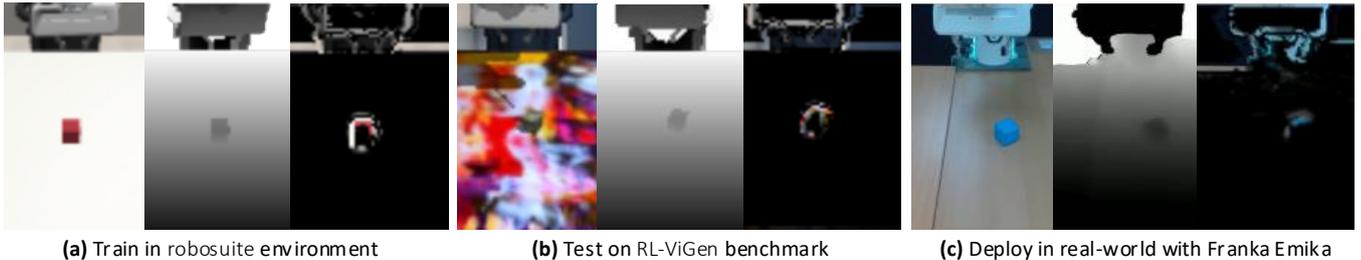


# DeGuV: Depth-Guided Visual Reinforcement Learning for Generalization and Interpretability in Manipulation

Tien Pham<sup>1</sup>, Xinyun Chi<sup>1</sup>, Khang Nguyen<sup>2</sup>, Manfred Huber<sup>2</sup>, and Angelo Cangelosi<sup>1</sup>



**Fig. 1:** (a) *DeGuV* is trained with robotic manipulation tasks in `robosuite` environment, (b) is evaluated on the RL-ViGen benchmark for generalization and interpretability verification, and (c) is deployed on a Franka Emika robot to demonstrate its zero-shot sim-to-real transferability. The RGB-D observations and the masked image produced by *DeGuV* are presented from left to right in each case.

**Abstract**—Reinforcement learning (RL) agents can learn to solve complex tasks from visual inputs, but generalizing these learned skills to new environments remains a major challenge in RL application, especially robotics. While data augmentation can improve generalization, it often compromises sample efficiency and training stability. This paper introduces *DeGuV*, an RL framework that enhances both generalization and sample efficiency. In specific, we leverage a learnable *masker network* that produces a mask from the depth input, preserving only critical visual information while discarding irrelevant pixels. Through this, we ensure that our RL agents focus on essential features, improving robustness under data augmentation. In addition, we incorporate contrastive learning and stabilize Q-value estimation under augmentation to further enhance sample efficiency and training stability. We evaluate our proposed method on the RL-ViGen benchmark using the Franka Emika robot and demonstrate its effectiveness in zero-shot sim-to-real transfer. Our results show that *DeGuV* outperforms state-of-the-art methods in both generalization and sample efficiency while also improving interpretability by highlighting the most relevant regions in the visual input. Our implementation is available at: <https://github.com/tiencapham/DeGuV>.

## I. INTRODUCTION

Vision-based reinforcement learning (RL) has demonstrated its impact in various applications, including video games [1]–[3], autonomous navigation [4]–[6], and robotics [7]–[9]. As RL agents often learn from simulation environments [10], they are precluded from generalizing unseen environments in which visual observations differ from those encountered during training. The distribution shift, caused by

variations in object textures, colors, and environmental radiance (Fig. 1), significantly degrades real-world performance. To better generalize these factors, many researchers have explored data augmentation techniques [11]–[18] to diversify training data. Despite enhancing robustness, these methods raise sample efficiency concerns [19], [20] and destabilize training processes [12] of visual RL models for robotics.

To address these concerns, we present *DeGuV*, an RL framework that improves generalization while maintaining sample efficiency. Our approach leverages depth input to mitigate the intrinsic variations caused by data augmentation. We ensure that RL agents focus only on the most informative regions by employing a depth-guided masking module to filter out task-irrelevant pixels in RGB images. Through this, we are able to reduce unnecessary variations in augmented training data, enhancing sample efficiency while minimizing the gap between training and evaluation distributions. To further reinforce visual perturbation invariance, we incorporate contrastive learning to facilitate the model to focus on consistent, task-relevant features across varying visual conditions. Lastly, we combine Stabilized Q-Value Estimation (SVEA) [12] to improve the training stability of *DeGuV*.

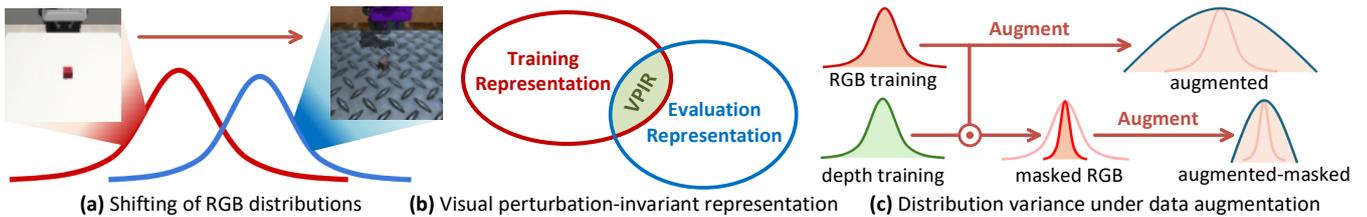
We evaluate the proposed method on four tasks and three evaluation modalities in the RL-ViGen benchmark [21]. Our results show that we achieve superior results compared to state-of-the-art methods in generalization while maintaining sample efficiency. Moreover, our model can enhance interpretability by visualizing attention through its learned masks, providing insight into the agent’s decision-making process. For real-robot experiments, we illustrate the feasibility of sim-to-real transfer in a zero-shot manner.

Our contributions are summarized as follows:

- We formulate the generalization problem inherent in visual RL under data augmentation, solving the predicament between the model’s generalization capability and learning complexity.

\*This work has been partially supported by Horizon Europe (and UKRI Horizon Guaranteed Fund) under the Marie Skłodowska-Curie grant agreement No 101072488 (TRAIL). This work was also in part supported by a project funded by the EPSRC Prosperity grant CRADLE (EP/X02489X/1) and by the Air Force Office of Scientific Research, USAF, under the CASPER++ Awards (FA8655-24-1-7047).

<sup>1</sup>Cognitive Robotics Lab, School of Computer Science, University of Manchester, UK; <sup>2</sup>Learning and Adaptive Robotics Lab, Department of Computer Science and Engineering, University of Texas at Arlington, USA; (corresponding email: canhantien.pham@manchester.ac.uk).



**Fig. 2:** Illustrations of (a) generalization problem in visual RL posed by distribution shifting, (b) visual perturbation-invariant space between training and evaluation representations, (c) changes of variances of RGB and depth distributions under data augmentation.

- We introduce an RL framework, leveraging depth input to reduce visual observation’s variation internally during training and evaluation, enhancing the trained policy’s generalization capability and sample efficiency.
- We conduct experiments on the generalization benchmark against state-of-the-art techniques. The results show improvements in generalization ability, sample efficiency, and interpretability.

## II. RELATED WORK

**Generalization in Reinforcement Learning:** Prior works have been investigated in enhancing generalization capabilities of RL using various techniques such as contrastive learning [22], [23], data augmentation [11]–[18], [24], domain adaptation [25], and incorporation of pre-trained image encoders [26]. While data augmentation and other techniques have shown effectiveness in improving the generalization ability of RL agents, they usually introduce additional complexity in optimal policy learning [12]. Chen *et al.* [27] introduced the Focus-then-Decide framework with SAM [28] to assist the agent in training efficiently under data augmentation. However, SAM requires an additional computing workload for both the training and inference processes, which worsens the computational load of RL models, especially for robotic applications. Yuan *et al.* [15] uses multi-view representation learning to endow the generalization ability of agents under various types of data augmentation. In this work, we aim to stabilize and improve RL agents’ performance by intrinsically reducing the variation in the augmented observation distribution.

**Interpretability by Visual Masking:** A few studies focus on enhancing RL agents’ generalization ability and interpretability by selectively masking input portions. Yu *et al.* [29] randomly obscure portions of the input and apply an auxiliary loss to recover the missing pixels. Bertoin *et al.* [14] proposes a Saliency-Guided Q-Network (SGQN) that self-supervisedly learns to generate the mask based on the gradient obtained during the training. SGQN is sensitive to hyperparameter `sgqn-quantile` (often set to 95%-98%), which determines how many pixels are masked. Tomar *et al.* [30] introduce InfoGating with U-Net to create the mask from the downstream loss. Grooten *et al.* [11] present MaDi, which generates the soft mask solely from the reward signal of the environment. However, the mask generated by InfoGating and MaDi depends on the perturbed observations. Furthermore, all the mentioned works use augmented visual observations to produce the mask, while we use depth input – a more stable distribution under augmentation.

## RGB-D Fusion for Visual Reinforcement Learning:

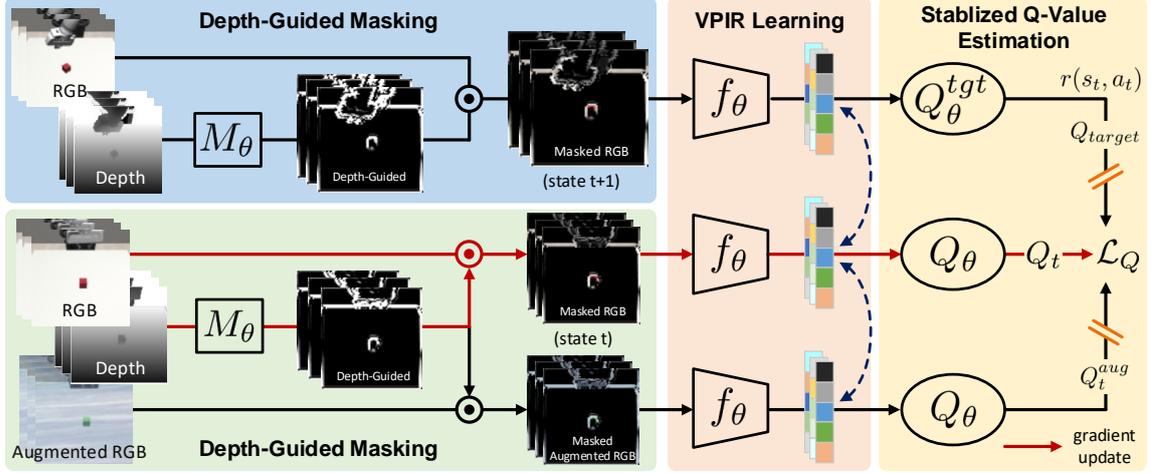
While most image-based RL studies focus on visual observations, some works incorporate visual and depth inputs to enrich environmental information and improve the performance of RL agents. The two data streams could be encoded separately and then fused by simple concatenation [31]–[33] or by a cross-attention mechanism [34]. Balakrishnan *et al.* [31] uses a pre-trained Twin Variational Autoencoder to extract the environment’s latent embedding simultaneously, which is used to train RL policies. James *et al.* [34] introduces a Q-Attention Agent that is used to extract cross-attention features, combined with a Next-Best Pose Agent to predict the next-best poses from RGB and point cloud inputs. Again, within this work, we only use depth images to produce the masks, which helps to reduce the variation in perturbed visual observations. The primary semantic information is still extracted from RGB observations.

## III. PROBLEM FORMULATION

Learning of generalizable policies in a Markov Decision Process (MDP) is formulated as an invariant representation learning problem, where learned policies could maintain their performance in unseen environments shifted from the visual training distribution. The interaction between environment and policy is constructed as an MDP:  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma \rangle$ , with  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space,  $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$  is the transition function,  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbf{R}$  is the reward function, and  $\gamma$  is the discount factor.

To overcome partial observability problems of static visual input, we wrap a state  $s_t$  to include  $k$  consecutive frames  $\{\mathbf{o}_t, \mathbf{o}_{t+1}, \dots, \mathbf{o}_{t+k}\}$ ,  $\mathbf{o}_i \in \mathcal{O}$ , where  $\mathcal{O}$  is the high-dimensional observation space. We let the state space  $\mathcal{S}$  is the one constructed as  $(P^{\text{RGB}} \cup P^{\text{D}})^k$  with  $P^{\text{RGB}}$  and  $P^{\text{D}}$  are RGB and depth image pixel space, respectively, from the observation space  $\mathcal{O}$ .

Our objective is to learn a policy  $\pi : \mathcal{S} \rightarrow \mathcal{A}$  that maximizes the discounted return  $\mathcal{R}_t = \mathbf{E}_{\Gamma \sim \pi} \left[ \sum_{t=0}^T \gamma^t r(s_t, a_t) \right]$  along a trajectory  $\Gamma = (s_0, s_1, \dots, s_T)$ , which are sampled from the observation space following the policy  $\pi$  from the initial state  $s_0$ . The policy  $\pi$  is parameterized by a collection of learnable parameters  $\theta$ , denoted as  $\pi_\theta$ . We aim to learn the parameters  $\theta$  so that  $\pi_\theta$  generalizes well across unseen visual MDPs, denoted as  $\overline{\mathcal{M}} = \langle \overline{\mathcal{S}}, \mathcal{A}, \mathcal{P}, r, \gamma \rangle$ , where the states  $\overline{s}_t \in \overline{\mathcal{S}}$  are constructed from the observations  $\{\overline{\mathbf{o}}_t, \overline{\mathbf{o}}_{t+1}, \dots, \overline{\mathbf{o}}_{t+k}\} \in \overline{\mathcal{O}}$ . The perturbed observation space  $\overline{\mathcal{O}}$  has a visual distribution shift, such as changes in color and brightness, from the original observation space  $\mathcal{O}$ .



**Fig. 3:** The training pipeline of *DeGuV* consists three main components: Depth-Guided Masking module takes in depth inputs to create the mask that erases non-relevant pixels in RGB image, VPIR Learning module employs contrastive loss in Eq. 6 to facilitate the RL agent to extract invariant representations, and Stabilized Q-Value Estimation is applied to stabilize the optimal policy training by avoiding non-stationary updates using `stop_grad` operators (`//`). Details of this training are provided in Alg. 1.

#### IV. METHODOLOGY

To learn the optimal policy with the continuous action space in robotic manipulation, we leverage the deep deterministic gradient policy with random exploration noise (DrQv2) [35] algorithm to approximate the optimal state-action value function  $Q^*$  using a parameterized critic  $Q_\theta$ . This is achieved by minimizing the Bellman residual [36]:

$$R = (r(s_t, a_t) + \gamma \max(Q_\theta^{\text{tgt}}(s_{t+1}, a'_t)) - Q_\theta(s_t, a_t)) \quad (1)$$

The actor, defined by a stochastic policy,  $\phi_\theta$ , optimizes the output of the critic network while ensuring high entropy. SAC often incorporates an optional shared encoder  $f_\theta$  for image-based environments. The critic and encoder use target networks initialized with the same parameters  $\theta_{\text{tgt}} = \theta$ , updated via an exponential moving average  $\theta_{\text{tgt}} \leftarrow (1 - \tau)\theta_{\text{tgt}} + \tau\theta$ . We formulate the optimal policy learning problem as a two-staged problem:

- 1) *Latent Representation Learning:* The agent learns to encode the latent representation  $z_t = f_\theta(s_t)$  of size  $N \ll \dim(\mathcal{S})$  where  $f_\theta : \mathcal{S} \rightarrow \mathcal{Z}$  is an encoder parameterized by  $\theta$  and  $\mathcal{Z}$  represents the latent space.
- 2) *Decision Making:* The agent learns to choose the optimal action from the latent representation  $a_t \sim \phi_\theta(a_t | z_t)$  to maximize the discounted return.

##### A. Data Augmentation

As the well-explored approach in prior works [11]–[15], data augmentation techniques, where observations are augmented to create a broader visual distribution during network training (Fig. 2c), can prevent the agent from overfitting to the training environment and improve generalization ability. Under data augmentation, the model is incentivized to learn task-relevant semantic information that could generalize well in visual perturbations, ensuring the trained policy’s generalization in various unseen scenarios (Fig. 2a and Fig. 2b). This leads to the Visual Perturbation-Invariant Representation’s definition (VPIR), formulated as follows:

**Definition.** Given an MDP  $\mathcal{M}$  that has a state  $s \in \mathcal{S}$  and its perturbed counterpart  $\bar{s} \in \bar{\mathcal{S}}$ , there is a set of latent representations  $z^*$  that could be encoded by the optimal policy to represent the task-relevant semantic information from both  $s$  and  $\bar{s}$ . Mathematically:

$$\exists z^* \in \mathcal{Z}, z^* = f_\theta(s) = f_\theta(\bar{s}), s \in \mathcal{S}, \bar{s} \in \bar{\mathcal{S}} \quad (2)$$

However, the data augmentation phases inevitably increase the state space size, creating additional complexity in the training process and reducing the sample efficiency. We observe that the augmented distribution’s variance is larger than the original variance:  $\text{Var}(P_{\text{aug}}^{\text{RGB}}) \gg \text{Var}(P^{\text{RGB}})$ .

##### B. Distractions

We denote  $p_{\text{distraction}}^{\text{RGB}}$  as the unrelated input pixels, the so-called visual distraction, to the task, which could be safely excluded during optimal policy learning. Meanwhile, the remaining pixels are called task-related pixels  $P_{\text{relevant}}^{\text{RGB}} = P_{\text{aug}}^{\text{RGB}} \setminus P_{\text{distraction}}^{\text{RGB}}$ . Therefore, VPIR is extracted from only the task-related pixels, we derive the relationship between their distribution variance as follows:

$$\text{Var}(P_{\text{aug}}^{\text{RGB}}) = \text{Var}(P_{\text{relevant}}^{\text{RGB}}) + \text{Var}(P_{\text{distraction}}^{\text{RGB}}) \quad (3)$$

When the distractions are properly removed, the variance created on those pixels will vanish in Eq. 3, reducing the complexity of optimal policy learning. We divide the *latent representation learning* into *visual distraction masking*, where the agent learns to erase distractions in visual states, and *learning latent representation from masked observations*, where the agent is trained to encode the latent representation from the masked state.

The visual distraction masking problem complements the MDP, identifying the irrelevant pixels to the task and Q-value estimation in the Bellman residual. Since the optimal  $Q^*$  remains unchanged by distractions, minimizing the Bellman residual (Eq. 1) simultaneously accomplishes learning of an optimal policy and masking visual distractions. This

---

**Algorithm 1: DeGuV Training**

---

**Input :**  $f_\theta, Q_\theta, \phi_\theta, M_\theta :=$  initialized networks  
 $Q_\theta^{tgt}, f_\theta^{tgt} \leftarrow Q_\theta, f_\theta :=$  target networks  
 $T :=$  total timesteps

**Output:**  $f_\theta, Q_\theta, \phi_\theta, M_\theta :=$  trained networks

```
1 function DeGuV_train( $f_\theta, Q_\theta, \phi_\theta, M_\theta, T$ )
2   for  $t \in [1, \dots, T]$  do
3      $\triangleright$  ----- act phase -----
4      $s_t^{\text{masked}} \leftarrow \text{DGMask}(s_t)$ 
5      $a_t \sim \phi_\theta(\cdot | f_\theta(s_t^{\text{masked}}))$ 
6      $s_{t+1}, r_t \sim \mathcal{P}(\cdot | s_t, a_t)$ 
7      $\mathcal{B} \leftarrow \mathcal{B} \cup (s_t, a_t, r_t, s_{t+1})$ 
8      $\triangleright$  ----- update phase -----
9      $\{s_b, a_b, r_b, s'_b\} \sim \mathcal{B}$ 
10     $s_b, s'_b \leftarrow \text{DGMask}(s_b), \text{DGMask}(s'_b)$ 
11     $s_b, s'_b \leftarrow [s_b, \tau(s_b, \nu)], [s'_b, \tau(s'_b, \nu')]$ 
12     $\triangleright$  ----- update actor -----
13     $\theta_\phi \leftarrow \theta_\phi - \nabla_{\theta_\phi} \mathcal{L}_\phi(s_b)$ 
14     $\triangleright$  ----- update critics -----
15    for  $N \in \{Q_\theta, f_\theta, M_\theta\}$  do
16       $\theta_N \leftarrow \theta_N - \nabla_{\theta_N} \mathcal{L}_Q(s_b, a_b, r_b, s'_b)$ 
17    for  $N \in \{Q_\theta\}$  do
18       $\theta_N^{tgt} \leftarrow (1 - \tau)\theta_N^{tgt} + \tau\theta_N$ 
19     $\triangleright$  ----- update auxiliary -----
20     $\theta_f \leftarrow \theta_f - \nabla_{\theta_f} \mathcal{L}_{\text{InfoNCE}}$ 
21  return  $f_\theta, Q_\theta, \phi_\theta, M_\theta$ 
```

---

means the agent is incentivized to ignore distractions and only focus on relevant pixels to achieve an optimal policy. However, learning to generate the mask from highly varied augmented visual distribution  $P_{\text{aug}}^{\text{RGB}}$  is more complex than learning from original visual distribution, which leads to unstable optimization. Meanwhile, the depth distribution is invariant to visual perturbations and remains stable under data augmentation (Fig. 2c), thus could serve as a reliable source for distraction masking procedure.

### C. Depth-Guided Masking

To reduce the variance of observation distribution caused by the data augmentation, we introduce a depth-guided masker  $M_\theta$  that leverages the depth input from the state to zero out the irrelevant RGB pixels, as stated in the following Eq. 4. The masker produces a scalar matrix in the  $[0, 1]$  range to determine its relevance to the task. As depth input should be invariant to data augmentation, the mask could be applied element-wise for both the original visual input and its augmented counterpart, leaving minimal variance in visual distribution needed for generalization. The depth-guided masking,  $\text{DGMask}$ , is formulated as follows:

$$(s_t^{\text{RGB}}, s_t^D) \leftarrow s_t \quad (4a)$$

$$s_t^{\text{masked}} \leftarrow s_t^{\text{RGB}} \odot M_\theta(s_t^D) \quad (4b)$$

The depth-guided masker  $M_\theta$  is a learnable network composed of convolutional layers with the ReLU activation function in between. The last layer is followed by a  $\text{Hardtanh}(\cdot)$  activation function to regularize its value in the range of  $[0, 1]$ . Our mask can eliminate irrelevant pixels using  $\text{Hardtanh}$  rather than merely reducing their values close to zero. As the masker is fully differentiable, it could

be trained during the backward step with the critic update. Additionally, the mask generated by  $M_\theta$  also enhances the interpretability of the model by highlighting the regions of the visual input that are most relevant for task completion.

### D. Visual Perturbation-Invariant Representation Learning

Contrastive learning is used to train the encoder, which extracts VPIR in a self-supervised manner. In this approach, given a query  $\mathbf{q}$ , the goal is to enhance the similarity between  $\mathbf{q}$  and its corresponding positive key  $\mathbf{k}^+$ , while simultaneously reducing the similarity between  $\mathbf{q}$  and each negative key  $\mathbf{k}^-$  in a training batch. We quantify the disparities between vectors  $\mathbf{q}$  and  $\mathbf{k}$  using cosine similarity:

$$\text{sim}(\mathbf{q}, \mathbf{k}) = \frac{\mathbf{q}^T \mathbf{k}}{\|\mathbf{q}\| \|\mathbf{k}\|} \quad (5)$$

The InfoNCE loss [37],  $\mathcal{L}_{\text{InfoNCE}}$ , is used to penalize the model for learning VPIR in contrastive learning with  $\tau$  as the temperature hyperparameter by the following expression:

$$-\log \left[ \frac{\exp\left(\frac{\text{sim}(\mathbf{q}^T, \mathbf{k}^+)}{\tau}\right)}{\exp\left(\frac{\text{sim}(\mathbf{q}^T, \mathbf{k}^+)}{\tau}\right) + \sum_{i=0}^M \exp\left(\frac{\text{sim}(\mathbf{q}^T, \mathbf{k}_i^-)}{\tau}\right)} \right], \quad (6)$$

Given a masked visual state  $s_t^{\text{masked}}$  and its augmented counterpart  $s_t^{\text{masked+aug}}$ , an encoder  $f_\theta$  needs to learn to extract the states into a VPIR  $z^*$  so that it can be generalized into the visual perturbed state. Therefore, VPIR can be given by  $z^* = f_\theta(s_t^{\text{masked}}) = f_\theta(s_t^{\text{masked+aug}})$ . We use the auxiliary InfoNCE loss defined in Eq. 6 with the query key  $\mathbf{q} = f_\theta(s_t^{\text{masked}})$ . The positive key is the representation extracted from the augmented-masked state  $\mathbf{k}^+ = f_\theta(s_t^{\text{masked+aug}})$ , and the negative key is extracted from other states at  $t + 1$  in the batch of samples. During the auxiliary update, the masker remains constant, and only the shared encoder  $f_\theta$  is updated. This strategy mitigates the occurrence of non-stationary gradients originating from varied augmented RGB and stable depth inputs, thereby stabilizing the optimization procedure. Specifically, the masker receives updates exclusively from the environment's reward signal. VPIR is derived from a relatively stable masked visual distribution, resulting in improved sample efficiency and enhanced agent generalization capabilities.

### E. Learning Objectives

The critic loss,  $\mathcal{L}_Q$ , is defined as a combination of Bellman residuals, as indicated in Eq. 1, from  $s_t^{\text{masked}}$  and its augmented counterpart  $s_t^{\text{masked+aug}}$ , as follows:

$$\mathcal{L}_Q(s_t, a_t, r_t, s_{t+1}) = \alpha R^2 + \beta R_{\text{aug}}^2, \quad (7)$$

where  $\alpha$  and  $\beta$  are the coefficients to balance the ratio of the two data streams.

The masker  $M_\theta$  is updated along with the policy network by optimizing  $\mathcal{L}_Q$  in Eq. 7. In addition, a `stop_grad` operation is applied after  $Q^{\text{aug}}$  estimation to avoid the impact of non-stationary gradients on Q-value estimation resulting from augmented observations. The target network  $Q_\theta^{\text{tgt}}$  is

**TABLE I:** Comparisons of episode returns ( $\mu, \sigma$ ) of *DeGuV* against other baselines on RL-Vigen benchmark [21] in four training and evaluation environments: train, easy, medium, and hard of four tasks Lift, Door, NutAssemblyRound, and TwoArmPegInHole.

Task	Baseline	DrQv2 [35]	CURL [22]	SGQN [14]	SVEA [12]	MaDi [11]	DeGuV (ours)
Lift	train	242.21 ± 187.82	125.65 ± 84.10	142.86 ± 43.10	234.25 ± 15.83	364.87 ± 34.66	<b>443.85 ± 34.26</b>
	easy	17.14 ± 70.08	74.16 ± 71.06	57.16 ± 35.00	118.84 ± 90.42	130.47 ± 131.68	<b>293.80 ± 160.07</b>
	medium	0.11 ± 0.11	9.70 ± 15.91	39.65 ± 20.92	12.74 ± 23.96	10.79 ± 17.23	<b>324.67 ± 110.79</b>
	hard	0.16 ± 0.47	13.73 ± 17.14	23.92 ± 15.86	12.24 ± 25.50	23.35 ± 20.56	<b>227.32 ± 134.84</b>
Door	train	387.35 ± 169.50	438.12 ± 74.14	441.61 ± 137.58	<b>475.82 ± 27.99</b>	444.41 ± 121.52	467.23 ± 70.82
	easy	365.66 ± 186.72	436.39 ± 78.44	440.87 ± 139.02	464.73 ± 60.45	405.72 ± 145.43	<b>476.47 ± 47.23</b>
	medium	160.44 ± 215.12	190.40 ± 178.89	236.71 ± 233.97	247.97 ± 217.57	204.19 ± 198.21	<b>333.32 ± 203.53</b>
	hard	113.39 ± 182.30	135.65 ± 184.39	126.73 ± 204.60	256.17 ± 210.97	135.65 ± 184.39	<b>447.24 ± 103.55</b>
Nut Assembly Round	train	74.72 ± 60.64	61.91 ± 54.22	106.10 ± 67.89	135.01 ± 63.30	<b>138.55 ± 59.55</b>	137.93 ± 56.53
	easy	50.63 ± 52.81	61.49 ± 55.95	67.23 ± 68.29	<b>130.98 ± 67.71</b>	127.31 ± 63.08	113.34 ± 67.50
	medium	6.35 ± 12.28	38.91 ± 39.01	14.93 ± 31.83	4.30 ± 7.32	0.71 ± 1.67	<b>105.82 ± 62.52</b>
	hard	13.62 ± 27.13	22.43 ± 30.84	7.10 ± 13.87	8.08 ± 14.17	0.6 ± 1.22	<b>104.50 ± 66.99</b>
TwoArm PegInHole	train	288.37 ± 20.61	314.99 ± 15.23	286.20 ± 27.59	<b>422.62 ± 34.07</b>	355.65 ± 5.04	381.07 ± 53.02
	easy	284.53 ± 20.85	310.93 ± 20.63	365.97 ± 72.08	365.97 ± 72.08	355.78 ± 5.55	<b>376.29 ± 47.29</b>
	medium	190.75 ± 21.97	196.27 ± 28.92	211.96 ± 40.97	126.96 ± 12.61	280.23 ± 18.42	<b>378.18 ± 43.28</b>
	hard	177.19 ± 35.54	186.40 ± 30.17	202.85 ± 45.38	123.12 ± 17.43	225.34 ± 41.66	<b>372.01 ± 42.27</b>
<b>Average</b>		148.29 ± 169.81	163.57 ± 160.04	168.05 ± 168.19	196.24 ± 180.74	200.23 ± 174.35	<b>311.44 ± 159.59</b>

updated using the exponential moving average, which is also isolated from  $\mathcal{L}_Q$  by a `stop_grad` operation. It is noted that *DeGuV* not only minimizes the distributional variance induced by data augmentation but also enhances VPIR learning through contrastive learning and stabilizes the optimization of  $\mathcal{L}_Q$ , thereby augmenting the agent’s generalization capabilities and improving sampling efficiency.

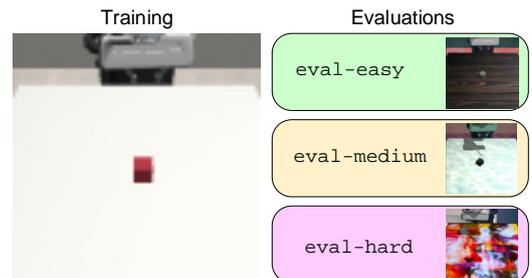
## V. EXPERIMENTS & EVALUATIONS

We evaluate our proposed approach in terms of generalization and data efficiency against three state-of-the-art RL algorithms focused on generalization: SGQN [14], SVEA [12], and MaDi [11]. As *DeGuV* is based on DrQv2 [35] and uses contrastive learning similar to CURL [22], we also include these two algorithms for comparisons to see their respective improvements.

### A. Experiment Setup

We evaluate our proposed approach in the `robosuite` simulation [38] using a virtual Franka Emika robot within the RL-Vigen benchmark [21]. The experiments are conducted on four distinct tasks: Lift, Door, NutAssemblyRound, and TwoArmPegInHole. We use `random_shift` [20], `random_overlay` [17] and `random_color_jitter` [39] augmentations: `random_shift` regularizes the encoder to prioritize salient features and enhance the agent’s data efficiency, `random_overlay` interpolates between the observed image and a randomly selected image for randomizing distributions containing distractions, and `random_color_jitter` creates the shifts of the contrast and hue color of observations.

We categorize the training and evaluation environments with difficulty level increases (`easy`, `medium`, `hard`). We train our models in a standard environment (`train`)



**Fig. 4:** Training and evaluation settings (Lift task). `eval_easy` includes changes in the background and object appearances, while `eval_medium` introduces moving light and arm color changes. `eval_hard` adds a more complex texture and video overlay.

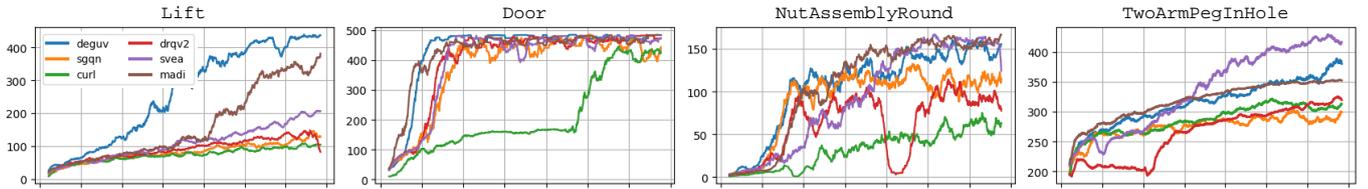
and evaluate them in evaluation environments with diverse visual perturbations, including texture, color, and brightness alterations. Note that each model is trained on 1,000,000 frame steps. For every 10,000 frame steps, we re-evaluate the training policy for 10 episodes with three random seeds. **Note:** All results are evaluated under out-of-distribution settings, as well as variation in task difficulty. As a result, the reported standard deviations are large, since the model typically either succeeds or fails, leading to substantial performance variability

**TABLE II:** Quantitative results on performance retention of *DeGuV* compared to other baselines over different evaluation modes.

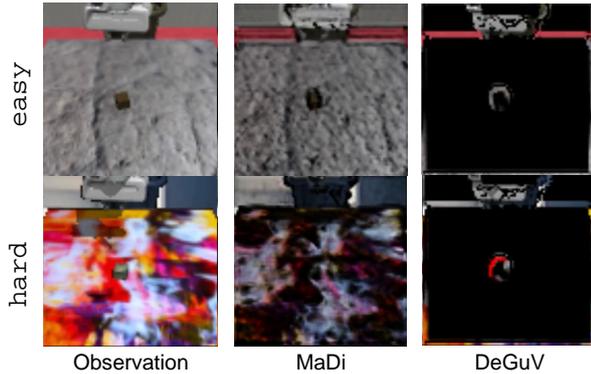
Mode	DrQv2	CURL	SGQN	SVEA	MaDi	DeGuV
easy	0.670	<b>0.892</b>	0.755	0.830	0.797	0.873
medium	0.290	0.441	0.424	0.227	0.321	<b>0.801</b>
hard	0.273	0.343	0.308	0.236	0.252	<b>0.801</b>
<b>Average</b>	0.411	0.559	0.495	0.431	0.457	<b>0.825</b>

### B. Generalization

Table I reports that *DeGuV* showcases its generalization across all tasks and evaluation modes, achieving an average



**Fig. 5:** Comparisons of episode returns of *DeGuV* compared to other baselines in four robotic tasks: *Lift*, *Door*, *NutAssemblyRound*, and *TwoArmPegInHole*, showing the our proposed method achieves better or similar data efficiency through the training process.



**Fig. 6:** Qualitative results of masked observations between *MaDi* and *DeGuV* from *easy* and *hard* of the *Lift* task.

episode return of 311.44, outperforming *MaDi* (200.23), *SVEA* (196.24), *SGQN* (168.05), *CURL* (163.57), and *DrQv2* (148.29) in this scheme. Moreover, we find that *DeGuV* retains 82.5% of its training performance on average, far surpassing the baselines, as shown in Table II. From these two experiments, we conclude that the other algorithms are not able to adapt to the unseen visual distribution, while *DeGuV* well manages this, thanks to depth-guided masking. Furthermore, our approach exhibits a significantly smaller performance drop in Table II, as the evaluation difficulty is more complex with visual perturbations, which highlights *DeGuV*’s ability to address the generalization problem. Fig. 5 additionally shows that *DeGuV* achieves superior data efficiency on the *Lift* task and is competitive compared to the baselines on *Door*, *NutAssemblyRound*, and *TwoArmPegInHole* tasks.

### C. Interpretability

Given that the mask serves as a tool for interpreting the agent’s decision-making rationale in response to an observation, we conduct analyses of the masks generated by *DeGuV* and *MaDi* in evaluation modalities, as shown in Fig. 6. Compared to *MaDi*, *DeGuV* generates masks that remove most visual distractions, leaving approximately identical masked observations in the two evaluation modes. In this example, *DeGuV* reveals merely 16.08% of the visual observations in both modes, whereas *MaDi* reveals 76.24% in *easy* and 53.77% in *hard* mode. This enhances generalization capabilities in intricate visual environments and ensures performance retention.

### D. Ablation Study

In this section, we conduct the experiments to verify the effectiveness of each component in *DeGuV*. We iteratively

remove the Depth-Guided Masking (Sec. IV-C), the Invariant Representation Learning (Sec. IV-D) and Stabilized-Q Learning (Sec. IV-E).

**TABLE III:** Experimental results showcasing the generalization performance of various ablations.

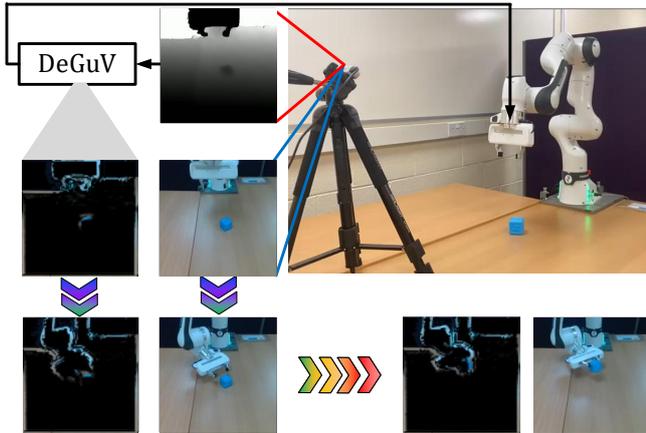
Ablation	Mode	train	easy	medium	hard
W/o Masking		<b>361.12</b> ± 142.14	37.32 ± 25.81	11.44 ± 8.32	12.21 ± 9.21
W/o VPIR Learning		287.15 ± 104.36	159.67 ± 77.03	148.66 ± 71.18	142.10 ± 73.59
W/o Stabilized-Q		291.44 ± 119.32	168.19 ± 87.24	144.91 ± 85.23	151.78 ± 82.11
<b>DeGuV</b>		357.52 ± 141.90	<b>314.98</b> ± 162.41	<b>285.50</b> ± 161.12	<b>287.77</b> ± 161.91

As shown in Table III, removing either Visual Perturbation-Invariant Representation Learning or Stabilized Q-learning leads to degraded training performance and reduced generalization capability. This degradation highlights the role of these two modules in enhancing training stability and learning compact representations that generalize well. In contrast, the training performance remains relatively stable without Depth-Guided Masking; however, the model suffers a substantial drop in out-of-distribution generalization performance. This result confirms the effectiveness of the masking module in extracting visual perturbation-invariant representations.

## VI. SIM-TO-REAL TRANSFERABILITY

We deploy our trained model from *robosuite* to the real-world scenario to verify its sim-to-real transferability. The Franka Emika robot is set up with the *Lift* task. We wrap *DeGuV* as a ROS2 package to interact with the robot via *franka\_ros2*. The ROS2 node processes images and controls the end-effector’s delta positions and orientations, along with the gripper control. We then use *panda-py* [40] to perform desired robotic actions. An Intel D435i RealSense RGB-D camera is attached to acquire RGB-D observations. Inherent noises in depth frames are reduced by RealSense ROS spatial, temporal and hole-filling post-processing filters.

Fig. 7 shows that the masked observation produced by *DeGuV* is similar to the masked observation in simulation, which only “highlight” important objects and locations that are related to the task, such as the robot gripper, hand, and cube. This real-world experiment confirms the effectiveness of *DeGuV* in generalization and interpretability. It also demonstrates the feasibility of sim-to-real transferability in



**Fig. 7:** Experiment set-up and results of original and masked observations for the Lift task on a real-robot system.

a zero-shot manner. Our demonstration video is available at: <https://youtu.be/-Gt5i6Wi5Fs>.

## VII. CONCLUSIONS AND FUTURE WORKS

In this paper, we propose *DeGuV*, targeting the generalization problem in visual reinforcement learning by leveraging the stable depth distribution to erase irrelevant pixels in the visual state. Our experiments show that the proposed method outperforms state-of-the-art algorithms on the RL-Vigen benchmark in generalizability. Our proposed method shows comparable sample efficiency during training and interpretability with masked observations compared to baselines. Finally, we demonstrate sim-to-real transferability, strengthening the generalization performance and potential for real robot deployment.

While depth input is invariant to visual perturbation, it usually contains noise, which is also worth investigating to improve *DeGuV*'s performance in sim-to-real transfer. Furthermore, long-horizon complex tasks remain challenging for our model in visual RL models. In the future, we will conduct more real robot experiments to validate the performance of sim-to-real transferability and investigate its potential in more challenging and long-horizon tasks.

## REFERENCES

- [1] S. Kapturowski, V. Campos, R. Jiang, N. Rakićević, H. van Hasselt, C. Blundell, and A. P. Badia, "Human-level atari 200x faster," *arXiv preprint arXiv:2209.07550*, 2022.
- [2] D. Schmidt and T. Schmied, "Fast and data-efficient training of rainbow: an experimental study on atari," *arXiv preprint arXiv:2111.10247*, 2021.
- [3] A. P. Badia, B. Piot, S. Kapturowski, P. Sprechmann, A. Vitvitskiy, Z. D. Guo, and C. Blundell, "Agent57: Outperforming the atari human benchmark," in *International conference on machine learning*. PMLR, 2020, pp. 507–517.
- [4] P. Mirowski, R. Pascanu, F. Viola, H. Soyer, A. J. Ballard, A. Banino, M. Denil, R. Goroshin, L. Sifre, K. Kavukcuoglu *et al.*, "Learning to navigate in complex environments," *arXiv preprint arXiv:1611.03673*, 2016.
- [5] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 3357–3364.
- [6] Z. Zhang, A. Liniger, D. Dai, F. Yu, and L. Van Gool, "End-to-end urban driving by imitating a reinforcement learning coach," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 15 222–15 232.
- [7] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas *et al.*, "Solving rubik's cube with a robot hand," *arXiv preprint arXiv:1910.07113*, 2019.
- [8] T. Haarnoja, B. Moran, G. Lever, S. H. Huang, D. Tirumala, J. Humpalik, M. Wulfmeier, S. Tunyasuvunakool, N. Y. Siegel, R. Hafner *et al.*, "Learning agile soccer skills for a bipedal robot with deep reinforcement learning," *Science Robotics*, vol. 9, no. 89, p. eadi8022, 2024.
- [9] R. Jangir, N. Hansen, S. Ghosal, M. Jain, and X. Wang, "Look closer: Bridging egocentric and third-person views with transformers for robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3046–3053, 2022.
- [10] C. Zhang, O. Vinyals, R. Munos, and S. Bengio, "A study on overfitting in deep reinforcement learning," *arXiv preprint arXiv:1804.06893*, 2018.
- [11] B. Grooten, T. Tomilin, G. Vasan, M. E. Taylor, A. R. Mahmood, M. Fang, M. Pechenizkiy, and D. C. Mocanu, "MaDi: Learning to Mask Distractions for Generalization in Visual Deep Reinforcement Learning," *The 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2024, uRL: <https://arxiv.org/abs/2312.15339>.
- [12] N. Hansen, H. Su, and X. Wang, "Stabilizing deep q-learning with convnets and vision transformers under data augmentation," in *Conference on Neural Information Processing Systems*, 2021.
- [13] J. Ha, K. Kim, and Y. Kim, "Dream to generalize: Zero-shot model-based reinforcement learning for unseen visual distractions," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 6, 2023, pp. 7802–7810.
- [14] D. Bertoin, A. Zouitine, M. Zouitine, and E. Rachelson, "Look where you look! saliency-guided q-networks for generalization in visual reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 30 693–30 706, 2022.
- [15] Z. Yuan, T. Wei, S. Cheng, G. Zhang, Y. Chen, and H. Xu, "Learning to manipulate anywhere: A visual generalizable framework for reinforcement learning," *arXiv preprint arXiv:2407.15815*, 2024.
- [16] R. Raileanu, M. Goldstein, D. Yarats, I. Kostrikov, and R. Fergus, "Automatic data augmentation for generalization in reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 5402–5415, 2021.
- [17] N. Hansen and X. Wang, "Generalization in reinforcement learning by soft data augmentation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 611–13 617.
- [18] A. Liang, J. Thomason, and E. Bryk, "Visarl: Visual reinforcement learning guided by human saliency," *arXiv preprint arXiv:2403.10940*, 2024.
- [19] M. Laskin, K. Lee, A. Stooke, L. Pinto, P. Abbeel, and A. Srinivas, "Reinforcement learning with augmented data," *Advances in neural information processing systems*, vol. 33, pp. 19 884–19 895, 2020.
- [20] I. Kostrikov, D. Yarats, and R. Fergus, "Image augmentation is all you need: Regularizing deep reinforcement learning from pixels," *arXiv preprint arXiv:2004.13649*, 2020.
- [21] Z. Yuan, S. Yang, P. Hua, C. Chang, K. Hu, and H. Xu, "RI-vigen: A reinforcement learning benchmark for visual generalization," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [22] R. Agarwal, M. C. Machado, P. S. Castro, and M. G. Bellemare, "Contrastive behavioral similarity embeddings for generalization in reinforcement learning," *arXiv preprint arXiv:2101.05265*, 2021.
- [23] M. Laskin, A. Srinivas, and P. Abbeel, "Curl: Contrastive unsupervised representations for reinforcement learning," in *International conference on machine learning*. PMLR, 2020, pp. 5639–5650.
- [24] A. Xie, L. Lee, T. Xiao, and C. Finn, "Decomposing the generalization gap in imitation learning for visual robotic manipulation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 3153–3160.
- [25] J. Xing, T. Nagata, K. Chen, X. Zou, E. Neftci, and J. L. Krichmar, "Domain adaptation in reinforcement learning via latent unified state representation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, 2021, pp. 10 452–10 459.
- [26] Z. Yuan, Z. Xue, B. Yuan, X. Wang, Y. Wu, Y. Gao, and H. Xu, "Pre-trained image encoder for generalizable visual reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 13 022–13 037, 2022.
- [27] C. Chen, J. Xu, W. Liao, H. Ding, Z. Zhang, Y. Yu, and R. Zhao, "Focus-then-decide: segmentation-assisted reinforcement learning," in

- Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 10, 2024, pp. 11 240–11 248.
- [28] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, “Segment anything,” *arXiv:2304.02643*, 2023.
- [29] T. Yu, Z. Zhang, C. Lan, Y. Lu, and Z. Chen, “Mask-based latent reconstruction for reinforcement learning,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 25 117–25 131, 2022.
- [30] M. Tomar, R. Islam, M. Taylor, S. Levine, and P. Bachman, “Ignorance is bliss: Robust control via information gating,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [31] K. Balakrishnan, P. Chakravarty, and S. Shrivastava, “An a\* curriculum approach to reinforcement learning for rgbd indoor robot navigation,” *arXiv preprint arXiv:2101.01774*, 2021.
- [32] S. Joshi, S. Kumra, and F. Sahin, “Robotic grasping using deep reinforcement learning,” in *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2020, pp. 1461–1466.
- [33] H. Friji, H. Ghazzai, H. Besbes, and Y. Massoud, “A dqn-based autonomous car-following framework using rgb-d frames,” in *2020 IEEE Global Conference on Artificial Intelligence and Internet of Things (GCAIoT)*. IEEE, 2020, pp. 1–6.
- [34] S. James and A. J. Davison, “Q-attention: Enabling efficient learning for vision-based robotic manipulation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1612–1619, 2022.
- [35] D. Yarats, R. Fergus, A. Lazaric, and L. Pinto, “Mastering visual continuous control: Improved data-augmented reinforcement learning,” *arXiv preprint arXiv:2107.09645*, 2021.
- [36] R. S. Sutton, “Learning to predict by the methods of temporal differences,” *Machine learning*, vol. 3, pp. 9–44, 1988.
- [37] A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018.
- [38] Y. Zhu, J. Wong, A. Mandlekar, R. Martín-Martín, A. Joshi, S. Nasiriany, Y. Zhu, and K. Lin, “robosuite: A modular simulation framework and benchmark for robot learning,” in *arXiv preprint arXiv:2009.12293*, 2020.
- [39] G. Ghiasi, Y. Cui, A. Srinivas, R. Qian, T.-Y. Lin, E. D. Cubuk, Q. V. Le, and B. Zoph, “Simple copy-paste is a strong data augmentation method for instance segmentation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 2918–2928.
- [40] J. Elsner, “Taming the panda with python: A powerful duo for seamless robotics programming and integration,” *SoftwareX*, vol. 24, p. 101532, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2352711023002285>