

Sim-and-Real Reinforcement Learning for Manipulation: A Consensus-based Approach

Wenxing Liu^{1,3}, Hanlin Niu^{1,3}, Wei Pan², Guido Herrmann³ and Joaquin Carrasco³

Abstract—Sim-and-real training is a promising alternative to sim-to-real training for robot manipulations. However, the current sim-and-real training is neither efficient, i.e., slow convergence to the optimal policy, nor effective, i.e., sizeable real-world robot data. Given limited time and hardware budgets, the performance of sim-and-real training is not satisfactory. In this paper, we propose a Consensus-based Sim-And-Real deep reinforcement learning algorithm (CSAR) for manipulator pick-and-place tasks, which shows comparable performance in both sim-and-real worlds. In this algorithm, we train the agents in simulators and the real world to get the optimal policies for both sim-and-real worlds. We found two interesting phenomena: (1) Best policy in simulation is not the best for sim-and-real training. (2) The more simulation agents, the better sim-and-real training. The experimental video is available at: <https://youtu.be/mcHJtNISTEQ>.

I. INTRODUCTION

As an essential component in robotic control, deep reinforcement learning (DRL) has been widely used in various applications [1], [2], [3]. The training process of DRL [4] builds the bridge between the environment state and the action, thereby maximizing the cumulative reward. Learning from the simulation is safer, cheaper and faster while learning from the real world is more dangerous, expensive and slower. If the simulation shows high fidelity, the training model in the simulation can be transferred directly to the real world. However, in many circumstances, the simulation cannot mimic the real world very well, which limits robot performance in the real world. To overcome this difficulty, we develop a sim-and-real training method to balance the relationship between the simulation and the real world. We use concepts from control engineering, i.e. consensus [5], [6], to accomplish sim-and-real training.

In this work, we propose a CSAR algorithm that combines consensus-based training with DRL in a sim-and-real environment, as shown in Fig. 1. We apply CSAR to a group of simulated agents together with a real agent each learning to carry out a pick-and-place task with a suction robot device. Compared to conventional sim-to-real training method, the

This work was supported by EPSRC project No. EP/S03286X/1, EPSRC RAIN project No. EP/R026084/1, EPSRC RNE project No. EP/P01366X/1 and UKAEA/EPSRC Fusion Grant 2022/2027 No. EP/W006839/1.

¹Wenxing Liu and Hanlin Niu are with Remote Applications in Challenging Environments (RACE), United Kingdom Atomic Energy Authority, Culham, UK. (Corresponding author: Hanlin Niu, hanlin.niu@ukaea.uk).

²Wei Pan is with the Department of Computer Science, The University of Manchester, Manchester, UK.

³Wenxing Liu, Hanlin Niu, Guido Herrmann and Joaquin Carrasco are with the Department of Electrical & Electronic Engineering, The University of Manchester, Manchester, UK.




Fig. 1: Pick-and-place objects with the CSAR approach

challenges of CSAR DRL are 1) information exchange between simulated and real robots, for instance, generating communication in a mixed environment, 2) data-efficient collection for training in a sim-and-real environment such as handling data from multiple robots simultaneously, 3) data pre-labelling for suctioning in a sim-and-real environment, for example, using aruco makers to locate suctioned objects. The main contributions of this paper can be concluded as follows:

- 1) A complete CSAR method is proposed for manipulators to learn pick-and-place tasks. By applying consensus-based training, the proposed method saves training time and reduces the number of required real robot training steps while maintaining a comparable suction success rate, which is cost-effective.
- 2) An end-to-end and lightweight neural network is proposed to train the suction policy, which uses raw 3D visual data directly without pre-labelling. The effectiveness and feasibility of the CSAR method are validated through simulation and real-world experiments.
- 3) We extend the consensus approach [7] from theory and simulations to a real-world pick-and-place problem and show the effectiveness of the proposed approach.

The structure of the paper proceeds as follows. Section II elucidates the related work. Section III details the CSAR algorithm. Experimental validation is given in Section IV to highlight the feasibility of our proposed algorithm. Section V summarizes this paper.

II. RELATED WORK

Sim-to-real: When a DRL model is transferred from simulation to the real world, the adoption problem becomes challenging as real-world environments contain unpredictable disturbances [8]. Fine-tuning has been widely used to bridge the gap between simulated and real environments [9], [10], [11]. However, fine-tuning usually takes a long time to perform parameter adaptation, which increases the experimental cost. Some recent works use only simulation but work well in

the real world. For instance, with only simulation, a distance function was trained in [12] between the current pose and the nearest optimal pose. In [13], a grasp quality network was proposed to evaluate robust grasp configuration based on the antipodal grasping sampling method. The key idea of these two papers is to use depth data rather than RGB images since depth images contain less information. Nevertheless, it is challenging for a depth camera to measure thin, dark colour objects because of their physical properties in the real world. Under this condition, performance cannot be guaranteed. Our approach focuses on reducing the gap between simulation and the real world, which is more general and flexible.

Sim-and-real: Sim-and-real training is a recent topic which focuses on adjusting the simulation according to real samples [14], [15]. A novel domain adaptation approach for robot perception was developed in [16] to close the sim-and-real gap by finding common features of real and synthetic data. In [17], the agent’s parameters in the simulation were updated to match the behaviour in the real world. Compared with [17], our approach is capable of dealing with situations that are hard to simulate precisely. In [18], one agent was used to select a simulated or real environment with a given probability and to interact. Transitions from all environments were stored in a common replay buffer to update training parameters. In contrast to [18], our method creates a sim-and-real environment directly with consensus, which avoids any form or transition but runs simulated and real agents in parallel. This is unique and has not been done before.

Reinforcement learning for manipulation: Reinforcement learning has been exploited to deal with robotic tasks [7], [19], [20]. In our work, we focus on improving training efficiency and saving real-world training costs of sim-and-real DRL for robotic manipulators. An efficient real-time hybrid path planning scheme was proposed in [21] to handle the uncertain dynamics of a robot manipulator by combining the probabilistic roadmap method with DRL. How to modulate the elementary movement of a robot arm through meta-parameters using reinforcement learning was proposed in [22]. In [23], a robotic manipulator was trained using DRL to solve the task of grasping an initially invisible object via a sequence of grasping and pushing actions. A high-precision peg-in-hole target task was selected in [24] for force-controlled robotic assembly with DRL. Specifically, the force and moment of the robotic manipulator end effector were chosen as the state. Nevertheless, the authors target solving specific low-level tasks such as motion planning in the work mentioned above. Our method pays attention to high-level tasks by treating each robot manipulator as an agent, which is more general and has a wider range of applications.

III. METHODOLOGY

We extend the consensus-based approach in [7], which only focuses on simulations, to sim-and-real scenarios. The proposed effective and efficient CSAR method can increase sim-and-real training speed as well as save real-world training costs with consensus-based training.

A. System Overview

Fig. 2 describes the overview of our proposed framework. The predefined workspace in the simulation is captured by a fixed simulated camera, which provides an ideal RGB-D image each time. Then the ideal RGB-D image is orthographically projected in the direction of gravity to construct the colour heightmap \bar{c}_t and the depth heightmap \bar{d}_t , which are the inputs of our framework. Both heightmaps are fed into the Q-function neural network to anticipate pixel-wise best suction position $[\bar{x}_t, \bar{y}_t]$. Given the specific use of these neural networks modelling the Q-function for pick and place success through suction gripping, we may call these “suction networks”. The suction height \bar{z}_t can be found from \bar{d}_t .

When it comes to the real world, the predefined workspace is captured by a fixed azure kinect camera. Compared with the ideal RGB-D image which is obtained from the simulated camera, the real-world RGB-D image contains more camera distortion [25]. Similarly, the real-world RGB-D image is orthographically projected in the direction of gravity to construct the colour heightmap \tilde{c}_t and the depth heightmap \tilde{d}_t which are also fed into the suction network to predict real-world pixel-wise best suction position $[\tilde{x}_t, \tilde{y}_t]$. The suction height \tilde{z}_t can be also acquired from \tilde{d}_t .

After performing predictions in both environments, consensus-based training is applied to the training parameters of each simulated or real agent. The suction process of each agent is carried out in parallel, which saves training time.

B. DRL Setup

1) *Action Space:* As stated in Section A, the action space a_t is a Cartesian motion command that consists of pixel-wise best suction position. In the simulated environment, $\bar{a}_t = [\bar{x}_t, \bar{y}_t, \bar{z}_t]$. Correspondingly, $\tilde{a}_t = [\tilde{x}_t, \tilde{y}_t, \tilde{z}_t]$ in the real world. The suction height \bar{z}_t and \tilde{z}_t can be acquired from \bar{d}_t and \tilde{d}_t .

2) *State Space:* As shown in Fig. 2, the state space s_t denotes the notation of the colour heightmap and depth heightmap of the captured RGB-D image. In the simulated environment, \bar{c}_t and \bar{d}_t are acquired by the fixed simulated camera. In the real environment, \tilde{c}_t and \tilde{d}_t can be obtained from the fixed azure kinect camera.

3) *Reward Space:* The distance μ_m in the simulated environment can be computed by

$$\mu_m = \sqrt{(\bar{x}_m - \tau_m)^2 + (\bar{y}_m - \sigma_m)^2} \quad (1)$$

where τ_m and σ_m denote x, y positions of the centre of the expected suctioned object of the m^{th} agent, respectively.

We assign suction reward $r_s = 1$ if the target is successfully suctioned, otherwise $r_s = 0$. Thus, the DRL reward \bar{r}_m for each agent in the simulation can be defined as

$$\bar{r}_m = \begin{cases} r_s r_0 & \text{if } \mu_m \leq \mu_{th} \\ r_s r_1 & \text{if } \mu_{th} < \mu_m \leq 2\mu_{th} \\ r_s r_2 & \text{if } 2\mu_{th} < \mu_m \leq 3\mu_{th} \\ r_s r_3 & \text{if } \mu_m > 3\mu_{th} \end{cases} \quad (2)$$

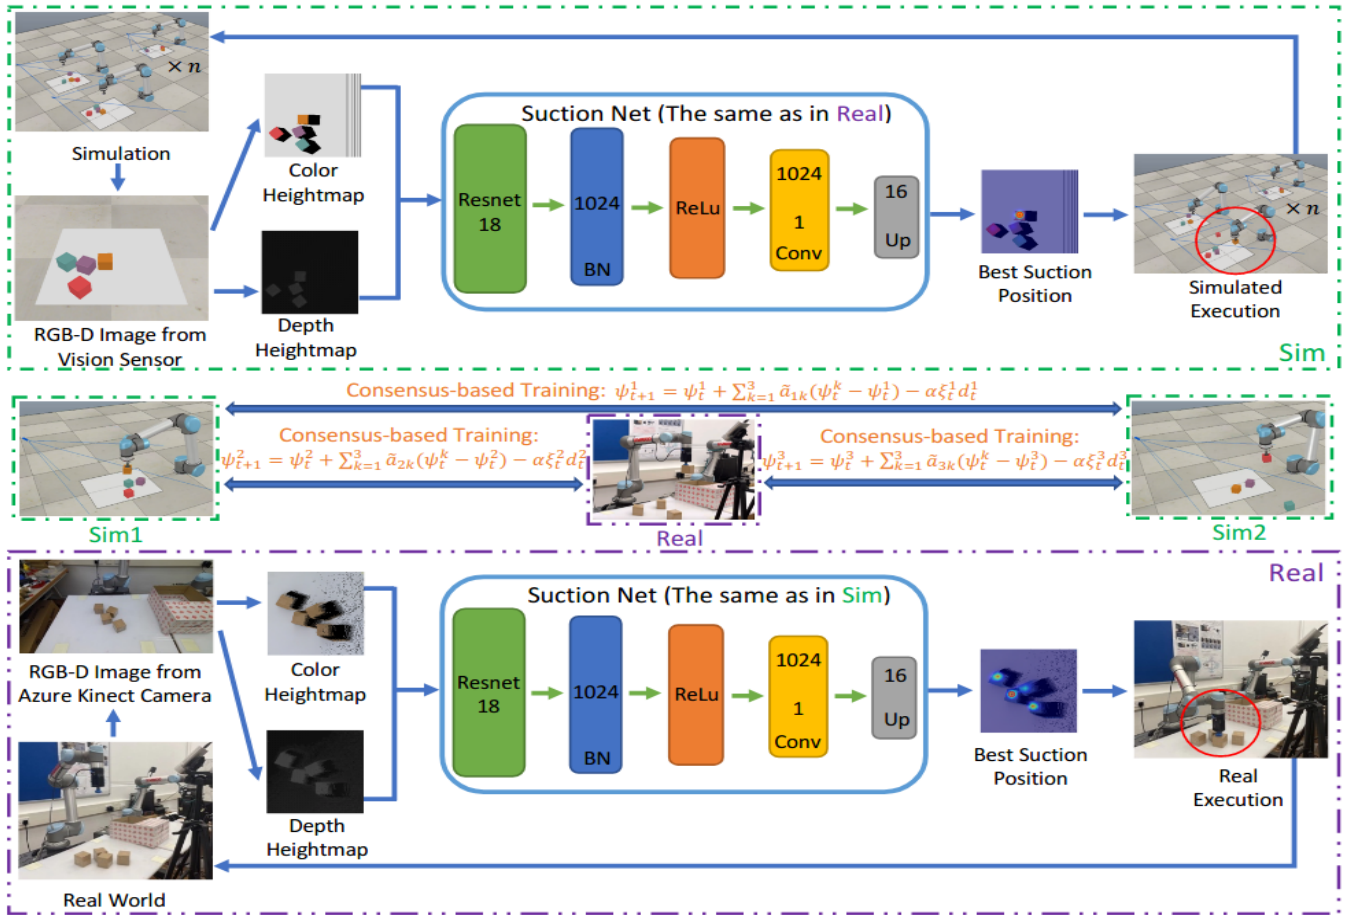


Fig. 2: Overview of the proposed DRL framework with consensus-based training in the sim-and-real environment (substantiation of Figure. 1). During each iteration, consensus-based training is applied to the training parameters of every suction net (multi-layer neural network modelling the Q-function for pick-and-place success through suction gripping). The suction executions occur simultaneously in both simulated and real environments. BN represents Batch Normalization. Conv stands for Convolution. Up represents Upsampling. More details can be found in Algorithm 1.

where \bar{r}_m stands for the reward of the m^{th} agent in the simulated environment, μ_{th} represents distance threshold of the m^{th} agent, r_0, r_1, r_2 and r_3 are the positive reward when μ_m is within the corresponding range.

The DRL reward \tilde{r}_m for each agent in the real environment is given by

$$\tilde{r}_m = r_s r_0 \quad (3)$$

4) *Neural Network Structure*: As stated in Fig. 2, the input of the suction net passes data through ResNet-18 [26] to extract concatenated features from the colour heightmap and the depth heightmap. The aforementioned features are fed into a Batch Normalization layer [27] with 1024 input features, a ReLu layer [27], a Convolution layer [27] with 1024 input channels, and 1 output channel, then are processed by a bilinear upsample layer [27] with a scale factor of 16. The output of the suction net has the same image size as the heightmap input, which is a dense pixel-wise map of different Q values. The pixel which has the maximum Q value represents the best suction position.

Remark 1: It should be noted that the suction net can be

substituted by any state-of-the-art neural network. Since we use a standard laptop for training, we purposely design a lightweight version of the suction net inspired by [28].

During each training iteration t , the training objective is to minimize the temporal difference error ξ_t [29]:

$$\xi_t = Q(\psi_t, s_t, a_t) - Y_t \quad (4)$$

where $Y_t = r_{t+1} + \gamma \max_a Q(\psi_t^-, s_{t+1}, a)$ and a represents all available actions, γ stands for the discount factor, Q represents the action-value function, r is the reward, ψ_t stands for the training parameters of the suction network at time t , ψ_t^- denotes the target training parameters.

5) *Loss function*: Inspired by [28], we use the Huber loss function [30] to train our proposed suction network in both simulated and real environments. The loss function Ω at the t^{th} iteration can be computed as follows:

$$\Omega_t = \begin{cases} \frac{1}{2}(\xi_t)^2 & \text{if } |\xi_t| < 1 \\ |\xi_t| - \frac{1}{2} & \text{otherwise} \end{cases} \quad (5)$$

Gradients are only passed through the single pixel on

Algorithm 1 CSAR: Consensus-based Sim-and-Real DRL

```

1: Initialize the  $m^{\text{th}}$  agent training parameter  $\psi_t^m$ , learning
   rate  $\alpha$ , RGB-D image  $\bar{g}_t^m$  from the simulation, initial
   RGB-D image  $\tilde{g}_t^m$  from the real world, discounted factor
    $\gamma$ , total training steps parameter  $T$ .
2: while  $t < T$  do
3:   Generate  $\bar{c}_t^m$  and  $\bar{d}_t^m$  from  $\bar{g}_t^m$ .
4:   Generate  $\tilde{c}_t^m$  and  $\tilde{d}_t^m$  from  $\tilde{g}_t^m$ .
5:   if object count  $O_t^m <$  empty threshold then
6:     Feed  $\bar{c}_t^m$  and  $\bar{d}_t^m$  into the  $m^{\text{th}}$  suction network
     to generate action-value function  $Q(\psi_t^m, \bar{s}_t^m, \bar{a}_t^m)$ .
7:     Feed  $\tilde{c}_t^m$  and  $\tilde{d}_t^m$  into the  $m^{\text{th}}$  suction network
     to generate action-value function  $Q(\psi_t^m, \tilde{s}_t^m, \tilde{a}_t^m)$ .
8:     if  $t > 2$  then
9:       Generate  $r_t^m$  with  $Q(\psi_{t-1}^m, \bar{s}_{t-1}^m, \bar{a}_{t-1}^m)$  and
        $Q(\psi_{t-1}^m, \tilde{s}_{t-1}^m, \tilde{a}_{t-1}^m)$ .
10:      Compute  $\xi_{t-1}^m$ :
11:         $Y_{t-1}^m = r_t^m + \gamma \max_a (Q(\psi_{t-1}^m, s_t^m, a^m))$ .
12:         $\xi_{t-1}^m = Q(\psi_{t-1}^m, s_{t-1}^m, a_{t-1}^m) - Y_{t-1}^m$ .
13:      For  $M$  agents, update the training parameters
        $\psi_t$  with consensus-based training:
14:         $\psi_t = \mathcal{C}(\psi_{t-1}, \mathcal{L}) - \alpha \Gamma_{t-1}$ .
15:      Sample a batch from the replay buffer  $R_p$  to
       implement experience replay.
16:    end if
17:    Perform suction execution in both simulated and
       real environments in parallel.
18:    Store  $(\bar{c}_t^m, \bar{d}_t^m, \bar{a}_t^m)$  and  $(\tilde{c}_t^m, \tilde{d}_t^m, \tilde{a}_t^m)$  in  $R_p$ .
19:  else
20:    Reposition objects.
21:  end if
22: end while

```

which the action is executed during each iteration t . All other pixels propagate with 0 loss [28].

C. Consensus-based Training

The Q-function of each simulated or real agent is trained through a consensus based algorithm. Hence, we wish to introduce at first the consensus network structure which facilitates that training process. The interaction topology of M agents can be depicted by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} represents a vertex set $\mathcal{V} = \{1, 2, \dots, M\}$ and \mathcal{E} stands for an edge set $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$. The edge $(j, m) \in \mathcal{E}$ if the j^{th} and m^{th} agents are connected with one another [31]. The adjacency matrix \mathcal{A} of \mathcal{G} can be described as $\mathcal{A} = [a_{jm}] \in \mathbb{R}^{M \times M}$, where $a_{jm} > 0$ if $(j, m) \in \mathcal{E}$, otherwise $a_{jm} = 0$. Hence, the Laplacian matrix \mathcal{L} of \mathcal{G} is defined as $\mathcal{L} = \mathcal{D} - \mathcal{A}$, where $\mathcal{D} = \text{diag}\{d_{11}, \dots, d_{MM}\} \in \mathbb{R}^{M \times M}$ and $d_{jj} = \sum_{j \neq m} a_{jm}$ [32]. For an undirected topology, \mathcal{L} is positive semi-definite. $\mathcal{L}\mathbf{1}_M = 0$, where $\mathbf{1}_M = [1, \dots, 1]^T$. If the graph \mathcal{G} has a spanning tree, the rank of \mathcal{L} should be $M - 1$ [32].

For an undirected graph \mathcal{G} , if $\hat{\chi}_m \in \mathbb{R}^n$ represents the updated training parameter of $\chi_m \in \mathbb{R}^n$ after a single

consensus step and χ_m stands for the row vector of the training parameter for agent m in the graph, the consensus training step of each agent m can be described as

$$\hat{\chi}_m = \chi_m + u_m \quad (6)$$

$$u_m = \sum_{k=1}^M a_{mk} (\chi_k - \chi_m) \quad (7)$$

where a_{mk} , the element of the graph adjacency matrix, is engendered by the undirected graph \mathcal{G} and u_m stands for the input of the agent m .

By integrating (7) and (6), the consensus algorithm can be used to update an agent m in the following scheme:

$$\begin{aligned} \hat{\chi}_m &= \chi_m + \sum_{k=1}^M a_{mk} (\chi_k - \chi_m) \\ &= \chi_m - \sum_{k=1}^M l_{mk} \chi_k \\ &= \mathcal{C}_m (\chi_k, l_{mk}) \end{aligned} \quad (8)$$

where l_{mk} is the element of the Laplacian matrix \mathcal{L} and \mathcal{C}_m represents the consensus protocol of the m^{th} agent. The training parameter update of all the M agents with a single consensus step can be summarised as

$$\begin{aligned} \hat{\chi} &= ((I_M - \mathcal{L}) \otimes I_n) \chi \\ &= \mathcal{C} (\chi, \mathcal{L}) \end{aligned} \quad (9)$$

where \mathcal{C} stands for the consensus protocol for all agents, I_M and I_n denote the $M \times M$ and $n \times n$ identity matrix, \mathcal{L} represents the Laplacian matrix. By repetitively computing (9), this consensus algorithm makes all agents converge to their weighted average [33].

D. Consensus-based Training with DRL

Given the consensus network structure in the previous subsection, the training algorithm for the training parameters ψ_t in (4) in the DRL is now introduced. As stated in [34], the process of updating ψ_t for the m^{th} agent is given as:

$$\psi_{t+1}^m = \psi_t^m - \alpha \xi_t^m \frac{dQ(\psi_t^m, s_t^m, a_t^m)}{d\psi_t^m} \quad (10)$$

where α represents the learning rate.

By applying (8), the training process of the CSAR algorithm can be summarised as:

$$\hat{\psi}_t^m = \psi_t^m + \sum_{k=1}^M \tilde{a}_{mk} (\psi_t^k - \psi_t^m) \quad (11)$$

$$\psi_{t+1}^m = \hat{\psi}_t^m - \alpha \xi_t^m d_t^m \quad (12)$$

where $d_t^m = \frac{dQ(\psi_t^m, s_t^m, a_t^m)}{d\psi_t^m}$.

Substituting (11) into (12), we can get

$$\psi_{t+1}^m = \psi_t^m + \sum_{k=1}^M \tilde{a}_{mk} (\psi_t^k - \psi_t^m) - \alpha \xi_t^m d_t^m \quad (13)$$

Let $\Gamma_t = [\xi_t^1 d_t^1, \xi_t^2 d_t^2, \dots, \xi_t^M d_t^M]^T$, for M agents, the update of the training parameters in our suction network in the t^{th} iteration can be illustrated as

$$\begin{aligned} \psi_{t+1} &= ((I_M - \mathcal{L}) \otimes I_n) \psi_t - \alpha \Gamma_t \\ &= \mathcal{C} (\psi_t, \mathcal{L}) - \alpha \Gamma_t \end{aligned} \quad (14)$$

Algorithm 1 summarizes our CSAR algorithm.

IV. EXPERIMENTS AND RESULTS

The feasibility of the CSAR algorithm is validated in this section. The system is implemented on a standard laptop with Nvidia GTX 2070 super and Intel Core i7 CPU (2.6 GHz) with 16 GB RAM. The experimental video is available at: <https://youtu.be/mcHJtNiStEQ>.

A. Experiment Setup

1) *Simulation*: Our system in the simulated environment is trained in Coppeliassim [35] with Bullet Physics 2.78 for dynamics, as demonstrated in Fig. 1. The simulation setup for each agent consists of a UR5 robot arm with a suction gripper [36]. The suctioned objects in the simulated environment are cubes with a side length of 5 cm. The motion planning task for each UR5 robot arm is accomplished by Coppeliassim [35] internal inverse kinematics. Simulated cameras are used to capture RGB-D images of each agent in a 0.448×0.448 m² workspace. The resolution of the simulated RGB-D images is 640×480 .

2) *Real World*: The setup for each agent in the real environment is composed of a UR5 robot arm with a Robotiq EPick vacuum gripper. The suctioned objects are cubes with a side length of 6.5 cm. To pick and place objects successfully with the suction gripper in the sim-and-real environment, the objects should have a flat surface and no overlap between objects placed in the workspace. We use a fixed Azure Kinect camera to acquire real-world RGB-D images with a resolution of 1280×720 . The location of the Azure Kinect camera is shown in Fig. 1, which can generate a top-down view in a 0.448×0.448 m² workspace.

3) *Reward*: Depending on the intrinsic and distortion of the Azure Kinect camera and the size of our suction gripper, we assign $r_0 = 2000$, $r_1 = 1000$, $r_2 = 100$, $r_3 = 1$ and $\mu_{th} = 0.005$ m in (2). These values can also be reconfigured for other robotic platforms.

4) *Neural Network*: The proposed framework is fully trained under self-supervision through the interactions between the UR5 robot arms and the sim-and-real environment. The learning rate α in (10) has a fixed value of 0.0001. The discounted factor γ listed in (4) is set to 0.5. The future reward discount is fixed at 0.5. The total training steps parameter T is initialized at 270. Algorithm 1 satisfies ϵ -greedy exploration strategy with ϵ initialized at 0.5 and annealed to 0.1 over training. The simulated camera and the Azure Kinect camera capture RGB-D images to generate colour and depth heightmaps, which are fed into the suction nets to predict pixel-wise best suction positions.

5) *Evaluation Metric*: The suction performance of the m^{th} agent can be evaluated using the suction success rate S_r^m , which is defined as follows:

$$S_r^m = \frac{N_s^m}{N_i^m} \times 100\% \quad (15)$$

where N_s^m represents the number of successful target suction of the m^{th} agent, N_i^m represents the number of iterations of the m^{th} agent.

We explore various training strategies to discover the most suitable training conditions for robots:

Sim-and-Real: Only simulation samples are used to train and optimise the model initially. When the suction success rate in the simulation reaches 0.5, we switch to the CSAR method with 3 **simulated** robots and 1 **real** robot.

Sim-to-Real: Only simulation samples are used to train and optimise the model at the beginning. When the suction success rate in the simulation reaches 0.5, we switch to real-world training with 1 real robot.

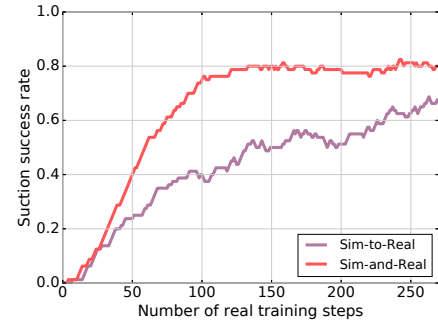


Fig. 3: Suction success rates of the real robot between “Sim-to-Real” and “Sim-and-Real” strategies




Fig. 4: Topology of the interaction of simulation and the real world: (a) 1 **simulated** robot and 1 **real** robot; (b) 2 **simulated** robots and 1 **real** robot; (c) 3 **simulated** robots and 1 **real** robot

B. Sim-and-Real is Better Than Sim-to-Real

Fig. 3 demonstrates the suction training success rate of the real robot using two different training strategies. The interaction topology of Sim-and-Real is shown in Fig. 4 (c). When applying the Sim-and-Real strategy, the suction success rate of the real robot reaches 80% at around 140 training steps,

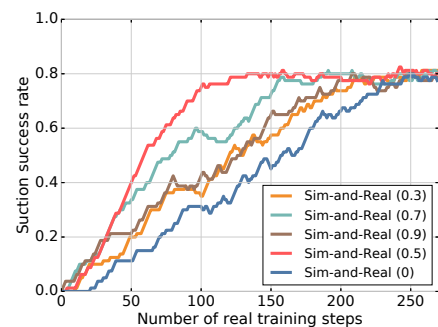


Fig. 5: Suction success rates of the real robot with different initial weights when applying the Sim-and-Real strategy. The number in brackets denotes the suction success rate from the pre-trained simulation model.

which outperforms the Sim-to-Real strategy. Since our policy for each robot is greedy deterministic, a robot may execute the same action repetitively if there is no environment change when using the Sim-to-Real training strategy. However, by applying consensus-based training, the simulated agent can be used to introduce noise indirectly into the sim-and-real environment, which prevents robots from getting stuck in the same action. In summary, applying the Sim-and-Real strategy leads to a faster training speed, which saves real-world training costs.

C. Best Policy in Simulation is Not the Best for Sim-and-Real Training

A striking observation from our experiment is that the best-obtained policy trained in simulation is not the best pre-trained model to start the co-training between simulated and real robots, as shown Fig. 5. When the suction success rate of the pre-trained simulation model is 0.5, the Sim-and-Real strategy achieves the best performance. When the suction success rate drops to 0.3, it takes longer for the real robot to solve the task. Surprisingly, when the suction success rate of the pre-trained simulation model is too high (0.7, 0.9), the performance deteriorates.

This is counterintuitive, as shown in the Sim-to-Real experiment, that the best policy obtained in the simulation is typically the one to be deployed. This observation suggests that the “mediocre” policy is the best for co-training. When the success rate of the pre-trained simulation model is too high, the sim-and-real framework will be initialised at a value that is close to the optimal simulation value. This will take longer to converge to the mixed optimality in a sim-and-real environment. As a result, applying the “mediocre” policy can reduce real robot training costs and save the pre-training time in the simulation.

D. The More Agents in Simulation, the Better for Sim-and-Real Training

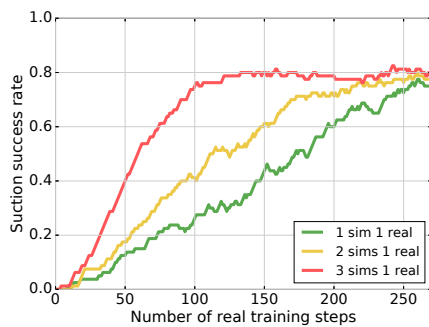


Fig. 6: Suction success rates of the **real** robot with different number of **simulated** robots using Sim-and-Real strategy

Readers may wonder why we use 3 **simulated** robots and 1 **real** robot during training. Therefore, we vary the number of simulated robots when using the Sim-and-Real strategy. Fig. 6 describes the suction success rate when using the Sim-and-Real strategy with different numbers of simulated robots. The interaction topology used in Fig. 6 is shown in Fig. 4.

It takes around 260 steps to make the real robot arrive at 80% suction success rate when using 1 **simulated** robot and 1 **real** robot strategy. In the case of 2 **simulated** robots 1 **real** robot, the required training steps descend to around 240. Only around 140 steps are required to maintain the same suction success rate when using the 3 **simulated** robots 1 **real** robot strategy. More simulated robots participating in the proposed framework can accelerate the training speed and exhibit good robustness in the sim-and-real environment, thus decreasing the number of required real robot training steps while maintaining a comparable suction success rate.

E. Generalisation of Real-world Unseen Objects

The Sim-and-Real strategy is capable of generalising to novel objects (Fig. 7) with a suction success rate of 80%. After training on cubes in both simulated and real environments, the CSAR training model can also be applied to pick and place novel objects such as cylinders and irregularly shaped objects with different heights, as shown in Fig. 8.




Fig. 7: Novel objects for validation: (a) Environment 1; (b) Environment 2; (c) Environment 3




Fig. 8: The demonstration of picking novel objects. More details can be seen in the video.

V. CONCLUSION

In this work, we propose a CSAR approach which is able to improve sim-and-real training speed and reduce real-world training costs. By implementing the Sim-and-Real strategy, the suction success rate of the real robot attains 80% at around 140 training steps, which outperforms the Sim-to-Real strategy. Applying the “mediocre” policy can not only reduce the number of required real robot training steps but also save the pre-training time in the simulation. More simulated robots participating in the CSAR method increase the training speed, thereby reducing real-world training expenses. The Sim-and-Real strategy is also capable of generalising to novel objects. The CSAR method is a straightforward generalization and practical verification of the team’s recently developed theory of a consensus-based RL approach [7]. In the future, an optimisation of the CSAR approach will be exploited to tackle more complicated scenarios.

REFERENCES

- [1] B. Beyret, A. Shafti, and A. A. Faisal, "Dot-to-dot: Explainable hierarchical reinforcement learning for robotic manipulation," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 5014–5019.
- [2] Y. Han, I. H. Zhan, W. Zhao, J. Pan, Z. Zhang, Y. Wang, and Y.-J. Liu, "Deep reinforcement learning for robot collision avoidance with self-state-attention and sensor fusion," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6886–6893, 2022.
- [3] R. Han, S. Chen, S. Wang, Z. Zhang, R. Gao, Q. Hao, and J. Pan, "Reinforcement learned distributed multi-robot navigation with reciprocal velocity obstacle shaped rewards," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 5896–5903, 2022.
- [4] R. S. Sutton, "Introduction: The challenge of reinforcement learning," in *Reinforcement Learning*. Springer, 1992, pp. 1–3.
- [5] K. Wu, J. Hu, B. Lennox, and F. Arvin, "Sdp-based robust formation-containment coordination of swarm robotic systems with input saturation," *Journal of Intelligent & Robotic Systems*, vol. 102, no. 1, pp. 1–16, 2021.
- [6] K. Wu, J. Hu, Z. Ding, and F. Arvin, "Finite-time fault-tolerant formation control for distributed multi-vehicle networks with bearing measurements," *IEEE Transactions on Automation Science and Engineering*, 2023.
- [7] W. Liu, H. Niu, I. Jang, G. Herrmann, and J. Carrasco, "Distributed neural networks training for robotic manipulation with consensus algorithm," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–15, 2022.
- [8] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 3803–3810.
- [9] R. Kaushik, K. Arndt, and V. Kyrki, "Safeapt: Safe simulation-to-real robot learning using diverse policies learned in simulation," *IEEE Robotics and Automation Letters*, 2022.
- [10] W. Chen, Y. Xu, Z. Chen, P. Zeng, R. Dang, R. Chen, and J. Xu, "Bidirectional sim-to-real transfer for gelsight tactile sensors with cyclegan," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6187–6194, 2022.
- [11] T. Bi, C. Sferrazza, and R. D'Andrea, "Zero-shot sim-to-real transfer of tactile control policies for aggressive swing-up manipulation," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5761–5768, 2021.
- [12] U. Viereck, A. Pas, K. Saenko, and R. Platt, "Learning a visuomotor controller for real world robotic grasping using simulated depth images," in *Conference on robot learning*. PMLR, 2017, pp. 291–300.
- [13] 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.
- [14] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, "Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 627–12 637.
- [15] K. Kang, S. Belkhale, G. Kahn, P. Abbeel, and S. Levine, "Generalization through simulation: Integrating simulated and real data into deep reinforcement learning for vision-based autonomous flight," in *2019 international conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 6008–6014.
- [16] E. Tzeng, C. Devin, J. Hoffman, C. Finn, X. Peng, S. Levine, K. Saenko, and T. Darrell, "Towards adapting deep visuomotor representations from simulated to real environments," *CoRR*, vol. abs/1511.07111, 2015.
- [17] Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox, "Closing the sim-to-real loop: Adapting simulation randomization with real world experience," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8973–8979.
- [18] S. Di-Castro Shashua, D. Di Castro, and S. Mannor, "Sim and real: Better together," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [19] T. G. Thuruthel, E. Falotico, F. Renda, and C. Laschi, "Model-based reinforcement learning for closed-loop dynamic control of soft robotic manipulators," *IEEE Transactions on Robotics*, vol. 35, no. 1, pp. 124–134, 2018.
- [20] S. K. Pradhan and B. Subudhi, "Real-time adaptive control of a flexible manipulator using reinforcement learning," *IEEE Transactions on Automation Science and Engineering*, vol. 9, no. 2, pp. 237–249, 2012.
- [21] J.-J. Park, J.-H. Kim, and J.-B. Song, "Path planning for a robot manipulator based on probabilistic roadmap and reinforcement learning," *International Journal of Control, Automation, and Systems*, vol. 5, no. 6, pp. 674–680, 2007.
- [22] J. Kober, E. Oztop, and J. Peters, "Reinforcement learning to adjust robot movements to new situations," in *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.
- [23] Y. Yang, H. Liang, and C. Choi, "A deep learning approach to grasping the invisible," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2232–2239, 2020.
- [24] T. Inoue, G. De Magistris, A. Munawar, T. Yokoya, and R. Tachibana, "Deep reinforcement learning for high precision assembly tasks," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 819–825.
- [25] J. Cai, H. Cheng, Z. Zhang, and J. Su, "Metagrasp: Data efficient grasping by affordance interpreter network," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 4960–4966.
- [26] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *Thirty-first AAAI conference on artificial intelligence*, 2017.
- [27] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in neural information processing systems*, vol. 32, 2019.
- [28] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. Funkhouser, "Learning synergies between pushing and grasping with self-supervised deep reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4238–4245.
- [29] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
- [30] S. Lambert-Lacroix and L. Zwald, "Robust regression through the huber's criterion and adaptive lasso penalty," *Electronic Journal of Statistics*, vol. 5, pp. 1015–1053, 2011.
- [31] K. Wu, J. Hu, B. Lennox, and F. Arvin, "Mixed controller design for multi-vehicle formation based on edge and bearing measurements," in *2022 European Control Conference (ECC)*. IEEE, 2022, pp. 1666–1671.
- [32] C. Godsil and G. F. Royle, *Algebraic graph theory*. Springer Science & Business Media, 2013, vol. 207.
- [33] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on automatic control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [34] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [35] E. Rohmer, S. P. Singh, and M. Freese, "V-rep: A versatile and scalable robot simulation framework," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 1321–1326.
- [36] J. Ge, H. Saeidi, M. Kam, J. Opfermann, and A. Krieger, "Supervised autonomous electrosurgery for soft tissue resection," in *2021 IEEE 21st International Conference on Bioinformatics and Bioengineering (BIBE)*. IEEE, 2021, pp. 1–7.