On perceptual representations

and how they interact with actions and physical representations

Jeannette Bohg – Interactive Perception and Robot Learning Lab – Stanford


FIRE CHIEF
Jenga

Exploiting Multi-Modality

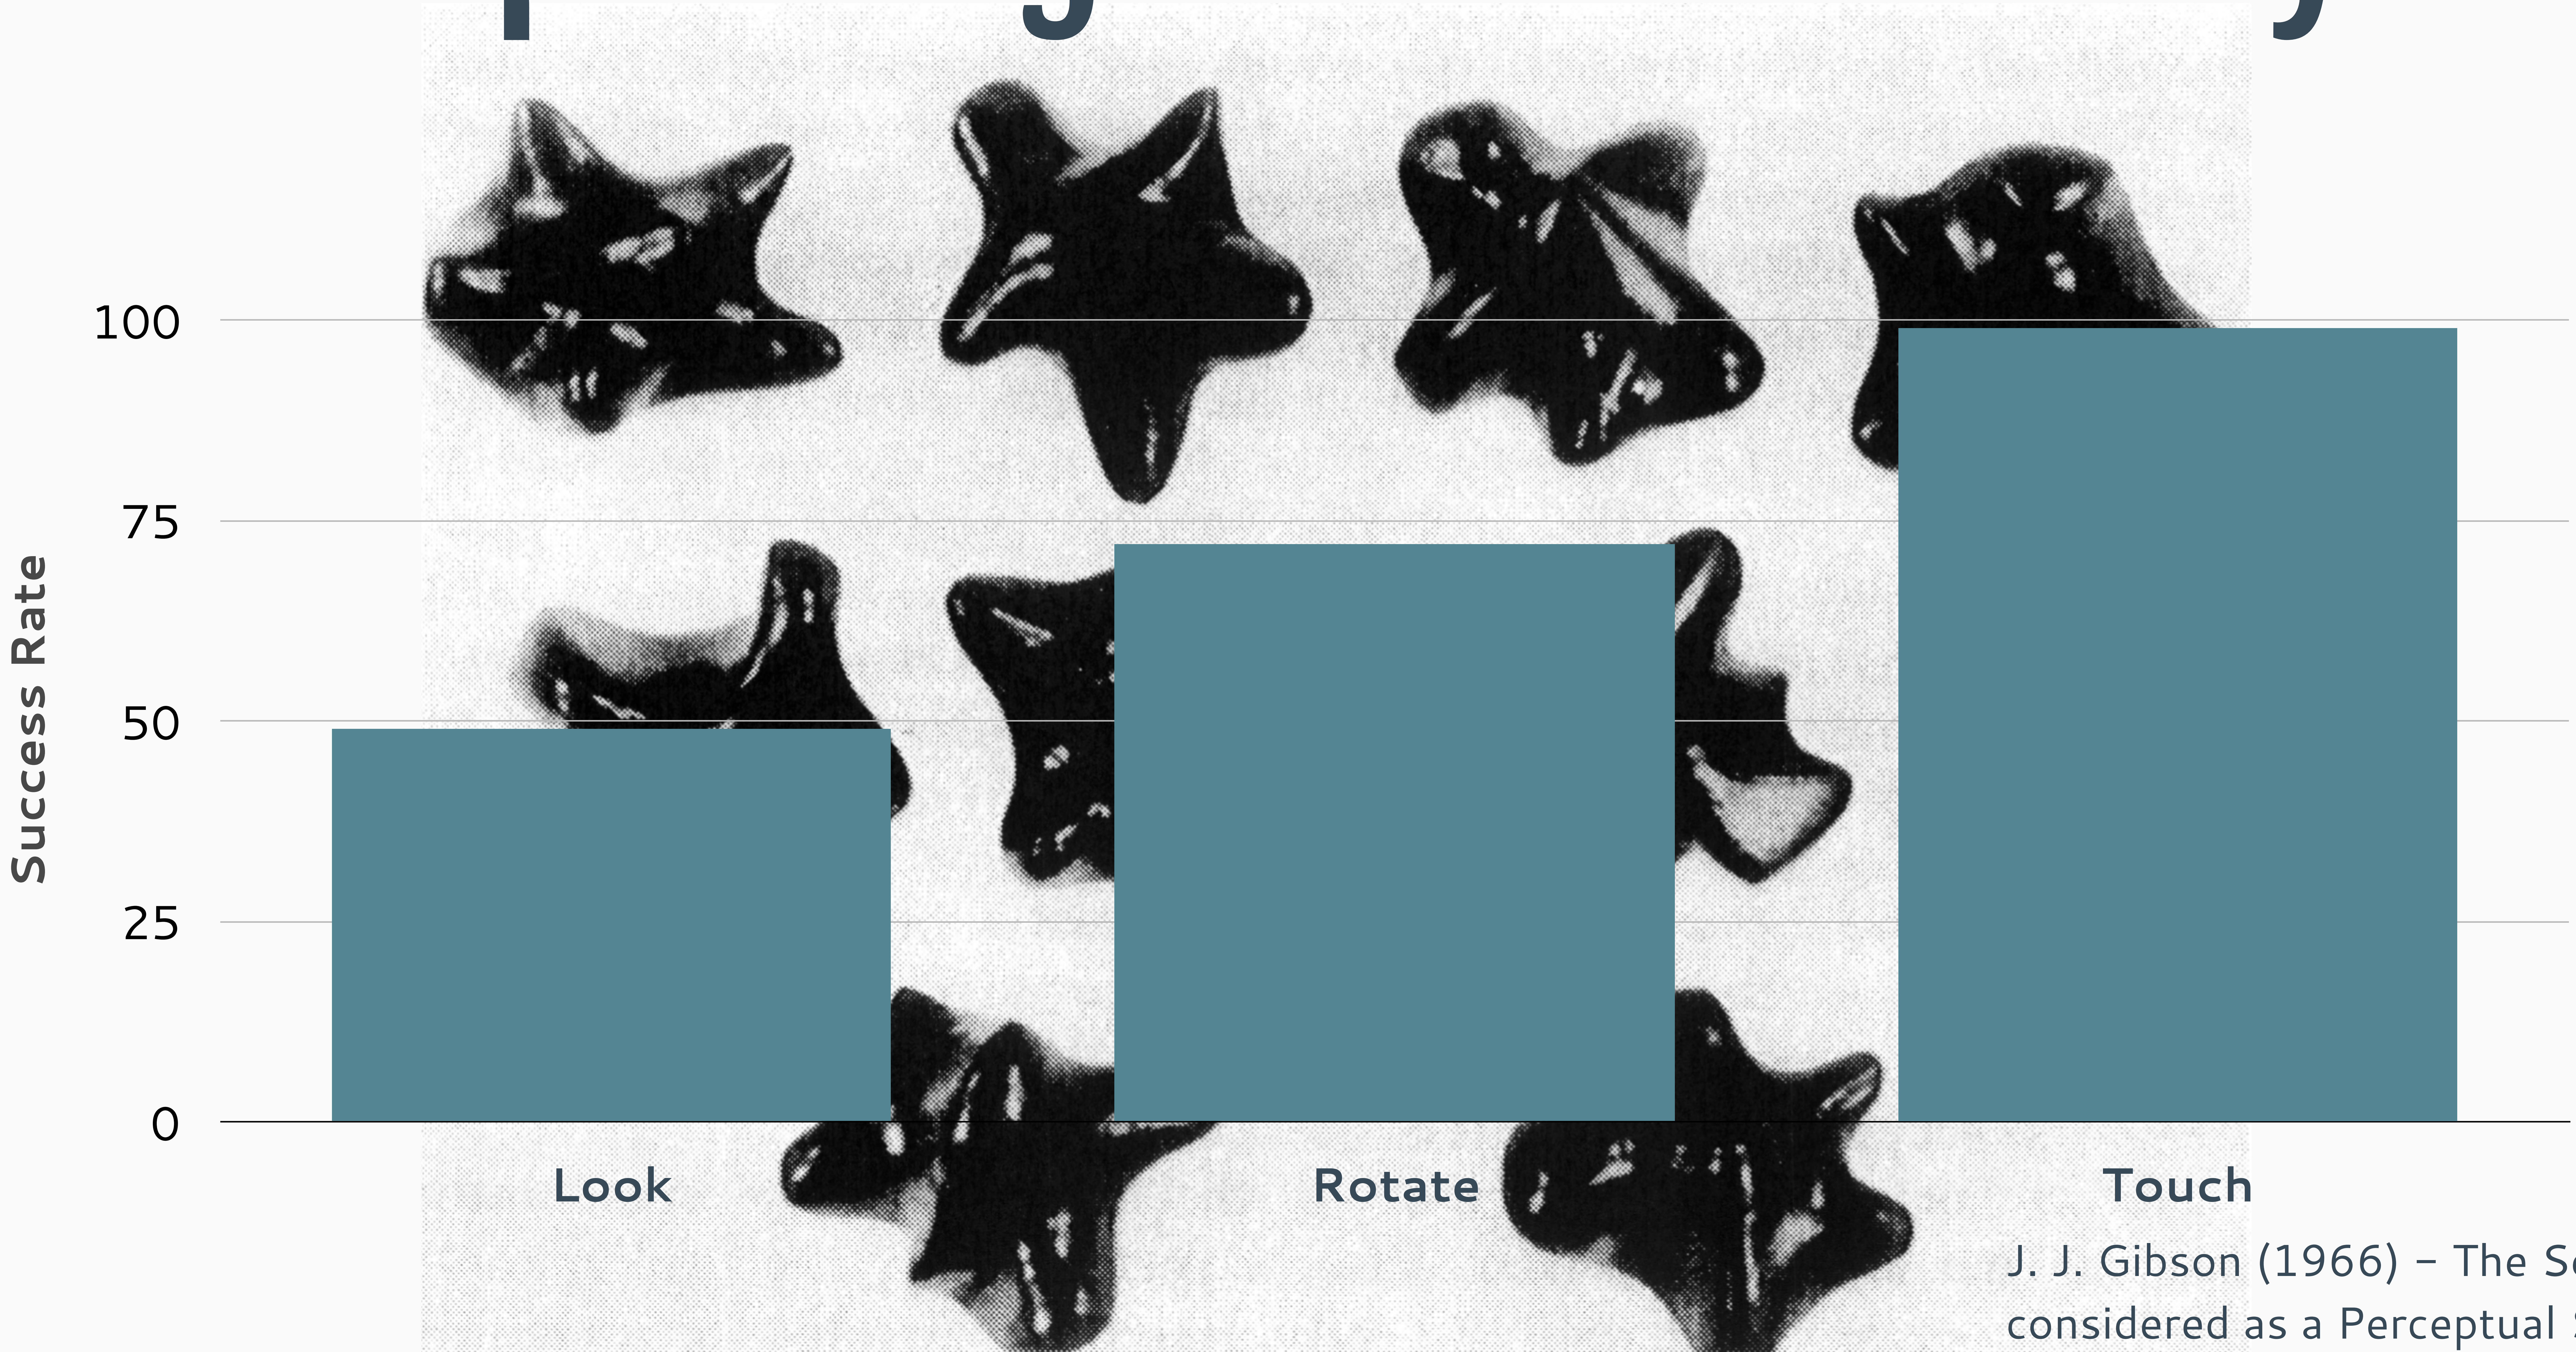

J. J. Gibson (1966) – The Senses considered as a Perceptual System.

Accumulation over Time

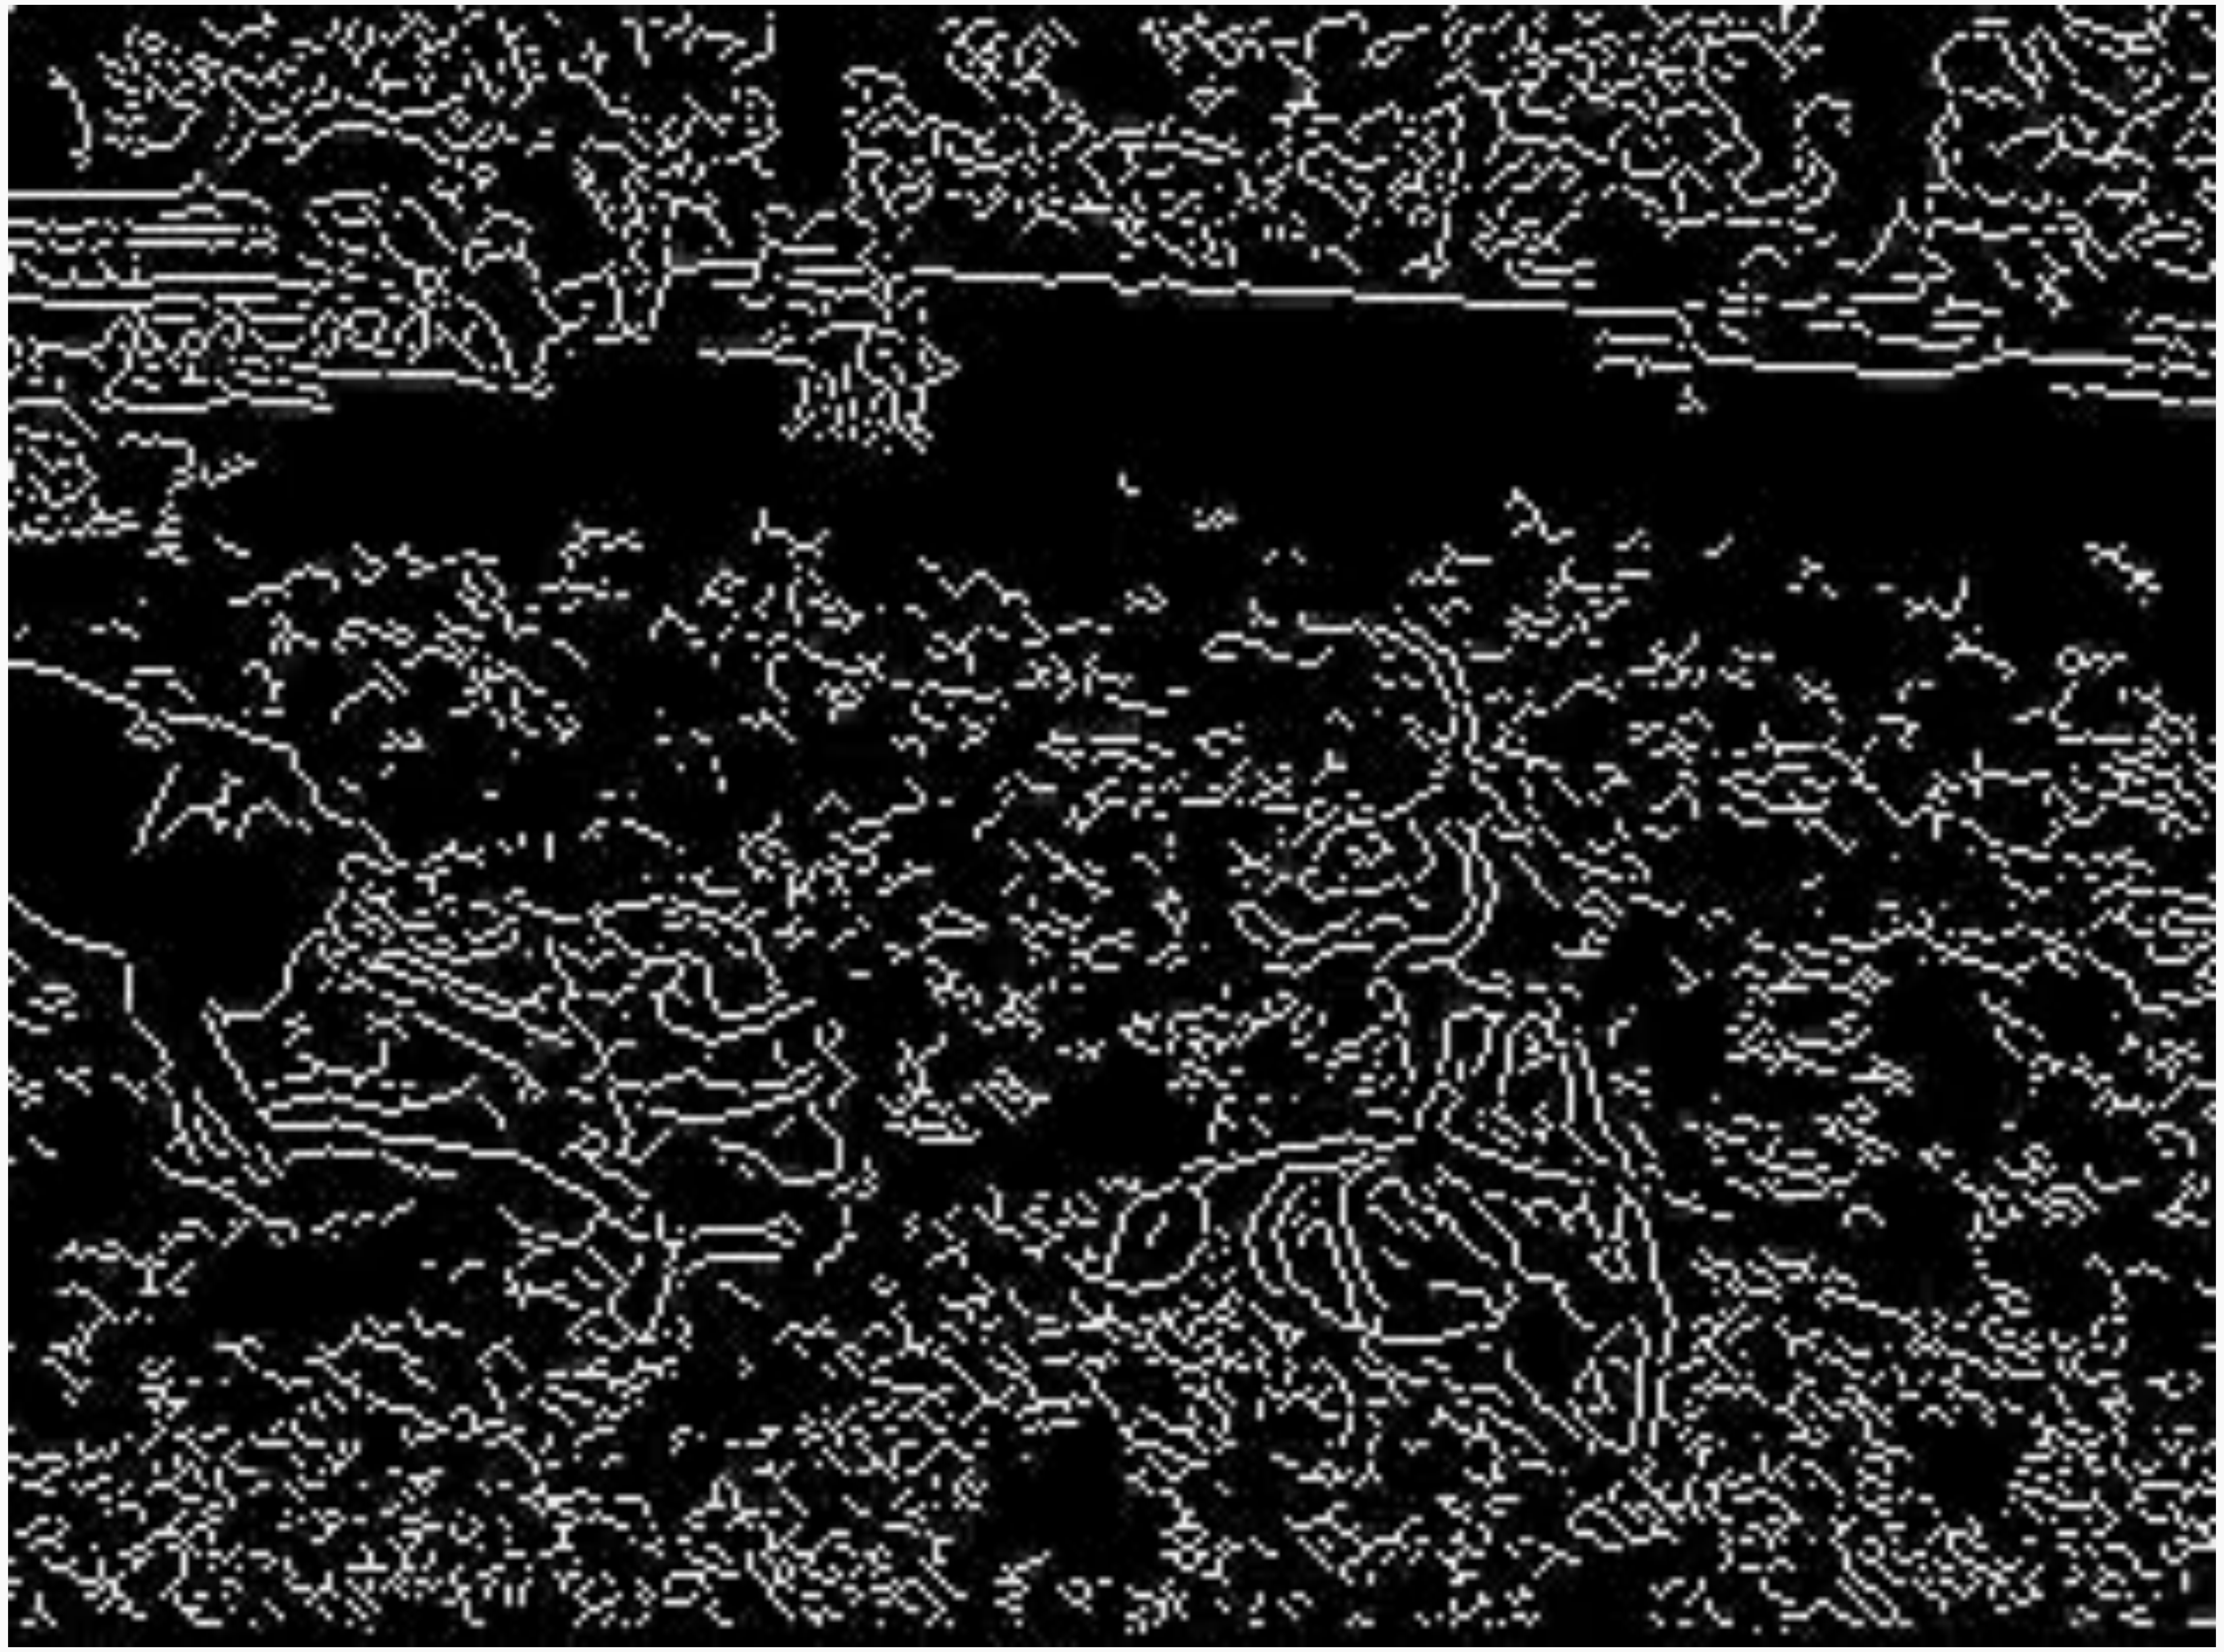

Thanks to Octavia Camps at Northeastern University, Boston

Concurrency of Egomotion and Sensing

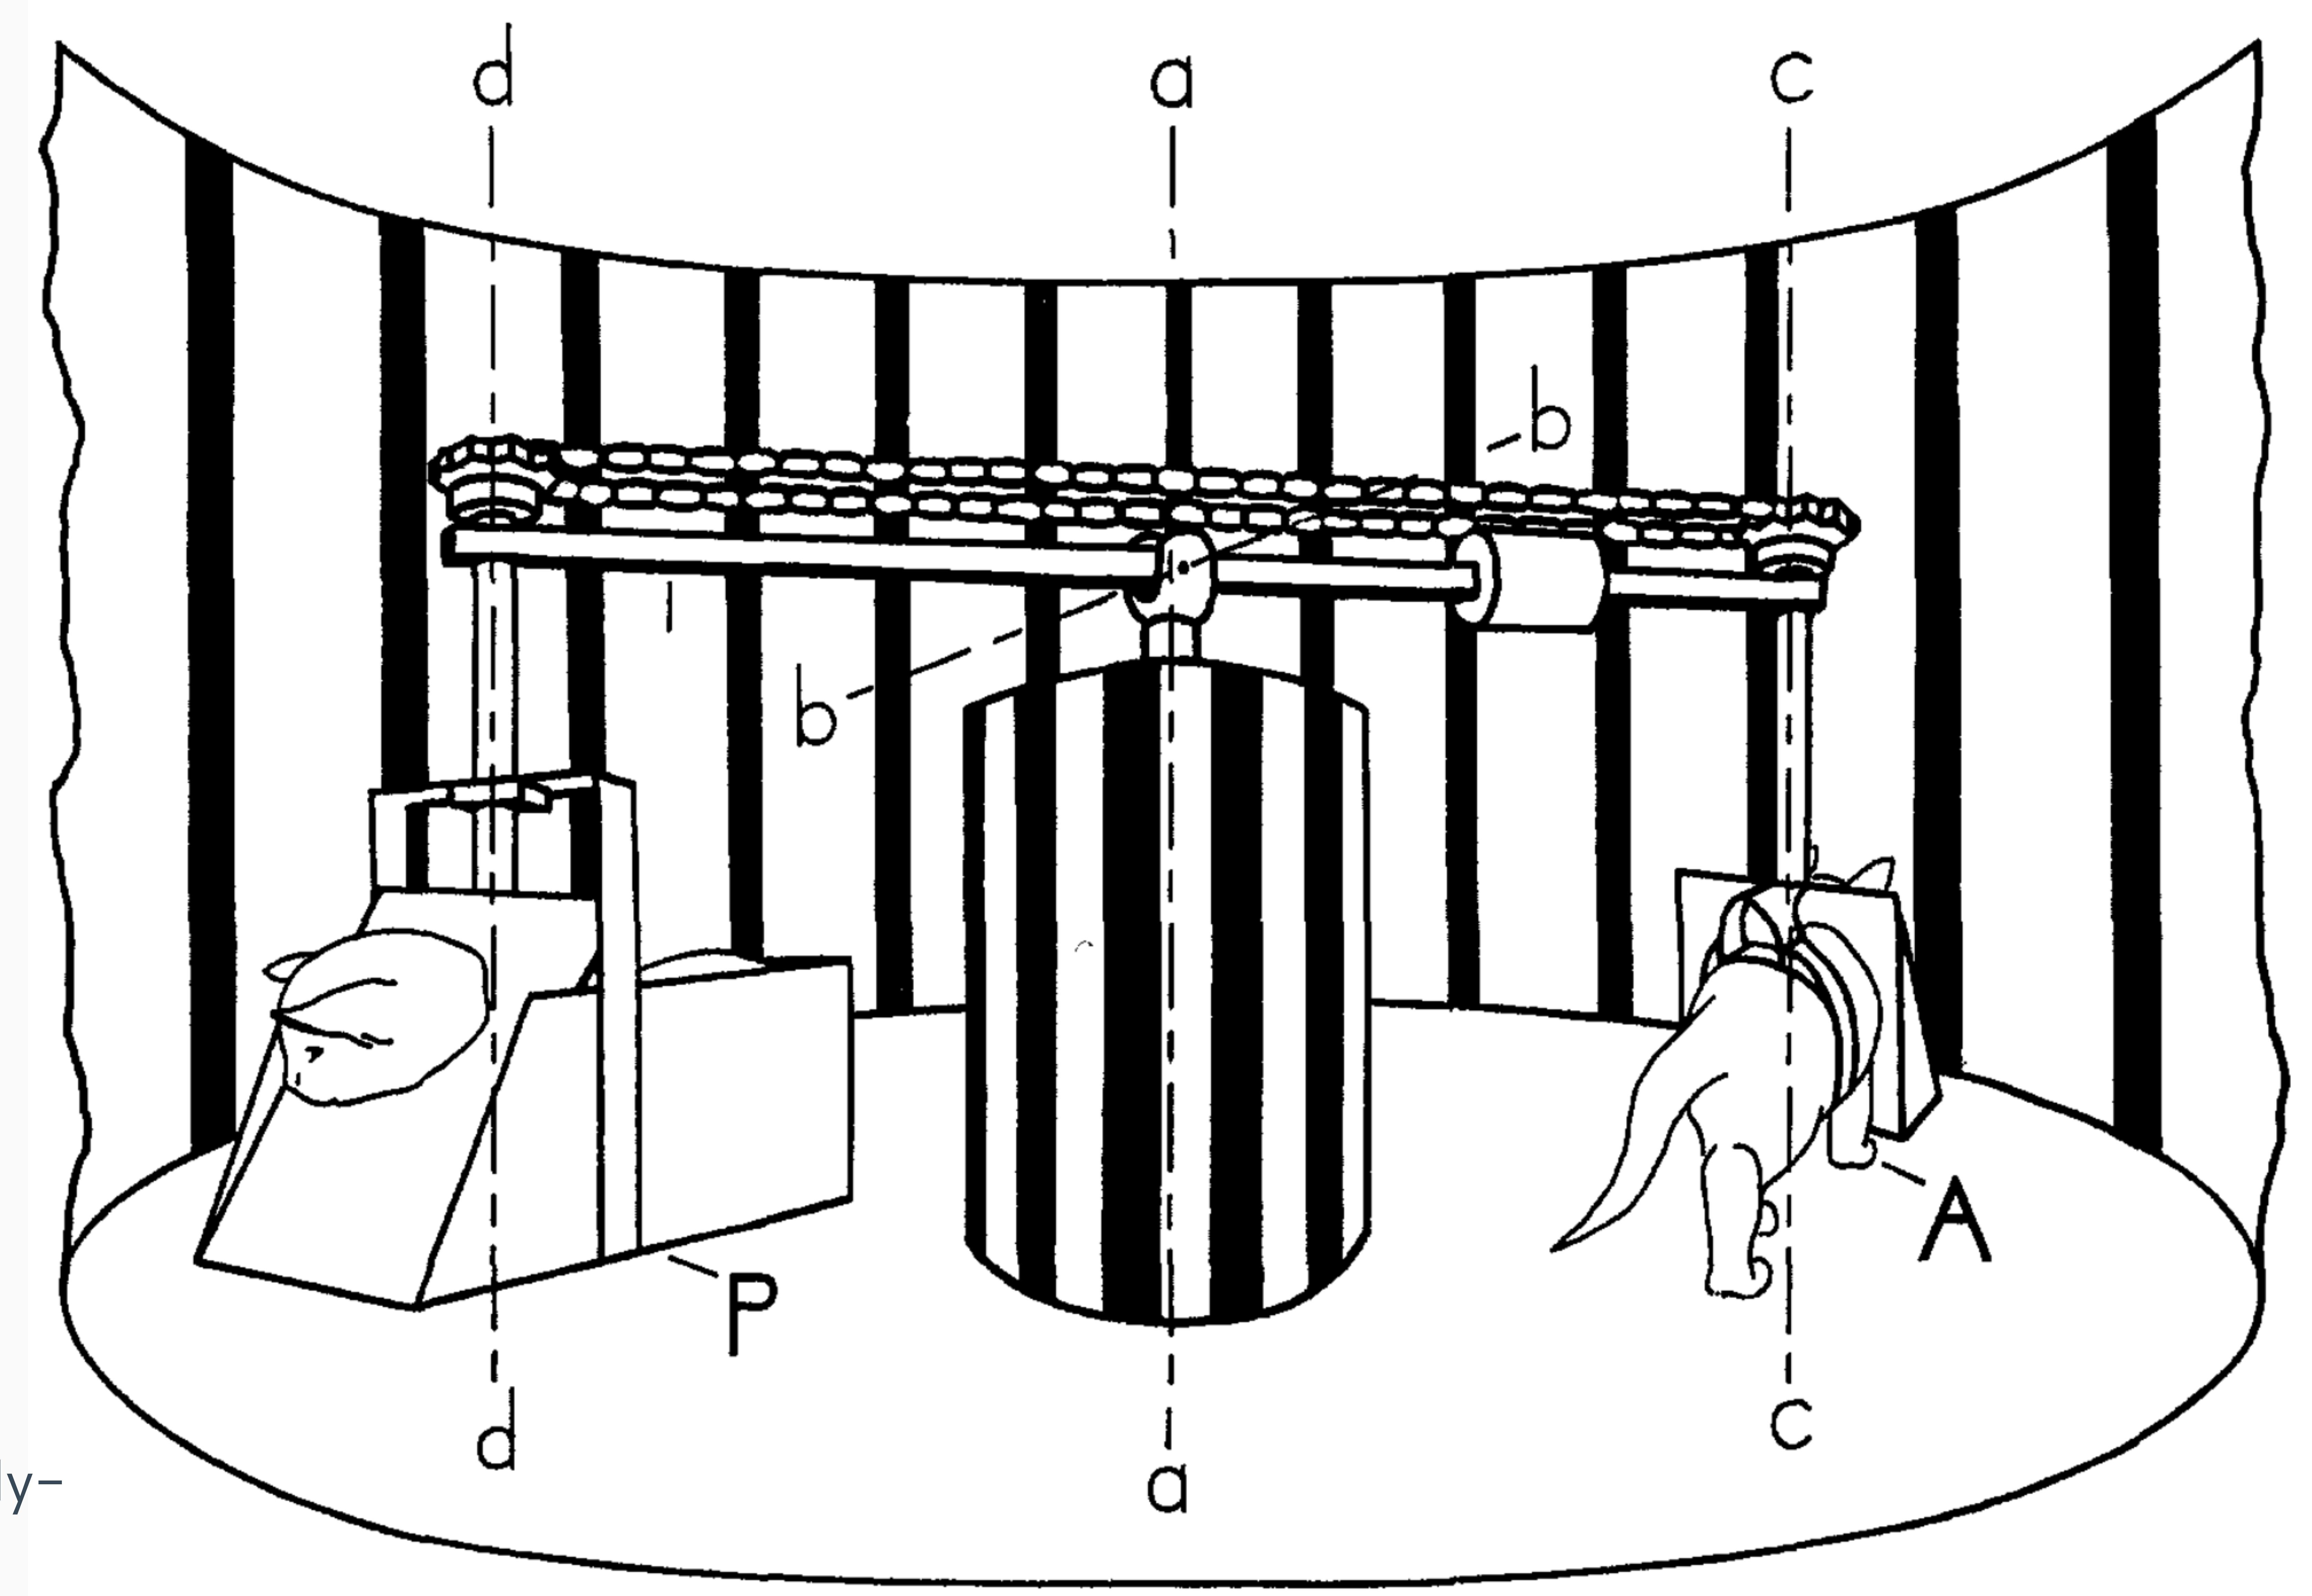

Held and Hein (1963). Movement-Produced Stimulation in the Development of Visually-Guided Behaviour

Interactive Perception

$$\underset{\text{Sensory Data}}{S} \times \underset{\text{Actions}}{A} \times \underset{\text{Time}}{t}$$

Control

Interactive Perception – Leveraging Action in Perception and Perception in Action.
Bohg, Hausman, Sankaran, Brock, Kragic, Schaal and Sukhatme. TRO '17.

Predictive Model

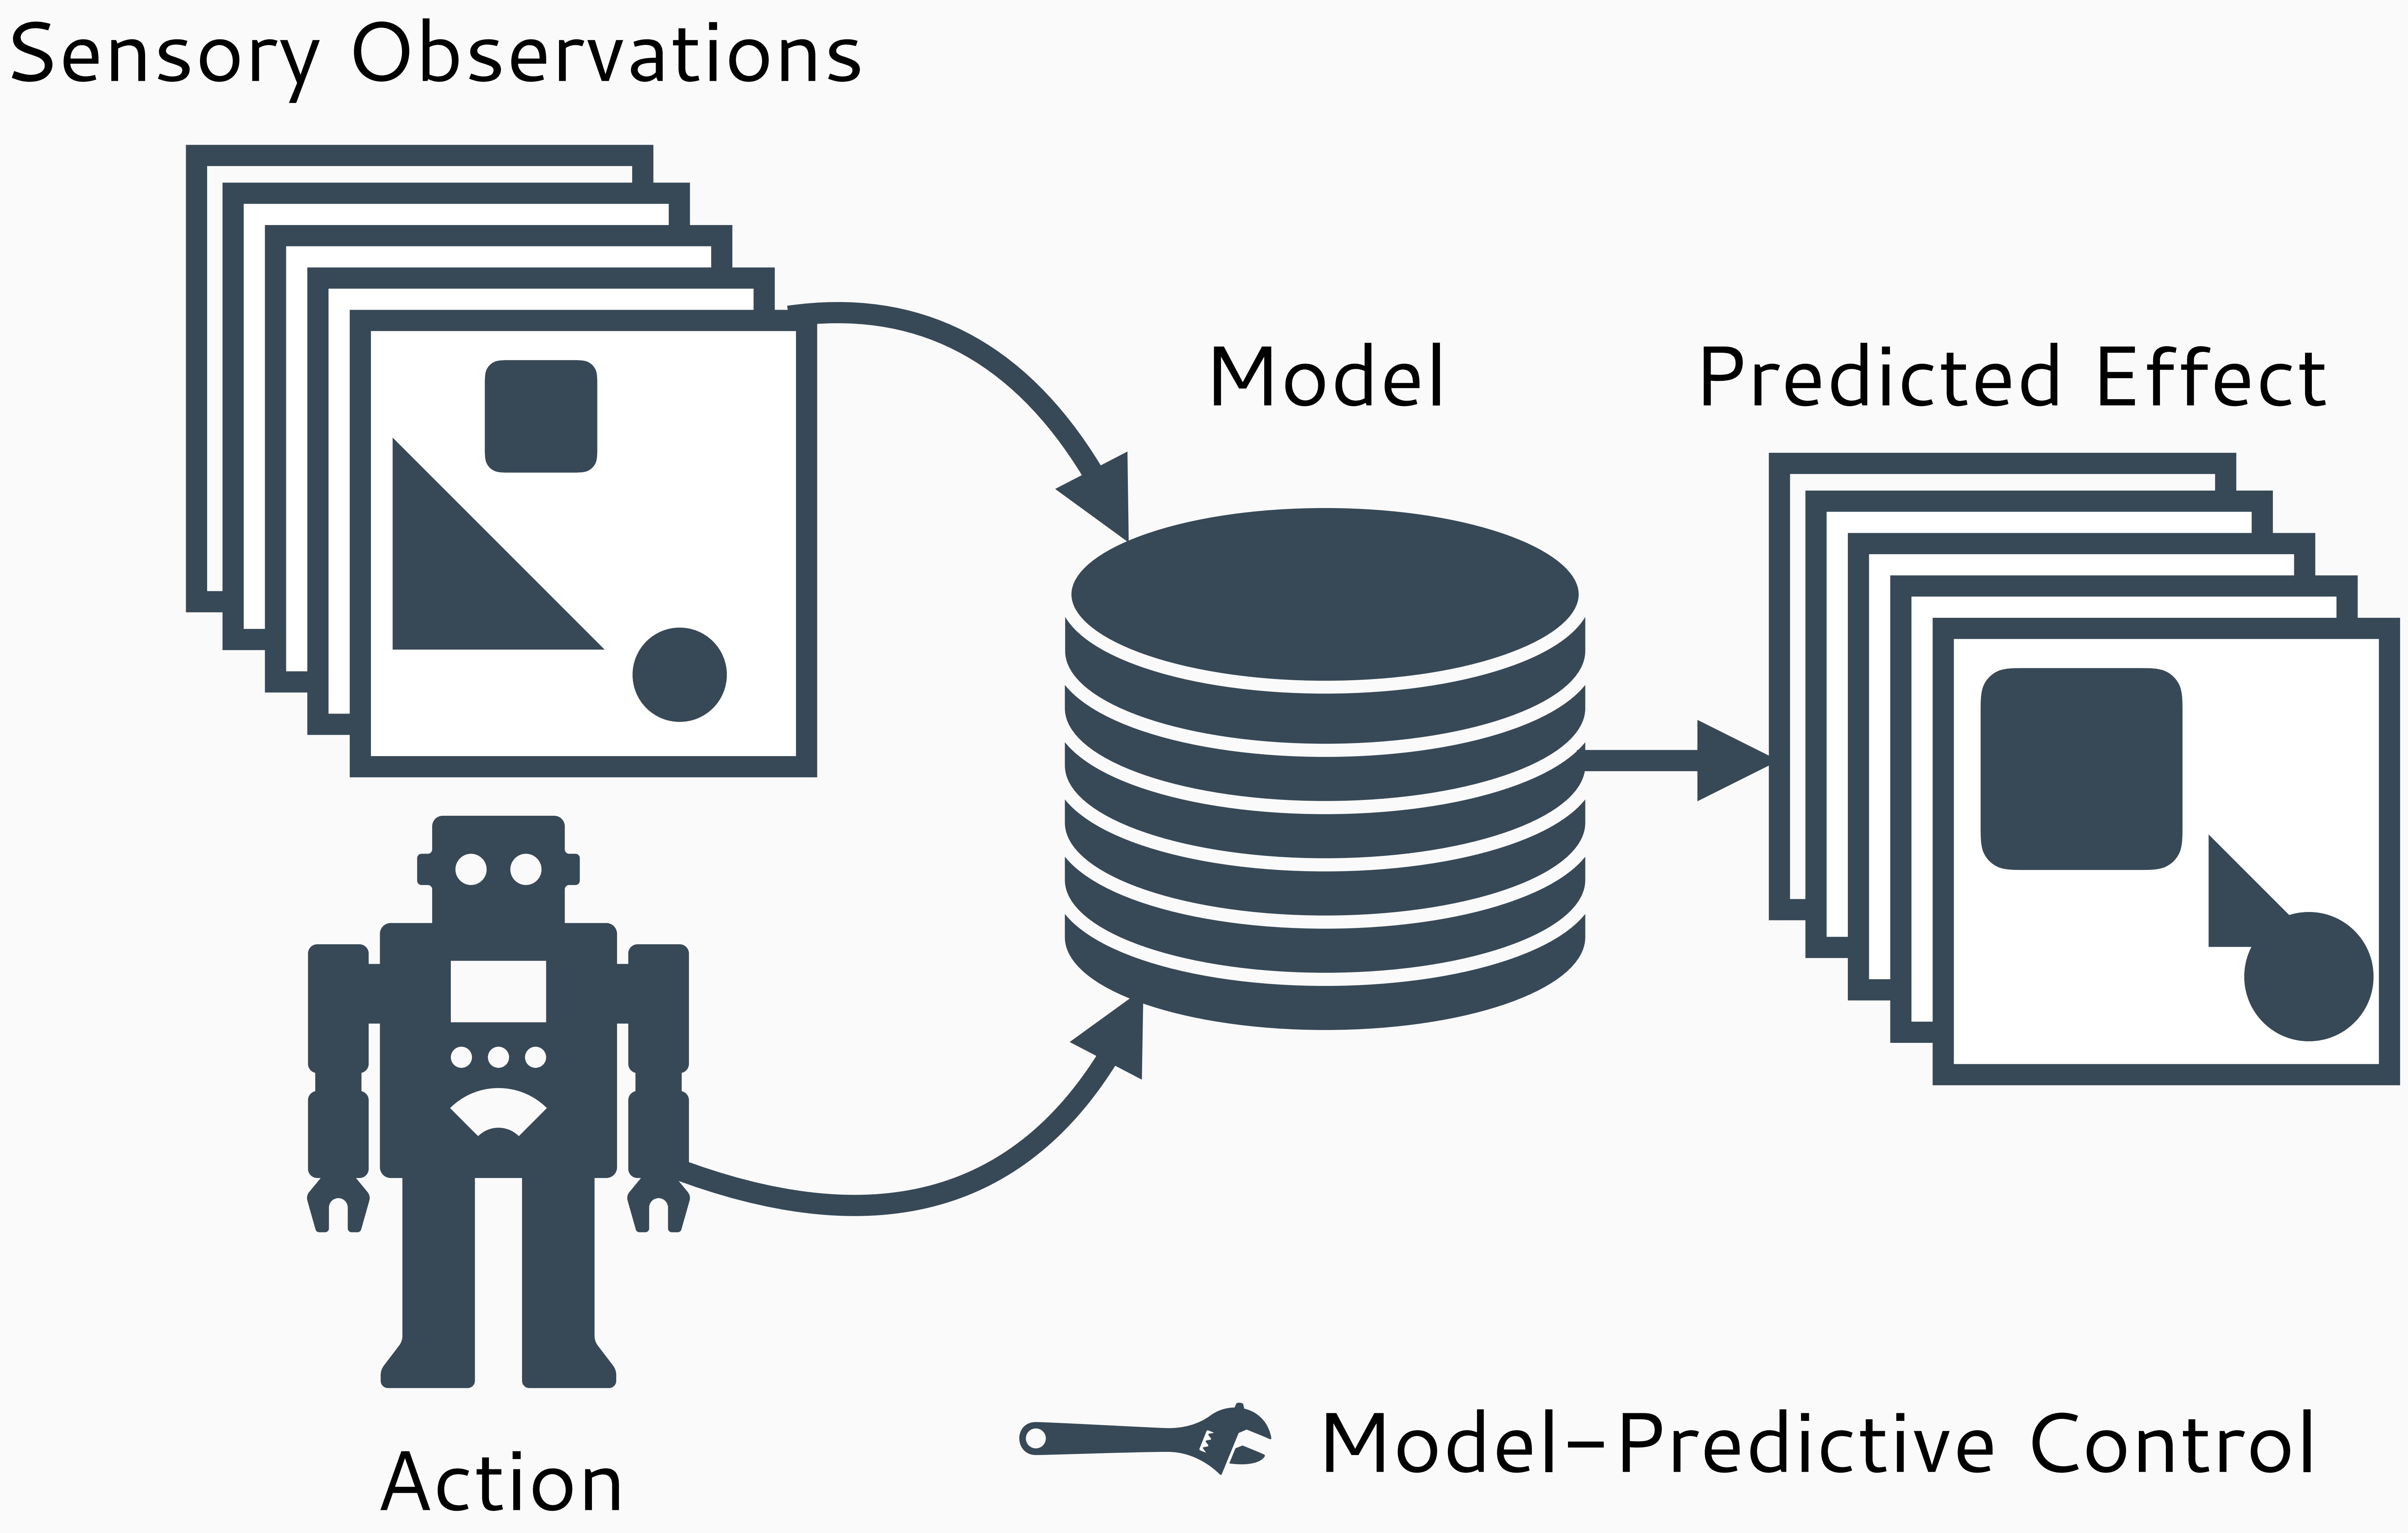

Modelling Physics

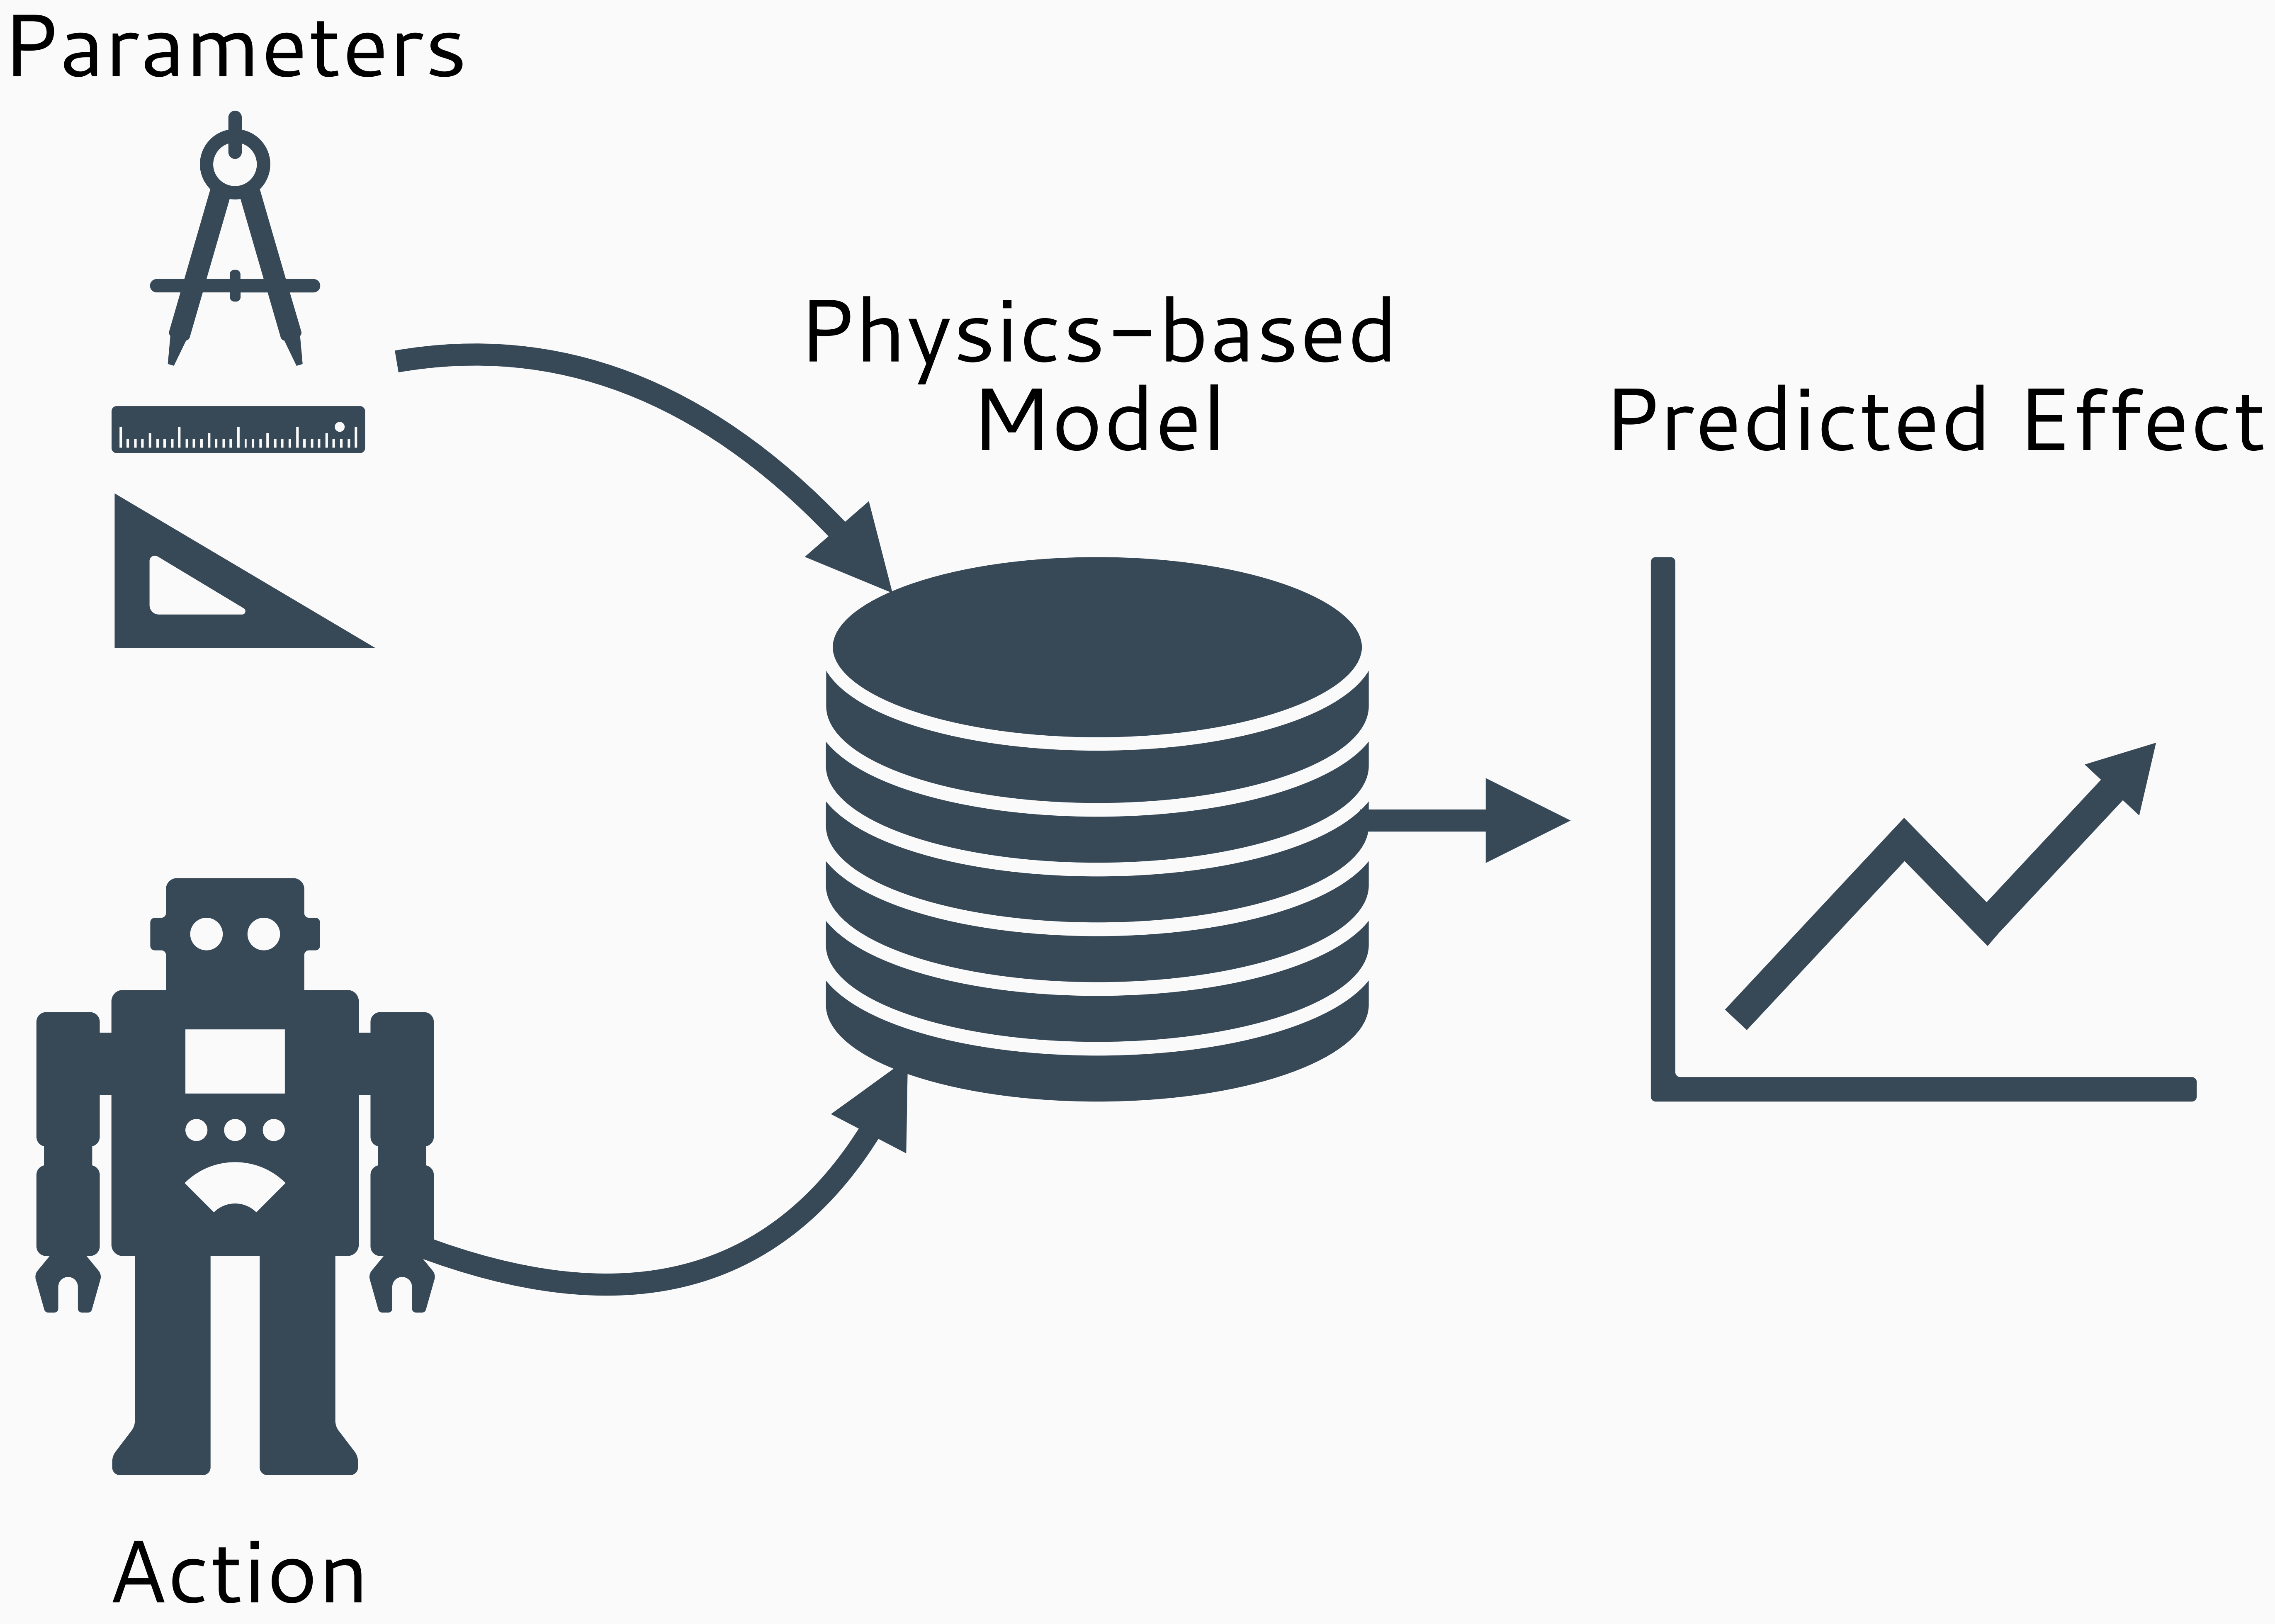

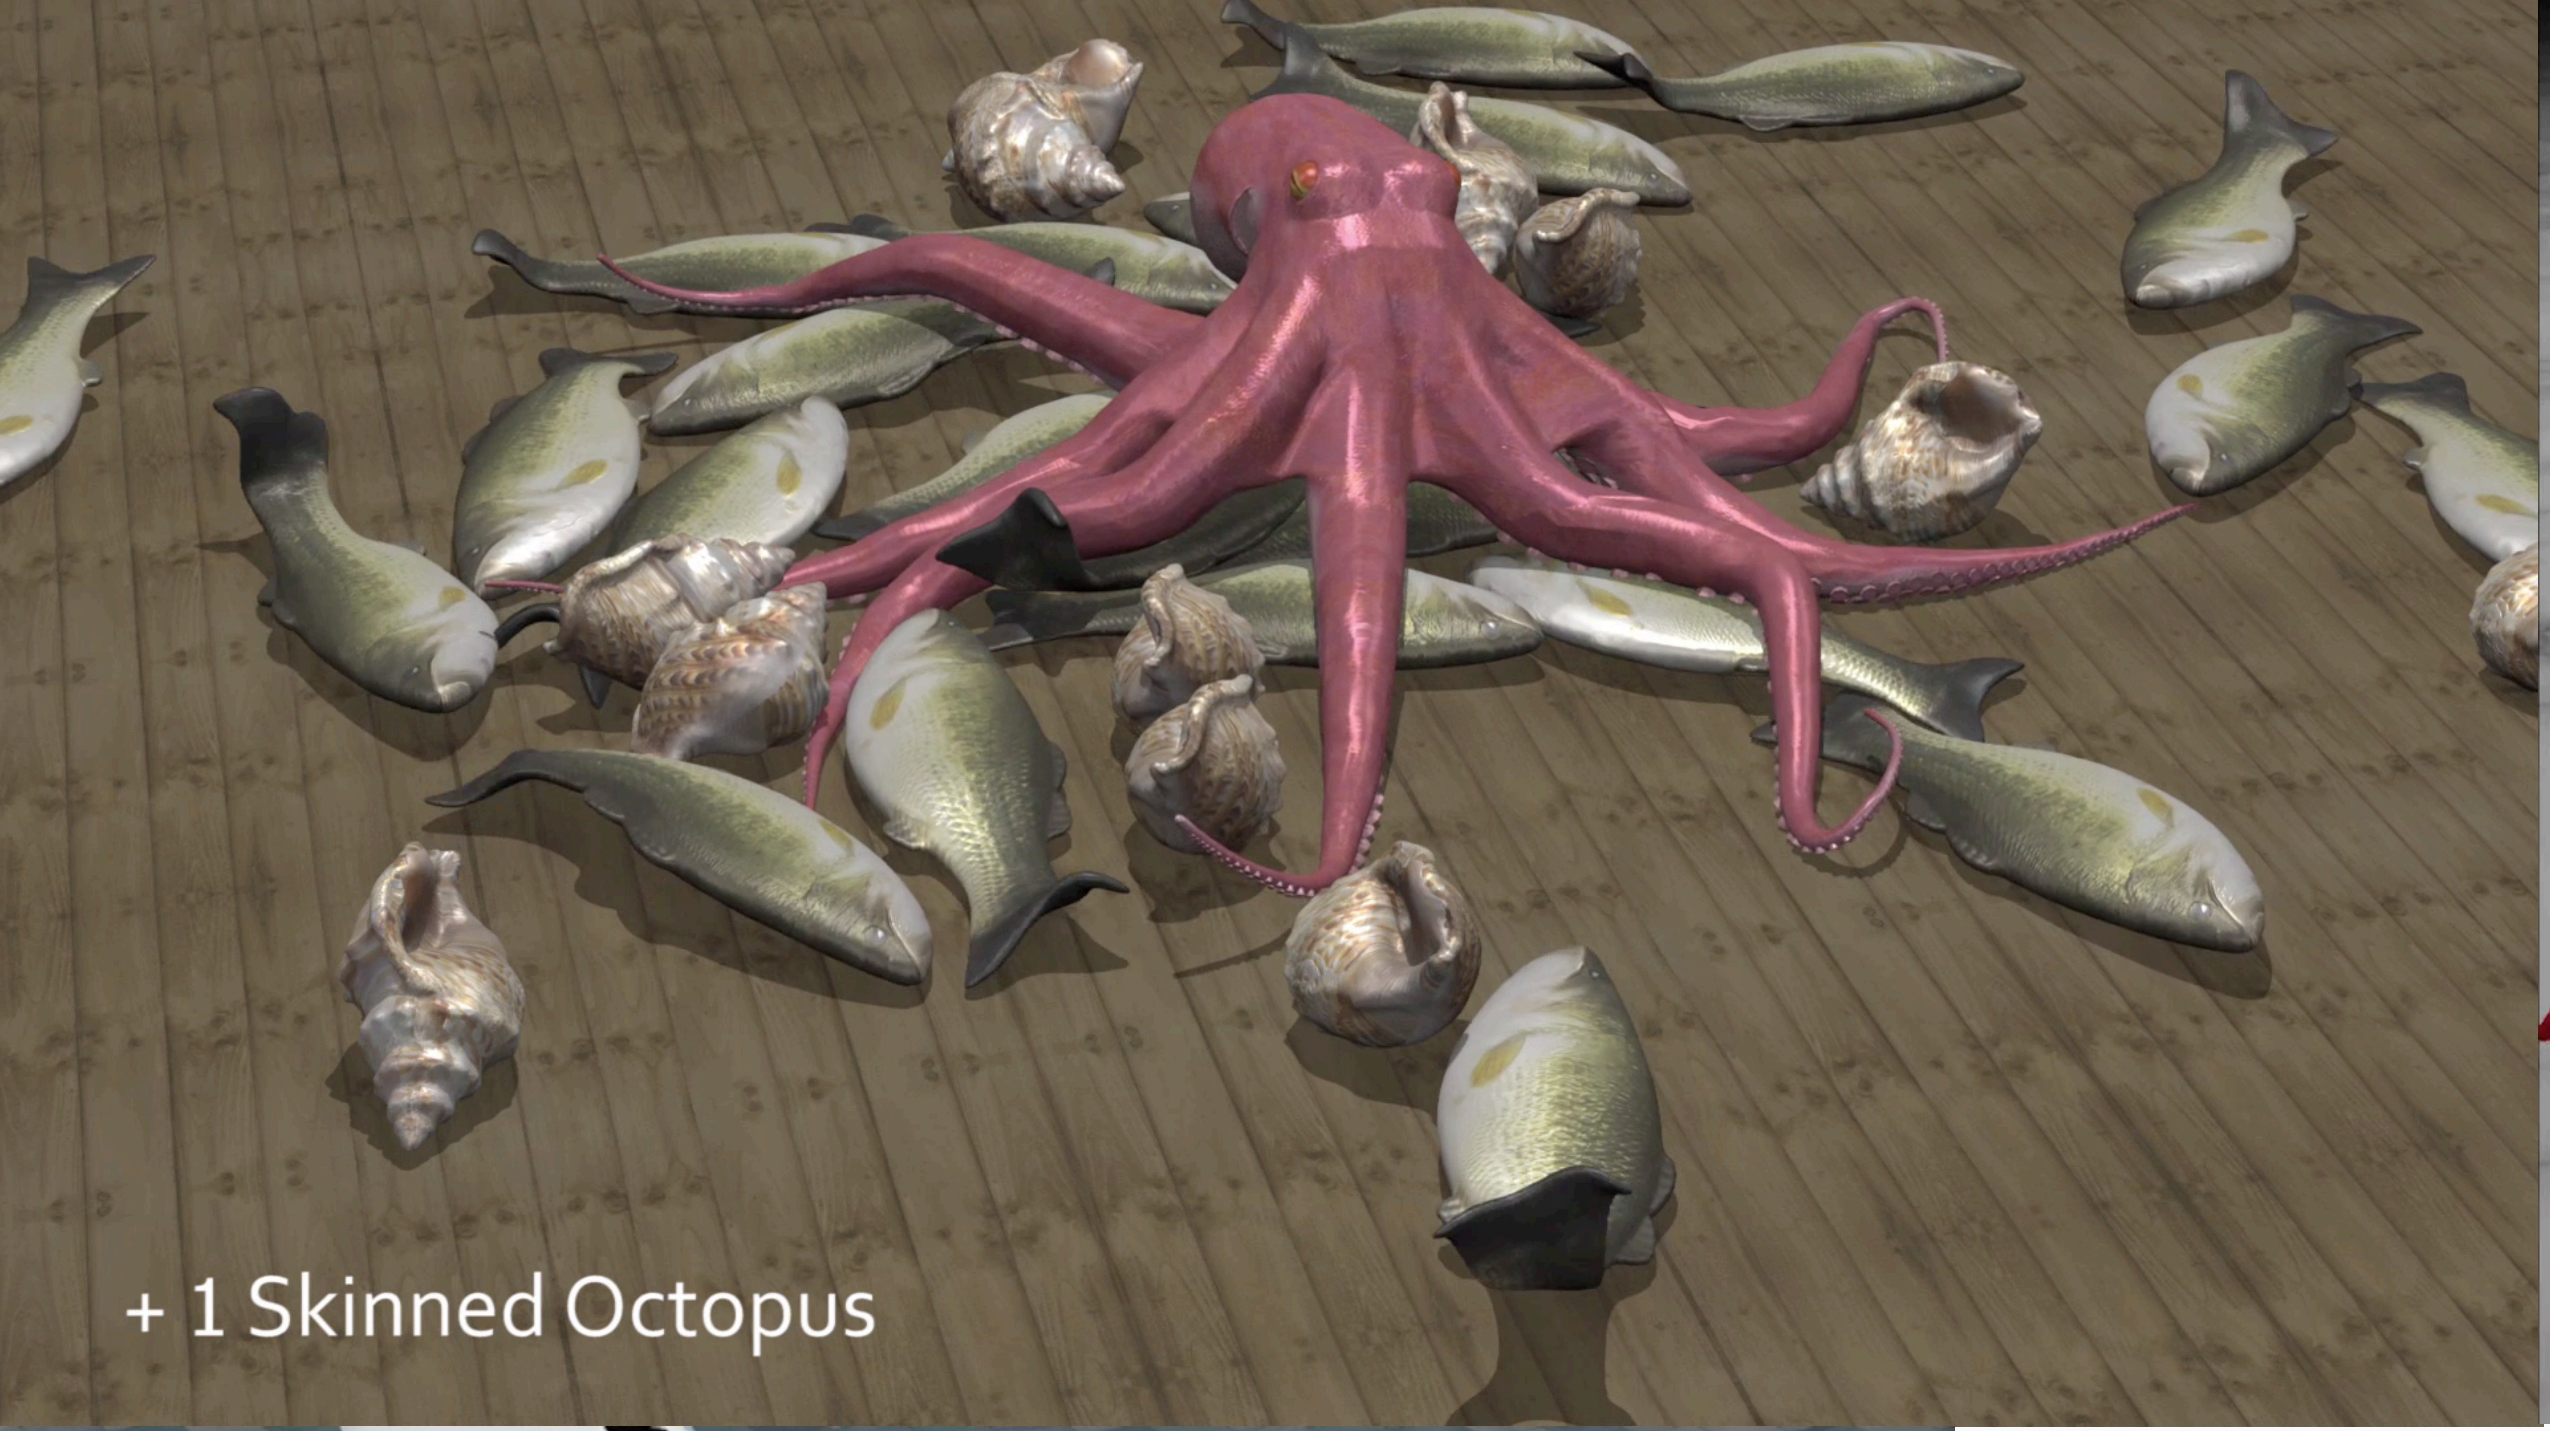

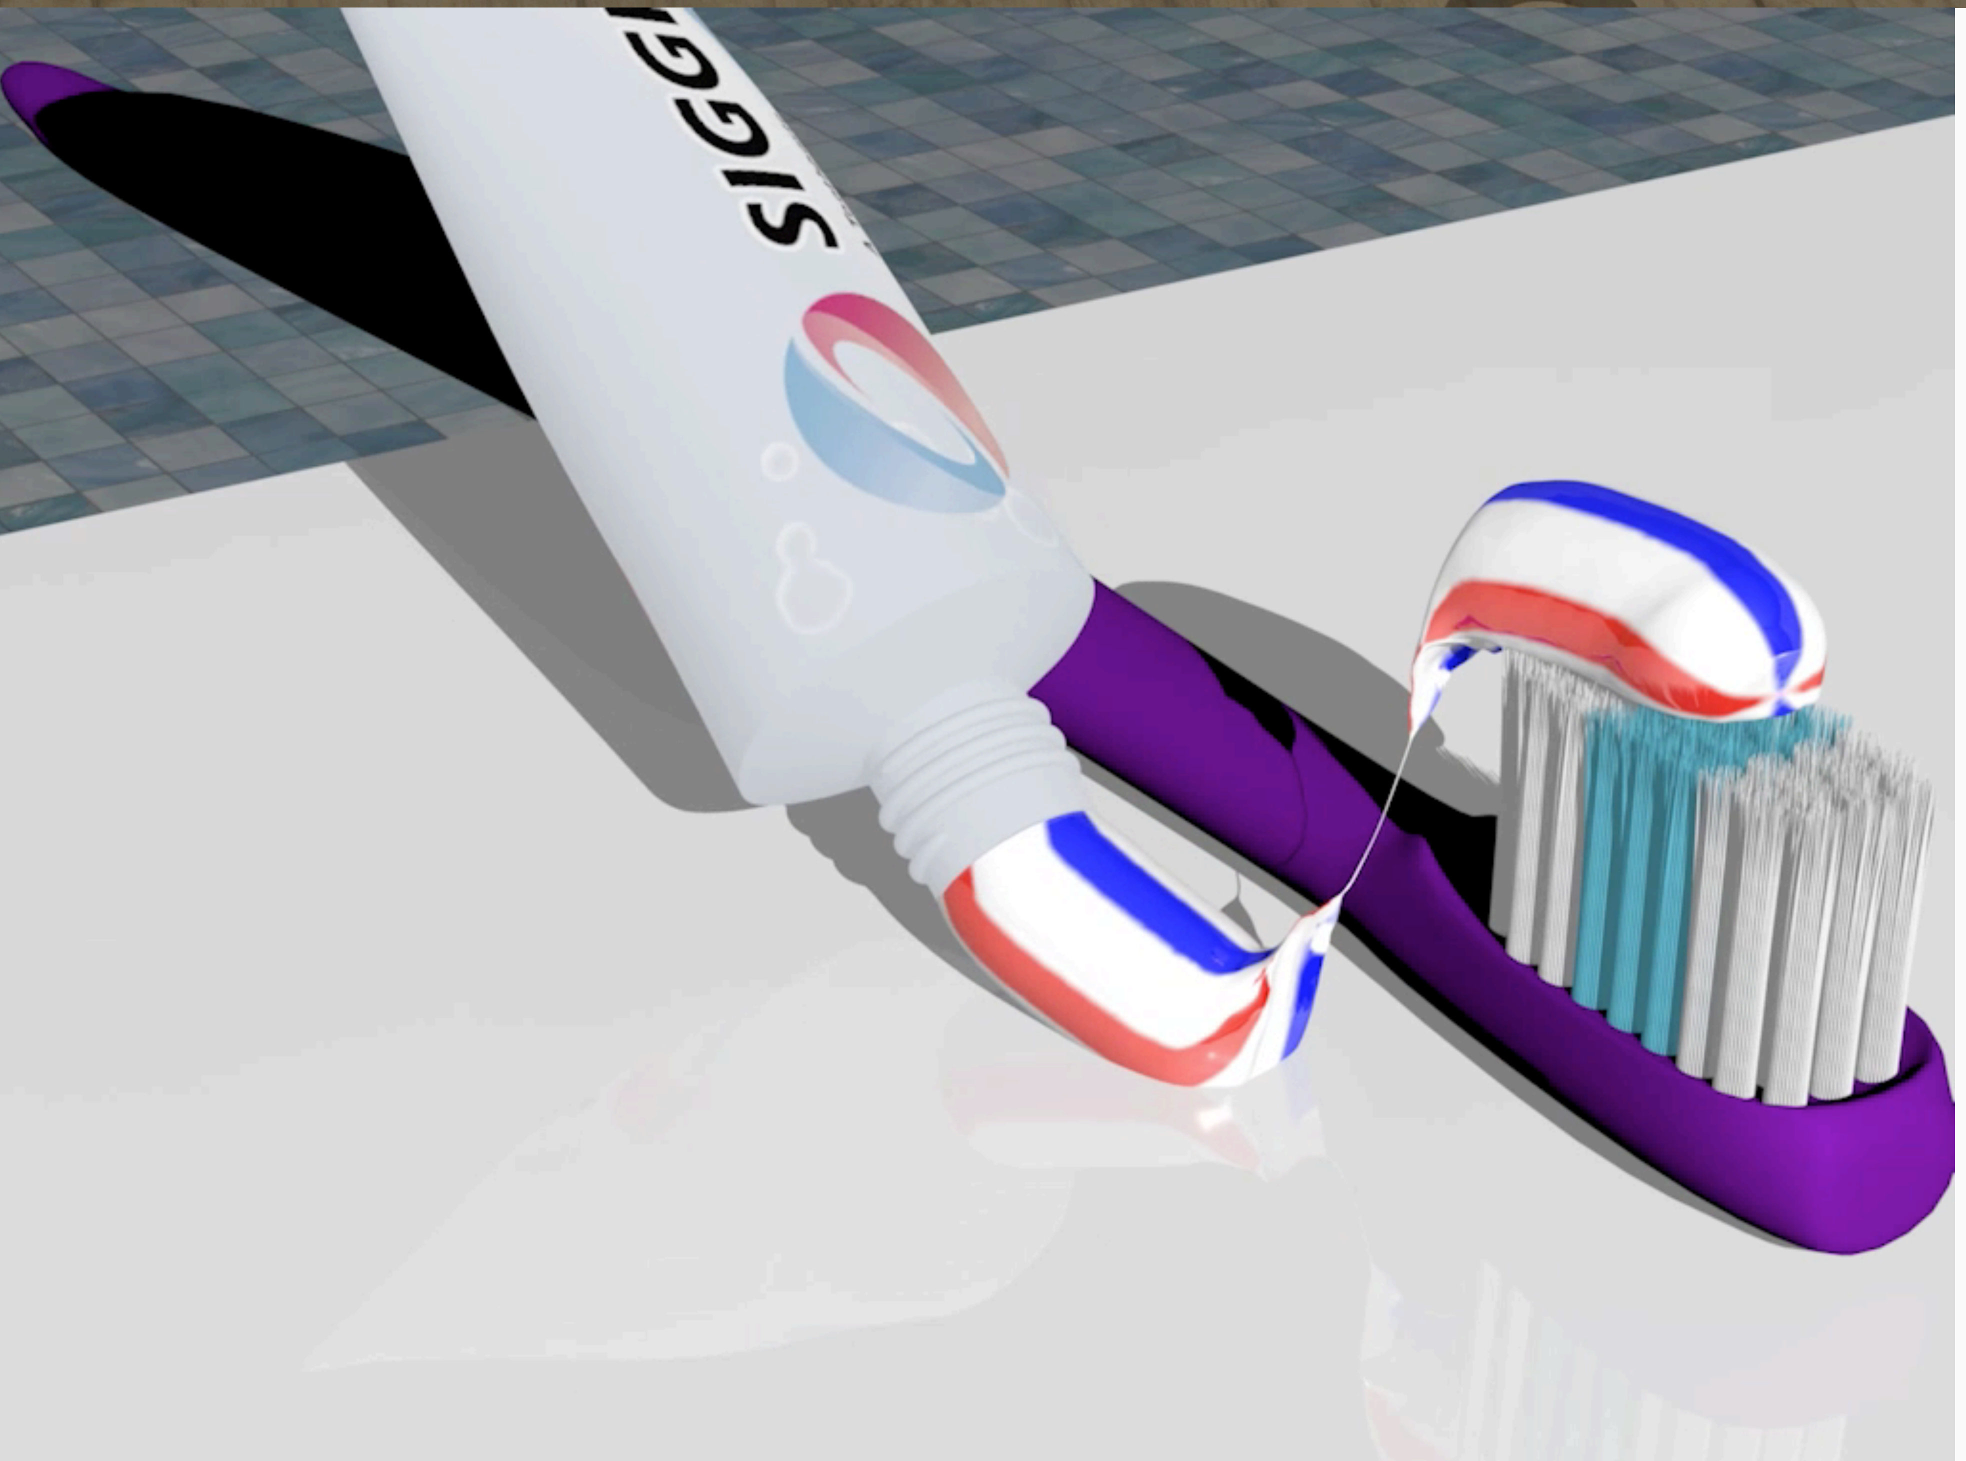

Video Credit: Ron Fedkiw@Stanford

Example:

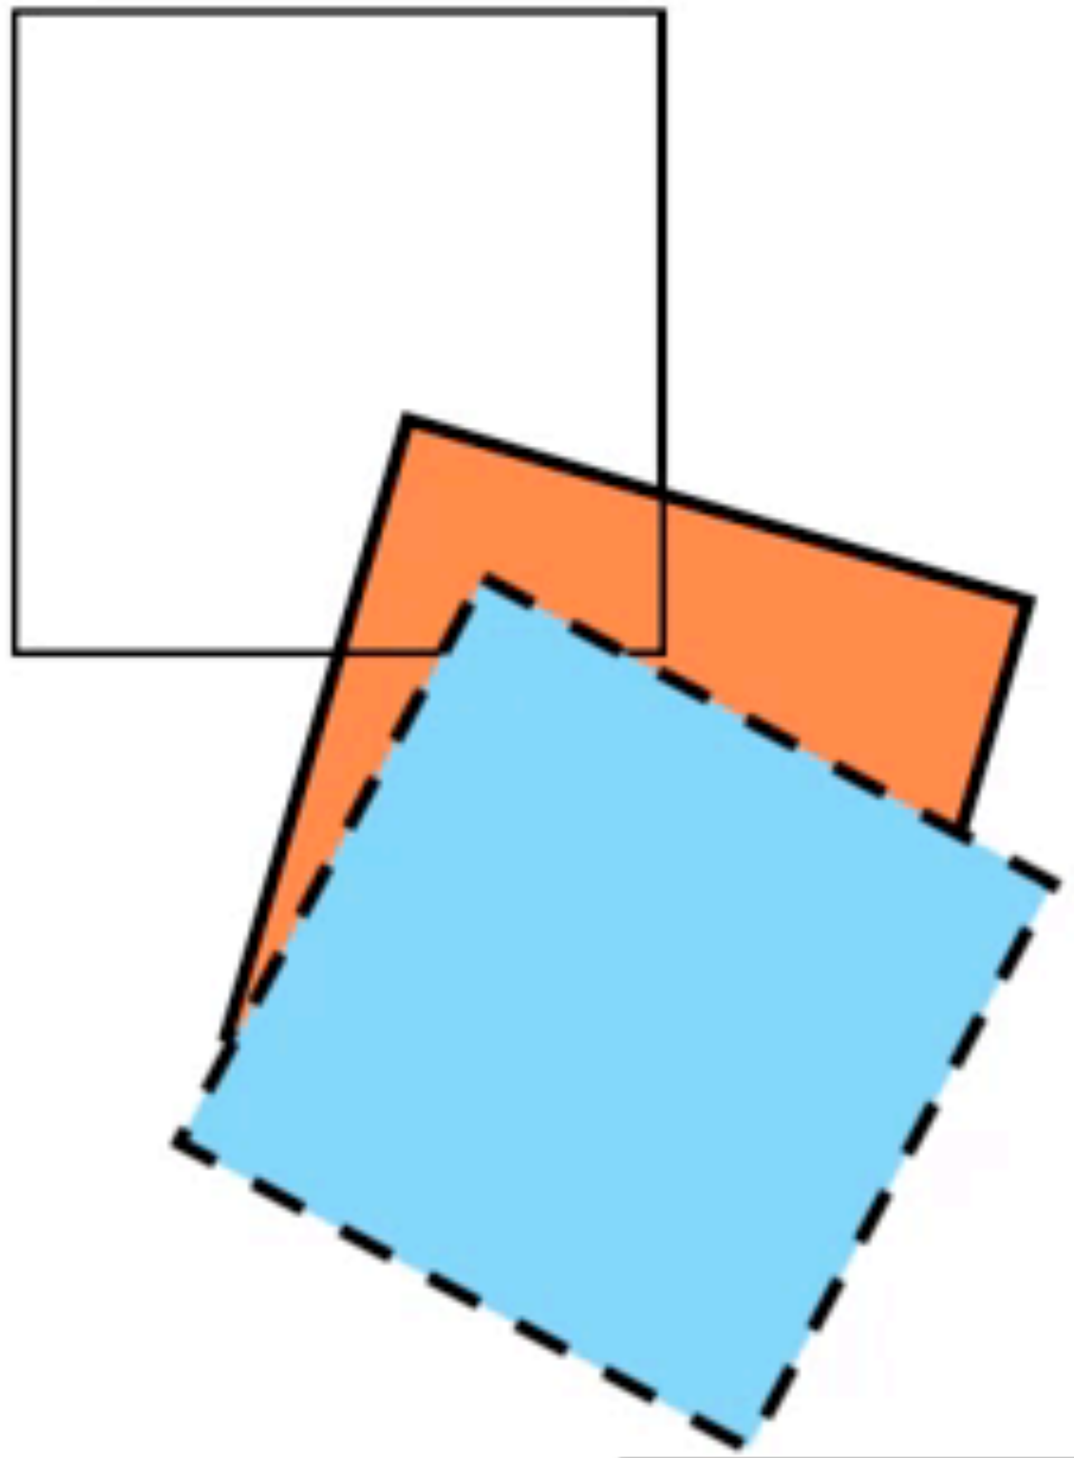

BIG difference!

Real motion

Simulated motion

(Lynch 1992)

More than a Million Ways to Be Pushed. A High-Fidelity Experimental Dataset of Planar Pushing. Yu et al. IROS 2016.

Challenges

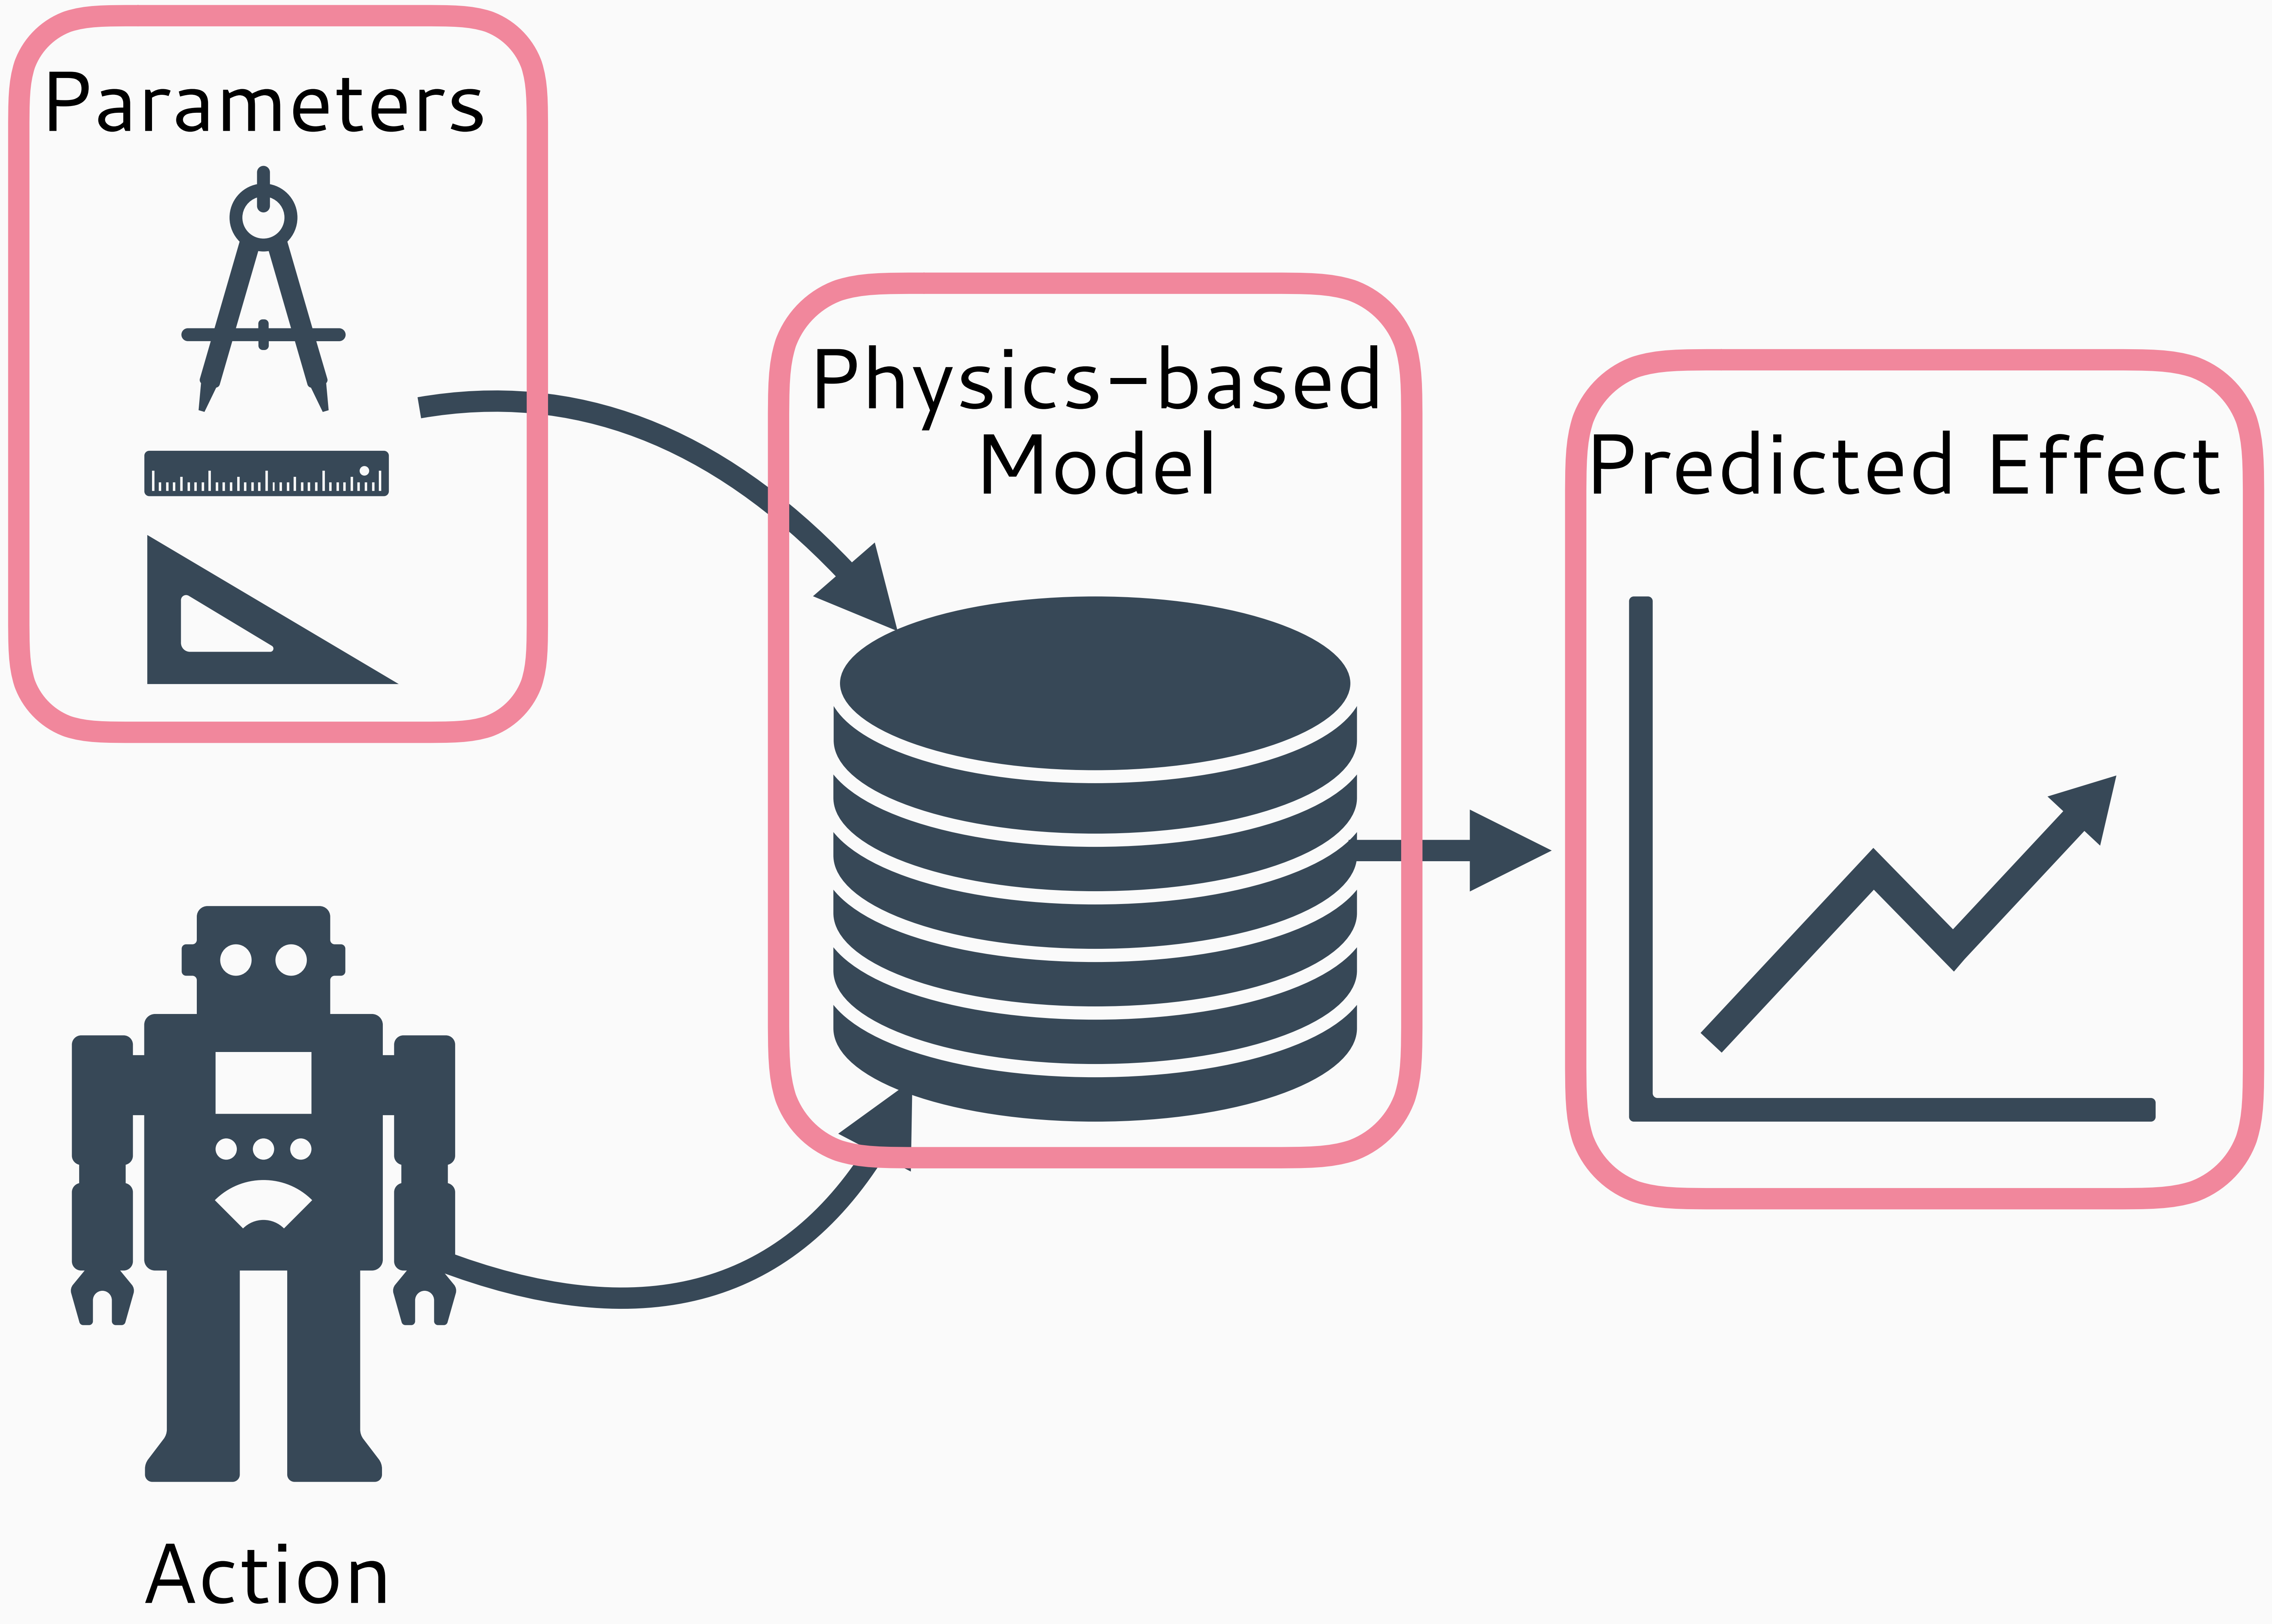

Data-Driven models

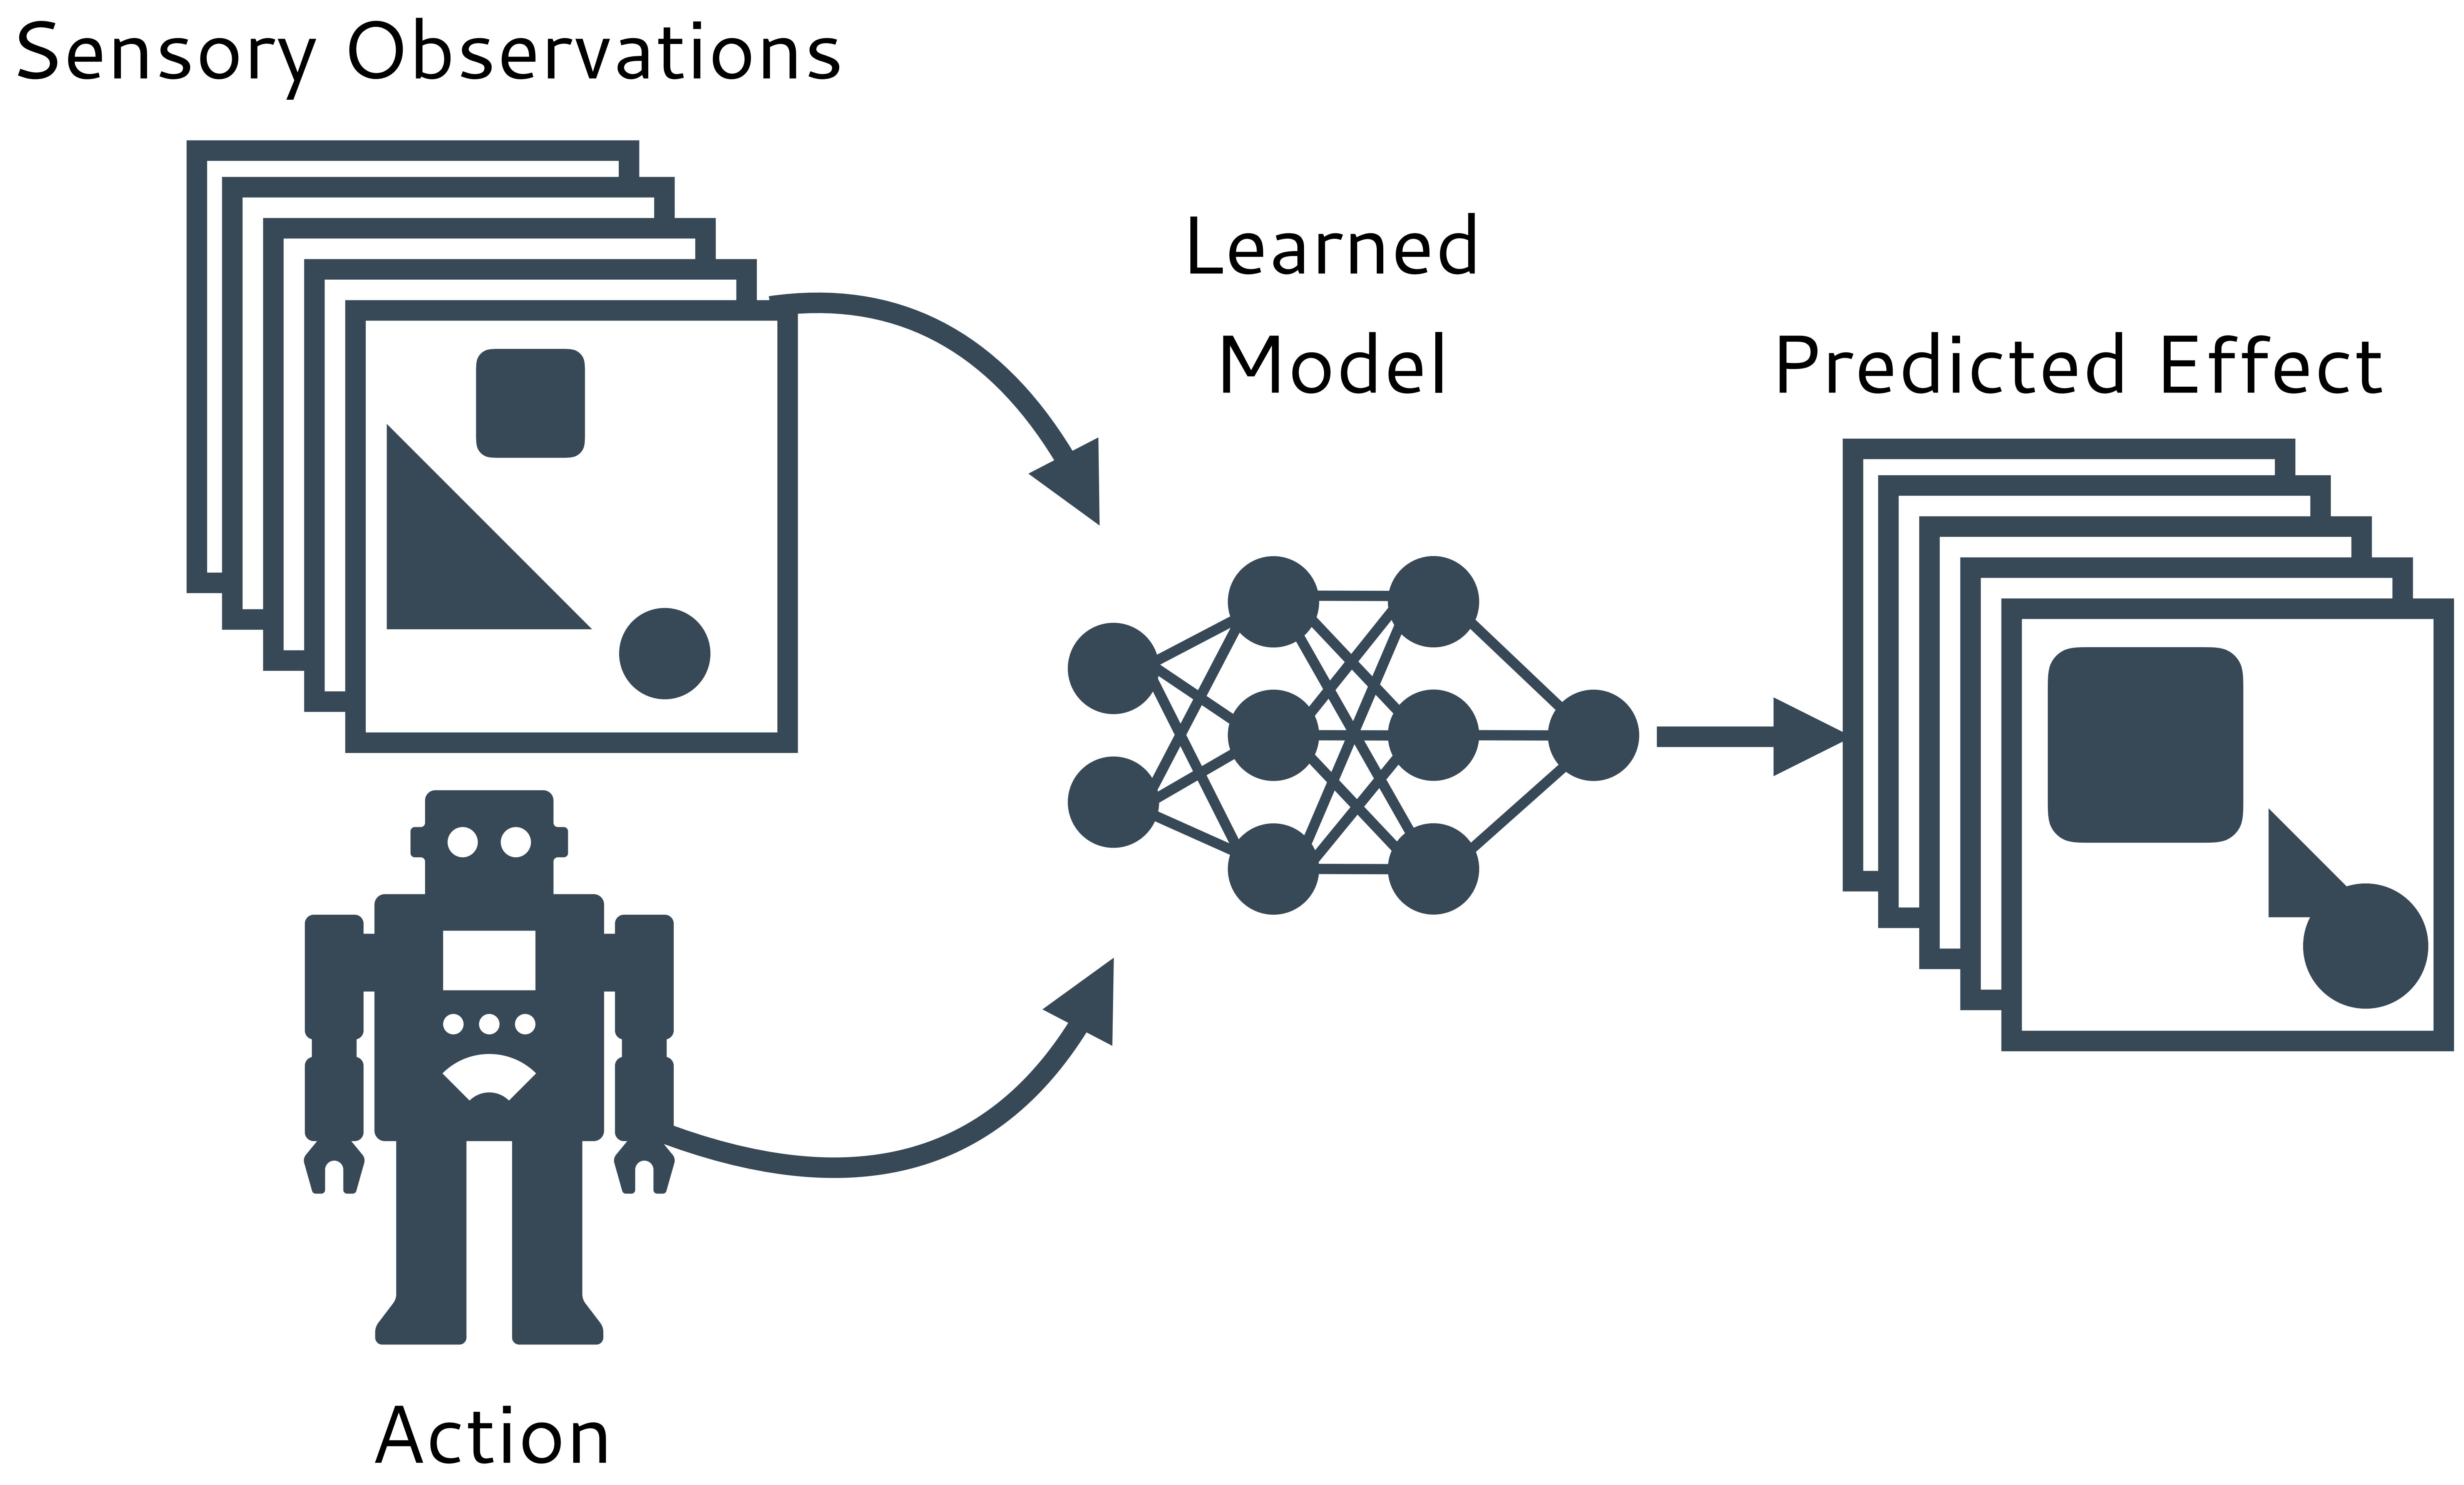

Learned Predictive Model

Unsupervised Learning for Physical Interaction through Video Prediction. Finn et al. NIPS 2016.

Our Hypothesis

Alina Kloss

Physics Models + Learning

Example:

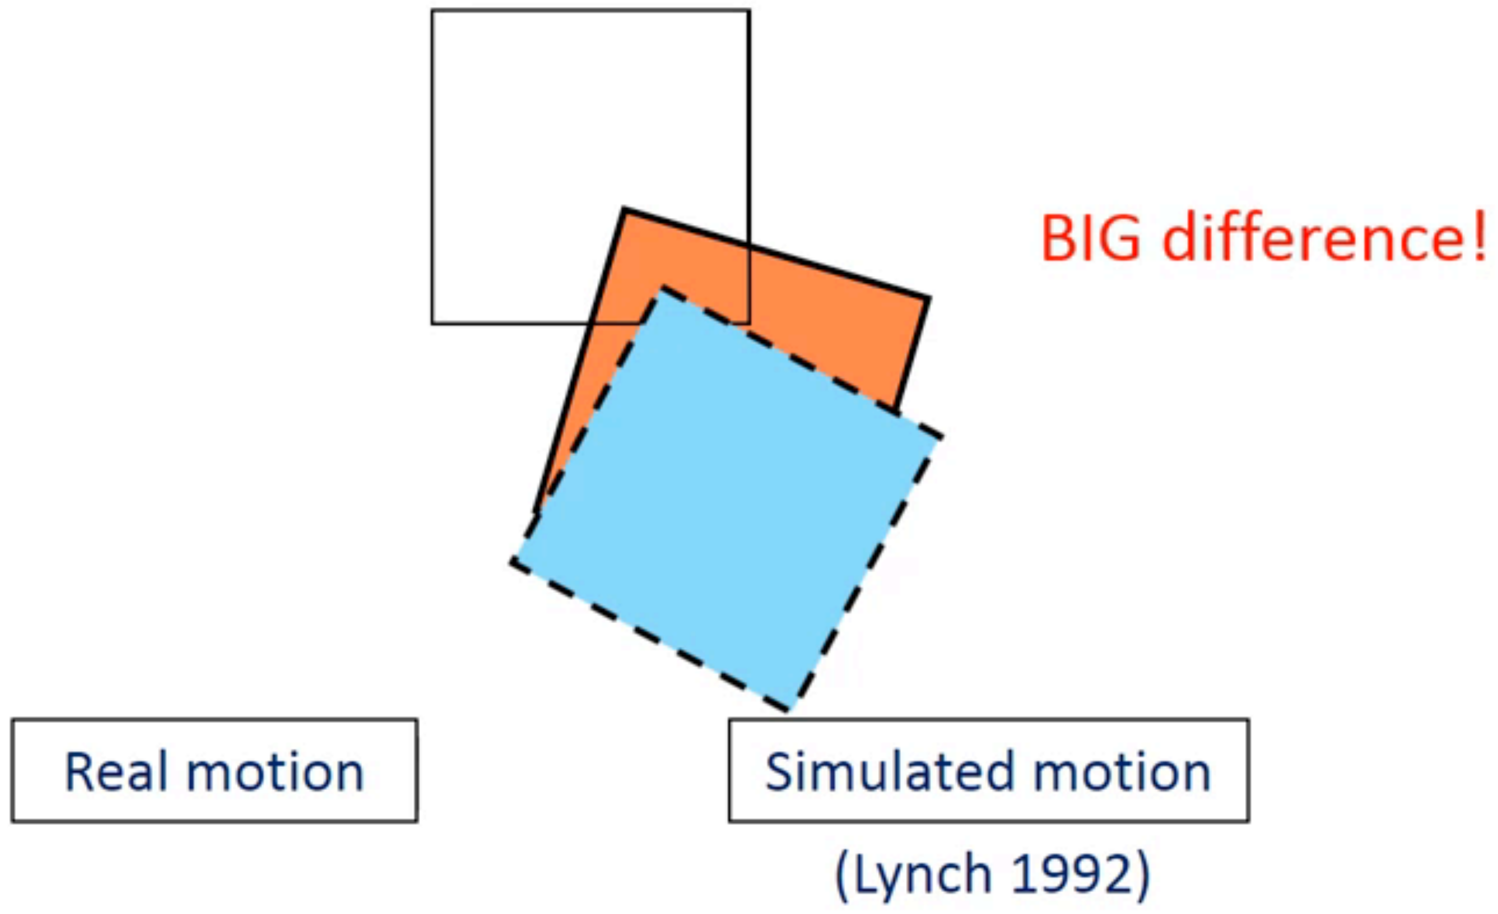

= Generalization

Alina Kloss et al, "Combining learned and analytical models for predicting action effects," 2017. Journal paper submitted to IJRR. Pre-print on arXiv.

Hybrid Model

Extrapolation

End-to-End

Loss on Effect

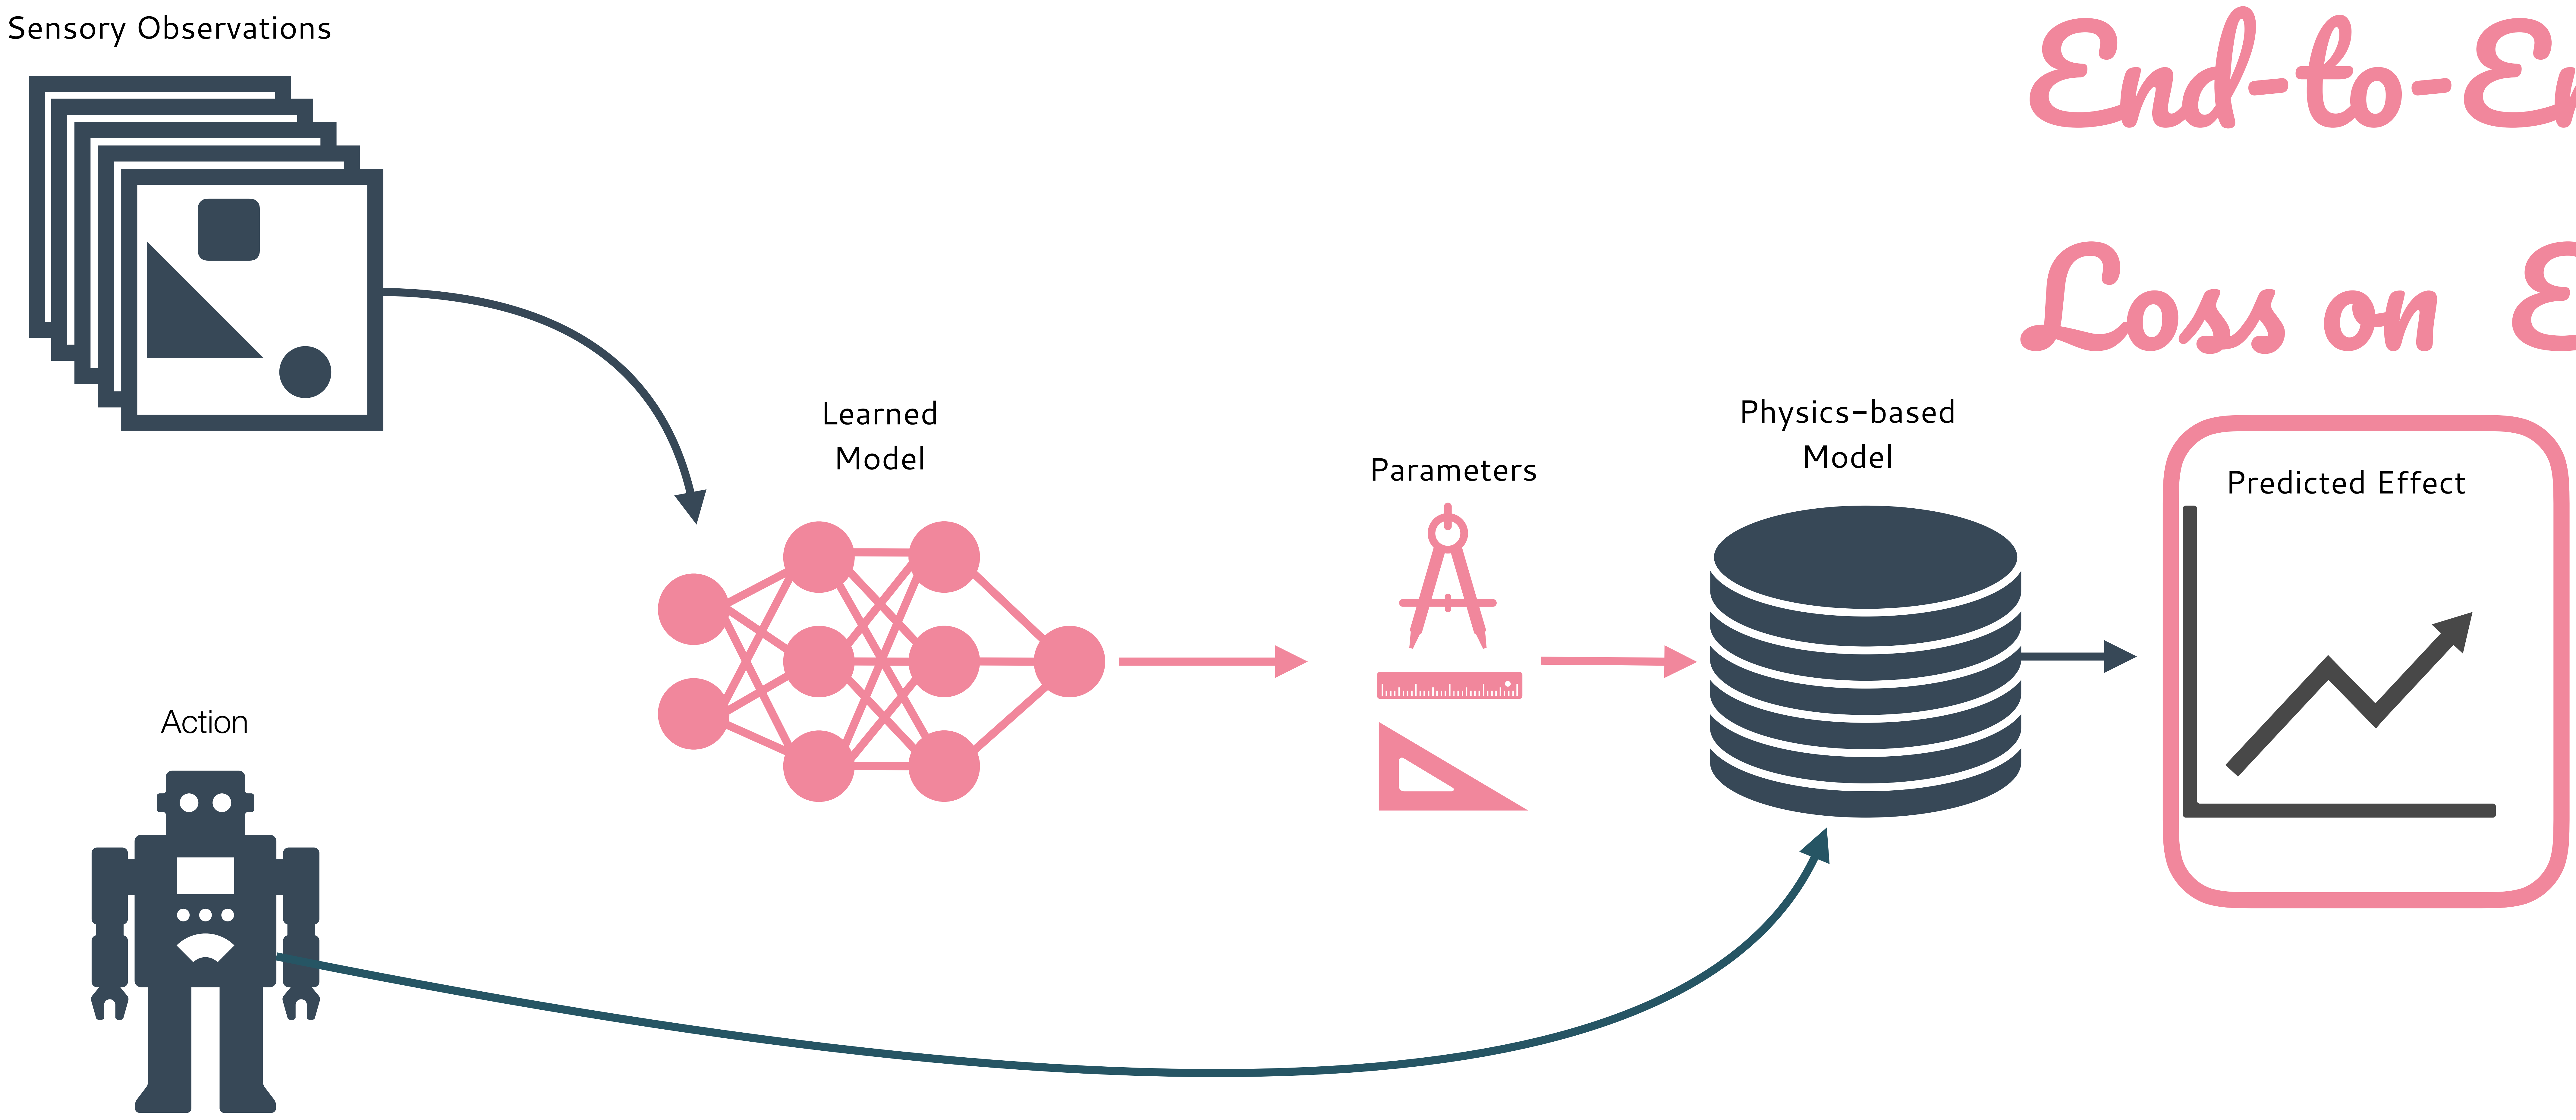

Testing Hypothesis on a Case Study

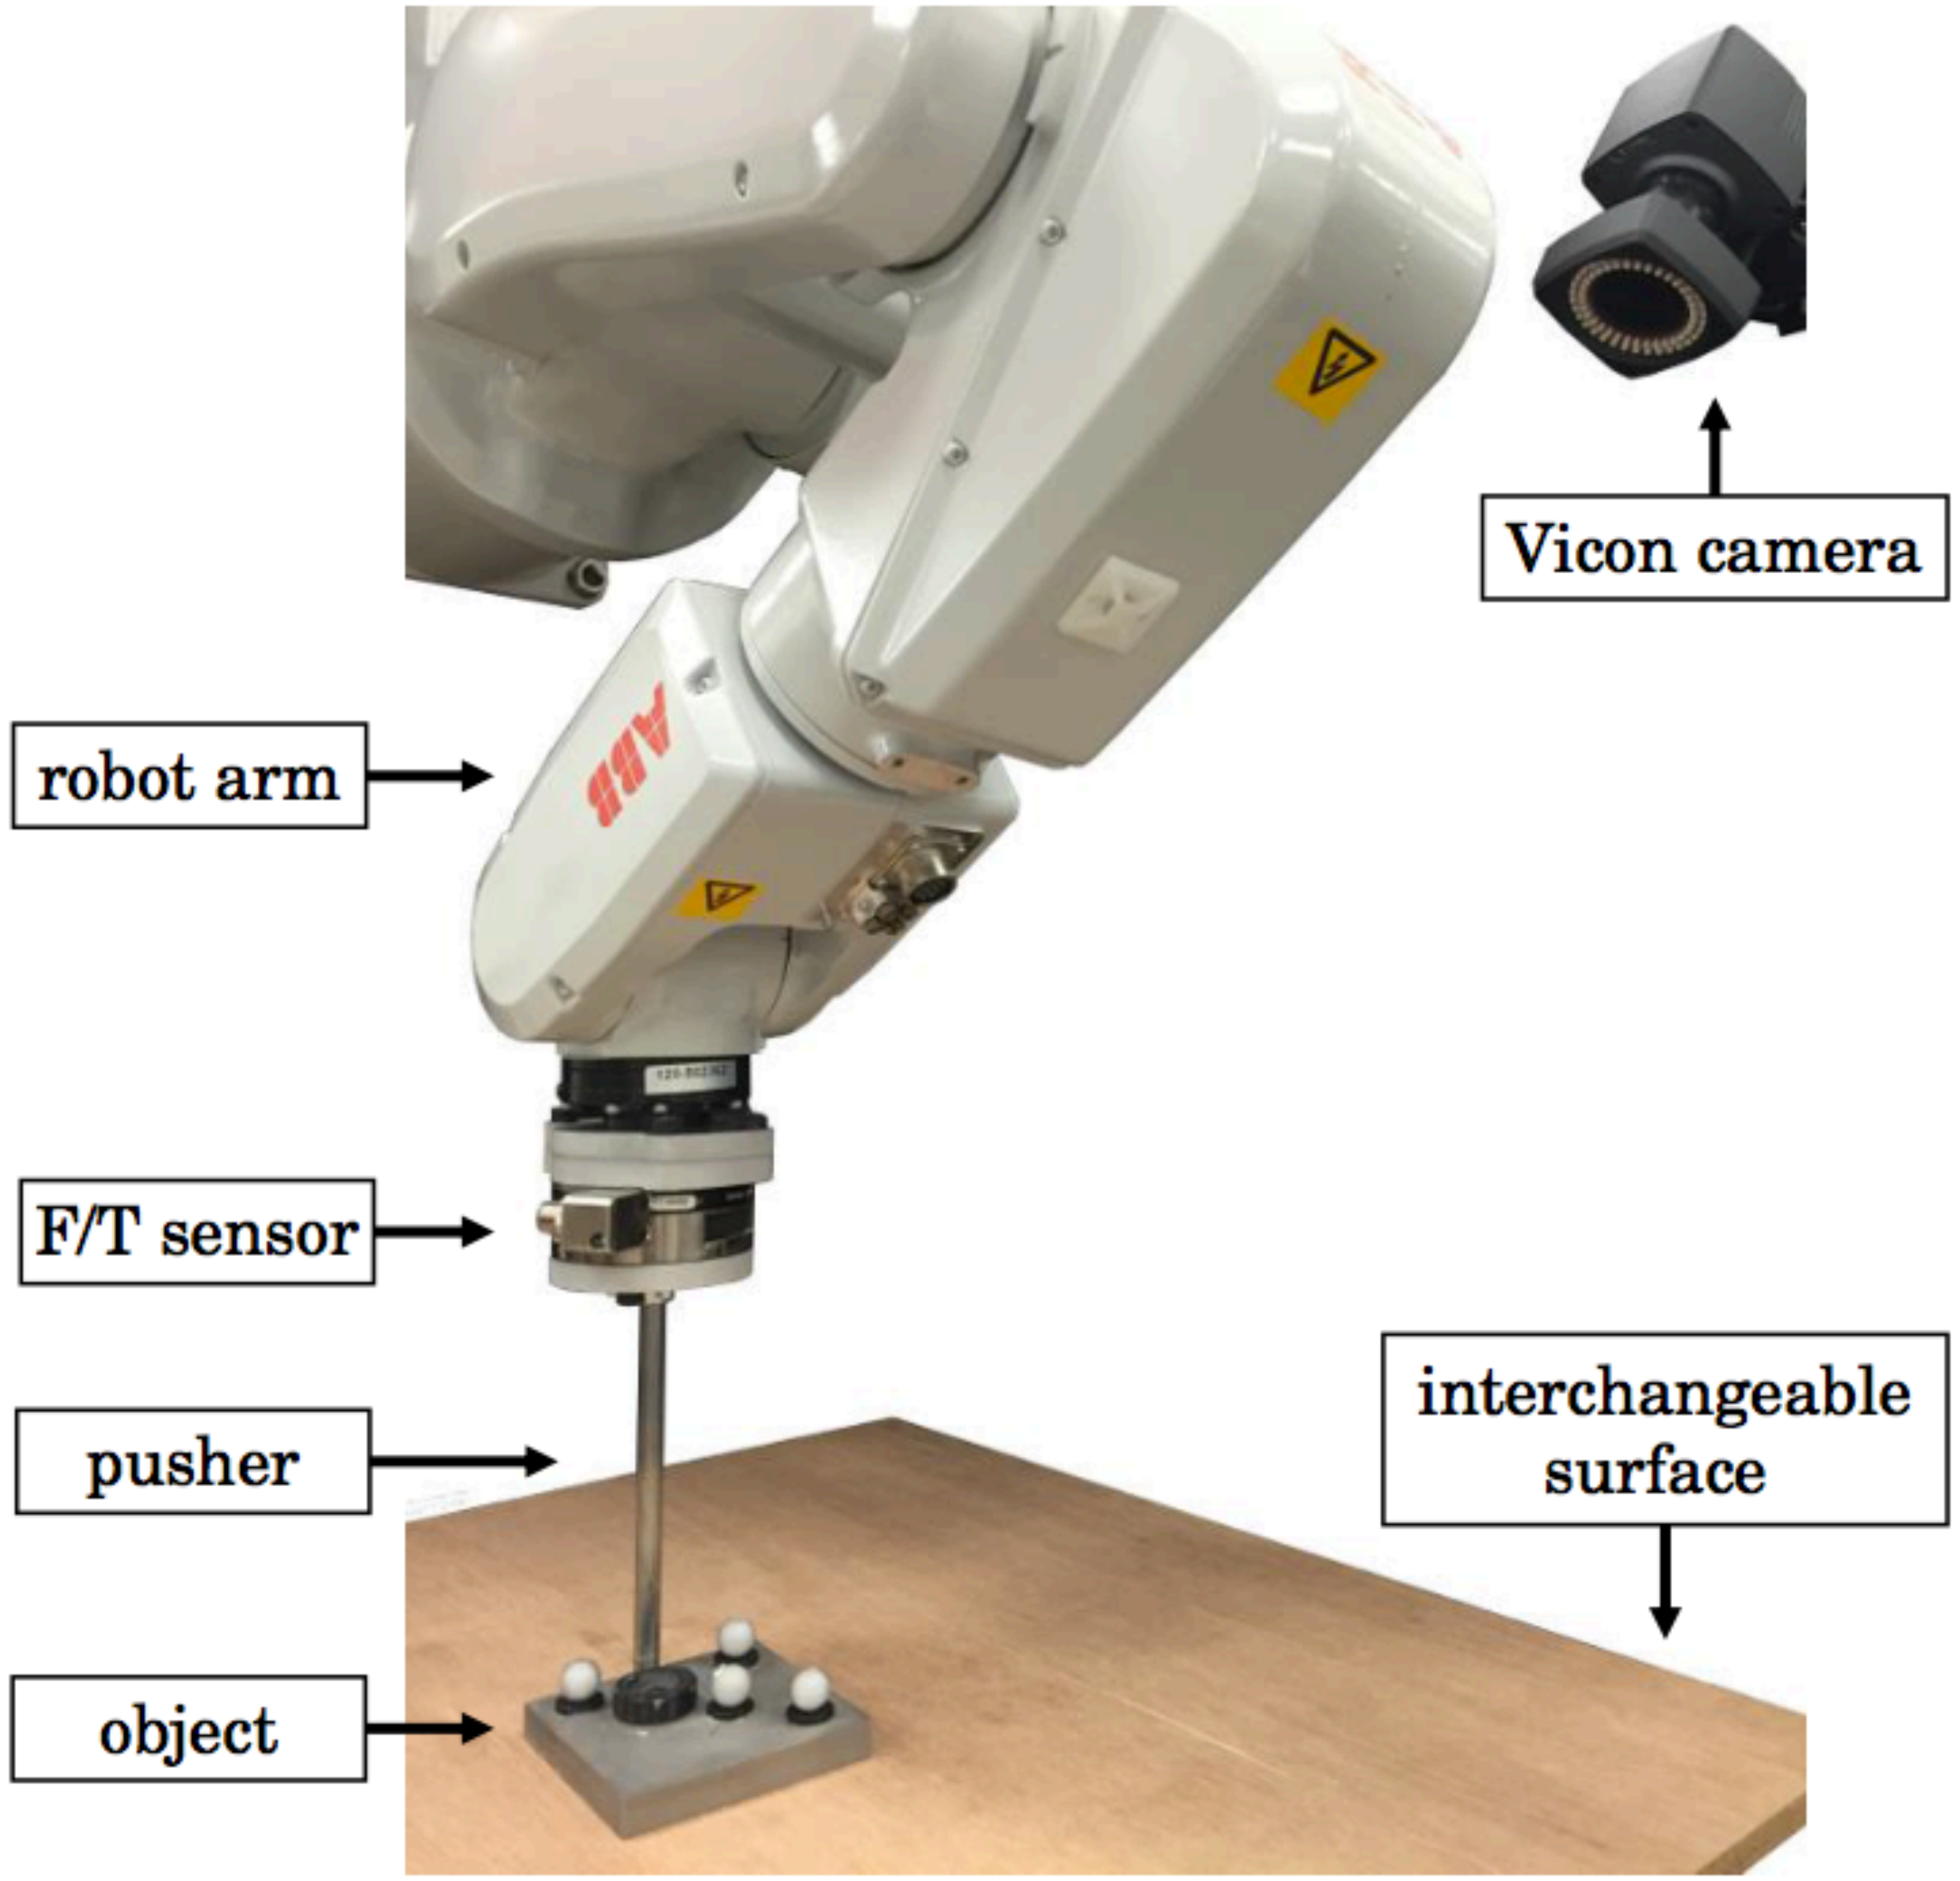

More than a Million Ways to Be Pushed. A High-Fidelity Experimental Dataset of Planar Pushing. Yu et al. IROS 2016.

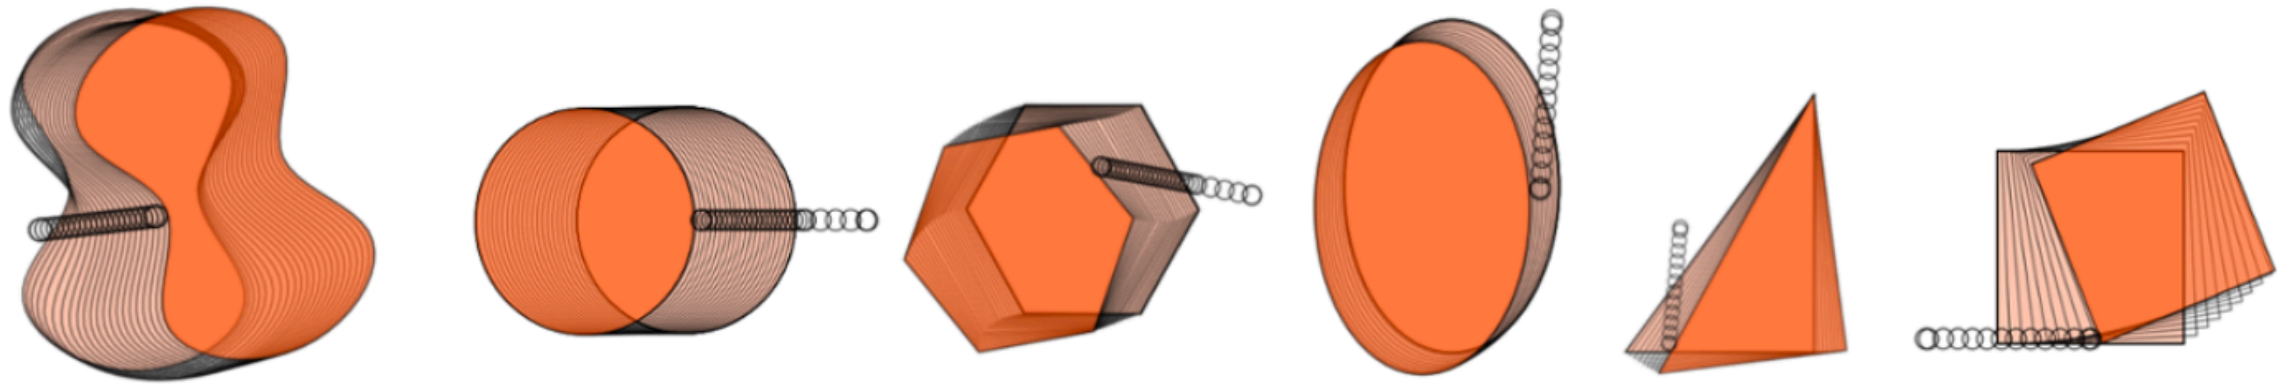

Analytic Model

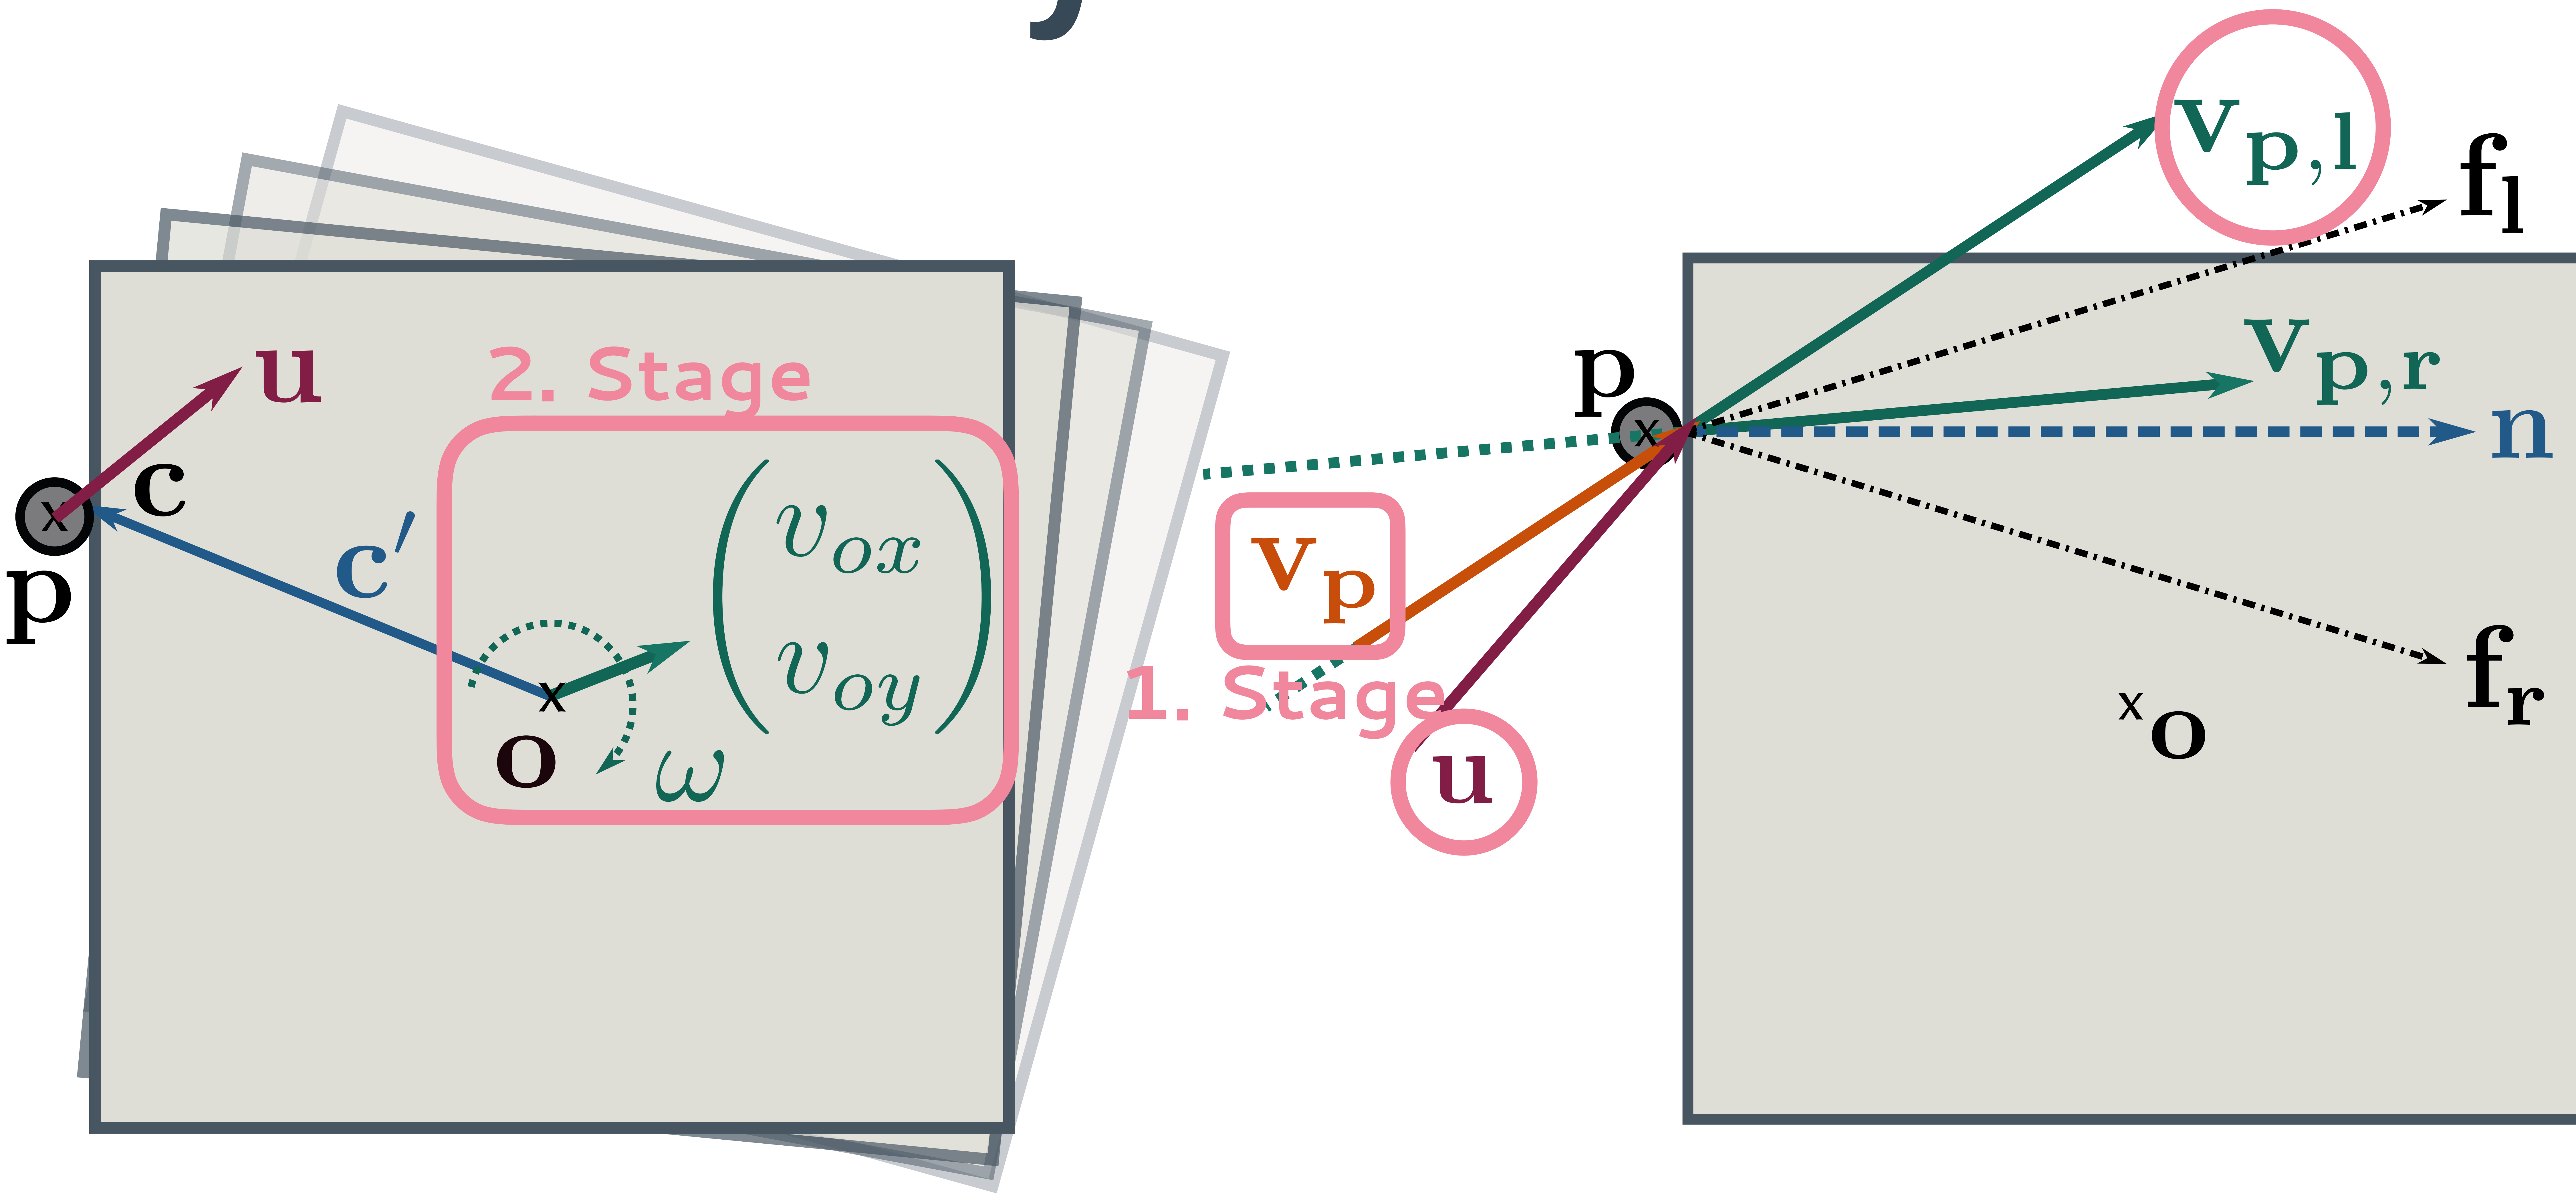

K. M. Lynch, H. Maekawa, and K. Tanie, "Manipulation and active sensing by pushing using tactile feedback," in Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems, vol. 1, Jul 1992, pp. 416–421

Compared Architectures

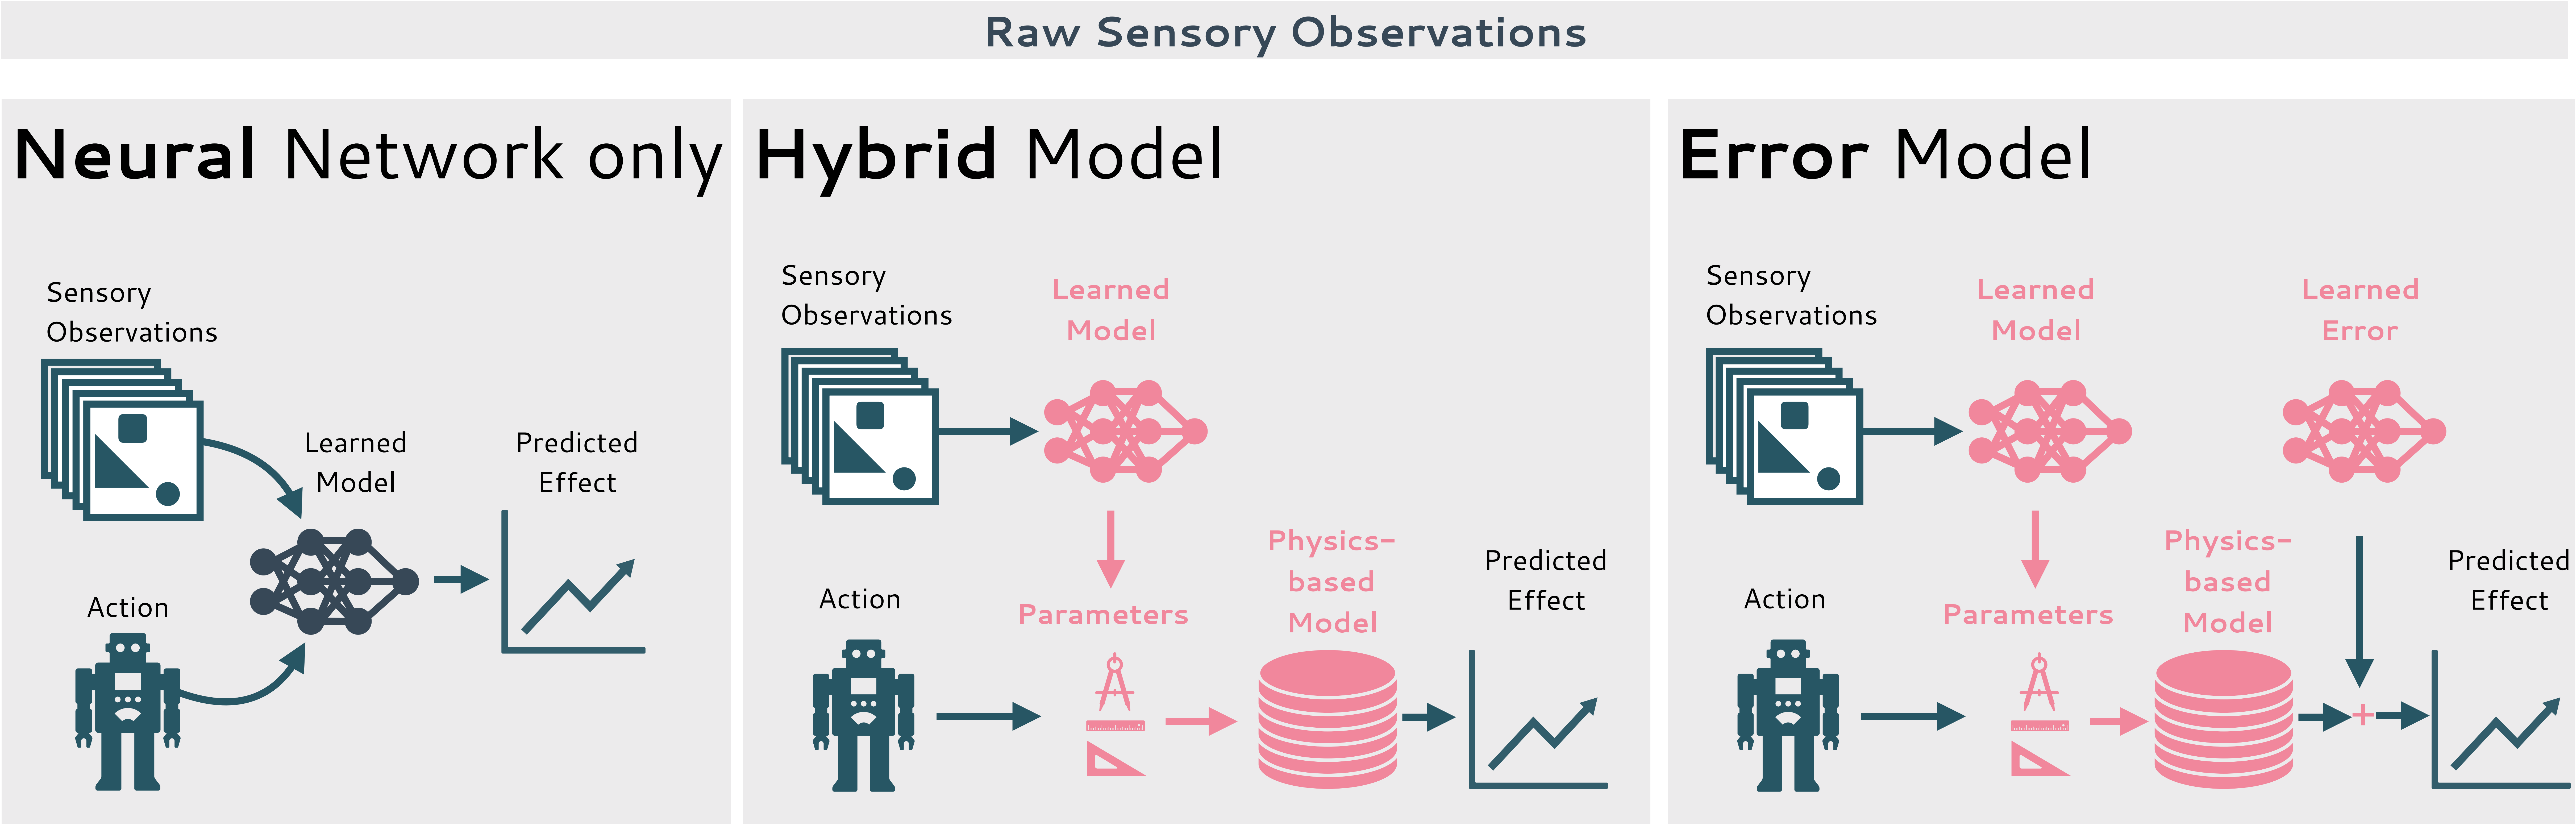

Training: End-to-End

Loss: Error between Predicted and Ground Truth Effect

Results

Improved Data Efficiency

Generalisation over actions and shapes

Alina Kloss et al, "Combining learned and analytical models for predicting action effects," 2017. Journal paper submitted to IJRR. Pre-print on arXiv.


Don’t throw away structure

Learn to extract given state representation from raw data

A Concrete Suggestion

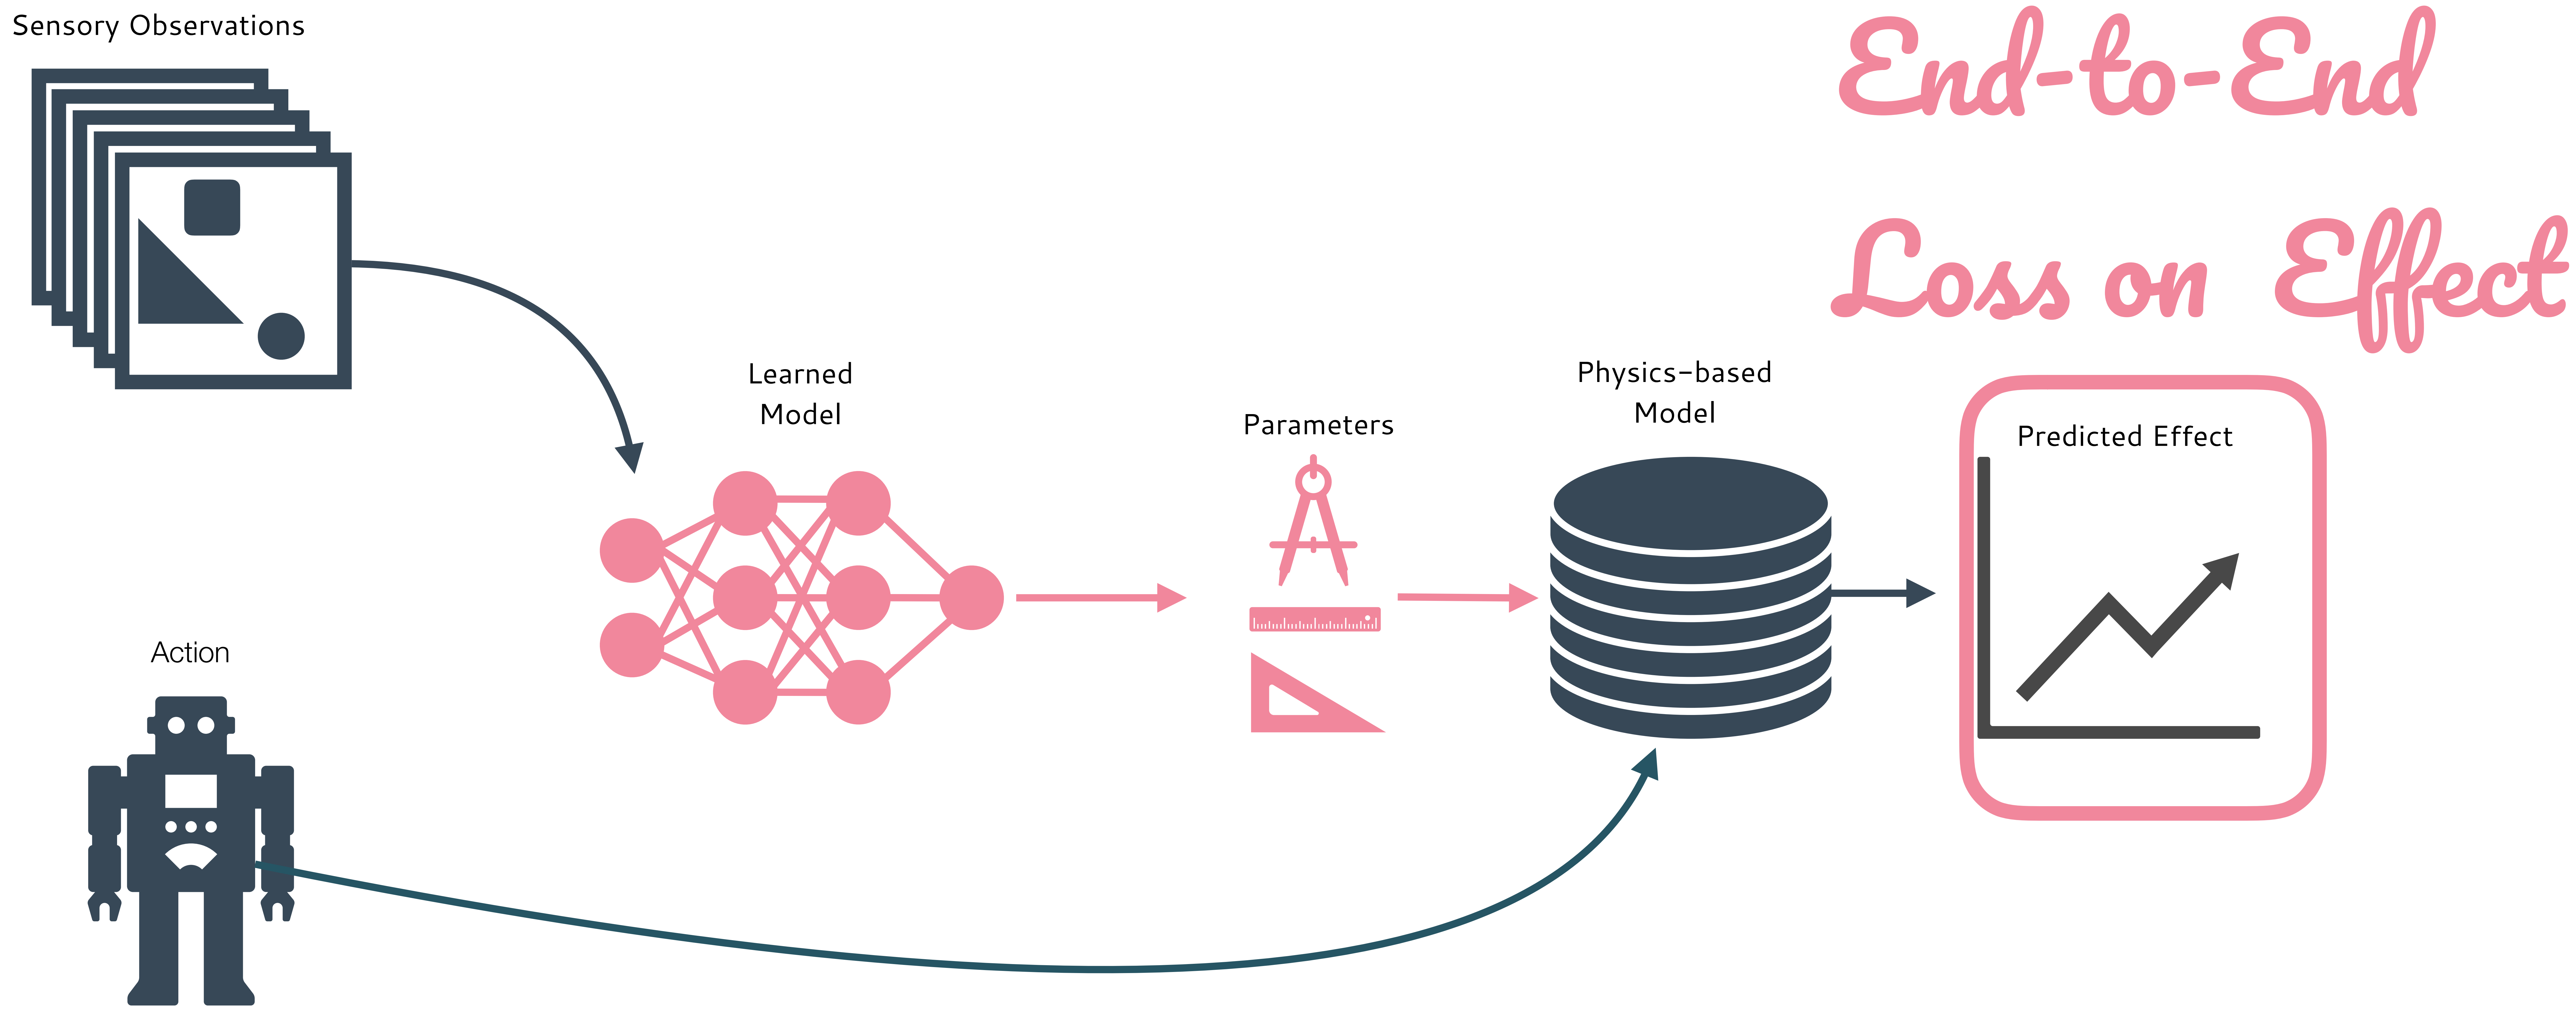

Interpretability

Example:

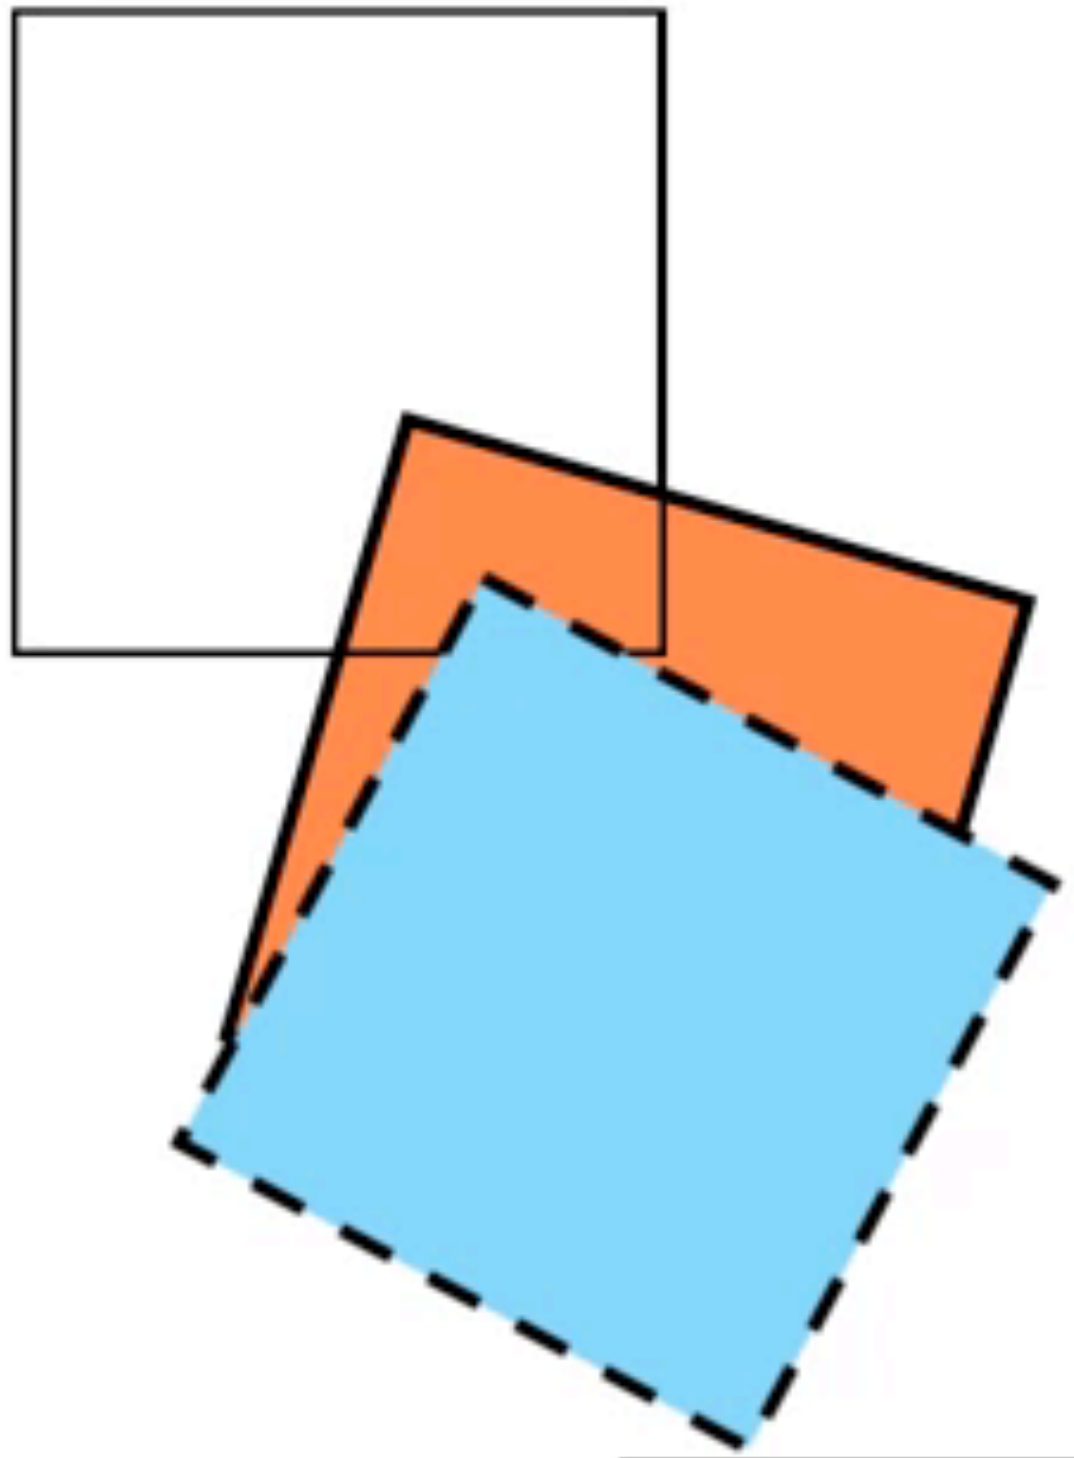

BIG difference!

Real motion

Simulated motion

(Lynch 1992)

More than a Million Ways to Be Pushed. A High-Fidelity Experimental Dataset of Planar Pushing. Yu et al. IROS 2016.

Compensation for Errors in Analytical Model

Wrong Friction Parameters of Analytical Model

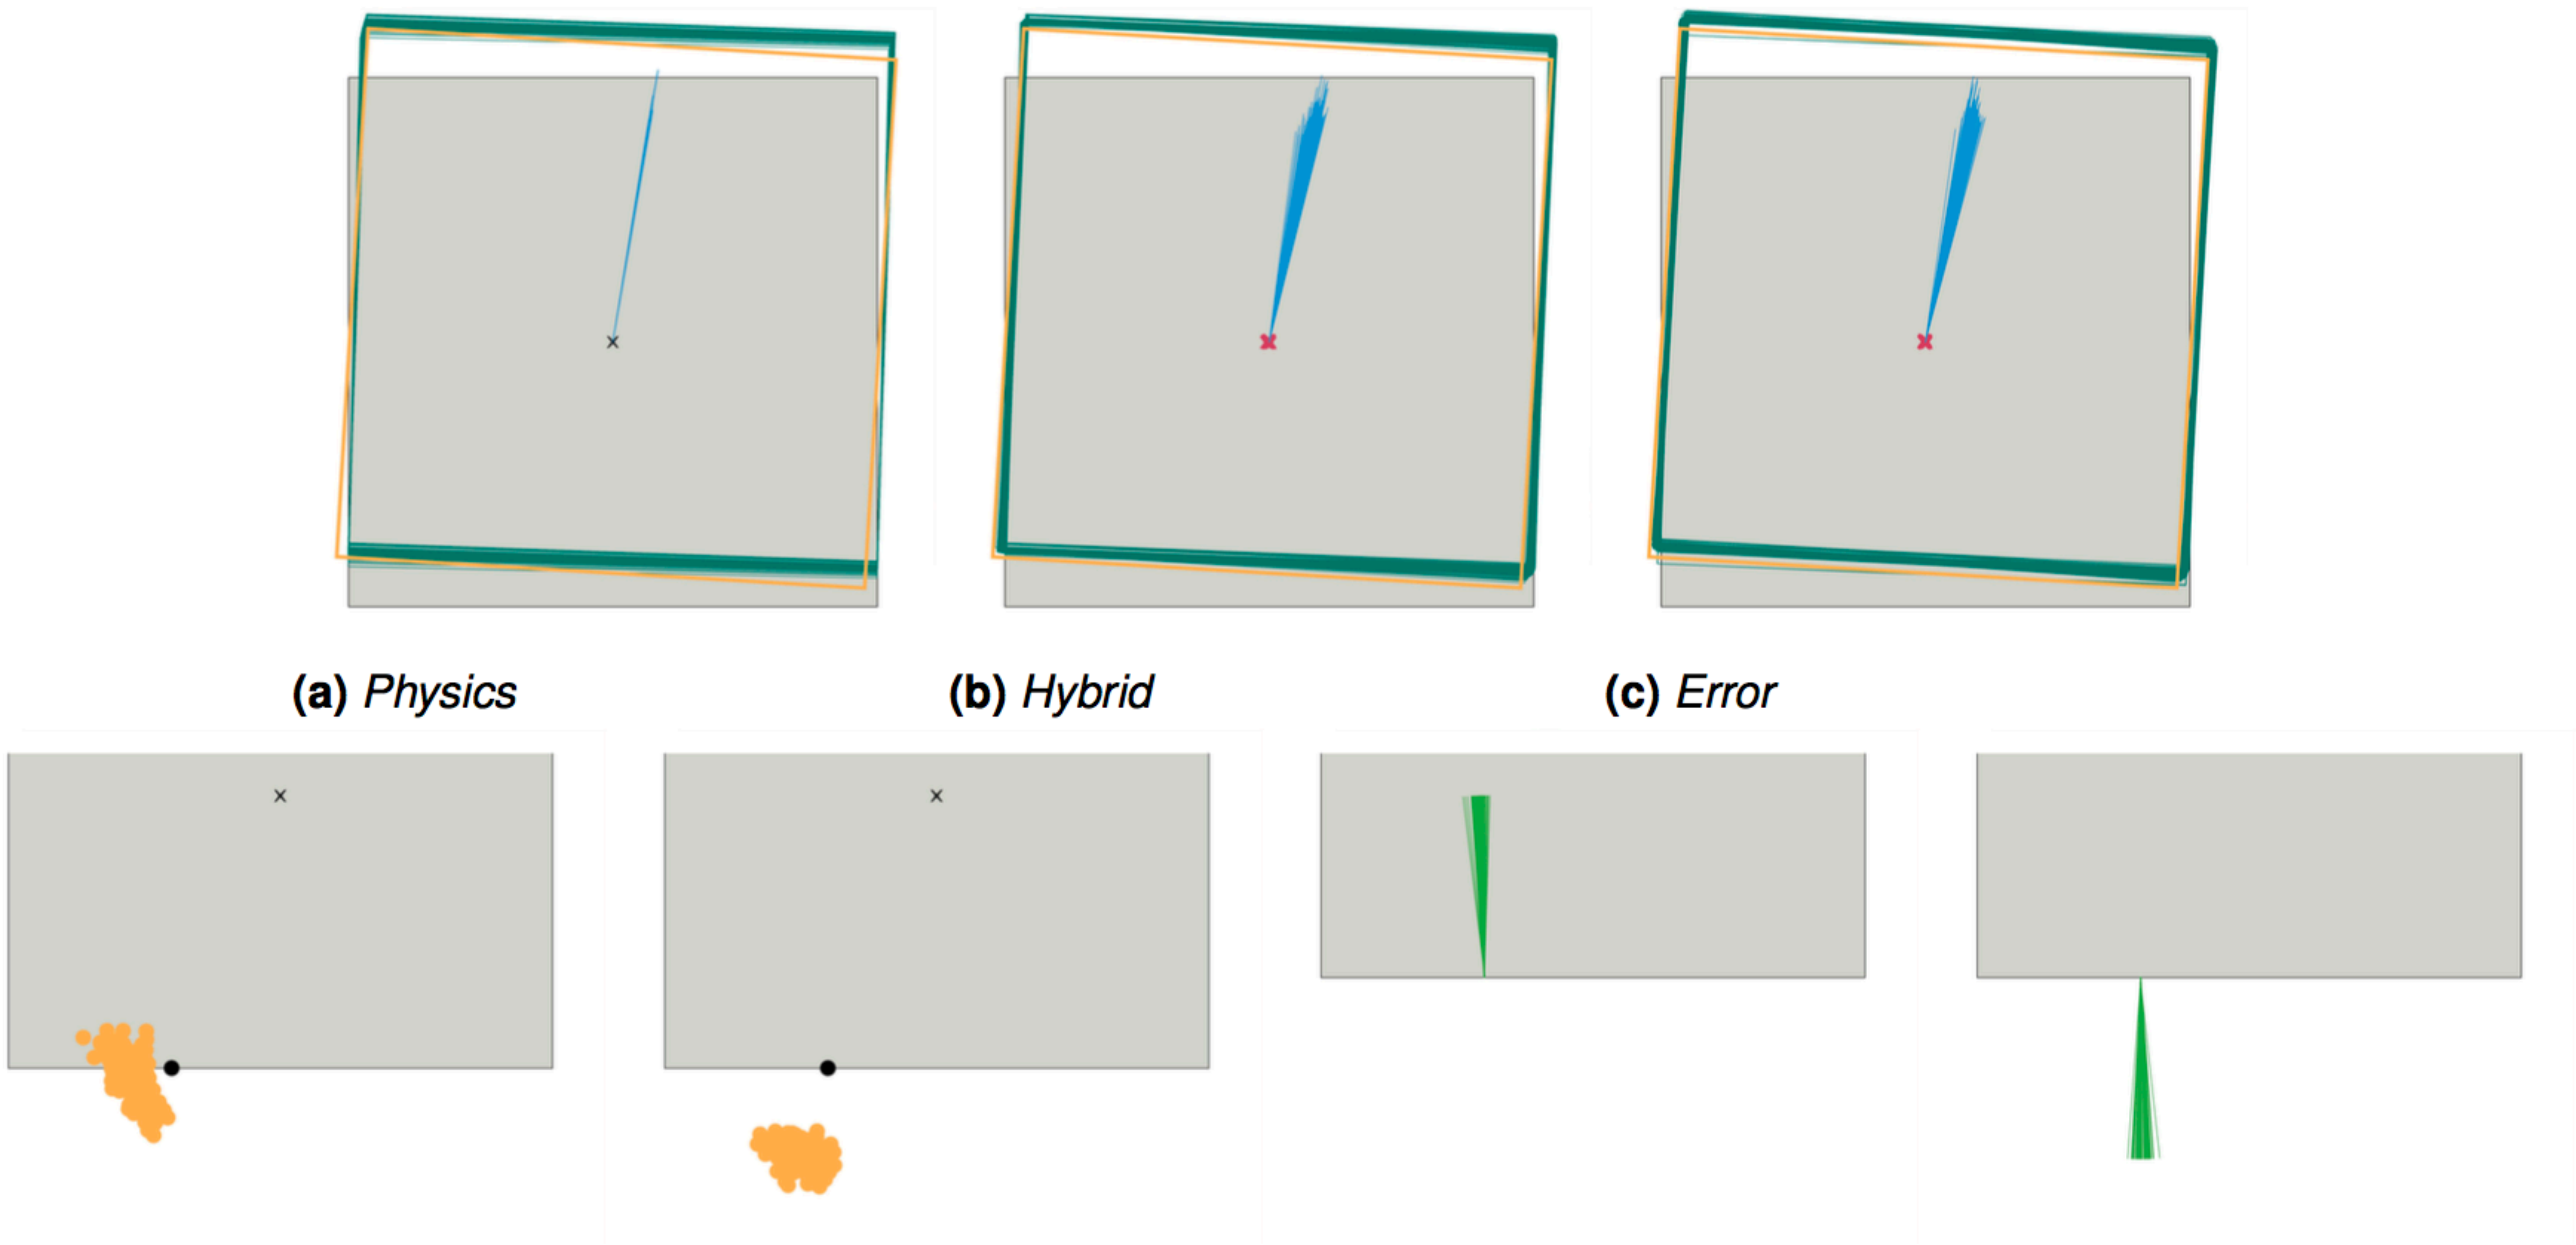

(a) *Physics* (b) *Hybrid* (c) *Error*

(d) Contact points predicted by *hybrid* (e) Contact points predicted by *error* (f) Contact normal predicted by *hybrid* (g) Contact normal predicted by *error*

Predictive Model

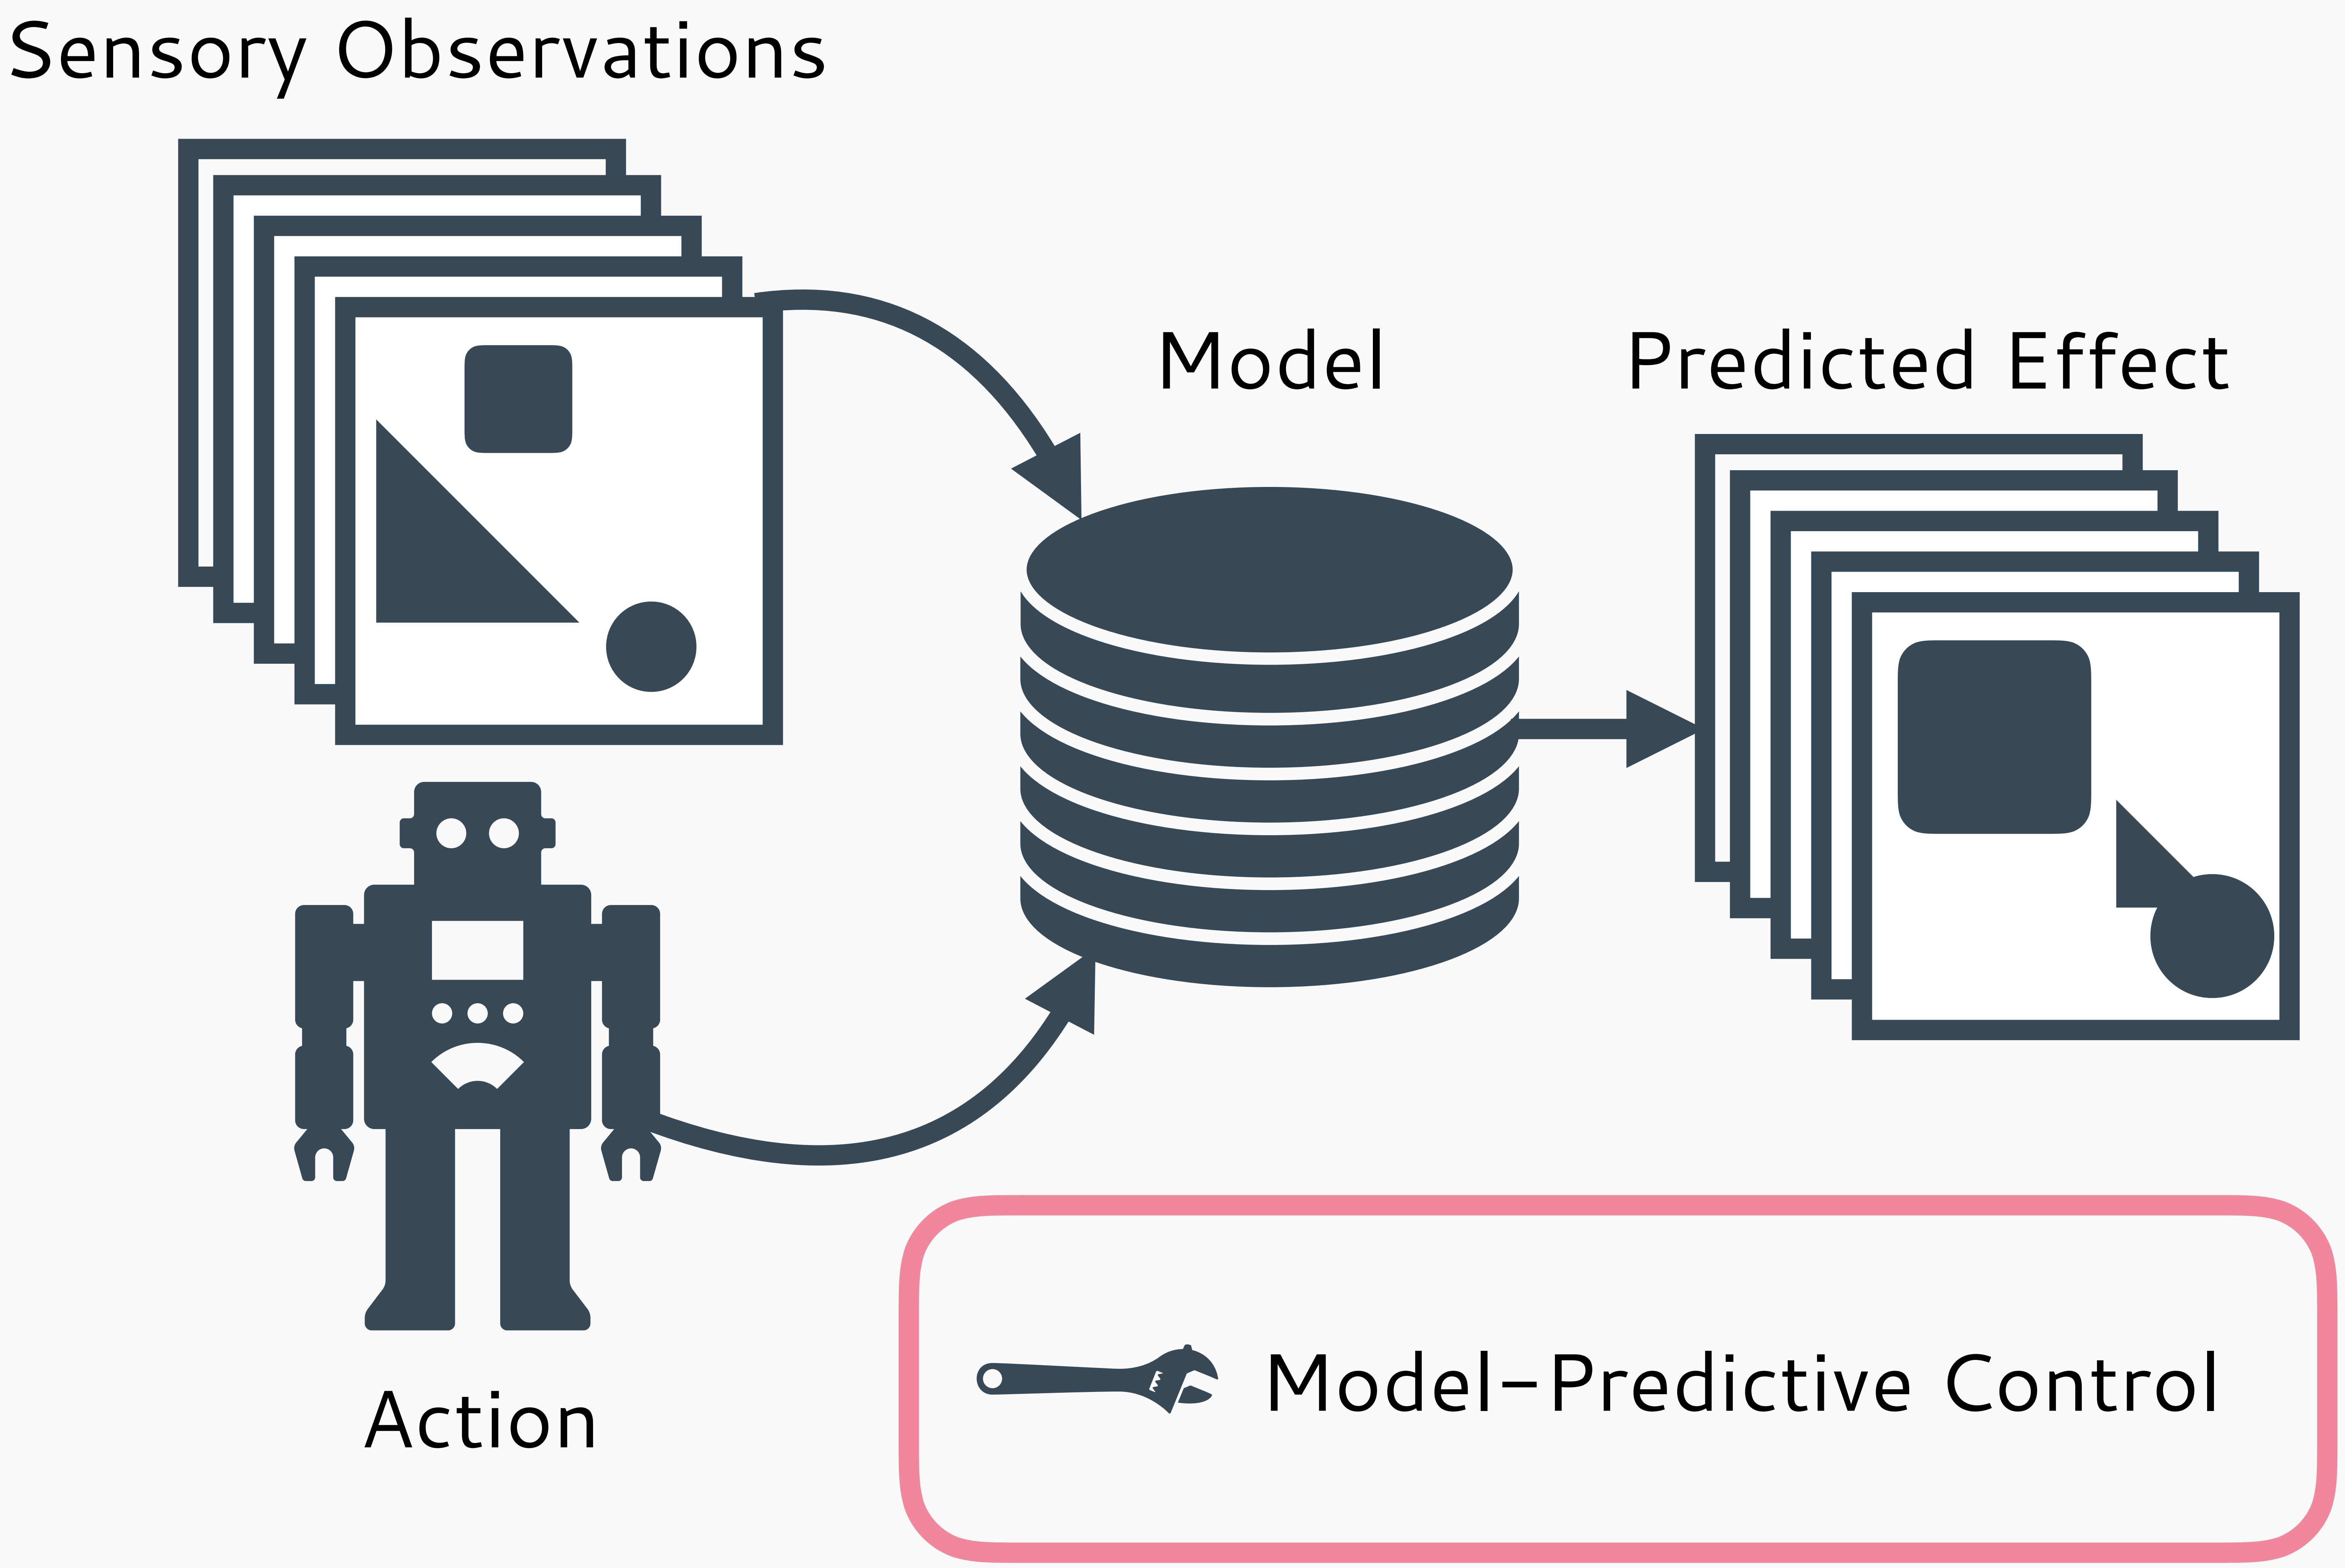

Poster

Thomas Baumeister

Predictive Model

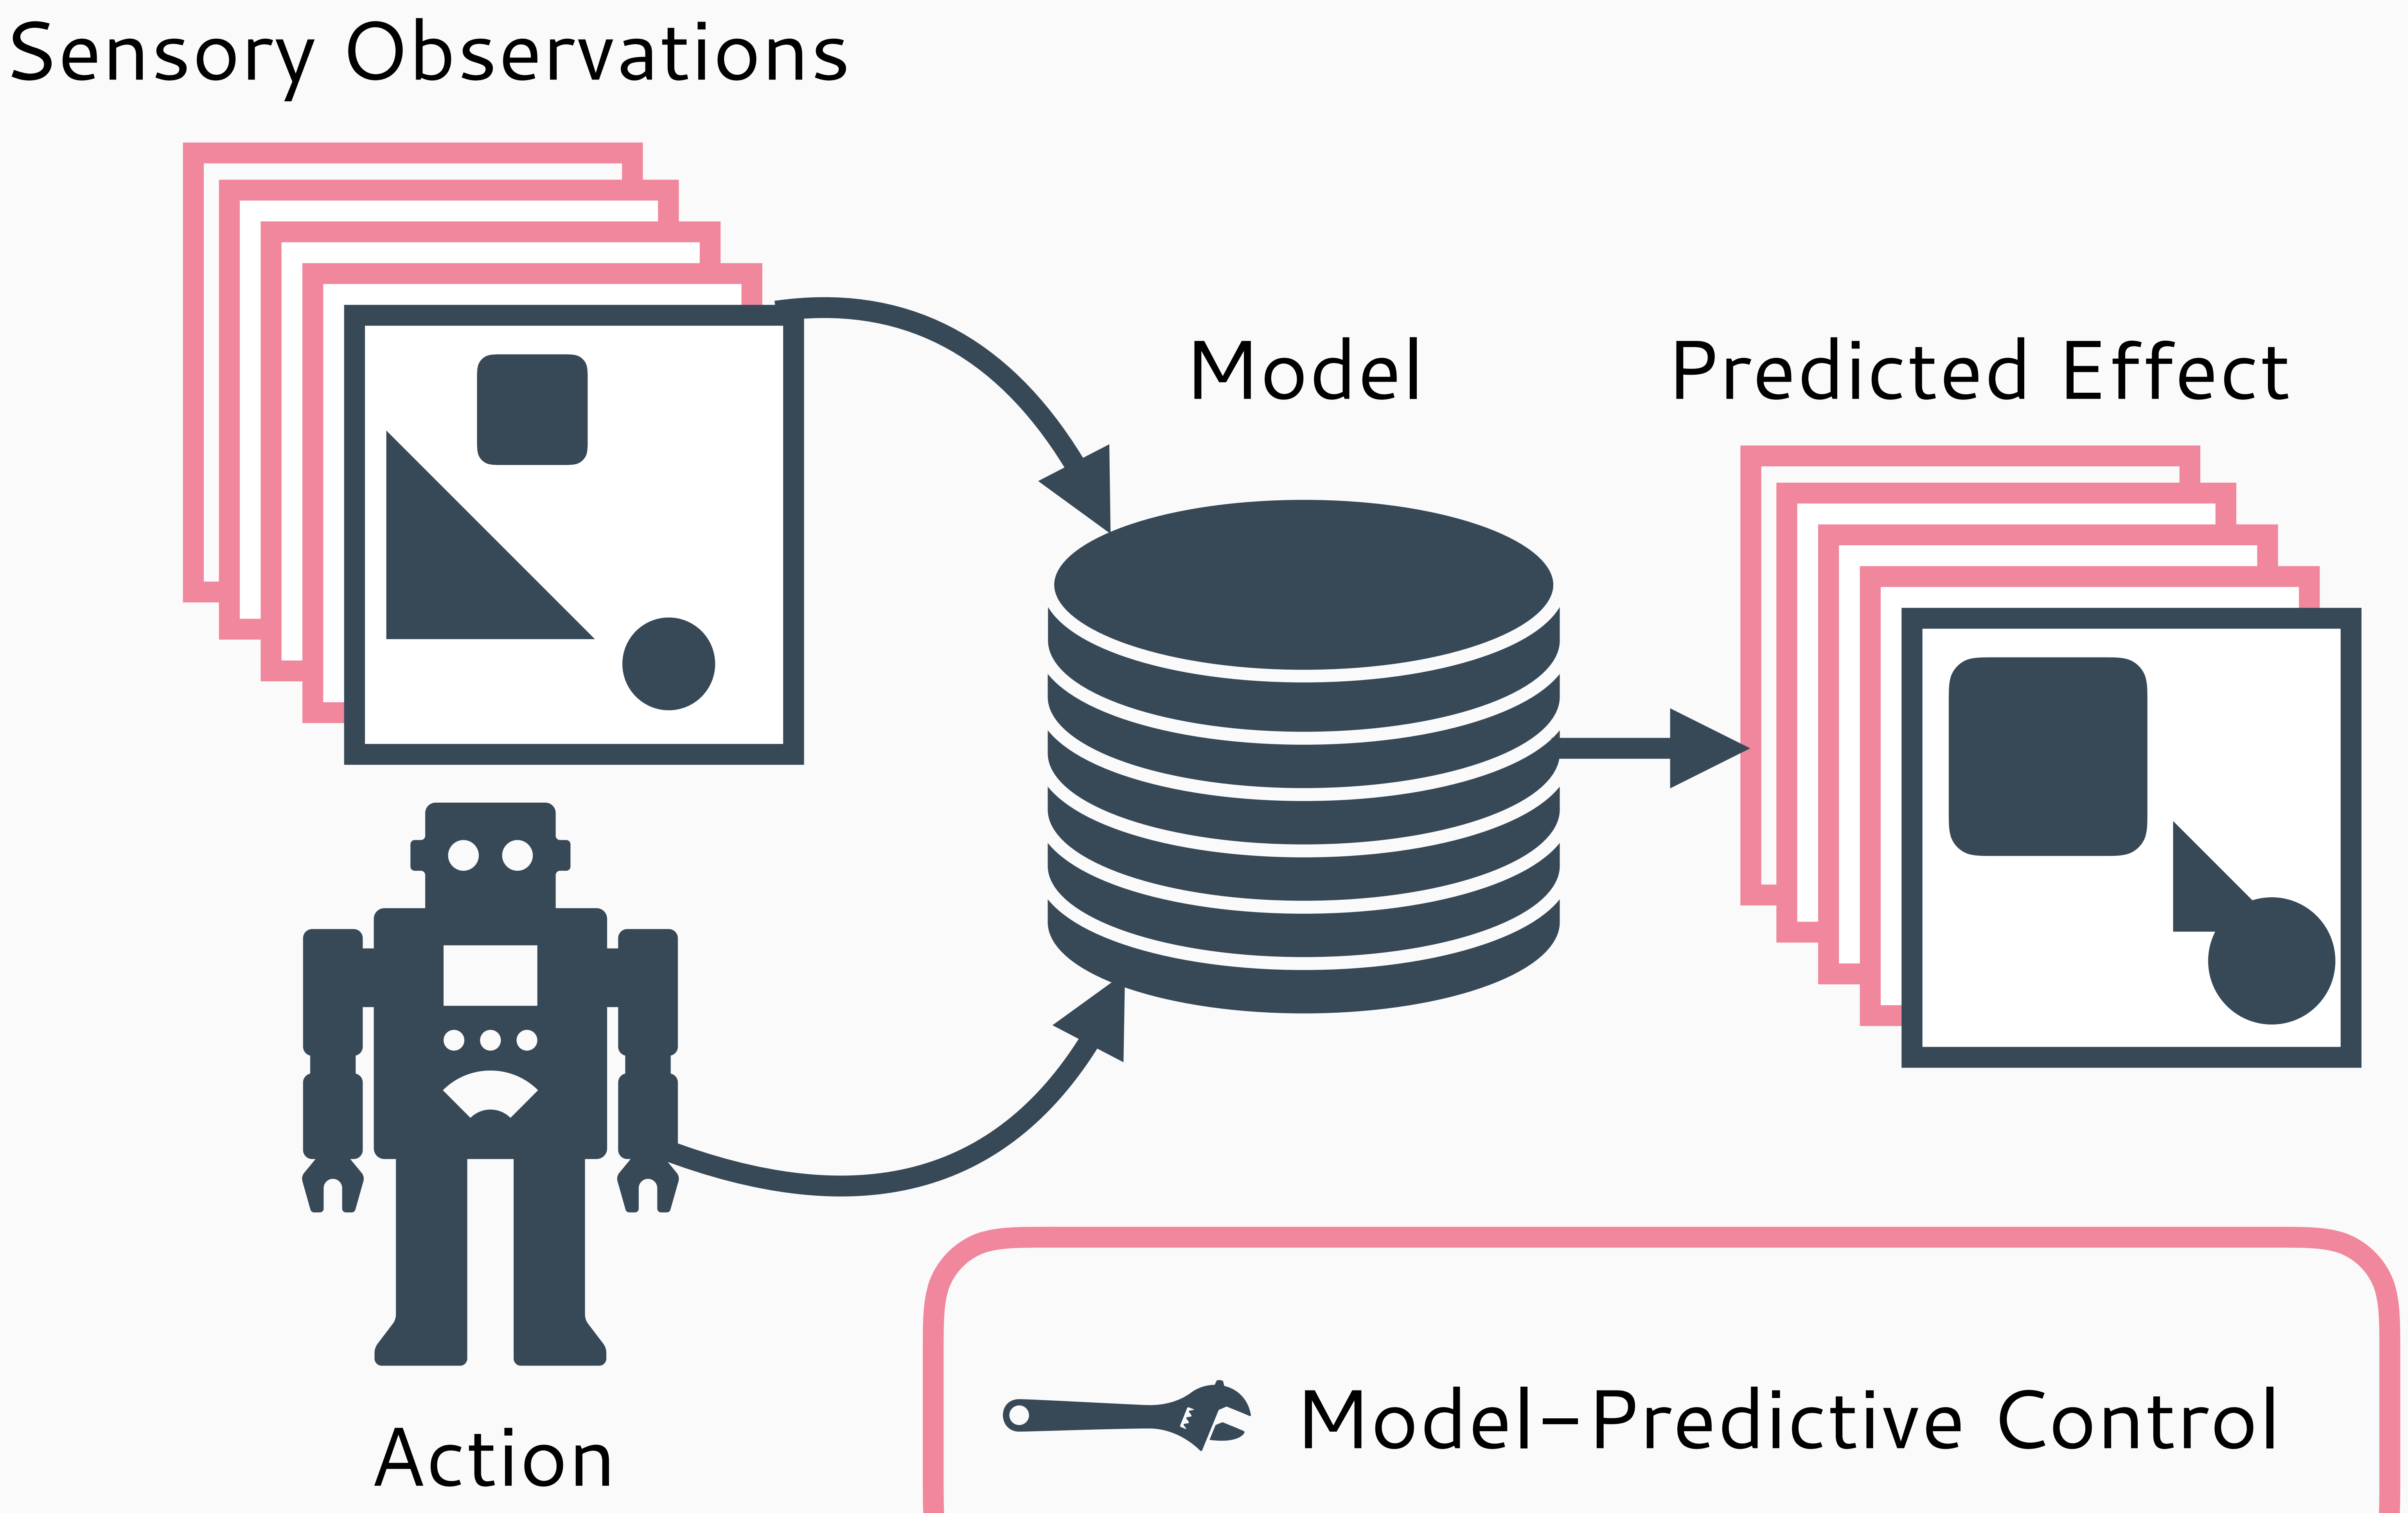

Model-Predictive Control

Poster

Thomas Baumeister

But what if?

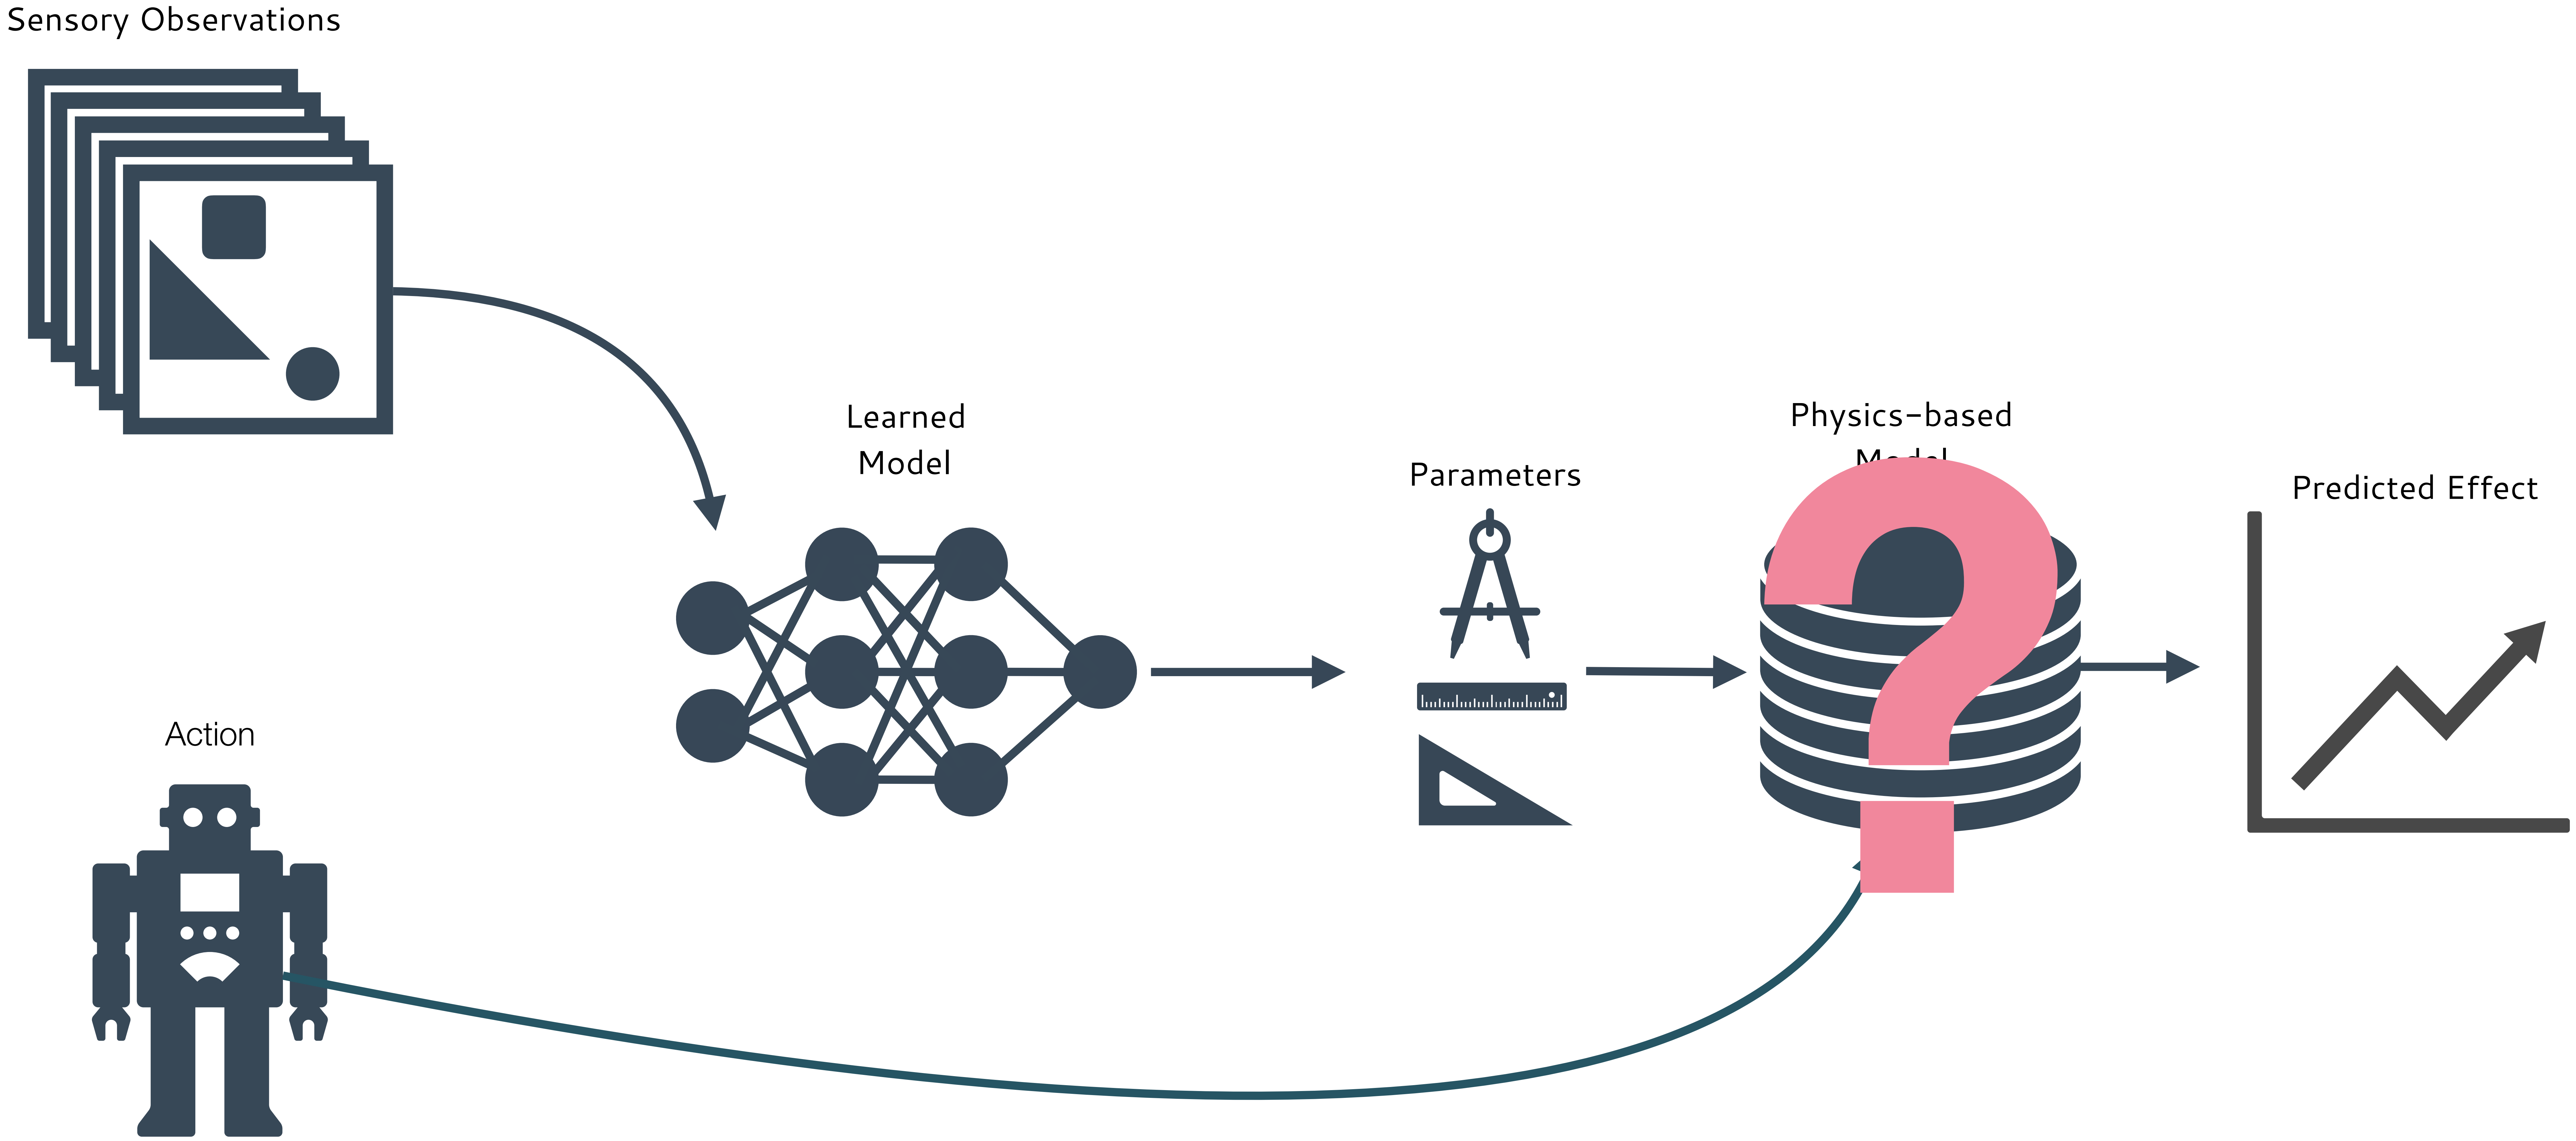
Sensory Observations
Learned
Model
Parameters
Physics-based
Predicted Effect
Action

Vision and Touch are complementary

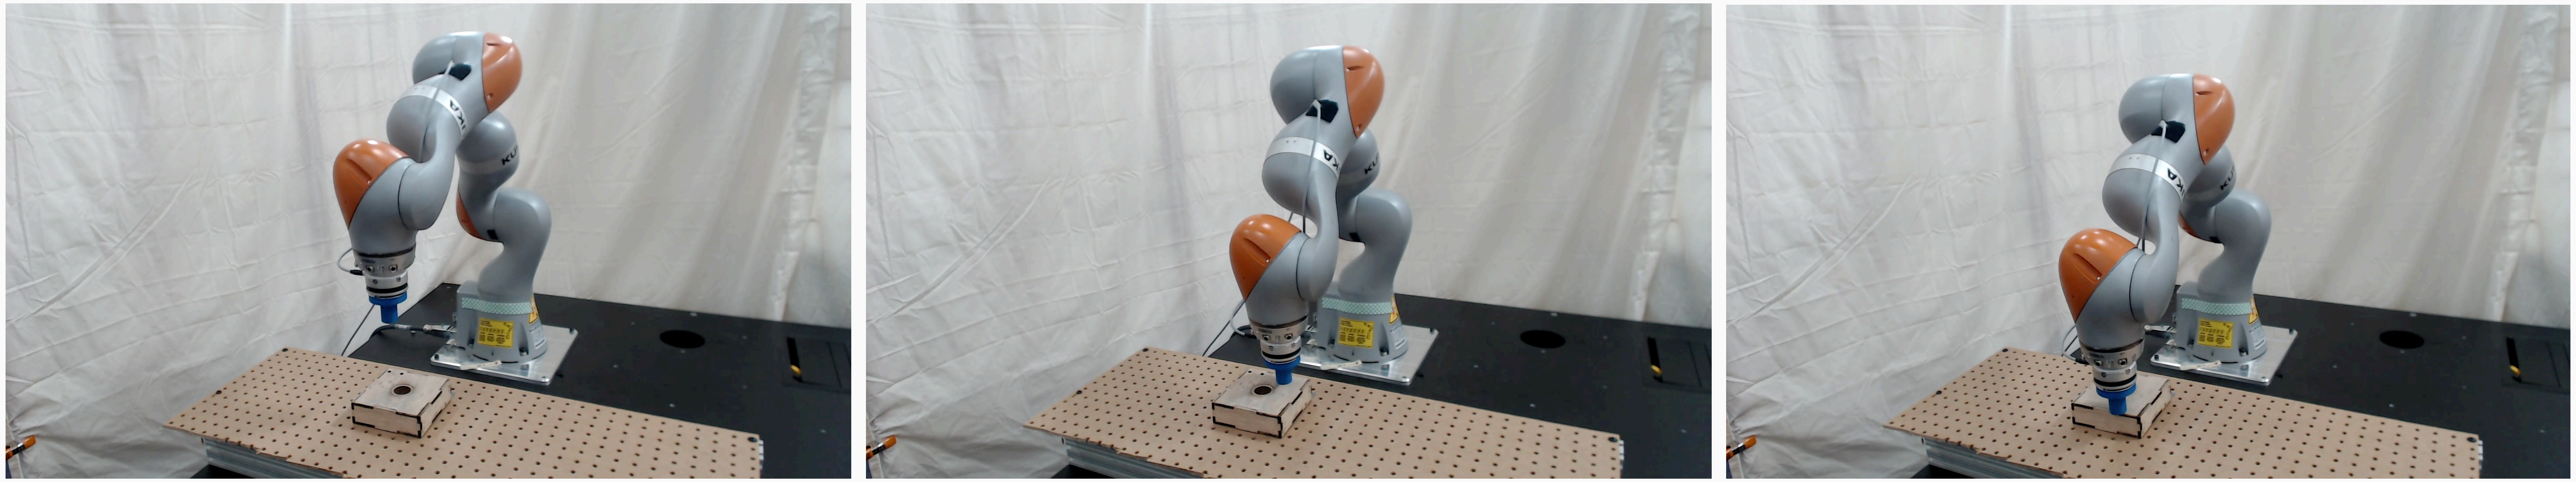

Reaching Alignment Insertion

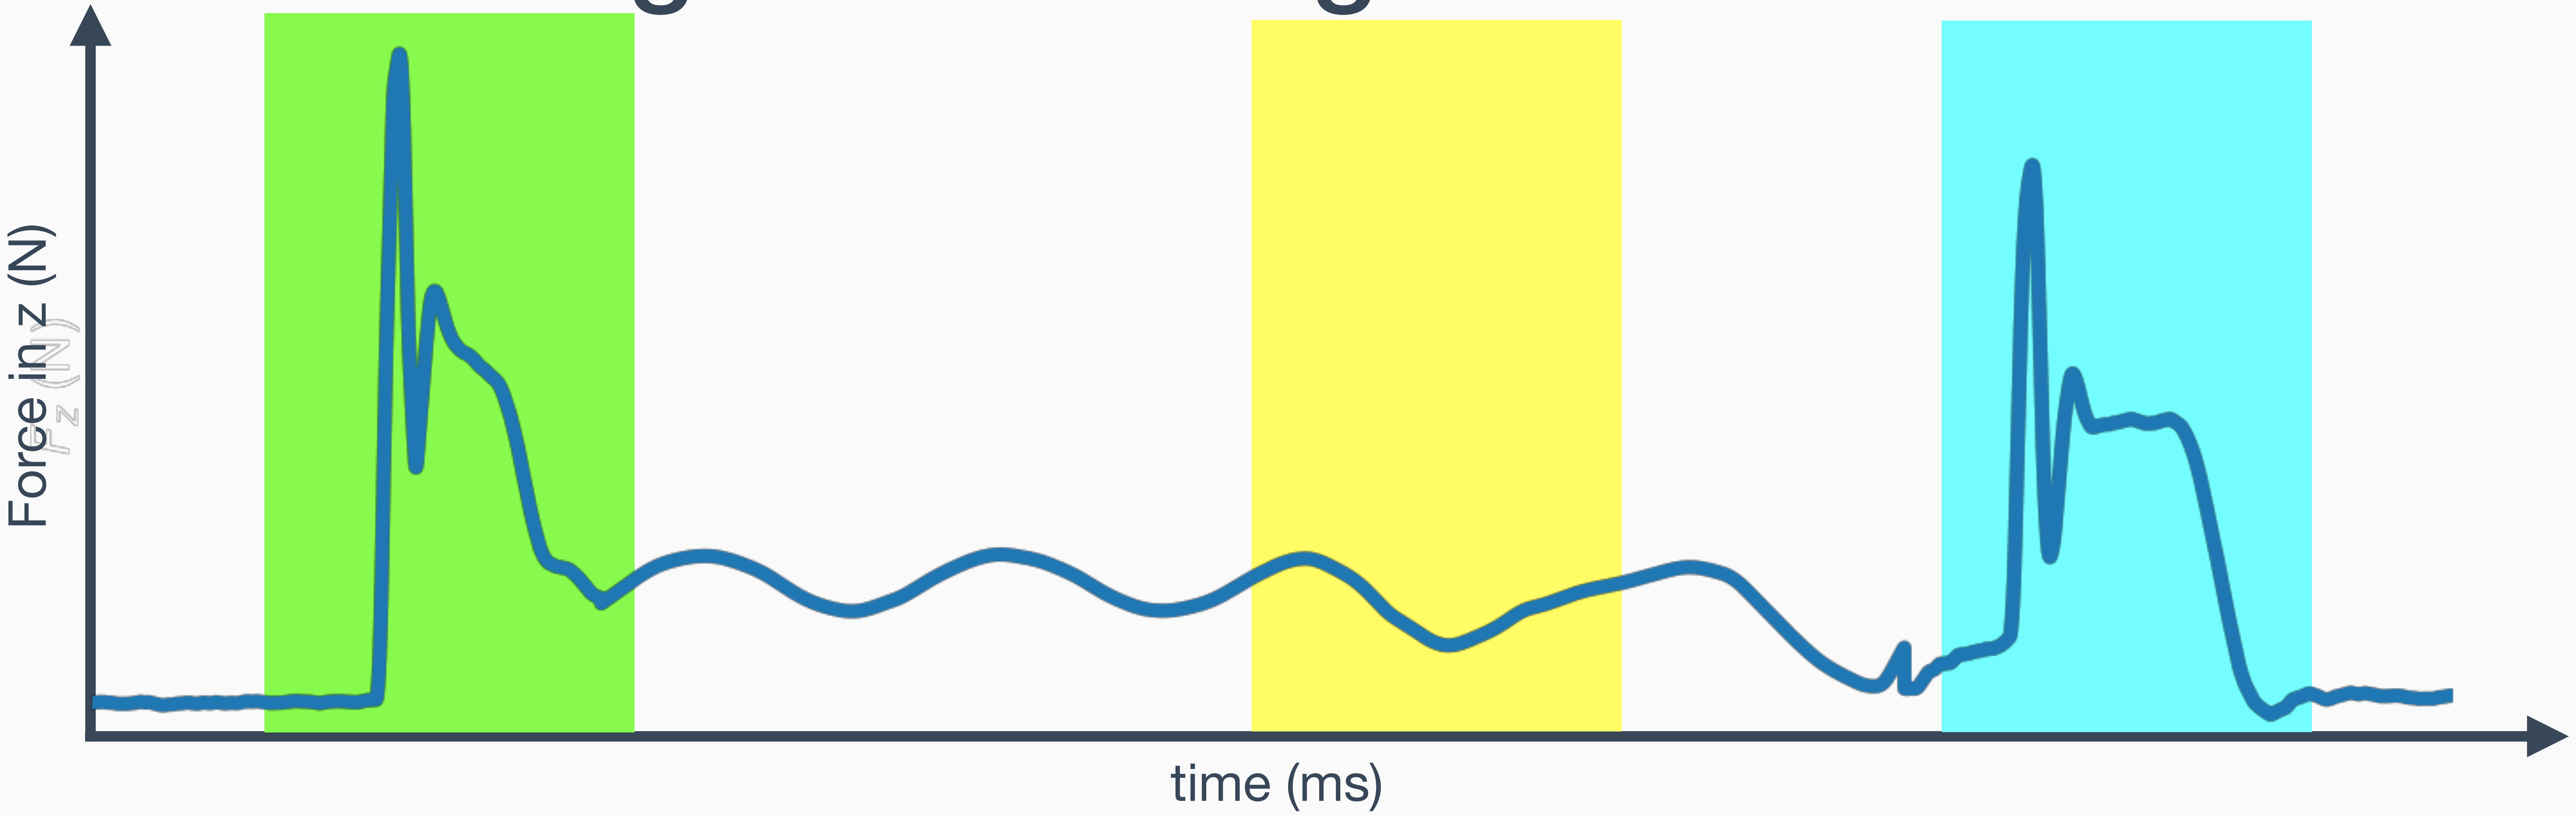

Multimodal Representation for Manipulation

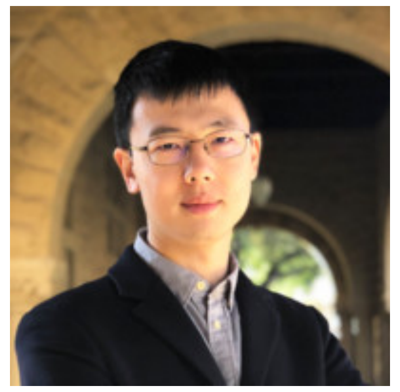

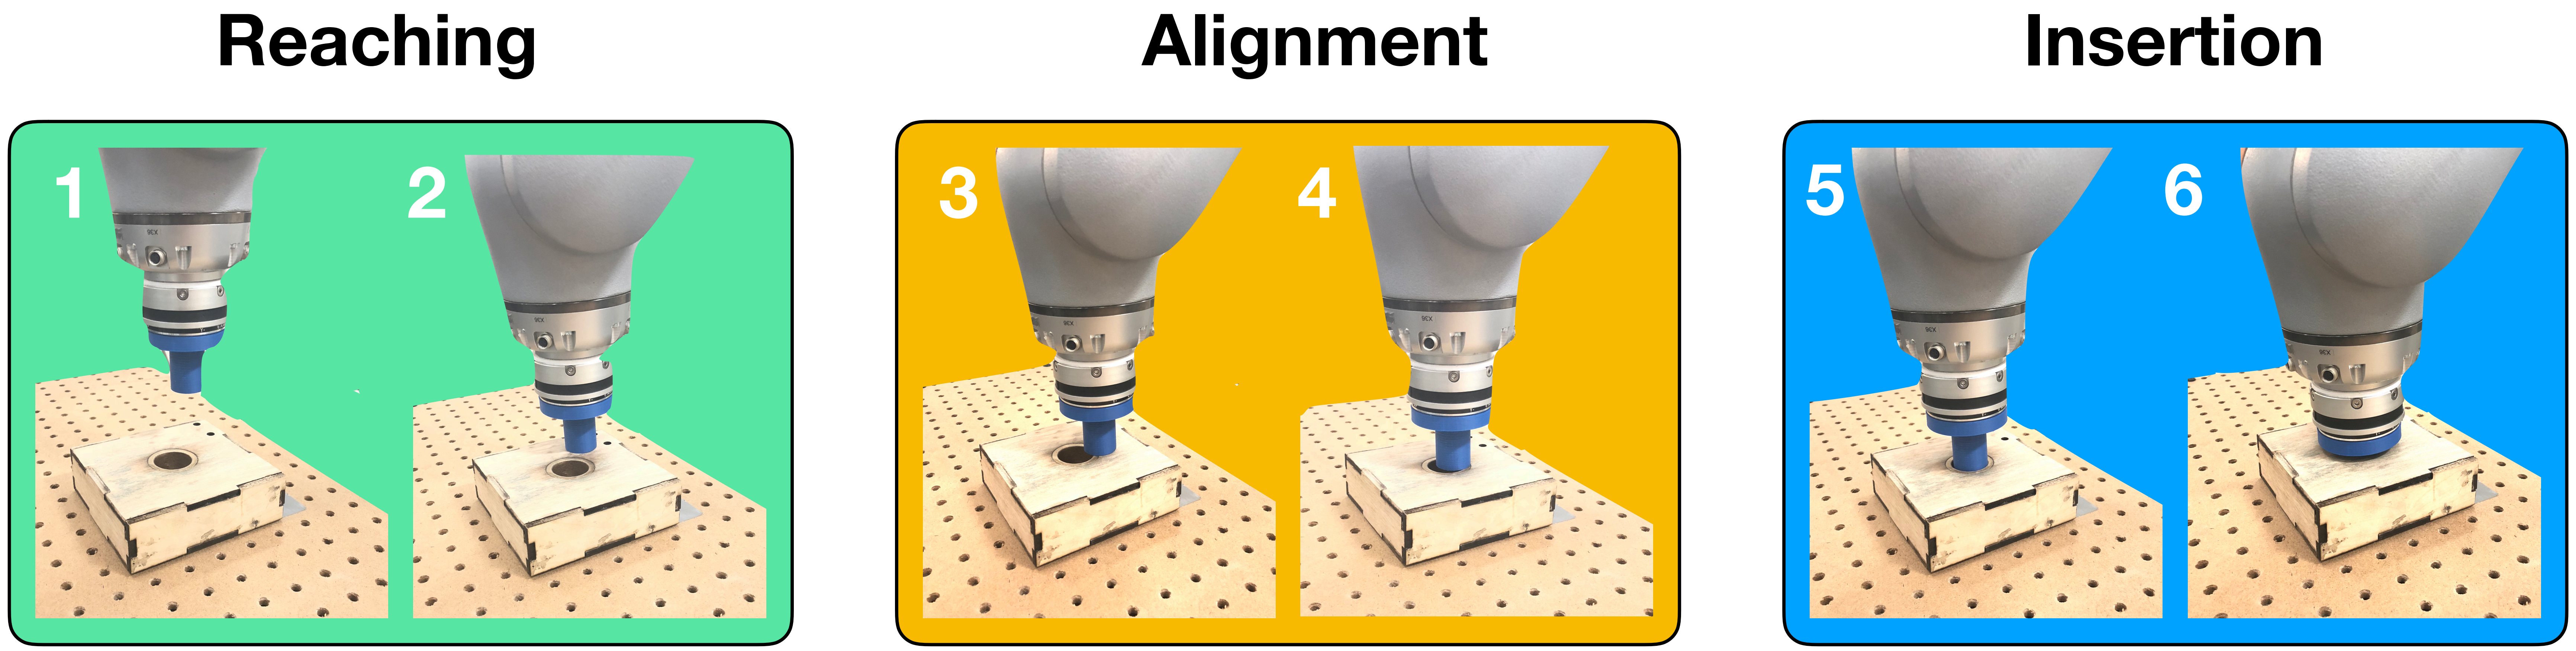

Peg Insertion

RGB Images + End Effector F/T

Michelle Lee, Yuke Zhu et al, “Making Sense of Vision and Touch: Self-Supervised Learning of Multimodal Representations for Contact-Rich Tasks” 2018. Submitted to ICRA. Pre-print on arXiv.

Learning a policy that leverages Vision & Touch

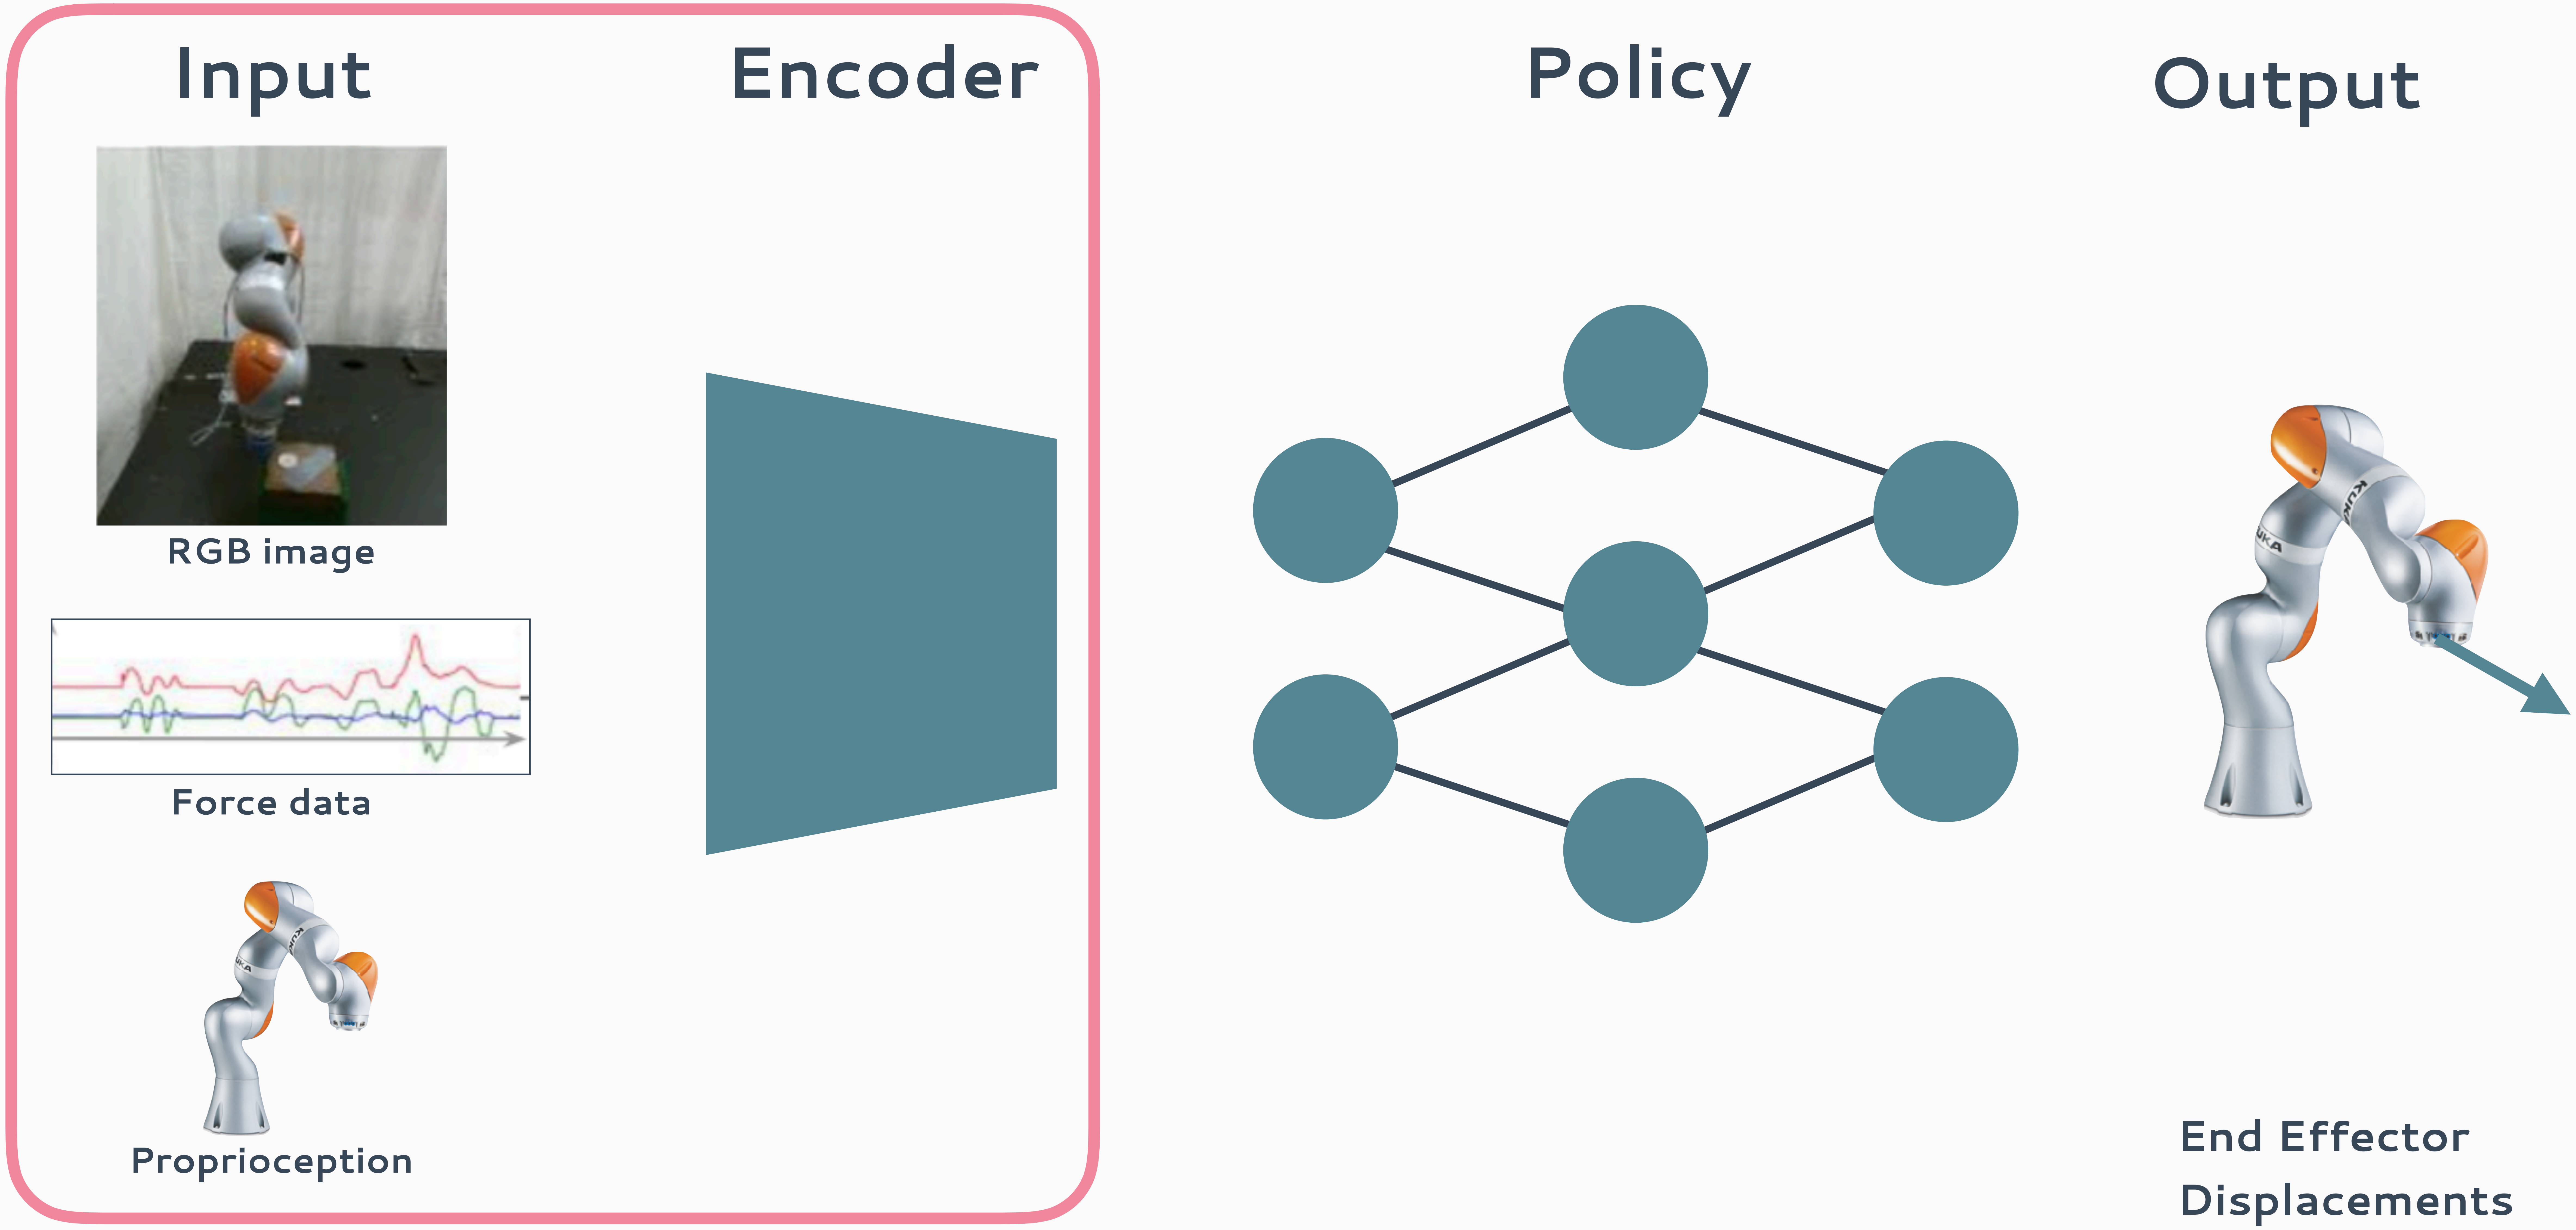

Representation Learning

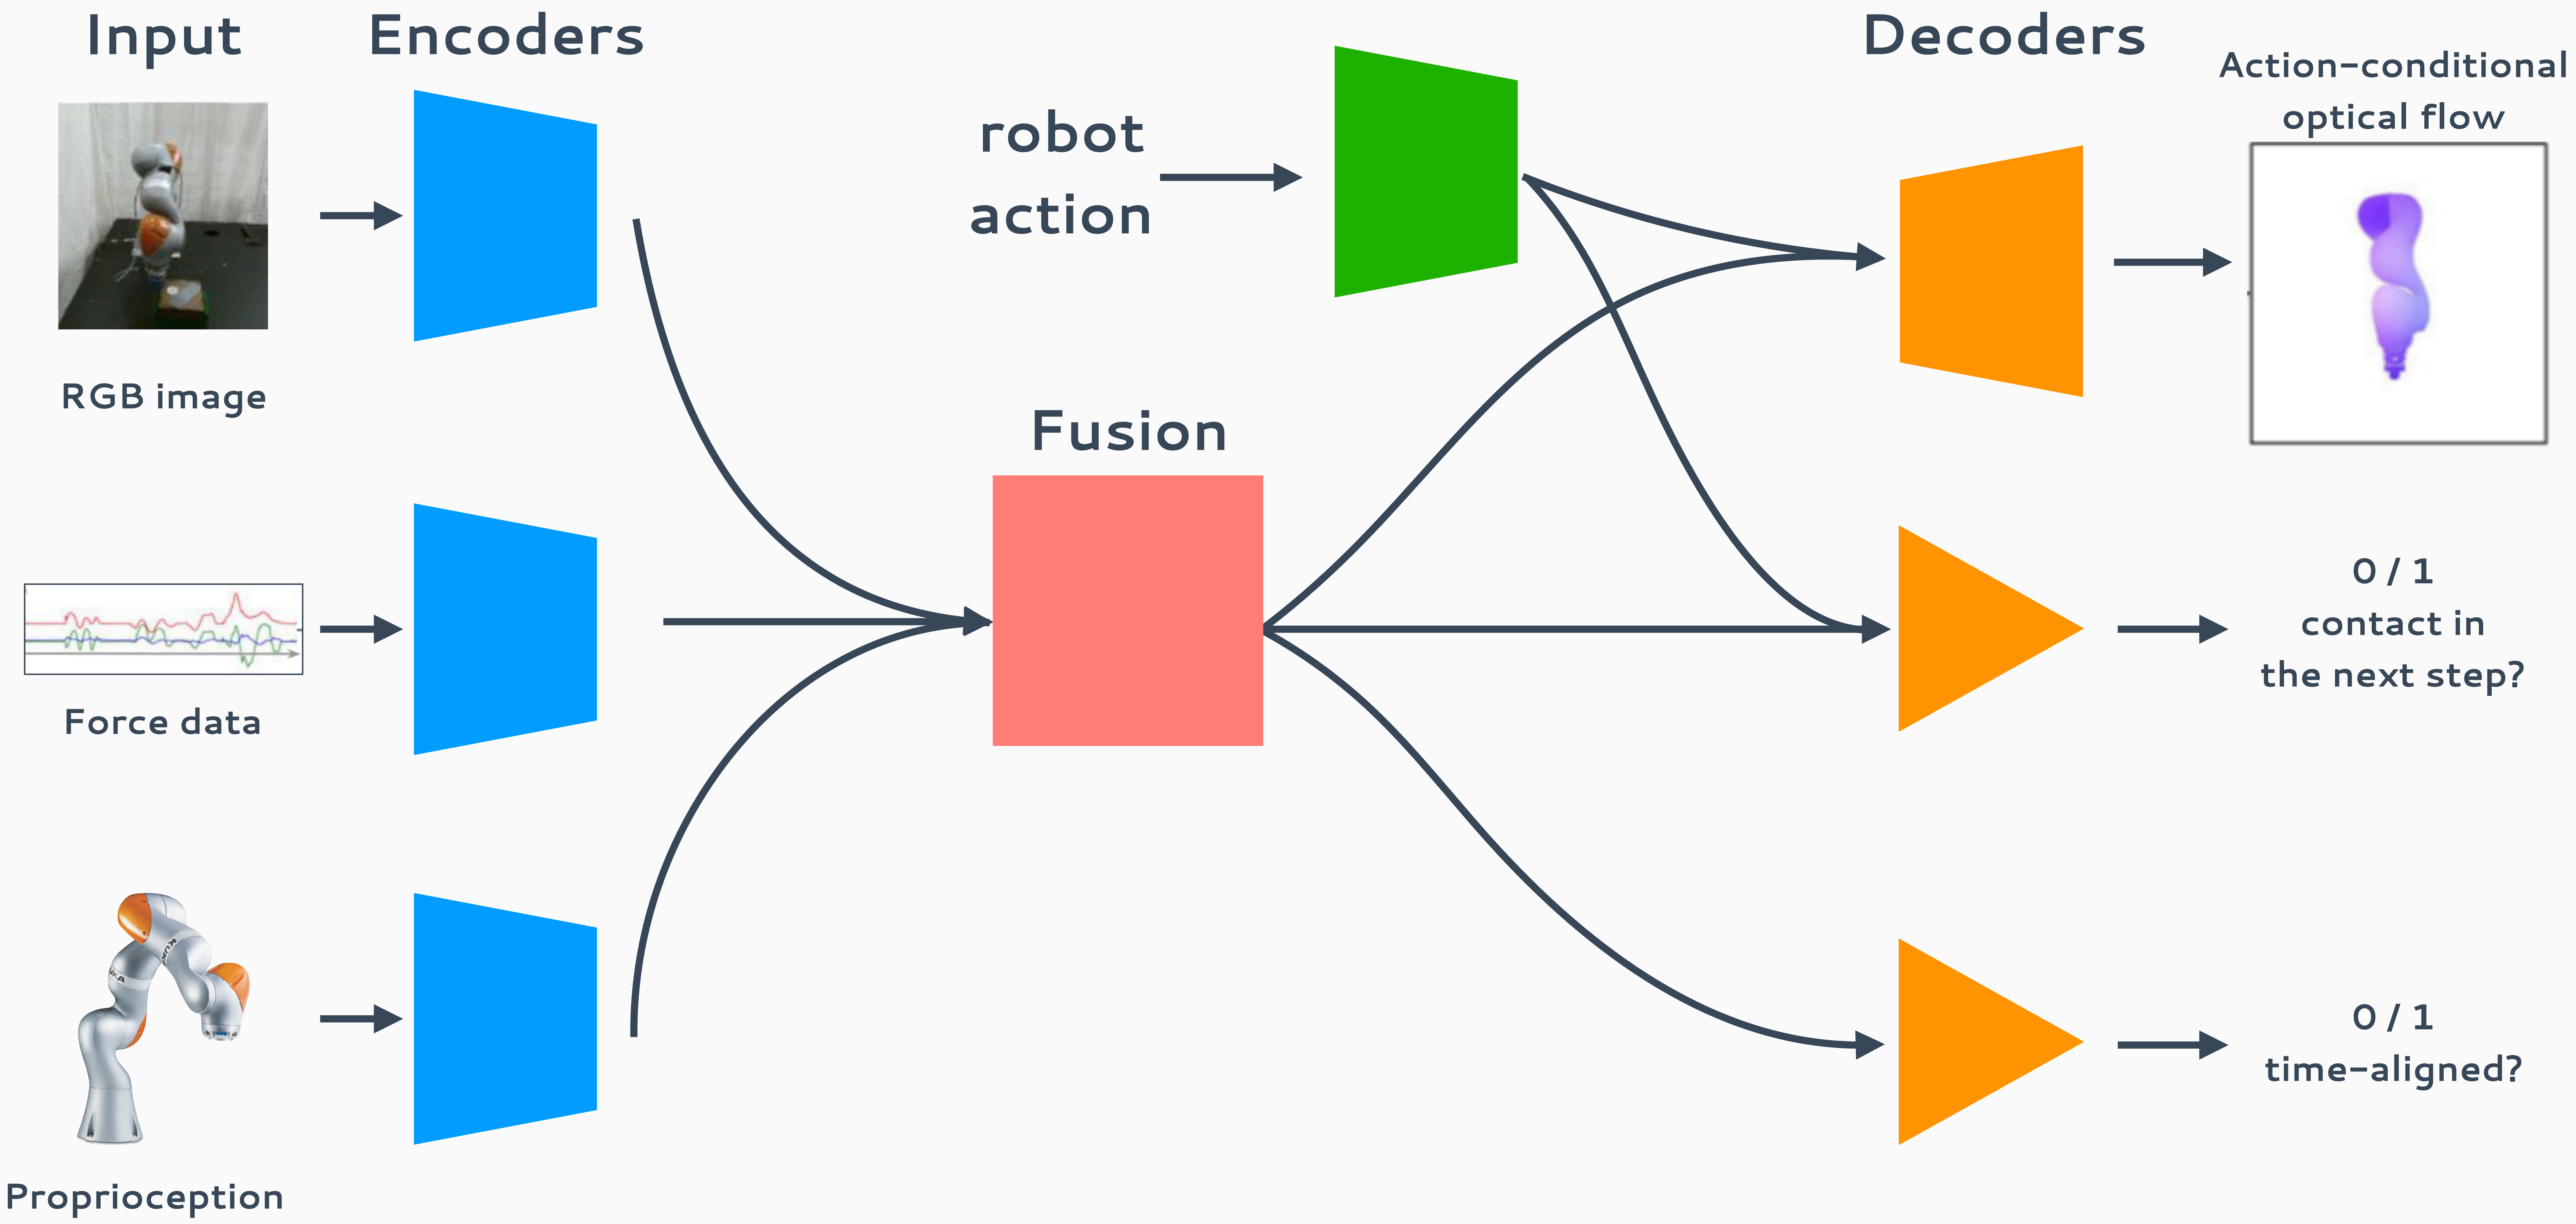

Learning a Policy based on this Representation

Input

RGB image

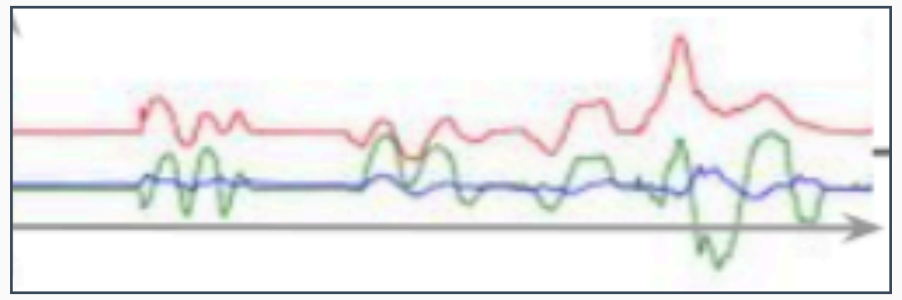

Force data

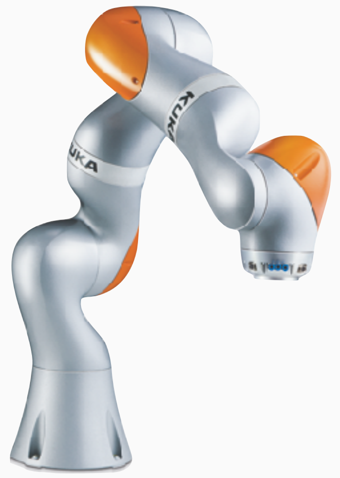

Proprioception

Encoder

Policy

TRPO

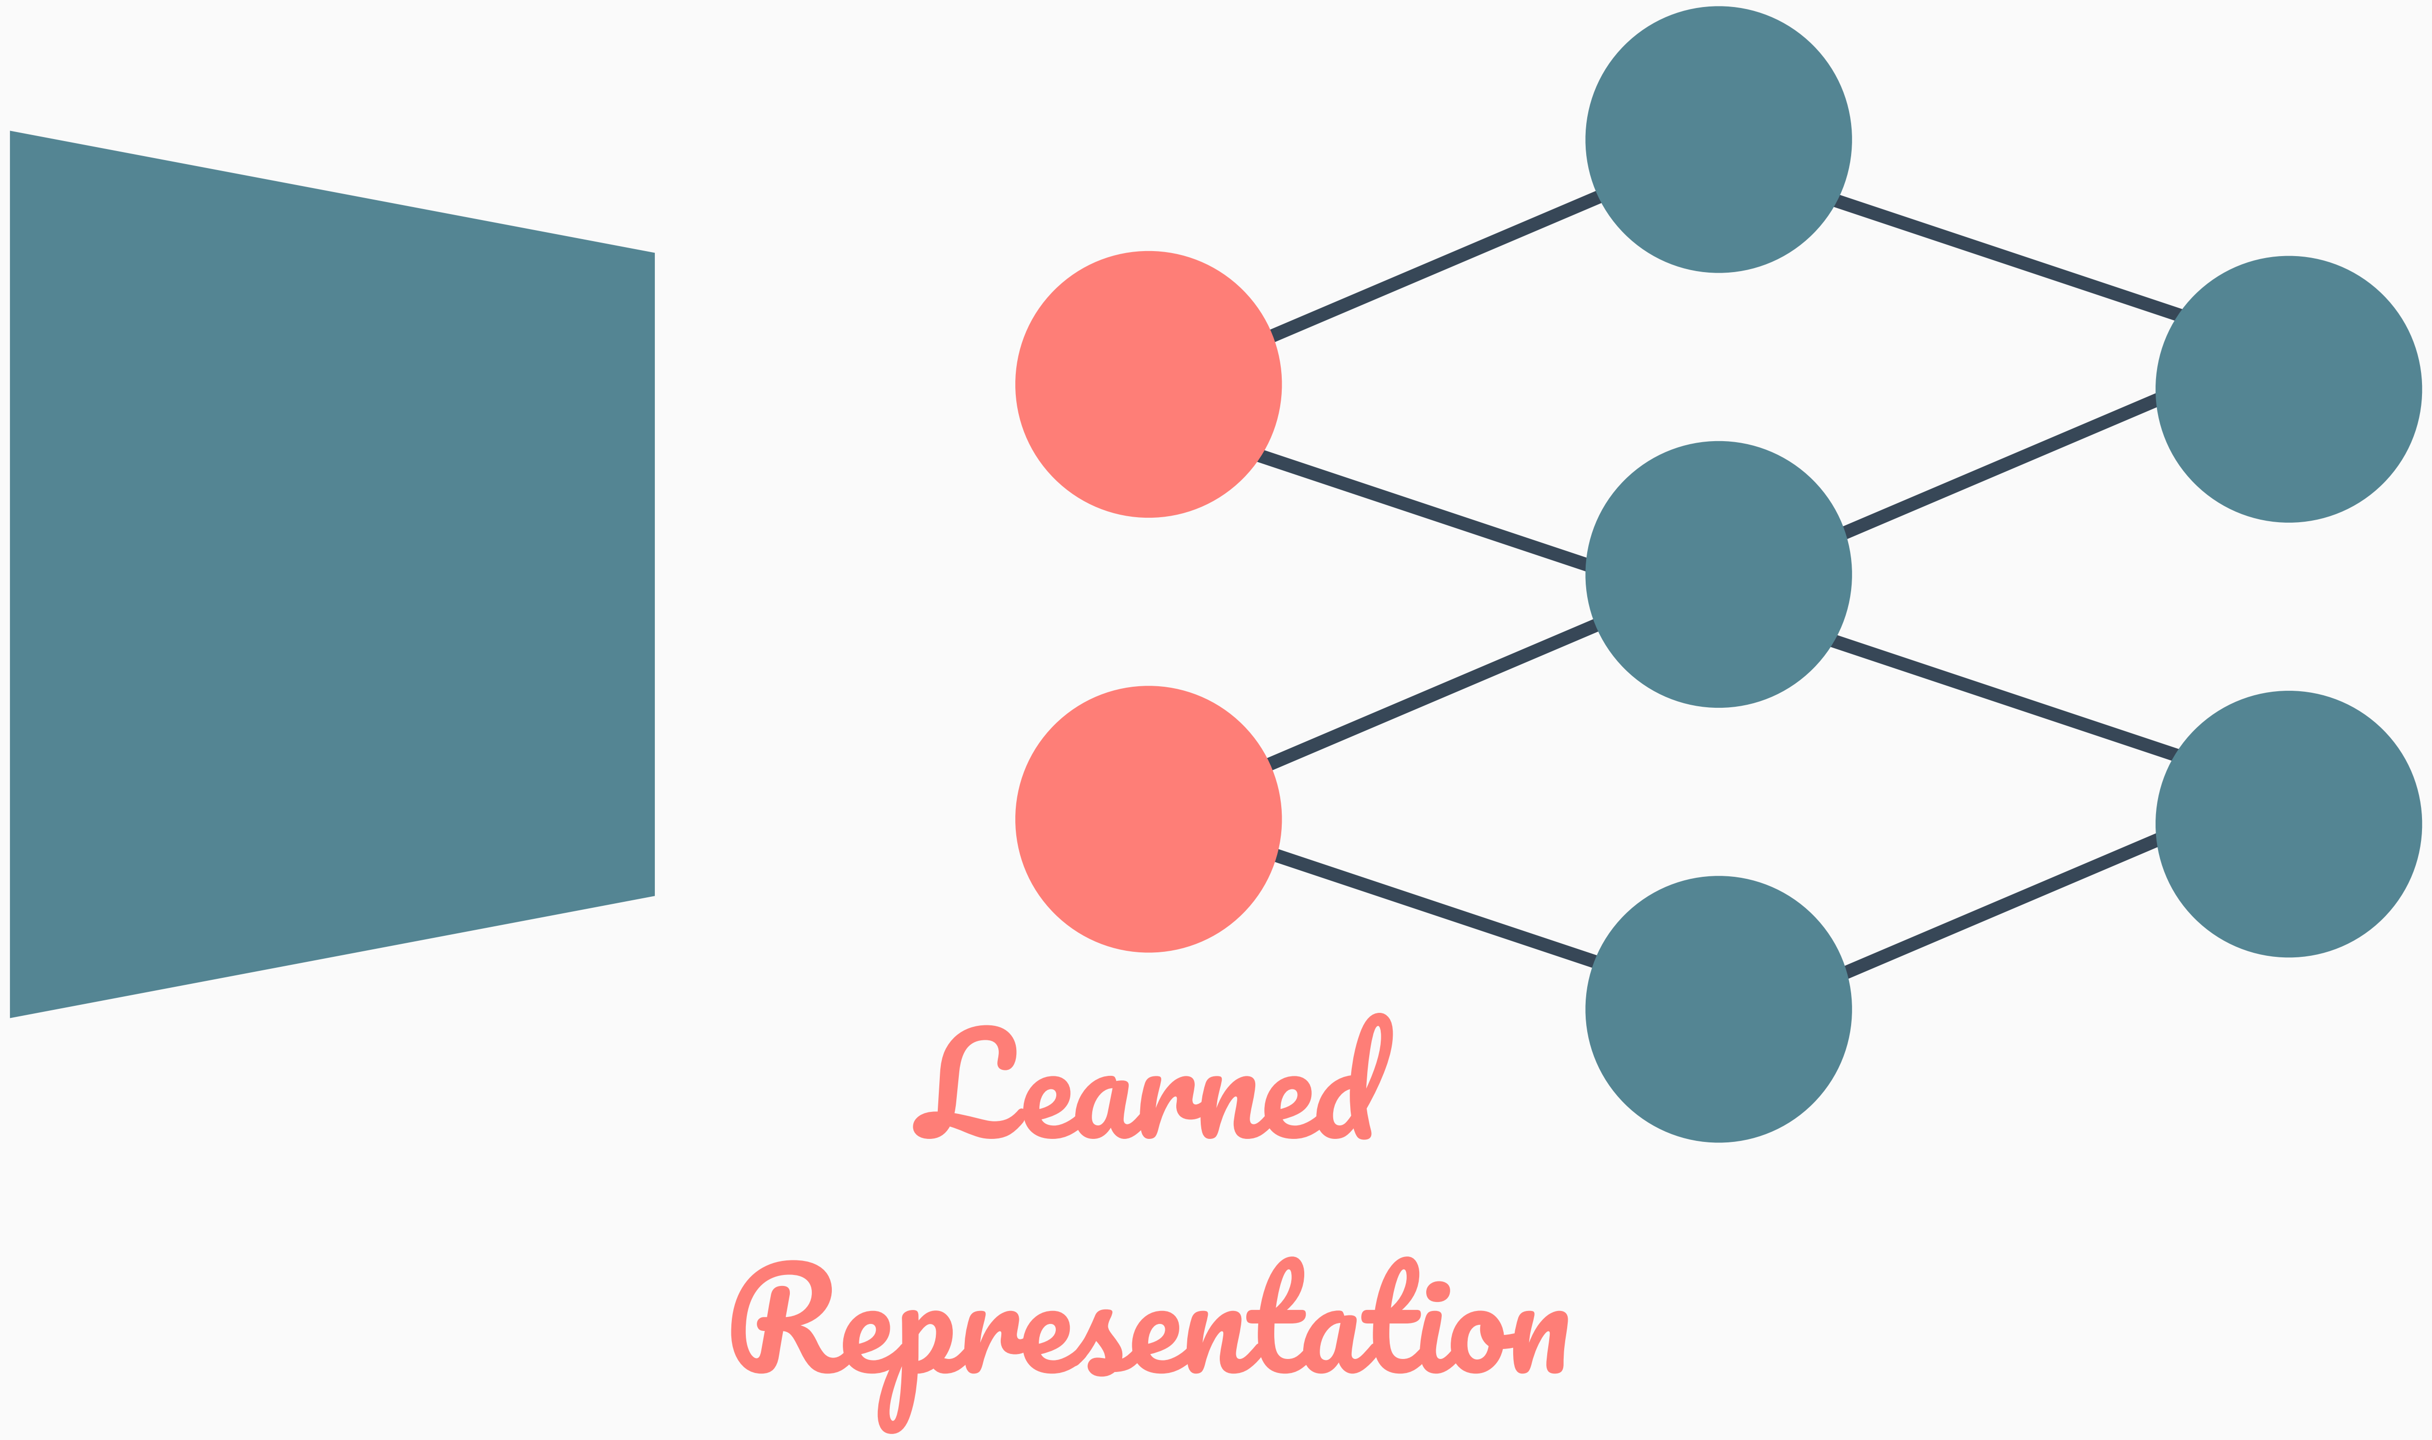

Output

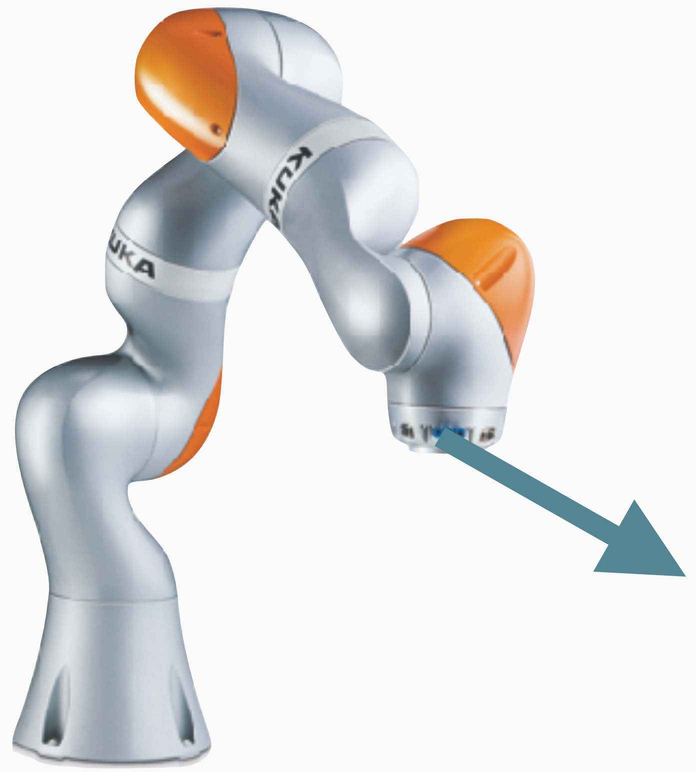

End Effector
Displacements

Ablation Results in Simulation

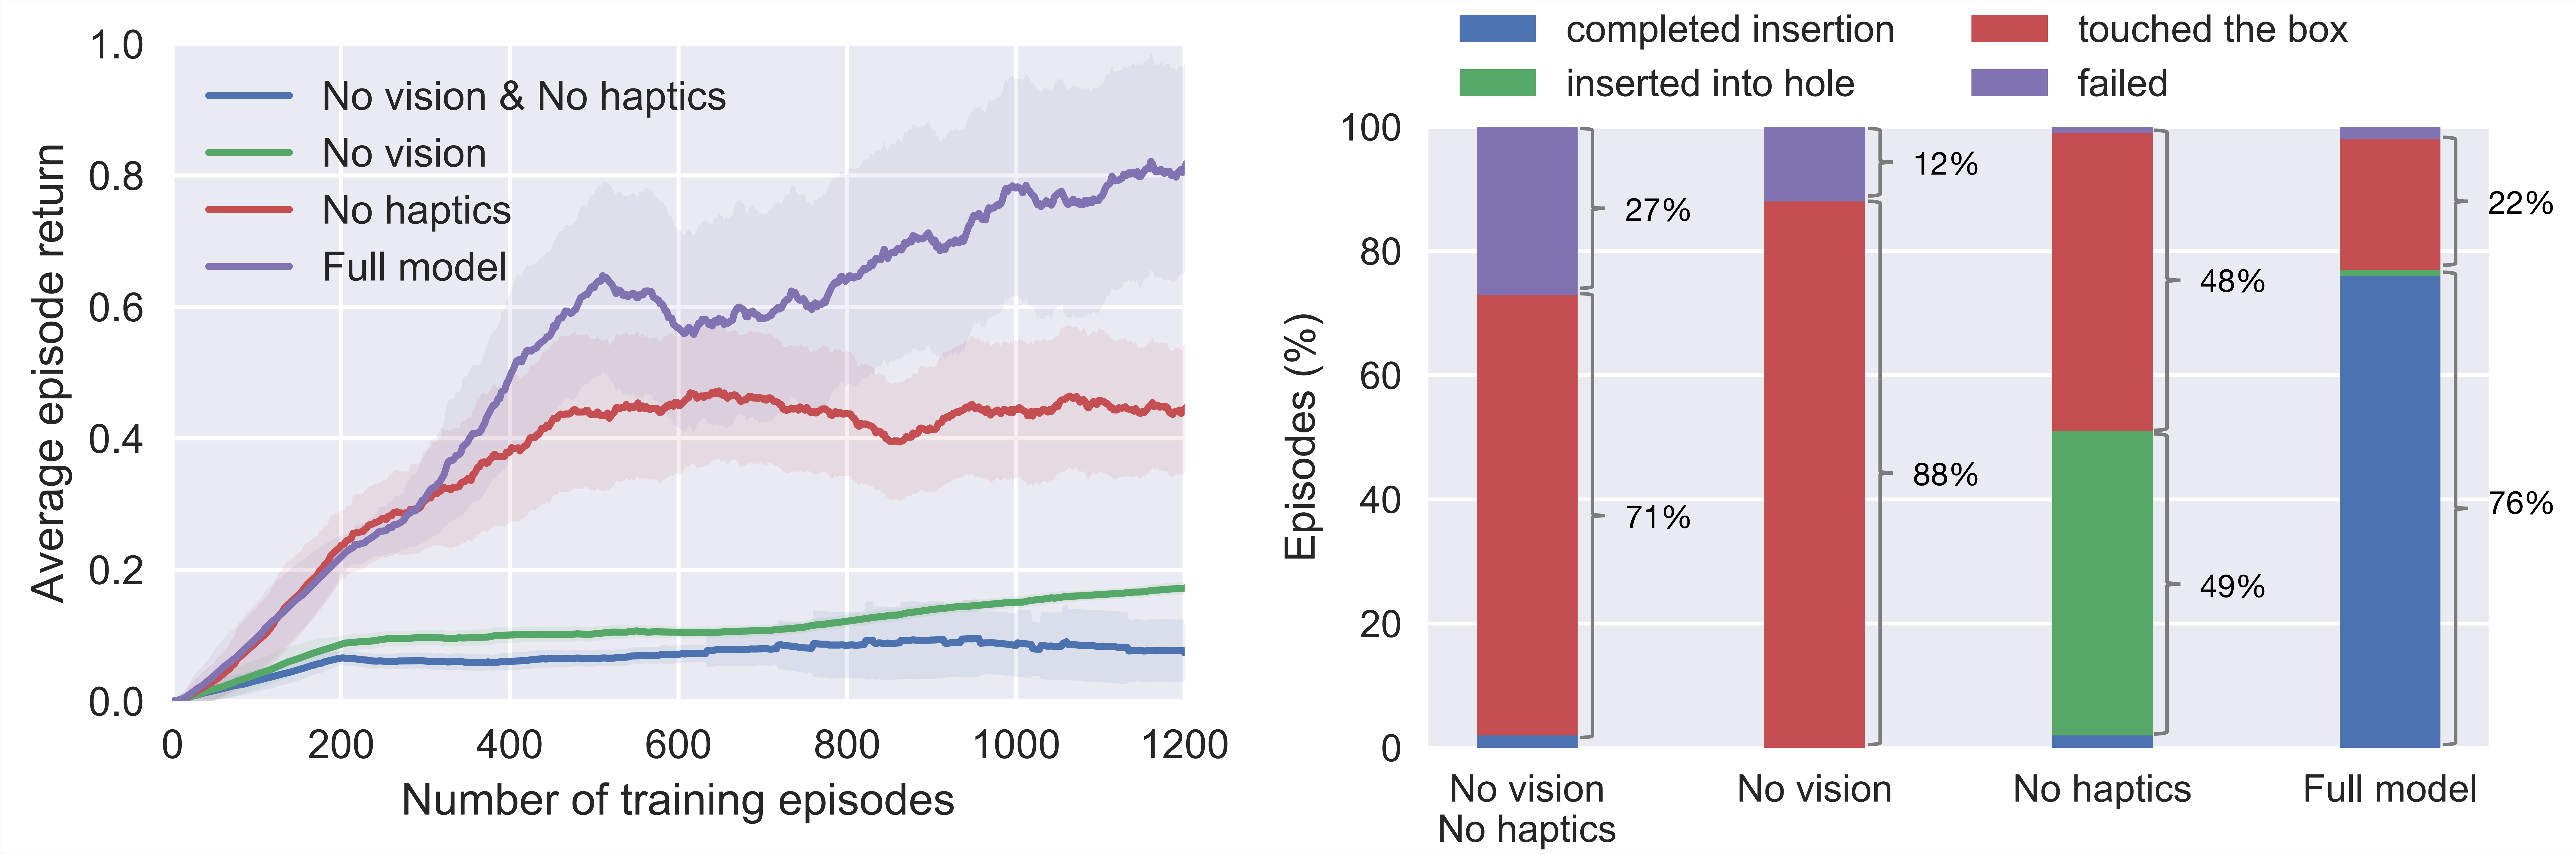

Generalisation Results in Real World

Our model works well with different peg shapes.

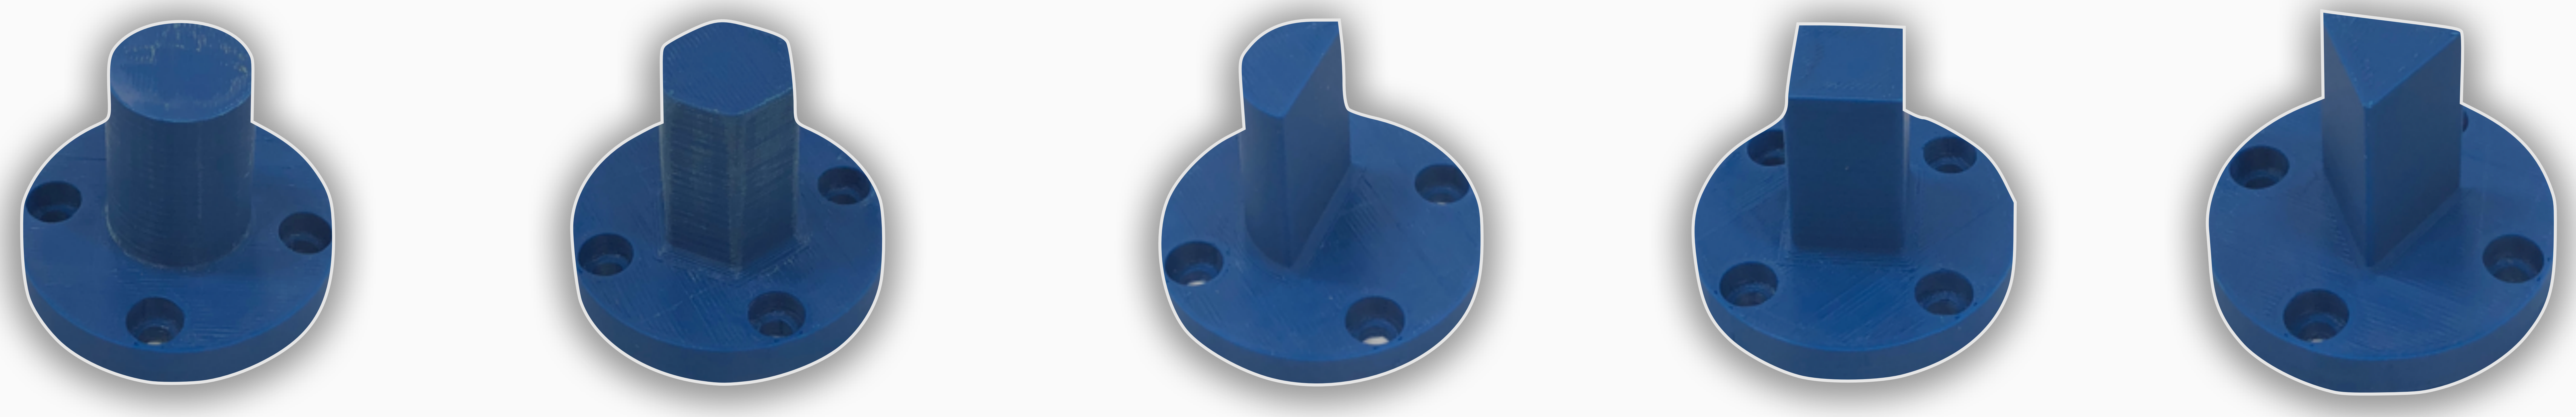

round	hexagonal	semicircular	square	triangular
(2.15mm)	(2.50mm)	(1.85mm)	(2.24mm)	(2.13mm)

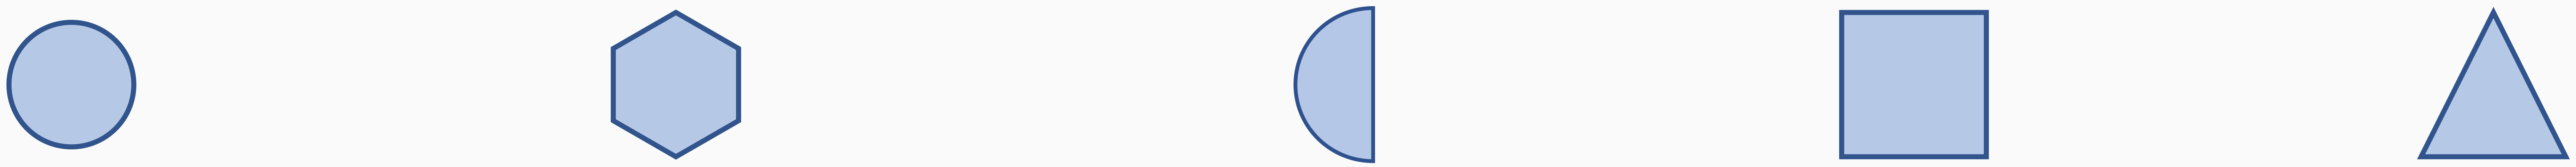

* Clearances are shown as numbers in parentheses.

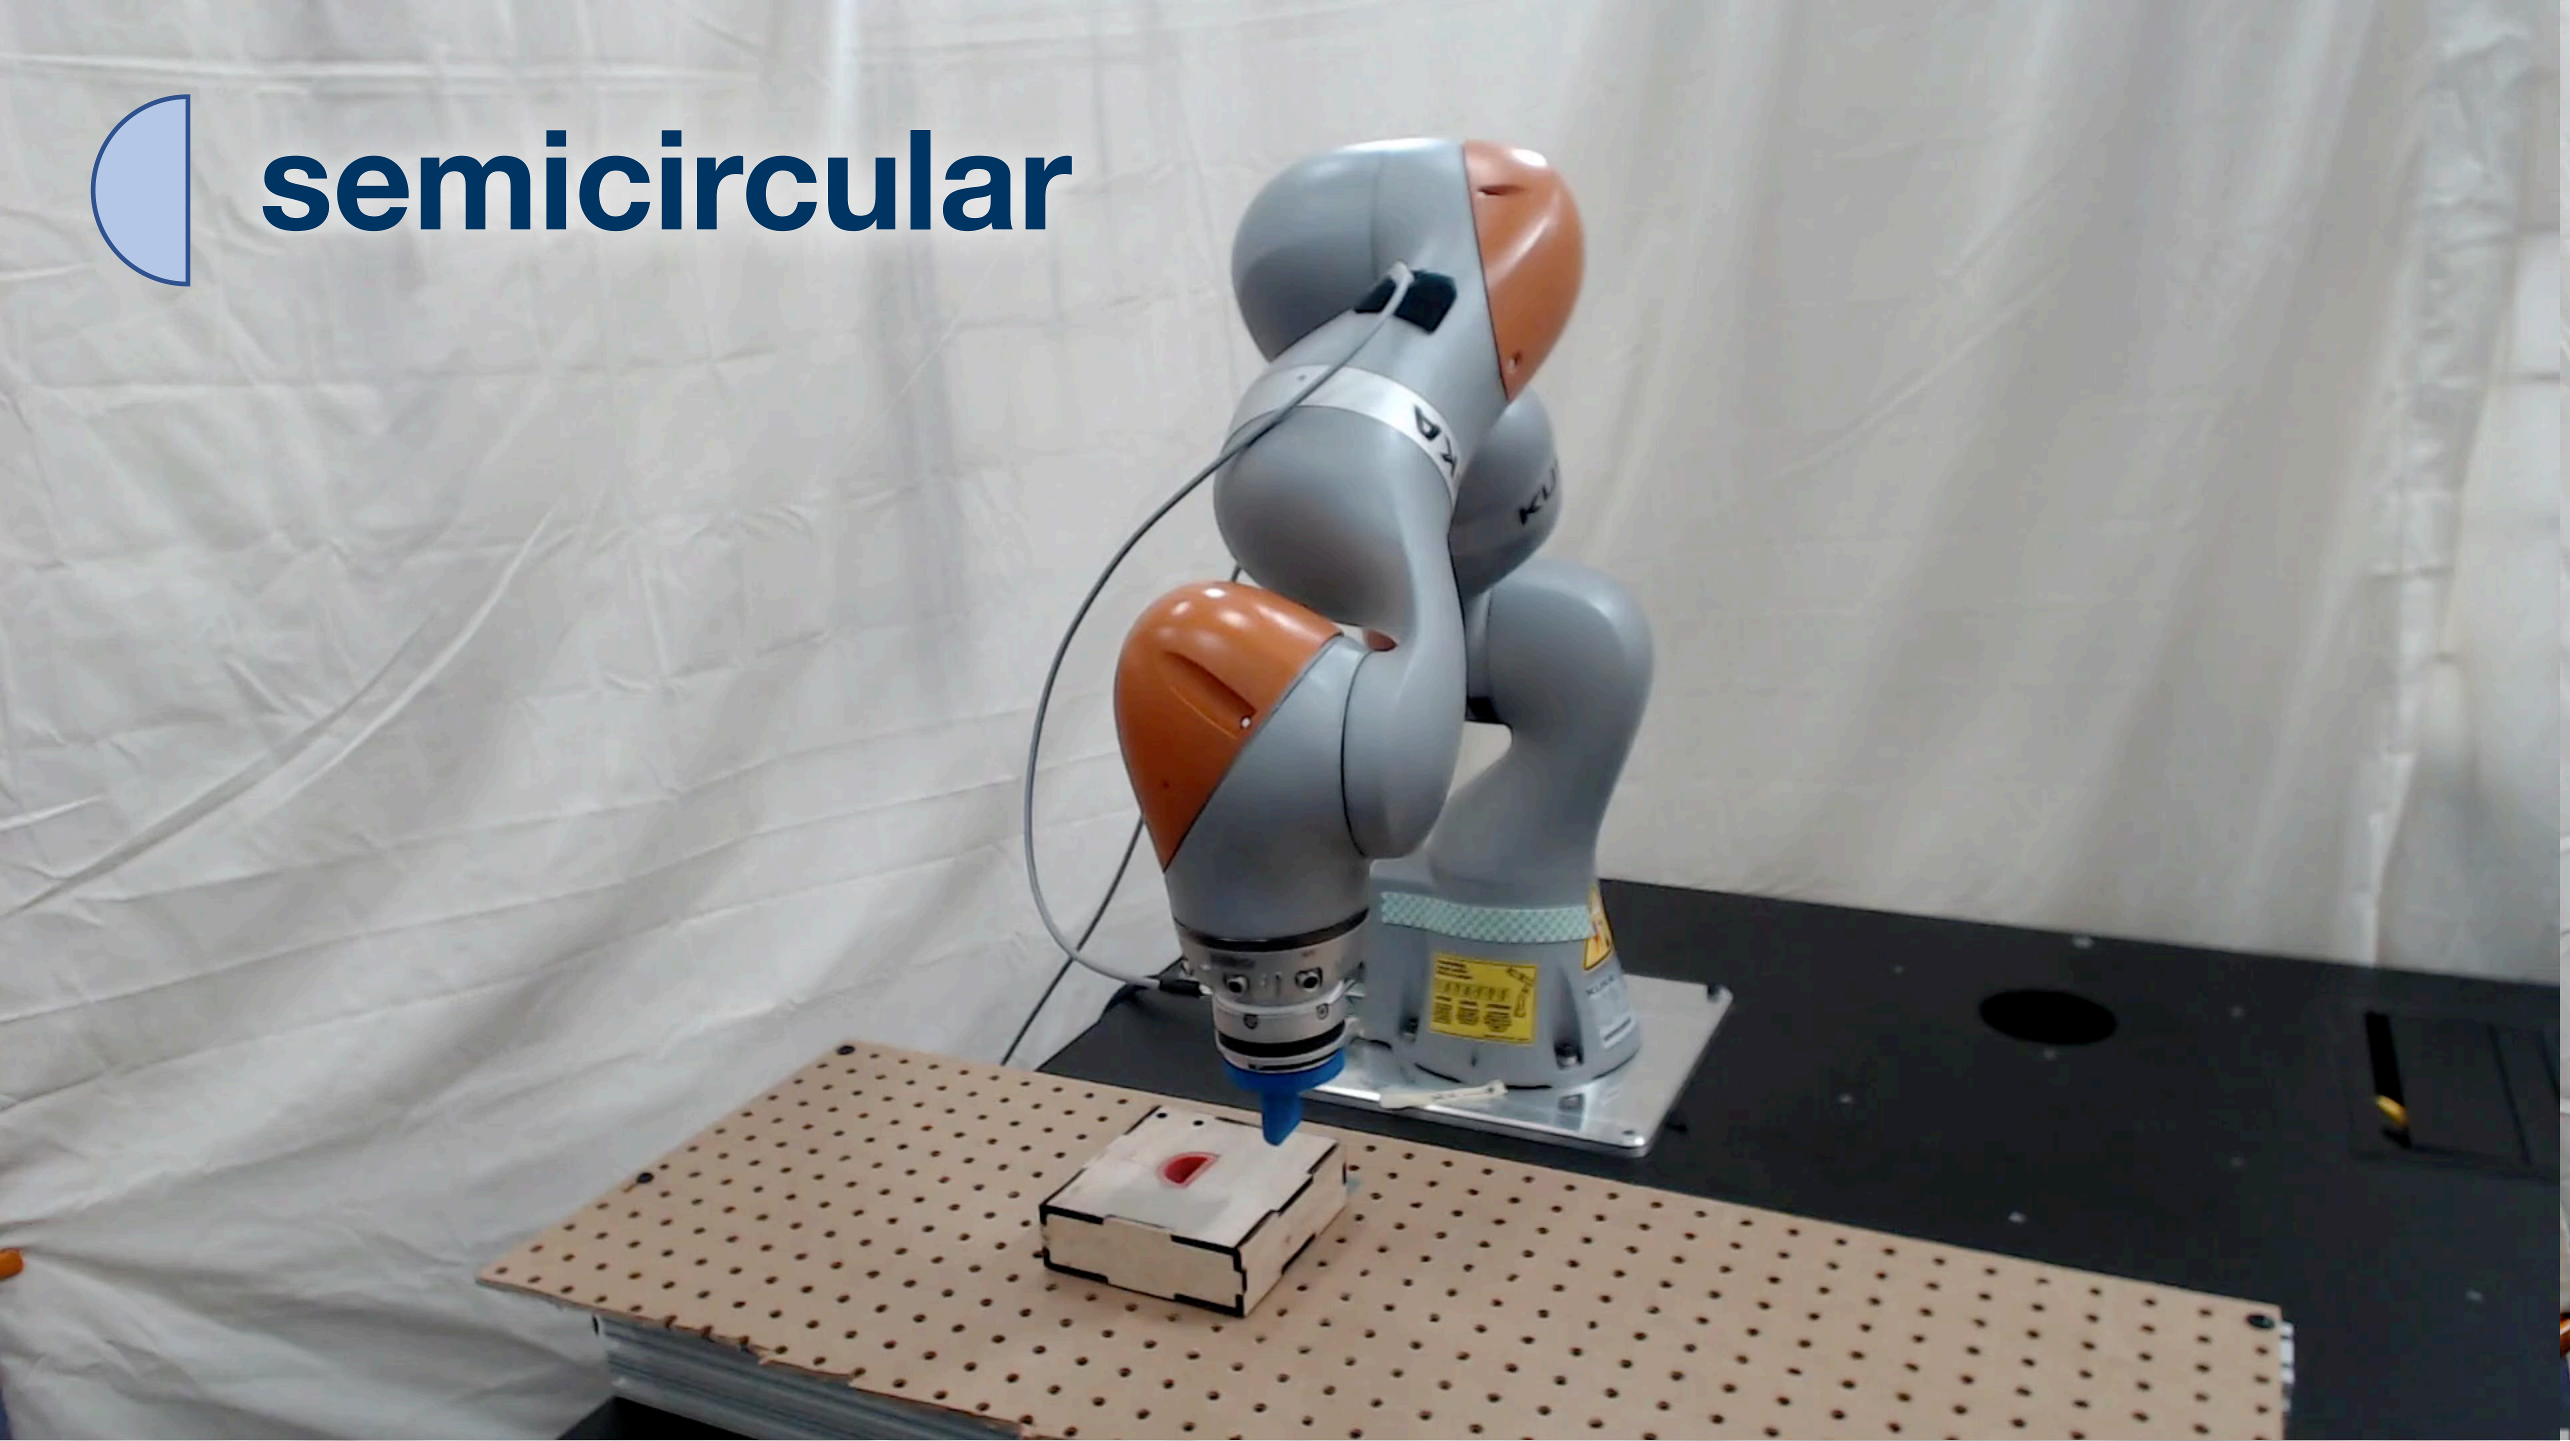

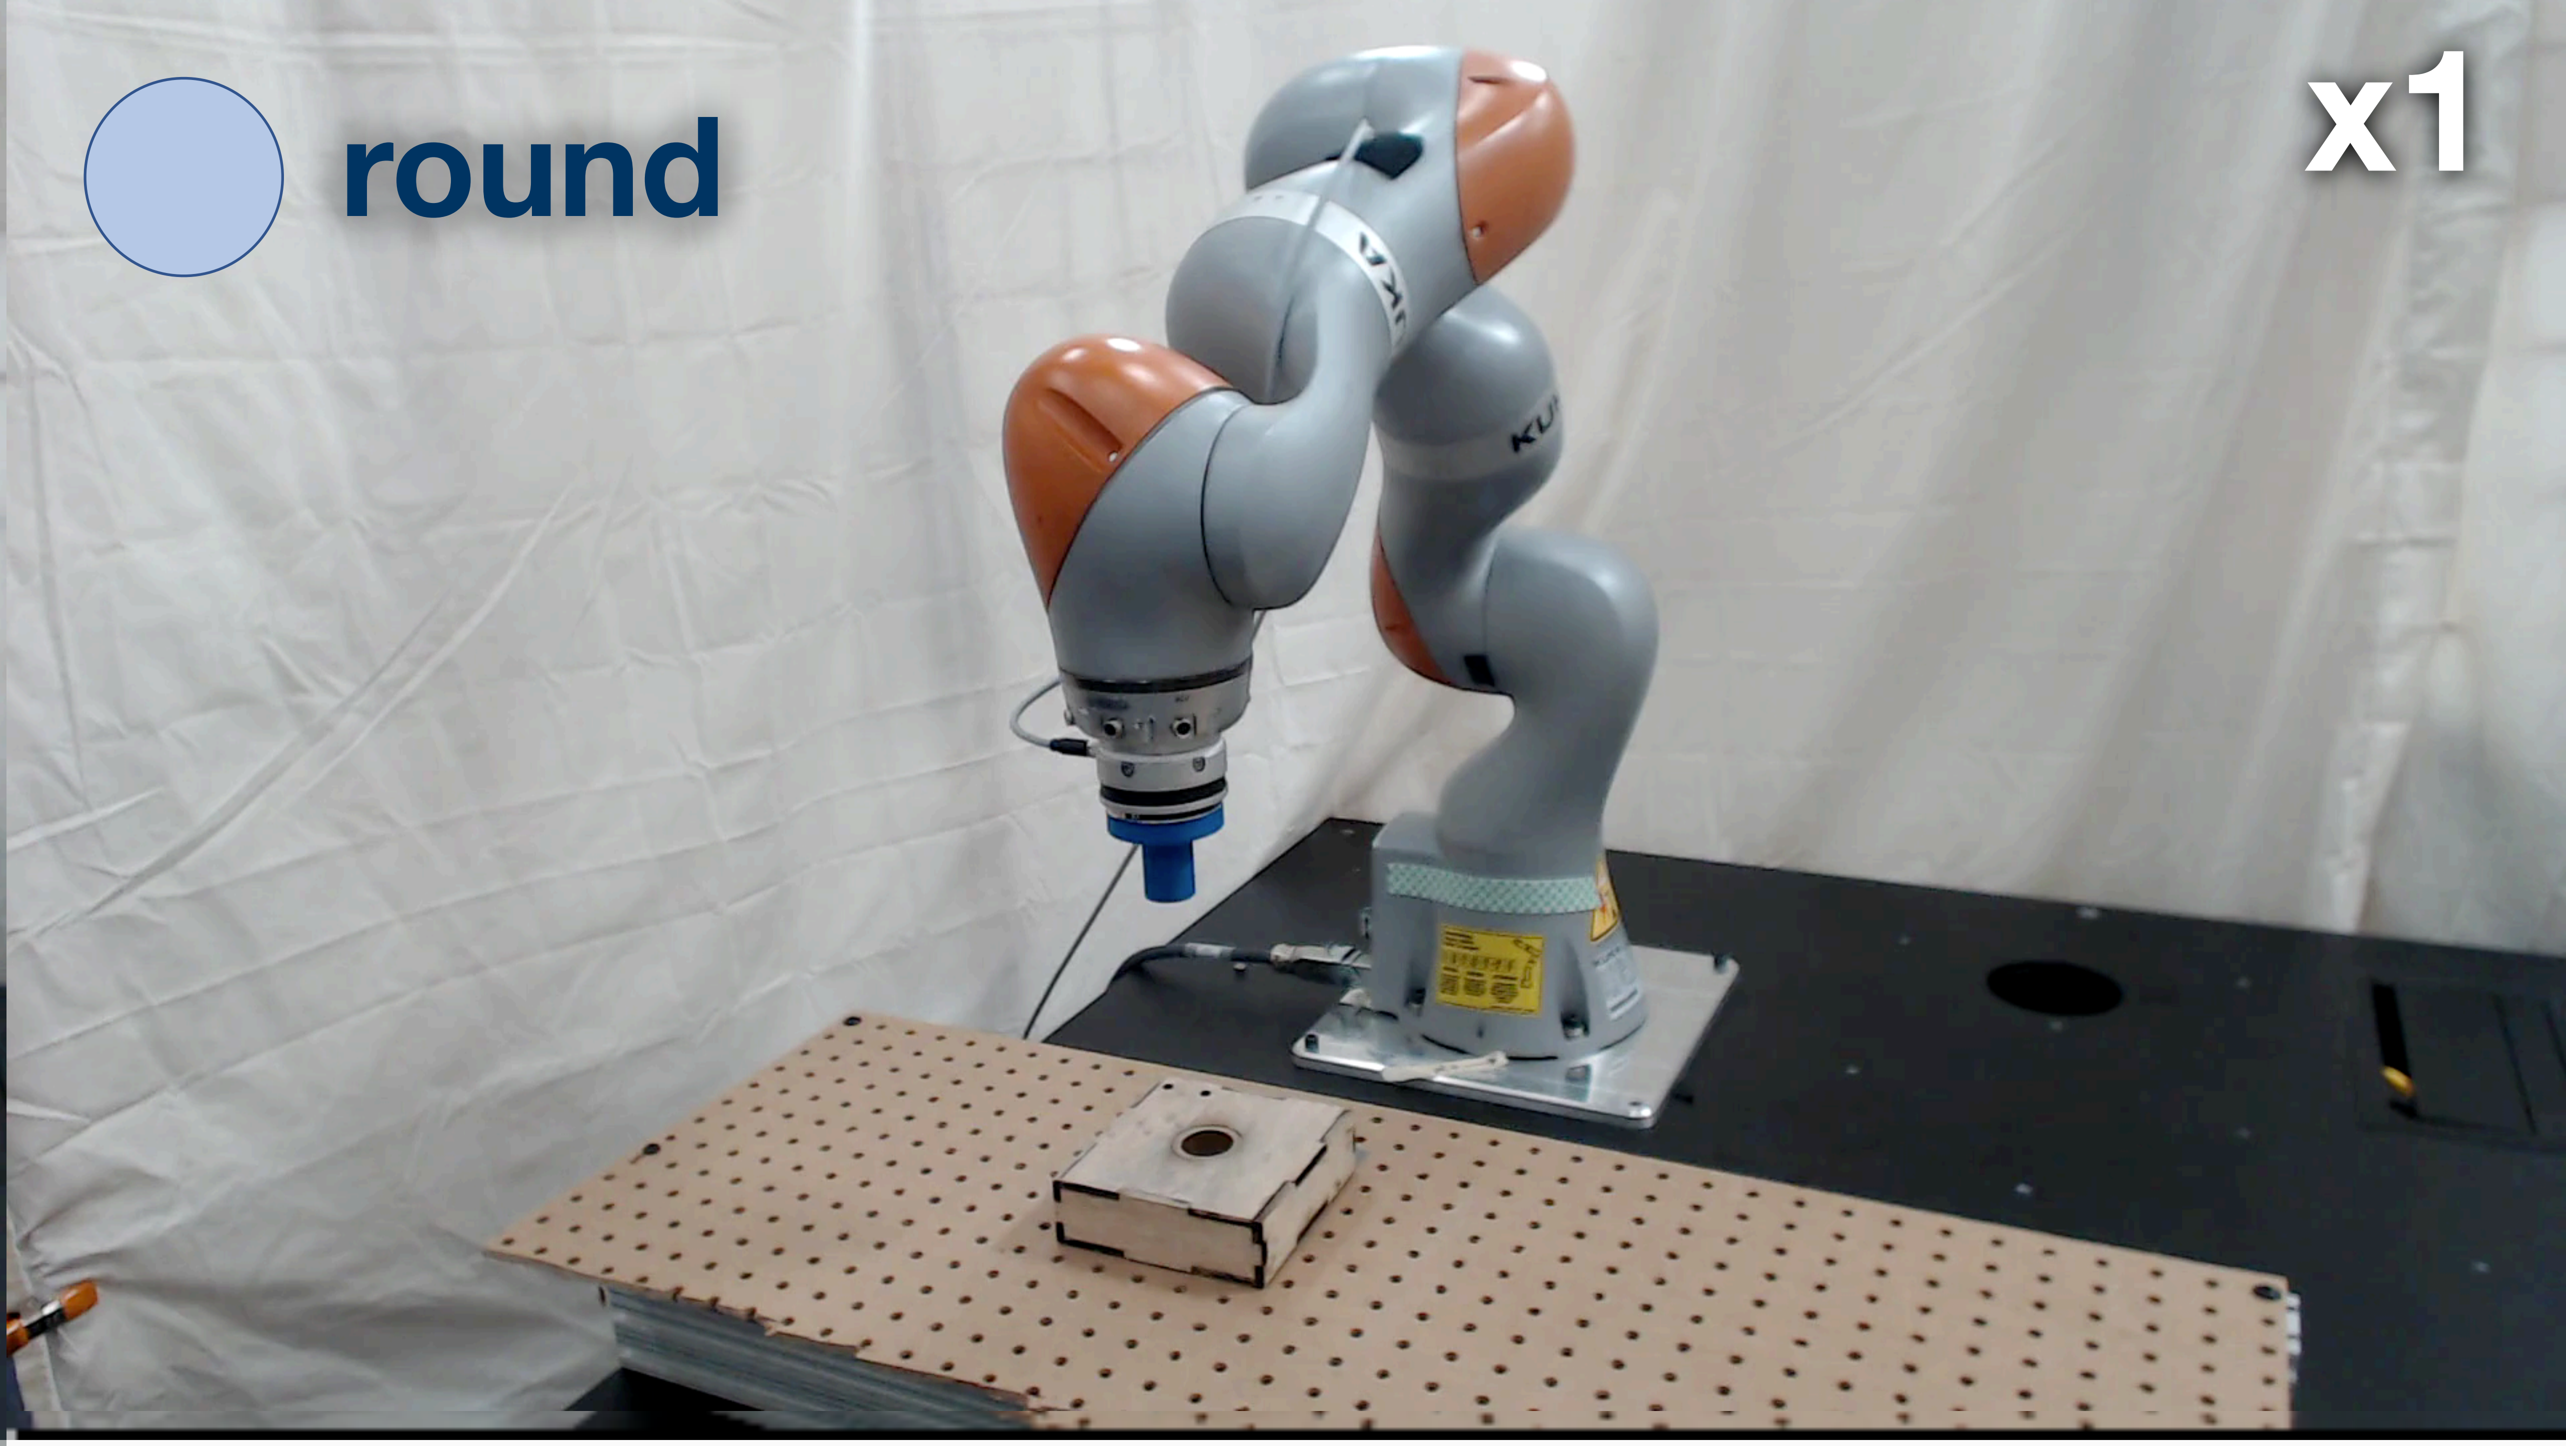

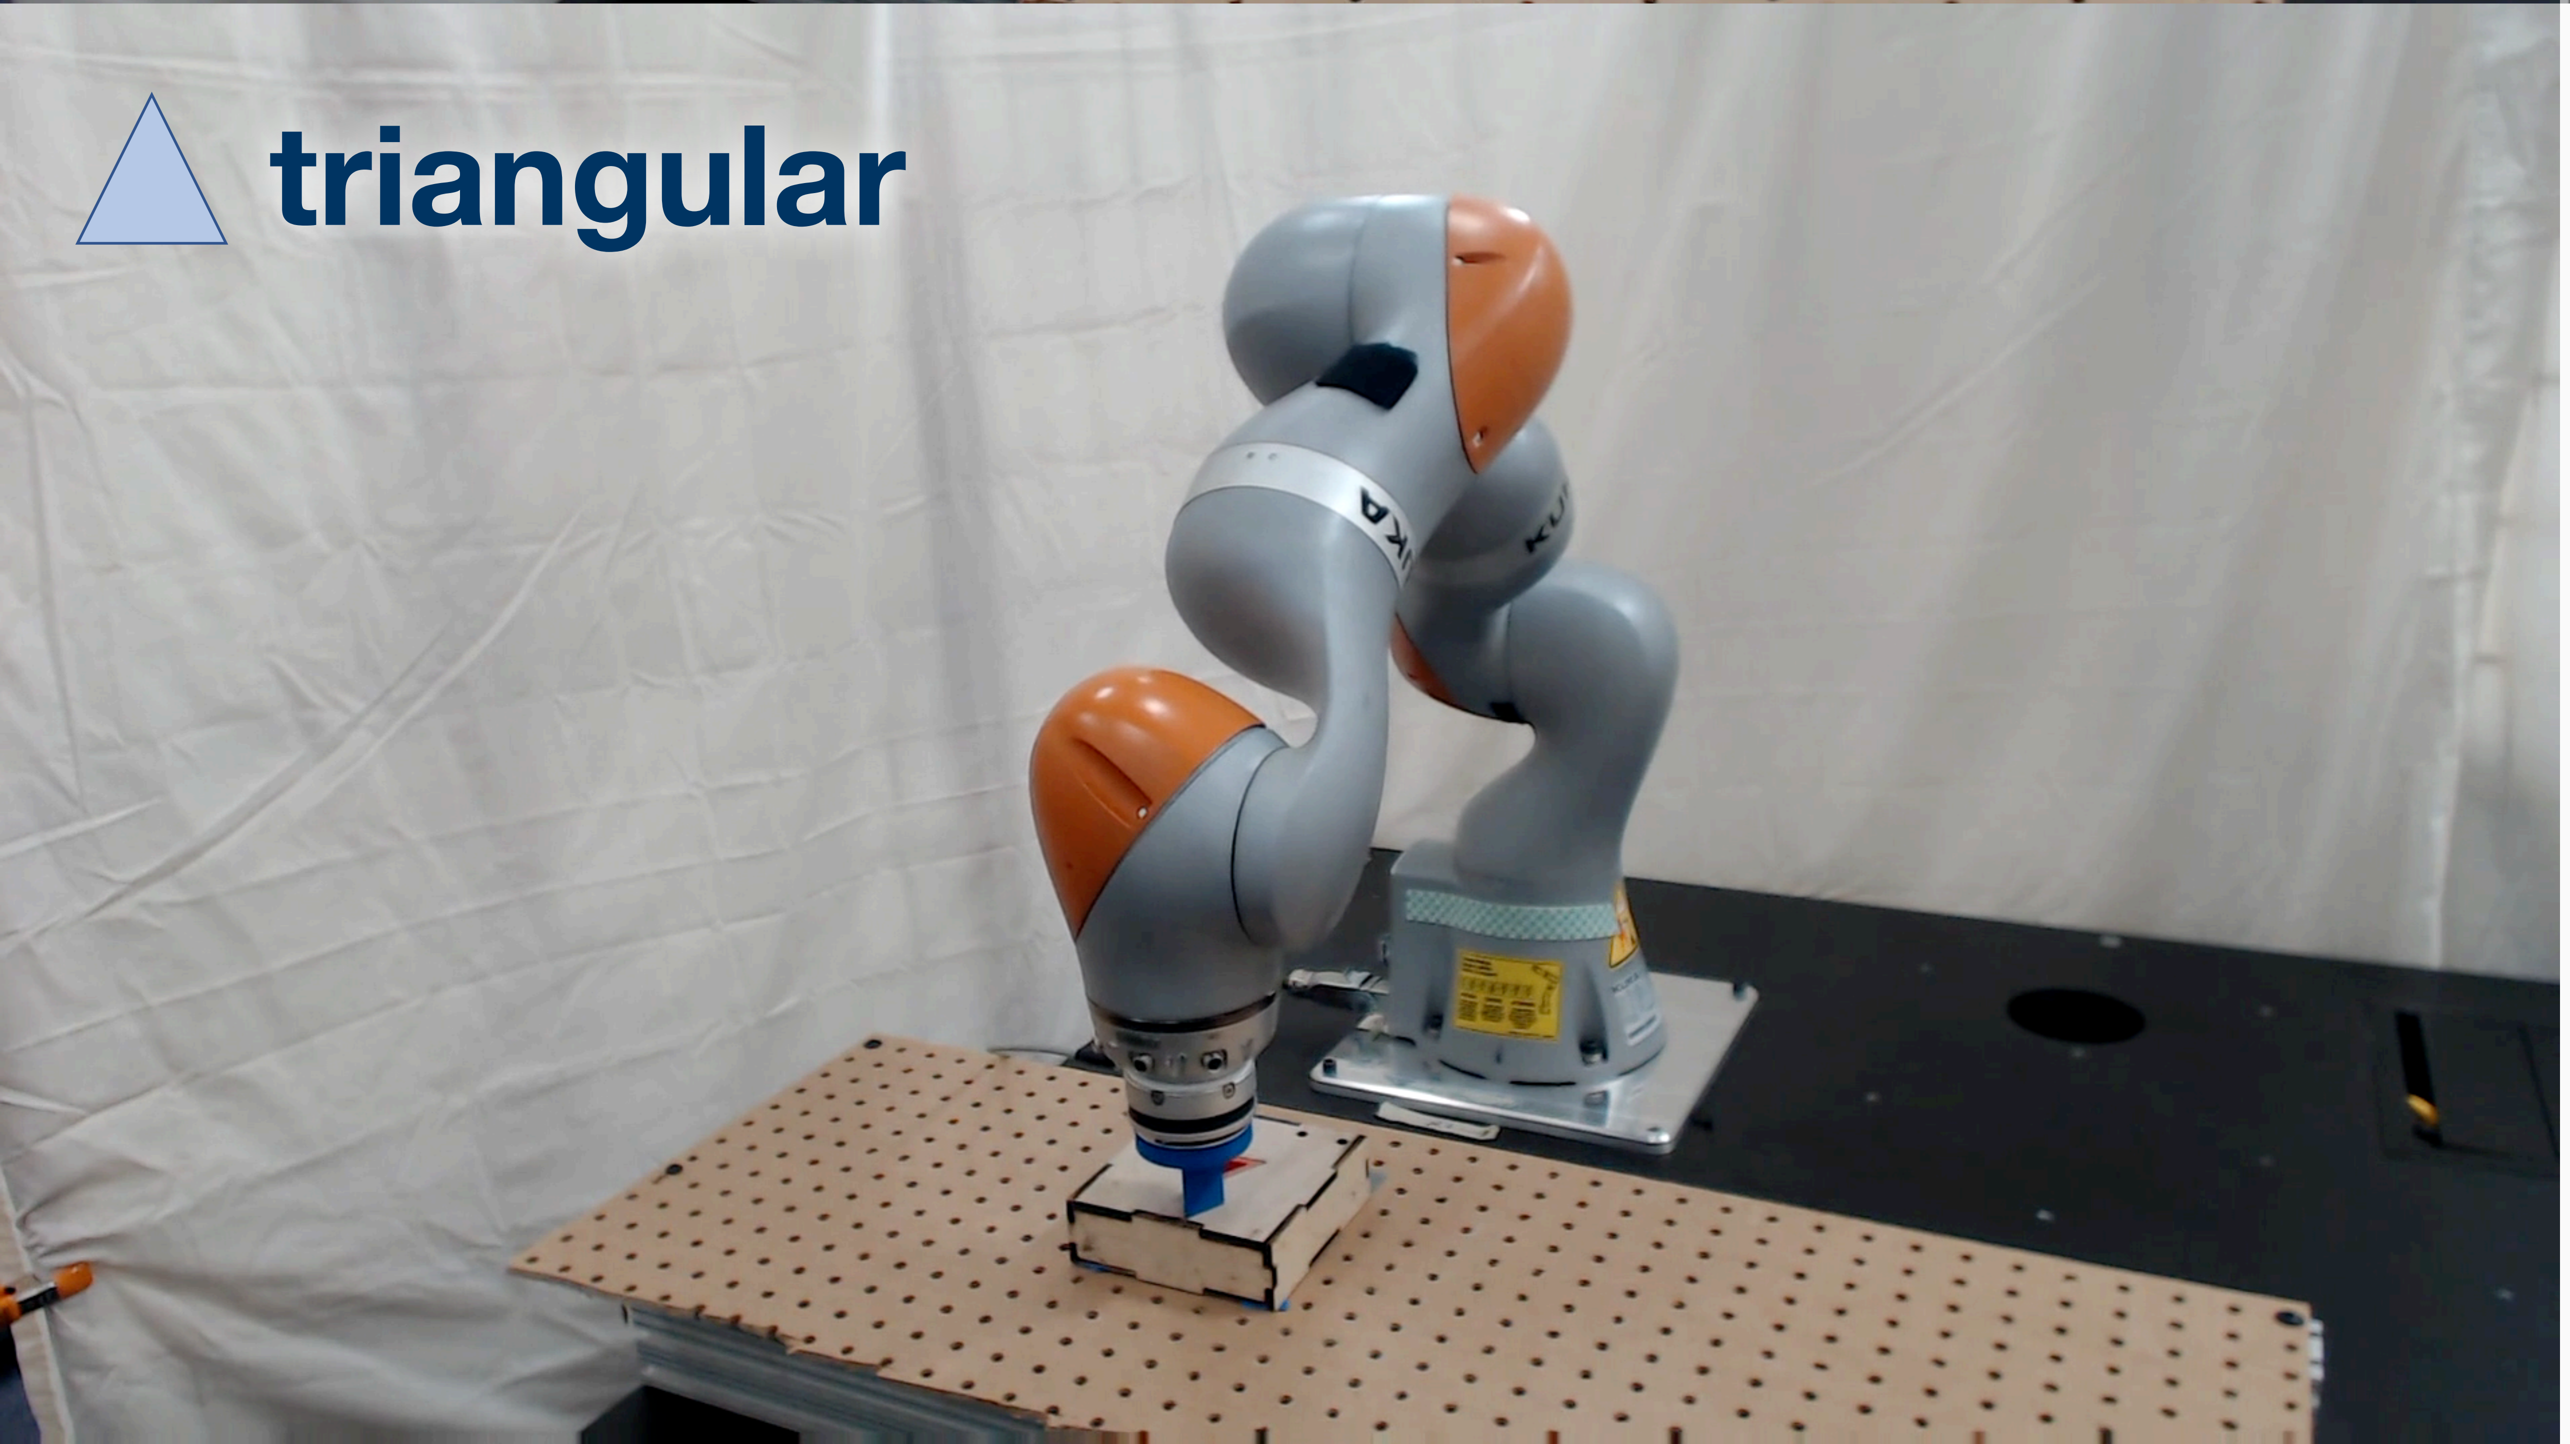

Each policy is trained with TRPO for 300 episodes (~5 hours) while fixing the multimodal representations.

We demonstrate that the learned representation and policy generalize well to new pegs.

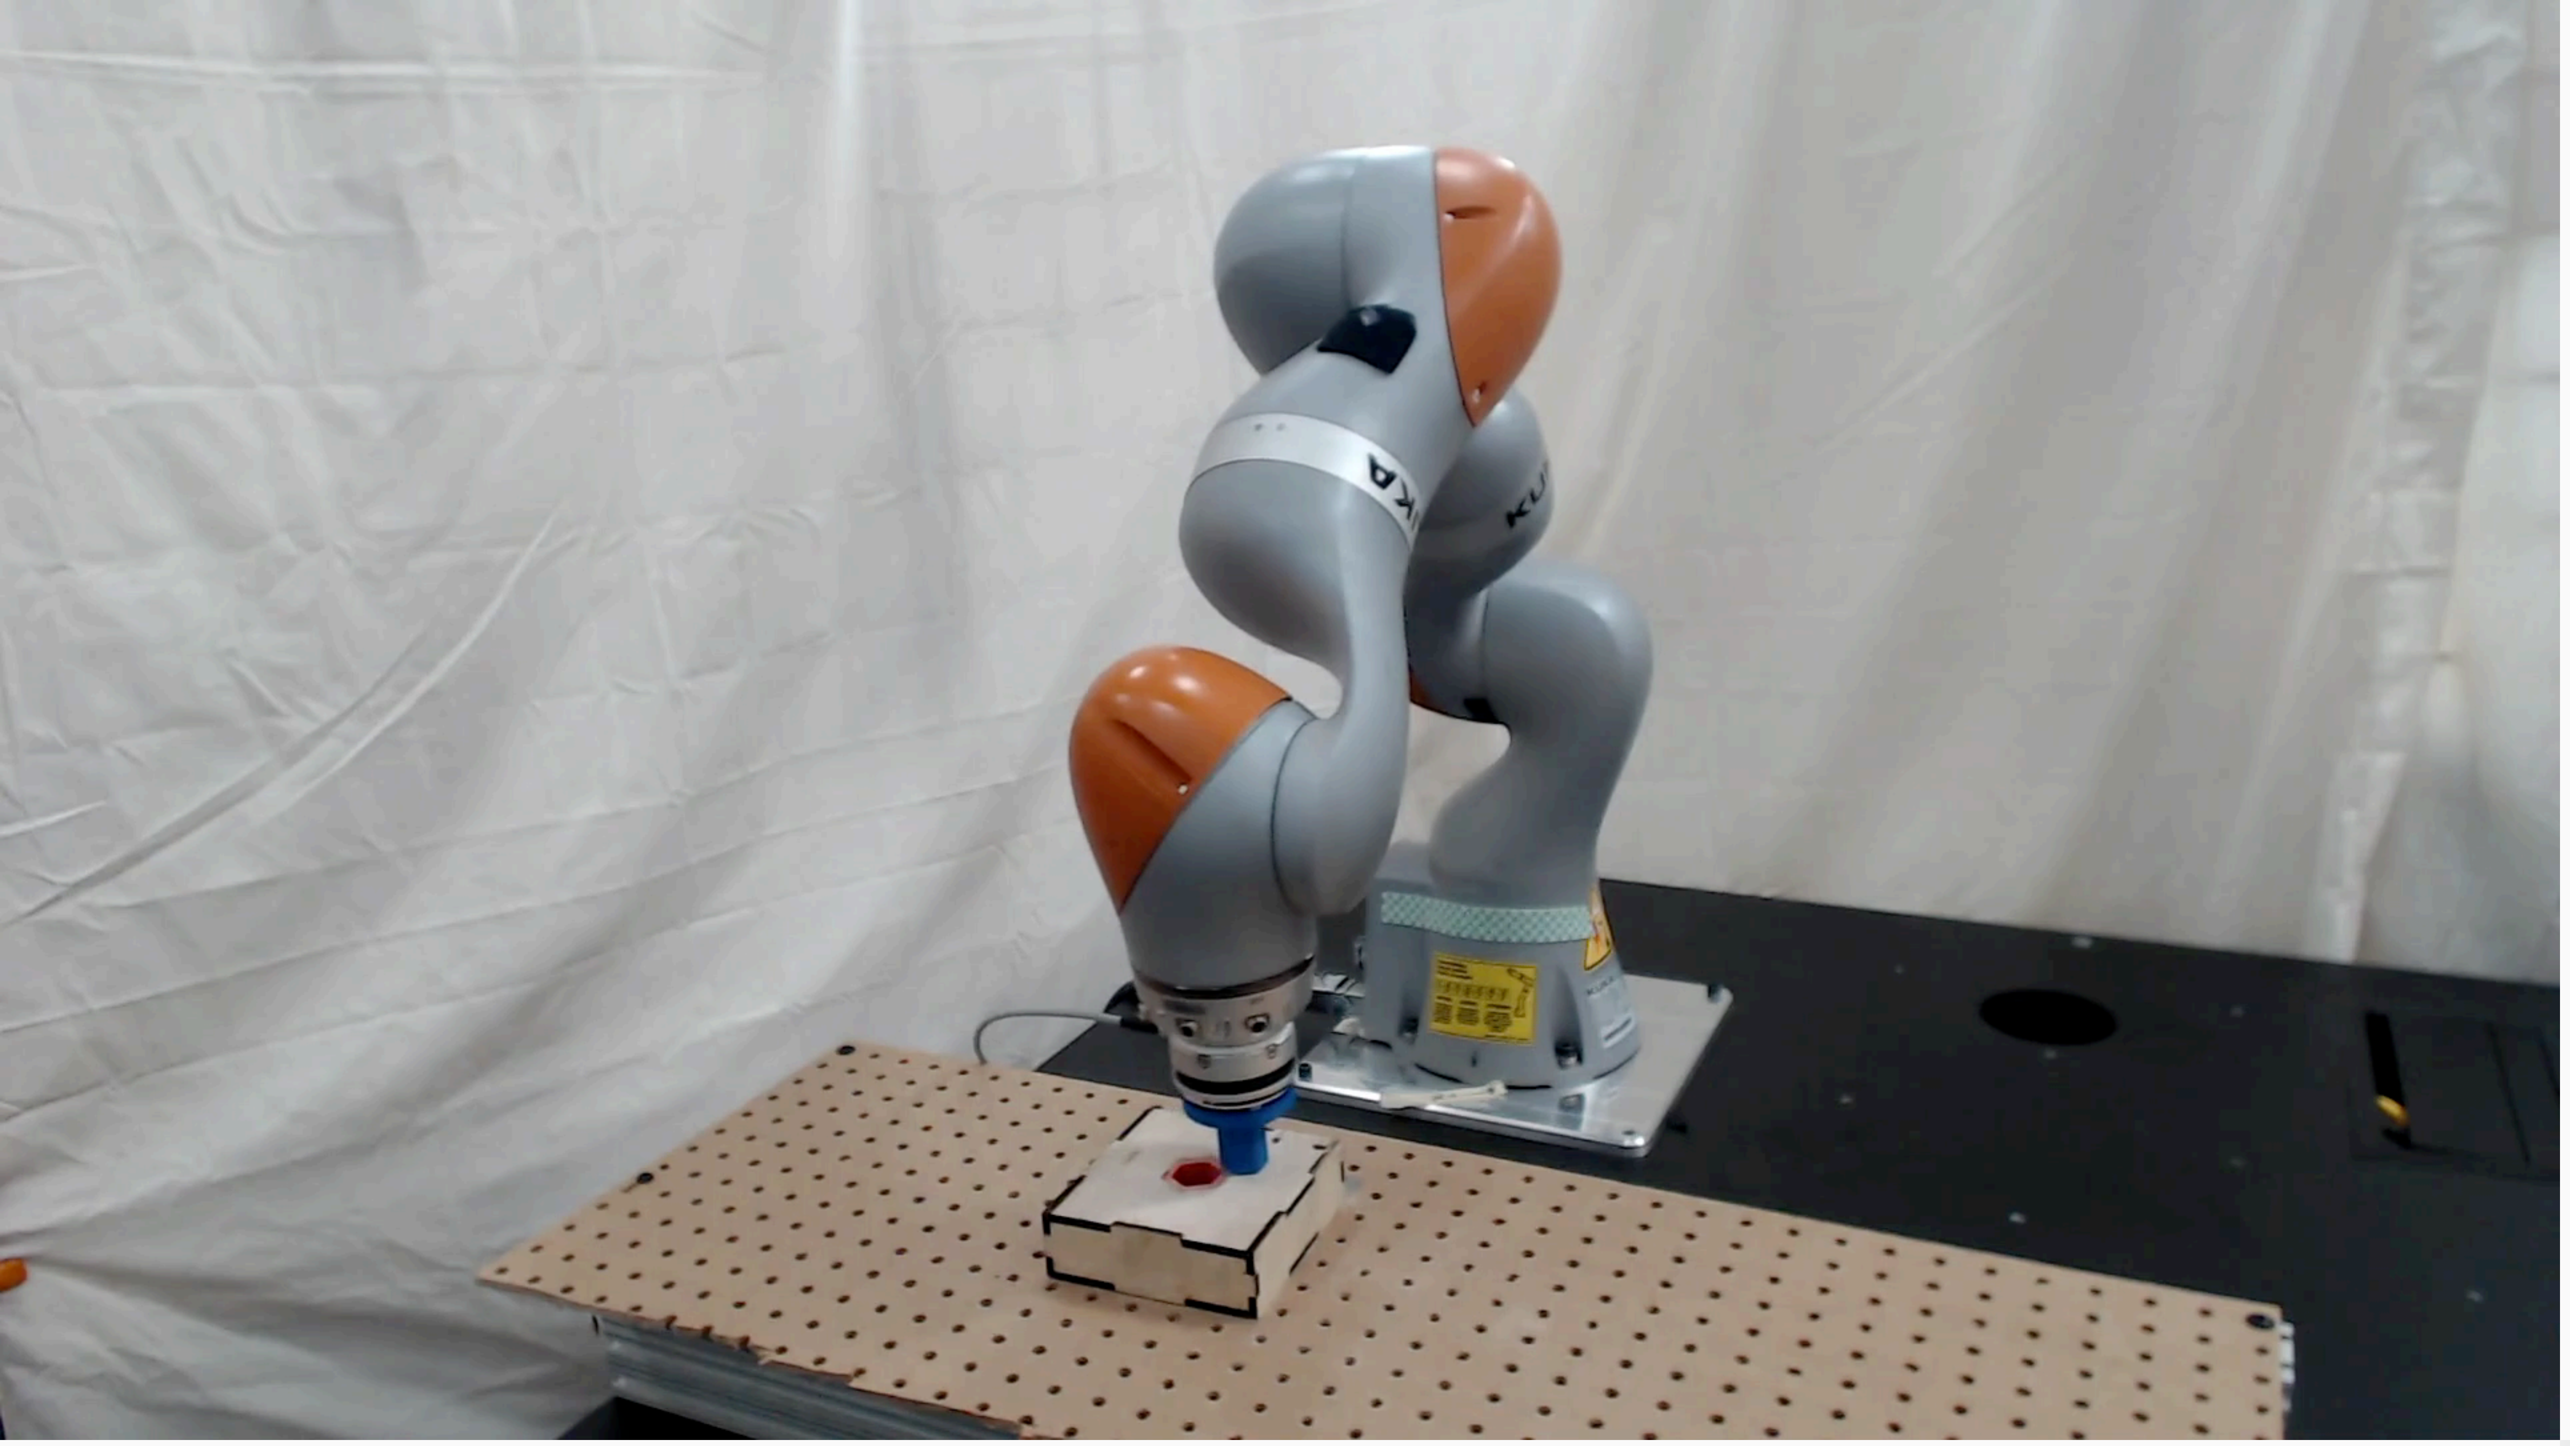

x1

Policy Transfer (62% success rate)

Representation trained on

Policy trained on

Policy evaluated on

Representation Transfer (92% success rate)

Representation trained on

Policy trained on

Policy evaluated on

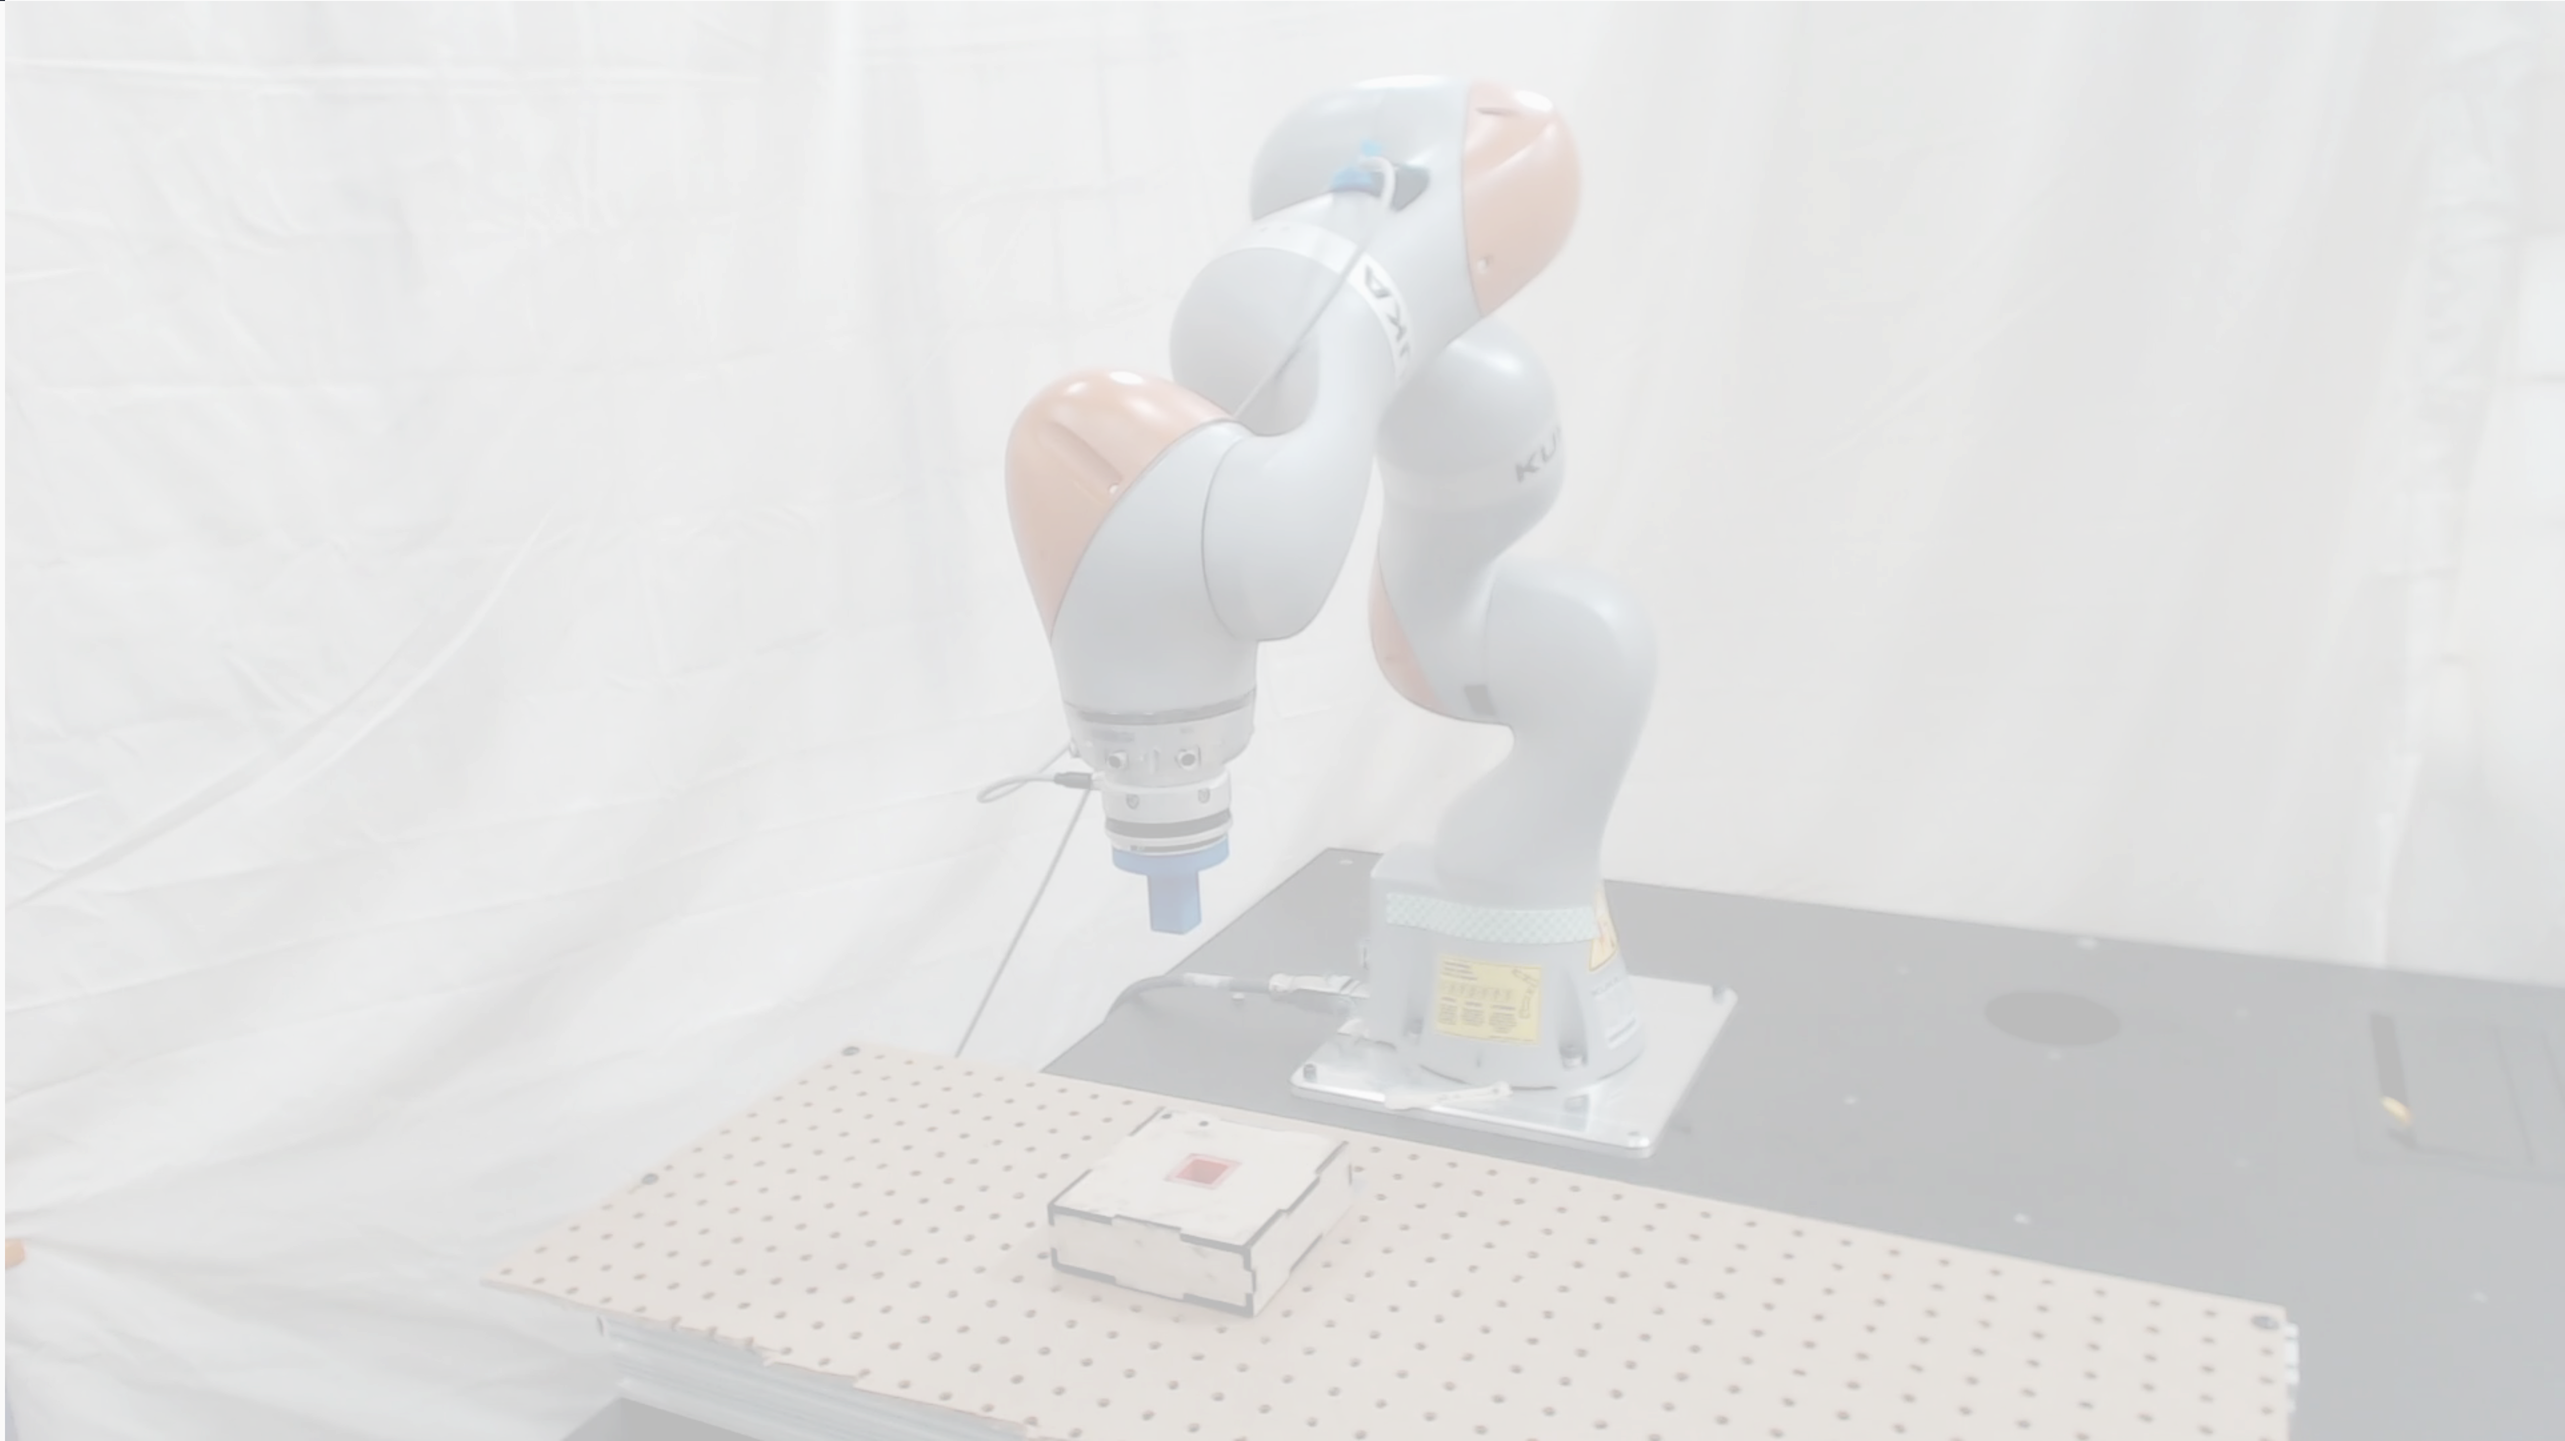

We demonstrate that our policy's robustness against four types of disturbances.

1. External Force
2. Haptics Perturbation
3. Target Movement
4. Camera Occlusion


External Force
The policy is able to recover from external pushes on the arm.
x1
Normalized Force
1.0
0.5
0
-0.5
-1.0
Force
Time
x
y
z

Haptics Perturbation

The policy is robust towards perturbations to the F/T sensors.

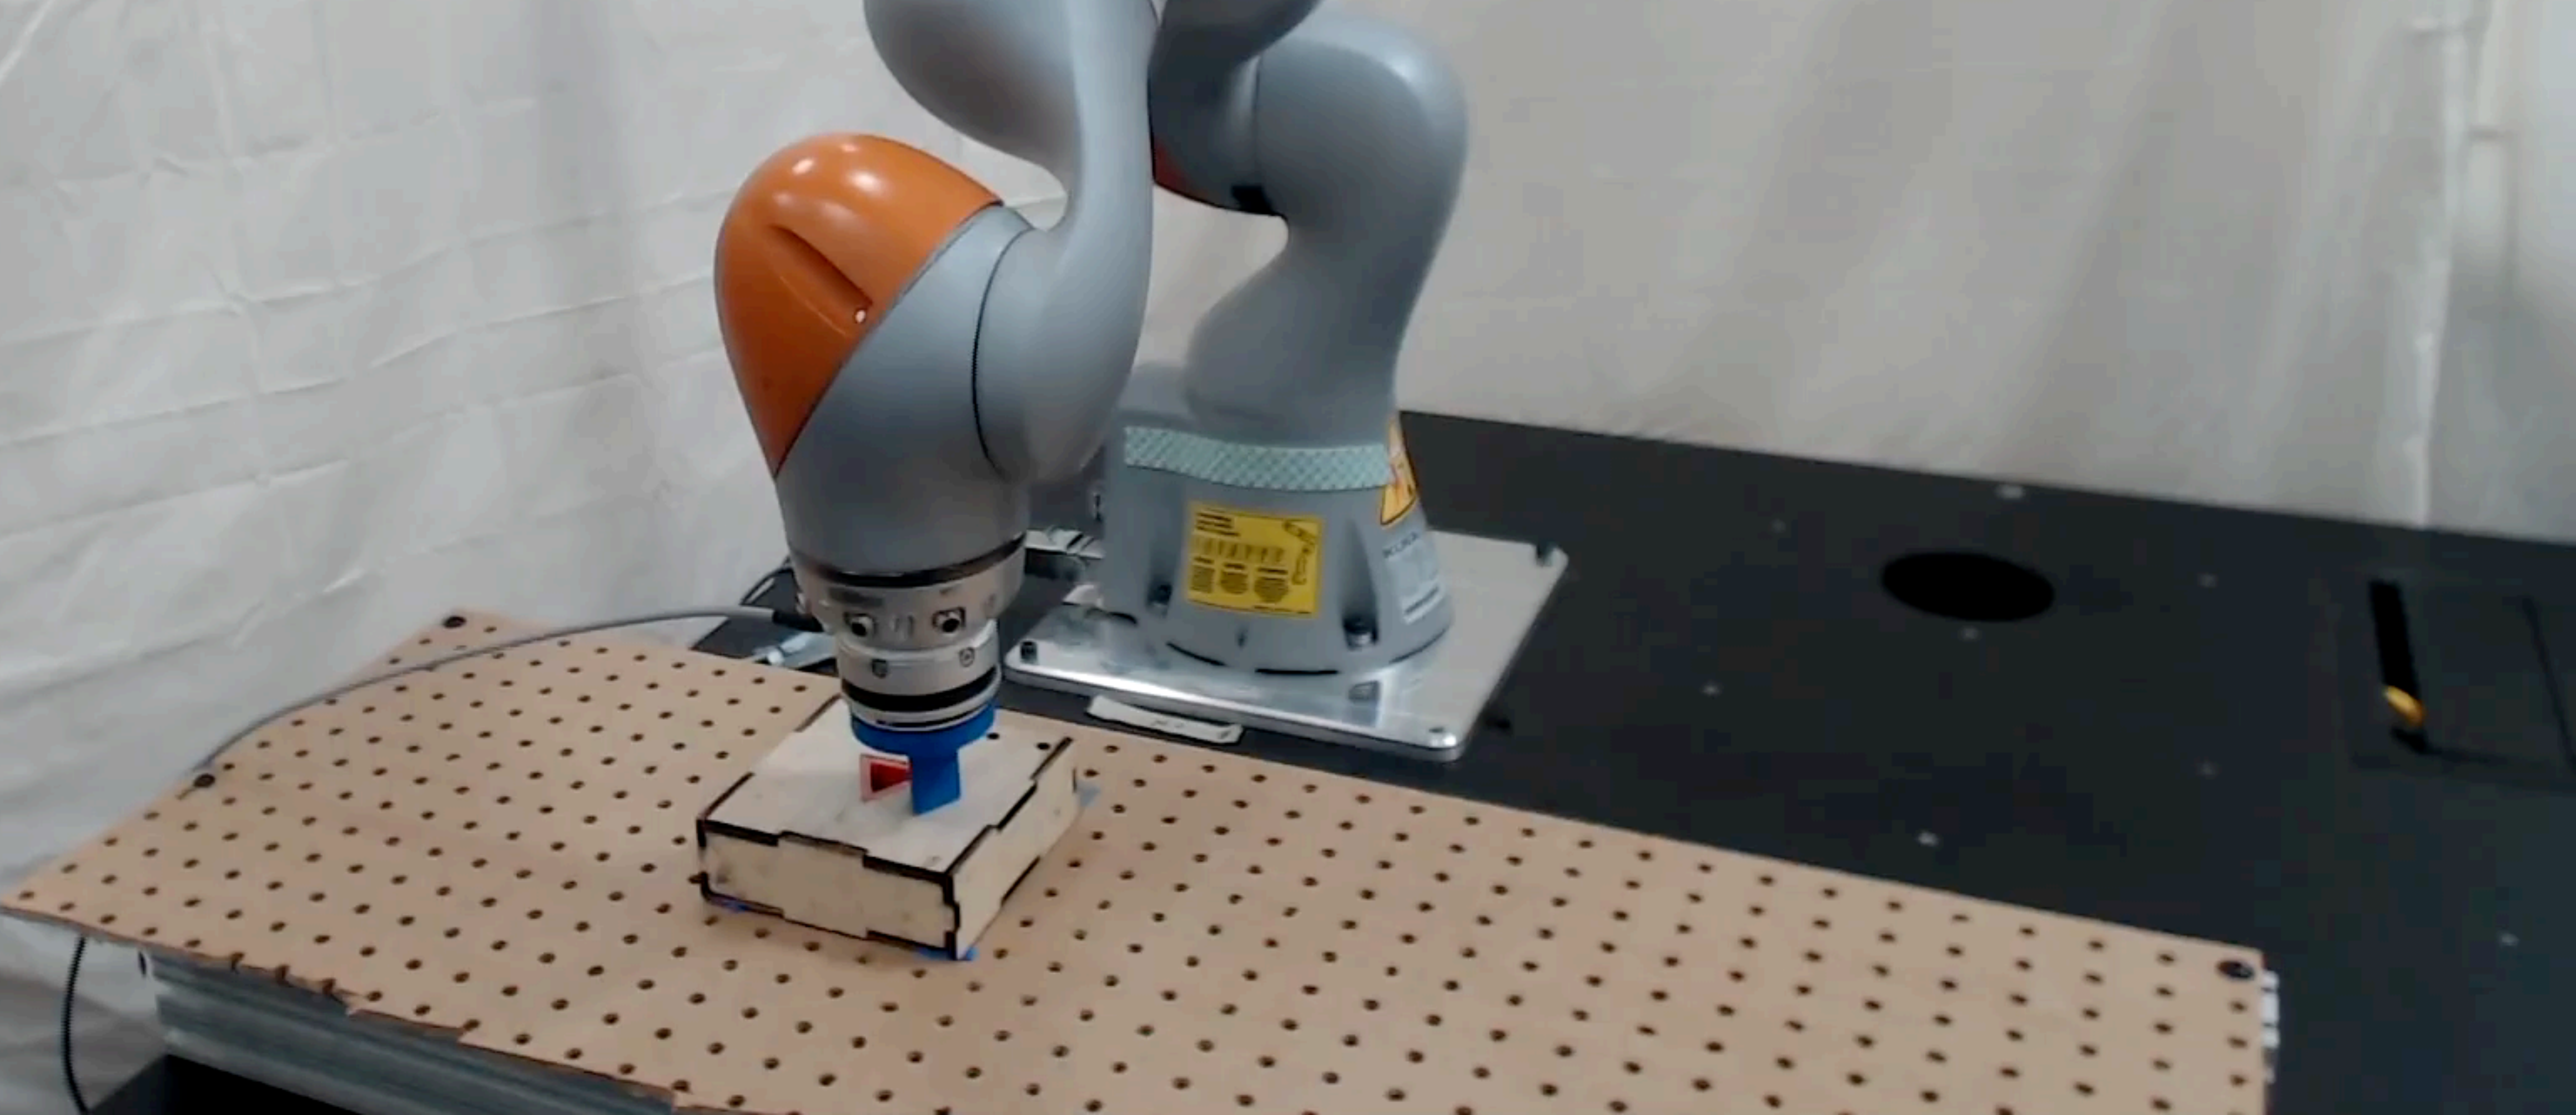

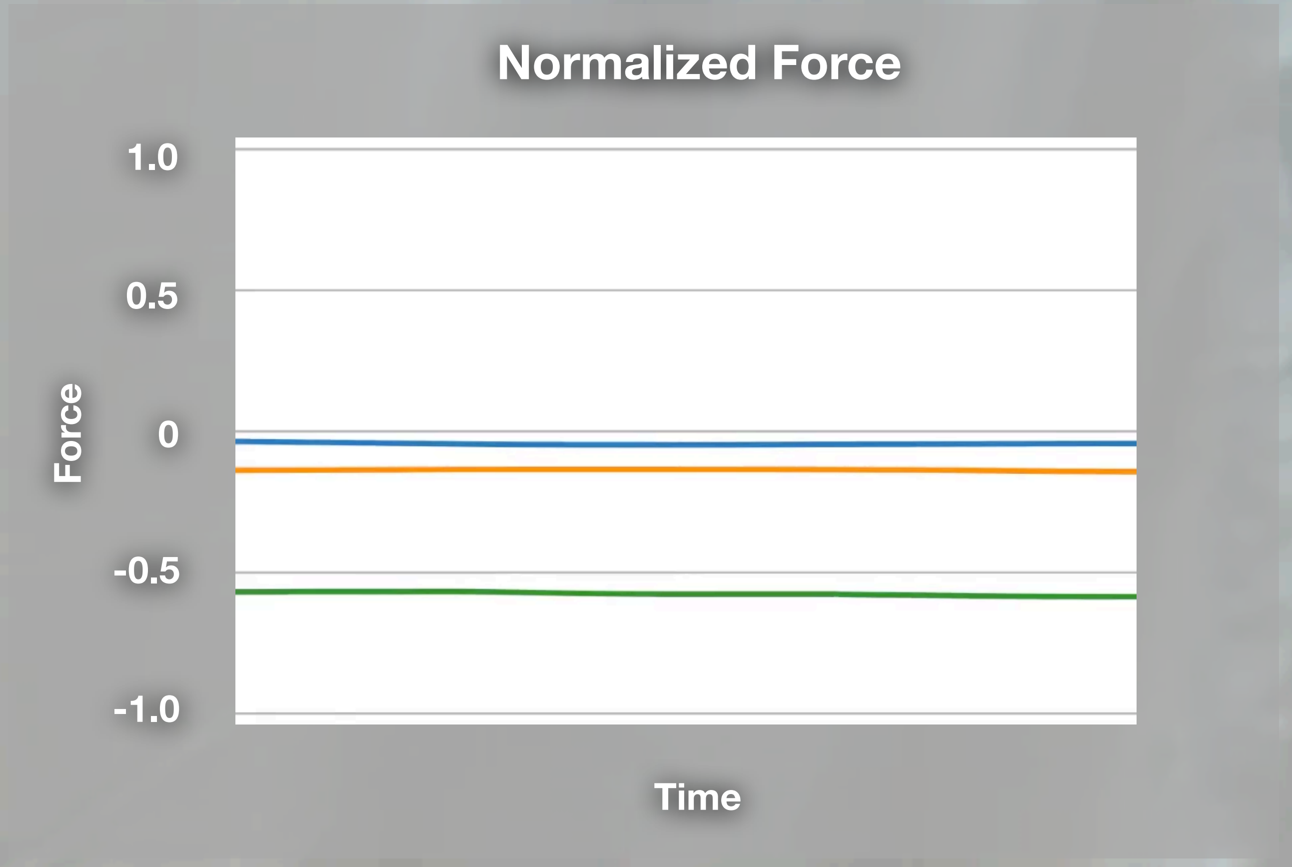

x1

Target Movement

The policy is able to handle small offsets of the position of the hole.

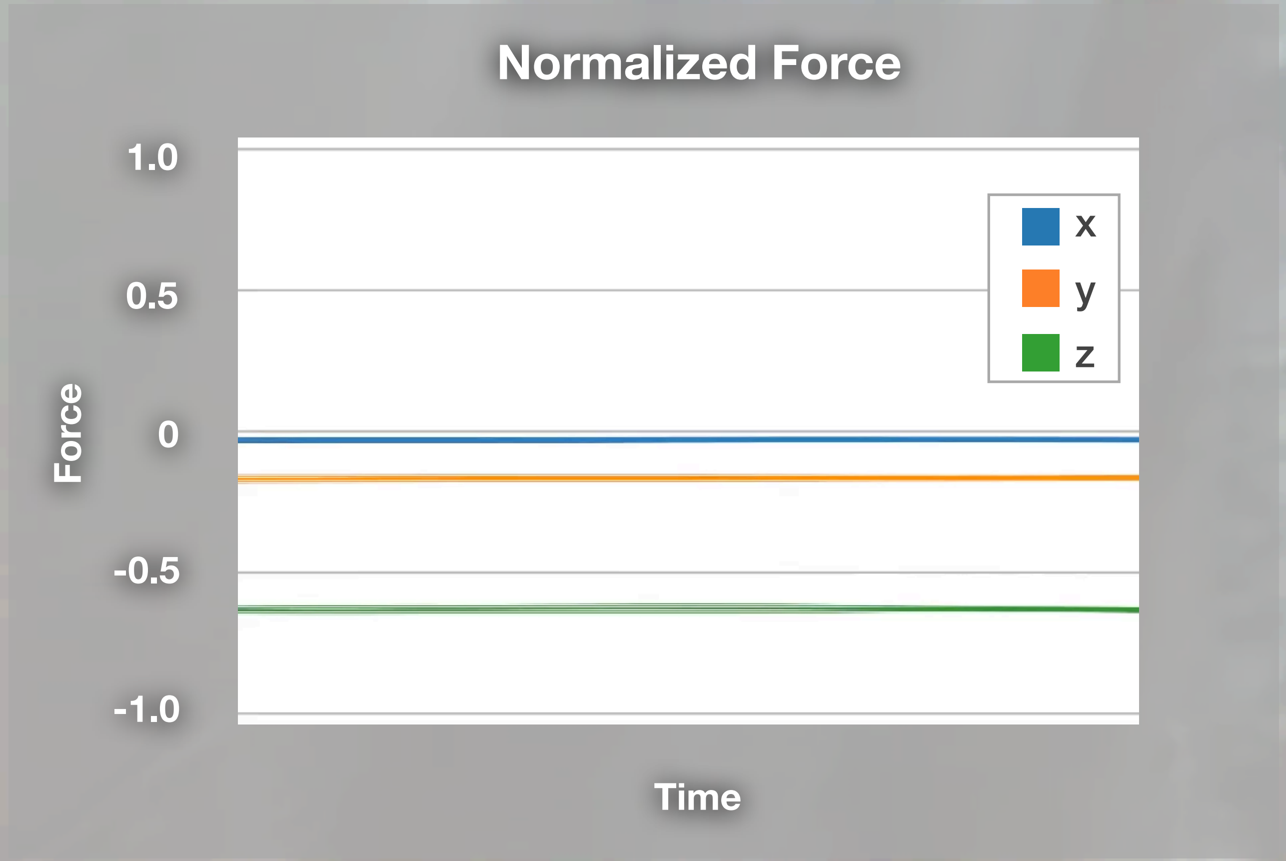

x1 * no object in environment is tracked


Camera Occlusion
The policy can complete insertion even with intermittent camera occlusion.
Agent View
x1
KUKA

Predictive Model

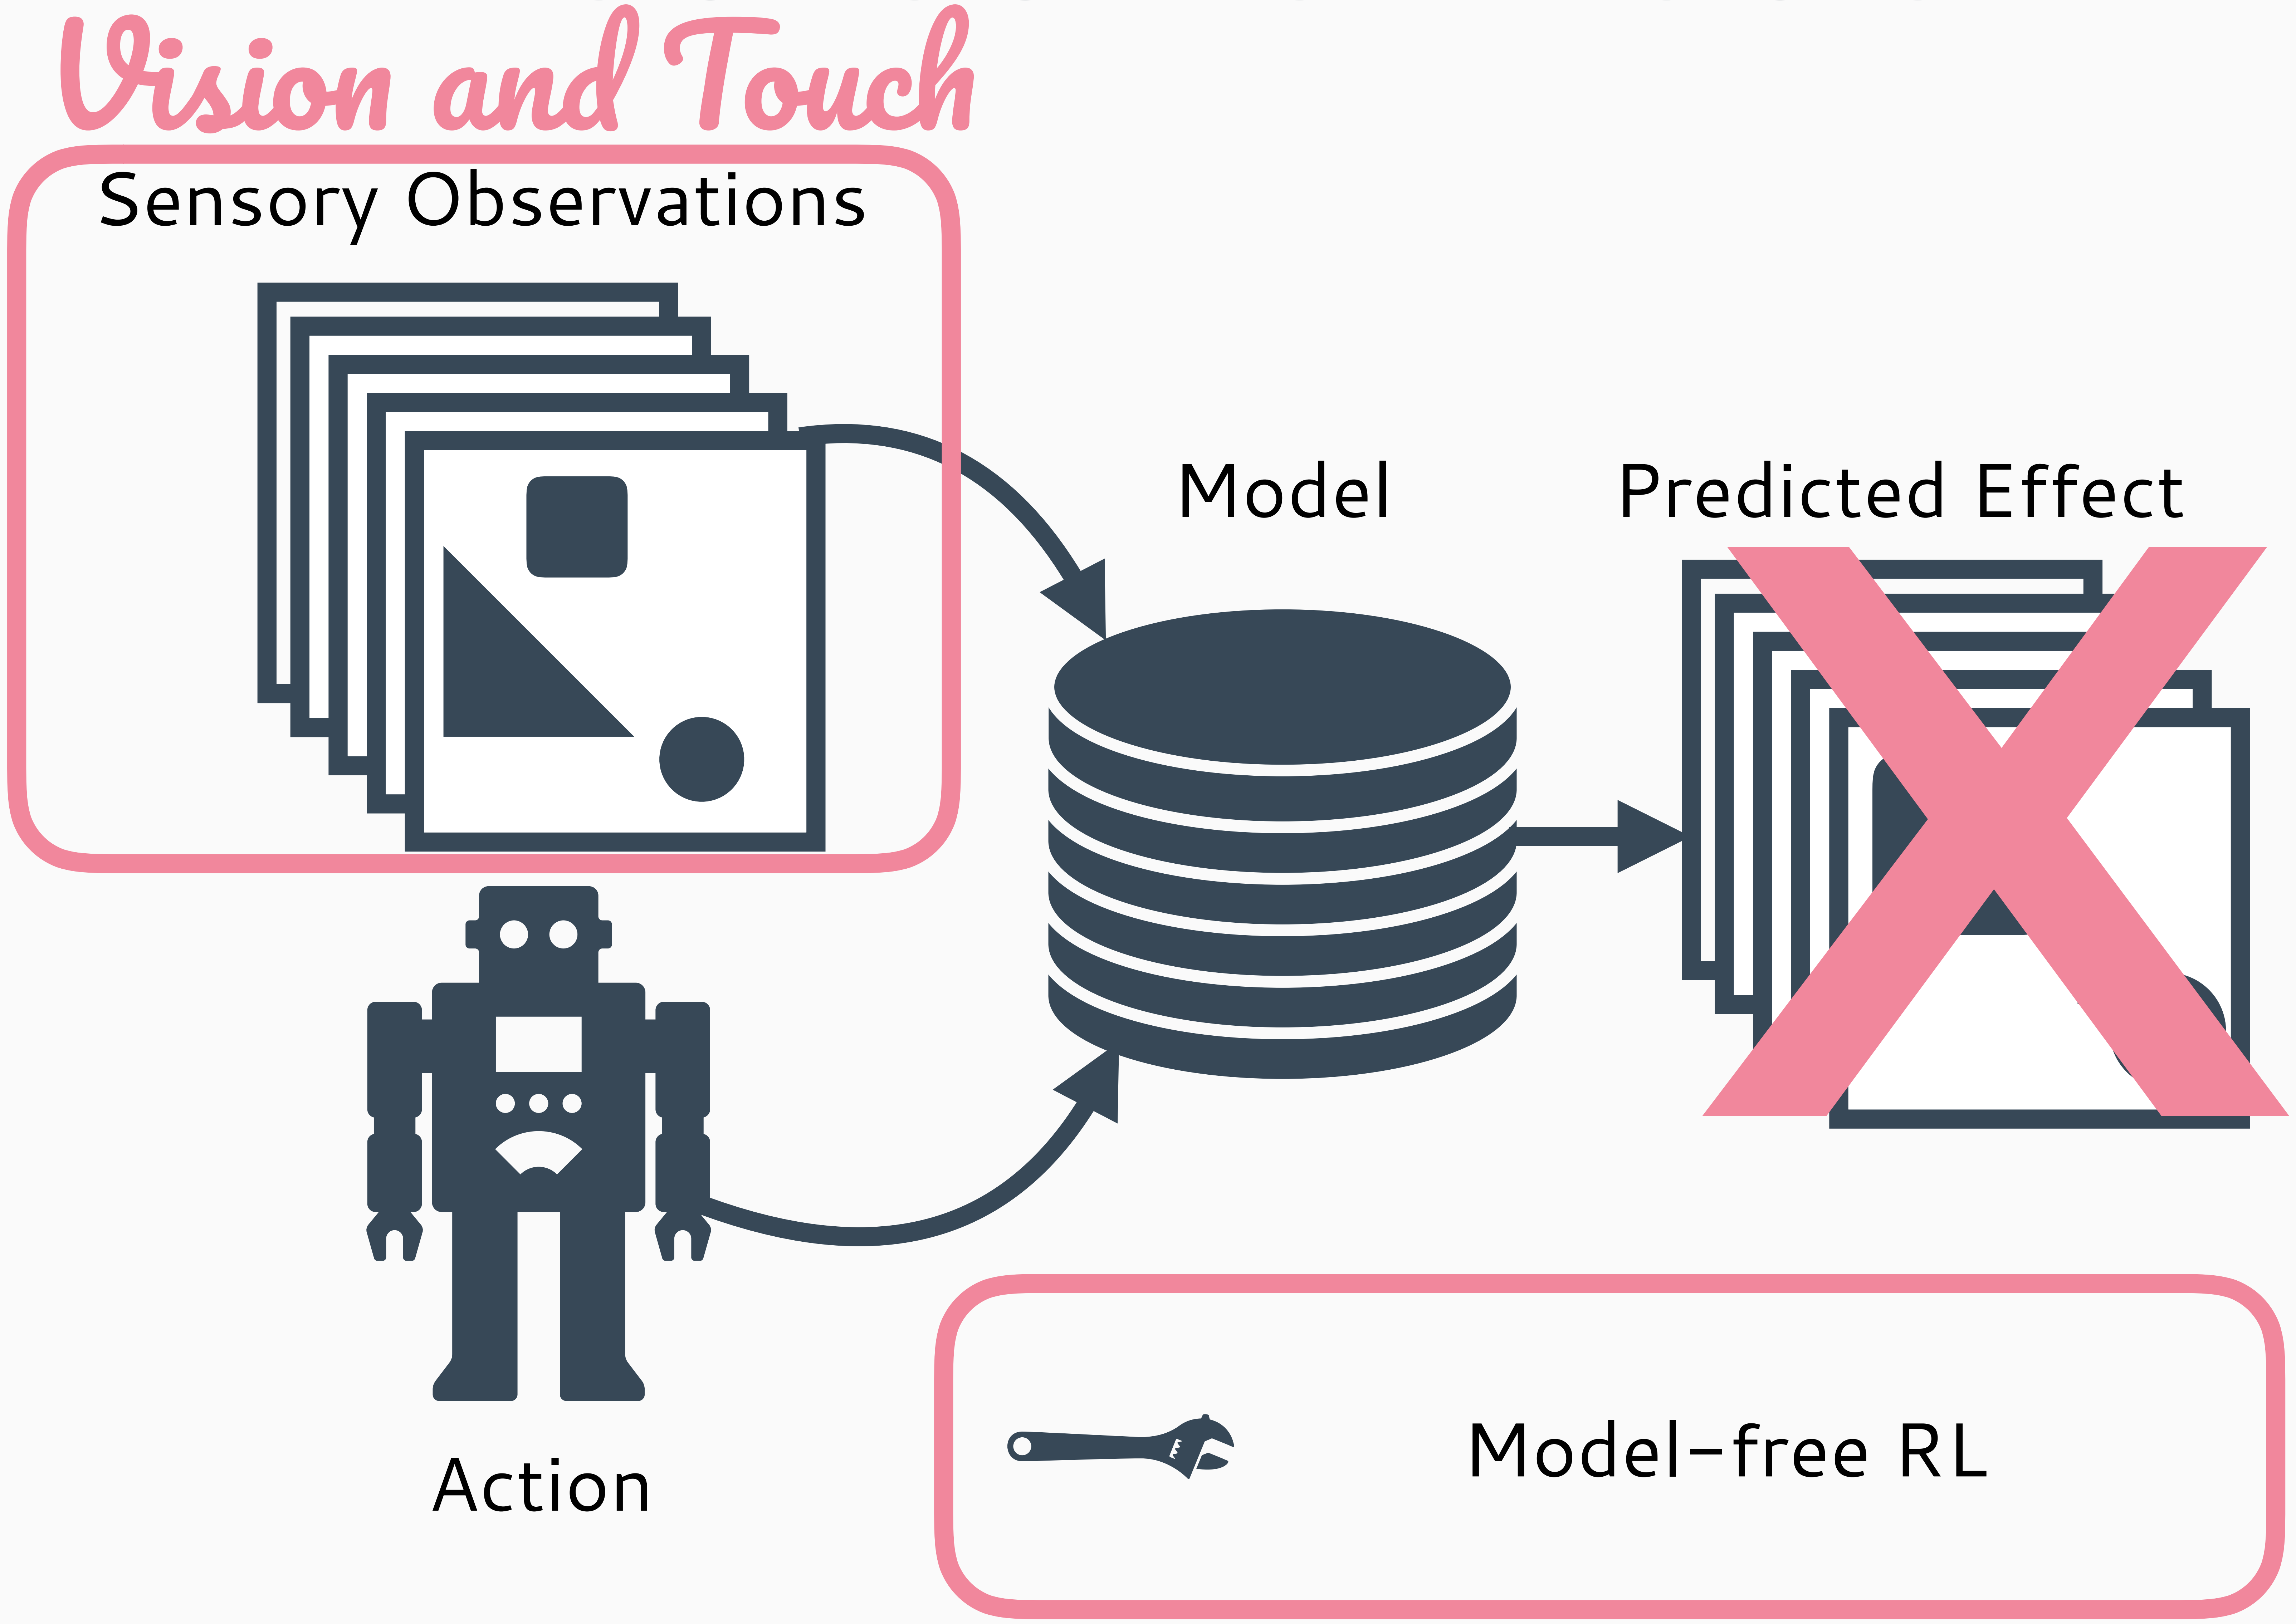

Conclusions

$$S \times A \times t$$

Sensory Data Actions Time

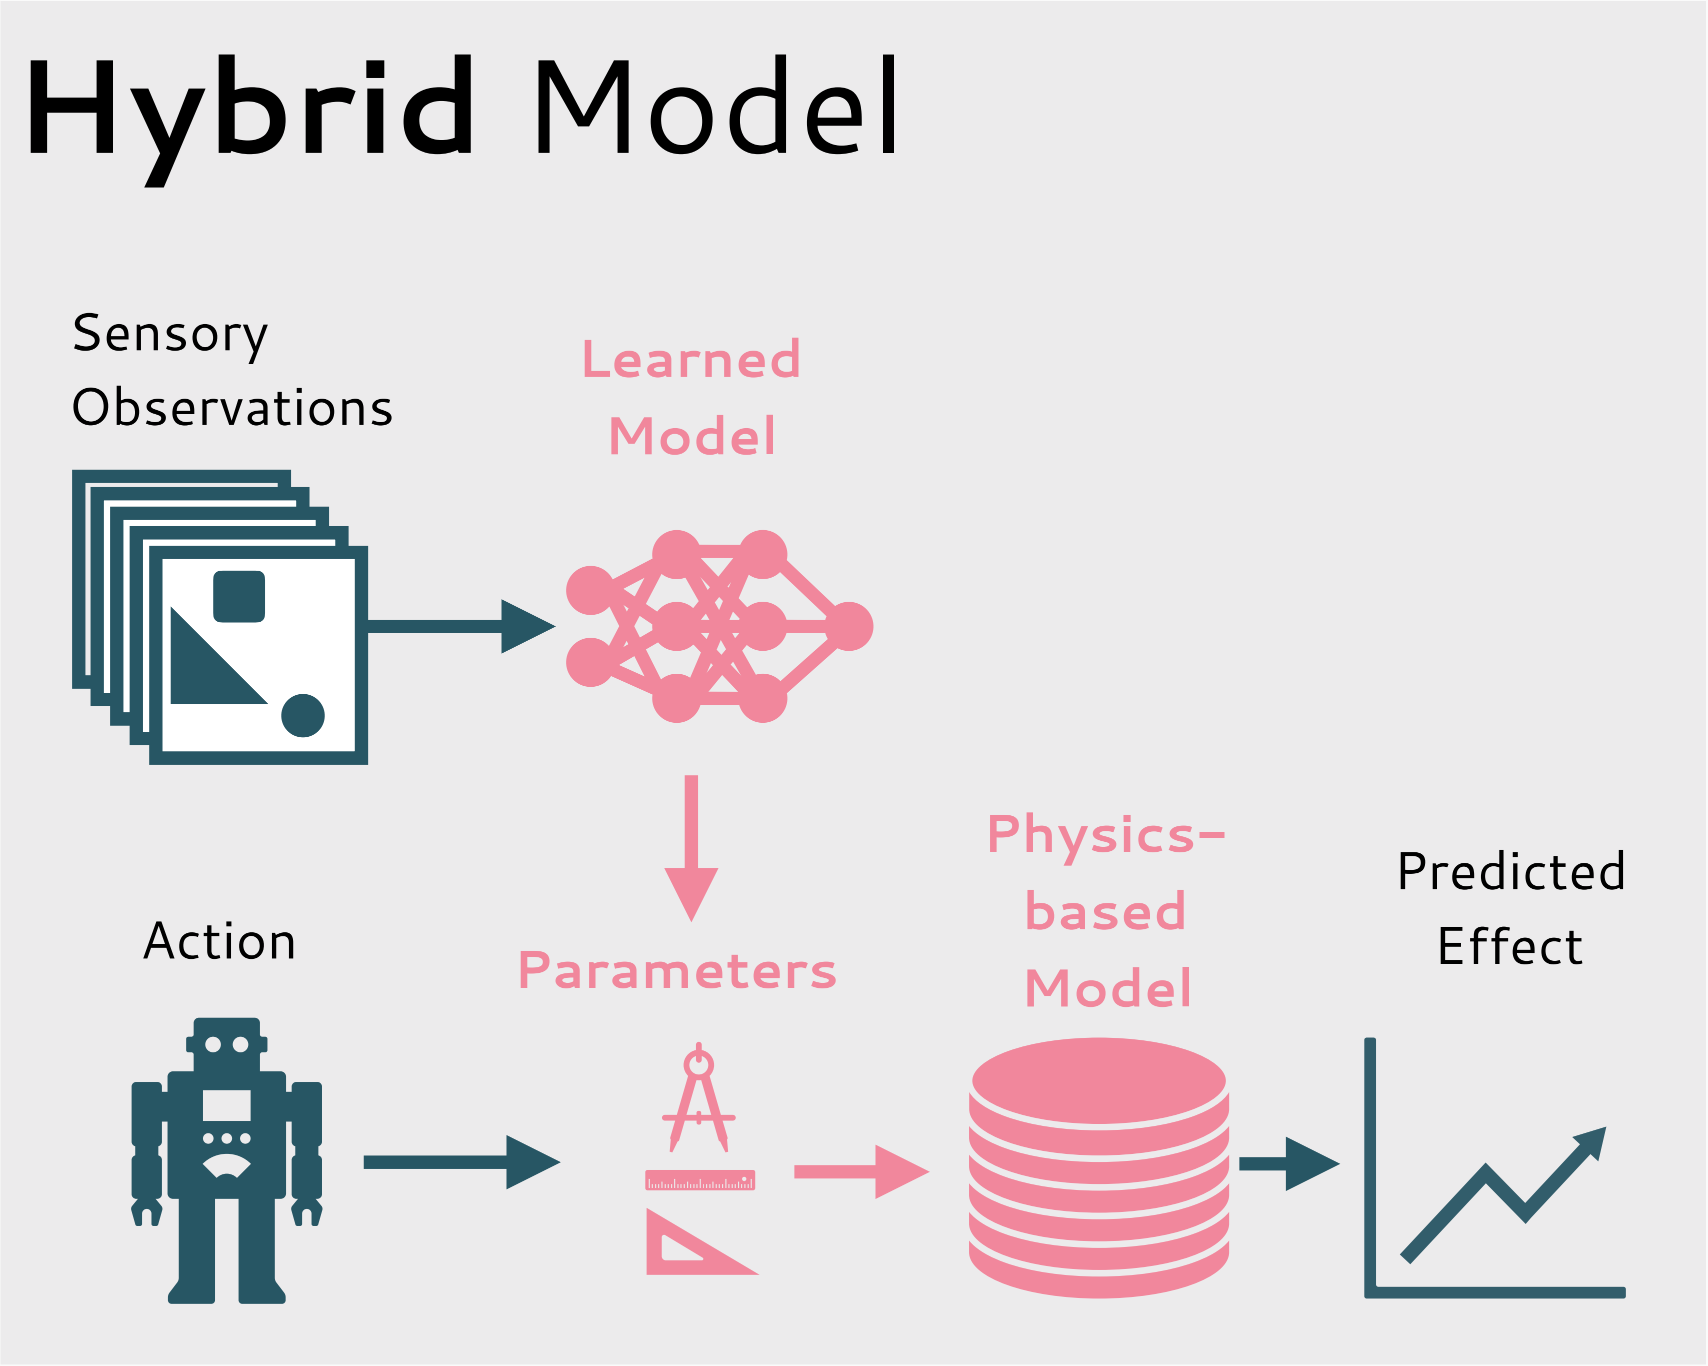

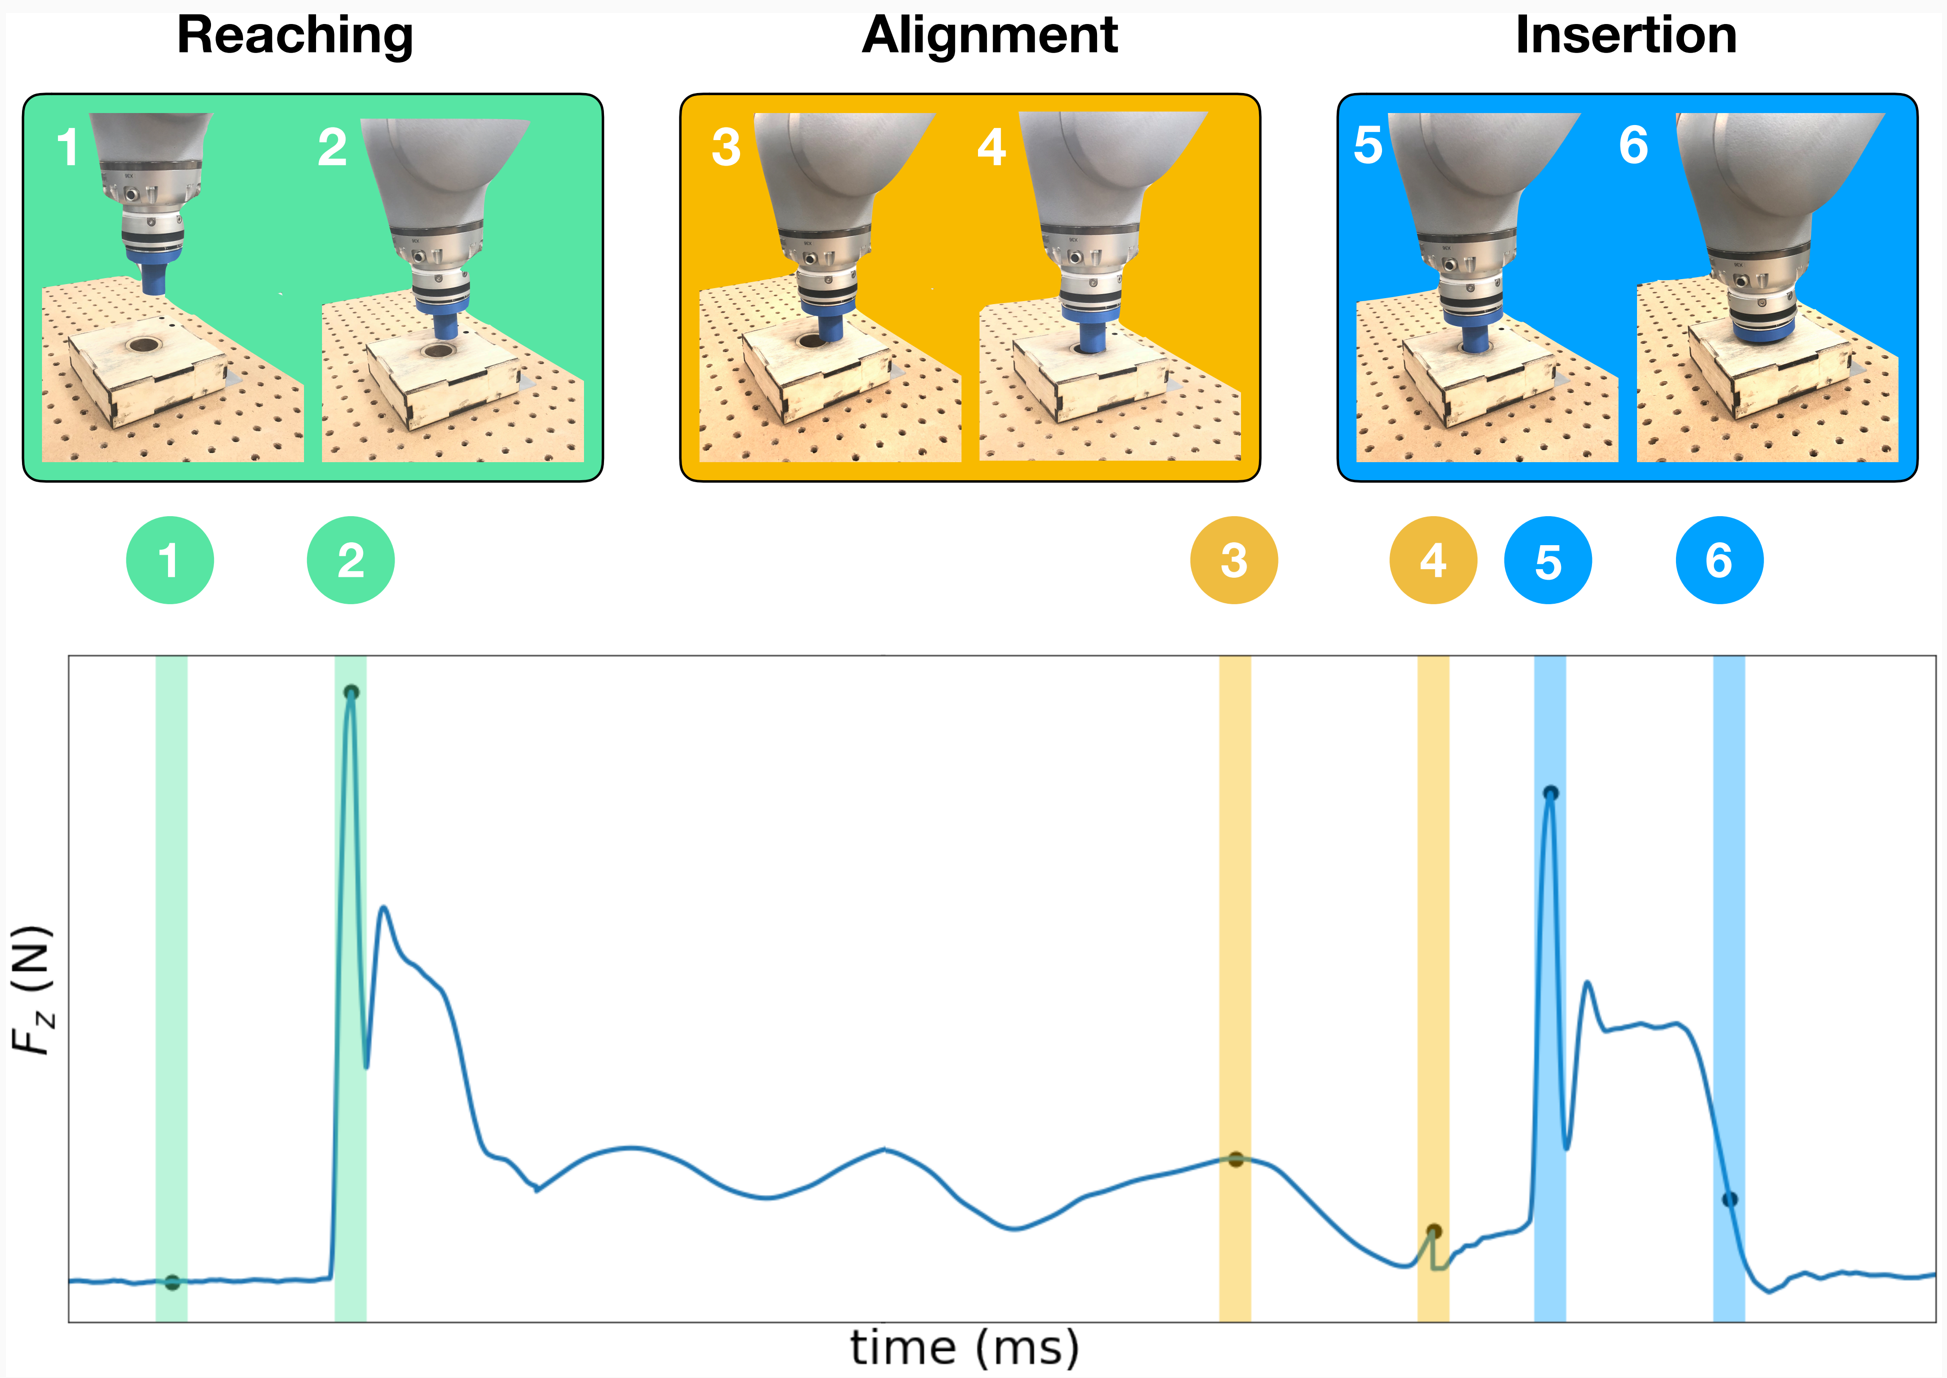

Bias–Variance Tradeoff

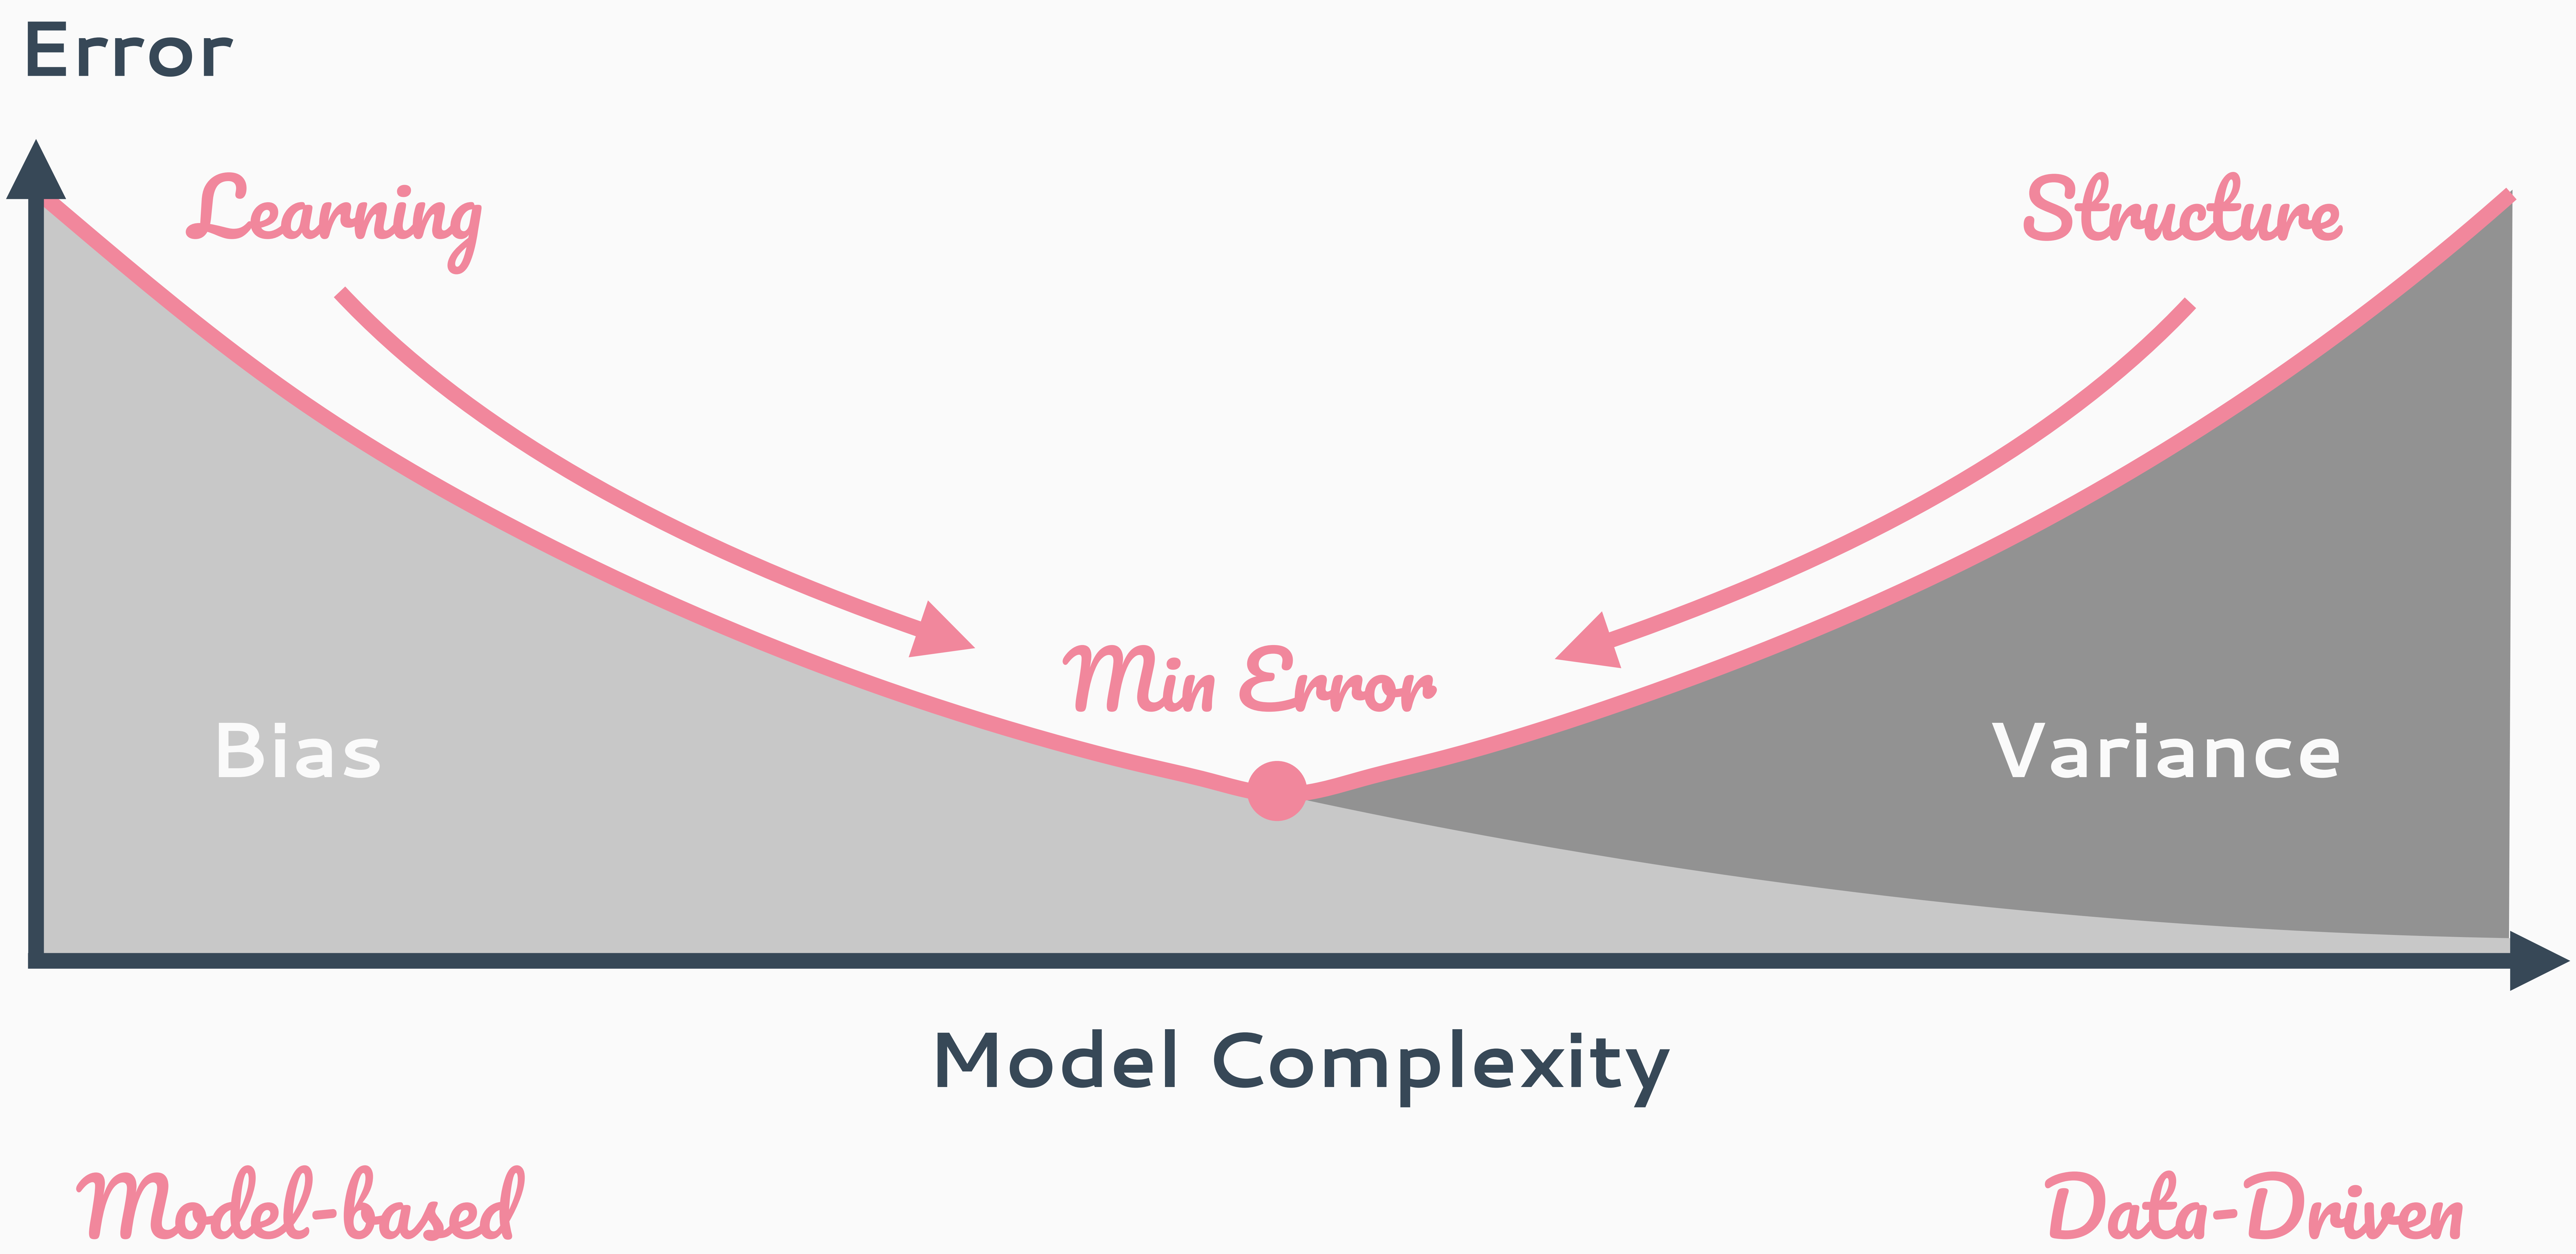

Thank you for your Attention!

MAX-PLANCK-GESELLSCHAFT

TOYOTA
RESEARCH INSTITUTE

JD.COM

Stanford
ARTIFICIAL
INTELLIGENCE

Detailed Model Architecture

Learned Model

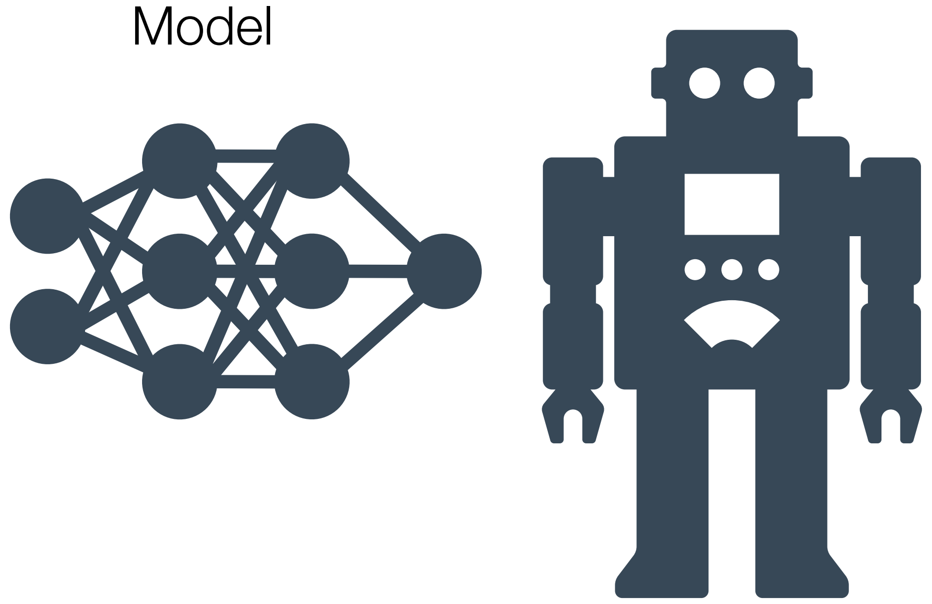

Physics-based Model

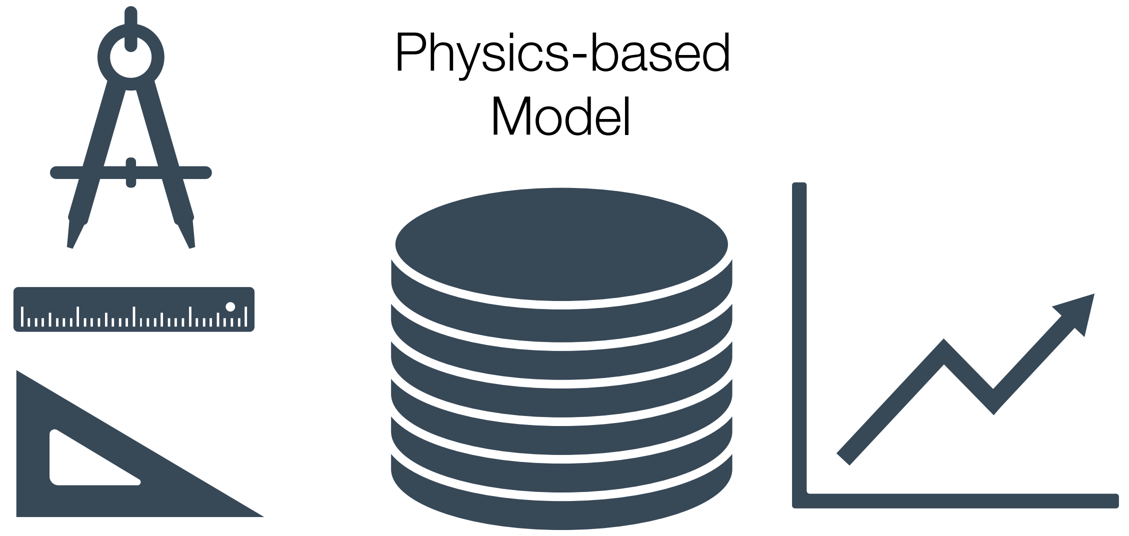

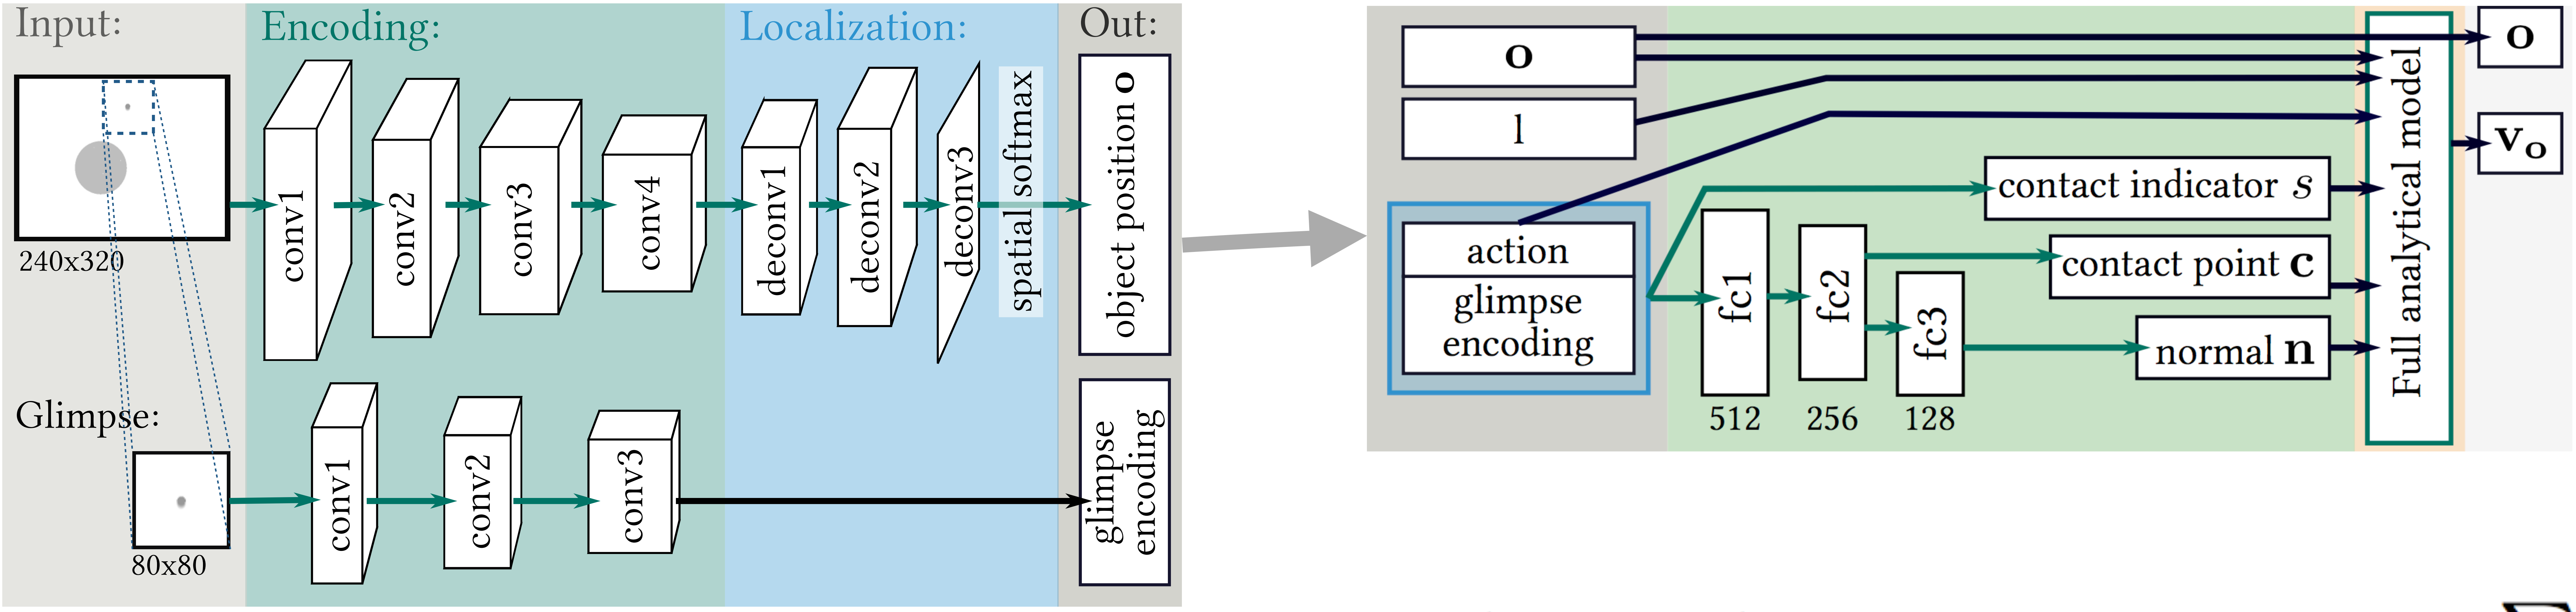

$$L(\hat{\mathbf{v}}_\mathbf{o}, \hat{\mathbf{o}}, \mathbf{v}_\mathbf{o}, \mathbf{o}) = trans + mag + rot + pos + \lambda \sum_w \parallel w \parallel$$

$$trans = \|\hat{\nu}_\mathbf{o} - \nu_\mathbf{o}\| \qquad mag = |\|\hat{\nu}_\mathbf{o} - \nu_\mathbf{o}\||$$

$$rot = \frac{180}{\pi}|\omega - \hat{\omega}| \qquad pos = \parallel \mathbf{o} - \hat{\mathbf{o}} \parallel$$

Testing Data Efficiency

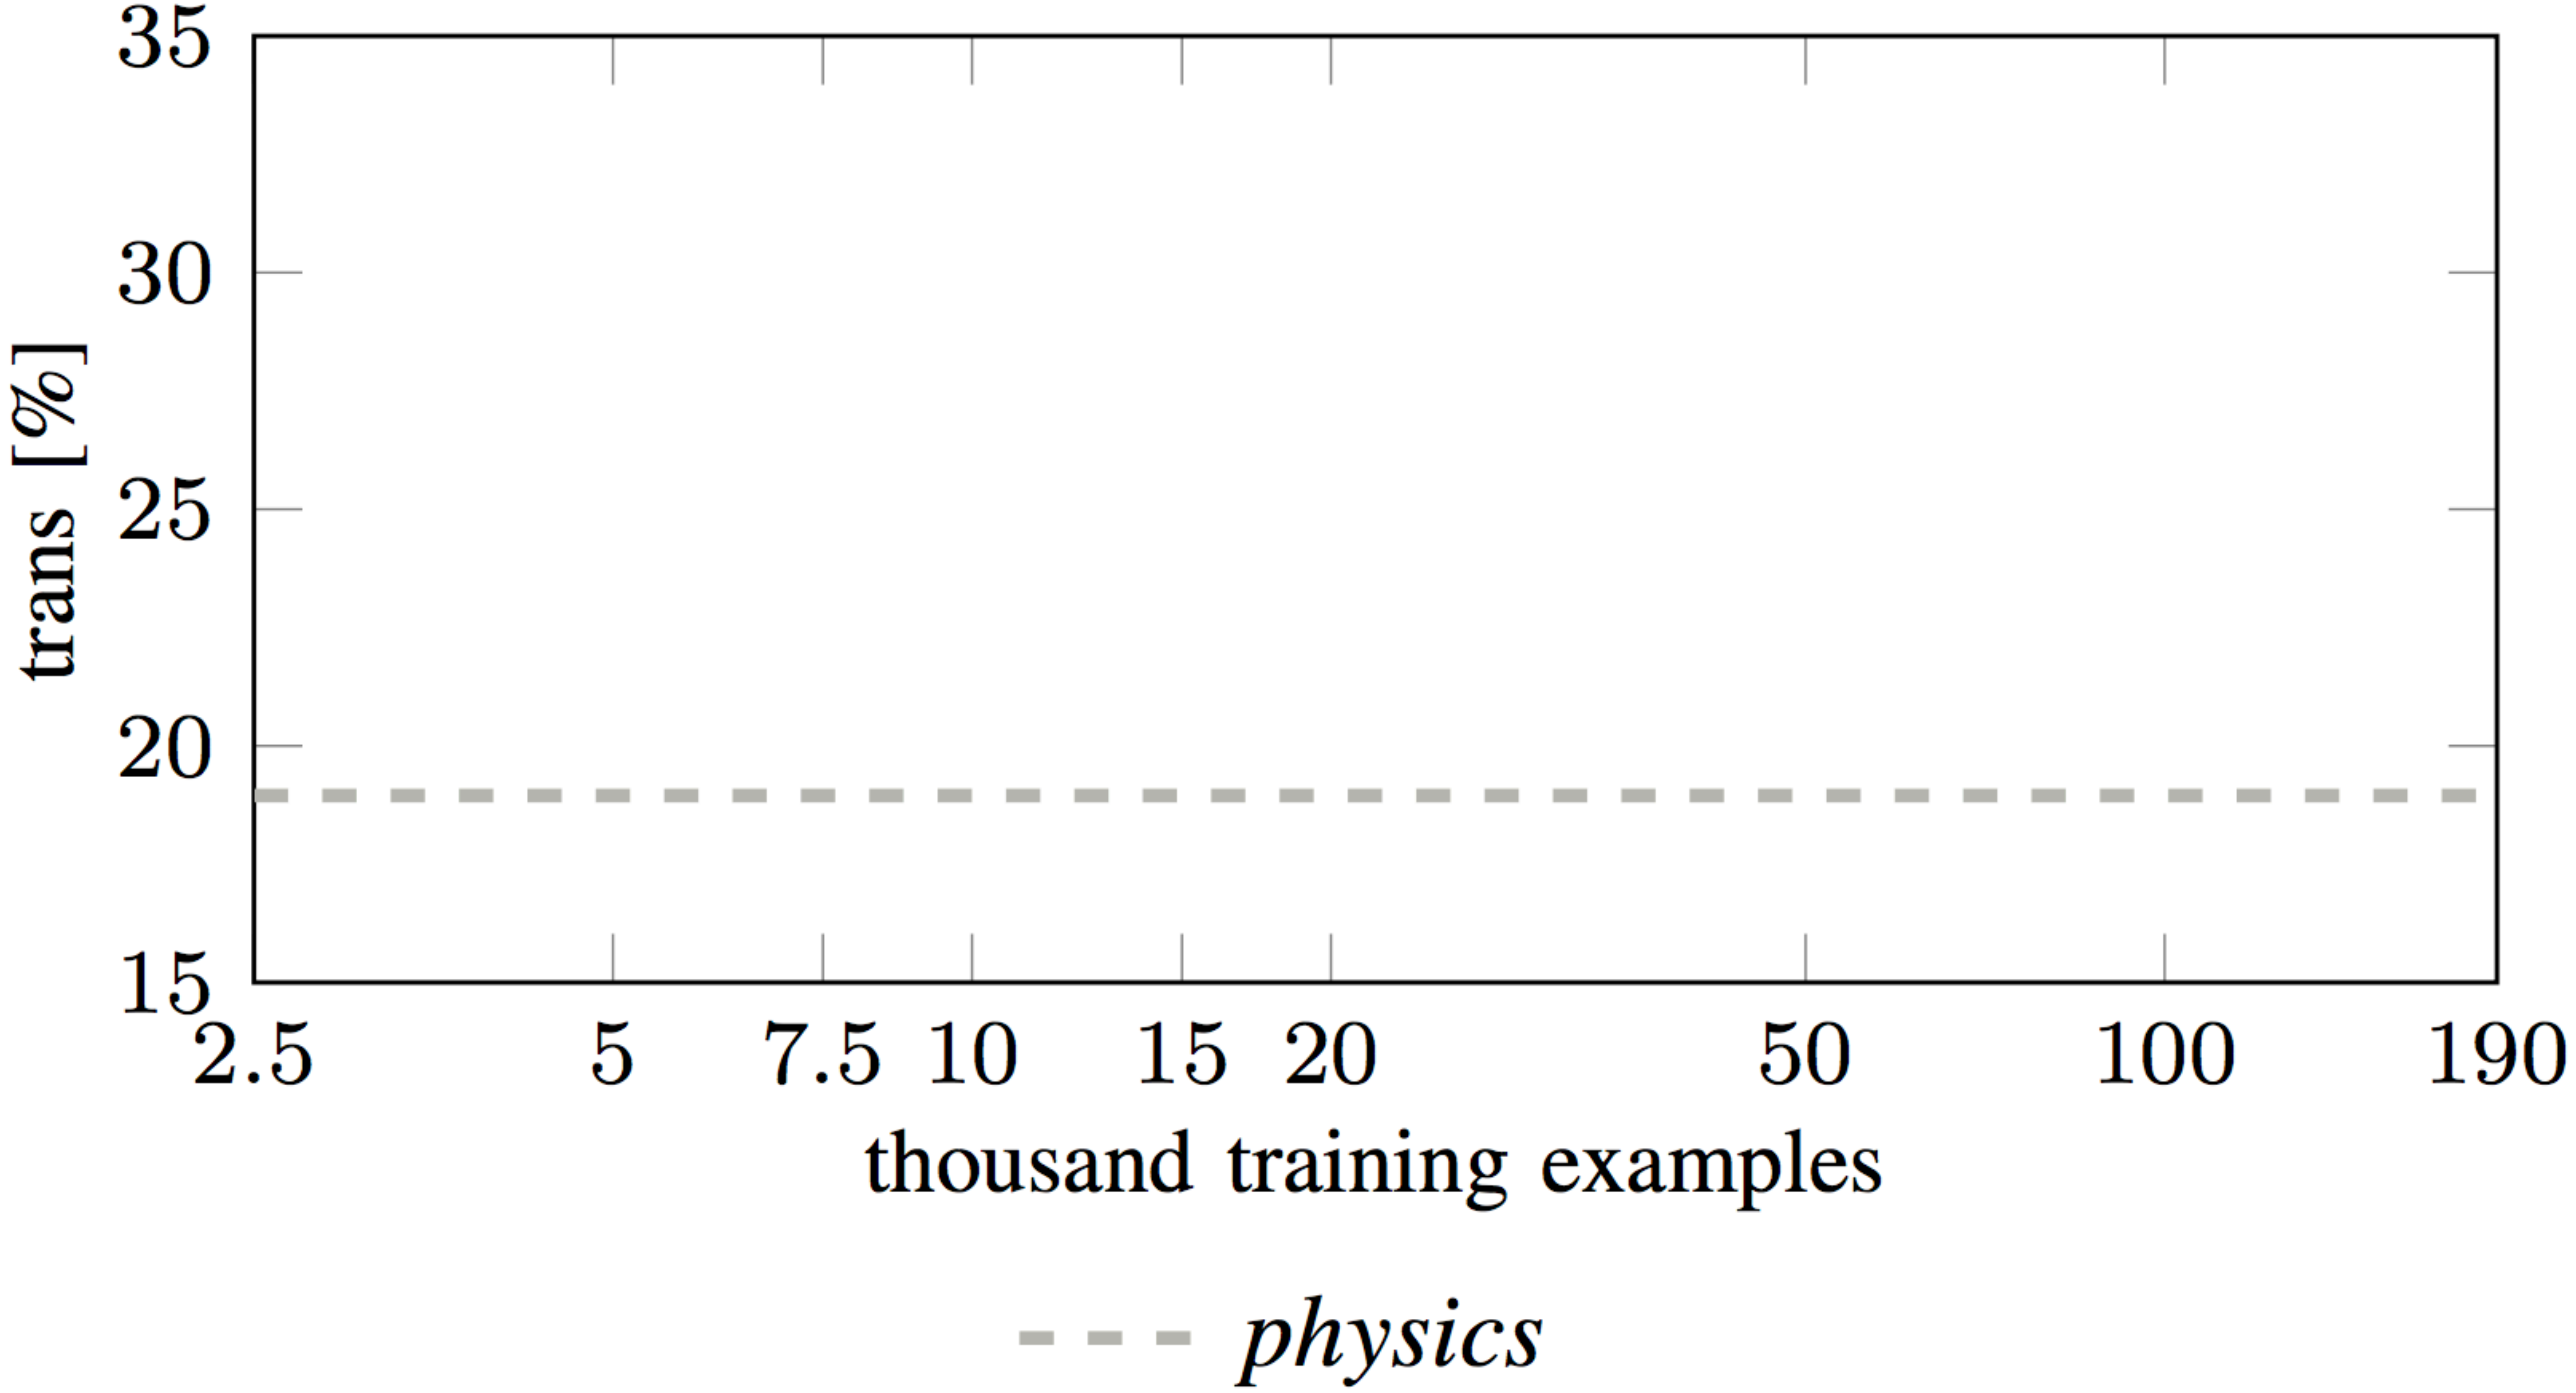

Testing Data Efficiency

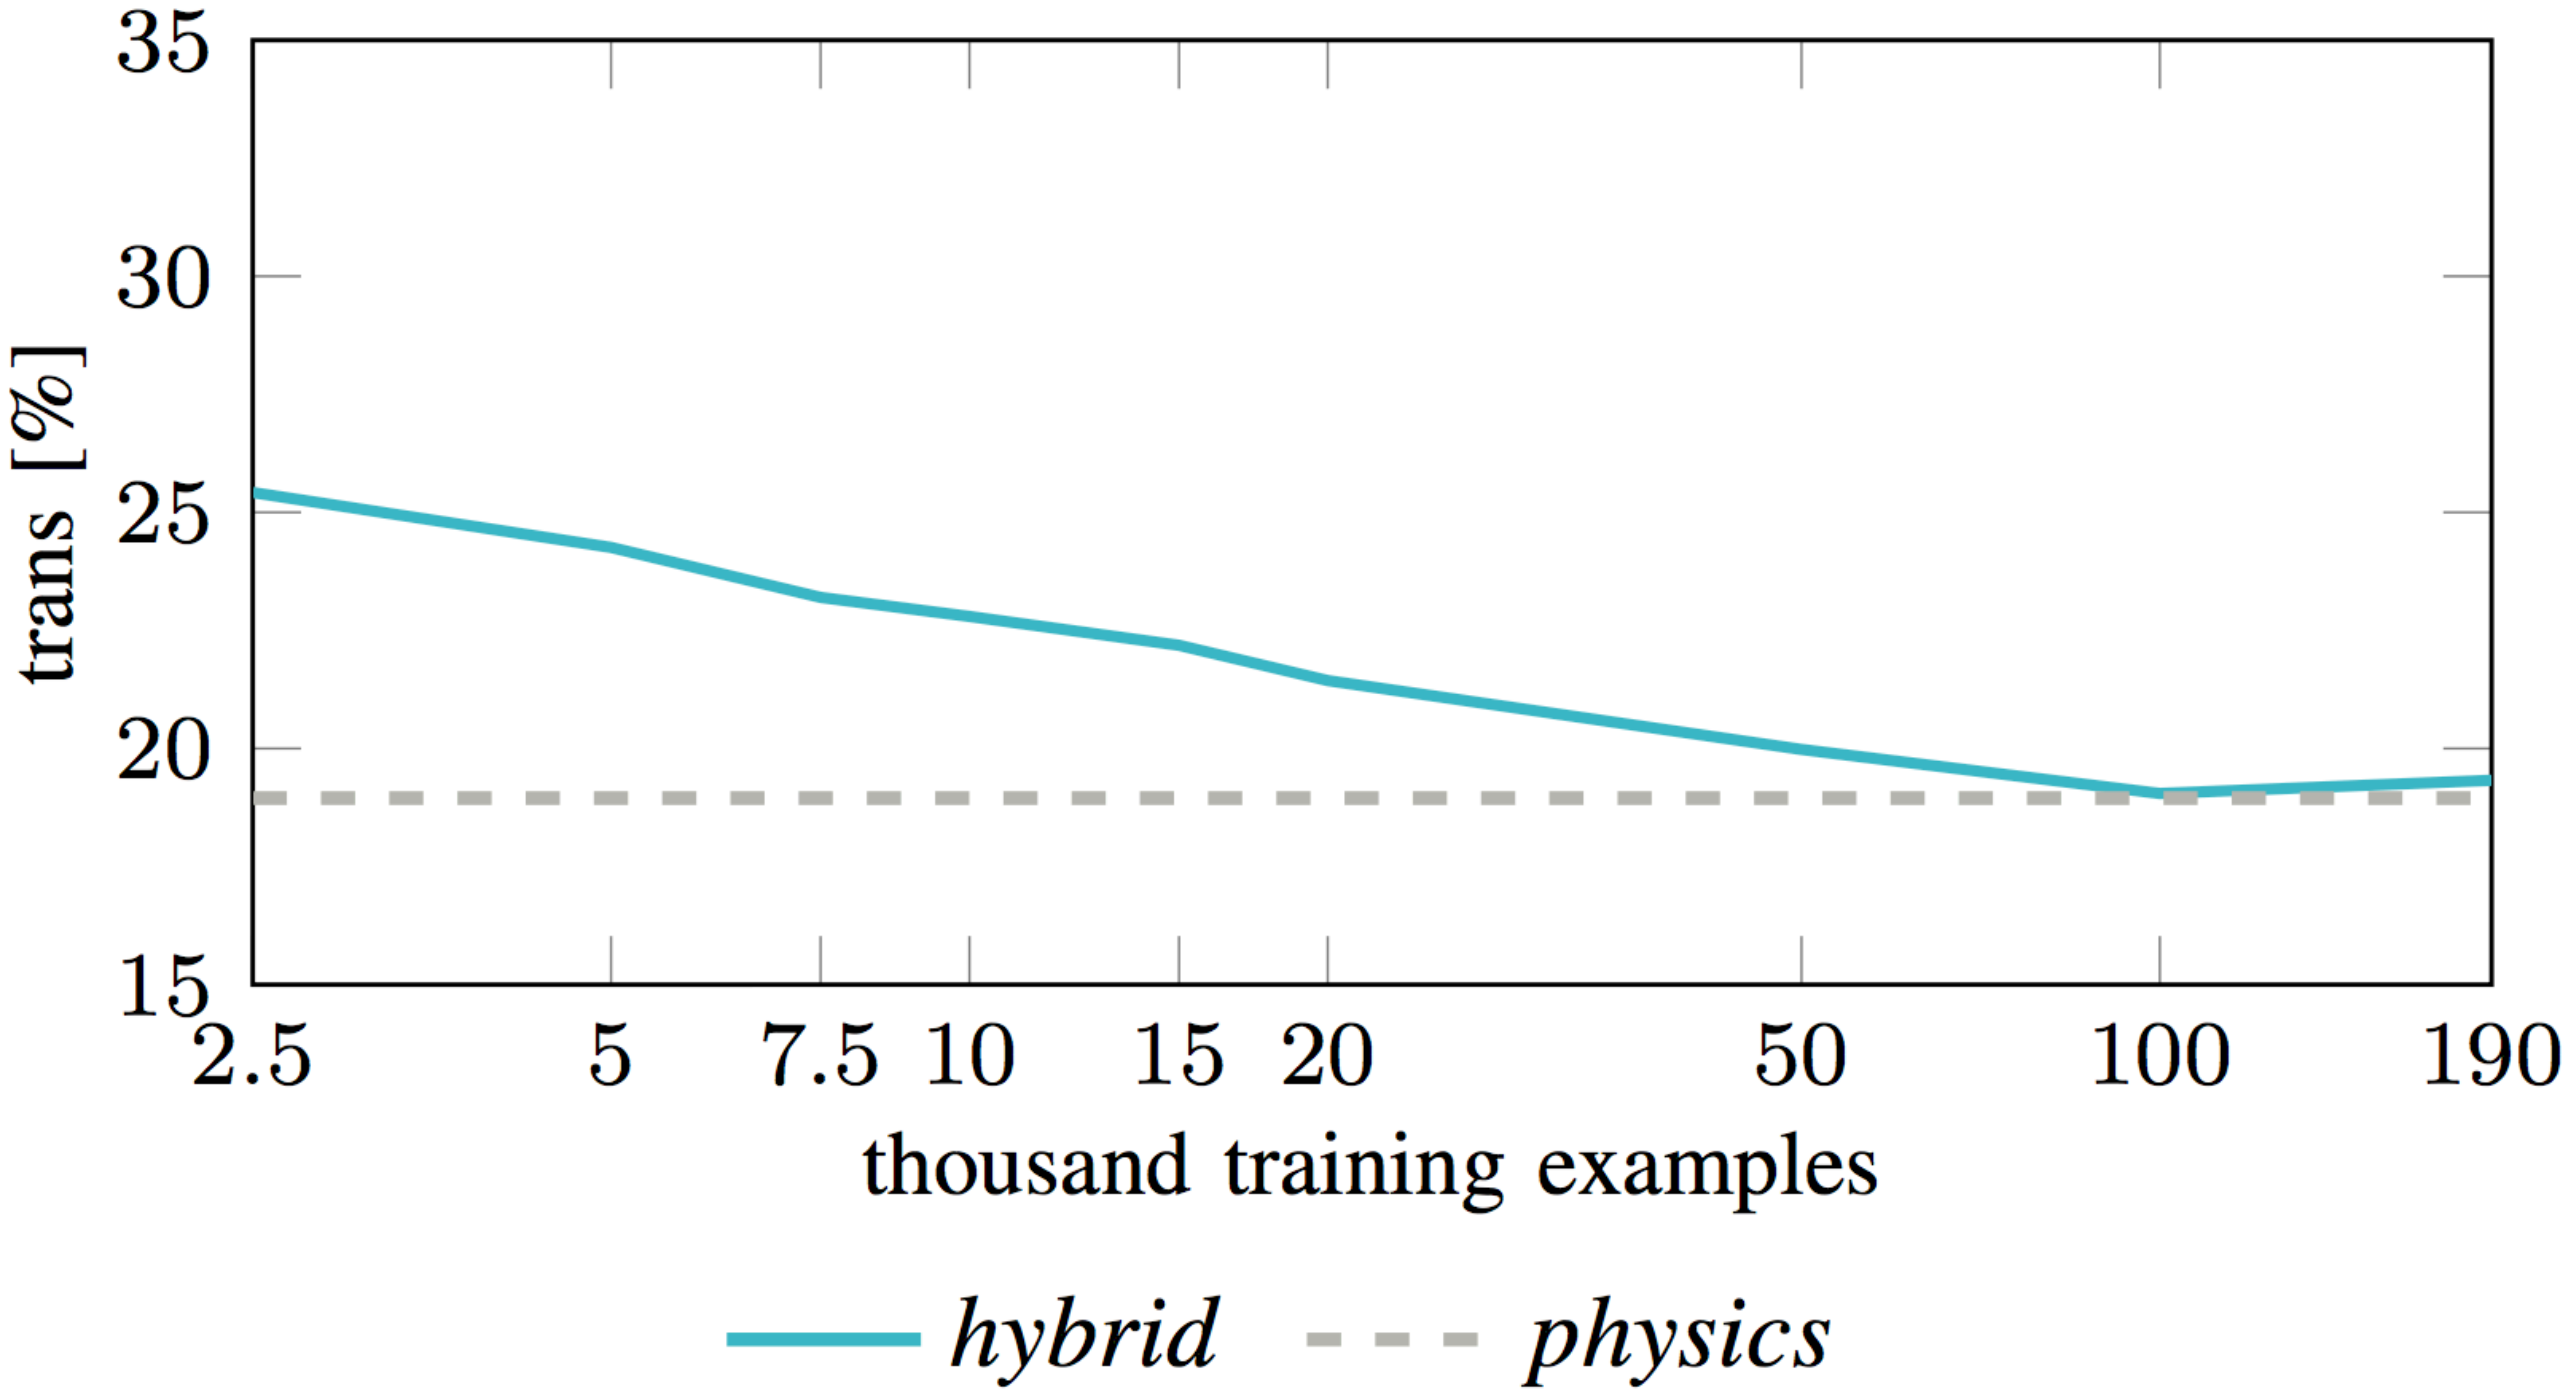

Testing Data Efficiency

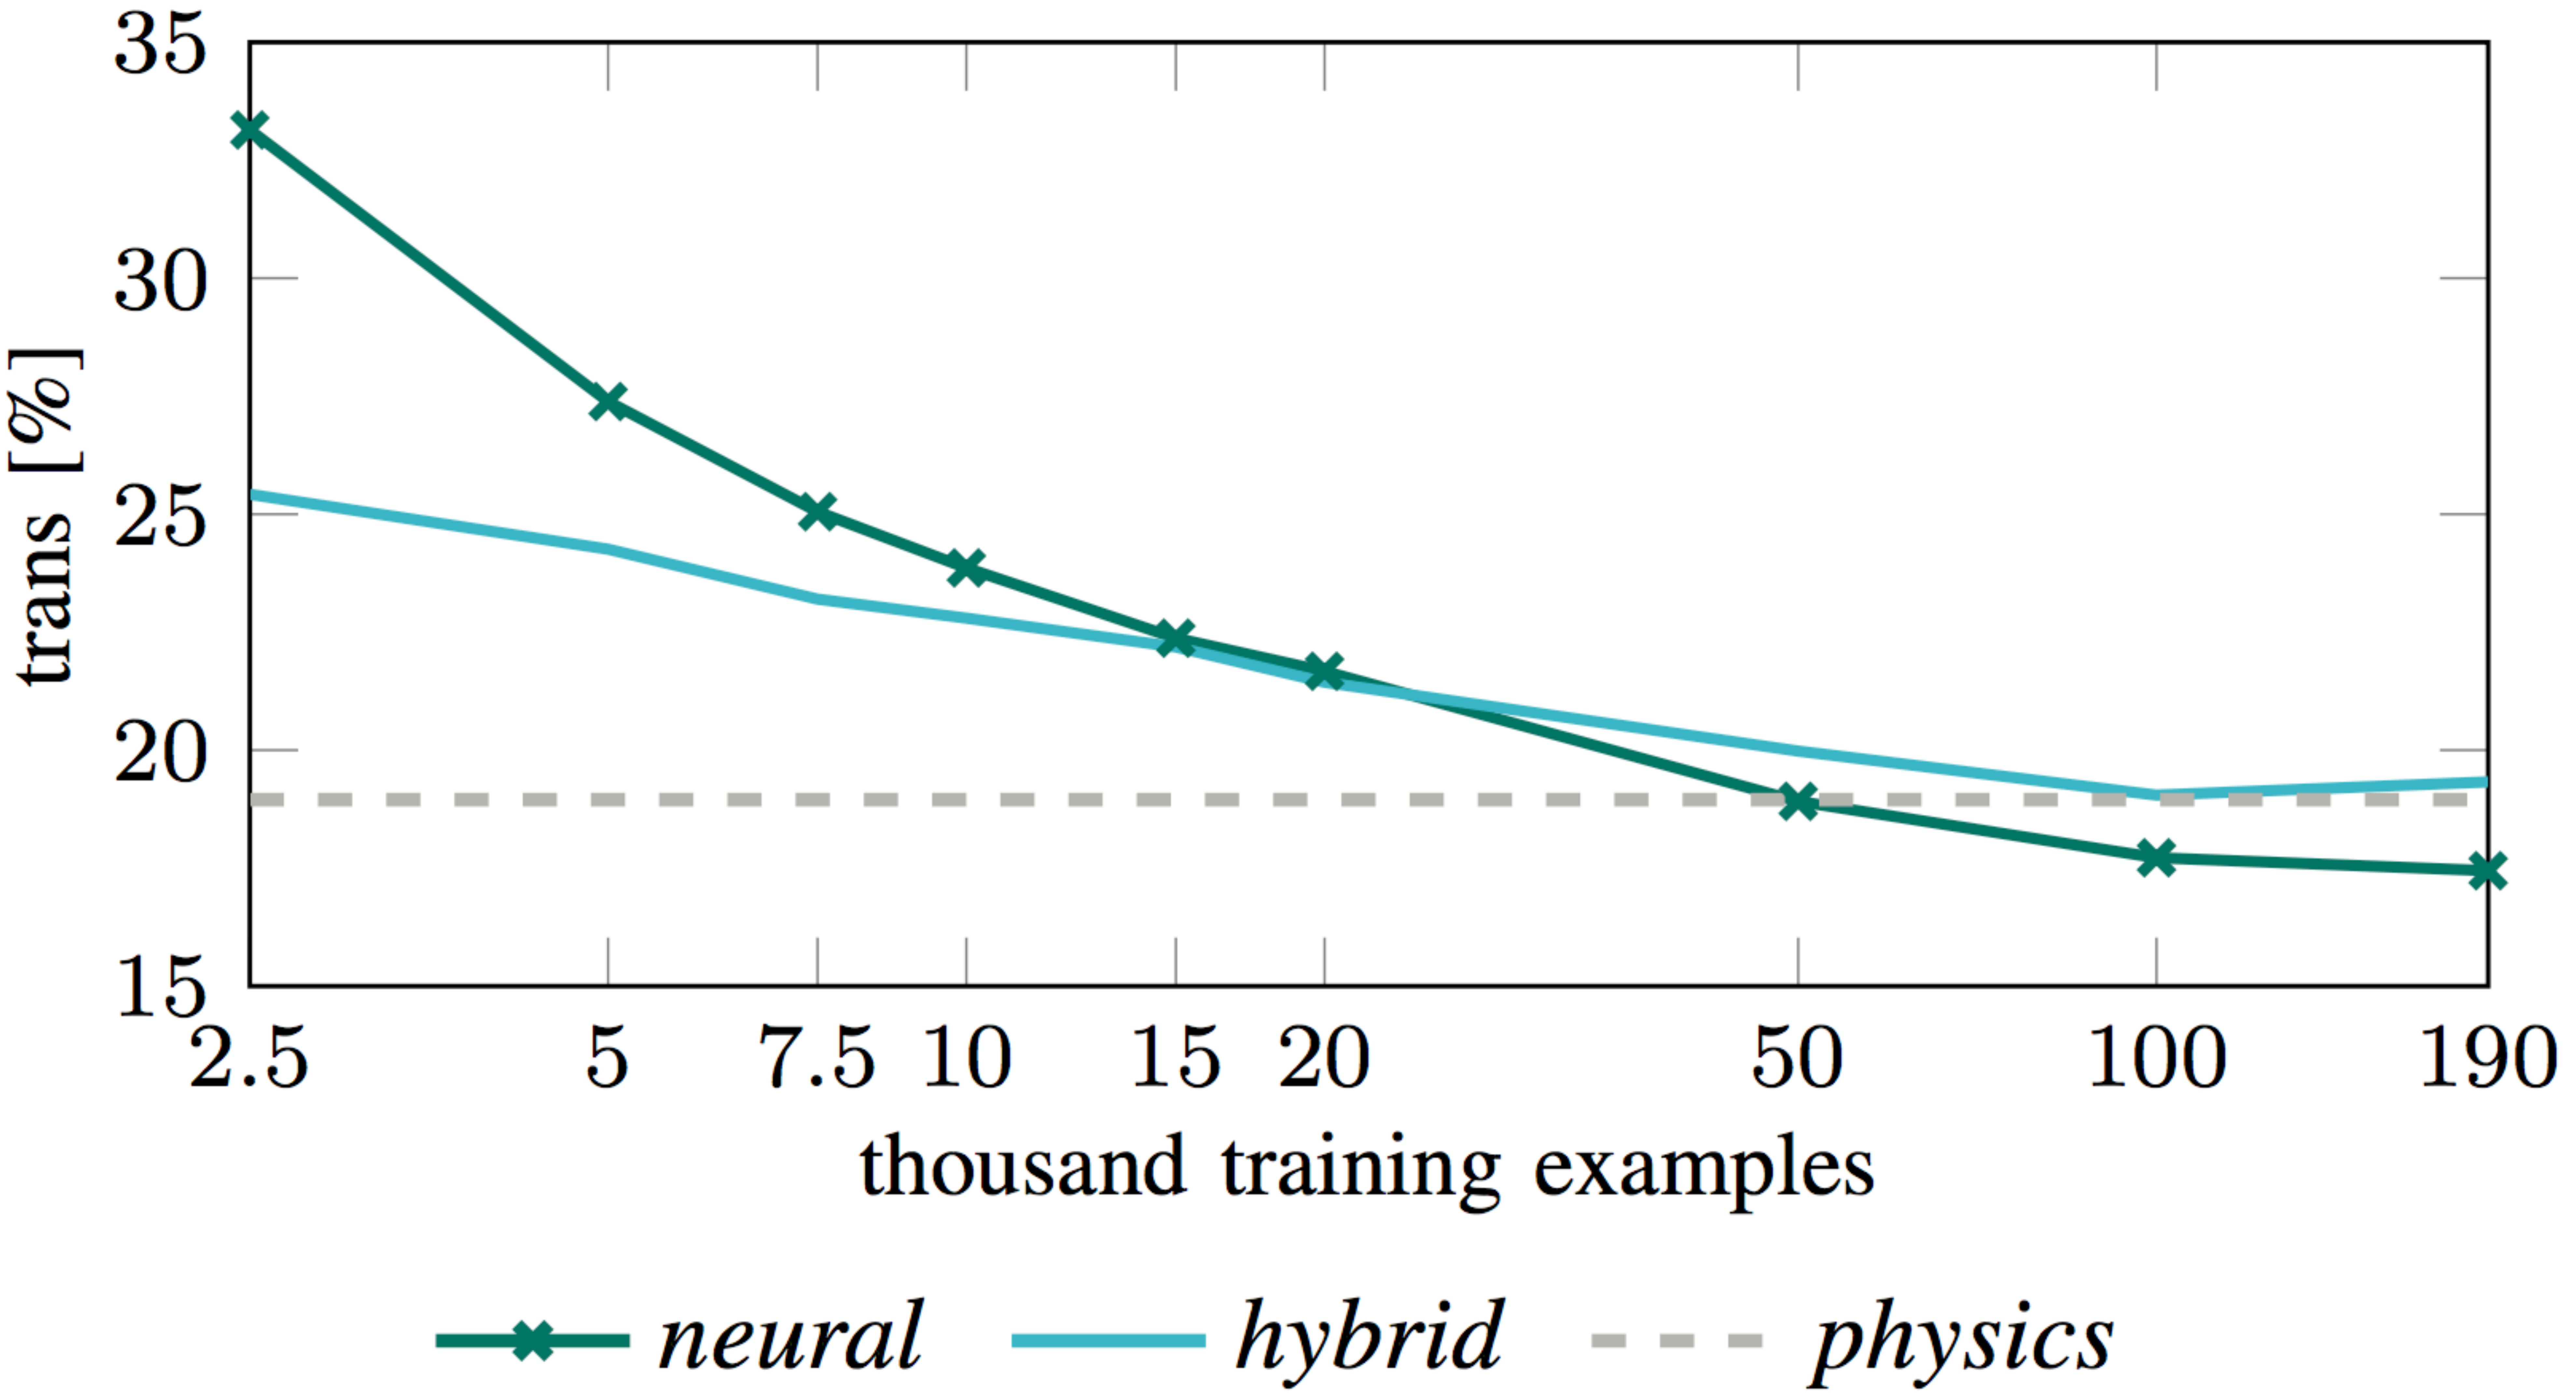

Testing Data Efficiency

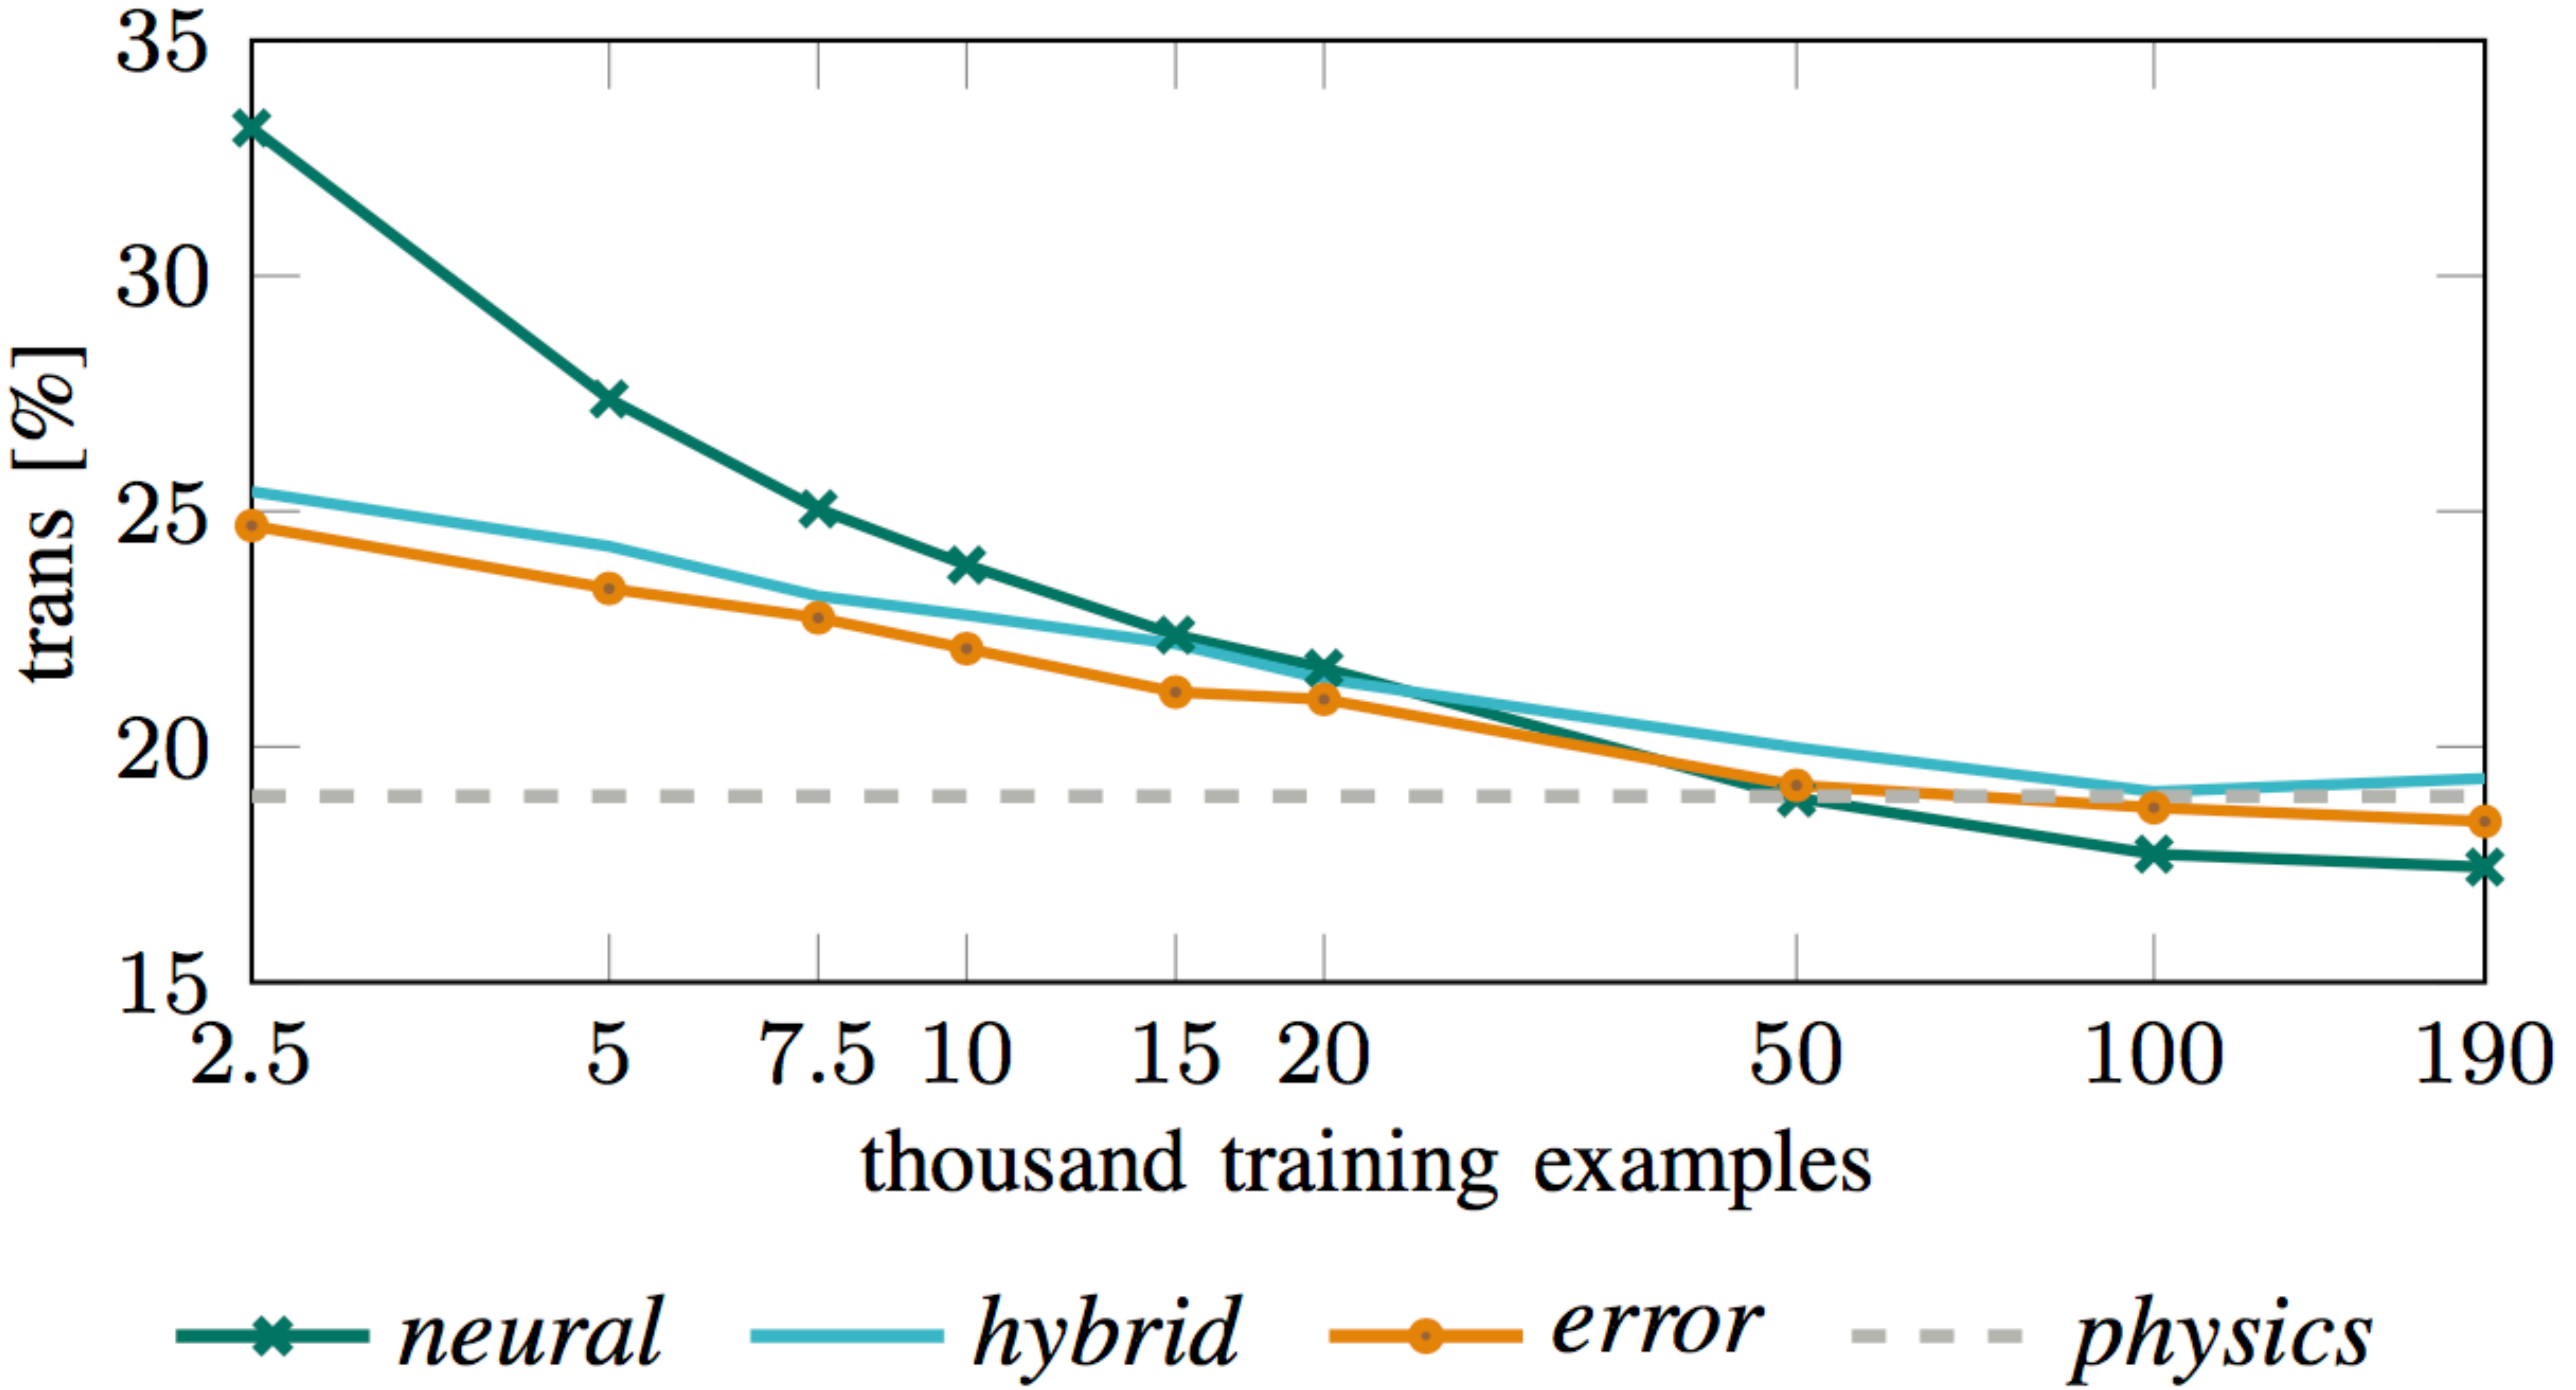

Testing Generalization

New Pushing Angles & Contact Points

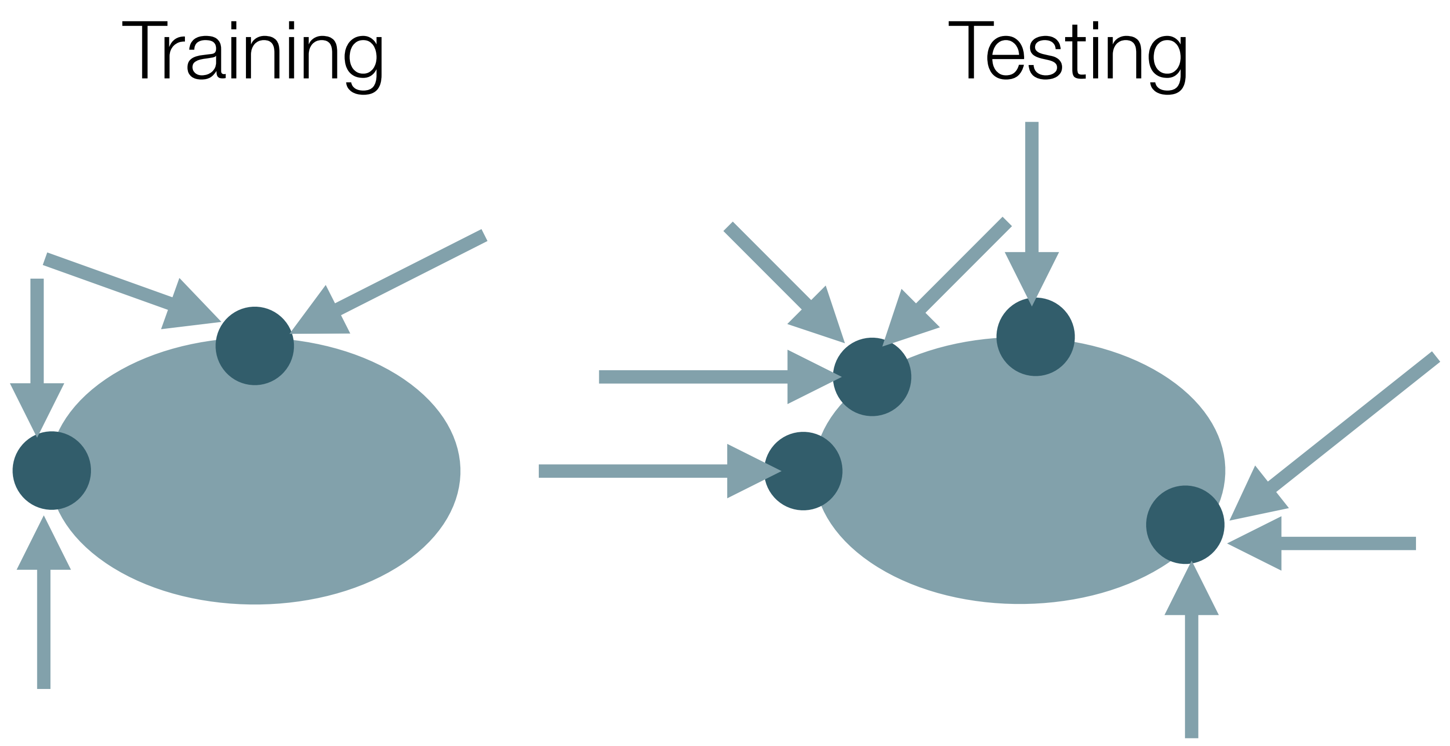

New Push Velocities

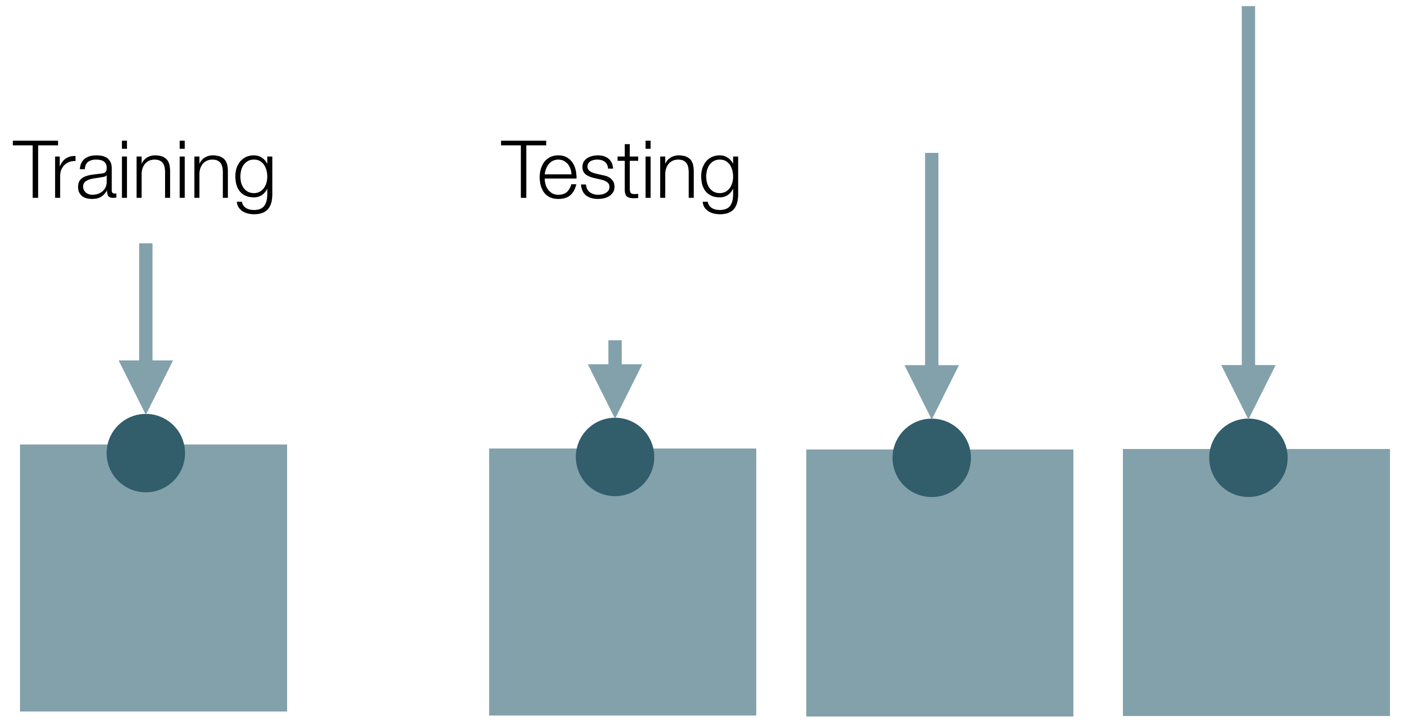

New Object Shapes

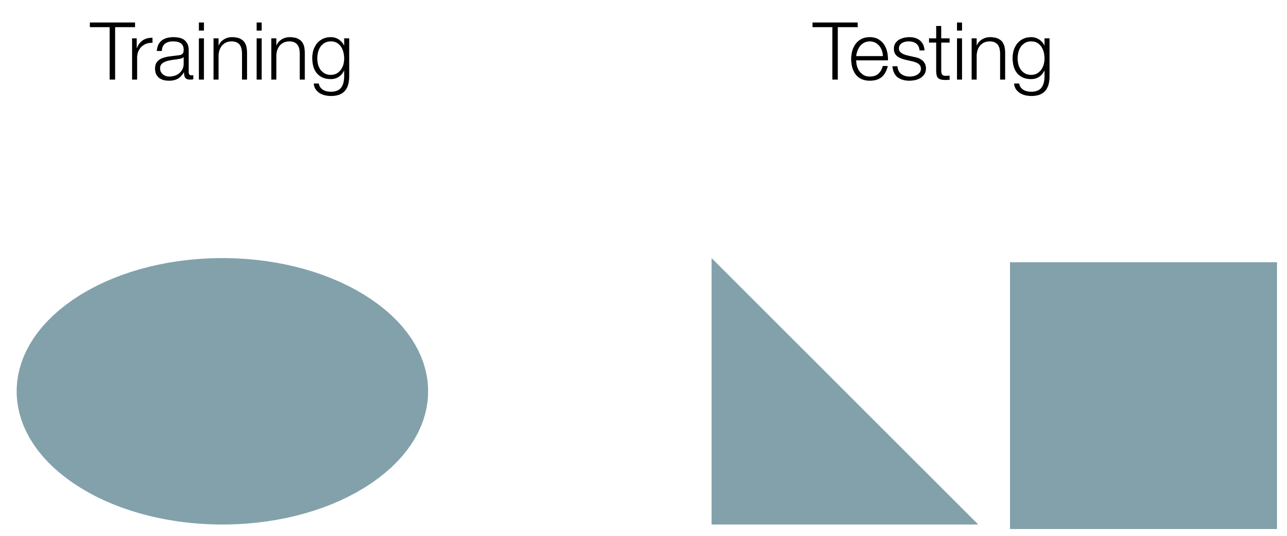

Interpolation

Extrapolation

Alina Kloss et al, "Combining learned and analytical models for predicting action effects," Submitted to IJRR. 2018. Pre-print on arXiv.

Generalization to new push velocities

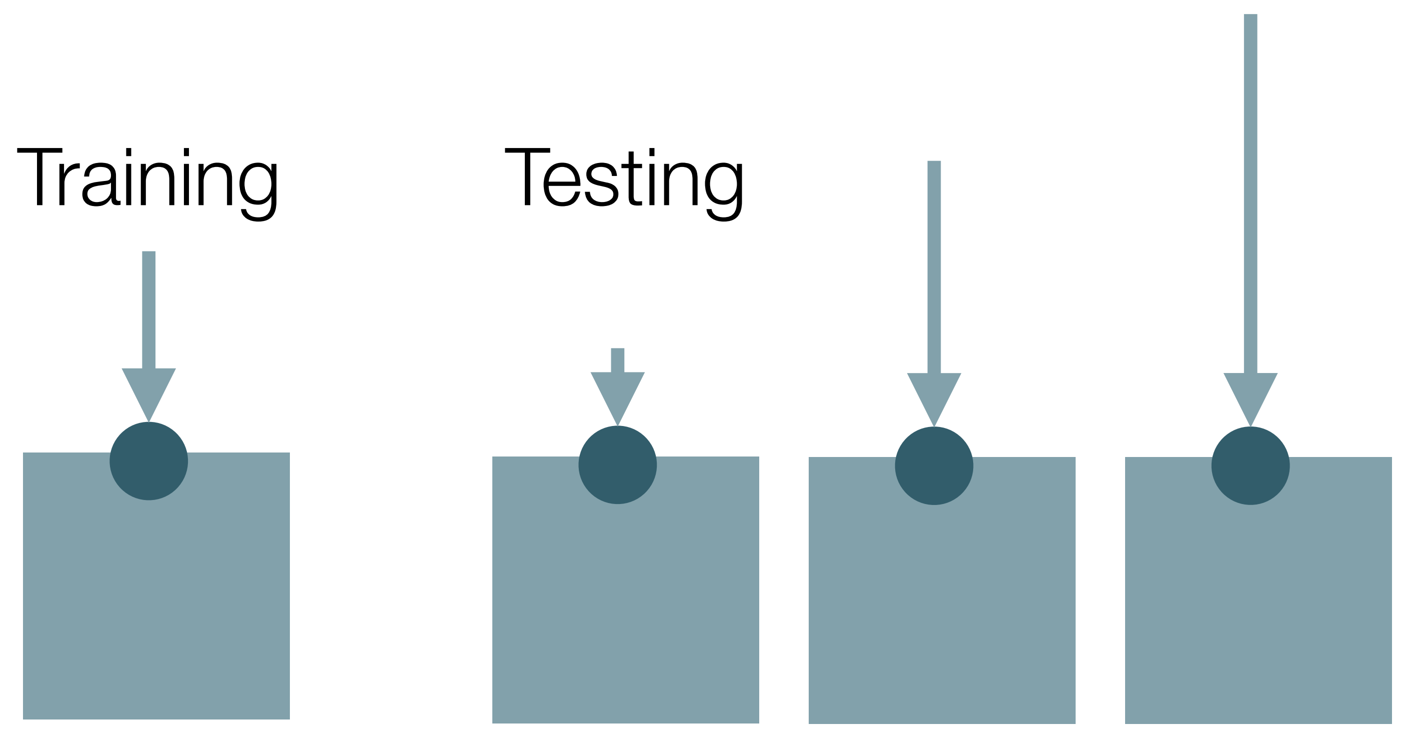

Extrapolation

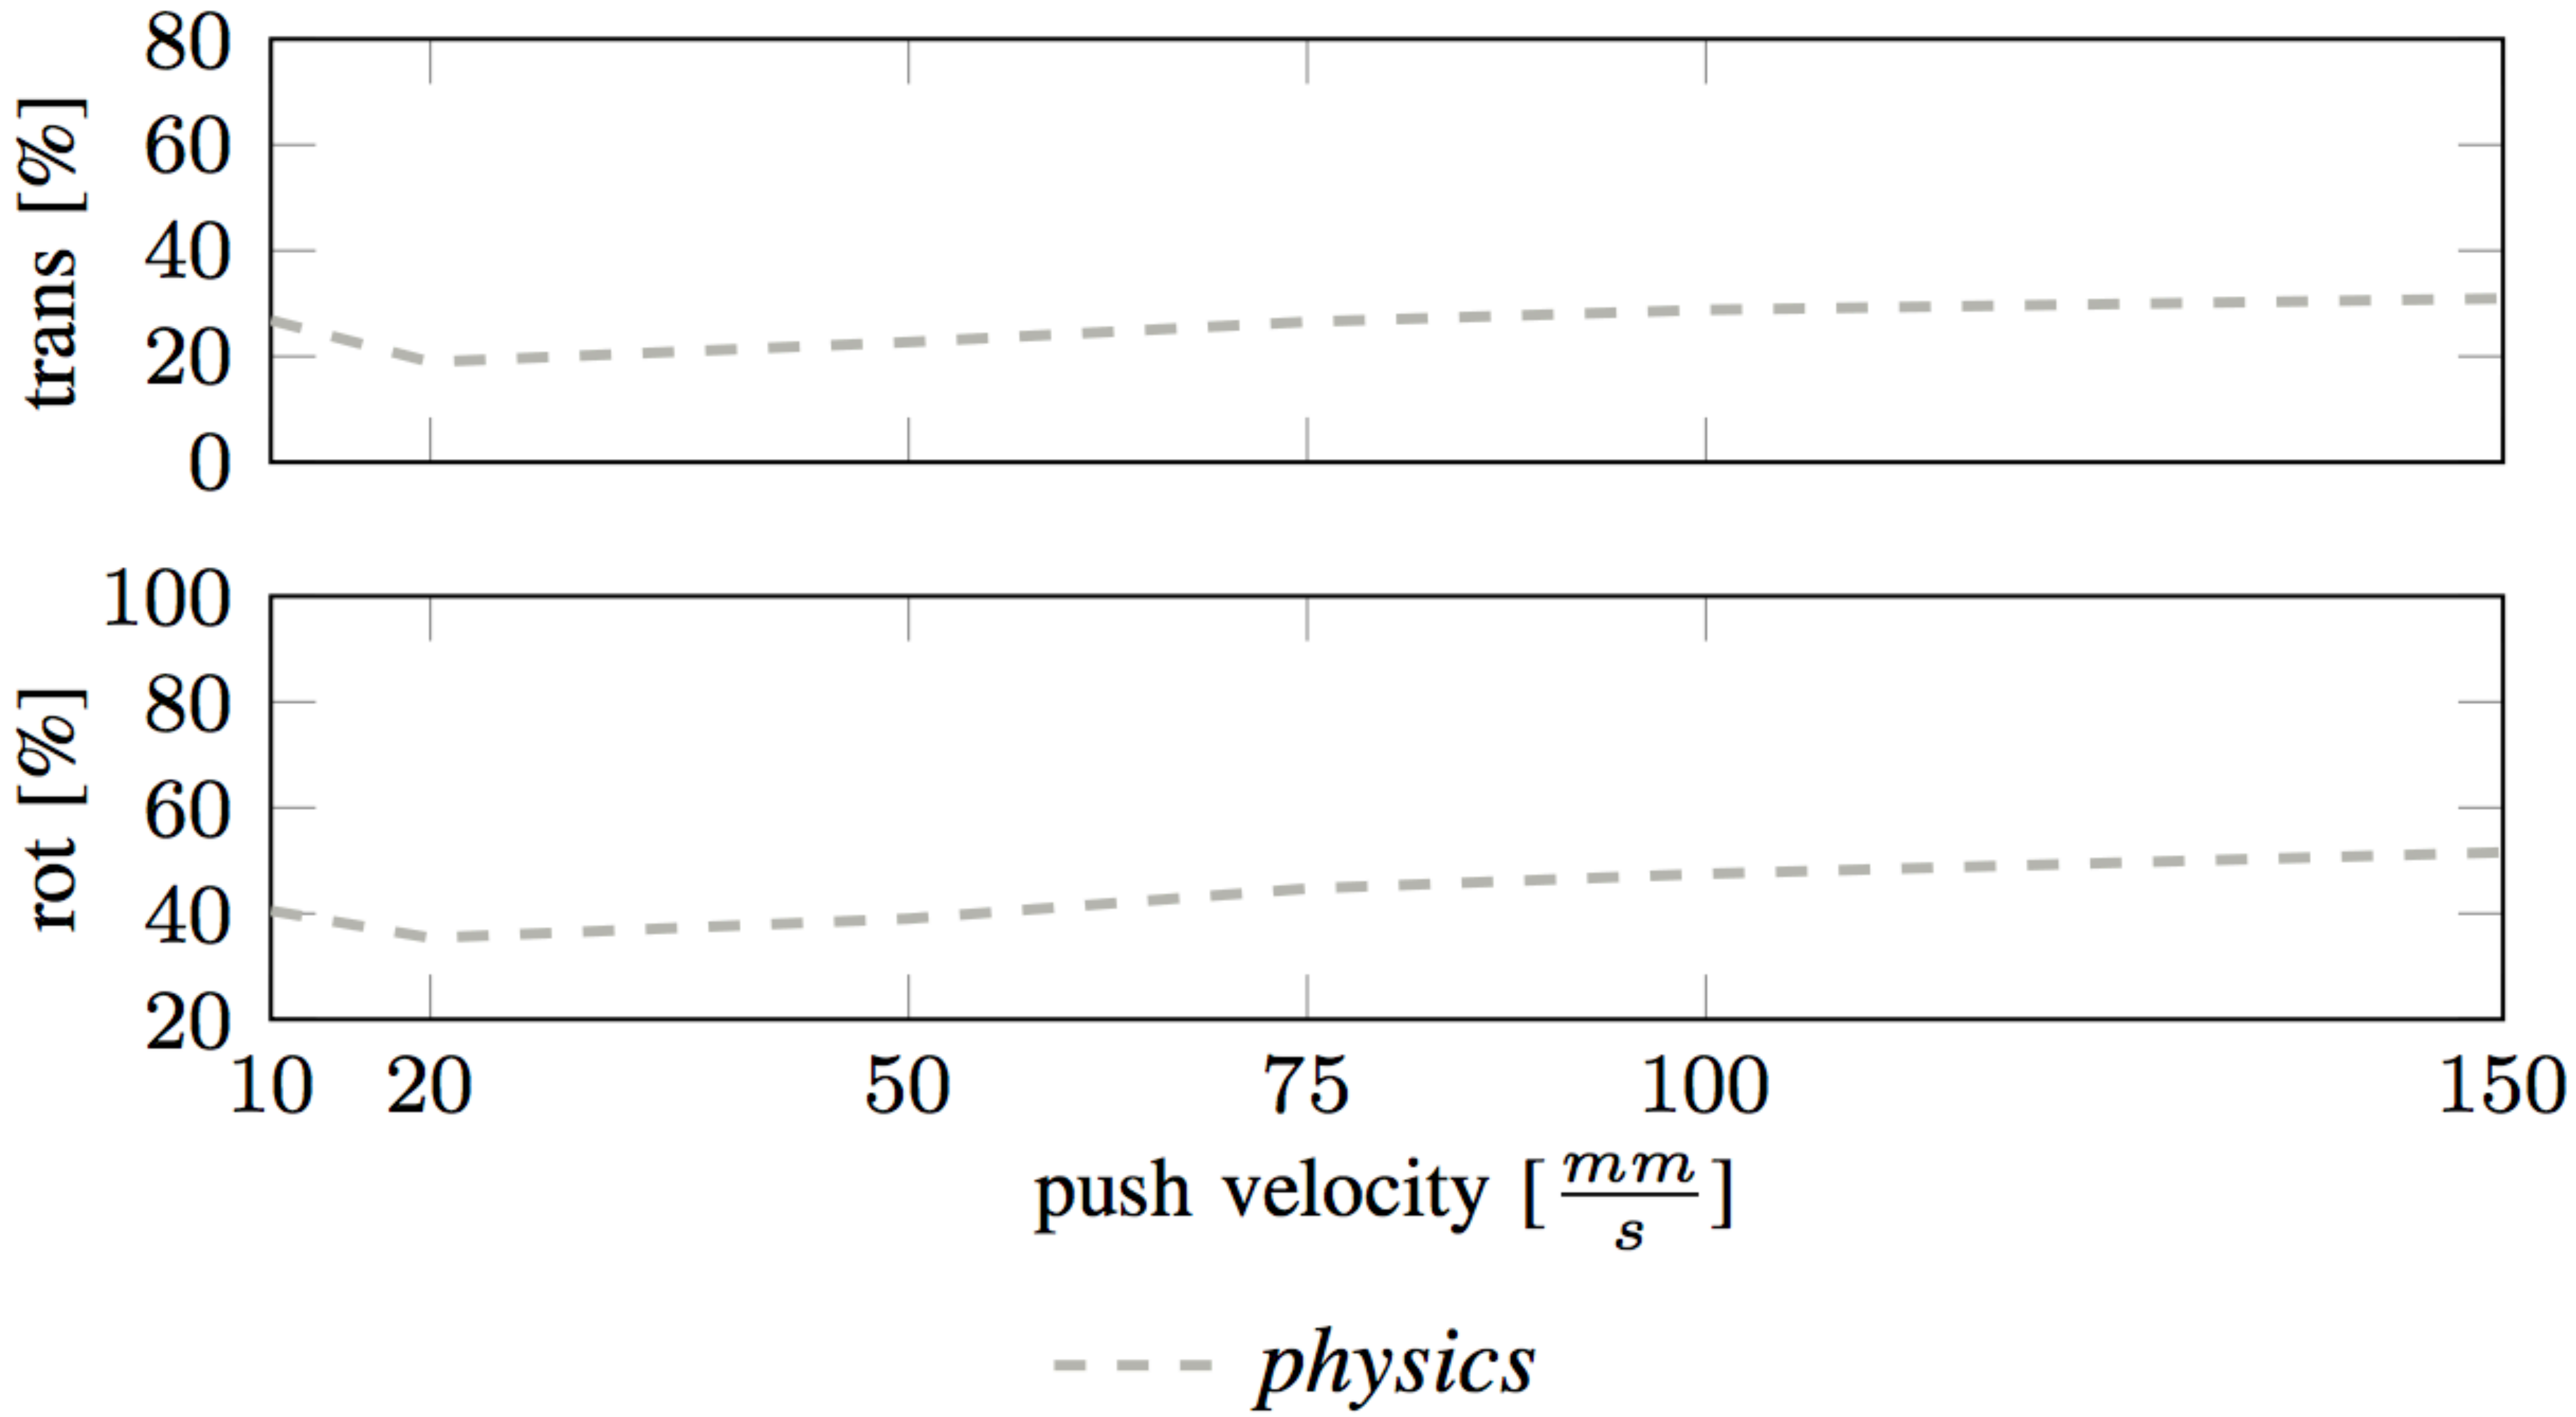

Generalization to new push velocities

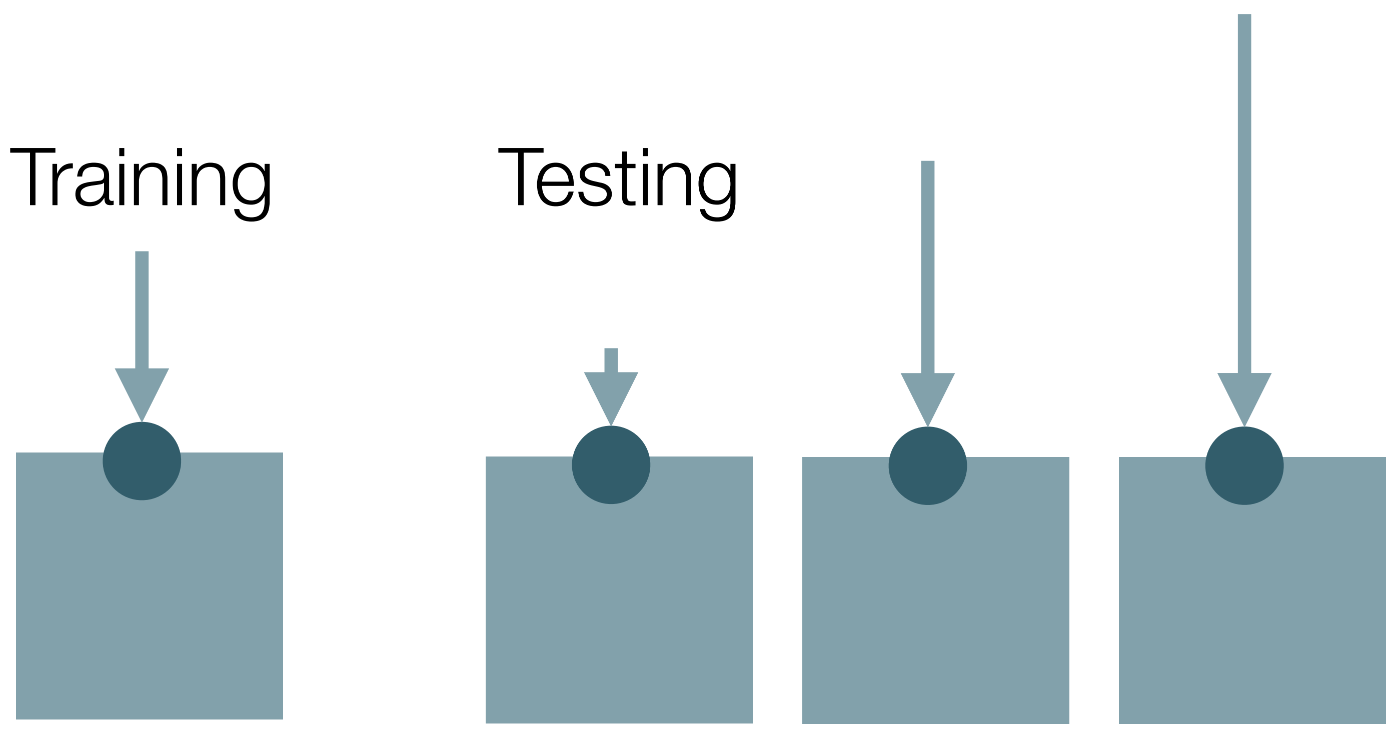

Extrapolation

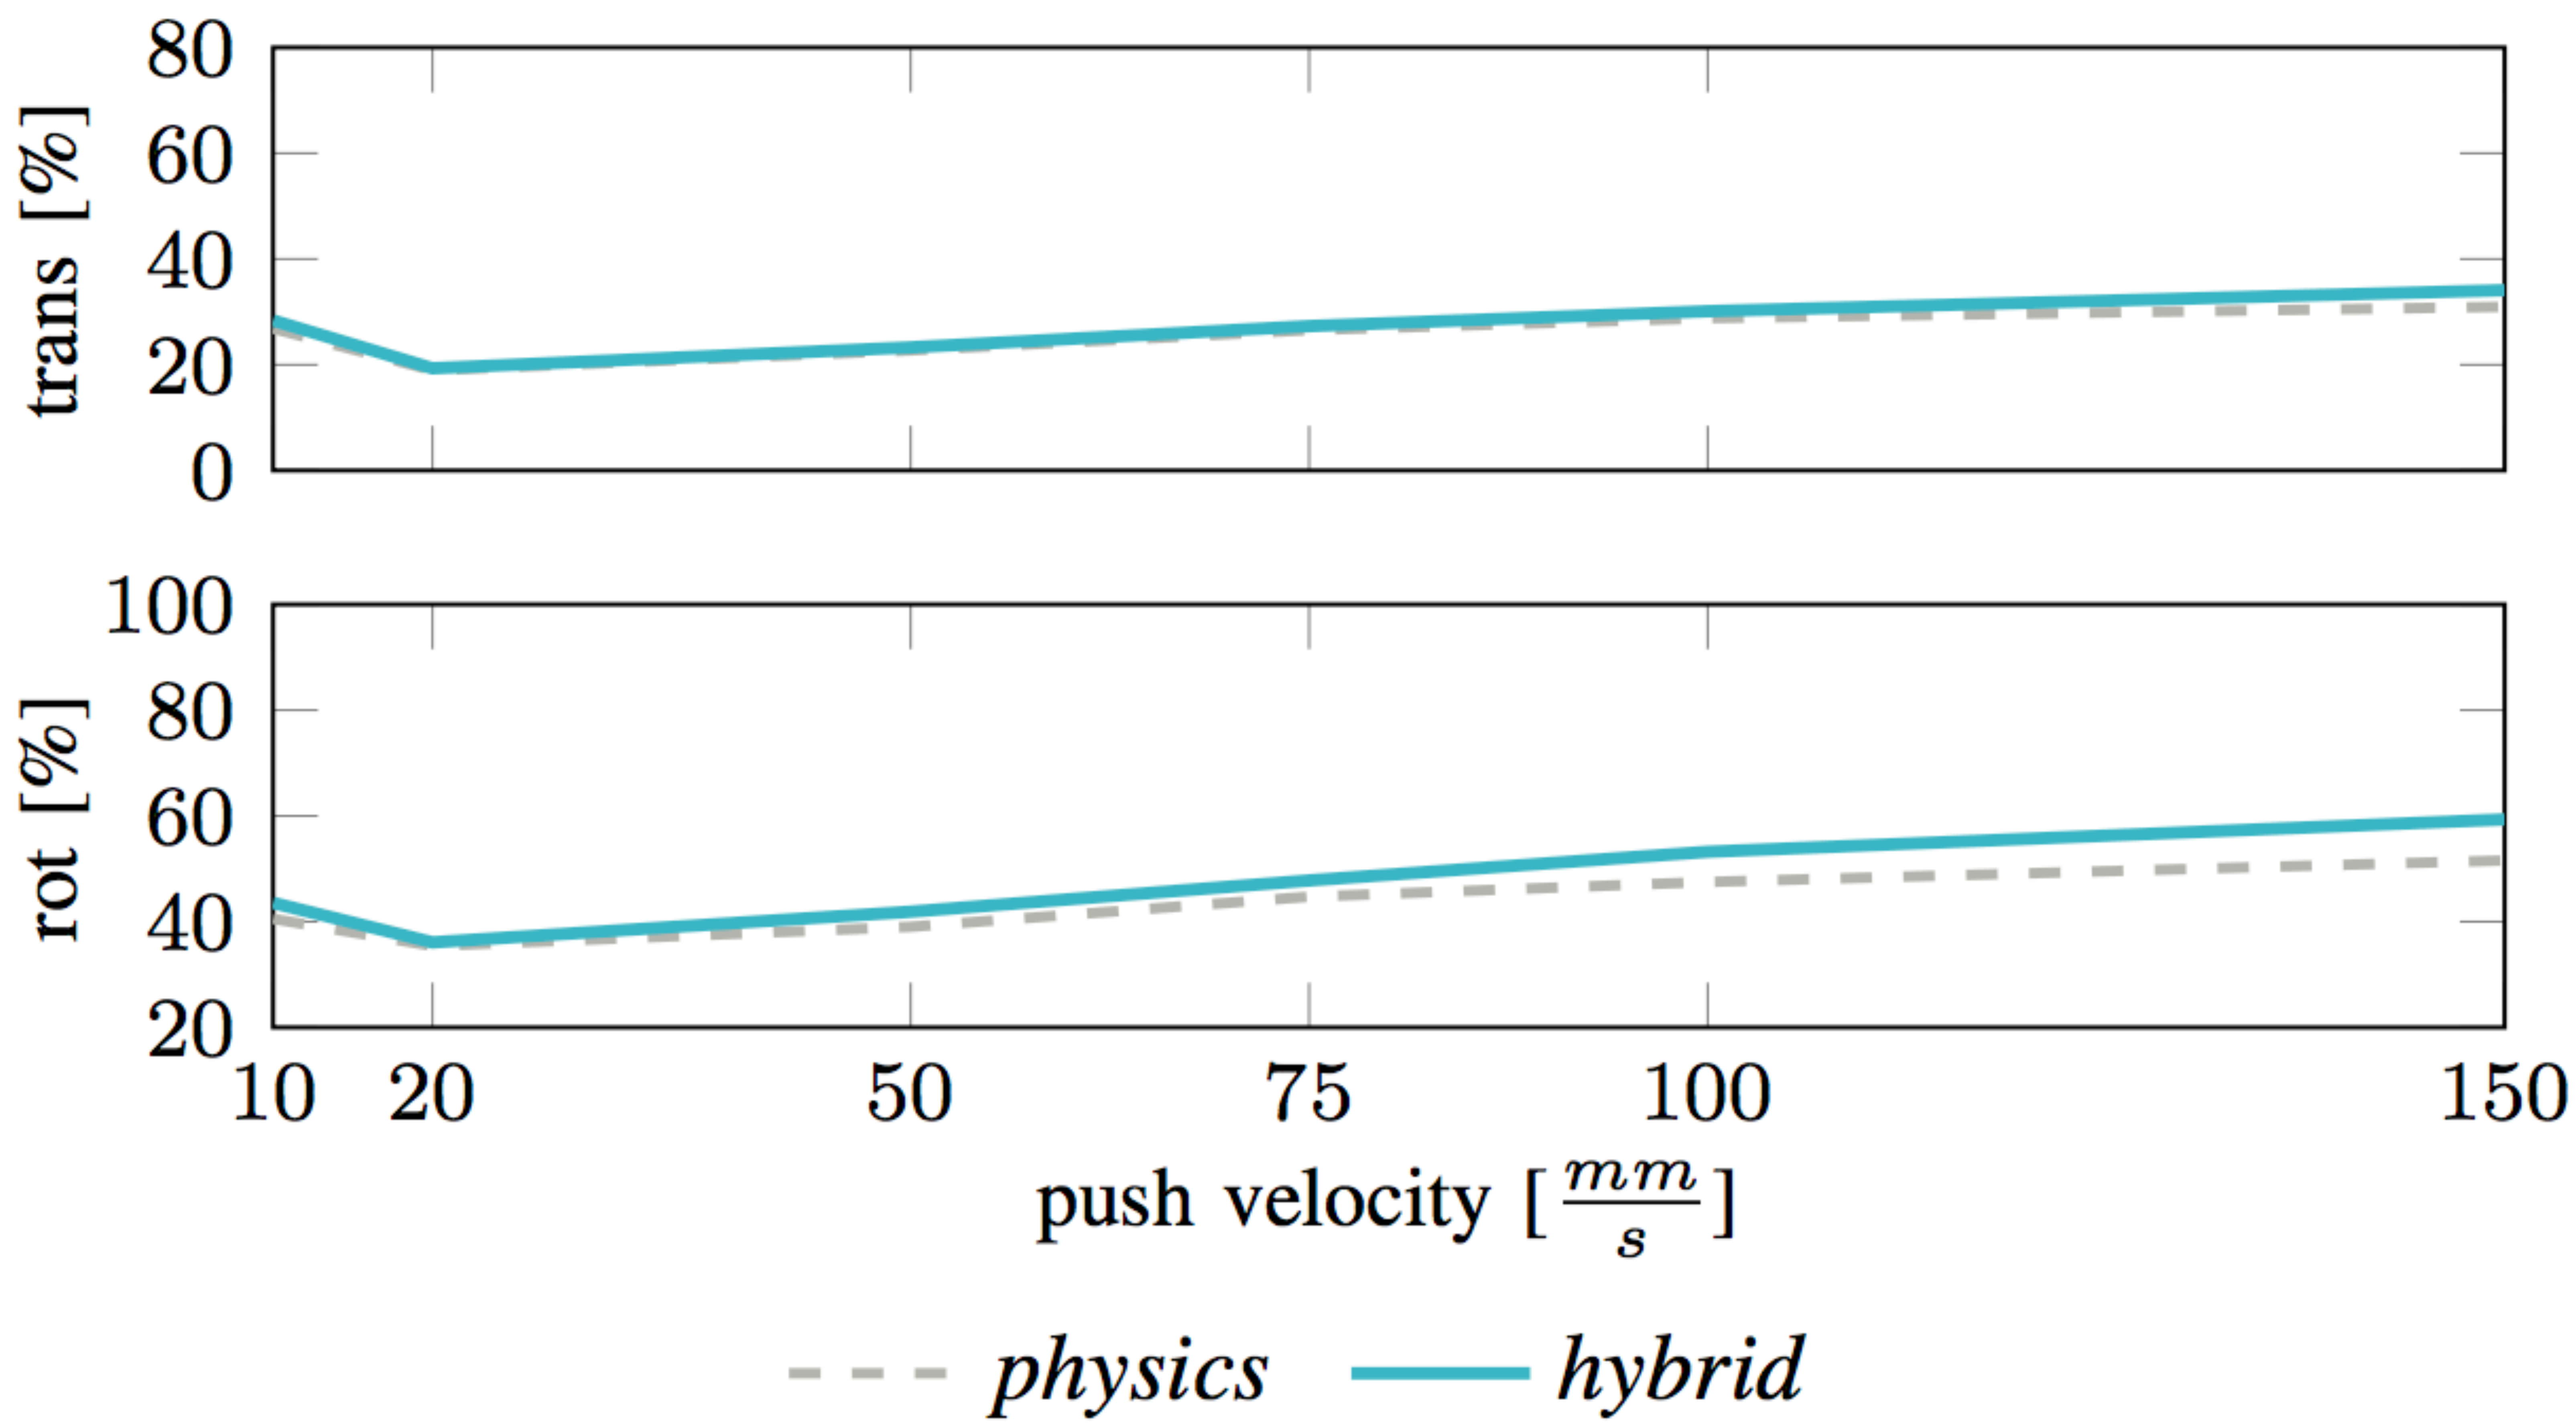

Generalization to new push velocities

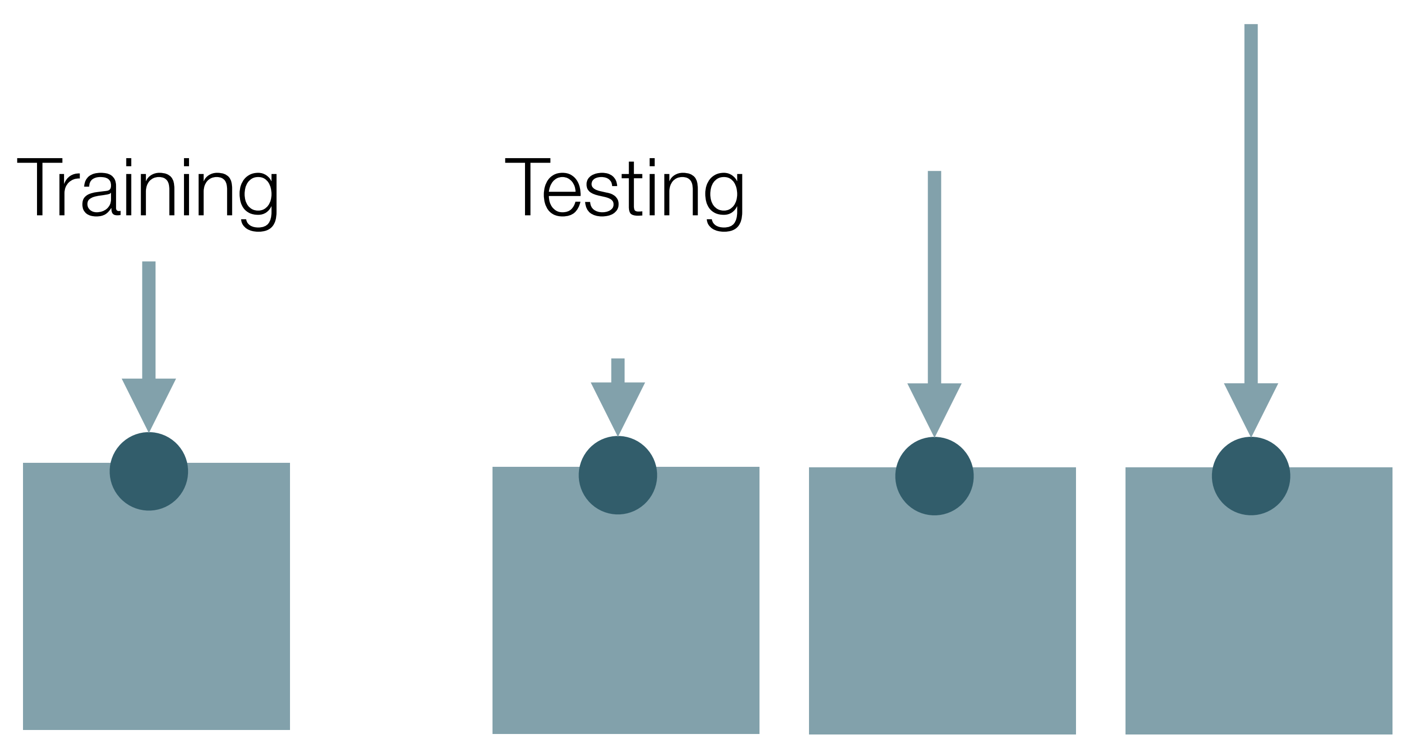

Extrapolation

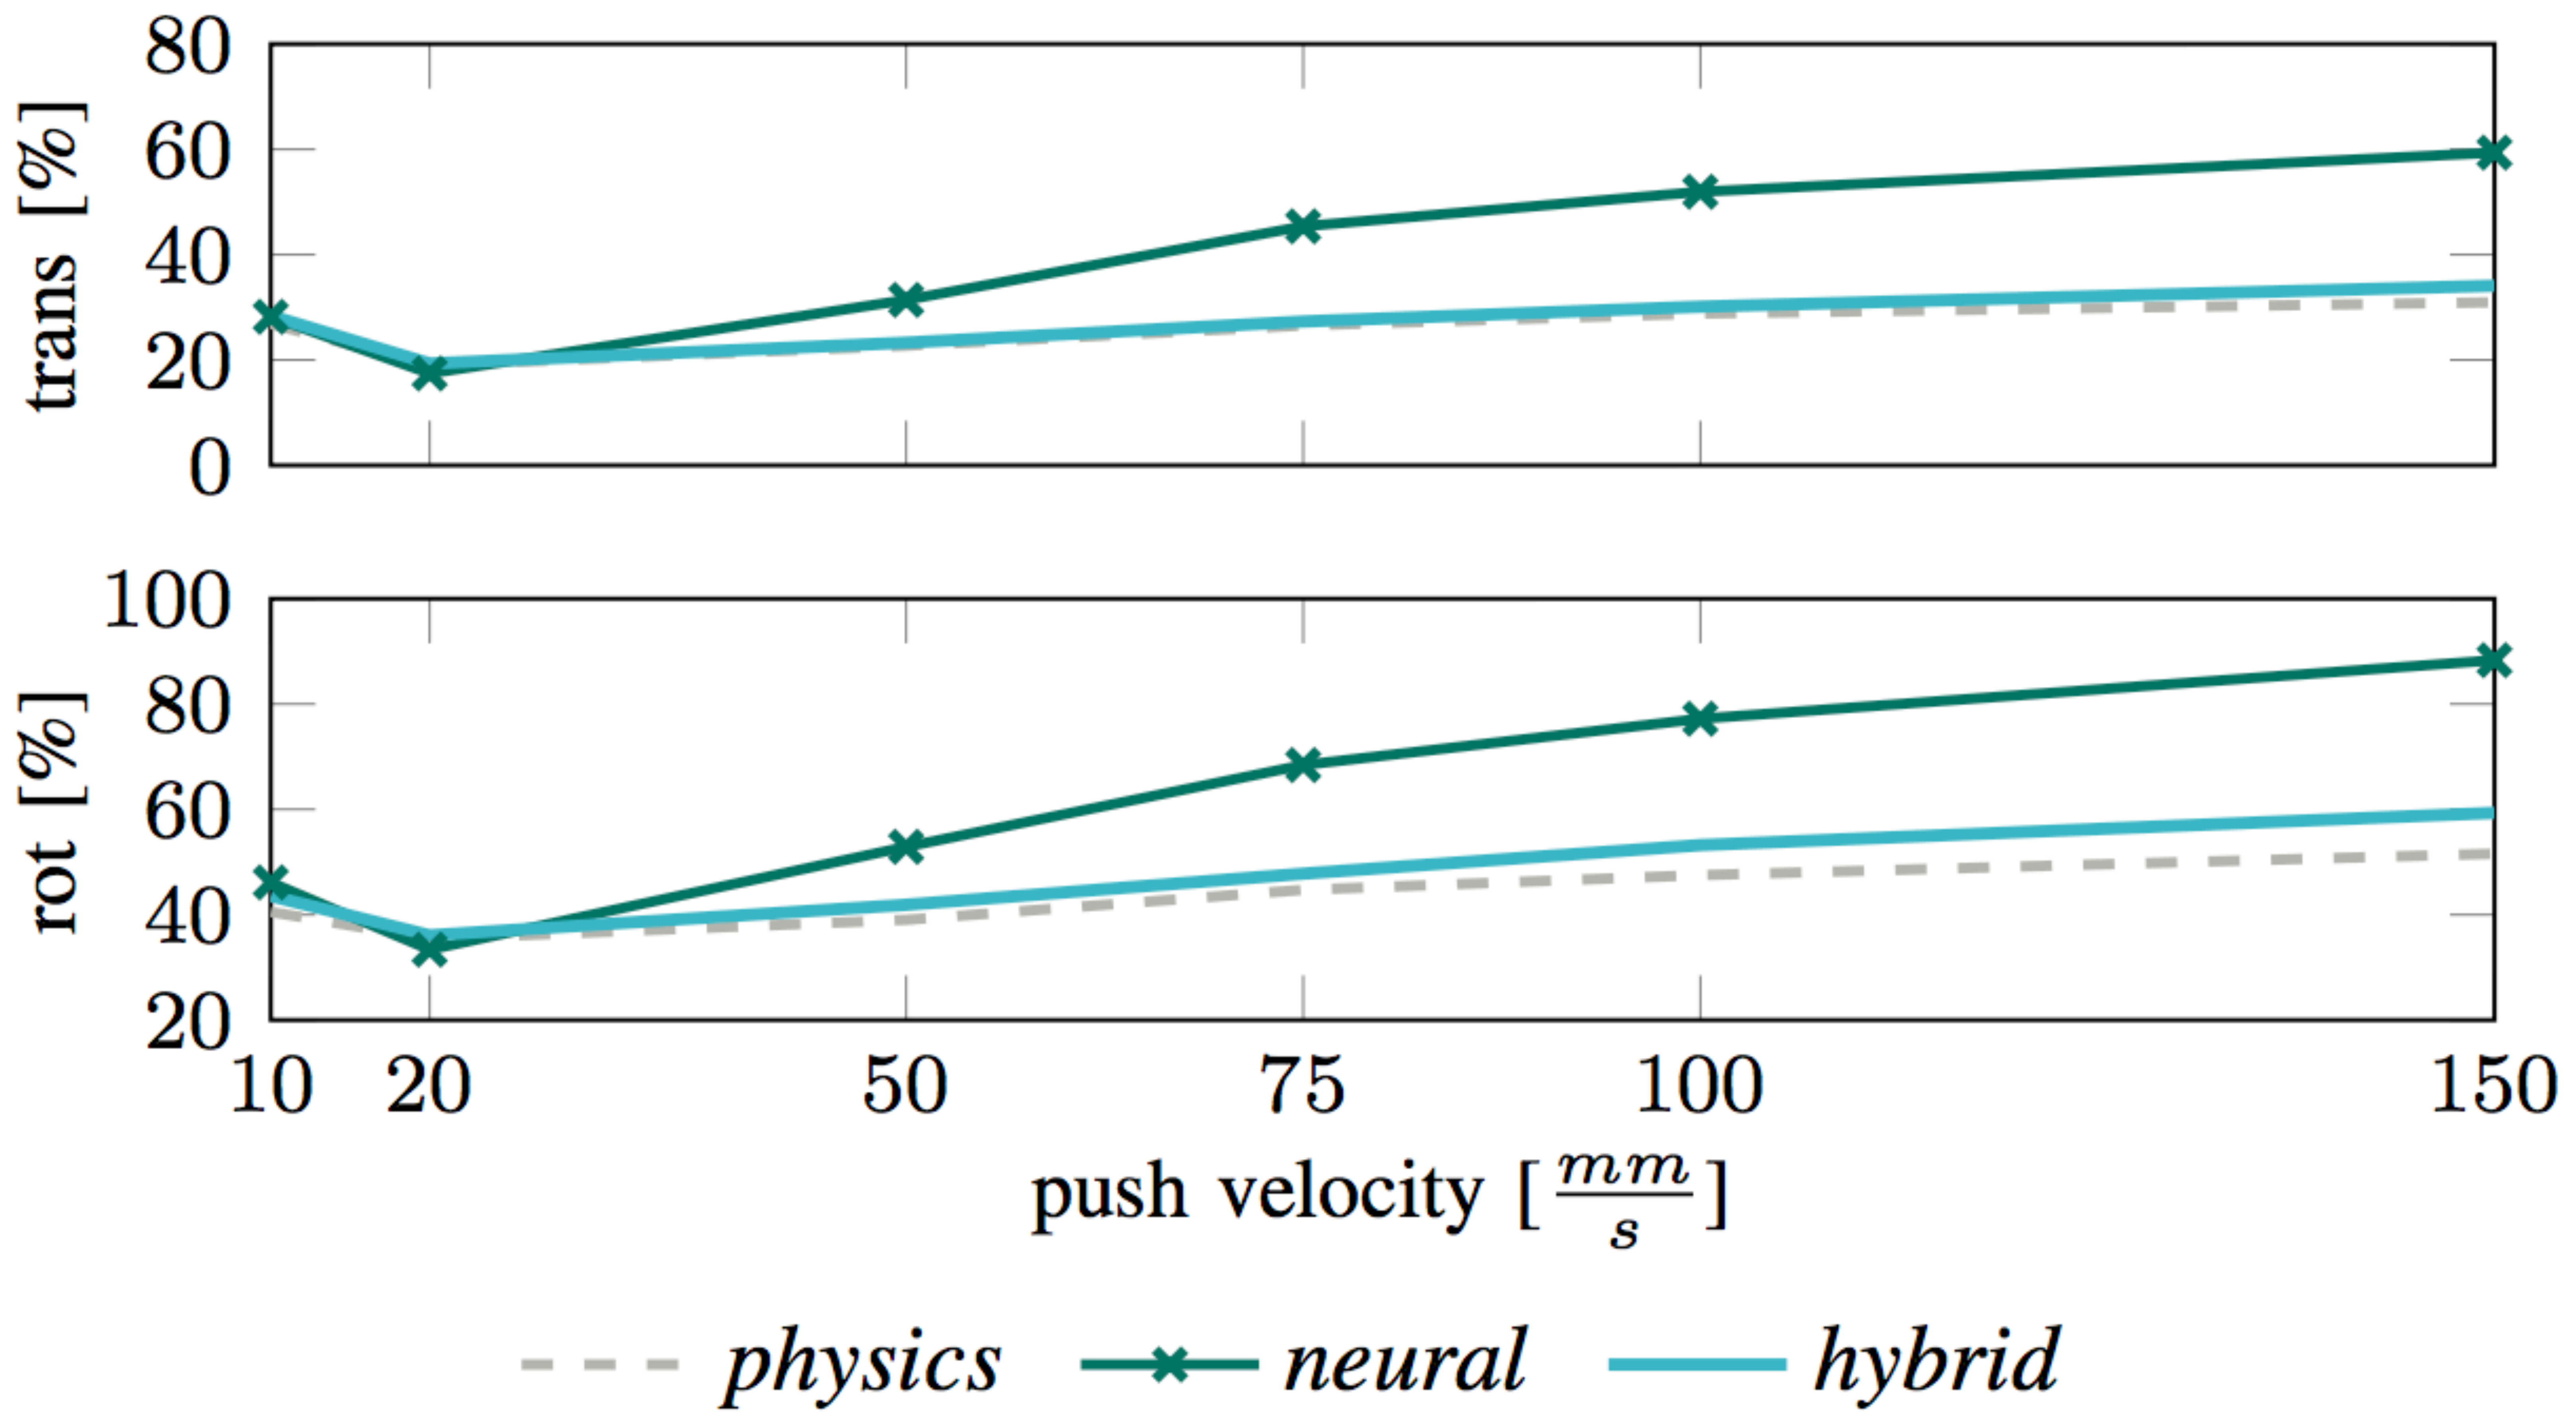

Generalization to new push velocities

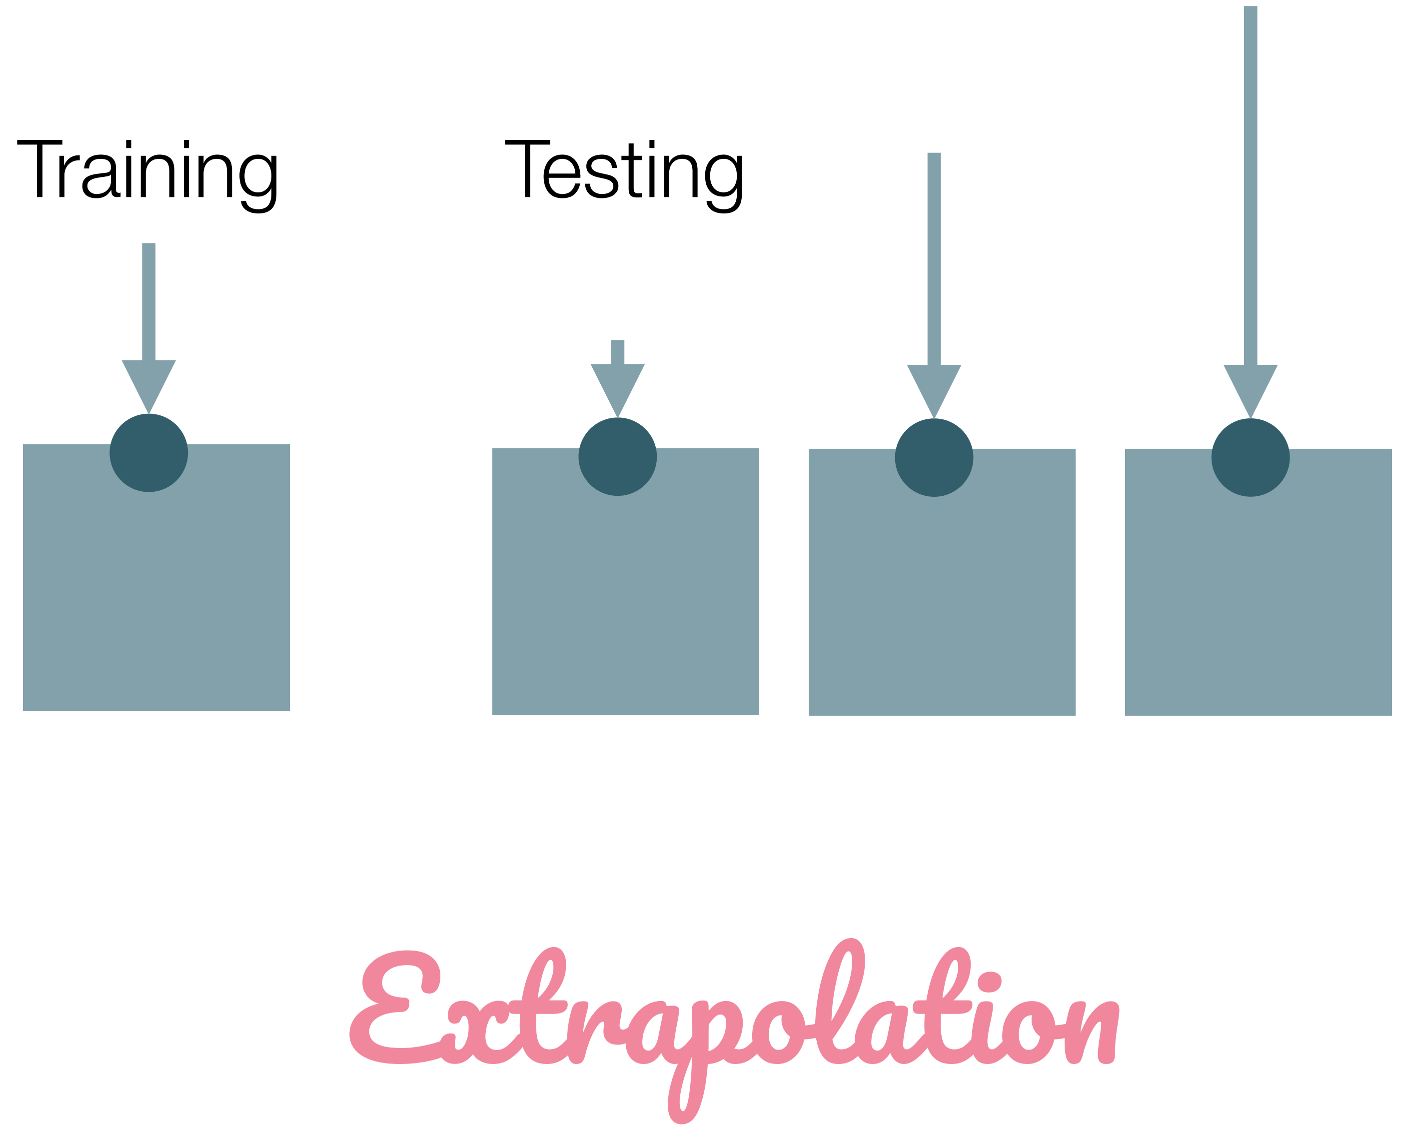

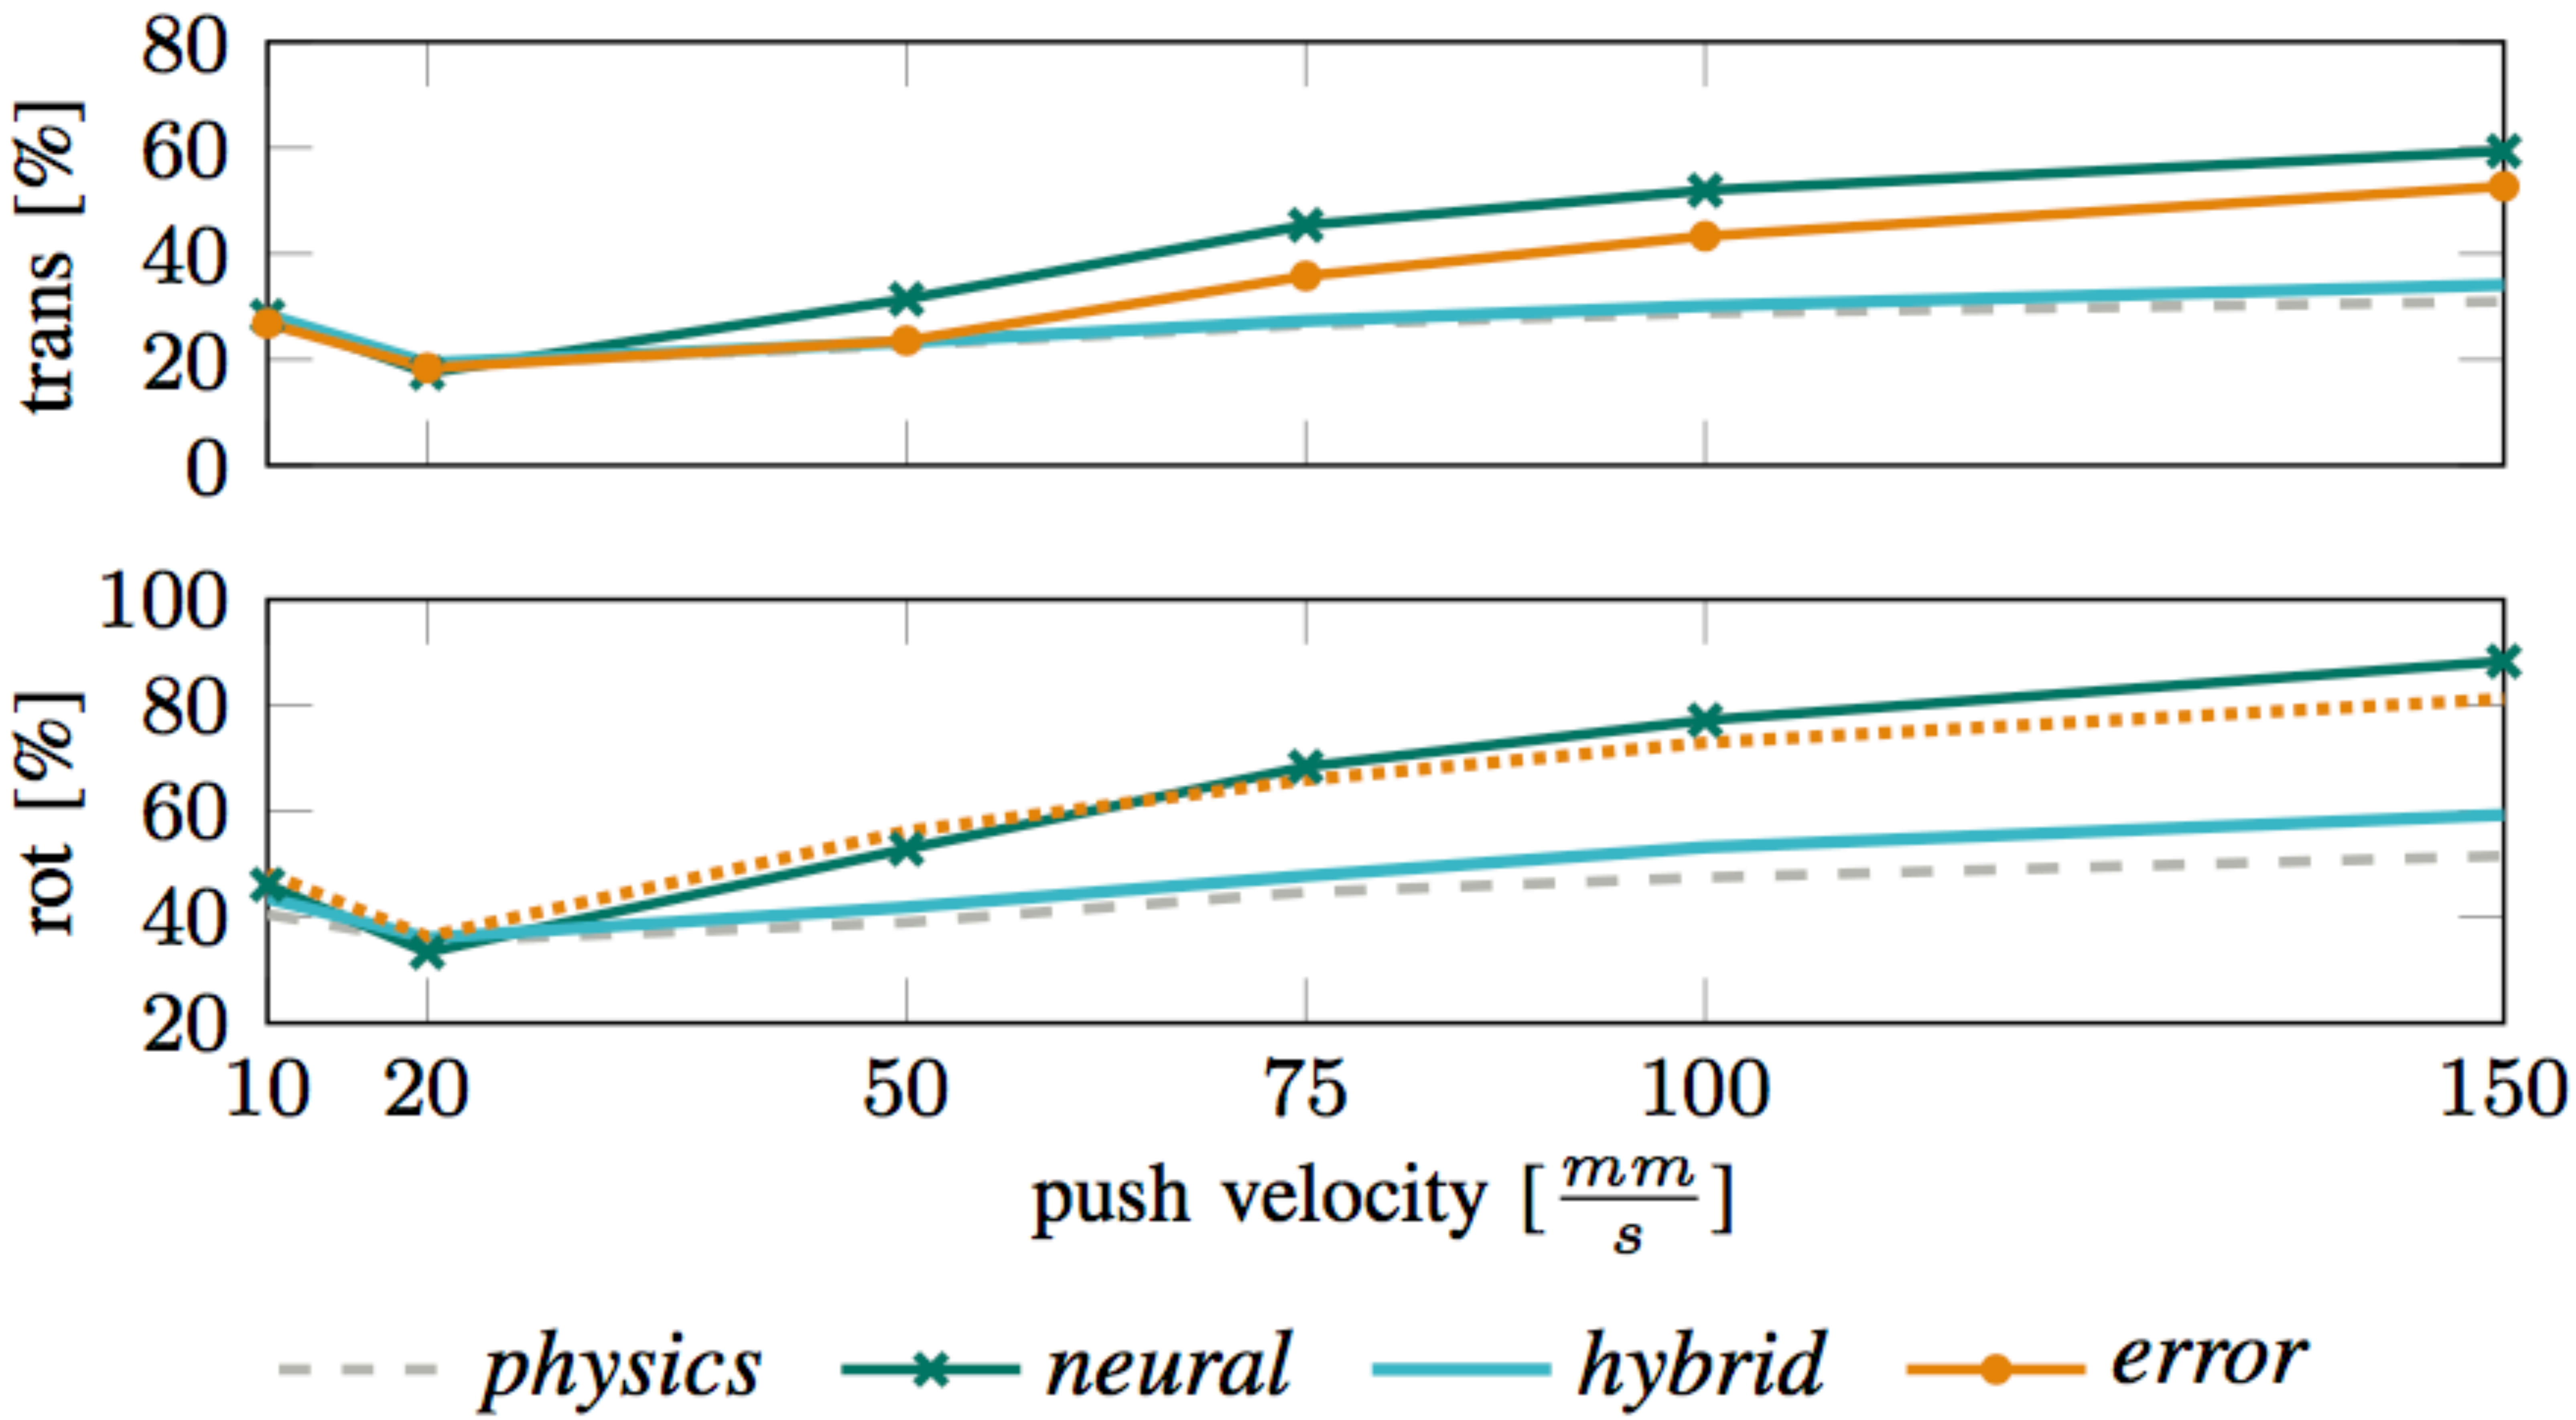

Action-Conditional Flow

current frame (with next action) | predicted optical flow | ground-truth optical flow

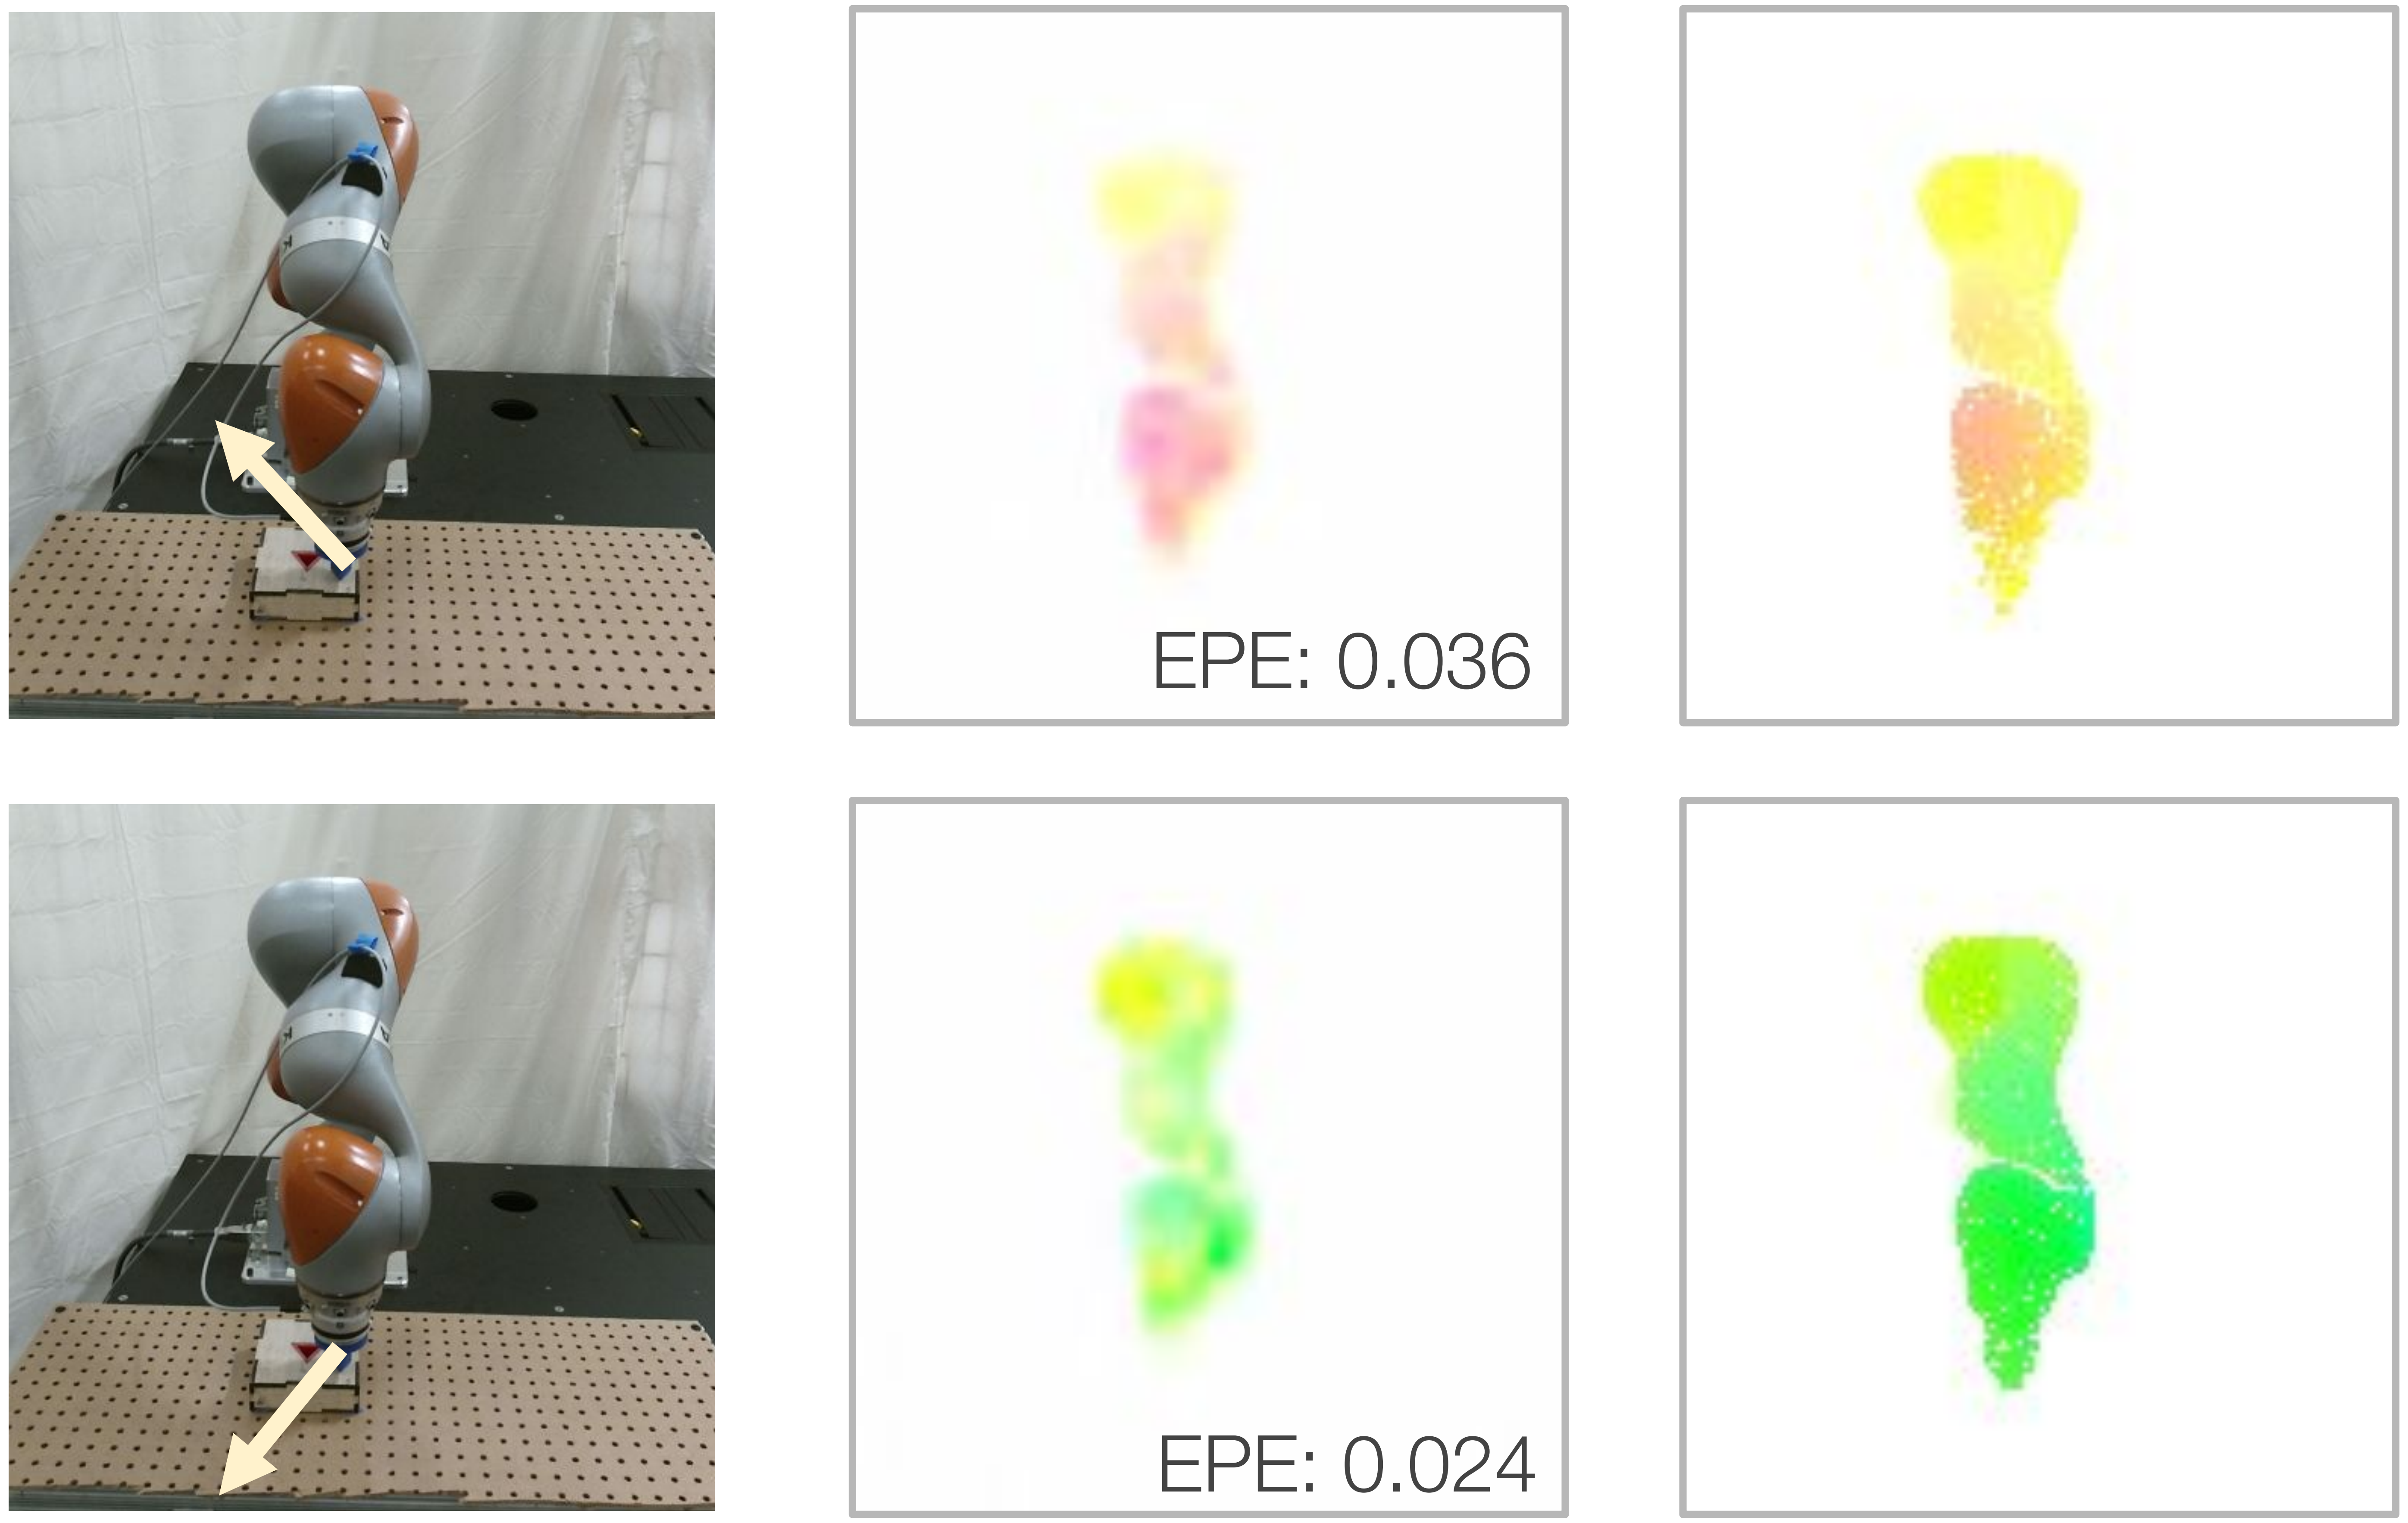

Controller Architecture