SIMULATION TO REALITY TRANSFER IN ROBOTIC LEARNING

Stan Birchfield, Principal Research Scientist
Jonathan Tremblay, Research Scientist

GTC San Jose, March 2019

ROBOTICS AT NVIDIA

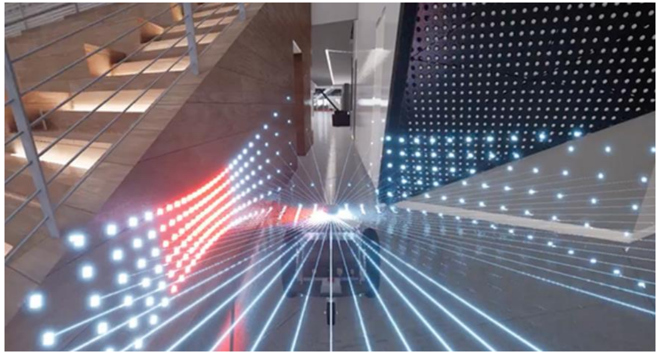

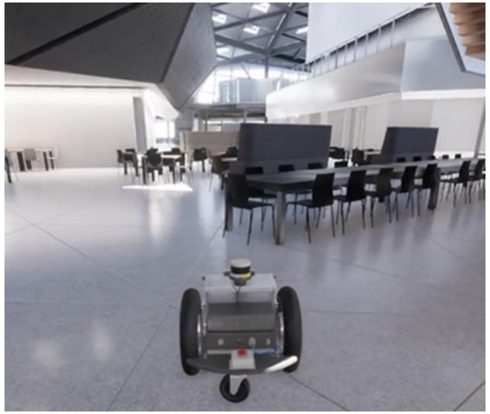

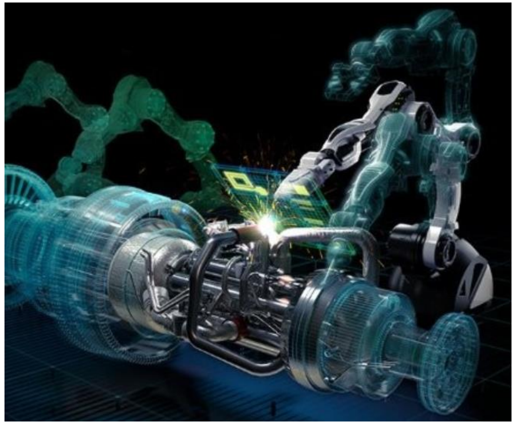

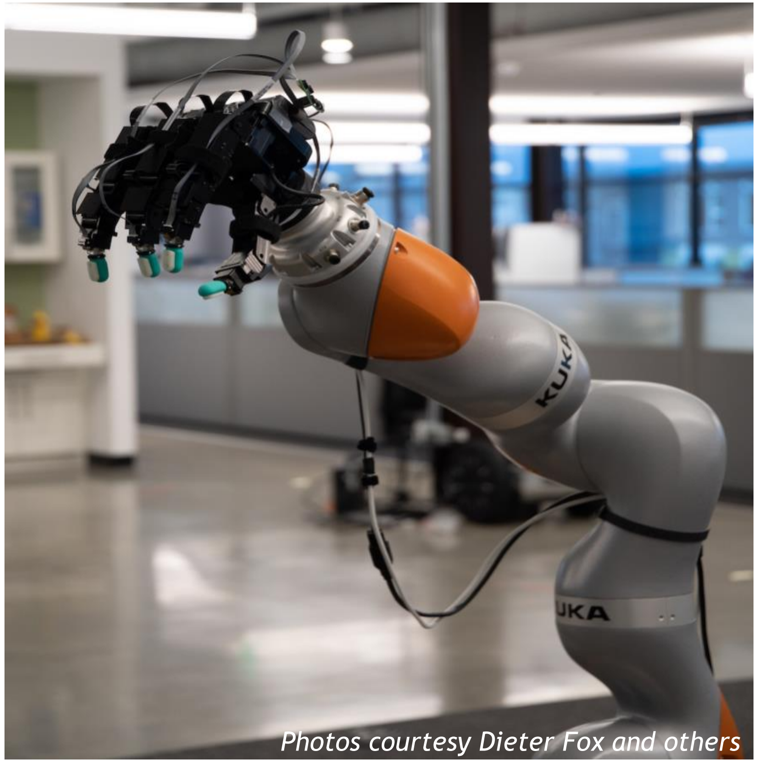

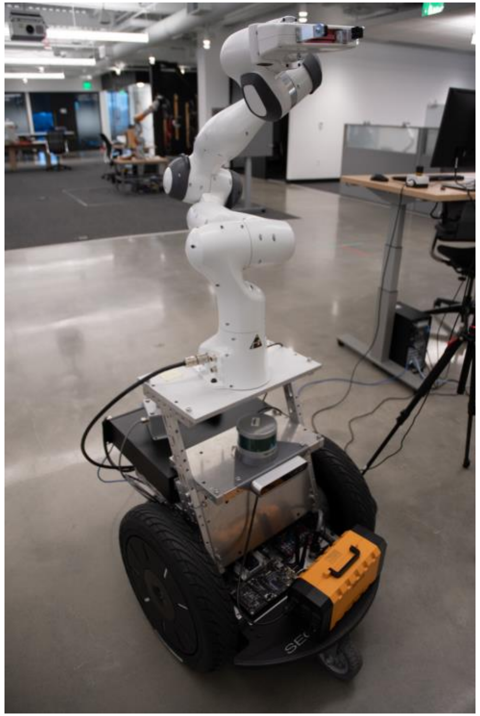

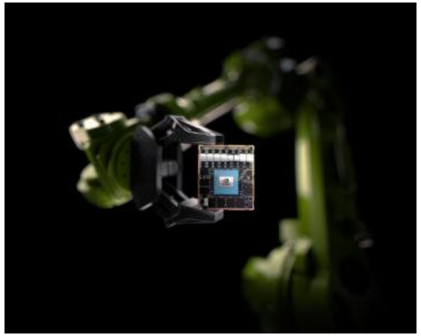

Photos courtesy Dieter Fox and others

OUR MISSION

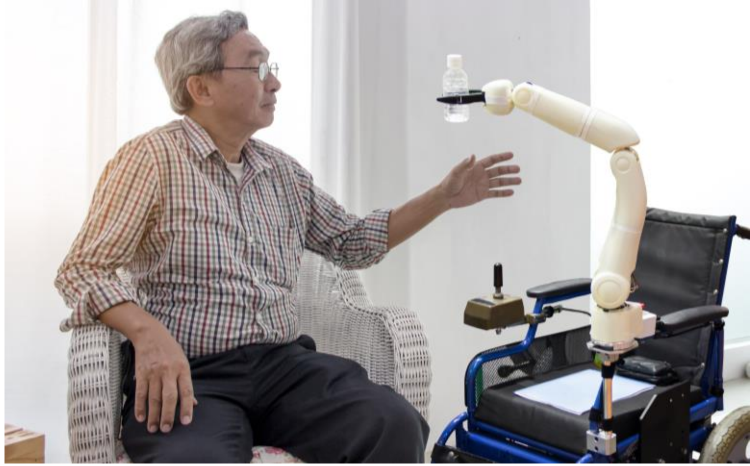

Drive breakthrough robotics research and development

Enable the next-generation of robots that safely work alongside humans, transforming industries such as

- manufacturing,
- logistics,
- healthcare,
- and more

Photo: Courtesy of Charlie Kemp/Georgia Tech

Slide courtesy Dieter Fox

CURRENT STATE OF ROBOTICS TECHNOLOGY

Navigation for fulfillment, delivery, assembly

Applications focus on

- getting from A to B without collision
- following specific trajectory

Slide courtesy Dieter Fox

HOW DO WE GET

FROM

TO

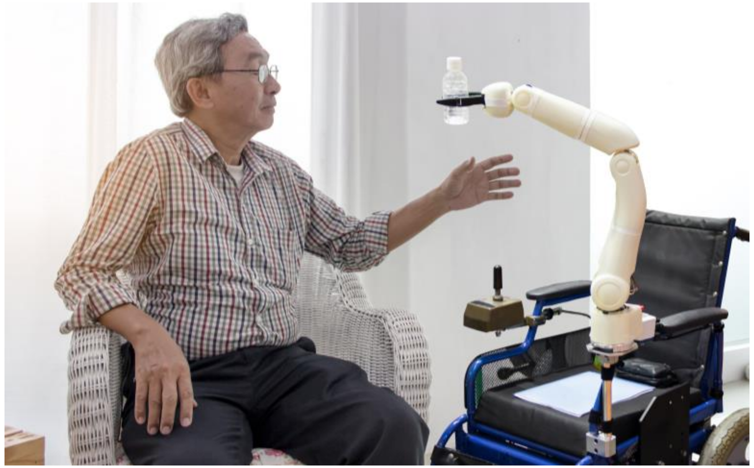

?

Natural user interfaces?

Compliant motion?

Better perception?

End-to-end learning?

Planning algorithms?

Tactile sensing?

Dexterous hands?

Cheaper H/W?

DEEP LEARNING REVOLUTION

Already happening

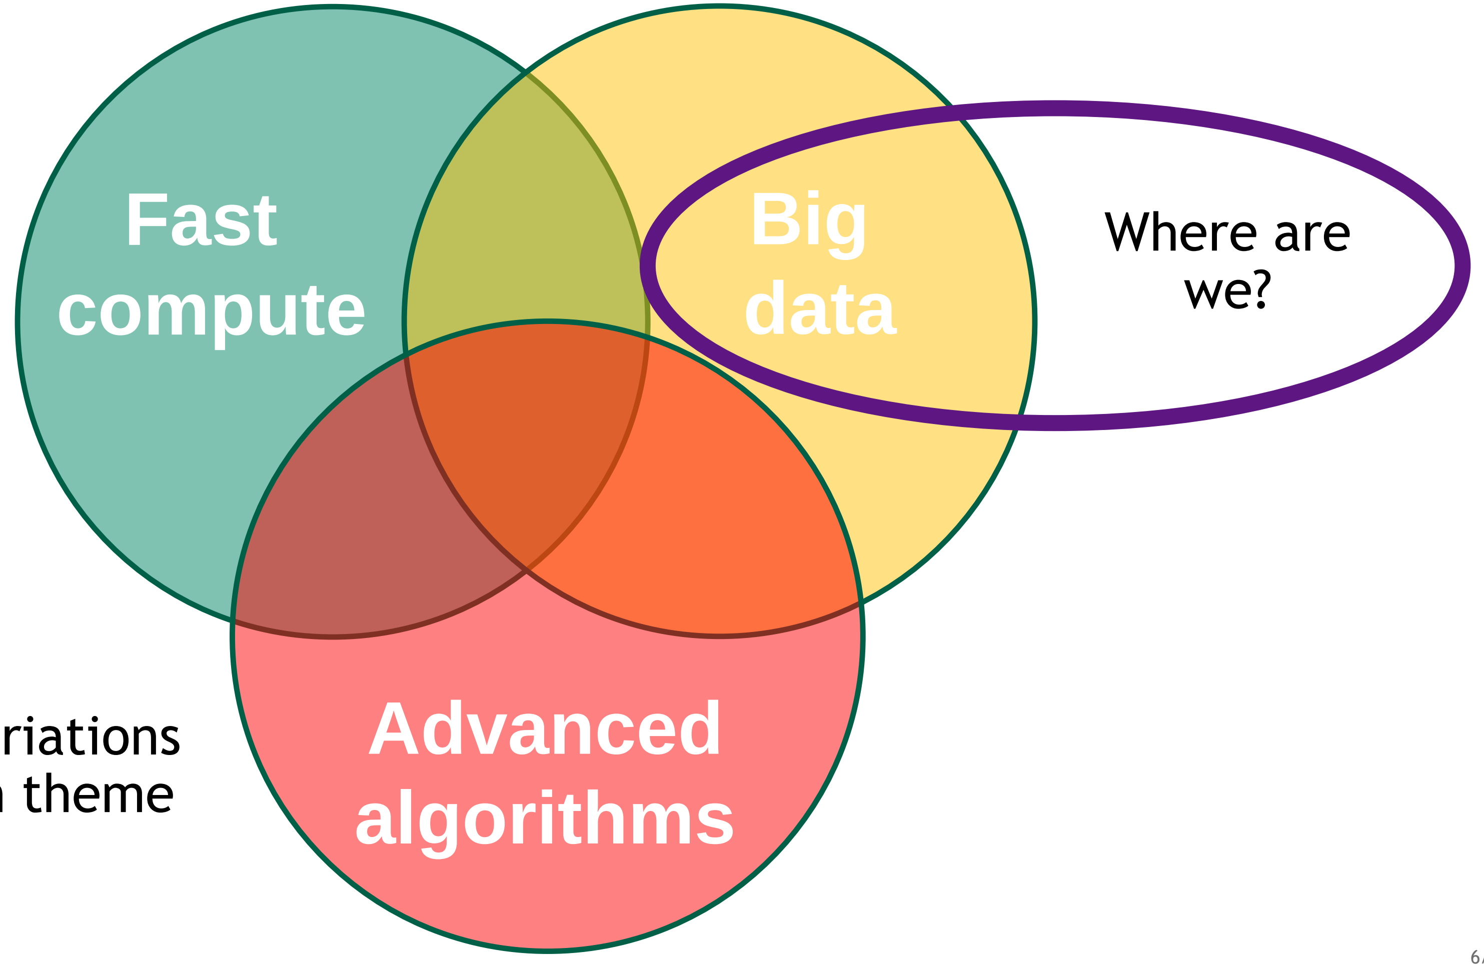

Variations on theme

VISION DATASETS

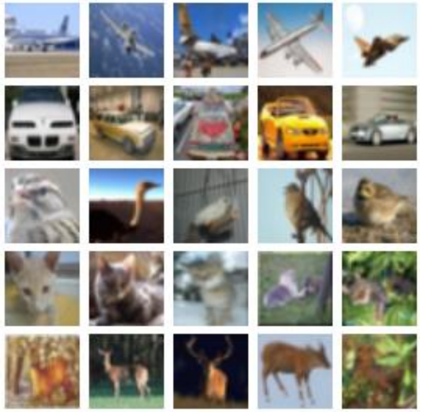

CIFAR
120k images

COCO
200k images

Pascal 3D+
30k images

Sintel
50k images

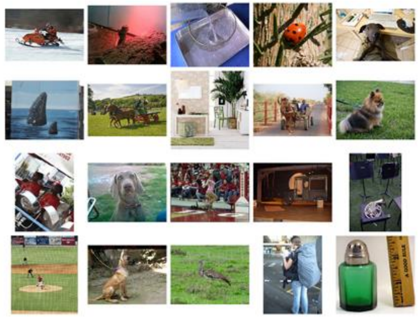

ImageNet
14M images
1M bounding boxes

T-LESS
50k images

RBO
90k images

ObjectNet3D
90k images

FlyingThings3D
20k images

• • •

ROBOTICS DATASETS

2D-3D-S

Penn Haptic Texture Toolkit
100 models

KITTI

iCubWorld

MPII Cooking

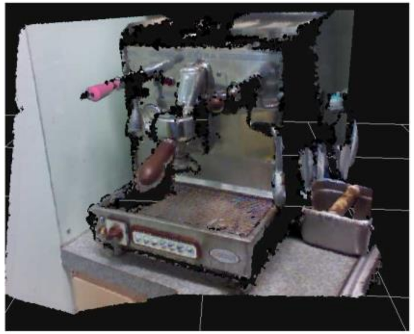

Robobarista
1k demonstrations

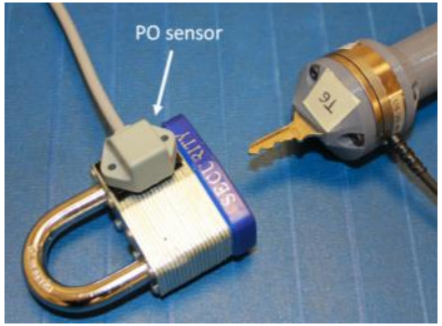

USF Manipulation
2k trials

SLAM

ScanNet

RoboTurk
2k demonstrations

UNIPI Hand
114 grasps

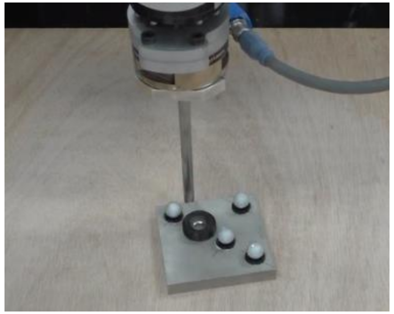

MIT Push
1M datapoints

SIMULATED ACTIONABLE ENVIRONMENTS

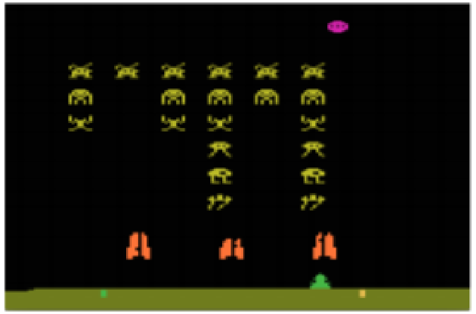

Arcade Learning Environment

Gibson

AI2-THOR

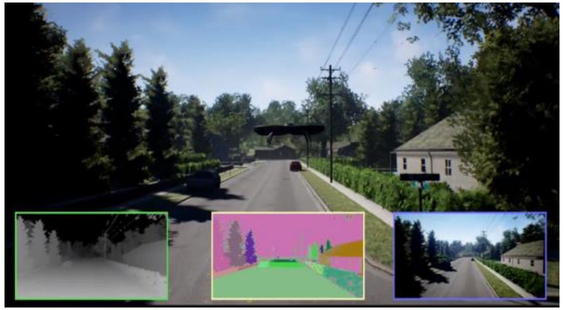

AirSim

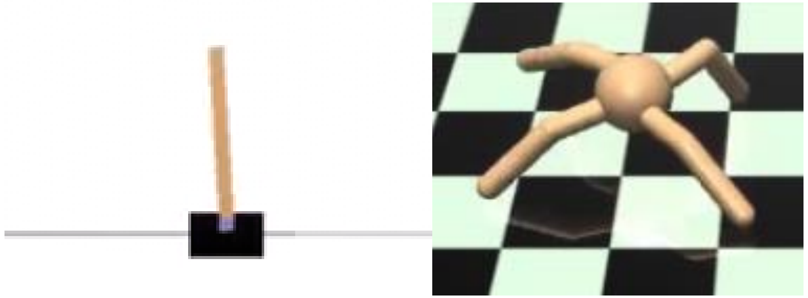

OpenAI Gym

Roboschool

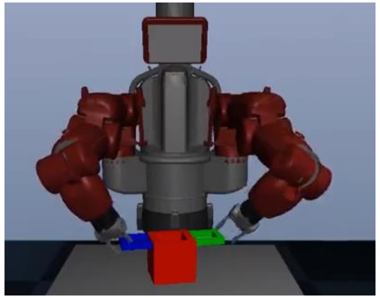

SURREAL

SIMULATION

Will simulation be ***the key*** that unlocks robot potential?

Three possibilities:

1. Simulation will ***never be good enough*** to be used
   "Software simulations are doomed to succeed." — Rod Brooks
2. Without simulation, interesting robotics problems ***cannot be*** solved
3. Eventually, simulation will mature to the point where
   1. Robotics will ***benefit*** from it (accelerate training, validate solutions, etc.)
   2. Some problems may ***require*** it due to their complexity

Simulation generates massive data with high consistency

AN ANALOGY

Then (Leslie Jones Collection/Boston Public Library)

Now (Public domain)

AN ANALOGY

(Photo by SuperJet International. CC BY-SA 2.0)

Training

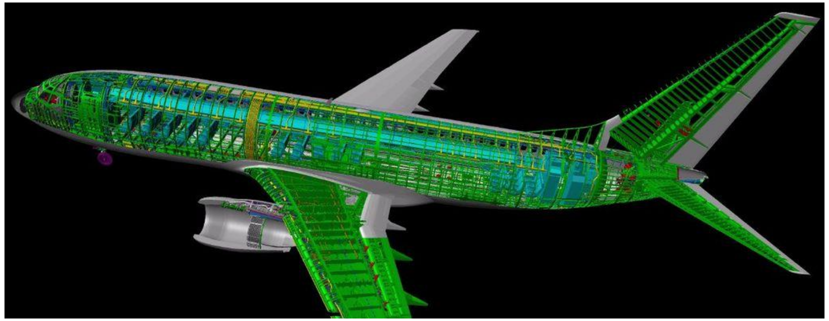

(Photo by Prana Fistianduta. CC BY-SA 3.0)

Design

(Photo by Marian Lockhart / Boeing)

Support

DEMOCRATIZATION

NVIDIA

PROBLEM STATEMENT

environment

Photorealistic

observations

agent

Physically realistic

actions

Simulation

Train

$$\pi : o \rightarrow a$$

Reality

Apply

LONG WAY TO GO

Early flight simulator
1983

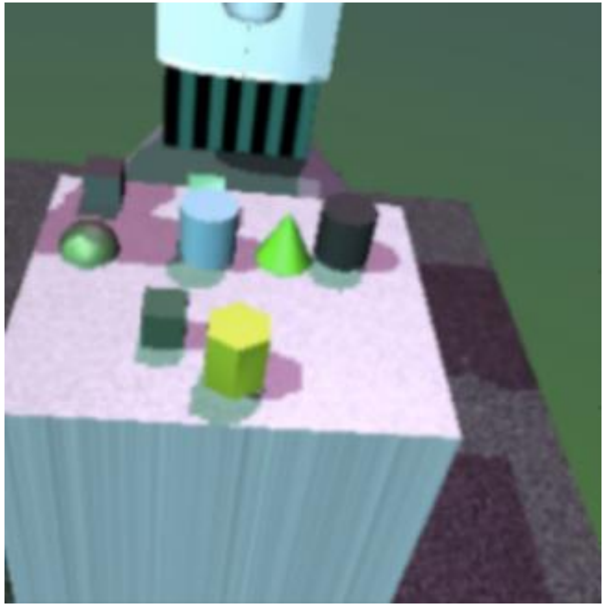

Early robot simulator
2017
[Tobin et al. 2017]

Today's robot simulators:

- Not photorealistic
- Not physically realistic

BUT PROGRESSING FAST

Photorealism
RTX ray tracing

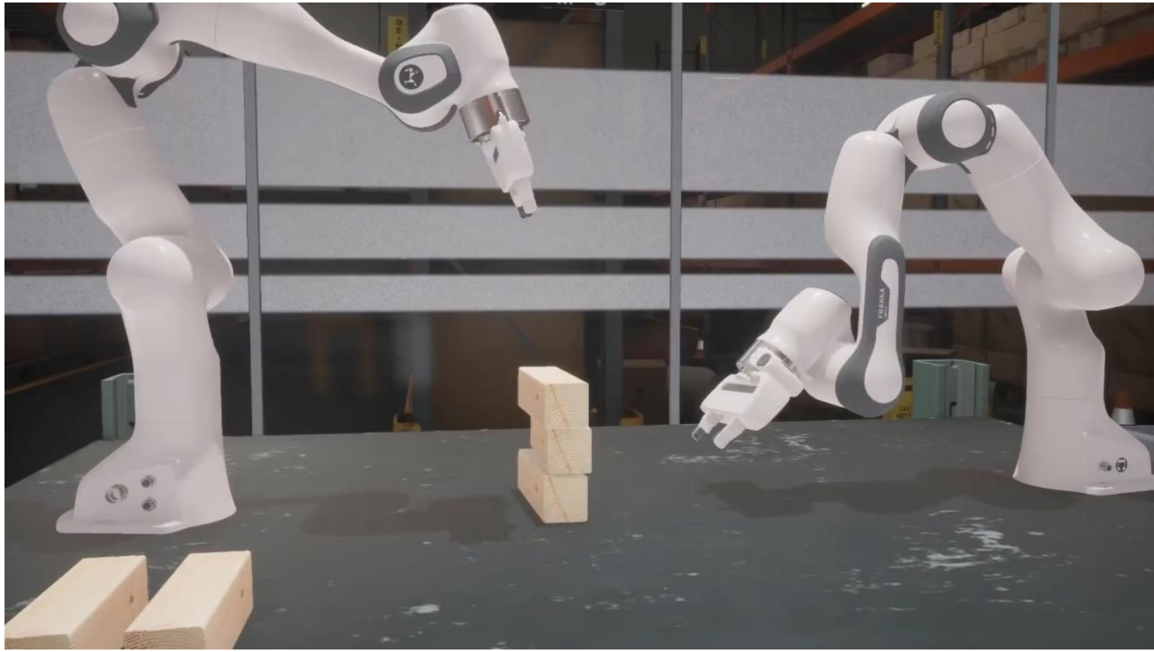

Physical realism
PhysX 4.0

REALITY GAP

Reality gap – discrepancy between simulated data and real data

Three ways to bridge reality gap:

1. Increase fidelity of simulator
    1. *Photo-realism* (light, color, texture, material, scattering, …;
       also tactile sensors, …)
    2. *Physical realism* (dimensions, forces, friction, collisions, …)
2. Learn mapping to bridge the gap
   Domain adaptation

[Dundar et al., 2018]

3. Make controller robust to imperfections
   Domain randomization, add noise during training, stochastic policy

SIM-TO-REAL SUCCESS

Locomotion

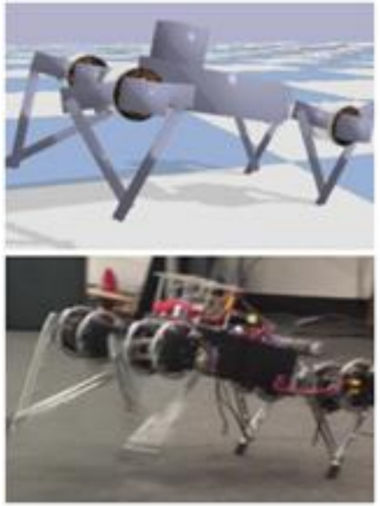

[Tan et al., 2018]

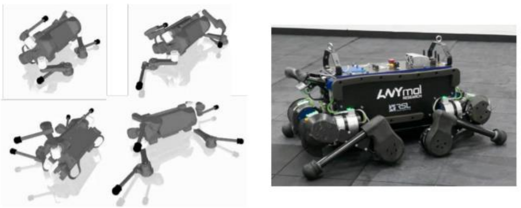

[Hwangbo et al., 2019; Lee et al., 2019]

Grasping / Manipulation

[James et al., 2017; Matas et al., 2018]

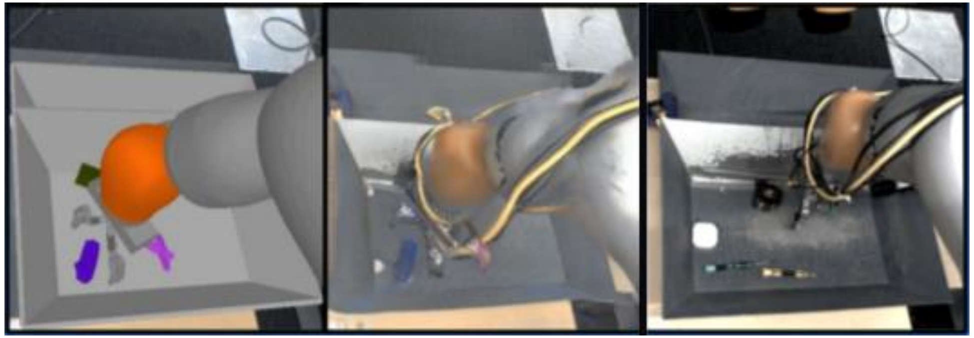

[Bousmalis et al., 2018]

Quadrotor flight

[Molchanov et al. 2019]

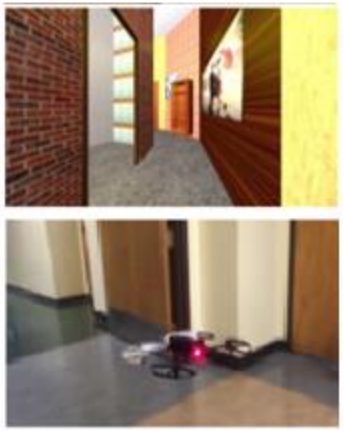

[Sadeghi et al. 2017]

SIM-TO-REAL AT NVIDIA

	Vision	Closed-loop control
Navigation		
Manipulation		

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

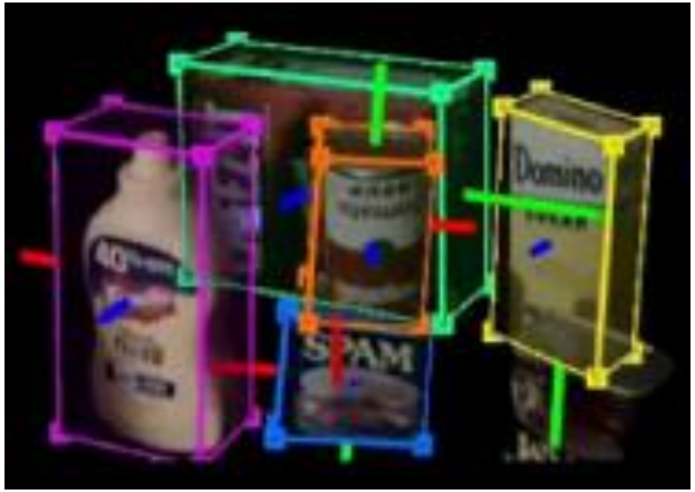

Vision

Closed-loop control

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

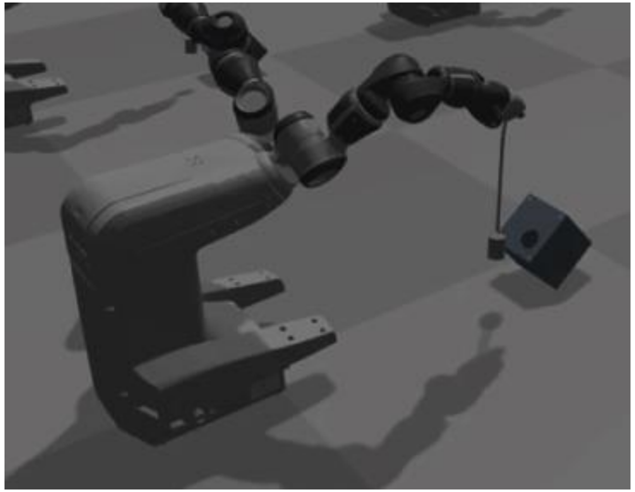

Vision

Closed-loop control

DOMAIN RANDOMIZATION

Domain randomization – Generate non-realistic randomized images

Idea – If enough variation is seen at training time, then real world will just look like another variation

Randomize:

- Object pose
- Lighting / shadows
- Textures
- Distractors
- Background

Training Deep Networks with Synthetic Data: Bridging the Reality Gap by Domain Randomization
J. Tremblay, A. Prakash, D. Acuna, M. Brophy, V. Jampani, C. Anil, T. To, E. Cameracci, S. Boochoon, S. Birchfield. *CVPR WAD 2018*

STRUCTURED DOMAIN RANDOMIZATION (SDR)

SDR – Generate randomized images with variety (as in DR) but with realistic structure

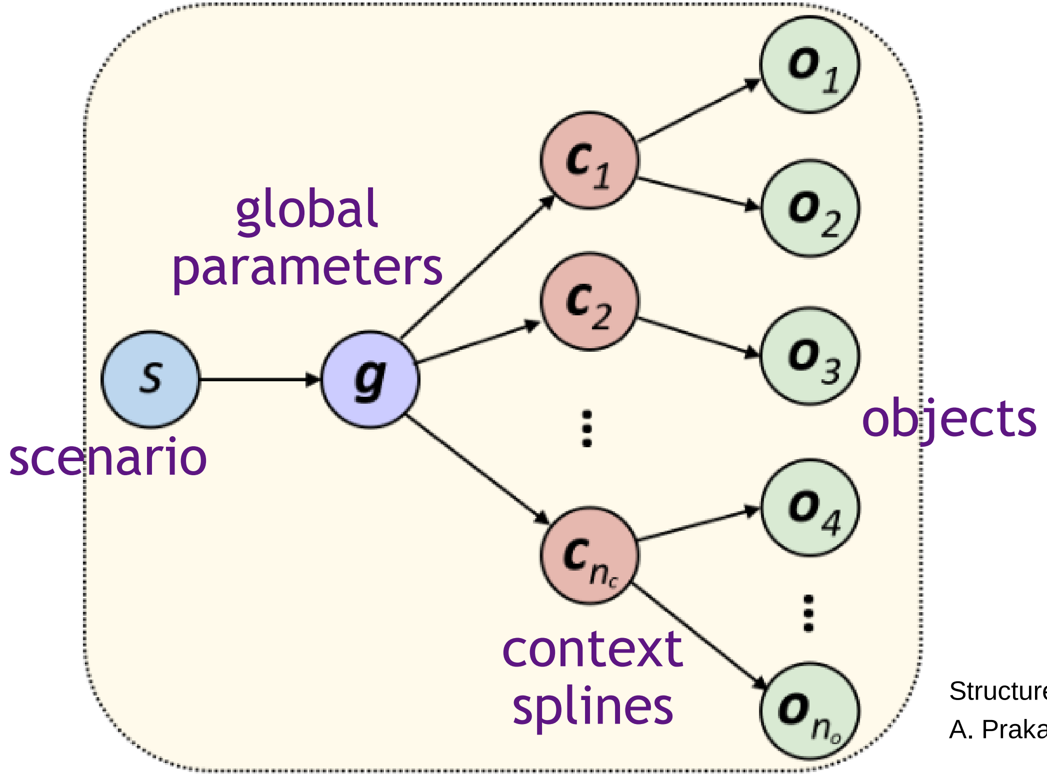

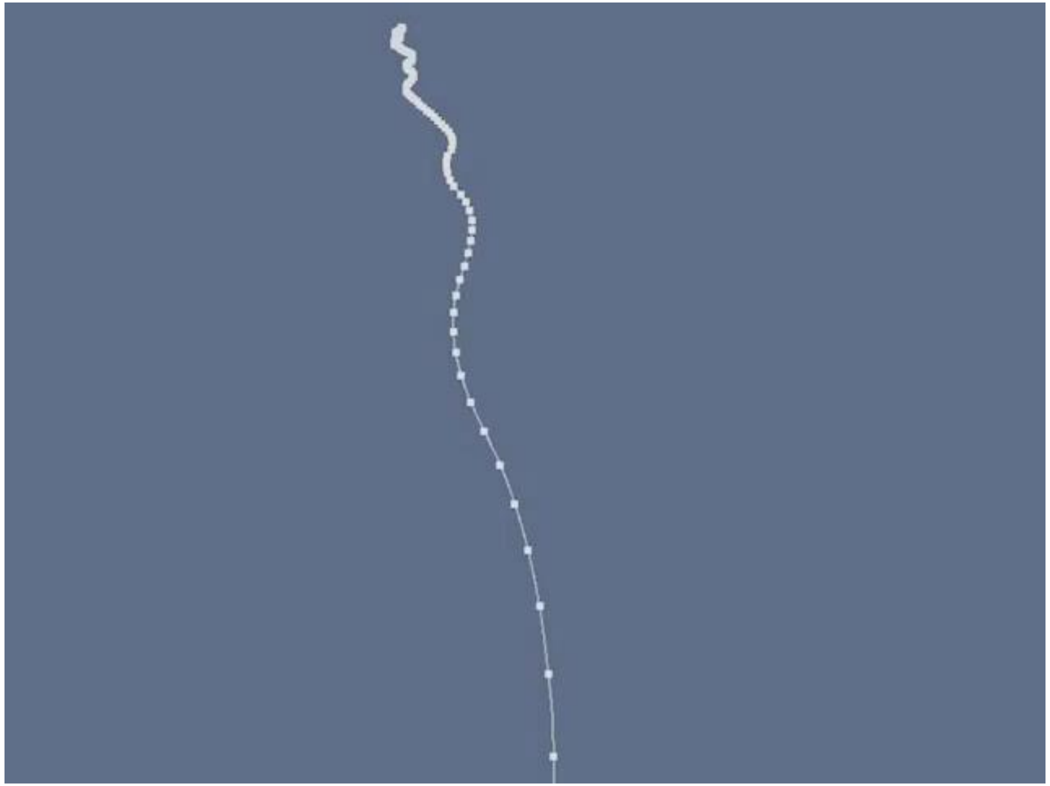

Structured Domain Randomization: Bridging the Reality Gap by Context-Aware Synthetic Data
A. Prakash, S. Boochoon, M. Brophy, D. Acuna, E. Cameracci, G. State, O. Shapira, S. Birchfield. *ICRA 2019*

SDR IMAGES

Not photorealistic, but structurally realistic

SDR RESULTS

Reality gap is large

Domain gap between real datasets is also large

SDR *25k* outperforms:

- *DR* *25k* (synthetic)
- *Sim 200k* (photorealistic synthetic)
- *VKITTI* *21k* (photorealistic synthetic with same content)
- *BDD100K* (real)

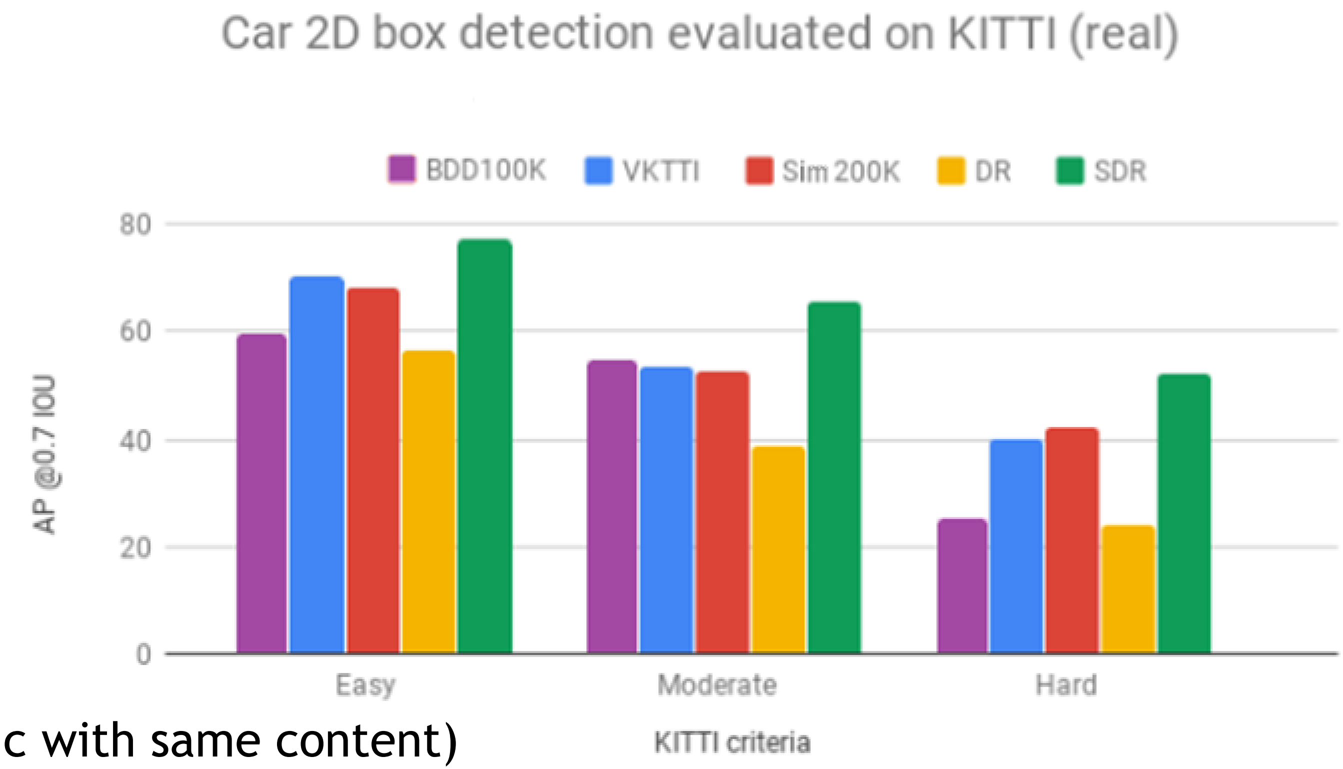

SDR RESULTS

KITTI

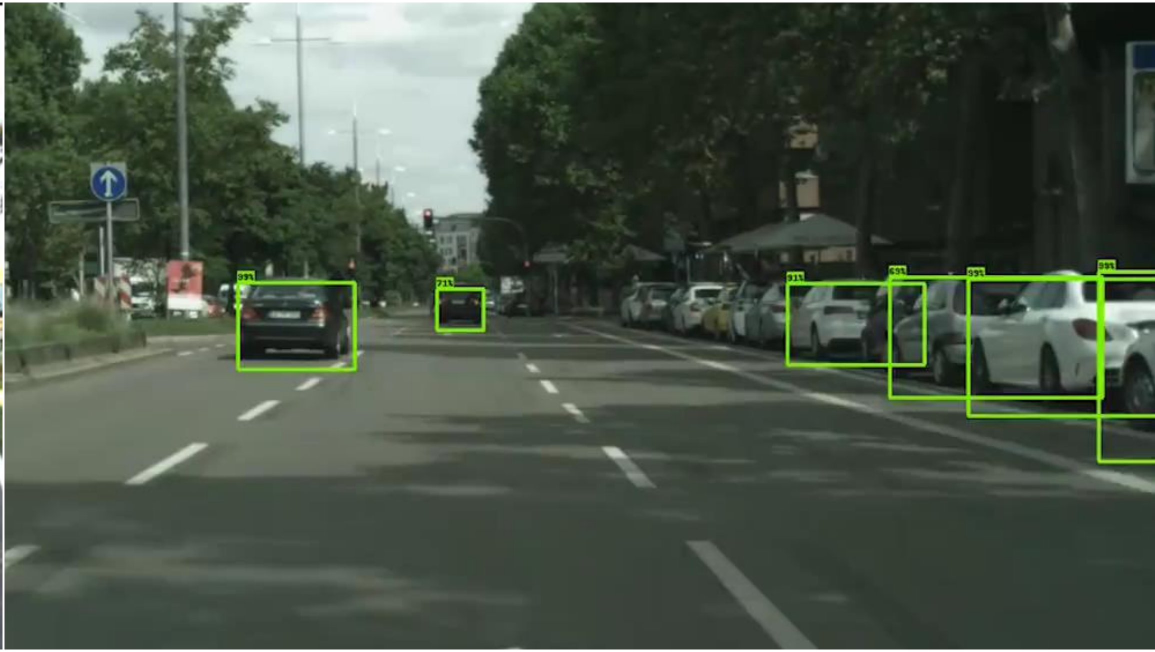

Cityscapes

Network has never seen a real image!

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

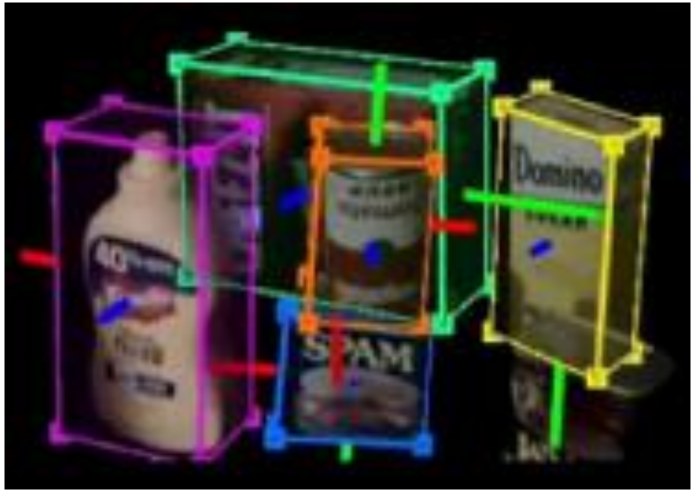

Vision

Closed-loop control

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

Vision

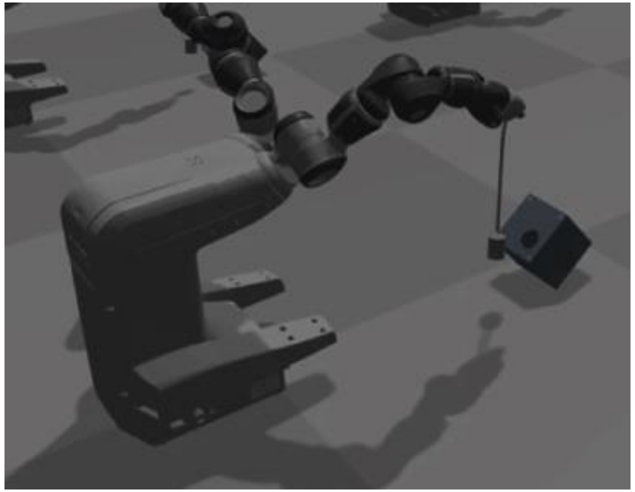

Closed-loop control

DRIVE SIM AND CONSTELLATION

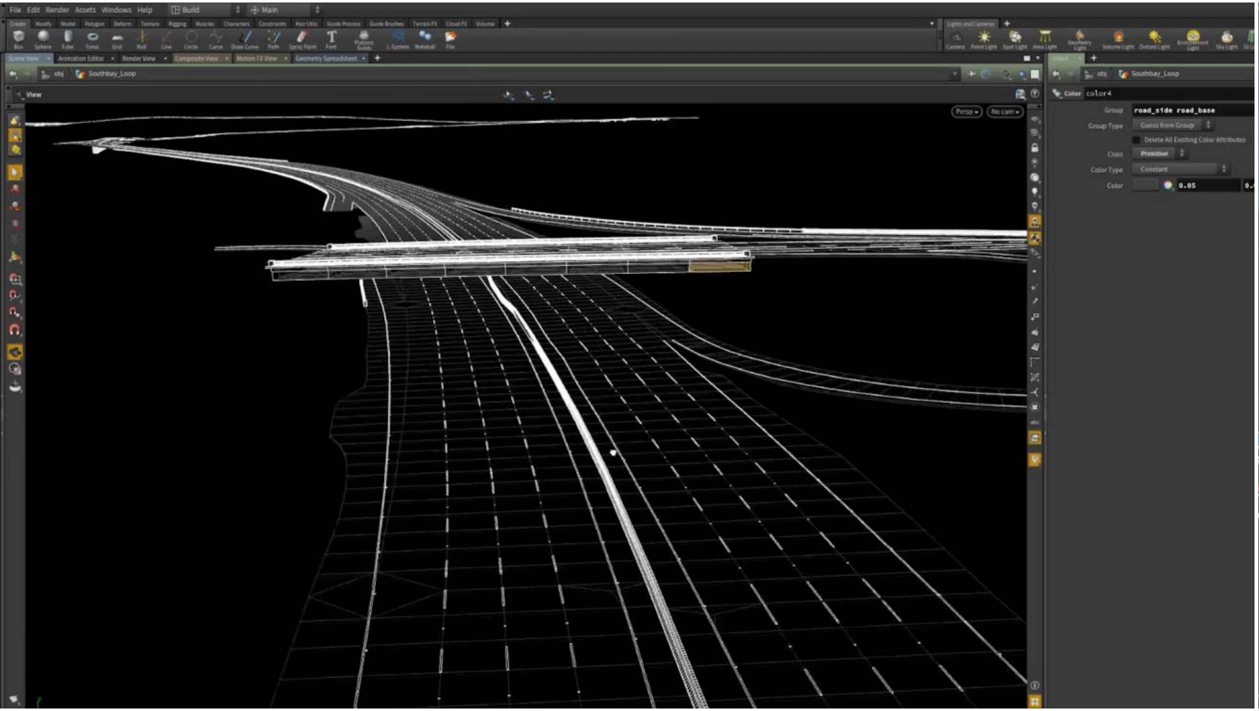

DRIVE Sim creates the virtual world

DRIVE Constellation runs simulation

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

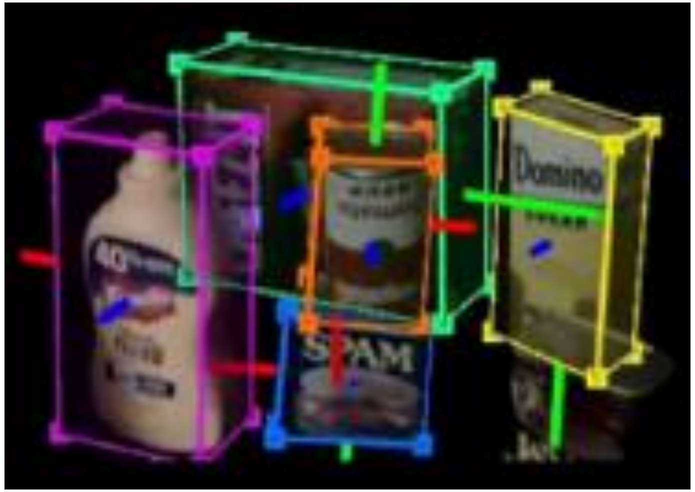

Vision

Closed-loop control

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

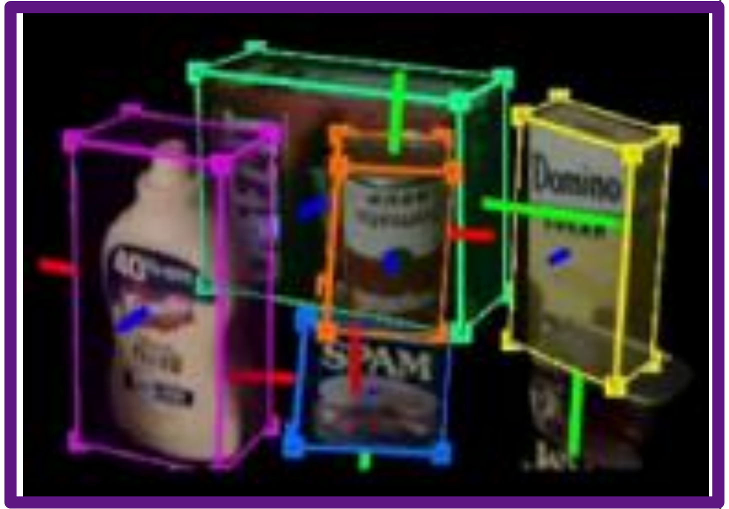

Vision

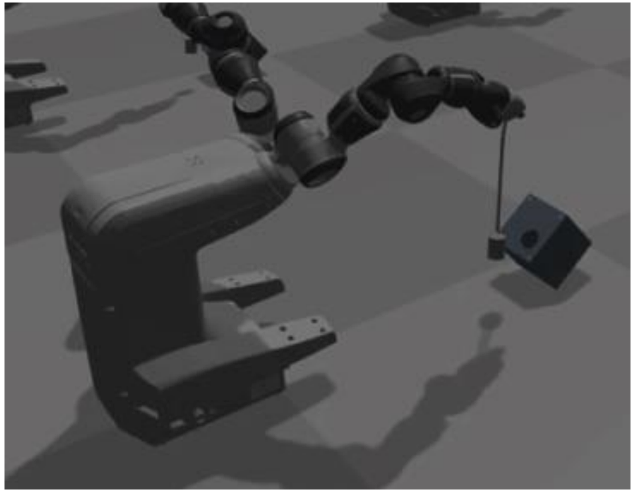

Closed-loop control

LEARNING HUMAN-READABLE PLANS

"Place the car on yellow."

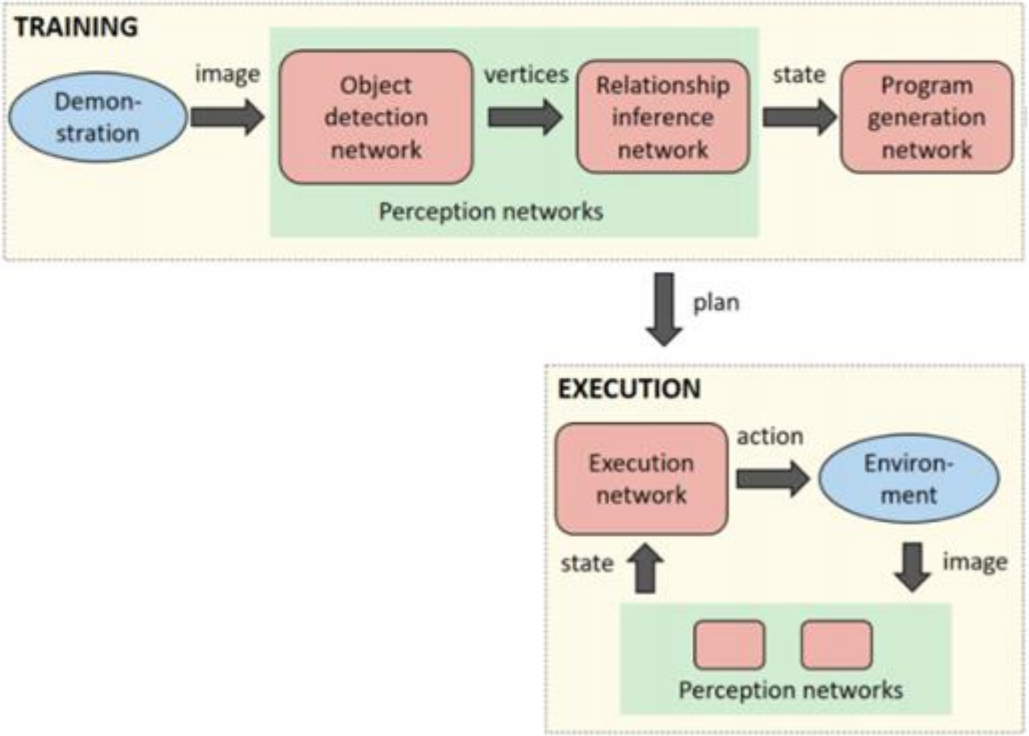

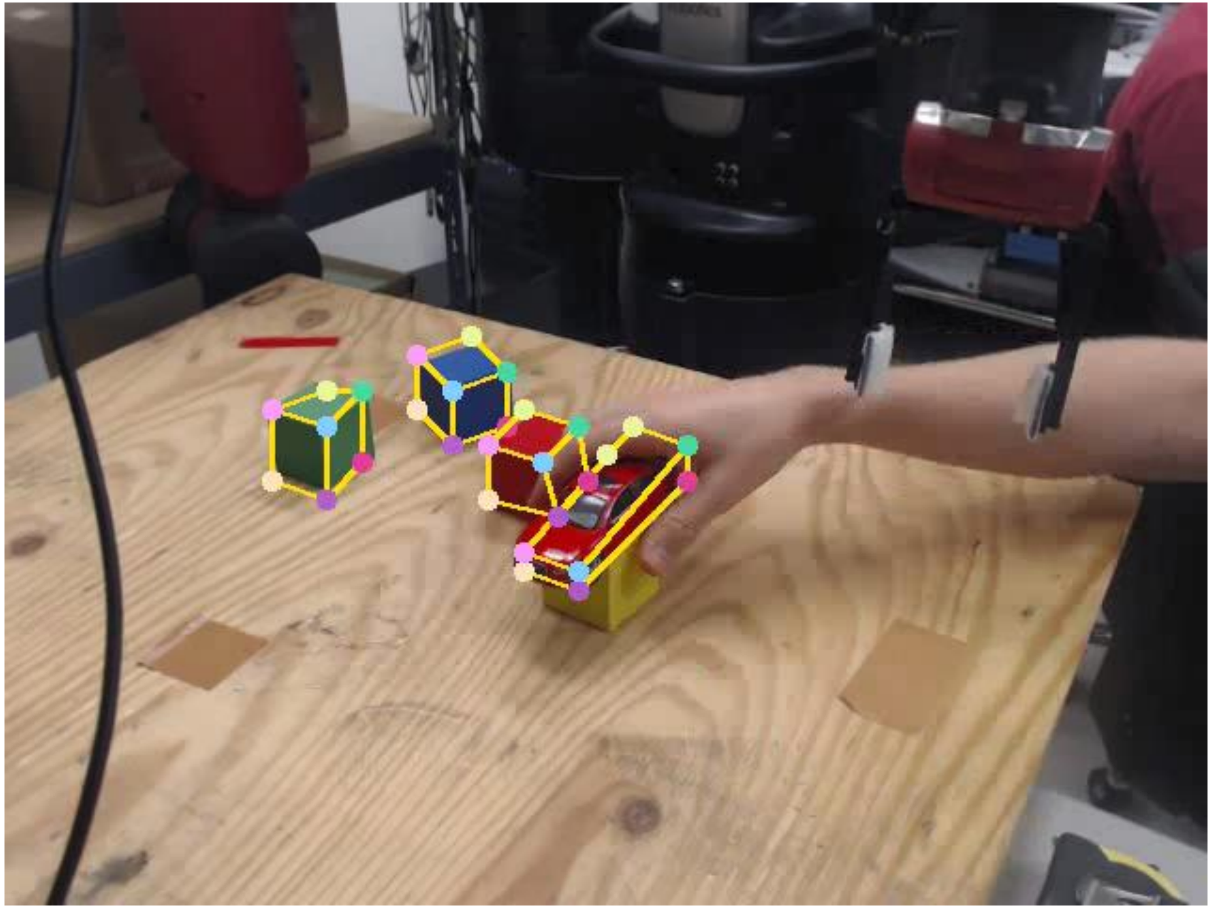

Synthetically Trained Neural Networks for Learning Human-Readable Plans from Real-World Demonstrations
J. Tremblay, T. To, A. Molchanov, S. Tyree, J. Kautz, S. Birchfield. *ICRA 2018*

DETECTING HOUSEHOLD OBJECTS

Does the technique generalize?

YCB objects [Calli et al. 2015]; subset of 21 used by PoseCNN [Xiang et al. 2018]

Baxter gripper

- parallel jaw
- 4 cm travel dist.

DEEP OBJECT POSE ESTIMATION (DOPE)

Design goals:

1. **Single RGB image**
2. **Multiple instances of each object type**
3. **Full 6-DoF pose**
4. **Robust to pose, lighting conditions, camera intrinsics**

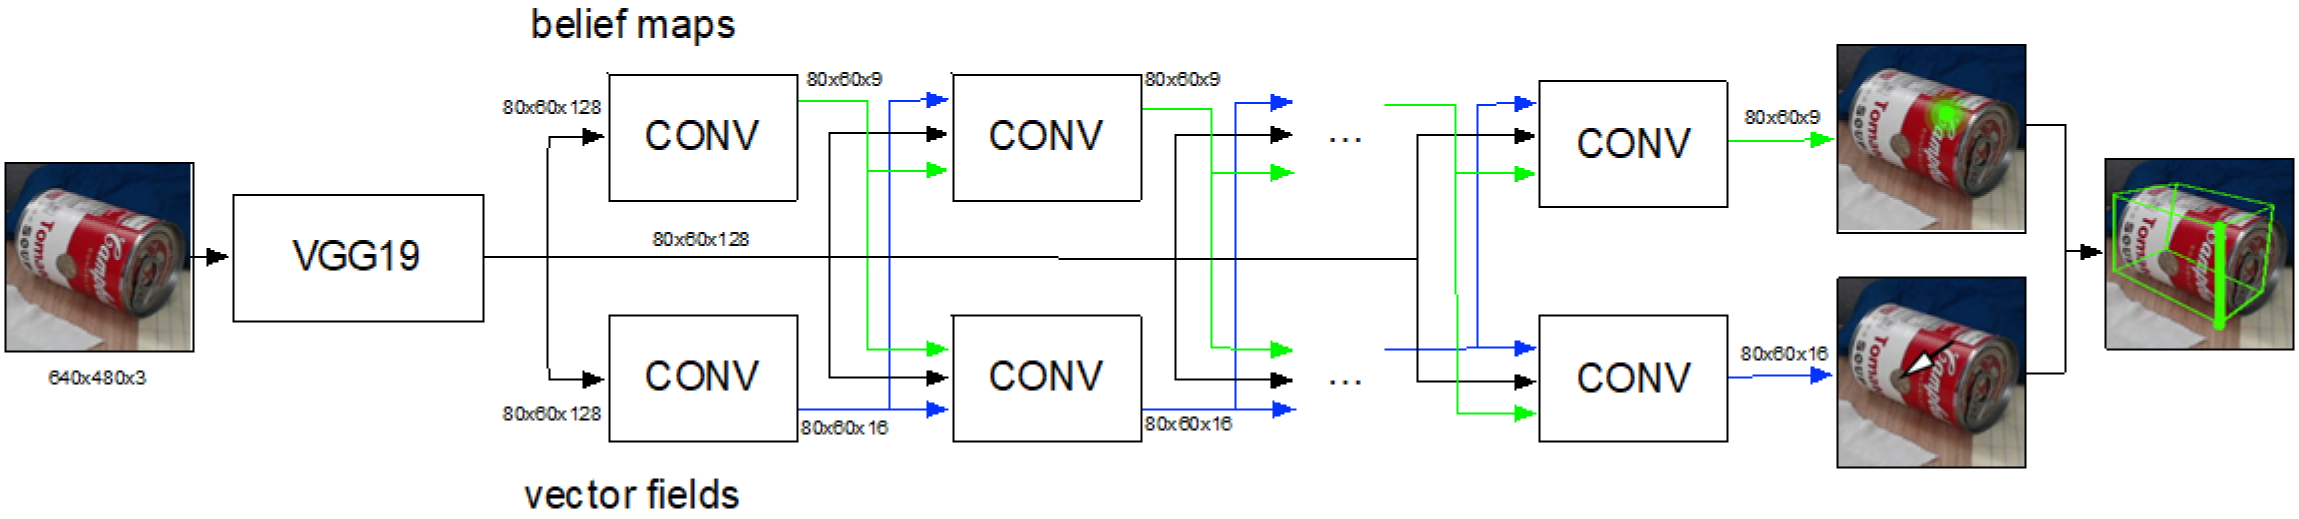

https://github.com/NVlabs/Deep_Object_Pose

Deep Object Pose Estimation for Semantic Robotic Grasping of Household Objects
J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, S. Birchfield. *CoRL 2018*

NDDS DATA SET SYNTHESIZER

- Data exporter using UE4
- Near photorealistic
- Domain randomization tool set
- Tutorial and documentation
- Export:
    - 2D bounding box
    - 3D pose
    - Keypoint location
    - Segmentation
    - Depth

https://github.com/NVIDIA/Dataset_Synthesizer

MIXING DR + PHOTOREALISTIC

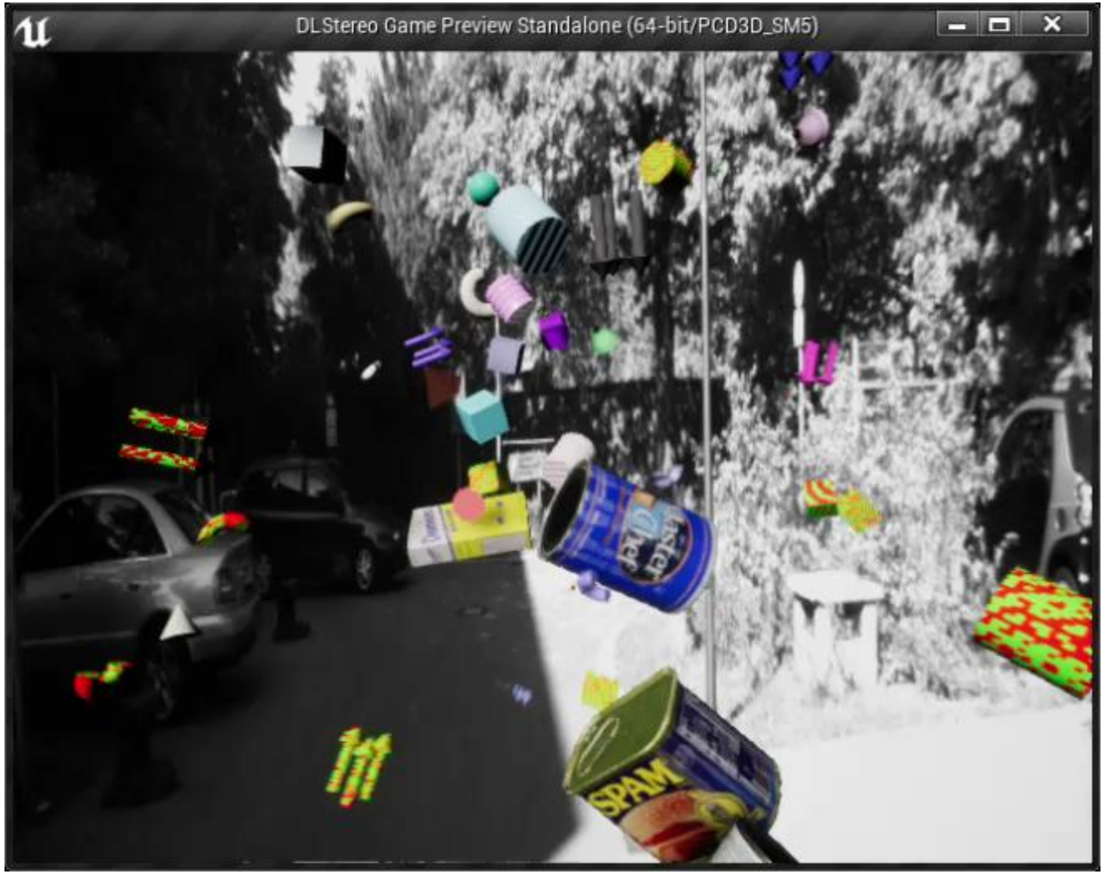

Falling Things: A Synthetic Dataset for 3D Object Detection and Pose Estimation. Tremblay *et al.* 2018

Together, these bridge the reality gap

ACCURACY MEASURED BY AREA UNDER THE CURVE

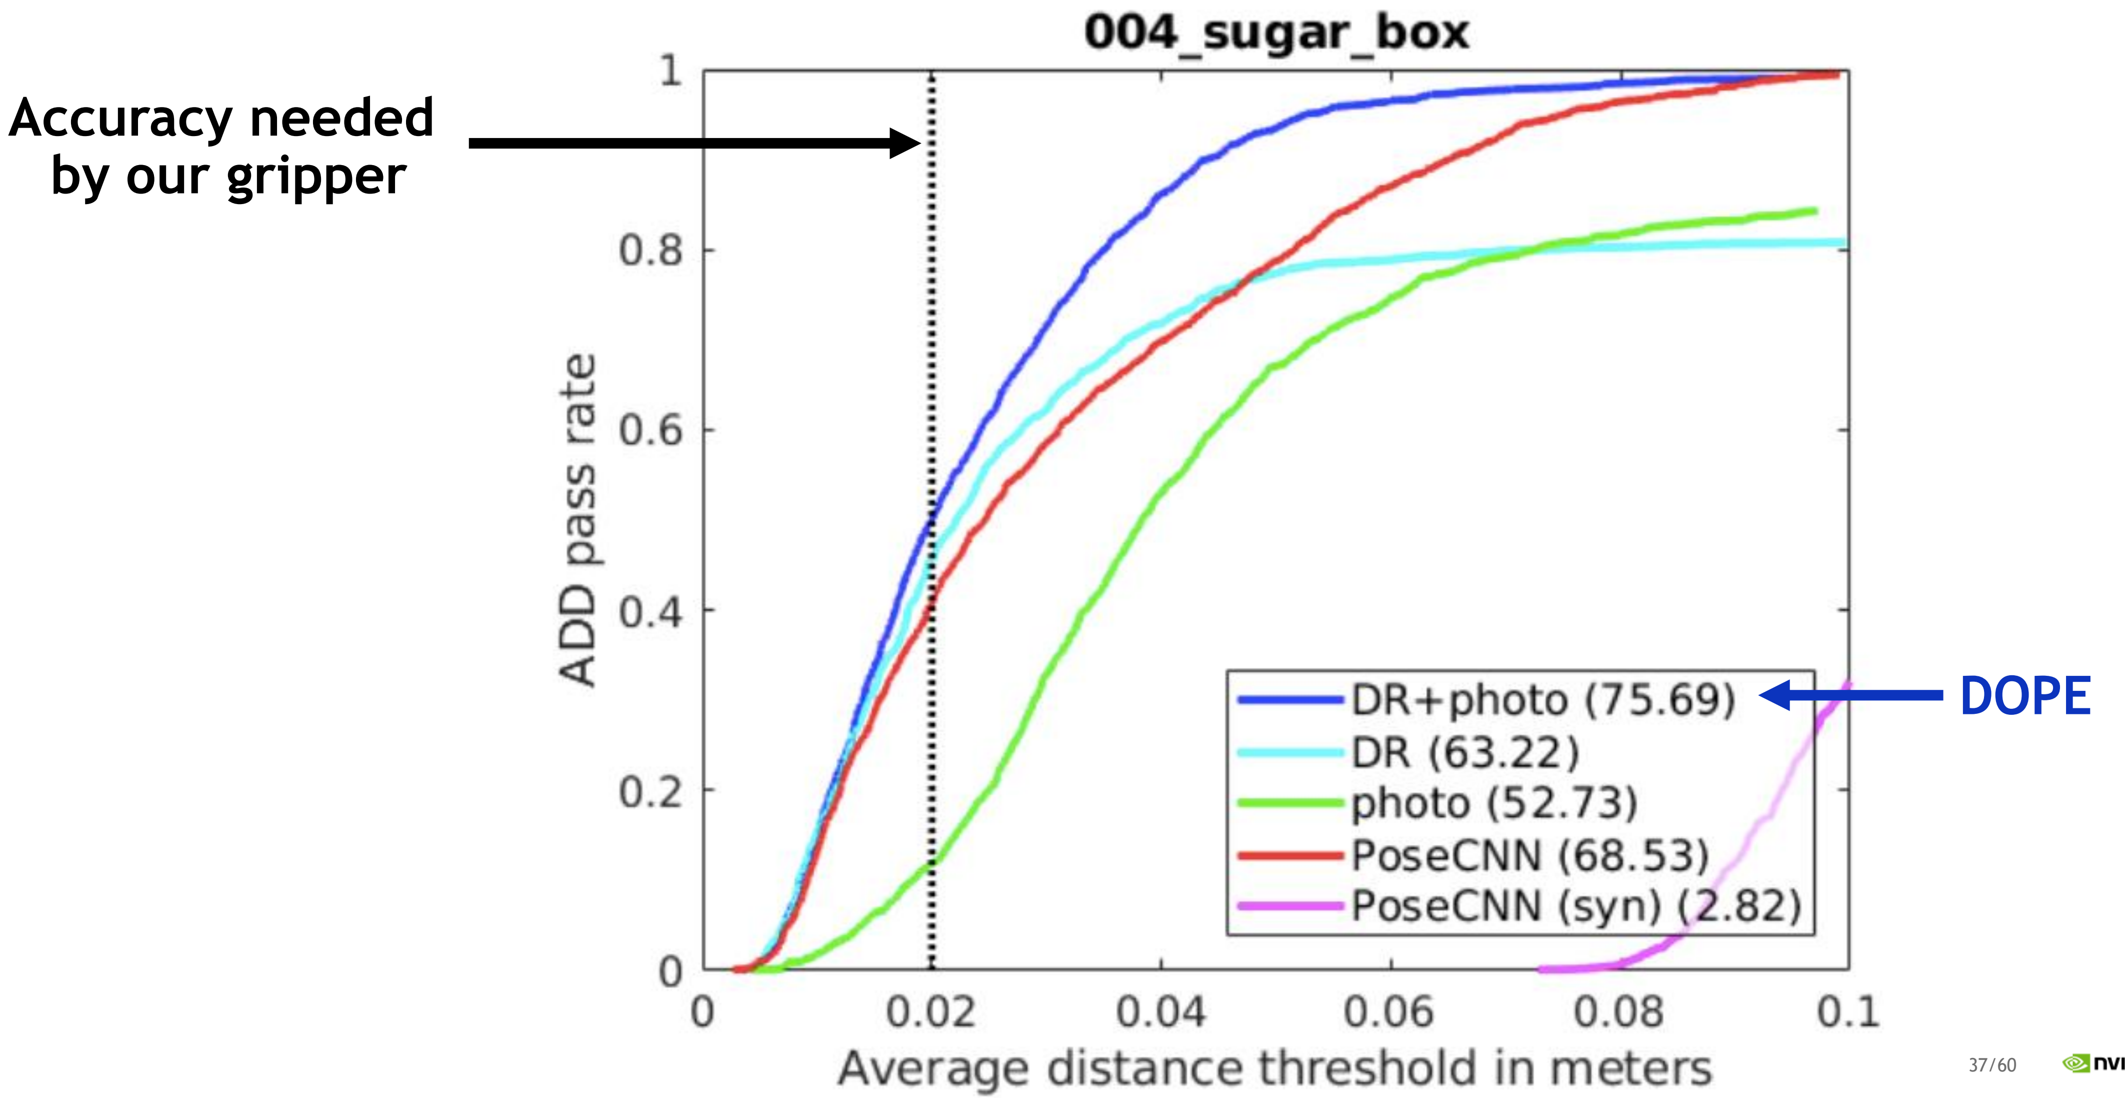

RESULTS ON YCB-VIDEO

	Cracker	Sugar	Soup	Mustard	Meat	Mean
DR	10.37	63.22	70.20	24.28	24.84	36.90
Photo	16.94	52.73	49.72	58.36	34.95	40.62
Photo+DR	**55.92**	**75.79**	**76.06**	**81.94**	39.38	**65.87**
PoseCNN (syn)	0	2.82	23.16	6.23	10.05	8.45
PoseCNN	51.51	68.53	66.07	79.70	**59.55**	65.07

Area under the curve for average distance threshold

DOPE trained only on synthetic data outperforms leading network trained on syn + real data

PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes
Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, Dieter Fox. RSS 2018

DOPE IN THE WILD

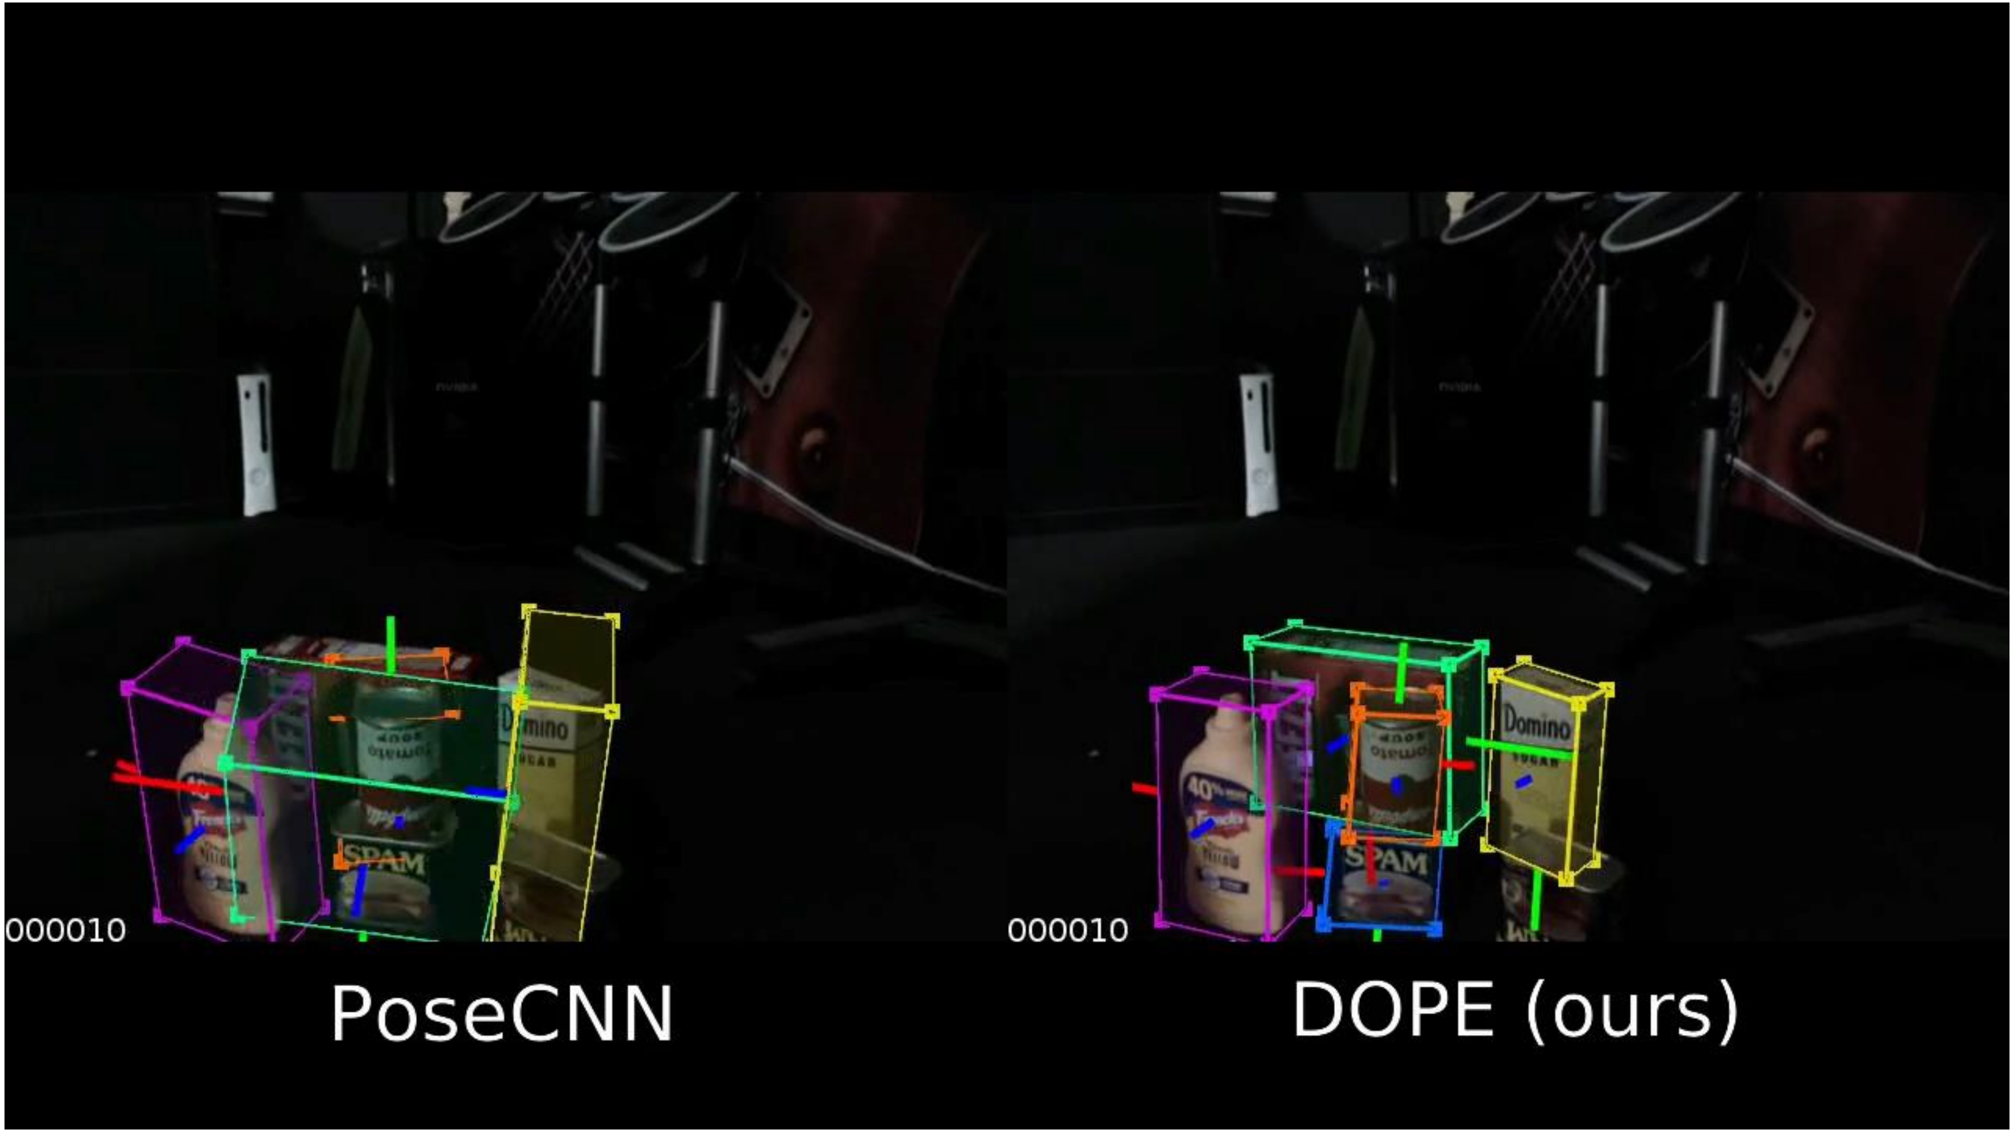

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

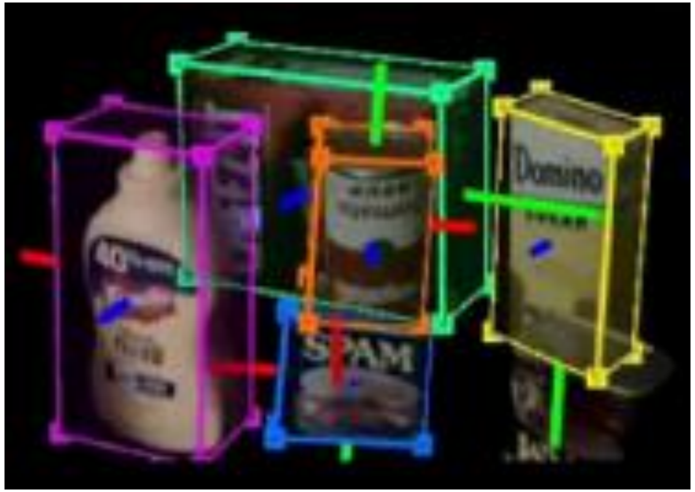

Vision

Closed-loop control

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

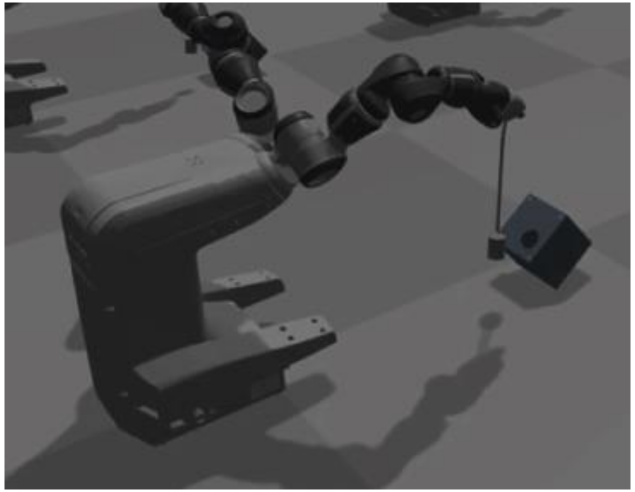

Vision

Closed-loop control

TRADITIONAL APPROACH

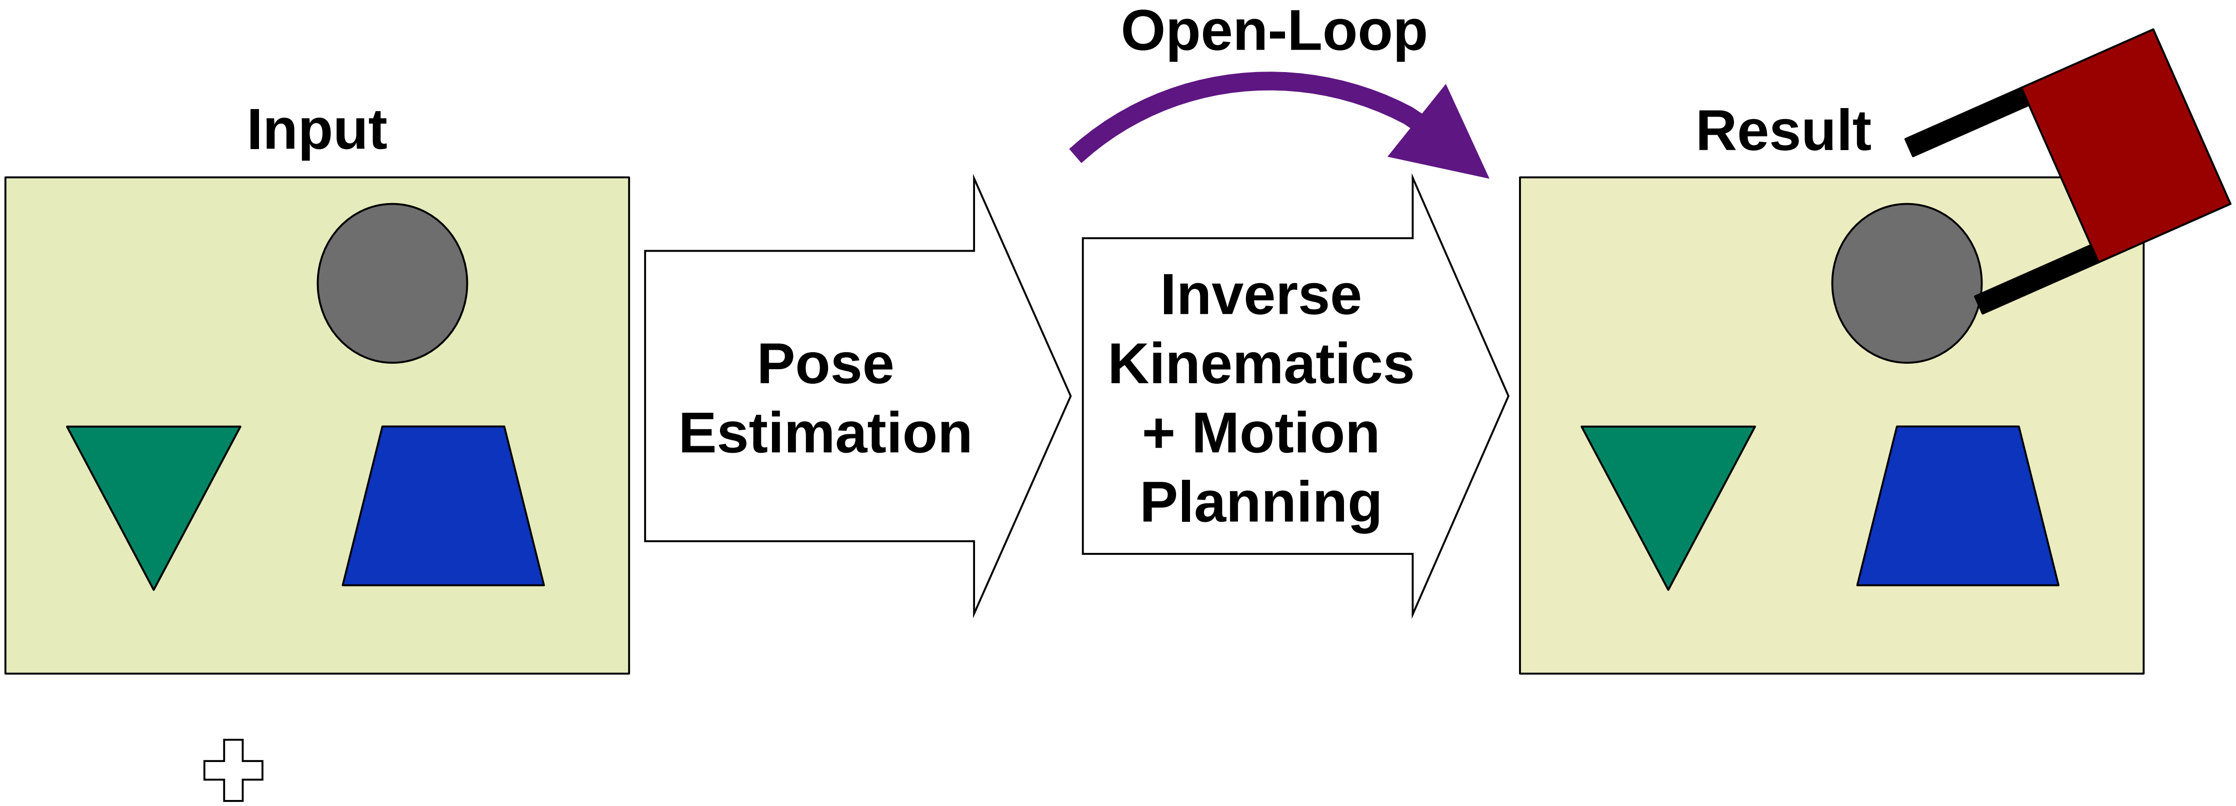

DOPE FOR ROBOTIC MANIPULATION

NVIDIA

DOPE ERRORS

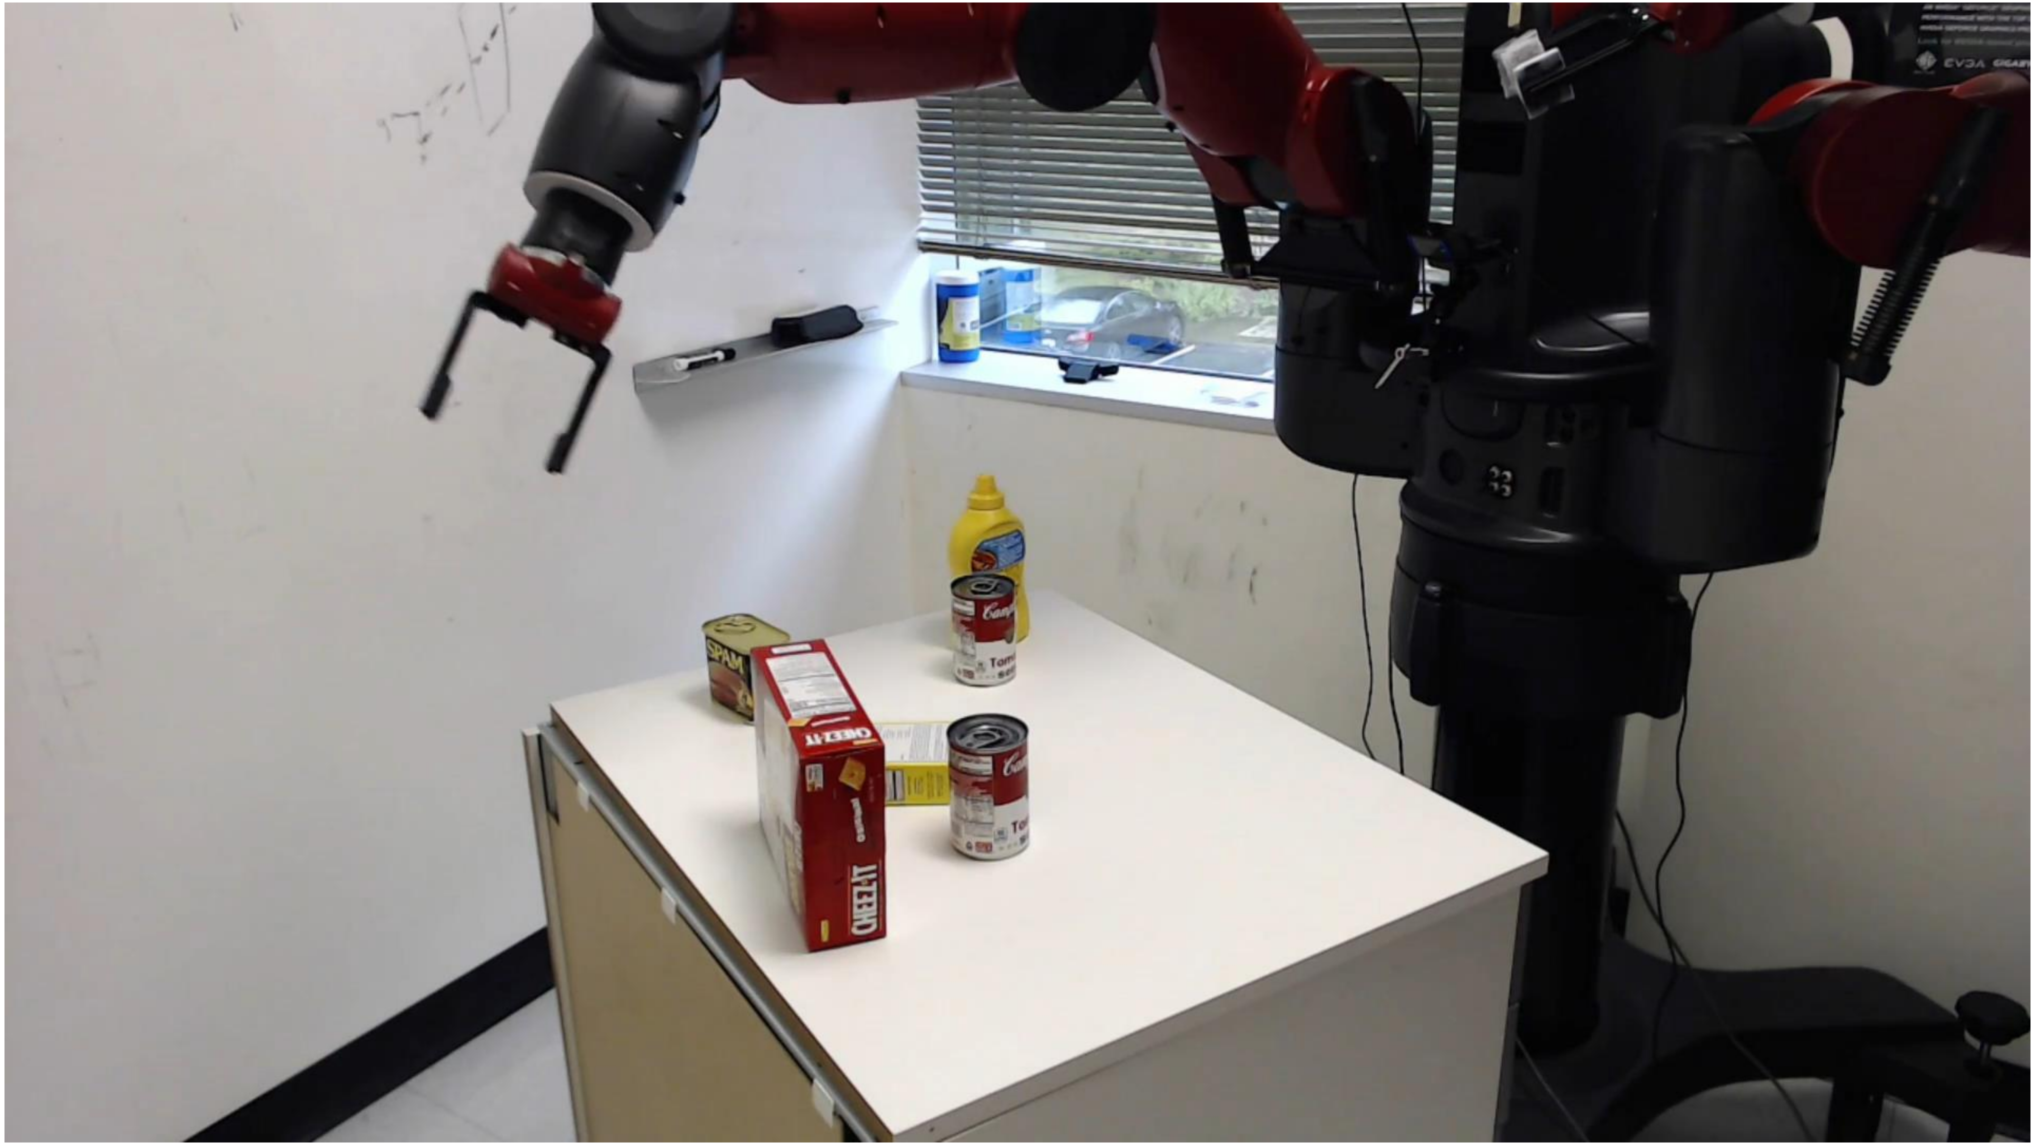

CLOSED-LOOP GRASPING

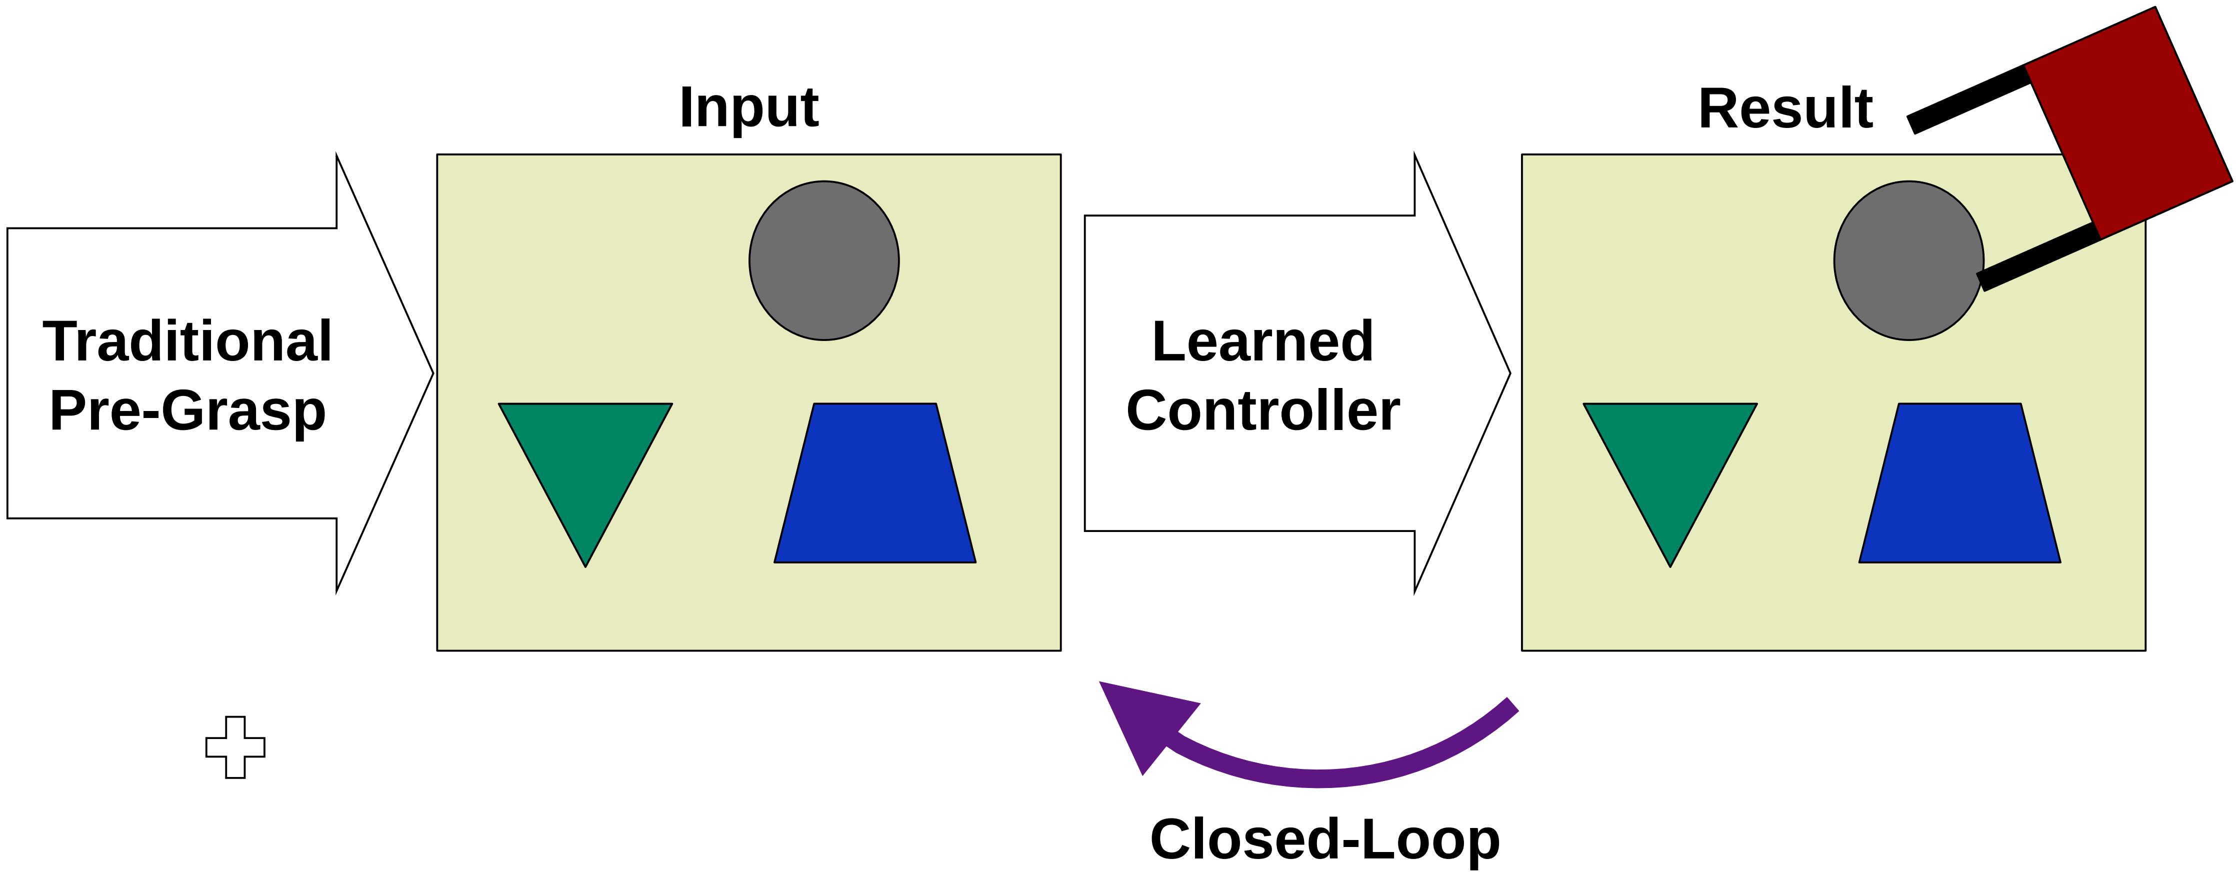

Feedback loop corrects errors in estimation / calibration

NVIDIA

ARCHITECTURE

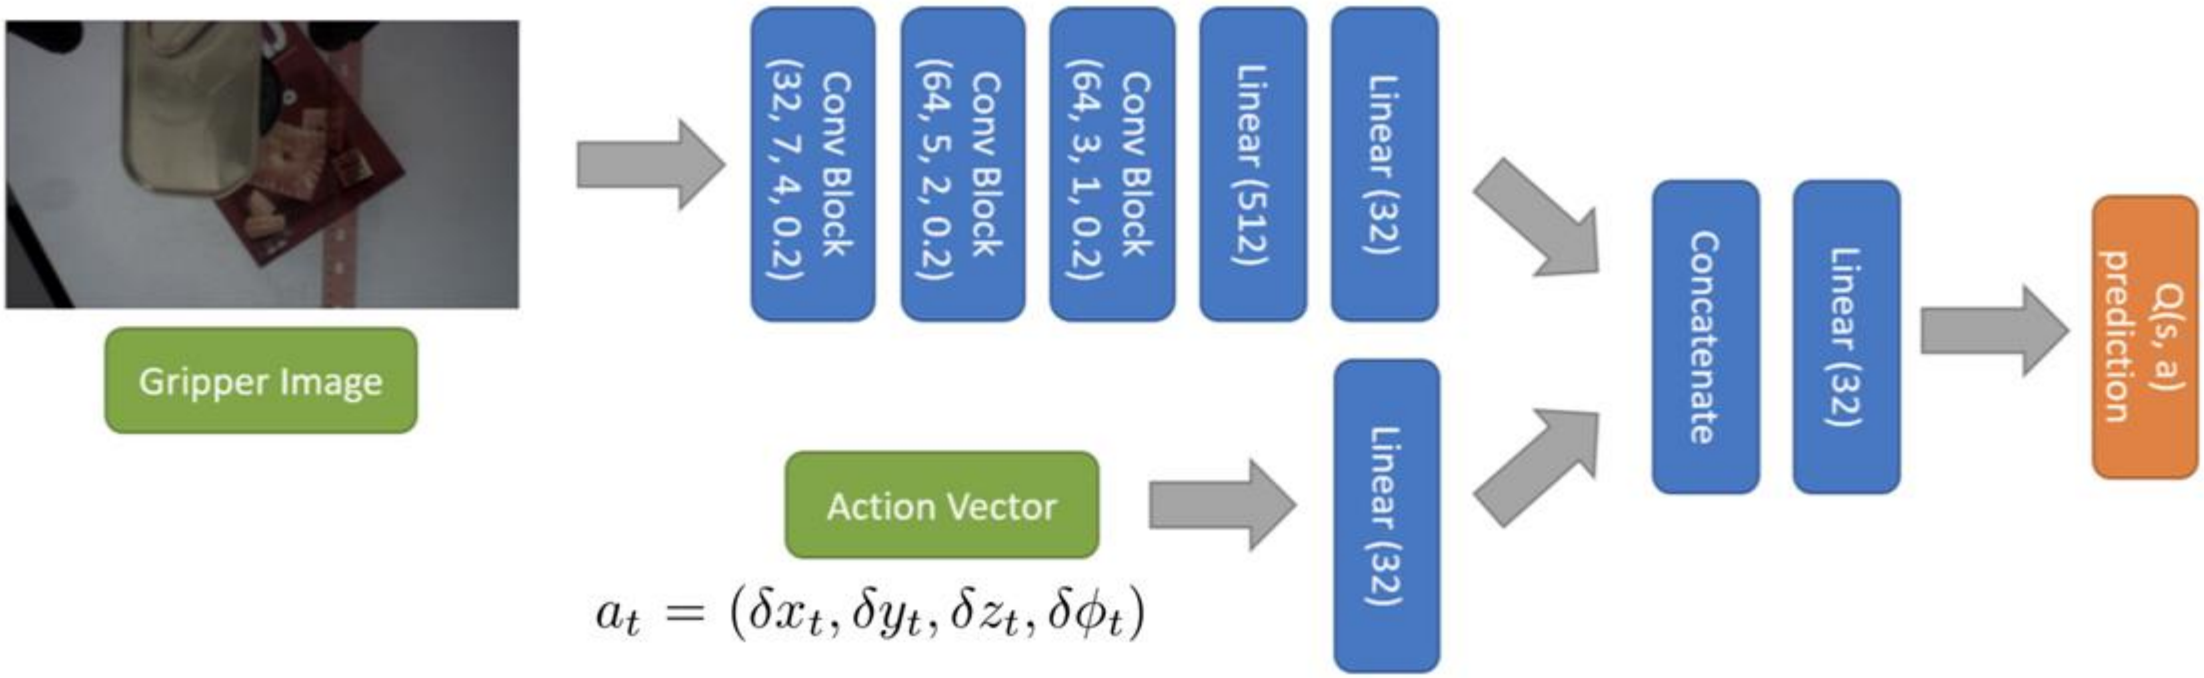

Trained via DDQN
(double deep Q-network)

Geometry-Aware Semantic Grasping of Real-World Objects: From Simulation to Reality.
S. Iqbal, J. Tremblay, T. To, J. Cheng, E. Leitch, D. McKay, S. Birchfield. *Submitted to IROS 2019*

SIMULATED ROBOT FARM

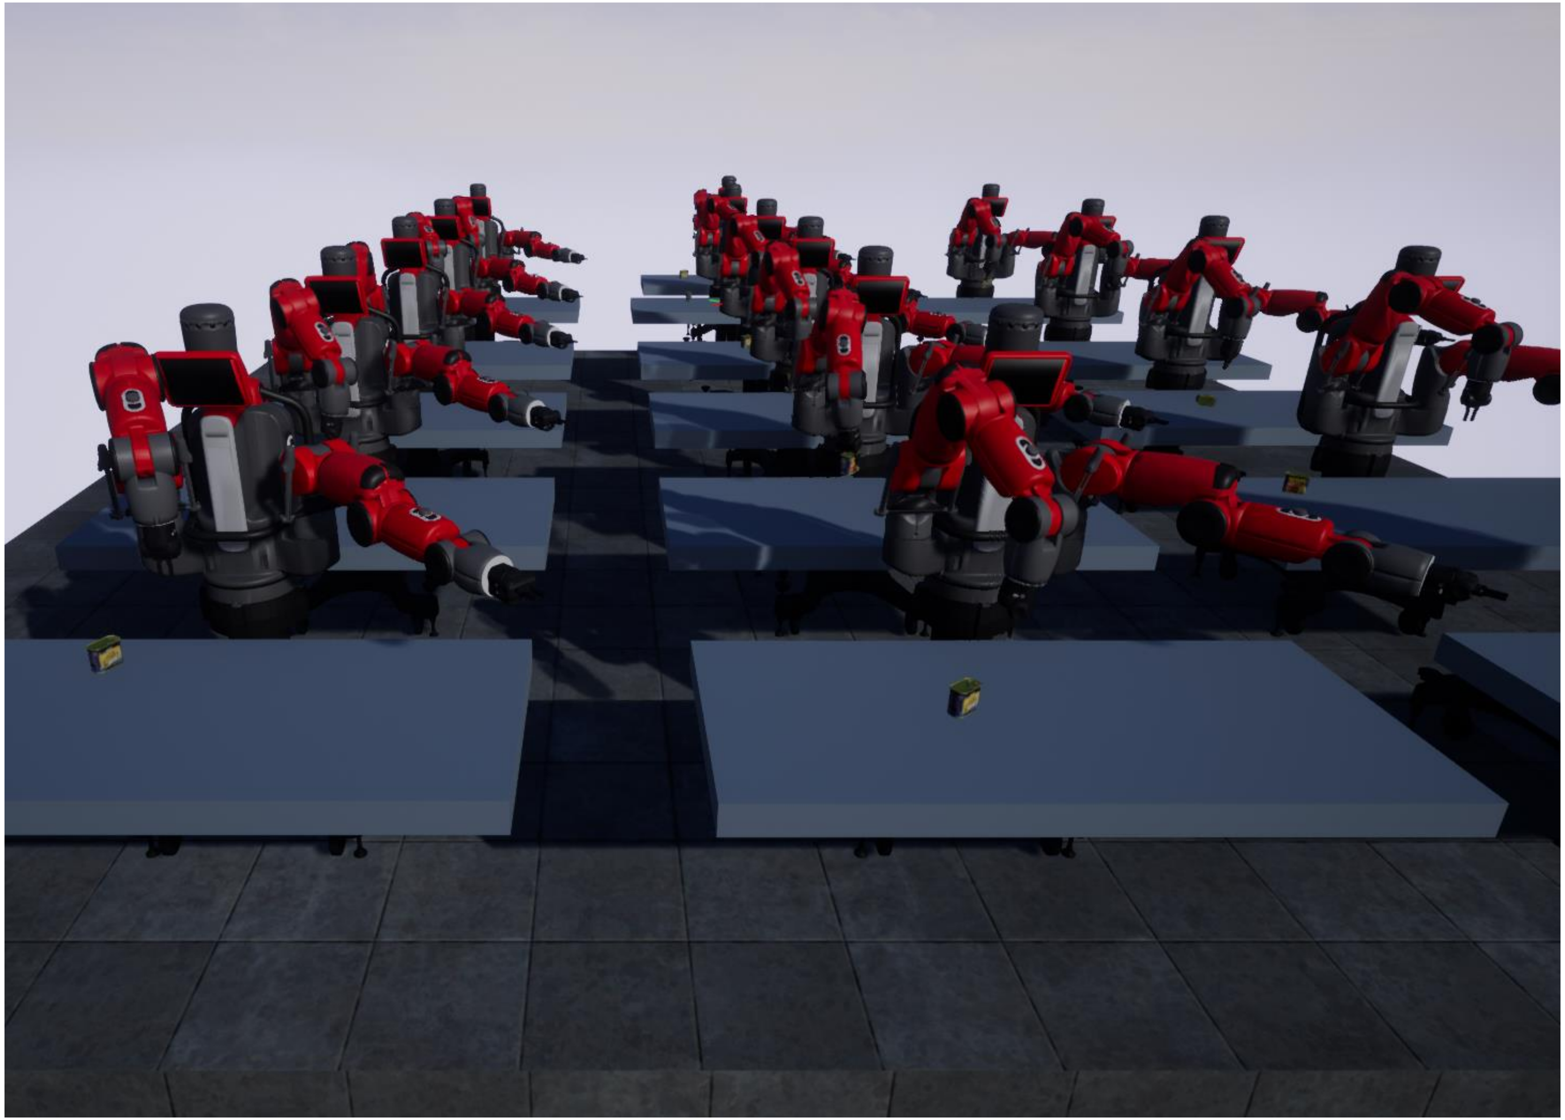

SIMULATED ROBOT FARM

RESULTS

LEARNING INVERSE DYNAMICS

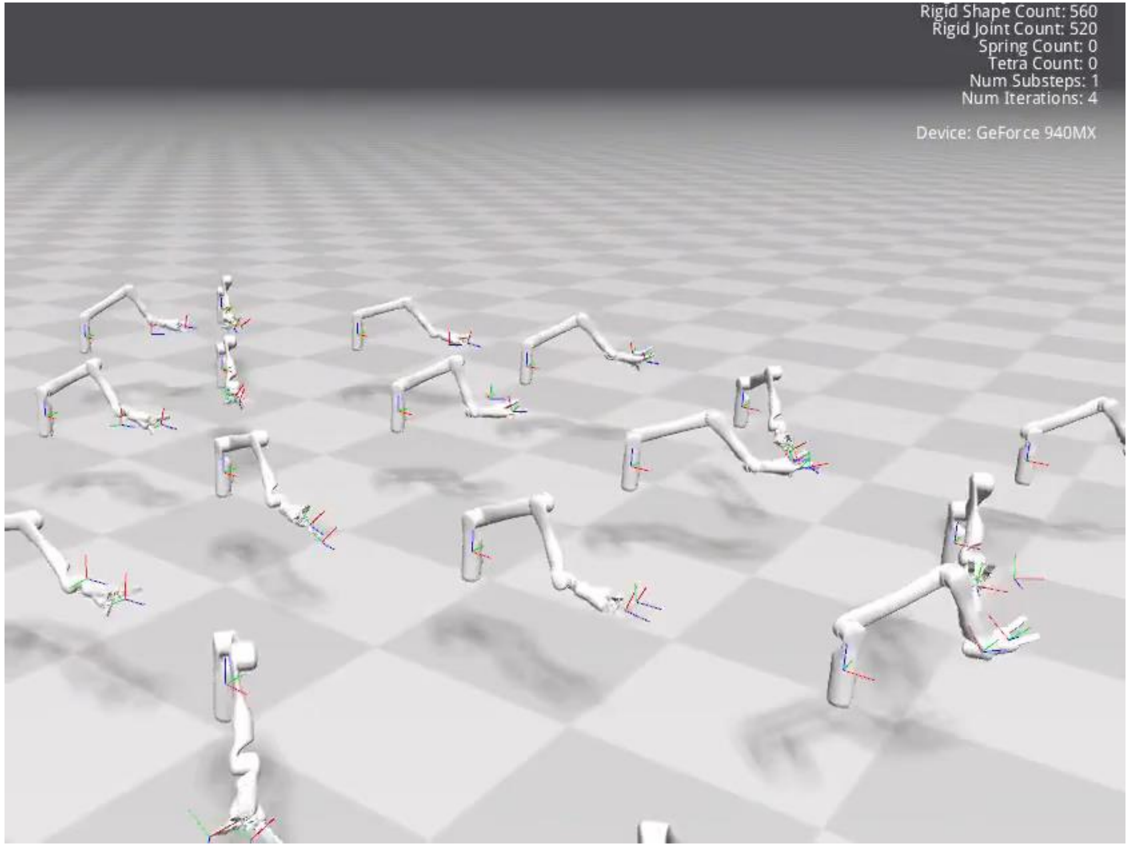

Simulation

Reality

Videos courtesy David Hoeller

REAL-TO-SIM

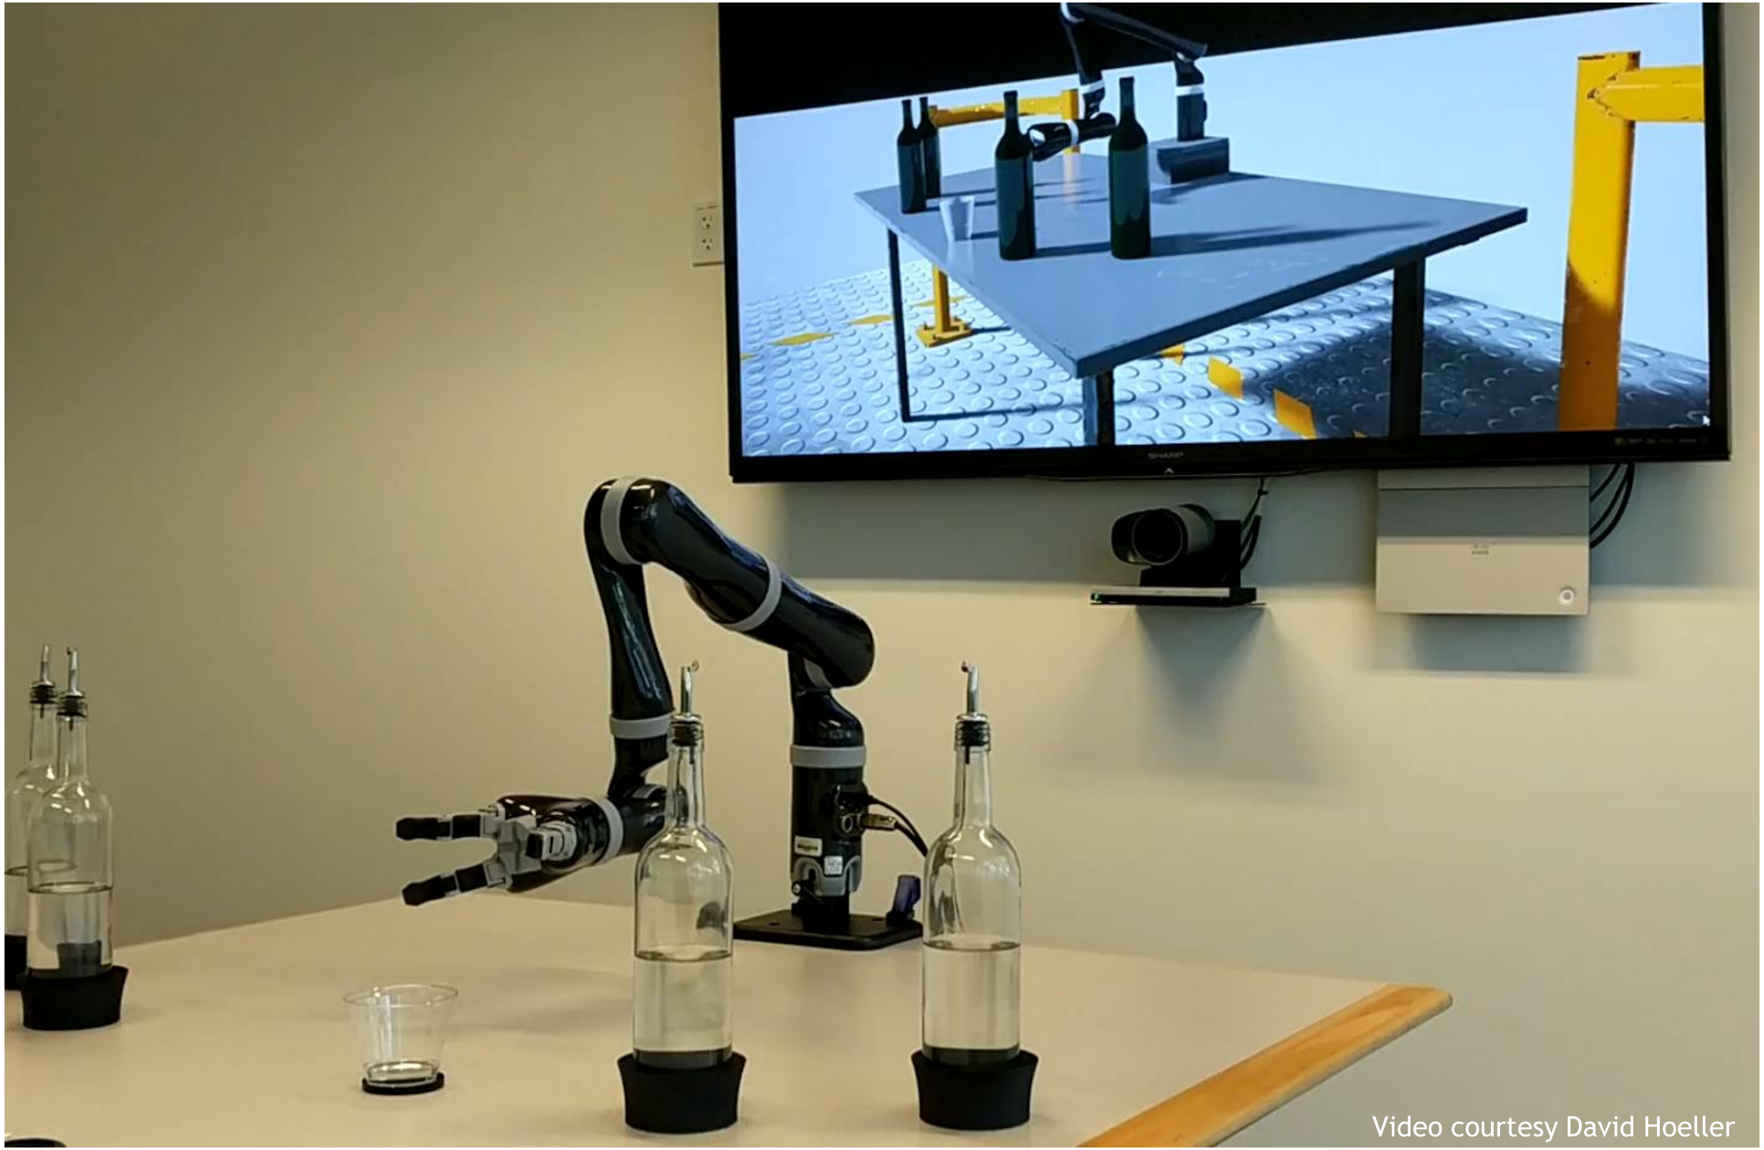

Video courtesy David Hoeller

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

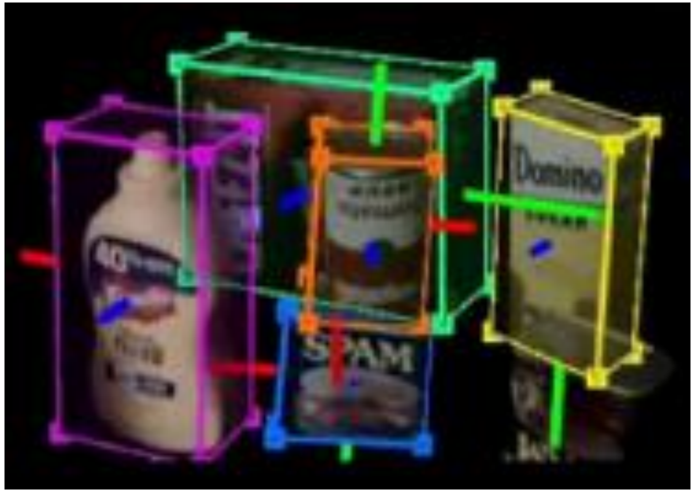

Vision

Closed-loop control

SIM-TO-REAL AT NVIDIA

Navigation

Manipulation

Vision

Closed-loop control

BAYES SIM

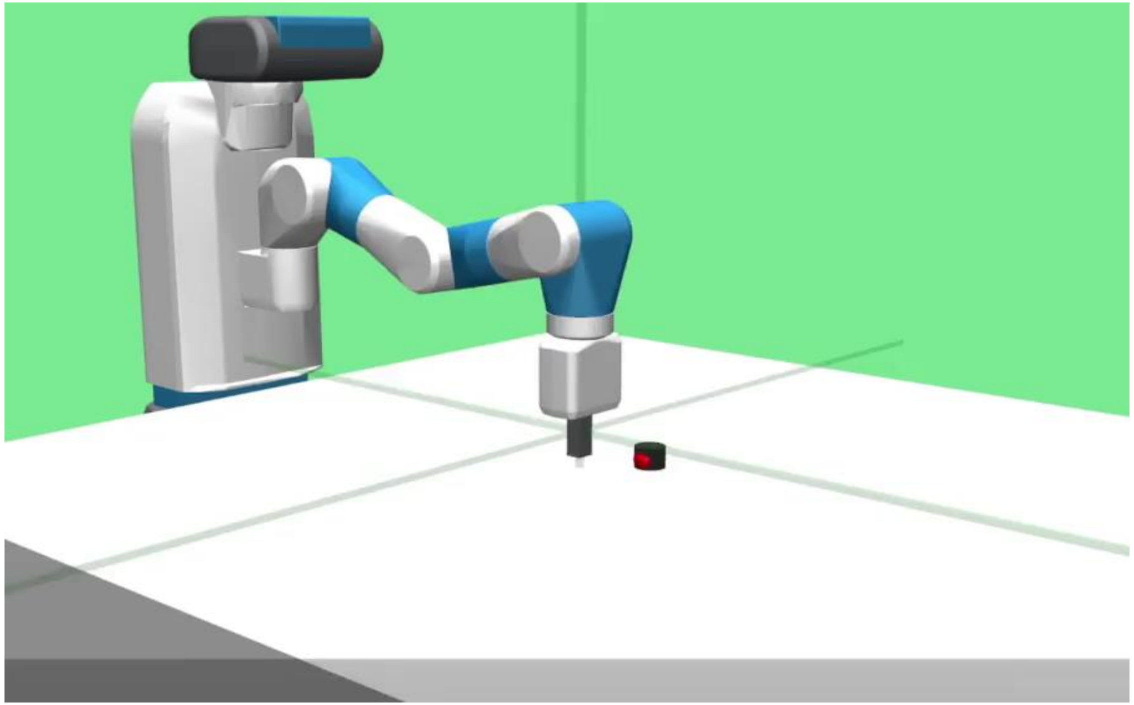

Training learns distribution of parameters

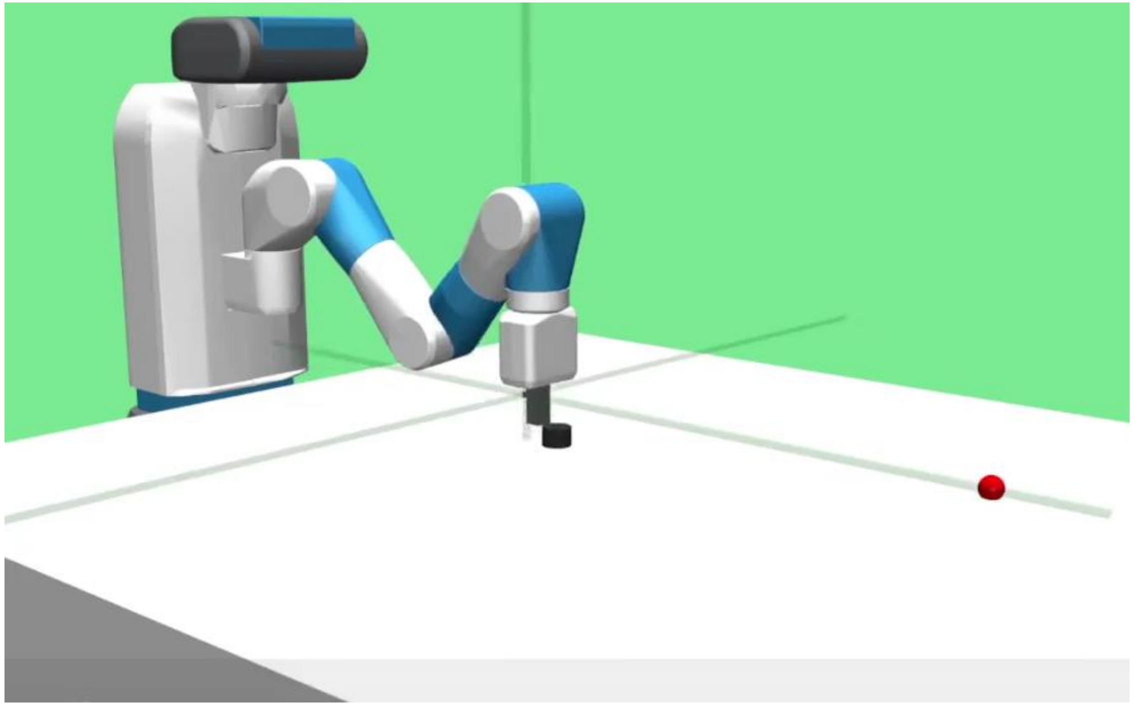

After training

BayesSim: Adaptive domain randomization via probabilistic inference for robotics simulators
F. Ramos, R. C. Possas, D. Fox. *Under review, 2019*

CLOSING THE SIM-TO-REAL LOOP

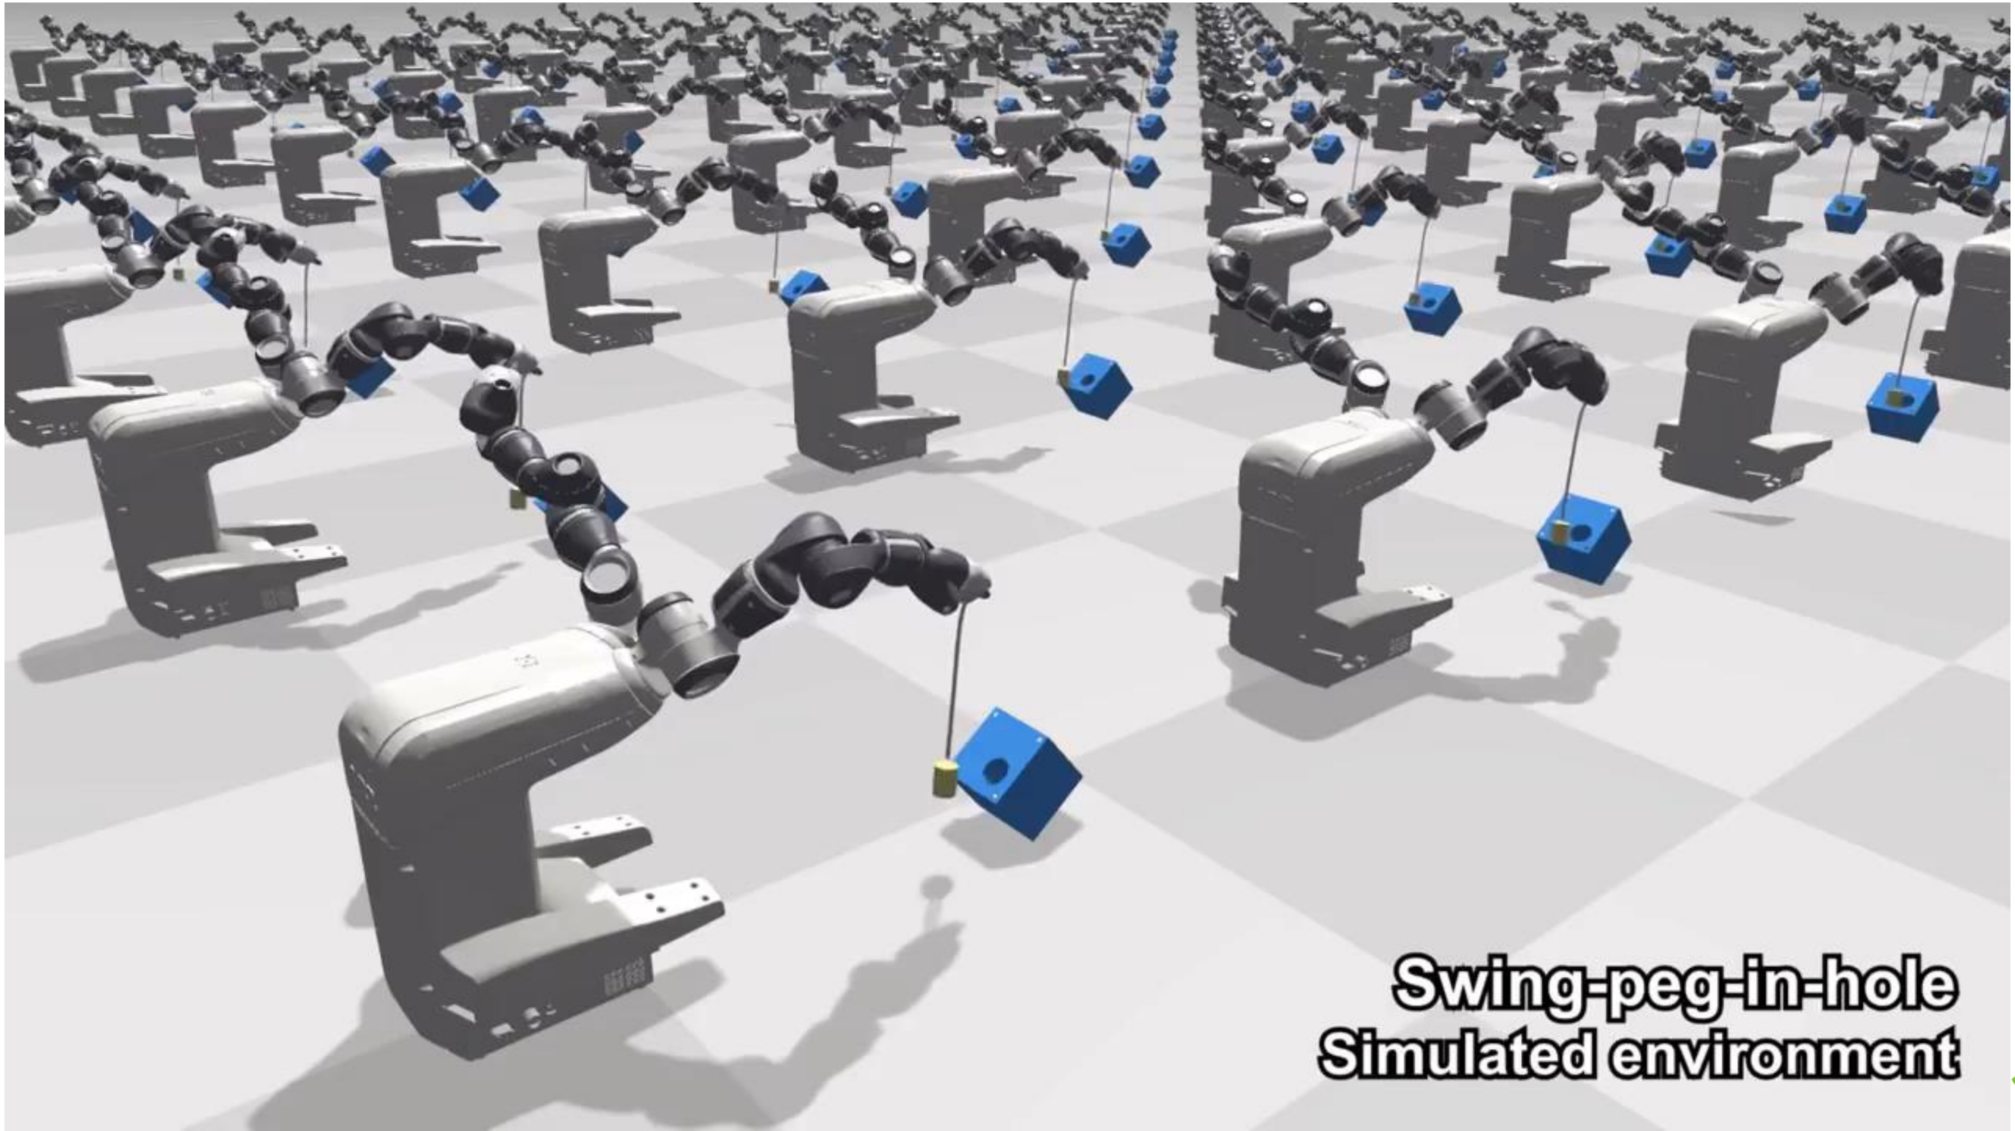

NVIDIA

CLOSING THE SIM-TO-REAL LOOP

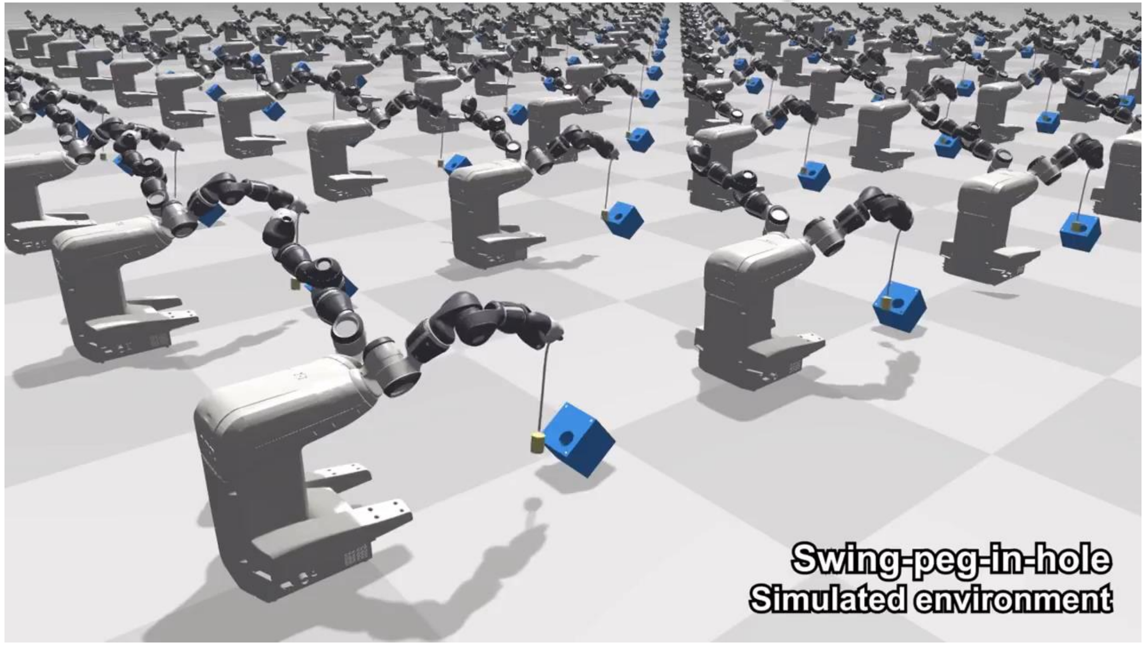

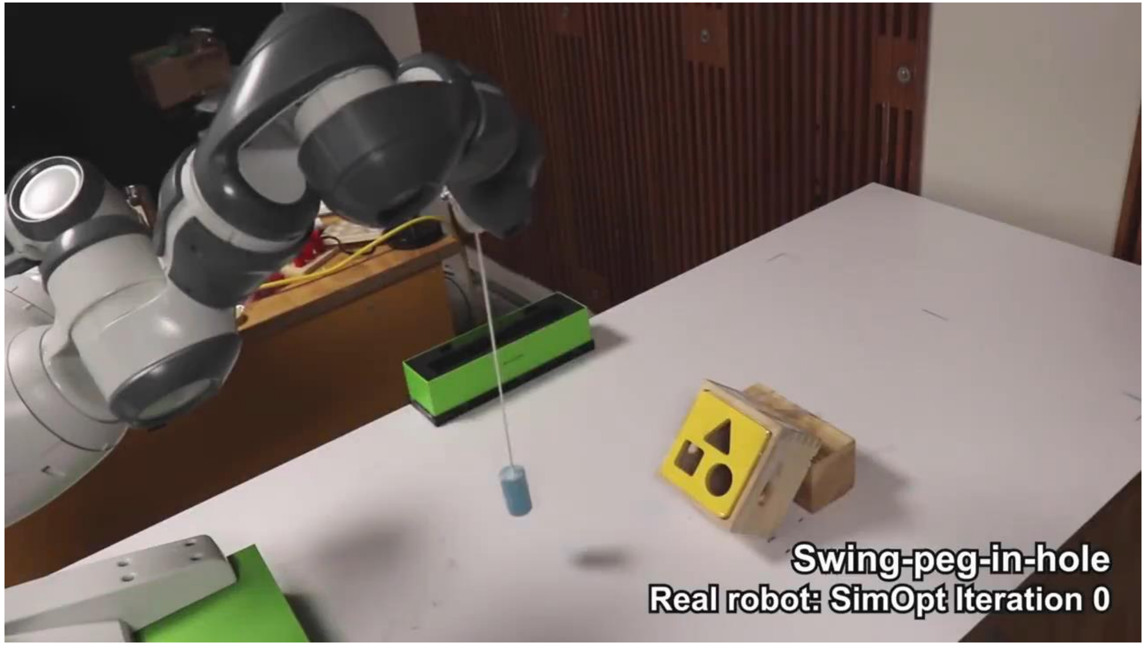

Closing the Sim-to-Real Loop: Adapting Simulation Randomization with Real World Experience
Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Issac, N. Ratliff, D. Fox. *ICRA 2019*

CLOSING THE SIM-TO-REAL LOOP

Simulation

Reality

SIM-TO-REAL LANDSCAPE

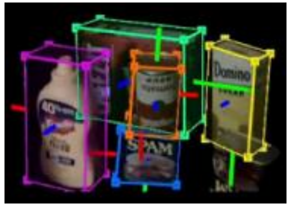

photorealism

physical realism

⋮ ⋰ ?

...

large-scale grasping
mobile manipulation
machine tending
in-hand manipulation
tactile sensing
object state changes
non-rigid objects
liquids
fast movement
generalization
...

CONCLUSION

Simulation will be key for robotics in

- Generating large amounts of labeled training data
- Quantitatively verifying policies / algorithms

Photorealism and ***physical realism*** are almost here

Many ***open problems***:

Tactile sensors?

Authoring content?

Super-real-time training?

Model verification?

Soft contact modeling?

Adaptation?

Scaling?

ACKNOWLEDGMENTS

Artem Molchanov
Shariq Iqbal
Thang To
Jia Cheng
Duncan McKay
Kirby Leung
Stephen Tyree
Jan Kautz
Dieter Fox

Ankur Handa
David Hoeller
Aayush Prakash
David Auld
Zvi Greenstein
Adam Moravanszky
Kier Storey
Nikolai Smolyanskiy
Alexei Kamenev

Vijay Baiyya
Jeffrey Smith
Johnny Costello
and many others

https://github.com/NVIDIA/Dataset_Synthesizer

https://github.com/NVlabs/Deep_Object_Pose

NVIDIA®