Increasing the Robustness of Manipulations by Monitoring Skill Executions

Oliver Kroemer
Carnegie Mellon University
Robotics Institute

Motivation

- Want robots to perform a variety of manipulation skills

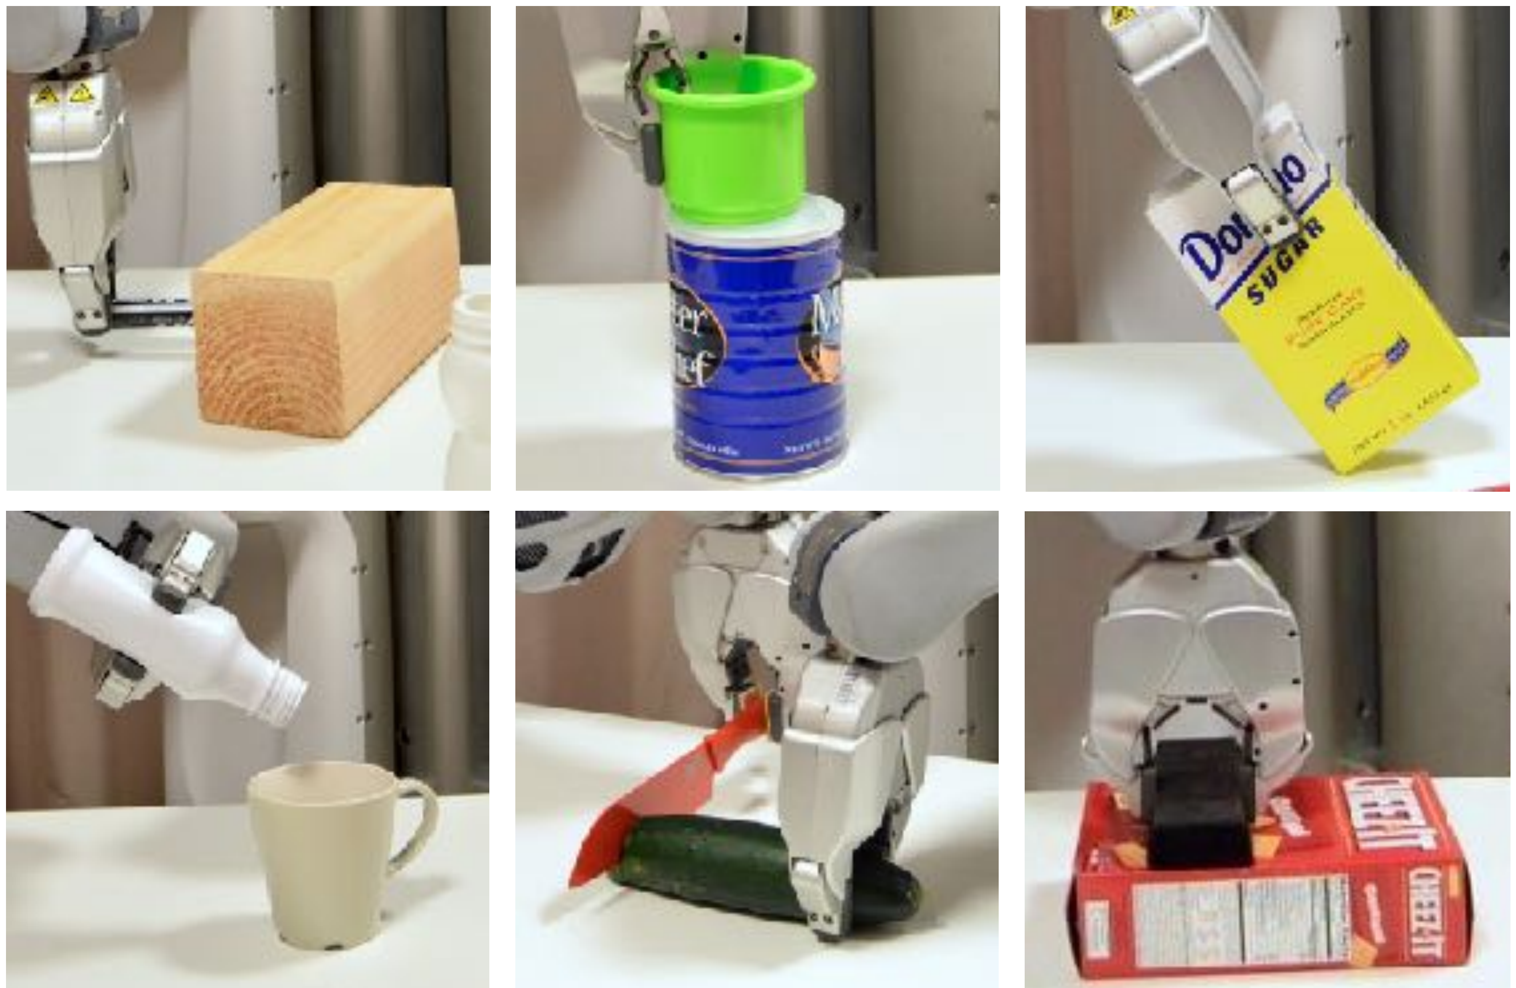

- Robots need to perform skills reliably despite variations
- Goal: Robots to learn robust and versatile manipulation skills

Manipulation Mode Structure

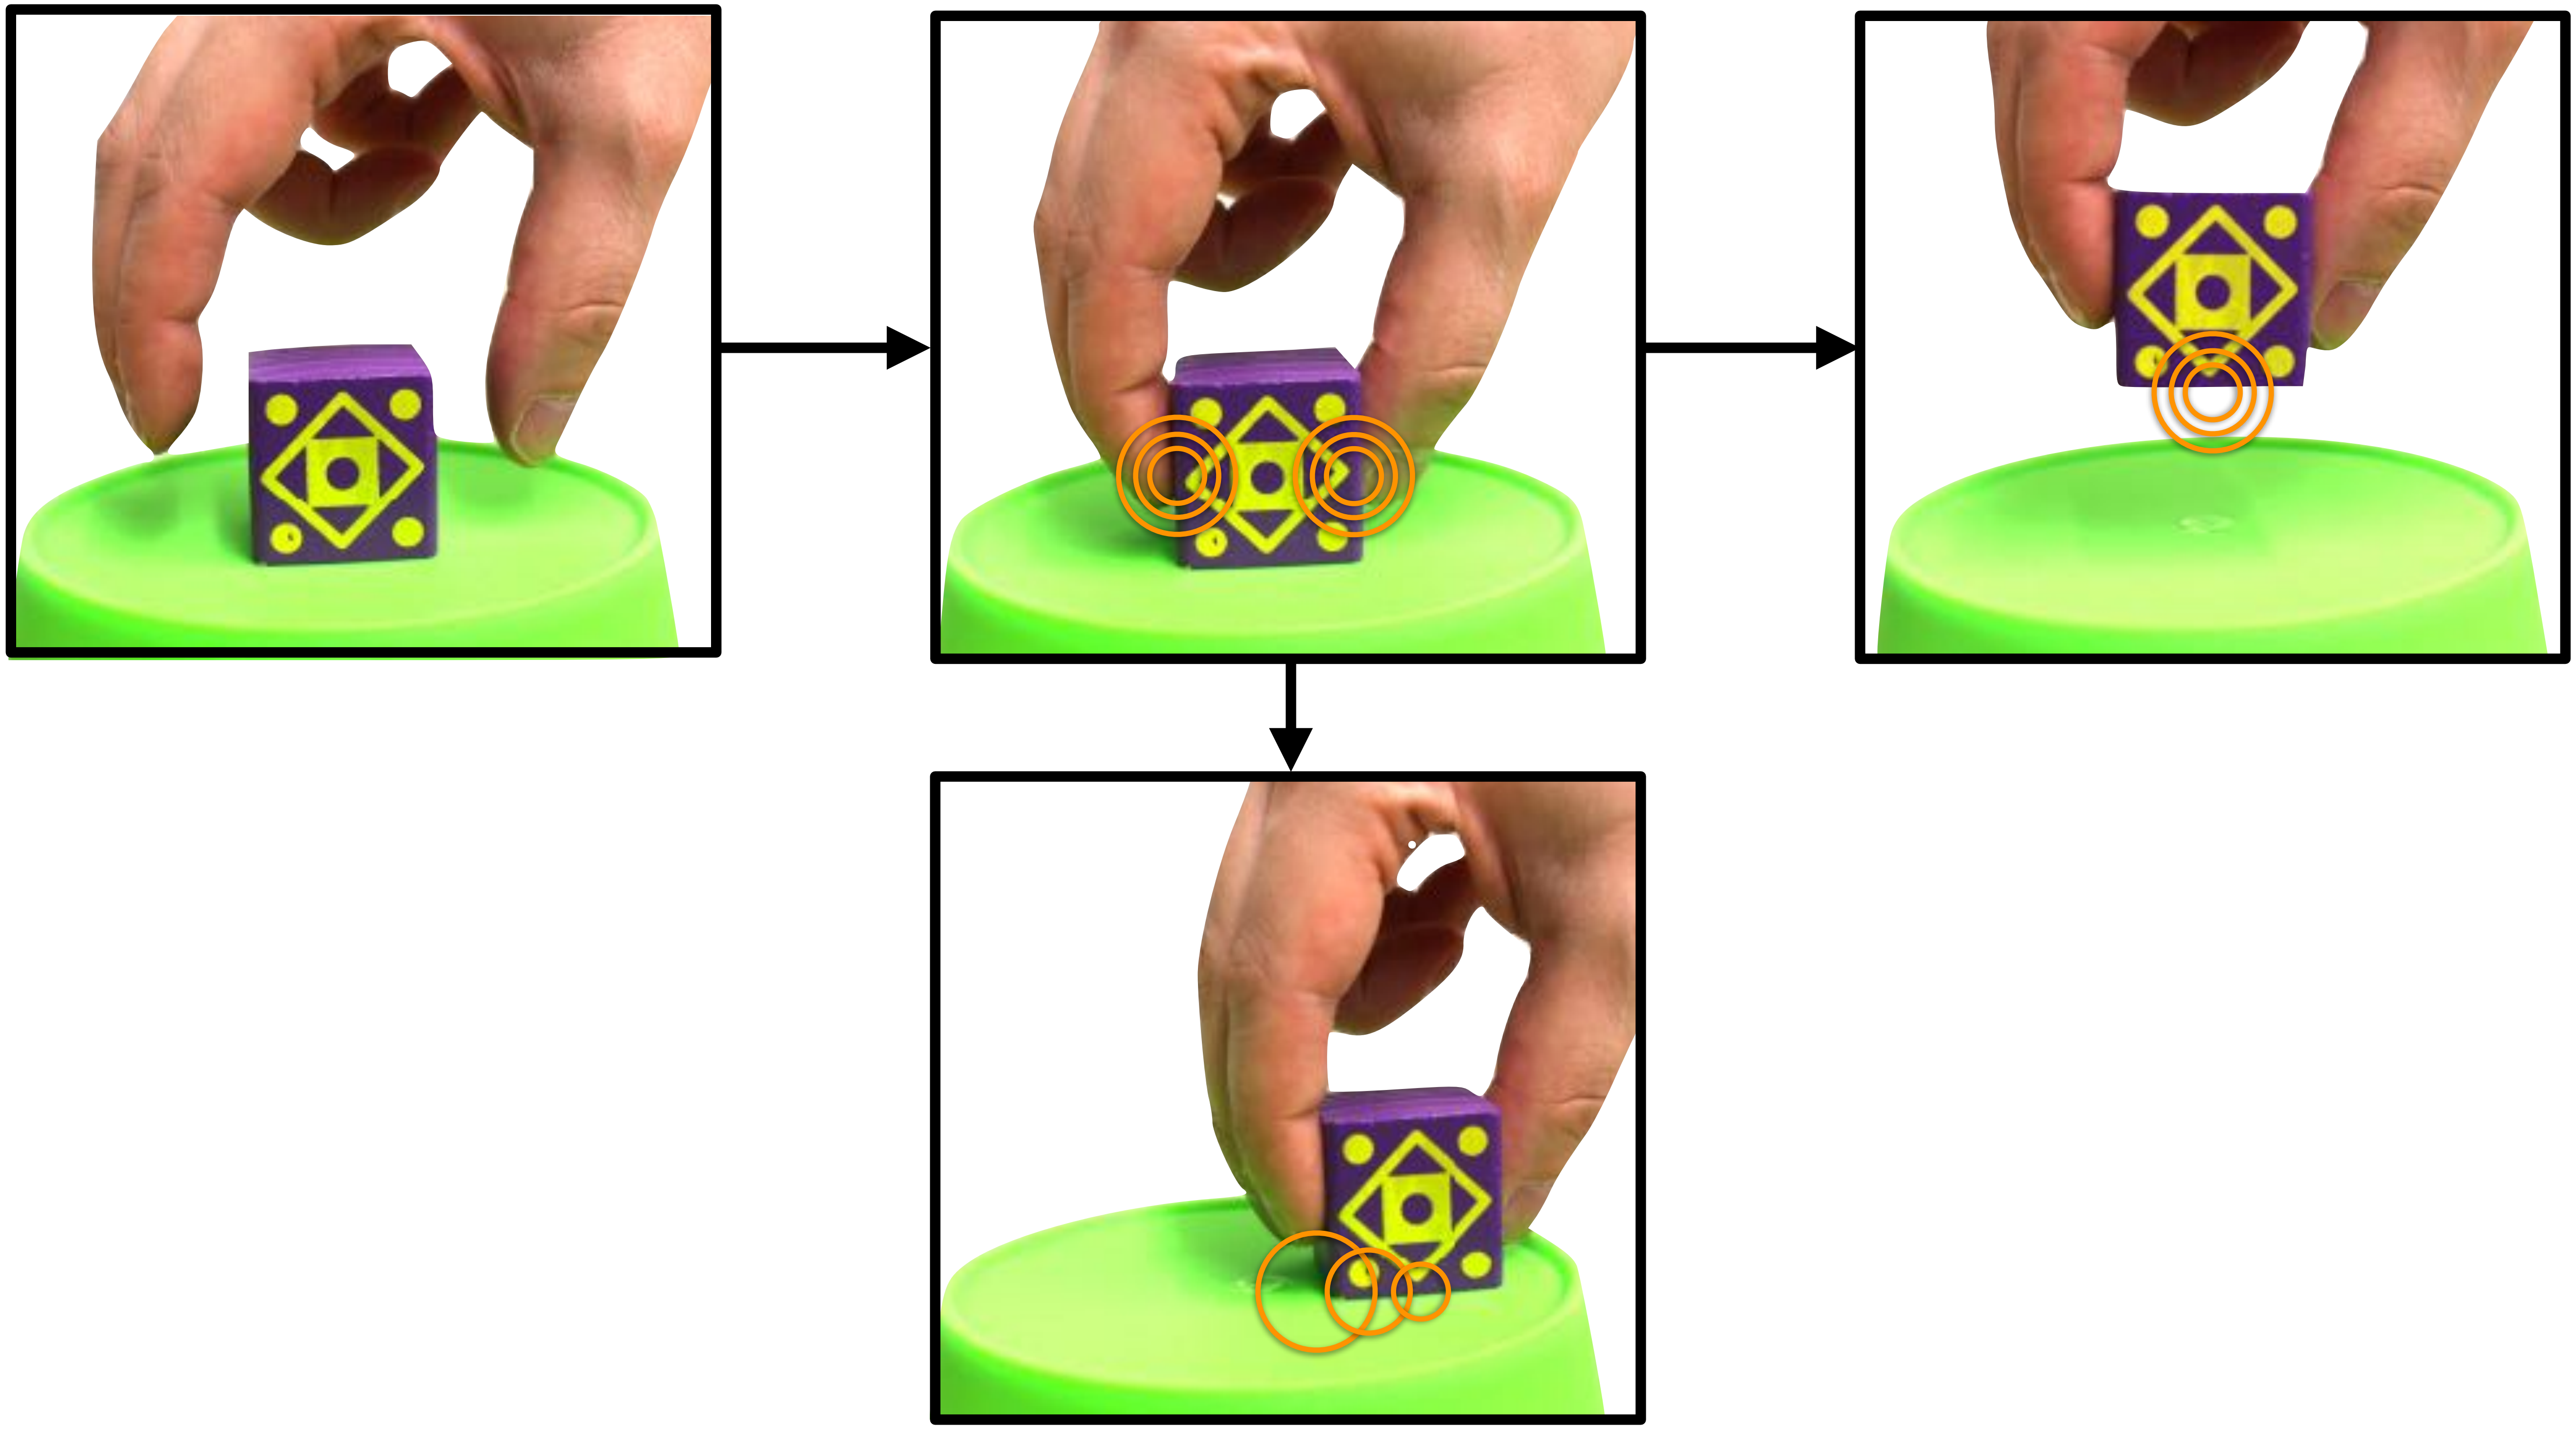

Sensory Feedback

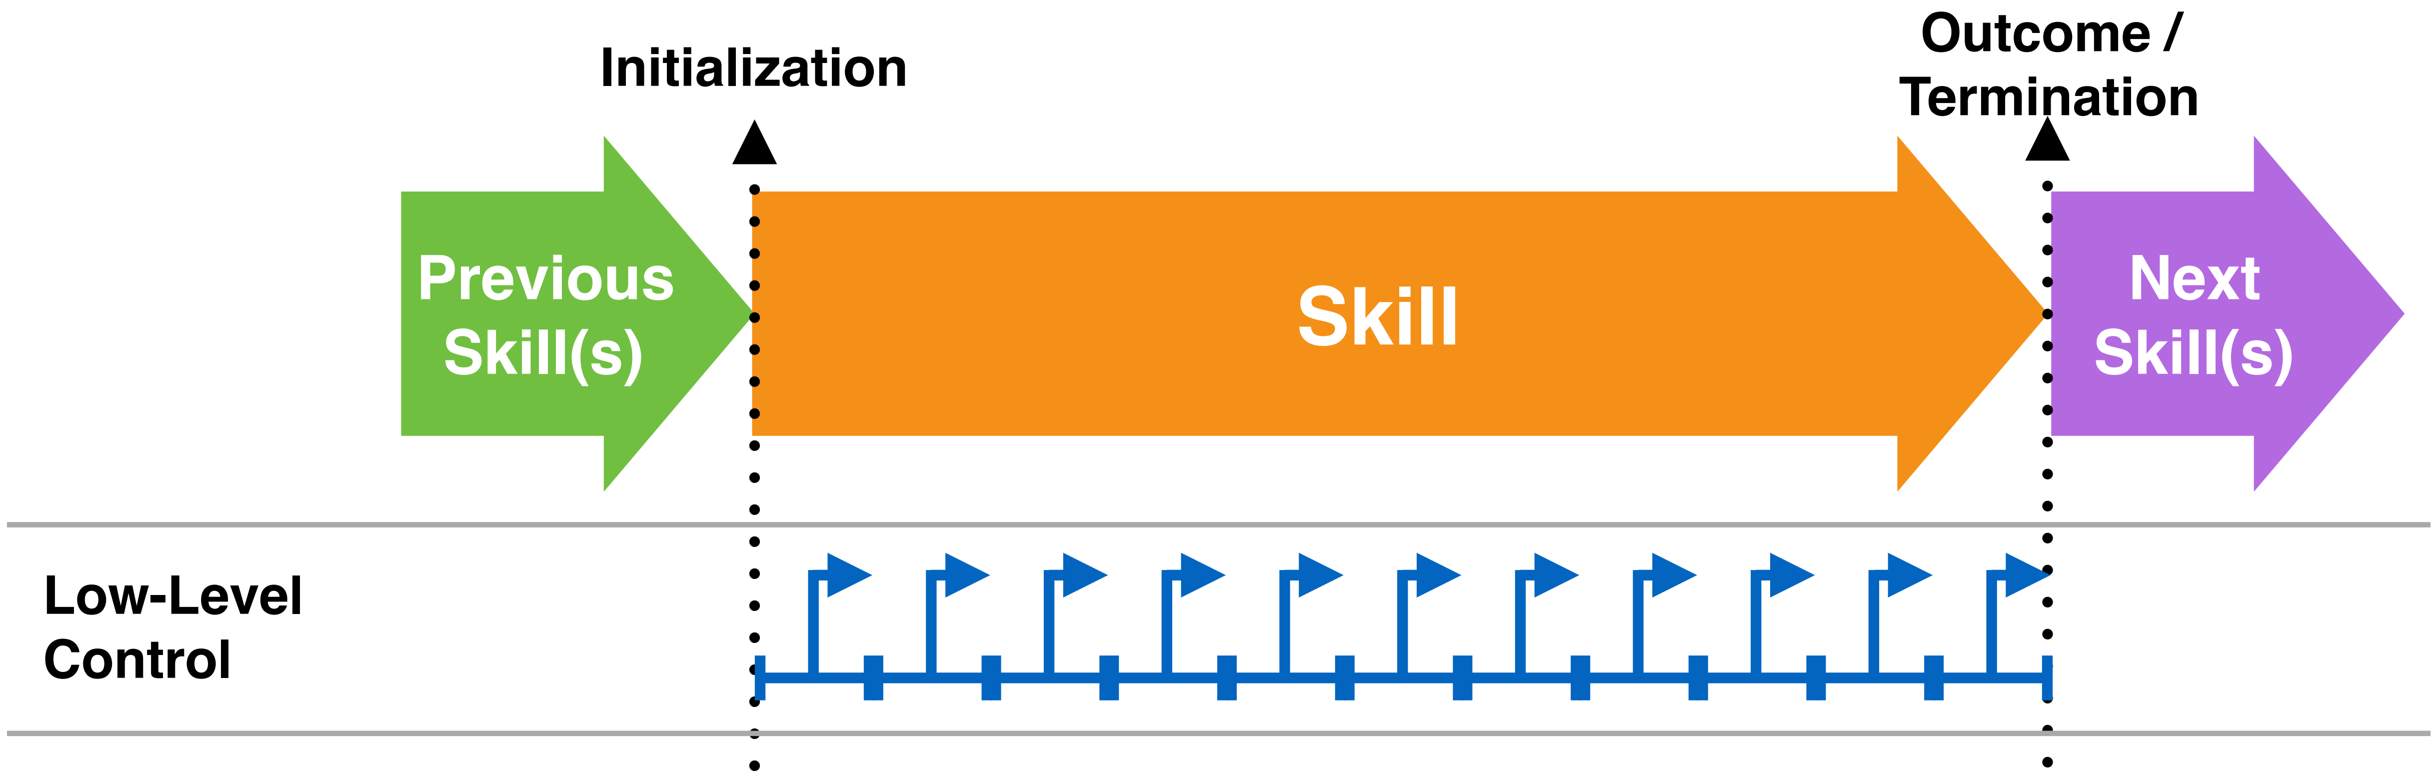

Continuous Feedback

- Can use tactile feedback to regulate forces during tasks

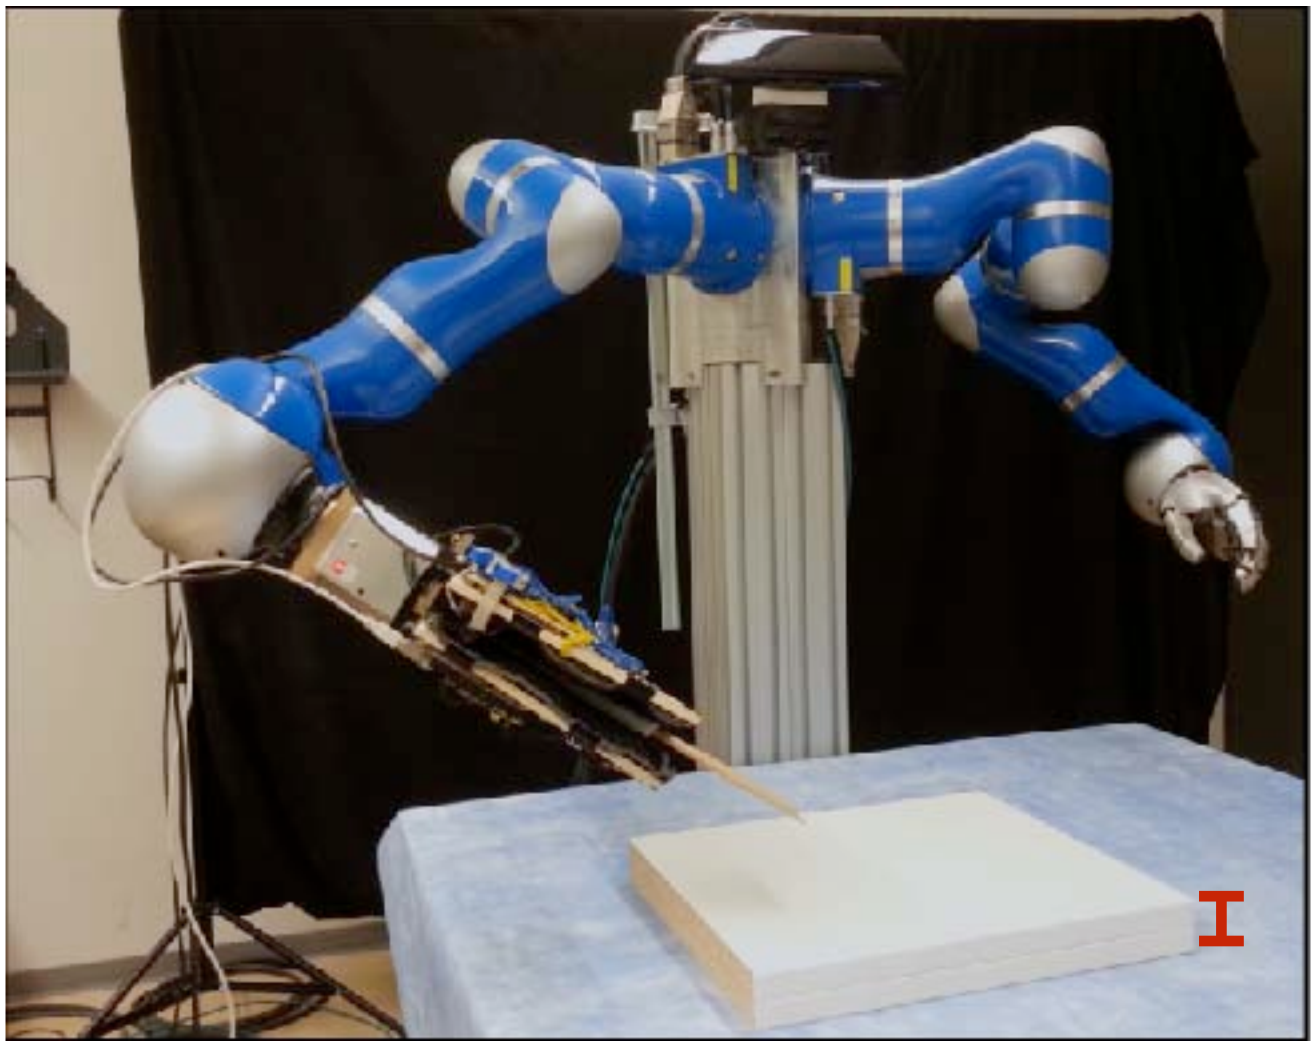

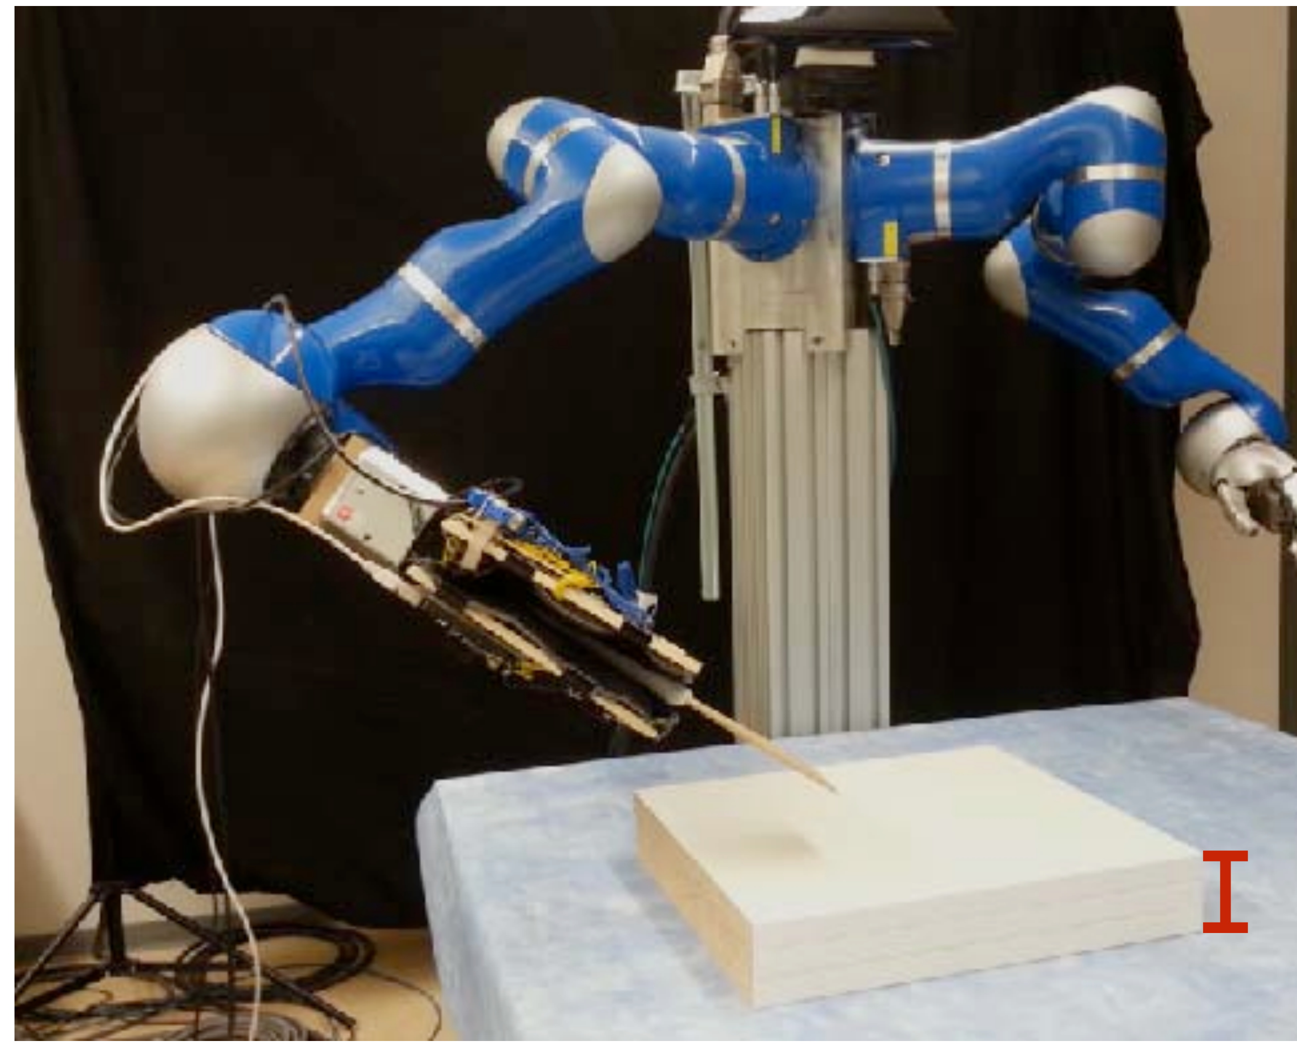

- Allows robot to compensate for perturbations

Sensory Feedback

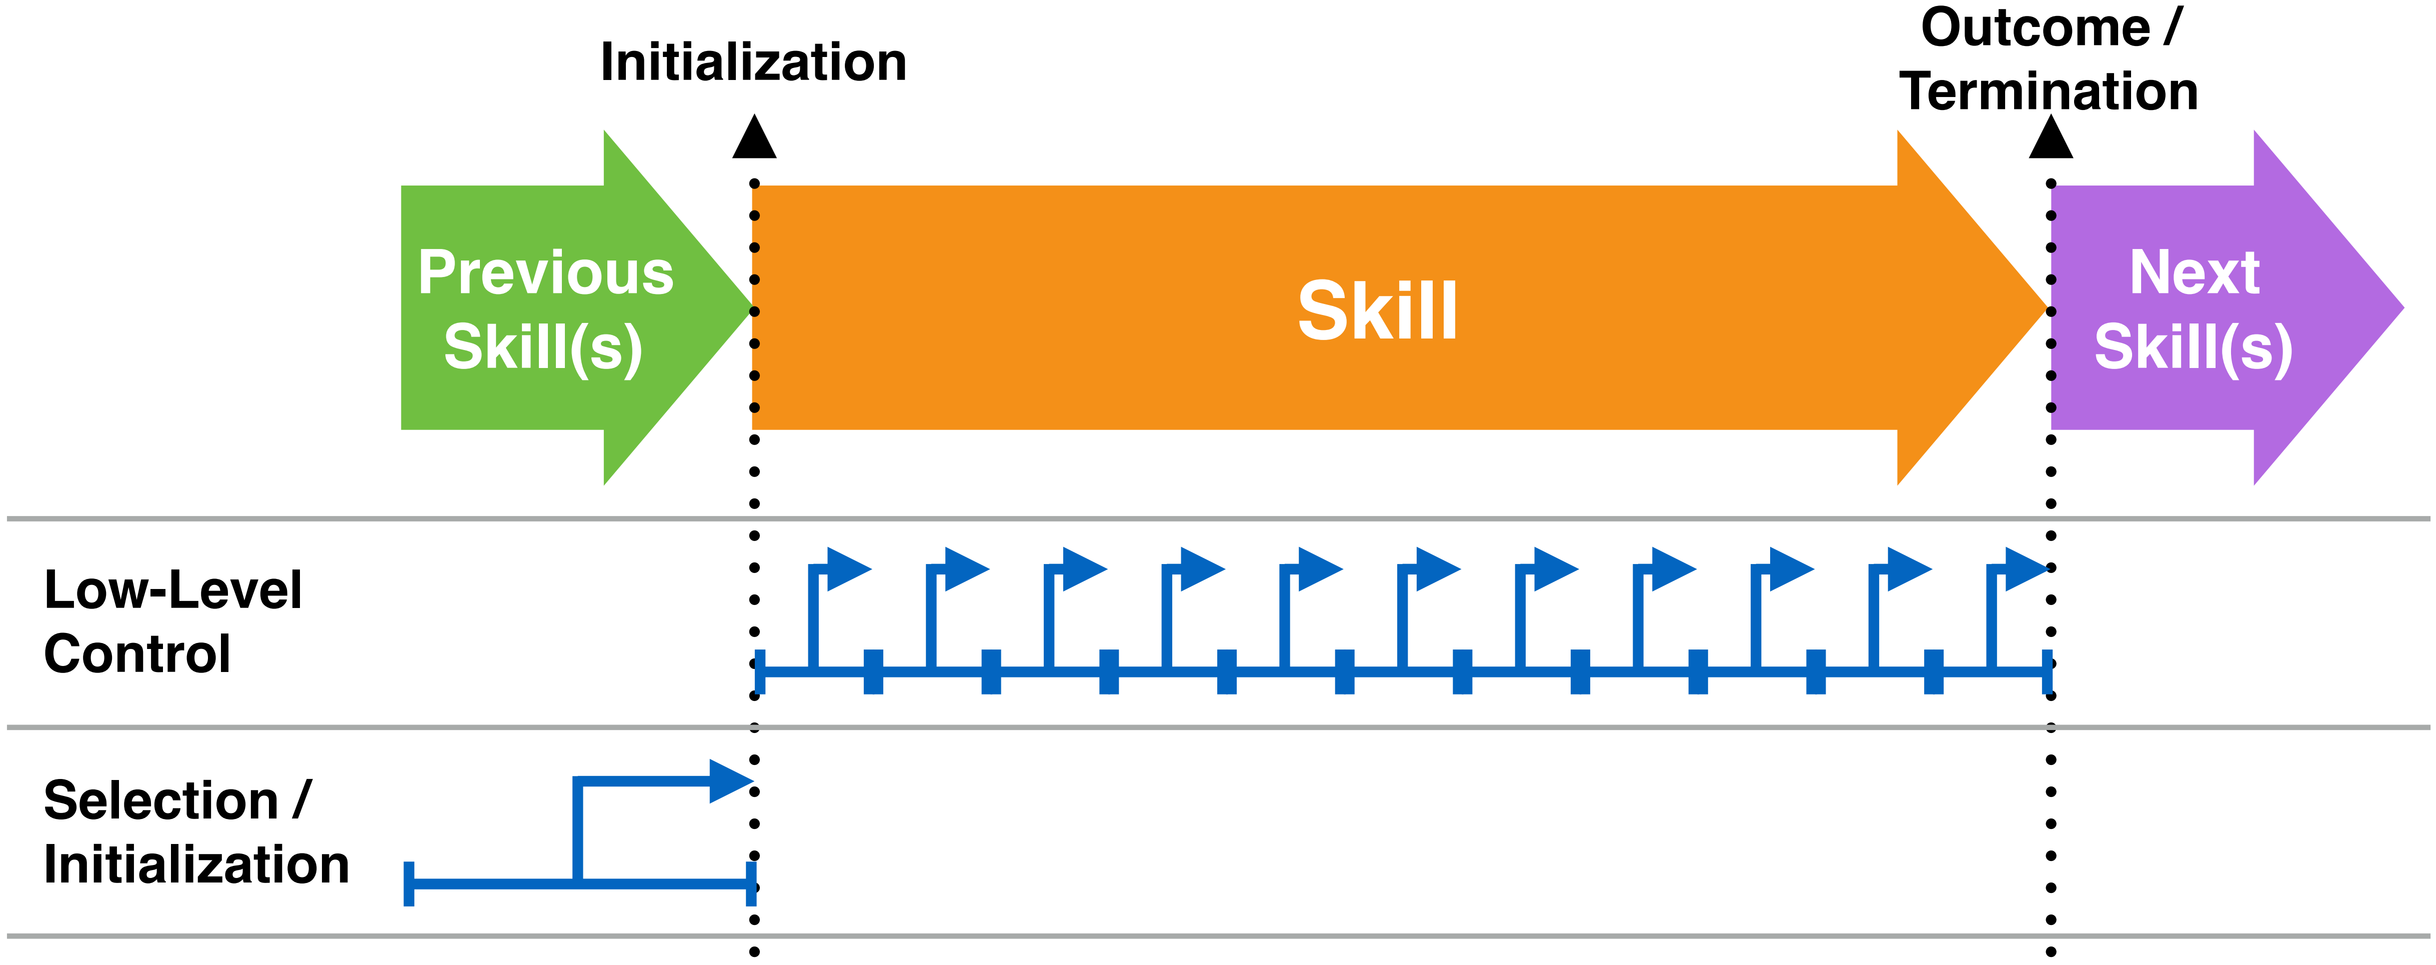

Skill Initialization and Preconditions

- Initialize skill parameters
  - Select skills parameters given current set of objects
- Check preconditions
  - Conditions in which skill will result in the intended effect

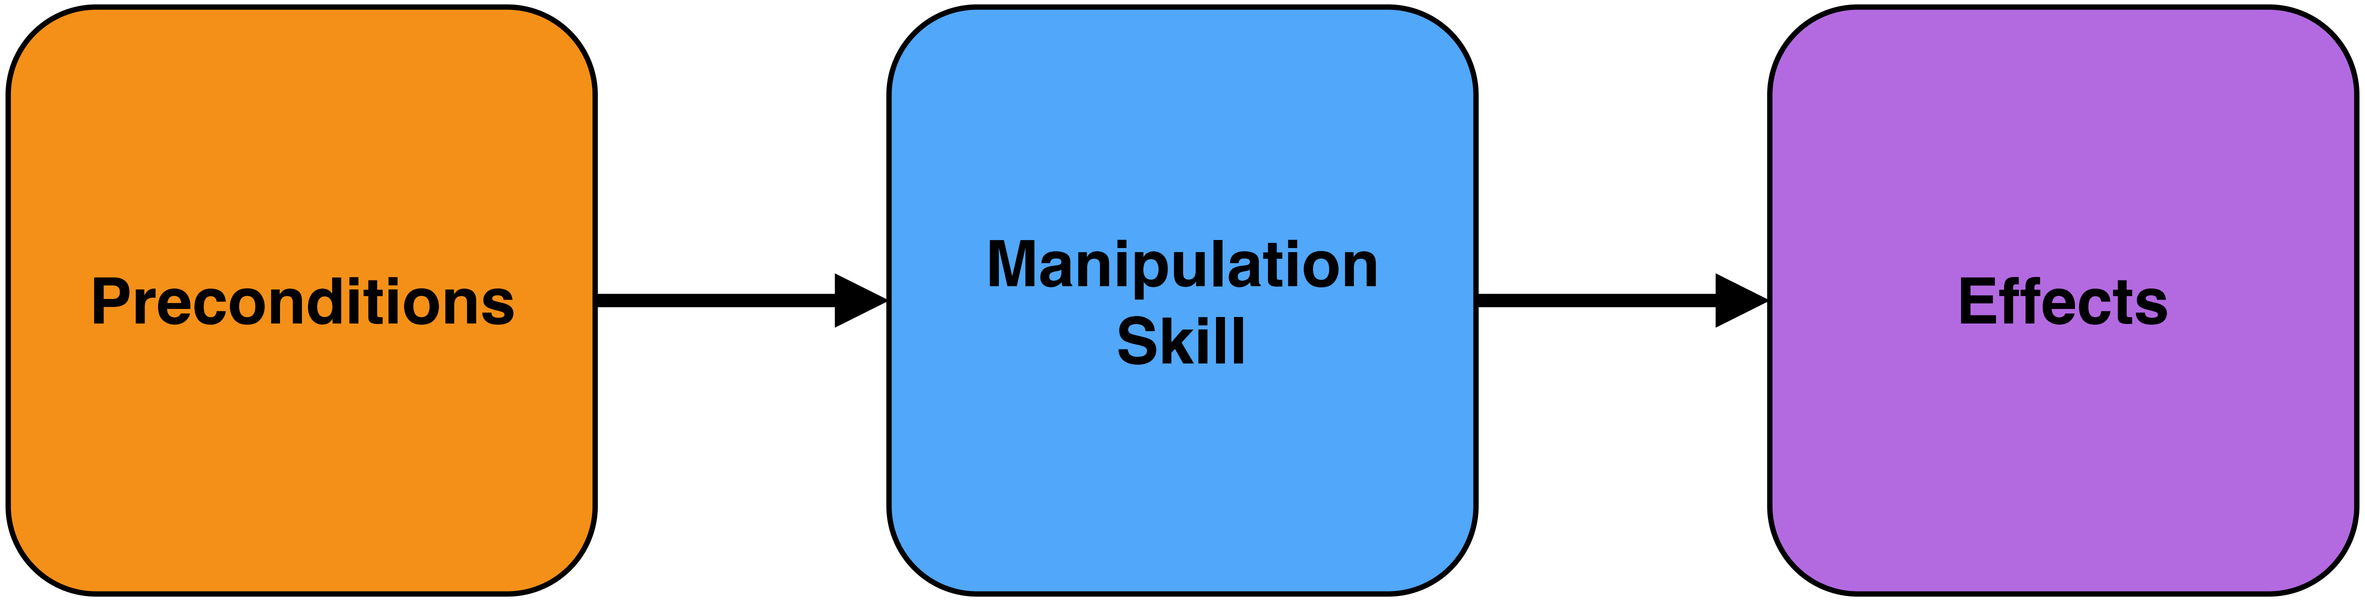

Learning Affordances of Objects

- Learning manipulations afforded by individual objects


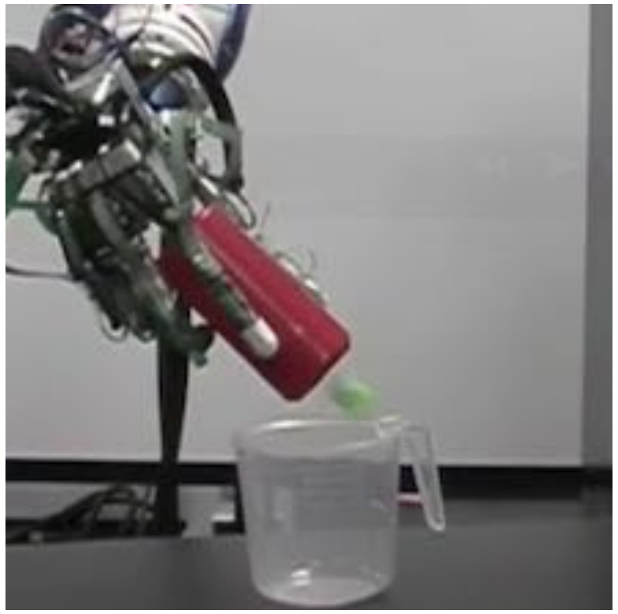
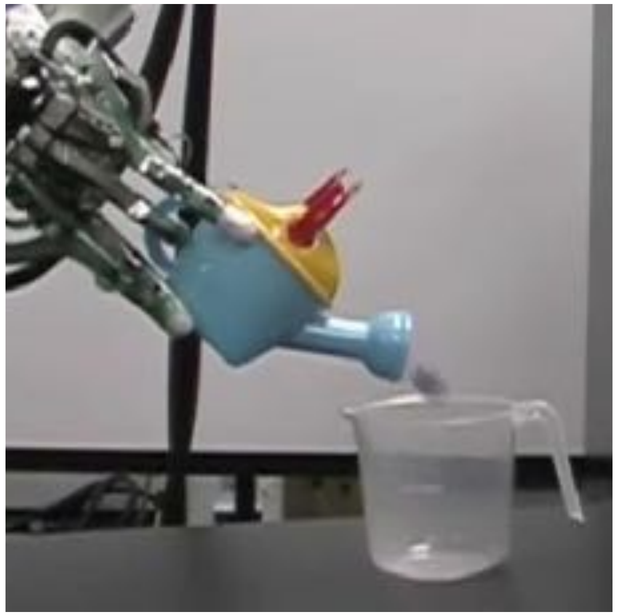

- Learn interactions afforded by pairs of objects




Learning Object Constellations

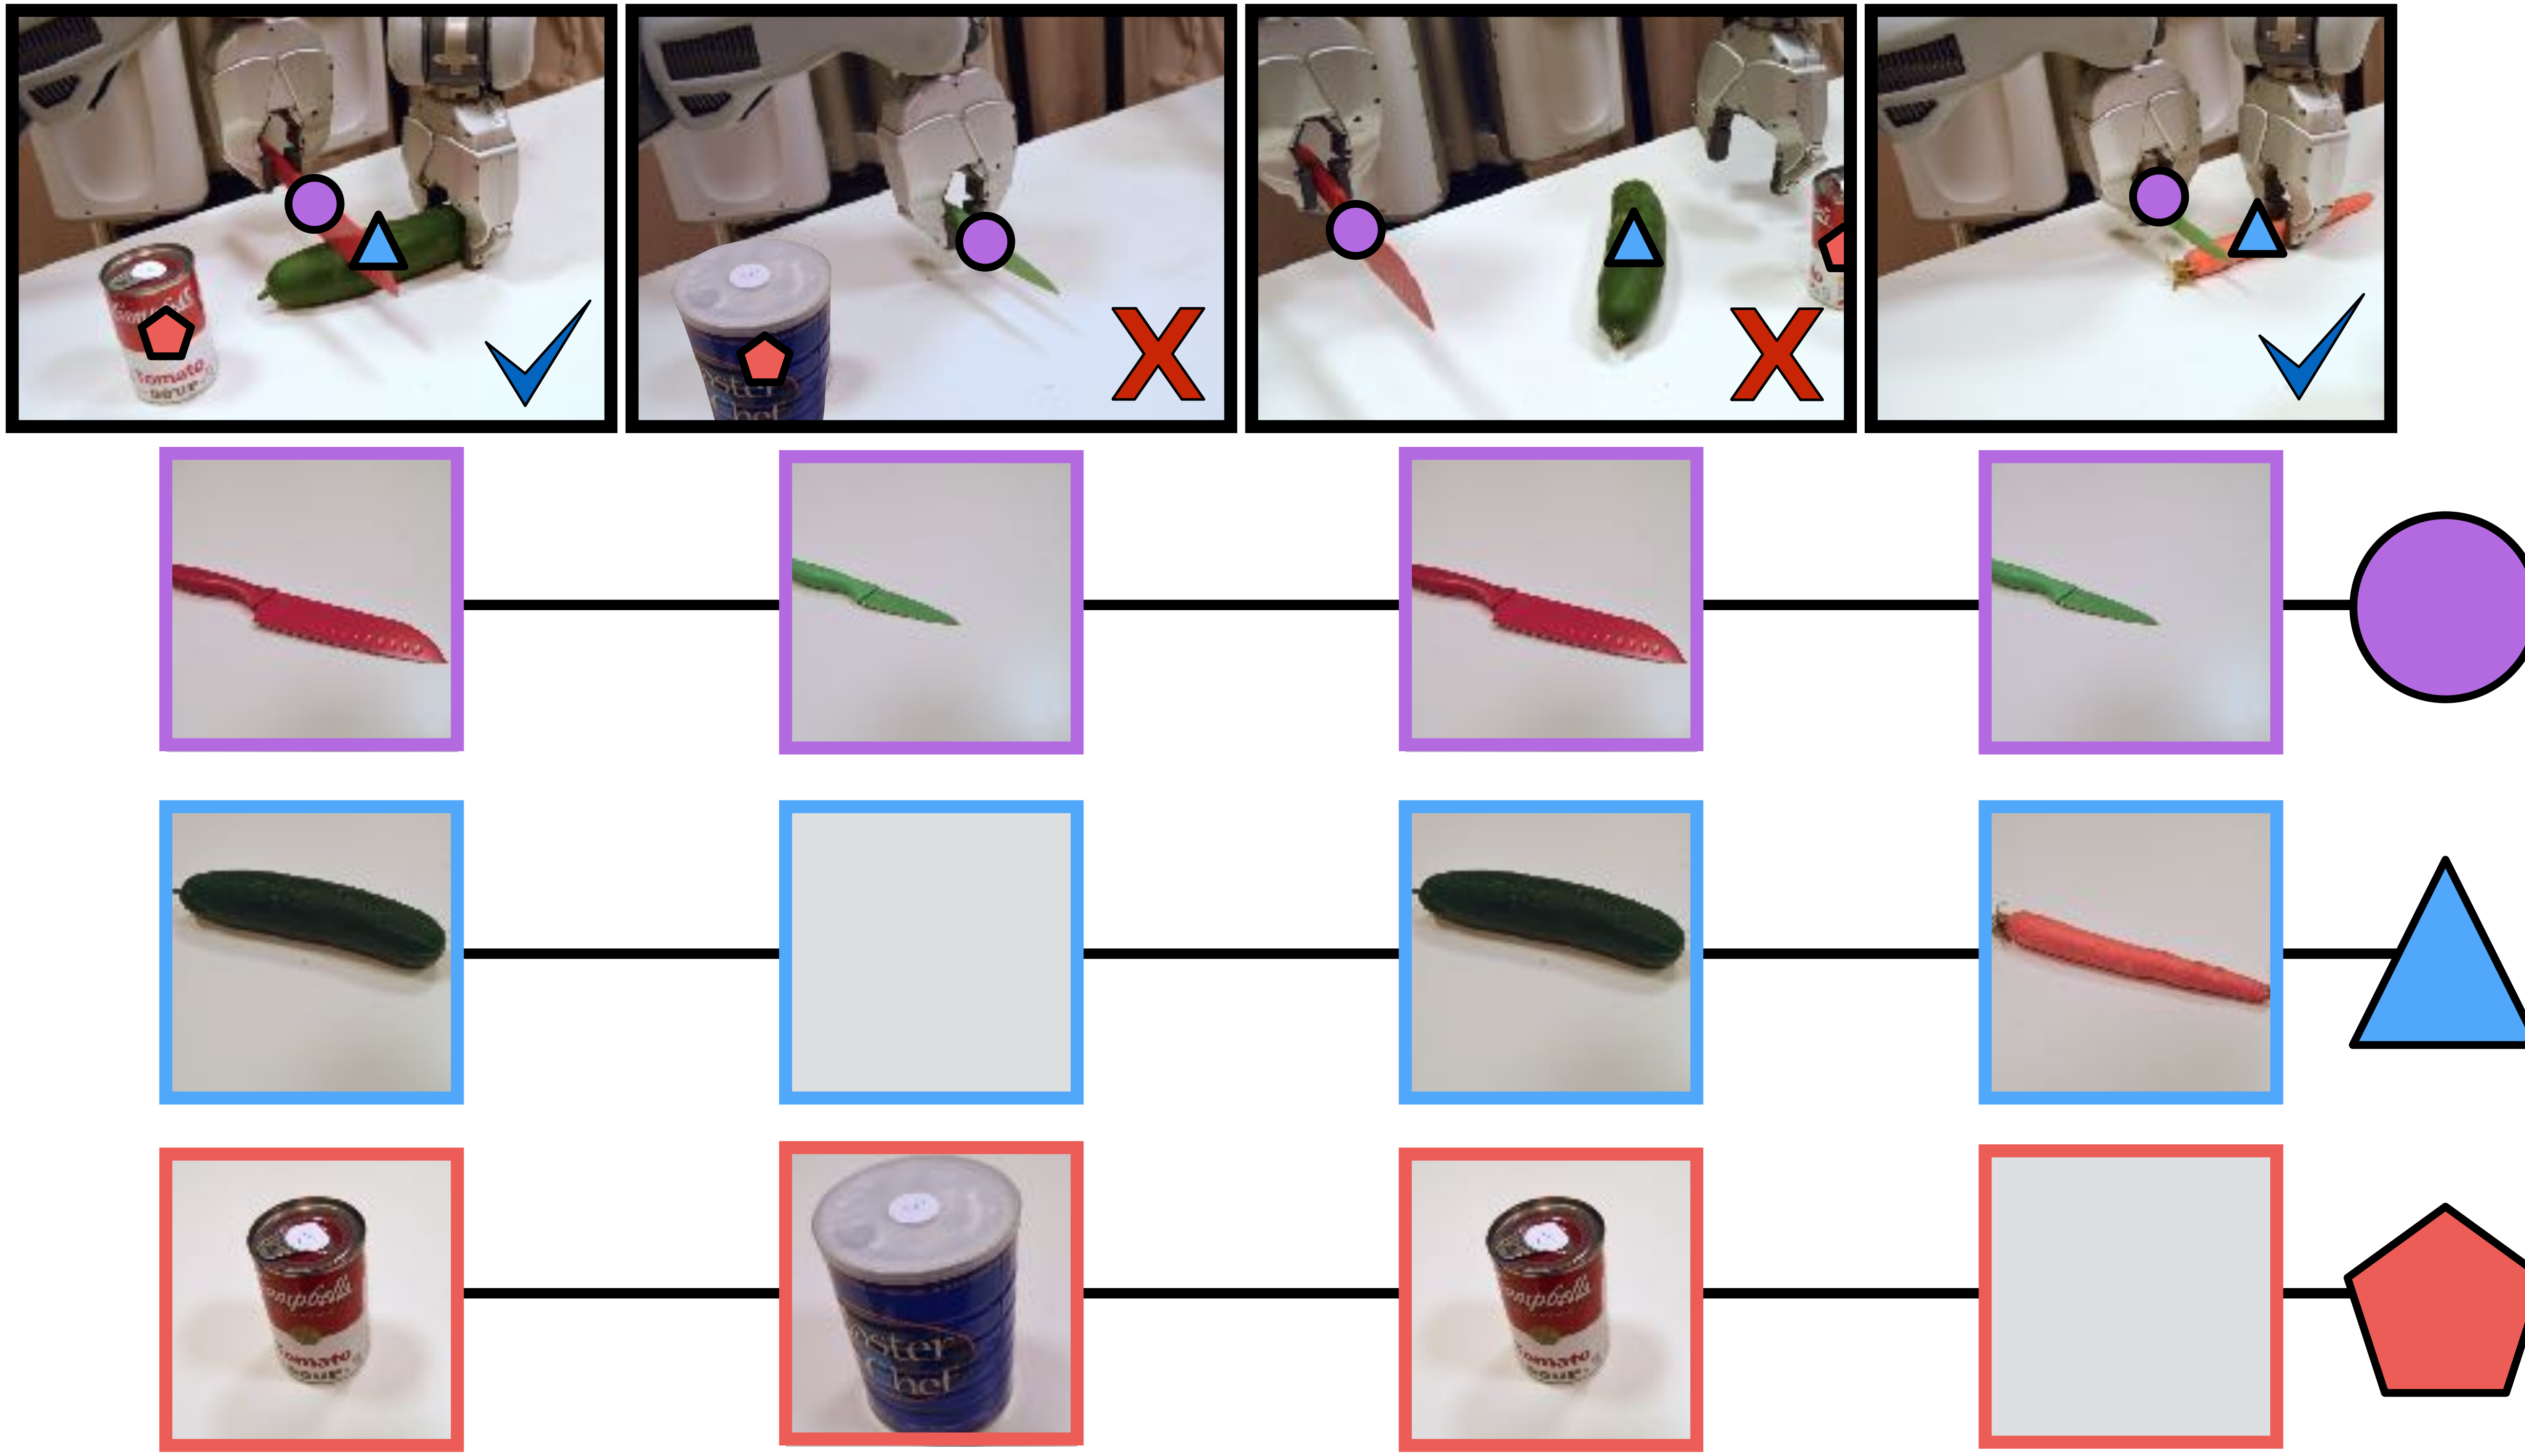

Object Constellations

Robot cannot differentiate between object scenes

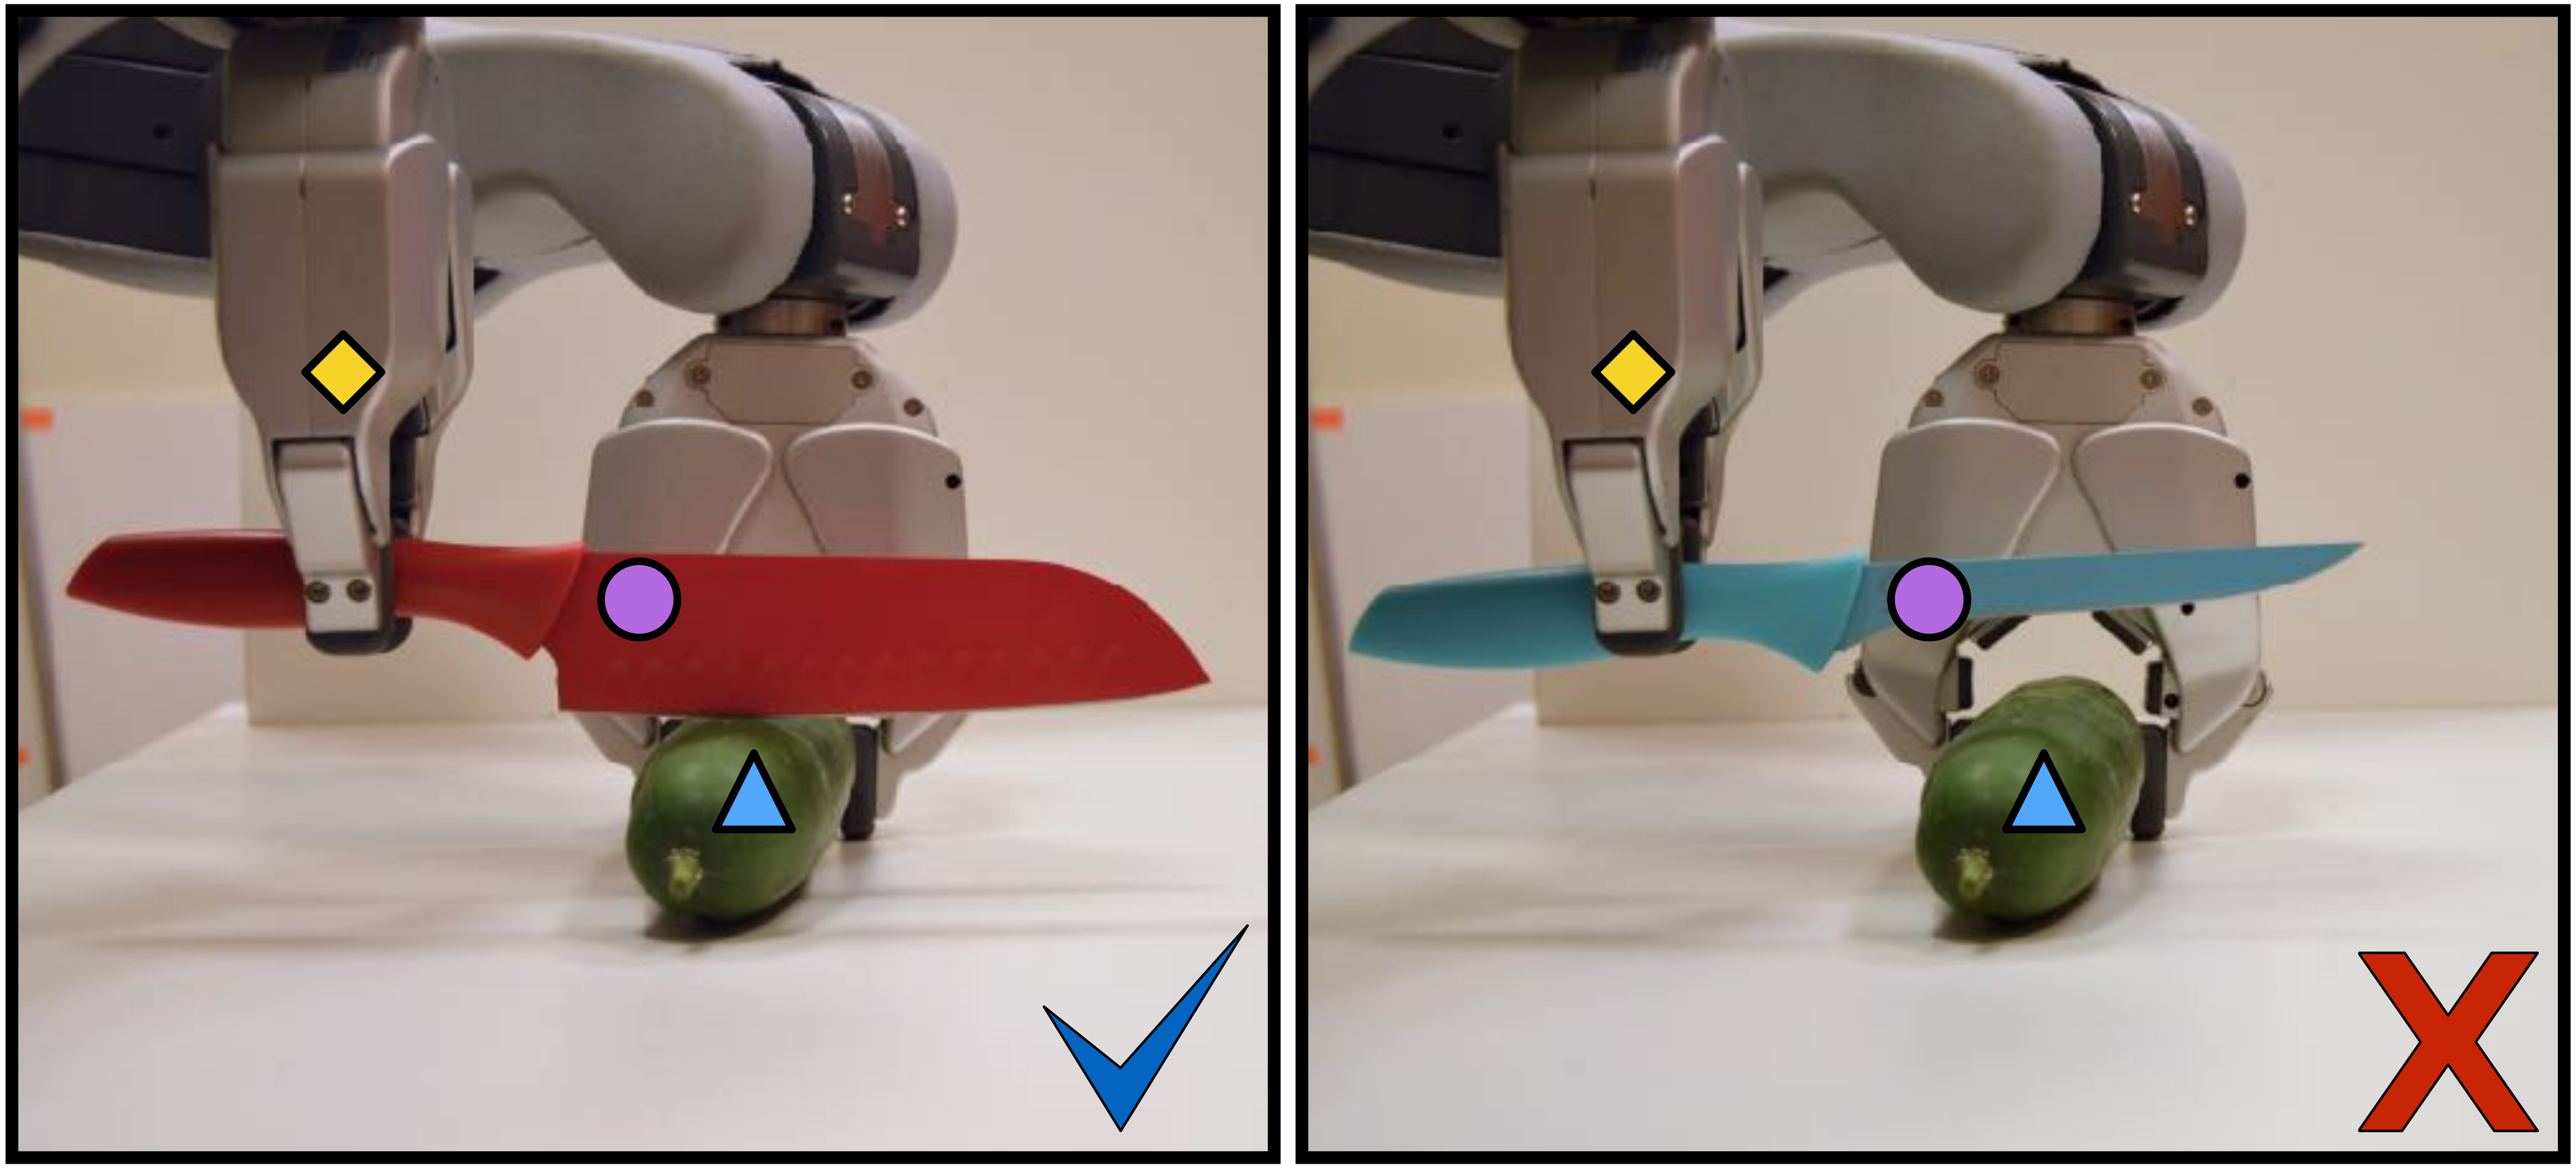

Object Constellations

- Parts and interactions provide additional details

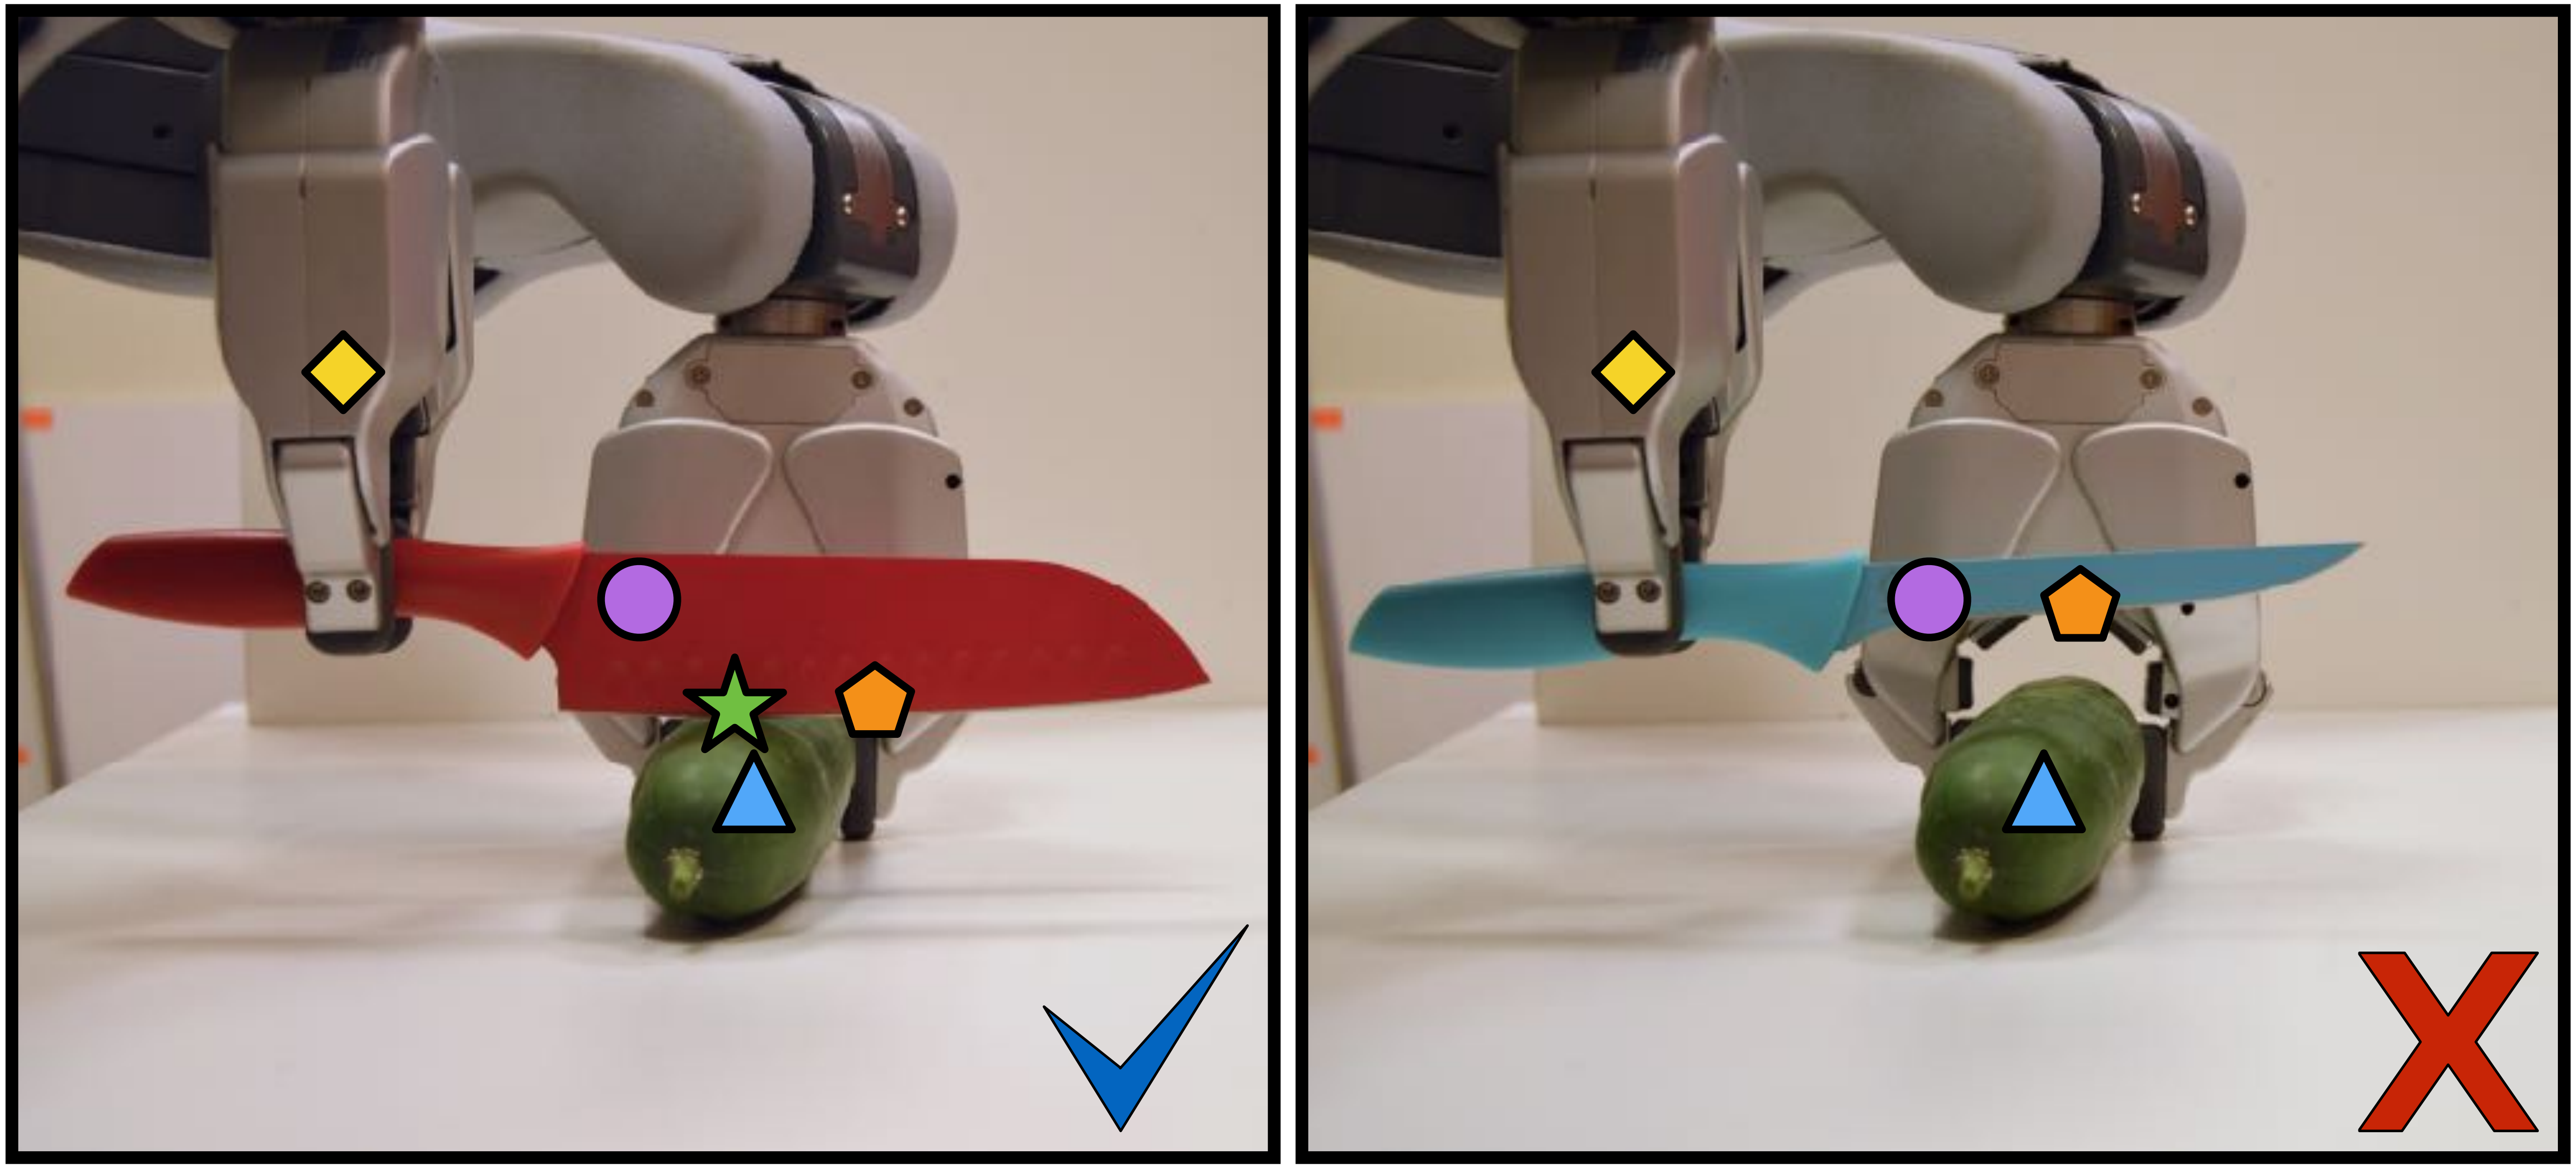

Learning Object Constellations

- Learn constellation preconditions using random forests

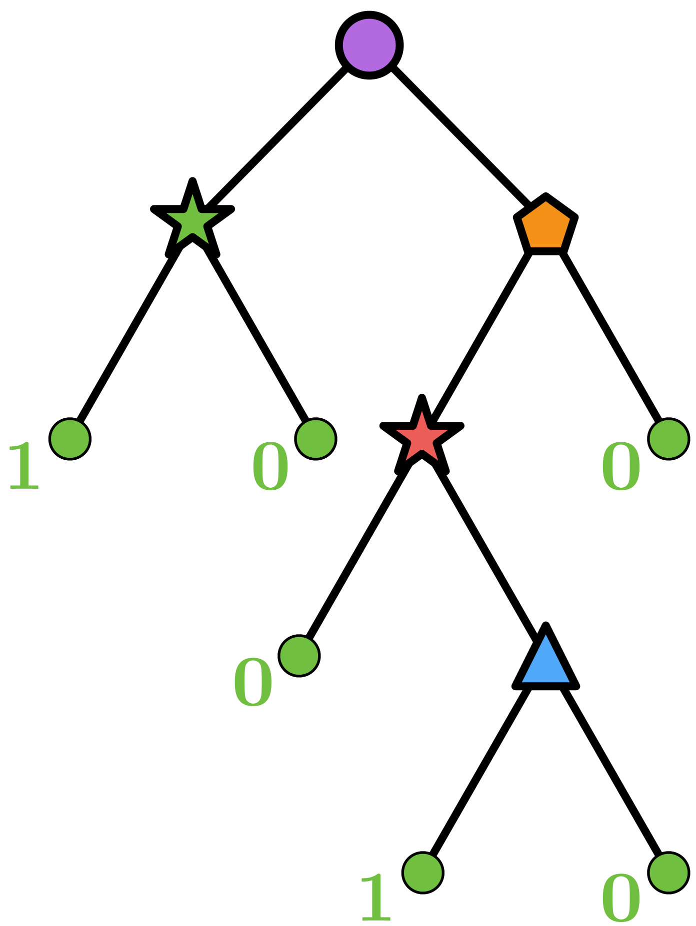

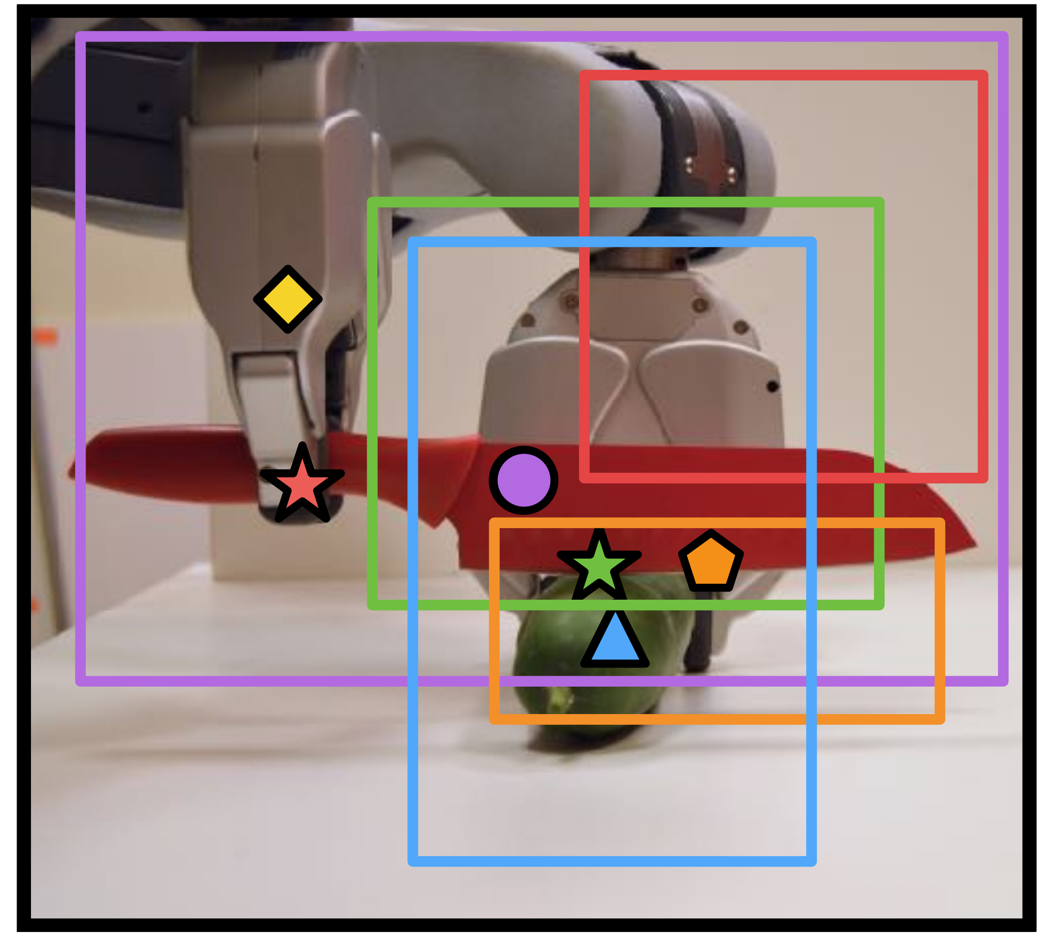

- Trees test for scene elements in different regions
- Combine predictions from ensemble of 100 trees

Results

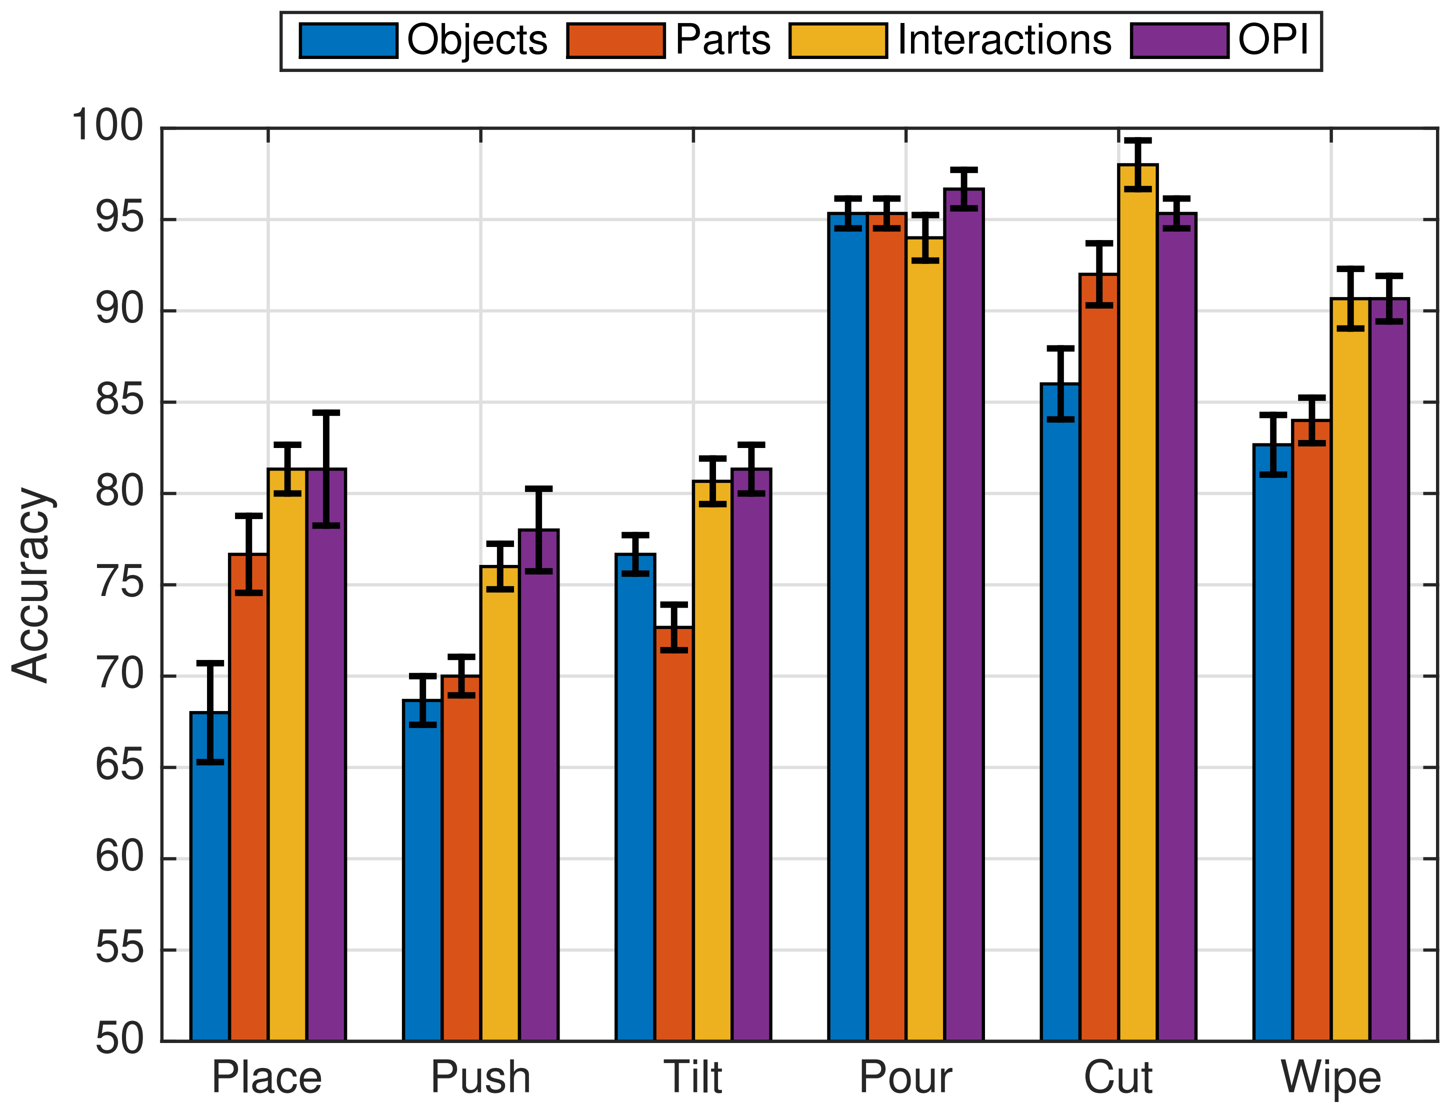

Object Only	Obj+Part+Interact
79.6%	87.2%

Preconditions over Skill Parameters

- Skill initialization and preconditions often intertwined
  - Determine if skill parameters will result in desired effect
- Sample and evaluate different skill parameters (e.g. goals)

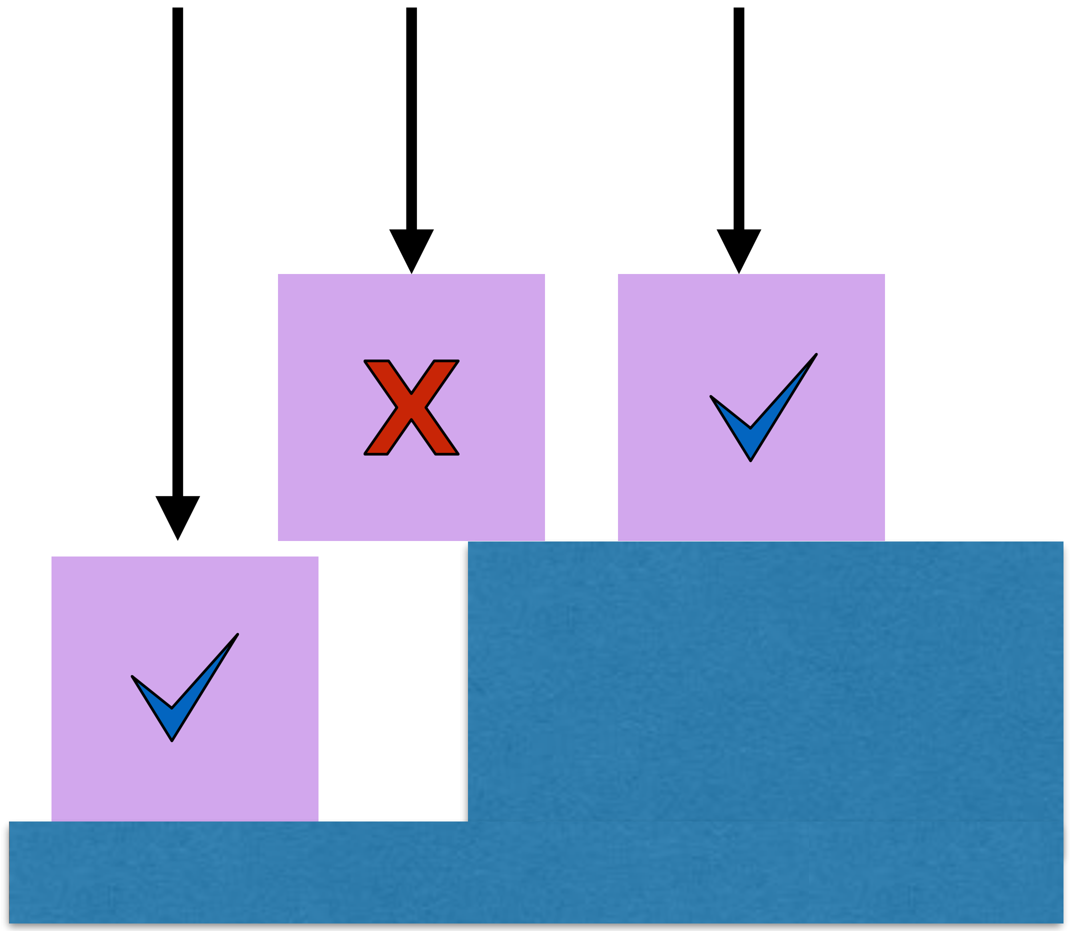

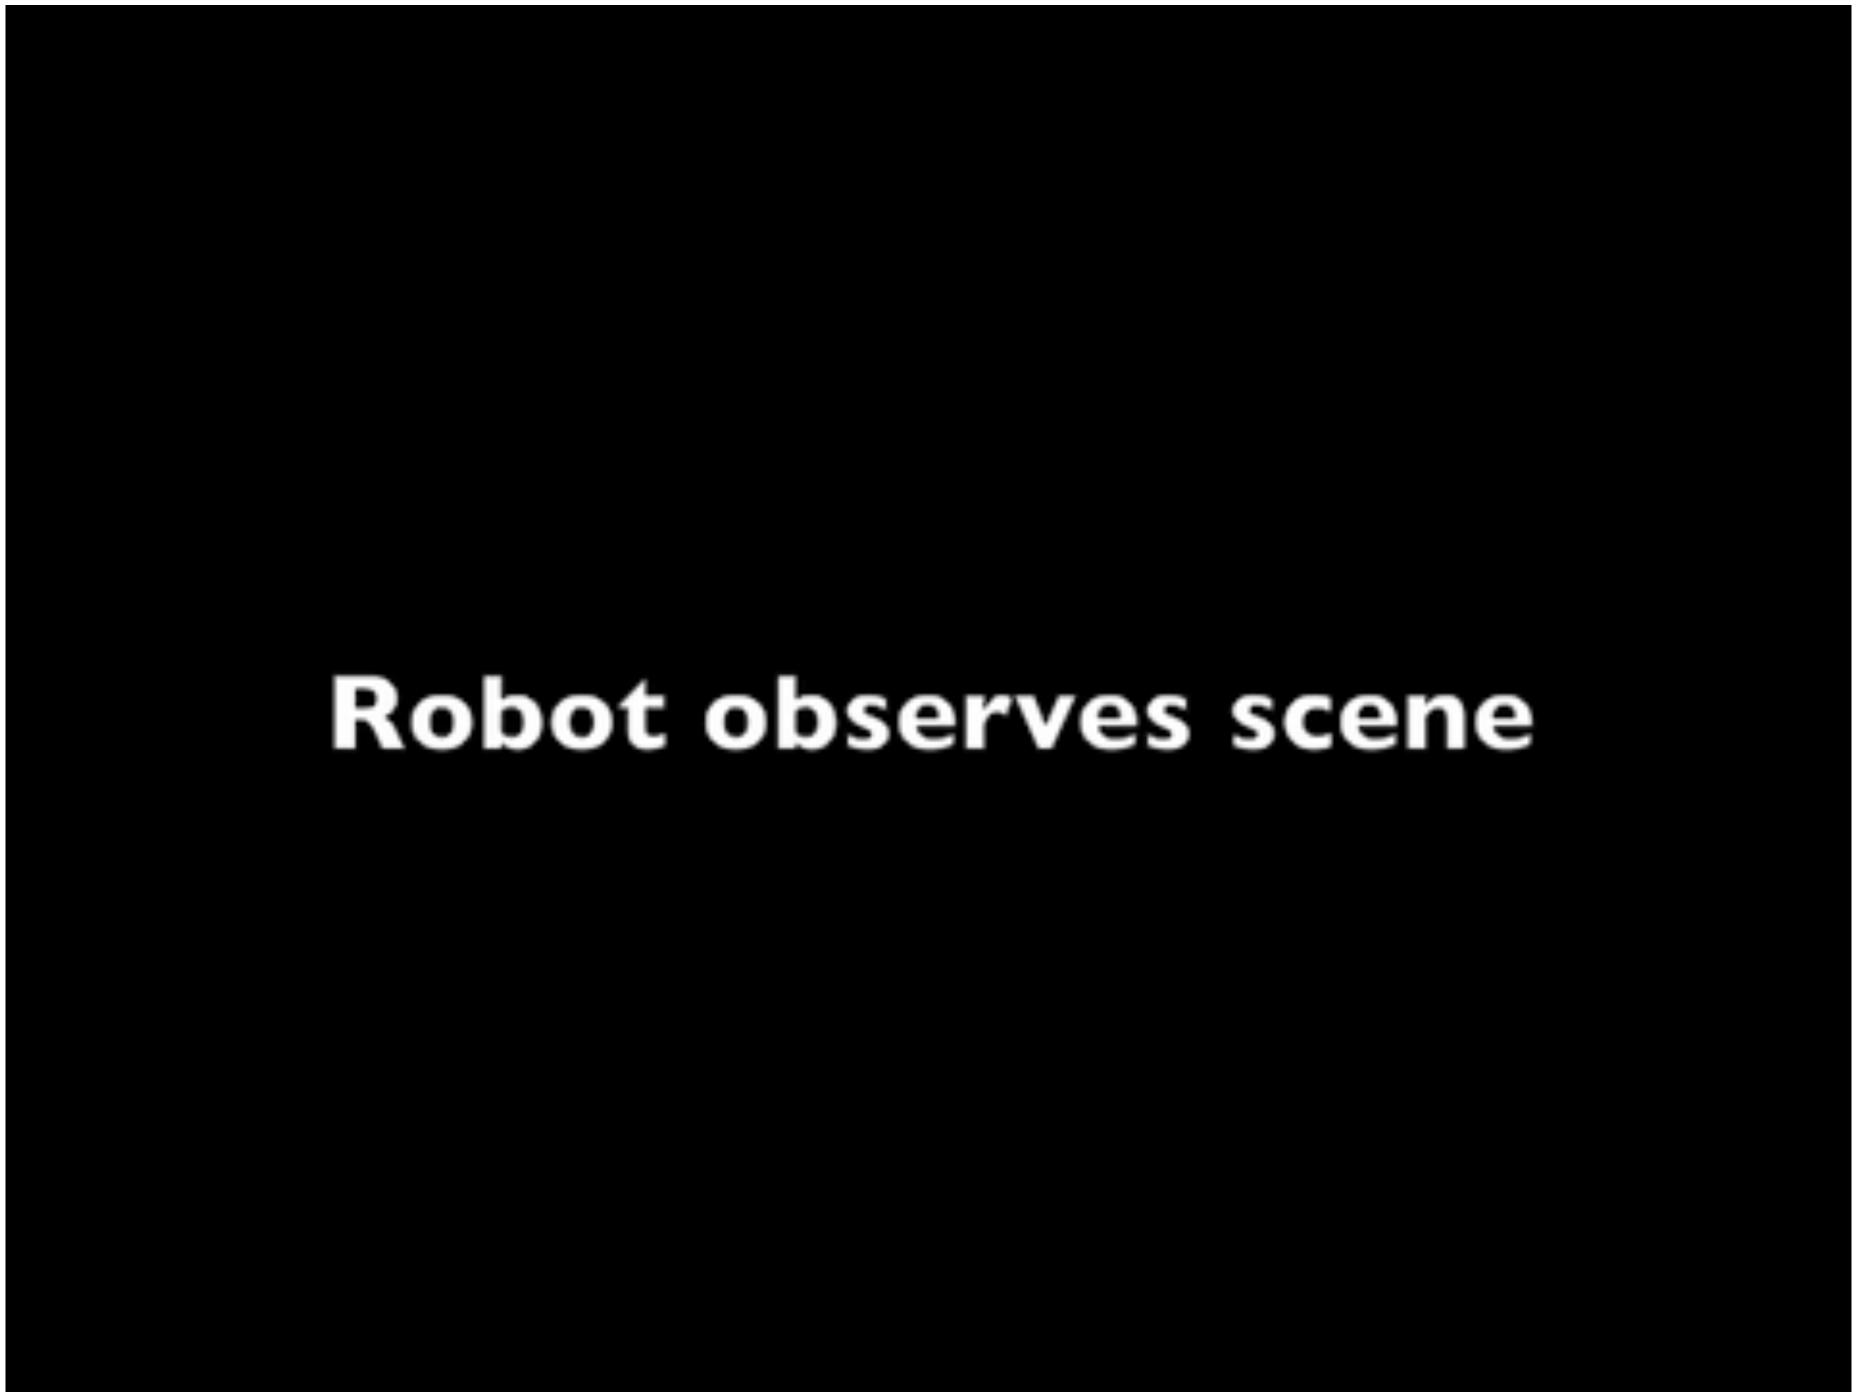
Robot observes scene

Block Stacking

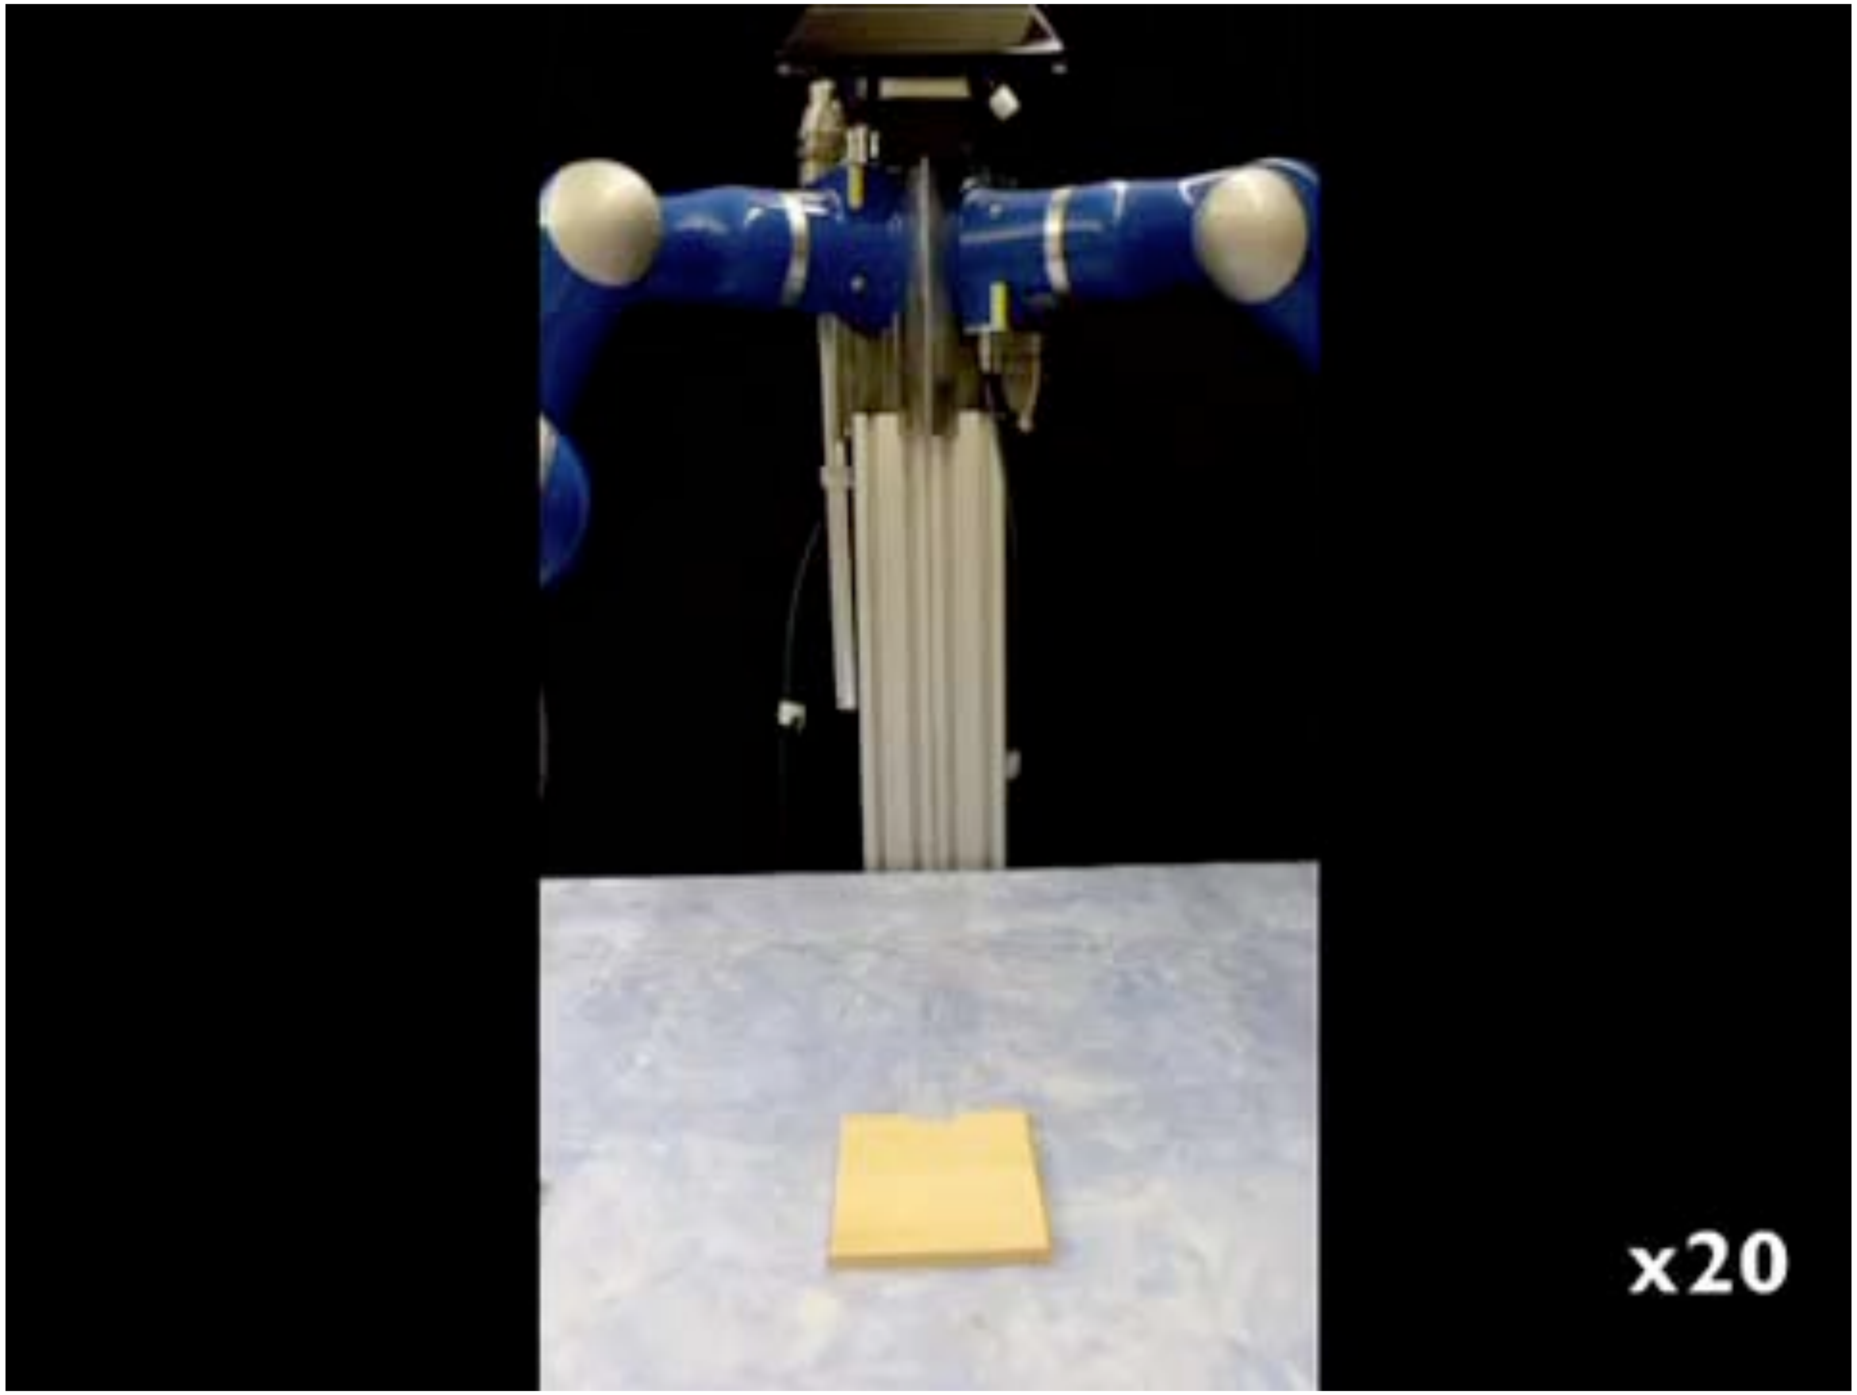

Sensory Feedback

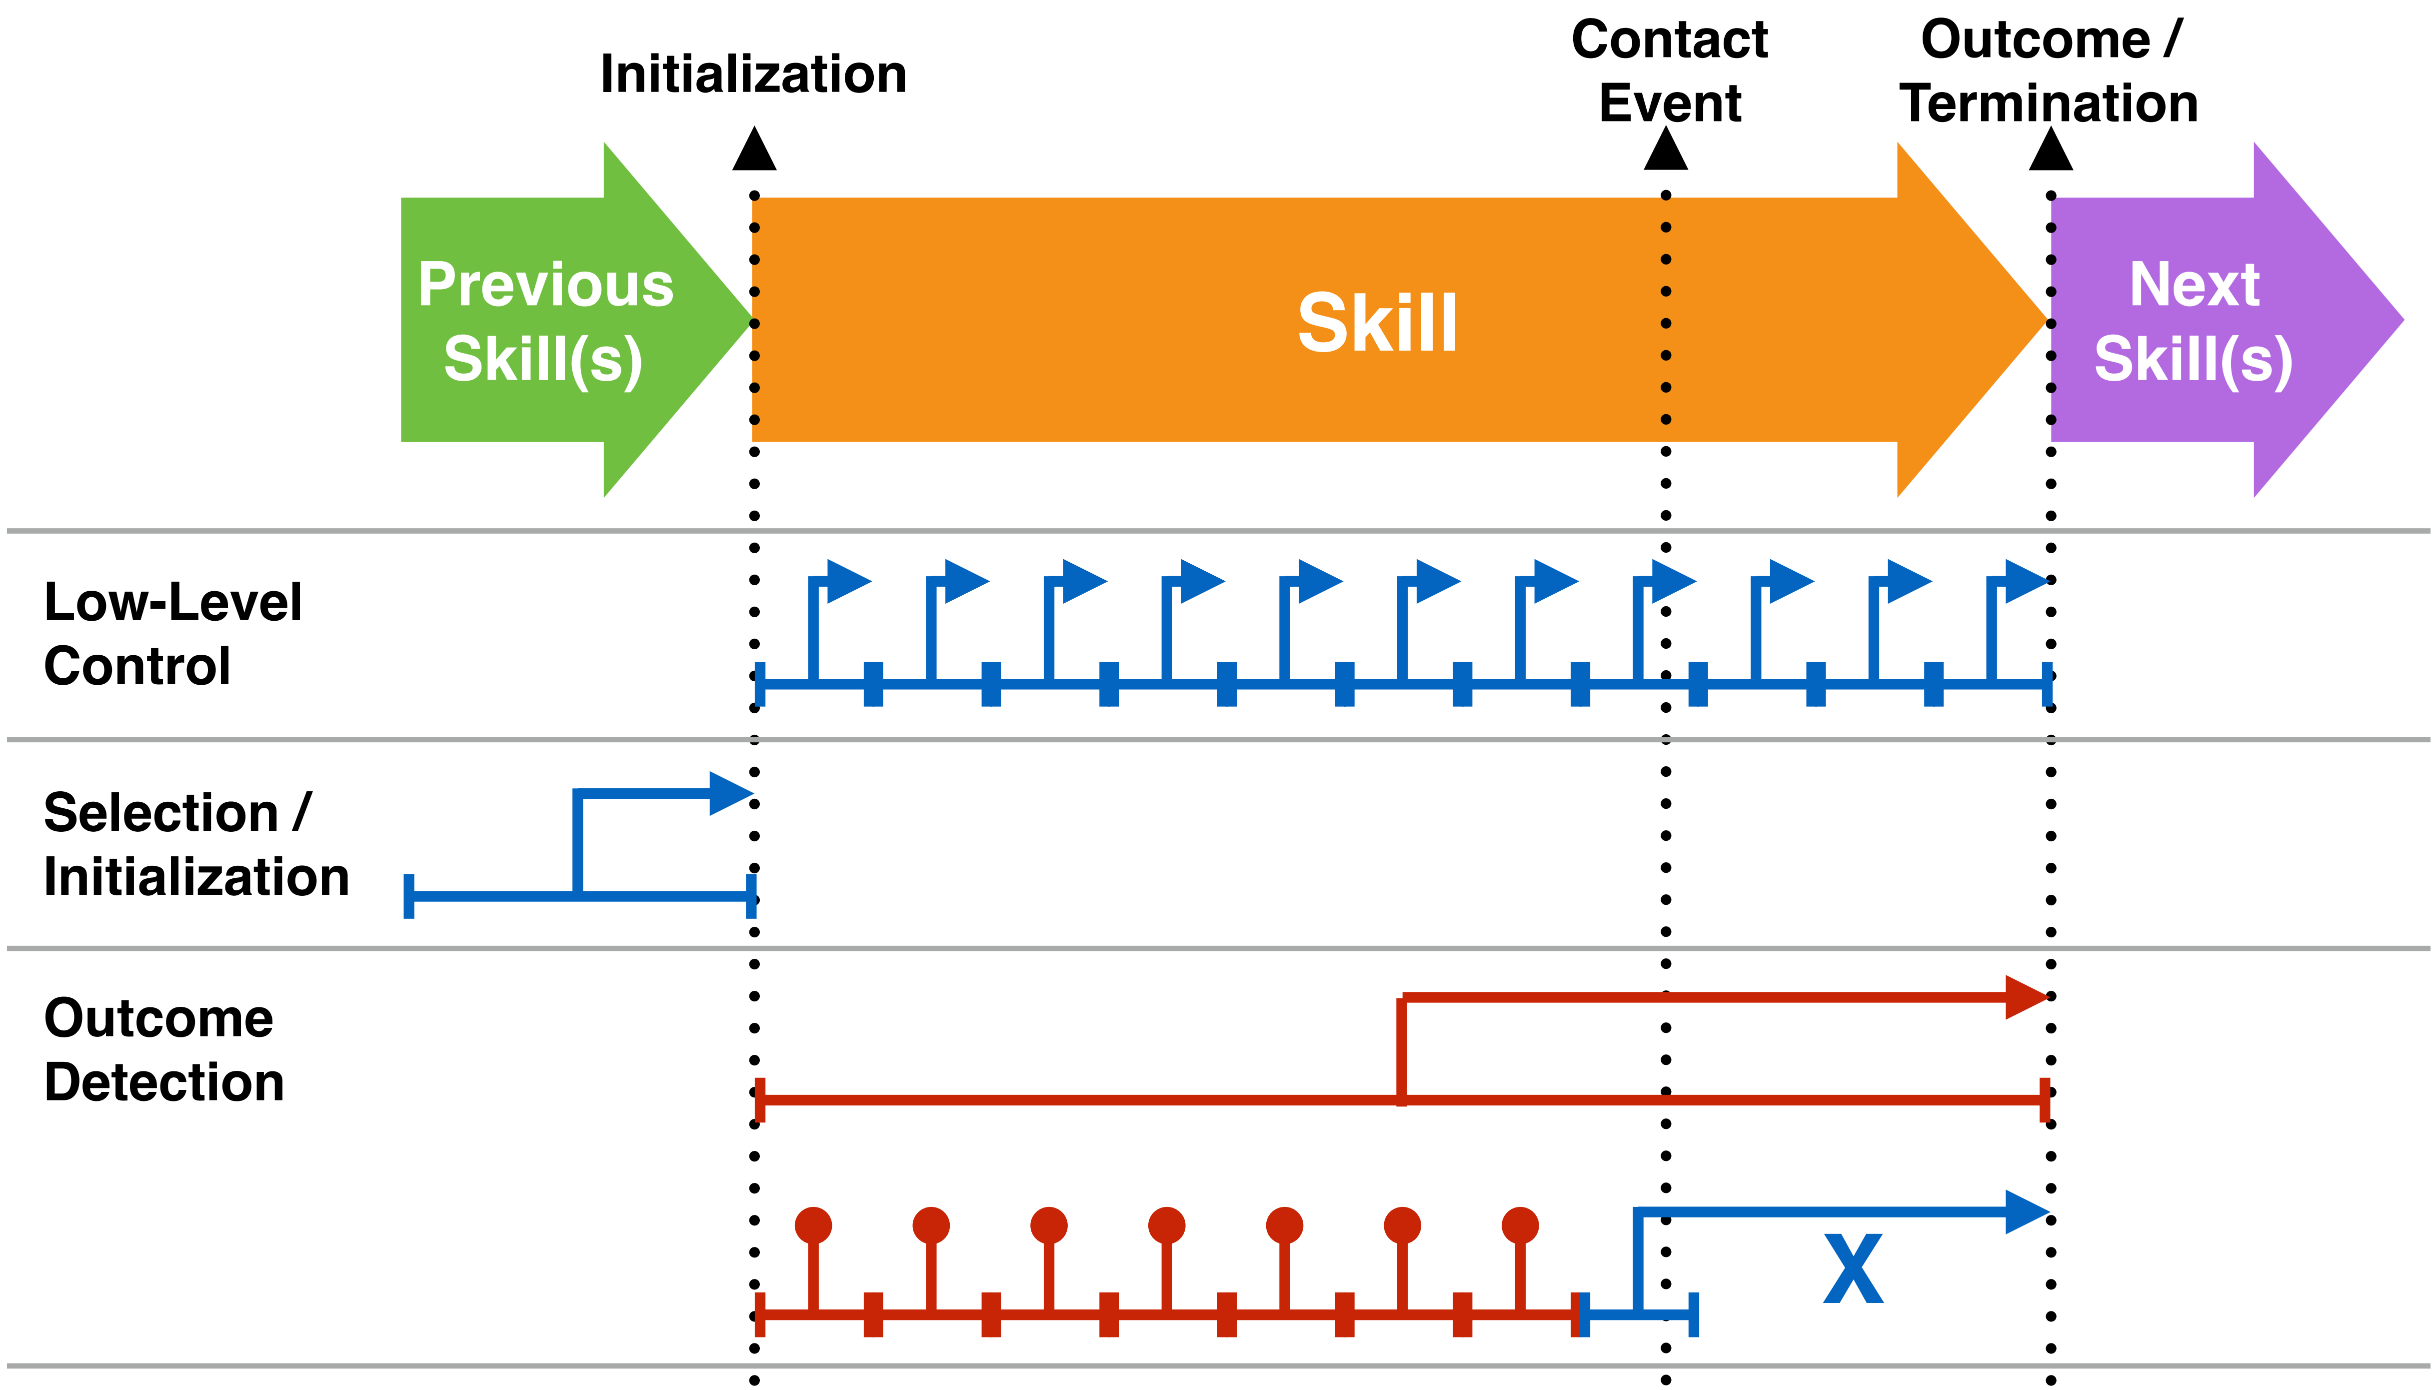

Outcome Detection

- Skill executions often terminate in salient sensory events

Test Case: Guided Peg-in-Hole

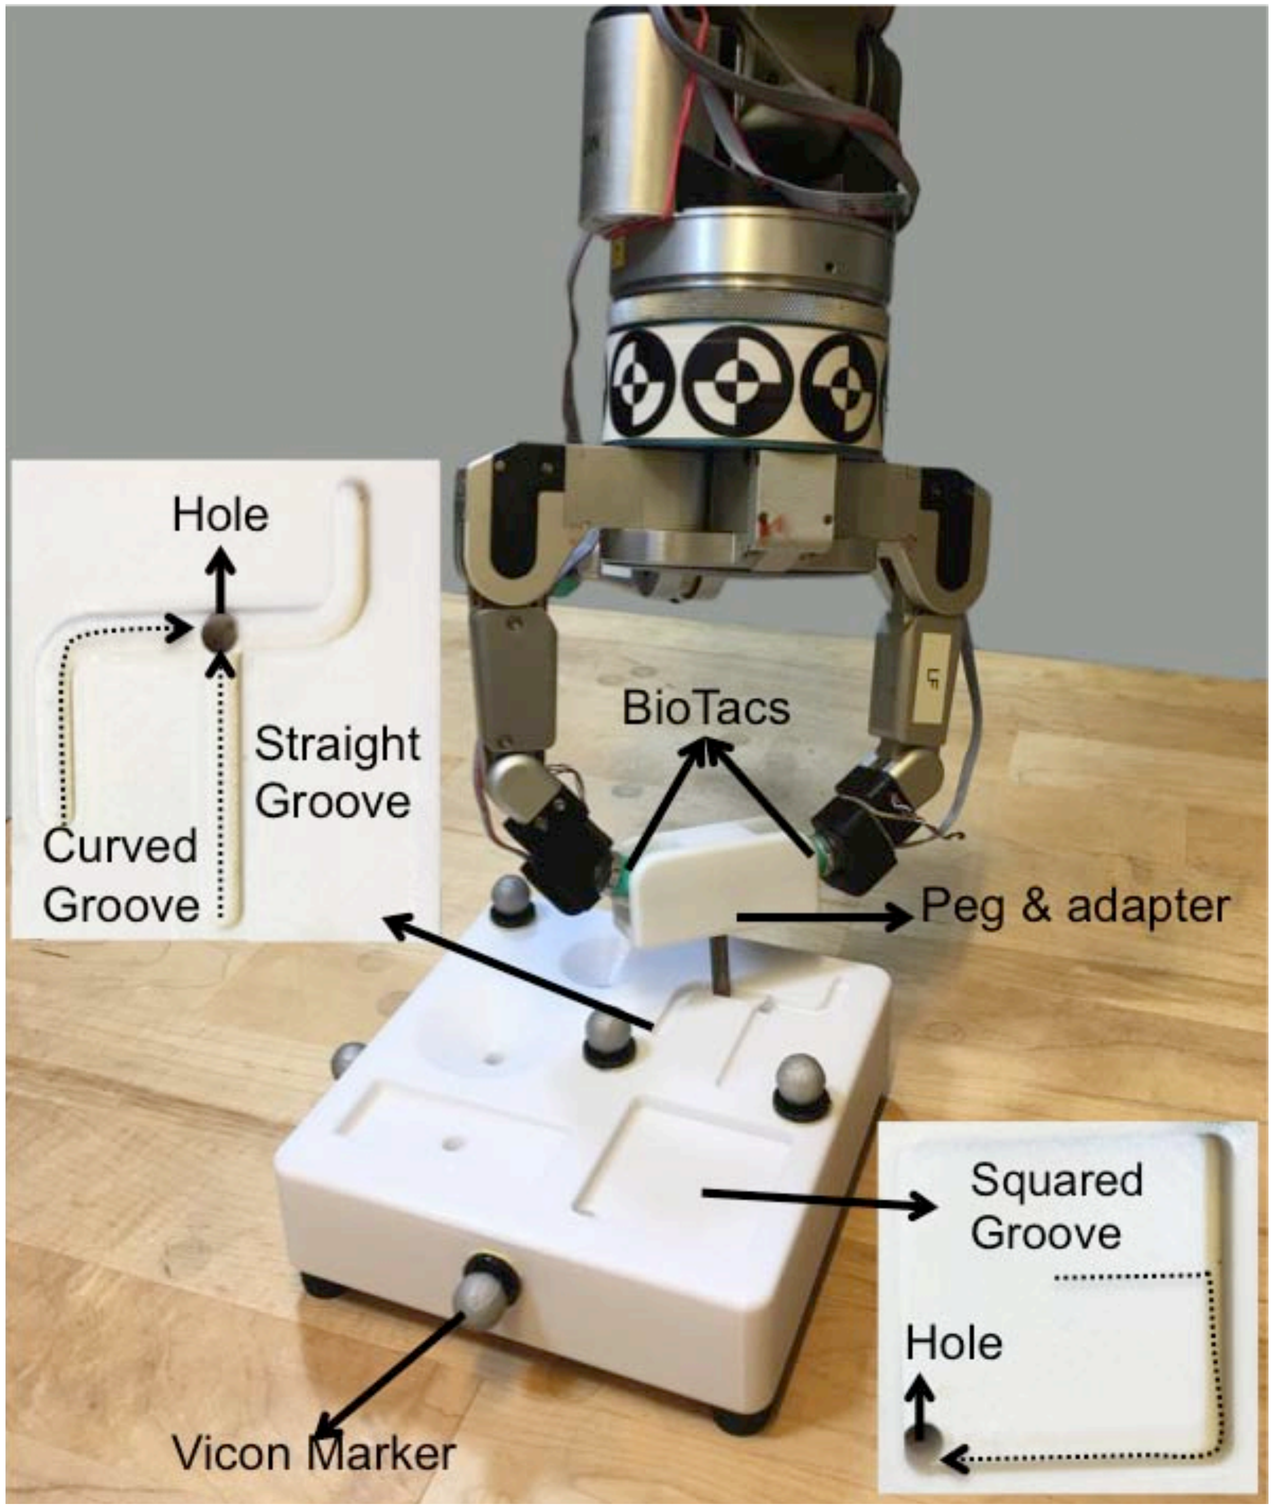

Z. Su, O. Kroemer, G. Loeb, G. S. Sukhatme and S. Schaal. "Learning to Switch between Sensorimotor Primitives using Multimodal Haptic Signals" SAB, 2016

Skill Demonstration

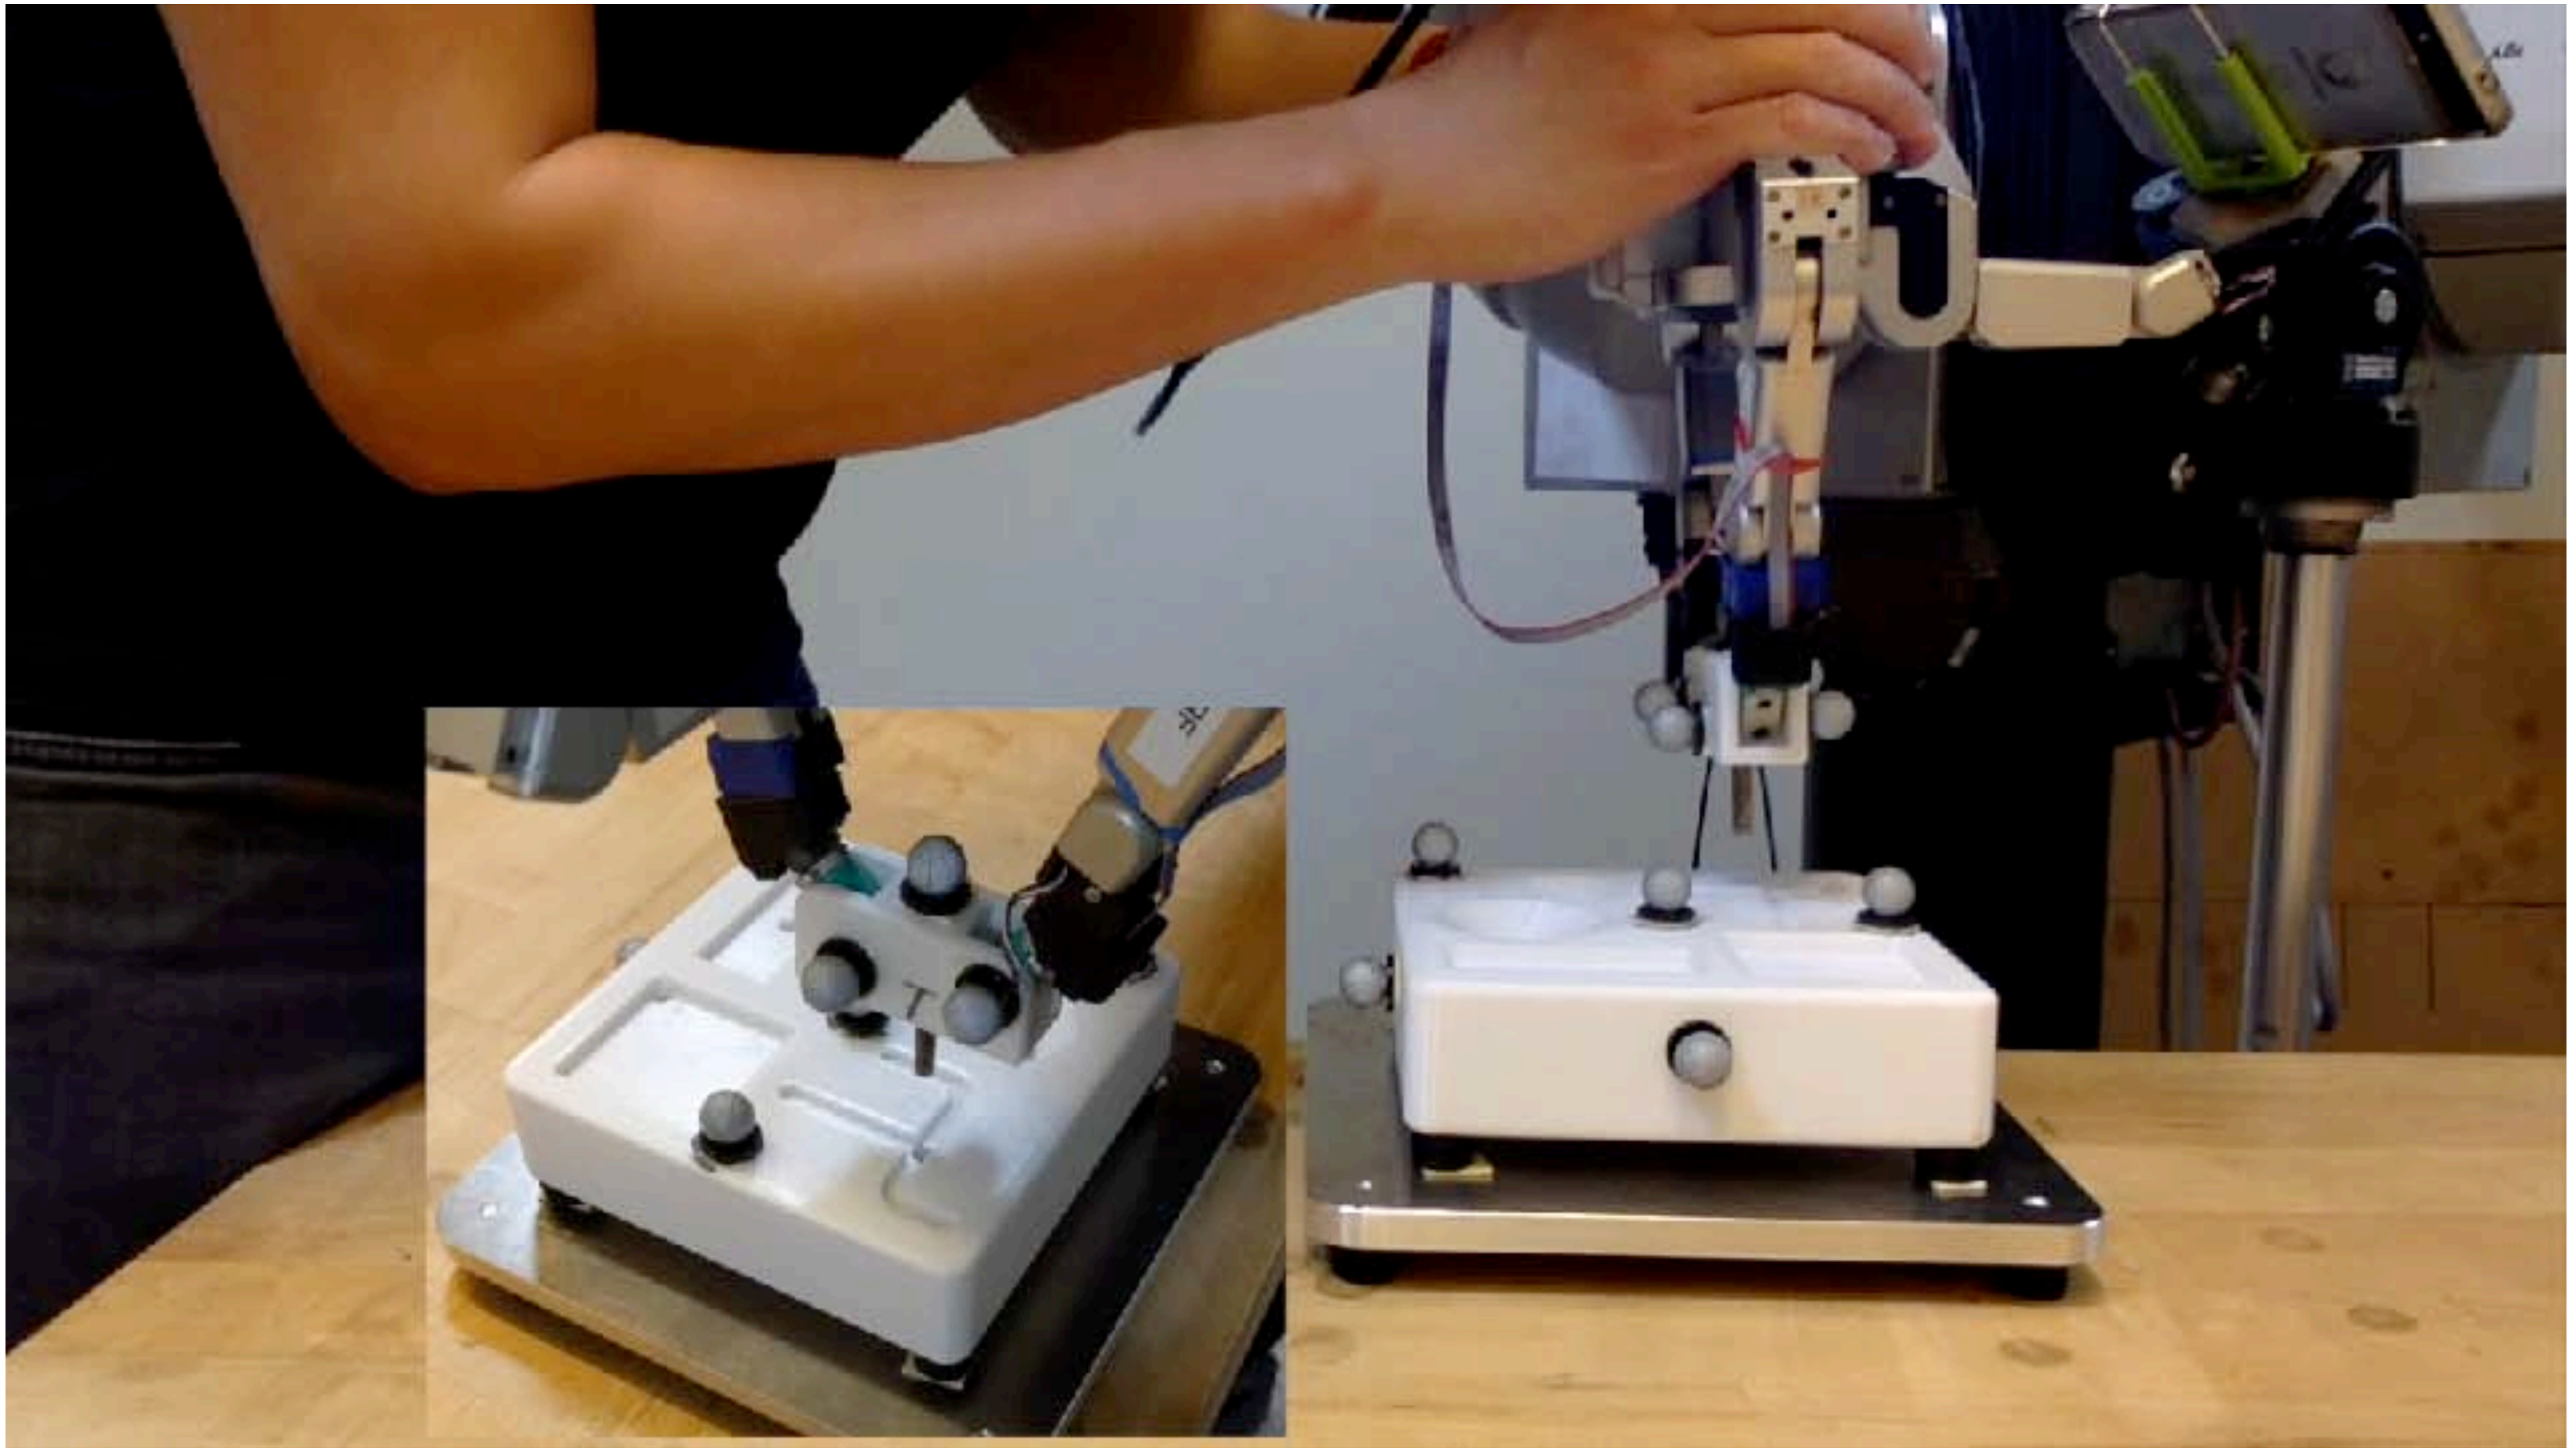

Segmentation and Goal Detection

- Want to learn skills that terminate in salient events

- Segment demos into skills using change point detection [Adams and McKay 2007]
  - Use position and tactile sensor signals for segmentation

- Each skill defines a desired goal state

- Learn a goal detector using ST-HMP features and SVMs [Bo et al. 2011, Madry et al. 2014]
  - Detect goals using low and high frequency tactile signals

Skill Executions

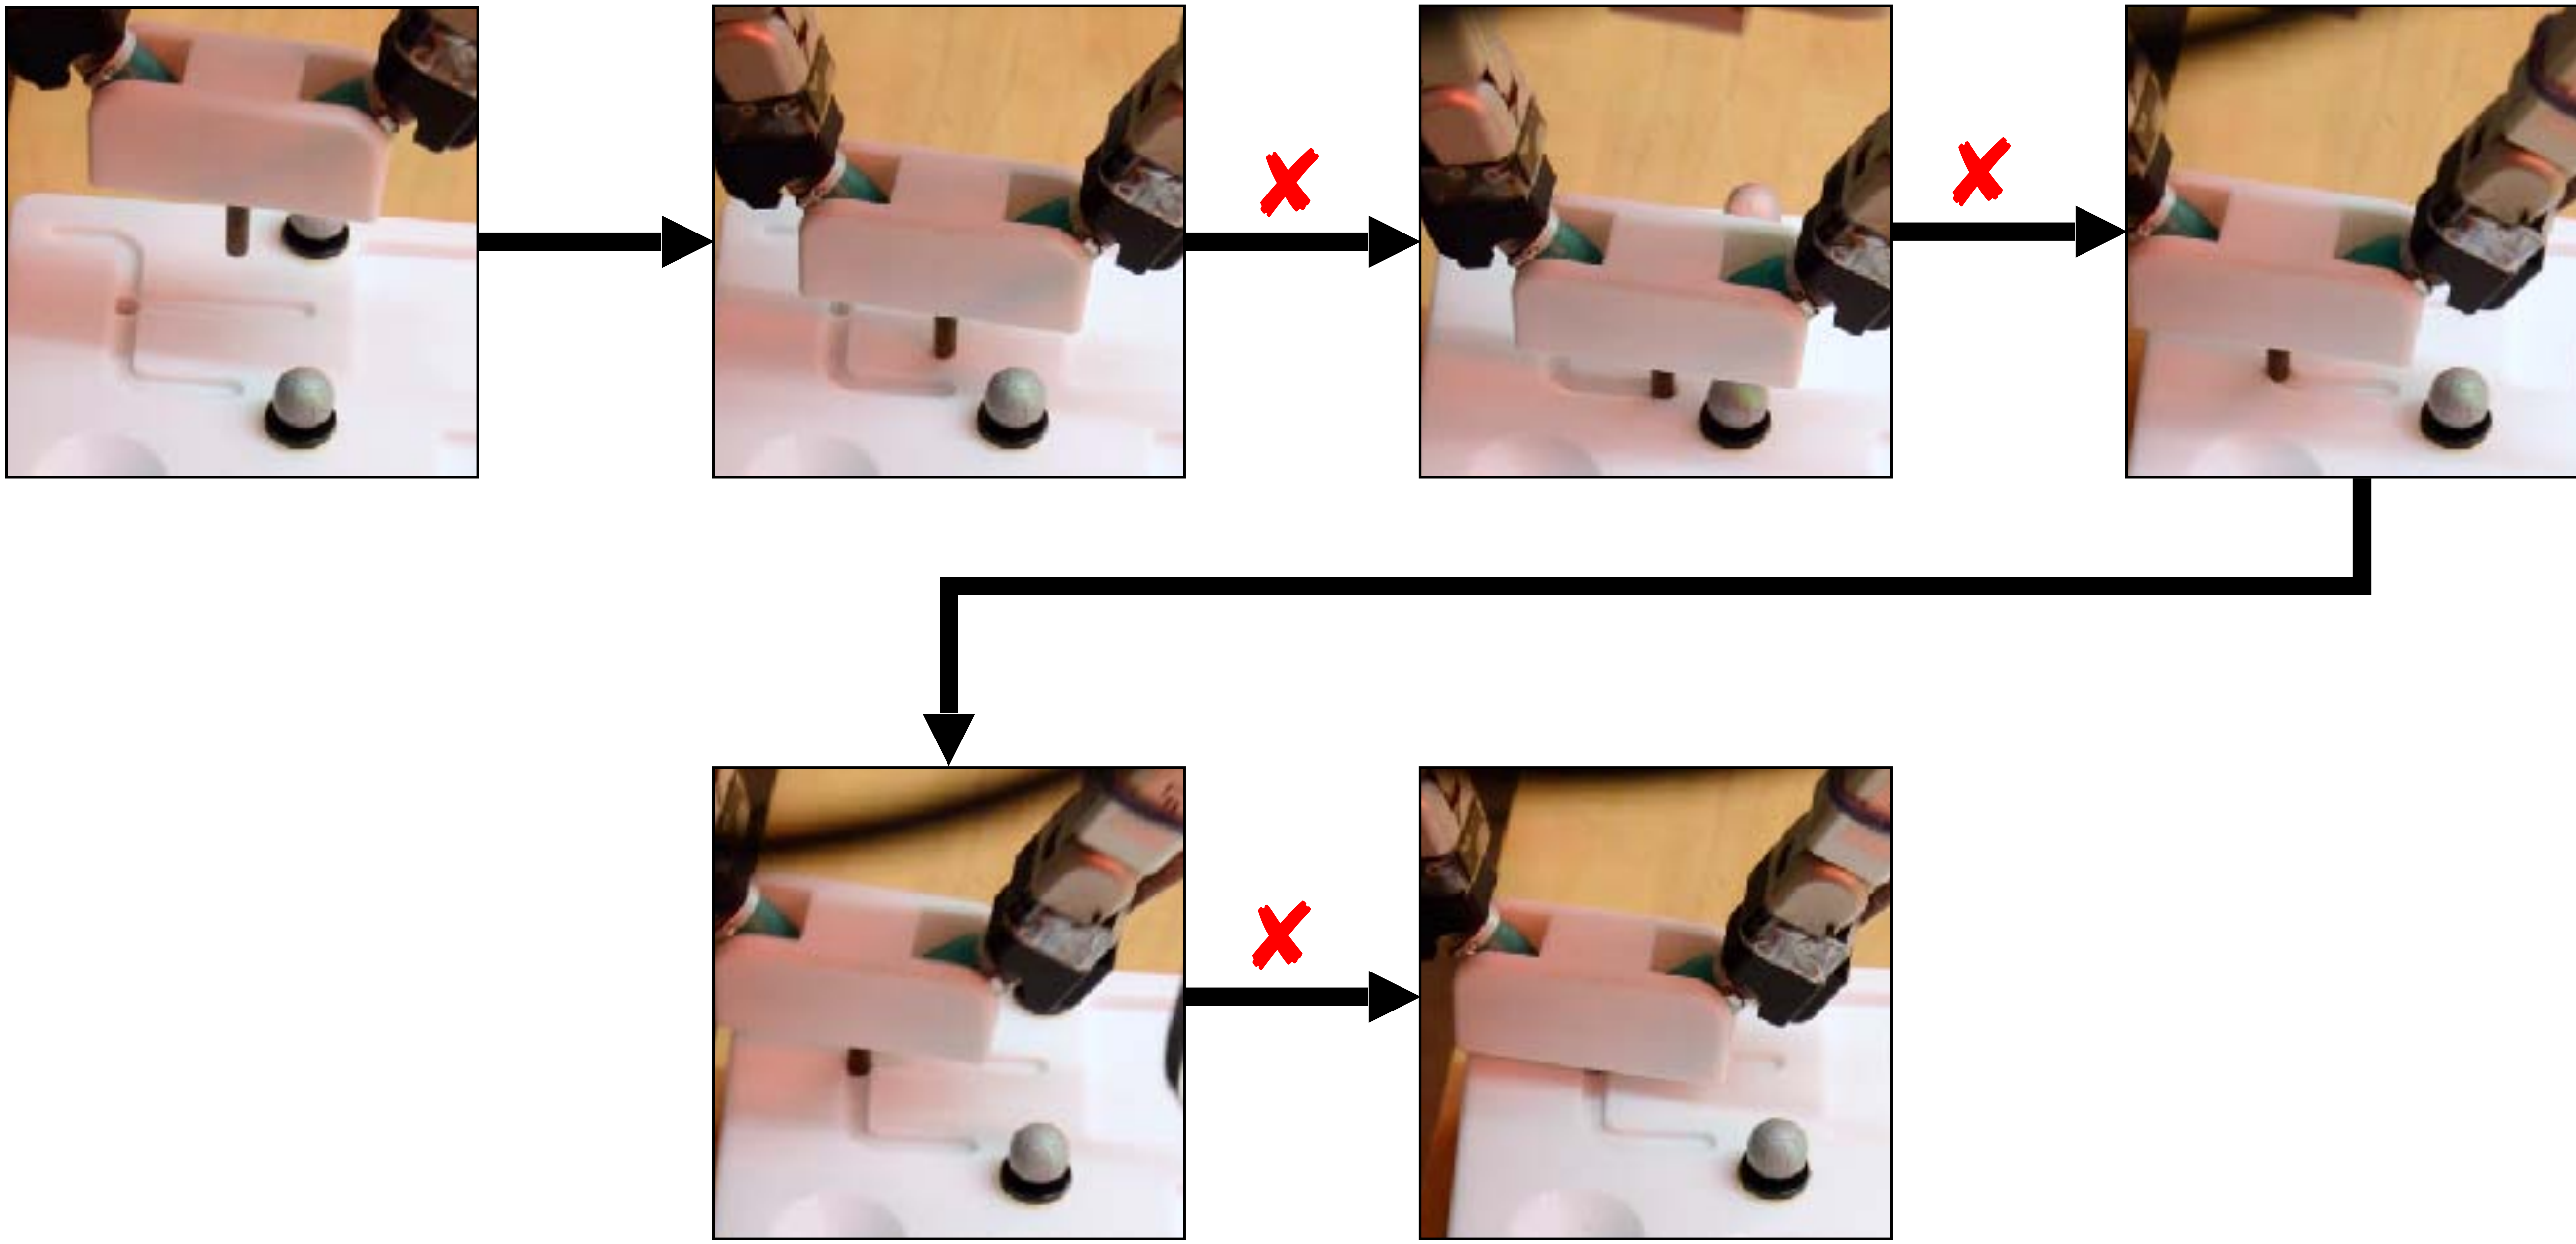

Reattempt motion from current location with slightly higher force

Failure to Reach the Groove

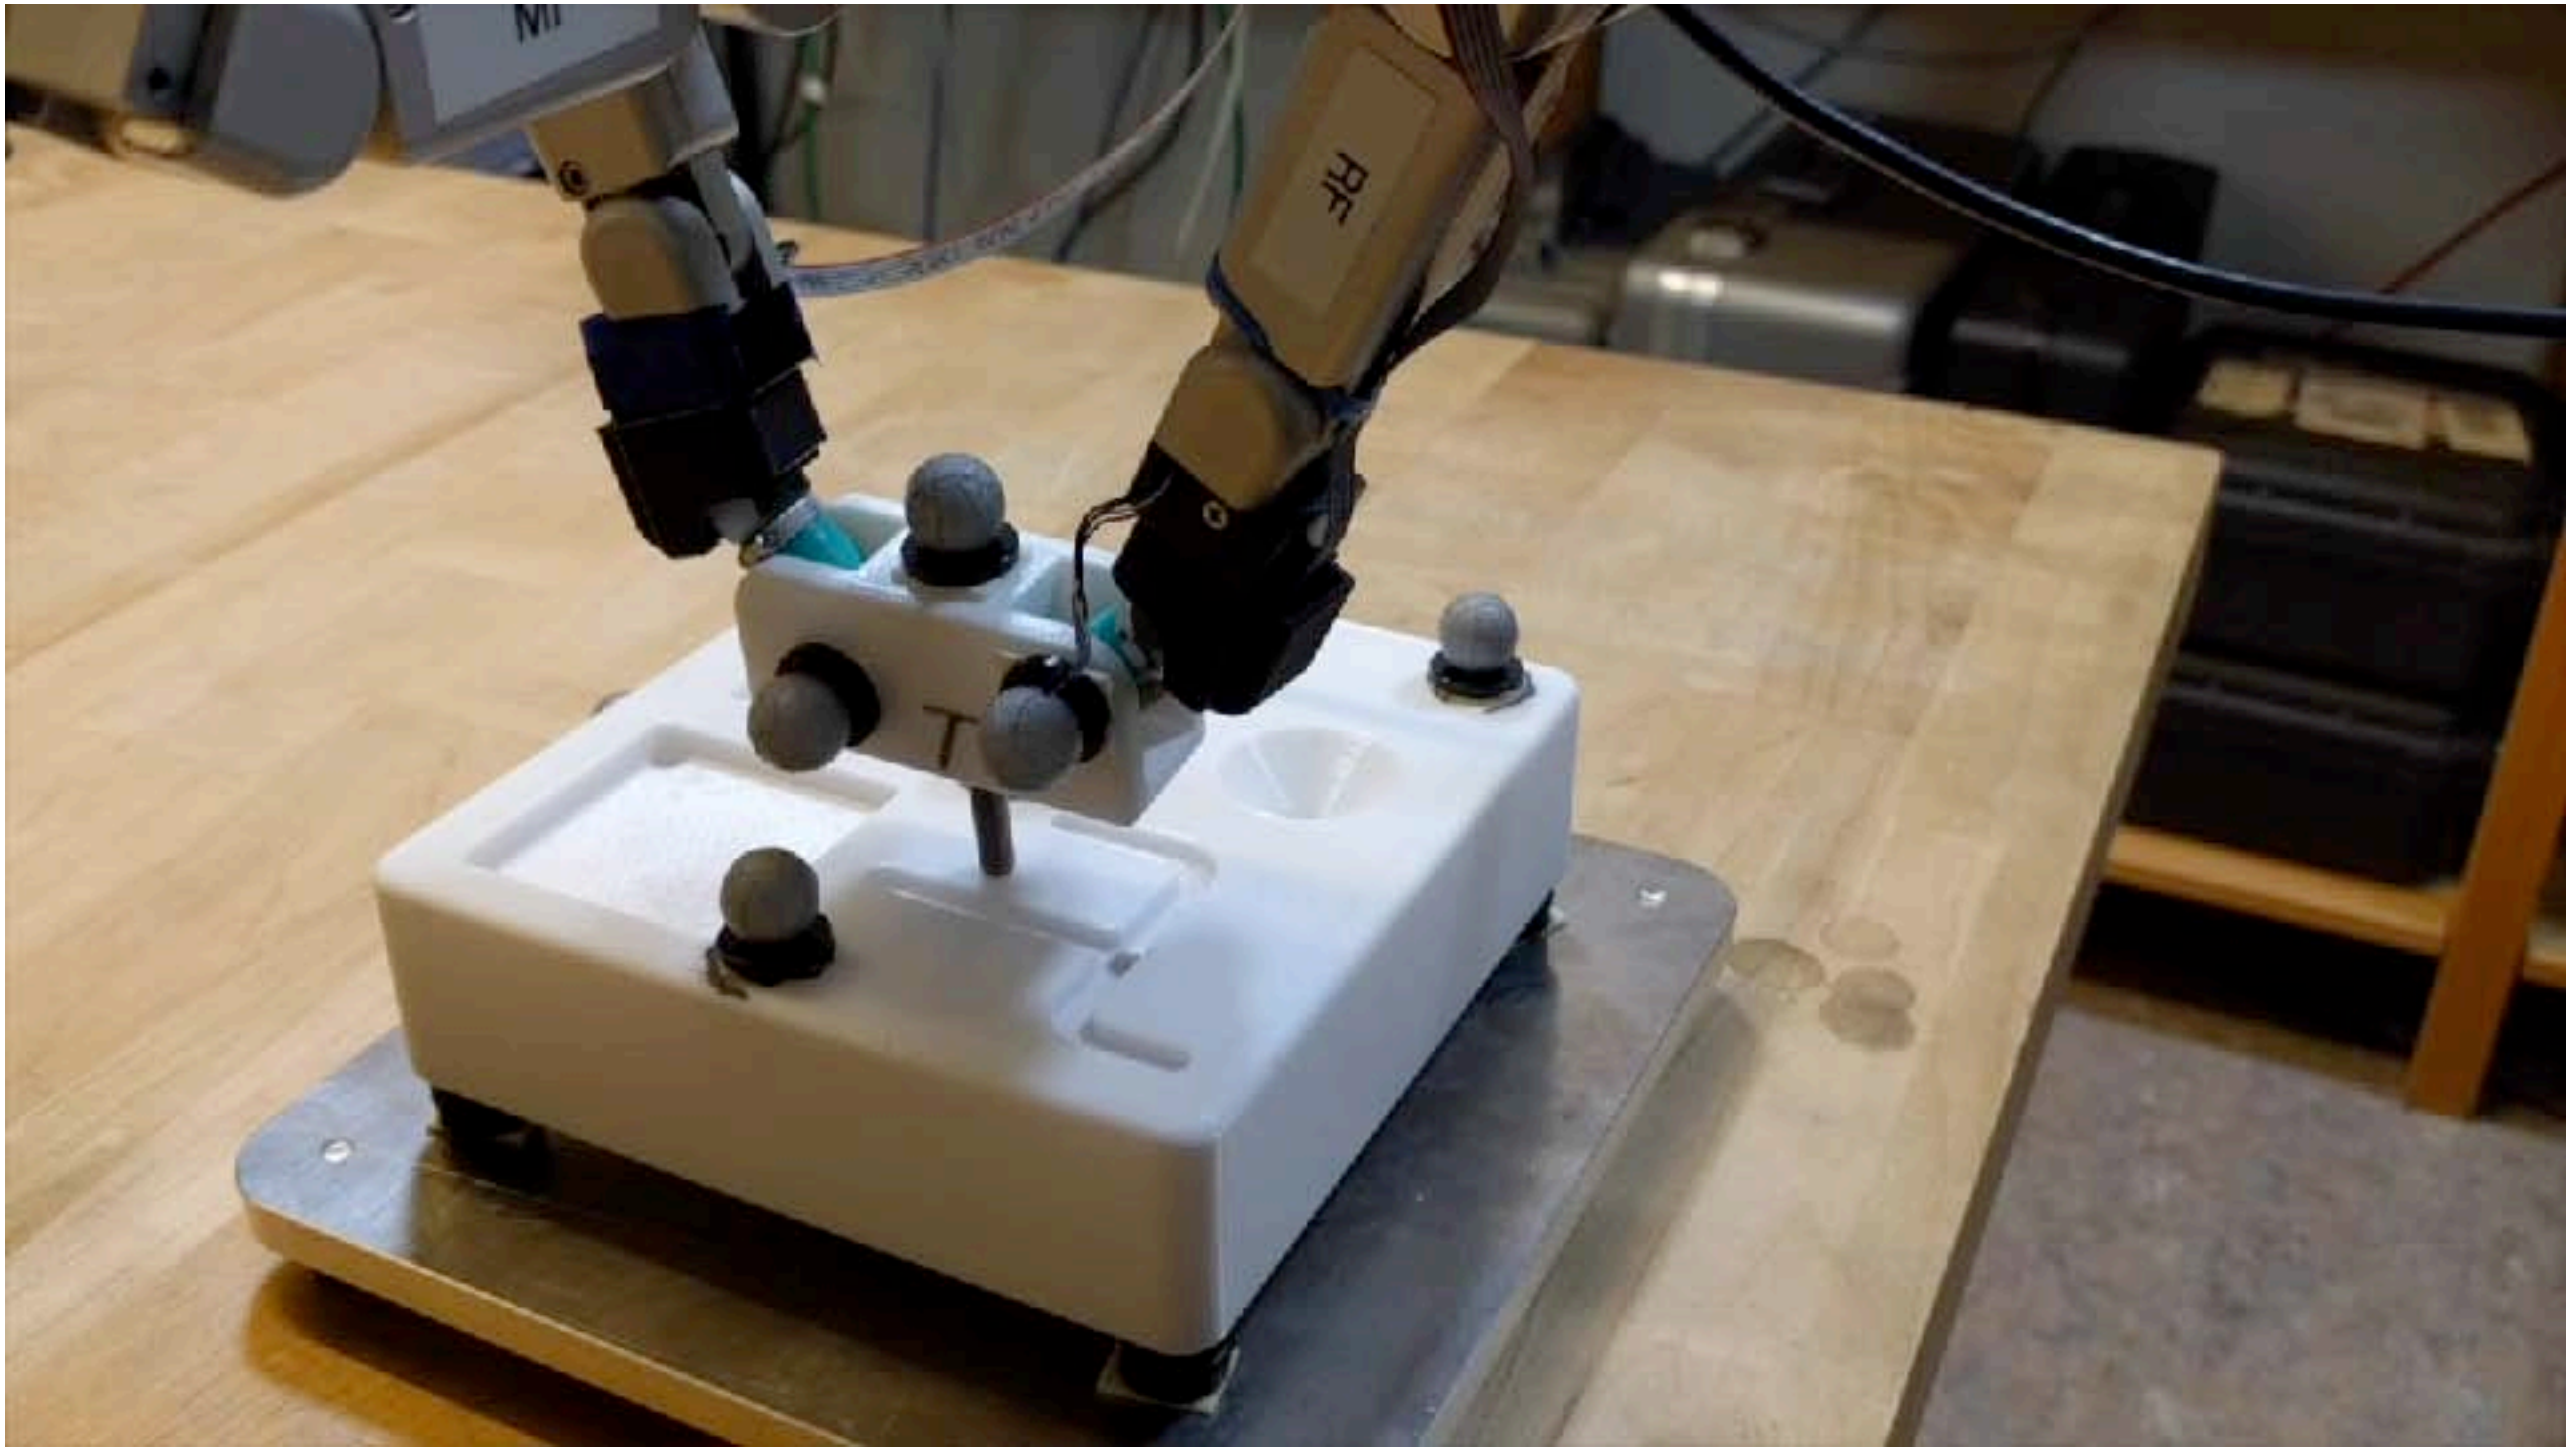

Detection and Correction

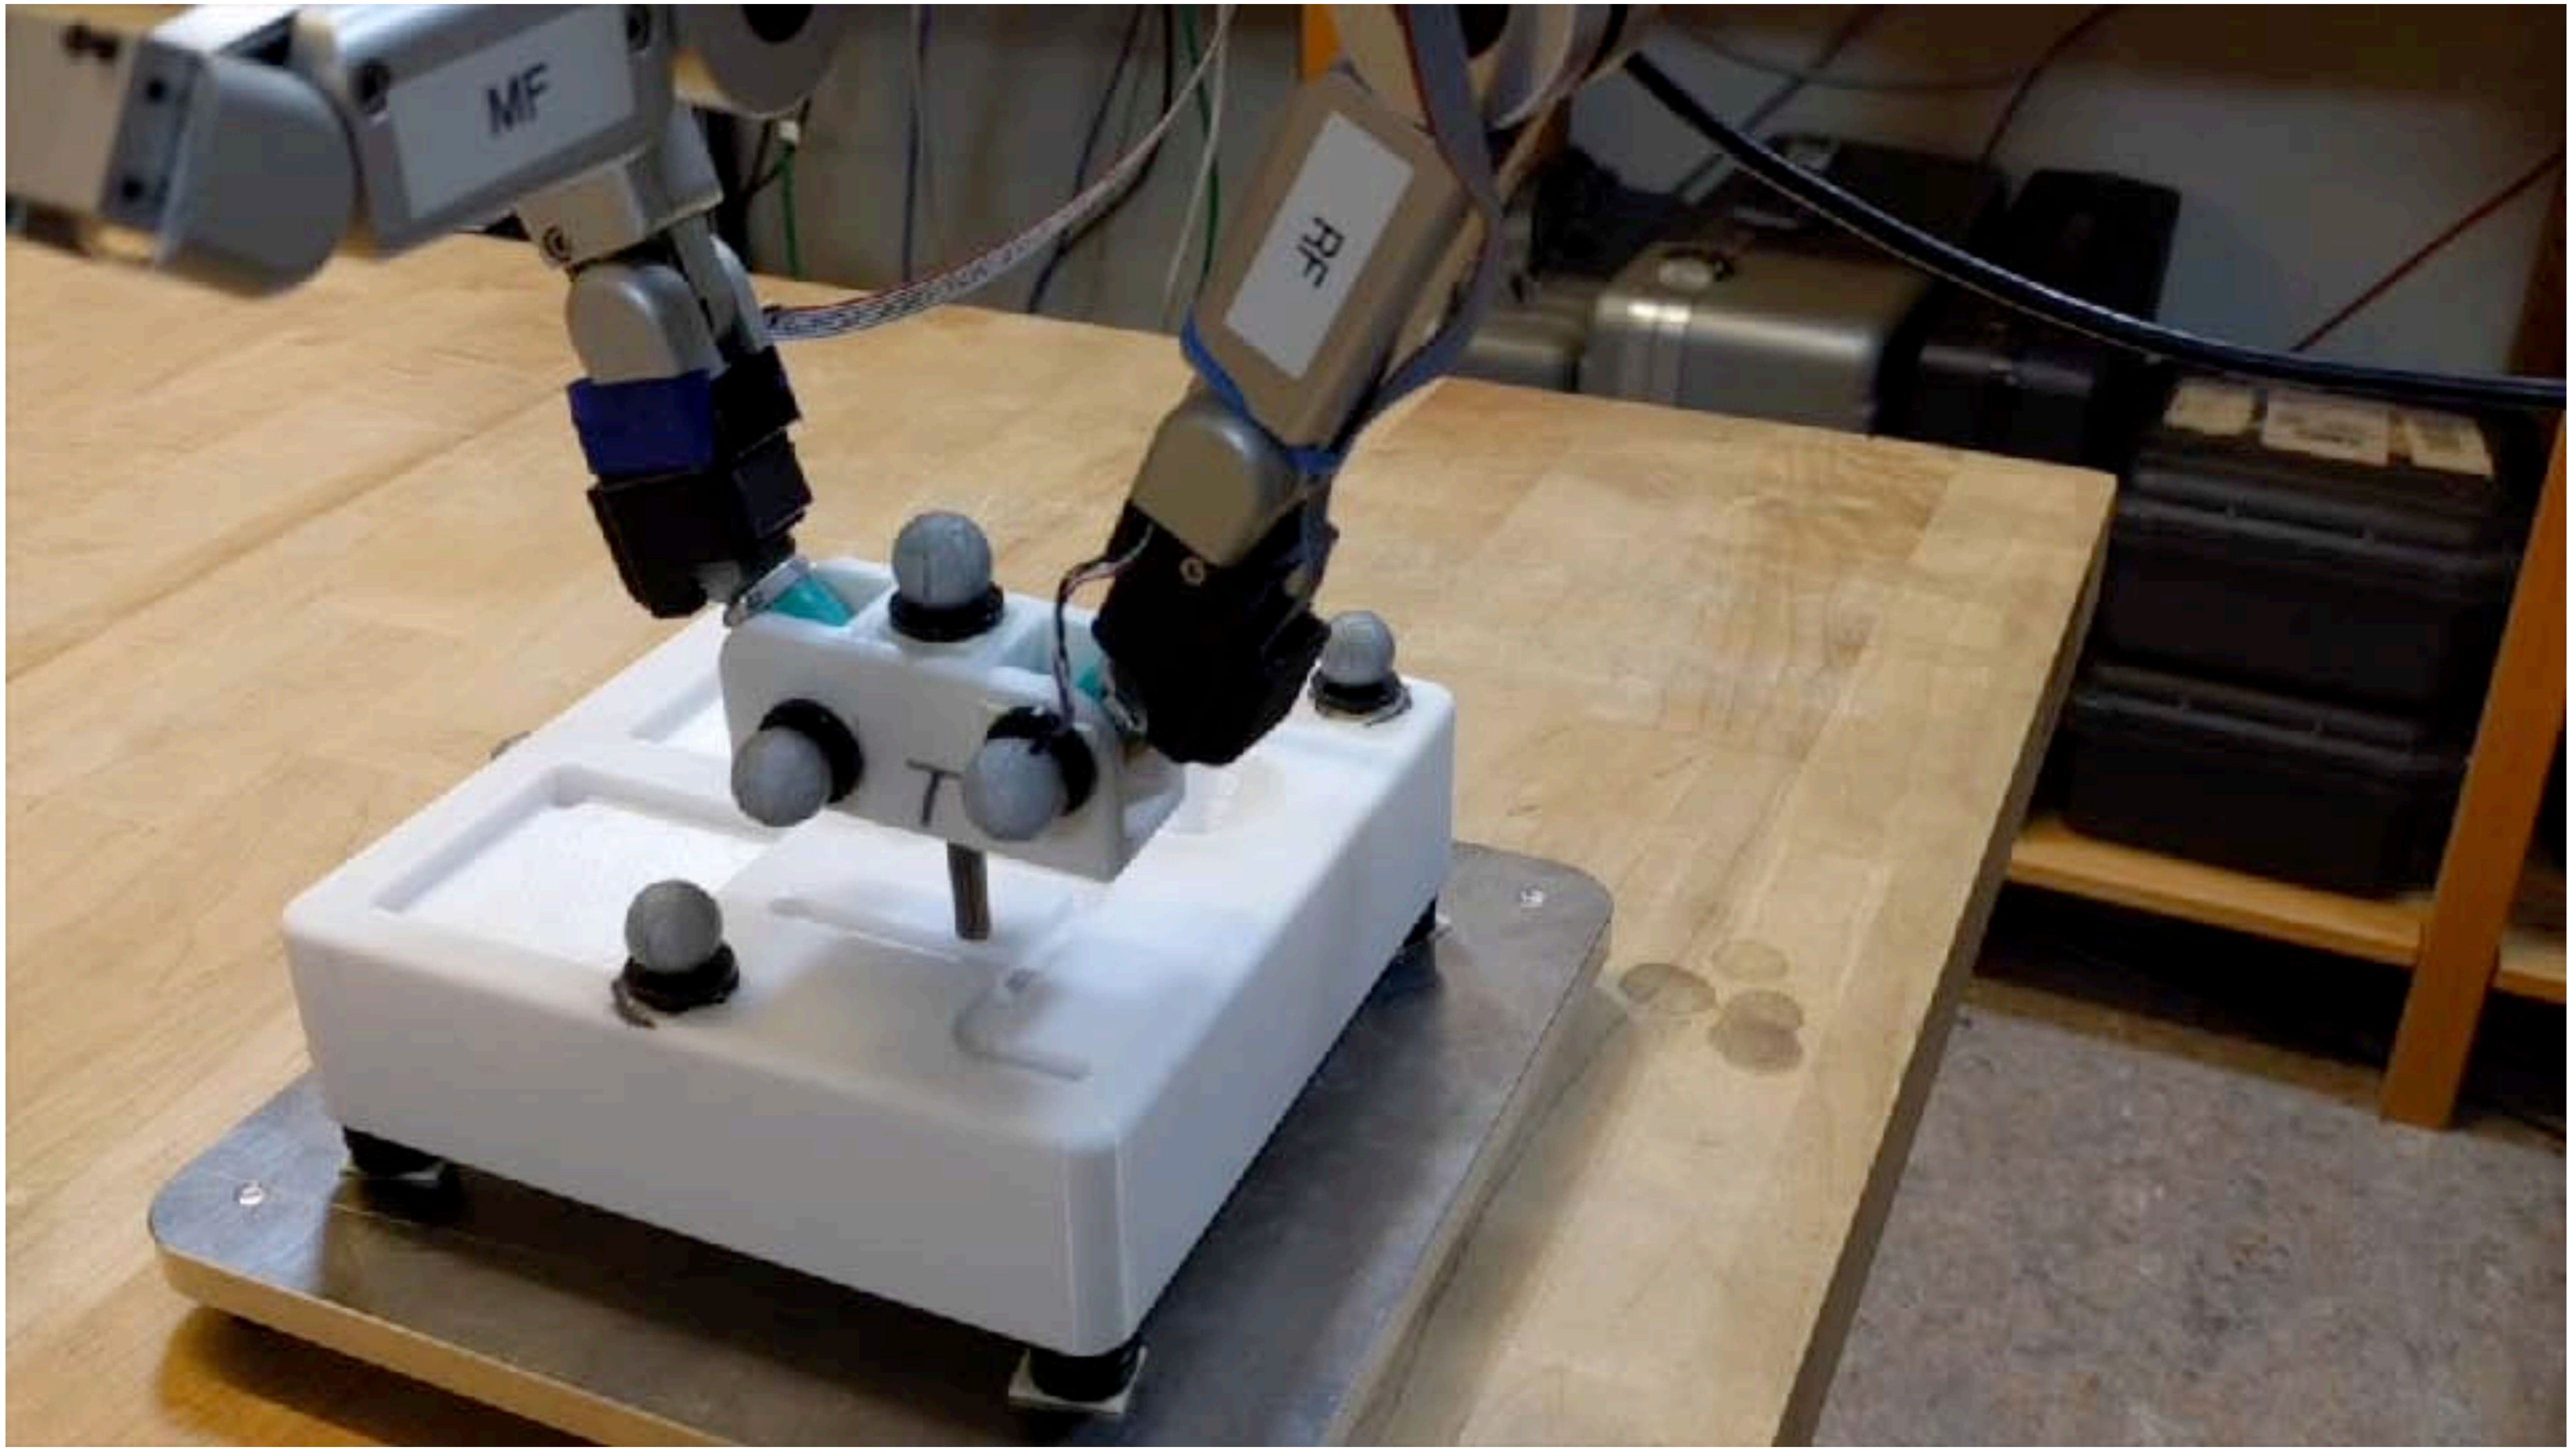

Failure to Reach the Corner

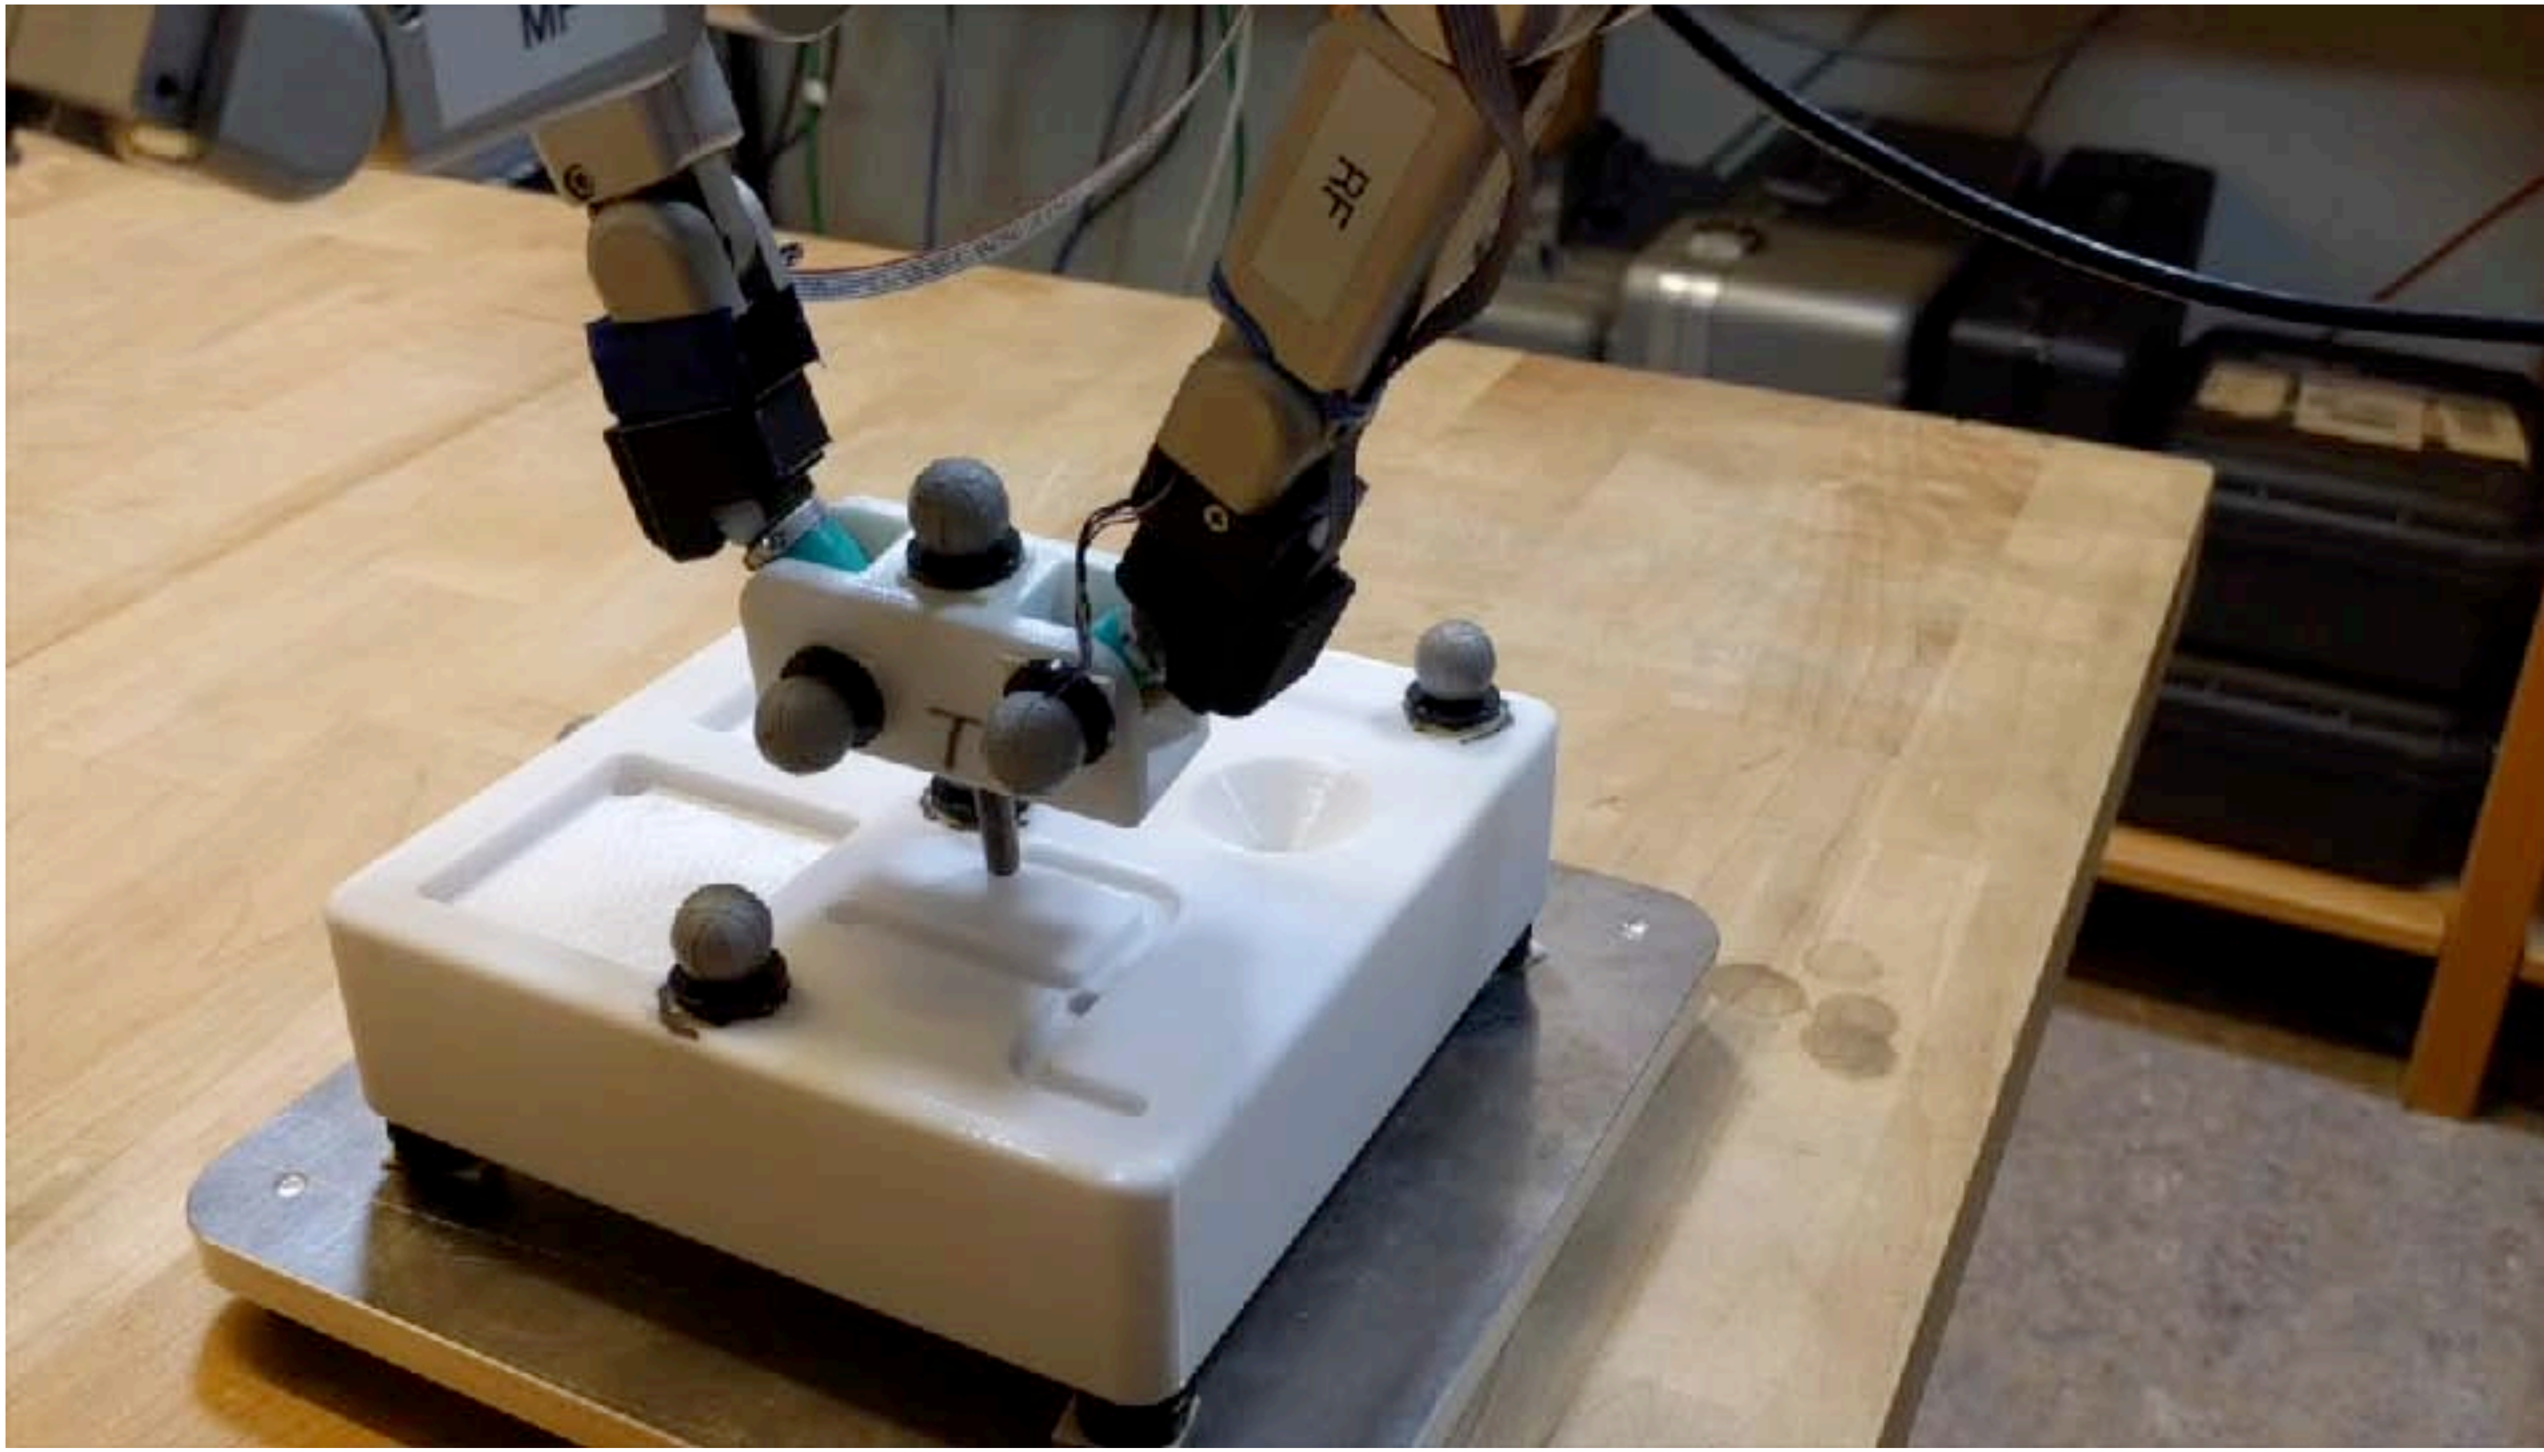

Detection and Correction

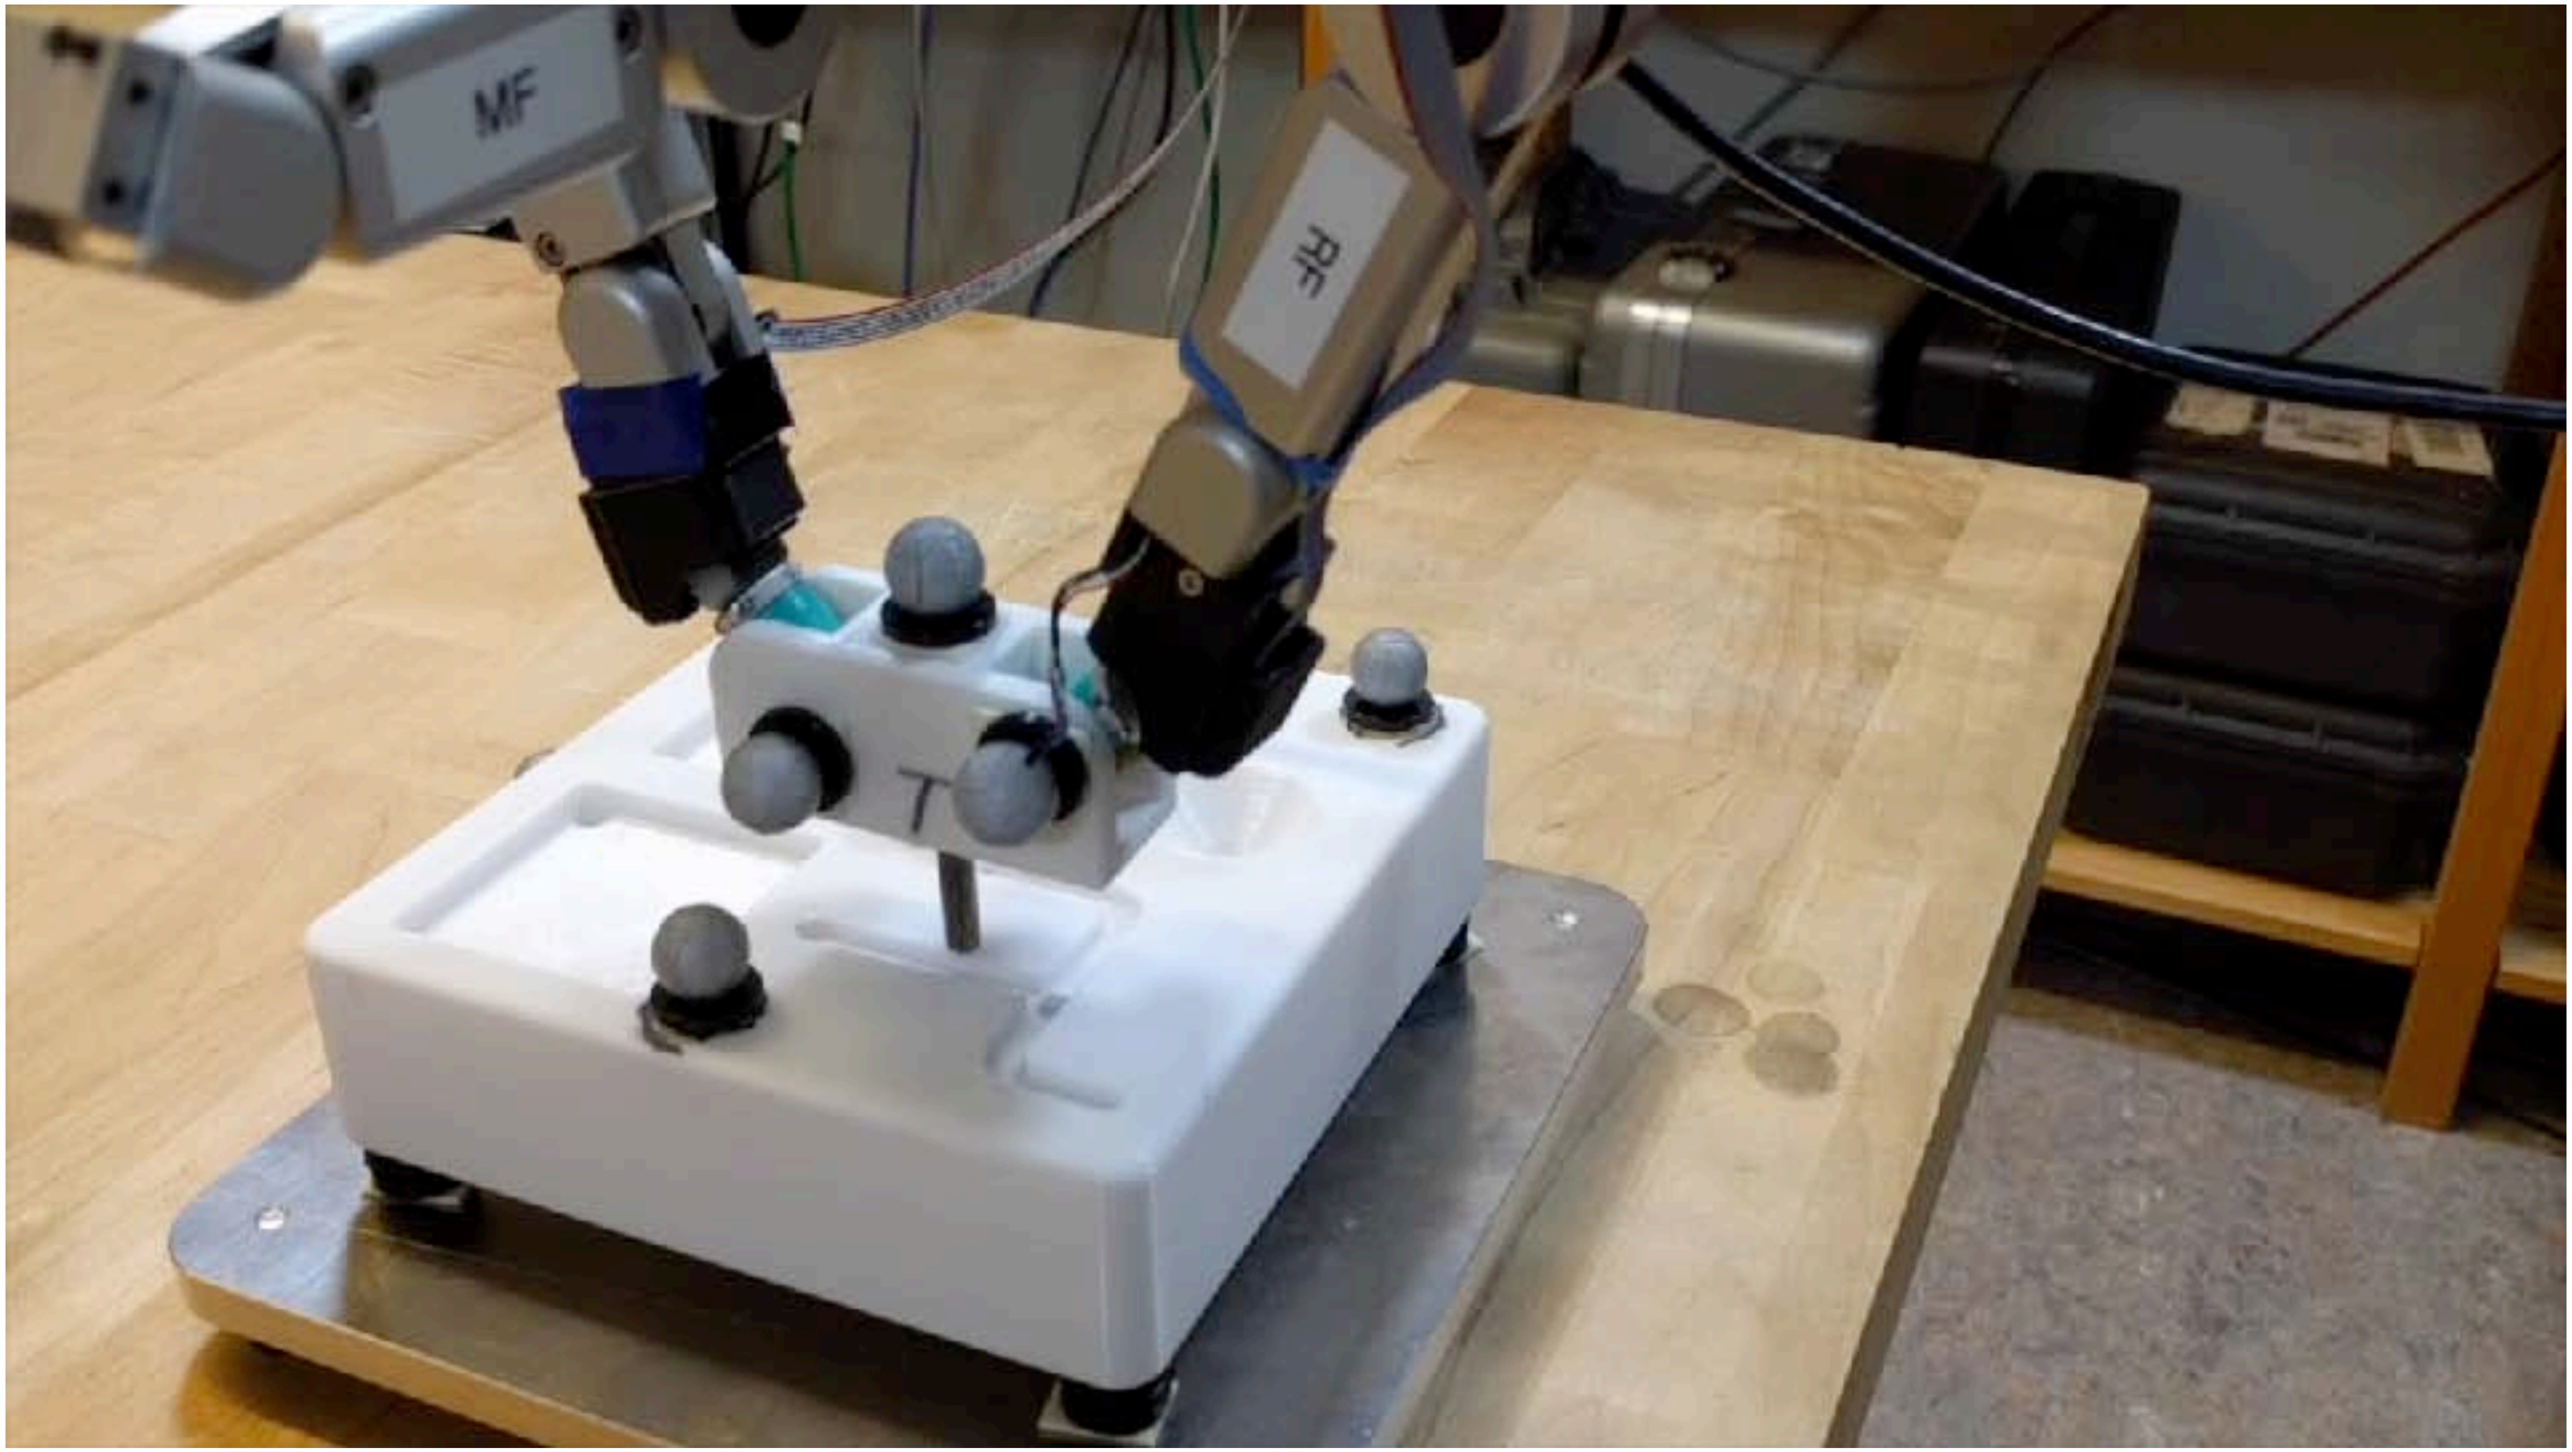

Skill Executions

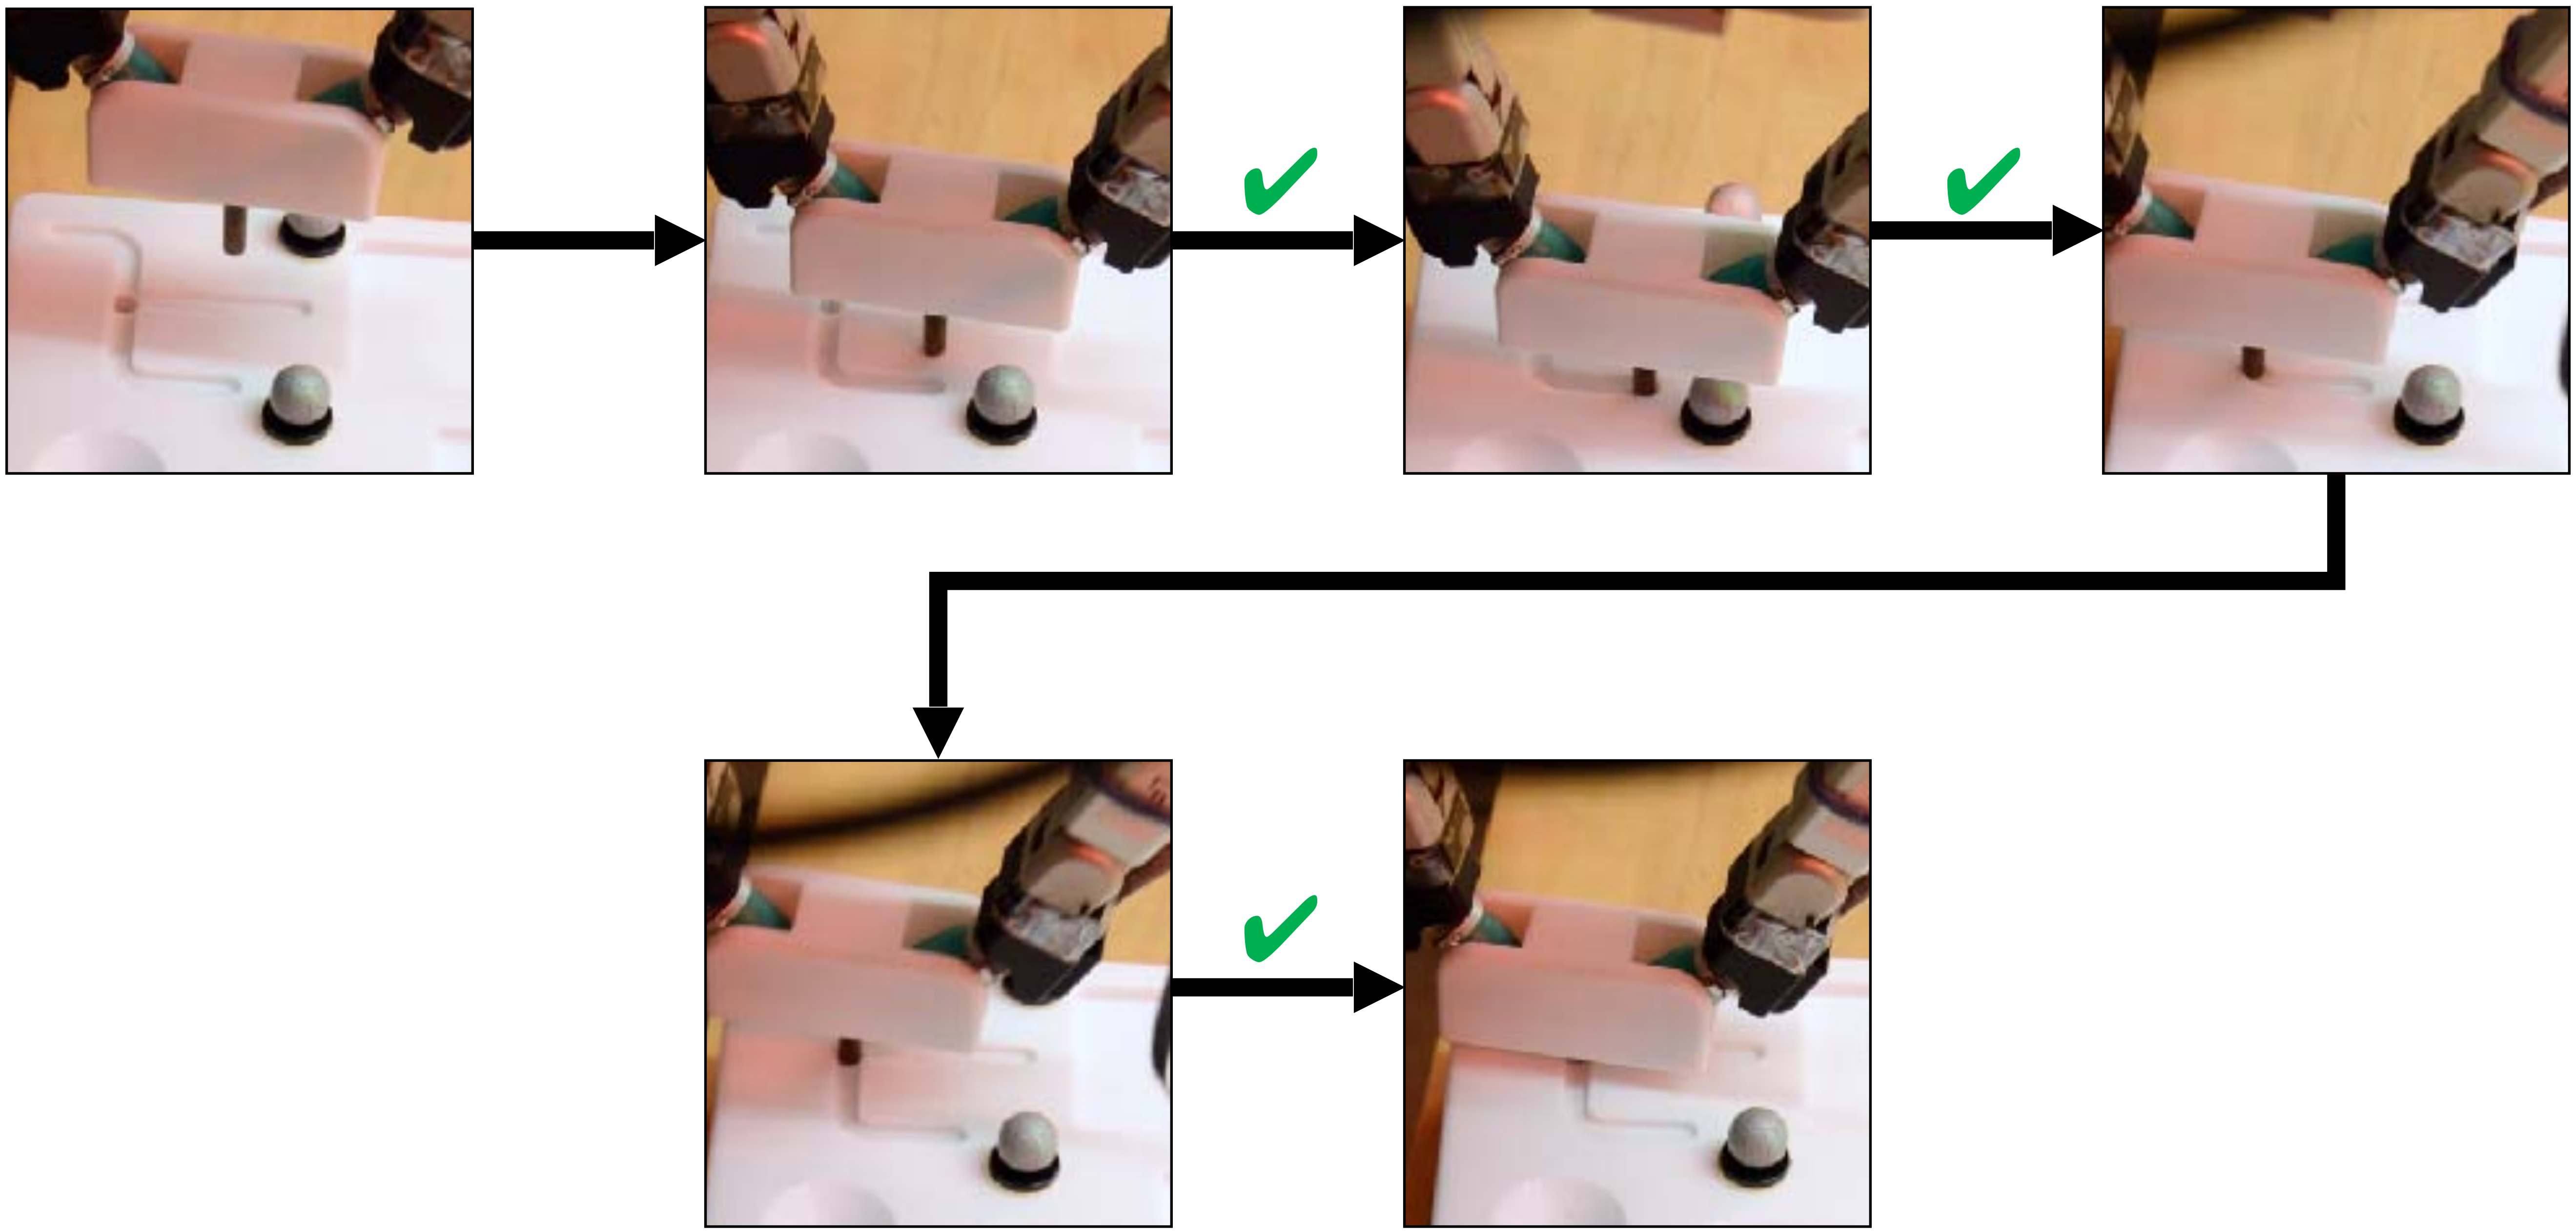

Reattempt motion from current location with slightly higher force

Sensory Feedback

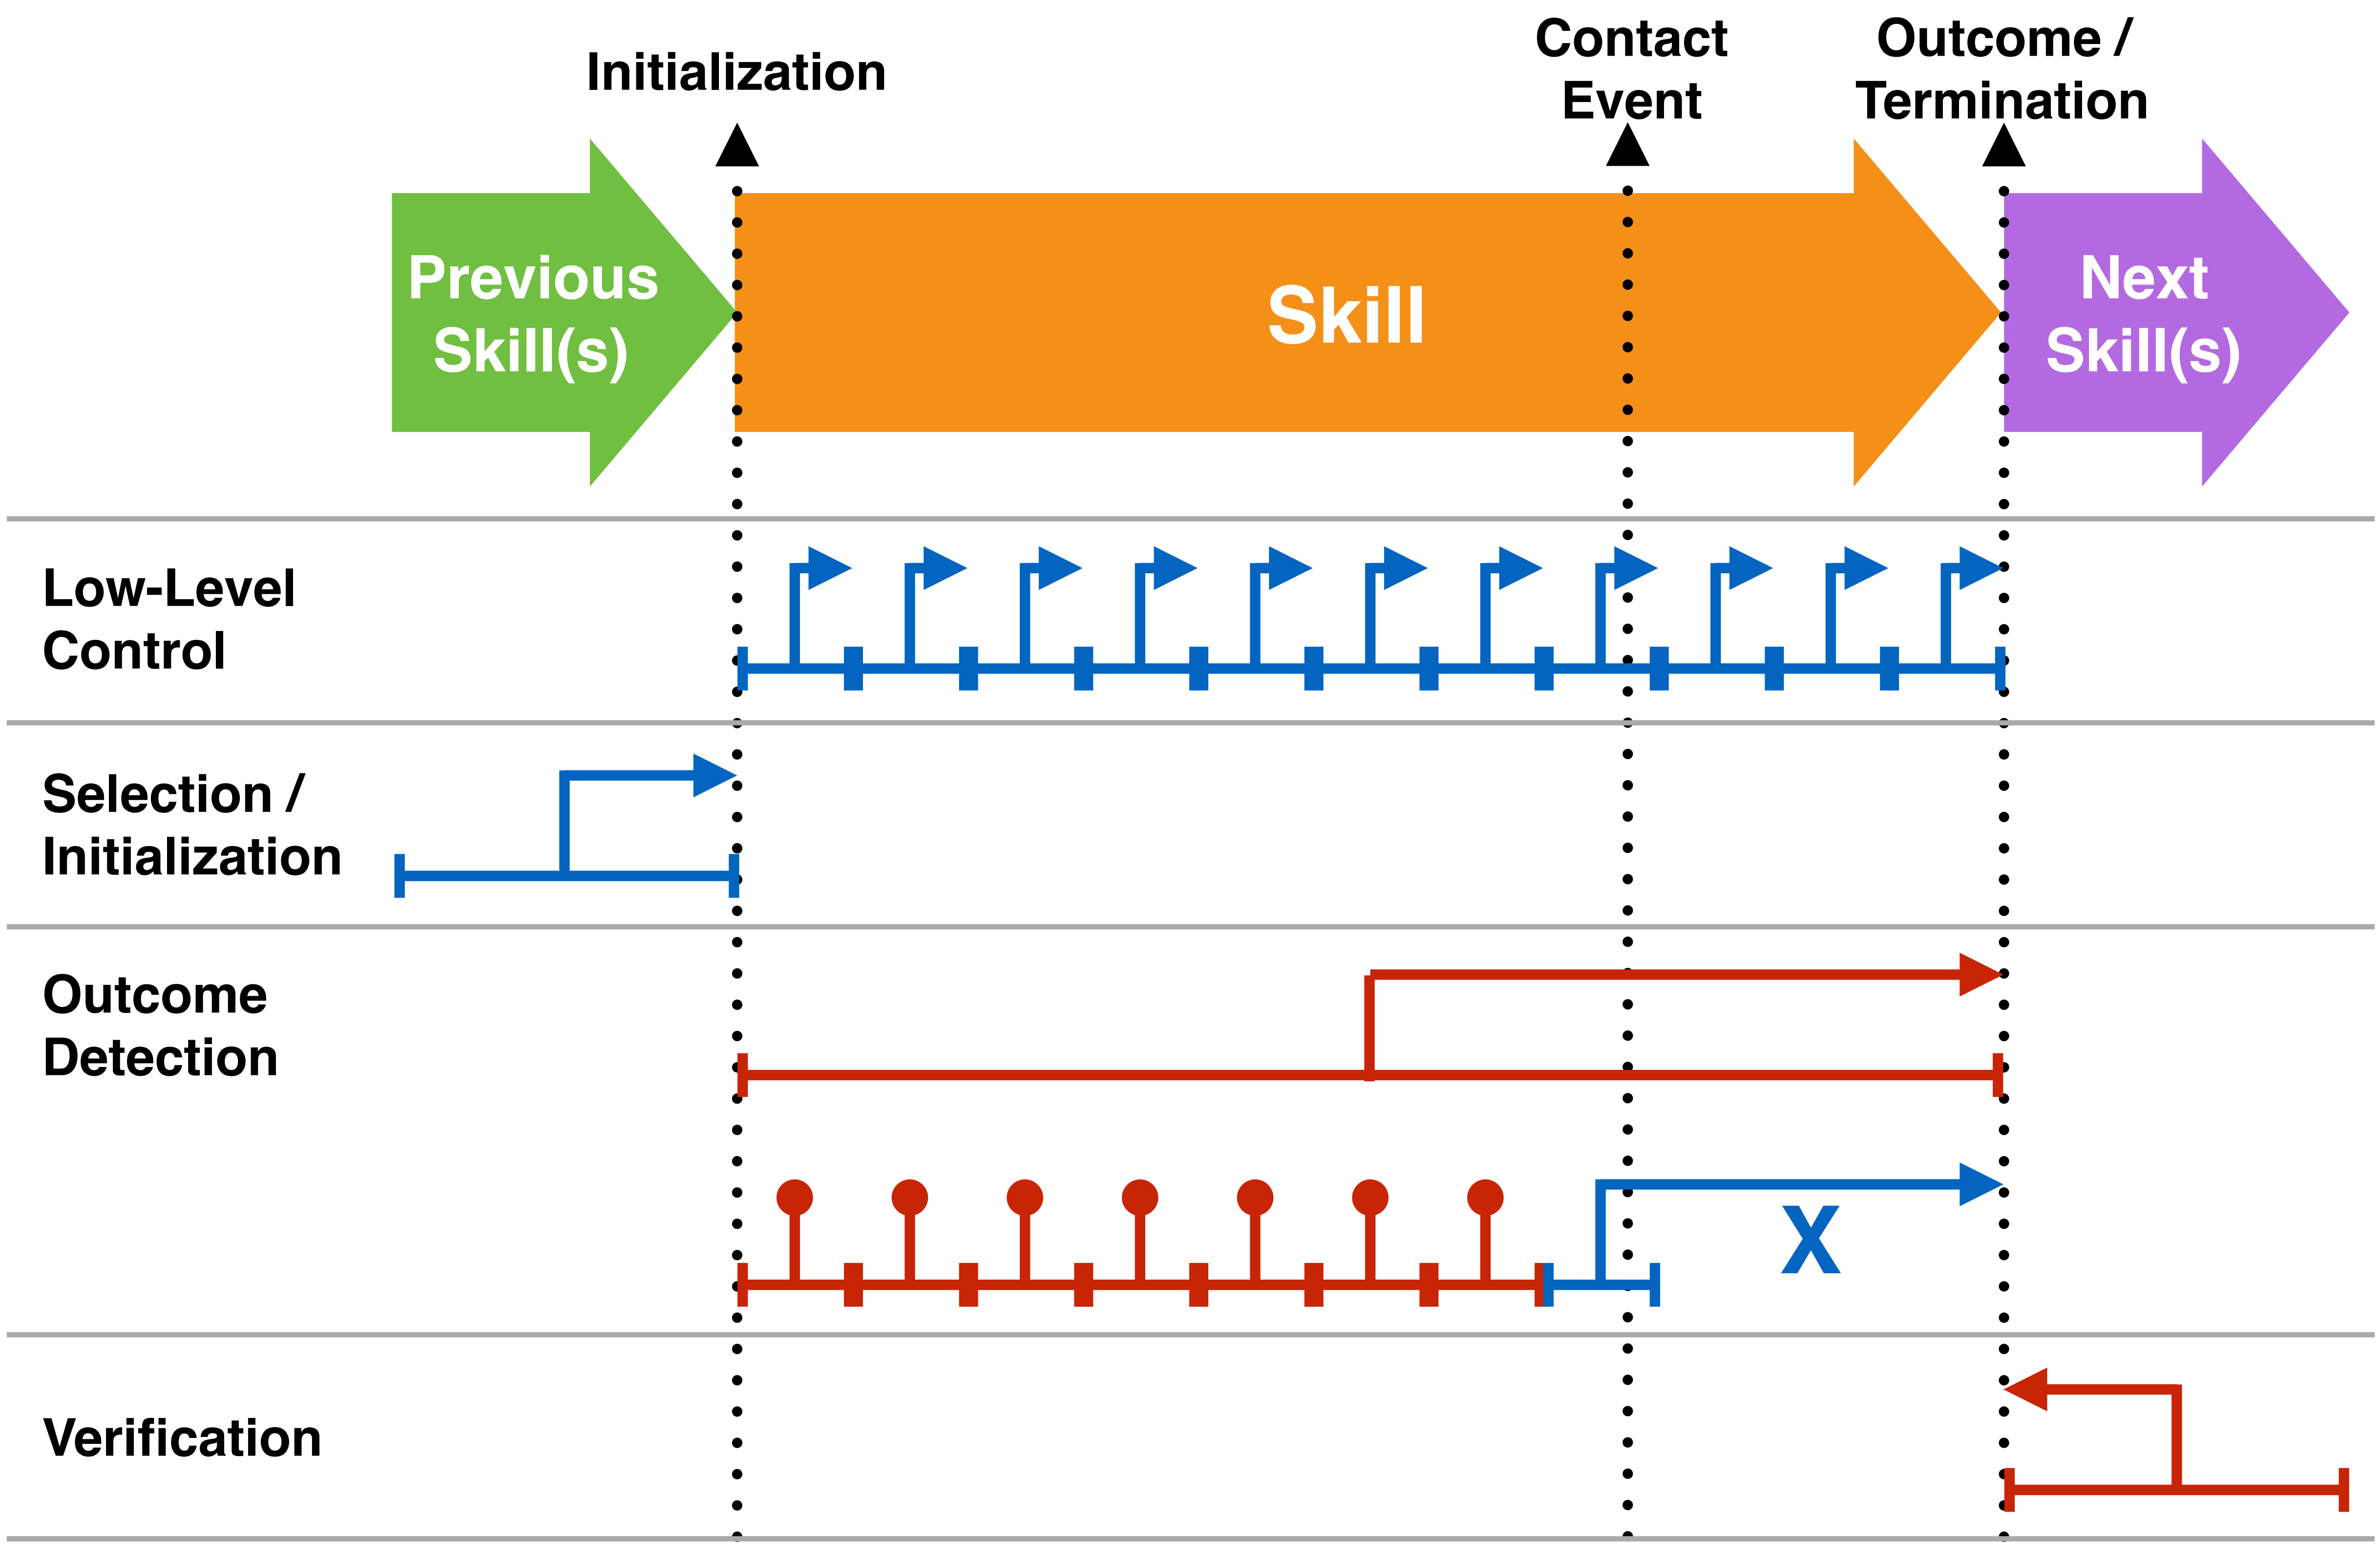

Use interactive perception to determine outcome

Outcome Verification

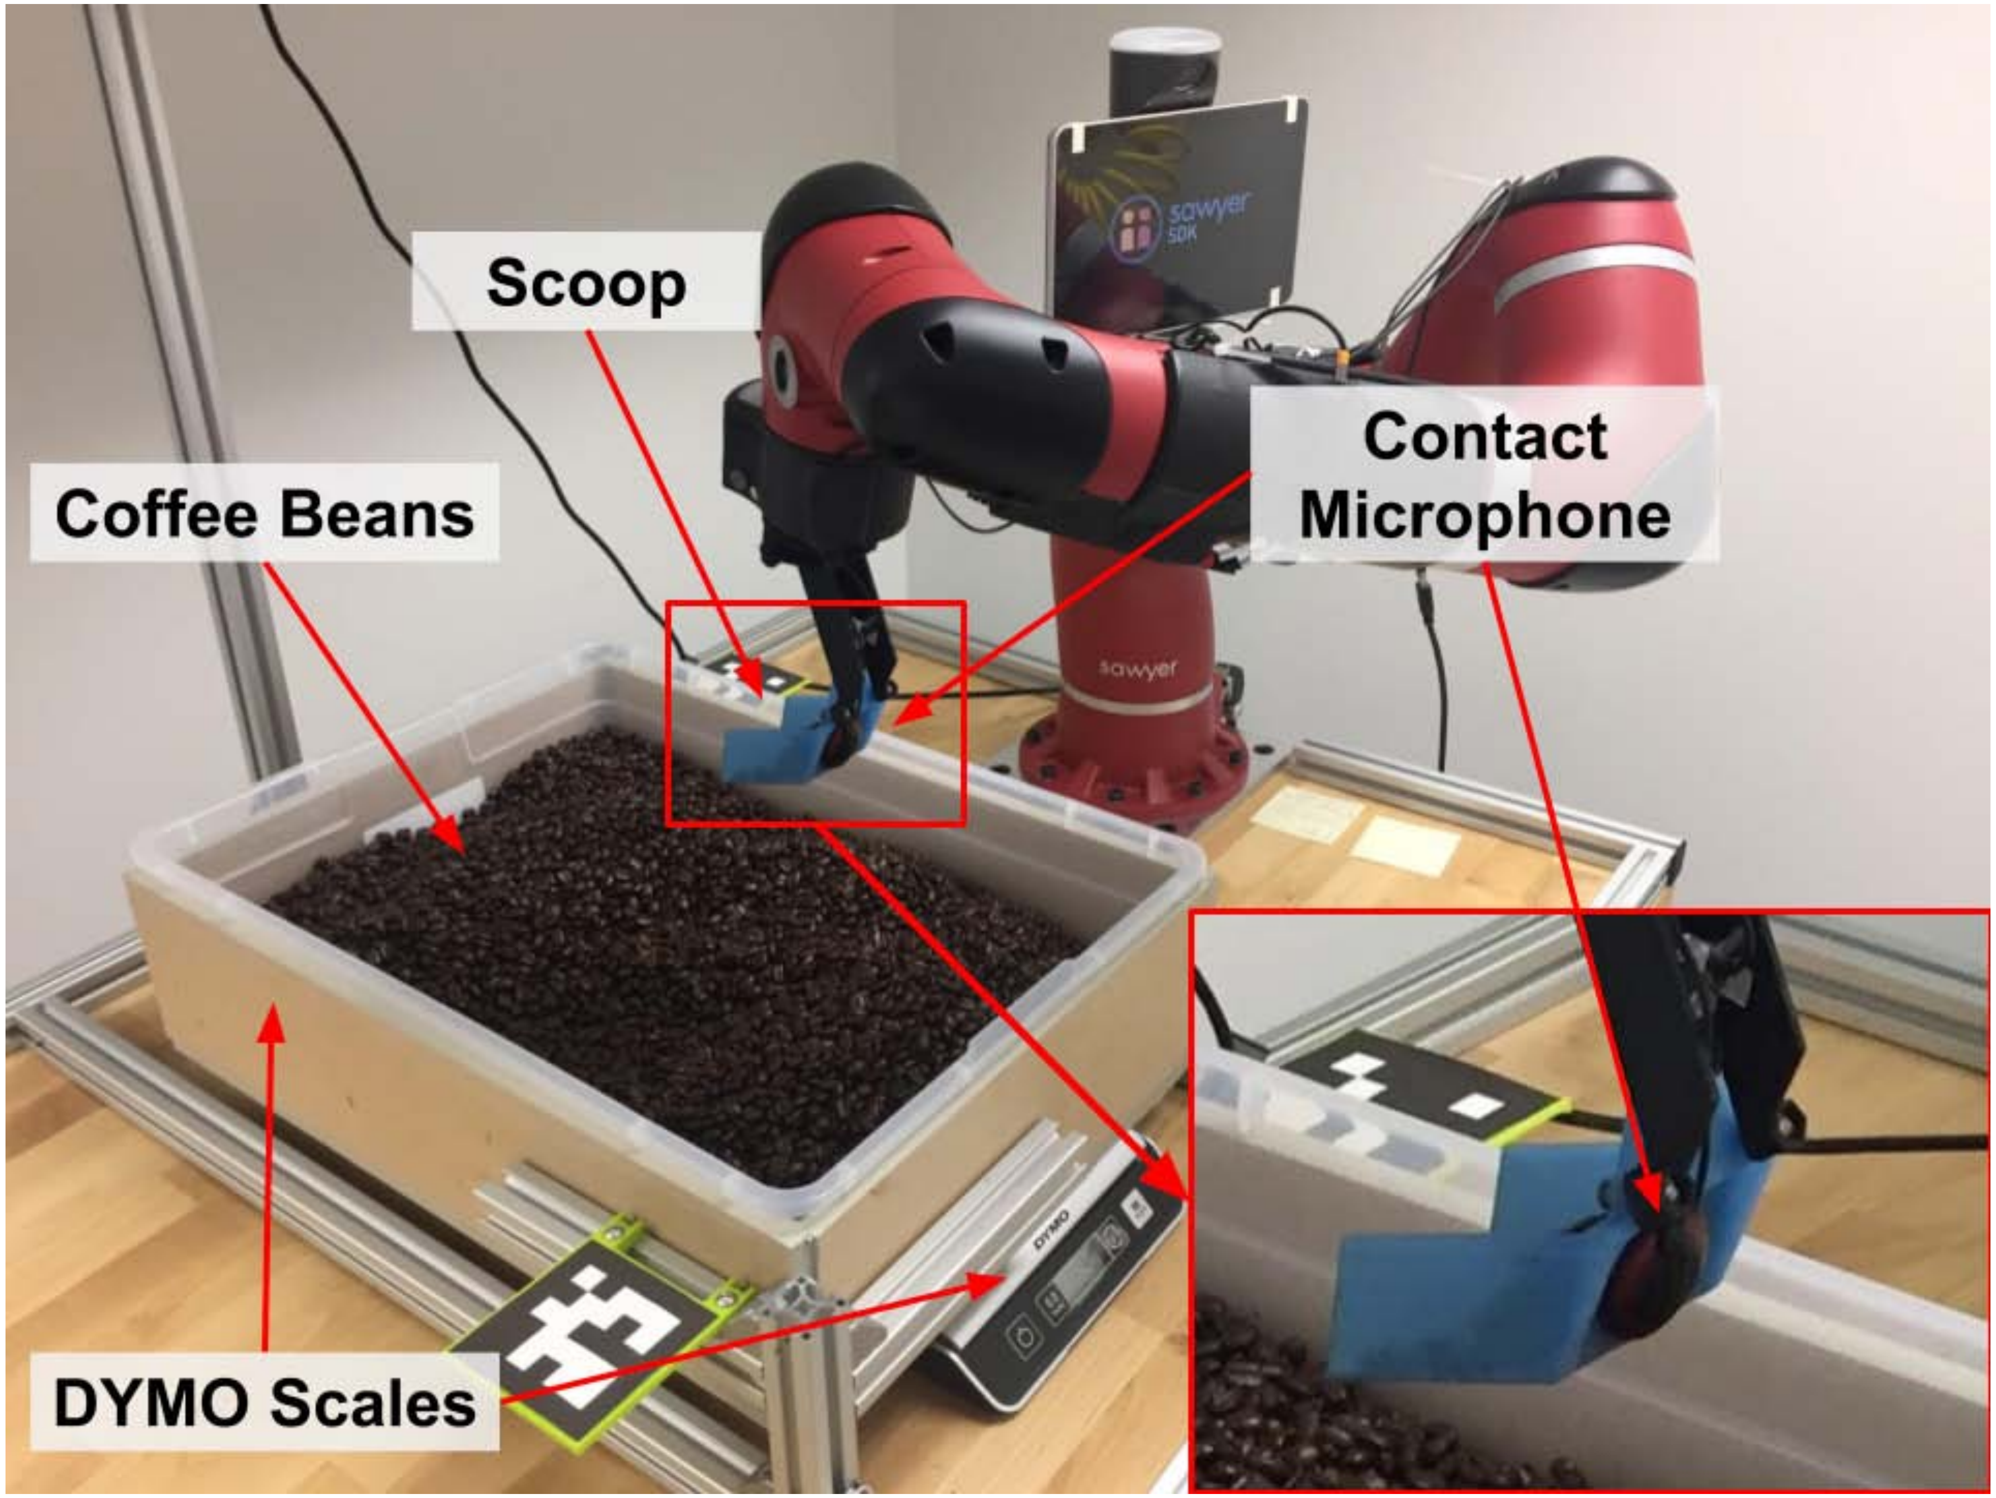

Scooping Skill

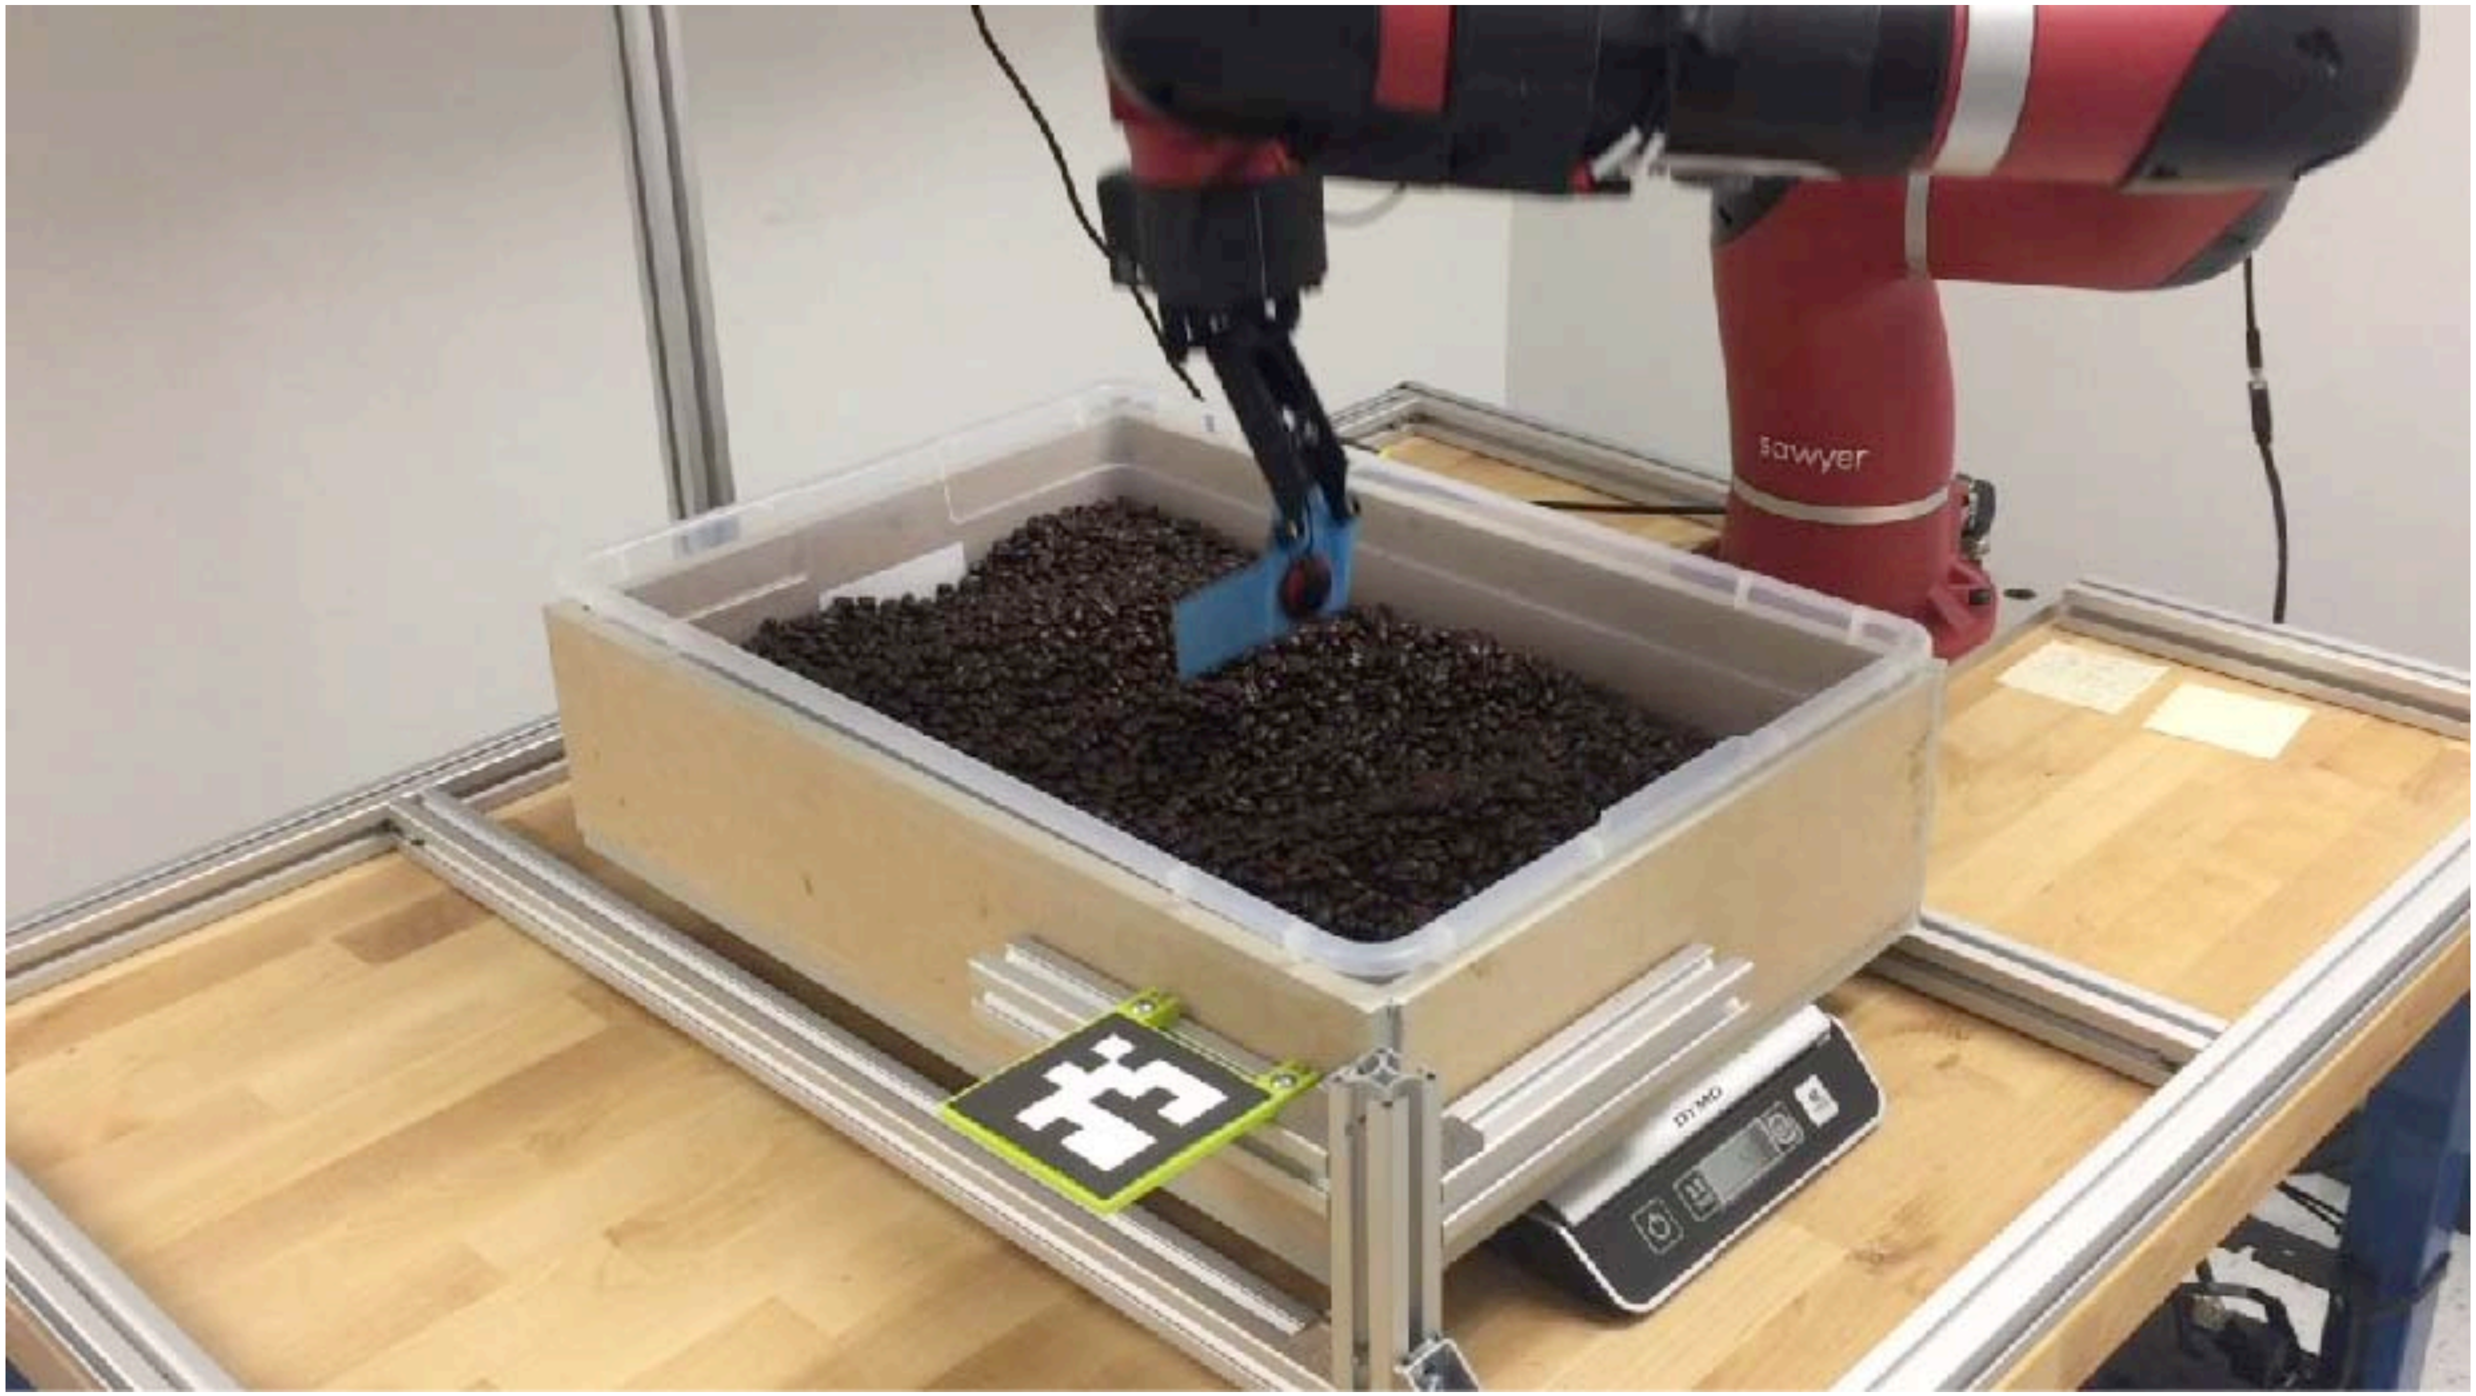

Shaking Verification

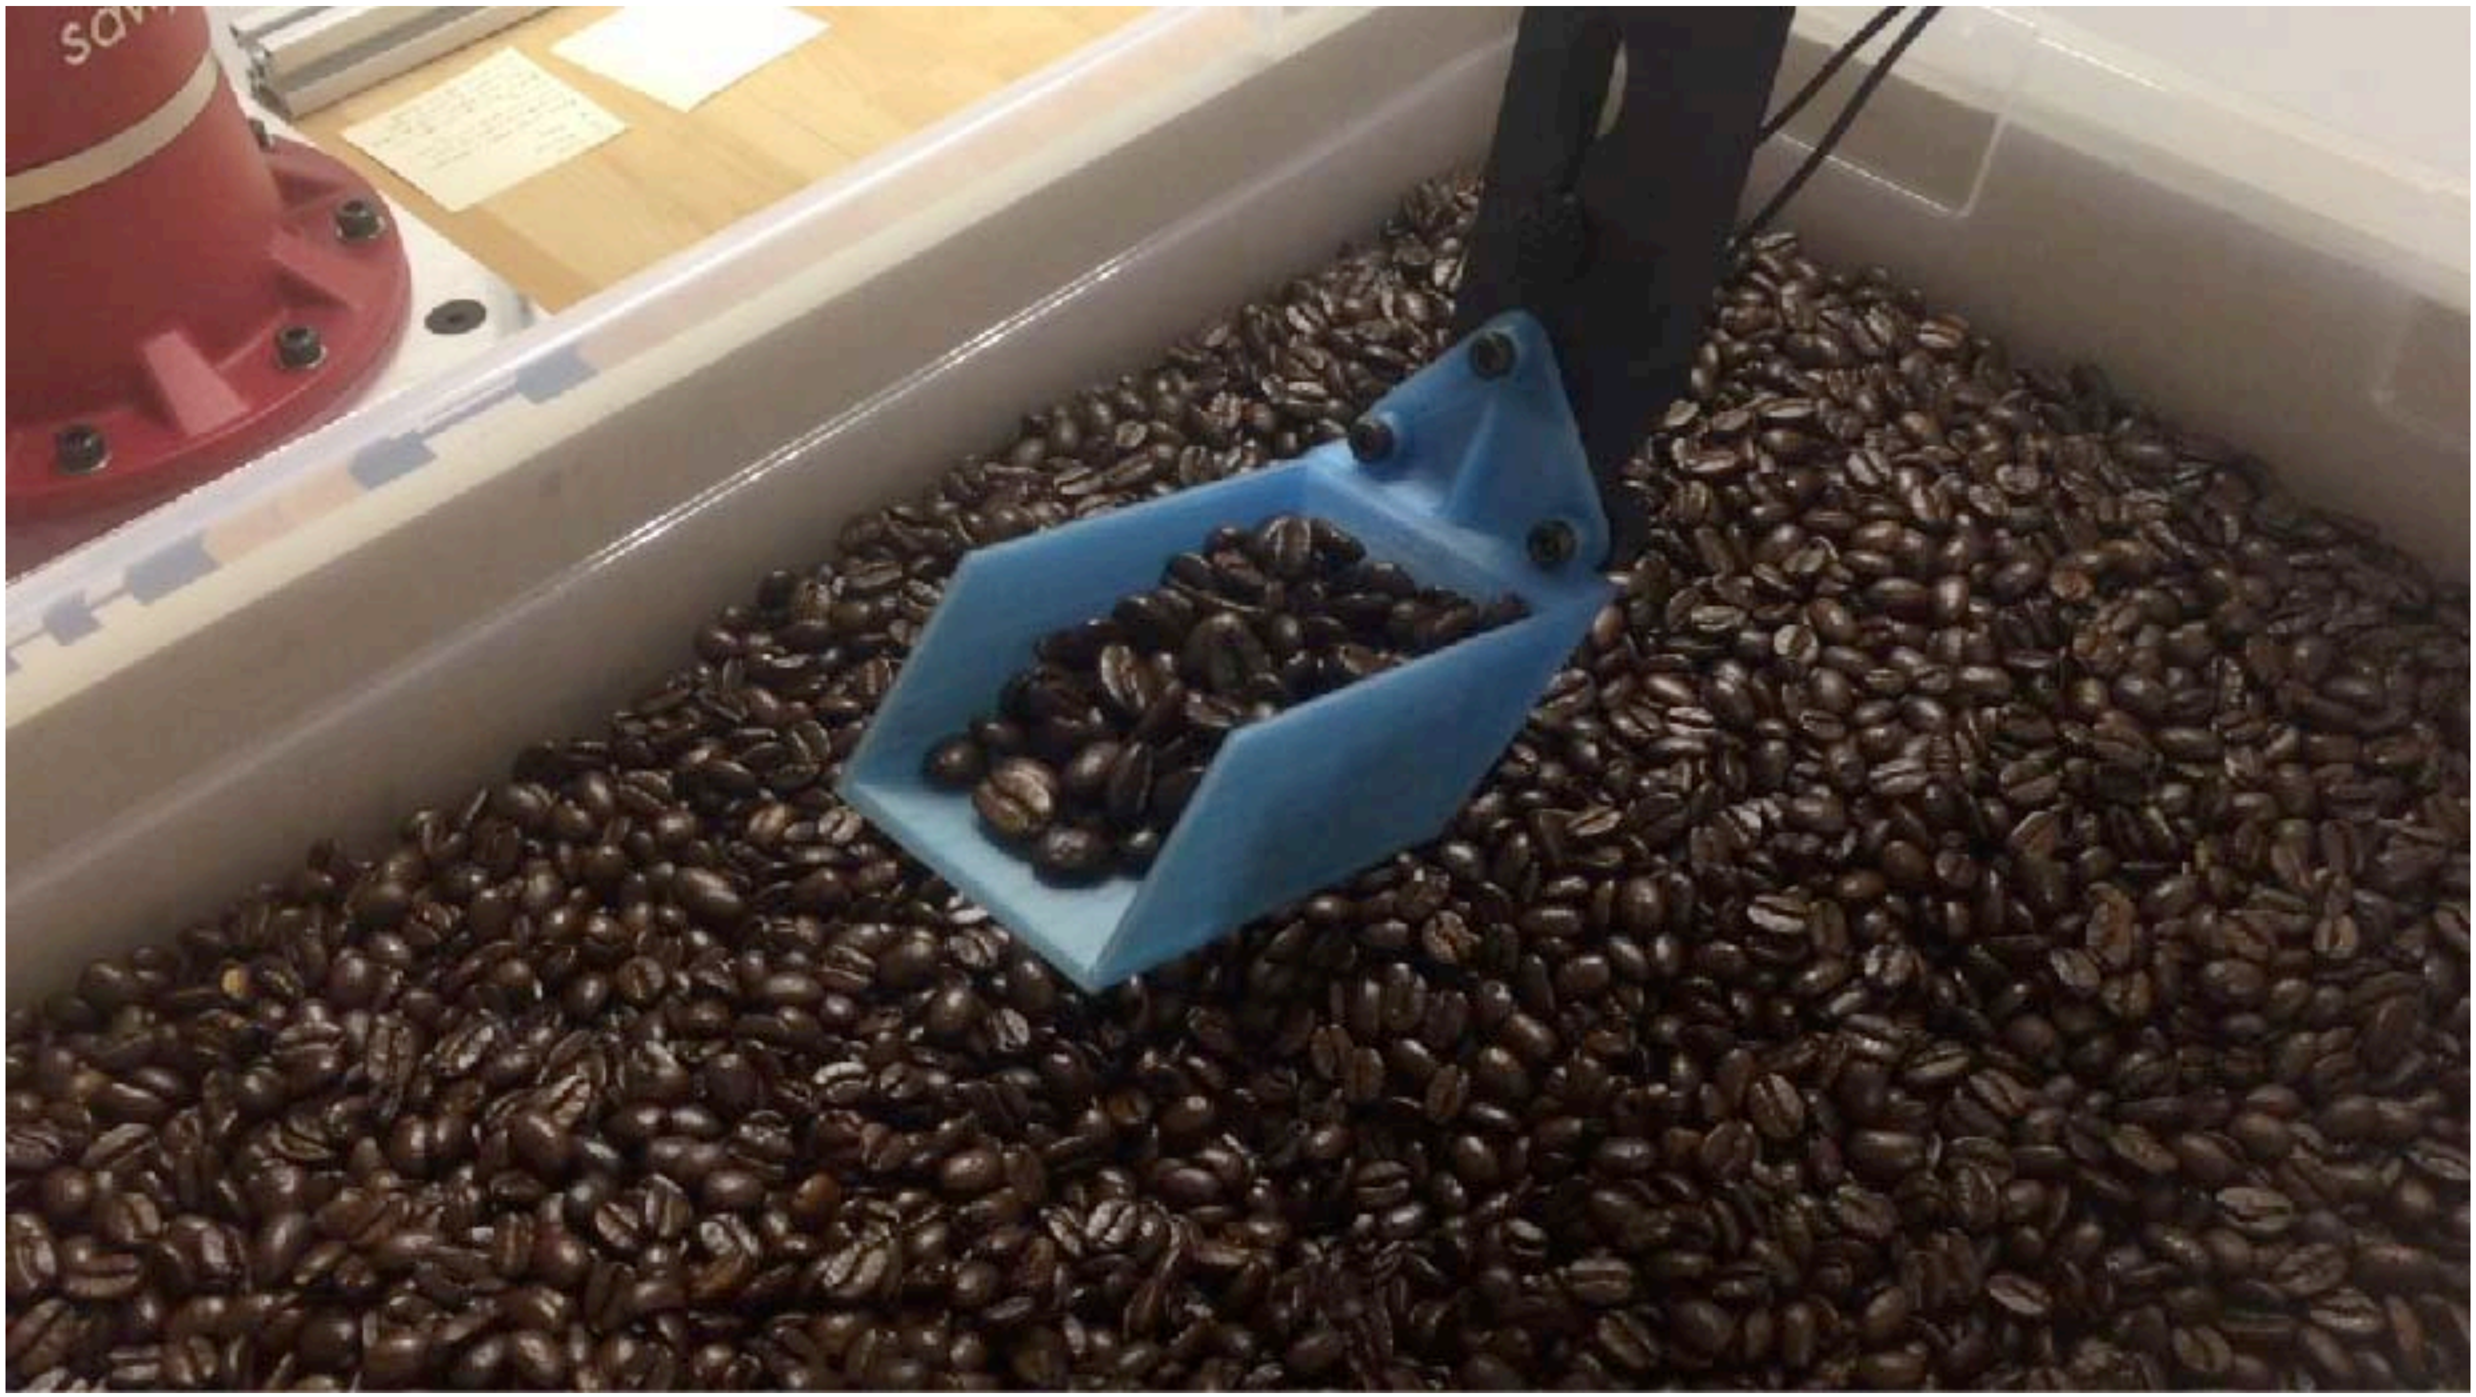

Learned network has an RMSE of less than 5 grams

(Pouring Termination)

Side note:

Audio signals can also be used to terminate pouring

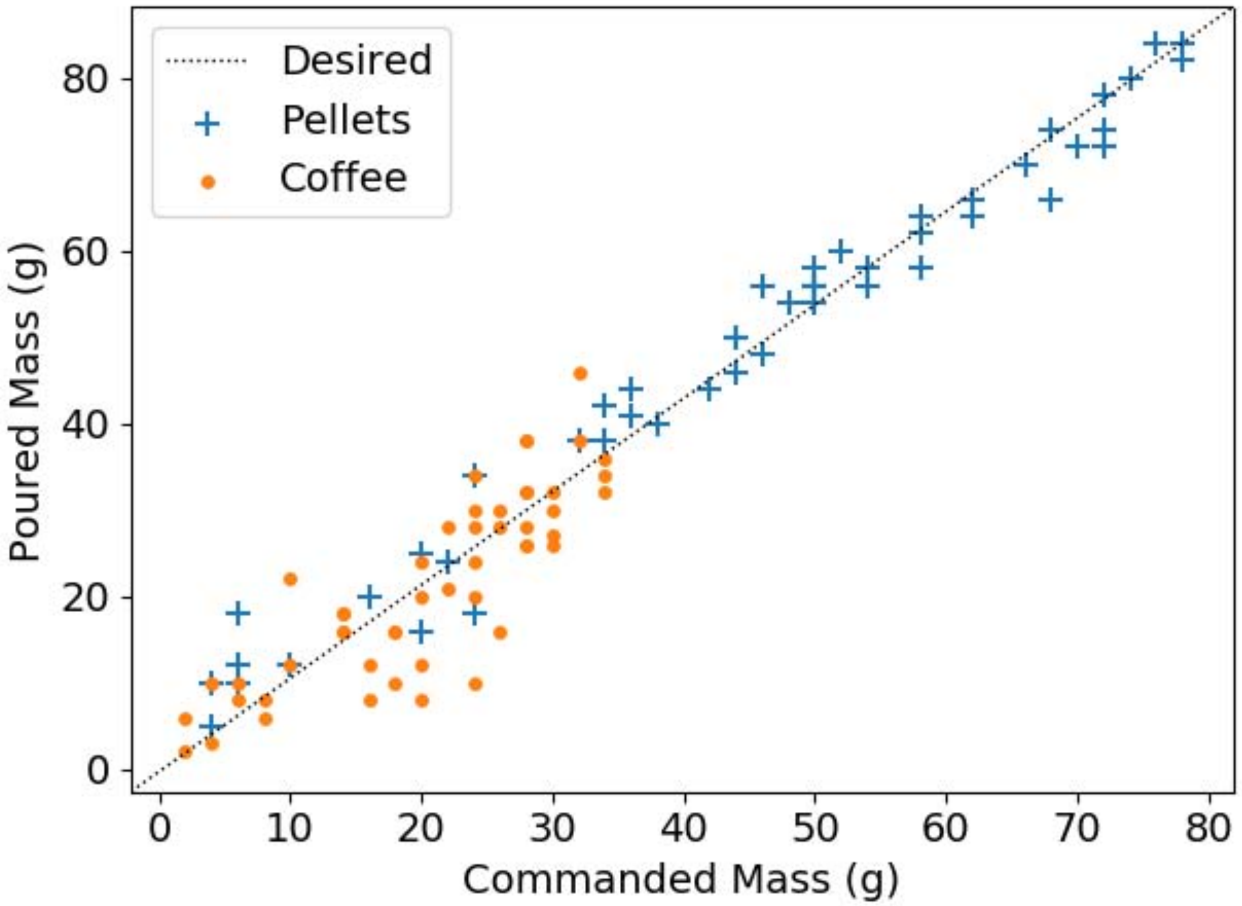

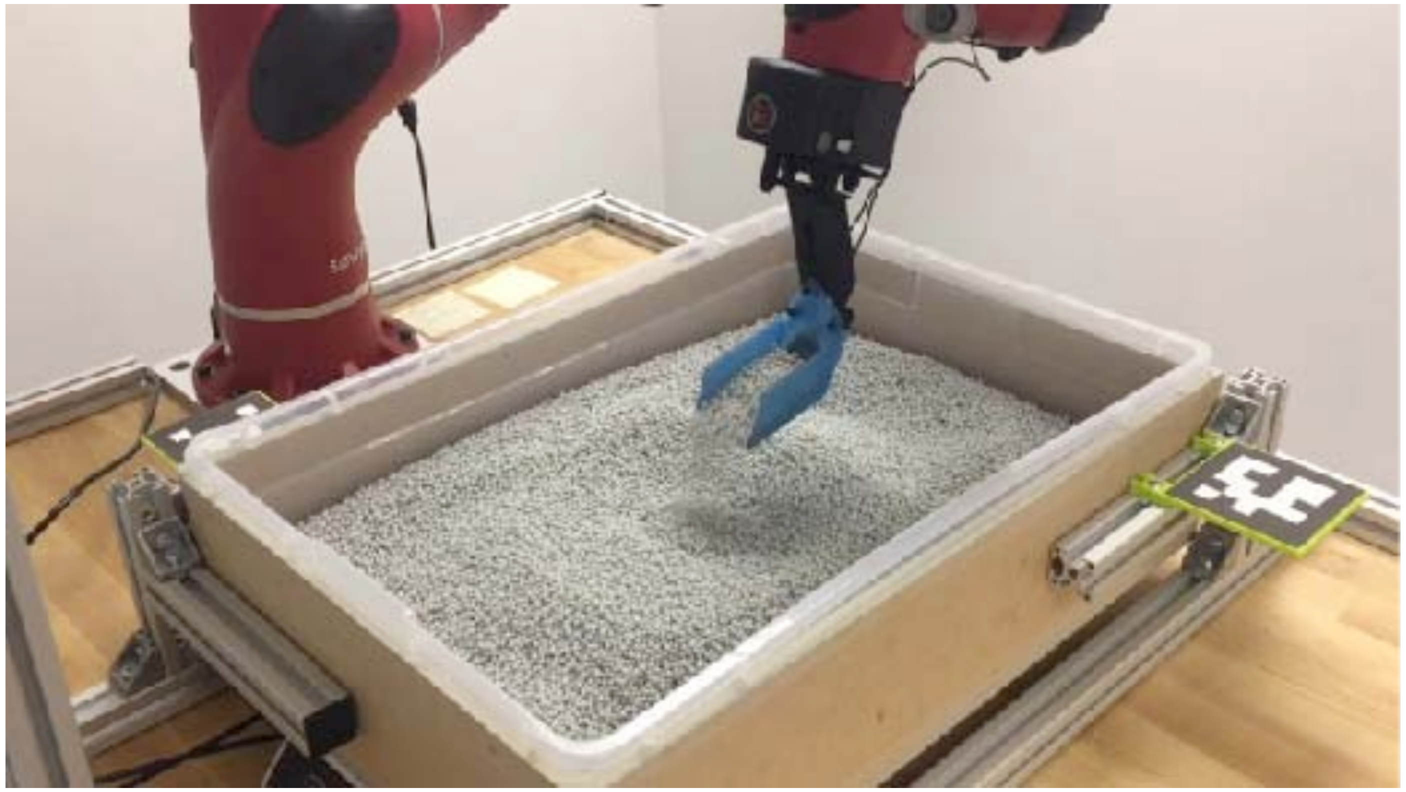

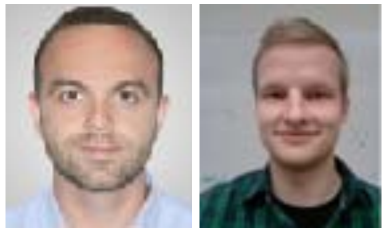

Verify Grasp with Shaking Action

- Also use verification to get self-supervised ground truth

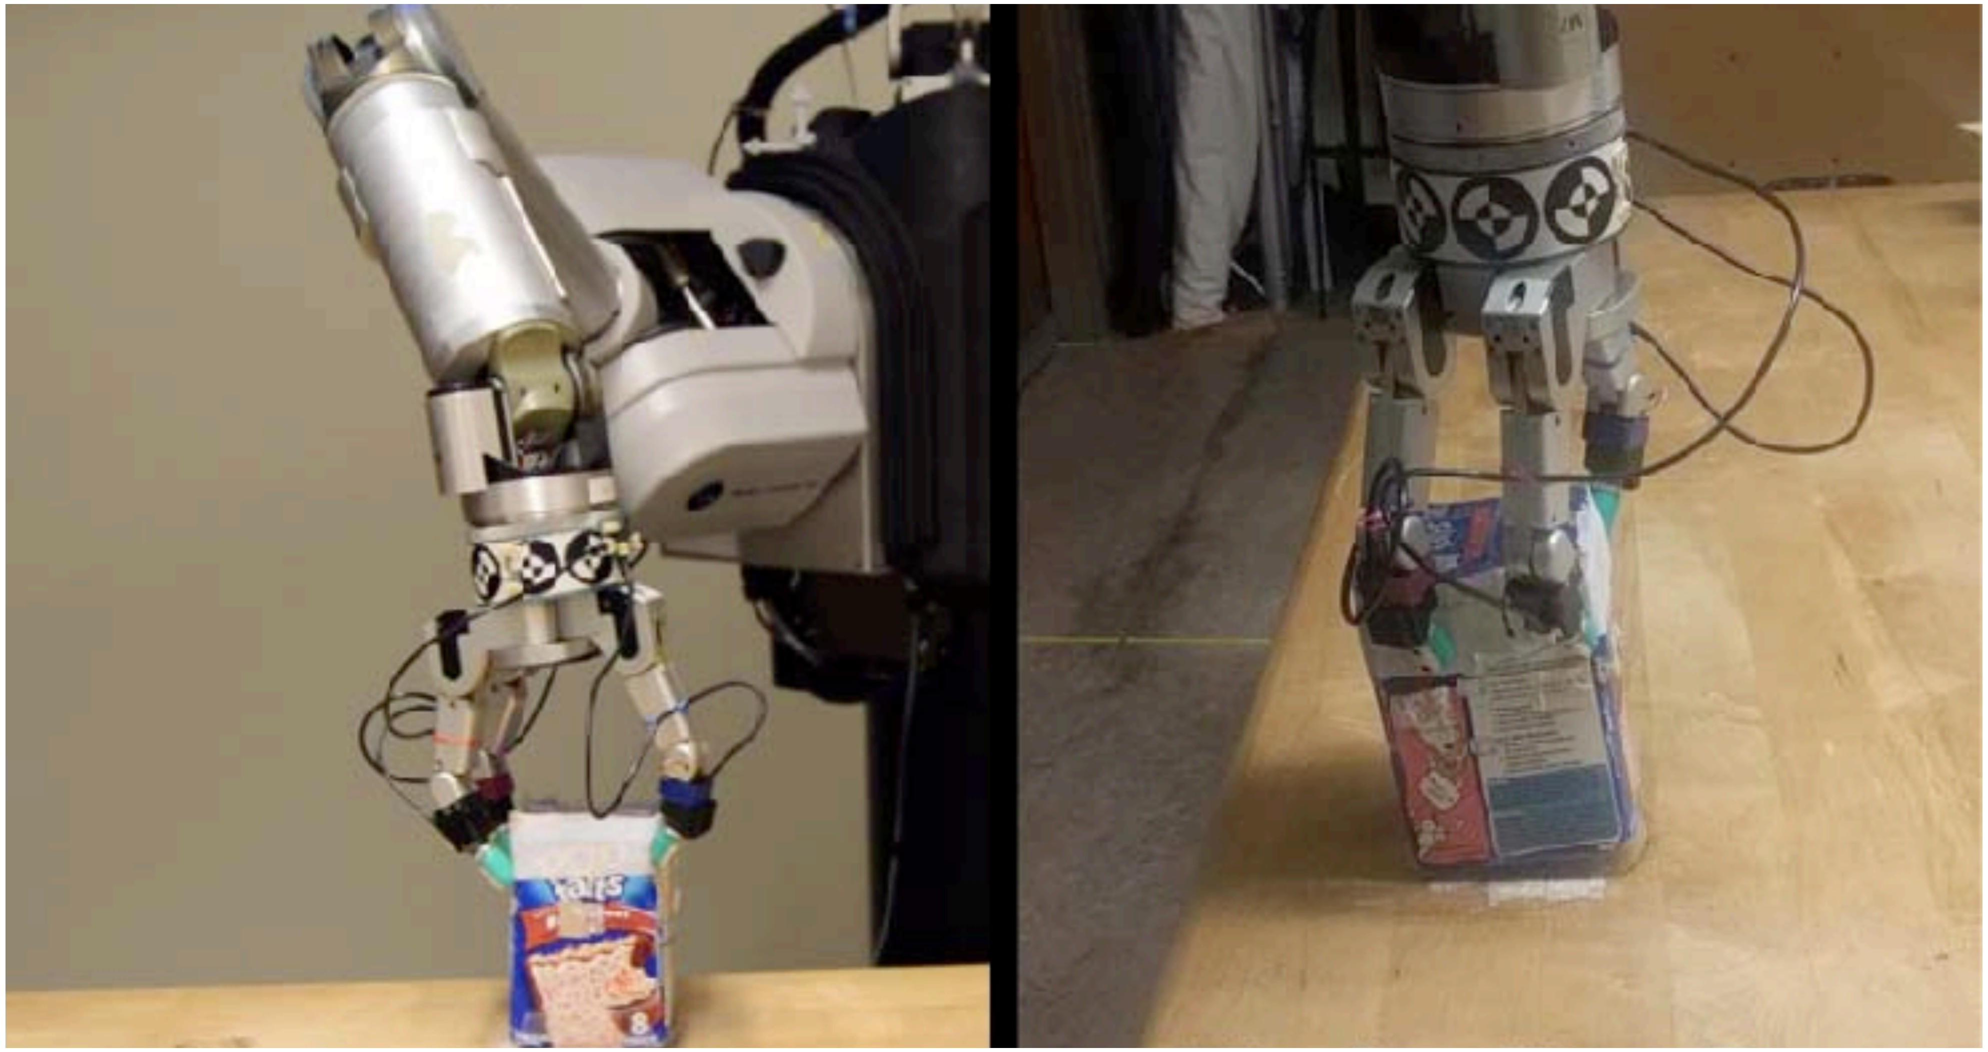

Y. Chebotar*, K. Hausman*, O. Kroemer, G. S. Sukhatme, S. Schaal. "Generalizing Regrasping with Supervised Policy Learning". ISER, 2016

Y. Chebotar*, K. Hausman*, Z. Su, G. S. Sukhatme,S. Schaal. "Self-Supervised Regrasping using Spatio-Temporal Tactile Features and Reinforcement Learning". IROS, 2016

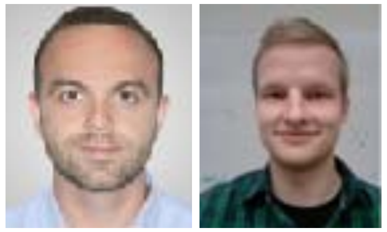

Self-Supervised Outcome Detection

- Learn to predict errors using self-supervised data

Lift object slightly

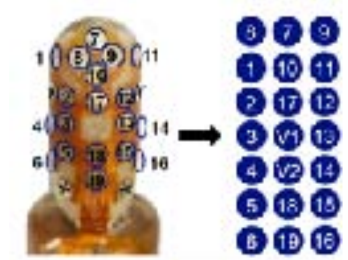

Predict grasp stability

Success

Continue with the grasp

Failure

?????

- Learn a classifier using ST-HMP features and SVM
- What to do if the grasp fails?
  - Use lift skill as a previous skill for initialising a regrasp

Regrasping

- Learn to predict errors and regrasp using tactile data

Lift object slightly

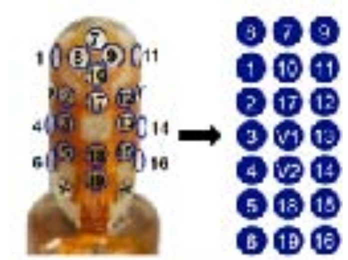

Predict grasp stability

Success

Continue with the grasp (Shake)

Failure

Adjust gripper pose based on **tactile signal**

- Adjust gripper position and orientation based on tactile


Reinforcement Learning

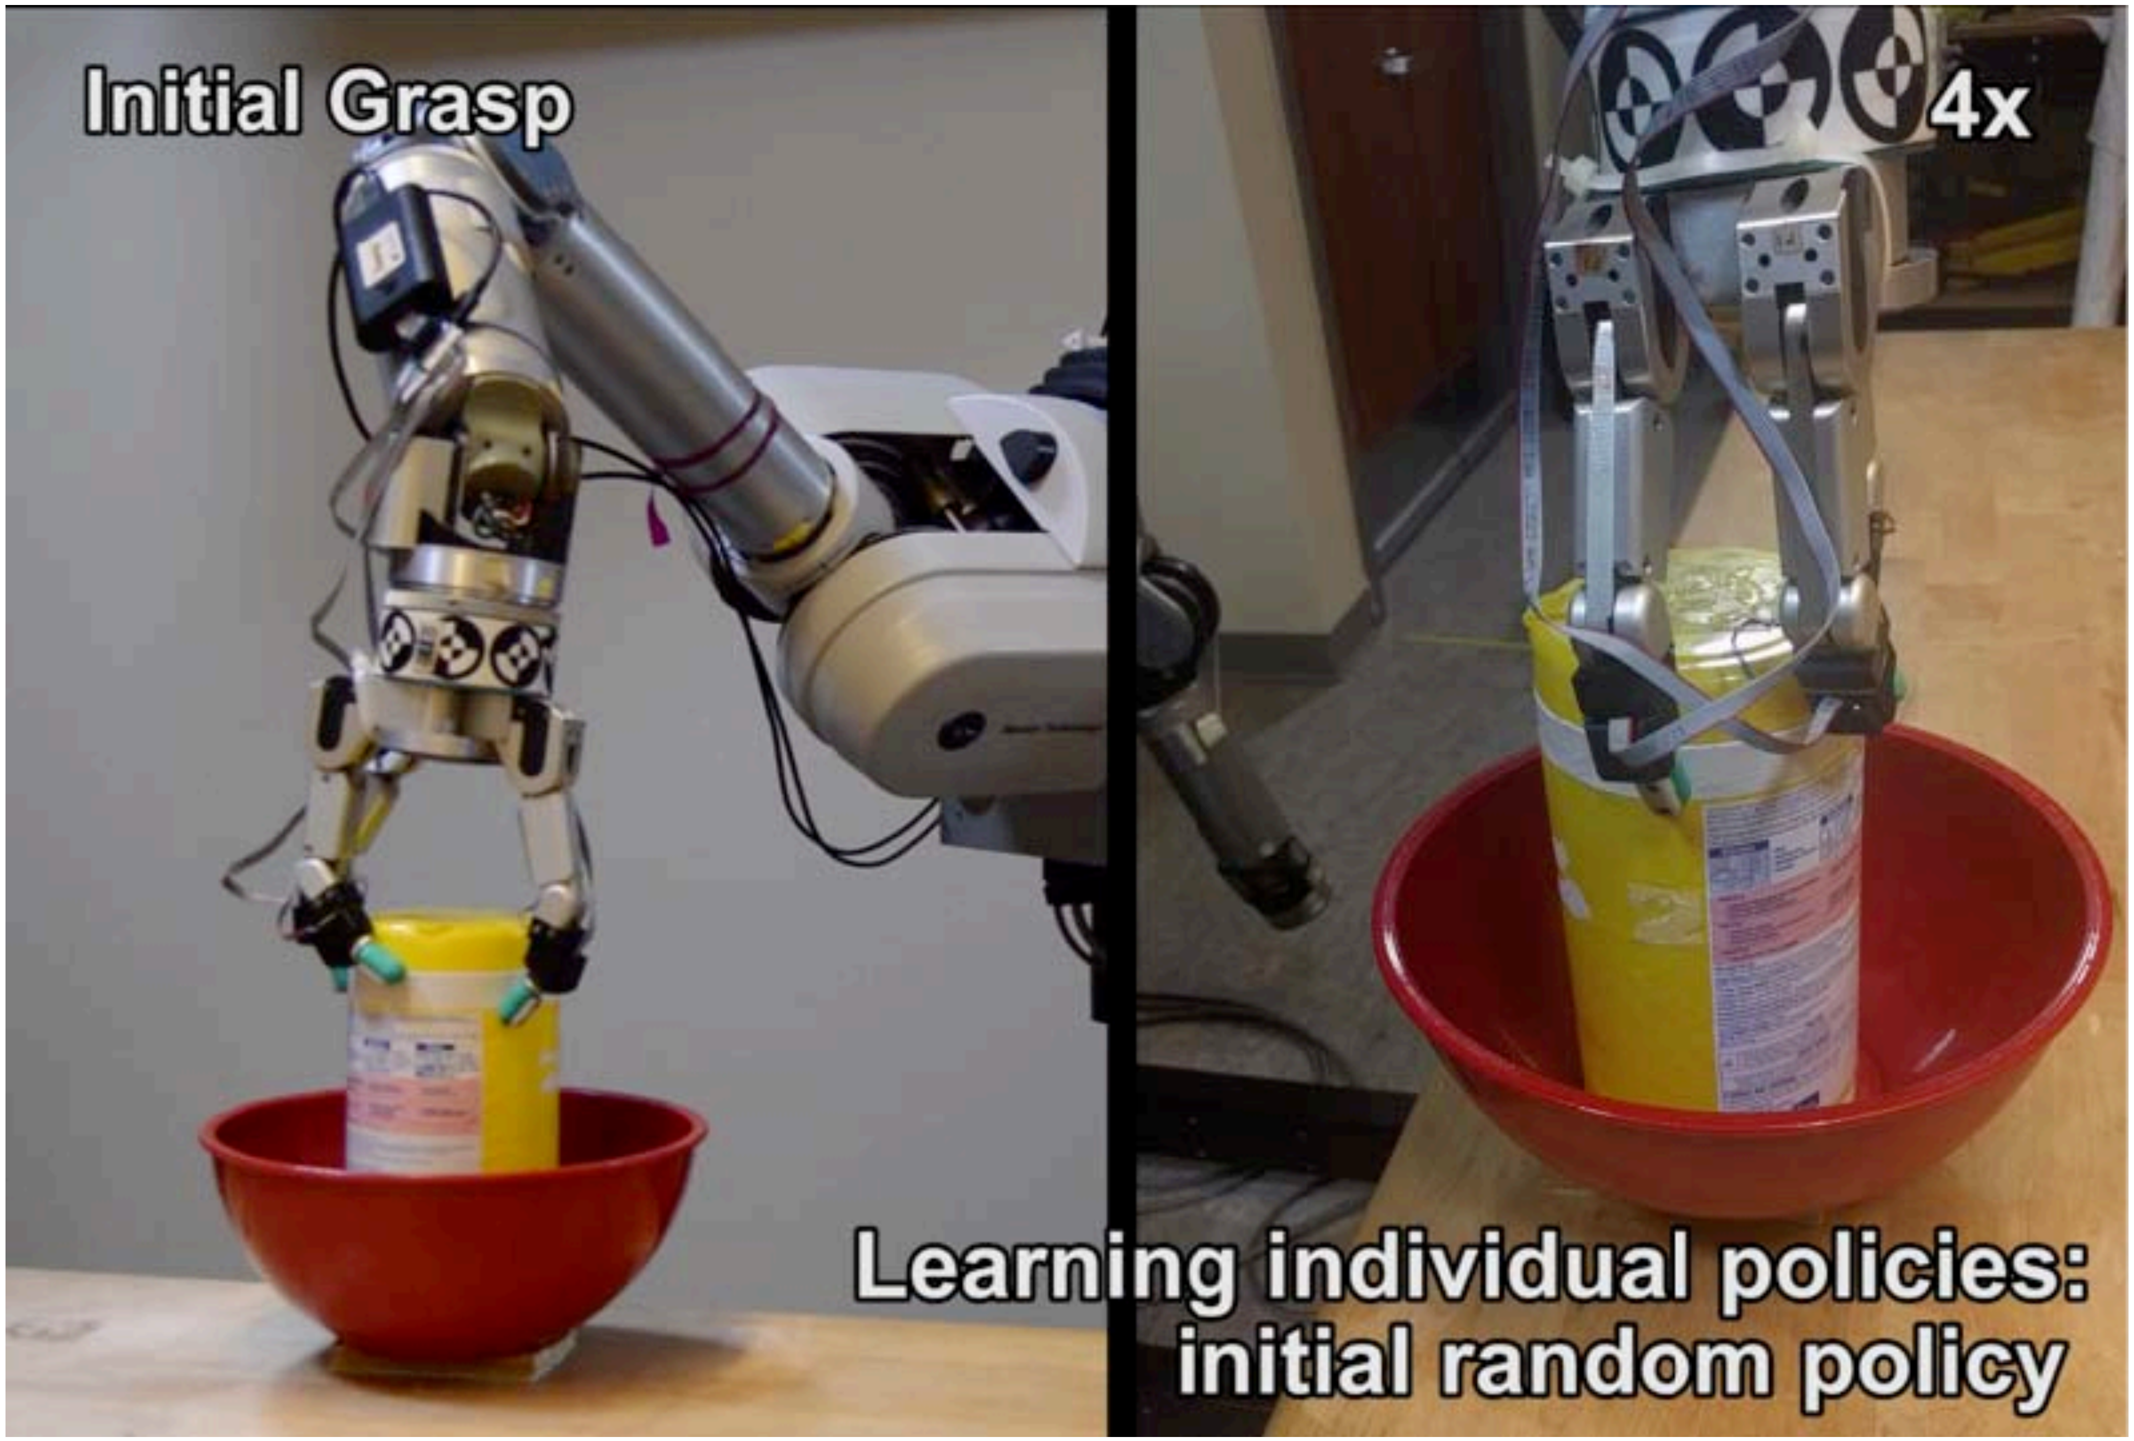

Learn regrasp policy parameters using relative entropy policy search

[Peters et al. 2010]

Regrasping Results

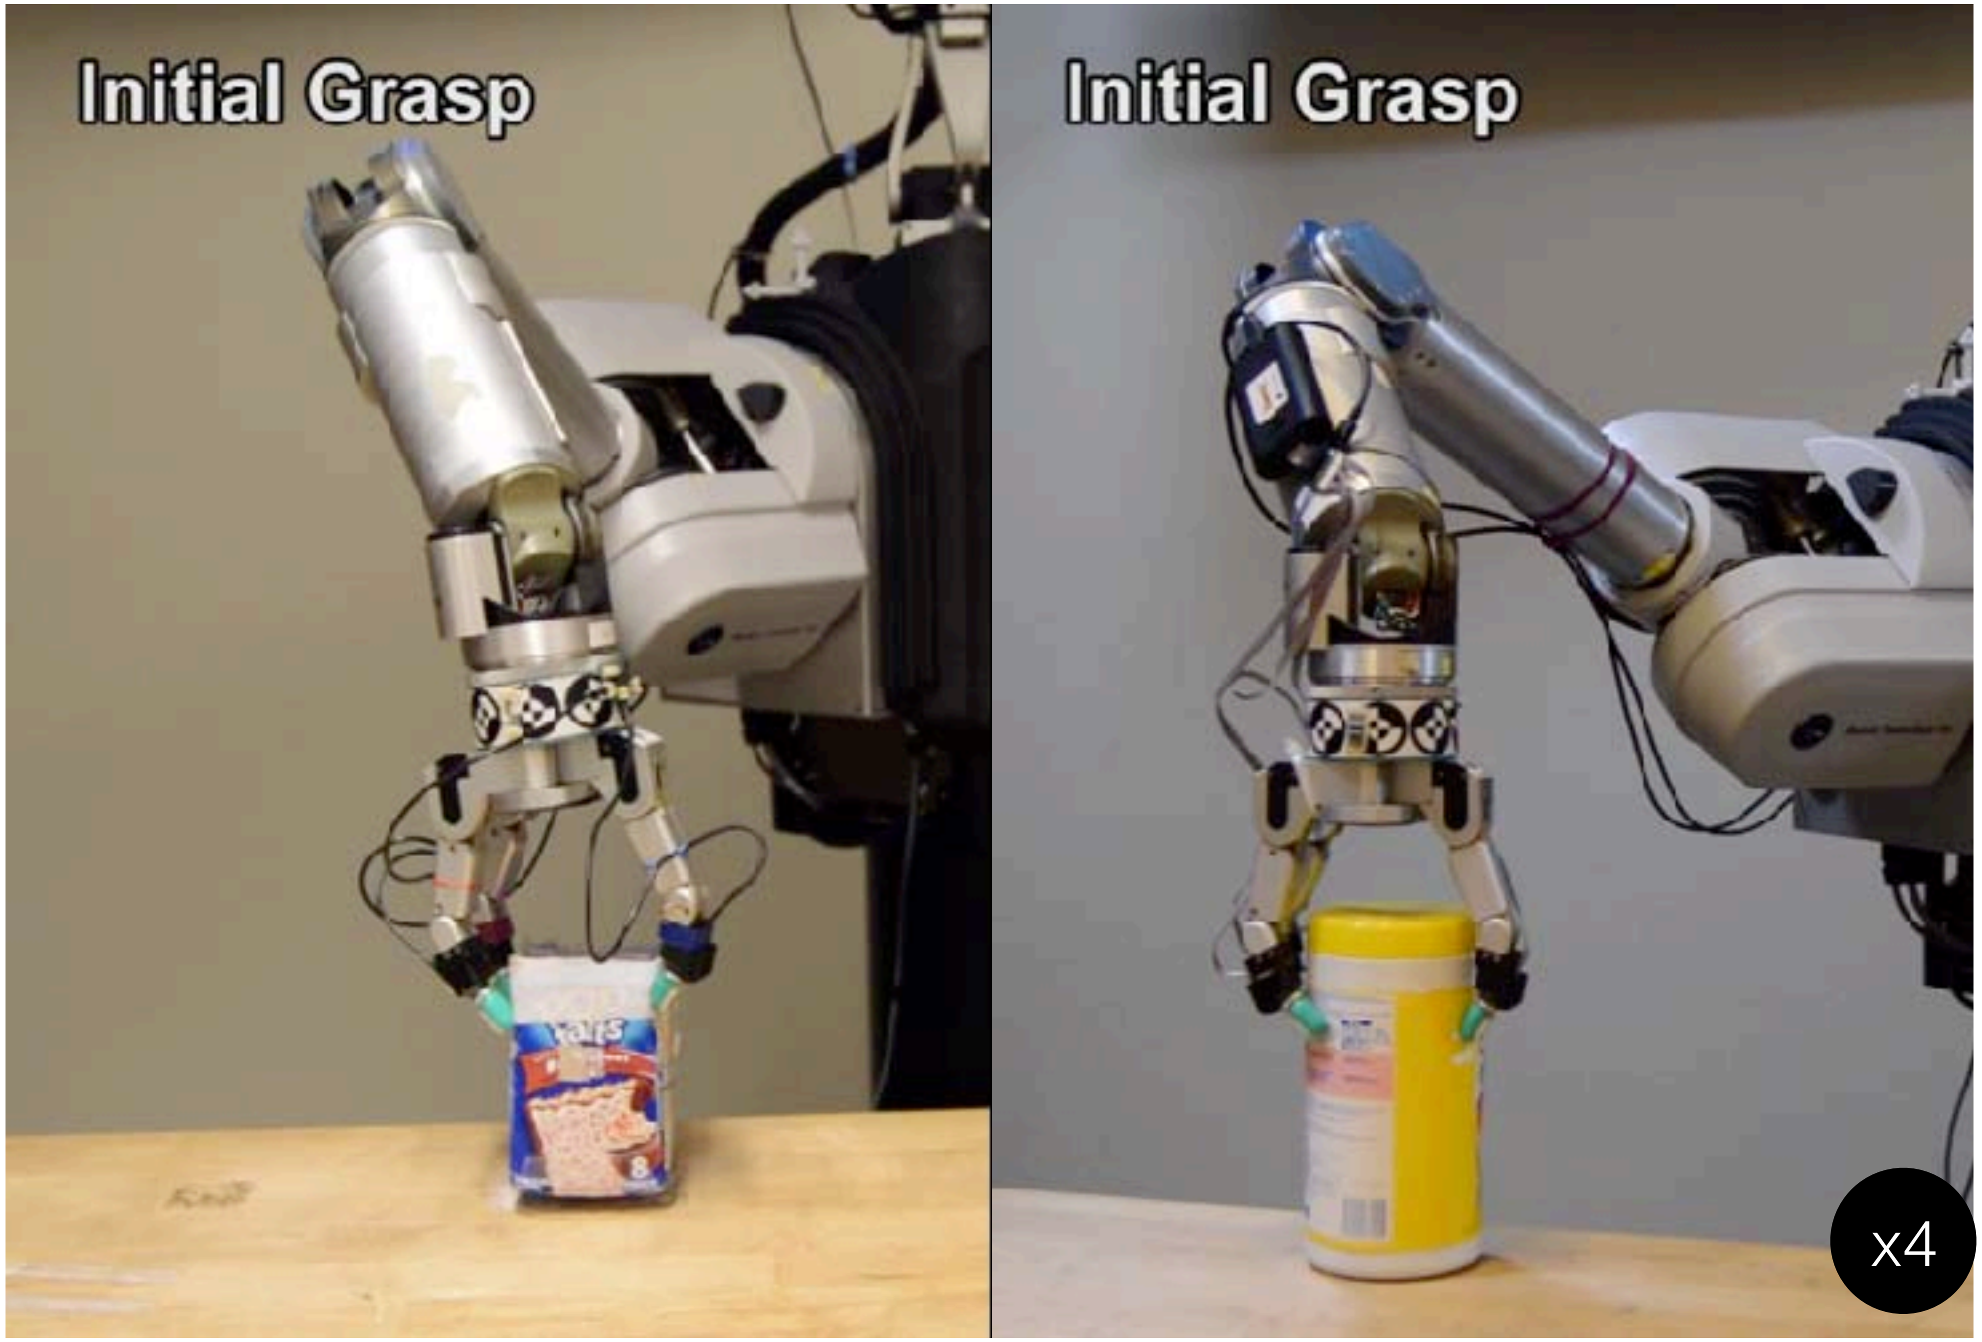

 

Regrasping Results

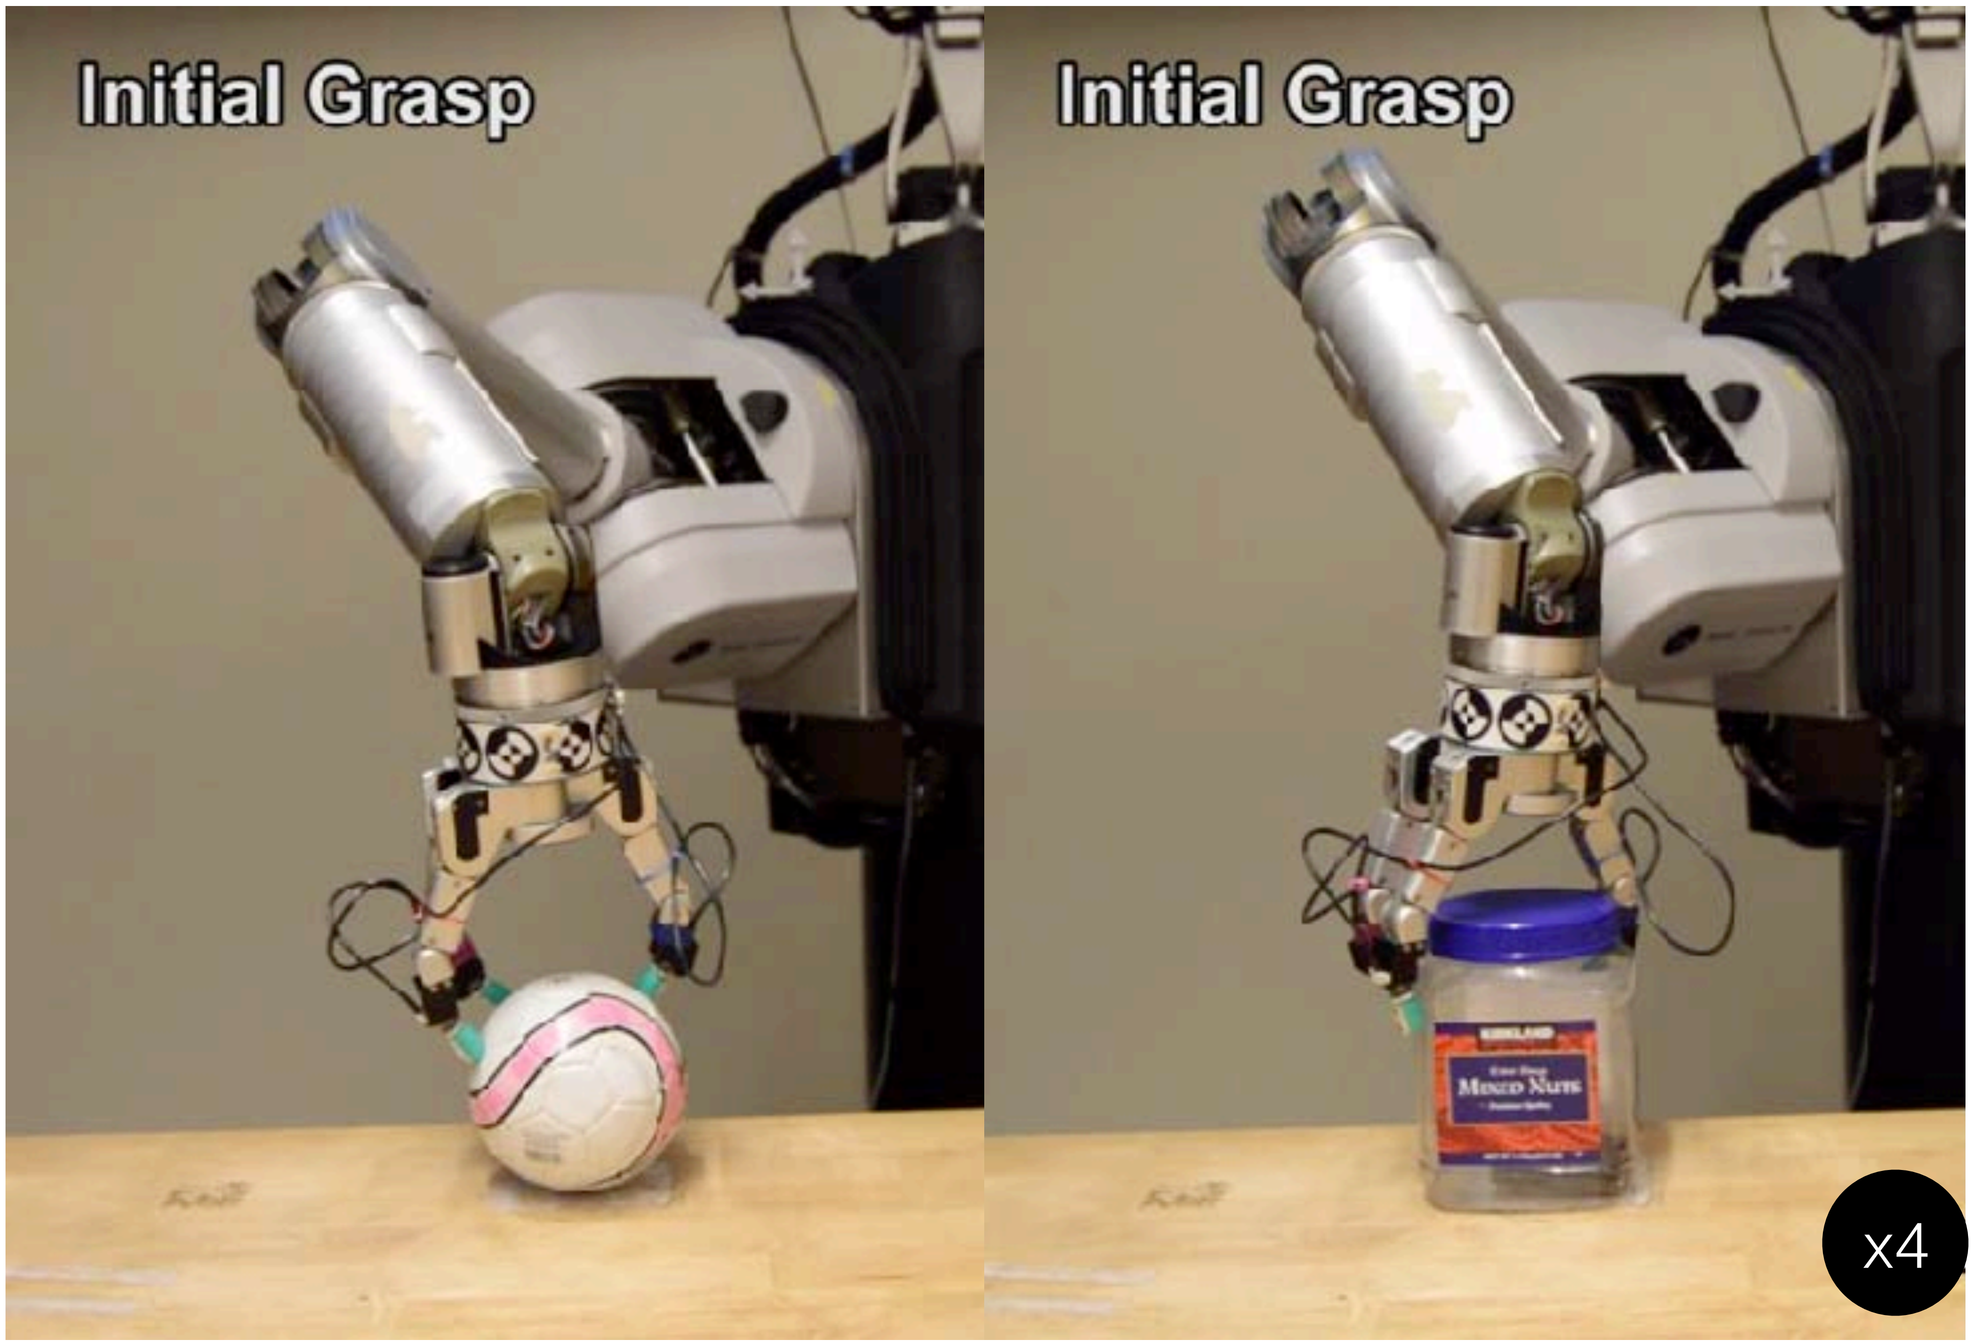

Sensory Feedback

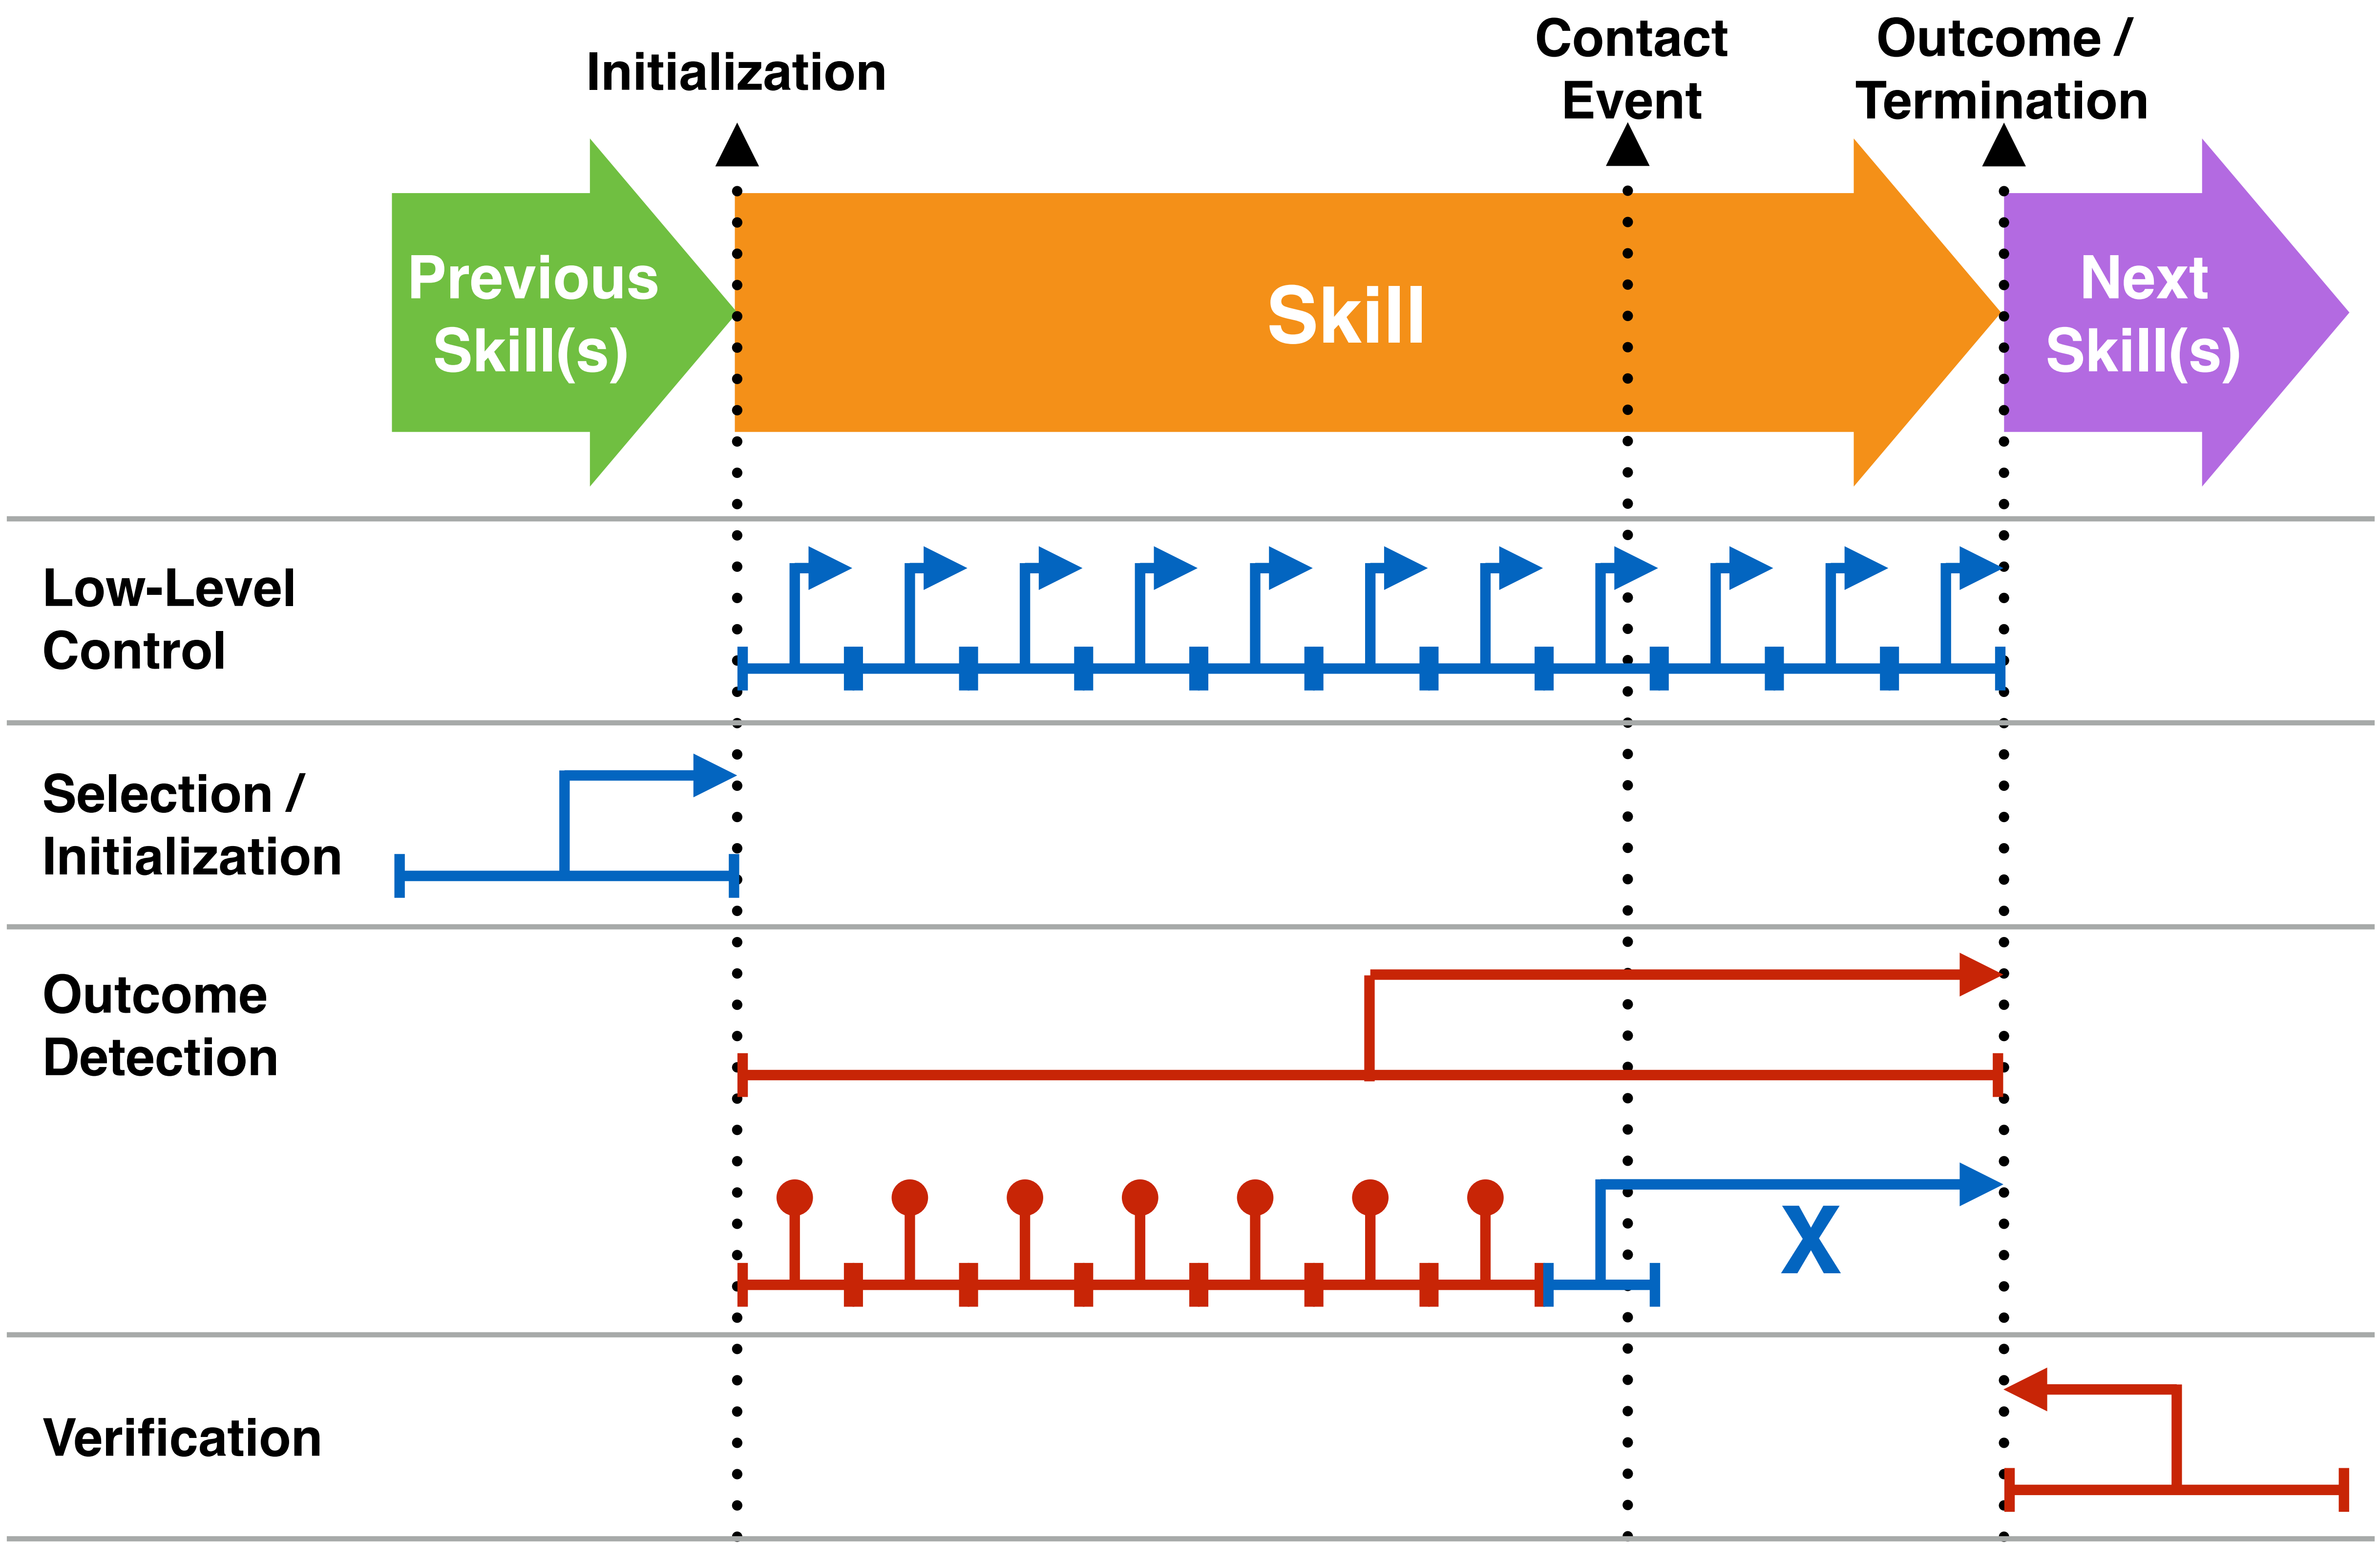

Key Challenge: Preconditions for Learning

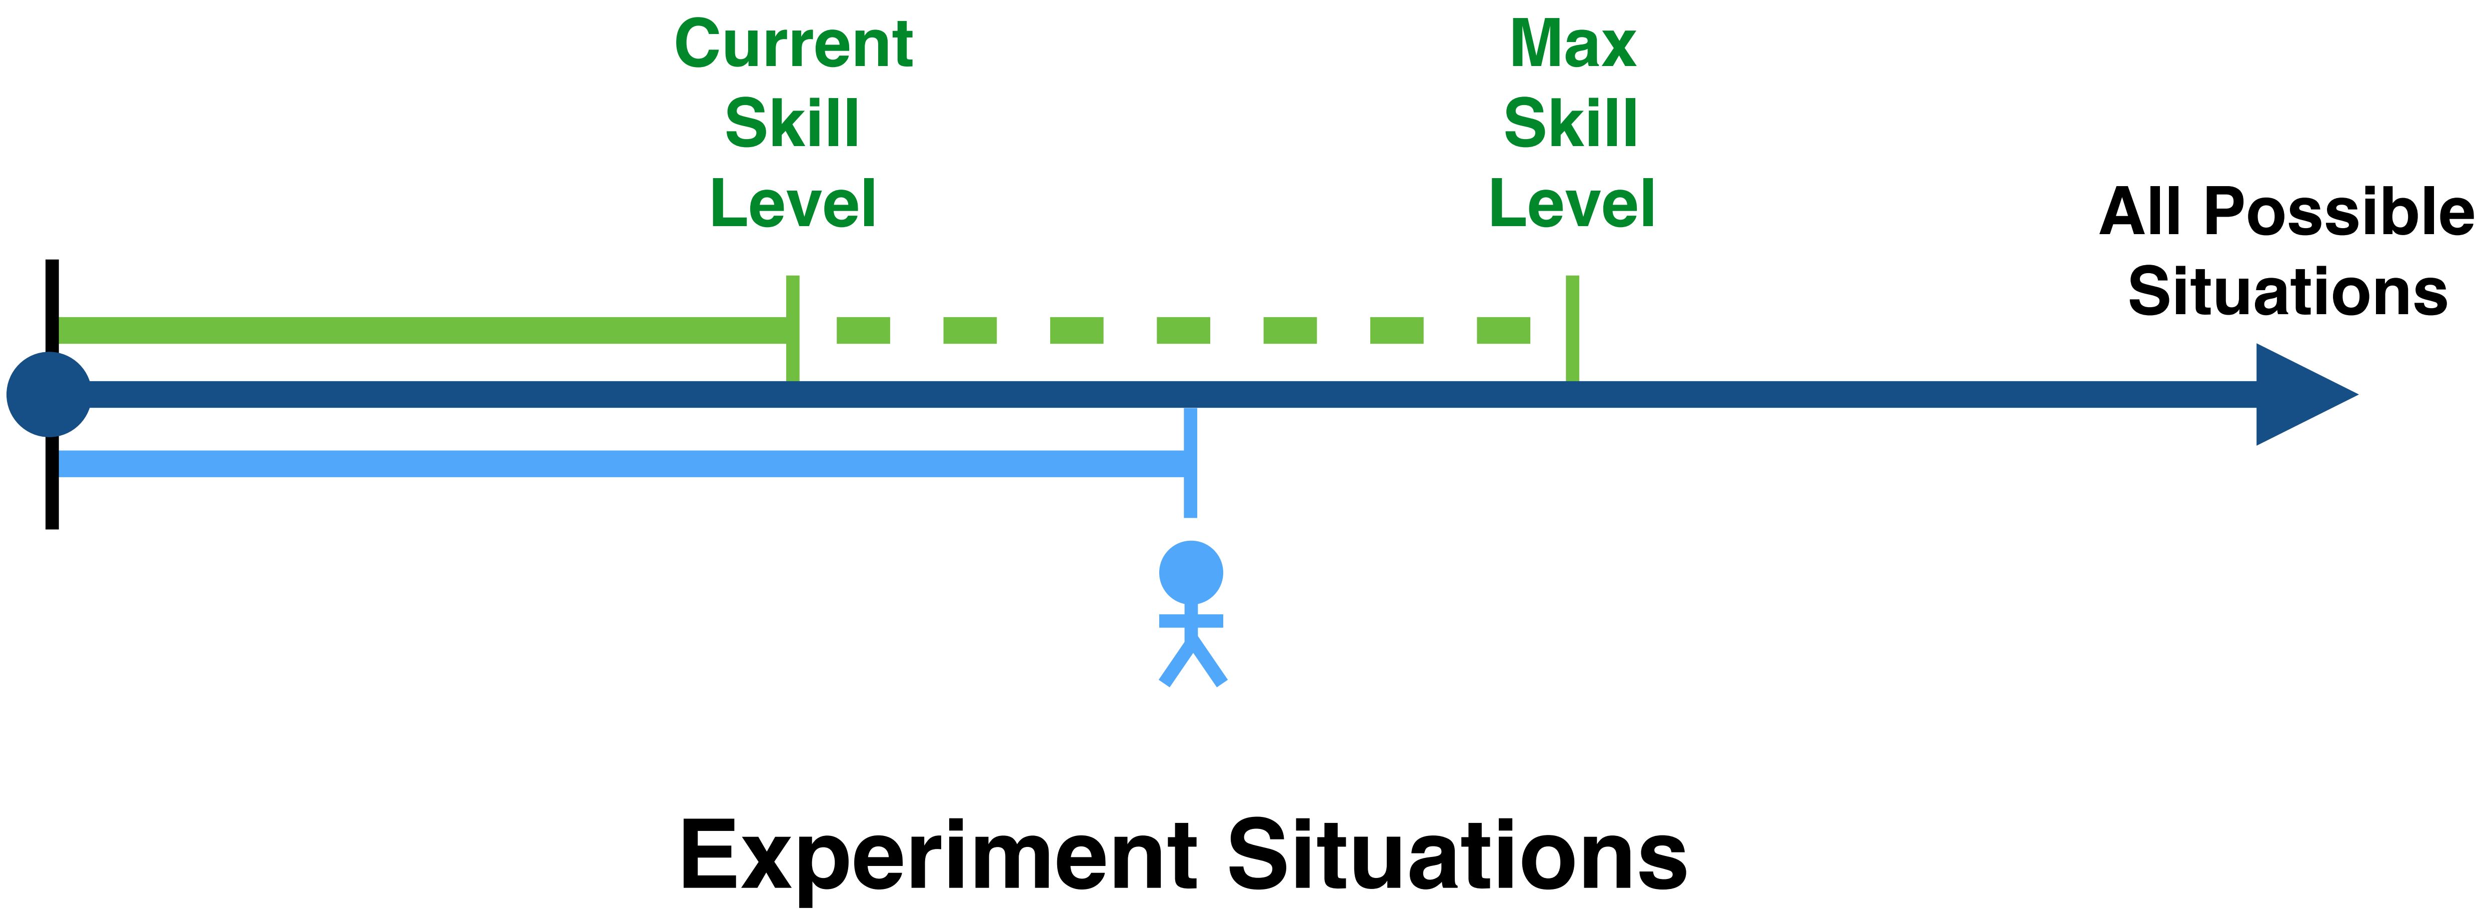

“Students”

Yevgen
Chebotar

Karol
Hausman

Zhe
Su

Samuel
Clarke

Conclusion

- Learn skills that exploit manipulations' mode structure
  - Goals and errors match mode transitions for more robust skills
- Selecting and initialising skill
  - Consider contacts for setting goals and checking preconditions
- Monitoring skills during executions
  - Detect salient sensory events during mode
- Verifying skill outcomes
  - Verify mode transitions based on changing dynamics
  - Use self-supervised
- Future challenge: Learn to predict what can be learned