

# UC San Diego

## UC San Diego Previously Published Works

### Title

Generalization in Reinforcement Learning by Soft Data Augmentation

### Permalink

<https://escholarship.org/uc/item/326235h0>

### Authors

Hansen, Nicklas

Wang, Xiaolong

### Publication Date

2021-06-05

### DOI

10.1109/icra48506.2021.9561103

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Generalization in Reinforcement Learning by Soft Data Augmentation

Nicklas Hansen<sup>\*†</sup>, Xiaolong Wang<sup>\*</sup>


**Abstract**—Extensive efforts have been made to improve the generalization ability of Reinforcement Learning (RL) methods via domain randomization and data augmentation. However, as more factors of variation are introduced during training, optimization becomes increasingly challenging, and empirically may result in lower sample efficiency and unstable training. Instead of learning policies directly from augmented data, we propose **SOft Data Augmentation (SODA)**, a method that decouples augmentation from policy learning. Specifically, SODA imposes a soft constraint on the encoder that aims to maximize the mutual information between latent representations of augmented and non-augmented data, while the RL optimization process uses strictly non-augmented data. Empirical evaluations are performed on diverse tasks from DeepMind Control suite as well as a robotic manipulation task, and we find SODA to significantly advance sample efficiency, generalization, and stability in training over state-of-the-art vision-based RL methods.<sup>1</sup>

## I. INTRODUCTION

Reinforcement Learning (RL) from visual observations has achieved tremendous success in many areas, including game-playing [1], robotic manipulation [2]–[4], and navigation tasks [5], [6]. While this combination of learning both representation and decision-making jointly with RL has led to great success, numerous studies have shown that RL agents fail to generalize to new environments [7]–[11], which is only exacerbated by the high-dimensional nature of images in practical applications.

To improve generalization in RL, domain randomization [12]–[15] has been widely applied to learn robust representations invariant to visual changes. The assumption made in these approaches is that the distribution of randomized environments sufficiently cover factors of variation encountered at test-time. However, as more factors of variation are introduced during training, RL optimization becomes increasingly challenging, and empirically may lead to lower sample efficiency and unstable training. If we constrain the randomization to be small, the distribution is unlikely to cover variations in the test environment. Finding the right balance between sample efficiency and generalization therefore requires extensive engineering efforts.

Similarly, recent studies also show that by using the right kinds of data augmentation, both sample efficiency and generalization of RL policies can be improved substantially [9], [16], [17]. Determining which data augmentations to use (and to which degree) is however found to be task-dependent, and a naïve application of data augmentation may lead



(a) Environments for DeepMind Control tasks. We consider 5 challenging continuous control tasks from this benchmark.

(b) Environments for robotic manipulation. The task is to push the yellow cube to the location of the red disc.

**Fig. 1. Generalization in RL.** Agents are trained in a fixed environment (denoted the *training* environment) and we measure generalization to unseen environments with (i) *random colors* and (ii) *video backgrounds*. To simulate real-world deployment, we additionally randomize camera, lighting, and texture during evaluation in the robotic manipulation task. Additional samples are shown in appendix I.

to unstable training and poor performance [9], [17]. It is therefore natural to ask whether we can obtain the benefits of data augmentation while interfering minimally with the RL optimization process.

In this paper, we propose **SOft Data Augmentation (SODA)**, a method that stabilizes training by decoupling data augmentation from policy learning. While previous work attempts to learn policies directly from augmented data, SODA uses strictly *non-augmented* data for policy learning, and instead performs auxiliary representation learning using *augmented* data. Through its auxiliary task, SODA learns to generalize by maximizing the mutual information between latent representations of augmented and non-augmented data. The proposed auxiliary task shares a learned encoder with the policy, and SODA therefore imposes a soft constraint on the learned representation. As the policy is learned only from non-augmented data, the difficulty of RL optimization is greatly reduced, while SODA still benefits substantially from data augmentation through representation learning.

We perform visual representation learning by projecting augmented observations into a compact latent space using the shared encoder and a learned projection, and similarly projecting non-augmented observations using a moving average of the learned network. The SODA objective is then to learn a mapping from latent features of the augmented observation to those of the original observation. With strong

<sup>\*</sup>University of California, San Diego, CA, USA

<sup>†</sup>Technical University of Denmark, Denmark

<sup>1</sup>Webpage: <https://nicklashansen.github.io/SODA/>

and varied data augmentation, SODA learns to ignore factors of variation that are irrelevant to the RL task, which greatly reduces observational overfitting [10]. Our visual representation learning task is a self-supervised learning task. Our approach is related to recent work on joint learning with self-supervision and RL [18]–[20], where two tasks are trained jointly on the same augmented observations. However, our method is fundamentally different. We decouple the training data flow by using non-augmented data for RL and using augmented data only for representation learning. Besides, instead of learning invariance by contrasting two augmented instances of the same image to a batch of negative samples [19], [21], [22], SODA learns to map augmented images to their non-augmented counterparts in latent space, without the need for negative samples. Because we strictly use data augmentation for representation learning and instead impose a soft constraint on the shared encoder through joint training, we call it *soft* data augmentation.

Empirical evaluations are performed on the novel *DM-Control Generalization Benchmark* (DMControl-GB) based on continuous control tasks from DeepMind Control suite (DMControl) [23], as well as a robotic manipulation task. We train in a fixed environment and evaluate generalization to environments that are unseen during training, as shown in Figure 1. Our method improves both sample efficiency and generalization over state-of-the-art vision-based RL methods in 9 out of 10 environments from DMControl-GB and all 3 types of test environments in robotic manipulation.

We highlight our main contributions as follows:

- We present SODA, a stable and efficient representation learning method for vision-based RL.
- We propose the *DMControl Generalization Benchmark*, a new benchmark for generalization in vision-based RL.
- We show that SODA outperforms prior state-of-the-art methods significantly in test-time generalization to unseen and visually diverse environments.

## II. RELATED WORK

**Learning Visual Invariance with Self-Supervision.** Self-supervision for visual representation learning has proven highly successful in computer vision [24]–[28]. Recently, contrastive learning for self-supervised pre-training has achieved very compelling results for many downstream visual recognition tasks [21], [22], [29]–[32]. For example, Chen et al. [21] perform an extensive study on different augmentations (e.g. random cropping, color distortion, rotation) and show that learned representations invariant to these transformations can make good initializations for downstream tasks. Instead of learning with a contrastive objective, Grill et al. [33] show that invariance can also be learned via a prediction error without using any negative samples. Our method also learns invariant feature representations, but rather than aligning two different augmented views of the same instance as in previous work, we instead learn to align the representation of augmented and non-augmented views. As in Grill et al., SODA does not require negative samples.

**Self-Supervised Visual Learning for RL.** Inspired by the success of self-supervised visual representation learning, researchers have proposed to perform joint learning with self-supervision and RL to improve sample efficiency and performance [18], [19], [34]. For example, Yarats et al. [18] show that jointly training an auto-encoder together with RL improves the sample efficiency of model-free RL. Srinivas et al. [19] further extend this idea by performing auxiliary contrastive learning together with RL, and nearly match state-based RL in sample efficiency. While the proposed contrastive learning task uses data augmentation in training, the same augmented data is used for RL as well. In contrast to these approaches, we separate the training data into two data streams and strictly use the non-augmented data for RL, which helps stabilize training, decrease sensitivity to the choice of data augmentation, and allows us to benefit from stronger data augmentation in the self-supervised learning. Recent work also attempts to close the generalization gap by adapting a learned policy using observations collected from a target environment [35]–[37]. Instead, SODA learns a *single* policy that generalizes *without* data from test environments.

**Data Augmentation and Randomization.** Our work is inspired by research on generalization by domain randomization [12]–[15], [38]. For example, Tobin et al. [13] render scenes with random textures for object localization in simulation, and show that the learned networks transfer to the real world. Similar generalization across environments have been shown in recent work on RL with data augmentation [9], [16], [17], [39]. In these works, policies are trained to be invariant to certain visual perturbations by augmenting observations. However, it is unpredictable and task-dependent how much randomization can be applied. Instead of training RL policies directly on augmented observations, we propose *soft* data augmentation that decouples training into two objectives with separate data streams.

## III. METHOD

We propose **Soft Data Augmentation (SODA)**, a novel approach to data augmentation in RL that exhibits significant improvements in sample efficiency and stability over prior work. SODA is a general framework for data augmentation that can be implemented on top of any standard RL algorithm. It aims to learn representations that effectively encode information shared between an observation and its augmented counterpart, while interfering minimally with the RL objective. In the following, we present the individual components of SODA in detail.

### A. Architectural Overview

SODA is implemented as a self-supervised auxiliary task that shares a common encoder  $f$  with an RL policy. For a given policy network parameterized by a collection of parameters  $\theta$ , we split the network and corresponding parameters sequentially into an encoder  $f_\theta$  and a policy  $\pi_\theta$  such that  $\mathbf{a} = \pi_\theta(f_\theta(o))$  outputs a distribution over actions (and any other algorithm-specific values) for an input observation  $o$ . SODA then consists of the following three components:

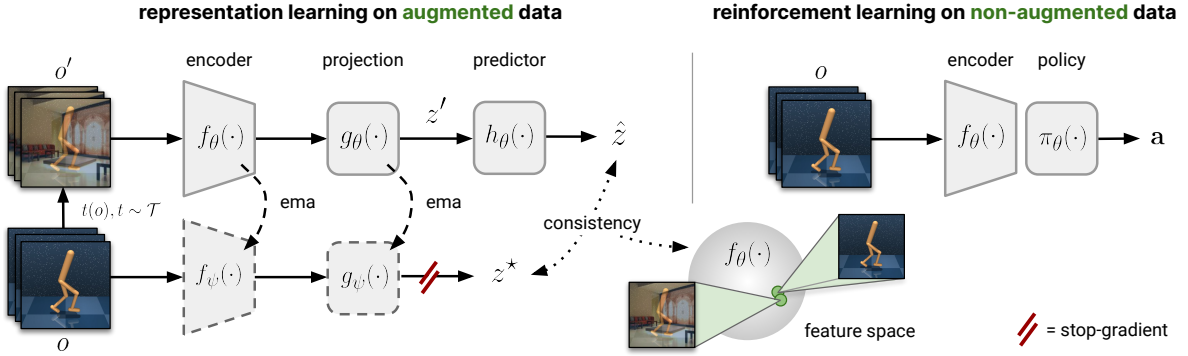


Fig. 2. **SODA architecture.** *Left:* an observation  $o$  is augmented to produce a view  $o'$ , which is then encoded and projected into  $z' = g_\theta(f_\theta(o'))$ . Likewise,  $o$  is encoded by  $f_\psi$  and projected by  $g_\psi$  to produce features  $z^*$ . The SODA objective is then to predict  $z^*$  from  $z'$  by  $h_\theta$  formulated as a consistency loss. *Right:* Reinforcement Learning in SODA. The RL task remains unchanged and is trained directly on the non-augmented observations  $o$ . *ema* denotes an exponential moving average.

the shared encoder  $f_\theta$ , a projection  $g_\theta$ , and a prediction head  $h_\theta$ . We additionally consider an architecturally identical encoder  $f_\psi$  and projection  $g_\psi$  where  $\psi$  denotes a collection of parameters separate from  $\theta$ . We denote  $f_\theta$  as the *online* encoder and  $f_\psi$  as the *target* encoder, and similarly for projections  $g_\theta, g_\psi$ . We then define the parameters  $\psi$  as an exponential moving average (EMA) of  $\theta$ , and update  $\psi$  with every iteration of SODA using the update rule,

$$\psi_{n+1} \leftarrow (1 - \tau)\psi_n + \tau\theta_n \quad (1)$$

for an iteration step  $n$  and a momentum coefficient  $\tau \in (0, 1]$ , such that only parameters  $\theta$  are updated by gradient descent [22], [33], [40]. Figure 2 provides an overview of the architecture, and the SODA task is detailed in the following.

### B. Representation Learning by SODA

Given an observation  $o \in \mathbb{R}^{C \times H \times W}$ , we sample a data augmentation  $t \sim \mathcal{T}$  and apply it to produce an augmented observation  $o' \triangleq t(o)$ . As such,  $o$  and  $o'$  can be considered different *views* of the same observation, where  $o$  is the original observation and  $o'$  is corrupted by some noise source. If we assume  $f : \mathbb{R}^{C \times H \times W} \rightarrow \mathbb{R}^{C' \times H' \times W'}$  such that  $H' \times W'$  is a feature map down-sampled from  $H \times W$ , then both projections  $g_\theta, g_\psi$  are learned mappings  $g : \mathbb{R}^{C' \times H' \times W'} \rightarrow \mathbb{R}^K$  where  $K \ll C' \times H' \times W'$ . Given a feature vector  $z' \triangleq g_\theta(f_\theta(o'))$ , the proposed SODA task is then to learn a mapping  $h_\theta : \mathbb{R}^K \rightarrow \mathbb{R}^K$  that predicts the target projection  $z^* \triangleq g_\psi(f_\psi(o))$  of the non-augmented observation  $o$ , as shown in Figure 2 (left). We optimize projection  $g_\theta$ , predictor  $h_\theta$ , and shared encoder  $f_\theta$  jointly together with the RL task(s), and employ a consistency loss

$$\mathcal{L}_{SODA}(\hat{z}, z^*; \theta) = \mathbb{E}_{t \sim \mathcal{T}} \left[ \|\hat{z}_o - z_o^*\|_2^2 \right] \quad (2)$$

where  $\hat{z} \triangleq h_\theta(z')$ , and  $\hat{z}_o \triangleq \hat{z} / \|\hat{z}\|_2$ ,  $z_o^* \triangleq z^* / \|z^*\|_2$  are  $\ell_2$ -normalizations of  $\hat{z}$  and  $z^*$ , respectively. We use a predictor  $h_\theta$  rather than mapping  $z'$  directly to  $z^*$  as we find it to improve expressiveness in the learned representations, which is consistent with prior work [21], [33], [34].

The policy  $\pi_\theta$  (including any algorithm-specific task heads) uses  $f_\theta$  to extract features and is optimized with no modifications to architecture nor inputs, i.e. we optimize the RL objective  $\mathcal{L}_{RL}$  directly on non-augmented observations  $o$  and update  $\pi_\theta, f_\theta$  by gradient descent, as shown in Figure 2 (right). During training, we continuously alternate between optimizing  $\mathcal{L}_{RL}$  and  $\mathcal{L}_{SODA}$ . At test-time, we only use the encoder  $f_\theta$  and policy  $\pi_\theta$ . The training procedure is summarized in Algorithm 1, and we further motivate the algorithmic design in section III-C.

---

### Algorithm 1 Soft Data Augmentation (SODA)

---

- $\theta, \psi$ : randomly initialized network parameters
  - $\omega$ : RL updates per iteration
  - $\tau$ : momentum coefficient
  - 1: **for** every iteration **do**
  - 2:   **for** update = 1, 2, ...,  $\omega$  **do**
  - 3:     Sample batch of transitions  $\nu \sim \mathcal{B}$
  - 4:     Optimize  $\mathcal{L}_{RL}(\nu)$  wrt  $\theta$
  - 5:   Sample batch of observations  $o \sim \mathcal{B}$
  - 6:   Augment observations  $o' = t(o)$ ,  $t \sim \mathcal{T}$
  - 7:   Compute online predictions  $\hat{z} = h_\theta(g_\theta(f_\theta(o')))$
  - 8:   Compute target projections  $z^* = g_\psi(f_\psi(o))$
  - 9:   Optimize  $\mathcal{L}_{SODA}(\hat{z}, z^*)$  wrt  $\theta$
  - 10:   Update  $\psi \leftarrow (1 - \tau)\psi + \tau\theta$
- 

### C. Data Augmentation as a Mutual Information Problem

SODA reformulates the problem of generalization as a representational consistency learning problem: the encoder  $f$  should learn to map different views of the same underlying state to similar feature vectors. This encourages the encoder to learn features that are shared across views, e.g. physical quantities and object interactions, and discard information that yields no predictive power such as background, lighting, and high-frequency noise. Given an observation  $o$  and a data augmentation  $t$ , we seek to encode  $o$  and  $o' = t(o)$  into compact feature vectors  $z', z^* \in \mathbb{R}^K$ , respectively, by learned mappings  $f_\theta, g_\theta$  and  $f_\psi, g_\psi$  such that the mutual information






Fig. 3. **Data augmentation.** We consider the following two data augmentations: *random convolution* (as proposed by [9], [16]) and *random overlay* (novel). See appendix III for additional data augmentation samples.

$I$  between  $o$  and  $o'$  is maximally preserved. The mutual information between  $z^*$  and  $z'$  is then given by

$$I(z^*; z') = \sum_{z'} \sum_{z^*} p(z^*, z') \log \frac{p(z^*|z')}{p(z^*)} \quad (3)$$

and is naturally bounded by the mutual information between  $o$  and  $o'$ , i.e.  $I(z^*; z') \leq I(o; o')$ . As  $I(z^*; z')$  is impractical to optimize directly, we approximate  $I$  as a simple consistency loss defined in (2) by the following intuition. If we assume  $f_\psi, g_\psi$  to maximally preserve information in  $o$ , then  $I(z^*; z') \propto I(o; o')$ , and minimizing  $\mathcal{L}_{SODA}$  thus maximizes  $I(z^*; z'|o')$ . In other words, by learning a mapping from  $z'$  to  $z^*$ , we maximally preserve information in  $o$  by learning to discard noise added by the data augmentation  $t(\cdot)$ . With strong and varied data augmentation,  $f_\theta$  learns to ignore factors of variation that are irrelevant to the RL task, and reduces observational overfitting [10], [41]. While the assumption of maximally preserved information may not hold in practice, we find Equation 2 to be a good enough approximation for expressive representation learning. Although the SODA objective does not explicitly prevent trivial solutions, e.g.  $f_\theta(\cdot) = \mathbf{0}$ , such behavior is implicitly discouraged, and the reason for that is two-fold: (i) trivial equilibria in  $\mathcal{L}_{SODA}$  are shown to be unstable due to updates in the targets  $f_\psi, g_\psi$  [33]; and (ii)  $f_\theta$  is jointly optimized between SODA and the RL objective(s), and trivial equilibria are in direct conflict with the RL objective since they entail that  $f_\theta$  then would preserve no information in  $o$ . While the predictor  $h_\theta$  is not strictly needed, we empirically find it to improve the sample efficiency of our method, which we conjecture is due to non-stationarity of the target projection  $z^*$ .

#### IV. EXPERIMENTS

We evaluate SODA on 5 tasks from DeepMind control suite (DMControl) [23], as well as in robotic manipulation. DMControl is comprised of diverse and challenging continuous control tasks and is a widely used benchmark for vision-based RL [9], [18], [19], [39], [42]. To quantify the generalization ability of SODA, we propose *DMControl Generalization Benchmark* (DMControl-GB)<sup>2</sup>, a new benchmark for generalization in vision-based RL, based on DMControl. Agents are trained in a fixed environment (denoted the *training environment*), and we measure generalization to two

distinct test distributions: (1) environments with randomized colors; and (2) environments with natural videos as background, first proposed by [42]. These test distributions correspond to the `color_hard` and `video_easy` benchmarks from DMControl-GB; we provide results for the full DMControl-GB benchmark in appendix I.

While DMControl-GB provides a good platform for benchmarking algorithms, we are ultimately interested in developing algorithms that solve real-world problems with vision-based RL. To better emulate real-world deployment scenarios, we also consider a robotic manipulation task on a robotic arm (in simulation). As in DMControl-GB, agents are trained in a fixed environment and evaluated on environments with randomized colors and video backgrounds, and we additionally perform random perturbations of camera, lighting, and texture to simulate real-world conditions during testing. Samples from DMControl-GB and robotic manipulation are shown in Figure 1, and additional samples are shown in appendix I.

##### A. Implementation Details


SODA is a general framework that can be implemented on top of any standard RL algorithm using parameterized policies. In this work we implement SODA using Soft Actor-Critic (SAC) [43] as the base algorithm. We adopt network architecture and hyperparameters from [42], which implements  $f_\theta$  as 11 convolutional layers and  $\pi_\theta$  as a multi-layer perceptron (MLP). Both  $g_\theta, h_\theta$  are implemented as MLPs with batch normalization [44] and we use a batch size of 256 for the SODA task. Projections are of dimension  $K = 100$ , and the target components  $f_\psi, g_\psi$  are updated with a momentum coefficient  $\tau = 0.005$ .  $\mathcal{L}_{RL}$  and  $\mathcal{L}_{SODA}$  are optimized using Adam [45], and in practice we find it sufficient to only make a SODA update after every second RL update, i.e.  $\omega = 2$ . In all tasks considered, observations are stacks of  $k$  consecutive frames ( $k = 3$  for DMControl-GB and  $k = 1$  for robotic manipulation) of size  $100 \times 100$  and with no access to state information. See appendix II for further implementation details and a list of hyper-parameters.

##### B. Data Augmentation

While SODA makes no assumptions about the data augmentation used, it is useful to distinguish between *weak* augmentations that may improve sample efficiency but contribute minimally to generalization, and *strong* augmentations that improve generalization at the cost of sample efficiency [46]. An example of the former is random cropping [9], [39], whereas examples of the latter are domain randomization [12], [13] and random convolution [9], [16]. Consistent with previous work [9], [19], [42], we apply temporally consistent random cropping (down to  $84 \times 84$ ) by default in SODA and all baselines, and we refer to the cropped images as non-augmented observations.

We additionally seek to stabilize training with *strong* augmentations for generalization, and consider the following two data augmentations: *random convolution*, where a randomly initialized convolutional layer is applied; and a novel *random*

<sup>2</sup>SODA and benchmark is open-sourced at: <https://github.com/nicklashansen/dmcontrol-generalization-benchmark>



**Fig. 4. Random convolution.** *Top*: average return on the training environment during training. *Bottom*: periodic evaluation of generalization ability measured by average return on the *random color* environment. SODA exhibits sample efficiency and convergence similar to SAC but improves generalization significantly. Average of 5 runs, shaded area is std. deviation.

*overlay* that linearly interpolates between an observation  $o$  and an image  $\varepsilon$  to produce an augmented view

$$t_{\text{overlay}}(o) = (1 - \alpha)o + \alpha\varepsilon, \quad \varepsilon \sim \mathcal{D} \quad (4)$$

where  $\alpha \in [0, 1]$  is the interpolation coefficient and  $\mathcal{D}$  is an unrelated dataset. In practice, we use  $\alpha = 0.5$  and sample images from Places [47], a dataset containing 1.8M diverse scenes. Samples of both data augmentations are shown in Figure 3, and additional samples are shown in Appendix III.


### C. Baselines

We compare SODA to the following baselines: (i) a baseline SAC that is equivalent to RAD [9] as we apply random cropping to all baselines; (ii) SAC trained with domain randomization on the *random colors* test distribution (denoted *SAC (DR)*); (iii) SAC using random convolution as data augmentation (denoted *SAC (conv)*); and (iv) SAC using random overlay as data augmentation (denoted *SAC (overlay)*). Likewise, our method is denoted *SODA (conv)* and *SODA (overlay)*. Data augmentation baselines perform RL on augmented observations, whereas SODA does not; we ablate this decision in section IV-D.

### D. Stability under Strong Data Augmentation

Domain randomization and data augmentation is known to decrease sample efficiency and can destabilize the RL training. Figure 4 shows training performance and generalization of baselines and SODA using random convolution. While SAC converges in training, its generalization remains poor. SODA improves generalization substantially without destabilizing the RL training, and exhibits a sample efficiency similar to the SAC baseline. Both SAC trained with domain randomization and *SAC (conv)* either fail to solve the tasks or converge to sub-optimal policies.

To further investigate the effect of soft data augmentation, we consider a variant of SODA in which random convolution



**Fig. 5. Soft data augmentation.** Average return on the training environment for *walker\_walk* and *walker\_stand* tasks. *Augment RL* corresponds to the *SAC (conv)* baseline, *Augment both* applies random convolution in both SODA and RL, and *Augment SODA* is the proposed formulation of SODA. Average of 5 runs, shaded area is std. deviation.

augmentation is applied in both representation learning and RL, and we summarize our findings in Figure 5. While using data augmentation in both objectives (*Augment both*) improves significantly over only applying data augmentation in RL (*Augment RL*), we find the proposed formulation of SODA (*Augment SODA*) superior in both sample efficiency and end-performance.

### E. Generalization to Unseen Environments

We evaluate the generalization ability of SODA on the challenging *random colors* and *video backgrounds* benchmarks from DMControl-GB (*color\_hard* and *video\_easy*, respectively), and additionally compare to a number of recent state-of-the-art methods: (i) CURL [19], a contrastive representation learning method; (ii) RAD [9], a study on data augmentation for RL that uses random cropping and is equivalent to the SAC baseline in section IV-C; and (iii) PAD [42], a method for self-supervised policy adaptation. Results are shown in Table I. We find SODA

TABLE I. **Generalization.** Average return of methods trained in a fixed environment and evaluated on: (*left*) DMControl-GB with natural videos as background; and (*right*) DMControl-GB with random colors. Mean and std. deviation of 5 runs.

DMControl-GB (generalization)	video backgrounds							random colors				
	CURL [19]	RAD [9]	PAD [42]	SAC (DR)	SAC (conv)	SODA (conv)	SAC (overlay)	SODA (overlay)	CURL [19]	RAD [9]	PAD [42]	SODA (overlay)
walker, walk	556 ±133	606 ±63	717 ±79	520 ±107	169 ±124	635 ±48	718 ±47	<b>768 ±38</b>	445 ±99	400 ±61	468 ±47	<b>692 ±68</b>
walker, stand	852 ±75	745 ±146	935 ±20	839 ±58	435 ±100	903 ±56	<b>960 ±2</b>	955 ±13	662 ±54	644 ±88	797 ±46	<b>893 ±12</b>
cartpole, swingup	404 ±67	373 ±72	521 ±76	524 ±184	176 ±62	474 ±143	718 ±30	<b>758 ±62</b>	454 ±110	590 ±53	630 ±63	<b>805 ±28</b>
ball_in_cup, catch	316 ±119	481 ±26	436 ±55	232 ±135	249 ±190	539 ±111	713 ±125	<b>875 ±56</b>	231 ±92	541 ±29	563 ±50	<b>949 ±19</b>
finger, spin	502 ±19	400 ±64	691 ±80	402 ±208	355 ±88	363 ±185	607 ±68	<b>695 ±97</b>	691 ±12	667 ±154	<b>803 ±72</b>	793 ±128




Fig. 6. **Random overlay.** *Top:* average return on the training environment during training. *Bottom:* periodic evaluation of generalization ability measured by average return on the *random color* environment. SODA offers better sample-efficiency than the novel SAC (*overlay*) baseline and similar generalization to SODA (*conv*) even though there is minimal visual similarity between random overlays and the random color environment. Average of 5 runs, shaded area is std. deviation.

to outperform previous methods in **9** out of **10** instances, and by as much as **81%** on *ball\_in\_cup\_catch* (video backgrounds). Further, SODA also outperforms the baselines from section IV-C in all tasks except *walker\_stand* on video backgrounds.

While SODA with random convolution data augmentation generalizes well to the random colors distribution, it generalizes comparably worse to videos. We conjecture that this is because random convolution captures less factors of variation than overlay (e.g. textures and local color changes), which is in line with the findings of [8], [42]. SODA (*overlay*) is found to generalize well to both test distributions even though there is minimal visual similarity between random overlays and the random color environment. This shows that SODA benefits from strong and varied data augmentation, and we leave it to future work to explore additional (soft) data augmentations. For completeness, training and generalization curves for the random overlay data augmentation are shown in Figure 6, and we provide results and comparisons on two additional benchmarks, *color\_easy* and *video\_hard*, from DMControl-GB in appendix I.

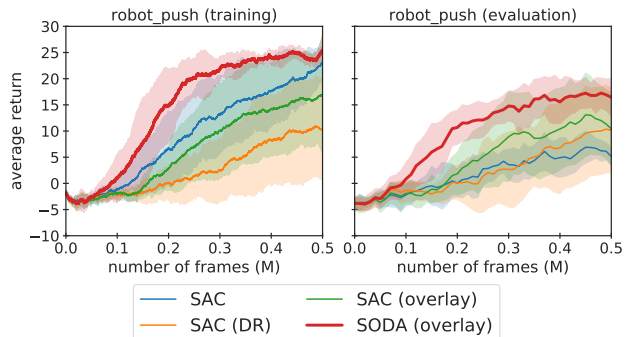


Fig. 7. **Robotic manipulation.** *Left:* average return on the training environment during training. *Right:* periodic evaluation of generalization measured by average return on the *random color* environment. SODA outperforms all baselines in both sample efficiency and generalization, and reduces variance. Average of 5 runs, shaded area is std. deviation.



**TABLE II. Generalization in robotic manipulation.** Average return of SODA and baselines in the training and test environments. Mean and std. deviation of 5 runs.

robot, push	SAC	SAC (DR)	SAC (conv)	SODA (conv)	SAC (overlay)	SODA (overlay)
training	14.5 ±5.2	9.3 ±7.5	0.2 ±1.7	14.4 ±4.6	15.8 ±7.3	<b>21.3</b> ±5.0
random colors	7.2 ±3.4	7.8 ±7.0	-3.1 ±1.4	1.4 ±3.4	13.5 ±4.4	<b>18.0</b> ±3.7
video bg.	0.9 ±6.0	7.0 ±7.4	-4.9 ±1.9	-4.5 ±1.1	7.6 ±7.1	<b>13.9</b> ±2.6

### F. Robotic manipulation

We additionally consider a robotic manipulation task in which the goal is for a robotic arm to push a yellow cube to the location of a red disc (as illustrated in Figure 1). Observations are  $100 \times 100$  images and agents are trained using dense rewards. Episodes consist of 50 time steps, and at each time step there is a positive reward of 1 for having the cube at the goal location (disc), and a small penalty proportional to the distance between the cube and goal. An average return of 25 therefore means that the cube is in goal position for more than half of all time steps.

As in DMControl-GB, we consider generalization to environments with (i) random colors; and (ii) video backgrounds. To simulate real-world deployment, we also randomize camera, lighting, and texture during testing. Training and generalization curves are shown in Figure 7, and results are summarized in Table II. We find SODA to outperform baselines in both sample efficiency and generalization to all considered test environments, which suggests that SODA may be a suitable method for reducing observational overfitting in real-world deployment of policies for robotic manipulation.

## V. CONCLUSION

We show empirically that SODA improves stability, sample efficiency, and generalization significantly over previous methods and a set of strong baselines. We dissect the success of SODA and find that it benefits substantially from (i) not learning policies on augmented data; (ii) its representational consistency learning formulation; and (iii) strong and varied data augmentation for representation learning.

### ACKNOWLEDGEMENTS

This work was supported, in part, by grants from DARPA LwLL, NSF 1730158 CI-New: Cognitive Hardware and Software Ecosystem Community Infrastructure (CHASE-CI), NSF ACI-1541349 CC\*DNI Pacific Research Platform, as well as gifts from Qualcomm and TuSim-ple.

### REFERENCES

[1] 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. 1

[2] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016. 1

[3] A. V. Nair, V. Pong, M. Dalal, S. Bahl, S. Lin, and S. Levine, "Visual reinforcement learning with imagined goals," in *Advances in Neural Information Processing Systems*, 2018, pp. 9191–9200. 1

[4] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine, "Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation," *ArXiv*, vol. abs/1806.10293, 2018. 1

[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 *ICRA*. IEEE, 2017, pp. 3357–3364. 1

[6] W. Yang, X. Wang, A. Farhadi, A. Gupta, and R. Mottaghi, "Visual semantic navigation using scene priors," in *ICLR*, 2019. 1

[7] S. Gamrian and Y. Goldberg, "Transfer learning for related reinforcement learning tasks via image-to-image translation," *ArXiv*, vol. abs/1806.07377, 2019. 1

[8] K. Cobbe, O. Klimov, C. Hesse, T. Kim, and J. Schulman, "Quantifying generalization in reinforcement learning," 2018. 1, 6

[9] M. Laskin, K. Lee, A. Stooke, L. Pinto, P. Abbeel, and A. Srinivas, "Reinforcement learning with augmented data," *arXiv preprint arXiv:2004.14990*, 2020. 1, 2, 4, 5, 6, 8, 9

[10] X. Song, Y. Jiang, S. Tu, Y. Du, and B. Neyshabur, "Observational overfitting in reinforcement learning," 2019. 1, 2, 4, 8

[11] K. Cobbe, C. Hesse, J. Hilton, and J. Schulman, "Leveraging procedural generation to benchmark reinforcement learning," *ArXiv*, vol. abs/1912.01588, 2019. 1

[12] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel, "Asymmetric actor critic for image-based robot learning," *arXiv preprint arXiv:1710.06542*, 2017. 1, 2, 4

[13] 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," *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep 2017. 1, 2, 4

[14] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018. 1, 2

[15] J. Yang, B. Petersen, H. Zha, and D. Faissol, "Single episode policy transfer in reinforcement learning," 2019. 1, 2

[16] K. Lee, K. Lee, J. Shin, and H. Lee, "A simple randomization technique for generalization in deep reinforcement learning," *ArXiv*, vol. abs/1910.05396, 2019. 1, 2, 4, 9

[17] R. Raileanu, M. Goldstein, D. Yarats, I. Kostrikov, and R. Fergus, "Automatic data augmentation for generalization in deep reinforcement learning," *ArXiv*, vol. abs/2006.12862, 2020. 1, 2

[18] D. Yarats, A. Zhang, I. Kostrikov, B. Amos, J. Pineau, and R. Fergus, "Improving sample efficiency in model-free reinforcement learning from images," 2019. 2, 4

[19] A. Srinivas, M. Laskin, and P. Abbeel, "Curl: Contrastive unsupervised representations for reinforcement learning," *arXiv preprint arXiv:2004.04136*, 2020. 2, 4, 5, 6, 8, 9

[20] A. Stooke, K. Lee, P. Abbeel, and M. Laskin, "Decoupling representation learning from reinforcement learning," *arXiv:2004.14990*. 2

[21] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," 2020. 2, 3

[22] K. He, H. Fan, Y. Wu, S. Xie, and R. B. Girshick, "Momentum contrast for unsupervised visual representation learning," *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9726–9735, 2020. 2, 3

[23] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. de Las Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, T. Lillicrap, and M. Riedmiller, "DeepMind control suite," <https://arxiv.org/abs/1801.00690>, DeepMind, Tech. Rep., Jan. 2018. 2, 4, 8

[24] X. Wang and A. Gupta, "Unsupervised learning of visual representations using videos," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 2794–2802. 2

[25] C. Doersch, A. Gupta, and A. A. Efros, "Unsupervised visual representation learning by context prediction," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1422–1430. 2

[26] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, "Context encoders: Feature learning by inpainting," in *Proceedings*


of the *IEEE conference on computer vision and pattern recognition*, 2016, pp. 2536–2544. 2

- [27] M. Noroozi and P. Favaro, “Unsupervised learning of visual representations by solving jigsaw puzzles,” in *European Conference on Computer Vision*. Springer, 2016, pp. 69–84. 2
- [28] R. Zhang, P. Isola, and A. A. Efros, “Colorful image colorization,” in *European conference on computer vision*. Springer, 2016, pp. 649–666. 2
- [29] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, “Unsupervised feature learning via non-parametric instance discrimination,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3733–3742. 2
- [30] A. van den Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” 2018. 2
- [31] Y. Tian, D. Krishnan, and P. Isola, “Contrastive multiview coding,” *arXiv preprint arXiv:1906.05849*, 2019. 2
- [32] I. Misra and L. van der Maaten, “Self-supervised learning of pretext-invariant representations,” in *CVPR*, 2020. 2
- [33] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. A. Pires, Z. D. Guo, M. G. Azar, B. Piot, K. Kavukcuoglu, R. Munos, and M. Valko, “Bootstrap your own latent: A new approach to self-supervised learning,” *ArXiv*, vol. abs/2006.07733, 2020. 2, 3, 4
- [34] M. Schwarzer, A. Anand, R. Goel, R. D. Hjelm, A. Courville, and P. Bachman, “Data-efficient reinforcement learning with momentum predictive representations,” *arXiv preprint arXiv:2007.05929*, 2020. 2, 3
- [35] T. Carr, M. Chli, and G. Vogiatzis, “Domain adaptation for reinforcement learning on the atari,” in *AAMAS*, 2019. 2
- [36] X. Chen, A. Ghadirzadeh, M. Björkman, and P. Jensfelt, “Adversarial feature training for generalizable robotic visuomotor control,” *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1142–1148, 2020. 2
- [37] R. Julian, B. Swanson, G. Sukhatme, S. Levine, C. Finn, and K. Hausman, “Never stop learning: The effectiveness of fine-tuning in robotic reinforcement learning,” *arXiv: Learning*, 2020. 2
- [38] F. Ramos, R. Possas, and D. Fox, “Bayessim: Adaptive domain randomization via probabilistic inference for robotics simulators,” *Robotics: Science and Systems XV*, Jun 2019. 2
- [39] I. Kostrikov, D. Yarats, and R. Fergus, “Image augmentation is all you need: Regularizing deep reinforcement learning from pixels,” 2020. 2, 4
- [40] T. Lillicrap, J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” *CoRR*, vol. abs/1509.02971, 2016. 3
- [41] C. Packer, K. Gao, J. Kos, P. Krähenbühl, V. Koltun, and D. Song, “Assessing generalization in deep reinforcement learning,” 2018. 4
- [42] N. Hansen, Y. Sun, P. Abbeel, A. A. Efros, L. Pinto, and X. Wang, “Self-supervised policy adaptation during deployment,” 2020. 4, 5, 6, 8, 9
- [43] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, and S. Levine, “Soft actor-critic algorithms and applications,” 2018. 4, 9
- [44] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” *ArXiv*, vol. abs/1502.03167, 2015. 4, 9
- [45] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2015. 4, 9
- [46] Y. Tian, C. Sun, B. Poole, D. Krishnan, C. Schmid, and P. Isola, “What makes for good views for contrastive learning,” *ArXiv*, vol. abs/2005.10243, 2020. 4
- [47] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba, “Places: A 10 million image database for scene recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 5, 9

## APPENDIX I

### DMCONTROL GENERALIZATION BENCHMARK

We open-source *DMControl Generalization Benchmark* (DMControl-GB), a benchmark for generalization in vision-based RL, based on the DeepMind Control suite [23]. In DMControl-GB, agents are trained in a fixed environment (denoted the *training* environment), and the generalization ability of agents is evaluated on unseen, visually diverse test



**Fig. 8. Samples from DMControl-GB.** Agents are trained in a fixed environment (denoted the *training* environment), and the generalization ability of agents is evaluated on the test distributions shown in (a-d). Environments (a-b) randomize the color of background, floor, and the agent itself, while (c-d) replaces the background with natural videos. In (c), only the skybox is replaced, whereas the entirety of the background is replaced in (d). Our main results, Figure 4, Figure 6, and Table I, consider the *color\_hard* and *video\_easy* benchmarks of DMControl-GB.

distributions. Samples from each of the proposed benchmarks are shown in Figure 8. We benchmark a number of recent, state-of-the-art methods for sample-efficiency and generalization: (i) CURL [19], a contrastive representation learning method; (ii) RAD [9], a study on data augmentation for RL that uses random cropping; and (iii) PAD [42], a method for self-supervised policy adaptation. All algorithms are implemented in a unified framework using standardized architecture and hyper-parameters, wherever applicable. DMControl-GB, SODA, and implementations of the baselines are open-sourced at <https://github.com/nicklashansen/dmcontrol-generalization-benchmark>.

#### A. Additional results on DMControl-GB

Our main results, Figure 4, Figure 6, and Table I, consider the *color\_hard* and *video\_easy* benchmarks of DMControl-GB. Results from our experiments on two additional benchmarks from DMControl-GB are shown in Table III. SODA outperforms all previous state-of-the-art methods in 9 out of 10 environments and achieves near-optimal generalization on the *color\_easy* benchmark. On the extremely challenging *video\_hard* benchmark, SODA shows substantial improvements over baselines, but there is still room for improvement on the majority of tasks. We conjecture that the considerable performance gap between the training environment and *video\_hard* is due to observational overfitting [10], and in our experiments we find that our baseline methods tend to rely on visual features such as



TABLE III. Average return of SODA and state-of-the-art methods on the `color_easy` and `video_hard` environments from DMControl-GB. Mean and std. deviation of 5 runs. Results for `color_hard` and `video_easy` are shown in Table I.

DMControl-GB ( <code>color_easy</code> )	CURL [19]	RAD [9]	PAD [42]	SODA (overlay)
walker, walk	645 ±55	636 ±33	687 ±119	<b>811</b> ±41
walker, stand	866 ±46	807 ±67	894 ±39	<b>960</b> ±4
cartpole, swingup	668 ±74	763 ±29	812 ±20	<b>859</b> ±15
ball_in_cup, catch	565 ±168	727 ±87	775 ±159	<b>969</b> ±3
finger, spin	781 ±139	789 ±160	<b>870</b> ±54	855 ±93
DMControl-GB ( <code>video_hard</code> )	CURL [19]	RAD [9]	PAD [42]	SODA (overlay)
walker, walk	58 ±18	56 ±9	93 ±29	<b>381</b> ±72
walker, stand	45 ±5	231 ±39	278 ±72	<b>771</b> ±83
cartpole, swingup	114 ±15	110 ±16	123 ±24	<b>429</b> ±64
ball_in_cup, catch	115 ±33	97 ±29	66 ±61	<b>327</b> ±100
finger, spin	27 ±21	34 ±11	56 ±18	<b>302</b> ±41

shadows that may be correlated with the reward signal, but are not necessarily present at test-time.

## APPENDIX II IMPLEMENTATION DETAILS

Table IV summarizes the hyper-parameters of SODA and baselines used in our DMControl-GB experiments. We use identical algorithms, architectures, and hyper-parameters for robotic manipulation, but opt for observations as stacks of  $k = 3$  frames in DMControl-GB and  $k = 1$  in robotic manipulation. We implement SODA and baselines using Soft Actor-Critic (SAC) [43] as the base algorithm, and adopt network architecture and hyperparameters from [42]. We apply the same architecture and hyperparameters in all methods, wherever applicable. As the action spaces of both DMControl-GB and the robotic manipulation task are continuous, the policy learned by SAC outputs the mean and variance of a Gaussian distribution over actions. Frames are RGB images rendered at  $100 \times 100$  and cropped to  $84 \times 84$  as in [9], [19], [42], such that inputs to the network are of size  $3k \times 84 \times 84$ , where the first dimension indicates the color channels and the following dimensions represent spatial dimensions. The same crop is applied to all frames in a stack. The shared encoder consists of 11 convolutional layers and outputs features of size  $32 \times 21 \times 21$ . In our SODA implementation, both  $g_\theta, h_\theta$  are implemented as MLPs with batch normalization [44] and we use a batch size of 256 for the SODA task. Projections are of dimension  $K = 100$ , and

TABLE IV. Hyper-parameters for the DMControl-GB tasks.

Hyper-parameter	Value
Frame rendering	$3 \times 100 \times 100$
Frame after crop	$3 \times 84 \times 84$
Stacked frames	3
Number of conv. layers	11
Number of filters in conv.	32
Action repeat	2 (finger-spin) 8 (cartpole-swingup) 4 (otherwise)
Discount factor $\gamma$	0.99
Episode time steps	1,000
Learning algorithm	Soft Actor-Critic
Number of training steps	500,000
Replay buffer size	500,000
Optimizer (RL/aux.)	Adam ( $\beta_1 = 0.9, \beta_2 = 0.999$ )
Optimizer ( $\alpha$ )	Adam ( $\beta_1 = 0.5, \beta_2 = 0.999$ )
Learning rate (RL)	1e-3
Learning rate ( $\alpha$ )	1e-4
Learning rate (SODA)	3e-4
Batch size (RL)	128
Batch size (SODA)	256
Actor update freq.	2
Critic update freq.	1
Auxiliary update freq.	1 (CURL) 2 (PAD/SODA) N/A (otherwise)
Momentum coef. $\tau$ (SODA)	0.005



(a) Random convolution data augmentation.



(b) Random overlay data augmentation.

Fig. 9. Samples from data augmentation. In this work, we consider the *random convolution* data augmentation shown in (a), as well as the novel *random overlay* shown in (b).

the target components  $f_\psi, g_\psi$  are updated with a momentum coefficient  $\tau = 0.005$ .  $\mathcal{L}_{RL}$  and  $\mathcal{L}_{SODA}$  are optimized using Adam [45], and in practice we find it sufficient to only make a SODA update after every second RL update, i.e.  $\omega = 2$ .

## APPENDIX III DATA AUGMENTATIONS

Section IV-B discusses two data augmentations that we consider in this work: (1) *random convolution* [9], [16], where a randomly initialized convolutional layer is applied; and (2) *random overlay*, where we linearly interpolate between an observation  $o$  and an image  $\varepsilon$  to produce an augmented view  $o' = t_{overlay}(o)$ , as formalized in Equation 4. In practice, we use an interpolation coefficient of  $\alpha = 0.5$  and sample images from Places [47], a dataset containing 1.8M diverse scenes. Samples from the two data augmentations are shown in Figure 9.