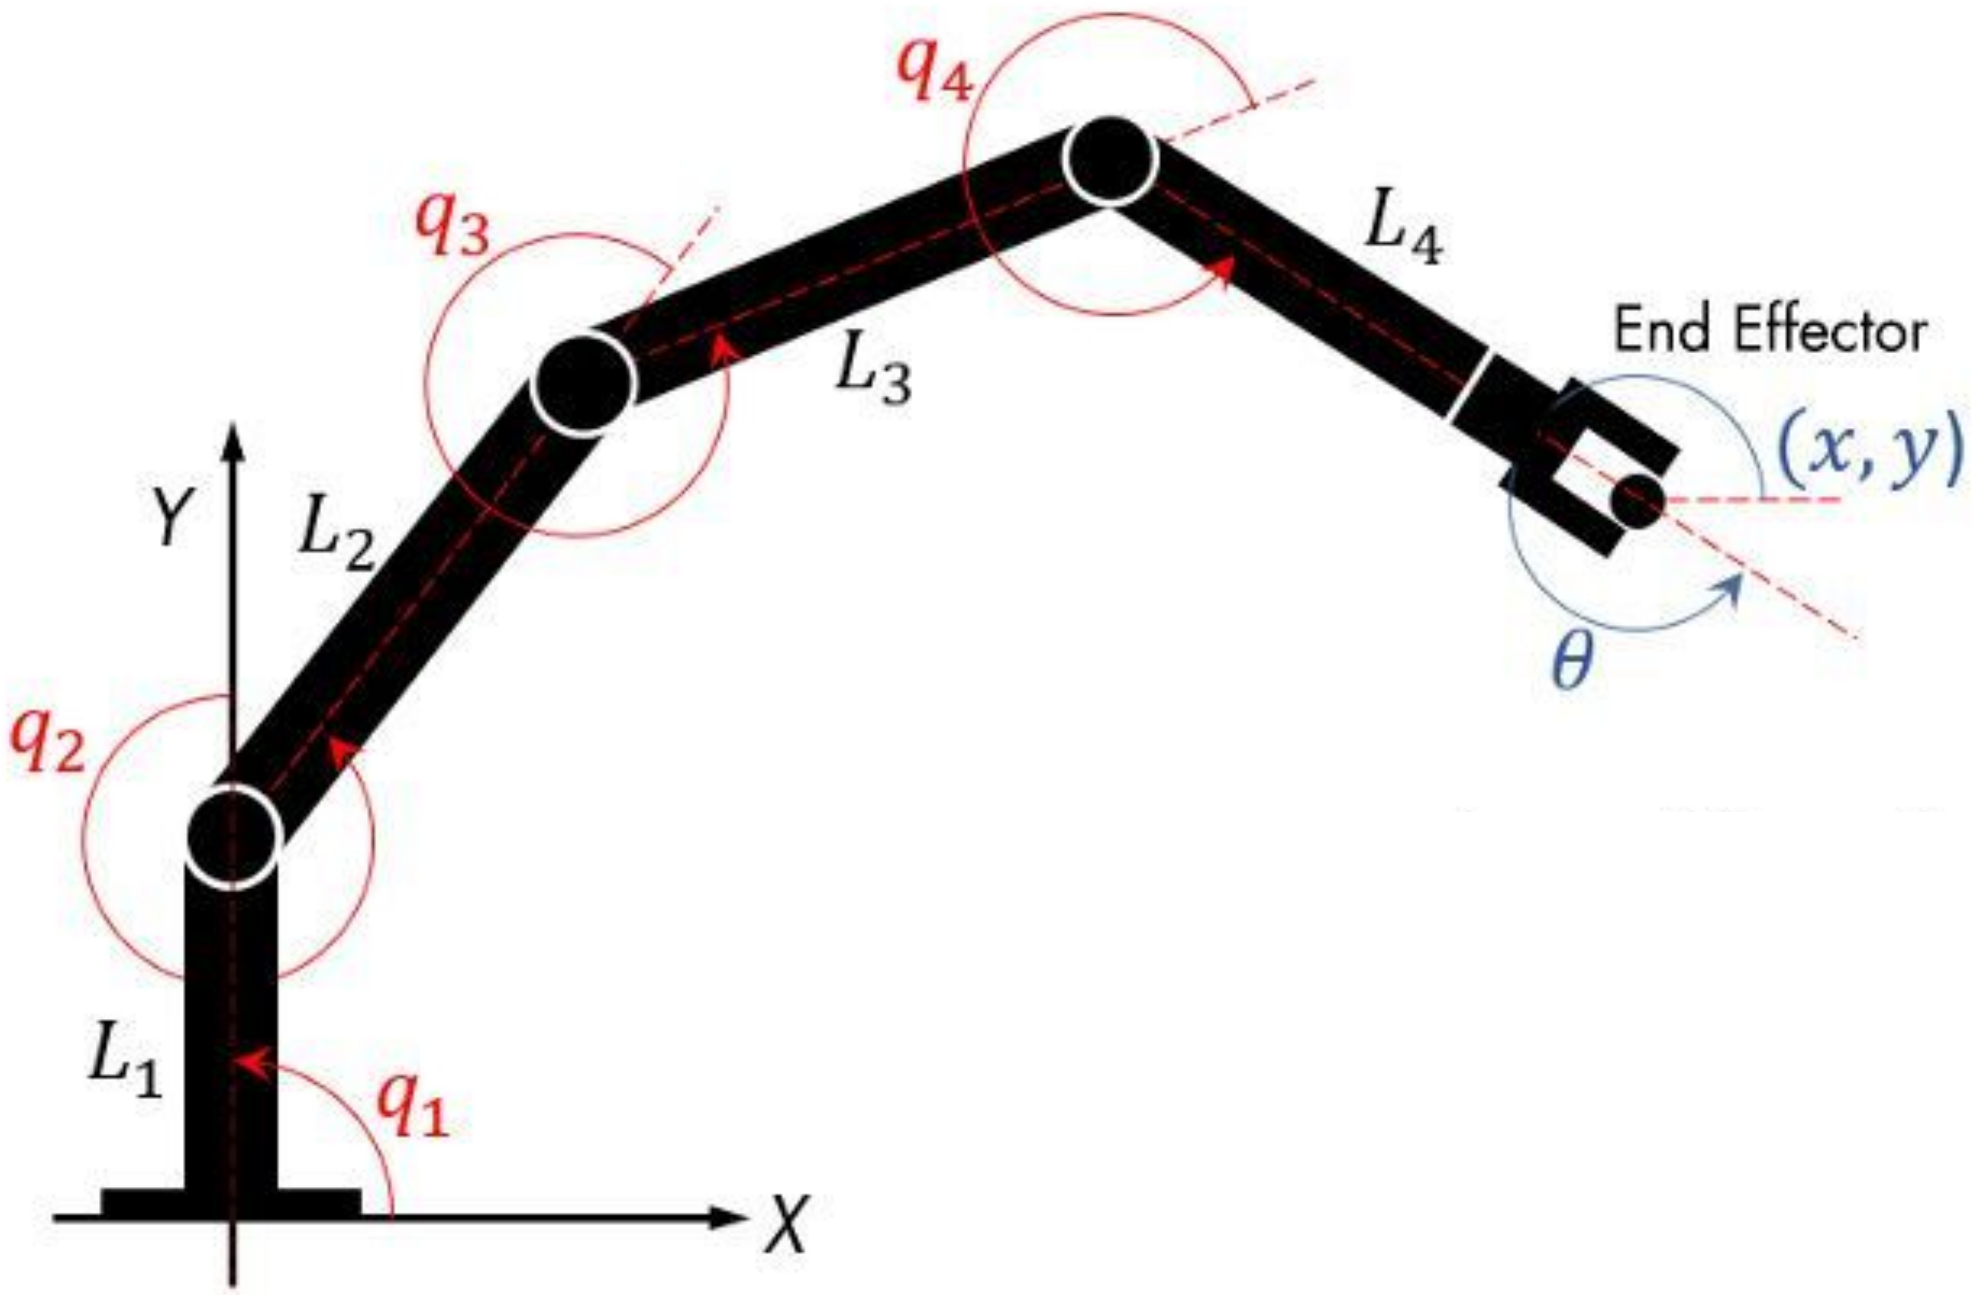

https://www.mathworks.com/discovery/inverse-kinematics.html

Forward Kinematics (FK)

- Given each joint position, determine the end effector pose
- Singular solution for each configuration

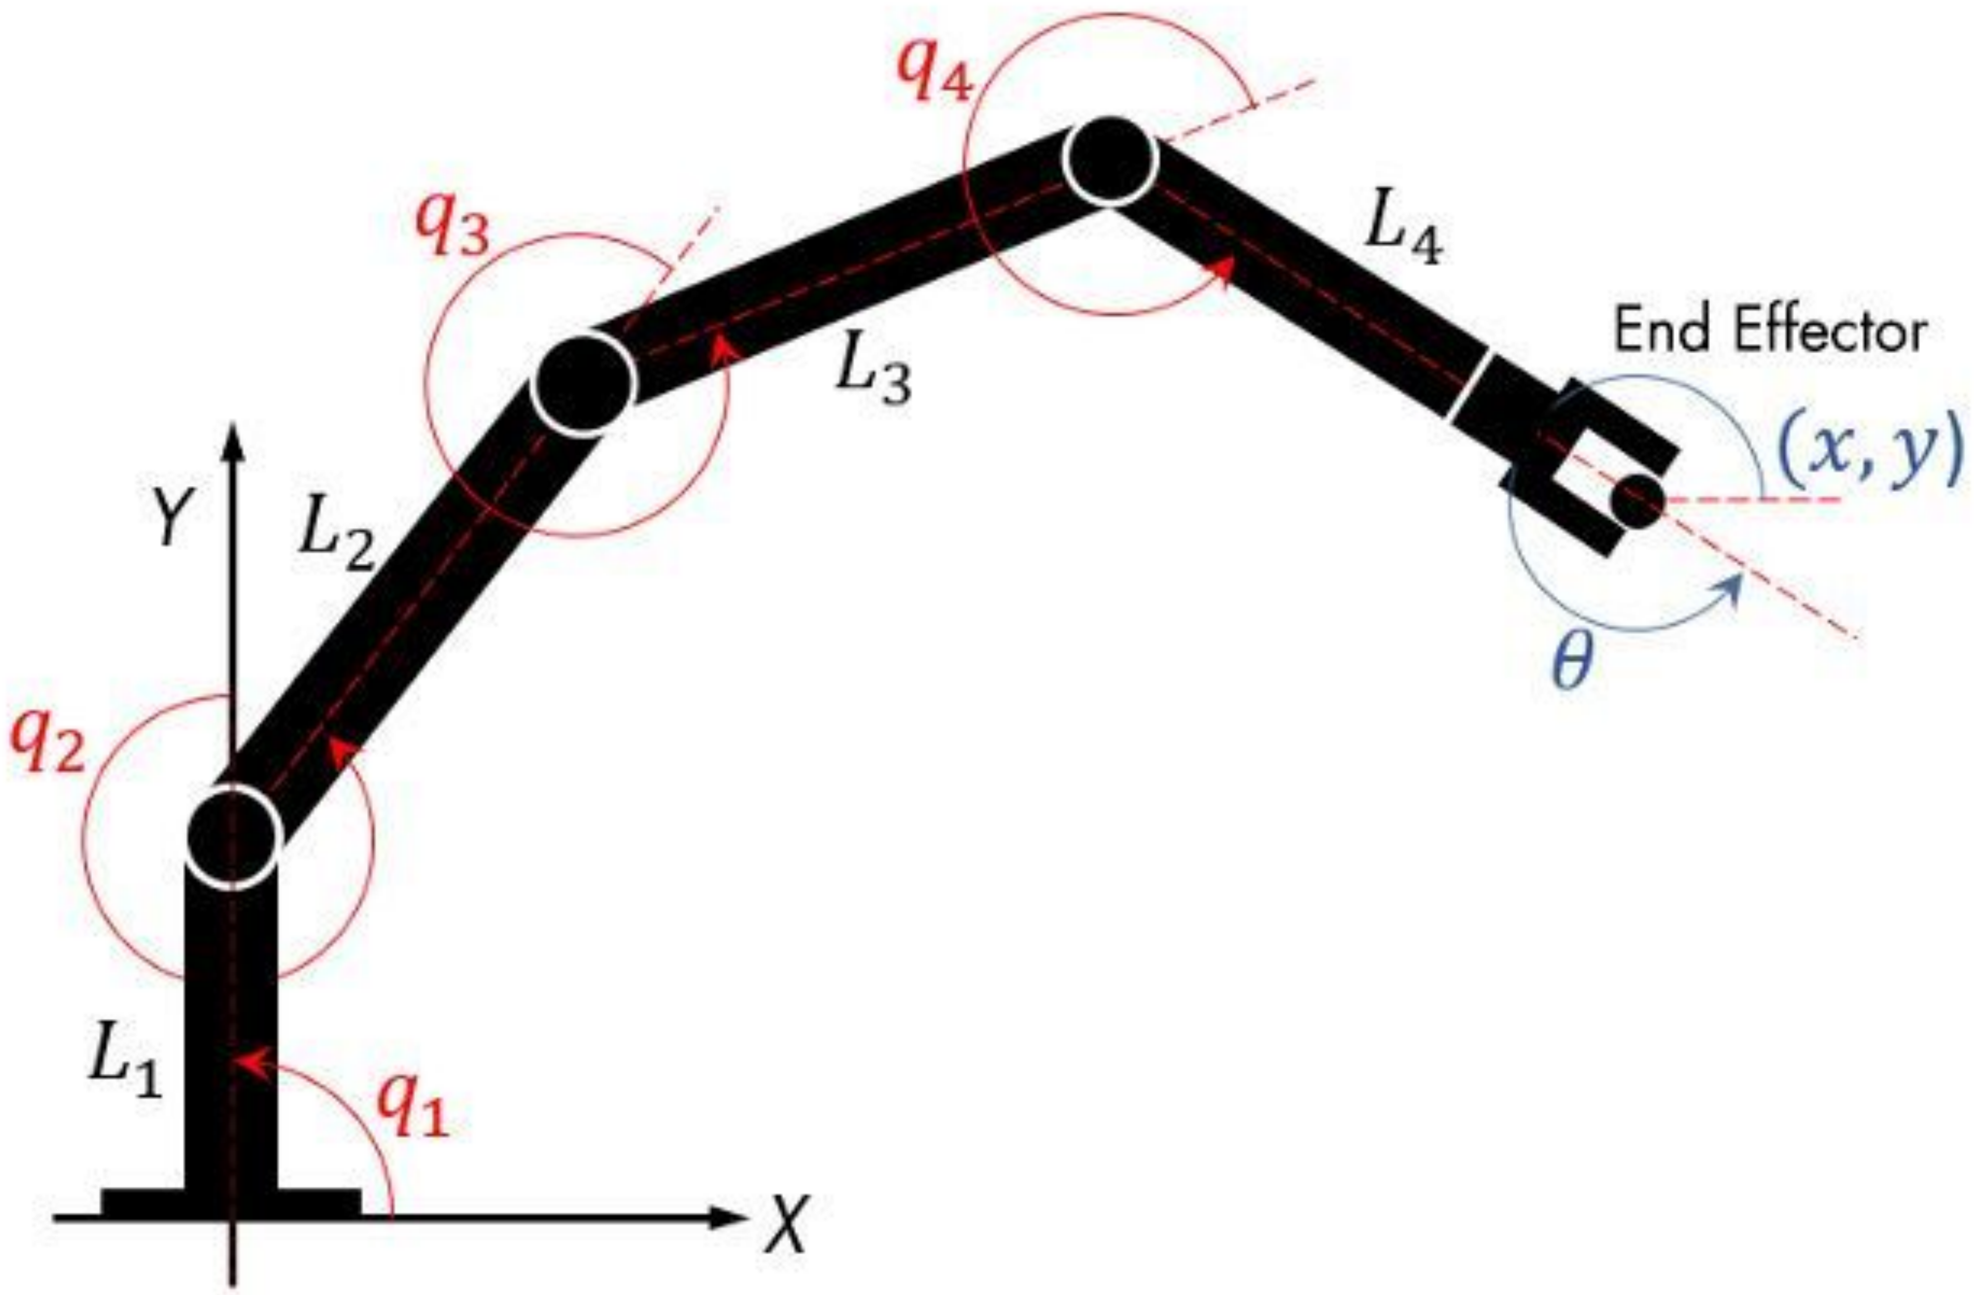

https://www.mathworks.com/discovery/inverse-kinematics.html

Forward Kinematics (FK)

- Given each joint position, determine the end effector pose
- Singular solution for each configuration

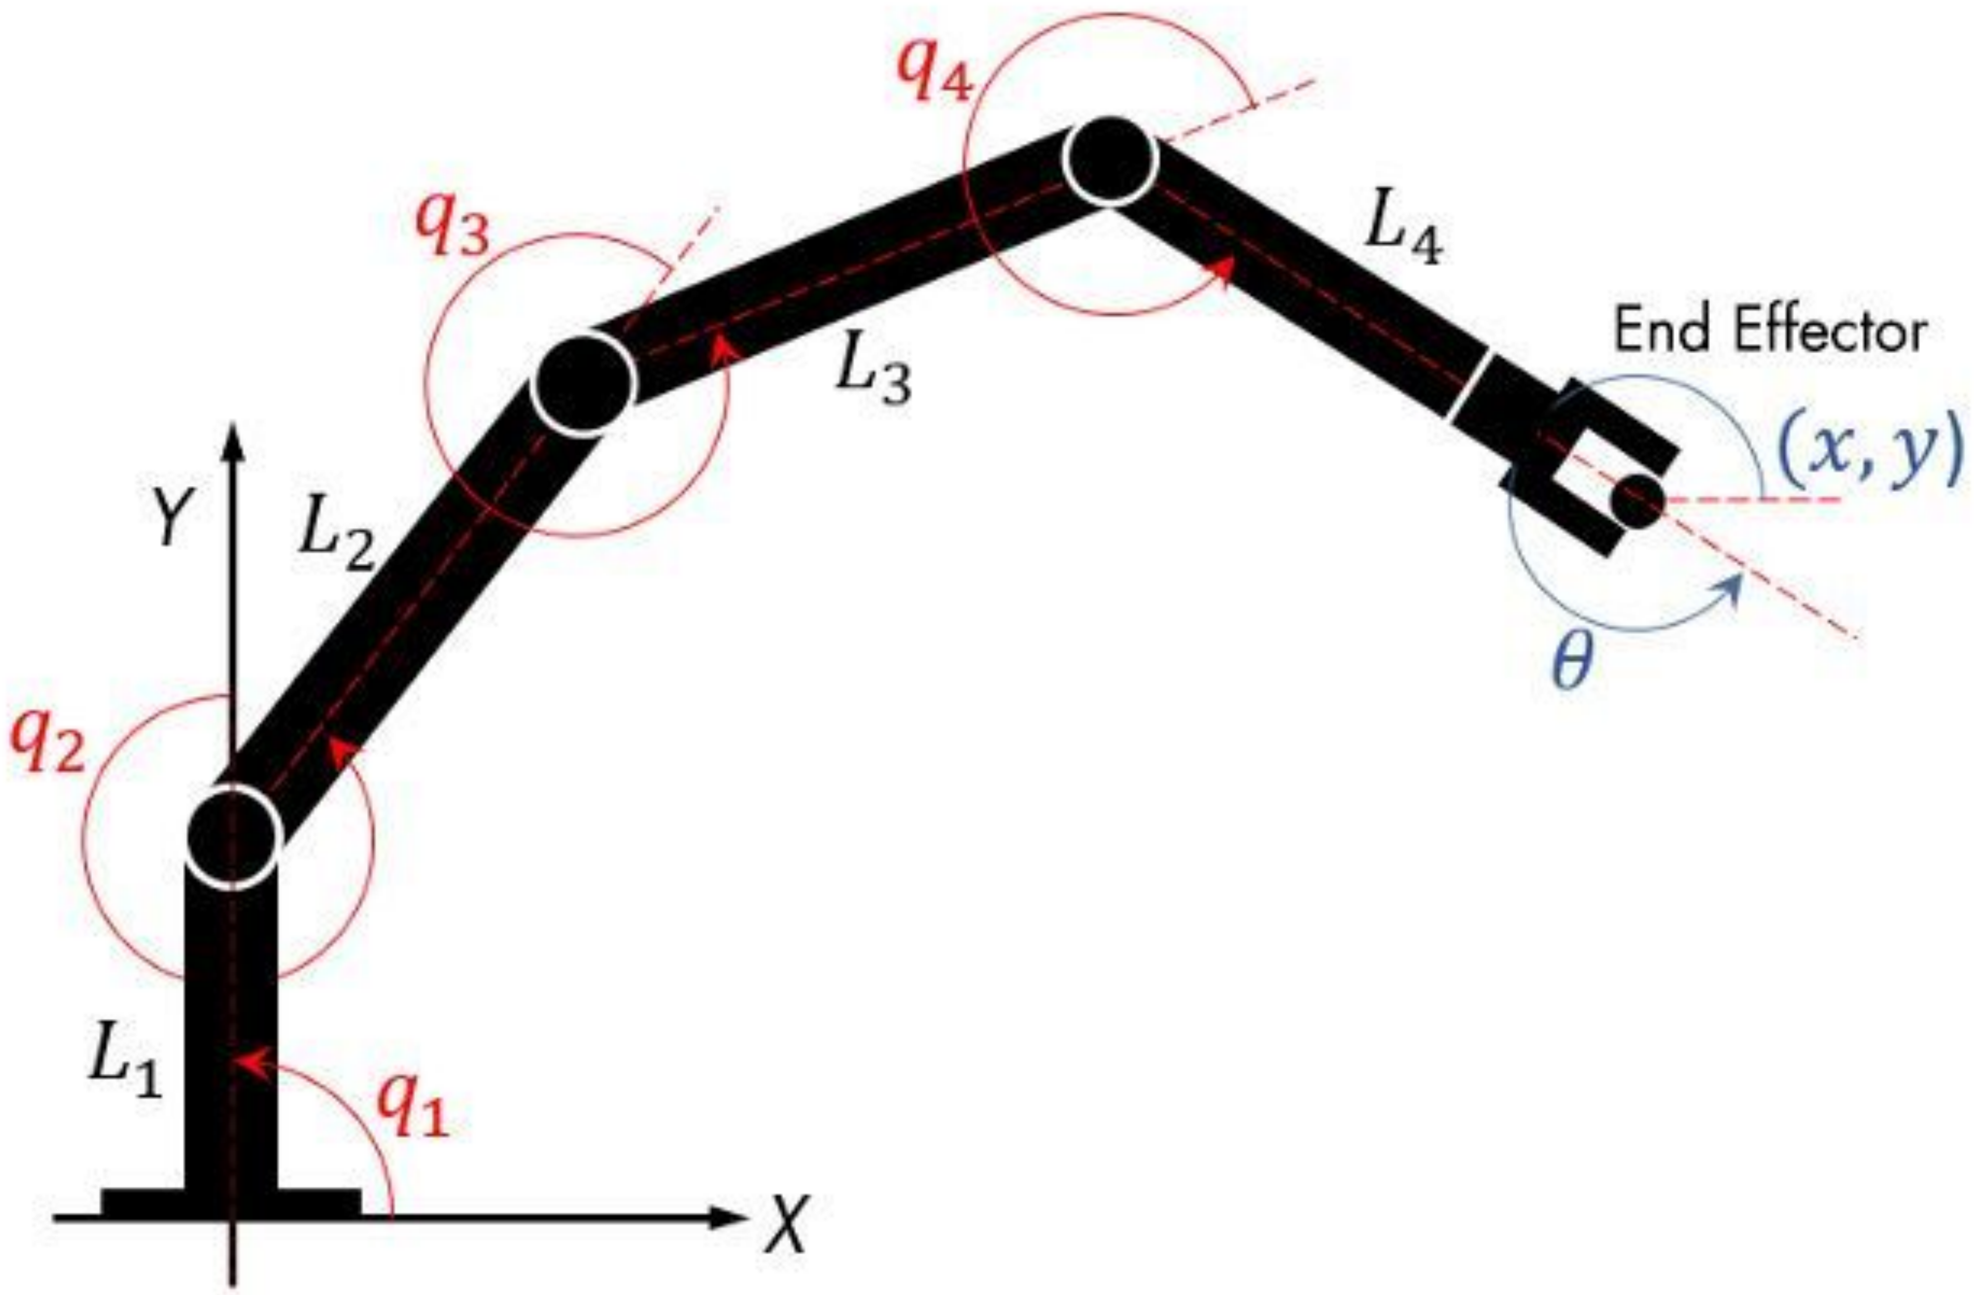

https://www.mathworks.com/discovery/inverse-kinematics.html

Forward Kinematics (FK)

- Each joint has its own coordinate frame
- Transformations between each join represented by homogeneous transformations

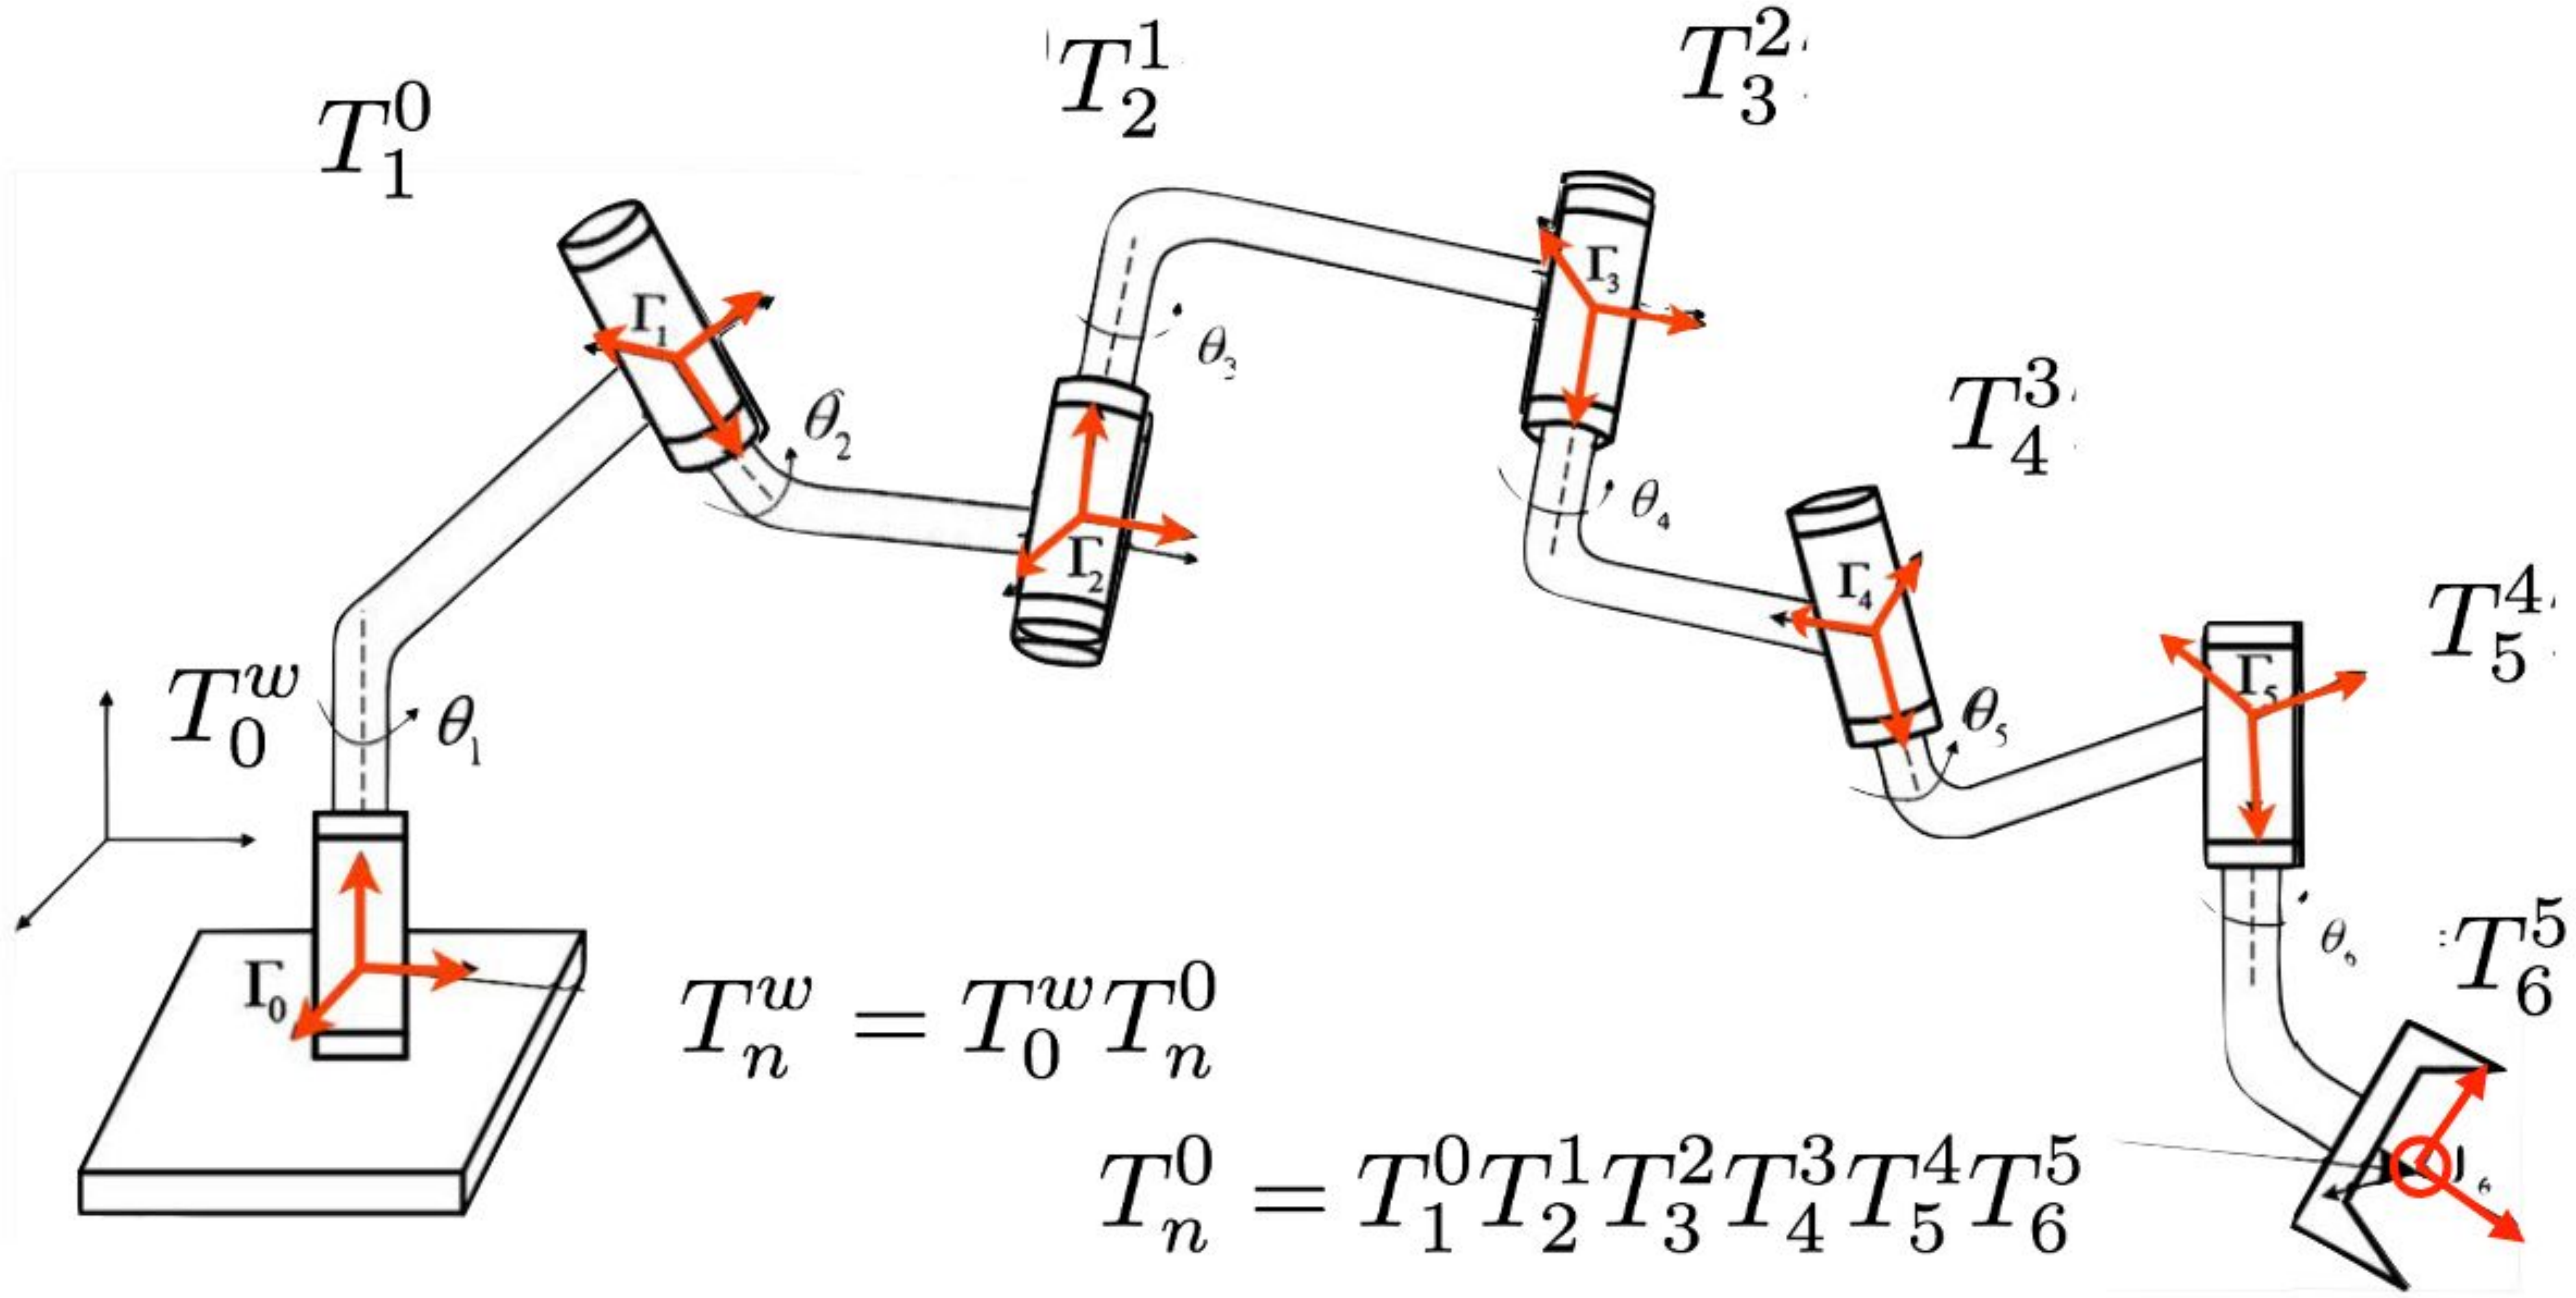

https://rpm-lab.github.io/CSCI5551-Spr24/assets/slides/lec06_manipulation_1_fk_decision_making.pdf

Inverse Kinematics (IK)

- What is “inverse kinematics”?

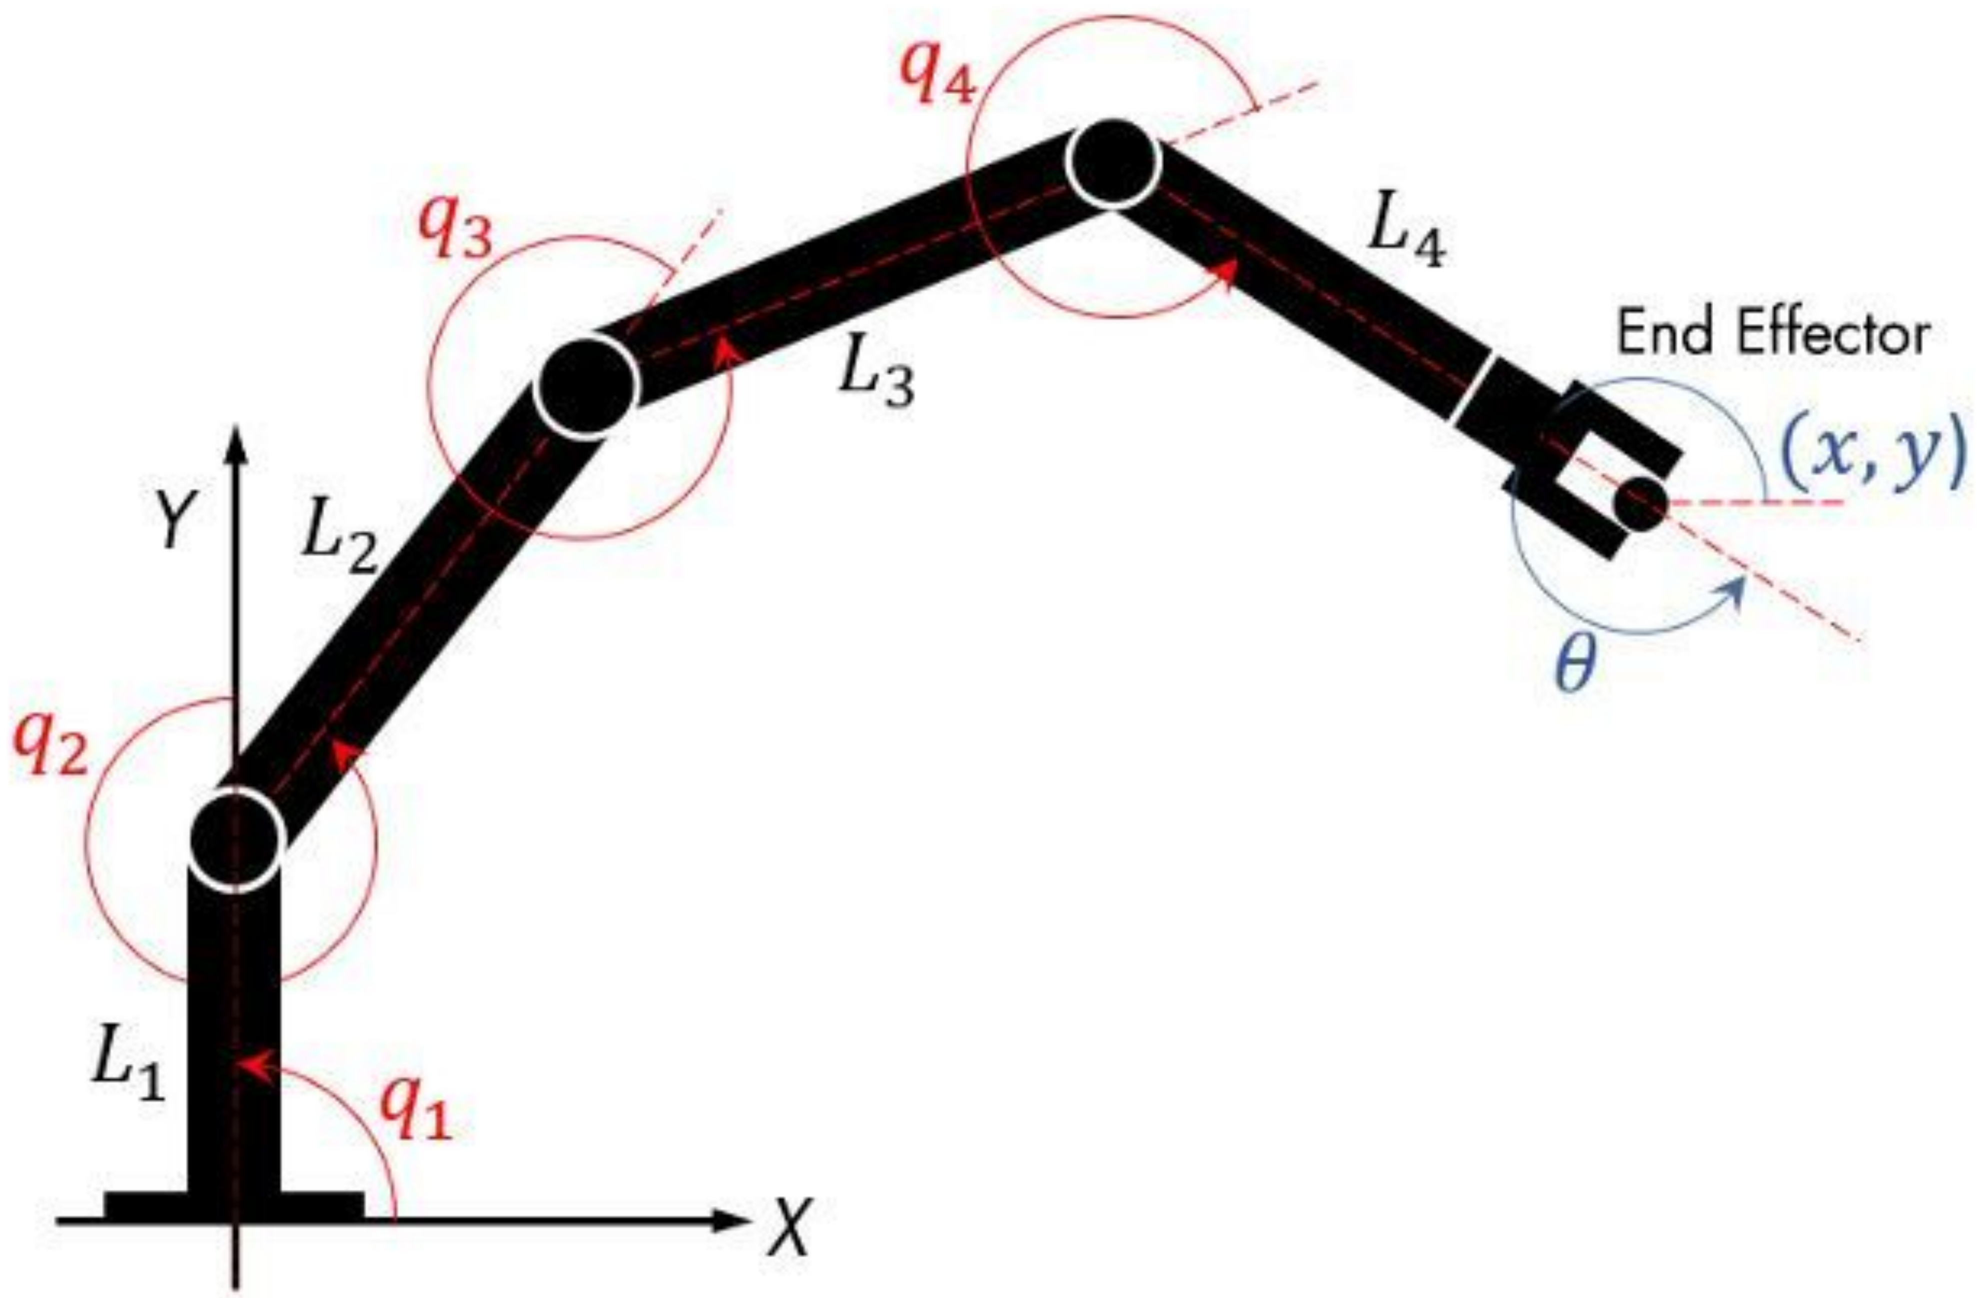

https://www.mathworks.com/discovery/inverse-kinematics.html

Inverse Kinematics (IK)

- Given the end effector pose, determine each joint position
- Can be many solutions for each joint position
- Summarized by these equations (numerical):

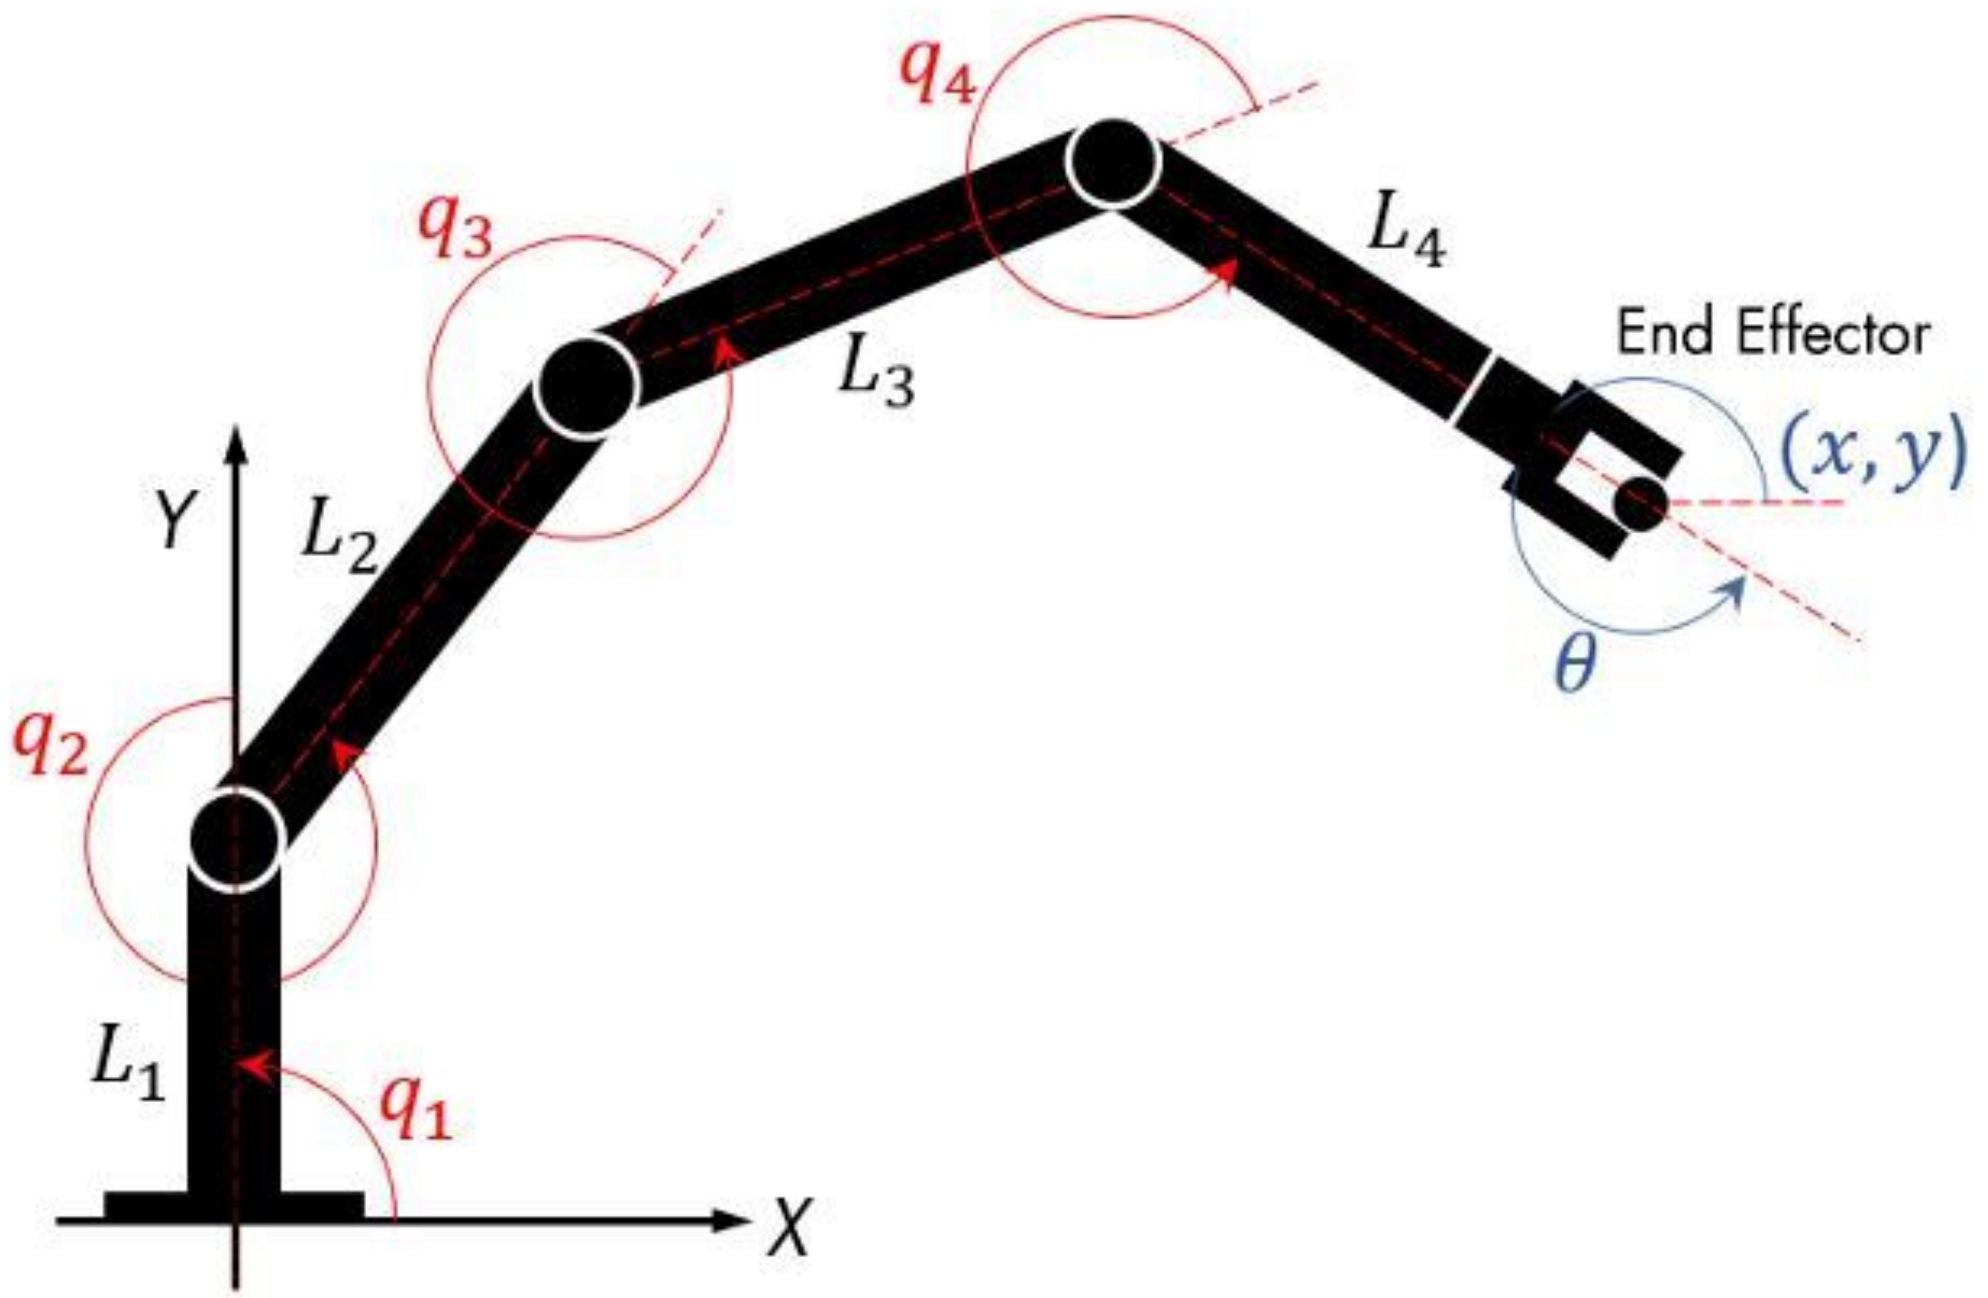

https://www.mathworks.com/discovery/inverse-kinematics.html

Inverse Kinematics (IK)

- Given the end effector pose, determine each joint position
- Can be many solutions for each joint position
- Summarized by these equations (numerical):

$$\Delta \mathbf{x}_n = \mathbf{x}_d - \mathbf{x}_n$$

Start with error from end point

https://rpm-lab.github.io/CSCI5551-Spr24/assets/slides/lec08_manipulation_3_ik_jacobian.pdf

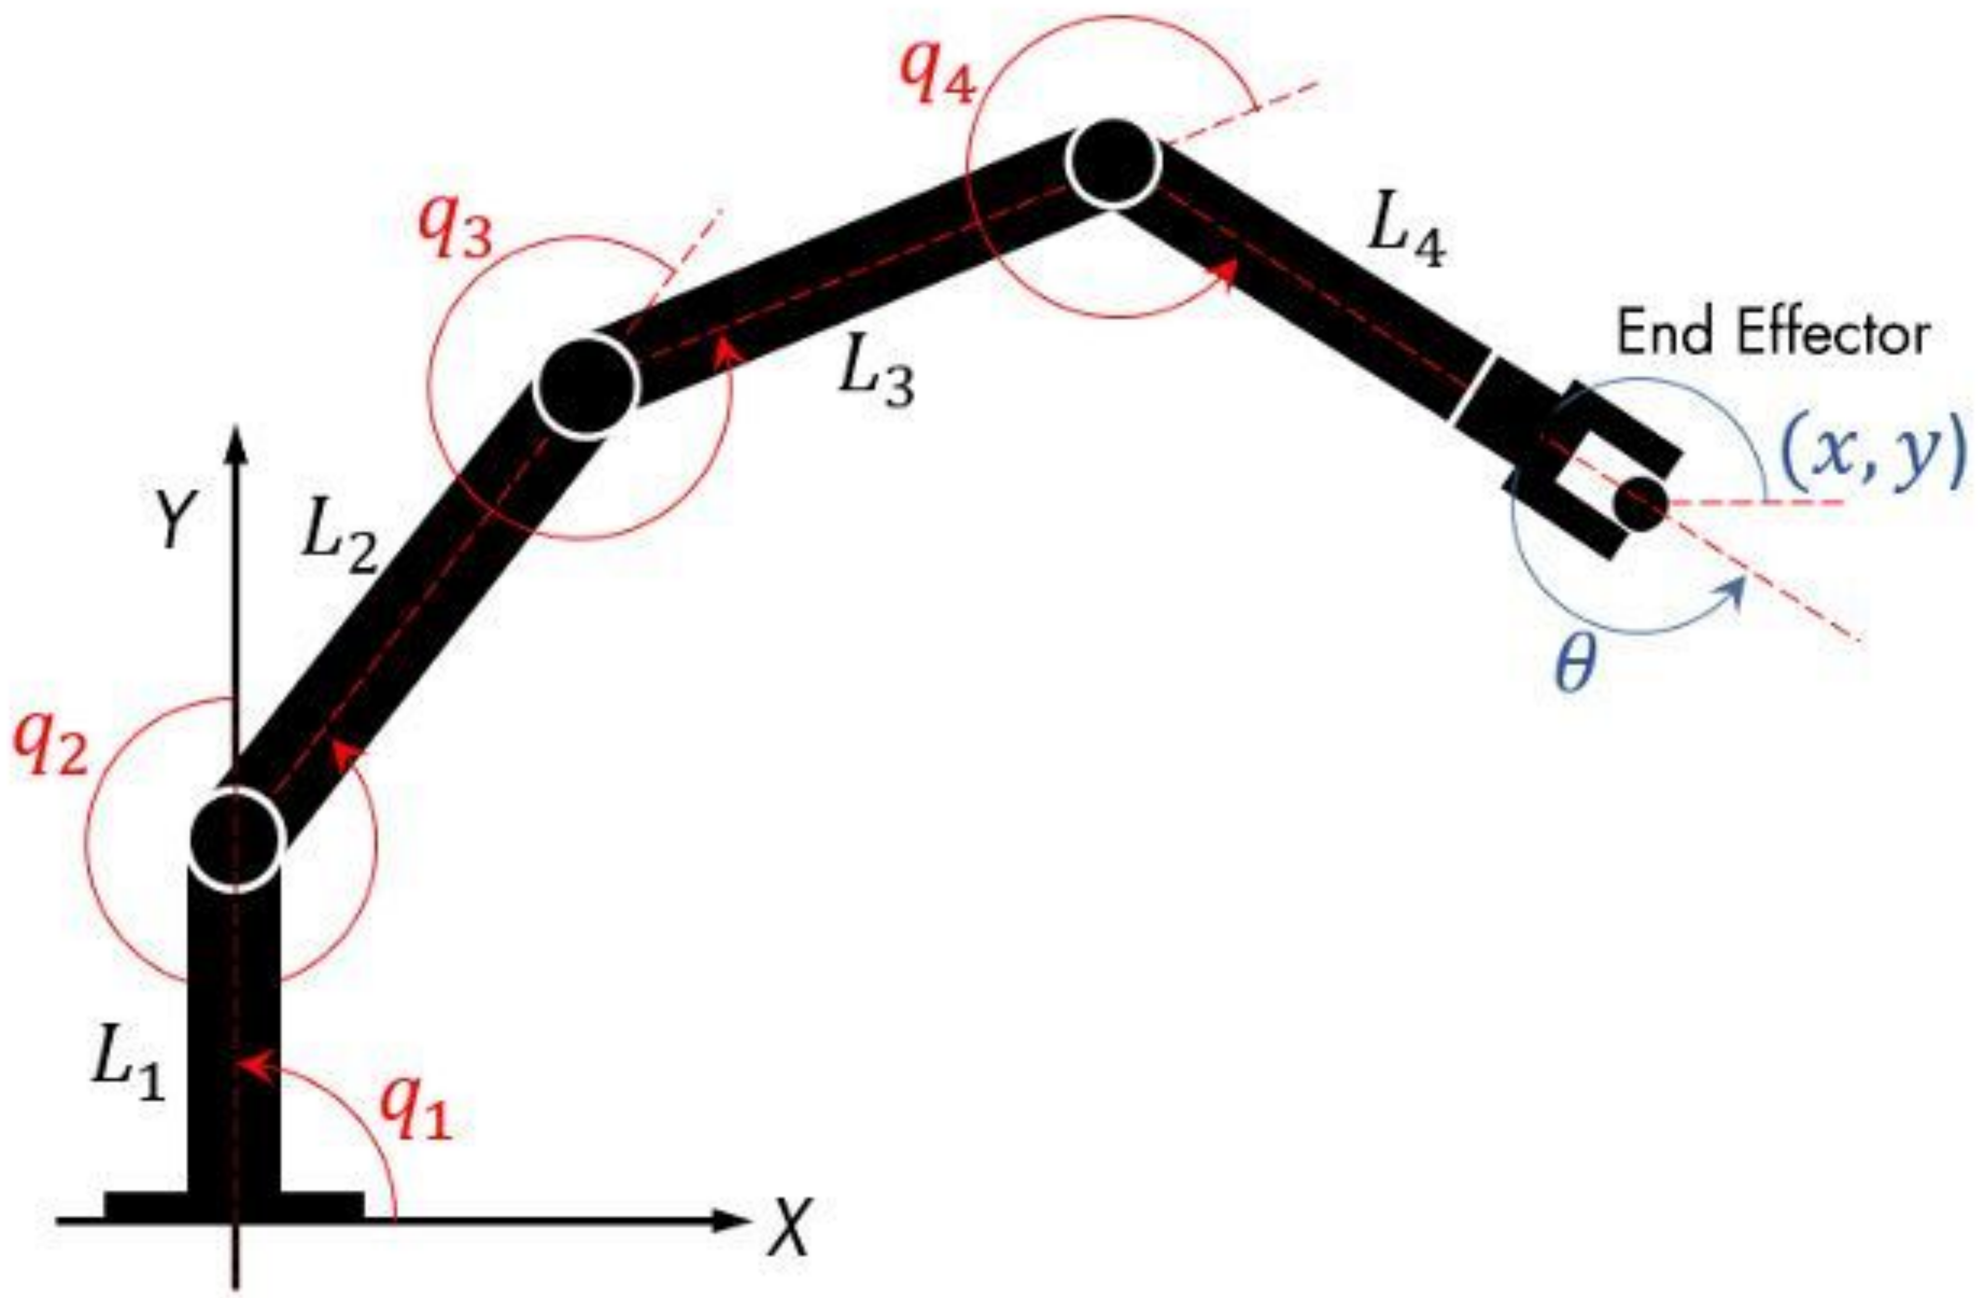

https://www.mathworks.com/discovery/inverse-kinematics.html

Inverse Kinematics (IK)

- Given the end effector pose, determine each joint position
- Can be many solutions for each joint position
- Summarized by these equations (numerical):

$$\Delta \mathbf{x}_n = \mathbf{x}_d - \mathbf{x}_n$$

$$\Delta \mathbf{q}_n = J(\mathbf{q}_n)^{-1} \Delta \mathbf{x}_n$$

Find the direction to move

https://rpm-lab.github.io/CSCI5551-Spr24/assets/slides/lec08_manipulation_3_ik_jacobian.pdf

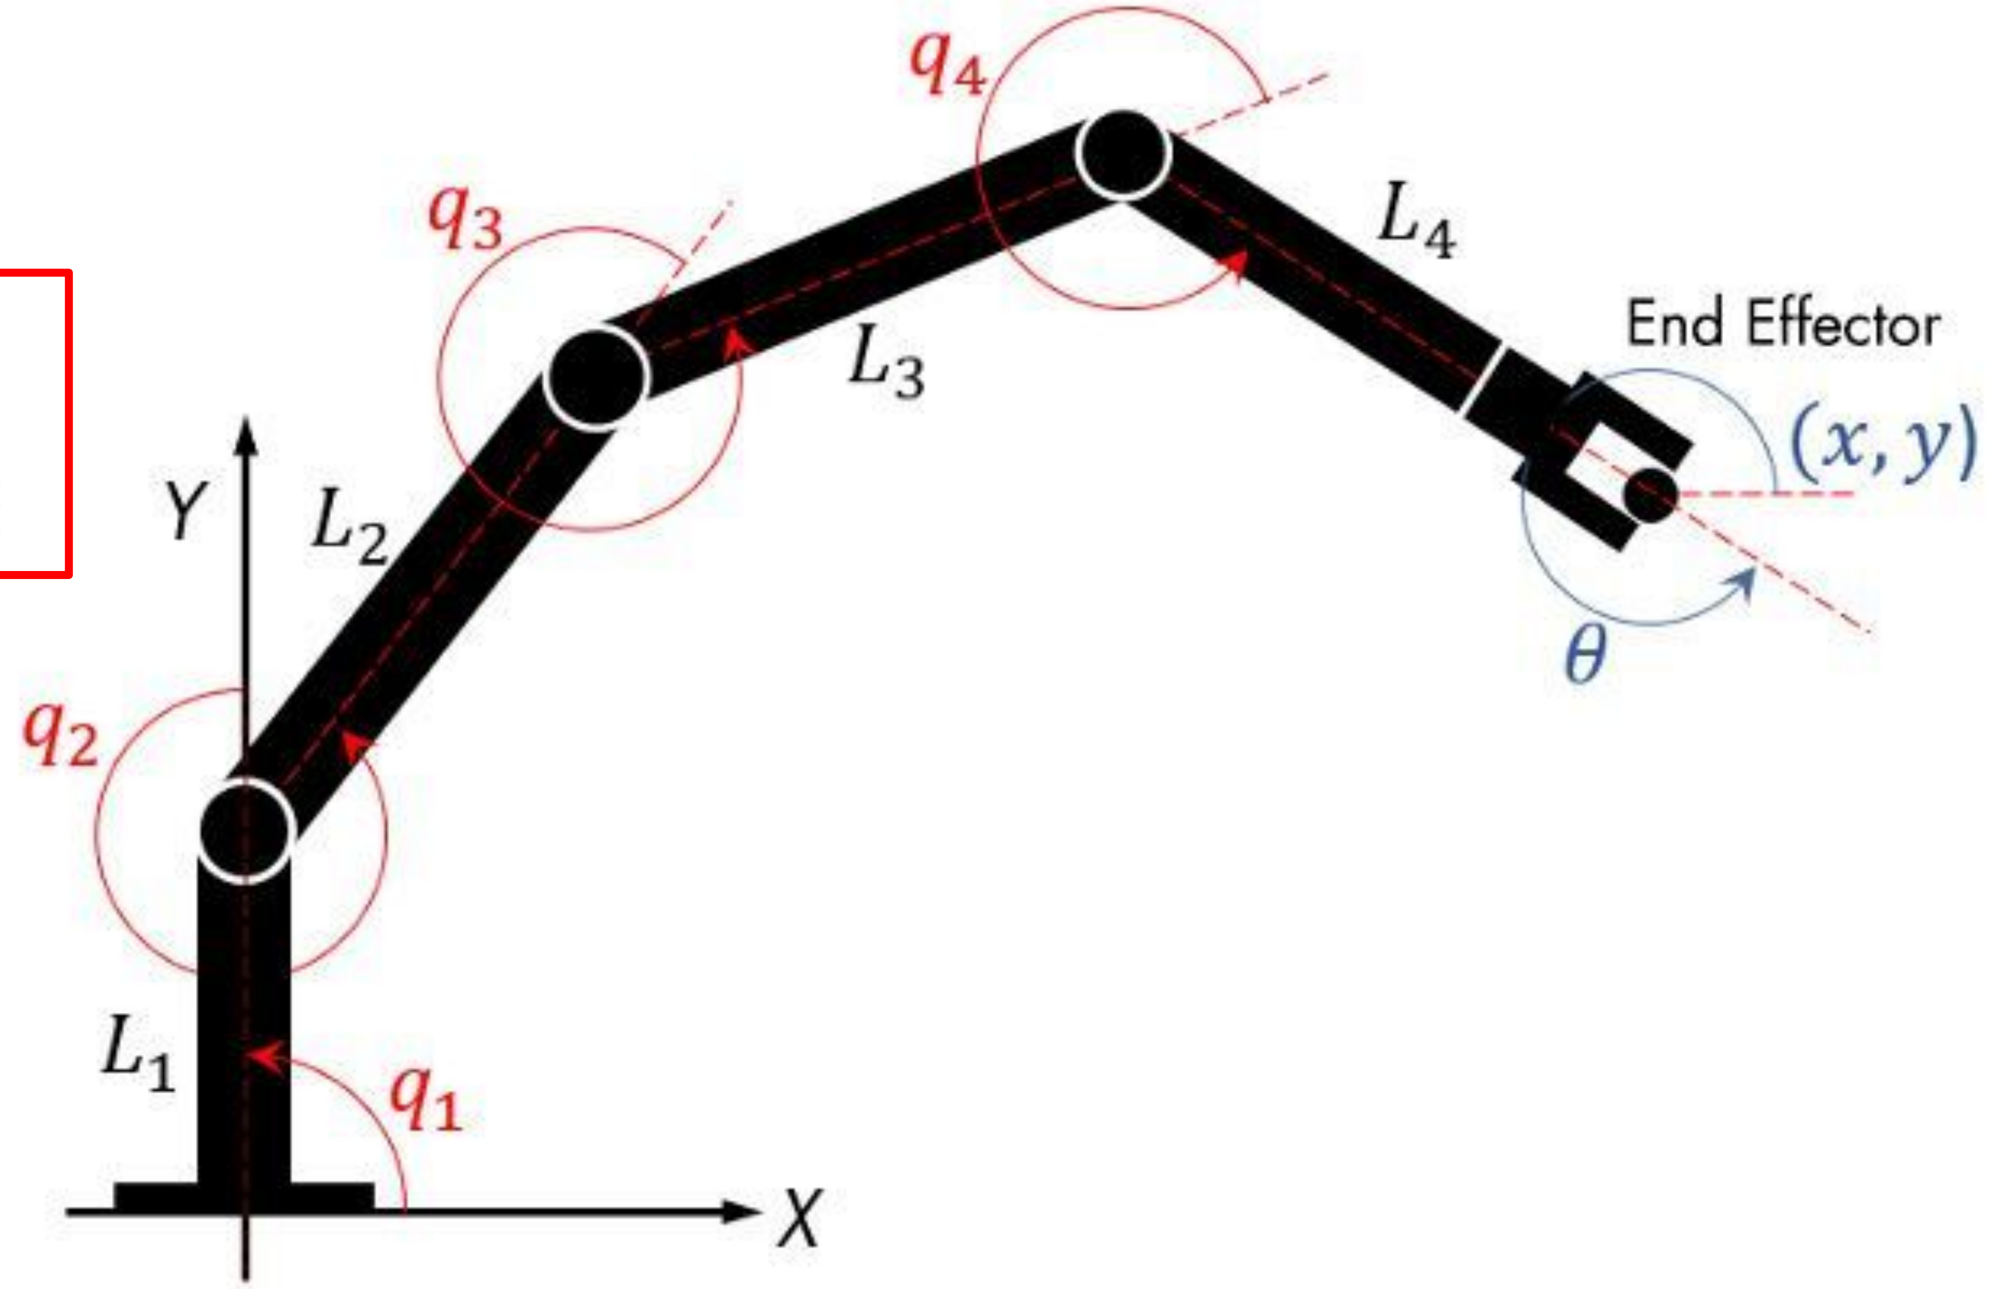

https://www.mathworks.com/discovery/inverse-kinematics.html

Inverse Kinematics (IK)

- Given the end effector pose, determine each joint position
- Can be many solutions for each joint position
- Summarized by these equations (numerical):

$$\Delta \mathbf{x}_n = \mathbf{x}_d - \mathbf{x}_n$$

$$\Delta \mathbf{q}_n = J(\mathbf{q}_n)^{-1} \Delta \mathbf{x}_n$$

$$\mathbf{q}_{n+1} = \mathbf{q}_n + \gamma \Delta \mathbf{q}_n$$

https://rpm-lab.github.io/CSCI5551-Spr24/assets/slides/lec08_manipulation_3_ik_jacobian.pdf

Take a step in that direction!

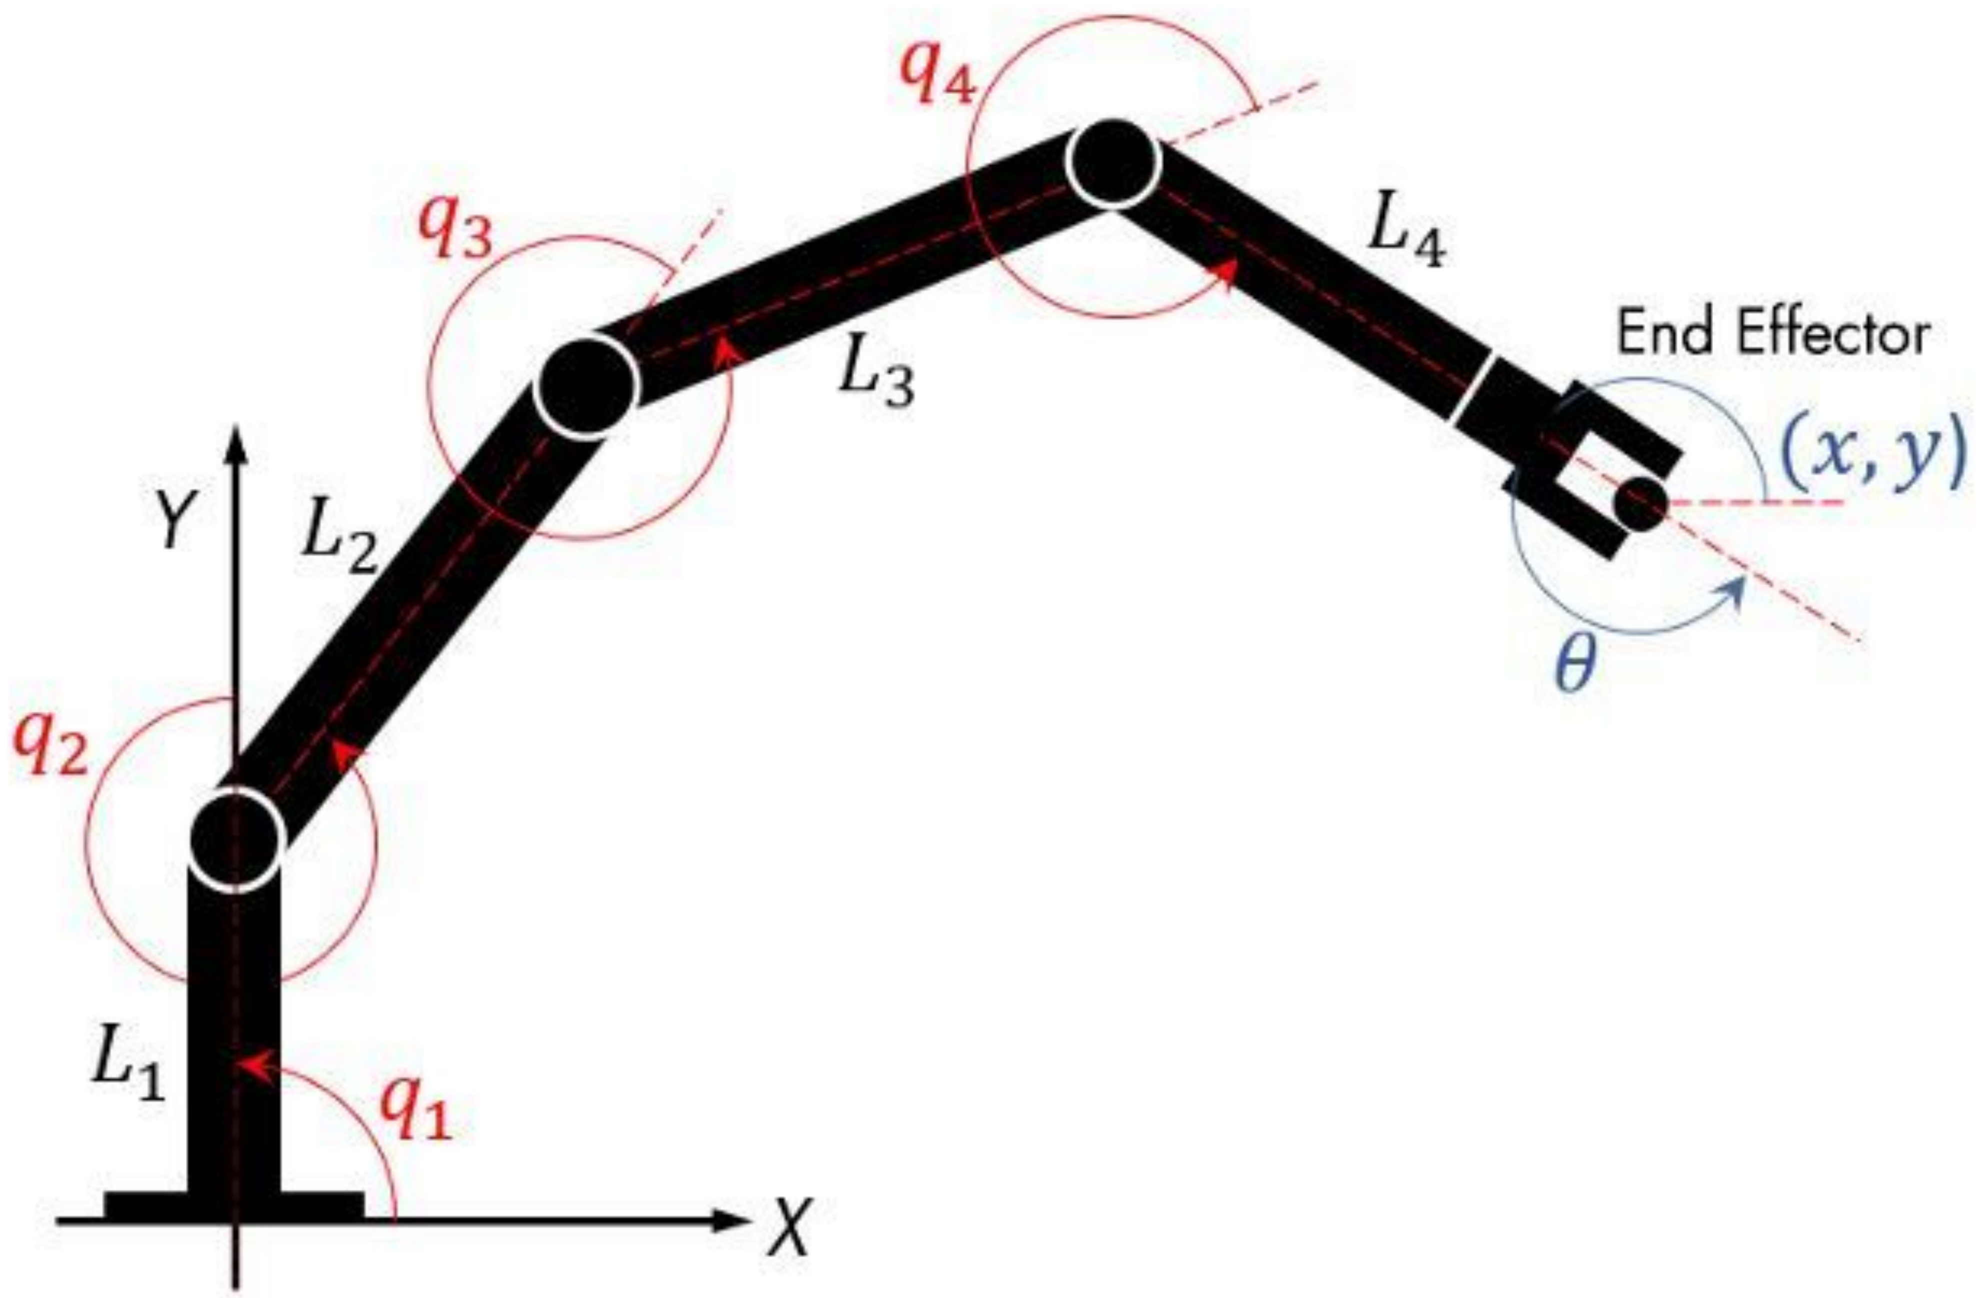

https://www.mathworks.com/discovery/inverse-kinematics.html

Forward vs Inverse Kinematics

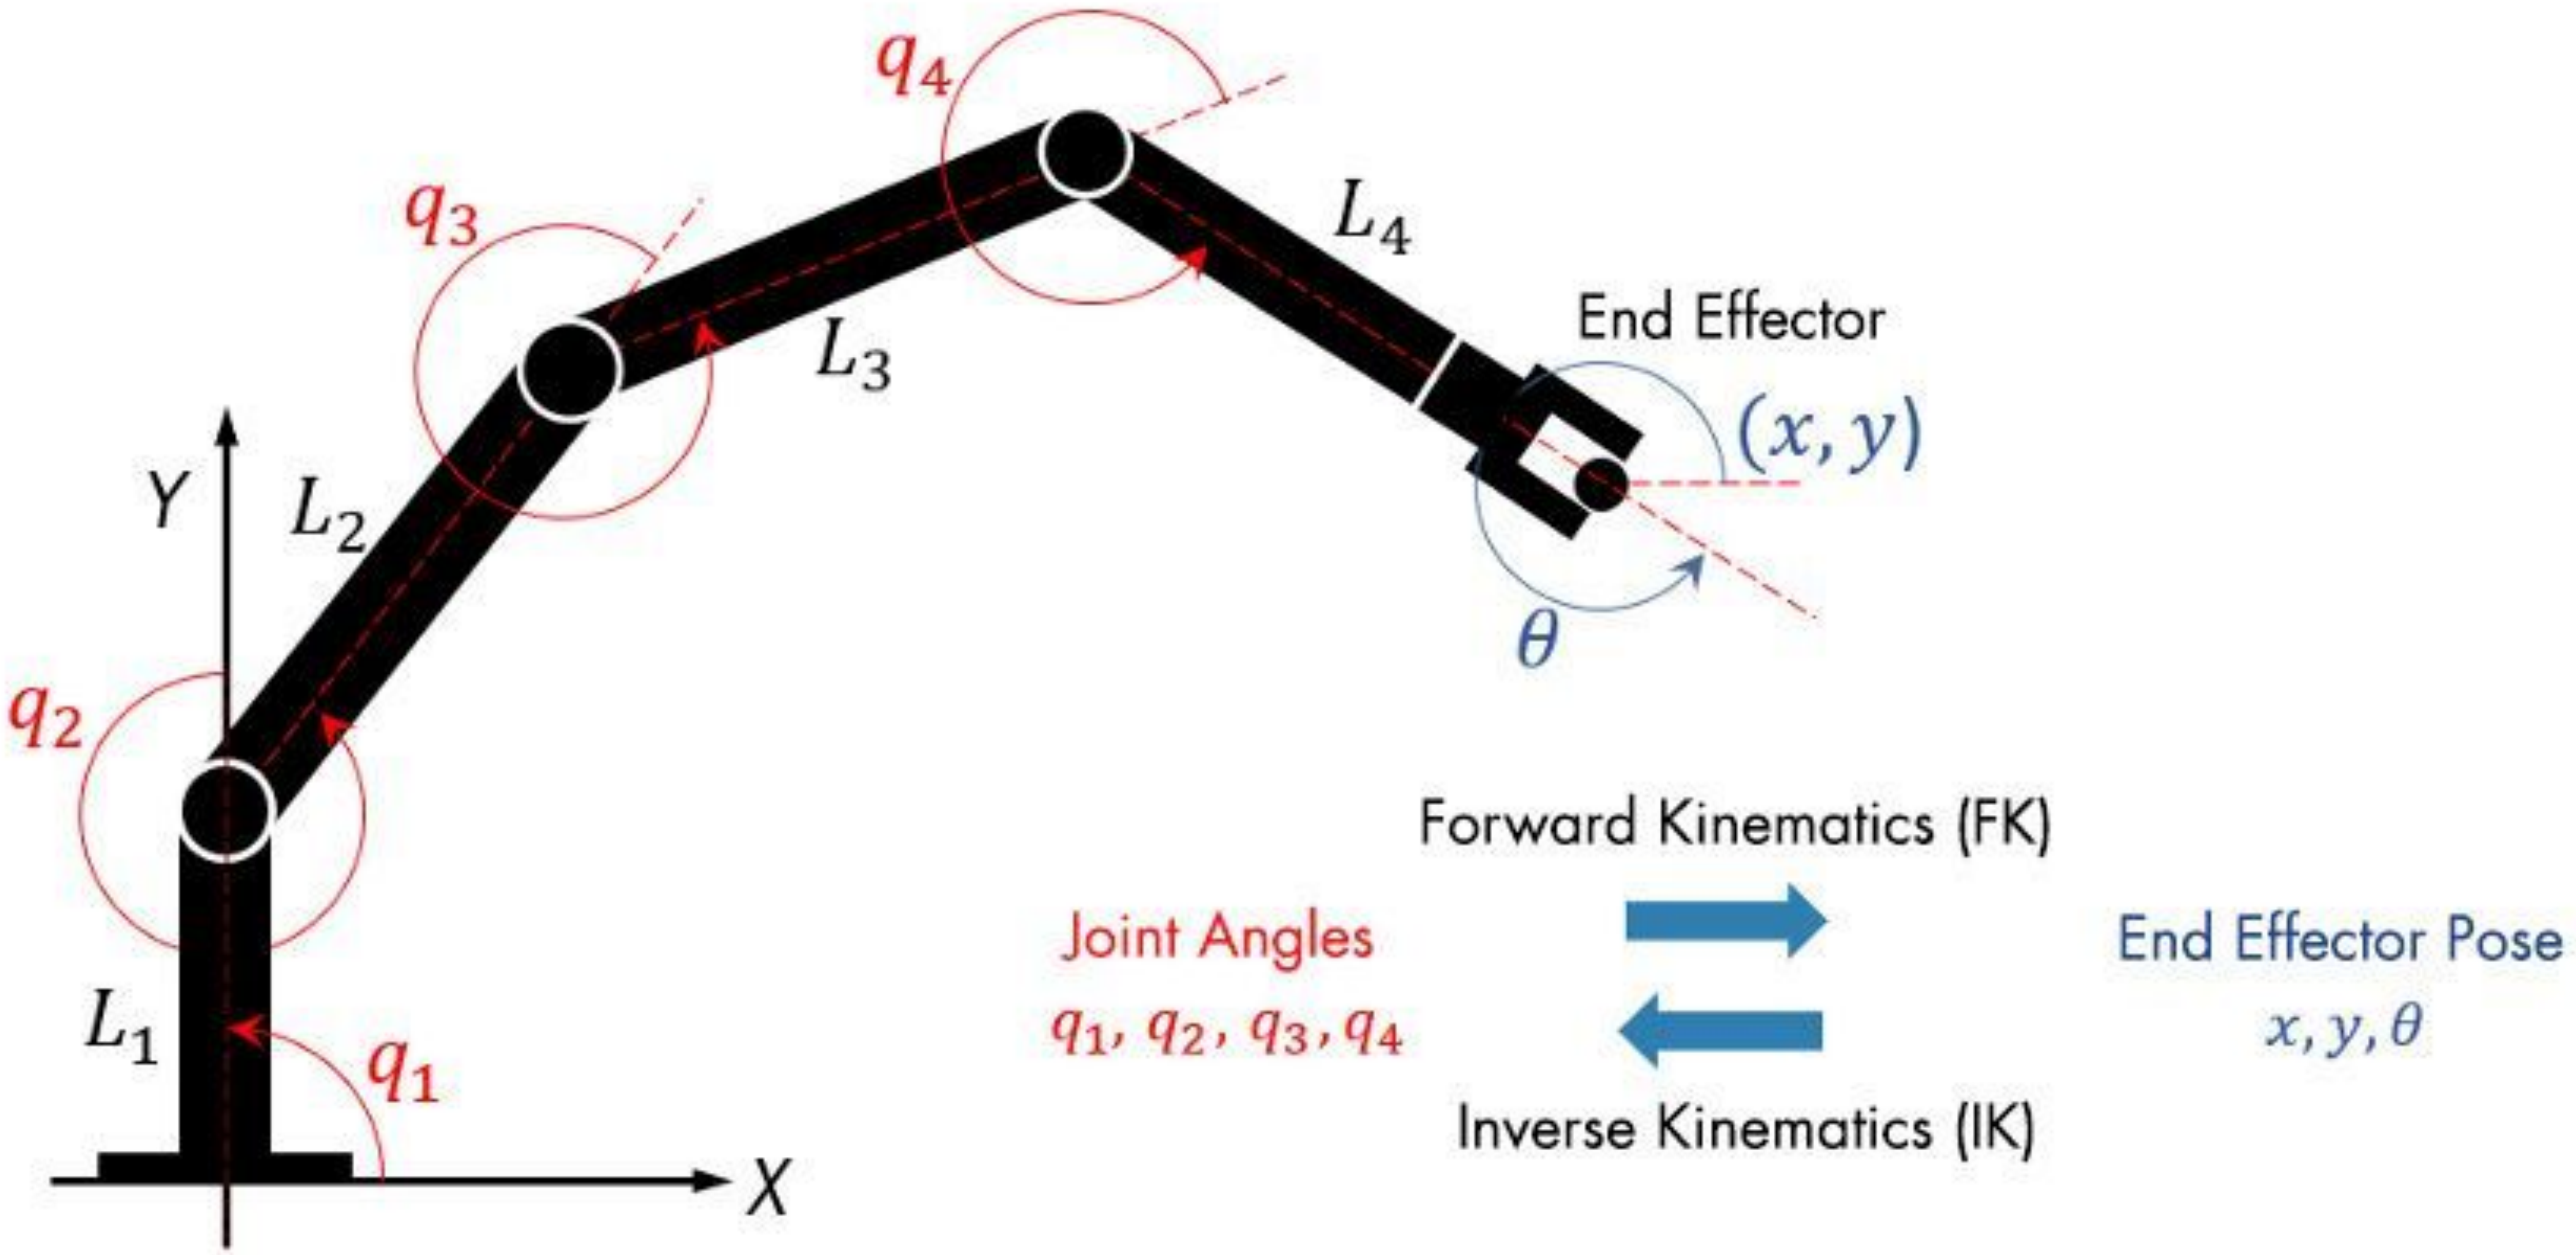

https://www.mathworks.com/discovery/inverse-kinematics.html

In the case of dual arm manipulation, is it as simple as applying IK to both arms?

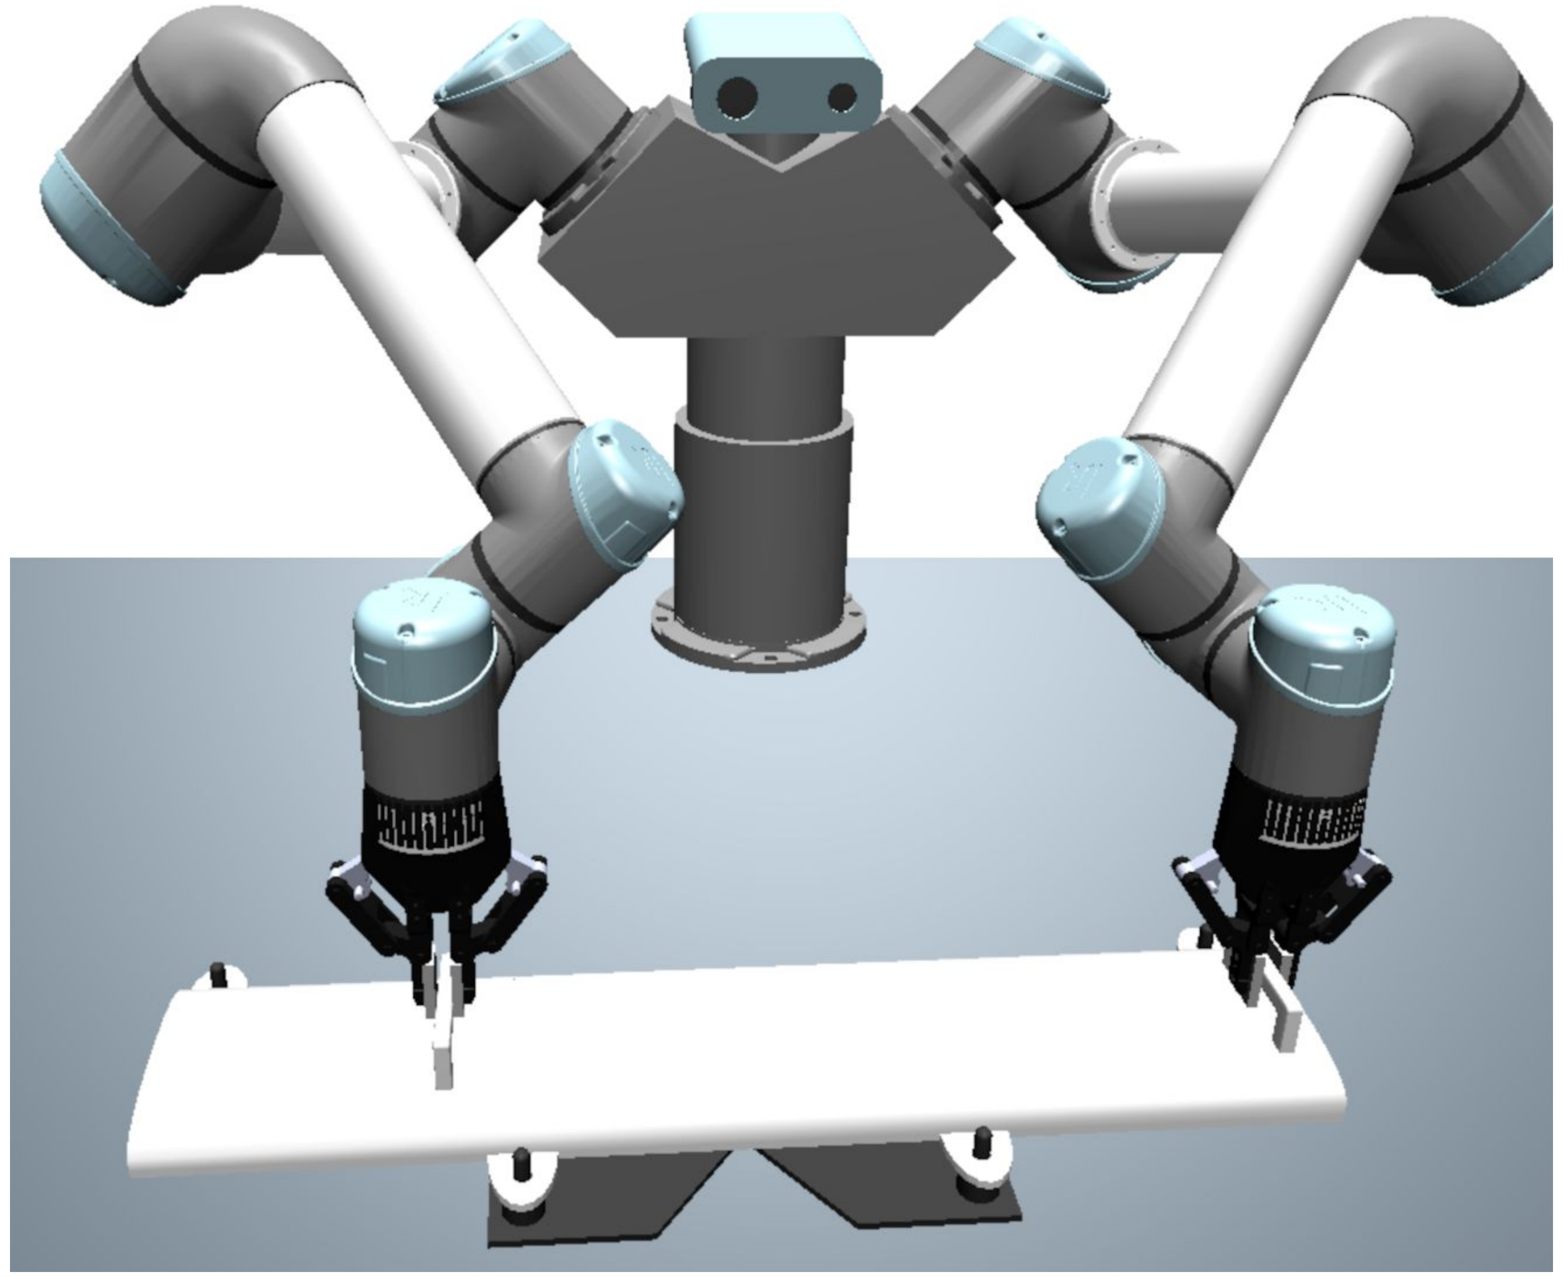

https://www.aimodels.fyi/papers/arxiv/peract2-benchmarking-learning-robotic-bimanual-manipulation-tasks

In the case of dual arm manipulation, is it as simple as applying IK to both arms?

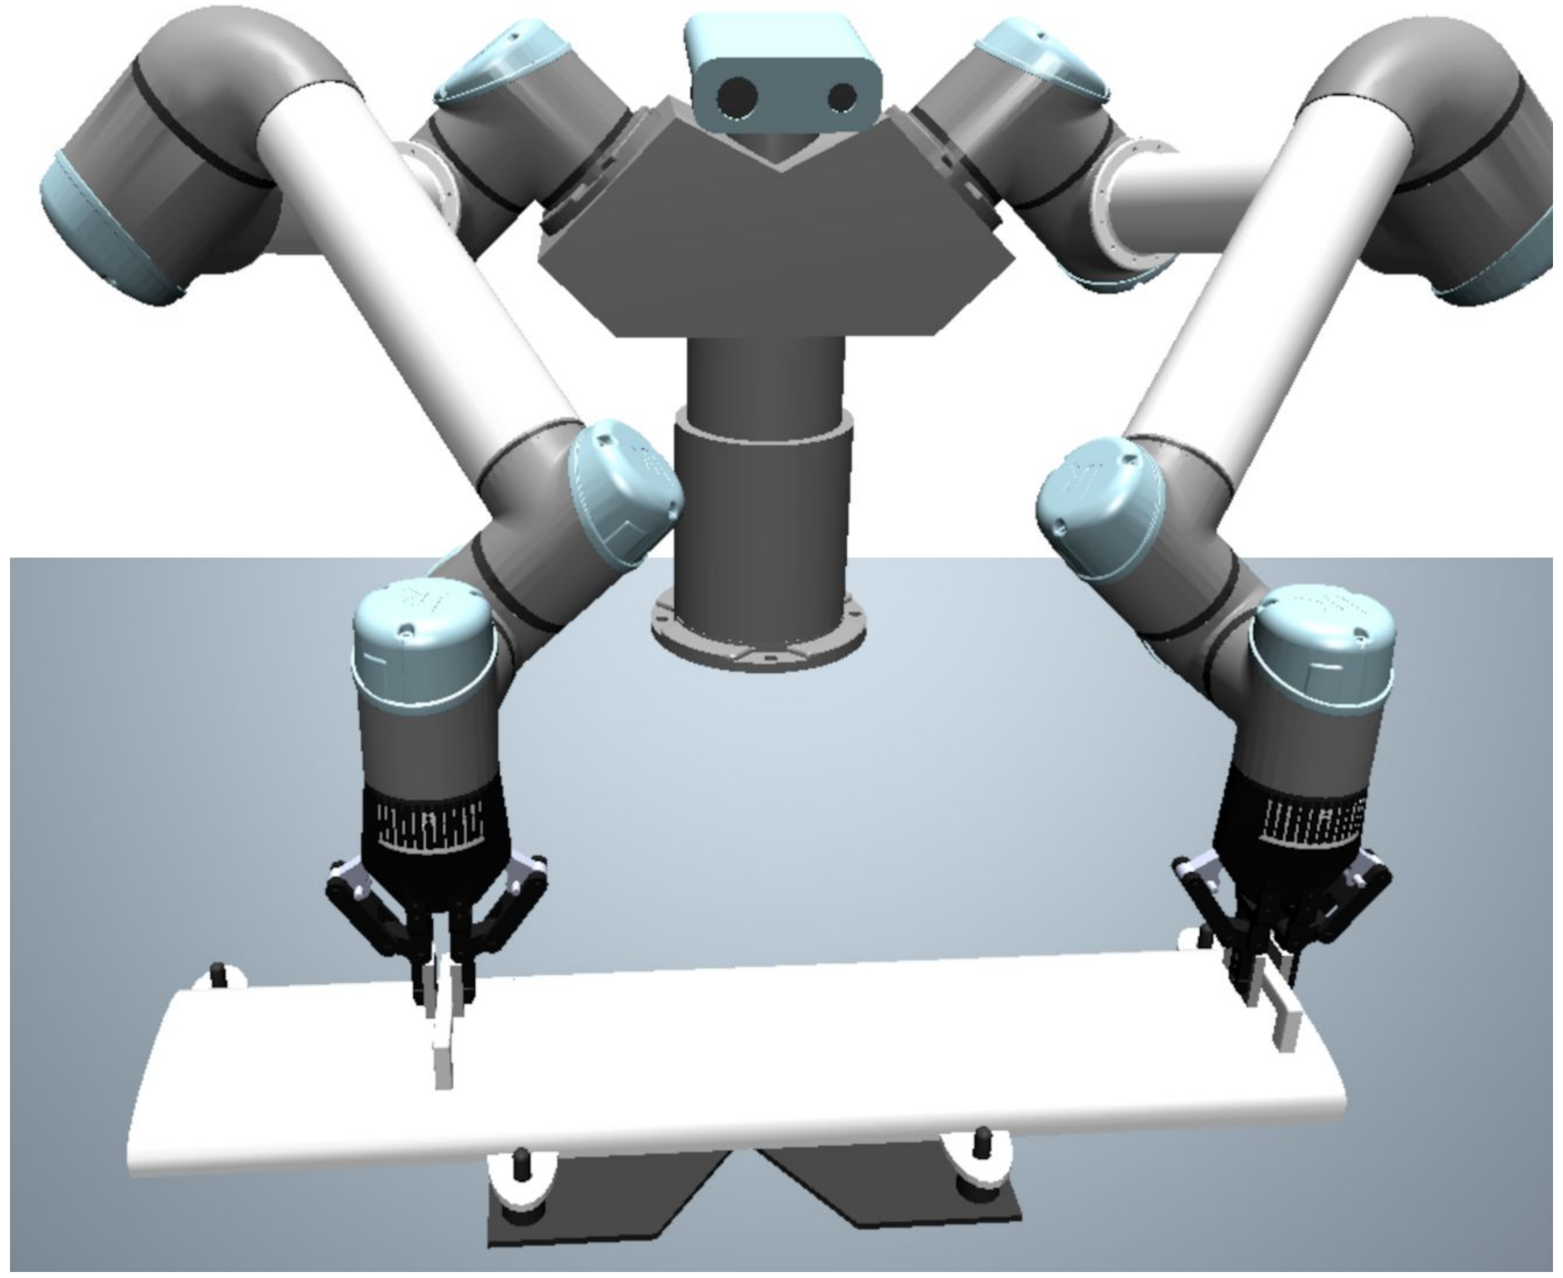

Maybe a bit more involved…

https://www.aimodels.fyi/papers/arxiv/peract2-benchmarking-learning-robotic-bimanual-manipulation-tasks

Bimanual manipulation kinematics
Challenges:

- The transform between the end effectors must remain fixed

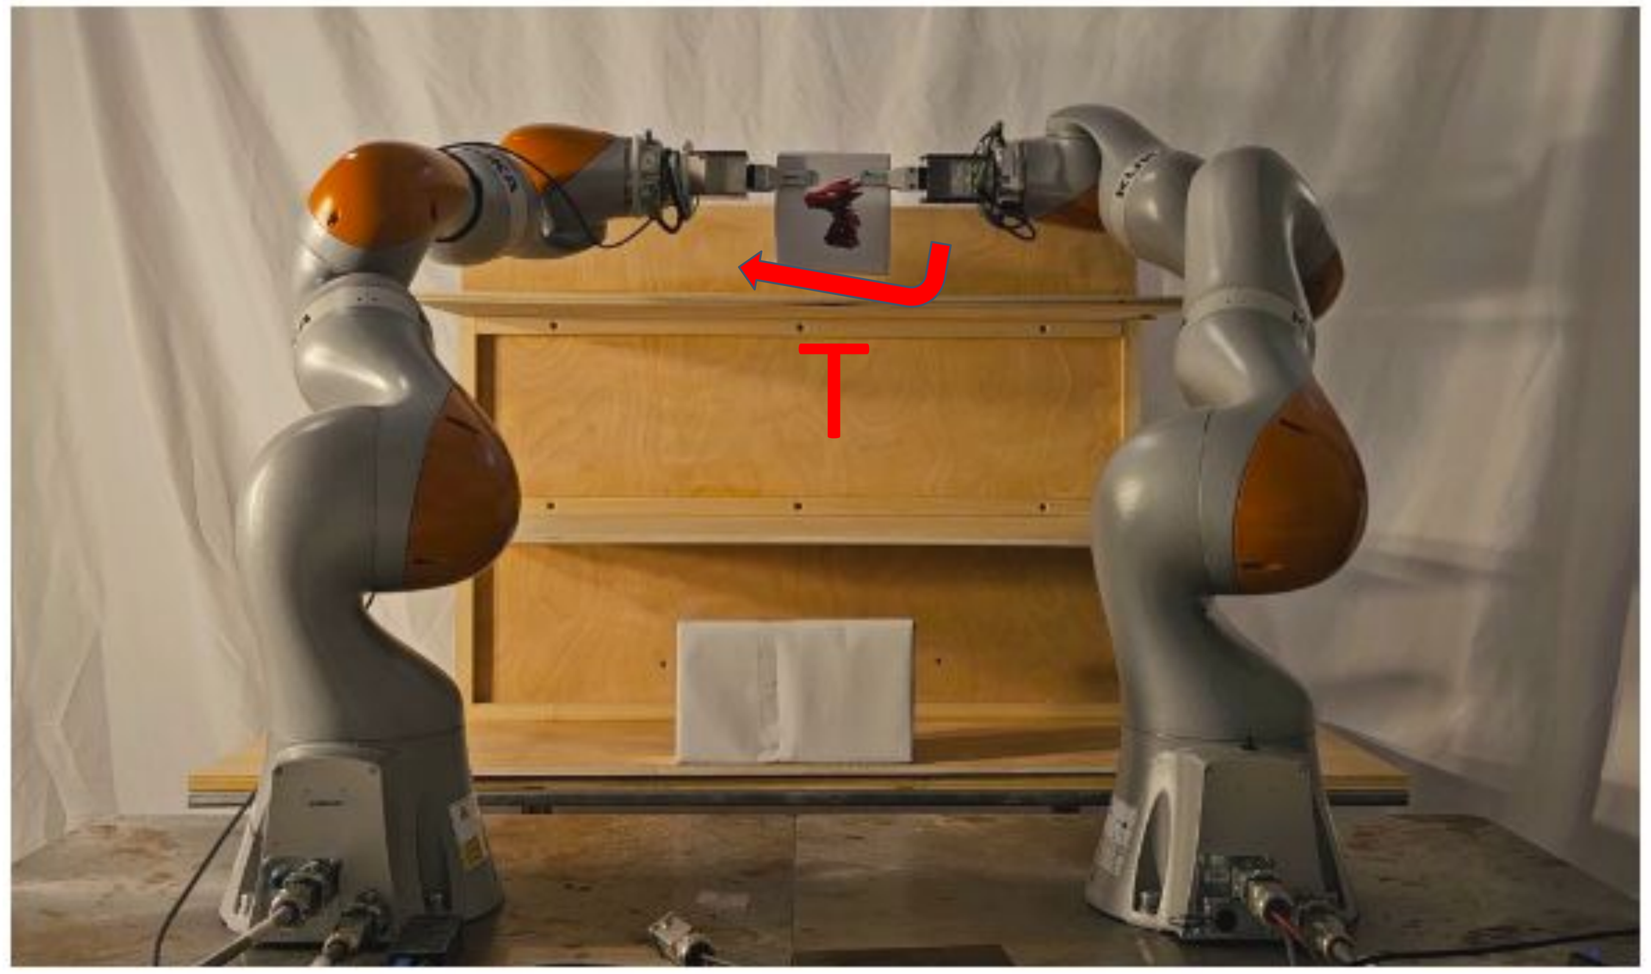

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

Bimanual manipulation kinematics challenges:

- The transform between the end effectors must remain fixed
- Configuration space becomes a non-linear mess with obstacles and other arm

  - Workspace: set of all reachable eeff points
  - Configuration space: all possible configurations for the robots joints

Obstacles in T^2

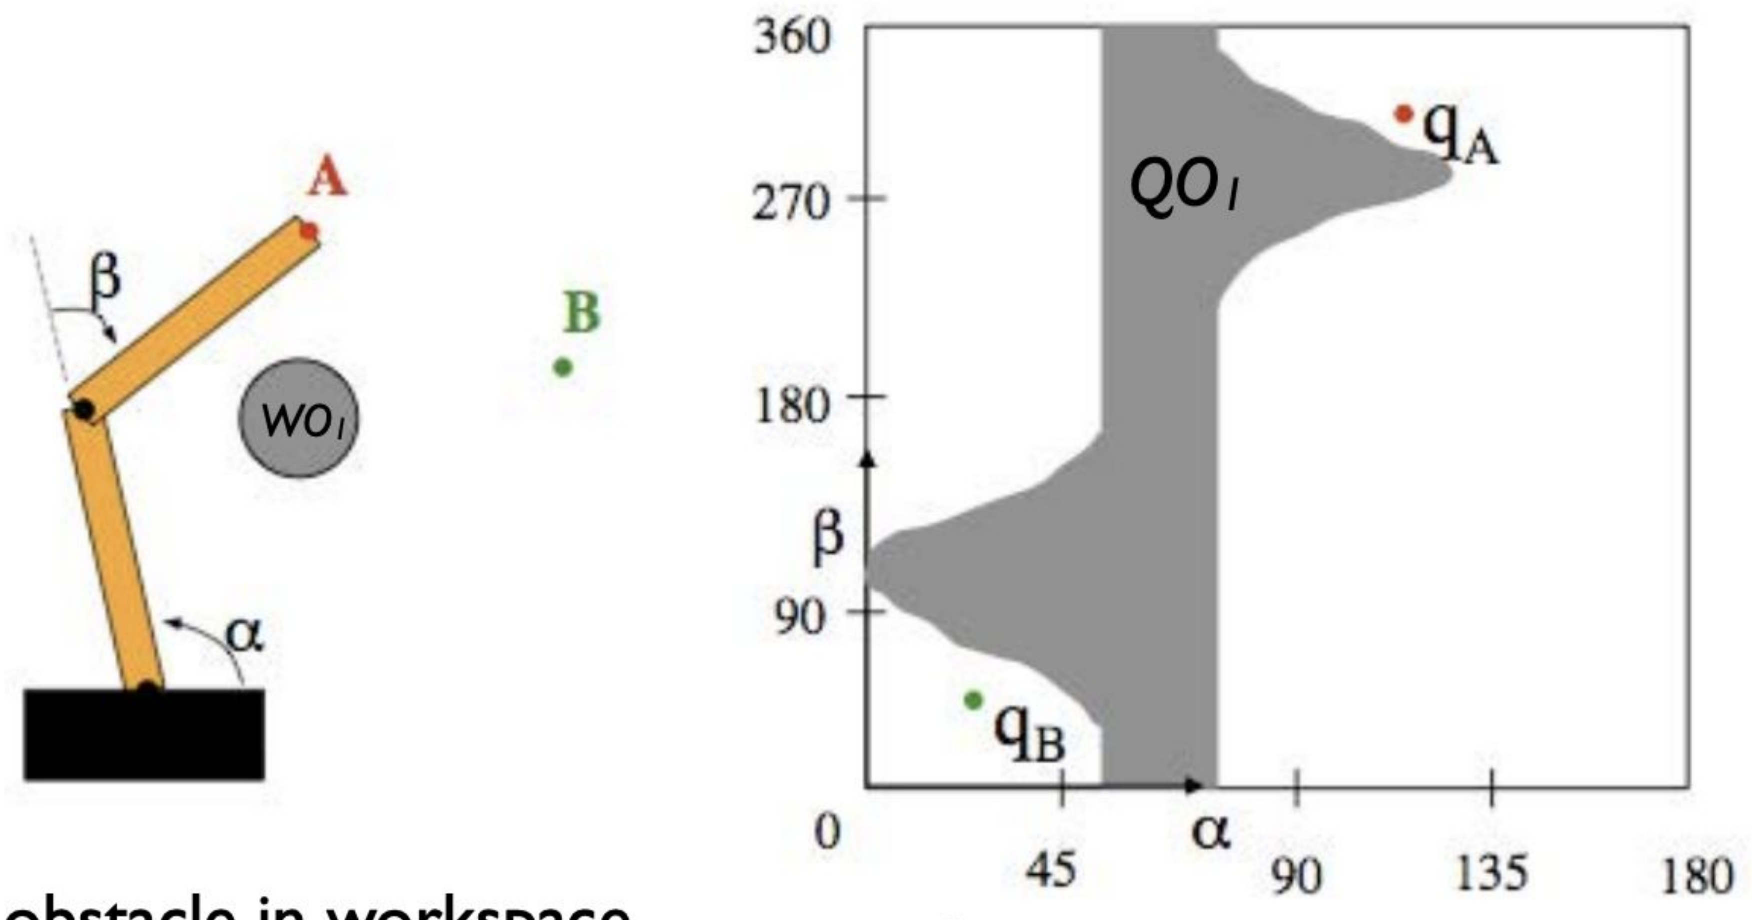

Circular obstacle in workspace

C-space representation

Slide borrowed from Michigan Robotics autorob.org

https://rpm-lab.github.io/CSCI5551-Spr24/assets/slides/lec11_planning_3_configuration_spaces.pdf

Bimanual manipulation kinematics challenges:

- The transform between the end effectors must remain fixed
- Configuration space becomes a non-linear mess with obstacles and other arm

    - Workspace: set of all reachable eeff points
    - Configuration space: all possible configurations for the robots joints

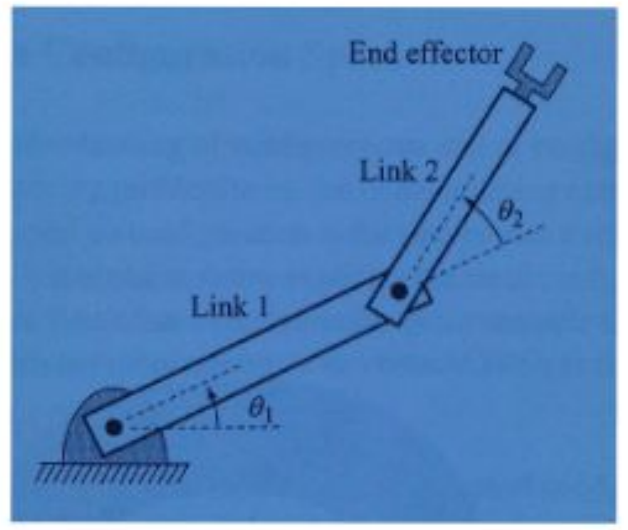

Workspace is w.r.t. end-effector position *(x,y)*

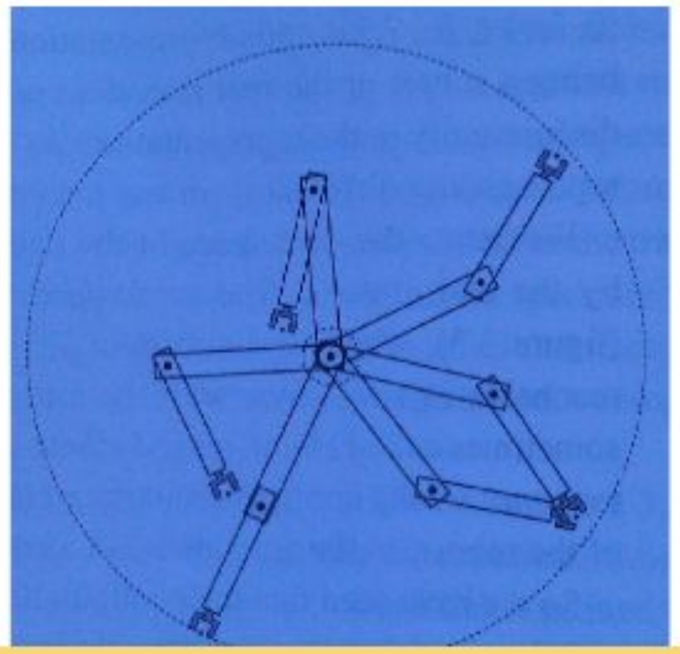

C-space is w.r.t. joint angles (Θ_1, Θ_2)

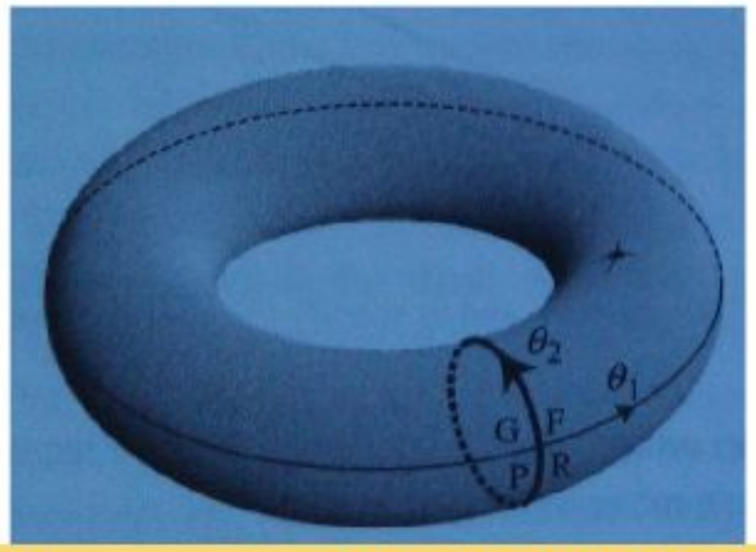

CSCI 5551 - Spring 2024

Slide borrowed from Michigan Robotics autorob.org

https://rpm-lab.github.io/CSCI5551-Spr24/assets/slides/lec11_planning_3_configuration_spaces.pdf

Bimanual manipulation kinematics challenges:

- The transform between the end effectors must remain fixed
- Configuration space becomes a non-linear mess with obstacles and other arm
- Numeric IK becomes very computationally expensive

Bimanual manipulation kinematics challenges:

- The transform between the end effectors must remain fixed
- Configuration space becomes a non-linear mess with obstacles and other arm
- Numeric IK becomes very computationally expensive
- How would we tackle all of these?

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

One neat approach:

One neat approach:

Constrained Bimanual Planning with Analytic Inverse Kinematics

Thomas Cohn, Seiji Shaw, Max Simchowitz, and Russ Tedrake

***Abstract*— In order for a bimanual robot to manipulate an object that is held by both hands, it must construct motion plans such that the transformation between its end effectors remains fixed. This amounts to complicated nonlinear equality constraints in the configuration space, which are difficult for trajectory optimizers. In addition, the set of feasible configurations becomes a measure zero set, which presents a challenge to sampling-based motion planners. We leverage an analytic solution to the inverse kinematics problem to parametrize the configuration space, resulting in a lower-dimensional representation where the set of valid configurations has positive measure. We describe how to use this parametrization with existing motion planning algorithms, including sampling-based approaches, trajectory optimizers, and techniques that plan through convex inner-approximations of collision-free space.**

I. Introduction

Enabling bimanual robots to execute coordinated actions with both arms is essential for achieving (super)human-like skill in automation and home contexts. There exists a variety of tasks that are only solvable when two arms manipulate in concert [1], such as carrying an unwieldy object, folding clothes, or assembling parts. In many manipulation tasks, one gripper can be used to provide fixture to the manipuland, while the other performs the desired action [2]; such tasks include opening a bottle, chopping vegetables, and tightening a bolt. Furthermore, some tools explicitly require two arms to use, such as hand mixers, rolling pins, and can openers.

To accomplish many of these desired tasks, the motion of the robot arms becomes subject to equality constraints imposed in task space. For example, when moving an object that is held by both hands, the robot must ensure that the transformation between the end effectors remains constant. Such task space constraints appear as complicated nonlinear equality constraints in configuration space, posing a major challenge to traditional motion planning algorithms.

In the existing literature, there are general techniques for handling task-space constraints in configuration-space planning. Sampling-based planners can project samples onto the constraint manifold [3] or use numerical continuation [4] to construct piecewise-linear approximations. Constraints can also be relaxed [5] or enforced directly with trajectory optimization [6]. In the case of certain bimanual planning problems, there is additional structure that is not exploited

This work was supported by Amazon.com, PO No. 2D-06310236, the MIT Quest for Intelligence, and the National Science Foundation Graduate Research Fellowship Program under Grant No. 2141064. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The authors are with the Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology, Cambridge, Massachusetts {tcohn,seijis,msimchow,russt}@mit.edu

Fig. 1: Hardware setup for our experiments. The two arms must work together to move an objects between the shelves, avoiding collisions and respecting the kinematic constraint.

by these general methods. For certain classes of robot arms, *analytic inverse kinematics* (analytic IK) can be used to map an end-effector pose (along with additional parameters to resolve kinematic redundancy) to joint angles in closed form. Such solutions are specific to certain classes of robot arms, but are a powerful tool to be leveraged if available. Fortunately, analytic IK *is* available for many popular robot arms available today, including the KUKA iiwa. See Figure 1.

If a robot must move an object that it is holding with both hands, we propose constructing a plan for one "controllable" arm, and then the other "subordinate" arm can be made to follow it via an analytic IK mapping. Configurations where the subordinate arm cannot reach the end-effector of the primary arm, or where doing so would require violating joint limits, are treated as obstacles. In this way, we parametrize the constraint manifold so that the feasible set has positive measure in the new planning space. Because we no longer have to consider the equality constraints, sampling-based planning algorithms can be applied without modification. We can also differentiate through the IK mapping, enabling the direct application of trajectory optimization approaches.

The remainder of this paper is organized as follows. First, we give an overview of the existing techniques used for constrained motion planning, and describe the available analytic IK solutions. Then, we present our parametrization of the constraint manifold for bimanual planning, and discuss its relevant geometric and topological properties. We describe the slight modifications which are necessary to adapt standard planning algorithms (including sampling-based planning and trajectory optimization) to operate in this framework. We then present a technique for generating

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

Analytic IK

- Rather than gradient descent - find closed form solution for joint angles instead…

Analytic IK

- Rather than gradient descent - find closed form solution for joint angles instead…
- Geometric algebra can be very difficult, many common configurations are already solved

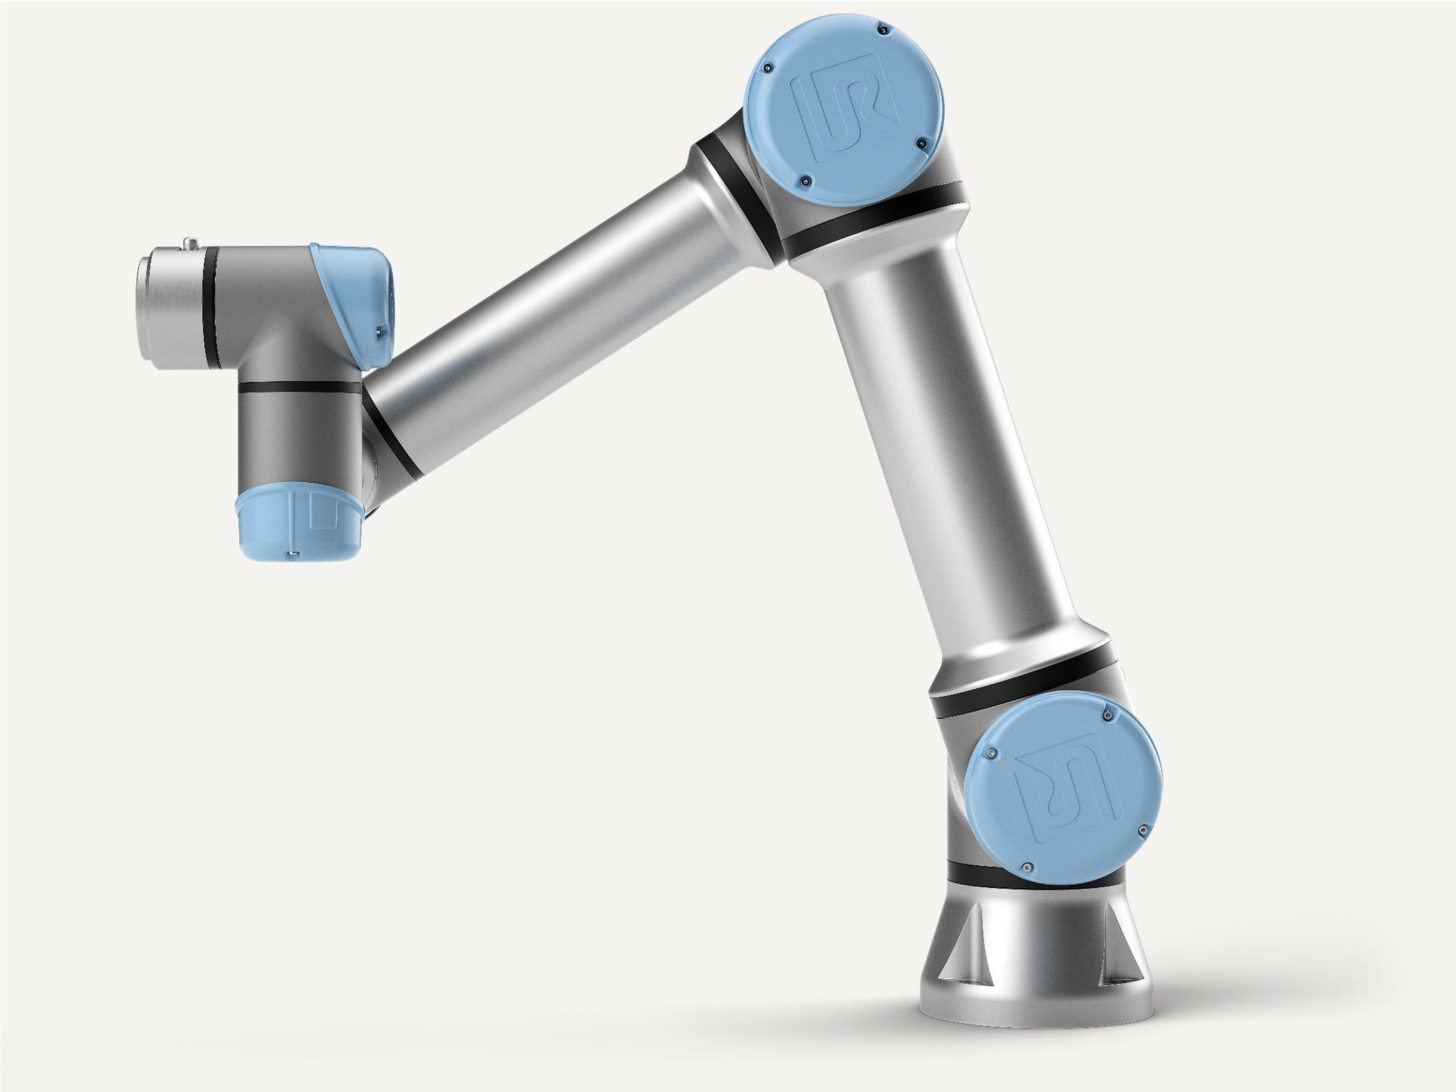

https://www.universal-robots.com/products/ur5-robot/

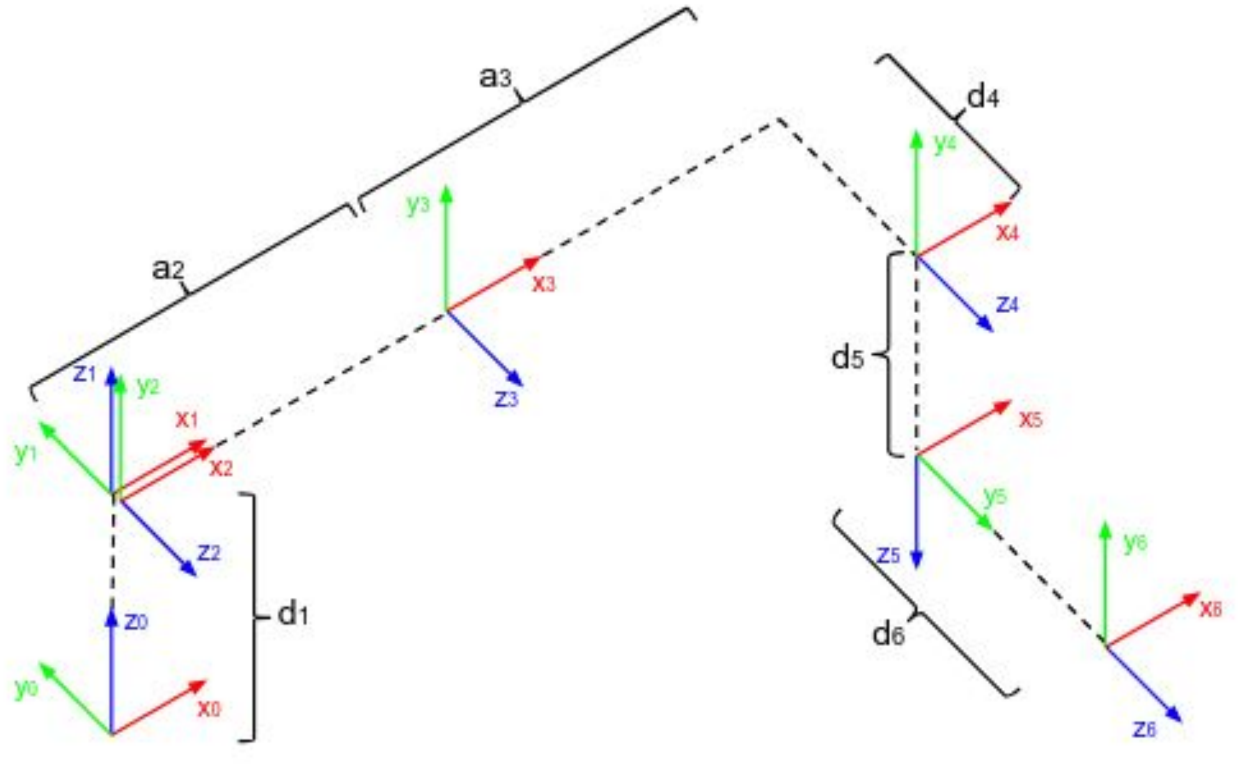

Figure 1: Coordinate frames for UR arm. Joints rotate around the z-axes and are pictured at $\theta_i = 0$ for $1 \le i \le 6$.

i	d_i	a_i	α_i
0	-	0	0
1	d_1	0	$\pi/2$
2	0	a_2	0
3	0	a_3	0
4	d_4	0	$\pi/2$
5	d_5	0	$-\pi/2$
6	d_6	-	-

Table 1: Denavit-Hartenberg parameters for the UR Arms

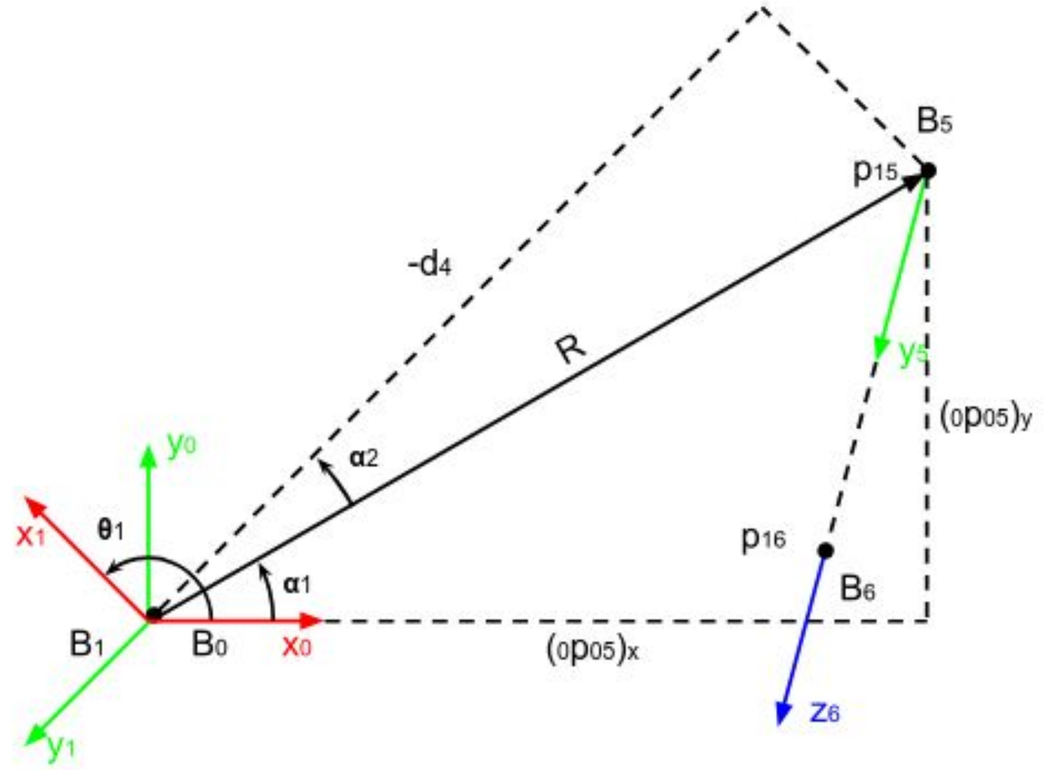

Figure 2: Illustration of the geometry of finding θ_1. This is essentially an overhead view of the robot, looking down on the x-y plane.

https://repository.gatech.edu/server/api/core/bitstreams/e56759bc-92c8-43df-aa62-0dc47581459d/content

Analytic IK

- Rather than gradient descent - find closed form solution for joint angles instead…
- Geometric algebra can be very difficult, many common configurations are already solved

https://www.kuka.com/en-se/products/robotics-systems/industrial-robots/lbr-iiwa

5	0	$-\pi/2$	d_{ew}	θ_5
6	0	$\pi/2$	0	θ_6
7	0	0	d_{wf}	θ_7

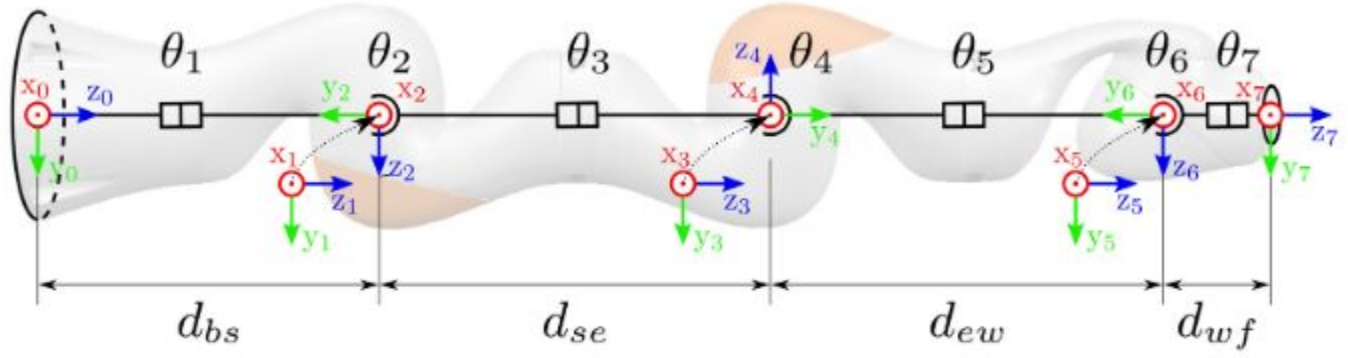

Fig. 1. Manipulator generic structure, joint variables and DH frames assigned. The 7-DoF manipulator model of LBR iiwa 7 R800 from KUKA AG is used to depict the shape of an anthropomorphic arm without offsets.

6.2. Trajectory example

To demonstrate the redundancy resolution strategy, we created an example linear trajectory to be performed by the robot manipulator. The robot is purposely positioned at a configuration near its mechanical joint limits. The robot starts at the joint angles in degrees:

$$\theta^c = \begin{bmatrix}-5.4101 & -26.4986 & -48.1542 & -61.6500 & 152.6198 & 114.4466 & 8.1812\end{bmatrix} \quad (37)$$

which correspond to the global configuration $GC = 3$, arm angle $\psi = 58.5882^\circ$ and pose:

$$^0\mathbf{T}_7^c = \begin{bmatrix}-0.2634 & -0.9112 & -0.3166 & -0.1174 \\ 0.3014 & -0.3895 & 0.8703 & -0.1464 \\ -0.9164 & 0.1338 & 0.3773 & 1.0203 \\ 0 & 0 & 0 & 1\end{bmatrix} \quad (38)$$

The manipulator performs a linear motion in Cartesian space, keeping the same end-effector orientation but translating its position along the z-axis ($^0\mathbf{R}_{7,z}$) of a distance of 0.25 m, ending at the target pose:

$$^0\mathbf{T}_7^d = \begin{bmatrix}-0.2634 & -0.9112 & -0.3166 & -0.1966 \\ 0.3014 & -0.3895 & 0.8703 & 0.0712 \\ -0.9164 & 0.1338 & 0.3773 & 1.1146 \\ 0 & 0 & 0 & 1\end{bmatrix} \quad (39)$$

The path is interpolated and a new pose is passed to the redundancy resolution algorithm (Fig. 10) every iteration. The global configuration remains unchanged throughout the trajectory, and the arm angle varies according to the parameters defined (α and K).[5]

https://www.sciencedirect.com/science/article/pii/S0094114X17306559

Analytic IK

- Rather than gradient descent - find closed form solution for joint angles instead…
- Look towards existing software solutions, OpenRAVE IK Fast

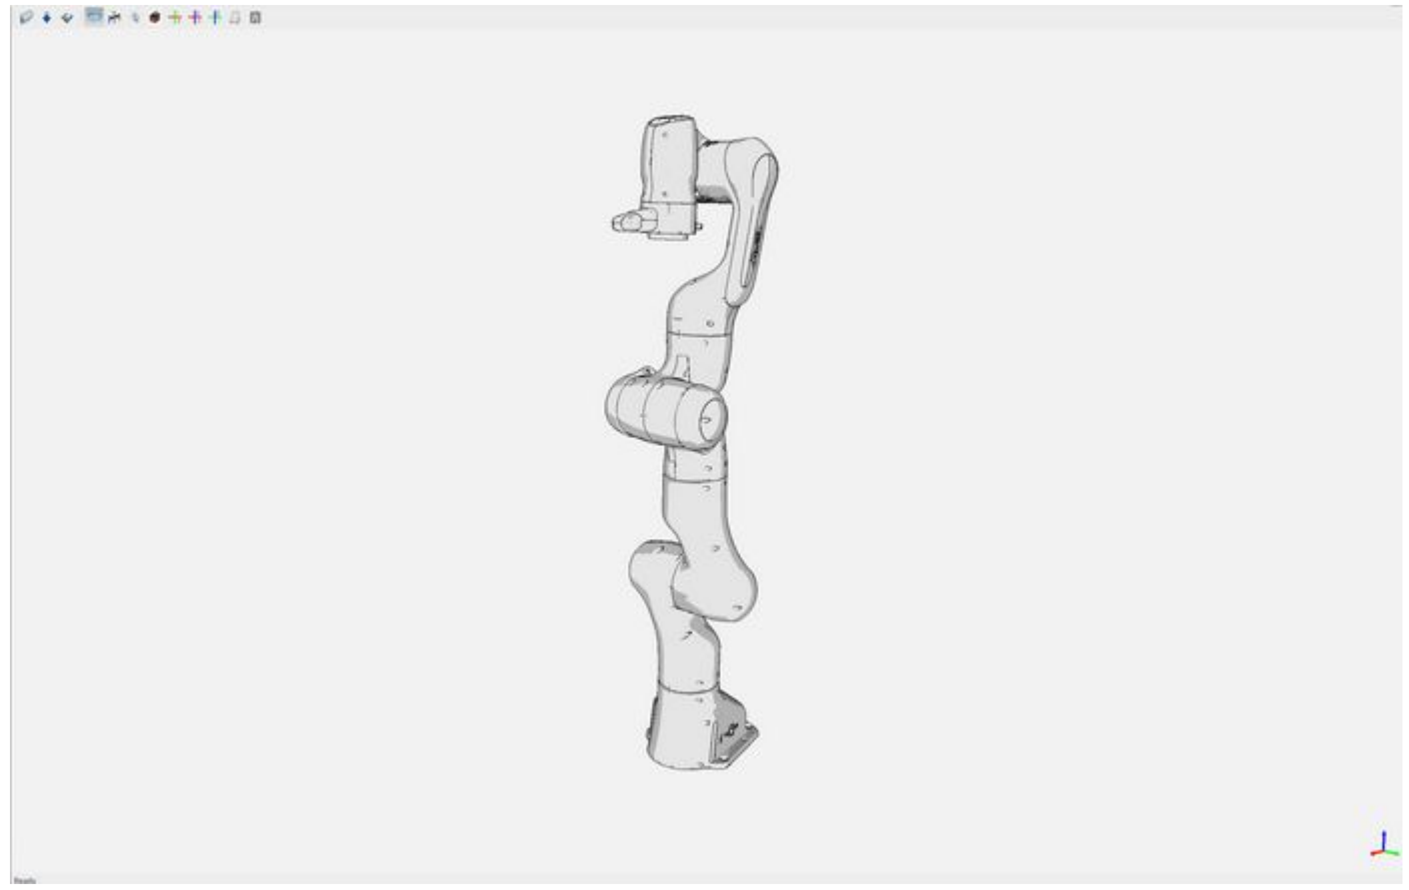

Self-motion

- TL;DR - There a a degree of freedom that exists by virtue of 7 DoF arms (such as the Kuka) that allows for movement without changing end-effector position.
- Show JS example…

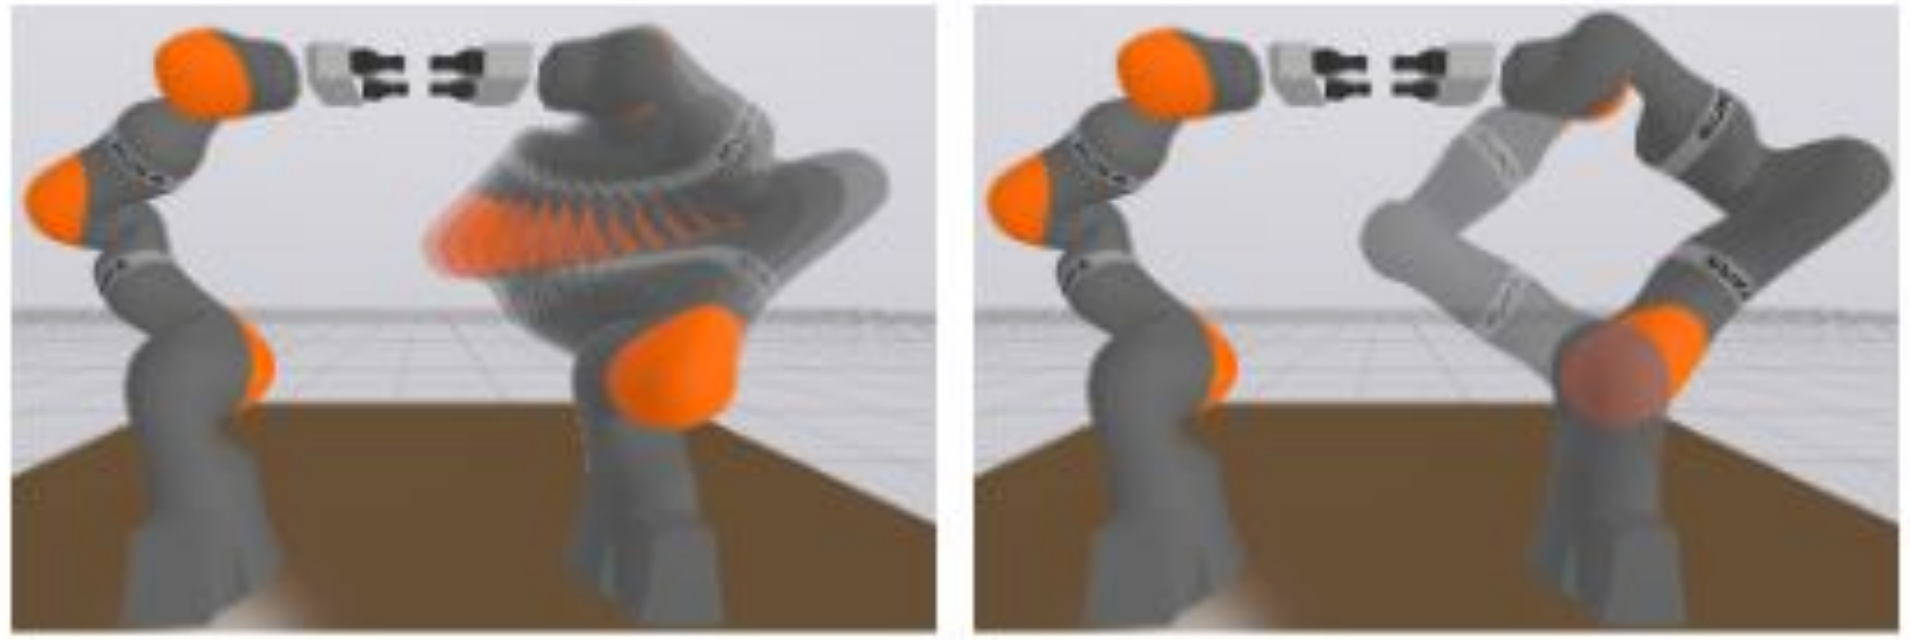

Fig. 3: Continuous (left) and discrete (right) self-motions of a 7DoF arm. The continuous self-motion yields an additional degree of freedom for the planner to consider, whereas the discrete self-motion is not utilized.

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

Constrained configuration space

- Constantly checking configuration space to determine the following:
  - Are there any collisions?
  - Is the transformation between end effectors the same?
- Constraint manifold - set of possible configurations that satisfy the systems constraints
  - Requires offline planning to build

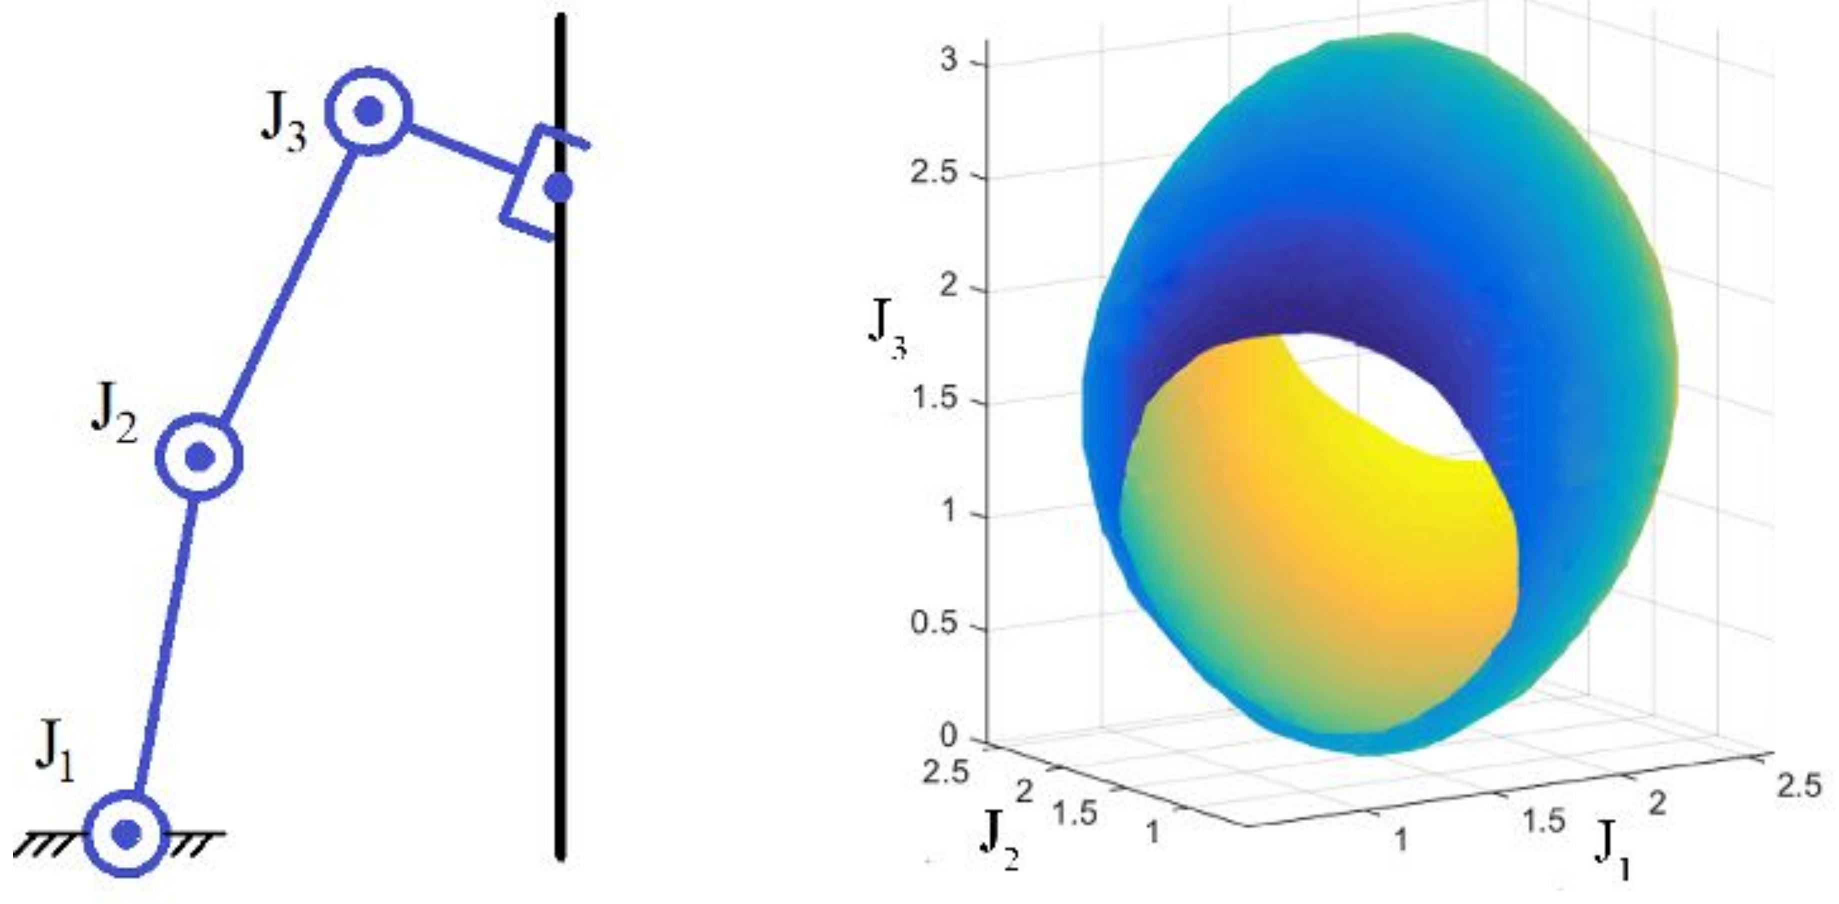

(a) 3-DOF robot **(b)** Constraint manifold

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

https://www.semanticscholar.org/paper/Learning-the-Metric-of-Task-Constraint-Manifolds-Zha-Liu/c3e11e5447a30b9f8ea16d73866bdd8ddccfecf6

Quick aside - Manifolds

- Set of points resembling Euclidean space
- Connectedness is defined

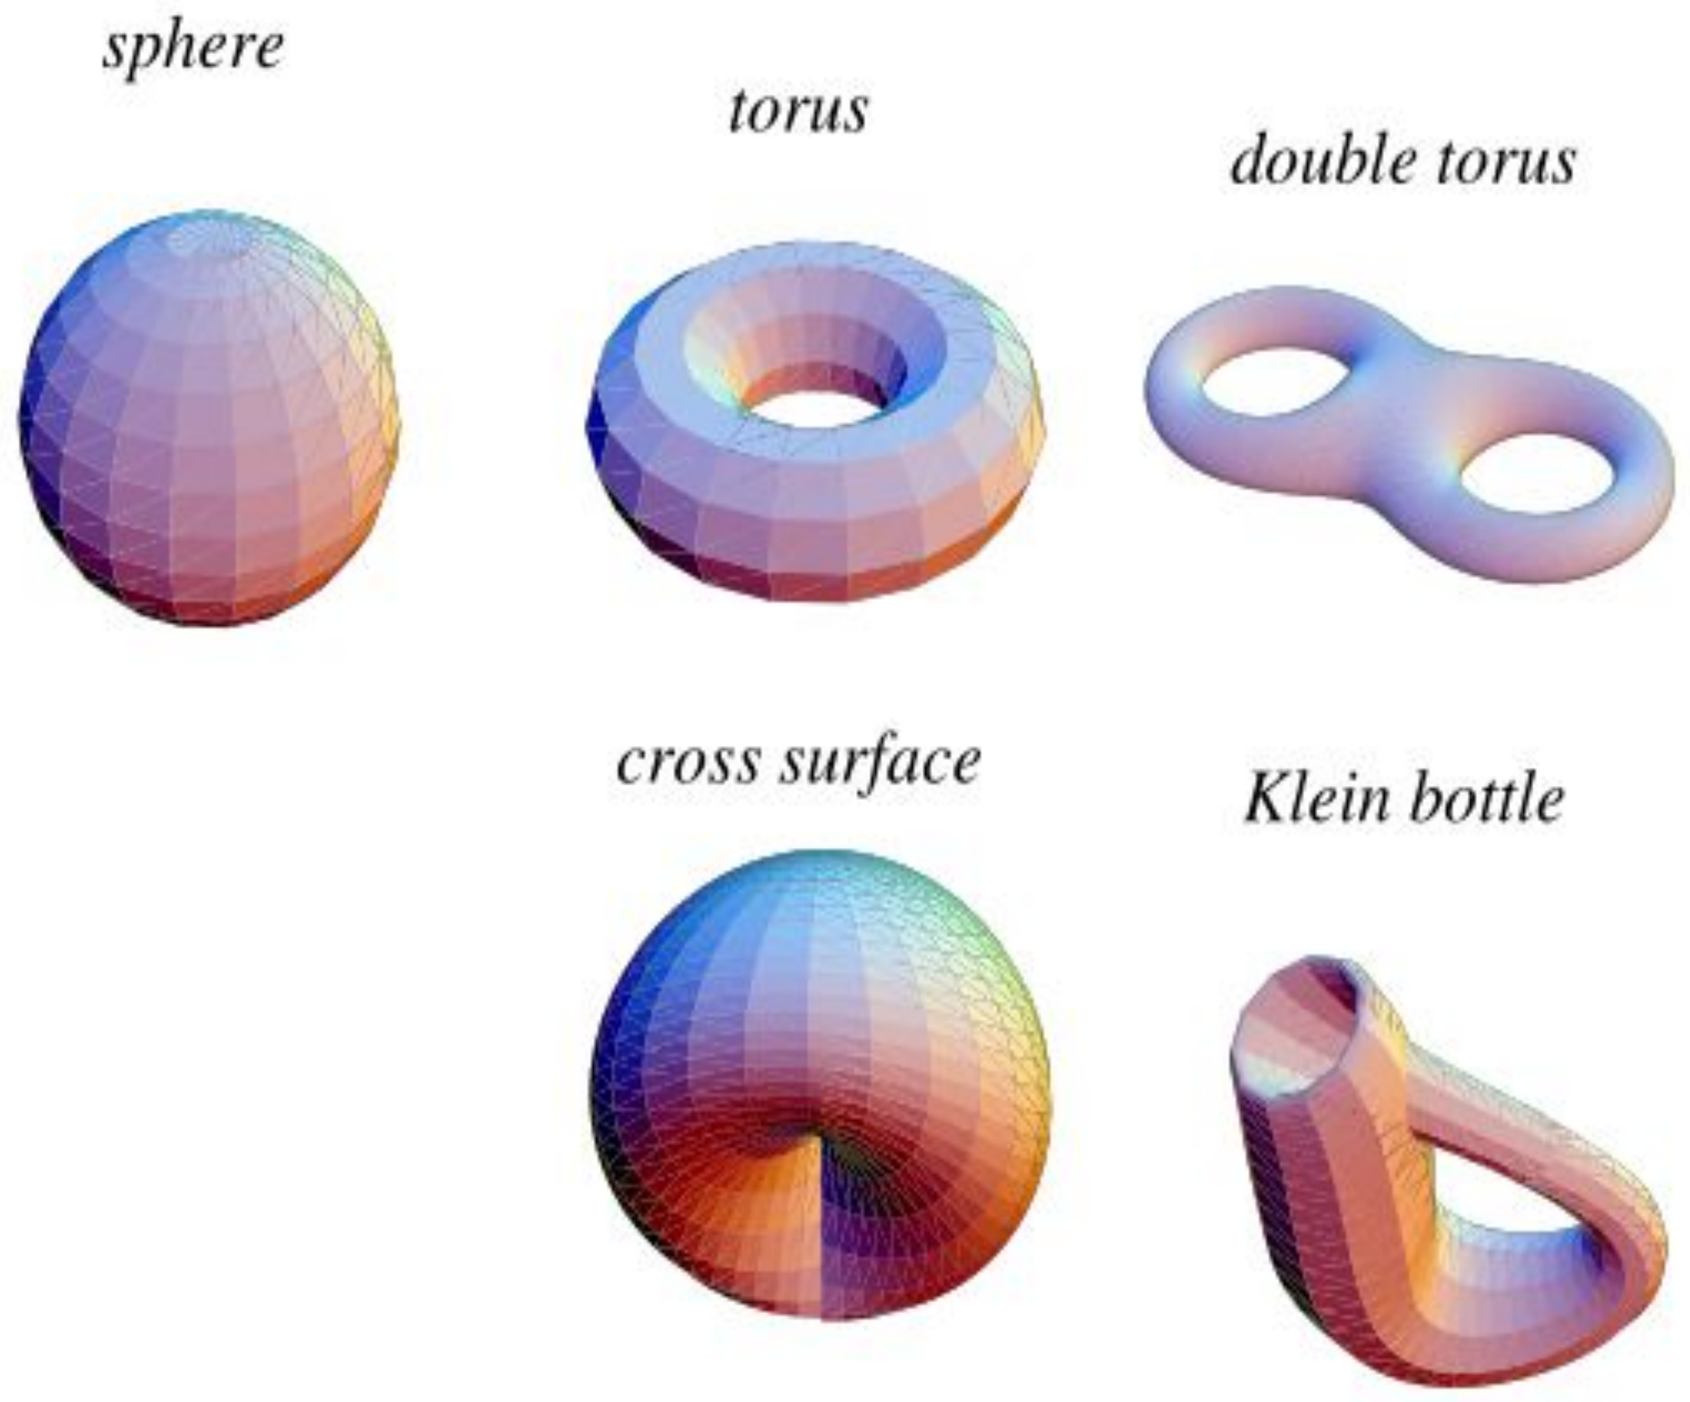

https://mathworld.wolfram.com/CompactManifold.html

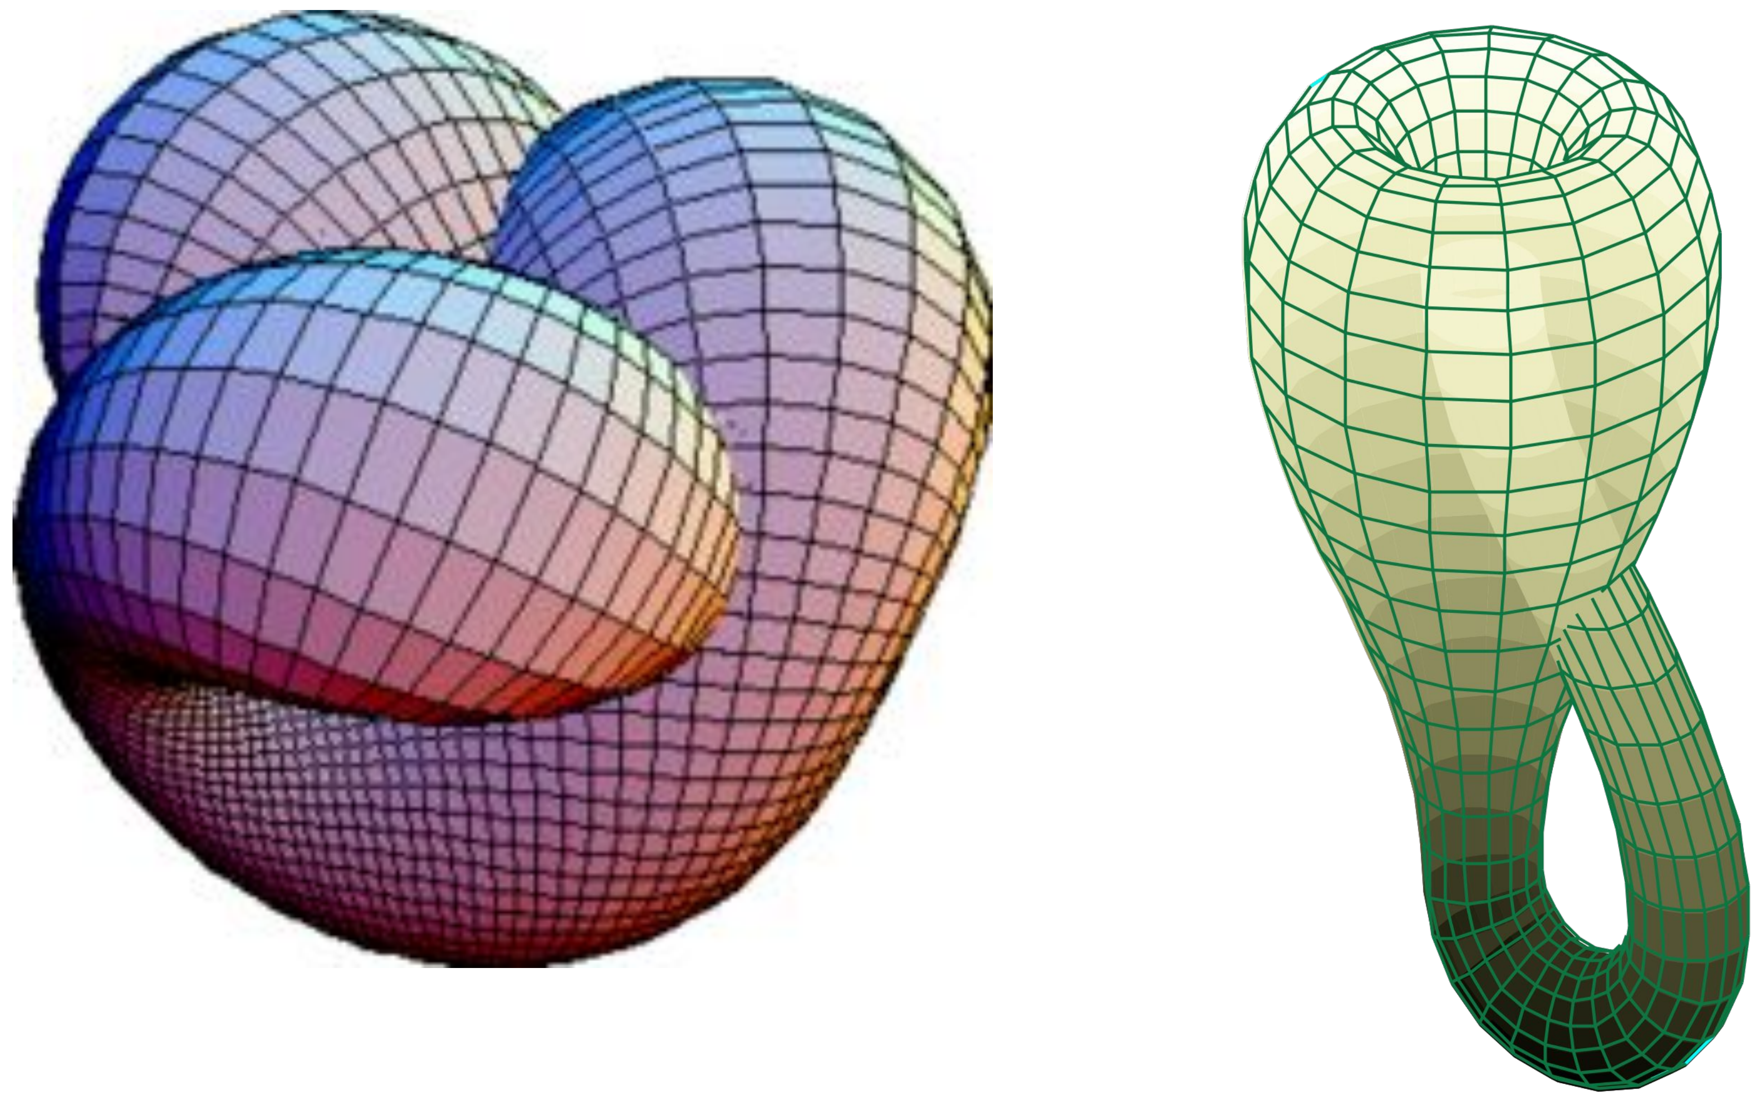

https://en.wikipedia.org/wiki/Manifold

Quick aside - Manifolds

- Set of points resembling Euclidean space
- Connectedness is defined

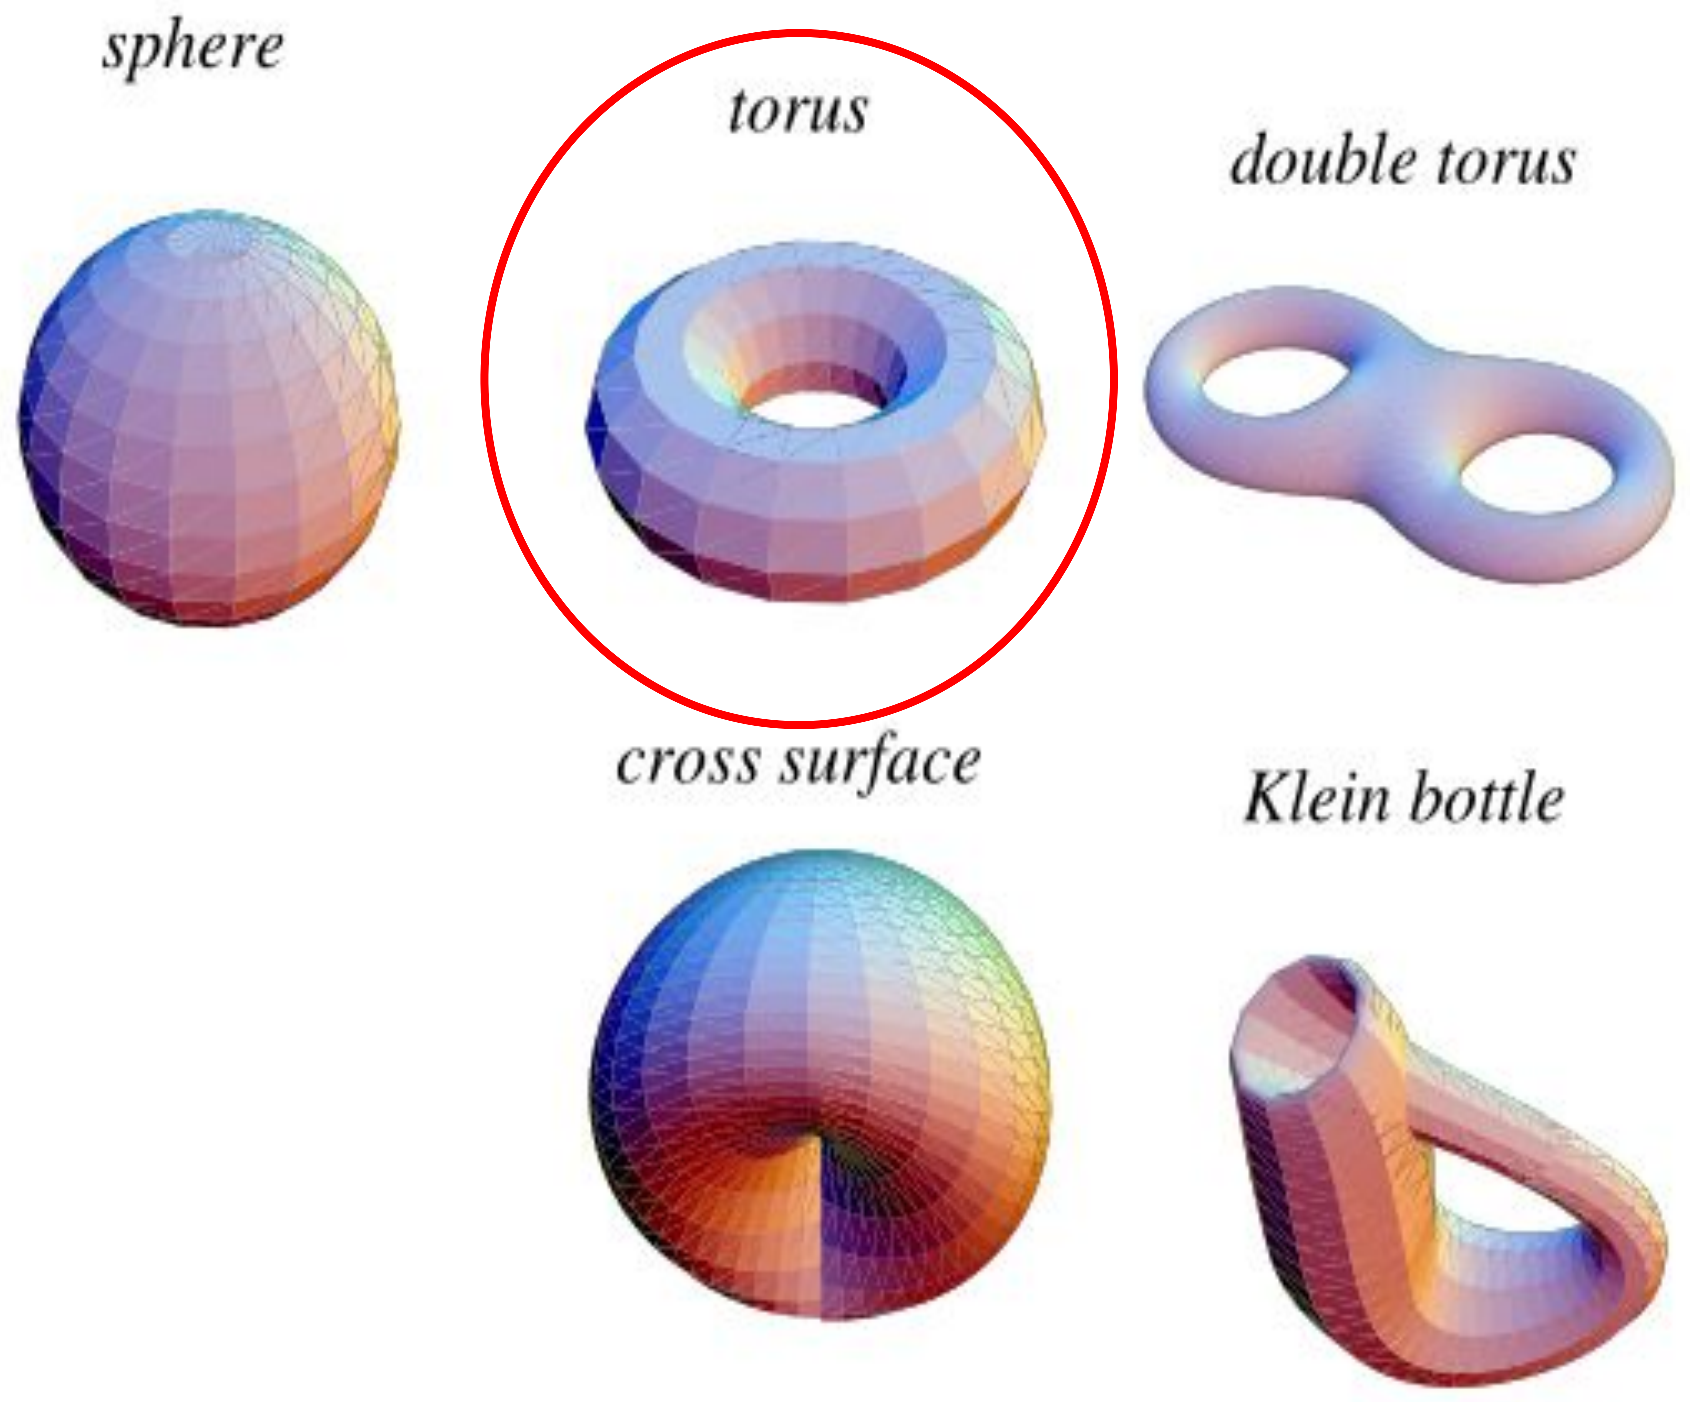

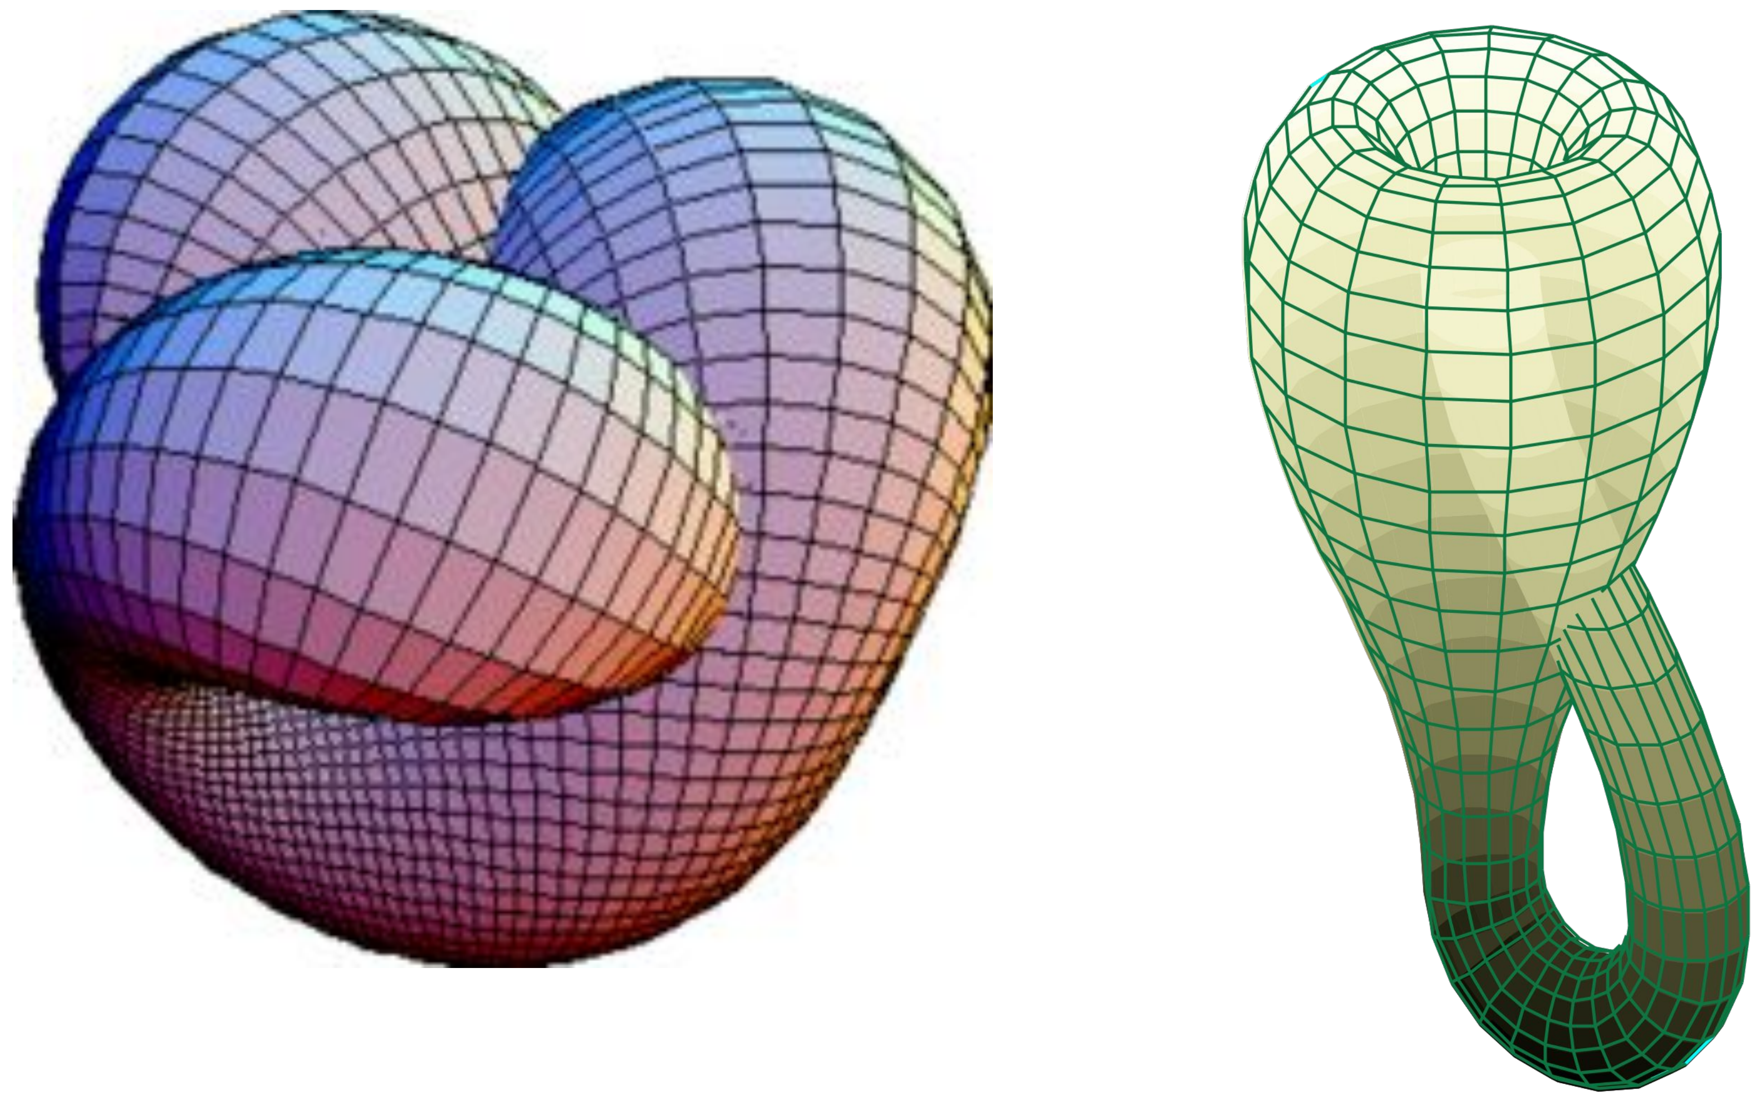

https://en.wikipedia.org/wiki/Manifold

https://mathworld.wolfram.com/CompactManifold.html

Path and motion planning

- Planning occurs in configuration space using constraint manifold

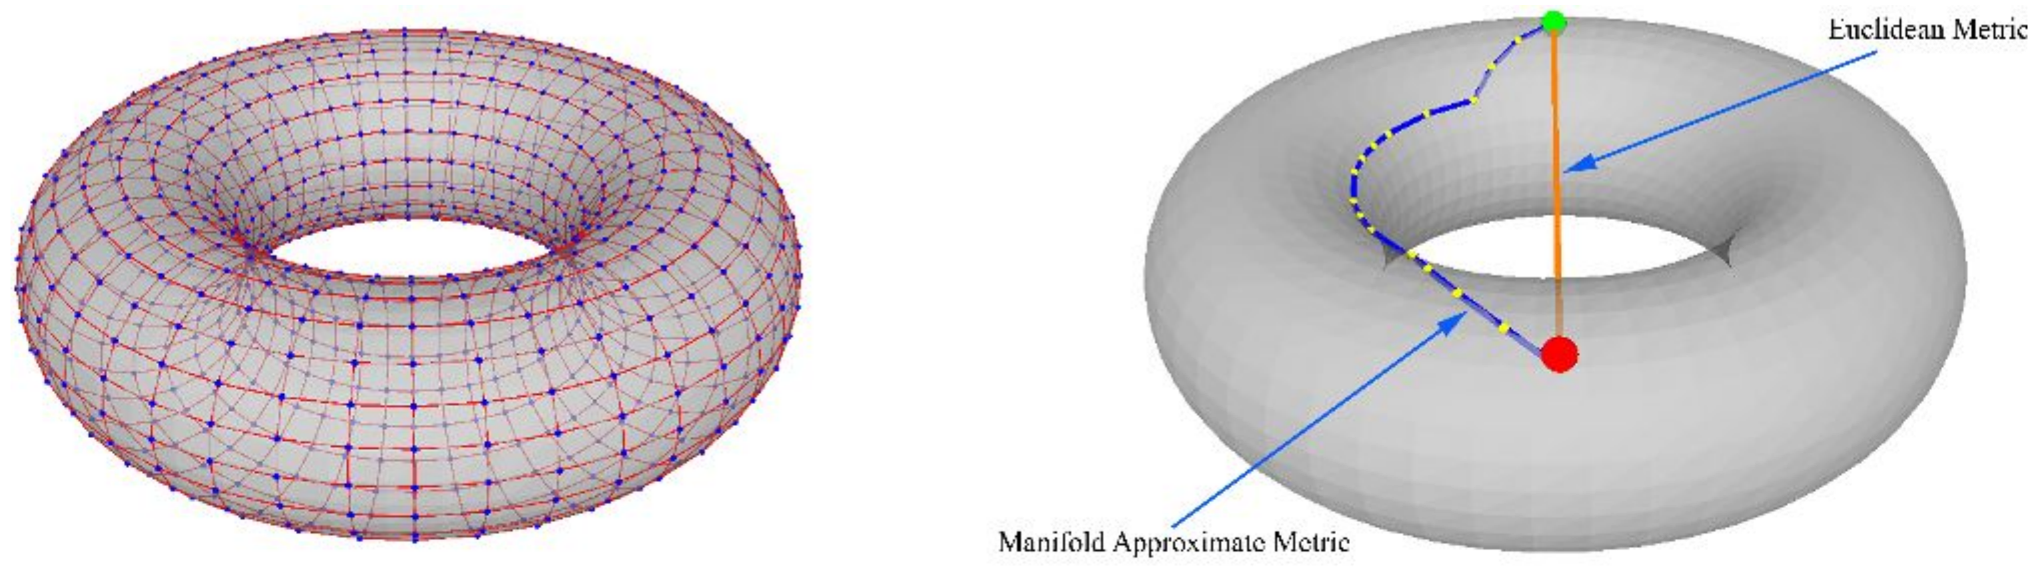

(a) Manifold approximate graph (b) Manifold approximate matric

Zha, Fusheng, Yizhou Liu, Wei Guo, Pengfei Wang, Mantian Li, Xin Wang, and Jingxuan Li. "Learning the metric of task constraint manifolds for constrained motion planning." *Electronics* 7, no. 12 (2018): 395.

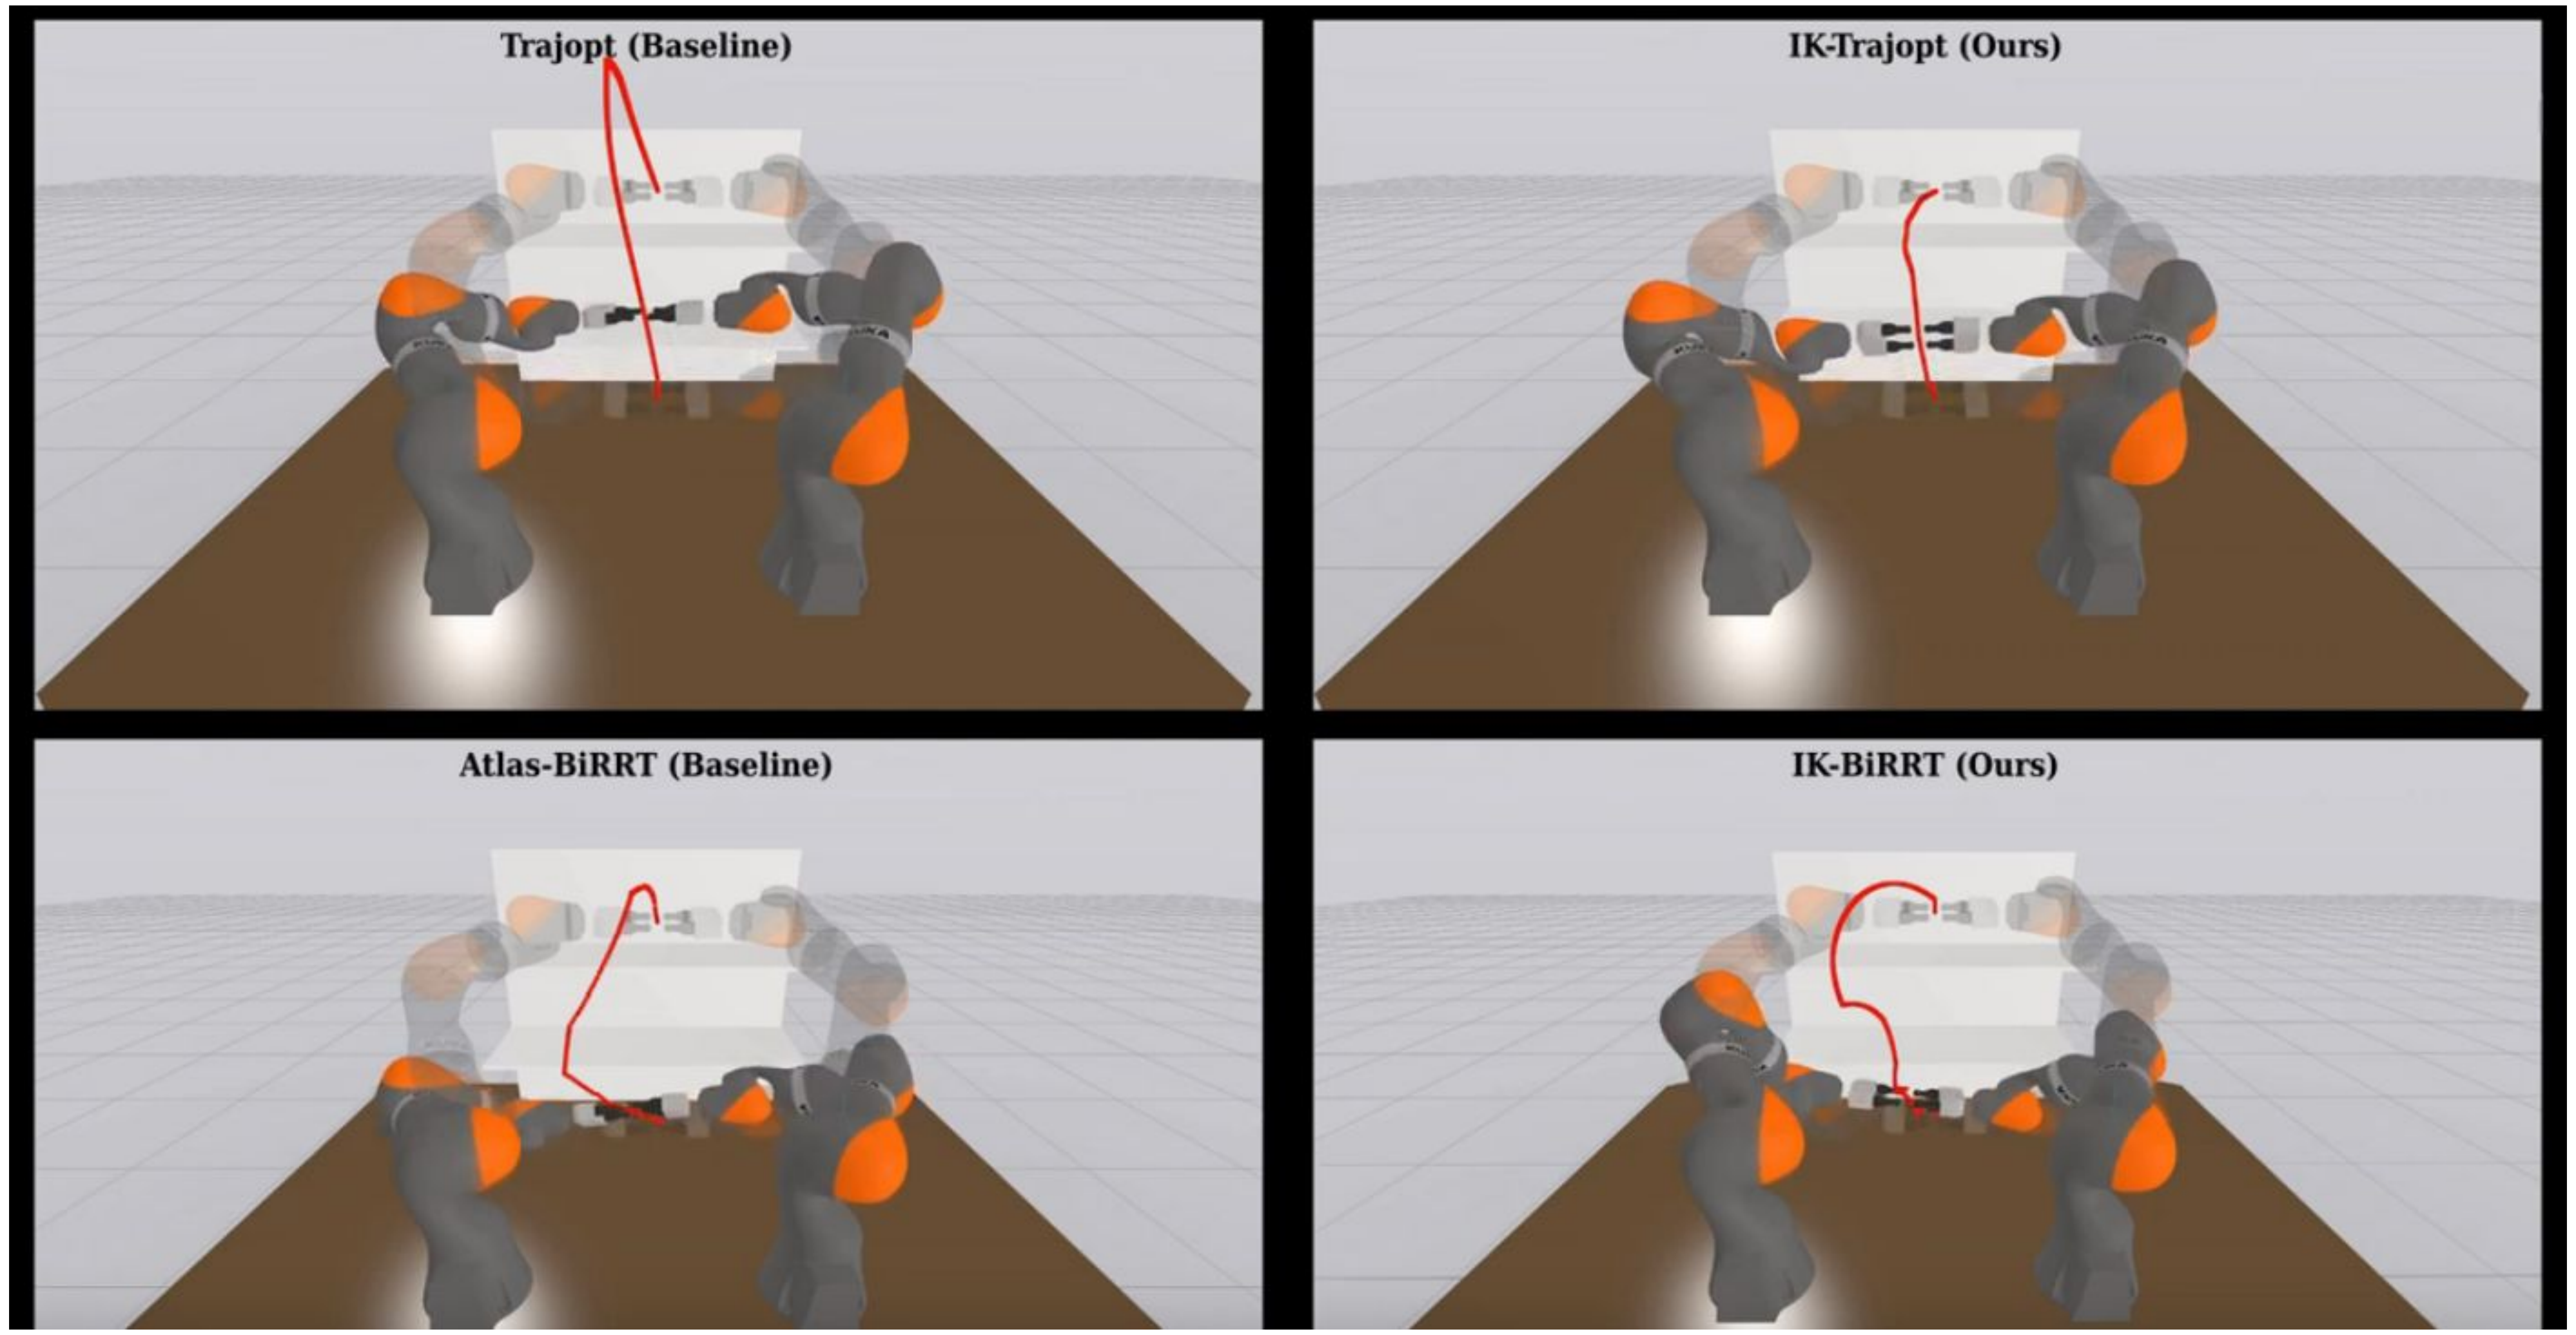

https://www.youtube.com/watch?v=vmujyn4EgTU

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

Path and motion planning

- Common path planning/graph algorithms can be used

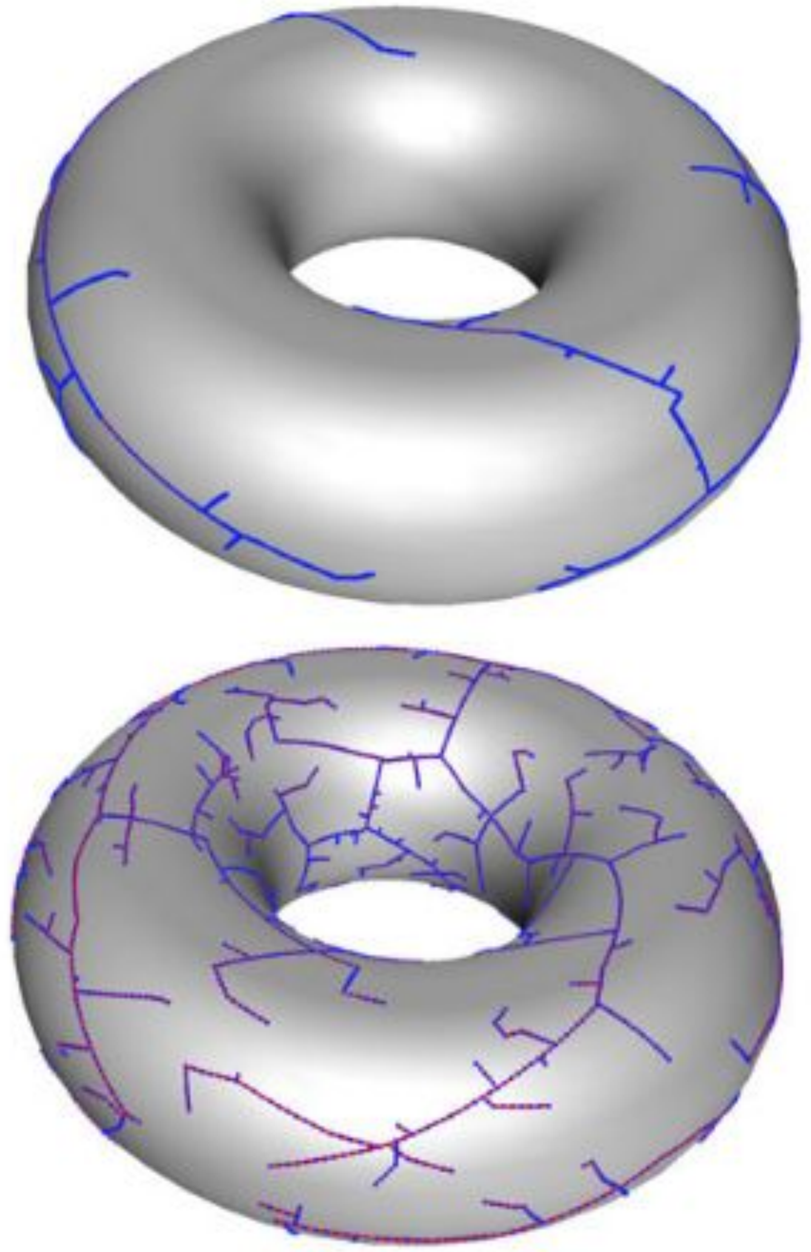

Fig. 2. Two RRTs of 500 samples built on a torus-like manifold. (Top) With an ambient space sampling method, the exploration focuses on the outer parts of the torus, and many samples do not produce a tree extension. (Bottom) With an AtlasRRT, the diffusion process is largely independent of the ambient space, which improves the coverage.

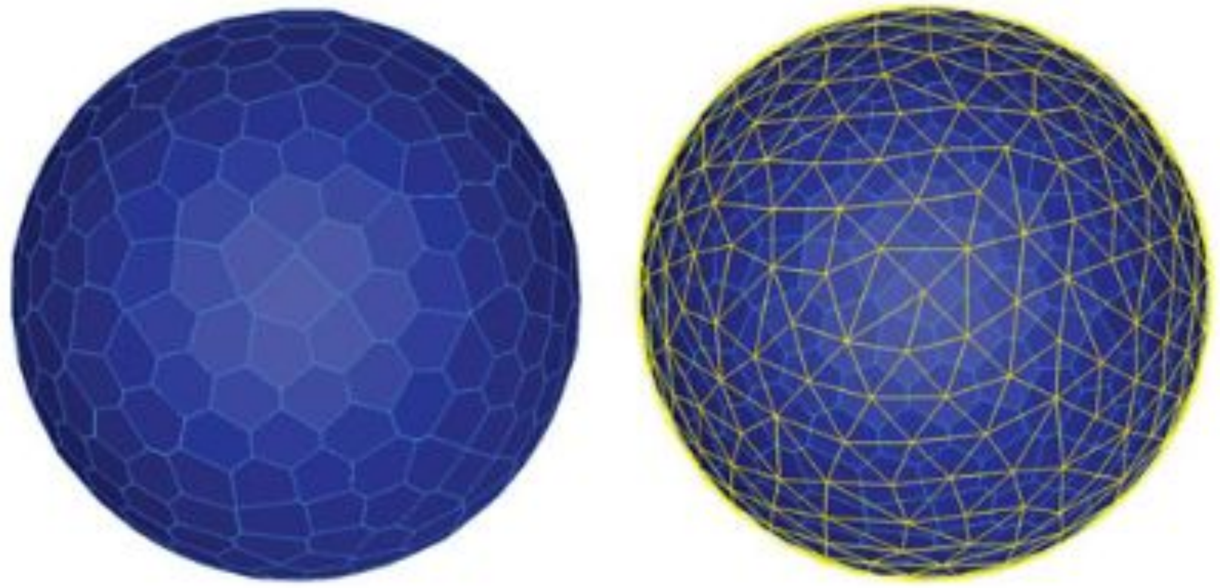

Fig. 6. (Left) Atlas of a sphere. Each polygonal patch corresponds to a given $\mathcal{P}_i$: a conservative approximation of the validity area for the associated chart. (Right) A roadmap can be extracted from the atlas where the nodes are the chart centers and where the edges are given by the neighborhood relations between charts. This roadmap could be used to devise collision free paths between any two given configurations.

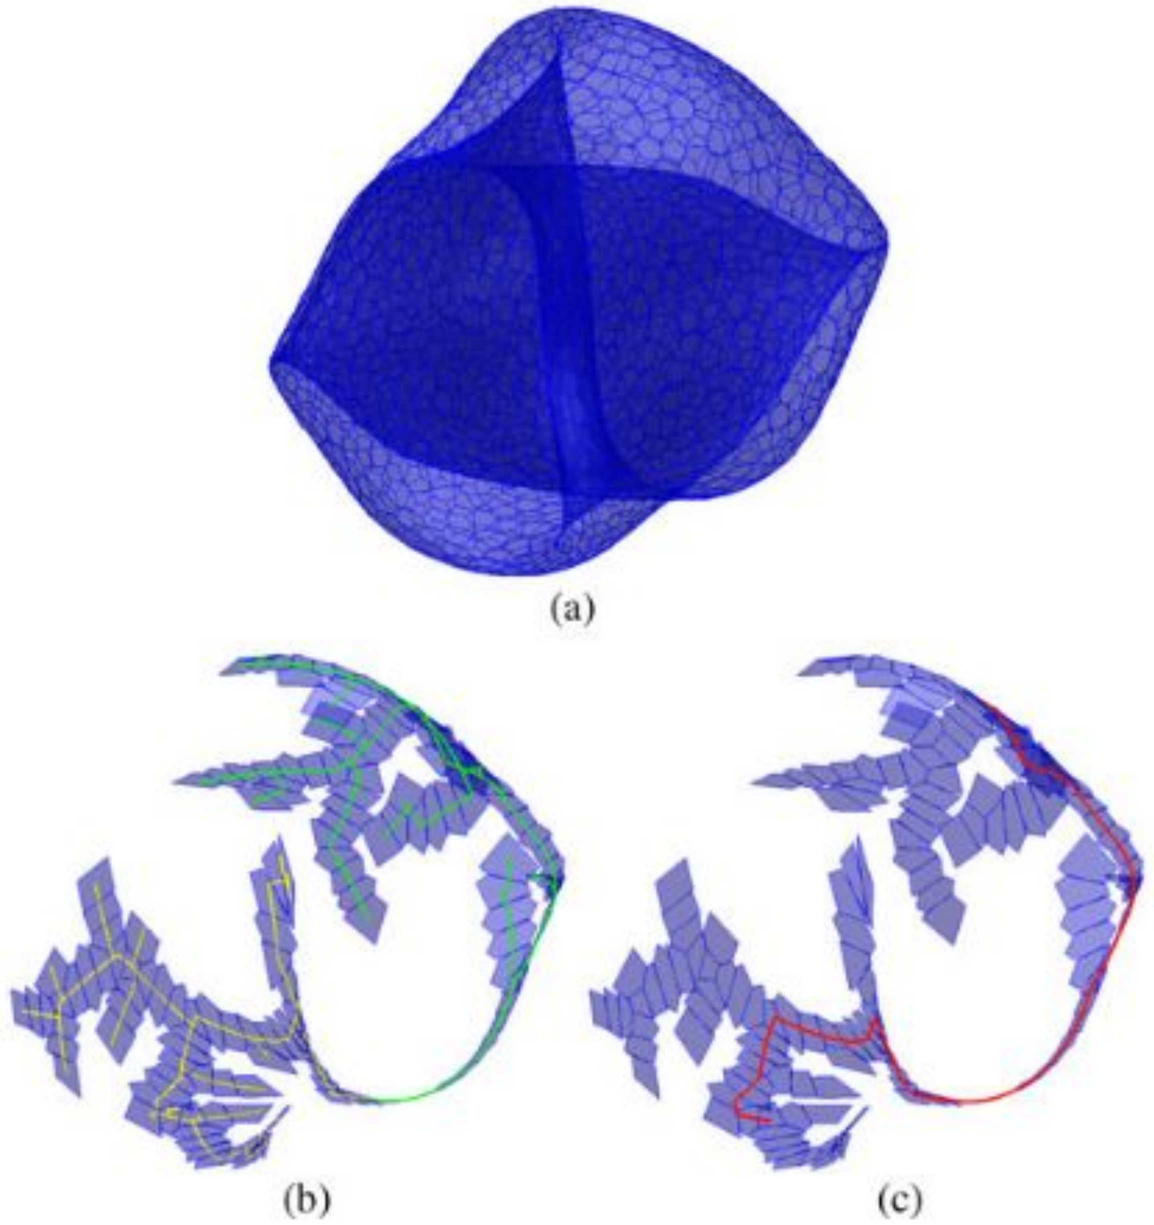

Fig. 1. Example of exploration with AtlasRRT. (a) Full atlas of the bidimensional configuration space of the cyclooctane. (b) AtlasRRT intertwines the construction of a bidirectional RRT with an atlas construction. The trees rooted at the start and goal configurations are represented in yellow and green, respectively. (c) When the two RRTs are connected, a solution path (represented in red) can be readily computed. Observe that only a small fraction of the full atlas is necessary to connect the query configurations.

https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6352929

L. Jaillet and J. M. Porta, "Path Planning Under Kinematic Constraints by Rapidly Exploring Manifolds," in IEEE Transactions on Robotics, vol. 29, no. 1, pp. 105-117, Feb. 2013, doi: 10.1109/TRO.2012.2222272.
keywords: {Manifolds;Kinematics;Path planning;Joints;Space exploration;Robot kinematics;Higher-dimensional continuation;kinematic constraints;manifolds;path planning},

Intuition for the papers method

- Constrain manifold is created using with obstacle collisions for each arm
- If the transformation between end effectors differs, that is treated as an obstacle
- The left arm is controlled
- Right arm follows
- Any path planning algorithm can be used for trajectories

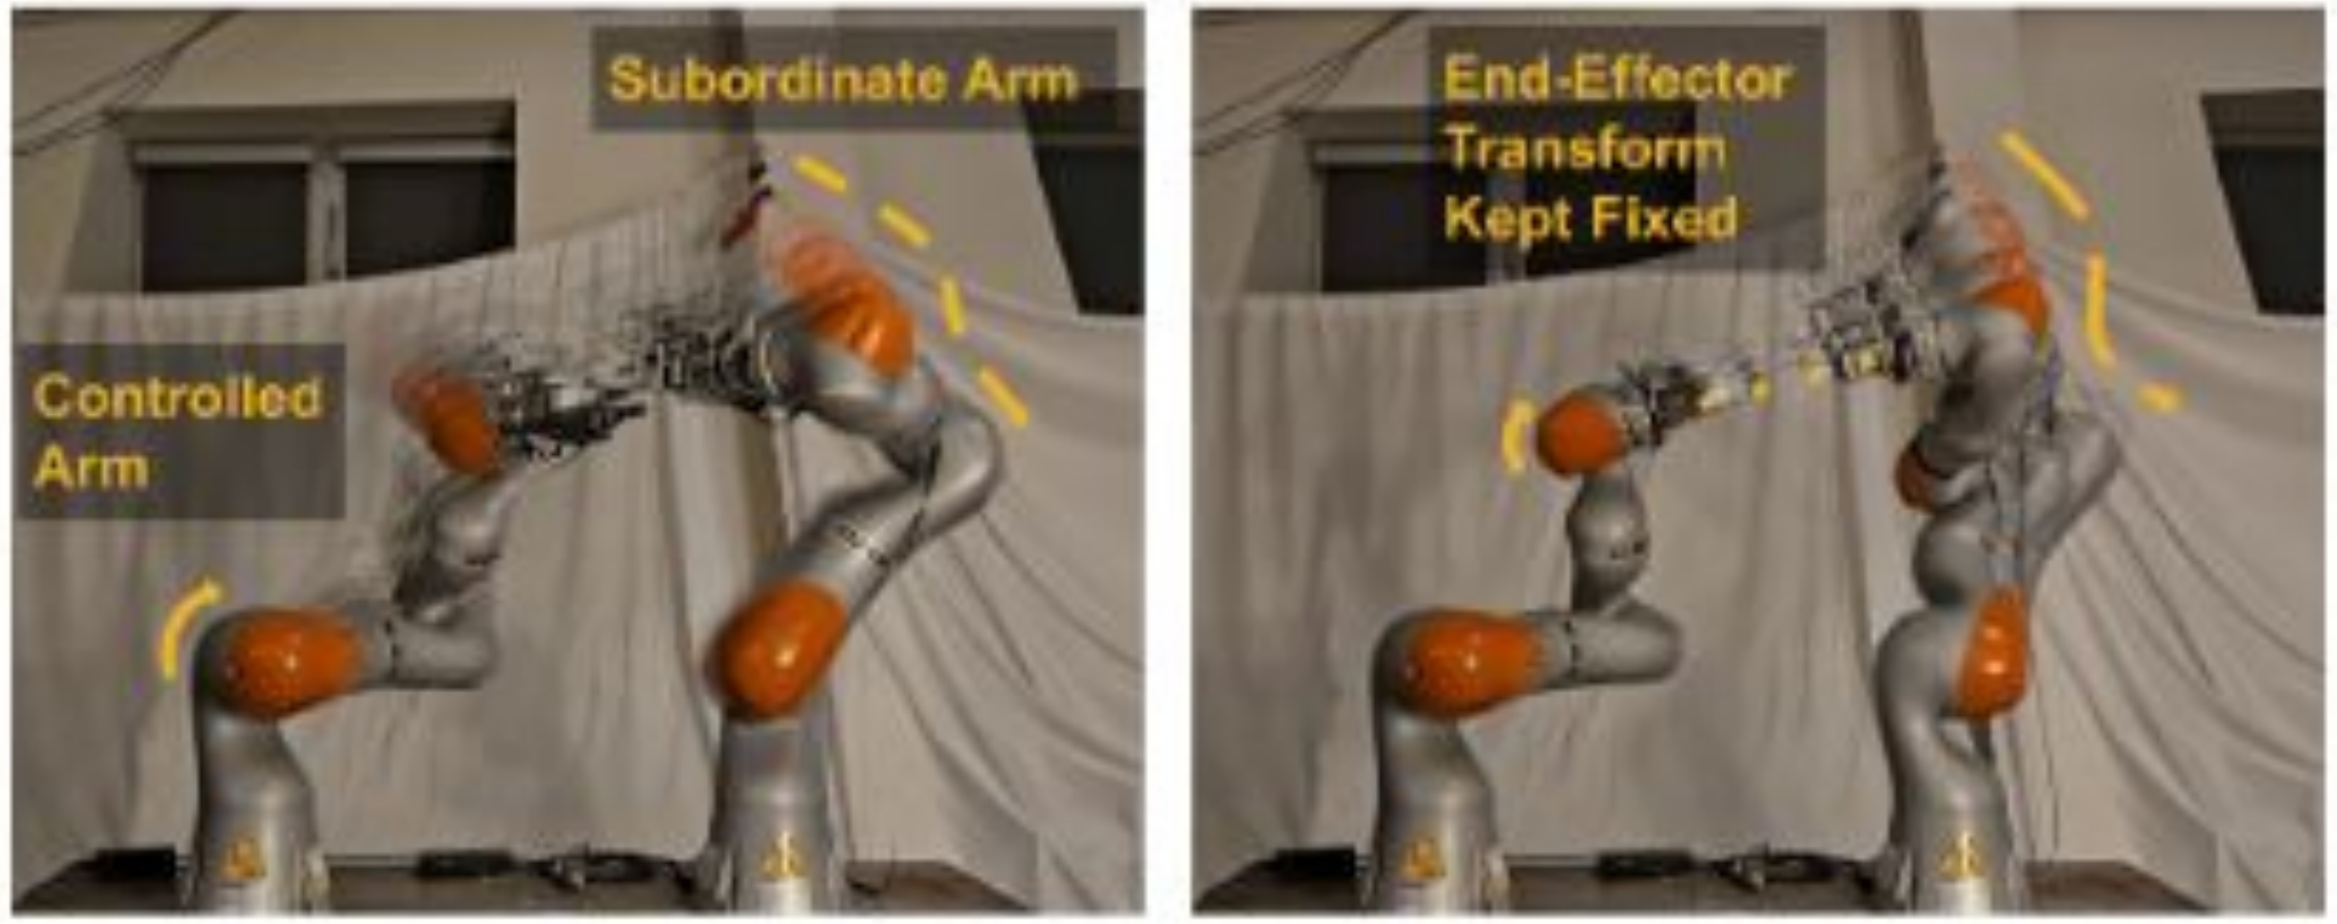

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

Results - setup

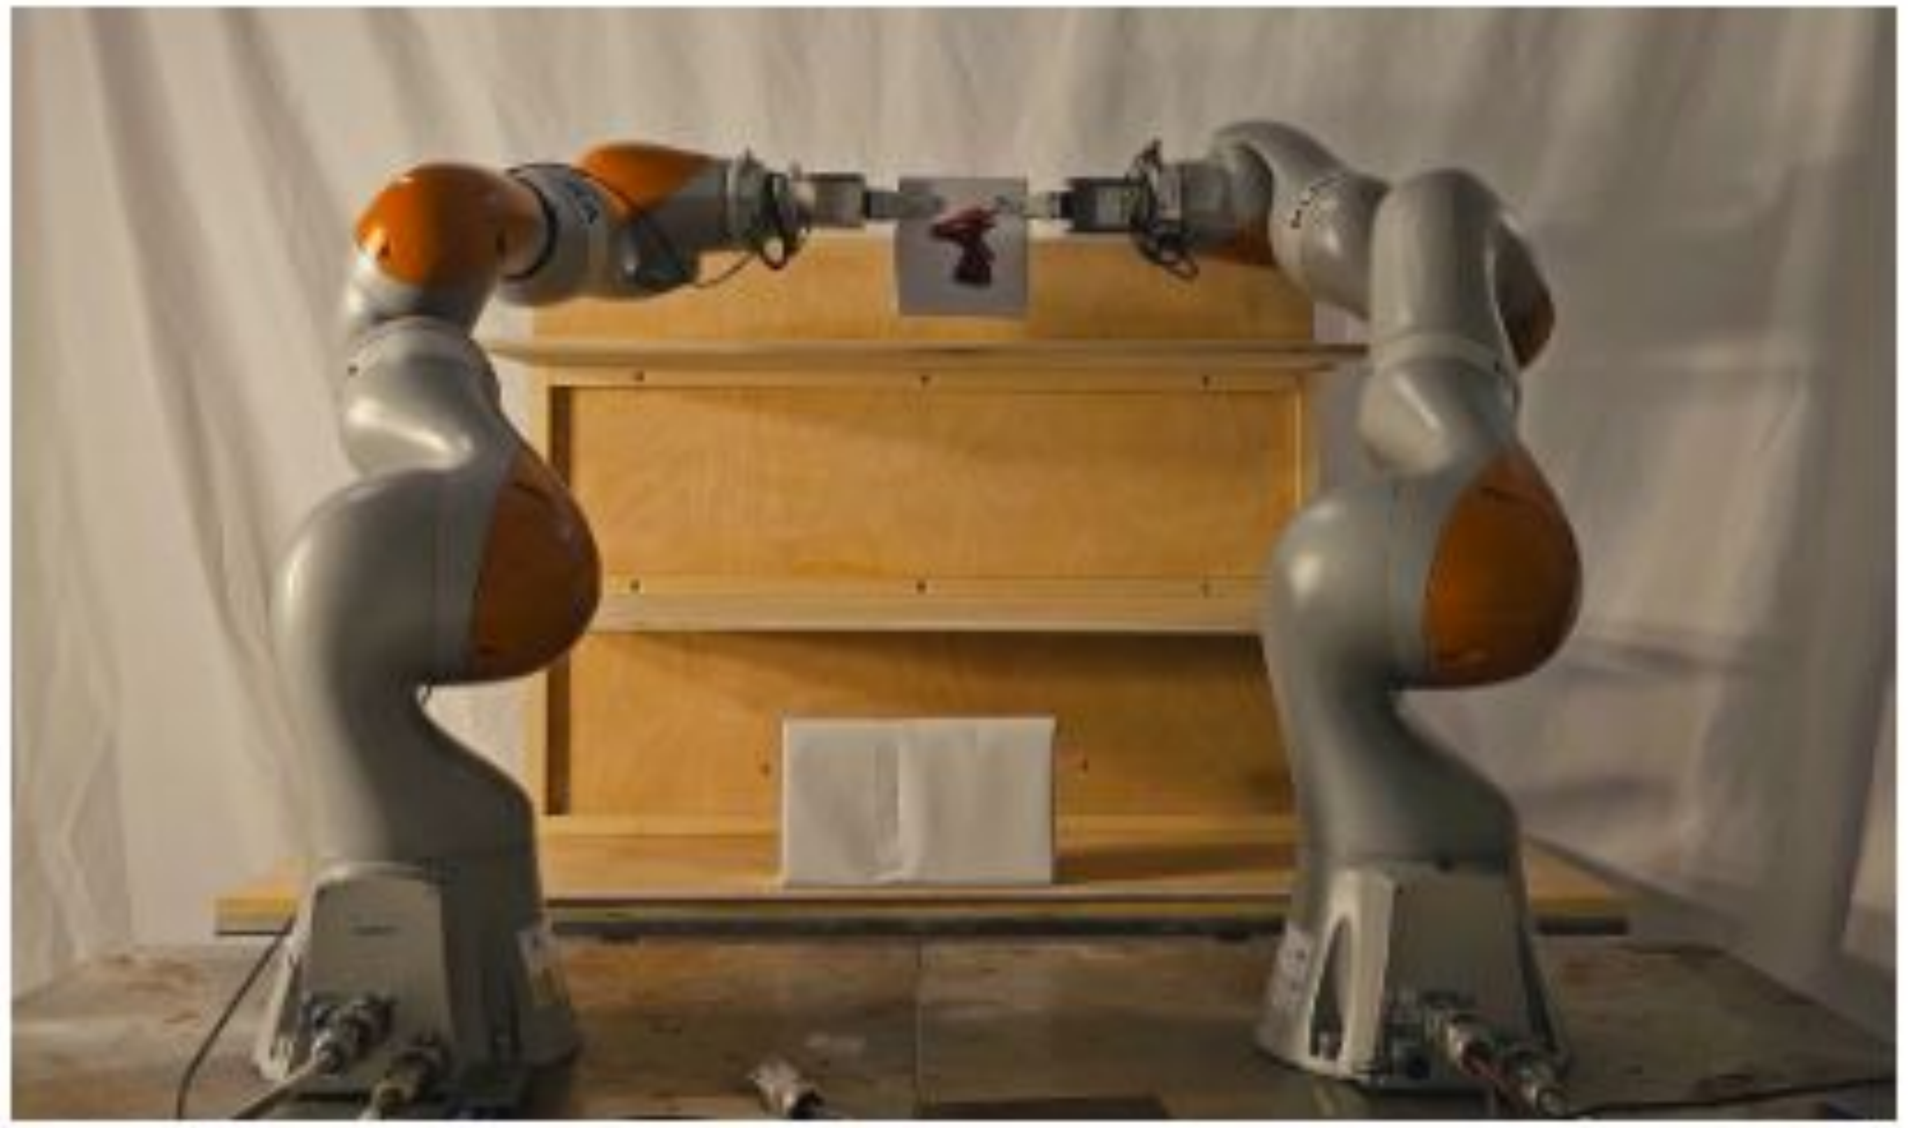

Fig. 1: Hardware setup for our experiments. The two arms must work together to move an objects between the shelves, avoiding collisions and respecting the kinematic constraint.

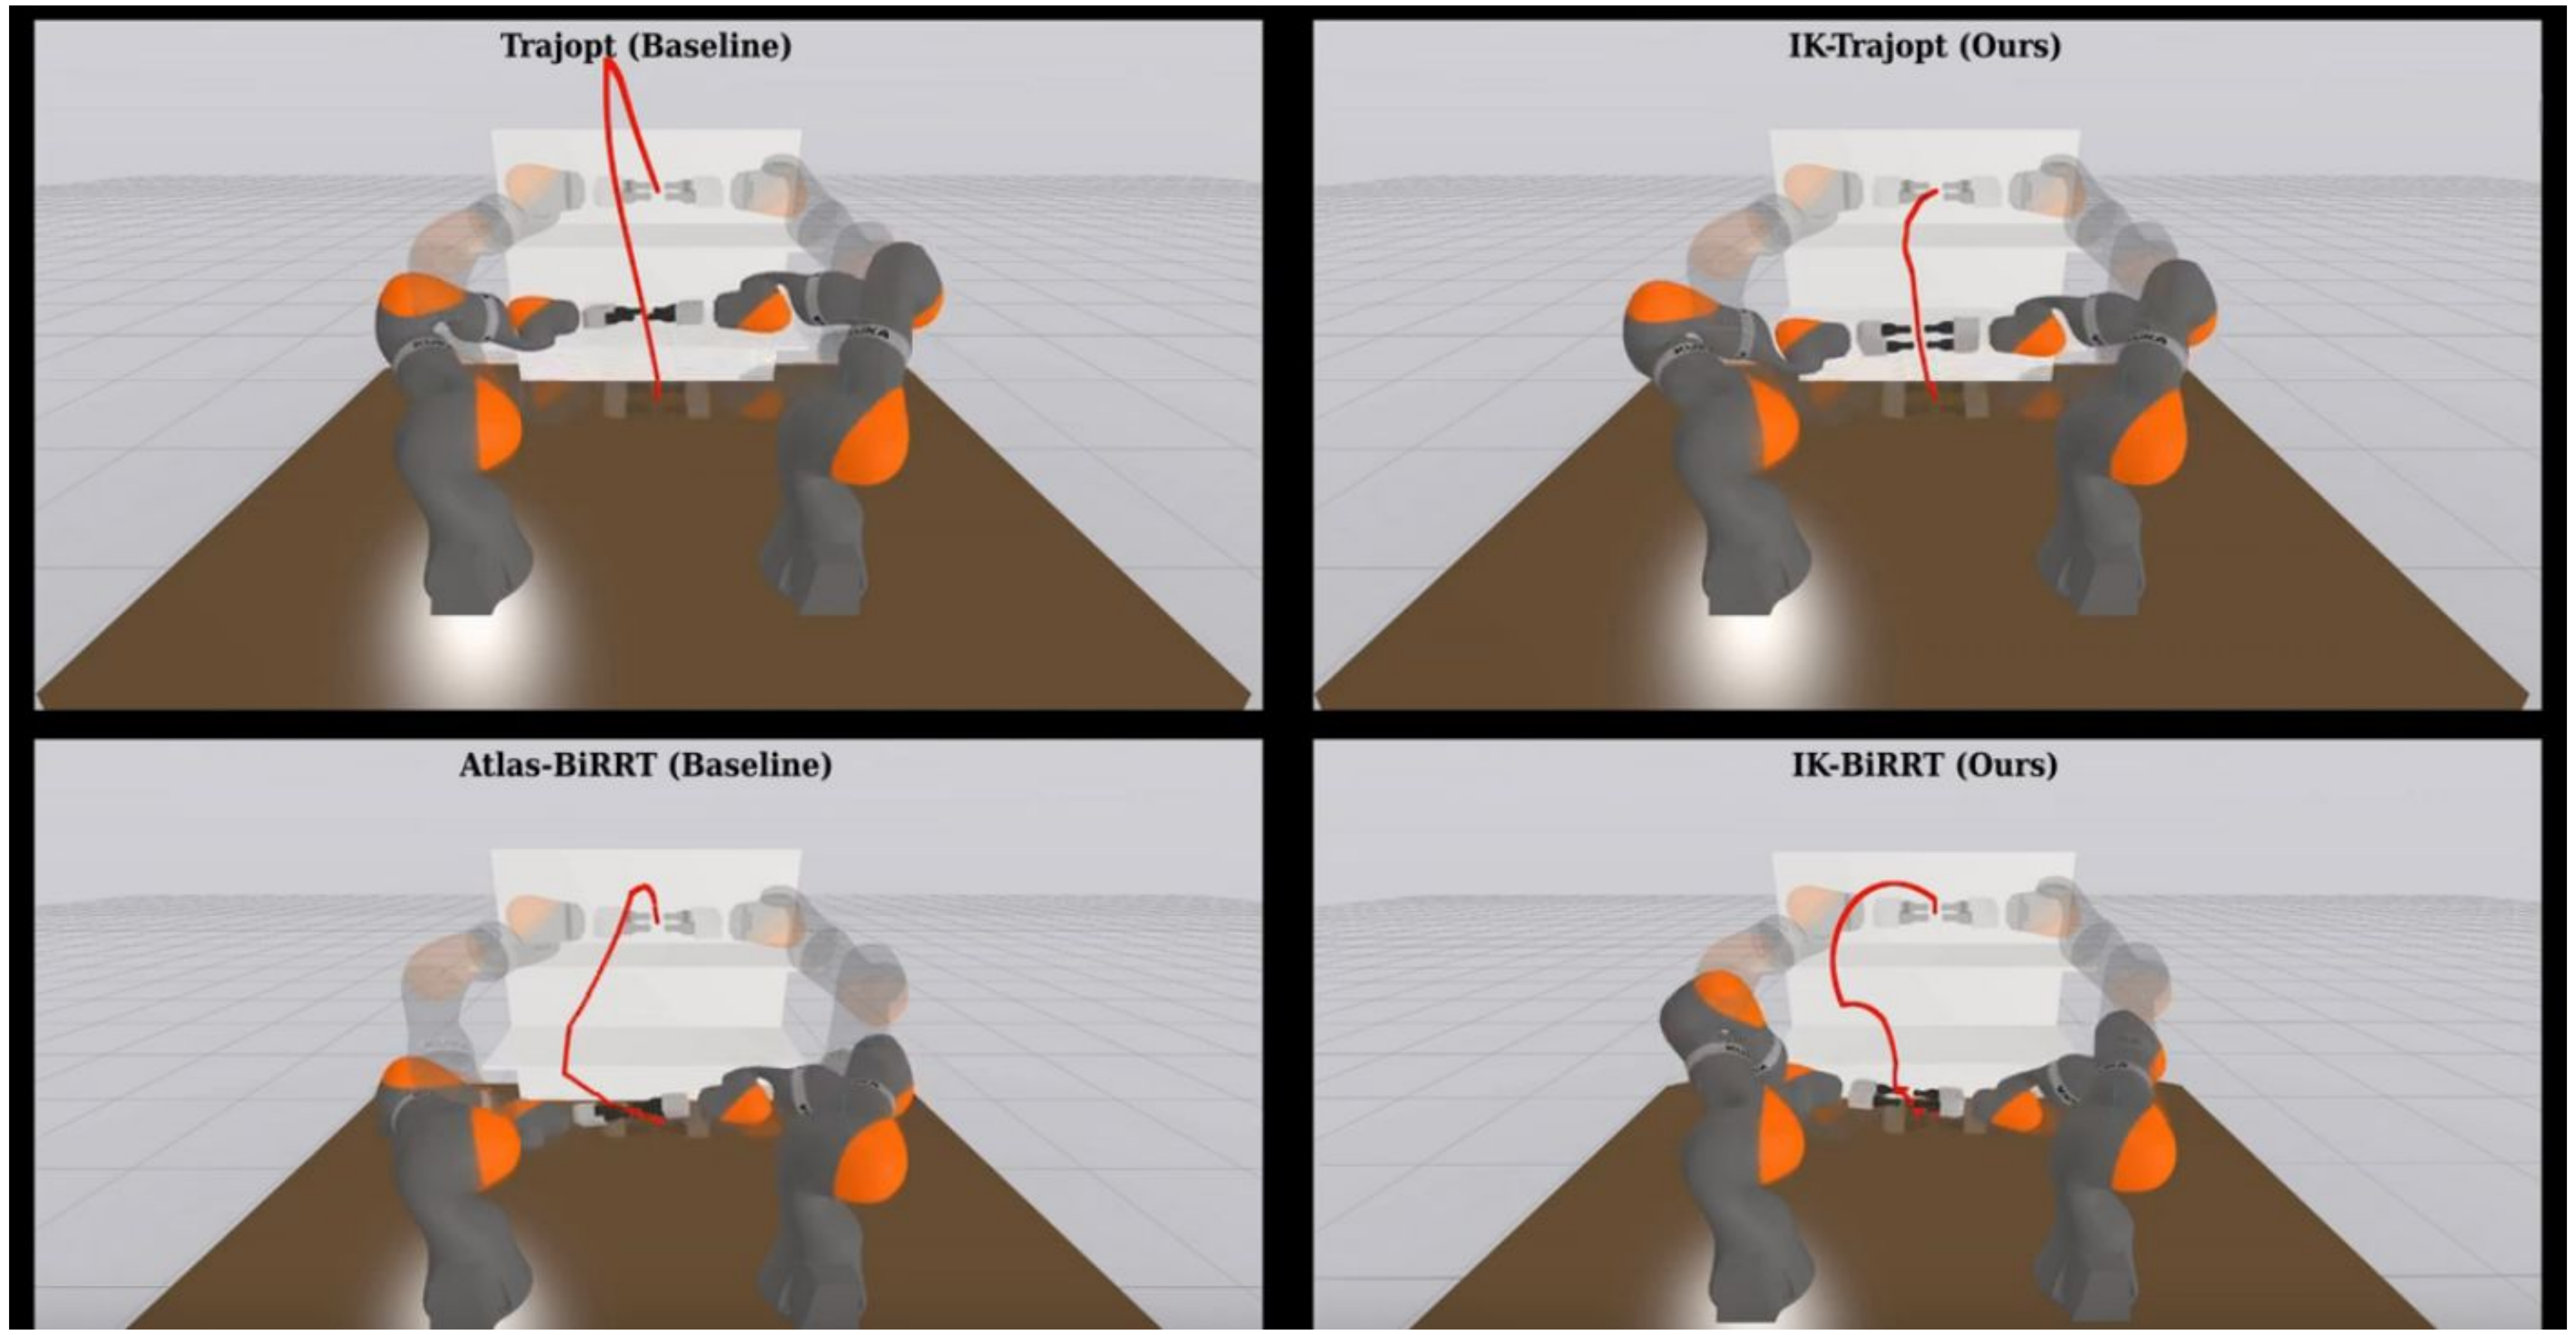

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

Results

Method	Top to Middle	Middle to Bottom	Bottom to Top
Trajopt	4.58*	2.85*	**4.35***
Atlas-BiRRT	4.72	5.04	6.61
Atlas-PRM	5.43	5.67	6.99
IK-Trajopt	4.24*	**1.81***	8.87
IK-BiRRT	9.91	8.69	11.42
IK-PRM	4.67	8.93	9.21
IK-GCS	**2.09**	3.32	5.62

TABLE I: Path lengths (measured in configuration space) for each method with various start and goal configurations. Paths marked with an asterisk were not collision-free.

path arc length (feet)

Method	Top to Middle	Middle to Bottom	Bottom to Top
Trajopt	10.37	5.36	7.25
Atlas-BiRRT	140.82	155.91	201.32
Atlas-PRM	0.69	0.86	0.96
IK-Trajopt	19.48	18.70	22.29
IK-BiRRT	49.42	52.53	54.10
IK-PRM	**0.46**	**0.64**	**0.61**
IK-GCS	3.41	2.32	3.32

TABLE II: Online planning time (in seconds) for each method with various start and goal configurations. Atlas-BiRRT runtimes were only averaged over successful runs (not including timeouts).

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

Results

Method	Top to Middle	Middle to Bottom	Bottom to Top
Trajopt	4.58*	2.85*	**4.35***
Atlas-BiRRT	4.72	5.04	6.61
Atlas-PRM	5.43	5.67	6.99
IK-Trajopt	4.24*	**1.81***	8.87
IK-BiRRT	9.91	8.69	11.42
IK-PRM	4.67	8.93	9.21
IK-GCS	**2.09**	3.32	5.62

TABLE I: Path lengths (measured in configuration space) for each method with various start and goal configurations. Paths marked with an asterisk were not collision-free.

path arc length (feet)

Method	Top to Middle	Middle to Bottom	Bottom to Top
Trajopt	10.37	5.36	7.25
Atlas-BiRRT	140.82	155.91	201.32
Atlas-PRM	0.69	0.86	0.96
IK-Trajopt	19.48	18.70	22.29
IK-BiRRT	49.42	52.53	54.10
IK-PRM	**0.46**	**0.64**	**0.61**
IK-GCS	3.41	2.32	3.32

TABLE II: Online planning time (in seconds) for each method with various start and goal configurations. Atlas-BiRRT runtimes were only averaged over successful runs (not including timeouts).

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

Results

https://www.youtube.com/watch?v=vmujyn4EgTU

Cohn, Thomas, Seiji Shaw, Max Simchowitz, and Russ Tedrake. "Constrained bimanual planning with analytic inverse kinematics." In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6935-6942. IEEE, 2024.

Policy learning approaches

Methods - Imitation Learning

Recall:

IBC: Implicit Behavior Cloning

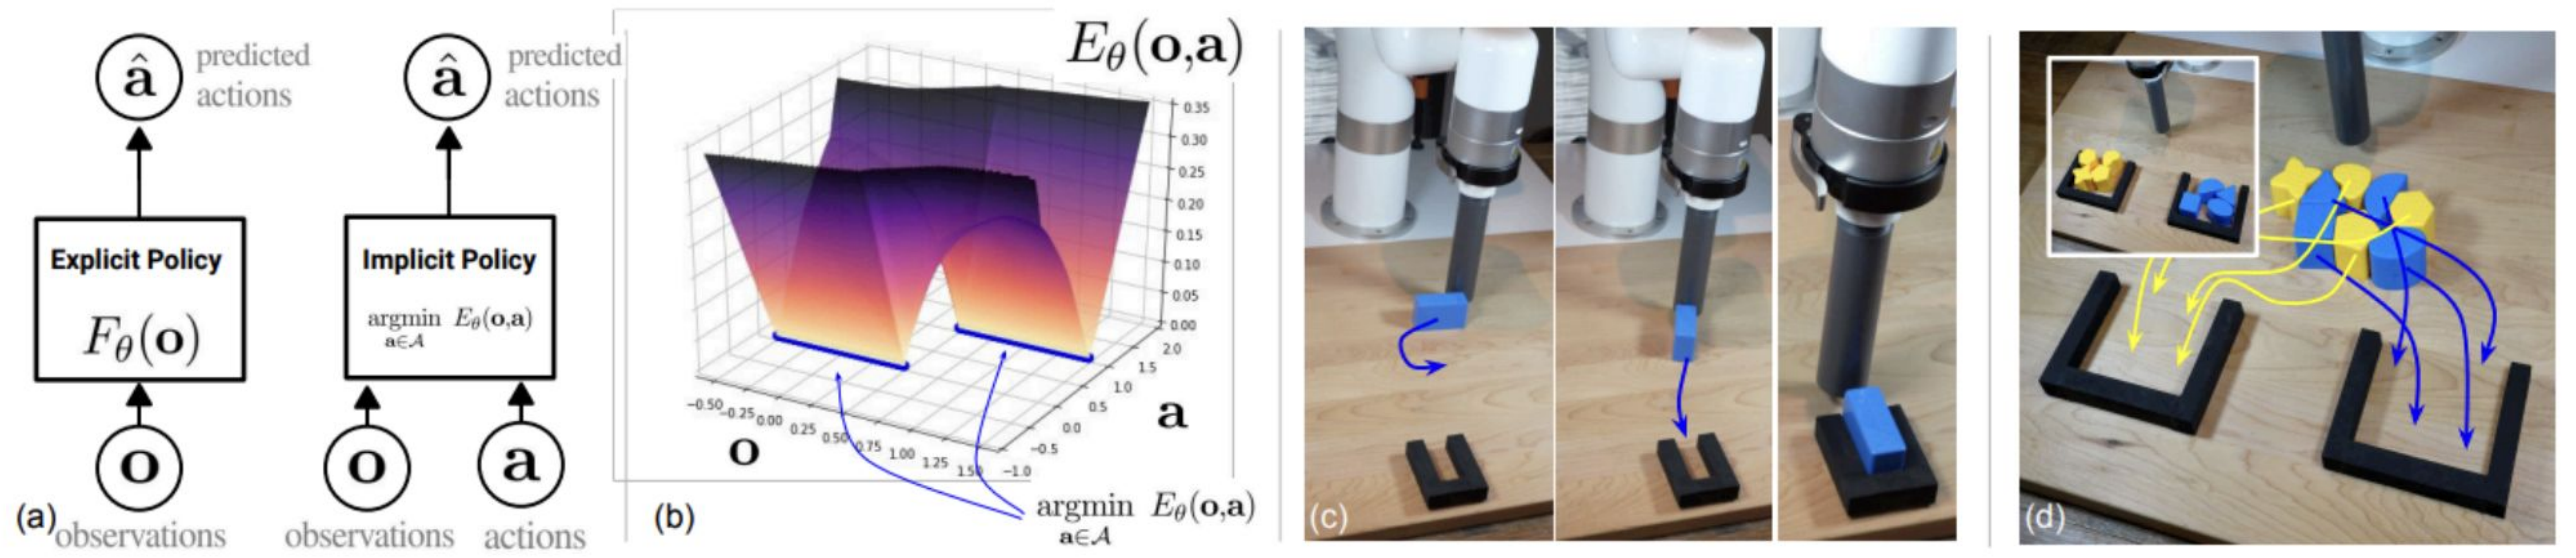

Figure 1. (a) In contrast to explicit policies, implicit policies leverage parameterized energy functions that take both observations (e.g. images) and actions as inputs, and optimize for actions that minimize the energy landscape (b). For learning complex, closed-loop, multimodal visuomotor tasks such as precise block insertion (c) and sorting (d) from human demonstrations, implicit policies perform substantially better than explicit ones.

Much more sample efficient than RL!

ALOHA Unleashed

T. Z. Zhao, J. Tompson, D. Driess, P. Florence, S. K. S. Ghasemipour, C. Finn, and A. Wahid. ALOHA unleashed: A simple recipe for robot dexterity. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=gvdXE7ikHI.

ALOHA Unleashed - Data

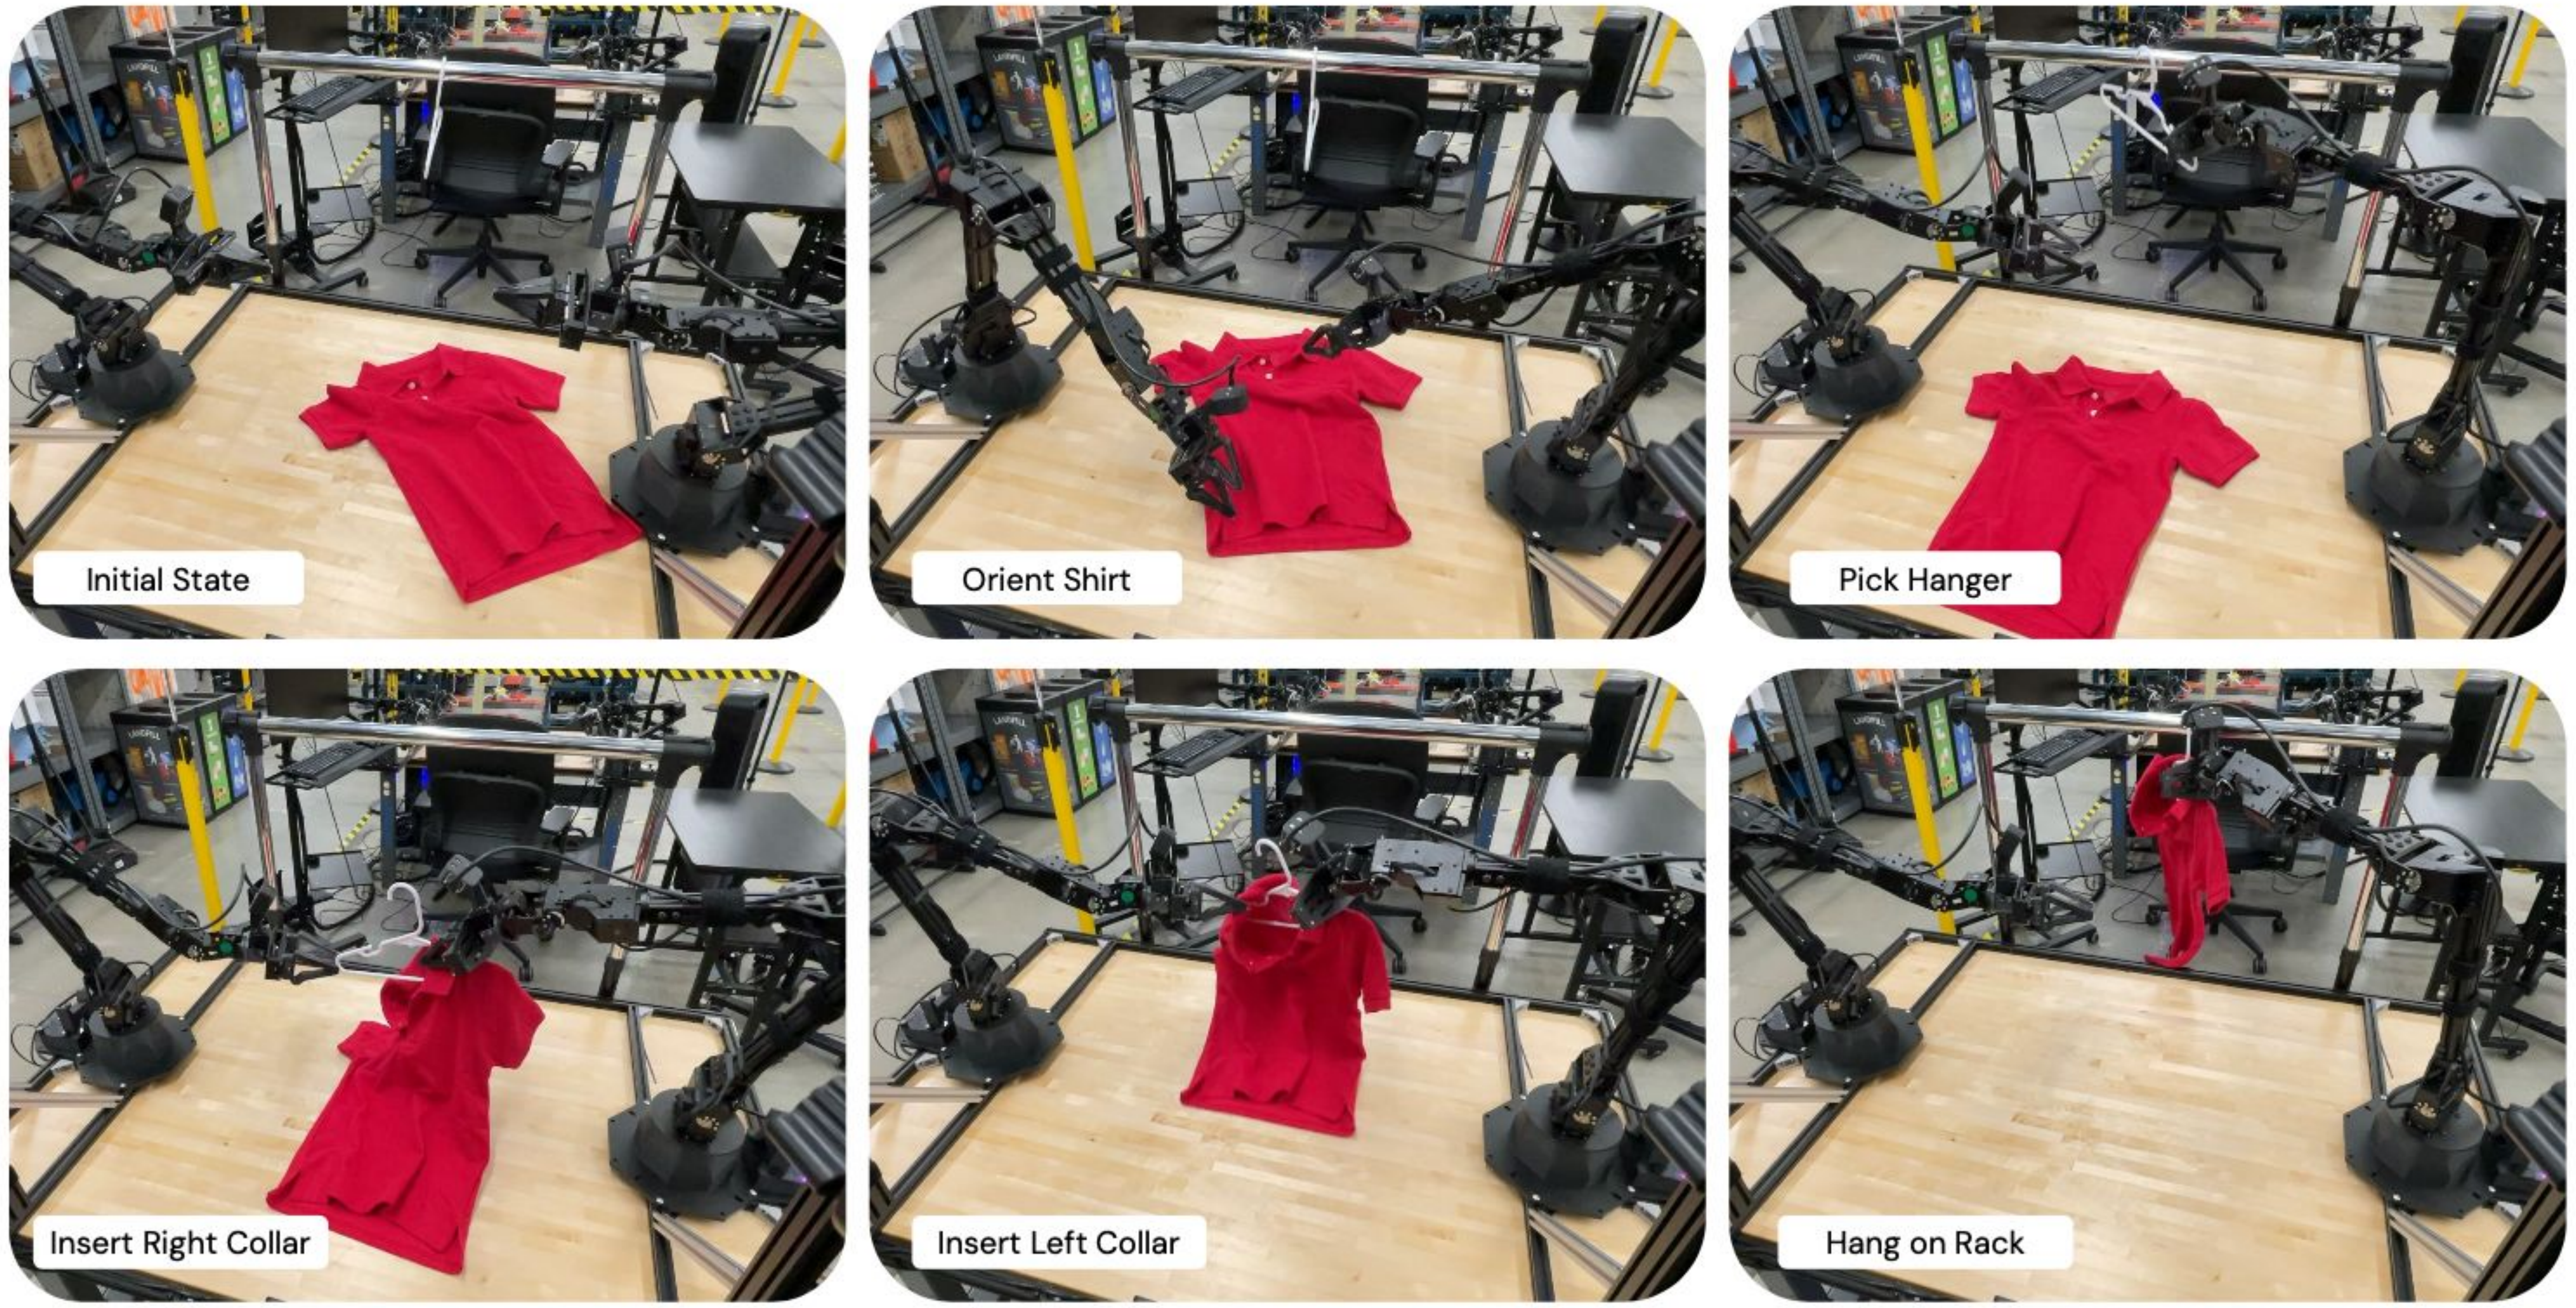

- ALOHA allows bimanual teleoperation for data collection
- 5 different tasks
- Tasks are somewhat long horizon and require precision and dexterity

T. Z. Zhao, J. Tompson, D. Driess, P. Florence, S. K. S. Ghasemipour, C. Finn, and A. Wahid. ALOHA unleashed: A simple recipe for robot dexterity. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=gvdXE7ikHI.

ALOHA Unleashed - Architecture

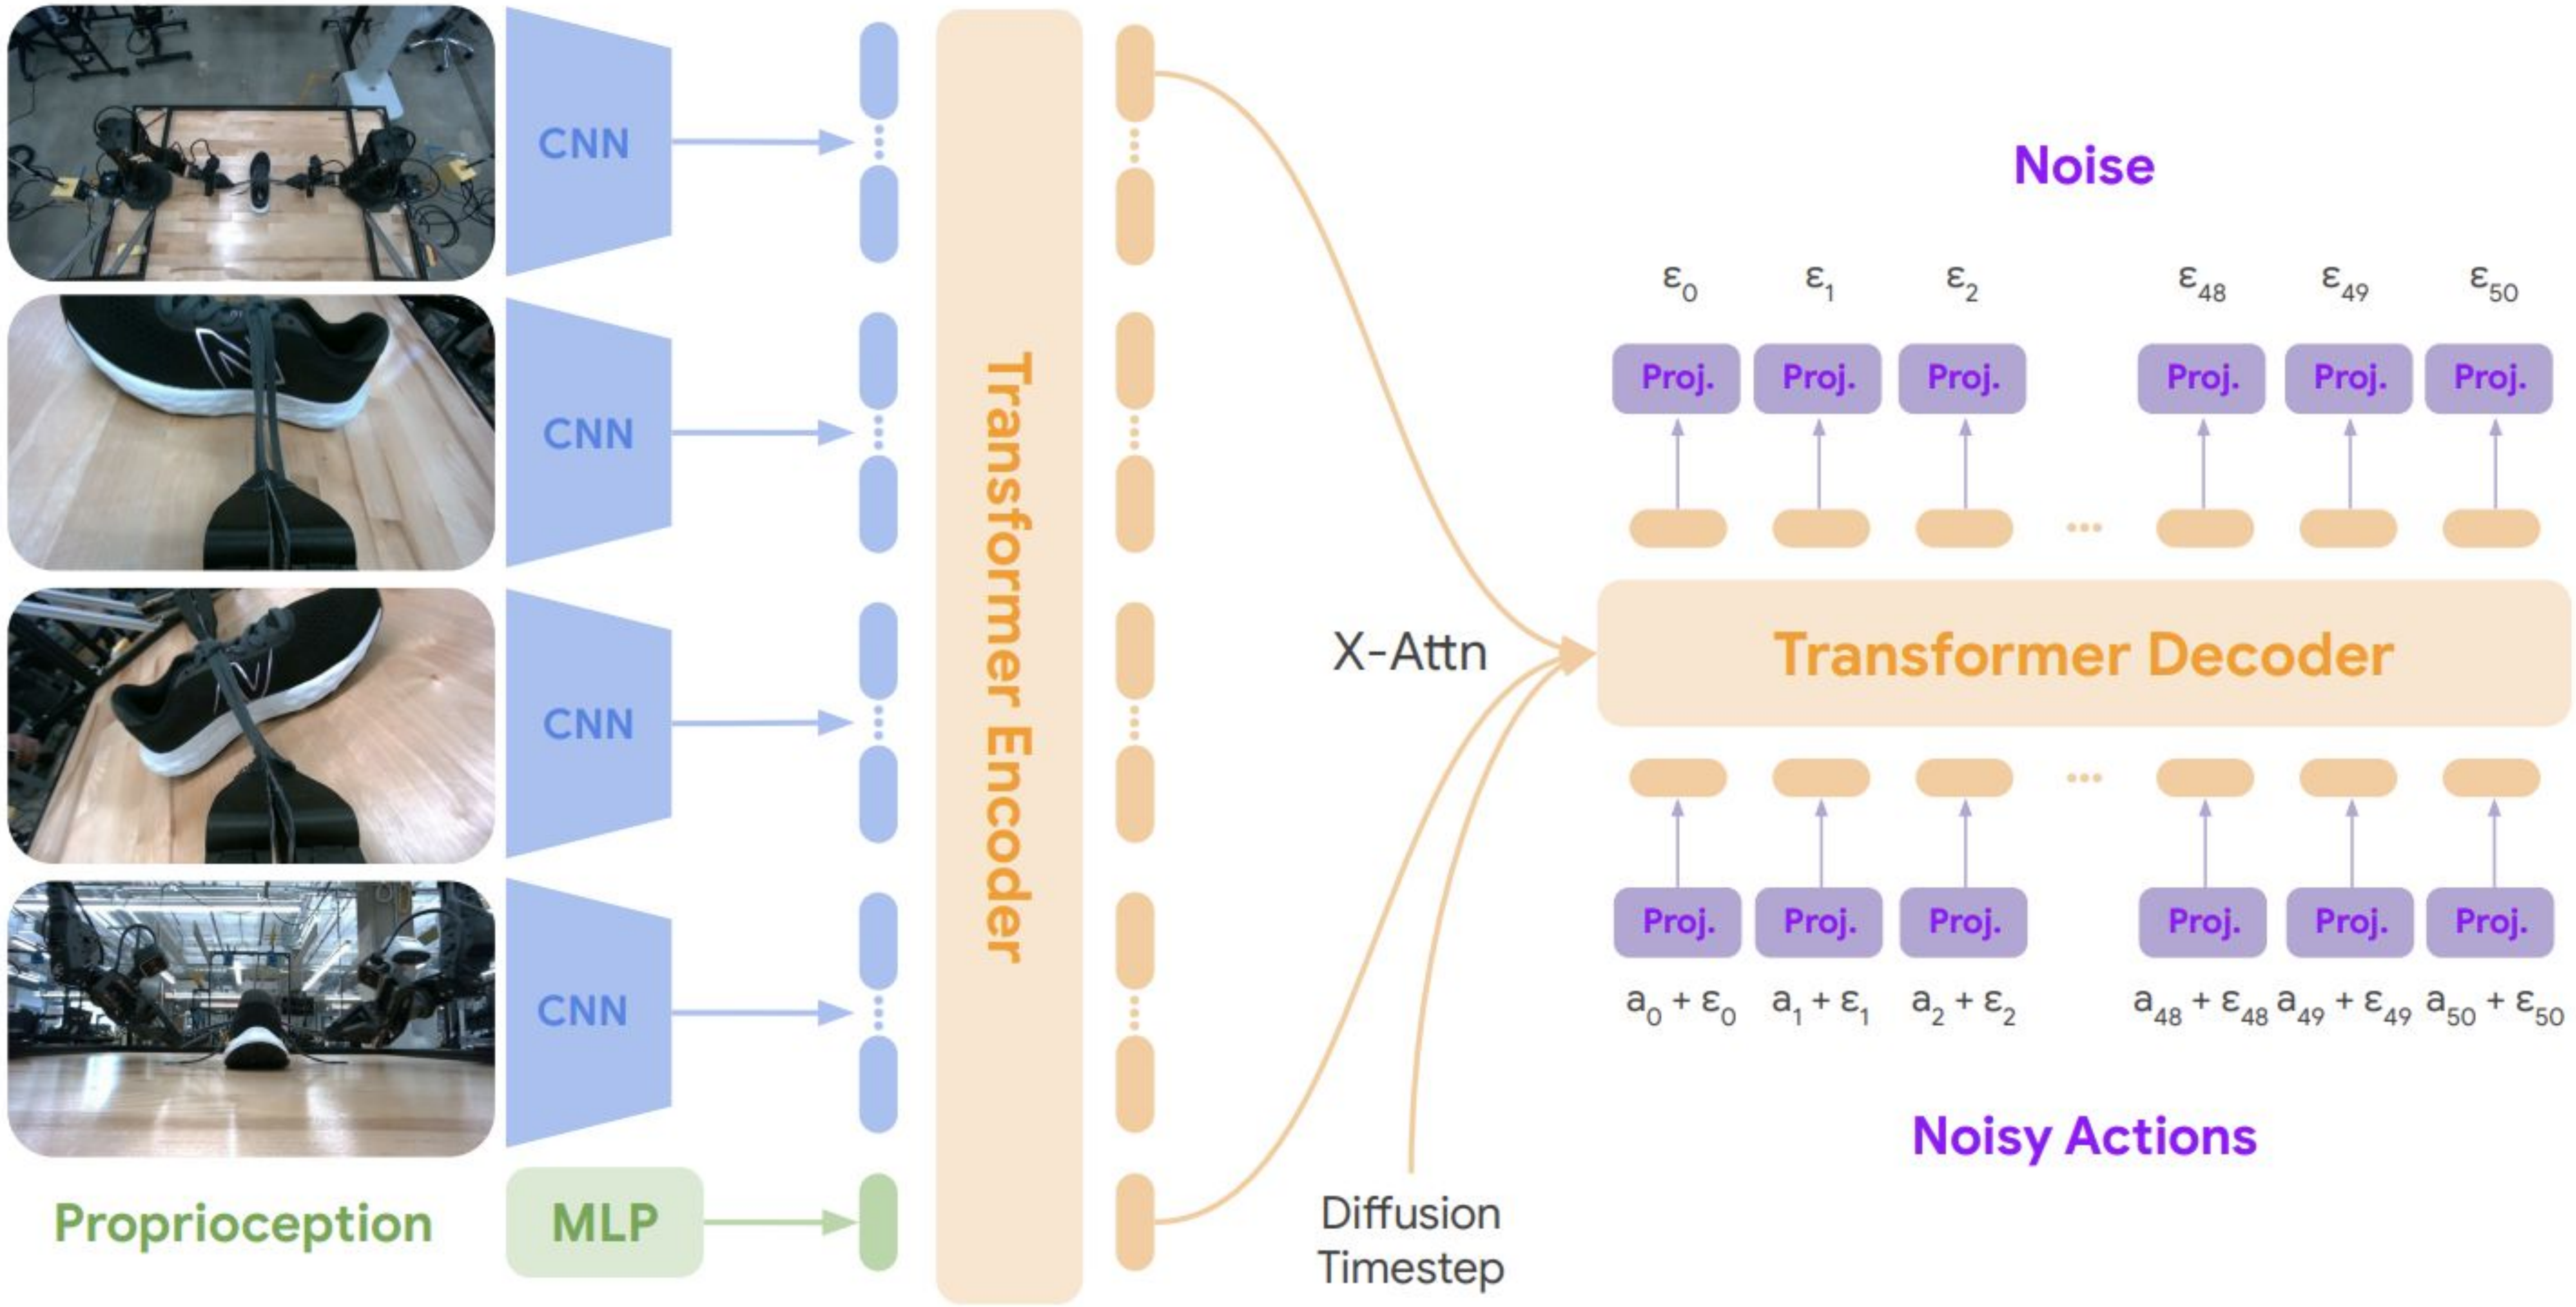

T. Z. Zhao, J. Tompson, D. Driess, P. Florence, S. K. S. Ghasemipour, C. Finn, and A. Wahid. ALOHA unleashed: A simple recipe for robot dexterity. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=gvdXE7ikHI.

ALOHA Unleashed - Architecture

- Encoder-decoder architecture with diffusion loss
- 4 cameras + proprioception
- CNNs are ResNet-50s
- 50 diffusion steps(ie: during inference decoder runs 50 times)
- This is just diffusion policy!

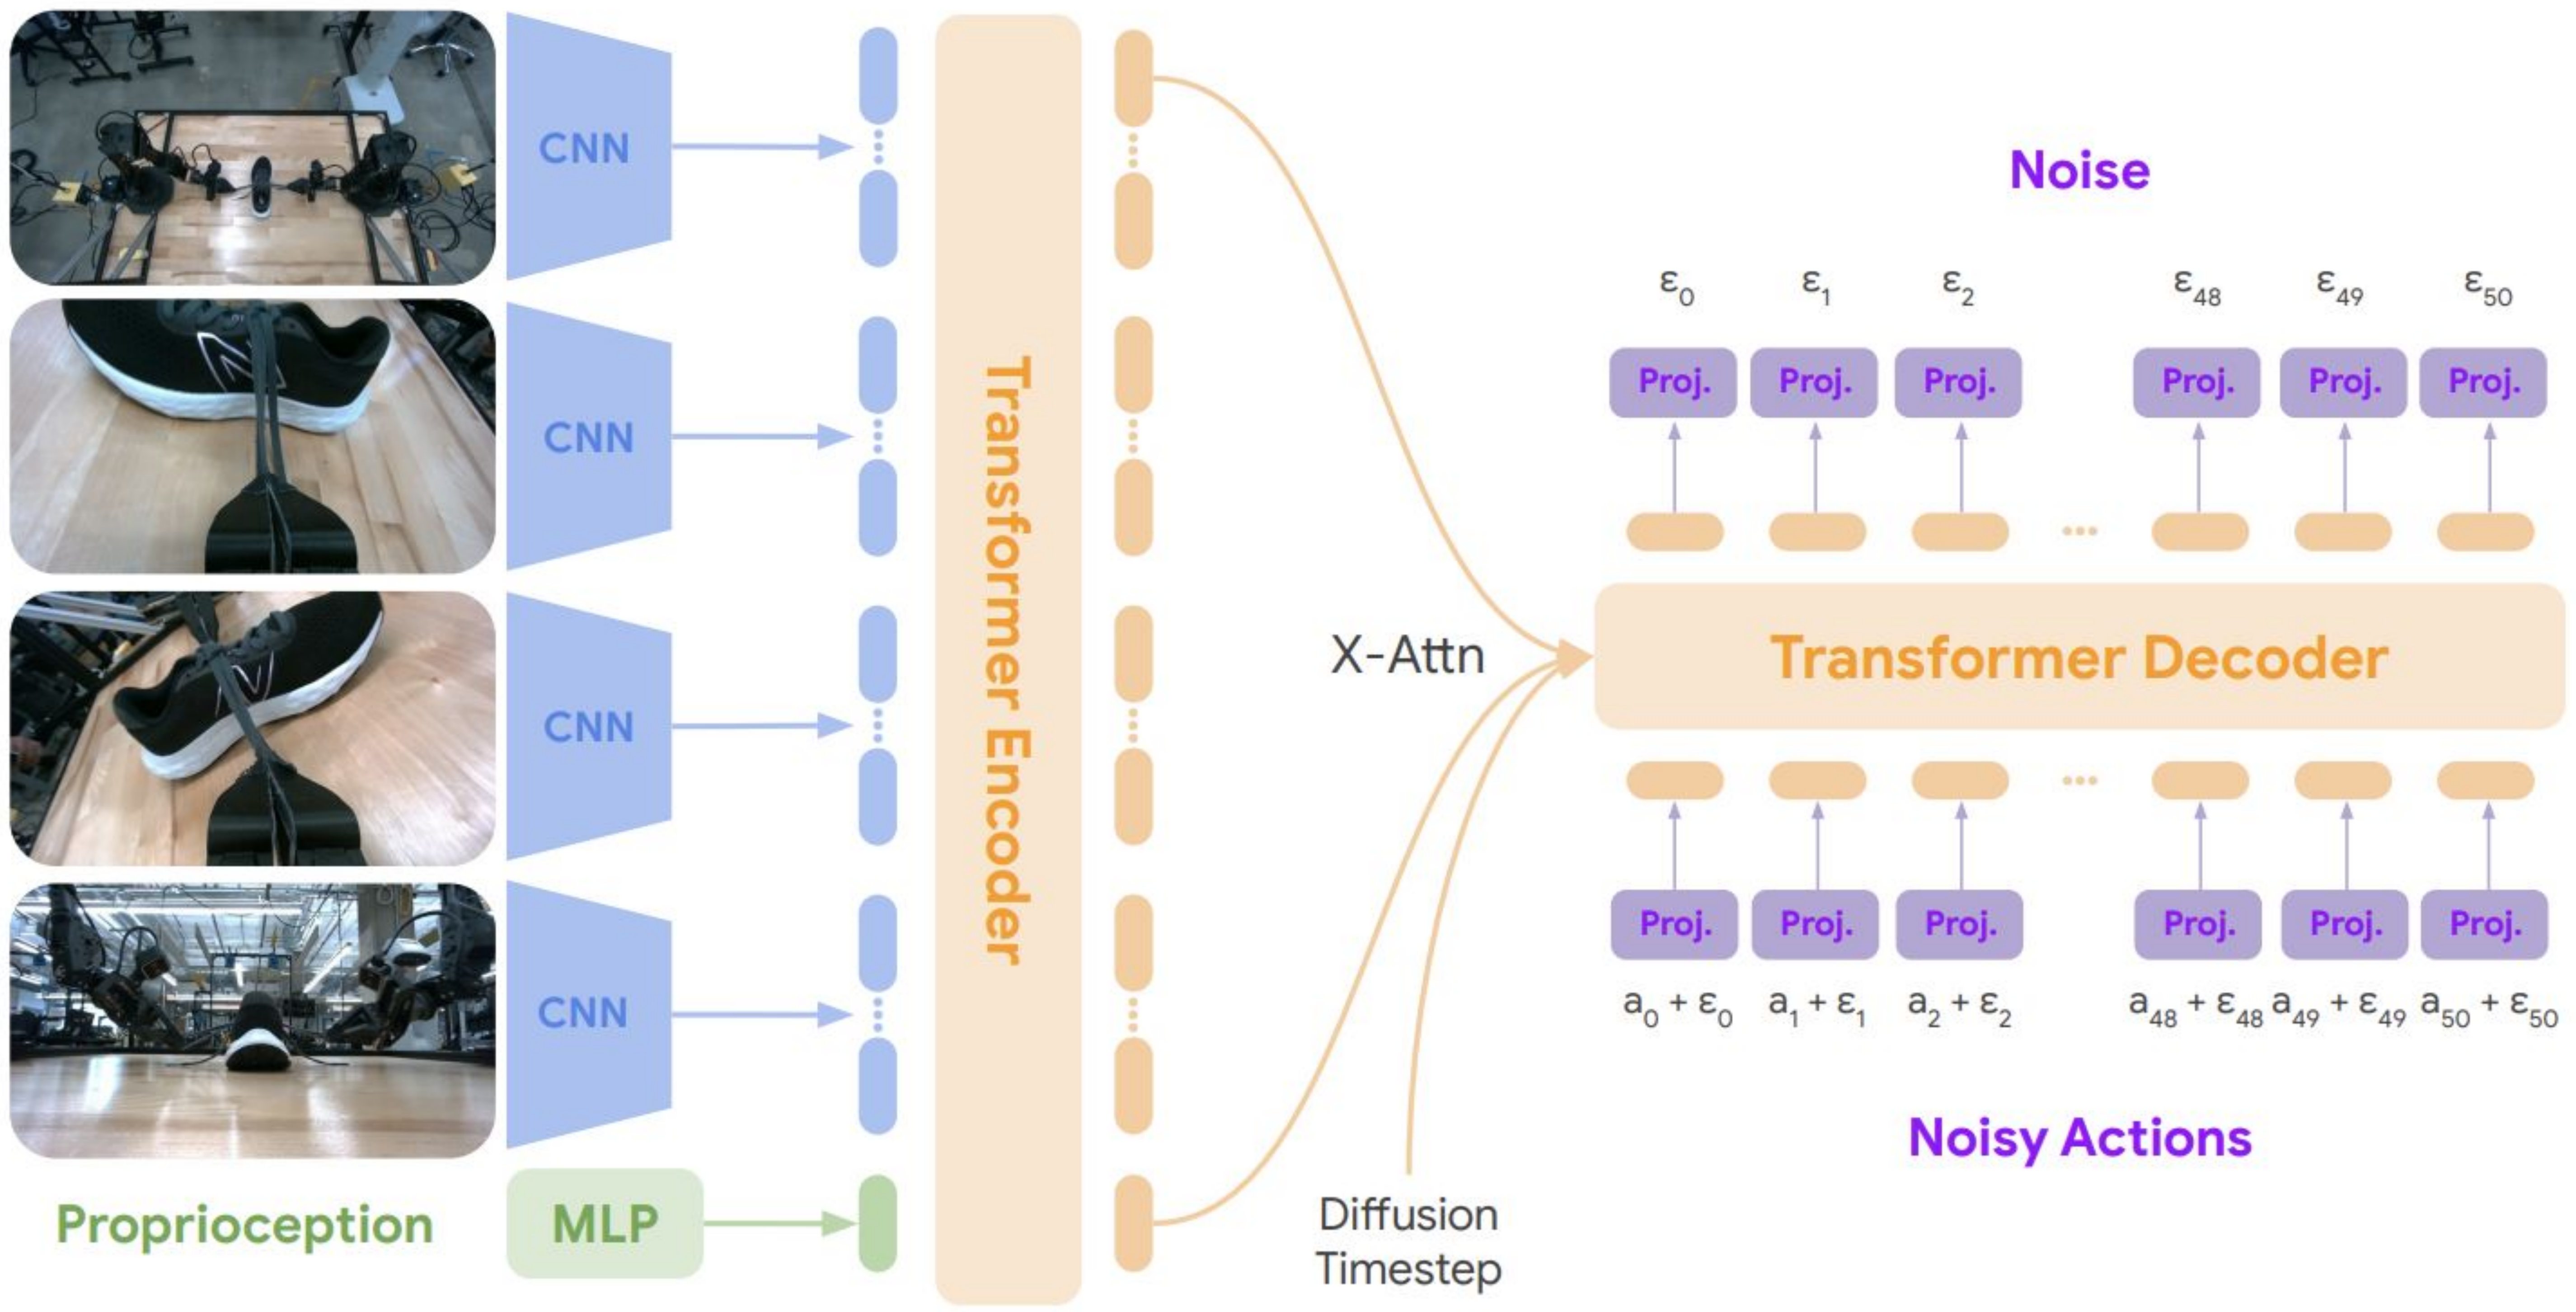

T. Z. Zhao, J. Tompson, D. Driess, P. Florence, S. K. S. Ghasemipour, C. Finn, and A. Wahid. ALOHA unleashed: A simple recipe for robot dexterity. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=gvdXE7ikHI.

ALOHA Unleashed - Results

- Messy demonstrations help the agent learn to recover from mistakes

Task	Success Rate	Number of Demonstrations
ShirtEasy	75%	8658 (5345 Easy; 3313 Messy)
ShirtMessy	70%	
LaceEasy	70%	5133 (2212 Easy; 2921 Messy)
LaceMessy	40%	
FingerReplace	75%	5247
GearInsert-1	95%	4005
GearInsert-2	75%	
GearInsert-3	40%	
RandomKitchen-Bowl	95%	3198 (216 In-Domain)
RandomKitchen-Bowl+Cup	65%	
RandomKitchen-Bowl+Cup+Fork	25%	

T. Z. Zhao, J. Tompson, D. Driess, P. Florence, S. K. S. Ghasemipour, C. Finn, and A. Wahid. ALOHA unleashed: A simple recipe for robot dexterity. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=gvdXE7ikHI.

ALOHA Unleashed - Results

- Messy demonstrations help the agent learn to recover from mistakes

Task	DP (S)	DP (XS-LowRes)	ACT (XS-LowRes)	Num Demos
SingleInsertion (sim)	72	58 ±3	32	522
DoubleInsertion (sim)	60	48 ±2	58	201
MugOnPlate (sim)	80	74 ±0	40	550

Task	DP (S)	ACT (150M)	Num Demos
ShirtMessy (real)	70	25	8658

Task	Success Rate	Number of Demonstrations
ShirtEasy ShirtMessy	75% 70%	8658 (5345 Easy; 3313 Messy)
LaceEasy LaceMessy	70% 40%	5133 (2212 Easy; 2921 Messy)
FingerReplace	75%	5247
GearInsert-1 GearInsert-2 GearInsert-3	95% 75% 40%	4005
RandomKitchen-Bowl RandomKitchen-Bowl+Cup RandomKitchen-Bowl+Cup+Fork	95% 65% 25%	3198 (216 In-Domain)

- Outperforms previous methods significantly

T. Z. Zhao, J. Tompson, D. Driess, P. Florence, S. K. S. Ghasemipour, C. Finn, and A. Wahid. ALOHA unleashed: A simple recipe for robot dexterity. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=gvdXE7ikHI.

ALOHA Unleashed - Results

- Ridiculous number of demonstrations required
- Messy demonstrations help the agent learn to recover from mistakes

Task	Success Rate	Number of Demonstrations
ShirtEasy ShirtMessy	75% 70%	8658 (5345 Easy; 3313 Messy)
LaceEasy LaceMessy	70% 40%	5133 (2212 Easy; 2921 Messy)
FingerReplace	75%	5247
GearInsert-1 GearInsert-2 GearInsert-3	95% 75% 40%	4005
RandomKitchen-Bowl RandomKitchen-Bowl+Cup RandomKitchen-Bowl+Cup+Fork	95% 65% 25%	3198 (216 In-Domain)

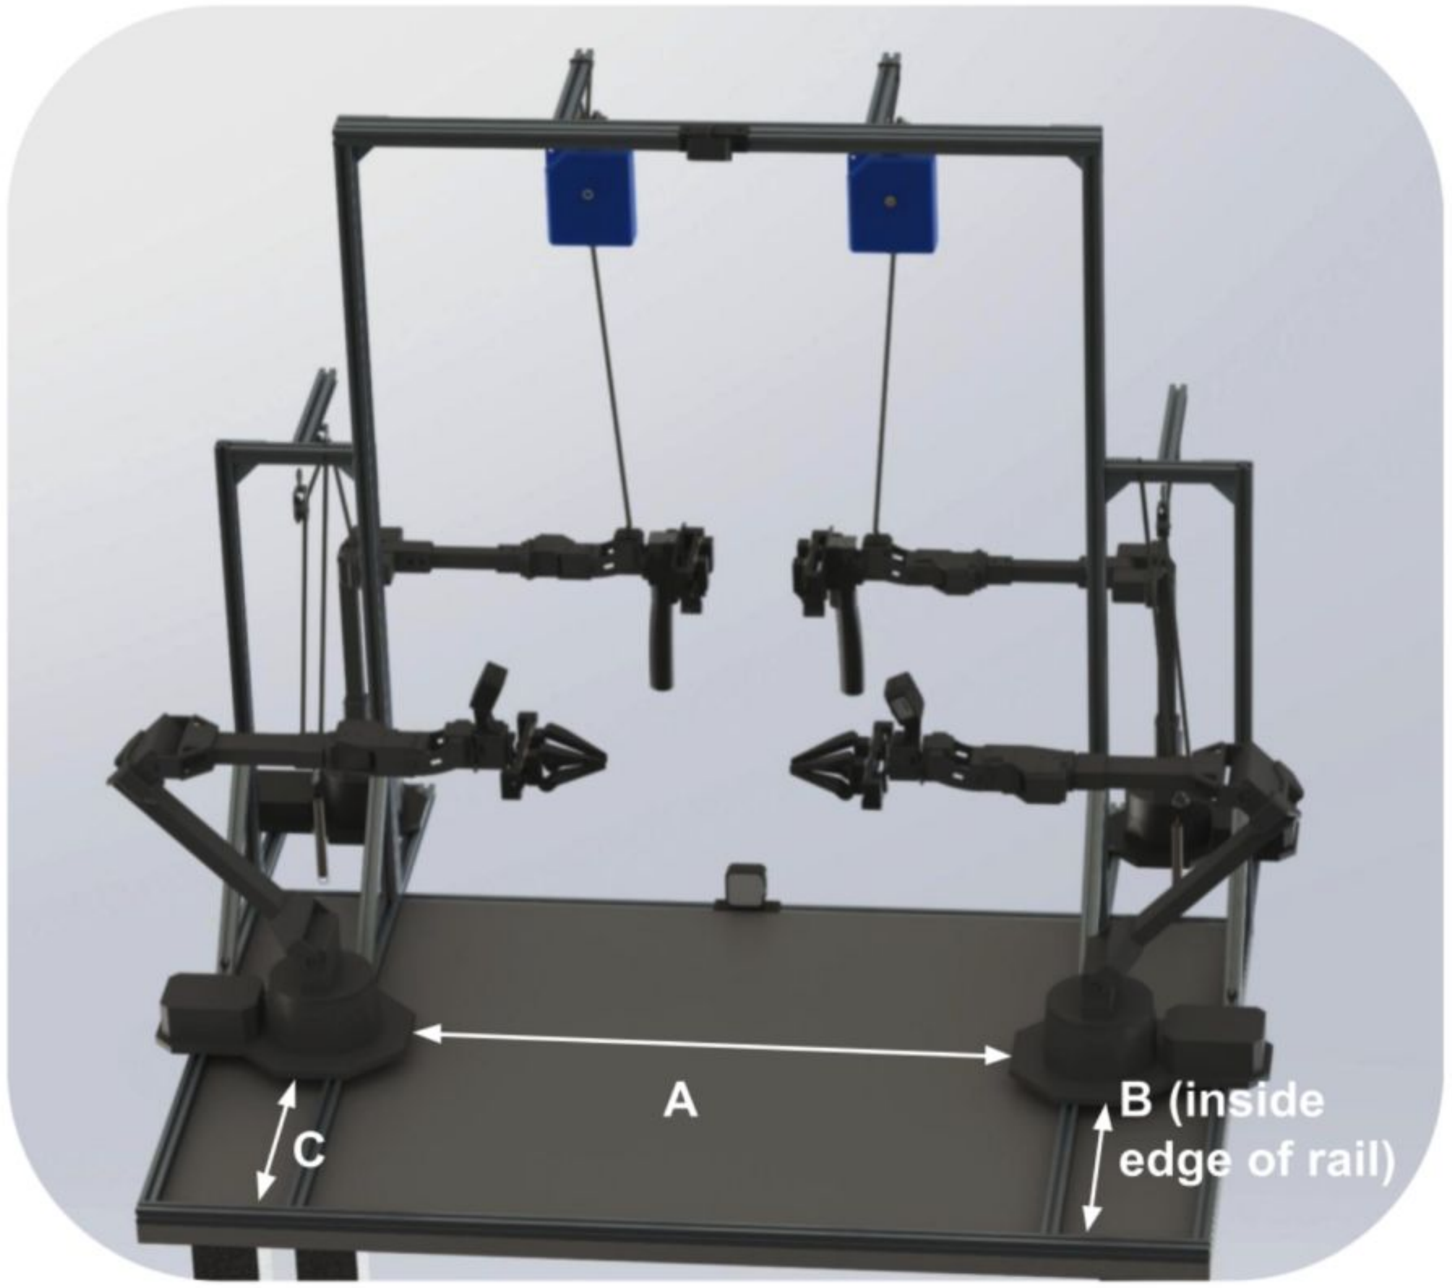

Ridiculous amount of demonstrations.

T. Z. Zhao, J. Tompson, D. Driess, P. Florence, S. K. S. Ghasemipour, C. Finn, and A. Wahid. ALOHA unleashed: A simple recipe for robot dexterity. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=gvdXE7ikHI.

VLAs - Briefly

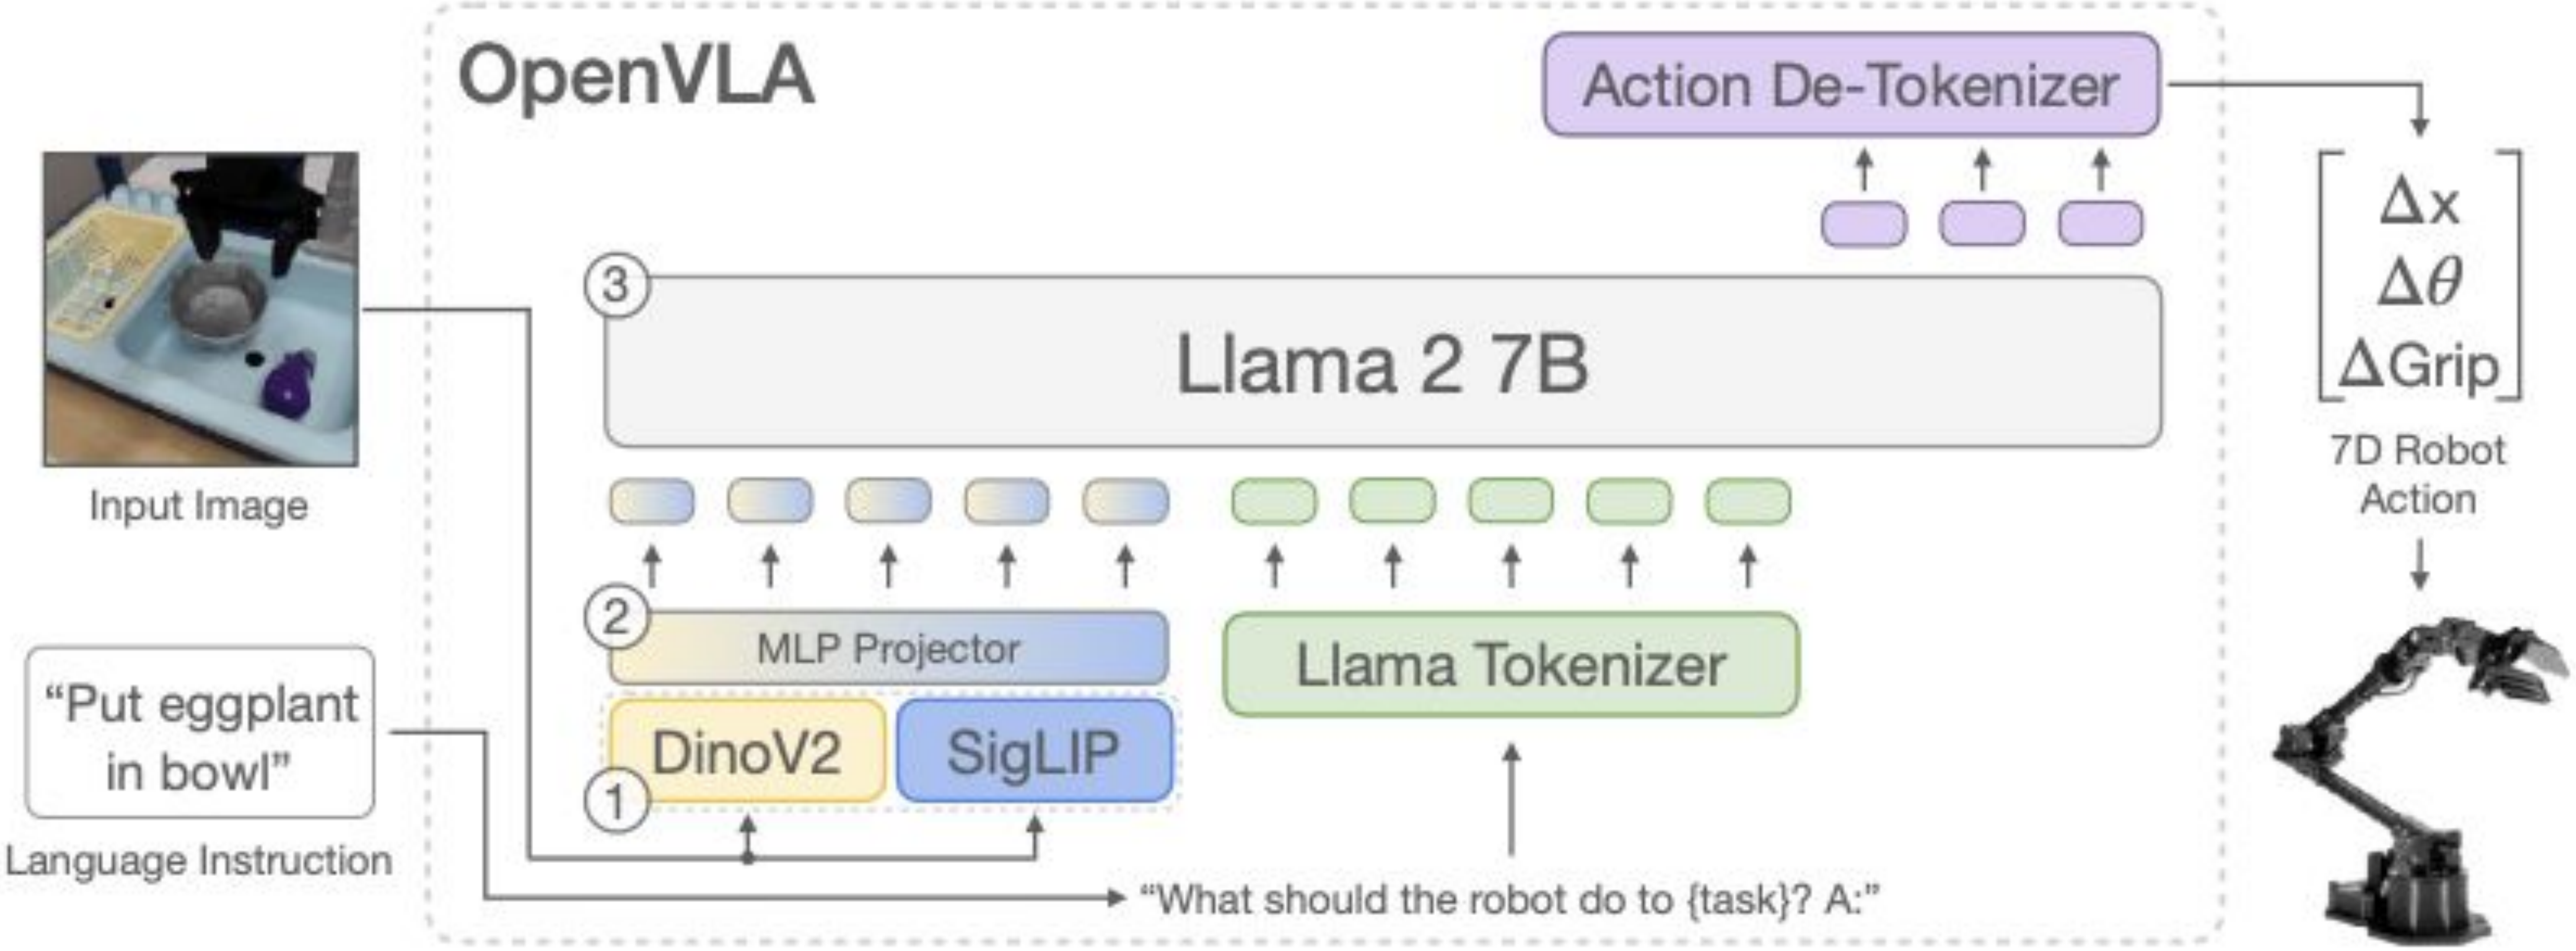

Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. arXiv preprint arXiv:2406.09246, 2024.

VLAs - Briefly

- Make use of LLMs
- Visual understanding from SigLIP and DinoV2

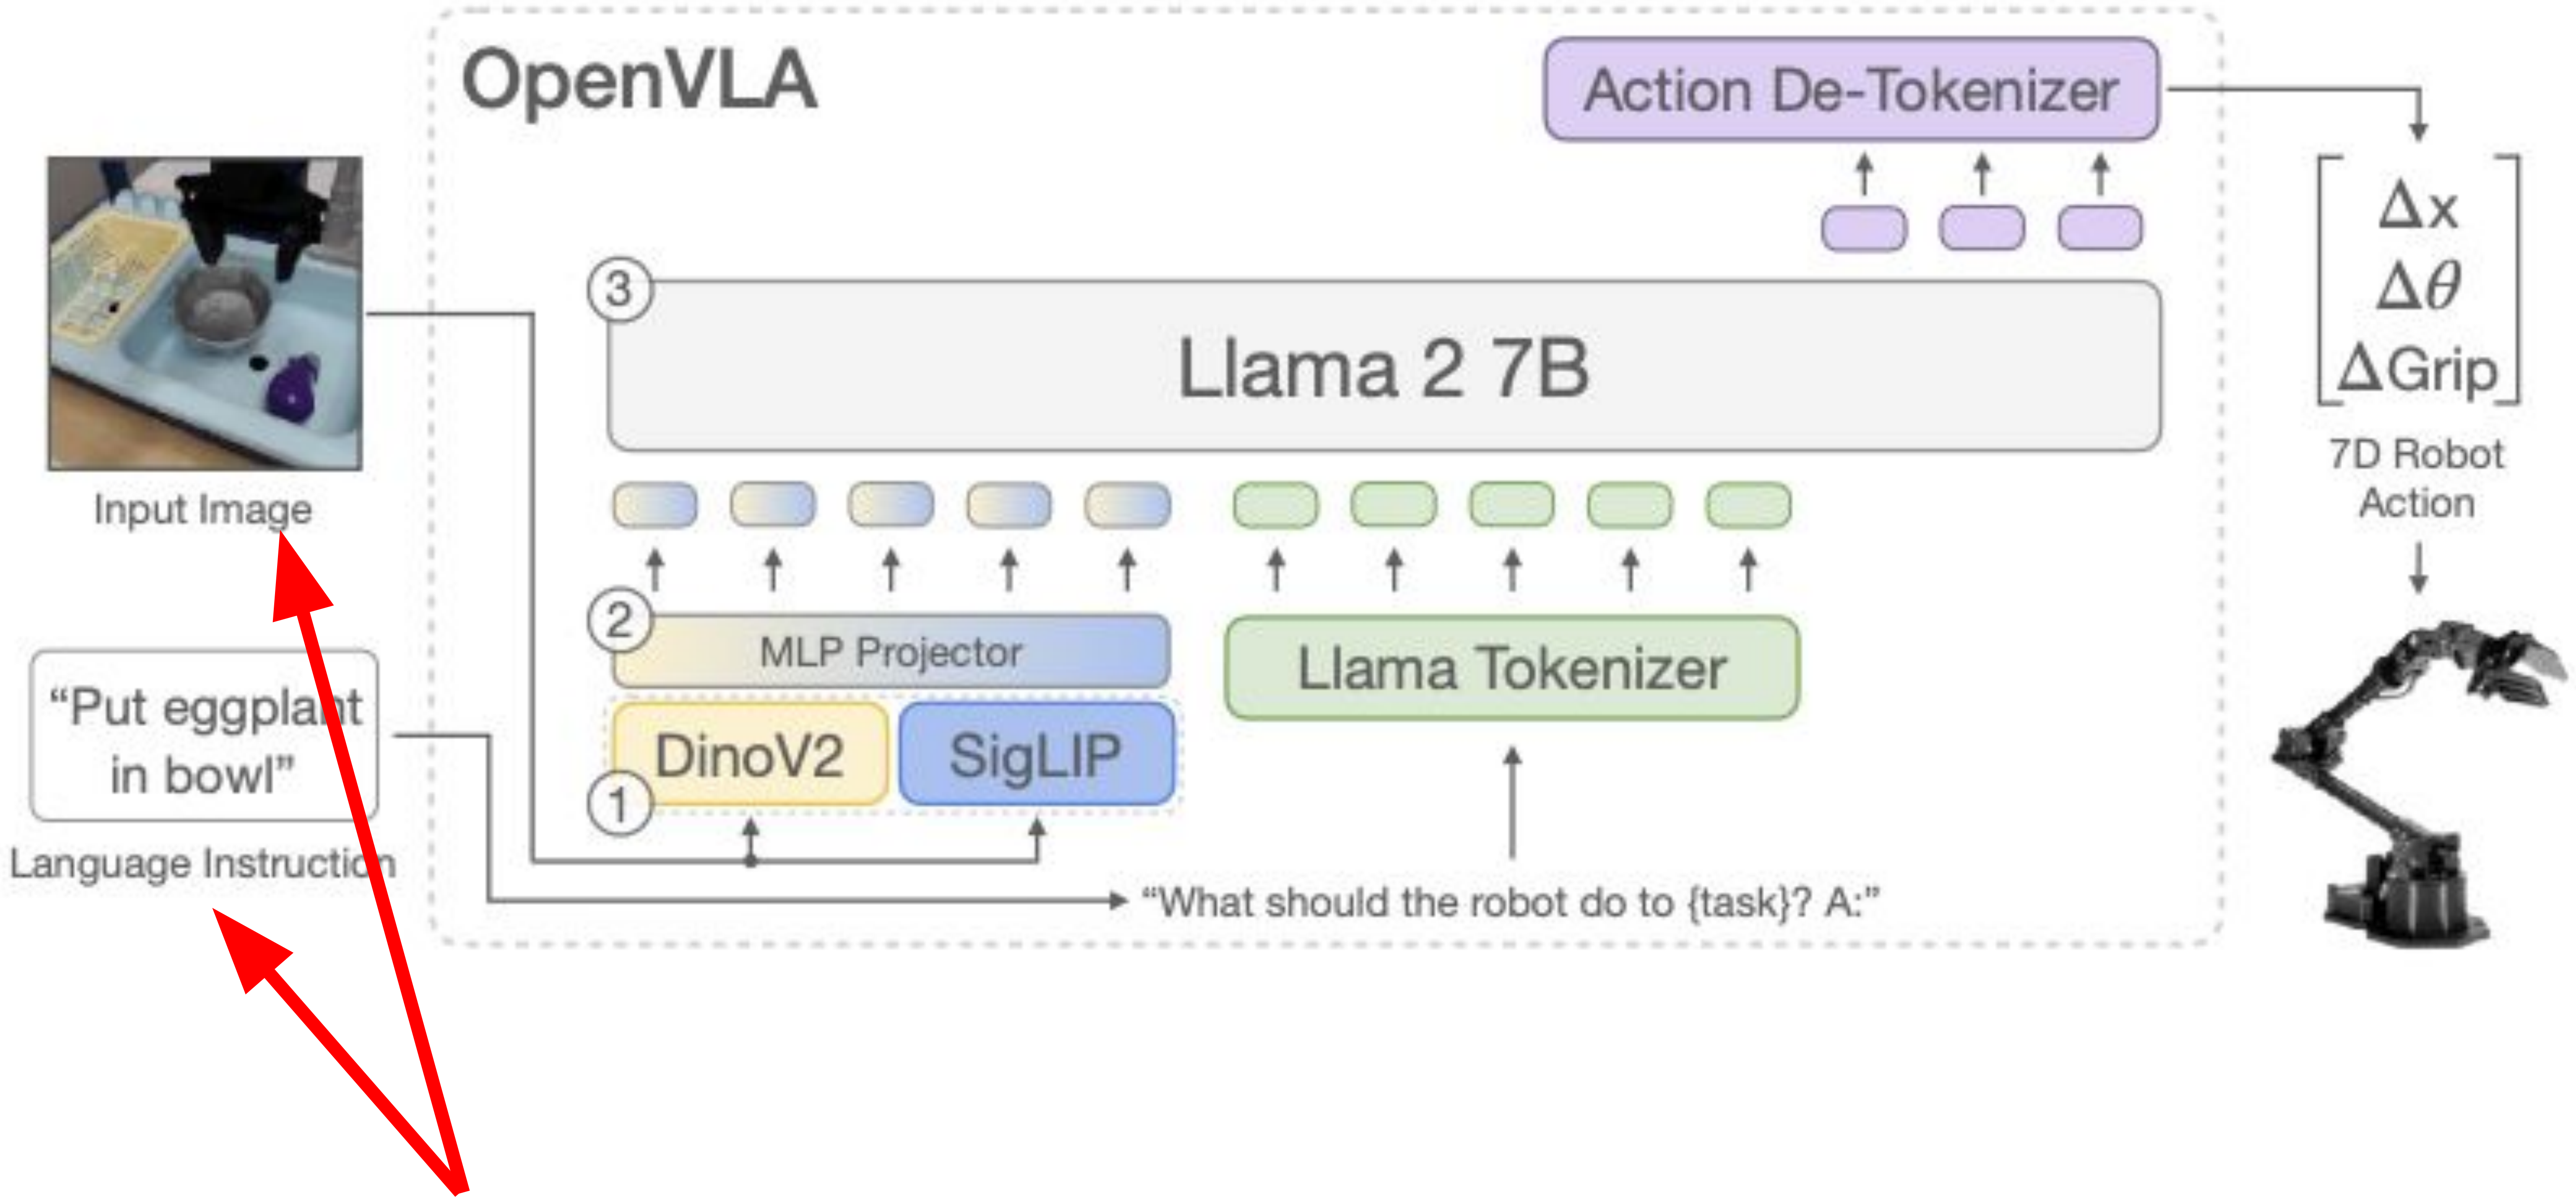

- Takes in visual observation + textual input

Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. arXiv preprint arXiv:2406.09246, 2024.

VLAs - Briefly

- Make use of LLMs
- Visual understanding from SigLIP and DinoV2

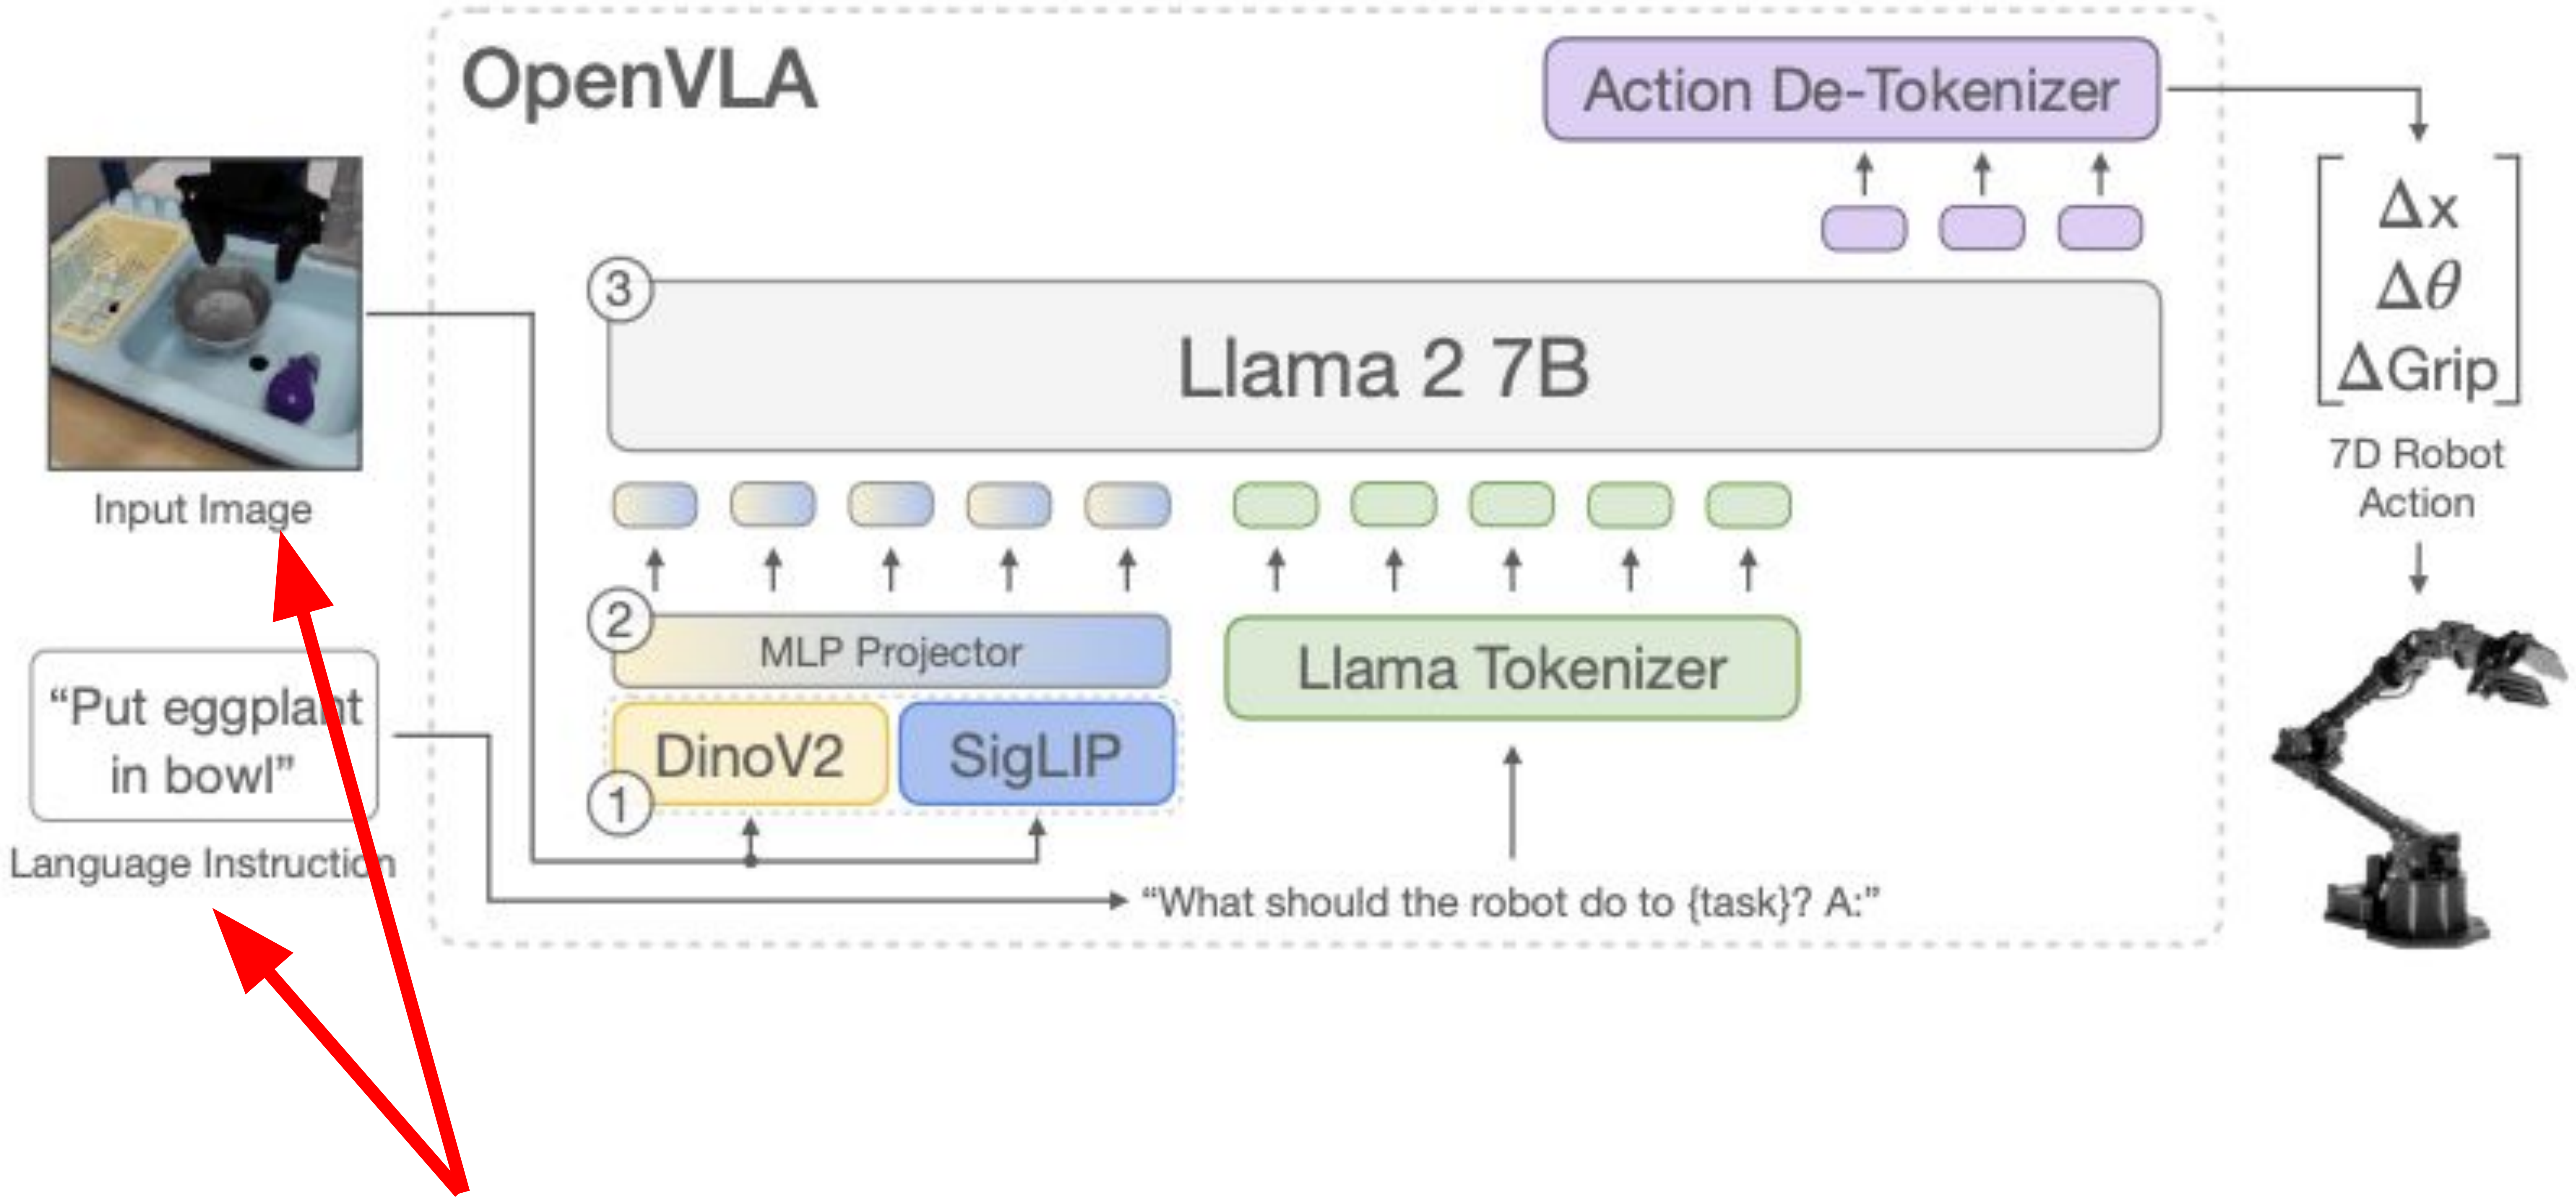

- The output translates to robot actions

- Takes in visual observation + textual input

Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. arXiv preprint arXiv:2406.09246, 2024.

π_0

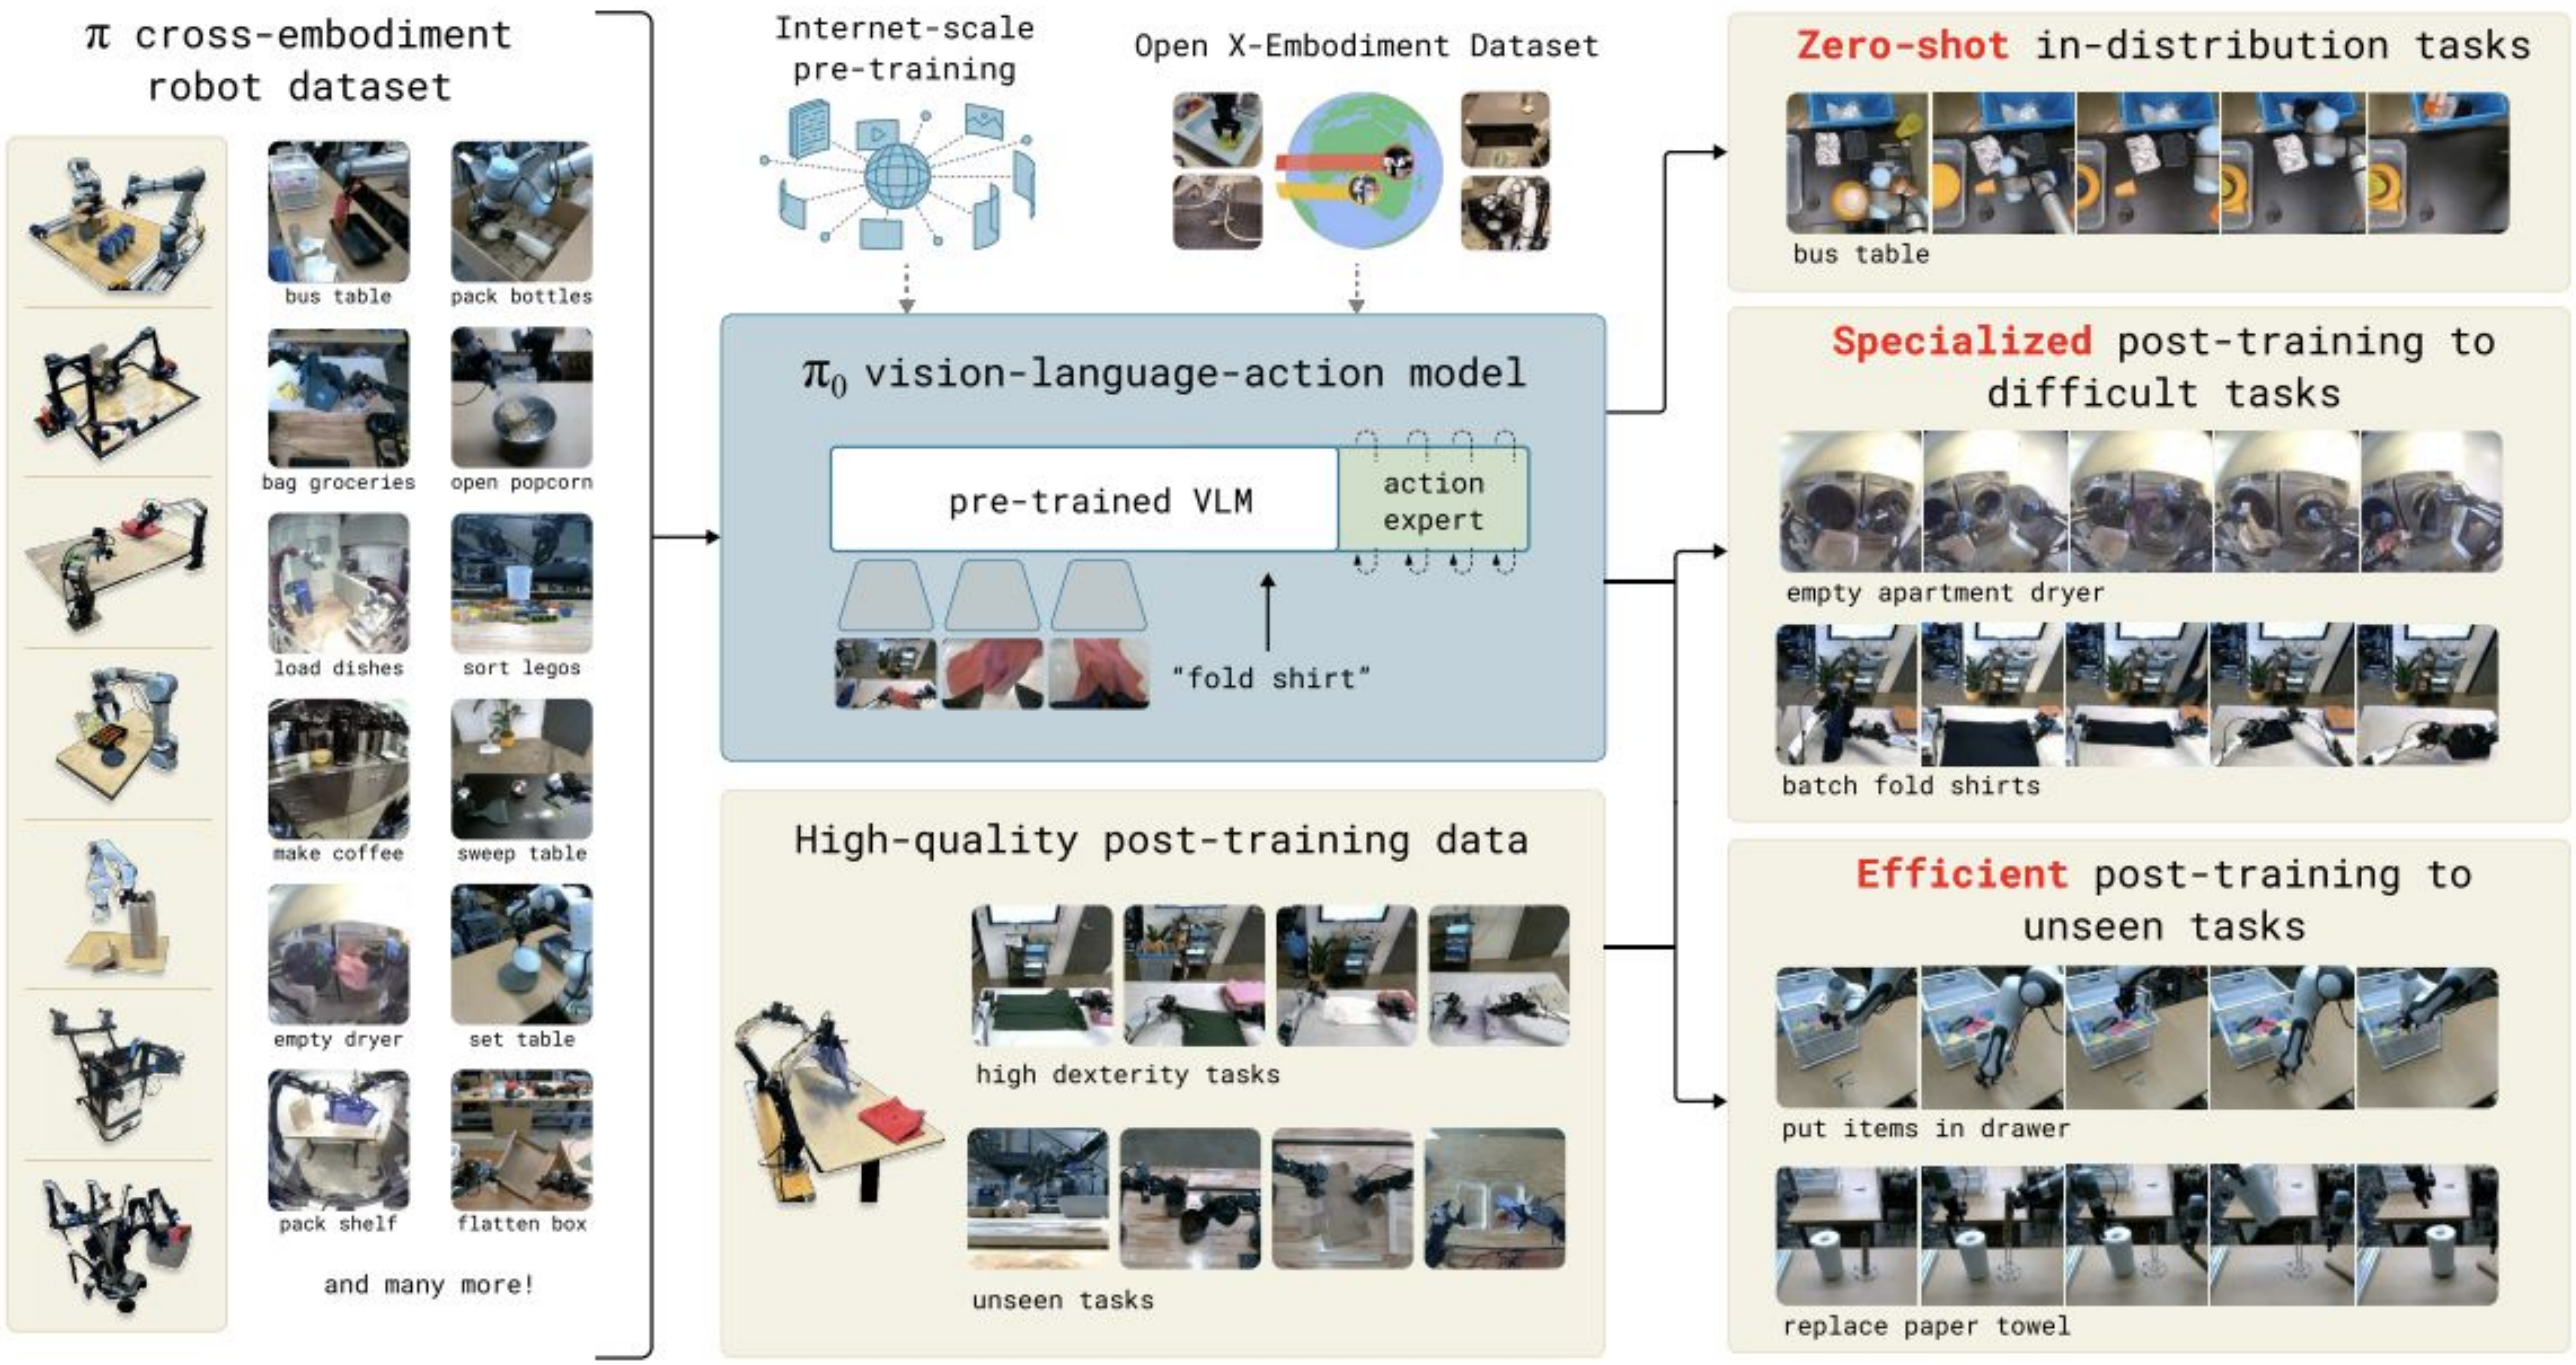

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. Physical Intelligence (2024). Available at https://www.physicalintelligence.company/download/pi0.pdf.

π_0

- A general framework for training generalist policies

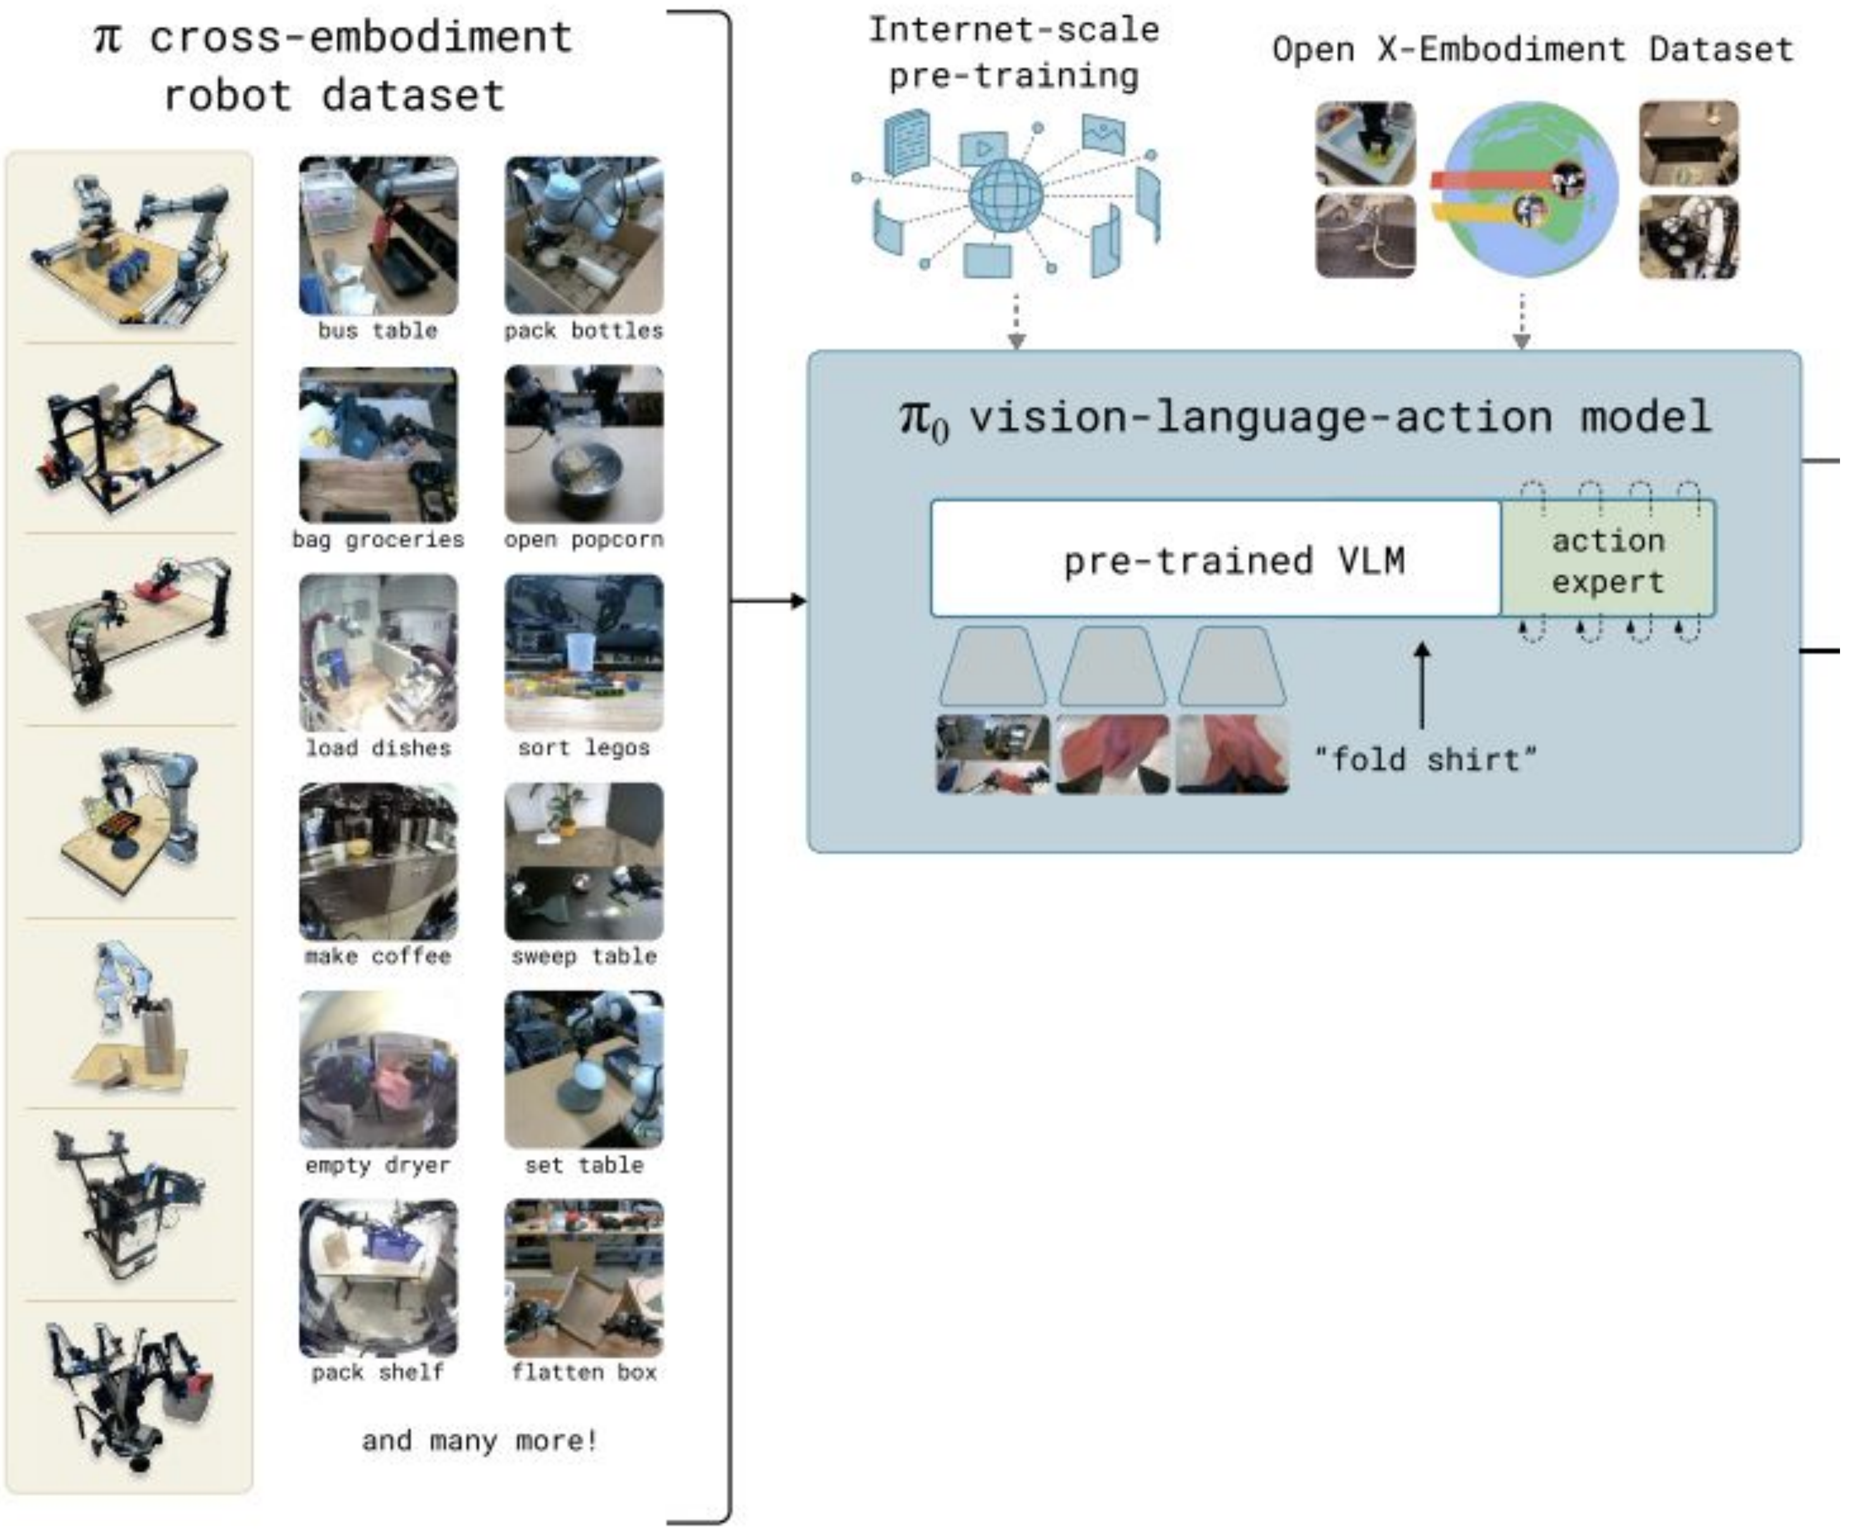

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. Physical Intelligence (2024). Available at https://www.physicalintelligence.company/download/pi0.pdf.

π_0

- A general framework for training generalist policies
- VLMs + diffusion variant(flow matching)

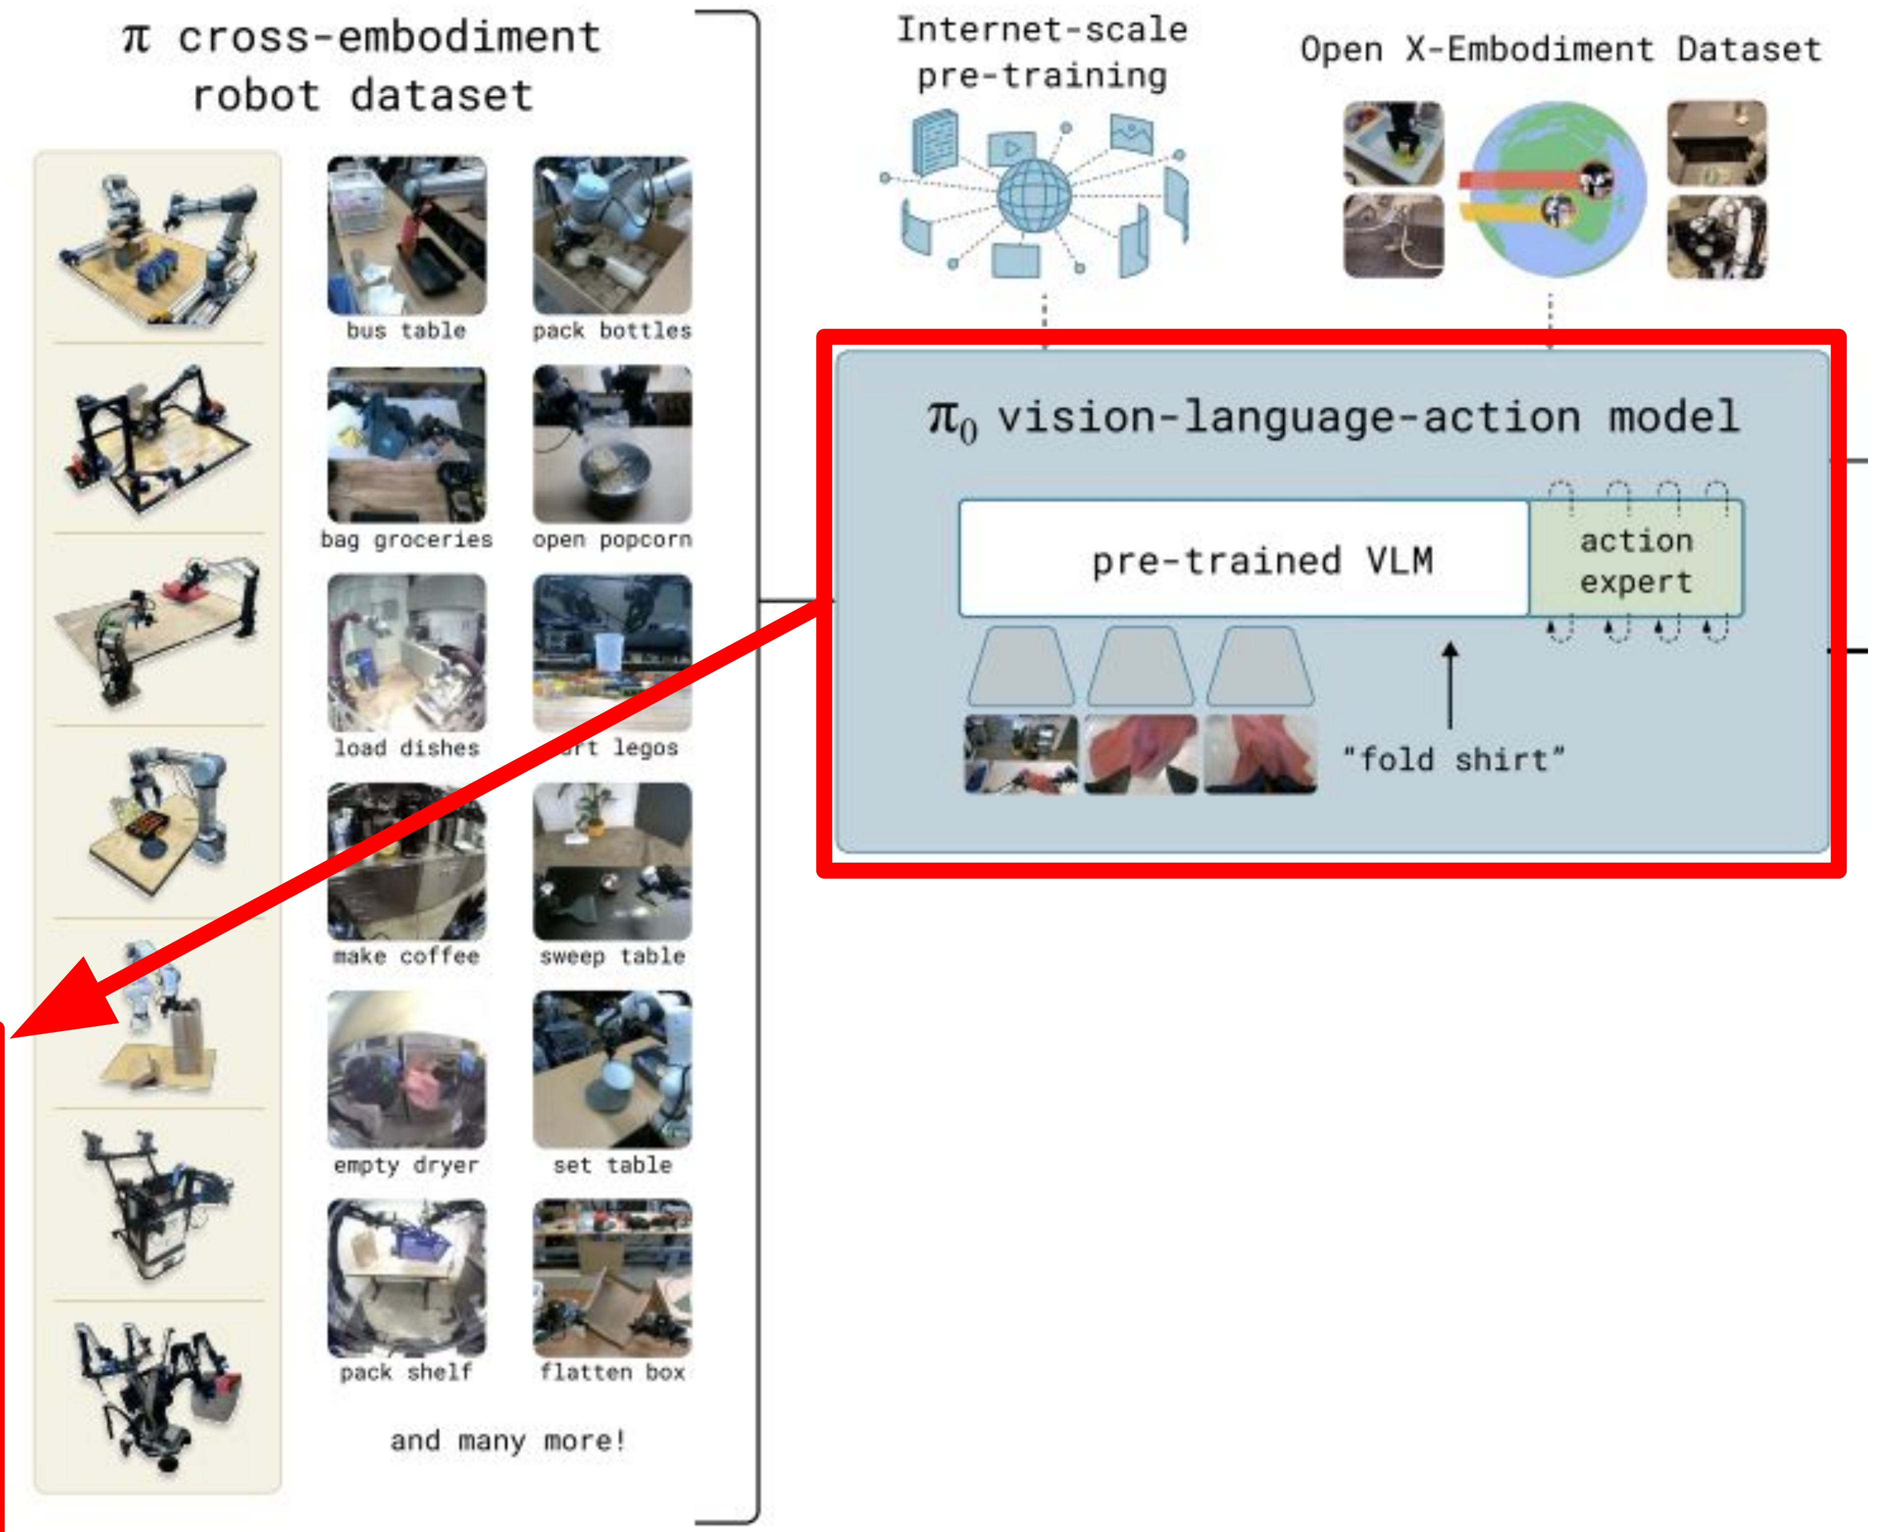

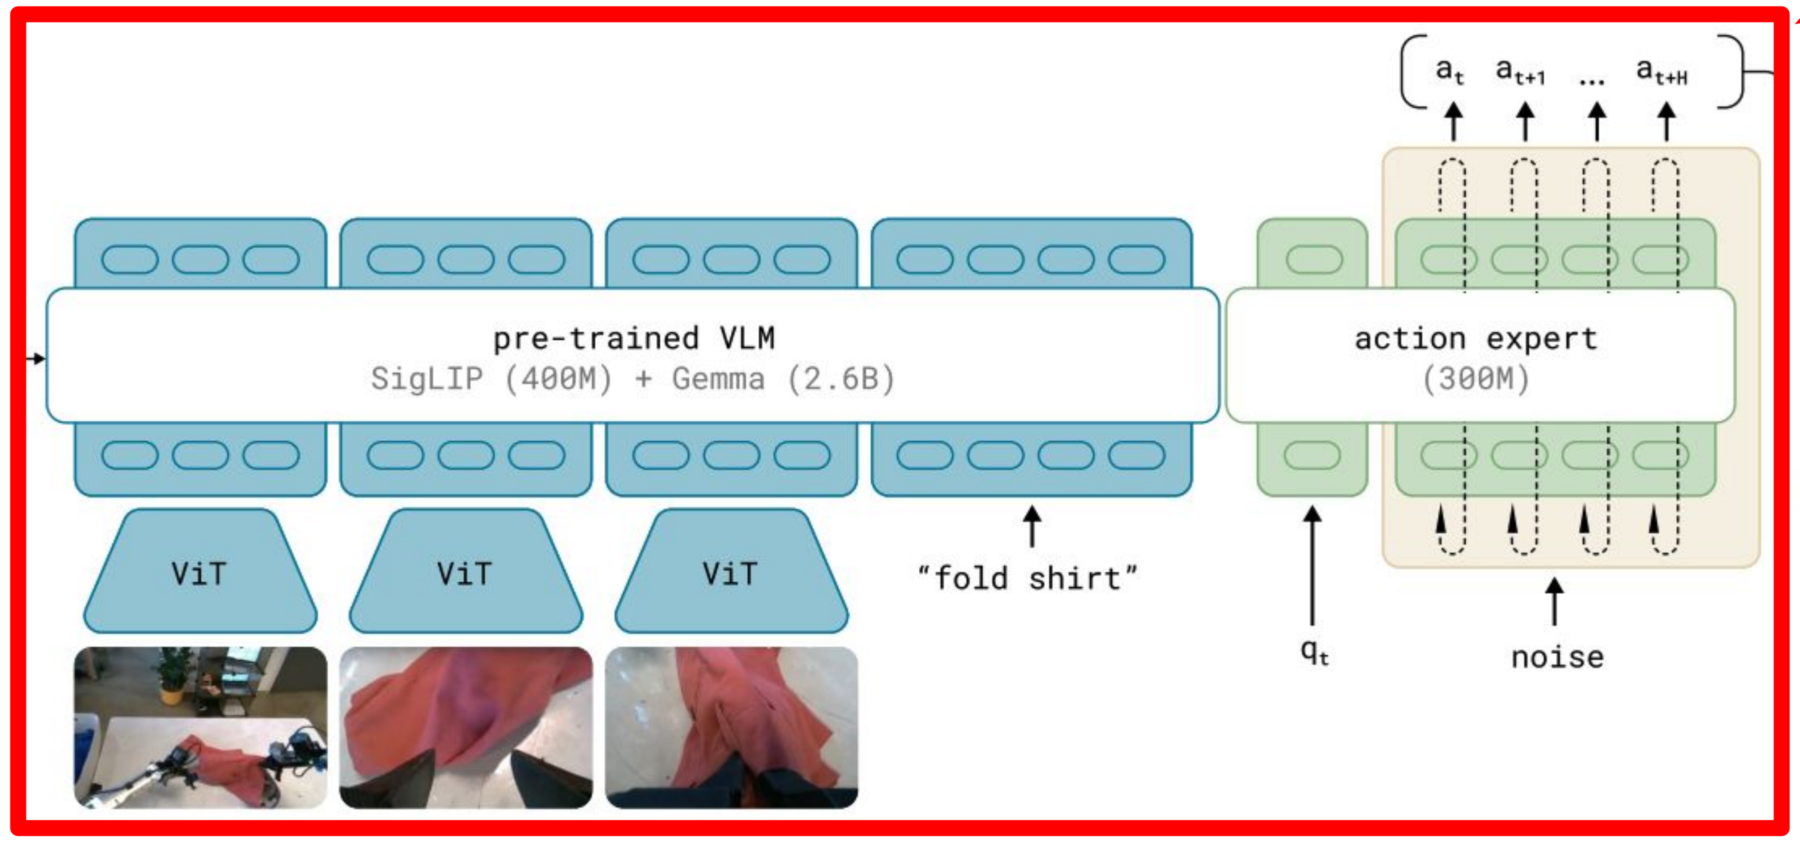

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. Physical Intelligence (2024). Available at https://www.physicalintelligence.company/download/pi0.pdf.

Results

- Outperforms previous methods(OpenVLA, Octo) by a large margin

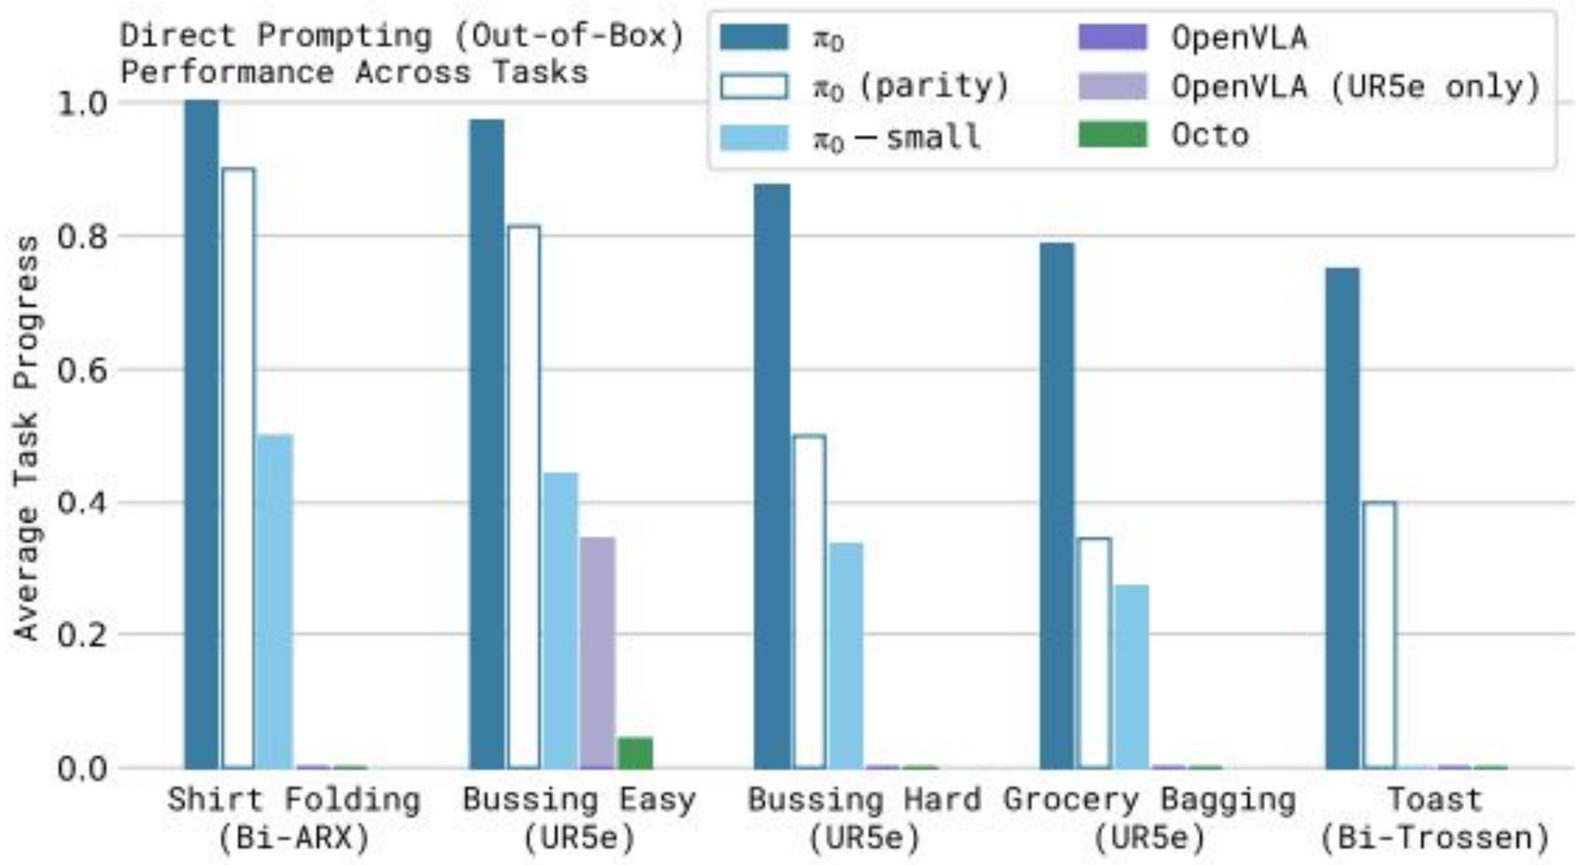

Performance out of the box

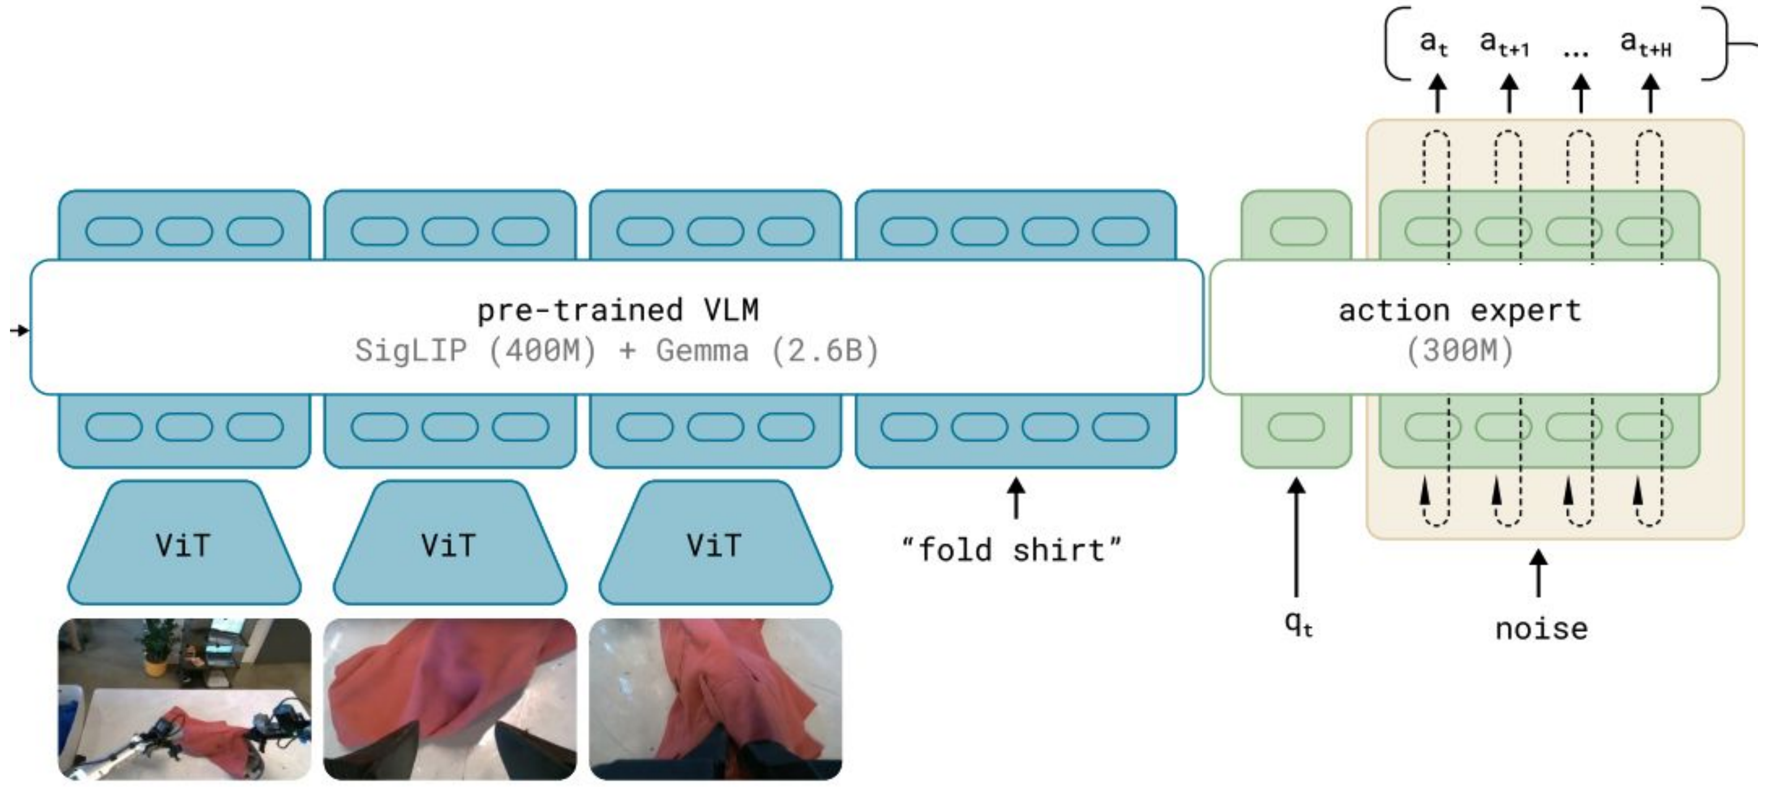

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. Physical Intelligence (2024). Available at https://www.physicalintelligence.company/download/pi0.pdf.

π_0

- A general framework for training generalist policies
- VLMs + diffusion variant(flow matching)

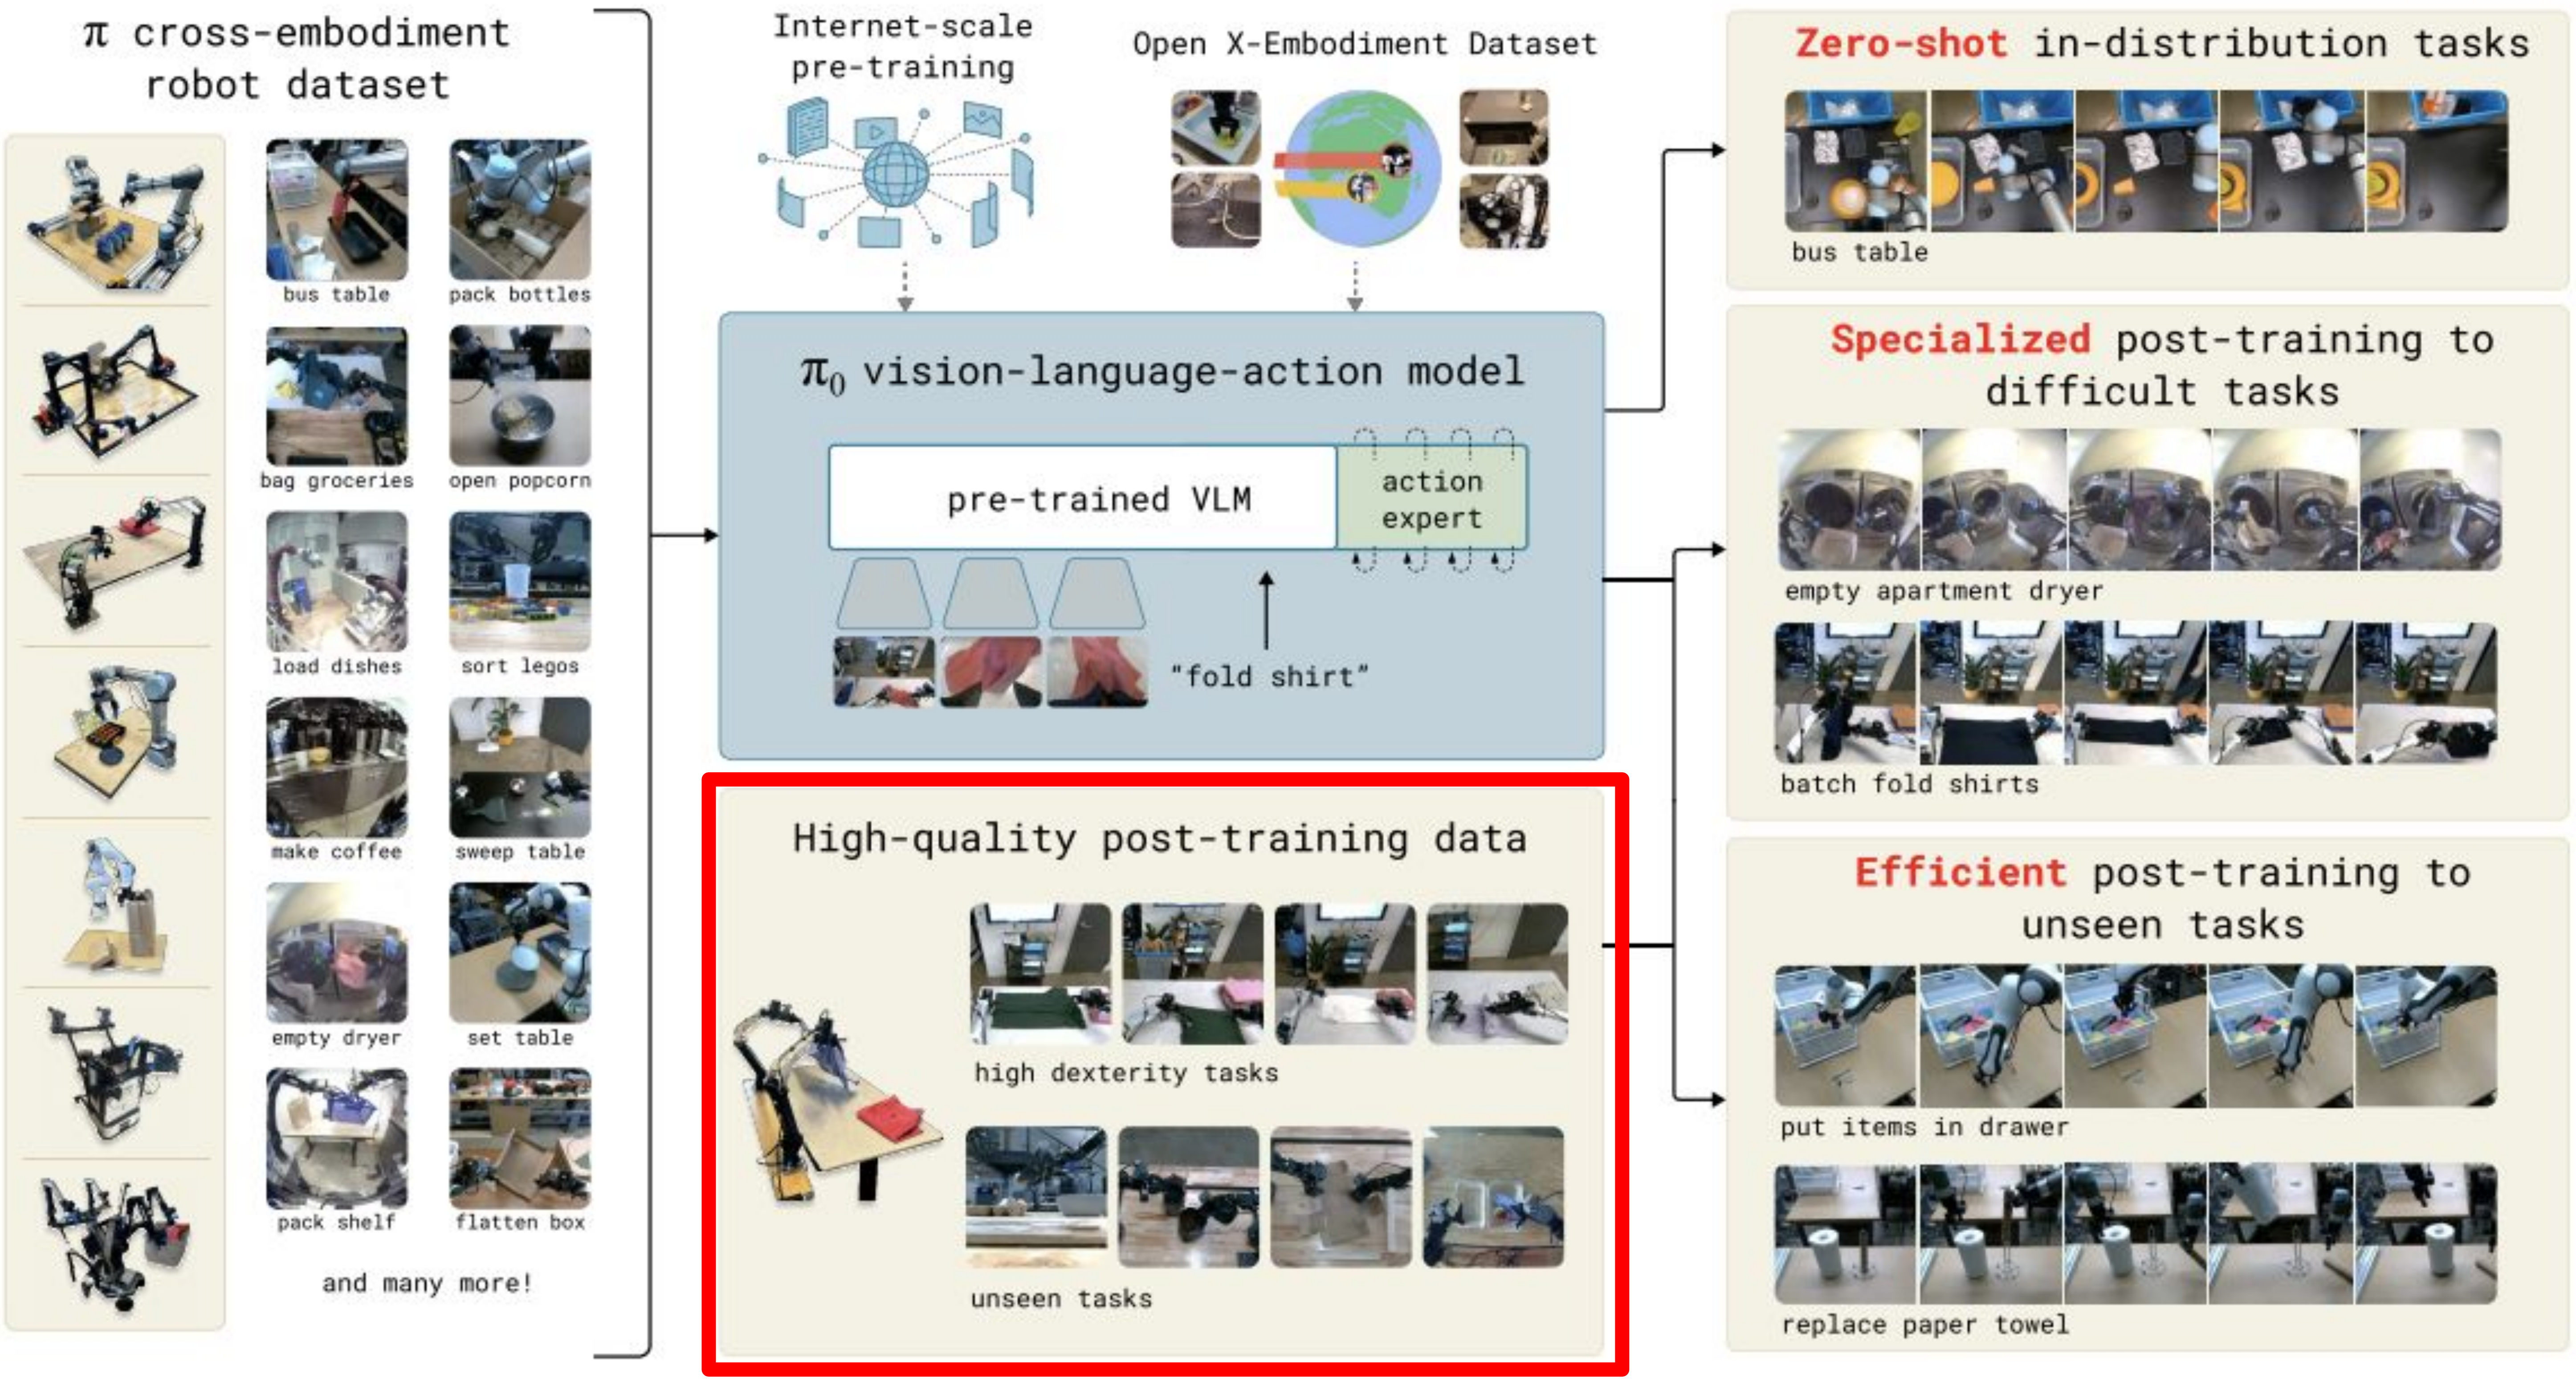

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. Physical Intelligence (2024). Available at https://www.physicalintelligence.company/download/pi0.pdf.

π_0

- Some difficult tasks needed 100s of hours of data
- Their dataset(below) contained 970M timesteps

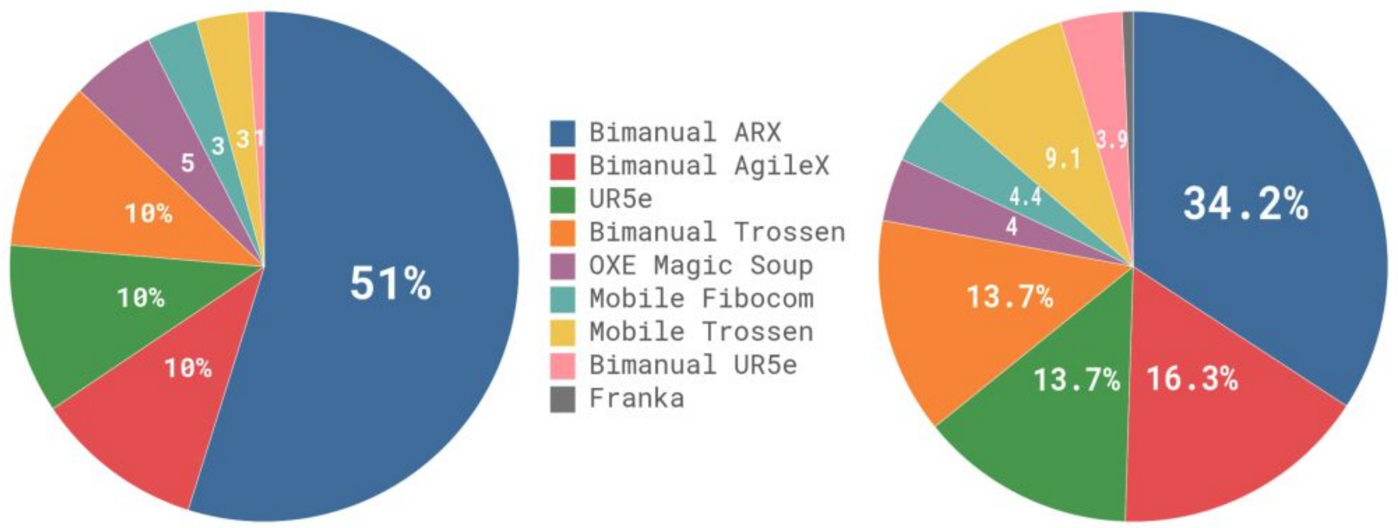

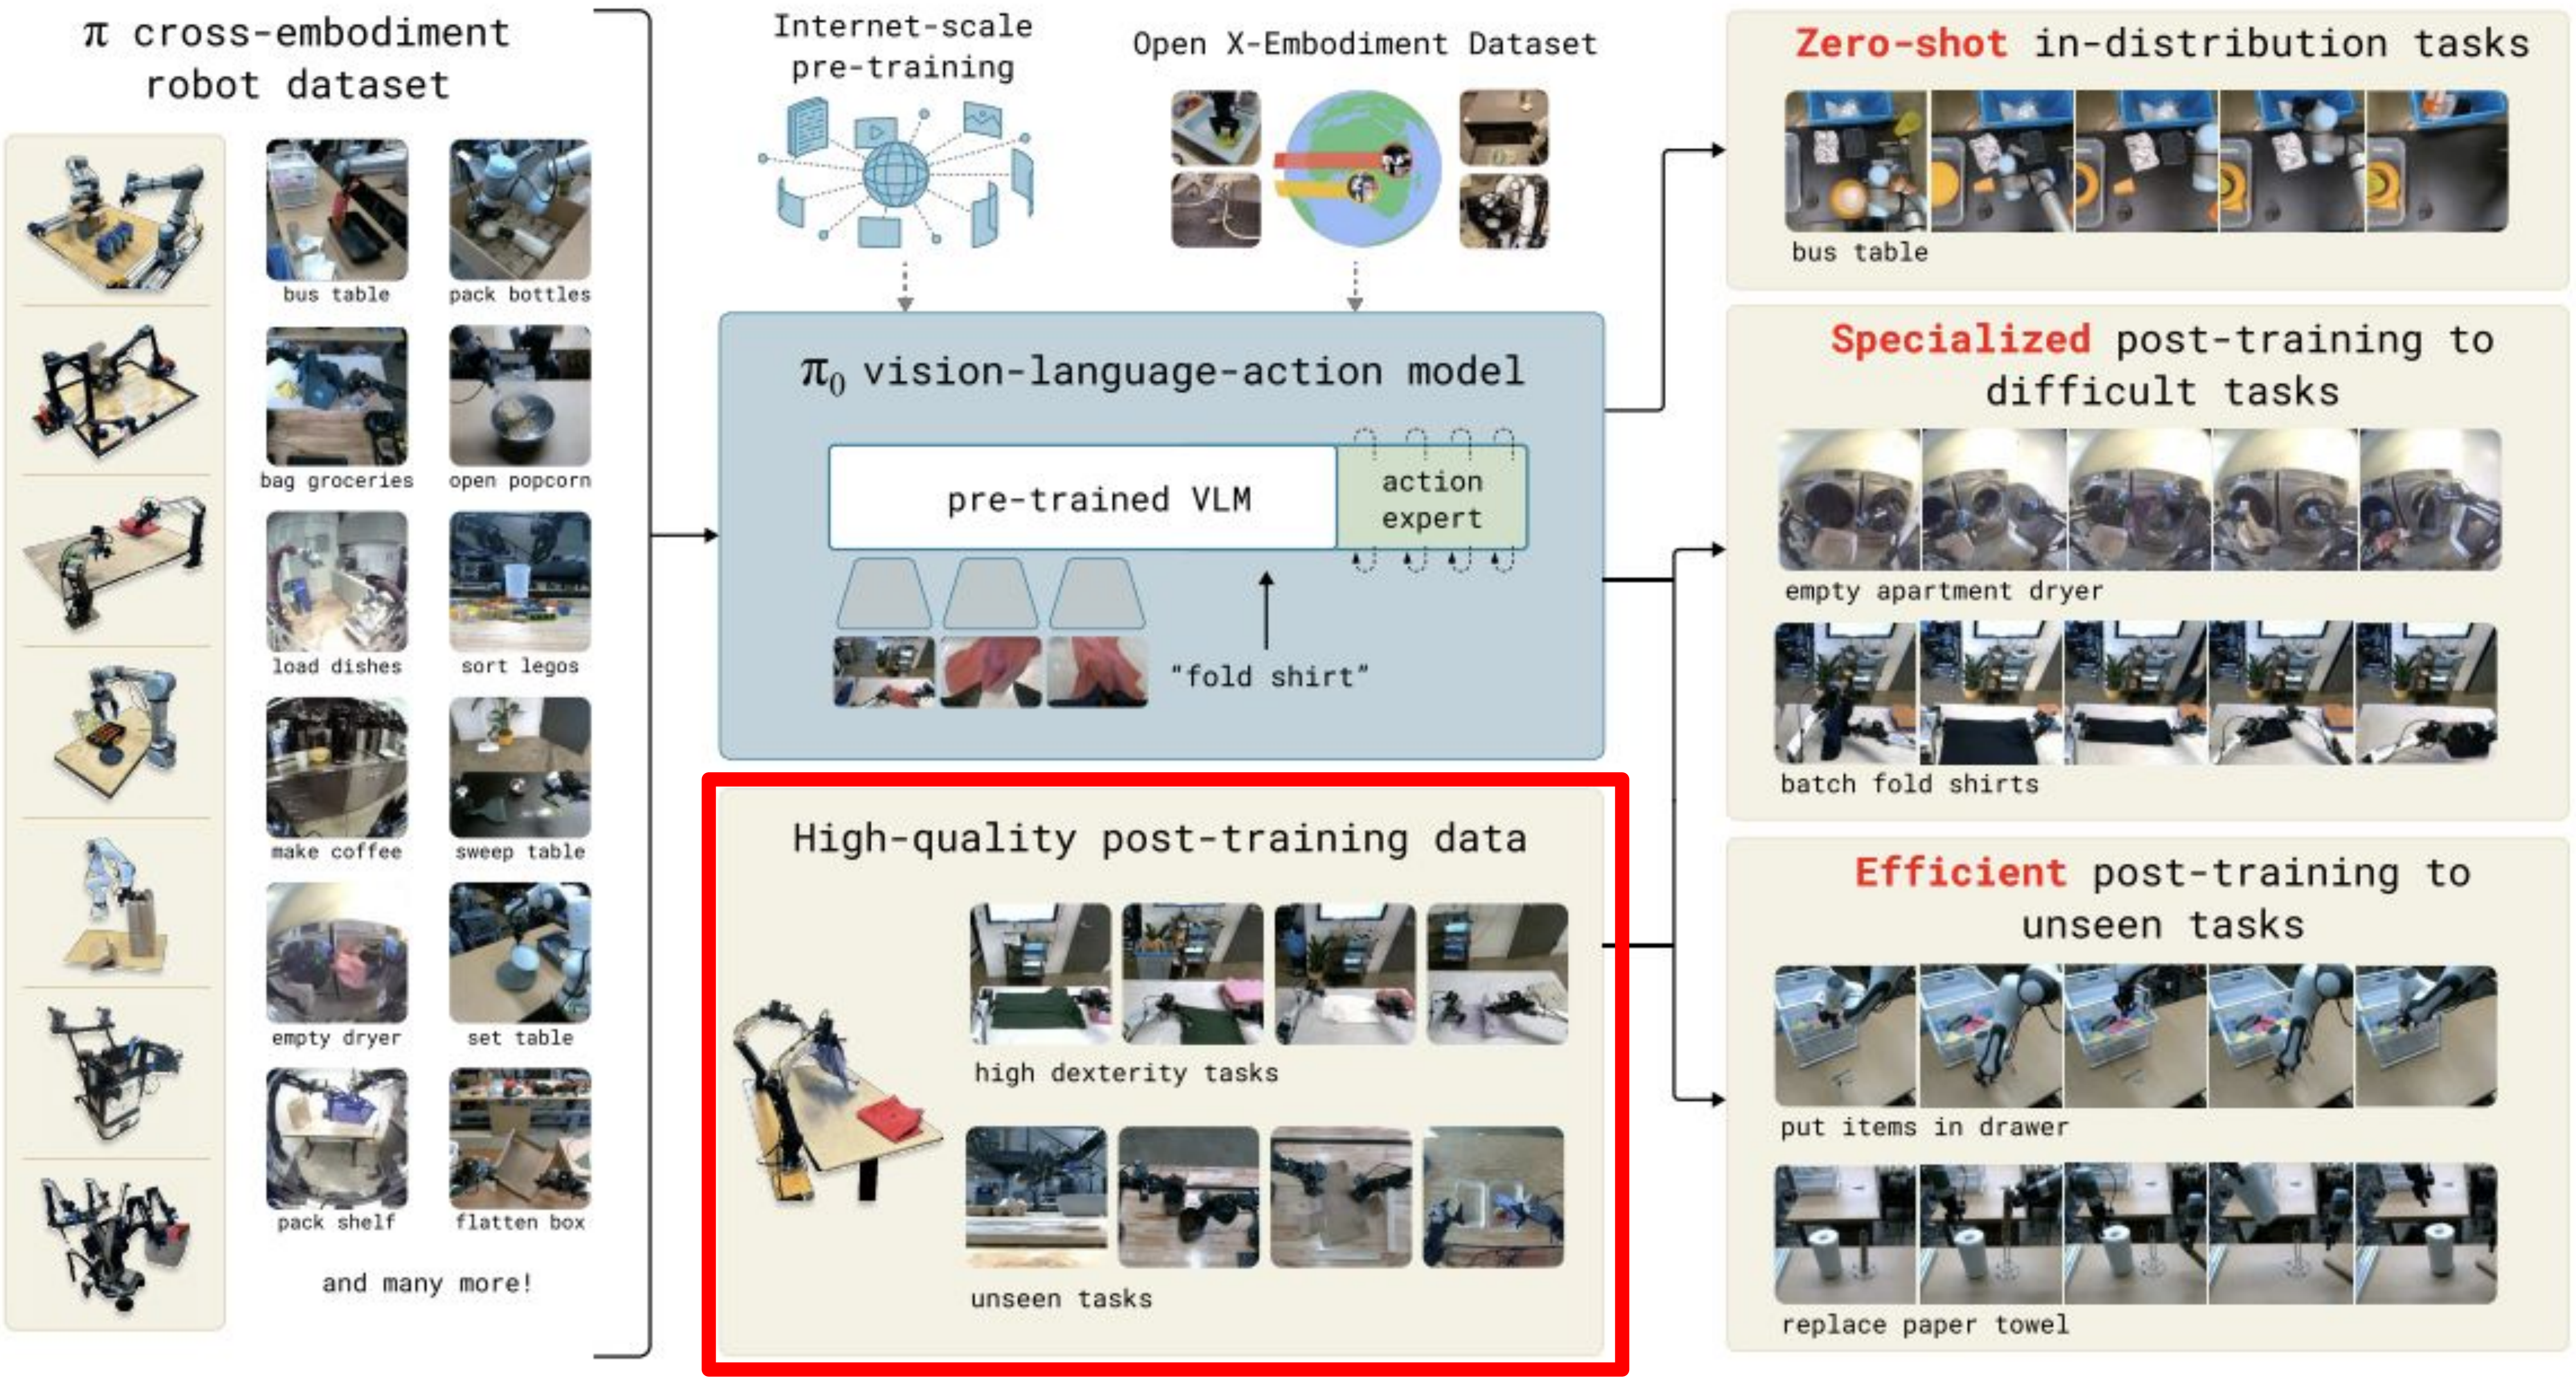

- Post training fine tunes for difficult tasks.

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. Physical Intelligence (2024). Available at https://www.physicalintelligence.company/download/pi0.pdf.

Results

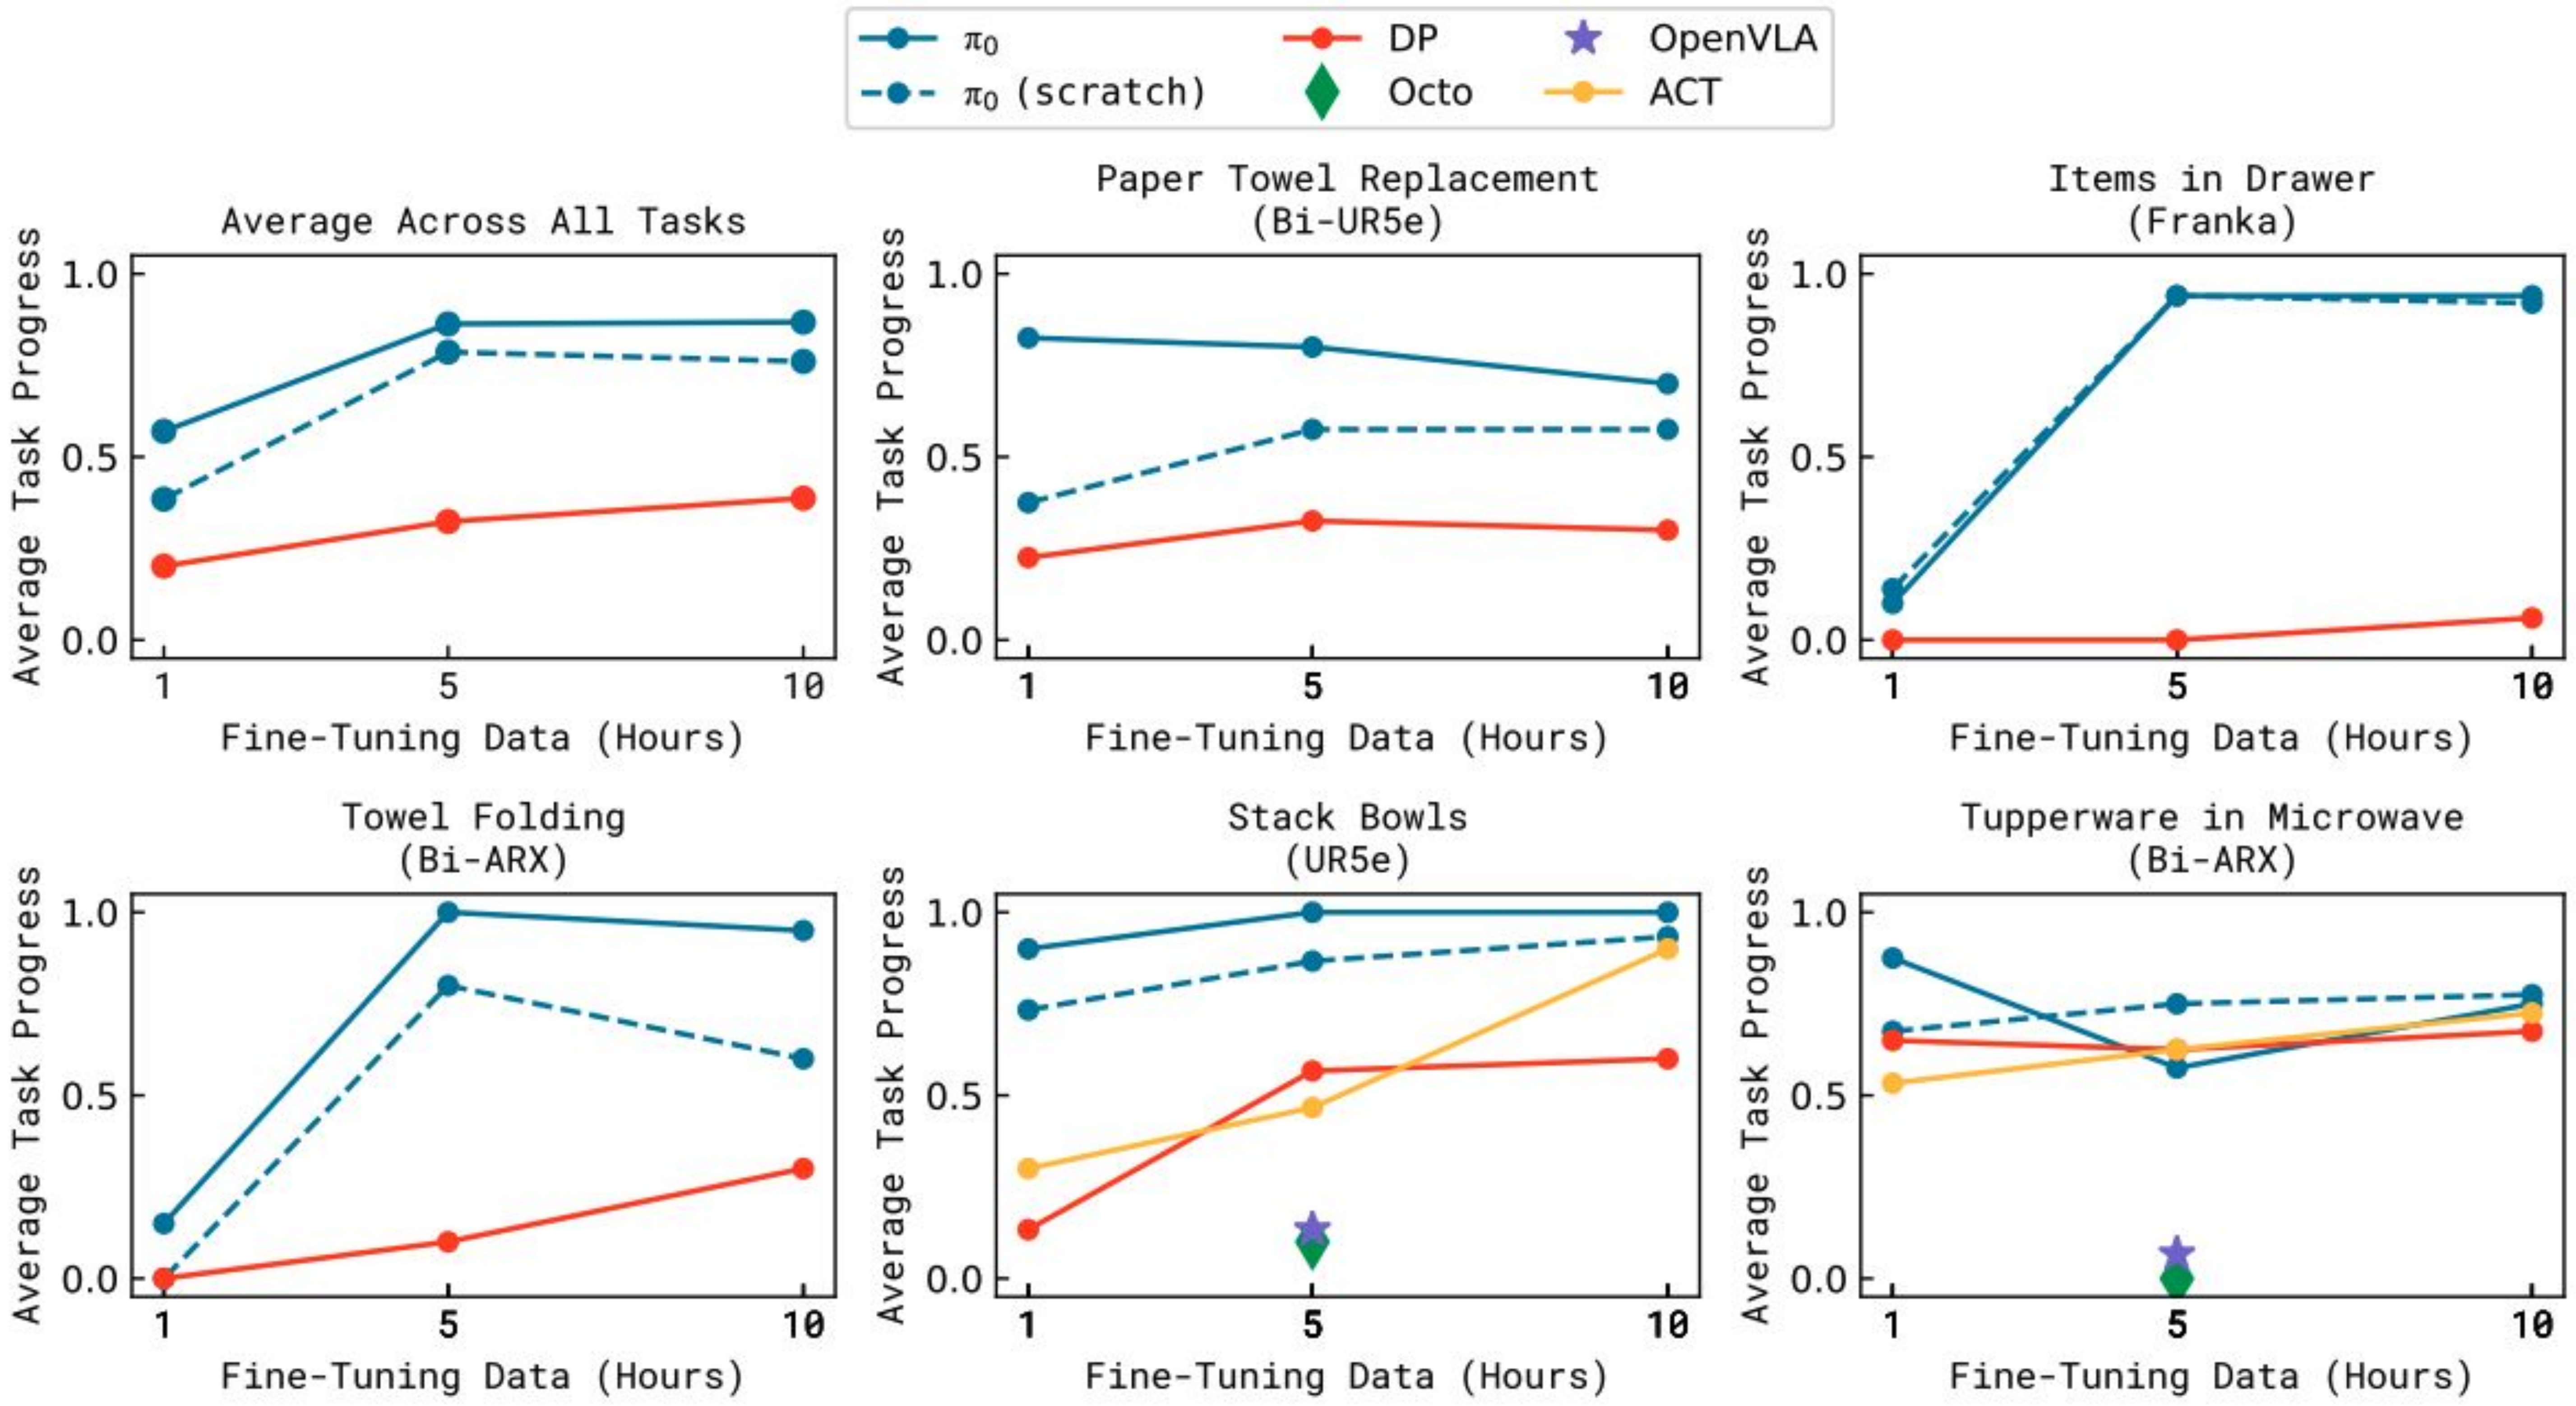

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. Physical Intelligence (2024). Available at https://www.physicalintelligence.company/download/pi0.pdf.

Results

- More complex tasks that bring together smaller tasks in pre-training

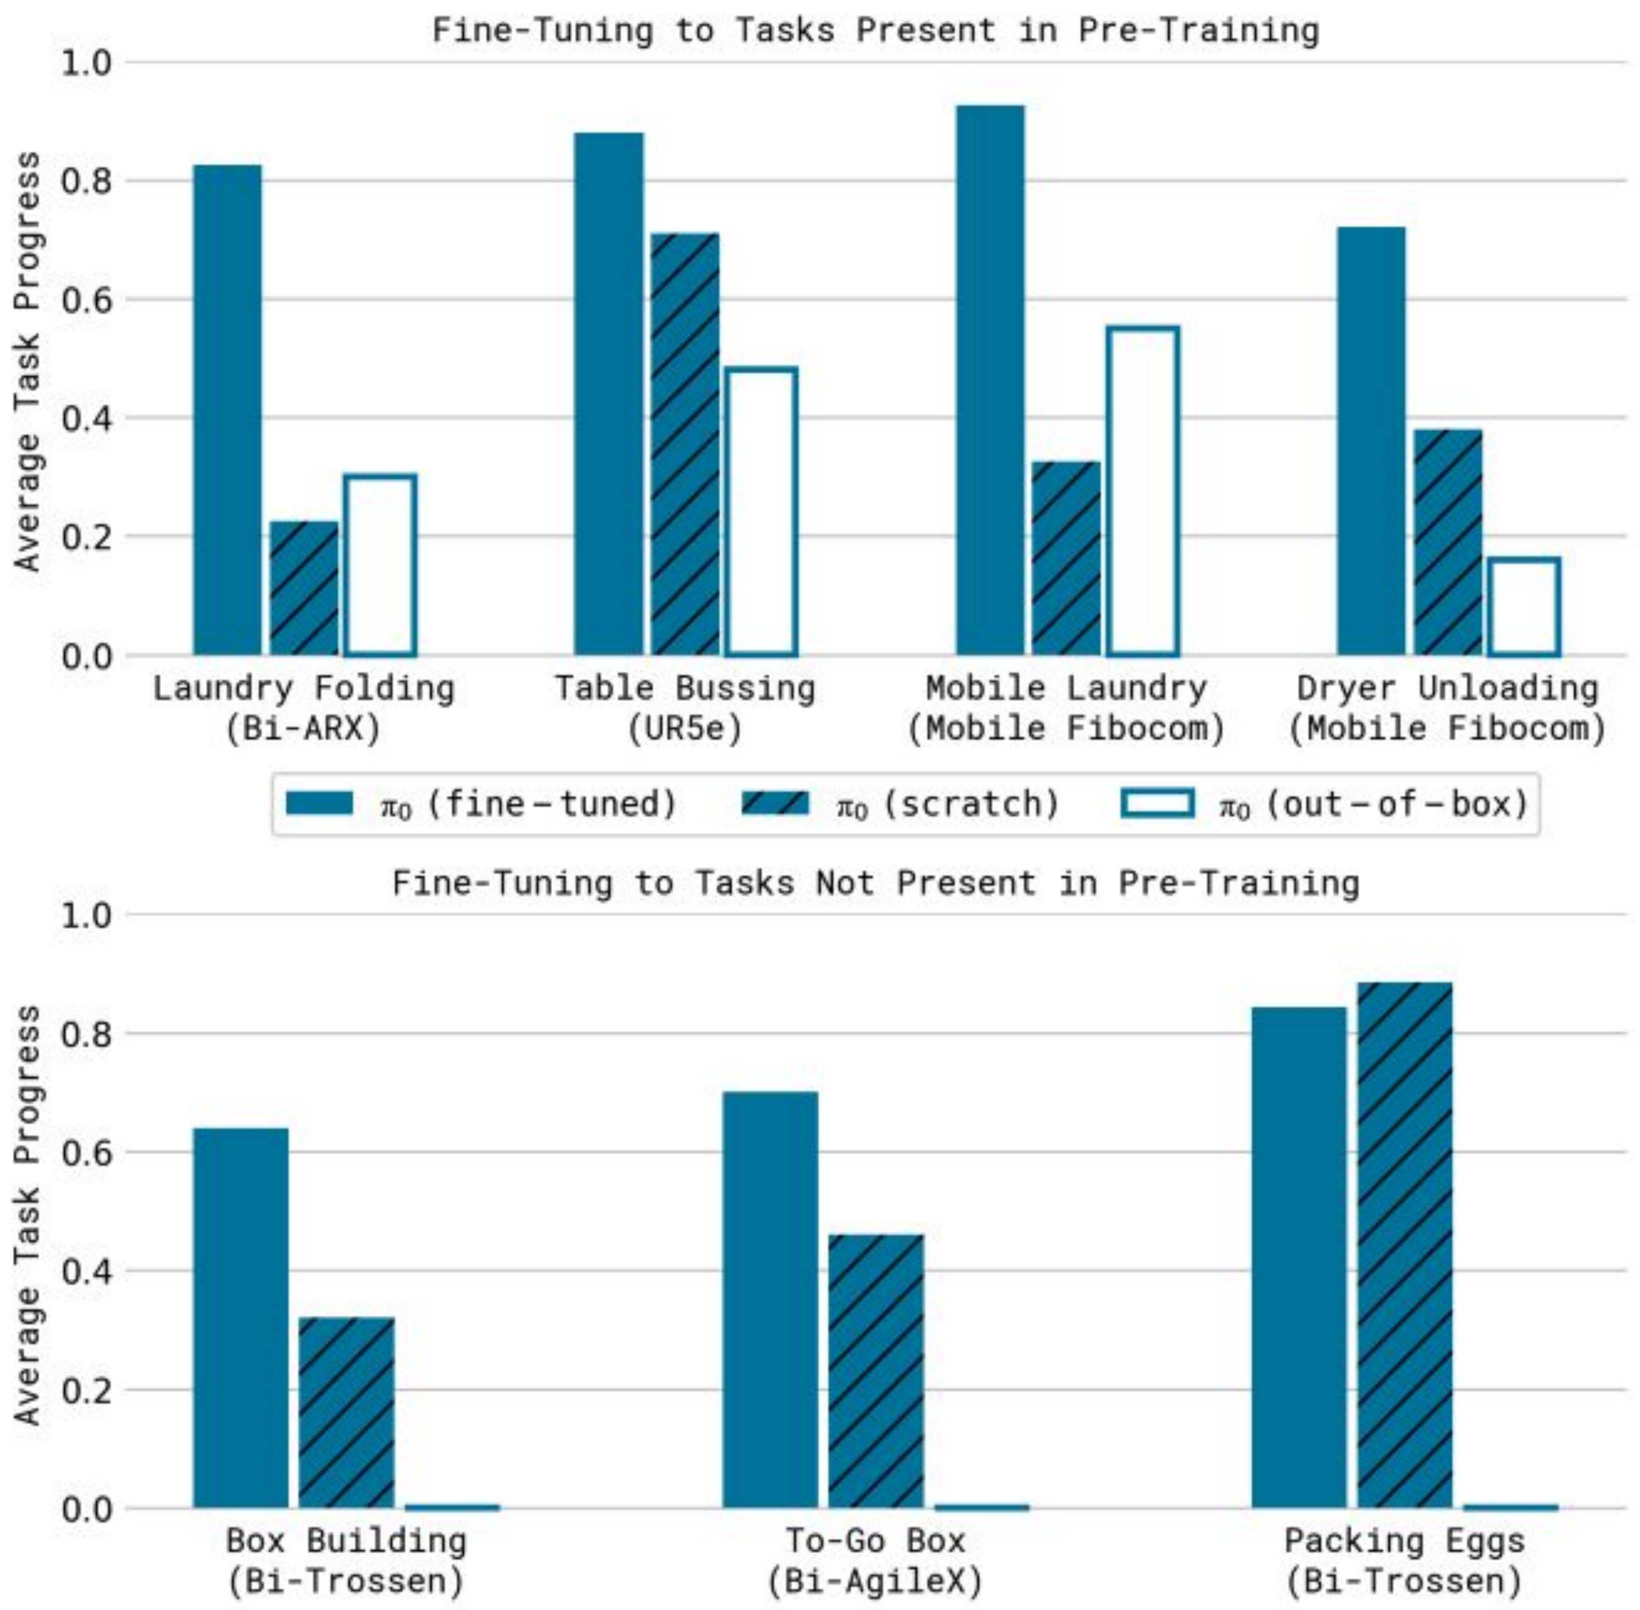

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. Physical Intelligence (2024). Available at https://www.physicalintelligence.company/download/pi0.pdf.

π_0 from Physical Intelligence

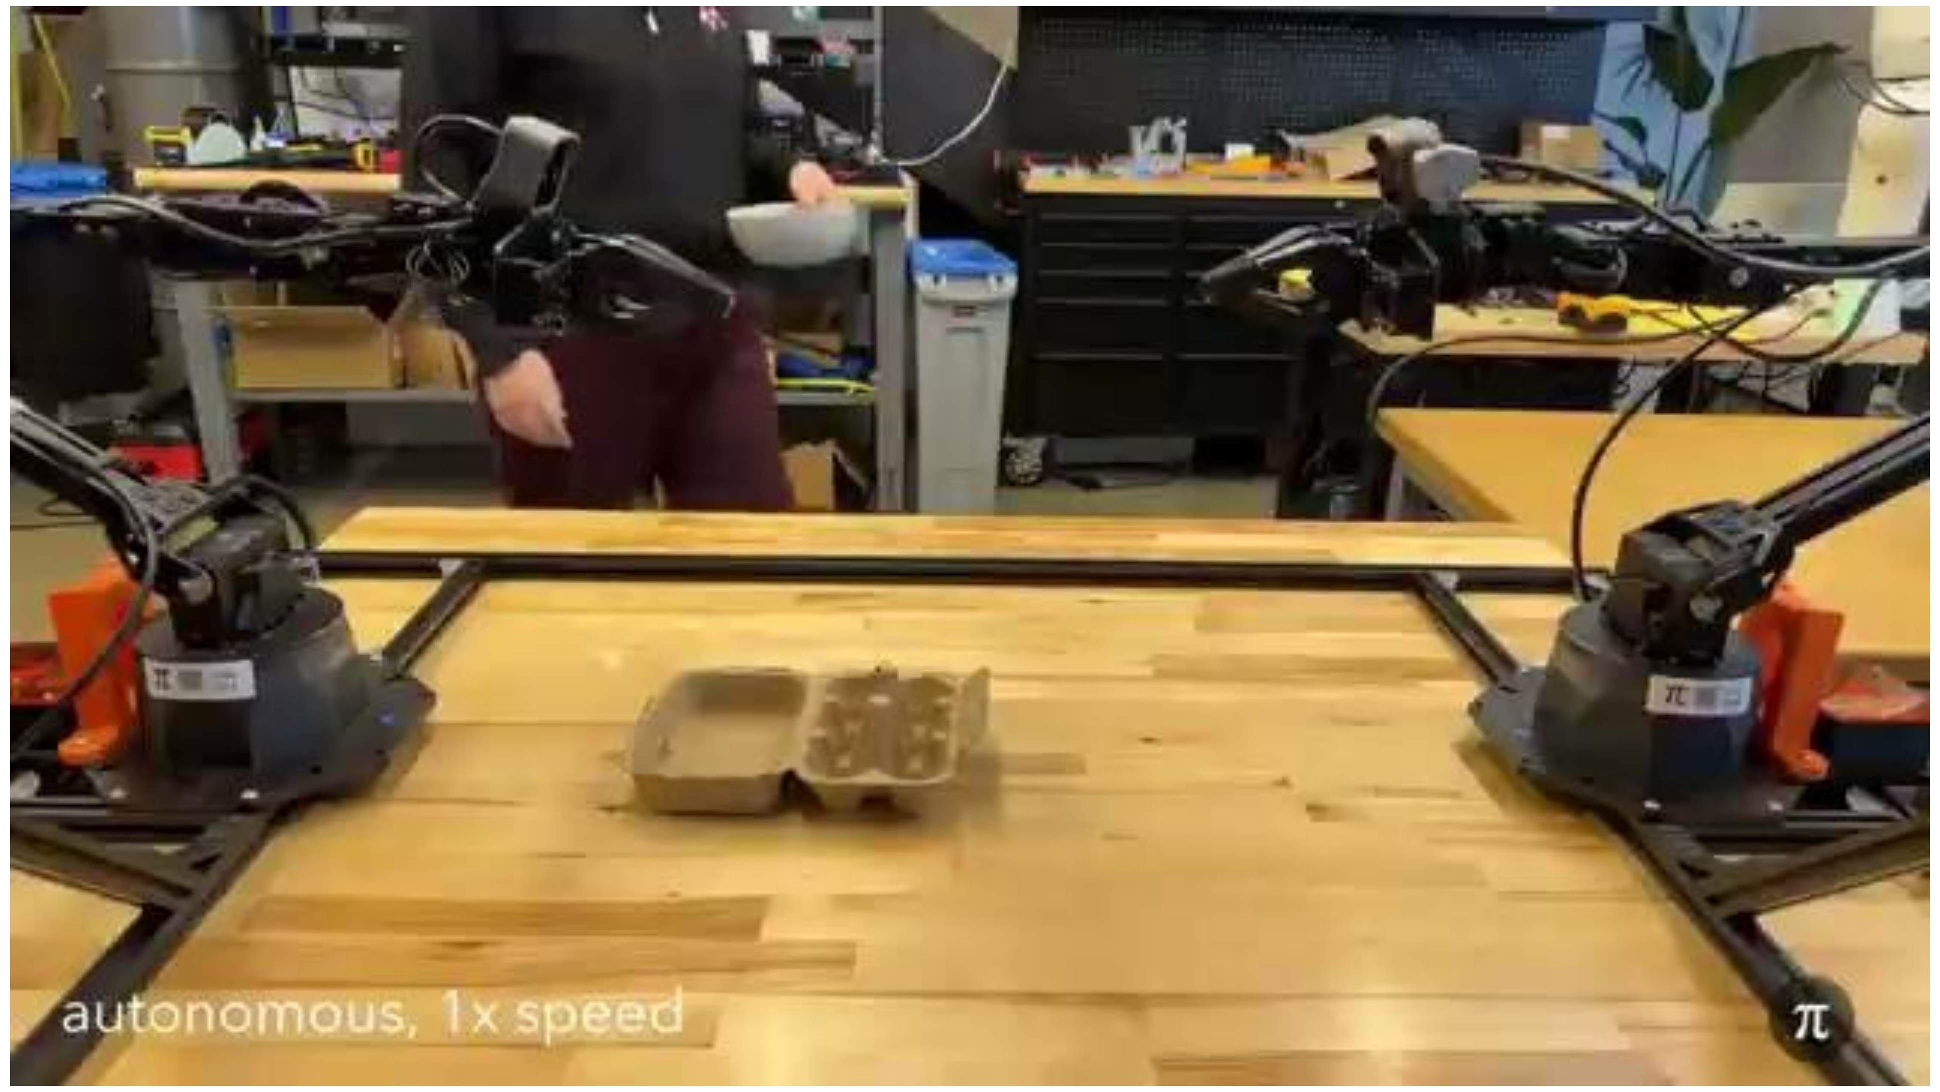

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. Physical Intelligence (2024). Available at https://www.physicalintelligence.company/download/pi0.pdf.

DR

Next Lecture:
Student Lecture 8
Foundational Models and Robot Manipulation

Reminder for Final Project Check-ins

Edstem post

12/04 Model Check-in: The presentations should include a discussion of the neural network models, loss functions, details on the training data and the test data, visualization of data and the amount of data, loss curves (train vs. test), and any other information you would like to share. Please upload your google-slides (not more than 4 slides per group) as "G#_model_training" in this folder. **Due 9am 12/04**

12/11 Evaluation Check-in: Using the trained model, how accurate is the task performance on the manipulation task? What scenarios are you experimenting with, etc.? How is your method compared to baseline(s)? What are your ongoing experiments? Please upload your google-slides (not more than 4 slides per group) as "G#_evaluation_baselines" in this folder. **Due 9am 12/11**

DeepRob

[Group 7] Lecture 7
Dual-arm Manipulation - Learning

by Ryan Roche, Matt Rajala, Adit Kadepurkar
University of Minnesota