LEARNING TO GRASP USING SIMULATION

Yunfei Bai
yunfeibai@google.com

Collaboration between X (Formally Google[X]), Robotics at Google (Research), and DeepMind

Jun 23, 2019

Outline

- Learning Vision Based Grasping
  - Self-Supervised Learning
  - Deep Reinforcement Learning
- Improving Data Efficiency
  - With Sim-to-Real
  - With Sim-to-Sim
- Learning New Tasks
- Learning New Object Representation

1

Learning Vision Based Grasping with Self-Supervised Learning

"Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection",

Sergey Levine, Peter Pastor, Alex Krizhevsky, Julian Ibarz, Deirdre Quillen

Goal: Learn to grasp arbitrary objects

- RGB monocular camera input
- Camera positioned “over the shoulder”
- Poor / non-existent camera calibration

Assumptions:
- Overhead grasps

Reward function:
- Gripper angle
- Image subtraction

1,100 objects used for training

Grasp Success Prediction Model

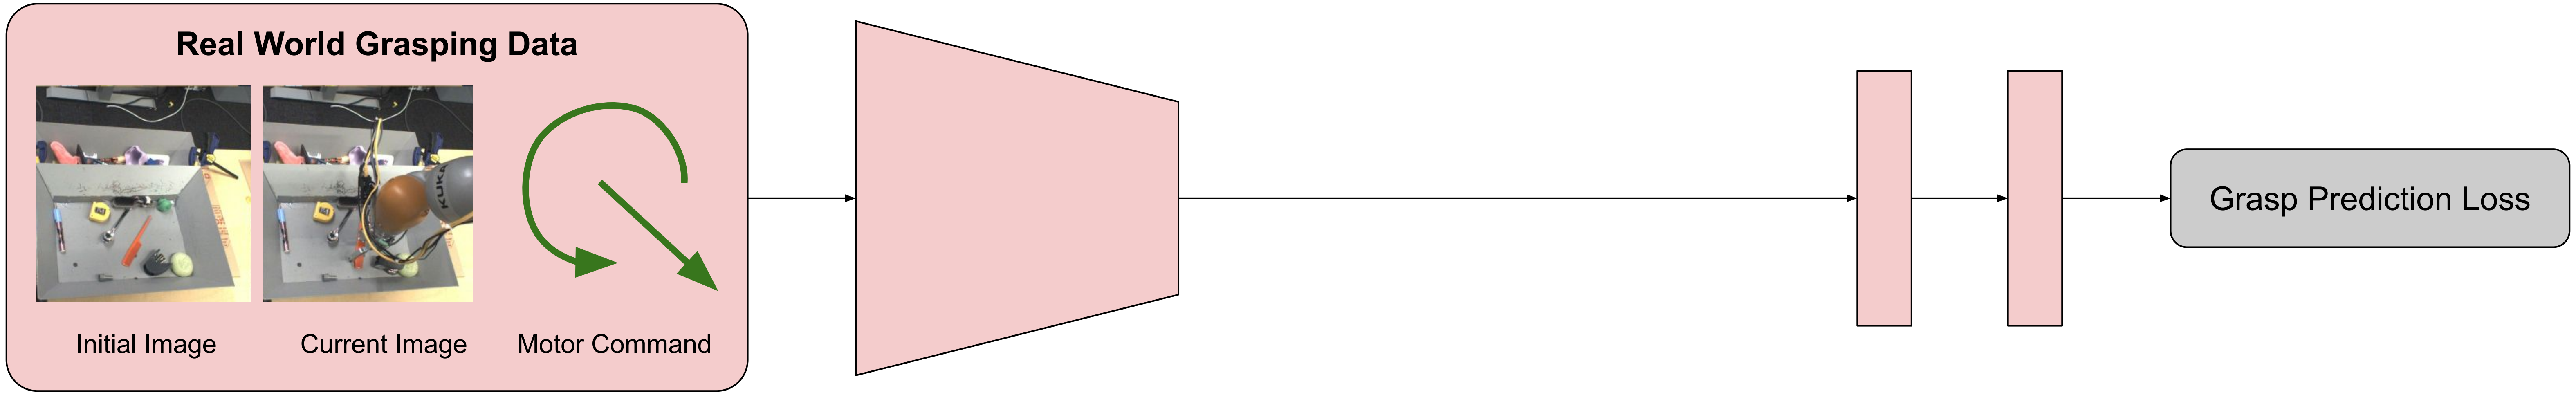

- Inference using Cross-Entropy Method (CEM)
- Replan every 300 to 500ms
- Hand-eye coordination

Grasp Success Rates with Increasing Amounts of Data

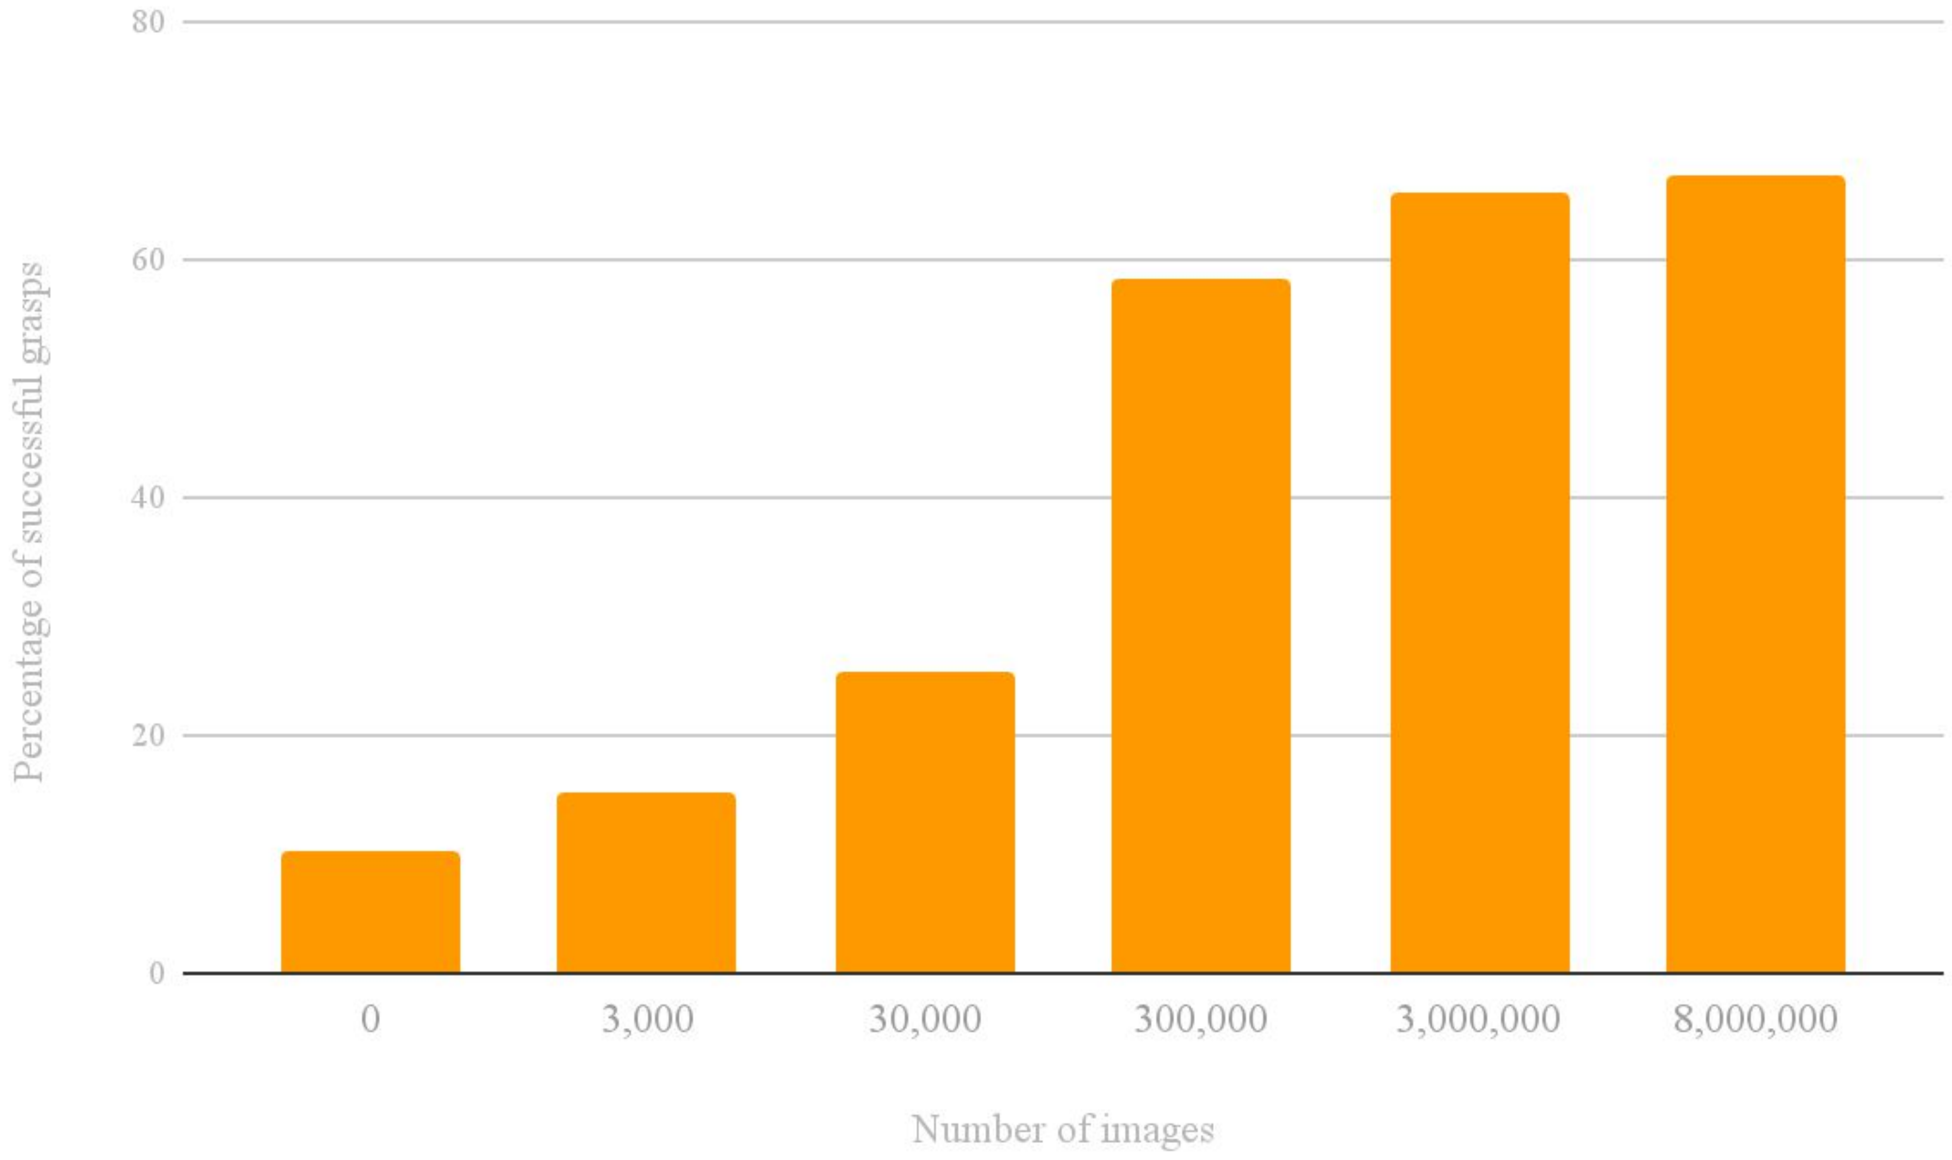

Two Directions for Improvement

- Break the ceiling of the grasp success rate to make it close to **100%**.
- Improve real world **data efficiency** and reduce the data collection time.

2

Learning Vision Based Grasping with Deep Reinforcement Learning

“QT-Opt: Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation”,

Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly,

Mrinal Kalakrishnan, Vincent Vanhoucke, Sergey Levine

Q Learning for Grasping

- Supervised learning
  - Optimize for the next step
- Reinforcement learning
  - Predict a few steps ahead

Q Learning for Grasping: QT-OPT

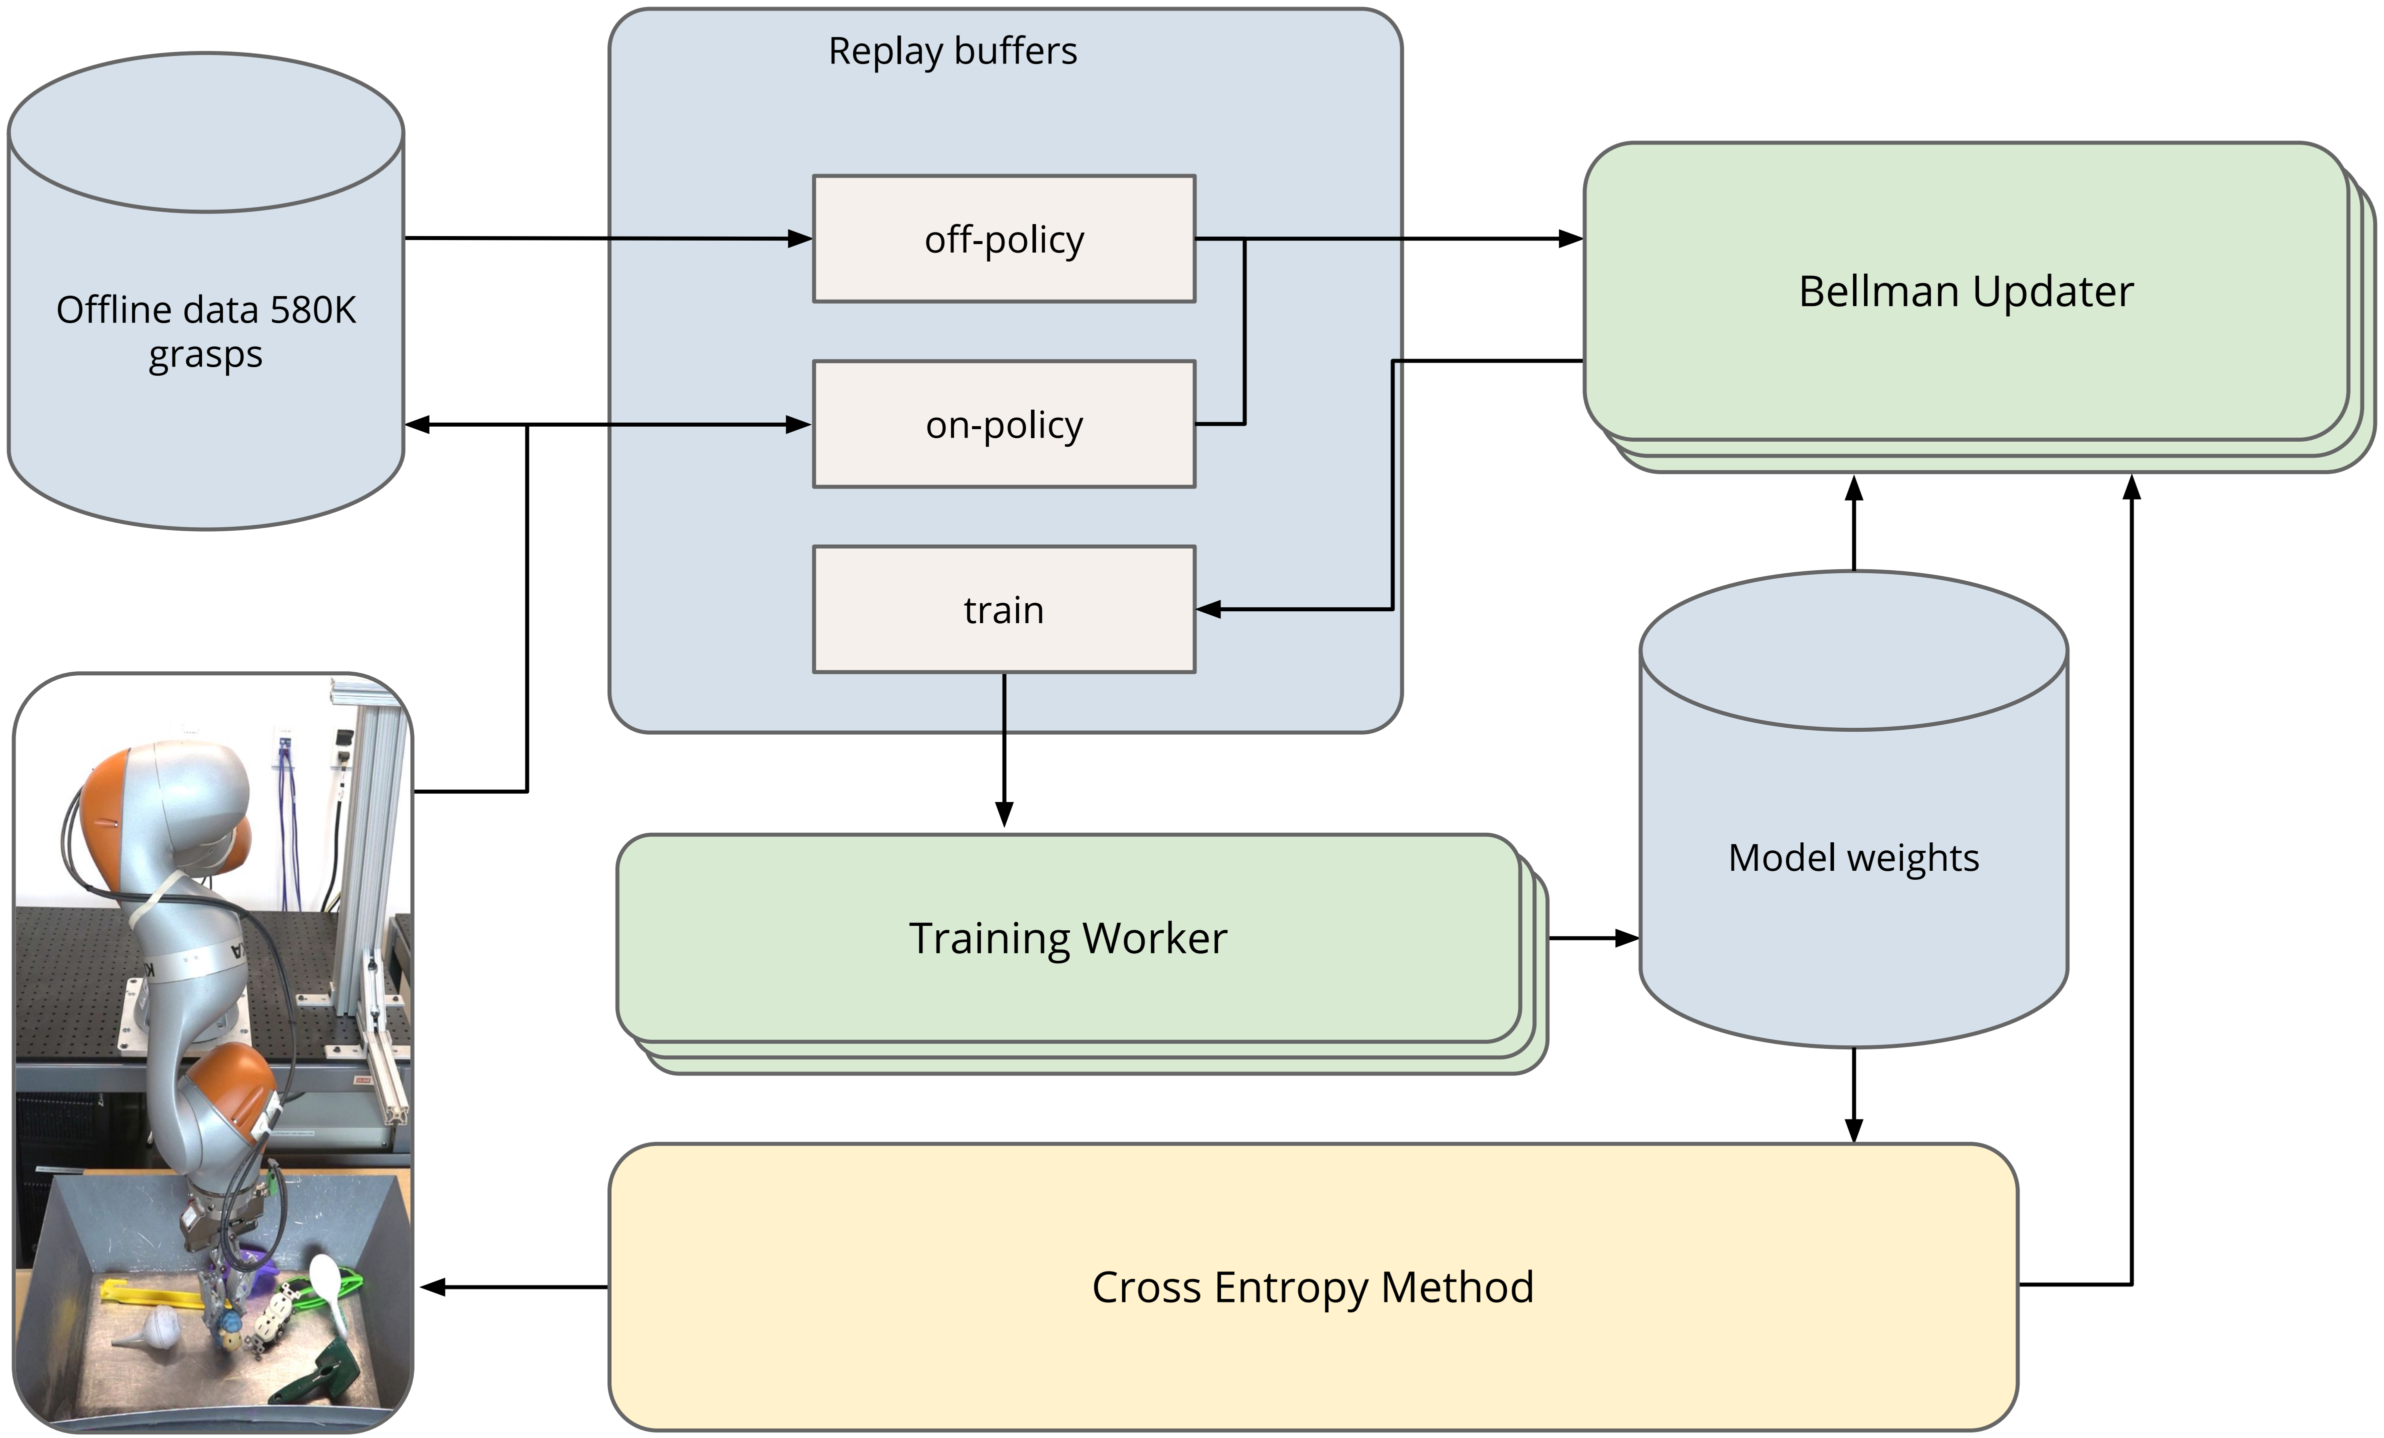

Q Network: Action Space Includes Open/Close Gripper and Termination

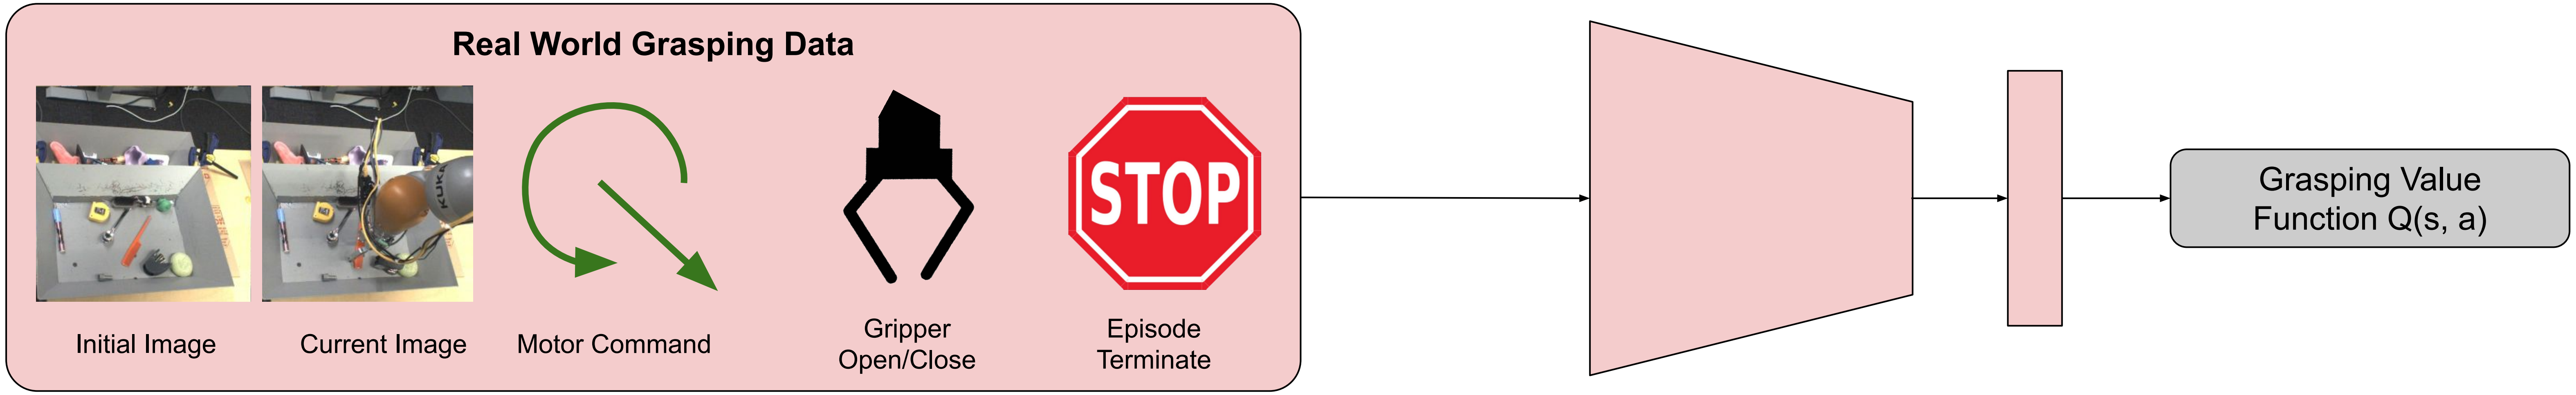


Levine et al. '16
QT-Opt (our method)
20x

QT-Opt achieves 96% grasp success, with higher data efficiency

3

Leveraging Sim-to-Real for Data Efficiency

"Using Simulation and Domain Adaptation to Improve Efficiency of Deep Robotic Grasping",

Konstantinos Bousmalis, Alex Irpan, Paul Wohlhart, Yunfei Bai, Matthew Kelcey, Mrinal Kalakrishnan, Laura Downs,

Julian Ibarz, Peter Pastor, Kurt Konolige, Sergey Levine, Vincent Vanhoucke

Motivation for Using Simulation

- 608,000 real-world grasps to achieve best performance
  - 7 KUKA robots running for 2-3 months
- Use simulation!
  - Easy to parallelize, reset, safe exploration, access to ground-truth for exploration policy, ...

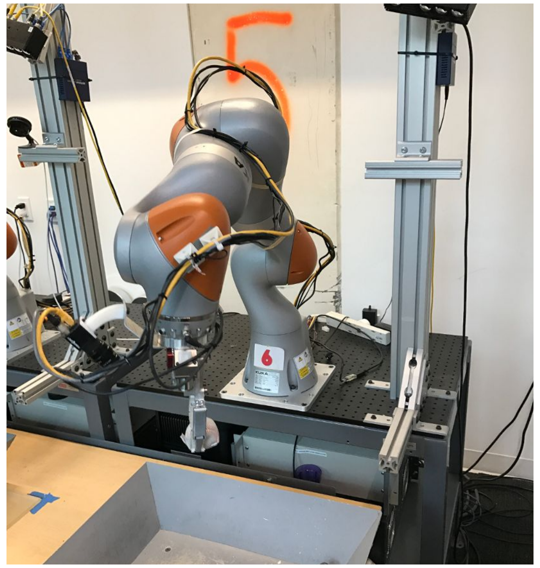

Real world

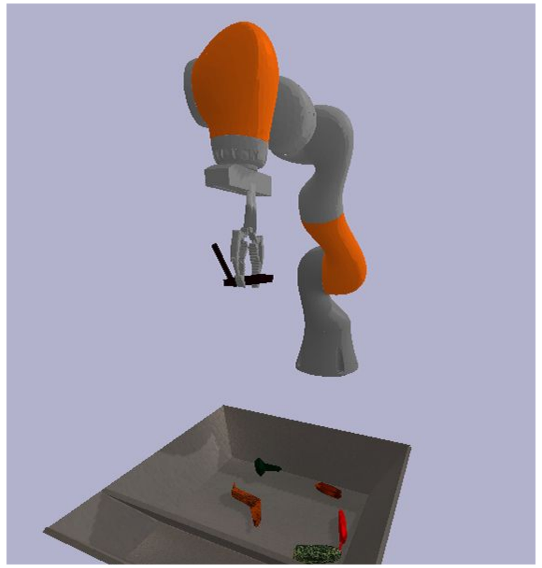

Simulation

Reality Gap

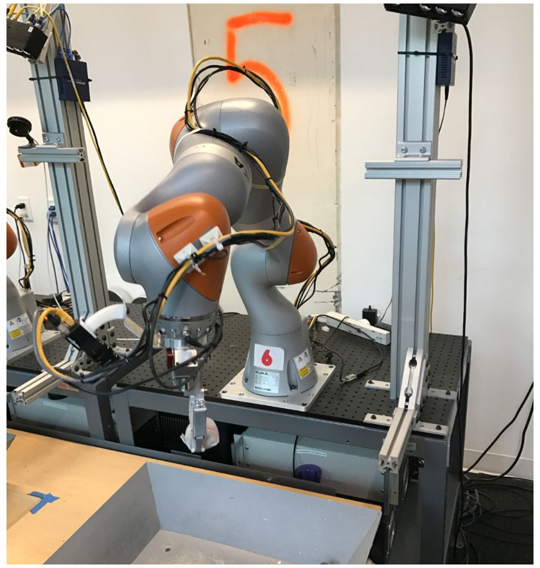

Real world

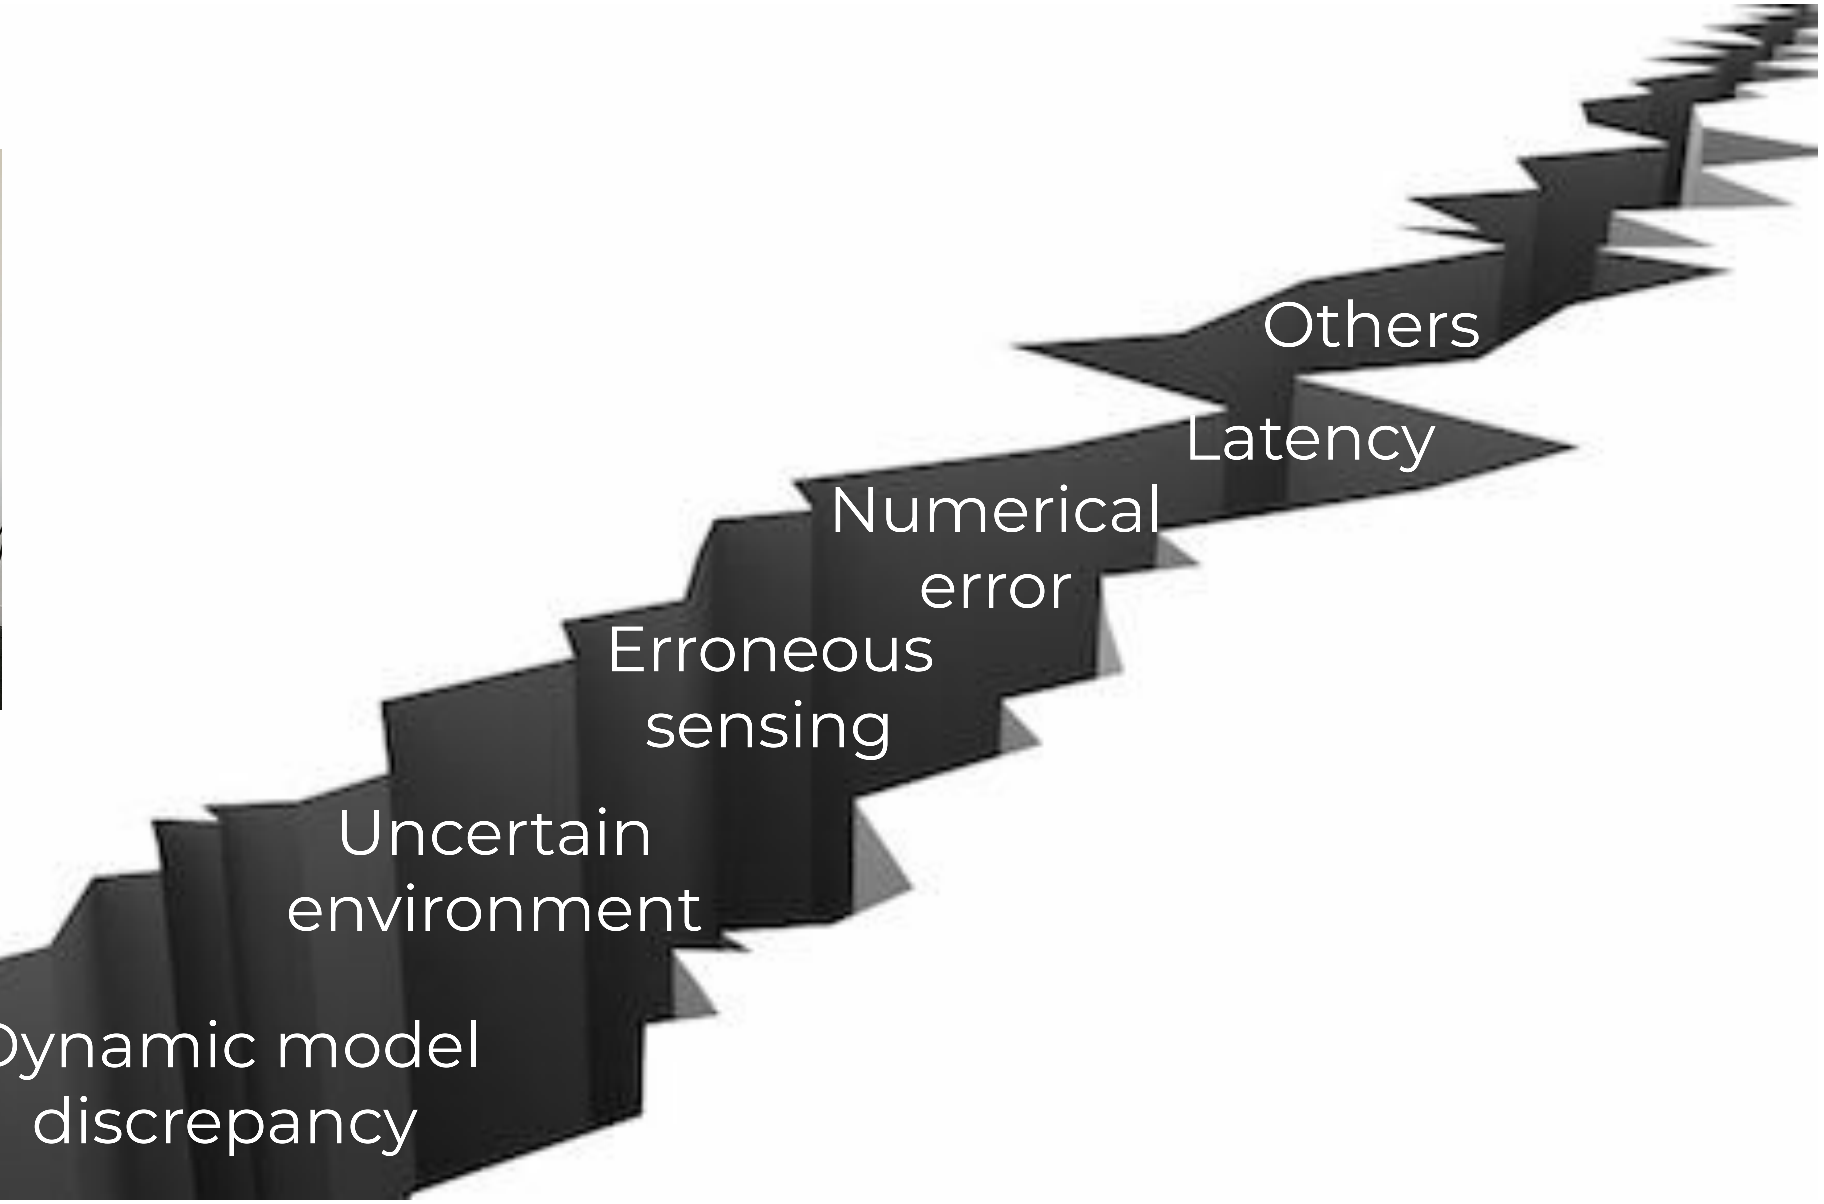

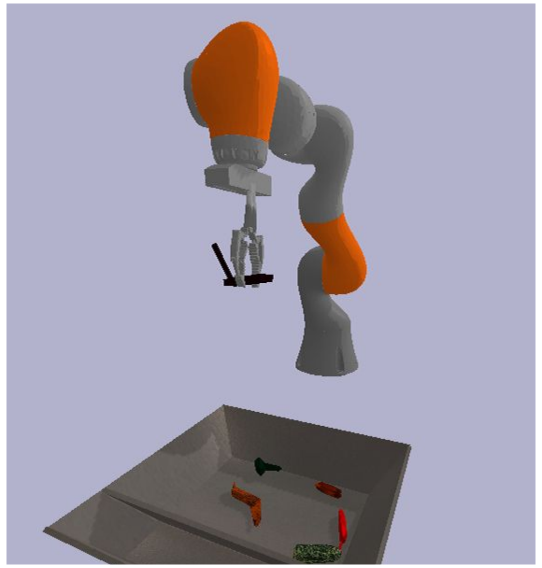

Simulation

Methods for Sim-to-Real Transfer

System Identification	Use statistical methods to build mathematical models of dynamical systems from measured data.
Domain Randomization	Vary texture, background, lighting, color, object shape, and dynamics in simulation.
Feature-level Domain Adaptation	Train features to be domain-invariant yet expressive, by using an adversarial loss.
Pixel-level Domain Adaptation	Train a generator network that converts simulated images to real images, by using an adversarial loss.

Sim-to-Real Transfer for Physics - System Identification

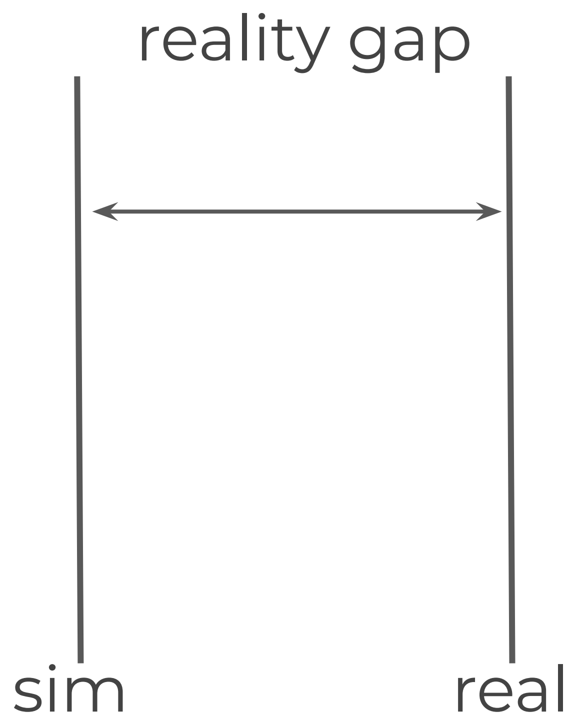

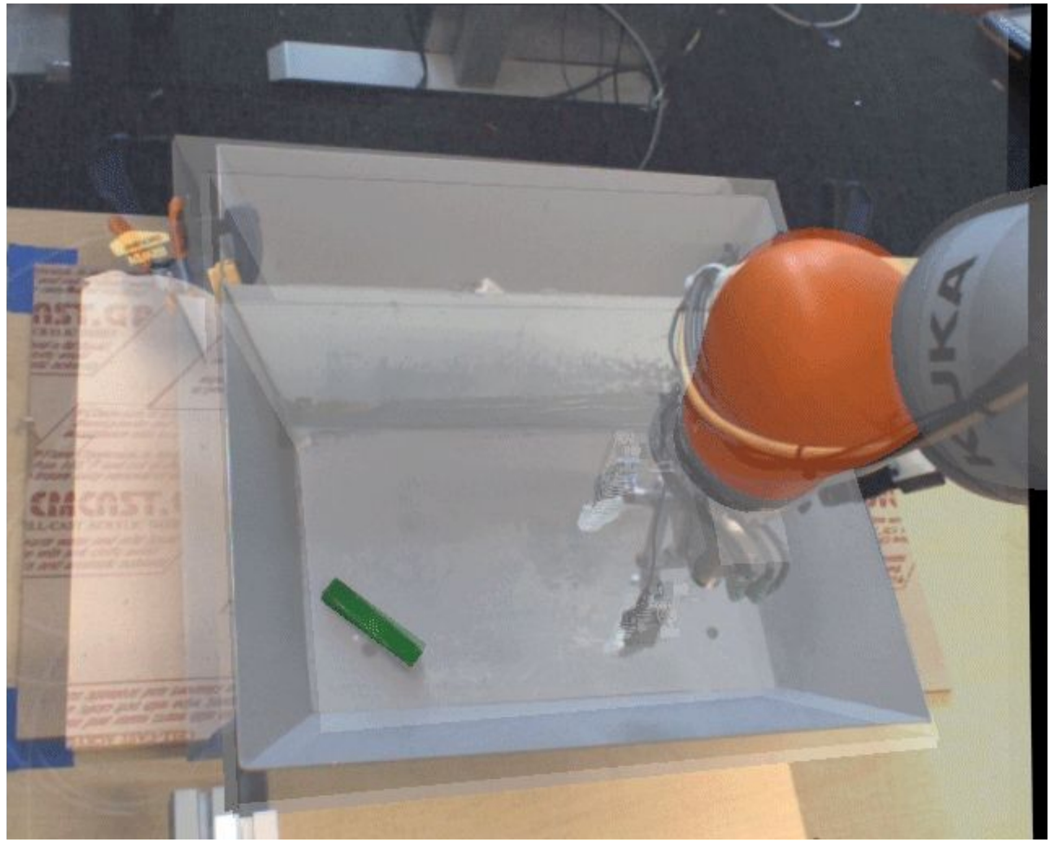

Real \ Sim	Successful grasp	Unsuccessful grasp
Successful grasp	855	56
Unsuccessful grasp	29	2080

Sim-to-Real Transfer for Physics - System Identification

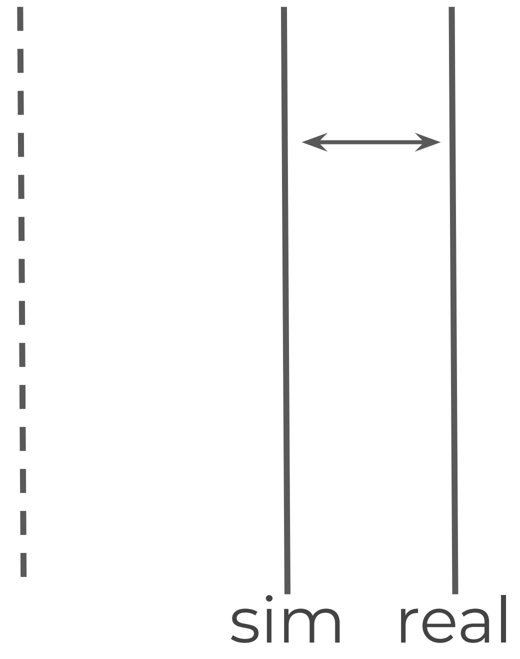

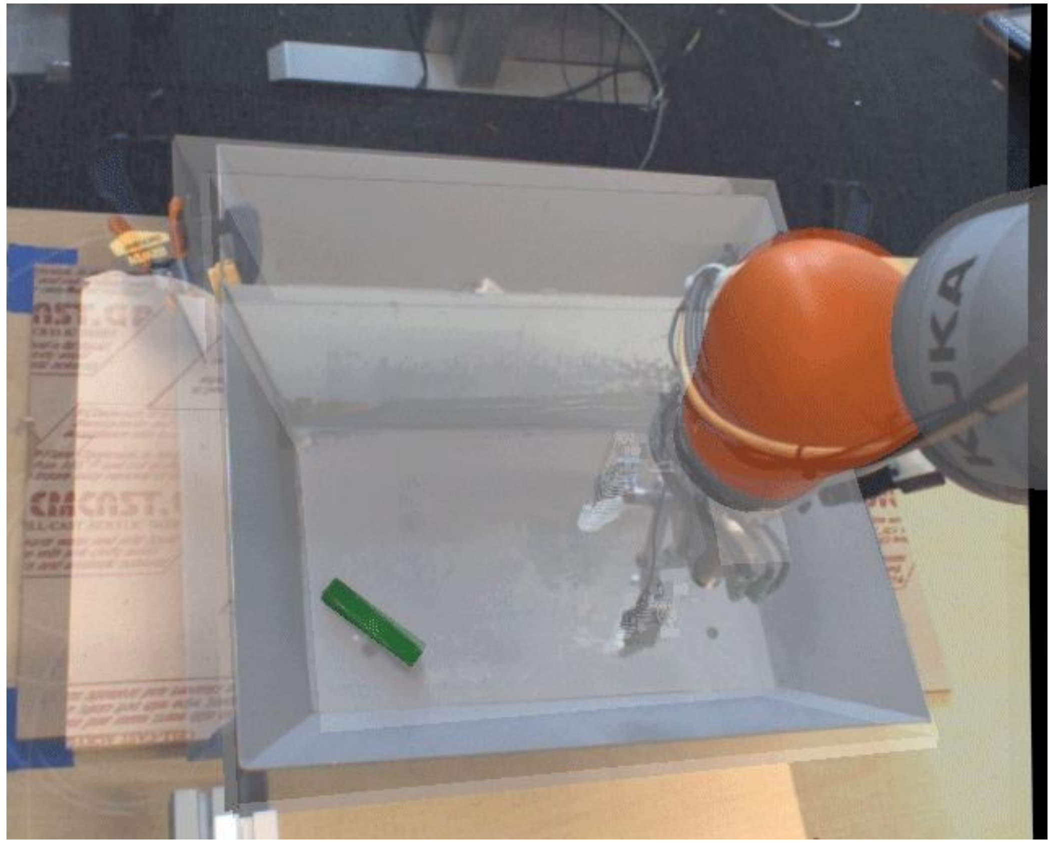

		Sim: Successful grasp	Sim: Unsuccessful grasp
Real	Successful grasp	855	56
Real	Unsuccessful grasp	29	2080

Sim-to-Real Transfer for Perception - Domain Randomization

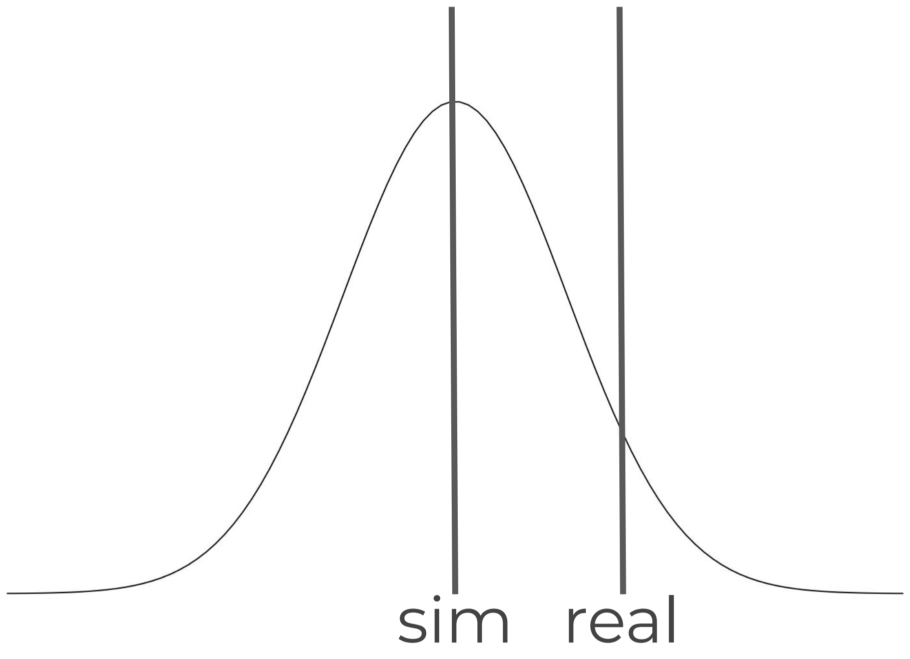

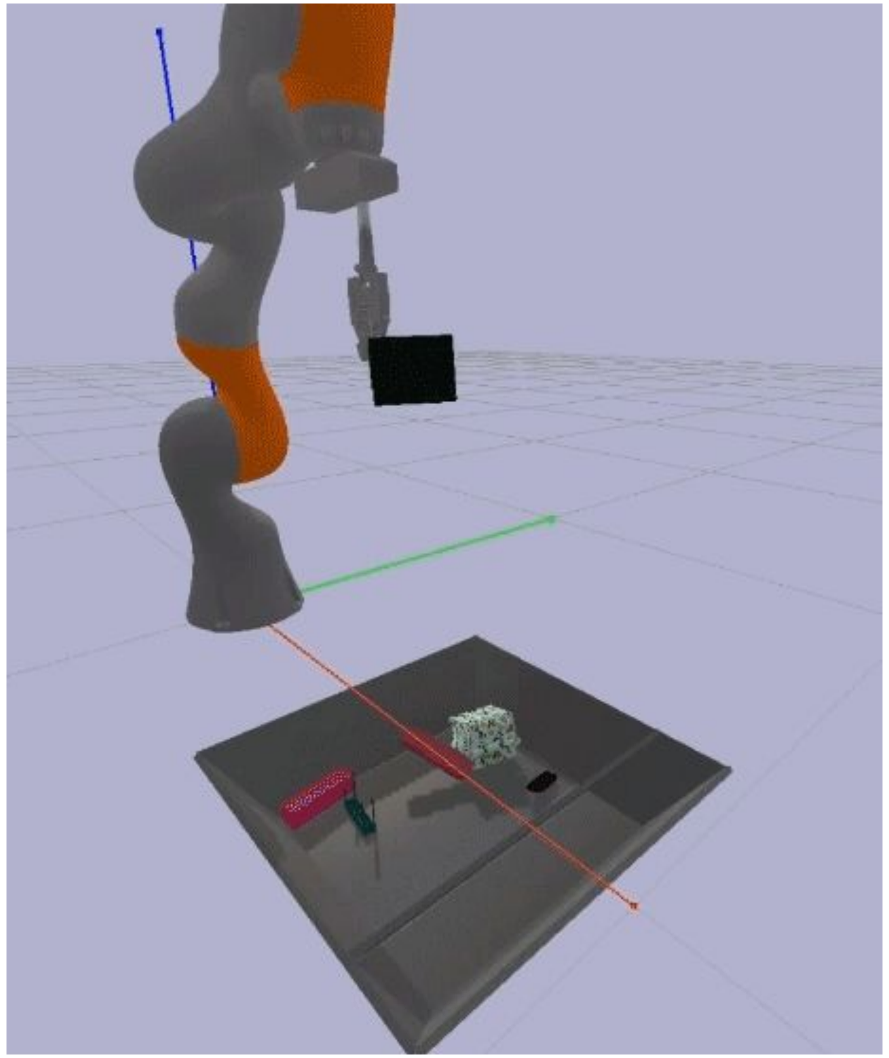

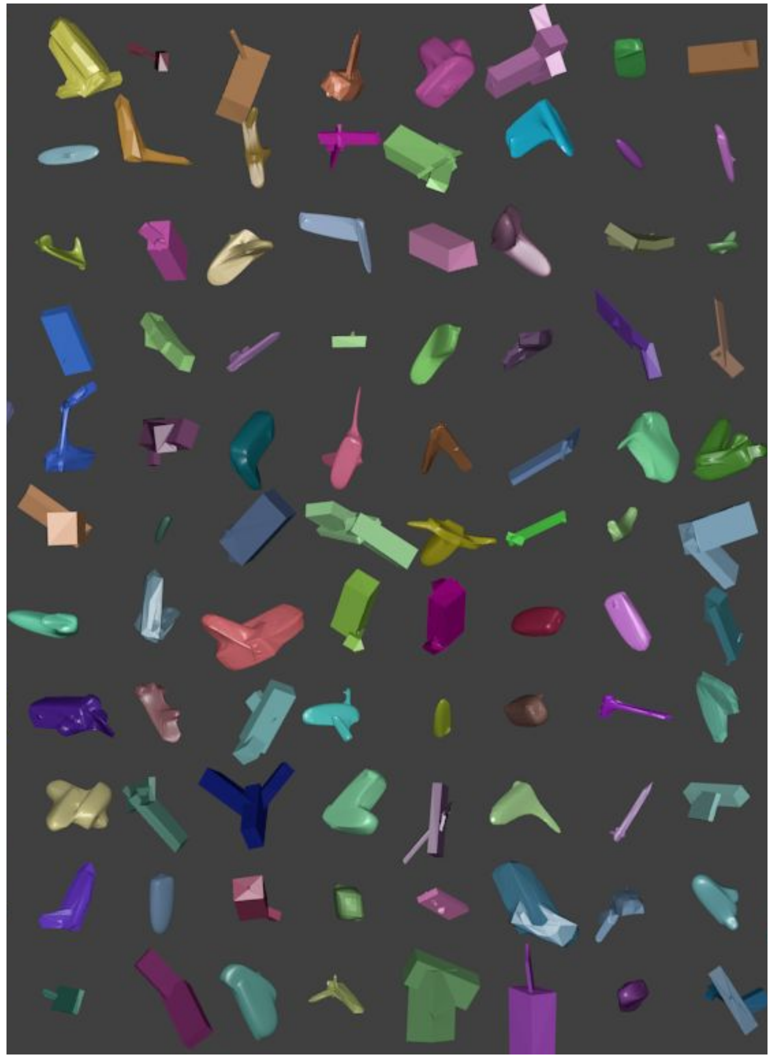

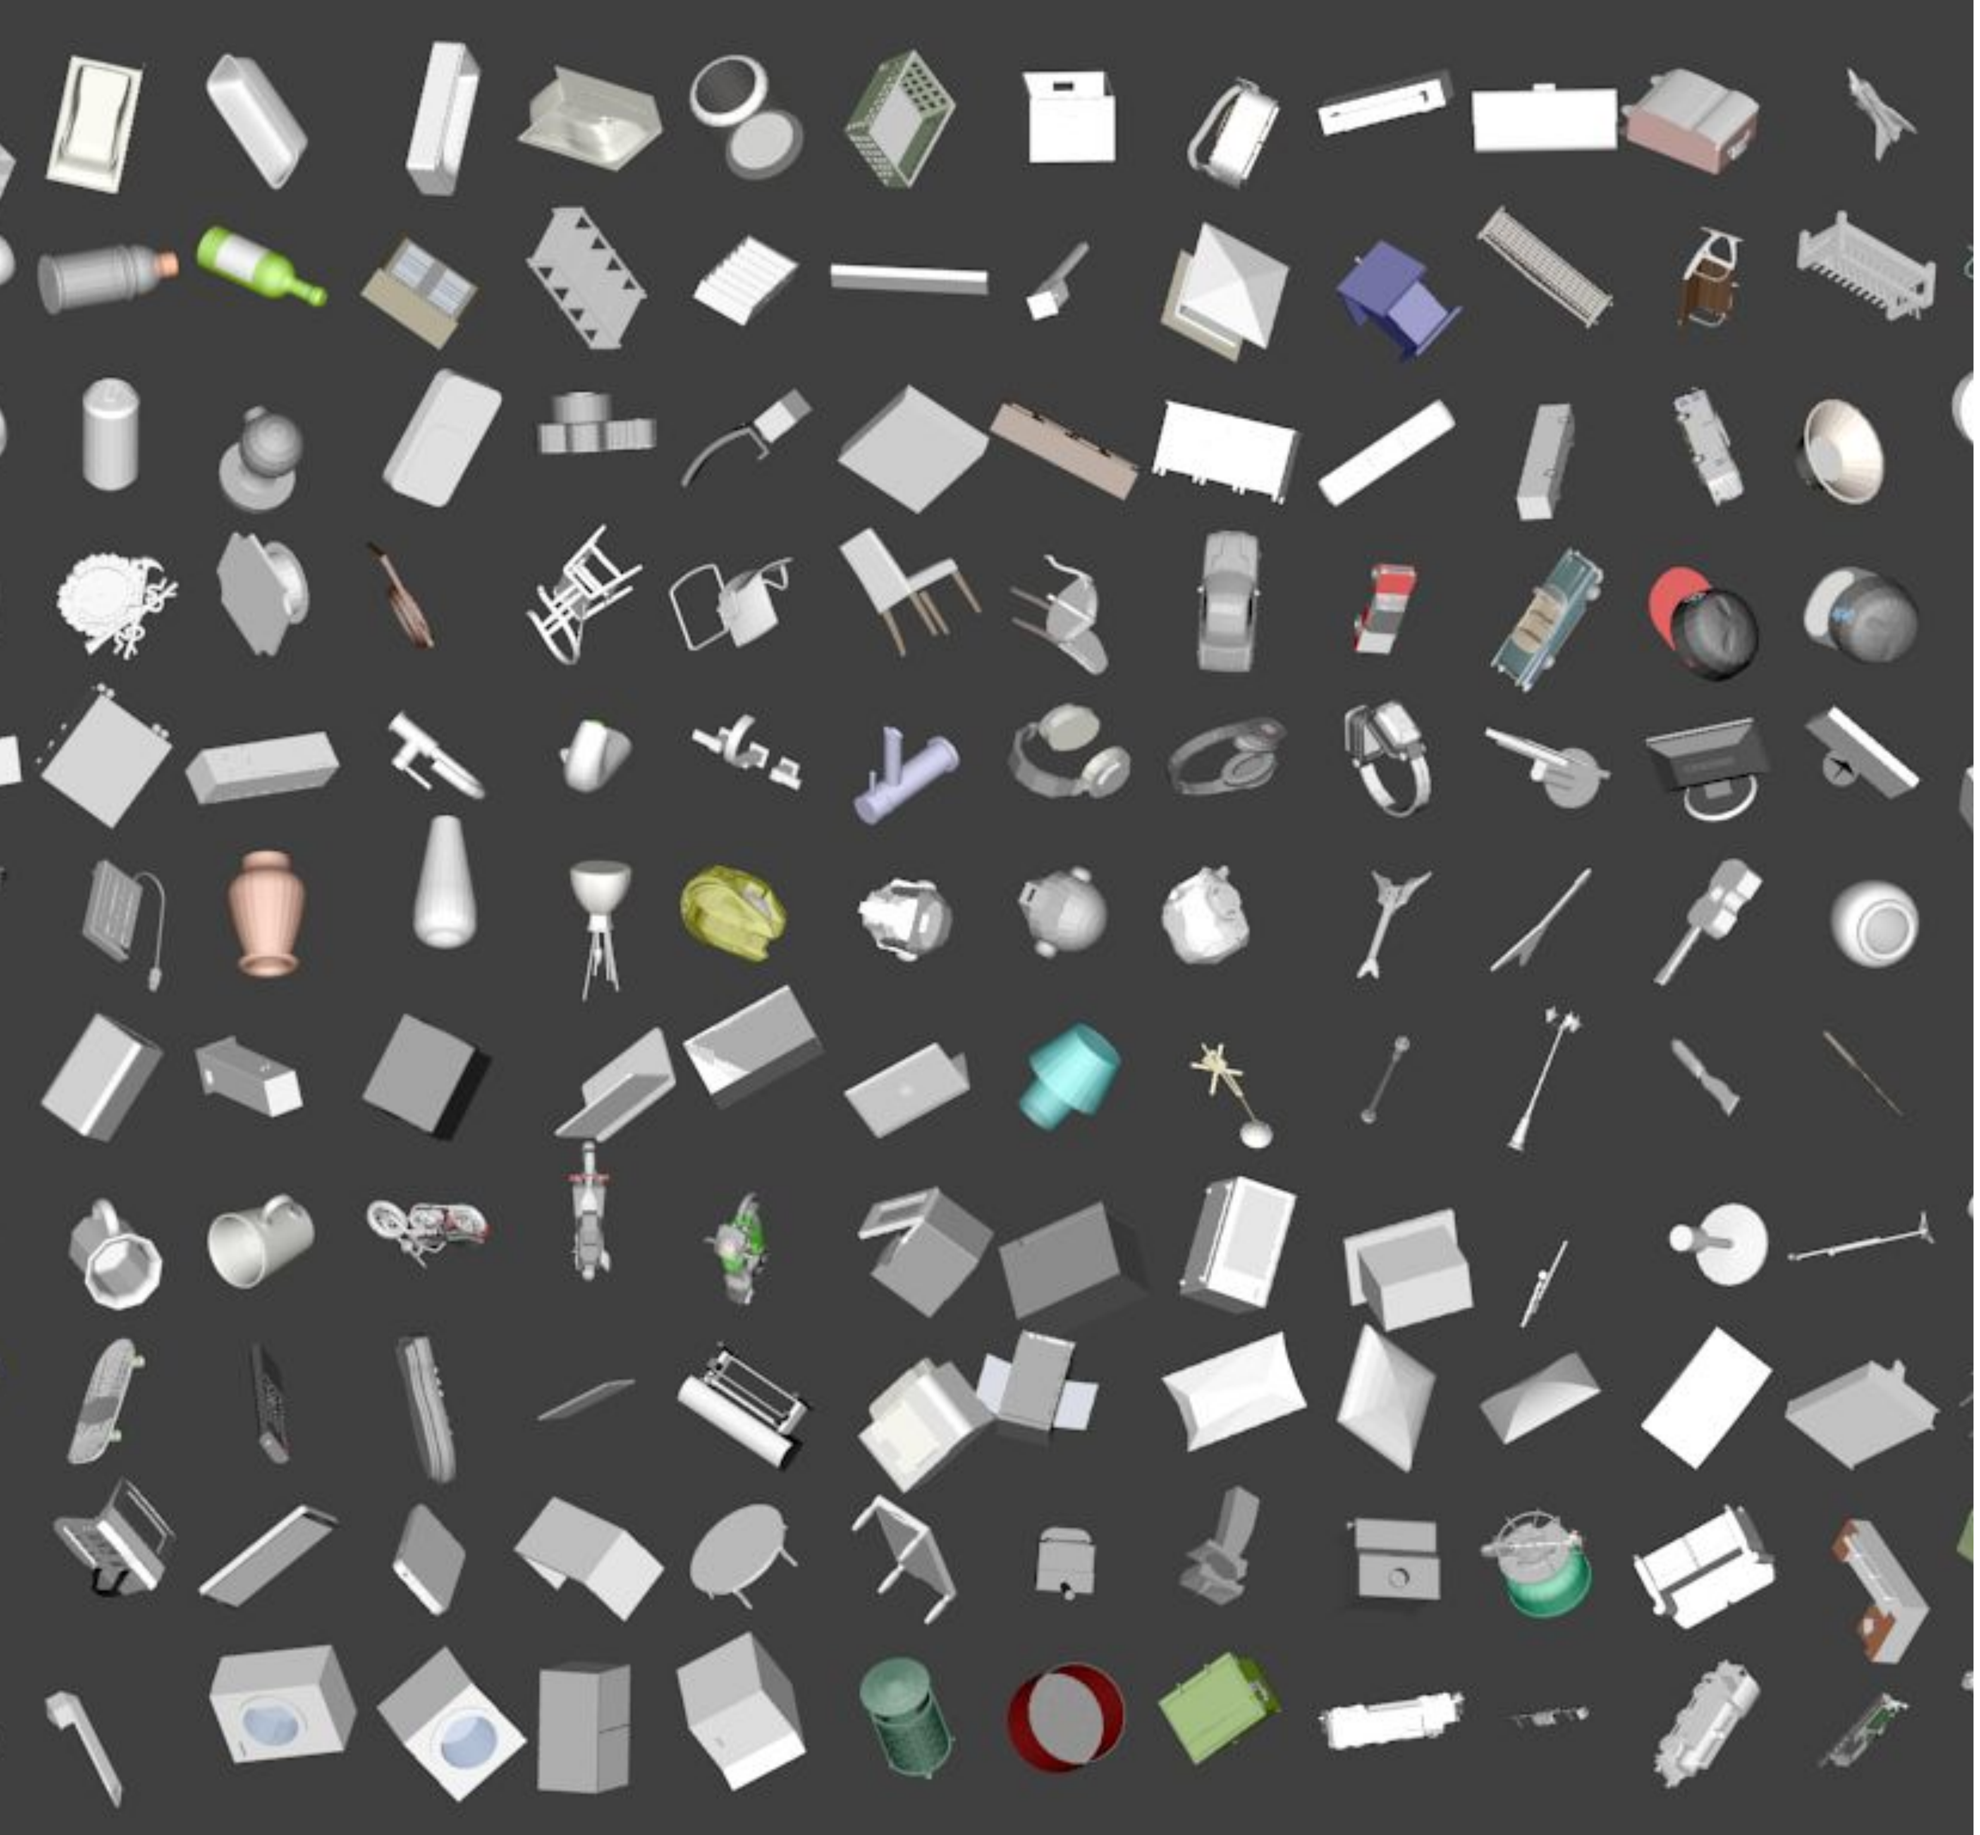

- 51,300 object models from ShapeNet.

Chang et al., CoRR 2015

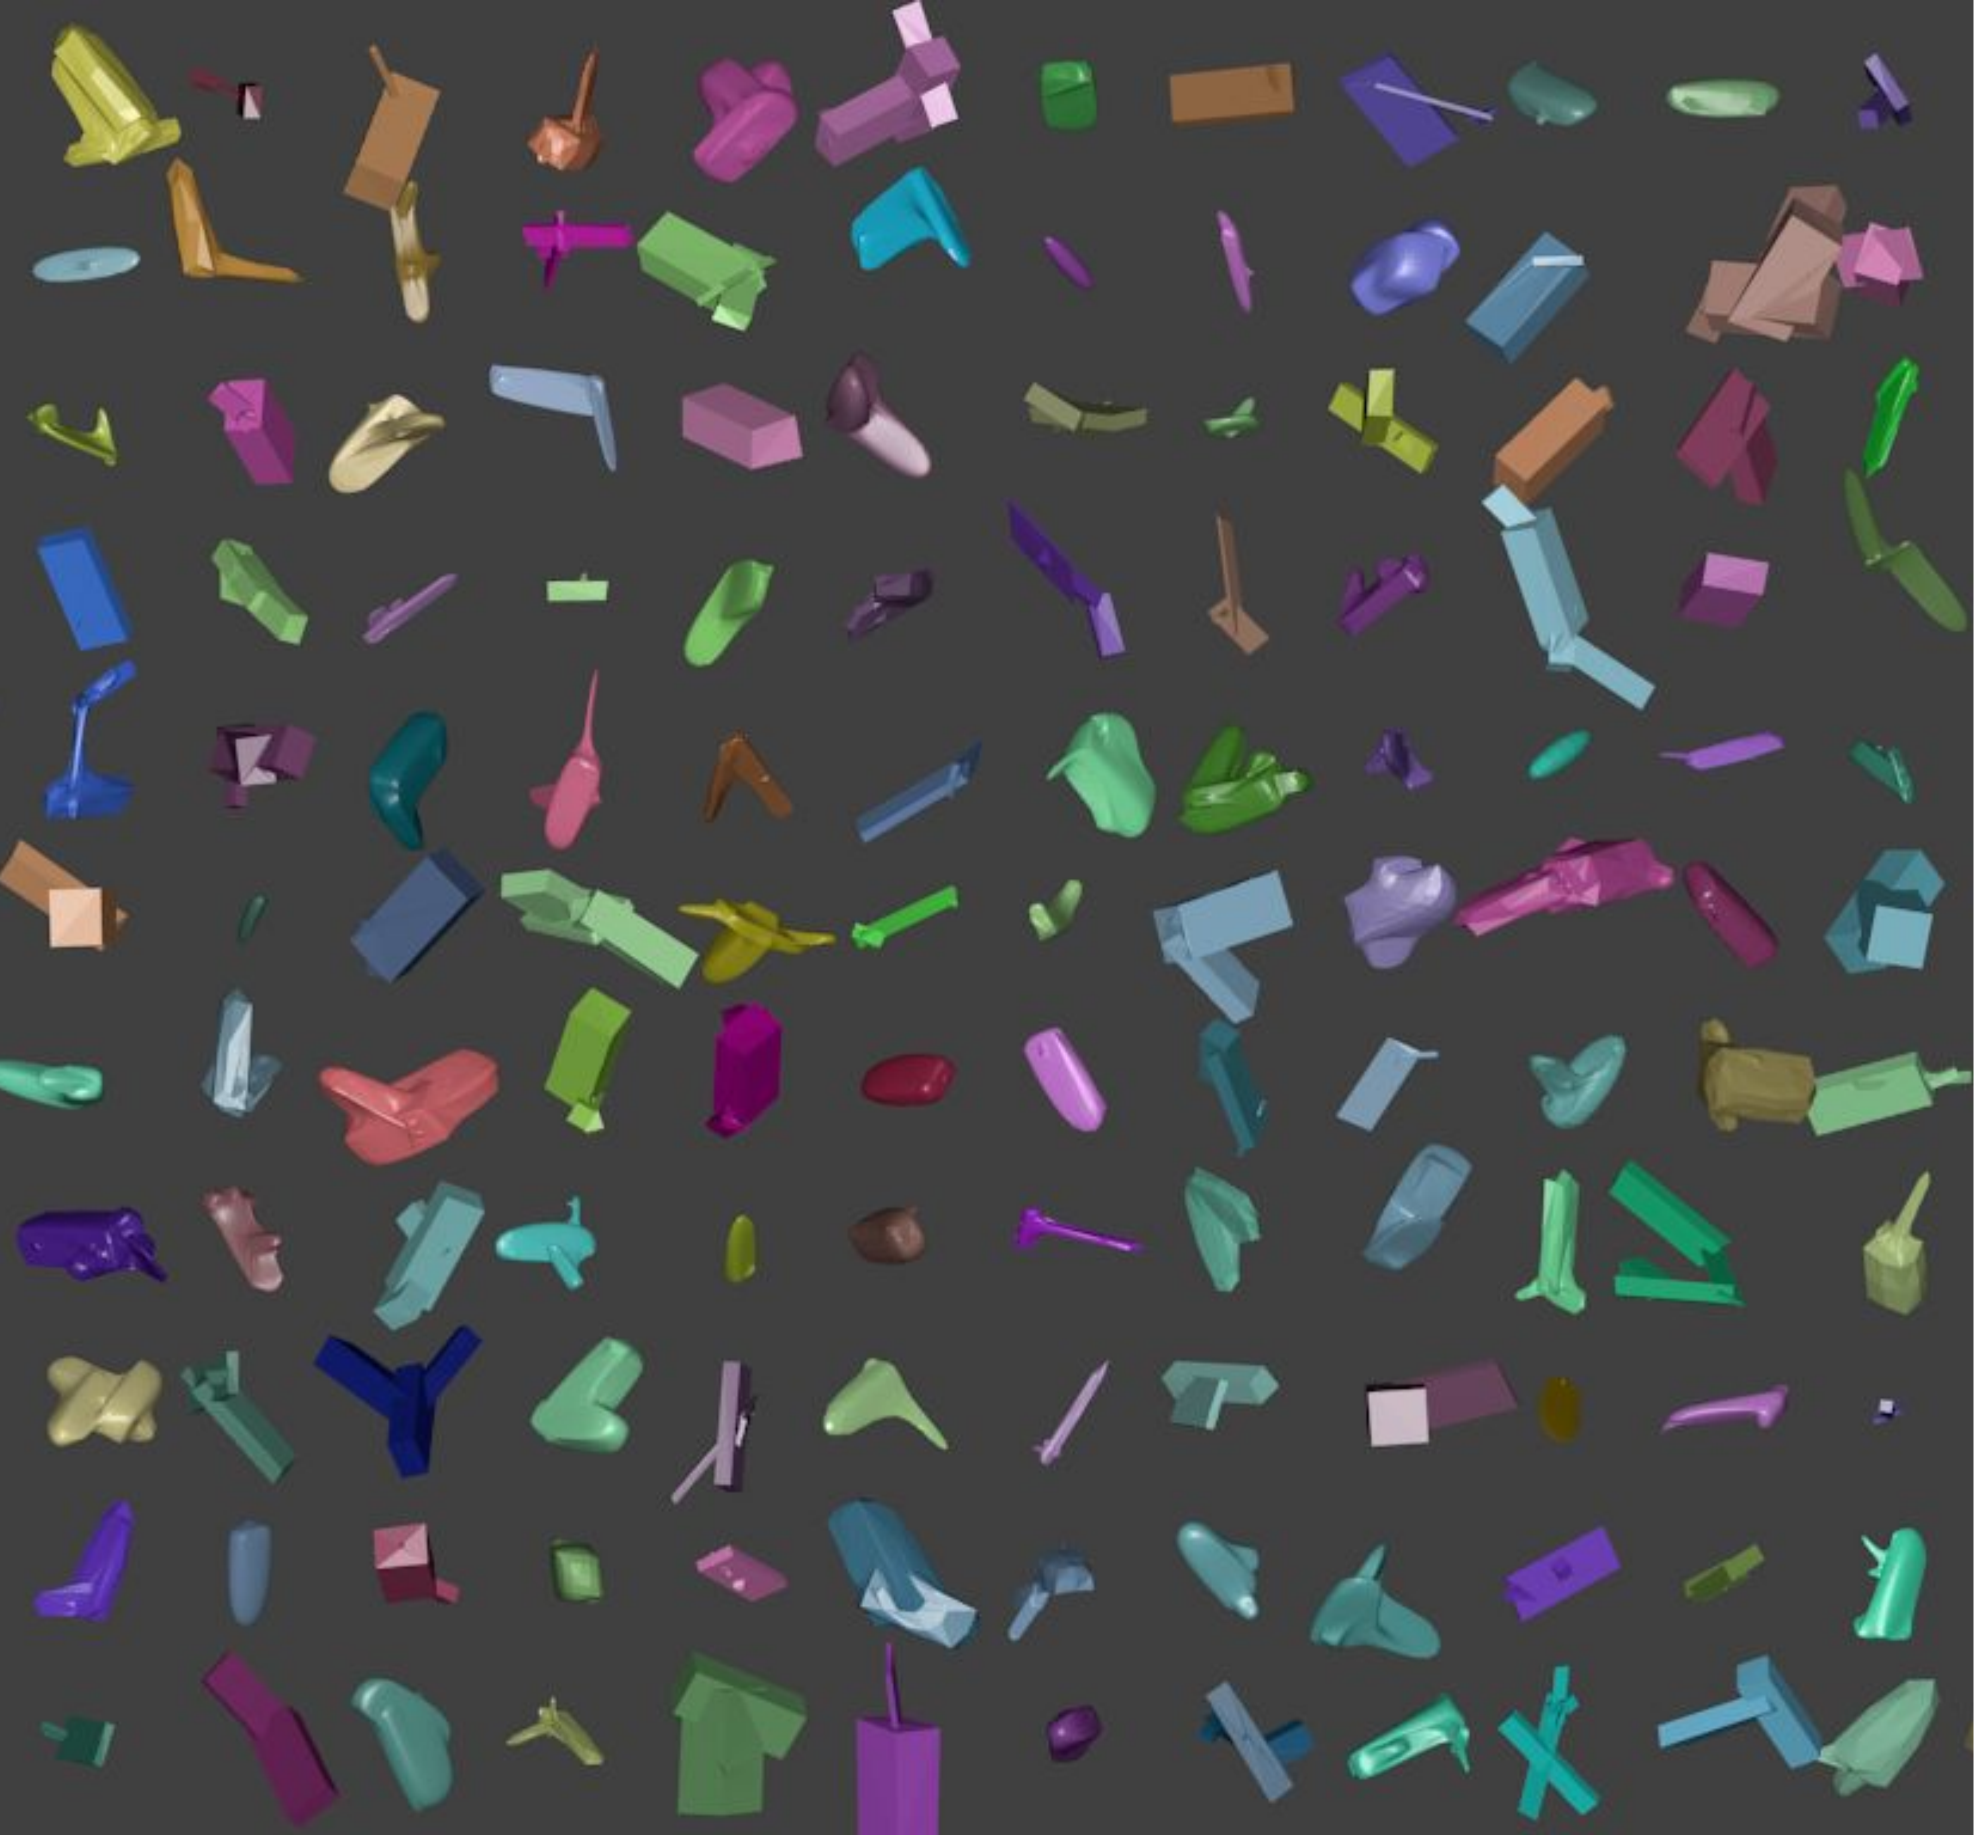

- 1,000 procedurally generated object models.

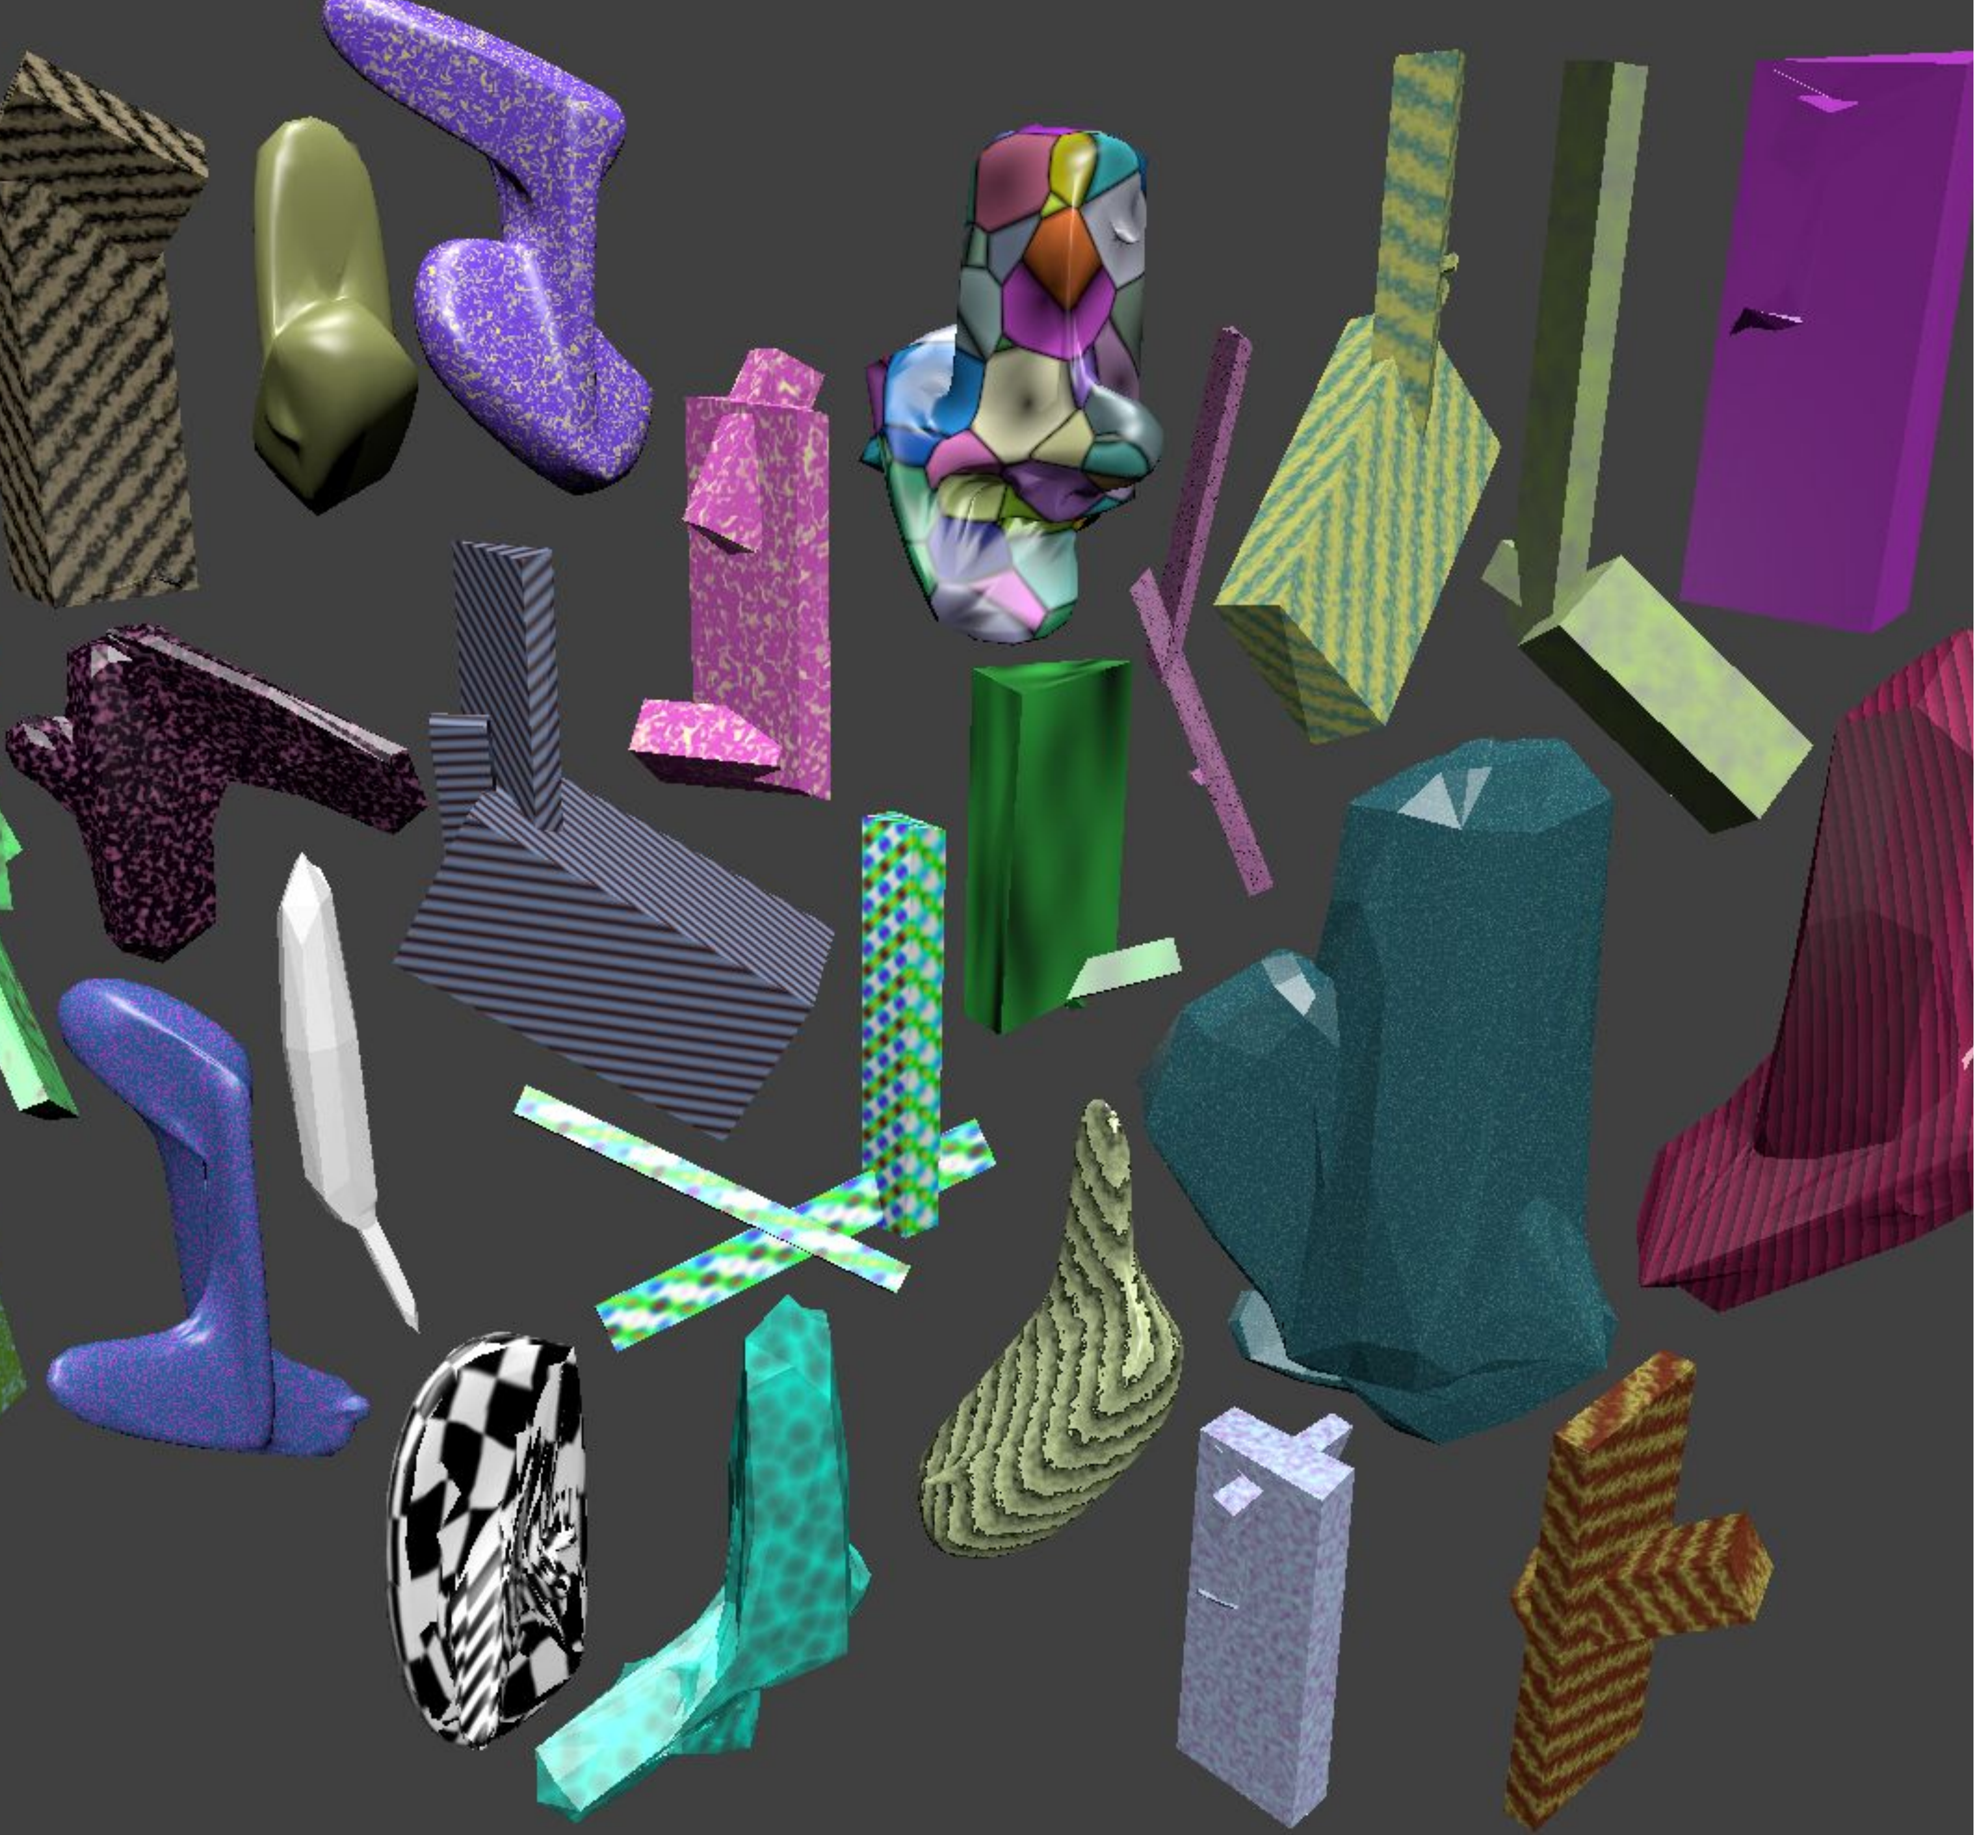

- Procedural objects with random textures.

Sim-to-Real Transfer for Perception - Feature Level Domain Adaptation

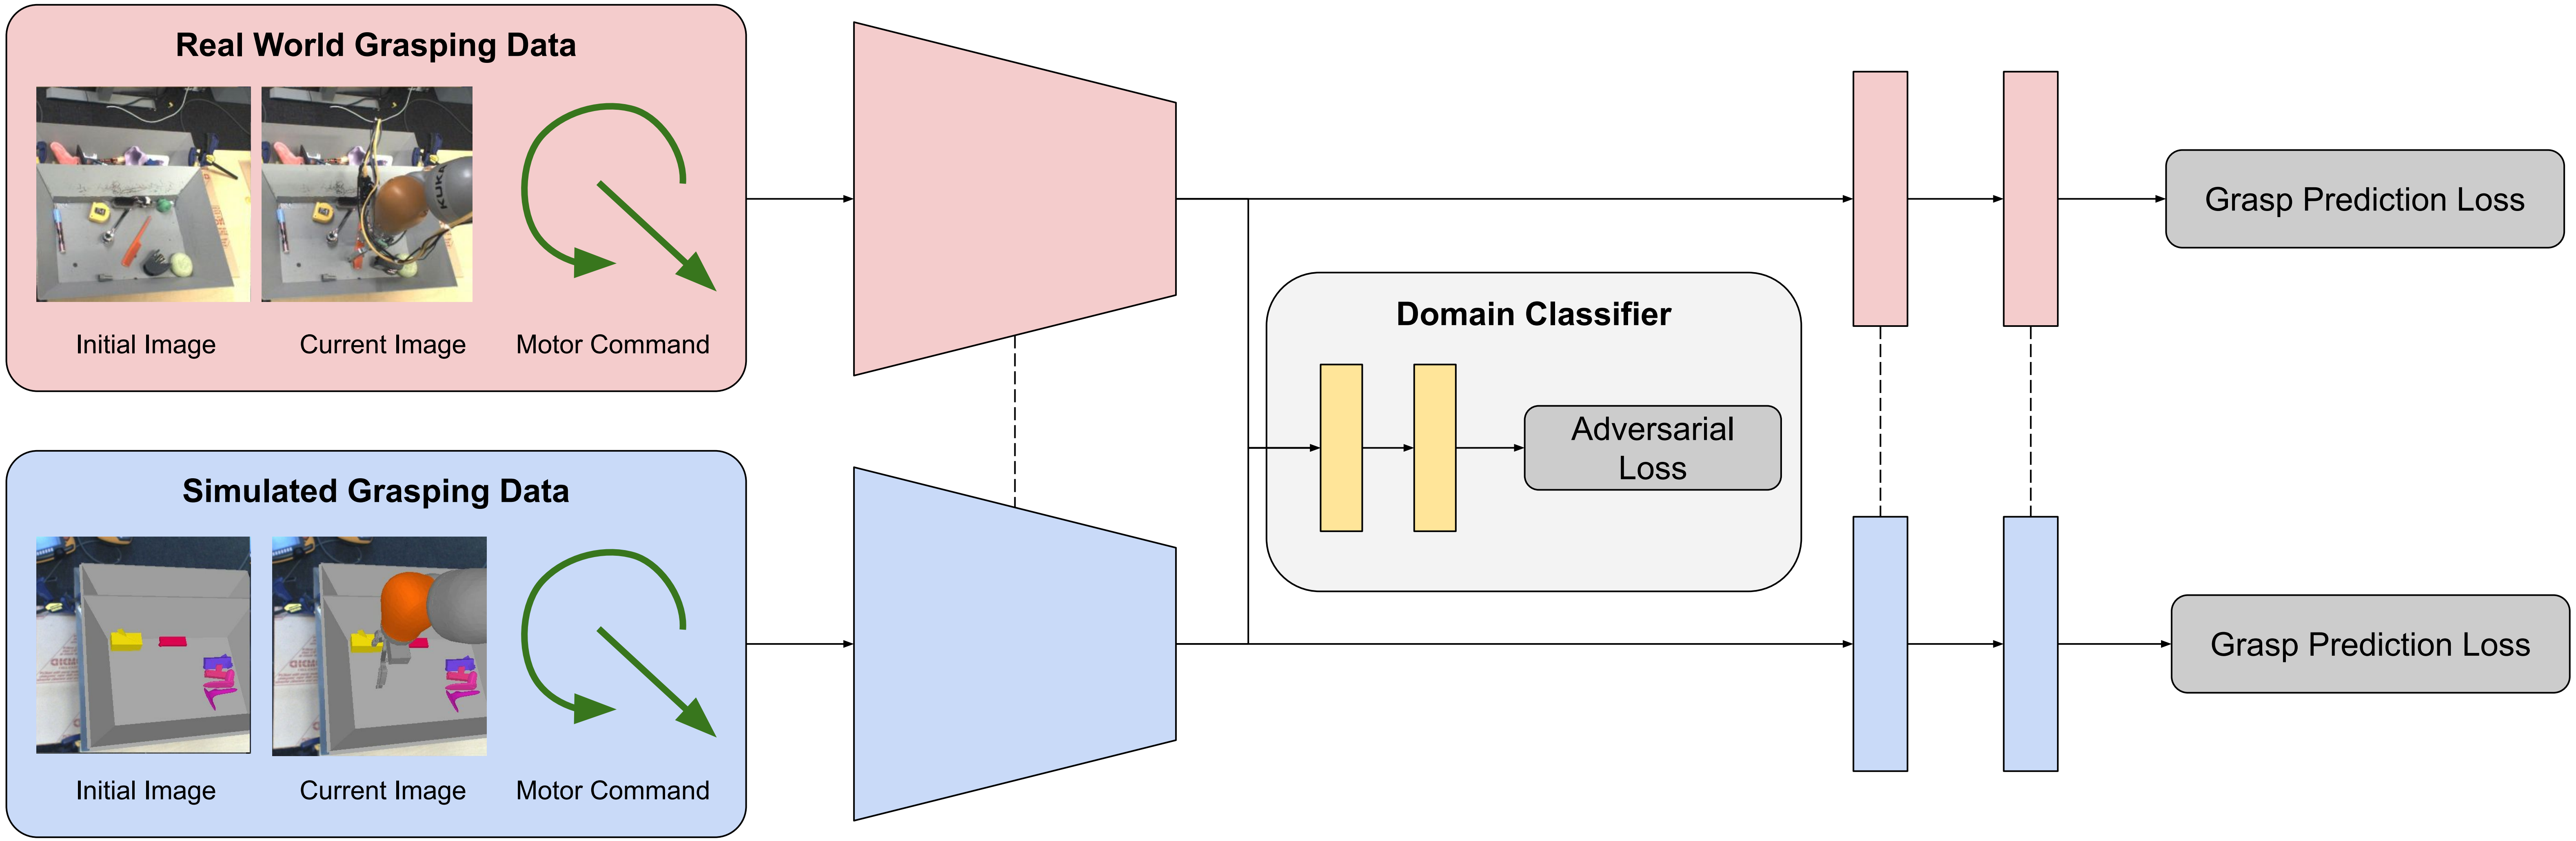

- Domain Adversarial Neural Network. Train features to be domain-invariant yet expressive, by using an adversarial loss. Learn features that confuse the domain classifier.

Based on Ganin et al., JMLR 2016

Sim-to-Real Transfer for Perception - Pixel Level Domain Adaptation

- Learn a generator (G) which converts synthetic images to real looking (fake) images.

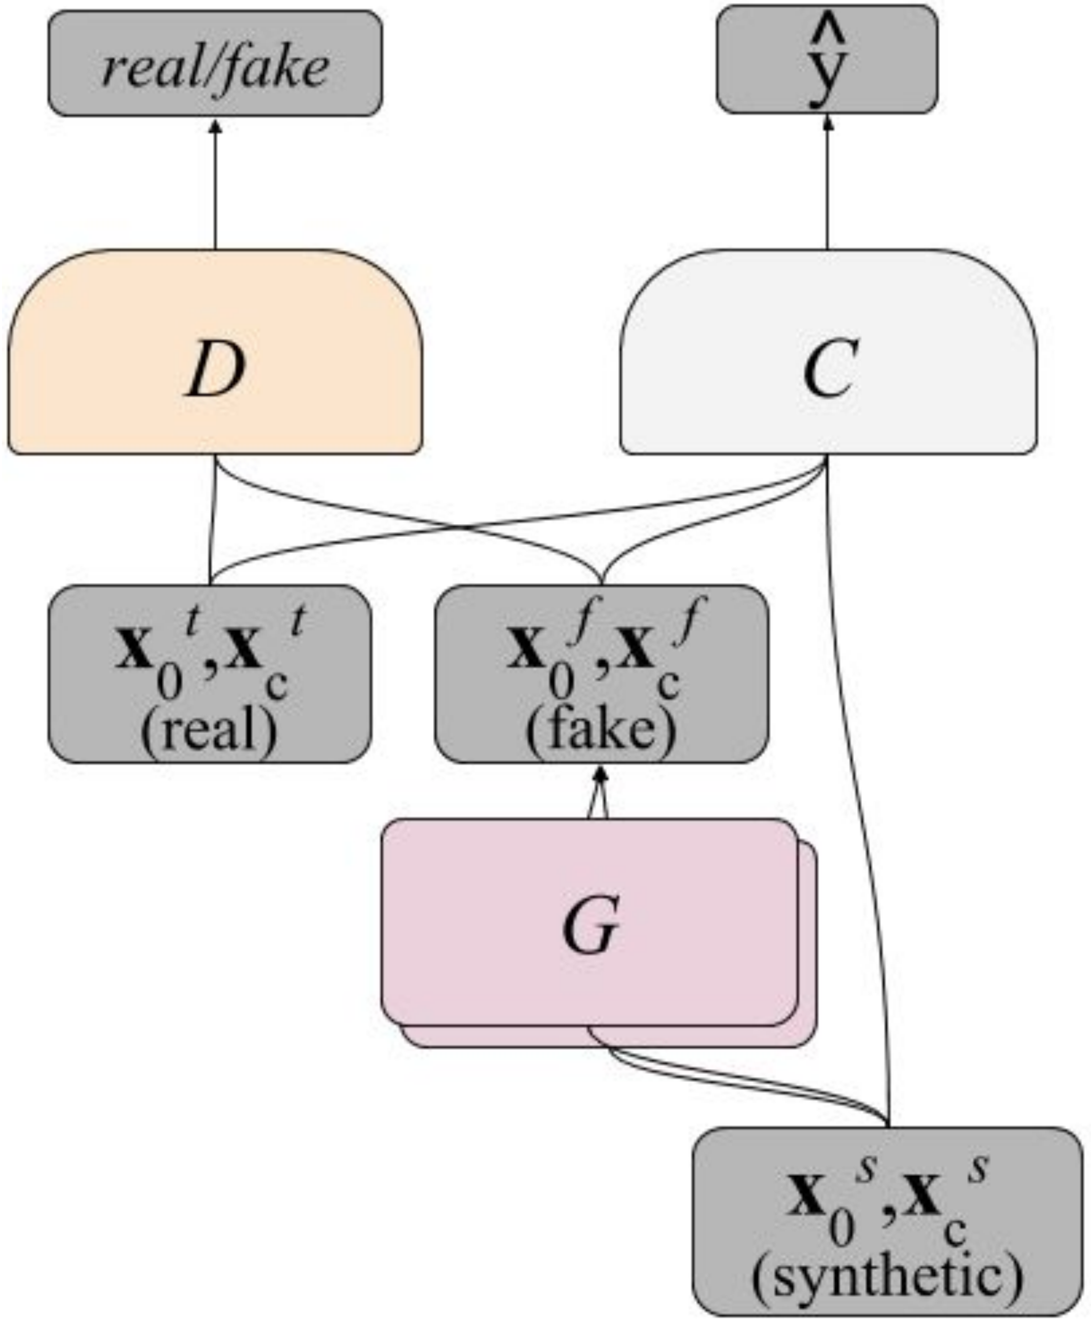

Sim-to-Real Transfer for Perception - Pixel Level Domain Adaptation

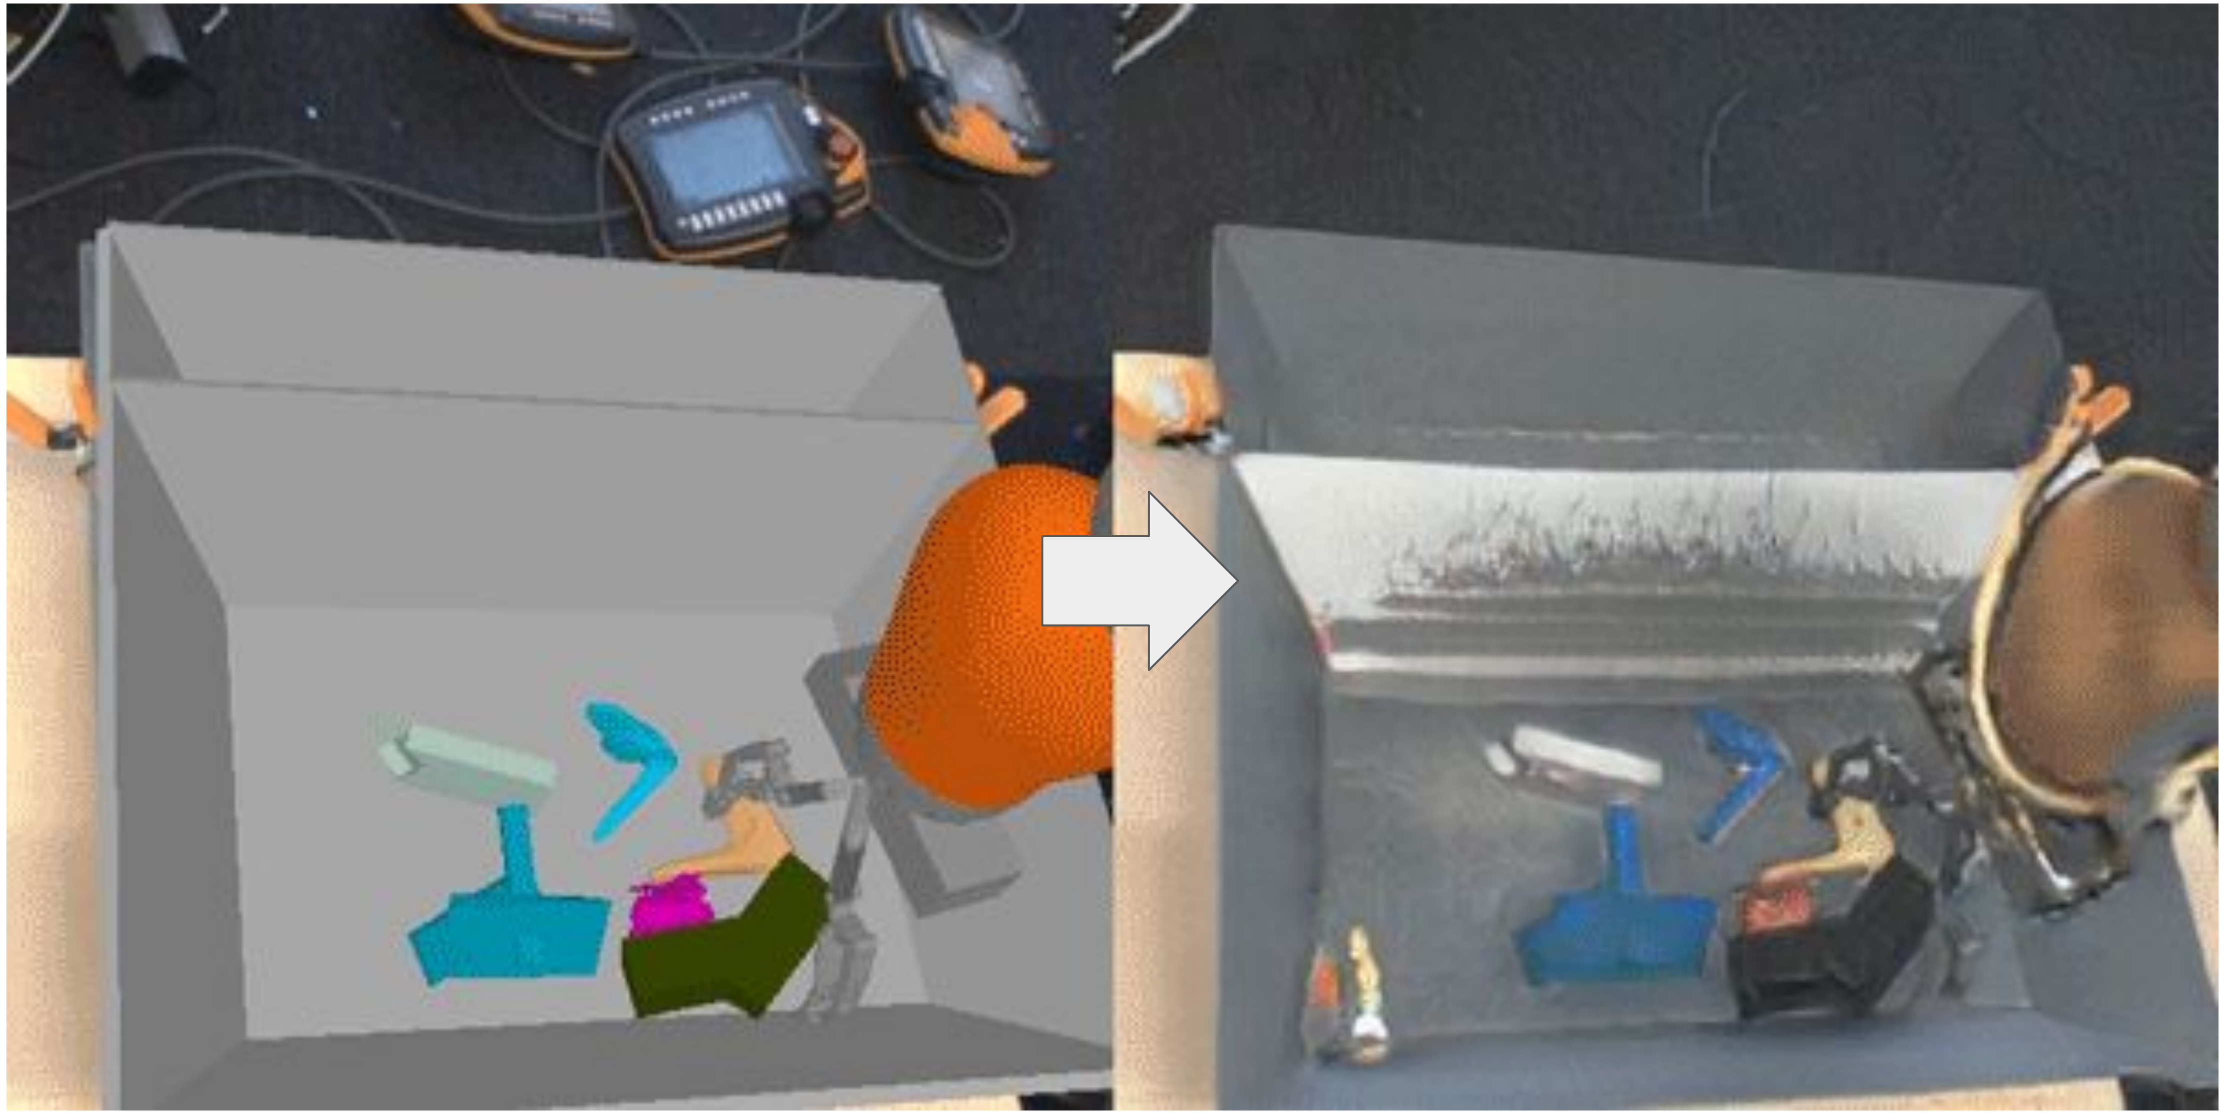

GraspGAN outperforms other techniques, providing more than 50x data efficiency.

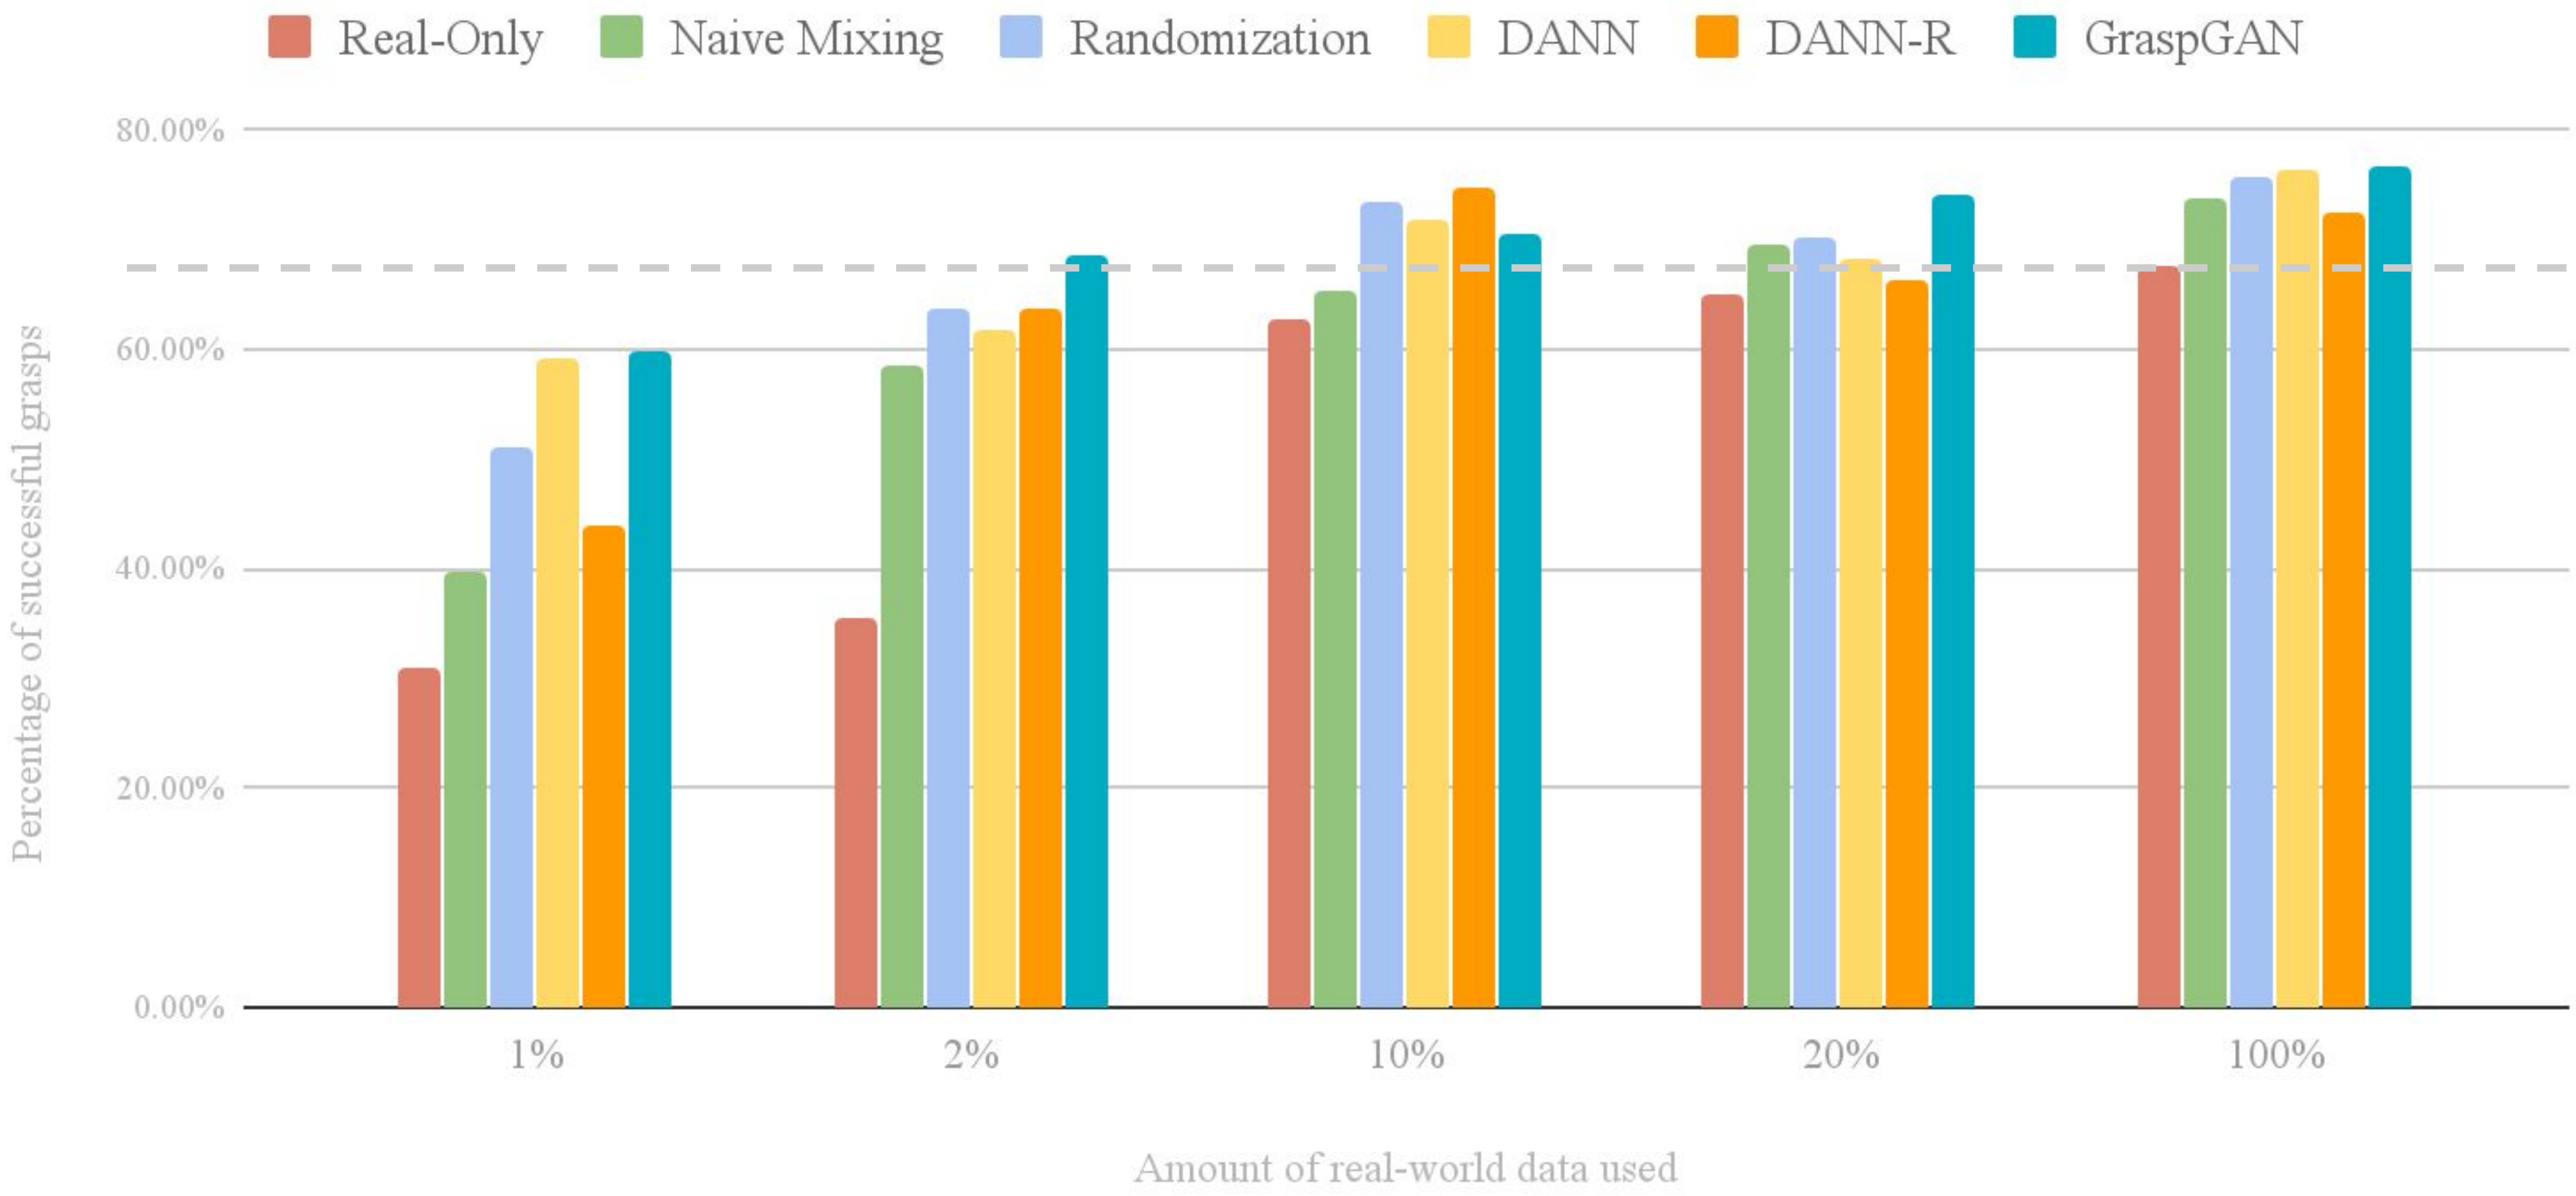

GraspGAN outperforms other techniques, providing more than 50x data efficiency.

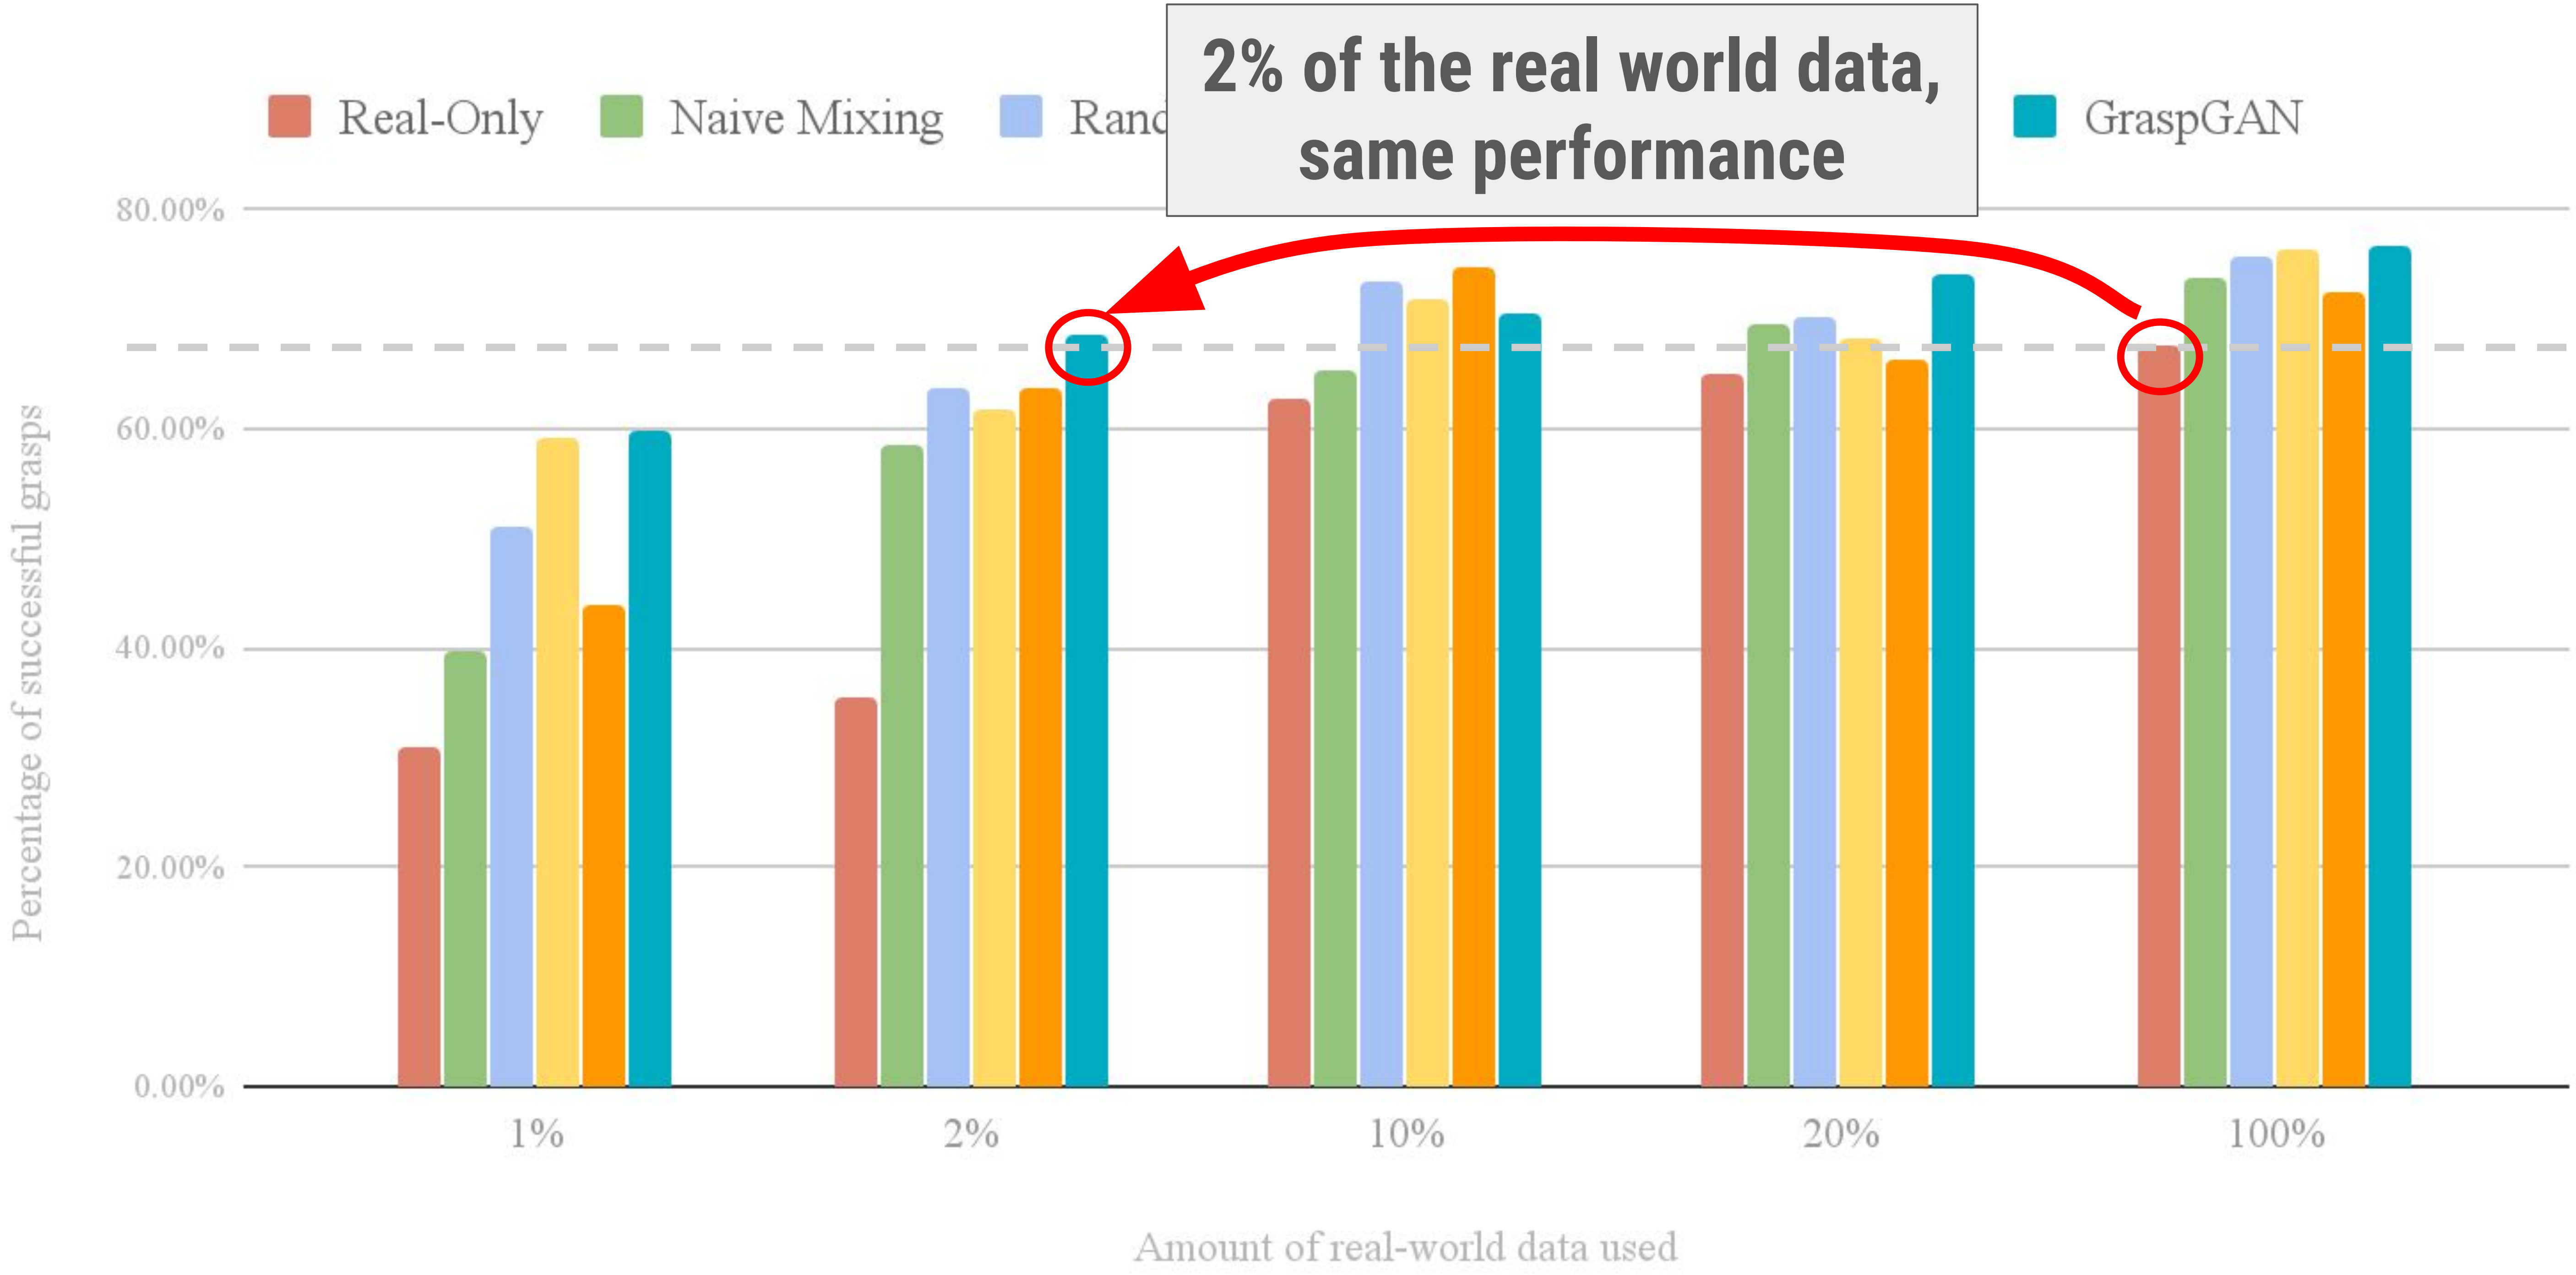

4

Solve Sim-to-Real via Sim-to-Sim

“Sim-to-Real via Sim-to-Sim: Data-efficient Robotic Grasping via Randomized-to-Canonical Adaptation Networks”,

Stephen James, Paul Wohlhart, Mrinal Kalakrishnan, Dmitry Kalashnikov, Alex Irpan, Julian Ibarz,

Sergey Levine, Raia Hadsell, Konstantinos Bousmalis

Randomized-to-Canonical Adaptation Networks

- **RCAN** is a real-to-sim image translator trained with domain randomization:
  - We define a “canonical” version of simulation and randomizations

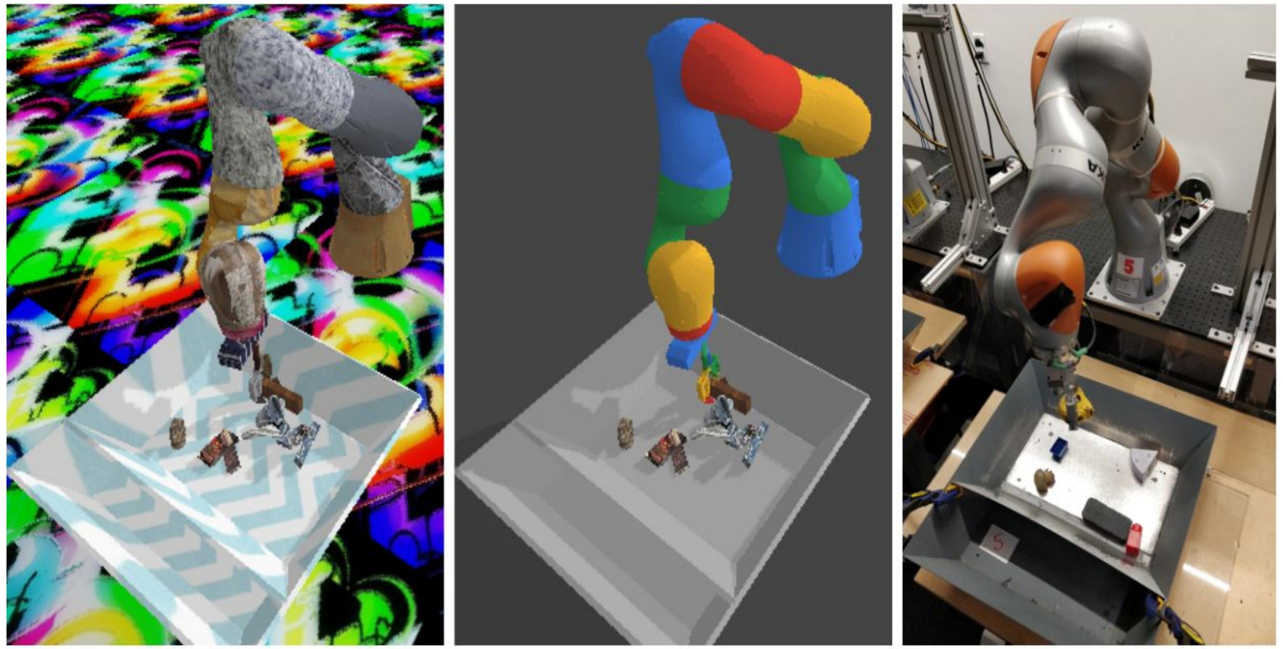

randomized canonical real

Randomized-to-Canonical Adaptation Networks

- **RCAN** is a real-to-sim image translator trained with domain randomization:
  - We define a “canonical” version of simulation and randomizations
  - We train a pix2pix model to convert randomized sim images to equivalent canonical versions

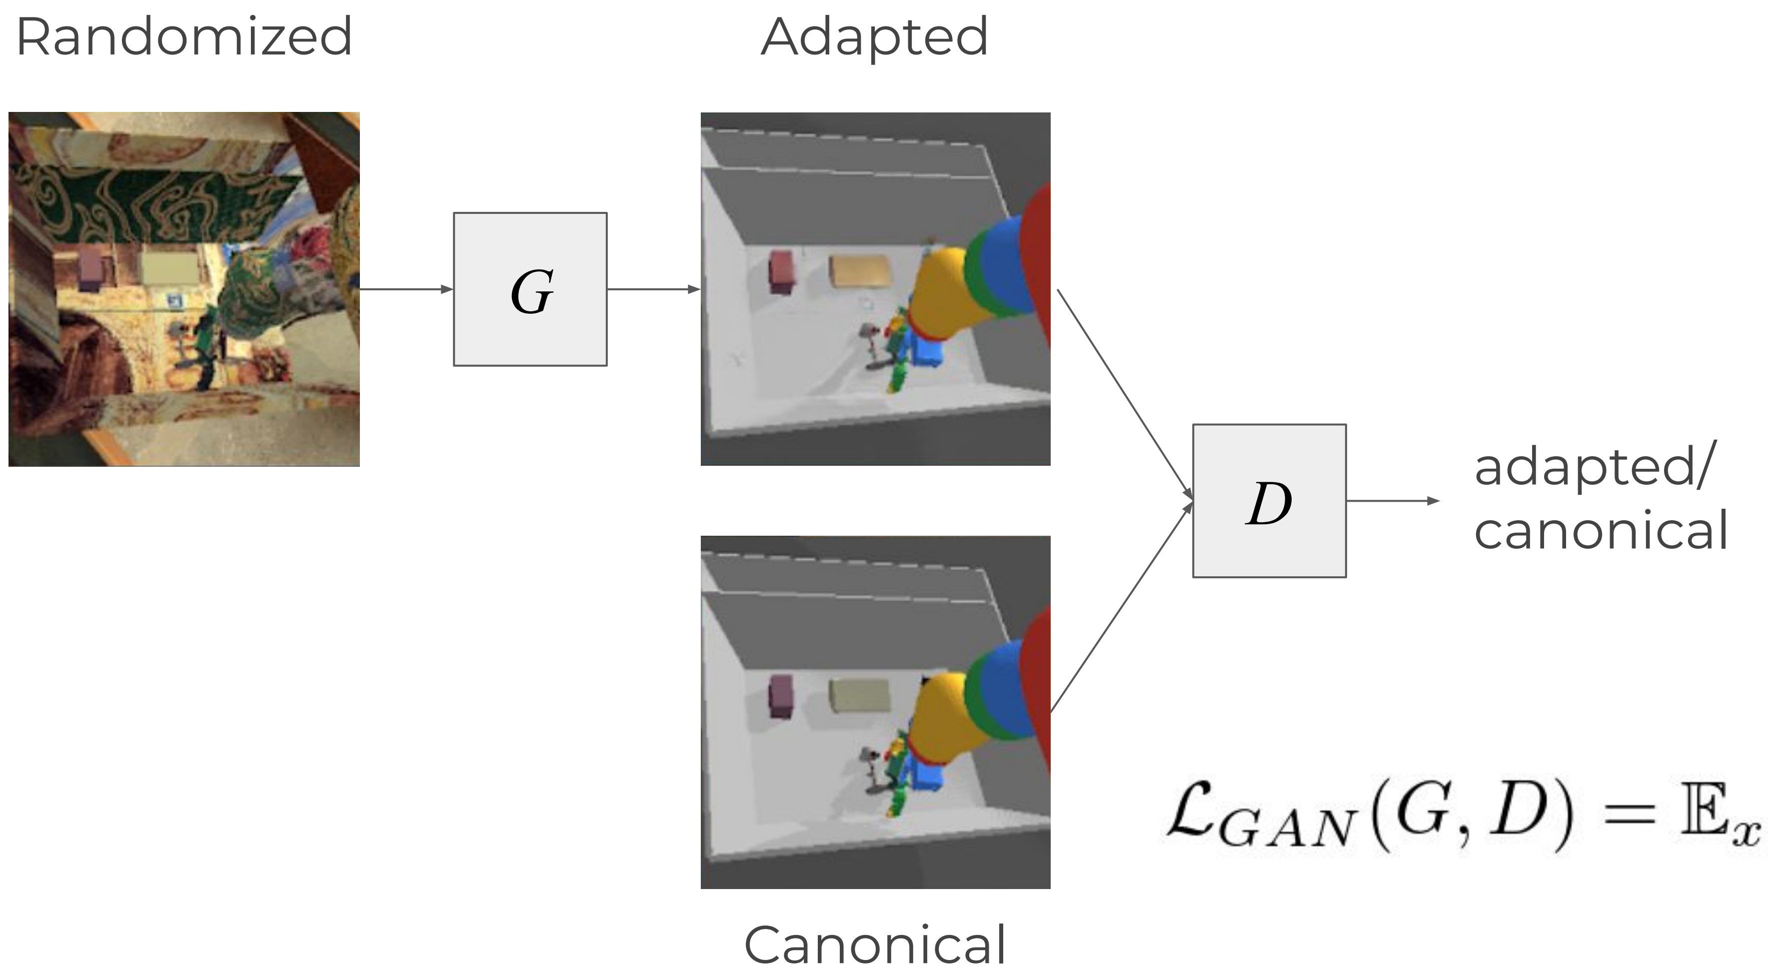

$$\mathcal{L}_{GAN}(G, D) = \mathbb{E}_x[\log D(x)] + \mathbb{E}_x[\log(1 - D(G_x(x)))]$$

Randomized-to-Canonical Adaptation Networks

- **RCAN** is a real-to-sim image translator trained with domain randomization:
  - We define a “canonical” version of simulation and randomizations
  - We train a pix2pix model to convert randomized sim images to equivalent canonical versions
  - In the real world, RCAN will then also be able to translate real images to canonical sim versions

Real Adapted

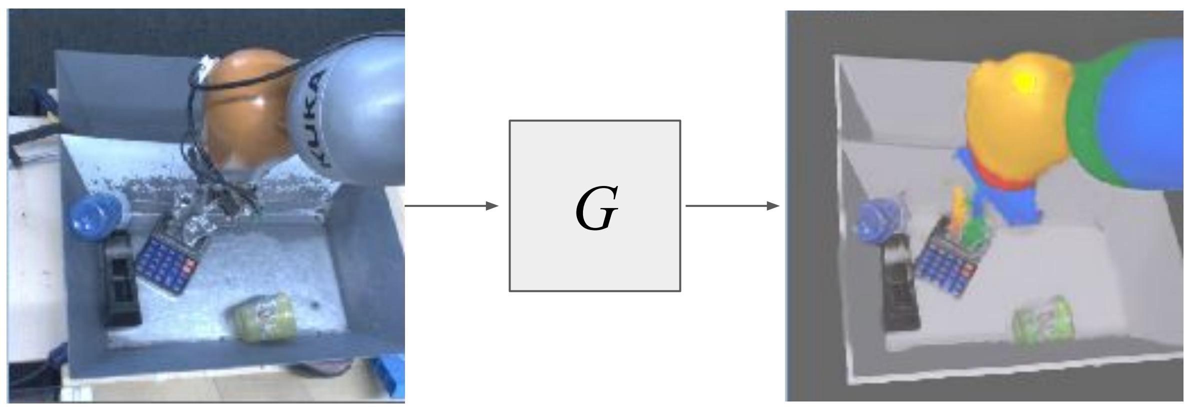

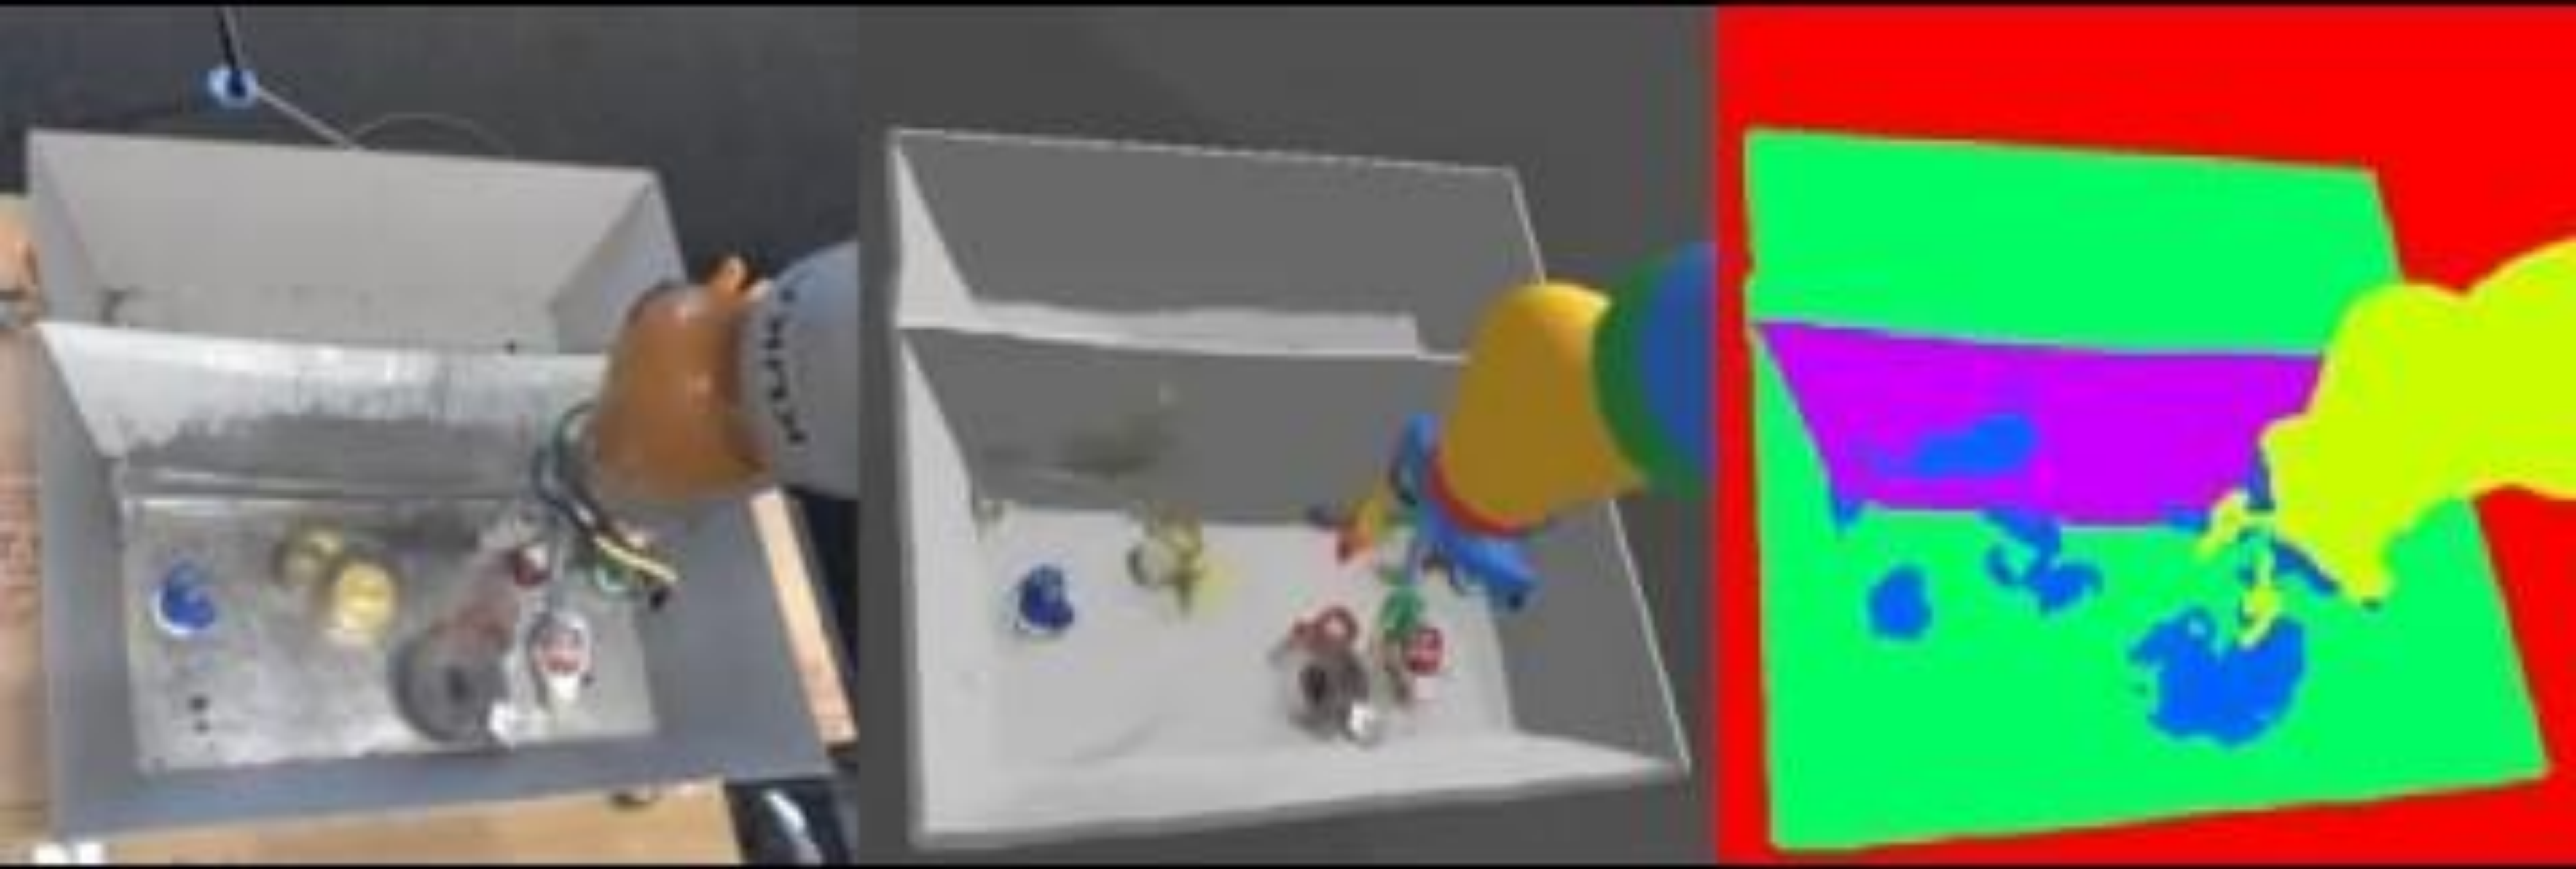

RCAN achieves the similar success rate (94% vs 96%), but without using 580k real word data

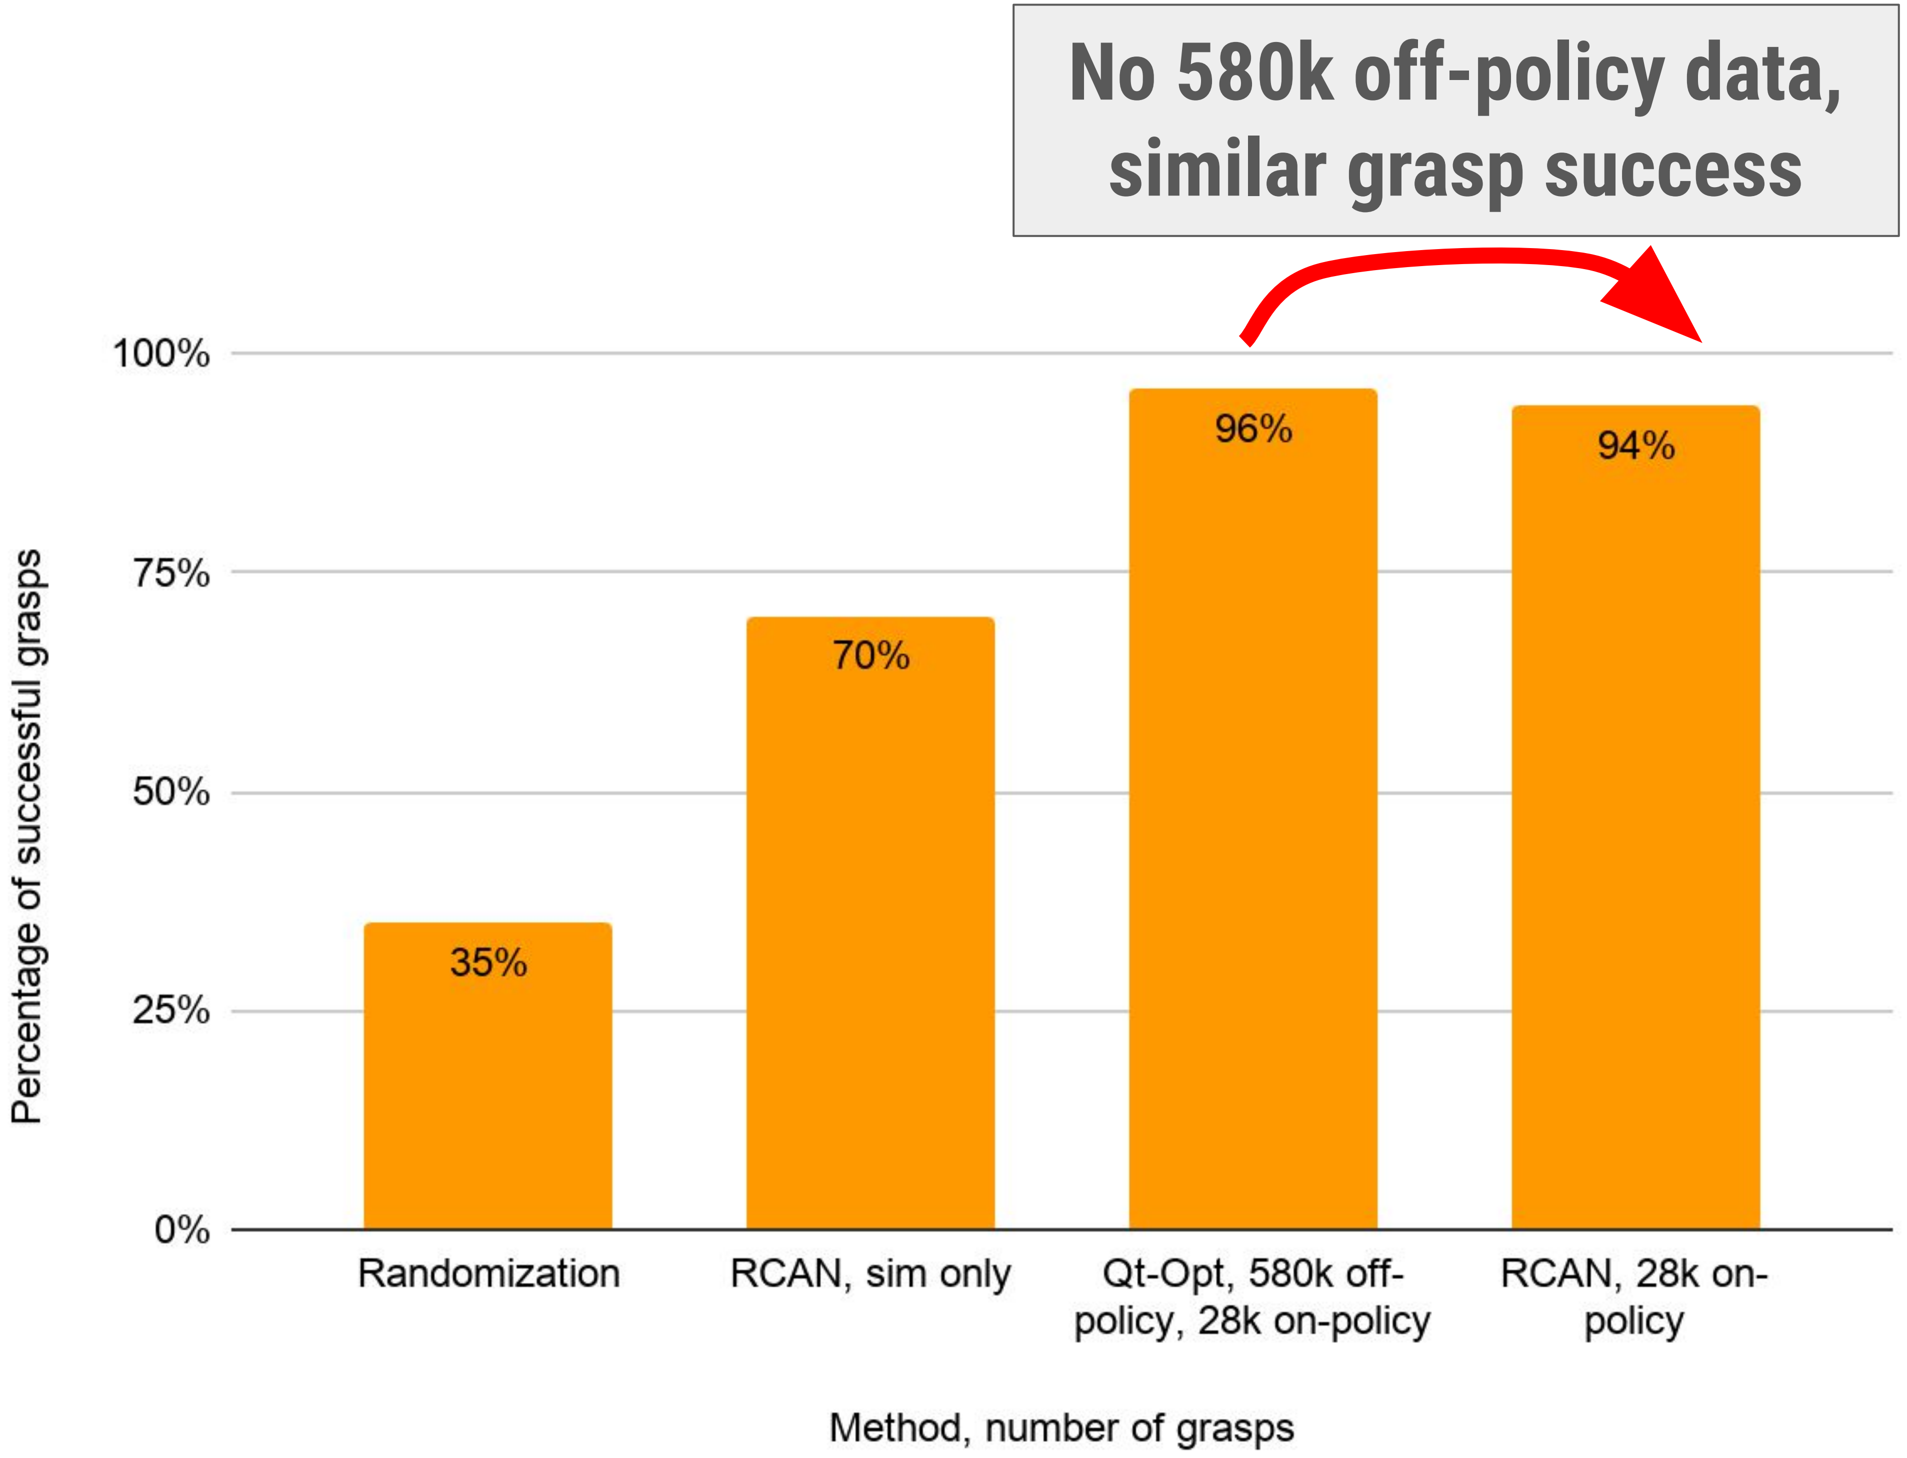

5

Learning New Tasks

“Multi-Task Domain Adaptation for Deep Learning of Instance Grasping from Simulation”,

Kuan Fang, Yunfei Bai, Stefan Hinterstoisser, Silvio Savarese, Mrinal Kalakrishnan

Sim-to-Real Transfer: Applying to a More Challenging Task

Initial Image

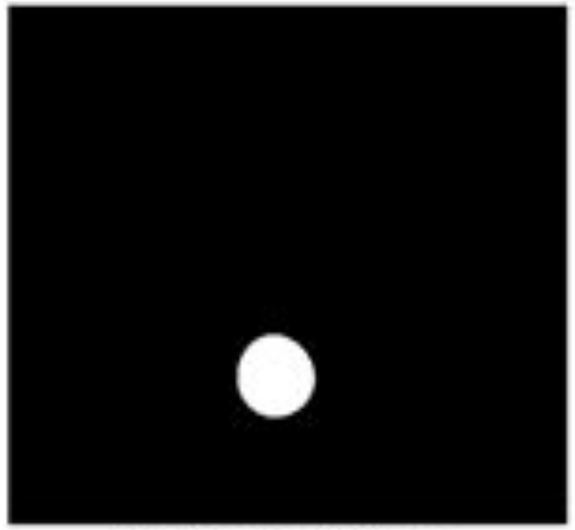

Target Mask

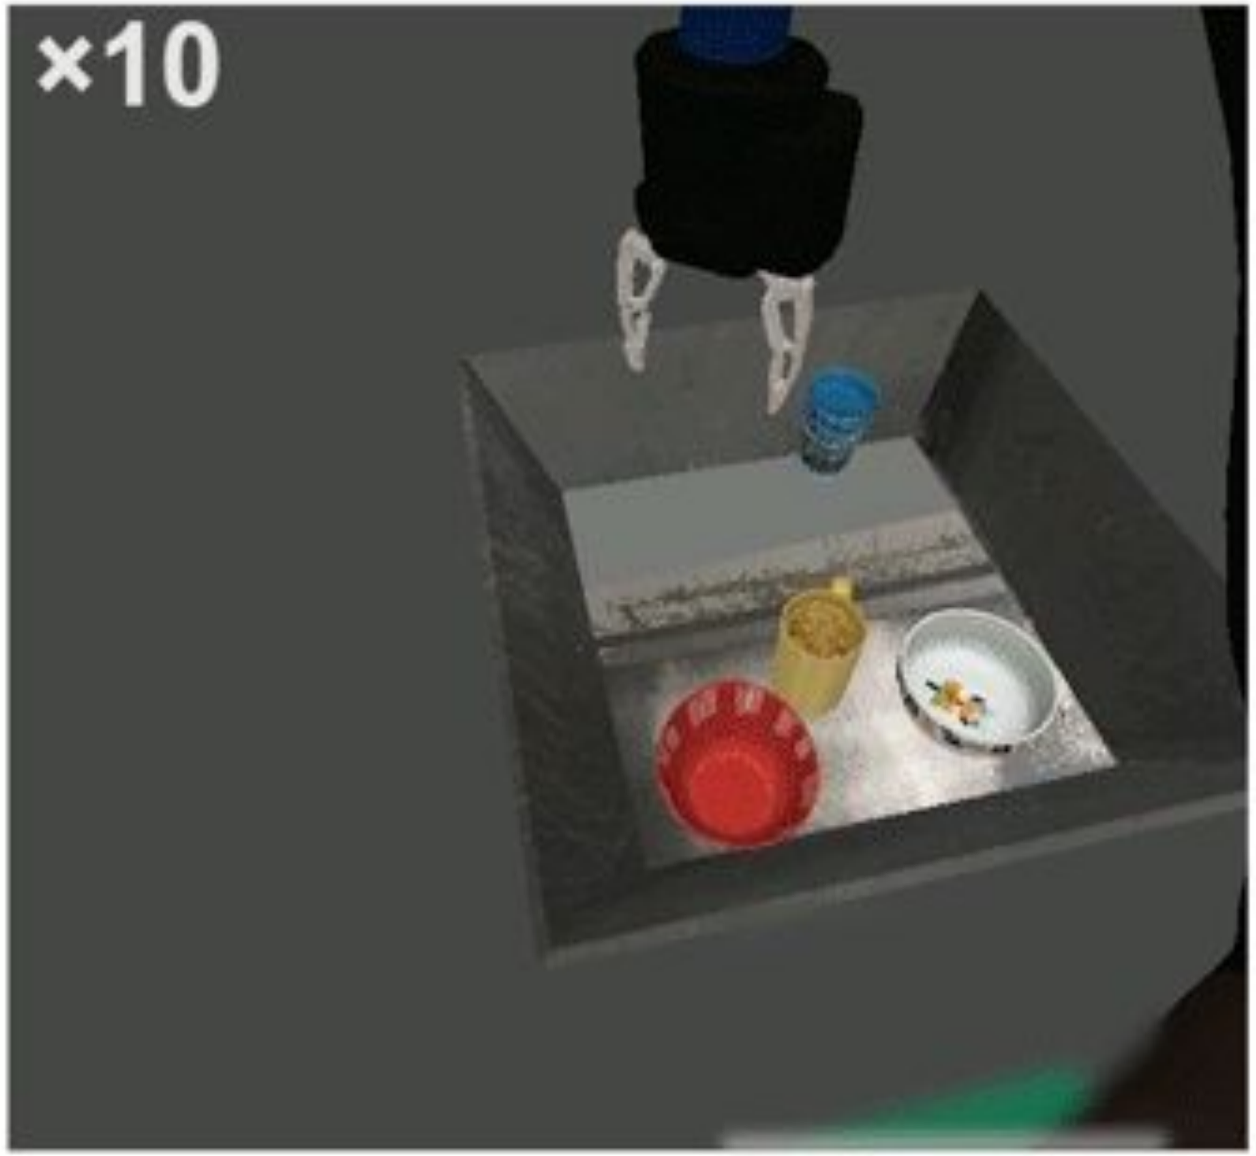

Instance Grasping Framework: Multi-Task Domain Adaptation

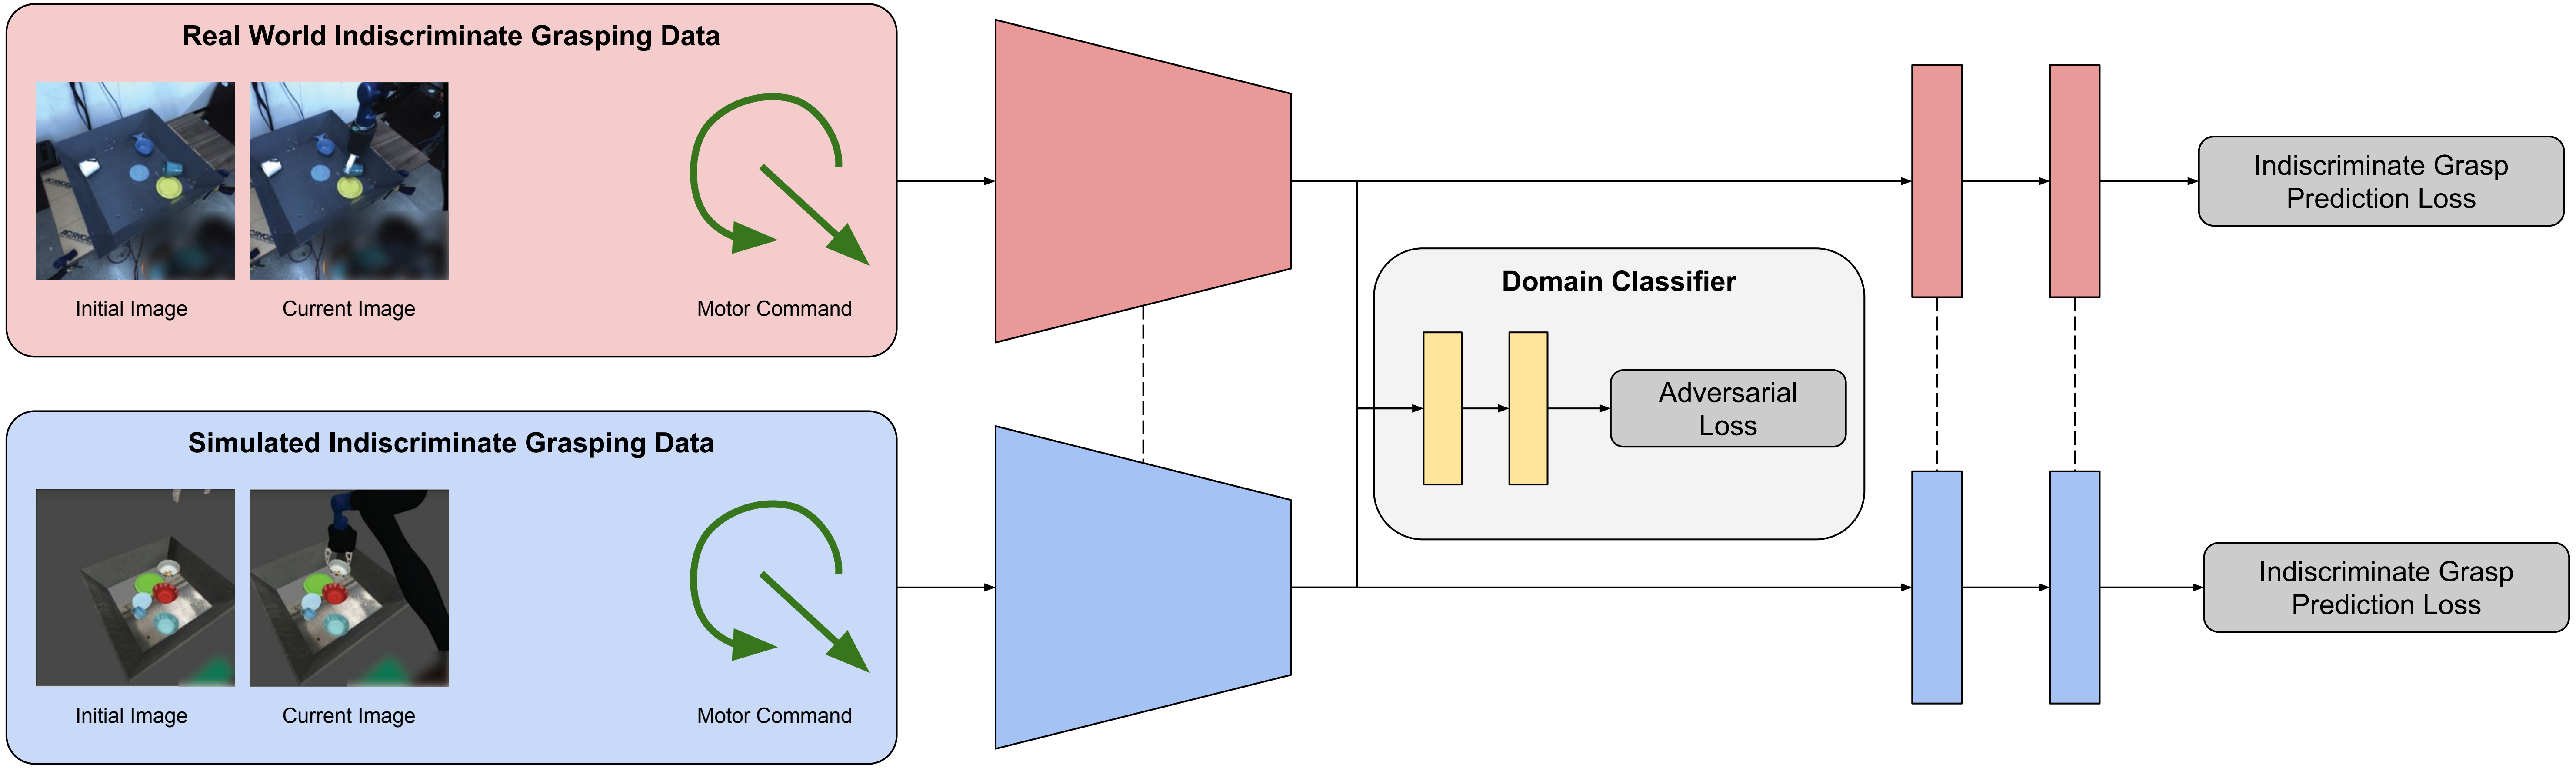

Instance Grasping Framework: Multi-Task Domain Adaptation

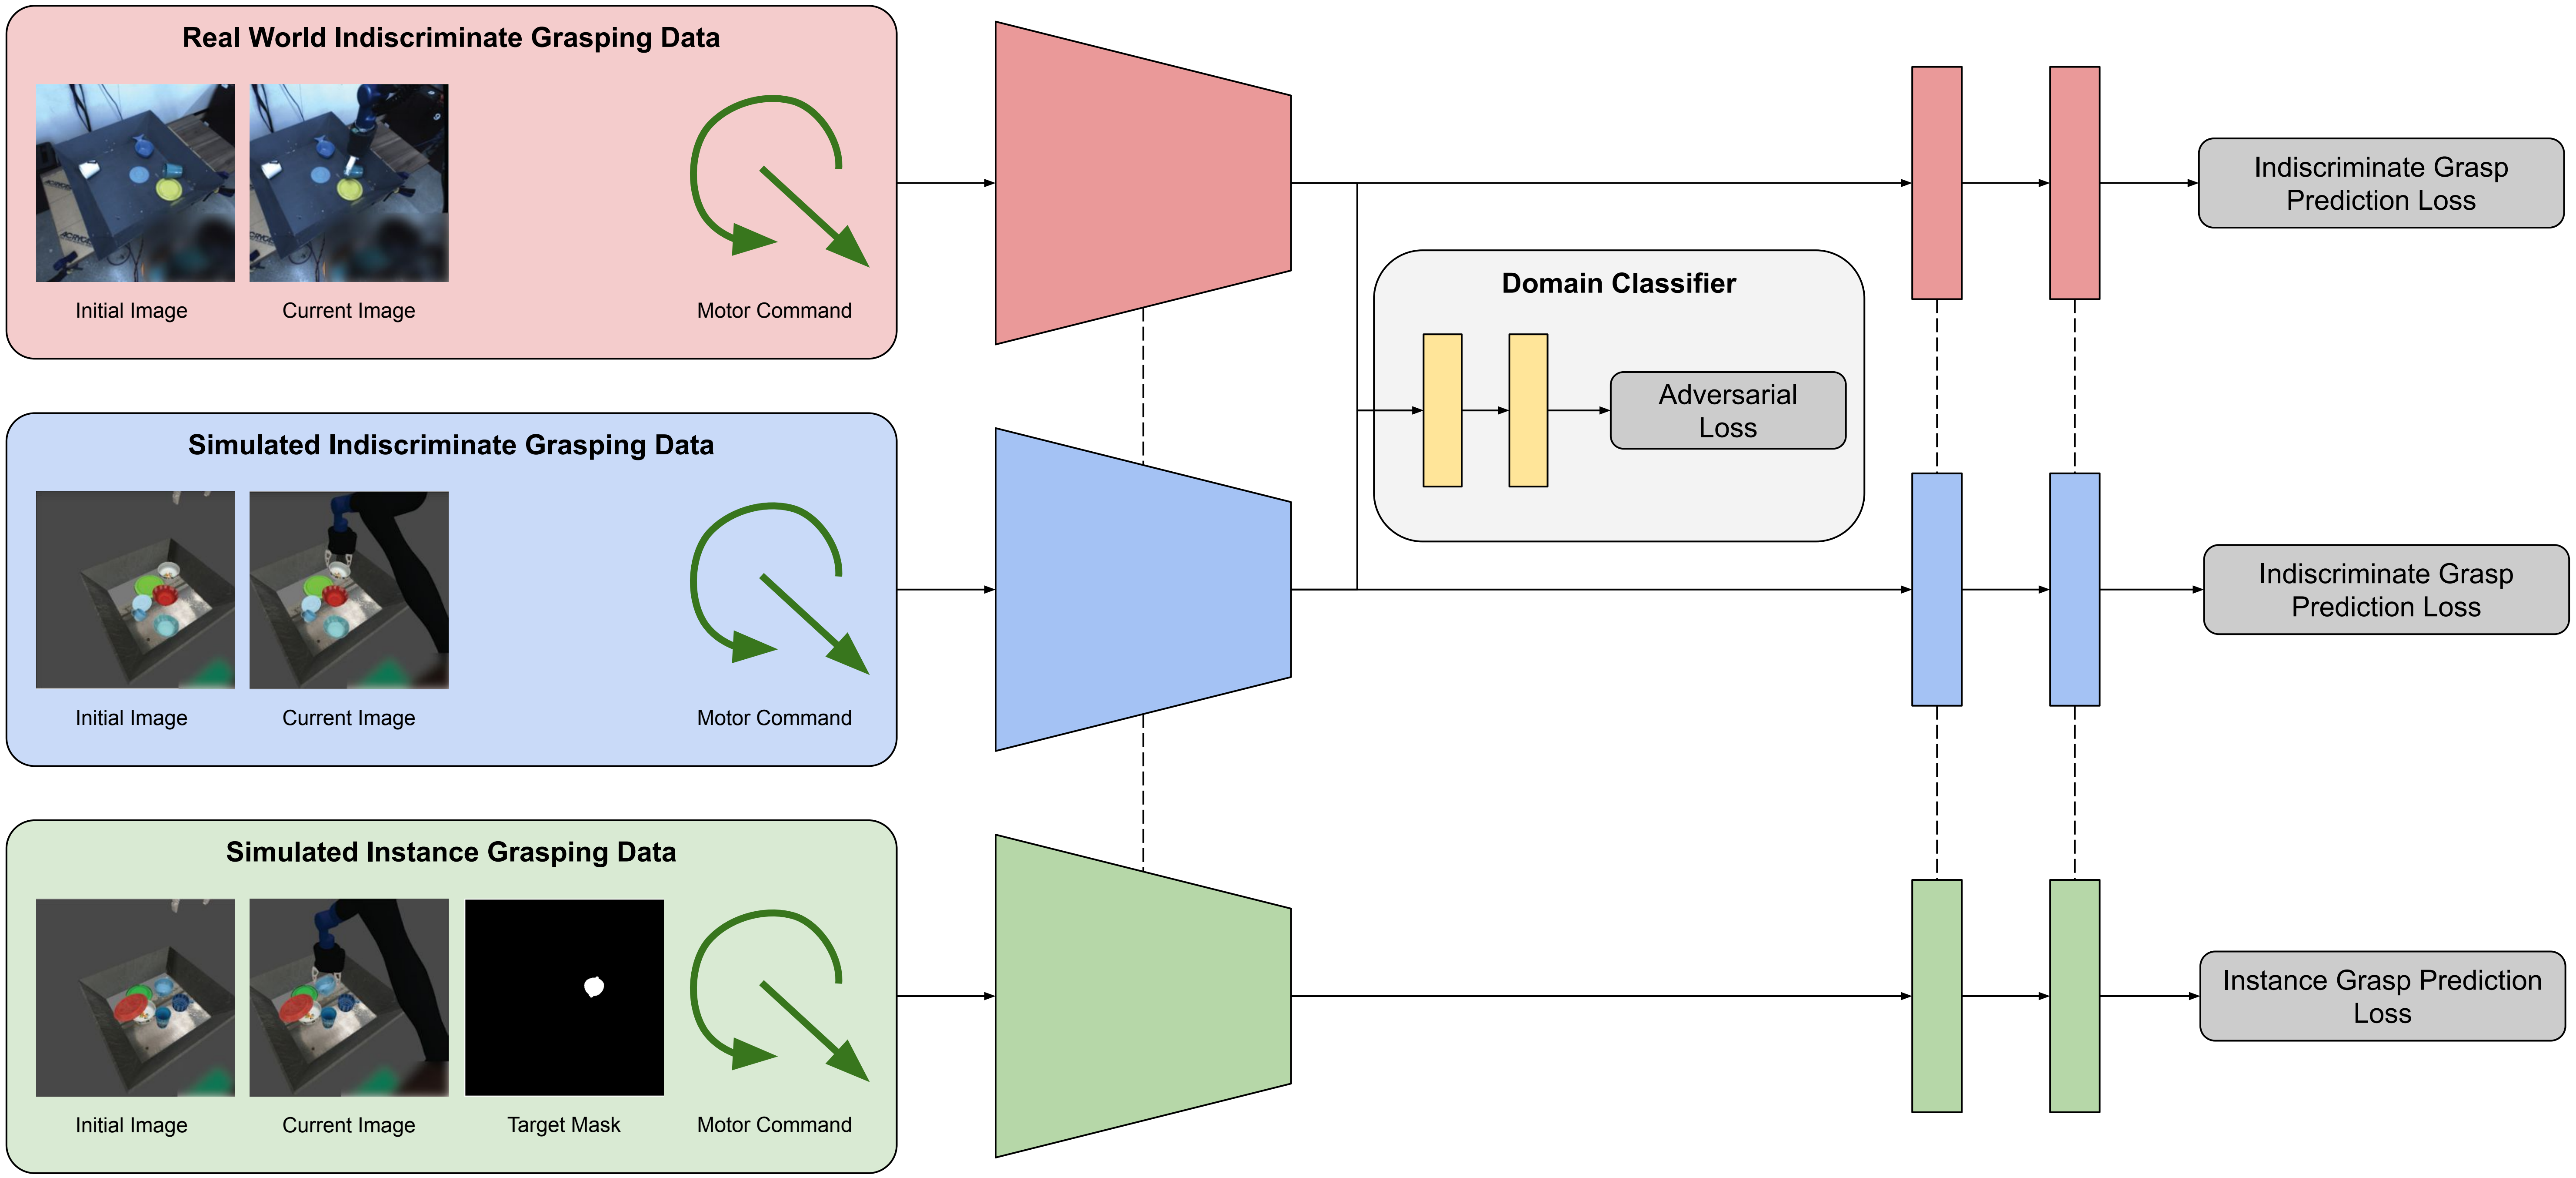

Instance Grasping Framework: Multi-Task Domain Adaptation

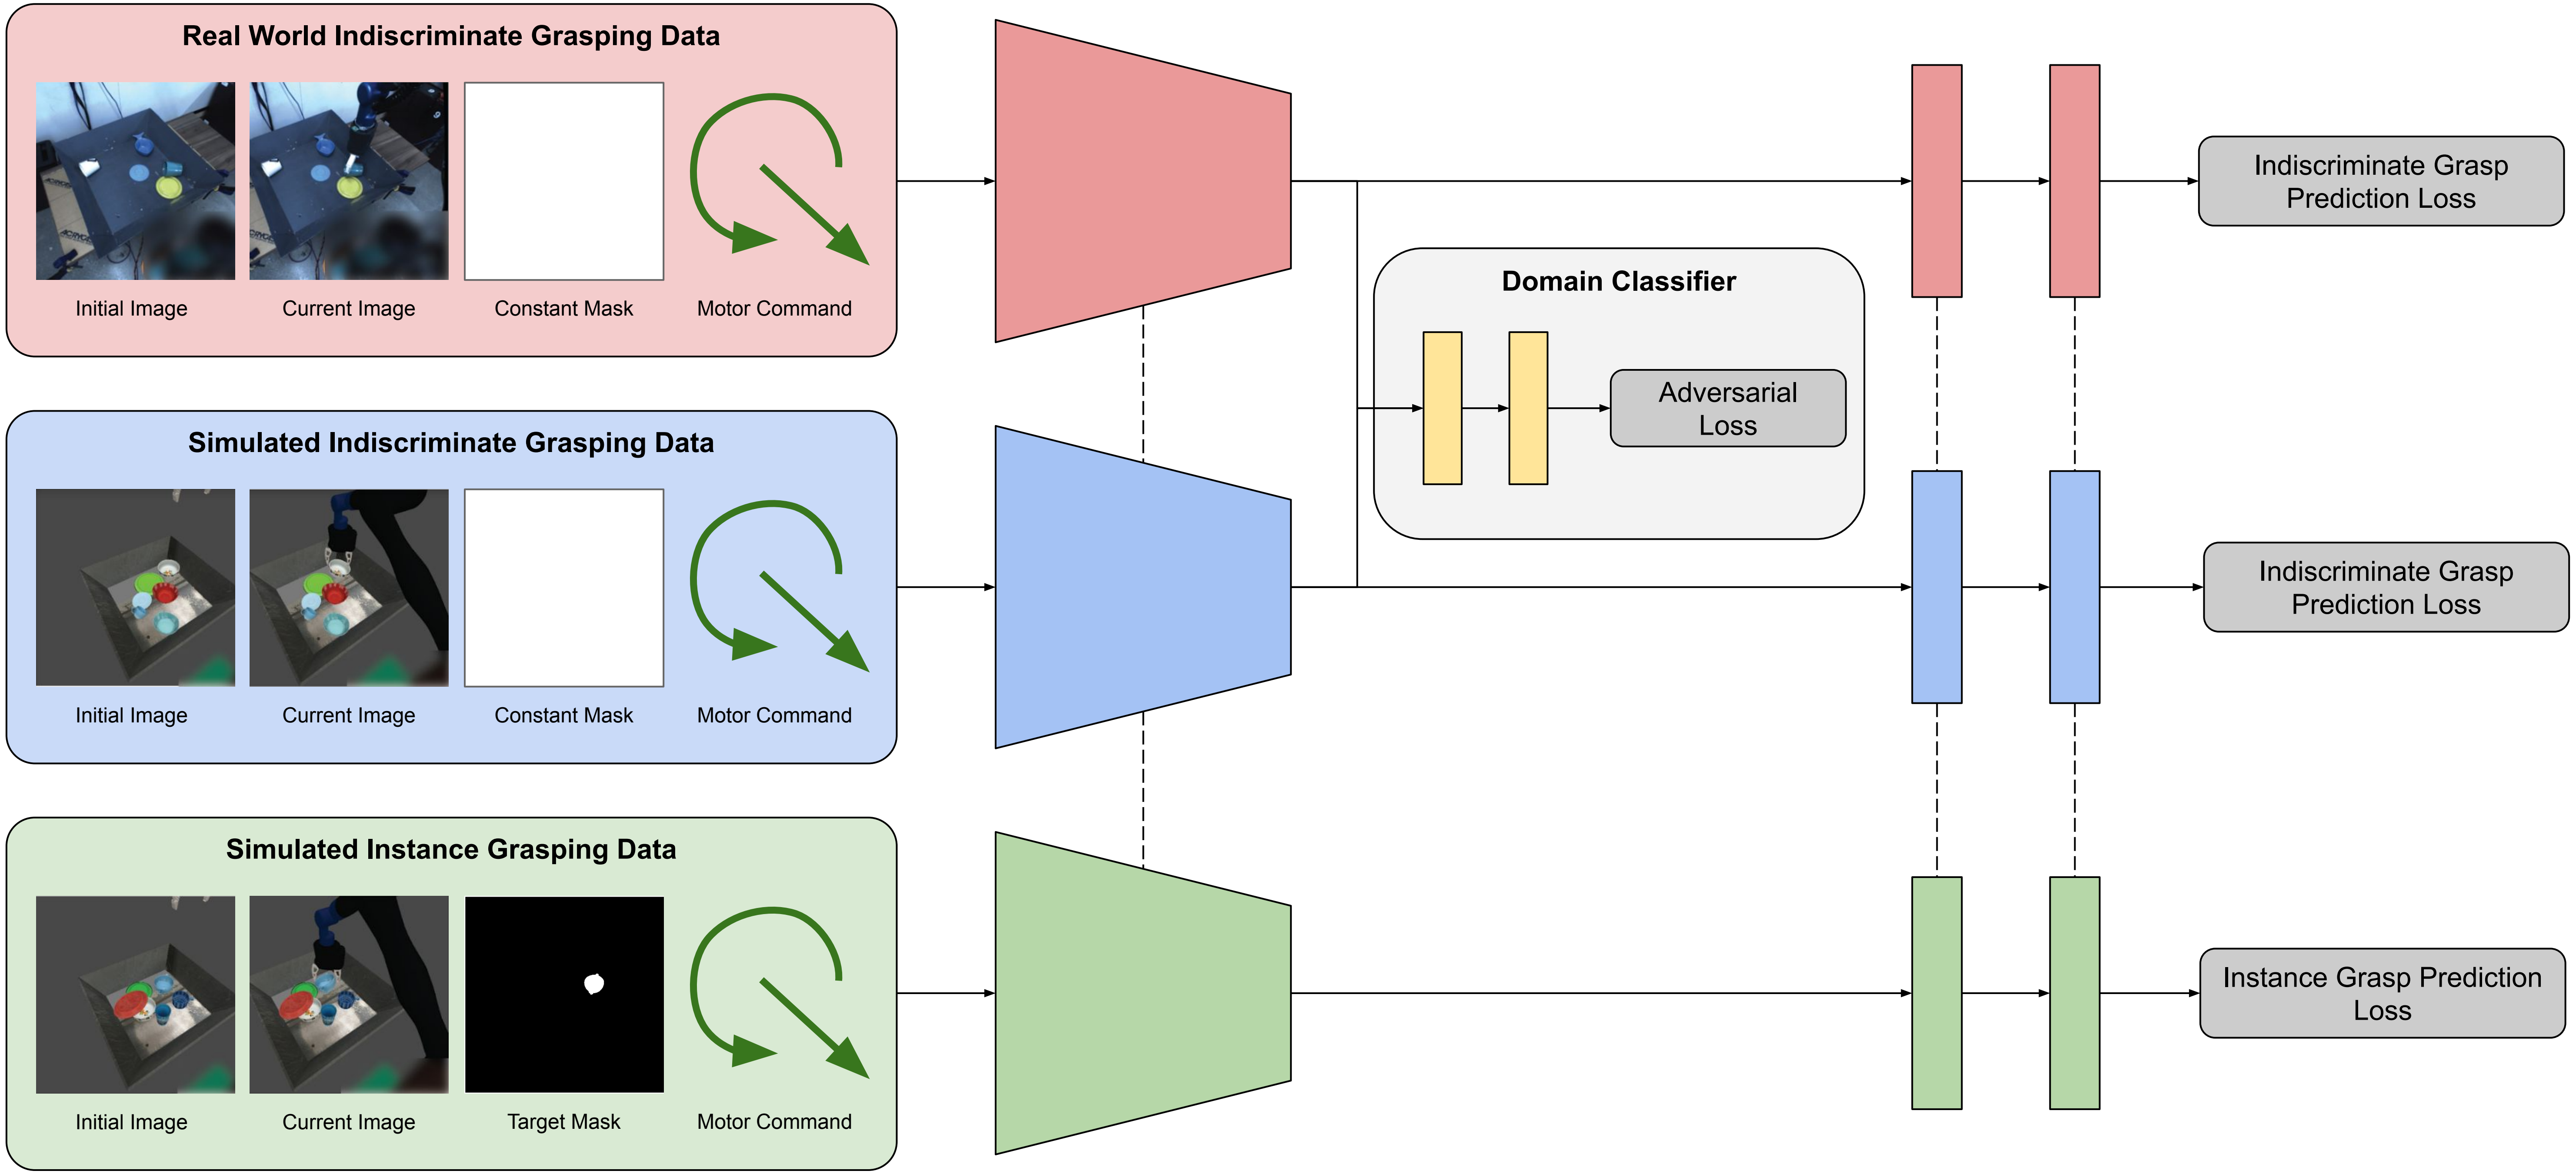

Evaluation in the Real World


Initial Image


Target Mask

6

Learning New Object Representation

“Data-Efficient Learning for Sim-to-Real Robotic Grasping using Deep Point Cloud Prediction Networks”,

Xinchen Yan, Mohi Khansari, Jasmine Hsu, Yuanzheng Gong, Yunfei Bai, Soren Pirk, Honglak Lee

Shape Prediction

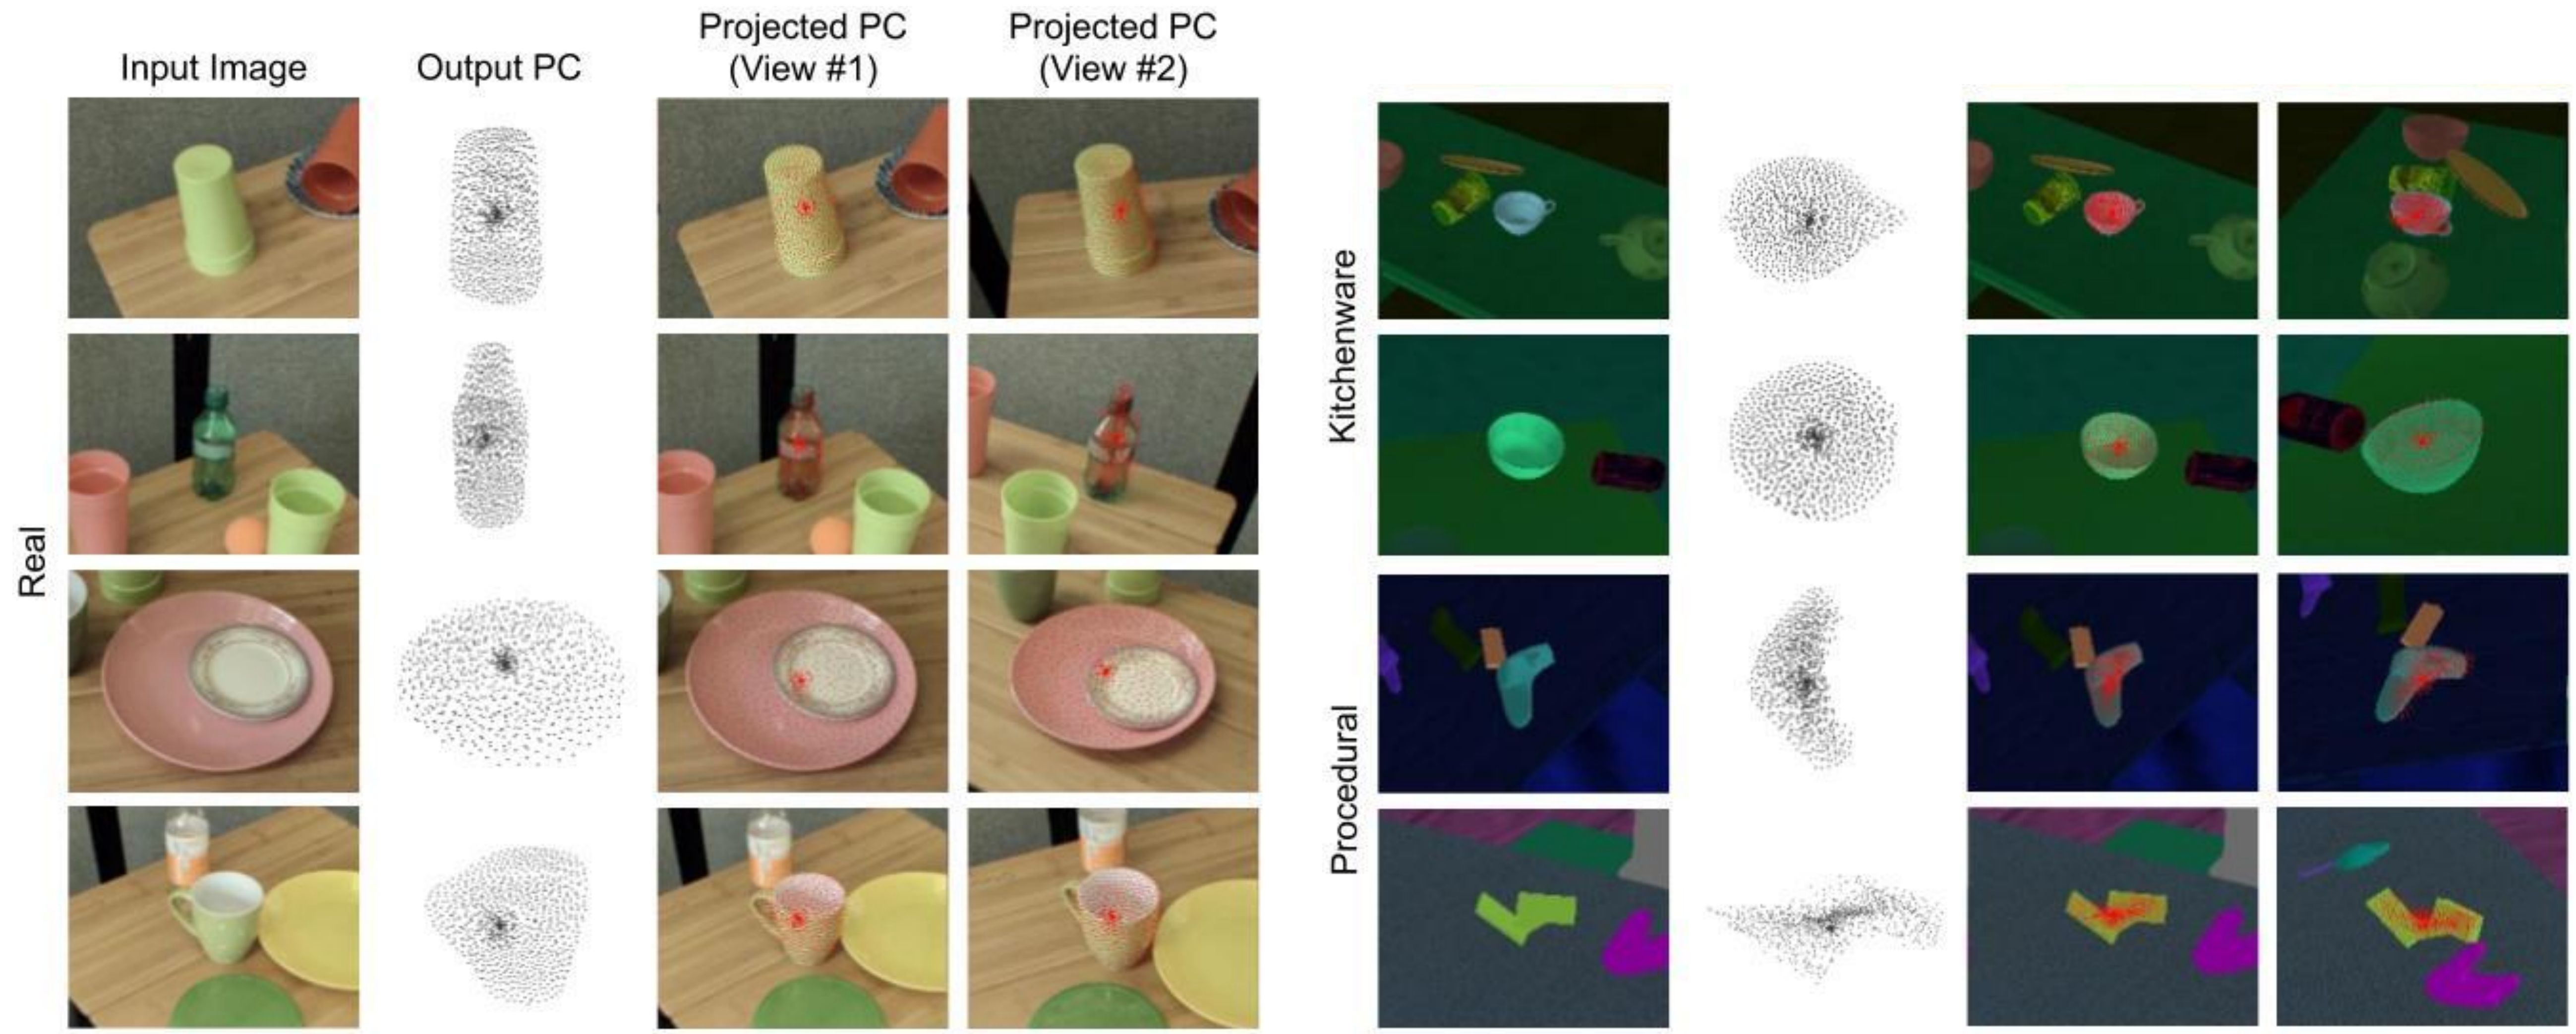

visualizations of point clouds generated with our point prediction network

Learning to Grasp Using Object Point Cloud

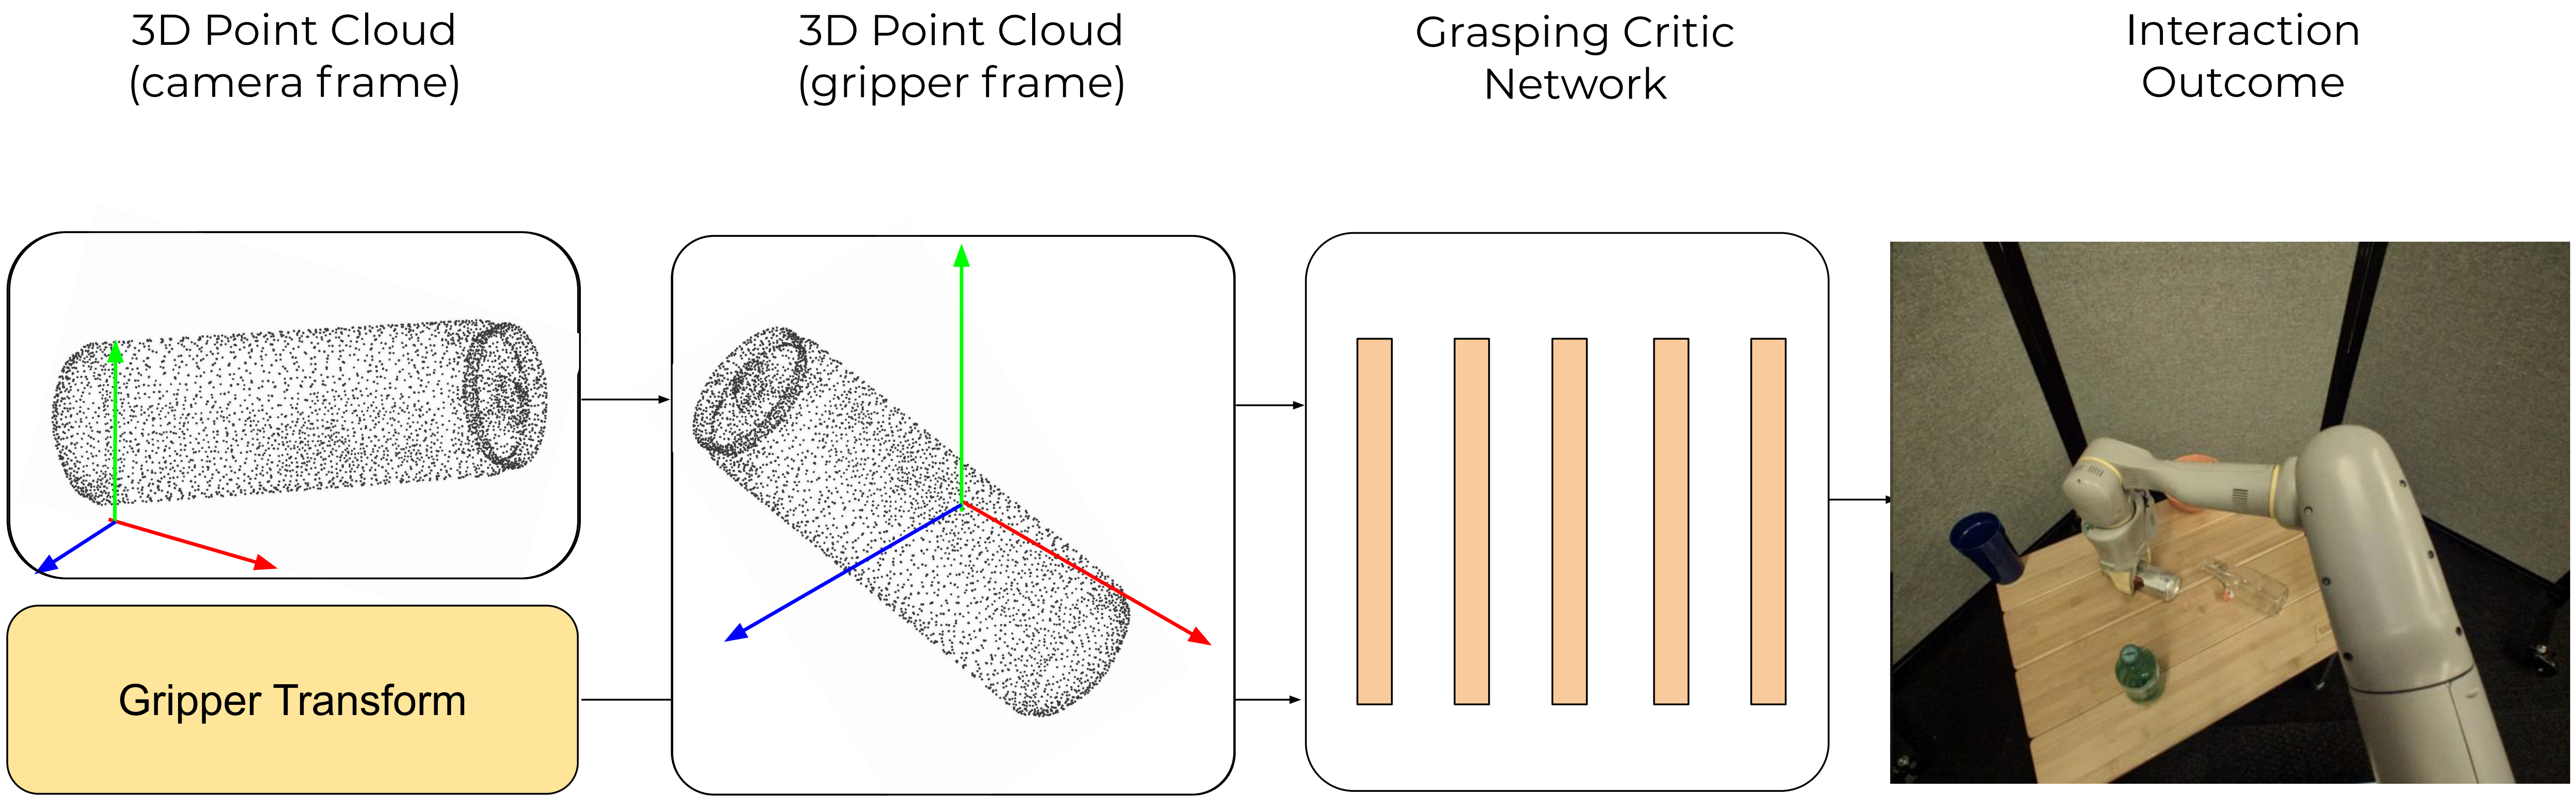

visualizations of point clouds generated with our point prediction network

Grasping Evaluations

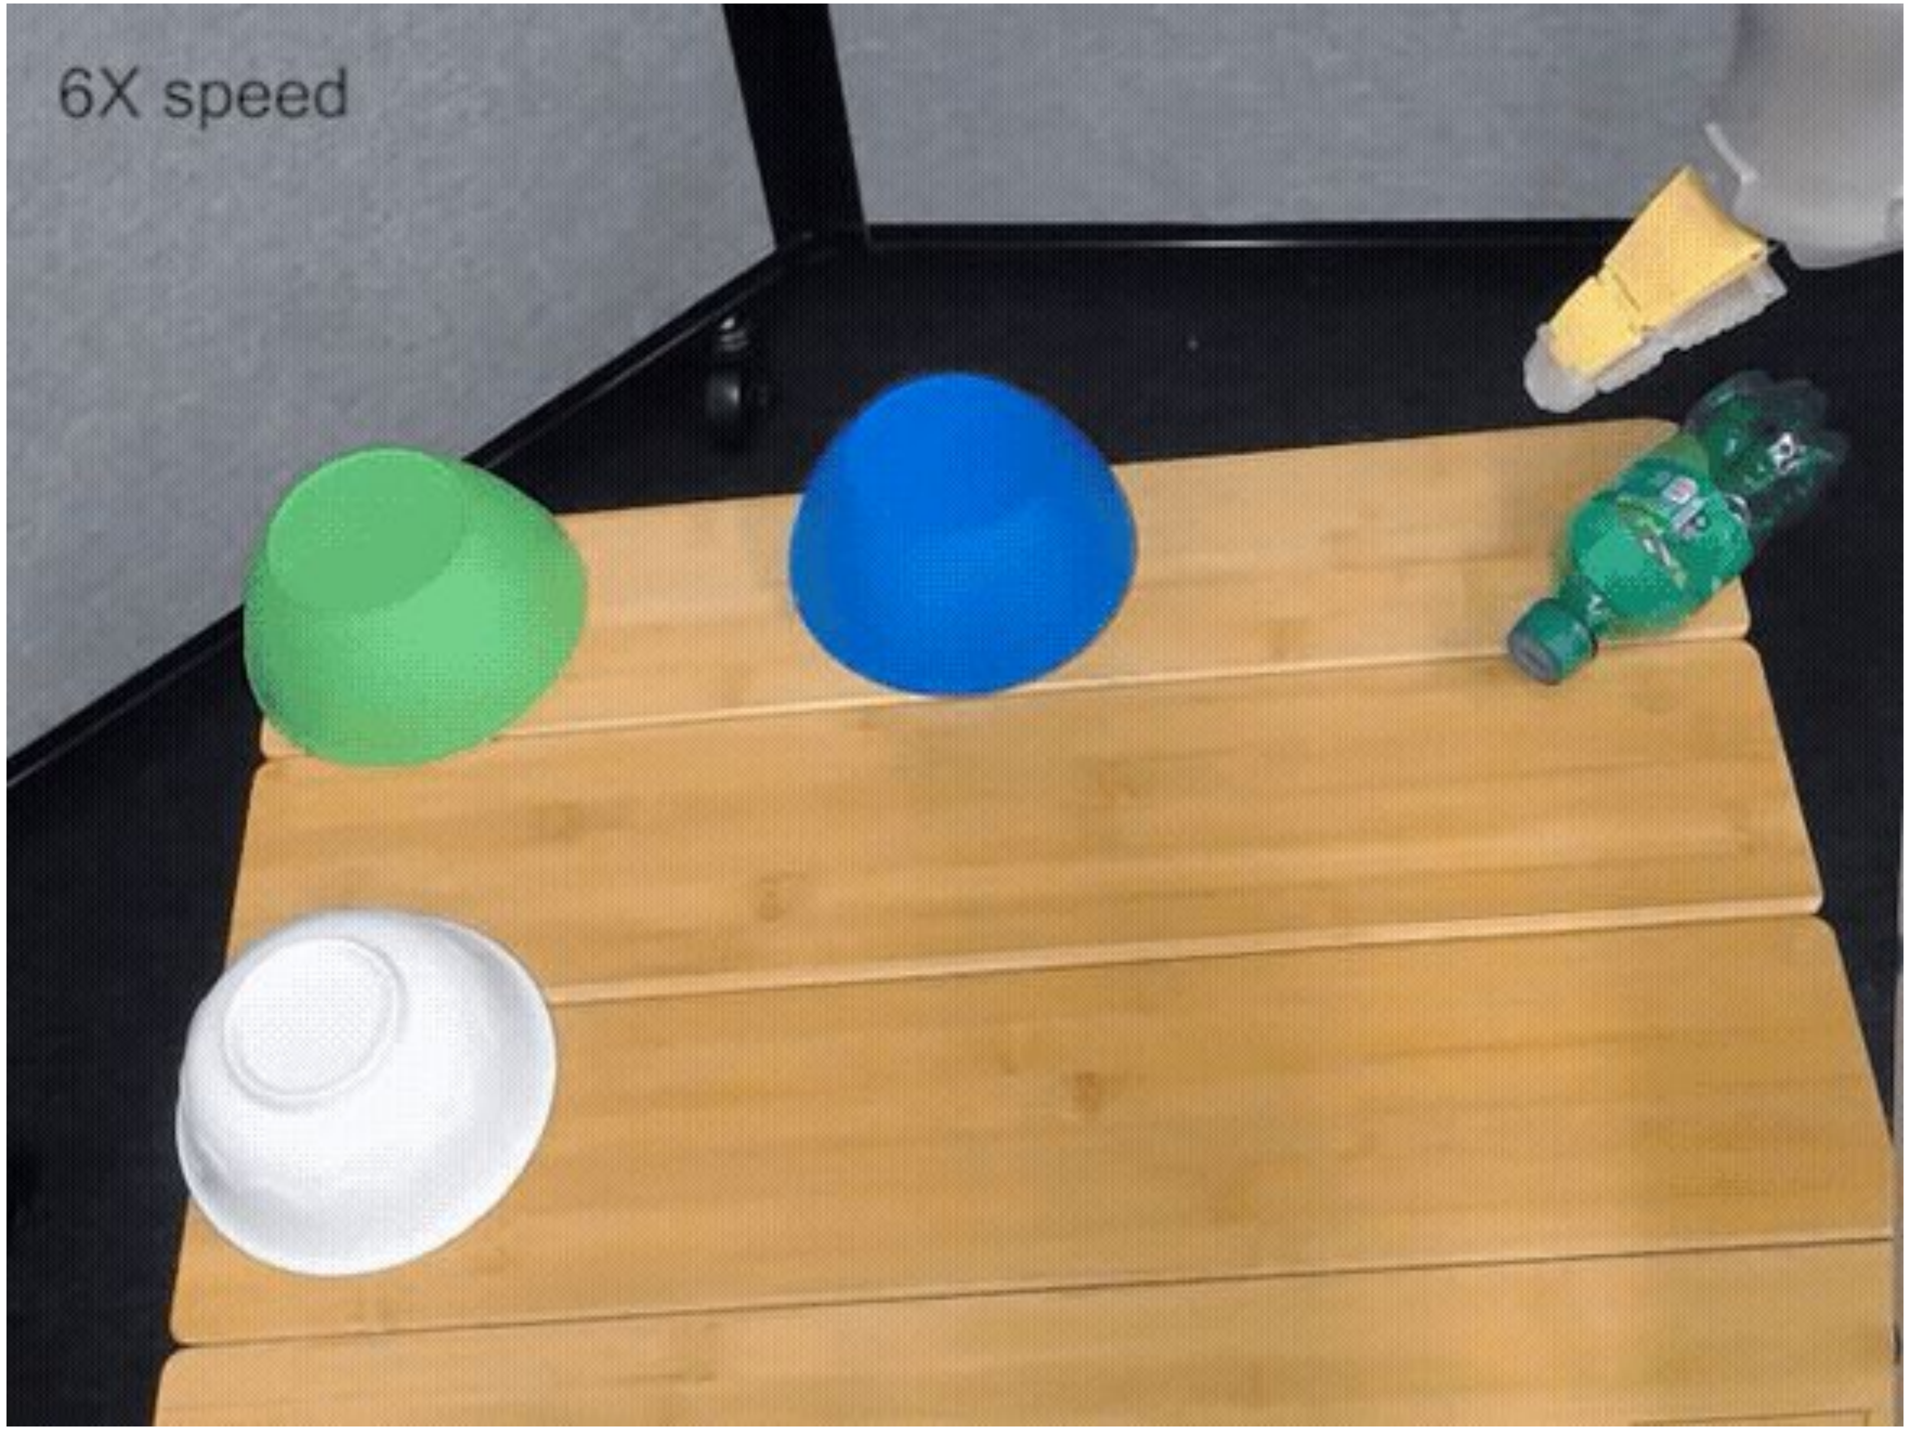

7

Conclusion

Summary

- Learn vision based grasping through **self-supervised learning** and **deep reinforcement learning**.

- Simulation helps reduce real world data requirements by **100x**, by solving sim-to-real transfer and **sim-to-sim transfer**.

- Simulation also enables us to learn **new related tasks**, and good **object representation** can facilitate sim-to-real transfer.

Reference

1. Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection
   S. Levine et al. IJRR 2017.

2. QT-Opt: Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation
   D. Kalashnikov et al. CoRL 2018.

3. Using Simulation and Domain Adaptation to Improve Efficiency of Deep Robotic Grasping
   K. Bousmalis et al. ICRA 2018.

4. Sim-to-Real via Sim-to-Sim: Data-efficient Robotic Grasping via Randomized-to-Canonical Adaptation Networks
   S. James et al. CVPR 2019.

5. Multi-Task Domain Adaptation for Deep Learning of Instance Grasping from Simulation
   K. Fang et al. ICRA 2018.

6. Data-Efficient Learning for Sim-to-Real Robotic Grasping using Deep Point Cloud Prediction Networks
   Under review

Credits

X

Yunfei Bai

Kuan Fang

Alexander Herzog

Stefan Hinterstoisser

Stephen James

Mrinal Kalakrishnan

Peter Pastor

Paul Wohlhart

Google Brain

Laura Downs

Ethan Holly

Julian Ibarz

Alex Irpan

Eric Jang

Dmitry Kalashnikov

Matthew Kelcey

Kurt Konolige

Alex Krizhevsky

Sergey Levine

Deirdre Quillen

Vincent Vanhoucke

DeepMind

Konstantinos Bousmalis

Raia Hadsell

Thank you!

Yunfei Bai
yunfeibai@google.com