sim2real domain adaptation for effecient robotic grasping

mat_kelcey@

"Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection" Levine et al

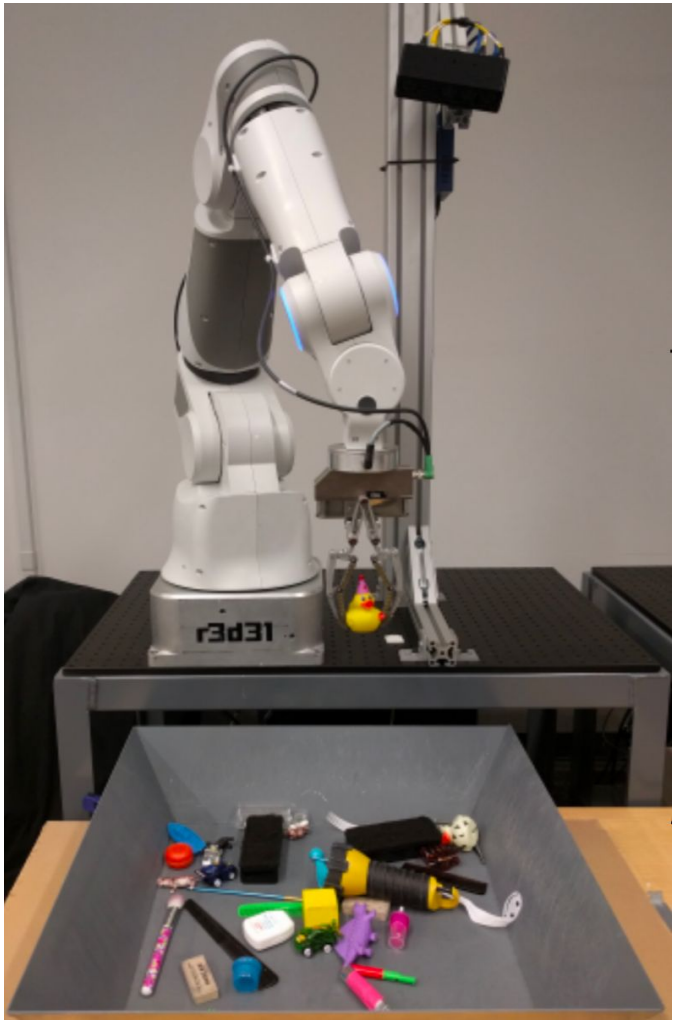
r3d31

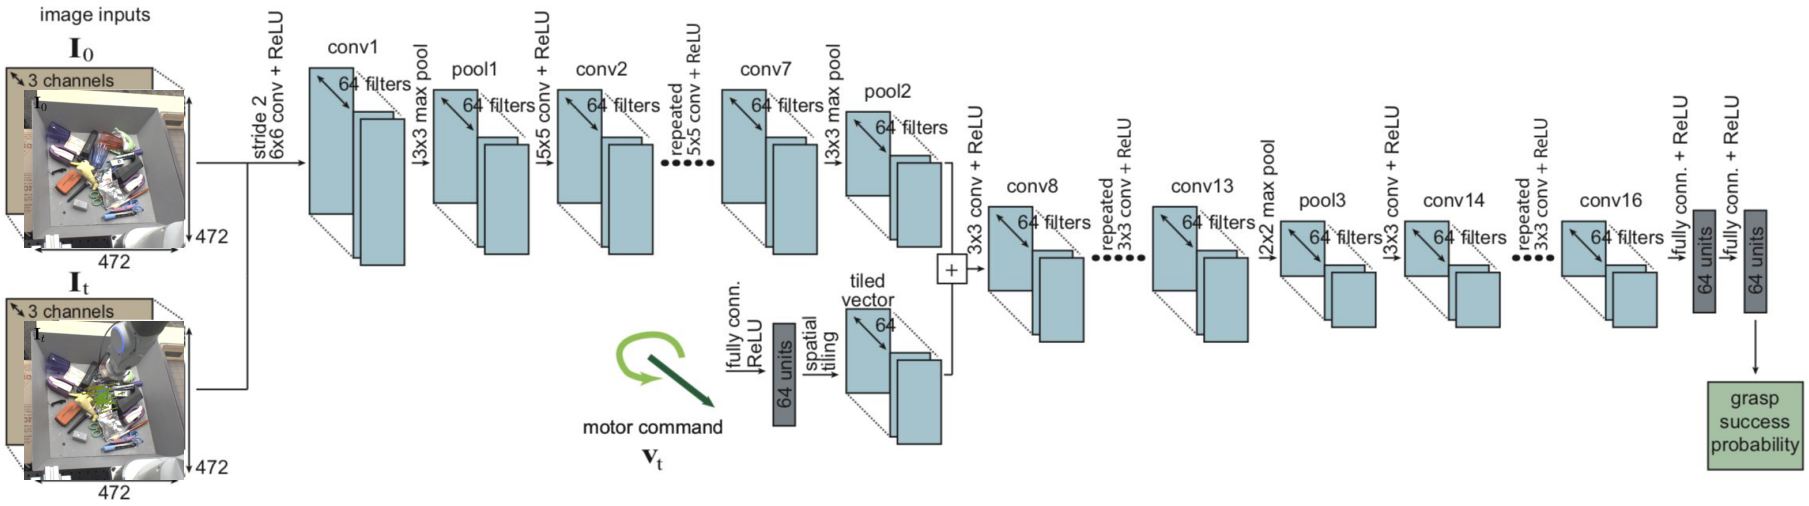

P(success | images, command)

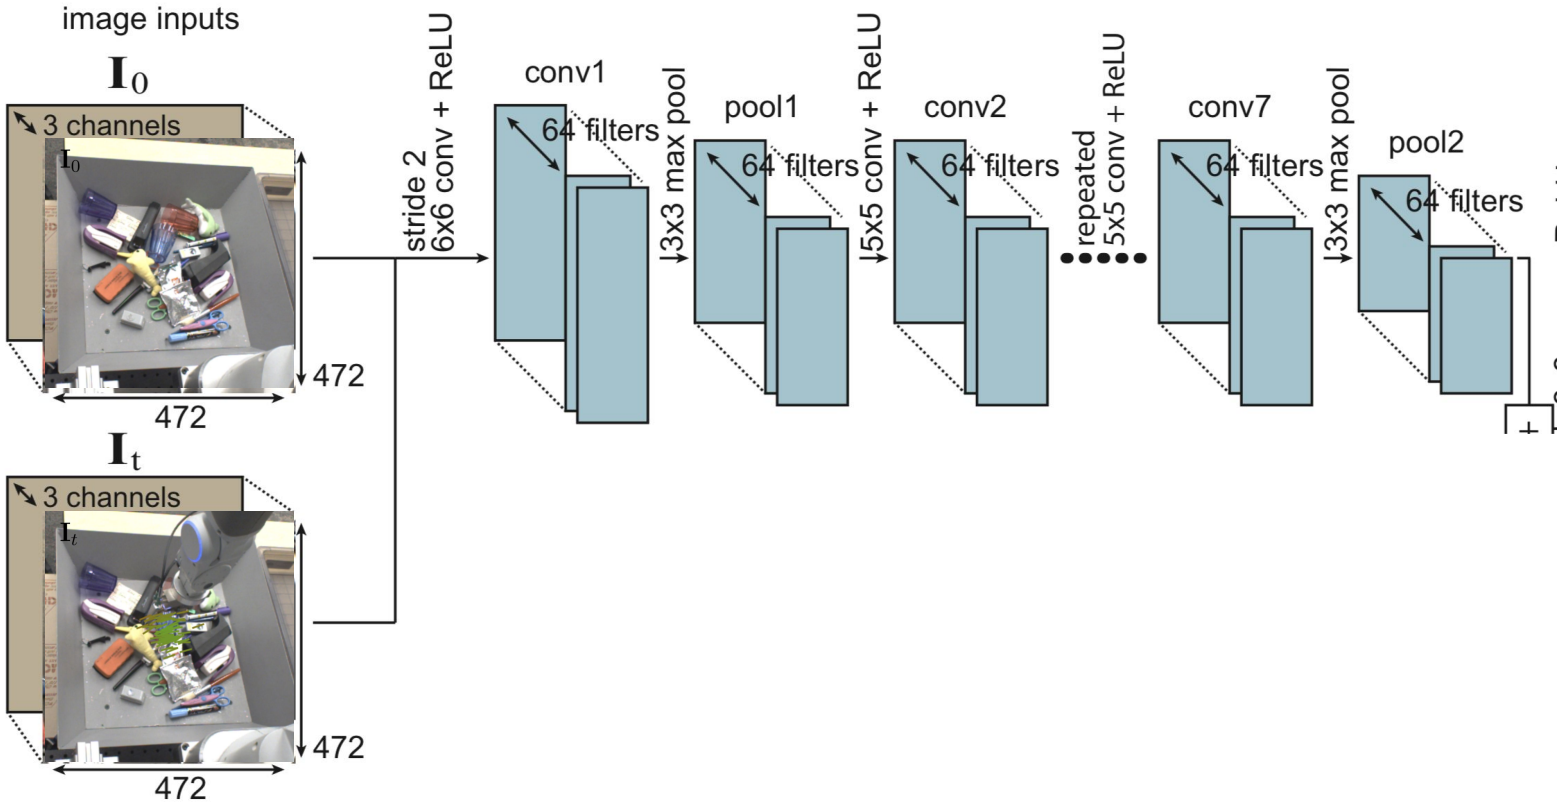
image inputs
$\mathbf{I}_0$
3 channels
$\mathbf{I}_0$
472
472
$\mathbf{I}_t$
3 channels
$\mathbf{I}_t$
472
472
stride 2
6x6 conv + ReLU
conv1
64 filters
3x3 max pool
pool1
64 filters
5x5 conv + ReLU
conv2
64 filters
repeated
5x5 conv + ReLU
conv7
64 filters
3x3 max pool
pool2
64 filters

motor command
$\mathbf{v}_t$
fully conn.
ReLU
64 units
spatial
tiling
tiled
vector
64

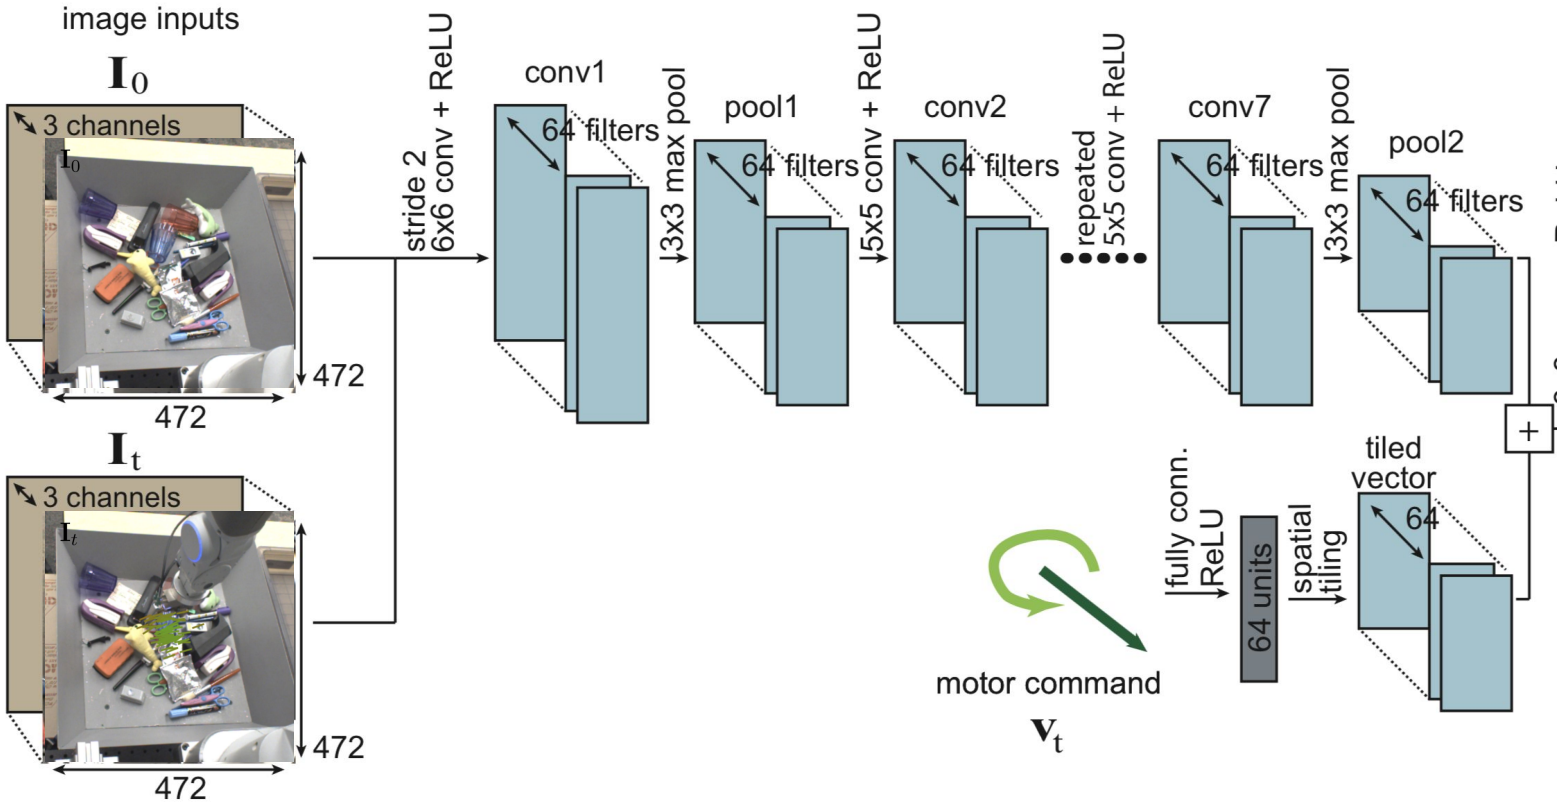
image inputs
$\mathbf{I}_0$
3 channels
$\mathbf{I}_0$
472
472
$\mathbf{I}_t$
3 channels
$\mathbf{I}_t$
472
472
stride 2
6x6 conv + ReLU
conv1
64 filters
3x3 max pool
pool1
64 filters
5x5 conv + ReLU
conv2
64 filters
repeated
5x5 conv + ReLU
conv7
64 filters
3x3 max pool
pool2
64 filters
+
tiled
vector
64
spatial
tiling
fully conn.
ReLU
64 units
motor command
$\mathbf{v}_t$

3x3 conv + ReLU
conv8
64 filters
repeated
3x3 conv + ReLU
conv13
64 filters
2x2 max pool
pool3
64 filters
3x3 conv + ReLU
conv14
64 filters
repeated
3x3 conv + ReLU
conv16
64 filters
fully conn. + ReLU
64 units
fully conn. + ReLU
64 units
grasp
success
probability

$P(\text{success} \mid \text{images}, \text{command})$

why not $\pi(\text{images}) \rightarrow \text{command}$?

P(success | images, command)

training

random grasps (images, commands, success)

P(success | images, command)

training

random grasps (images, commands, success)

testing

at inference time though we only have images

run quick optimisation to find command to maximise P(success | images, command)

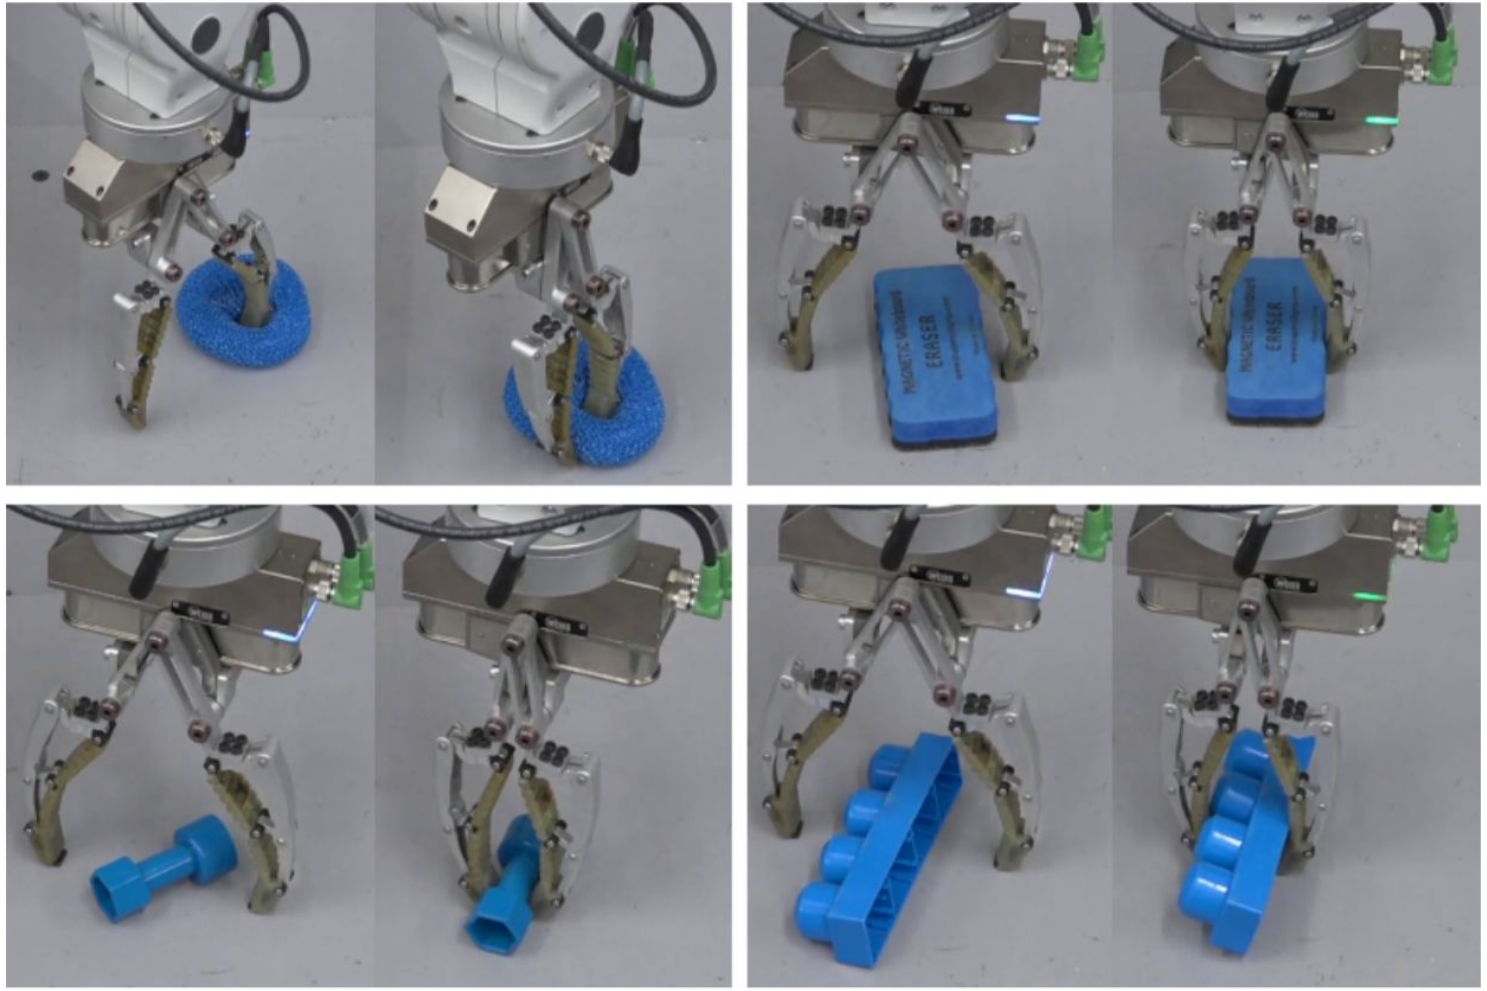
ERASER
ERASER

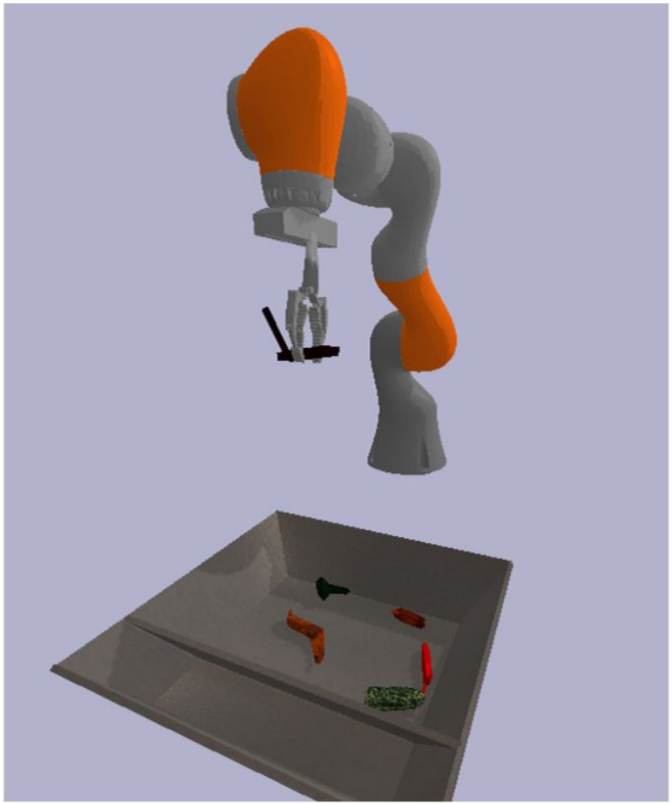

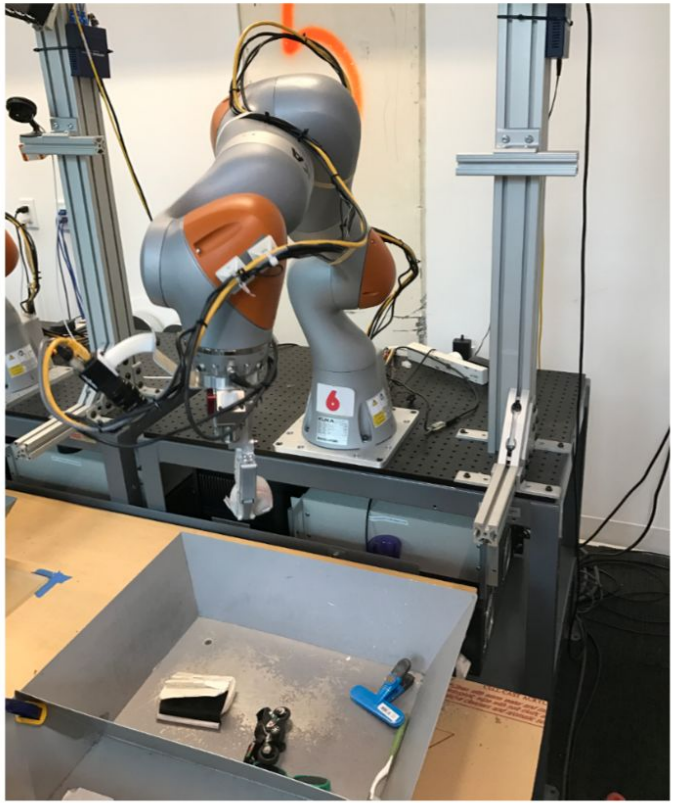

"Using Simulation and Domain Adaptation to Improve Efficiency of Deep Robotic Grasping" Bousmalis et al.

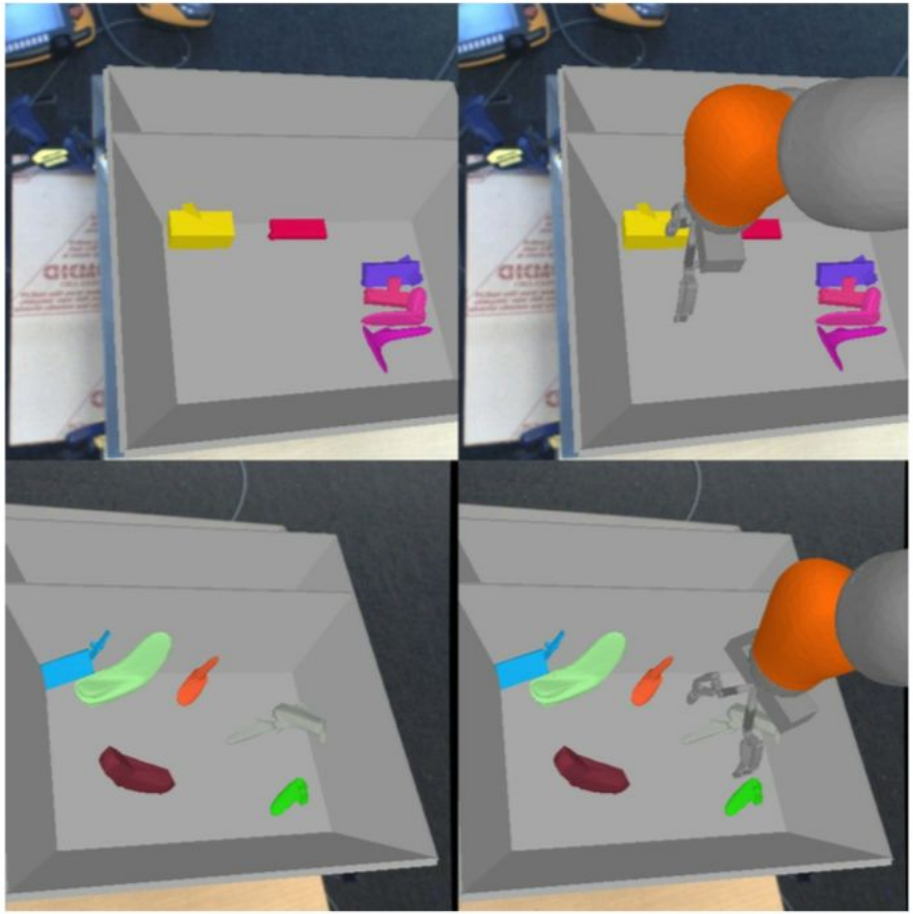

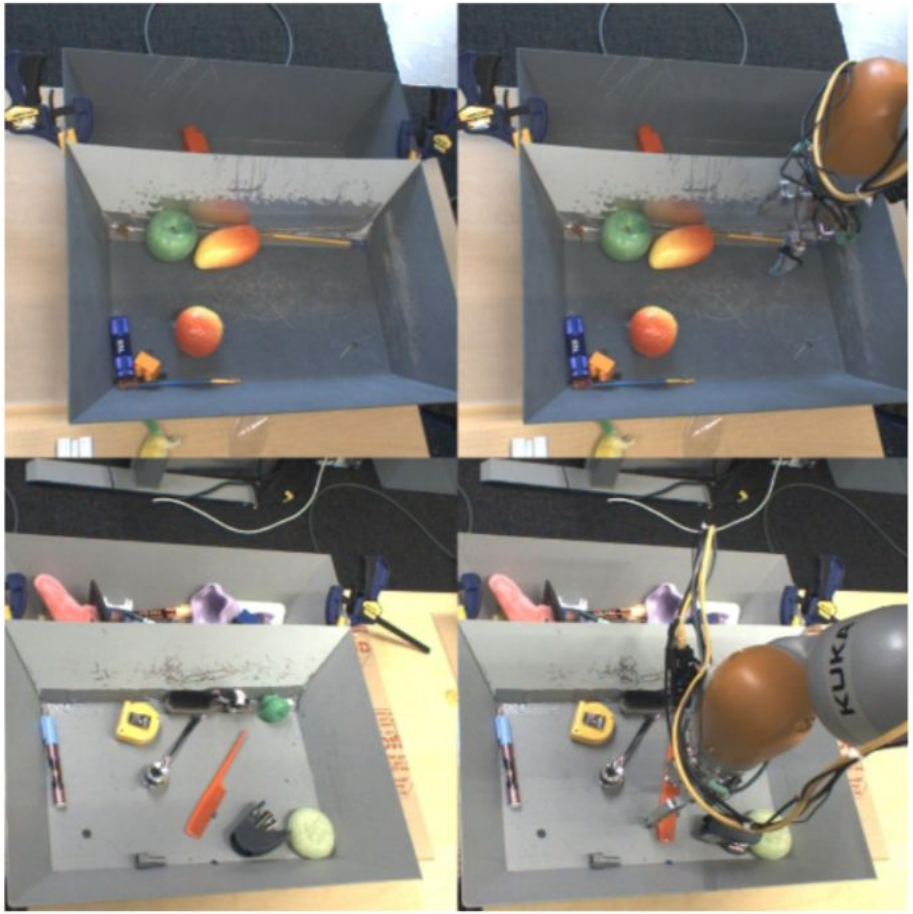

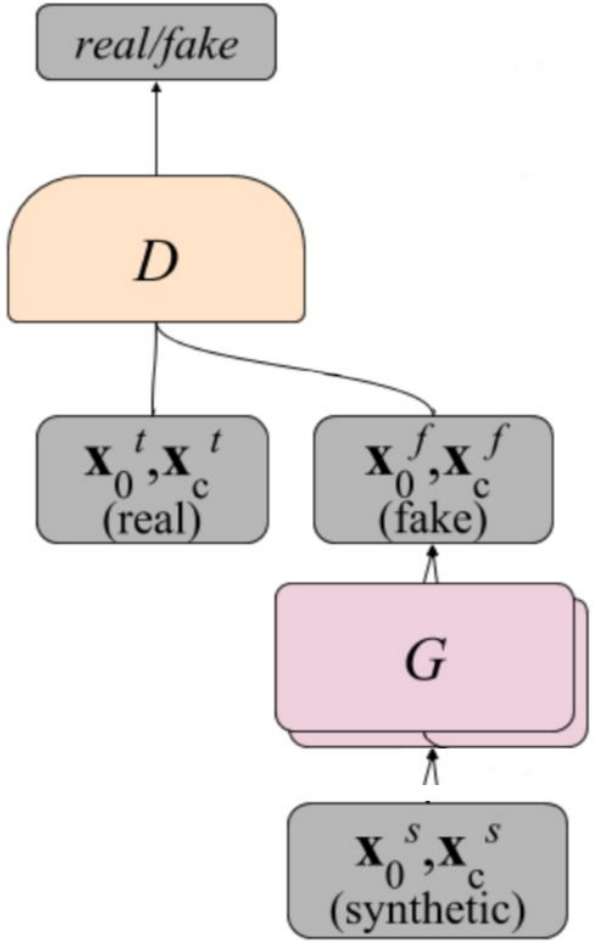
real/fake
D
$\mathbf{x}_0^t, \mathbf{x}_c^t$
(real)
$\mathbf{x}_0^f, \mathbf{x}_c^f$
(fake)
G
$\mathbf{x}_0^s, \mathbf{x}_c^s$
(synthetic)

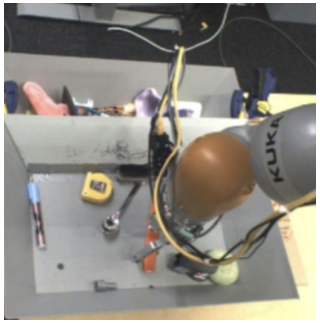

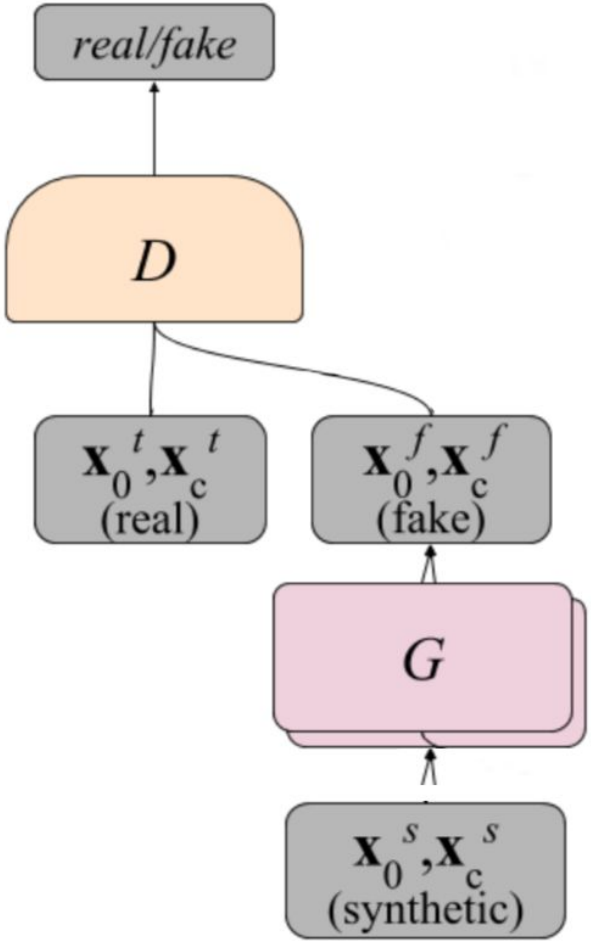
real/fake
D
$\mathbf{x}_0^t, \mathbf{x}_c^t$
(real)
$\mathbf{x}_0^f, \mathbf{x}_c^f$
(fake)
G
$\mathbf{x}_0^s, \mathbf{x}_c^s$
(synthetic)

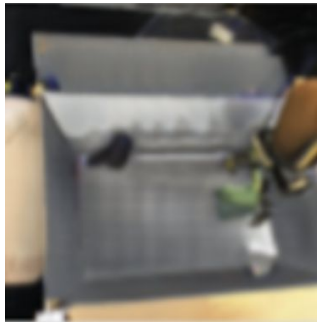

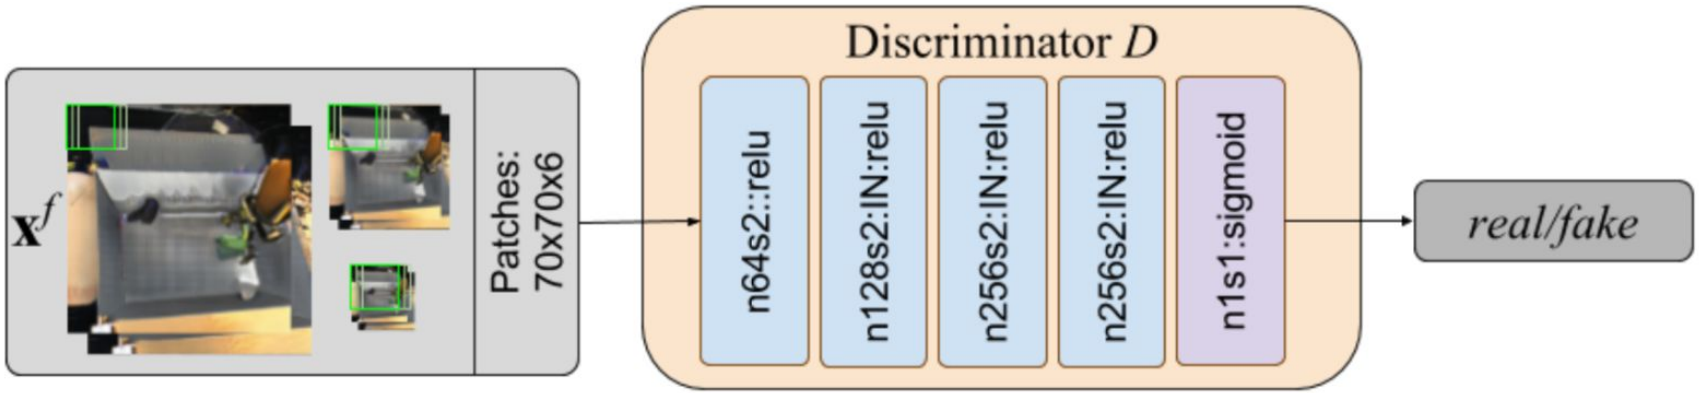
$\mathbf{x}^f$
Patches:
70x70x6
Discriminator D
n64s2::relu
n128s2:IN:relu
n256s2:IN:relu
n256s2:IN:relu
n1s1:sigmoid
real/fake

(b)

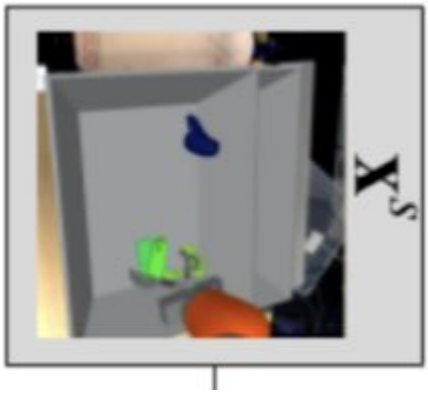
x^s

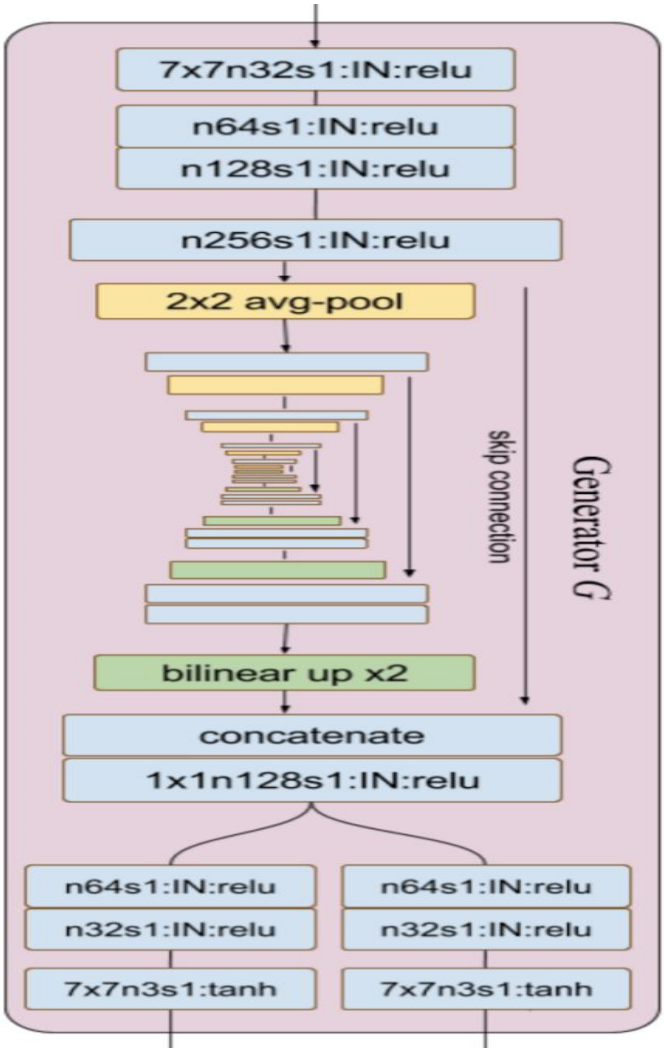
7x7n32s1:IN:relu
n64s1:IN:relu
n128s1:IN:relu
n256s1:IN:relu
2x2 avg-pool
skip connection
Generator G
bilinear up x2
concatenate
1x1n128s1:IN:relu
n64s1:IN:relu
n32s1:IN:relu
7x7n3s1:tanh
n64s1:IN:relu
n32s1:IN:relu
7x7n3s1:tanh

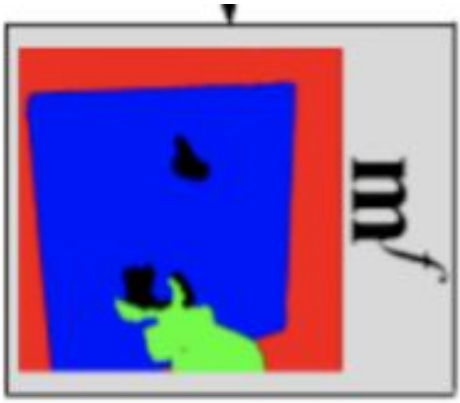
m^s


x^f

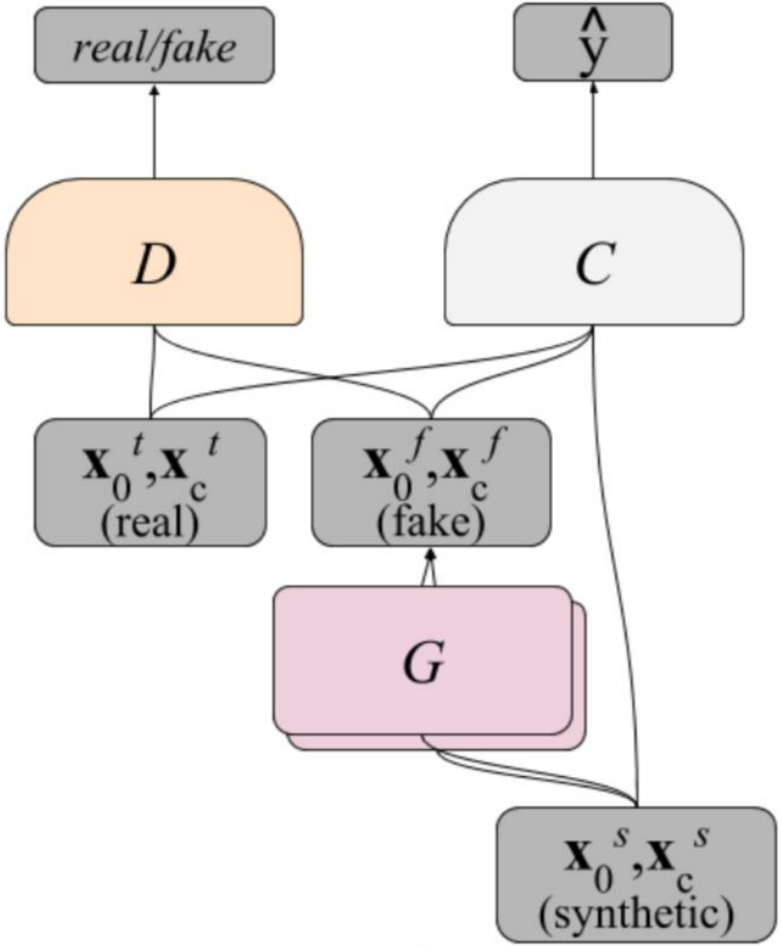
real/fake
ŷ
D
C
$\mathbf{x}_0^t, \mathbf{x}_c^t$ (real)
$\mathbf{x}_0^f, \mathbf{x}_c^f$ (fake)
G
$\mathbf{x}_0^s, \mathbf{x}_c^s$ (synthetic)

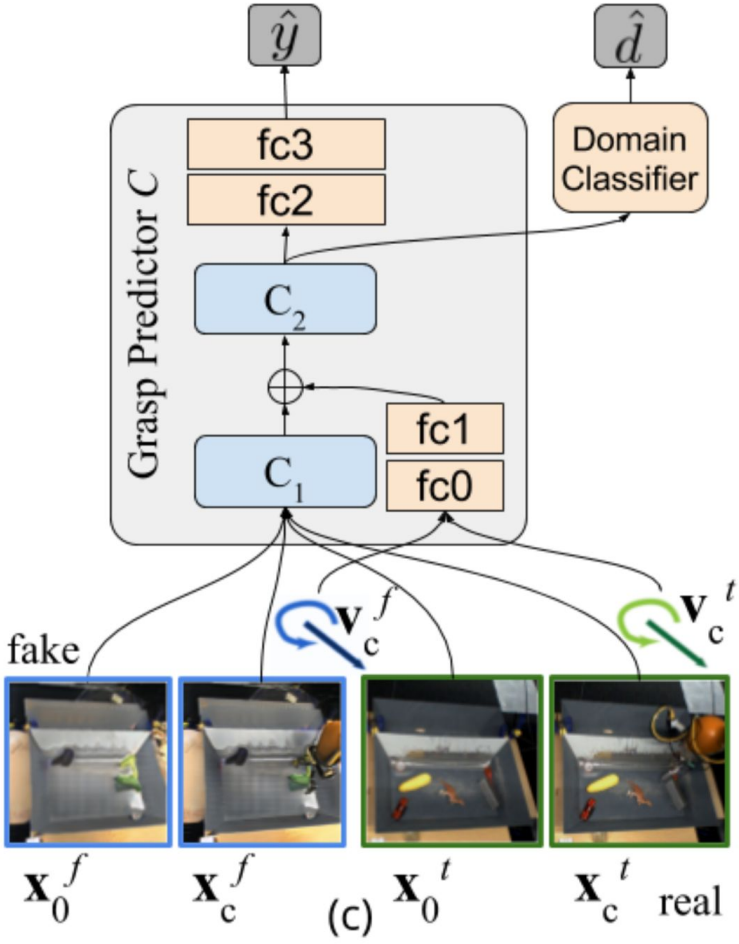
$\hat{y}$
$\hat{d}$
Grasp Predictor C
fc3
fc2
C_2
fc1
fc0
C_1
Domain Classifier
fake
$\mathbf{v}_c^f$
$\mathbf{v}_c^t$
$\mathbf{x}_0^f$
$\mathbf{x}_c^f$
$\mathbf{x}_0^t$
$\mathbf{x}_c^t$
real

(c)

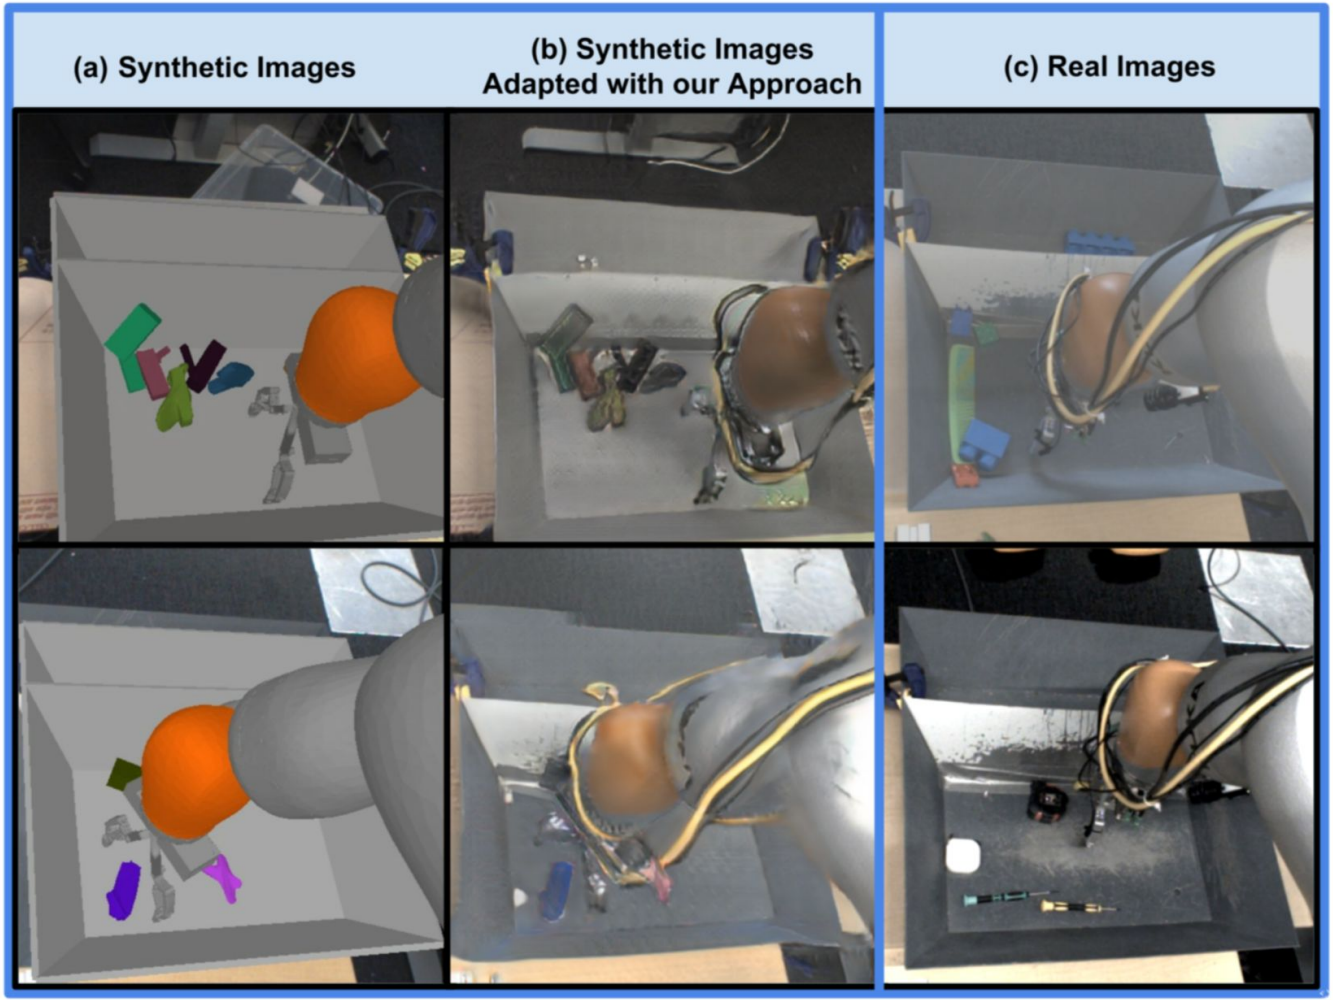
(a) Synthetic Images
(b) Synthetic Images Adapted with our Approach
(c) Real Images

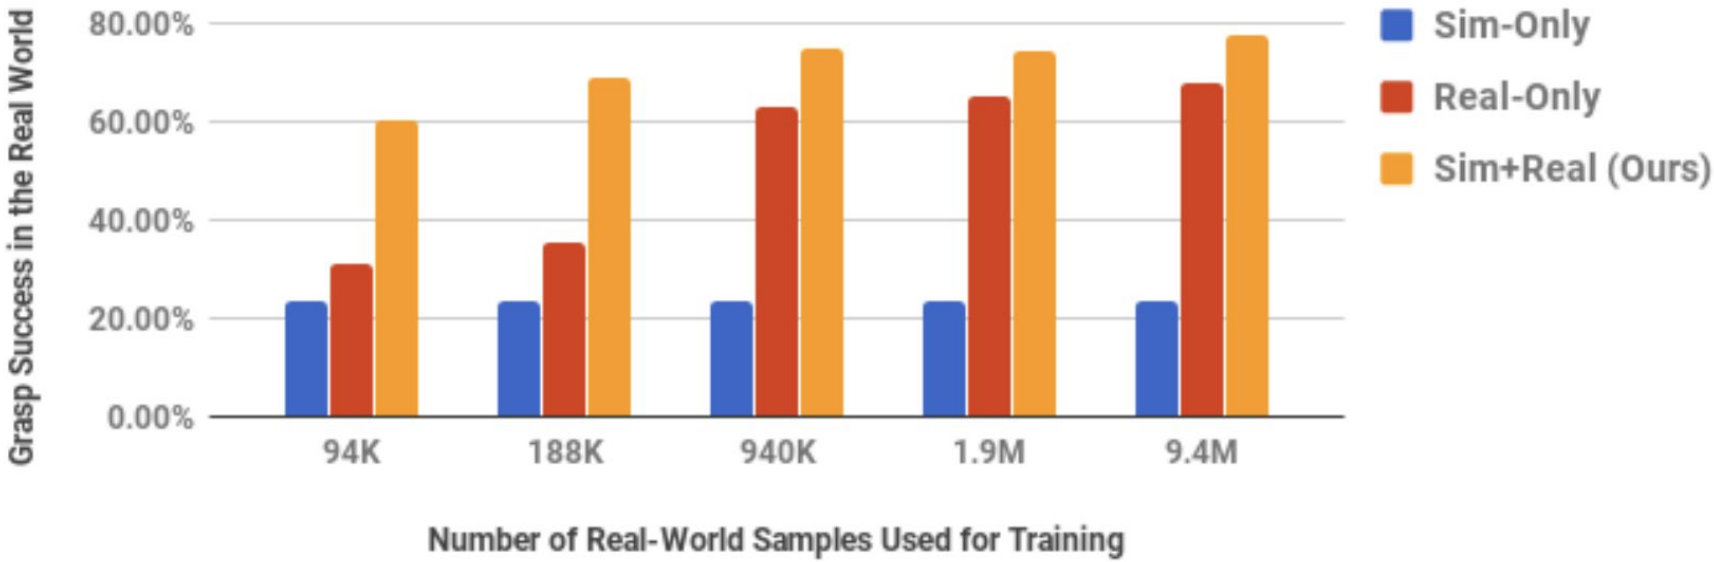
Grasp Success in the Real World
80.00%
60.00%
40.00%
20.00%
0.00%
94K
188K
940K
1.9M
9.4M
Number of Real-World Samples Used for Training
Sim-Only
Real-Only
Sim+Real (Ours)

Datasets

https://sites.google.com/site/brainrobotdata/home

Includes

- grasping, pushing, pouring
- 100,000s RGBD images + robot info (camera calibration, arm poses, etc)

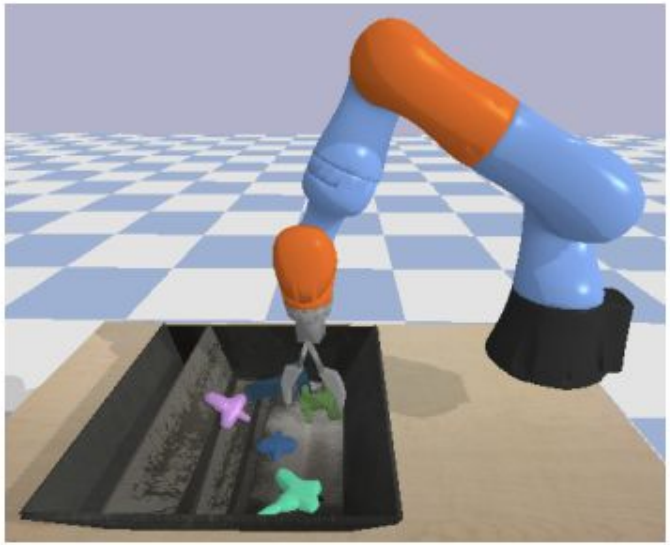

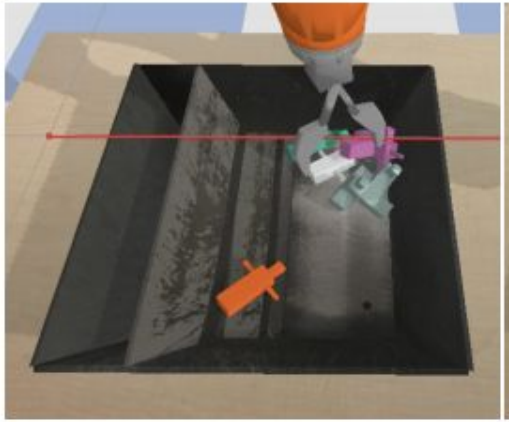

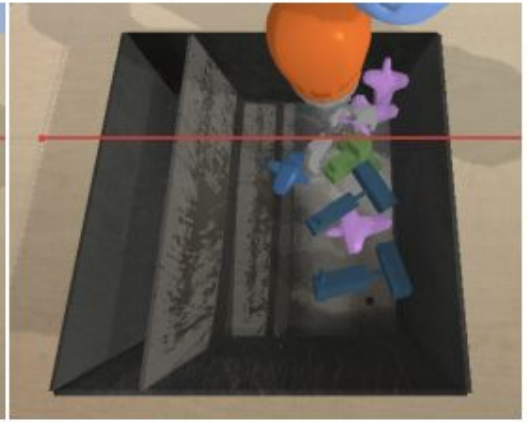

"Deep Reinforcement Learning for Vision-Based Robotic Grasping: A Simulated Comparative Evaluation of Off-Policy Methods" Quillen et al.