

# Multi-Task Domain Adaptation for Deep Learning of Instance Grasping from Simulation

Kuan Fang<sup>1,2</sup>, Yunfei Bai<sup>1</sup>, Stefan Hinterstoisser<sup>1</sup>, Silvio Savarese<sup>2</sup>, Mrinal Kalakrishnan<sup>1</sup>

**Abstract**—Learning-based approaches to robotic manipulation are limited by the scalability of data collection and accessibility of labels. In this paper, we present a multi-task domain adaptation framework for instance grasping in cluttered scenes by utilizing simulated robot experiments. Our neural network takes monocular RGB images and the instance segmentation mask of a specified target object as inputs, and predicts the probability of successfully grasping the specified object for each candidate motor command. The proposed transfer learning framework trains a model for instance grasping in simulation and uses a domain-adversarial loss to transfer the trained model to real robots using indiscriminate grasping data, which is available both in simulation and the real world. We evaluate our model in real-world robot experiments, comparing it with alternative model architectures as well as an indiscriminate grasping baseline.

## I. INTRODUCTION

With recent progress in deep learning and reinforcement learning, large scale learning-based methods have been used in robotic manipulation [1], [2], [3]. These methods enable us to replace manually designed perception and control pipelines with end-to-end neural networks learned from large training datasets. However, it is challenging to apply these methods on complex, real-world manipulation tasks for two reasons: First, collecting data through repeated robot experiments in the real world is time consuming. Second, evaluation and annotation of ground truth success labels can be expensive and sometimes infeasible. Although some previous works have designed experimental platforms to automatically collect and label data for non-semantic tasks such as indiscriminately grasping any object from the tray [1], it can quickly get more challenging in terms of data collection when switching to tasks with different settings.

One alternative solution to large scale data collection in robotics is to run robot experiments in simulation [4], [5], [6], [7], [8], [9]. Simulated robot experiments are parallelizable and can easily be reset. In addition, simulation also provides access to the ground-truth states of the robot and environment, making it straightforward to generate labeled data. However, learning from only simulated data does not yield robust real-world performance due to the reality gap between the real world and the simulation in terms of robotic perception and physics. In general, due to the phenomenon known as *domain shift* [10], a model trained on data from one domain does not generalize well to another target domain. A common approach to mitigate the domain shift is to use domain adaptation [11] to transfer the learned knowledge

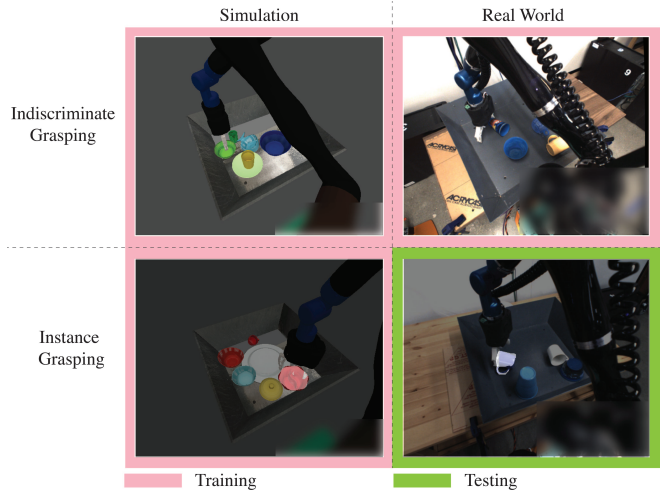


Fig. 1: **Multi-task domain adaptation for instance grasping.** Our multi-task domain adaptation framework trains a deep neural network for instance grasp prediction from simulation, while utilizing both real and simulated indiscriminate grasping data to mitigate the domain shift from simulation to the real world. Robot hardware details not relevant to this paper are blurred out in images to preserve confidentiality.

from one domain to another. In our context of transferring the learned policy from simulation to the real world (sim-to-real transfer), this transfer can be achieved by feeding both real and simulated data into the neural network and applying a similarity loss to regularize the difference between extracted features from the two domains. During this process, one usually still needs to collect sufficient amount of real-world data in order to learn the target task.

In this paper, we aim to address the problem of sim-to-real transfer where there is no labeled data available in the real world. Specifically, we address this problem in the context of *instance grasping*, where the goal is to grasp a particular object instance from a cluttered scene. It is more challenging in terms of both perception and control compared to what we refer to as *indiscriminate grasping*, where the robot can grasp any objects from the workspace indiscriminately. More importantly, unlike recent works on indiscriminate grasping [1], [2], to get labeled real-world data for instance grasping, we would need sophisticated perception system to robustly identify and track the target object across time. To overcome this hurdle, we propose to only use simulation to collect labeled data for instance grasping. Furthermore, to resolve the domain shift issue, we propose a multi-task domain adaptation framework, which trains a deep neural

<sup>1</sup>X, Mountain View, CA 94043 USA. <sup>2</sup>Stanford University, Stanford, CA 94305 USA. This research was conducted during Kuan’s internship at X.

network for instance grasp prediction from simulation, while utilizing both real and simulated indiscriminate grasping data to mitigate the domain shift (Fig. 1). The intuition of using indiscriminate grasping data is that indiscriminate and instance grasping share similar properties regarding both robot action space and perception. This framework shares model parameters between indiscriminate grasp prediction and instance grasp prediction, and trains both tasks simultaneously. Meanwhile, it enforces the extracted features to be transferable between the simulation and the real world through domain adaptation.

Our method is demonstrated on robot instance grasping using RGB images as inputs, which are captured from a moving monocular RGB camera mounted on the robot arm. We use Mask R-CNN [12] to get segmentation masks for object instances, one of which is then sampled to specify the target object for grasping. Our method only relies on a single segmentation mask obtained at the beginning of the grasp, rather than requiring the segmentation mask at each time-step. We evaluate our method by showing the robot grasping a variety of household dishware objects, as well as its generalization to unseen objects during instance grasp training in simulation.

Our main contributions are:

- A multi-task domain adaptation framework that trains a model for instance grasping in simulation and uses a domain-adversarial loss to transfer the trained model to real robots using indiscriminate grasping data, which is available both in simulation and the real world.
- A deep neural network architecture which takes monocular RGB images and the instance segmentation mask of a specified target object as inputs, and predicts the probability of successfully grasping the specified object for each candidate motor command.
- We demonstrate our method on a robot platform with a mounted moving monocular RGB camera. We showed instance grasping with a variety of household dishware objects, and the generalization to unseen objects of instance grasp training in simulation.

The rest of this paper is laid out as follows. In Section II we discuss related work. In Section III, we introduce the necessary background in end-to-end grasp prediction and domain adaptation upon which our work is based. We describe our main algorithmic contributions in Section IV. Experimental method and results on real robots are presented in Section V. Finally, we discuss potential directions for future work in Section VI.

## II. RELATED WORK

**Robotic grasping.** Grasping is one of the most fundamental robot manipulation tasks. Both geometry-based [13], [14], [15] and data-driven methods [16] have been proposed to address the problem of robotic grasping. Recent works [1], [2] have used deep learning to train models that can predict grasp success directly from visual sensor inputs, which has the benefit of generalization. Other works focused on semantic grasping [17], where the robot picks up an object

of a specified class. In our work, we propose a deep learning algorithm for instance grasping, and we use simulation to generate labeled instance grasping data.

**Transfer learning from simulation to real robots.** Deep learning methods for robotic manipulation usually require large amounts of data on the order of a million trials [1]. Recent works have considered using simulation to train robotic manipulation to scale up the data collection. Saxena *et al.* [4] used rendered objects to learn a vision-based grasping model. Rusu *et al.* [5] introduced progressive neural networks to adapt an existing deep reinforcement learning policy trained for a reaching task in simulation to the real world. However, due to differences in the dynamics and observation models (reality gap), solving the problem of sim-to-real transfer is the key to ensuring that policies trained in simulation can work in the real world [18].

One way for addressing sim-to-real transfer is to choose input modalities where the reality gap is not significant. One option is to use the depth image which abstracts away many appearance properties of real-world objects that are hard to render. A number of recent works used simulated depth images to learn indiscriminate grasping with deep convolutional neural networks (CNNs) [8], [9], then deploy the trained policy on the real robot with a calibrated fixed depth camera or a wrist mounted depth camera for indiscriminate grasping. Our work tackles instance grasping problem which is much more difficult. In addition, we use only monocular RGB images, without relying on either camera calibration or object geometry information. We believe there is considerable value in studying grasping systems which use RGB images only, considering the relatively low cost of monocular RGB cameras and the limitations of depth cameras (*e.g.*, unsuitable for bright outdoor lighting, transparent object, *etc.*).

Some recent works have shown the success of using domain randomization for transferring a deep neural network trained on simulated RGB images to the real world for solving end-to-end robot manipulation or mobility tasks [19], [6], [7]. These works extended prior works on data augmentation in computer vision [20], and applied randomization to parameters such as camera position, lighting, and texture in simulation. However, only simple geometries like cubes or free space motions are considered in those prior methods. In contrast, we address the problem of grasping diverse naturalistic household objects like bowls and cups, which are much more complex in terms of both perception and contact physics. Furthermore, we also demonstrate sim-to-real transfer with generalization to new objects that were never seen during training in simulation.

Domain adaptation is a process that allows a learning model trained with samples from a source domain, in our case the simulation domain, to generalize to a target domain, the real world. Feature-level domain adaptation focuses on learning domain-invariant features, either by learning a transformation of fixed, pre-computed features between source and target domains [21], [22], [23], [24] or by learning a domain-invariant feature extractor, often as a CNN [25], [26], [27], [28]. Tzeng *et al.* [29] perform sim-to-real transfer for




Fig. 2: **The pipeline of obtaining the segmentation mask for a target object.** Given the initial camera image (left), we forward the Mask R-CNN to predict the instance segmentation masks (middle). The detection bounding boxes and segmentation masks of different object instances are shown in different colors, with the assigned object IDs shown on the upper left corner. Then one target object is sampled and its segmentation mask is defined as the target mask (right).

a manipulation task using a adversarially trained domain discriminator, by doing nearest neighbor in the learned feature space to find weak pairings. A recent study [30] used combined feature-level and pixel-level domain adaption for learning robotic hand-eye coordination for grasping a variety of everyday objects, by training primarily with simulation data and small amount of labeled real-world data. Our work uses feature-level domain adaptation, particularly we propose a multi-task domain adaptation method to address the lack of labeled data for instance grasping in the real world.

**Instance segmentation and object detection.** Object detection and segmentation have been a long standing computer vision problem. Several recent deep learning techniques [12], [31], [32], [33], [34], [35], [36], [37] have made significant progress in this direction. In our work, we use Mask R-CNN [12] trained on synthetic images to segment objects from the background in real-world images. This gives us an easy way to specify the object to grasp without manually labeling objects in the image. Our proposed method could potentially also be trained to use a bounding box or a few pixels on the object as input instead of the instance mask. However, exploration of these alternative input modalities is left for future works.

### III. BACKGROUND

#### A. End-to-End Grasp Prediction

In this work, we follow the formulation and real-world experimental setup of end-to-end grasp prediction described in [1]. The goal is to choose a sequence of robot actions step by step given the current input RGB image, in order to control the robot gripper to grasp any object from the cluttered scene. We denote the grasp prediction network as  $g(I_0, I_t, v_t; \theta)$ , where  $I_0$  and  $I_t$  are the input RGB images from the robot camera at the initial time step and the current time step,  $v_t$  is a given robot command sampled from the action space,  $\theta$  represents the neural network parameters. At each time step during grasping, the network predicts the probability of grasp success for a set of sampled actions, and chooses the action with the highest predicted probability of success. We run the cross-entropy method (CEM) [38] for 3 iterations to find action sampling distributions centered around regions of high grasp success probability. In each

iteration, we fit a single Gaussian distribution to the top-ranked motor commands. A grasp is executed when the predicted grasp success probability is above a threshold.

The dataset for training the grasp prediction network is collected by controlling the gripper to repeatedly attempt to grasp objects from a tray with randomly selected objects and drop them back. This process starts with exploring a random policy at the beginning, and then switches to the CEM with the updated network parameters. In the real world, the ground truth success label is automatically determined by a hard coded perception system after each grasping episode. This is done by taking images of the tray and comparing pixel differences before and after dropping the object. In simulation, we determine the success label by checking the object position. The grasp prediction network is then trained using a log loss against the ground truth success labels.

#### B. Adversarial Loss for Domain Adaptation

We use the adversarial loss [25] during training to penalize the domain shift from simulation to the real world similar as in the recent study [30]. The adversarial loss regularizes the neural network weights  $\theta$  through a domain classifier  $D(I_0, I_t, v_t; \theta, \phi)$ , where  $\phi$  represents the parameters of the domain classifier. In our problem setup, the target domain is the real world and the source domain is simulation. During training, data from the two different domains are fed into the grasp prediction network. The domain classifier takes as input the intermediate feature extracted from the neural network, and is trained to predict which domain the feature is from, by maximizing the binomial cross-entropy [25] with respect to  $\phi$ . Meanwhile we minimize the adversarial loss with respect to  $\theta$ :

$$\mathcal{L}_{adversarial} = \sum_{i=0}^{N_S+N_R} d^i \log \hat{d}^i + (1-d^i) \log(1-\hat{d}^i) \quad (1)$$

where  $d^i \in \{0, 1\}$  is the ground truth domain label for each input  $I_0^i, I_t^i, v_t^i$  of sample  $i$ ,  $\hat{d}^i = D(I_0^i, I_t^i, v_t^i; \theta, \phi)$  is the predicted domain label,  $N_S$  and  $N_R$  are the sizes of the simulated and real datasets.




Fig. 3: **Multi-task domain adaptation framework.** The framework is composed of three grasp prediction towers and a domain classifier. Each tower takes training data from one of the three task domains: real-world indiscriminate grasping, simulated indiscriminate grasping and simulated instance grasping. The neural networks for the three towers share parameters as denoted by the dashed lines. In this framework, the neural network is trained to predict the instance grasp success probability with transferable features between the simulation and the real world by minimizing the four losses simultaneously with respect to the network parameters.

#### IV. MULTI-TASK TRANSFER LEARNING FROM SIMULATION FOR INSTANCE GRASPING

##### A. End-to-End Instance Grasp Prediction

The goal of instance grasp prediction is to predict the probability of successfully grasping the specified object for each candidate motor command. The target object is specified by a segmentation mask  $M_0$  given at the beginning of the episode before grasping starts. To predict the success probability for instance grasping, we define the instance grasp prediction network  $g(M_0, I_0, I_t, v_t; \theta)$  by extending the formulation for indiscriminate grasping [1]. In addition to RGB images and robot commands, the instance grasp prediction network also uses the segmentation mask  $M_0$  of the target object as input. In order to predict the instance grasp success given each sampled action, the network is required to understand the relative location between the gripper and the target object by fusing this information together.

For training, we only collect instance grasping data in simulation. It is crucial to keep the labels of successful and failed trials balanced during training. Directly running a random policy for data collection yields very few successful trials. To improve sample efficiency, we use a hindsight data collection trick inspired by [39]. Here we first run indiscriminate grasping data collection as described in [1]. For each successful indiscriminate grasping trial, we generate one successful instance grasping trial by taking the mask of the grasped object as the input target mask, and one failed instance grasping trial by using a mask sampled from other objects. As a result, about 20% of the grasping trials in

our training dataset are labeled as success for the task of instance grasping. During test time, instance grasping results are manually evaluated by human annotators.

Our pipeline to get a sampled segmentation mask is shown in Fig. 2. Given an initial RGB image  $I_0$  at the test time, we first run Mask R-CNN[12] to predict the instance segmentation masks, where each mask is assigned to a unique object instance ID. Then we sample one of the object IDs and get the corresponding segmentation mask  $M_0$  for the target object. Note that we only rely on the *initial* segmentation mask as the input to the instance grasp prediction neural network. Since the robot camera is attached to a pan-tilt unit mounted on the first link of the arm, it results viewpoint changing as the arm moves around. Therefore,  $M_0$  does not align with the target object in the following RGB images  $I_t$ . However, we still find  $M_0$  is effective to provide information about the location of the target object, by fusing both the segmentation and RGB images as input to the neural network, as discussed in Section IV-C.

##### B. Multi-Task Domain Adaptation

Our multi-task domain adaptation framework is shown in Fig. 3. The framework is composed of a domain classifier and three grasp prediction towers with training data from three task domains: real-world indiscriminate grasping data  $\mathbf{X}_R^A$ , simulated indiscriminate grasping data  $\mathbf{X}_S^A$ , and simulated instance grasping data  $\mathbf{X}_S^B$ . Here for simplicity of notation, we use  $S$  and  $R$  to represent the simulation and the real world, and use  $A$  and  $B$  to represent indiscriminate grasping and instance grasping respectively. The three grasp predic-

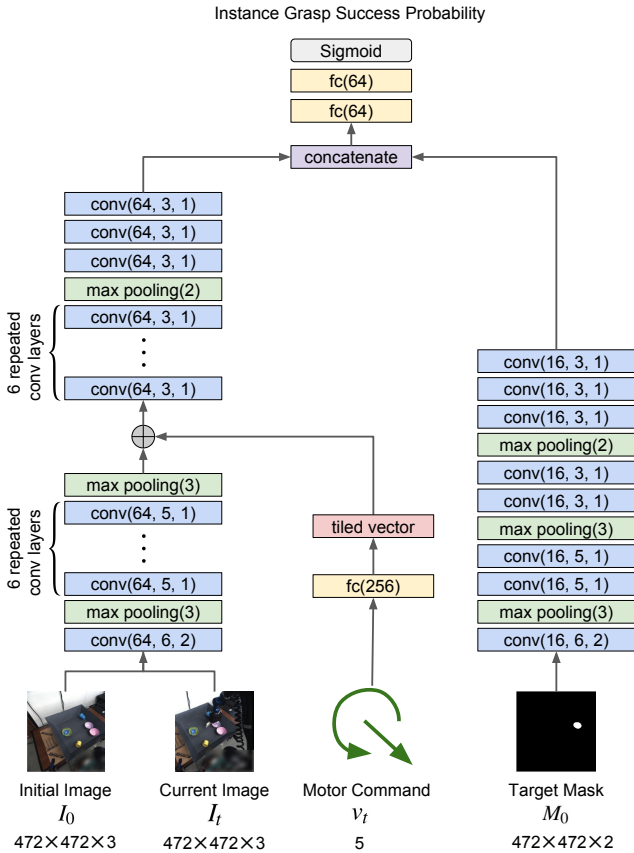


Fig. 4: **Instance Grasp Prediction Network.** Taking the initial and current camera images, the candidate motor command and the target mask as inputs, the network predicts the instance grasp success probability for the candidate motor command. Convolutional layer (conv) parameters are shown as (number of filters, kernel size, stride). Fully-connected layer (fc) parameters are shown as (number of channels). Max pooling layer parameters are shown as (stride).

tion towers are trained simultaneously by minimizing the following loss with respect to the neural network parameters  $\theta$ :

$$\mathcal{L} = \mathcal{L}_S^B + \alpha \mathcal{L}_R^A + \beta \mathcal{L}_S^A + \lambda \mathcal{L}_{adversarial} \quad (2)$$

where  $\lambda$ ,  $\alpha$  and  $\beta$  are weights of the loss terms. The instance grasp prediction loss  $\mathcal{L}_S^B$  trains the neural network to predict instance grasp success probability in the simulation. The two indiscriminate grasp prediction loss  $\mathcal{L}_R^A$  and  $\mathcal{L}_S^A$  train the neural network to extract meaningful features for grasp prediction both in the simulation and the real world. The adversarial loss  $\mathcal{L}_{adversarial}$  regularizes the neural network parameters through the domain classifier. By attempting to confuse the domain classifier, the neural network learns to extract features that are transferable between the simulation and the real world.

The three towers use the same instance grasp prediction network architecture and share all network parameters. Since the segmentation mask input is only available in the instance grasping task, we feed in a constant mask to the two indiscriminate grasping towers. The constant mask marks every




Fig. 5: **Representative synthetic image for training the Mask R-CNN.** Each image consist of a real background image and the objects rendered on top of it. The ground truth detection bounding box and the segmentation mask are overlaid on the rendered object.

pixel as a candidate object to be grasped. The interpretation of the mask is that any object covered by the mask can be grasped by the robot in order to achieve success, and this interpretation is consistent across both the instance grasping and indiscriminate grasping tasks. This enables us train on the three task domains with the same input modalities and avoid training additional layers only for the simulated instance grasping data.

### C. Neural Network Architecture

We show the architecture details of our instance grasp prediction network  $g(M_0, I_0, I_t, v_t; \theta)$  in Fig. 4. The original RGB images and target masks are  $640 \times 512$  and we cropped them into  $472 \times 472$  as input to the network. The candidate gripper motor command vector has 5 dimensions: a 3D translation vector, and a sine-cosine encoding of the change in orientation. We adopt the same layers from [1] to extract low level features of RGB images and motor commands. To extract the target mask features, we introduce a separate stream of convolutional layers. This mask stream maps the target mask to a convolutional feature which matches the spatial size of the features from the RGB stream. The features from the two streams are concatenated at the final convolutional layer and then two fully connected layers merge the information to predict instance grasp success probability.

To have fast convergence rate during training, convolutional layers and fully connected layers usually need batch normalizations [40]. However, as discussed in [30], sharing batch normalization parameters across domains can significantly harm the performance during testing in a domain adaptation framework. This is due to the inconsistent behaviors of batch normalization layers during training and testing. To resolve this problem, we remove all batch normalization layers in the grasp prediction network. Instead we use instance normalization [41] for convolutional layers and layer normalization [42] for fully connected layers, which also guarantees fast convergence rate but have consistent behaviors between training and testing.




Fig. 6: **Examples of our 3D models used for data collection in simulation.** We use totally 130 3D models scanned from real-world dishware objects including bowls, mugs, cups, teapots *etc.* The object textures are rendered in simulation.

## V. EXPERIMENTS

### A. Instance Segmentation Prediction with Mask R-CNN

In order to provide instance detection and segmentation, we apply Mask R-CNN [12] to the initial input image. Mask R-CNN is trained only with synthetically generated images [43], *i.e.* images which consist of a real background image and the objects rendered on top of it (shown in Fig. 5). We ensure that the training set uniformly covers the whole pose space of interest, including all possible translations and rotations under which we could potentially observe the objects. Although occlusion is not explicitly covered in the training set, the trained Mask-RCNN demonstrates the capability of detecting and segmenting object instances from cluttered scenes.

### B. Data Collection and Training

To collect real-world indiscriminate grasping data, we use an automatic data collection platform similar to [1]. We ran indiscriminate grasping on 9 Jaco robot arms, with a random policy as well as a CEM policy using a trained model, to collect 100,000 grasping trials on a total of around 100 unique dishware objects. We use standard data augmentation methods including random cropping and random image distortion to increase data diversity.

To collect both instance and indiscriminate grasping data in simulation, a basic virtual environment is built based on the Bullet physics simulator [44] and a simple software renderer shipped with Bullet is used for rendering. We scanned 130 3D models of household dishware objects (*e.g.* mugs, bowls, water bottles) as the objects to be grasped in simulation (shown in Fig. 6). The environment emulates the same Jaco robot arm setup by simulating the physics of grasping and by rendering what the moving camera mounted on the robot would perceive: the robot arm and gripper, the bin that contains the objects, and the dishware objects to be grasped. We collect 1 million indiscriminate grasping trials using a random policy and the CEM policy on iteratively trained models, with 1 to 6 objects in the tray per episode. The data is then labeled post-hoc for instance grasping as described in Sec. IV-A.



Fig. 7: **Our testing objects in the real world.** In each grasping trial, a combination of 5 objects is chosen and put into the tray in front of the robot. The target object is sampled from the detected objects in the tray using the Mask-RCNN.

We train on 10 GPUs for 300k iterations with a learning rate of 0.0001 which is decayed by 0.94 every 50k iterations, and a momentum of 0.9. We use a batch size of 32 for data from each of the three task domains. The training starts with only the three grasp prediction losses and the adversarial loss is added after 50k iterations. The weights of the loss terms in IV-B are chosen as  $\alpha = 1$ ,  $\beta = 1$ ,  $\gamma = 4$ . All of the hyperparameters are selected through experiments in [30] without tuning on any instance grasping validation datasets.

### C. Evaluation of Instance Grasping

To evaluate the instance grasping performance in the real world, we run 250 trials for each method across 5 Jaco robot arms. In each trial we place a combination of 5 objects in the tray. We use the same 10 object combinations chosen from 16 dishware objects with various colors and shapes as shown in Fig. 7. The objects chosen are present in the real-world *indiscriminate* grasping dataset, however they are held out from the *instance* grasping dataset collected in simulation. Table I shows the detail of the results. Our framework achieves 60.8% success rate for instance grasping. Among the failed instance grasping attempts, 26 attempts ended up with grasping the wrong object, 72 attempts failed grasping anything from the tray. Further details and videos can be found at <https://sites.google.com/view/multi-task-domain-adaptation/>

### D. Ablative Analysis

We run a number of ablations with the same evaluation settings to analyze the instance grasping results of multi-task domain adaptation framework. Detailed comparisons are shown in Table I and Fig. 8.

#### **Indiscriminate vs. Instance Grasp Prediction Network:**

The instance grasp prediction network controls the gripper to grasp the specified target object. While the indiscriminate grasp prediction network does not distinguish objects. To show this difference, we run indiscriminate grasping trials using a network trained only for indiscriminate grasp prediction with the sim-to-real domain adaptation algorithm for single task [30]. With this indiscriminate policy the robot gripper

TABLE I: **Performance of instance grasping in real-world experiments.** For each method, we run 250 trials across 5 Jaco robot arms. We show numbers of successful instance grasps, wrong object grasps, failed grasps, and the instance grasping success rate. Our method achieves the highest instance grasping success rate among the four as described in V-C and V-D.

Method	Successful Instance Grasps	Grasped Wrong Objects	Failed Grasps	Instance Grasping Success Rate
Indiscriminate	34	169	47	13.6%
Two-Tower	87	31	132	34.8%
No Constant Mask	146	24	80	58.4%
Our Method	152	26	72	60.8%

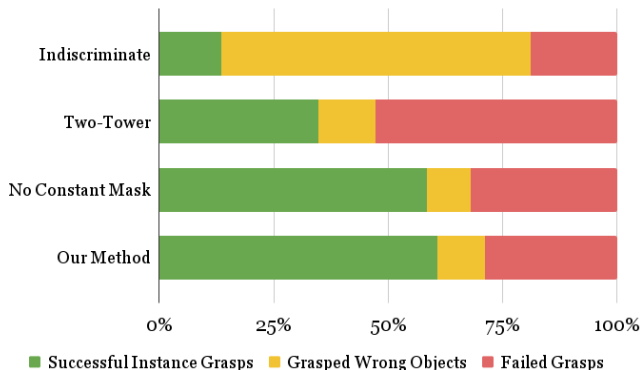


Fig. 8: **Failure Analysis.** We show the breakdown of successful instance grasps, wrong object grasps, and failed grasps for each of the methods described in V-C and V-D.

successfully grasps objects from the tray in 81.2% of the trials. Since no targets are specified, many grasps only aim for easy objects. To show the discrepancy of the objectives between the two tasks, we also sample target masks and evaluate if the grasp hits the target in each trial as in the instance grasping task. Only 13.6% of the trials end up grasping the correct target objects.

**Two-Tower vs. Three-Tower:** The transferring ability of the adversarial loss comes from training the network to confuse the adversarial loss. The domain shift can be difficult to resolve when using only the real world indiscriminate grasping data and simulated instance grasping data, since both the objective and the perception for the two task domains are different. To demonstrate this, we take the two-tower domain adaptation framework from [25] using data only from these two task domains. The adversarial loss is applied to the real world indiscriminate grasping data and simulated instance grasping data, without using simulated indiscriminate grasping data to bridge the gap. The instance grasp success rate of the two-tower framework is 34.8%. Comparing with the three-tower model, the two-tower model failed in grasping any objects in more trials. This suggests that our framework transfers the learned instance grasping policy better than the two-tower framework.

**Using Constant Masks vs. Training Additional Layers:** When dealing with the different input modalities between indiscriminate and instance grasping, one alternative to using constant masks is using different network architectures for the two tasks. However, additional layers have to be used for instance grasping since the extracted features can be very different with and without the information of the target

mask. We thus modify our framework by only sharing the convolutional layer parameters of the RGB stream between the instance and indiscriminate networks, but train new fully connected layers for instance grasping. With layers only trained in simulation, this framework achieves 58.4% grasping success rate for instance grasping. This shows the performance gains of introducing constant masks for using same network architectures.

## VI. CONCLUSION

In this work, we introduce a multi-task domain adaptation framework for instance grasping in cluttered scenes using simulated robot experiments. Our framework only collects labeled instance grasping data from simulation, and uses a domain-adversarial loss to transfer the trained model to real robots using indiscriminate grasping data from the simulation and the real world. We also presented a deep neural network, which takes monocular RGB images and instance segmentation mask as input, and predicts the probability of successfully grasping the specified object for each candidate motor command. We demonstrated our method on a real robot platform for grasping household dishware objects that were never seen during the training in simulation. And we showed that our model outperforms alternative model architectures and indiscriminate grasping baseline.

One limitation of using Mask R-CNN to predict segmentation masks is that it does not work well when objects are occluded by the robot gripper, which happens frequently during grasping. In addition, we cannot run Mask R-CNN for every single time step due to its computational expenses. Therefore, we only predict the segmentation mask for the initial RGB image. While we find our model can extract effective information from the single segmentation mask, it would be interesting to see how the instance grasping performance can be further improved by updating the segmentation masks for the later time steps.

In our work, we use the segmentation mask as the input modality for specifying the target object. There can be other options for specifying the target objects such as detection bounding boxes and user-specified pixels on the object. It will be straightforward to extend our method to these alternative input modalities. Although we choose instance grasping and indiscriminate grasping in our problem setup, our framework can potentially be applied to other tasks that share similar properties, *e.g.* robot picking and placing. Incorporating our framework to other robot tasks will be an exciting avenue for future works.

## ACKNOWLEDGMENT

We thank Adrian Li, Peter Pastor, Ian Wilkes, Kurt Konolige for contributions to the software and hardware development for the grasping infrastructures. We thank Alex Irpan, Paul Wohlhart, Konstantinos Bousmalis for discussions on training and implementation of the domain classifier. We thank John-Michael Burke for the supports of the real-world robot experiments.

## REFERENCES

- [1] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International Journal of Robotics Research*, 2016.
- [2] P. L. and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in *IEEE International Conference on Robotics and Automation*, 2016, pp. 3406–3413.
- [3] D. Kappler, J. Bohg, and S. Schaal, "Leveraging Big Data for Grasp Planning," in *IEEE International Conference on Robotics and Automation*, 2015.
- [4] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic grasping of novel objects using vision," *The International Journal of Robotics Research*, vol. 27, no. 2, pp. 157–173, 2008.
- [5] A. A. Rusu, M. Vecerik, T. Rothörl, N. Heess, R. Pascanu, and R. Hadsell, "Sim-to-real robot learning from pixels with progressive nets," *arXiv preprint arXiv:1610.04286*, 2016.
- [6] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," *arXiv preprint arXiv:1703.06907*, 2017.
- [7] S. James, A. J. Davison, and E. Johns, "Transferring end-to-end visuomotor control from simulation to real world for a multi-stage task," *arXiv preprint arXiv:1707.02267*, 2017.
- [8] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," *arXiv preprint arXiv:1703.09312*, 2017.
- [9] U. Viereck, A. t. Pas, K. Saenko, and R. Platt, "Learning a visuomotor controller for real world robotic grasping using easily simulated depth images," *arXiv preprint arXiv:1706.04652*, 2017.
- [10] J. Quiñero-Candela, M. Sugiyama, A. Schwaighofer, and N. Lawrence, "Covariate shift and local learning by distribution matching," 2008.
- [11] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [12] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," *arXiv preprint arXiv:1703.06870*, 2017.
- [13] A. Bicchi, "On the Closure Properties of Robotic Grasping," *The Intl. Journal of Robotics Research*, 1995.
- [14] A. Rodriguez, M. T. Mason, and S. Ferry, "From Caging to Grasping," *The International Journal of Robotics Research*, vol. 31, no. 7, pp. 886–900, 2012.
- [15] Y. Bai and C. K. Liu, "Dexterous manipulation using both palm and fingers," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 1560–1565.
- [16] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 289–309, 2014.
- [17] E. Jang, S. Vijaynarasimhan, P. Pastor, J. Ibarz, and S. Levine, "End-to-end learning of semantic grasping," in *Conference on Robot Learning*, 2017.
- [18] M. E. Taylor and P. Stone, "Transfer learning for reinforcement learning domains: A survey," *Journal of Machine Learning Research*, vol. 10, no. Jul, pp. 1633–1685, 2009.
- [19] F. Sadeghi and S. Levine, "Cad2rl: Real single-image flight without a single real image," *arXiv preprint arXiv:1611.04201*, 2016.
- [20] P. Y. Simard, D. Steinkraus, J. C. Platt, et al., "Best practices for convolutional neural networks applied to visual document analysis," in *ICDAR*, vol. 3, 2003, pp. 958–962.
- [21] B. Sun, J. Feng, and K. Saenko, "Return of frustratingly easy domain adaptation," in *AAAI*, 2016.
- [22] B. Gong, Y. Shi, F. Sha, and K. Grauman, "Geodesic flow kernel for unsupervised domain adaptation," in *CVPR*. IEEE, 2012, pp. 2066–2073.
- [23] R. Caseiro, J. F. Henriques, P. Martins, and J. Batista, "Beyond the shortest path: Unsupervised Domain Adaptation by Sampling Subspaces Along the Spline Flow," in *CVPR*, 2015.
- [24] R. Gopalan, R. Li, and R. Chellappa, "Domain Adaptation for Object Recognition: An Unsupervised Approach," in *ICCV*, 2011.
- [25] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *Journal of Machine Learning Research*, vol. 17, no. 59, pp. 1–35, 2016.
- [26] M. Long and J. Wang, "Learning transferable features with deep adaptation networks," *ICML*, 2015.
- [27] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell, "Deep domain confusion: Maximizing for domain invariance," *Preprint arXiv:1412.3474*, 2014.
- [28] K. Bousmalis, G. Trigeorgis, N. Silberman, D. Krishnan, and D. Erhan, "Domain separation networks," in *Proc. Neural Information Processing Systems (NIPS)*, 2016.
- [29] E. Tzeng, C. Devin, J. Hoffman, C. Finn, P. Abbeel, S. Levine, K. Saenko, and T. Darrell, "Adapting deep visuomotor representations with weak pairwise constraints," *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.
- [30] K. Bousmalis, A. Irpan, P. Wohlhart, Y. Bai, M. Kelcey, M. Kalakrishnan, L. Downs, J. Ibarz, P. Pastor, K. Konolige, S. Levine, and V. Vanhoucke, "Using simulation and domain adaptation to improve efficiency of deep robotic grasping," *arXiv preprint arXiv:1709.07857*, 2017.
- [31] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 580–587, 2014.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Spatial pyramid pooling in deep convolutional networks for visual recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, pp. 1904–1916, 2014.
- [33] R. B. Girshick, "Fast r-cnn," *2015 IEEE International Conference on Computer Vision (ICCV)*, pp. 1440–1448, 2015.
- [34] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [35] J. Redmon, S. K. Divvala, R. B. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 779–788, 2016.
- [36] E. Shelhamer, J. Long, and T. Darrell, "Fully convolutional networks for semantic segmentation," *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3431–3440, 2015.
- [37] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE transactions on pattern analysis and machine intelligence*, 2017.
- [38] R. Y. Rubinstein and D. P. Kroese, "The cross-entropy method: A unified approach to monte carlo simulation, randomized optimization and machine learning," *Information Science & Statistics*, Springer Verlag, NY, 2004.
- [39] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba, "Hindsight experience replay," *CoRR*, vol. abs/1707.01495, 2017.
- [40] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International Conference on Machine Learning*, 2015, pp. 448–456.
- [41] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," *arXiv preprint arXiv:1607.08022*, 2016.
- [42] J. Ba, R. Kiros, and G. E. Hinton, "Layer normalization," *CoRR*, vol. abs/1607.06450, 2016.
- [43] S. Hinterstoisser, V. Lepetit, P. Wohlhart, and K. Konolige, "On pre-trained image features and synthetic images for deep learning," *arXiv preprint arXiv:1710.10710*, 2017.
- [44] E. Coumans and Y. Bai, "pybullet, a python module for physics simulation, games, robotics and machine learning," <http://pybullet.org/>, 2016–2017.