

TEMPORAL DISENTANGLEMENT OF REPRESENTATIONS FOR IMPROVED GENERALISATION IN REINFORCEMENT LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Reinforcement Learning (RL) agents are often unable to generalise well to environment variations in the state space that were not observed during training. This issue is especially problematic for image-based RL, where a change in just one variable, such as the background colour, can change many pixels in the image, which can lead to drastic changes in the agent’s latent representation of the image, causing the learned policy to fail. To learn more robust representations, we introduce **TE**mporal **D**isentanglement (TED), a self-supervised auxiliary task that leads to disentangled image representations exploiting the sequential nature of RL observations. We find empirically that RL algorithms utilising TED as an auxiliary task adapt more quickly to changes in environment variables with continued training compared to state-of-the-art representation learning methods. Since TED enforces a disentangled structure of the representation, we also find that policies trained with TED generalise better to unseen values of variables irrelevant to the task (e.g. background colour) as well as unseen values of variables that affect the optimal policy (e.g. goal positions).

1 INTRODUCTION

Real-world environments are often not static and deterministic, but can be subject to changes, both incremental or with sudden effects. Reinforcement Learning (RL) algorithms need to be robust to these changes and adapt quickly. Moreover, since many real-world robotics applications rely on images as inputs (Vecerik et al., 2021; Hämmäläinen et al., 2019; Chebotar et al., 2019), RL agents need to learn robust representations of images that remain useful after a change in the environment. For example, a simple change in lighting conditions can change the perceived colour of an object, but this should not affect the agent’s ability to perform a task.

One of the reasons RL agents fail to generalise to unseen values of environment variables, such as colours and object positions, is that they overfit to variations seen in training (Zhang et al., 2018). The issue is especially problematic for image-based RL, where a change in one environment variable can mean the agent is presented with a very different set of pixels for which trained RL policies are often no longer optimal. This failure to generalise occurs for both variables that are irrelevant to the optimal policy, such as background colours, and relevant variables, such as goal positions (Kirk et al., 2022). In practice, this often results in agents needing to adapt their policy after a change to only one variable. We show experimentally that agents often cannot recover the optimal policy after the environment changes because it is too difficult to ‘undo’ the overfitting.

One approach to tackle the generalisation issue is to use domain randomisation during training to maximise the environment variations observed (Cobbe et al., 2019; Chebotar et al., 2019). However, in practice, we may be unaware of what variations an agent might see in the future. Even if all possible future variations are known, training on this full set is often sample inefficient and may result in a sub-optimal policy as the agent learns to compensate for many different values. It can also be impractical to generate all variations during training in the real world or even in simulation due to limitations on laboratory setups or in simulation suites. An alternative to domain randomisation is to learn robust representations that generalise to unseen variations (Kirk et al., 2022), but learning such representations remains an open problem (Lan et al., 2022). Some approaches aim to learn a

representation that is invariant to image distractors (Zhang et al., 2020; 2021) to improve generalisation to unseen values of only variables irrelevant to the optimal policy. These approaches do not enforce a structure on the remaining task-relevant variables in the representation for generalisation to unseen values of relevant variables.

A promising direction towards robust RL is to learn a disentangled representation of observations. Disentanglement techniques aim to learn robust representations for both task relevant and irrelevant variables by separating distinct, informative factors of variation in an image into the unknown ground truth factors that generated the image (Bengio et al., 2013). When one factor of variation changes to a previously unseen value, such as a new colour, changing many pixels in the image, only a subset of features in a disentangled representation will change. The RL agent will still be able to rely on the remaining unchanged features in the representation to adapt quickly, allowing generalisation and performance recovery similar to state-based RL. Higgins et al. (2017b) show that a disentangled representation improves generalisation in RL. However, this approach requires a trained or hard-coded policy to collect independent and identically distributed (i.i.d.) data to pre-train a VAE (Higgins et al., 2017a) of it, and it has since been proven that it is theoretically impossible to learn a disentangled representation from i.i.d. data alone (Locatello et al., 2019).

We introduce a self-supervised auxiliary task for learning disentangled representations for the robust encoding of images, which we call Temporal Disentanglement (TED). Unlike previous work, TED can be implemented with only minimal changes to existing RL algorithms and allows for lifelong learning of disentangled representations. In contrast to Higgins et al. (2017b), our approach uses the non-i.i.d. temporal data from consecutive timesteps in RL to learn the disentangled representation online. We provide experimental results across a variety of tasks from the DeepMind Control Suite (Tunyasuvunakool et al., 2020), Panda Gym (Gálloci et al., 2021) and Procgen (Cobbe et al., 2020) environments. For our experiments, we train on a subset of some of the environment variables (such as colours), then evaluate generalisation on a test environment with unseen values of the variables and continue training to demonstrate adaptation to the new environment. Our results demonstrate that TED improves generalisation of a variety of base RL algorithms on unseen environment variables that are relevant or irrelevant to the task, [while state-of-the-art baselines that achieve equally good training performance still fail to adapt and, in some cases, are unable to recover after overfitting to the training environment](#). We also evaluate a disentanglement metric (Higgins et al., 2017a) to demonstrate that our approach increases the extent to which the learned representation has disentangled the factors of variation in the image observations compared to baselines.

2 RELATED WORK

2.1 GENERALISATION IN IMAGE-BASED REINFORCEMENT LEARNING

Image augmentation. Image augmentation artificially increases the size of the dataset by adding image perturbations to improve robustness of image representations. Laskin et al. (2020a) apply a variety of image augmentation techniques such as translation, cutouts and cropping; Yarats et al. (2021) average over multiple augmentations; Hansen & Wang (2021) maximise mutual information between the representations of augmented and non-augmented images; and Hansen et al. (2021) use both augmented and non-augmented images to stabilise Q-value estimation. However, image augmentation approaches to generalisation can still fail when the agent experiences stronger types of variation after training (Kirk et al., 2022). In our experiments, we show that TED can be used alongside image augmentation techniques to further improve generalisation while benefiting from the augmentation.

Learning invariant representations. Invariance techniques aim to learn a representation that ignores distractors in the image. Zhang et al. (2020) use causal inference techniques assuming a block MDP structure, Zhang et al. (2021) use a bisimulation metric. Li et al. (2021) use domain adversarial optimisation to learn a representation invariant to visual distractors. These approaches all aim to generalise to unseen values of irrelevant variables, e.g. background colours. In contrast, TED uses disentanglement to enforce a structured representation that applies to both relevant and irrelevant variables.

Encoding inductive biases. Some approaches use an auxiliary task to encode inductive biases. Laskin et al. (2020b) learn a representation to maximise similarity between different augmentations of the

same observation, Mazouze et al. (2020) maximise similarity between observations at successive timesteps, and Agarwal et al. (2021) enforce a structured representation for task relevant variables using policy similarity metrics. These approaches are based on enforcing a similarity constraint between pairs of observations to learn informative features but they do not enforce any structure to the representation, whereas TED encourages a disentangled structure due to the form of the classifier without a similarity constraint. van den Oord et al. (2018) learn representations that capture information predictive of the future, Schwarzer et al. (2021) pretrain an encoder and re-tune on task specific data, and Jaderberg et al. (2017) uses a combination of multiple auxiliary tasks for representation learning, but these approaches also do not require a disentangled structure. To encode the inductive bias of disentanglement, Higgins et al. (2017b) train a VAE of the using data from a pre-trained or hardcoded agent. In contrast, our approach uses the temporal structure of data available in RL to learn a disentangled representation online with the RL policy.

2.2 DISENTANGLED REPRESENTATIONS

Unsupervised learning. Many variations of the Variational Autoencoder (VAE) (Kingma & Welling, 2014) aim to improve disentanglement of the learned representation, such as β -VAE (Higgins et al., 2017a) (Burgess et al., 2017) and the Factor-VAE (Kim & Mnih, 2018). Locatello et al. (2019) prove that it is theoretically impossible to learn disentangled representations from i.i.d. data alone. To bypass the impossibility result, many recent approaches extend VAE to use some form of supervision. Shu et al. (2020) provide weak supervision using a labelled grouping of images, while Locatello et al. (2020) generate pairs of images with limited factors of variation between pairs. In contrast, our approach exploits non-i.i.d. temporal observations available in RL to learn a disentangled representation without labelling or artificially generating images.

Independent Component Analysis (ICA). Hyvärinen & Morioka (2016) introduce Time-Contrastive Learning to learn a disentangled representation from time-series data with stationary factors. Hyvärinen & Morioka (2017) disentangle time-series data with stationary factors, introducing Permutation Contrastive Learning (PCL) to train a classifier to discriminate between temporal and non-temporal inputs. We consider RL episodes back-to-back as a time-series. Features that are randomised at the start of an episode are stationary, and due to the agent learning and adapting behaviour over time, features controlled by the agent are non-stationary. The TED classification objective is based on the PCL classifier structure to encourage disentanglement, but unlike PCL, TED addresses the combination of stationary and non-stationary features in RL.

3 PRELIMINARIES

We assume the environment is a fully-observable Markov Decision Process (MDP), defined as the tuple $M = (S; A; P; R; \gamma)$, where S is a set of states, A is a set of actions, $P : S \times S \times A \rightarrow [0; 1]$ is the state-transition function, $R : S \times A \rightarrow \mathbb{R}$ is the reward function, and $\gamma \in [0; 1]$ is the discount factor. An RL agent chooses an action $a_t \in A$ at time t based on its current state $s_t \in S$ and its policy $\pi(a_t | s_t)$. The agent then transitions to the next state according to the state-transition probability $P(s_{t+1} | s_t; a_t)$, and receives a reward $r_t = R(s_t; a_t)$. The goal of an RL agent is to learn a policy to maximise the expected discounted cumulative rewards, $E_P[\sum_{t=0}^{\infty} \gamma^t r_t]$.

In this work, the agent's observation $o_t \in O$ at time t is image pixels, a high-dimensional representation of the true underlying environment state $s_t \in S$. We assume the environment has a factored state representation where each factor corresponds to an environment variable. The components of the state vector are the unobserved ground truth factors of variation. We consider one observation to be a stack of consecutive frames to ensure all features, such as velocities, can be extracted from o_t . We assume o_t is generated by an invertible, non-linear transformation of the state factors s_t , $h : S \rightarrow O$, such that $o_t = h(s_t)$. The aim of disentanglement is to learn a representation $z_t \in Z$ that recovers the independent ground truth factors s_t from the observations o_t by approximating the inverse of the transformation, such that $z_t \approx f(o_t)$. To simplify notation, we will usually denote $z_t = f(o_t)$ as $z(o_t)$. We will use z^i to denote the i -th component of the vector z .

Figure 1: TED architecture: The classifier is trained to discriminate between temporal and non-temporal samples to encourage the encoder to disentangle the temporal structure in the image observations. The `//` indicates that the gradient flow is stopped. `2 f e; e^0 g and t-2 f t + 1; t^0; t^0 g` depending on the observation being processed.

4 TEMPORAL DISENTANGLEMENT

We introduce Temporal Disentanglement (TED) as an auxiliary task for representation learning with RL algorithms. Our goal is to improve generalisation to both relevant and irrelevant features when their testing variation is unknown a priori. TED aims to disentangle the factors of variation that generate an observation, such as background colour and the trajectory of an object across the frame stack. We will provide an overview of architecture in Section 4.1 then discuss the details of the TED auxiliary task in Section 4.2.

4.1 ARCHITECTURE OVERVIEW

The high-level architecture for TED with a generic RL algorithm is depicted in Figure 1. TED is designed to encourage the encoder $O : Z \rightarrow \mathbb{R}^d$ to recover the temporal structure determined by observations at consecutive timesteps t and $t+1$, such that a classifier $g : Z \times Z \rightarrow \mathbb{R}$ can discriminate between temporal and non-temporal pairs of observation encodings. We structure the classifier to compare each feature in the representation separately to encourage the encoder to disentangle the temporal structure into the ground truth factors that generated the observation. We simultaneously perform dimensionality reduction such that $\dim(Z) = \dim(O)$.

TED requires a batch of transitions B containing consecutive observations originating from various episodes. The classifier loss is applied as an auxiliary loss to any base RL algorithm. For off-policy algorithms, the batch can be sampled from the replay buffer following the sampling procedure for the base RL algorithm. For on-policy base algorithms, the batch can be created using multiple parallel environments to ensure transitions from different episodes, which is common in practice, for which $B = D$ in the subsequent explanations. Image augmentations can be applied to the observations where required by the base algorithm, and TED uses the same augmented images as the base algorithm. Both the TED classification loss and the relevant RL loss are used to learn the encoder parameters. For training stability, we also use a target encoder (He et al., 2020; Laskin et al., 2020b) $f : O \rightarrow Z$ for the next observation at a given timestep $t+1$, where $o_{t+1}^e = f(o_t^e)$. Only temporal transitions are used for RL, but non-temporal pairs are created where 1 to train the classifier, which we describe in more detail in the next section.

4.2 TEMPORAL DISENTANGLEMENT CLASSIFIER

We use the TED auxiliary loss to train a classifier to discriminate between temporal and non-temporal pairs of observation encodings. TED encourages the encoder to learn a representation that uncovers the temporal structure in the data to enable distinguishing whether two observations occurred consecutively or not.

Classifier inputs. For each batch of transitions $B = D$, we create three types of observation-pair batches for classifier input: 1) temporal samples, 2) different episode, non-temporal samples and 3) same episode, non-temporal samples. A single transition contains $(o_t^e, o_{t+1}^e) \in \mathbb{R}^d \times \mathbb{R}^d$ where e is the episode index corresponding to transition. The temporal sample (o_t^e, o_{t+1}^e) consists of two consecutive observations. Due to the episodic nature of RL, we use two types of non-temporal

Algorithm 1 TED update step

```

Input: batch of transition  $\mathcal{B} = \{f; \dots; (o_t^e; o_{t+1}^e); \dots; g; D\}$ 
Create batch  $\mathcal{Z}_{\text{obs}}$  of representations for each observation  $z_t^e = f(o_t^e)$ 
Create batch  $\mathcal{Z}_{\text{nextObs}}$  of representations for each next observation  $z_{t+1}^e = f(o_{t+1}^e)$ 
for each transition  $i \in \mathcal{B}$  do
  Create temporal sample  $x_t = (z_t^e; z_{t+1}^e), X = X \cup \{x_t\}$ 
  Sample  $z_{t_0}^e \in \mathcal{Z}_{\text{nextObs}}$  such that  $t_0 \in e$ 
  Create non-temporal different episode sample  $x_{t_0}^0 = (z_t^e; z_{t_0}^e), X^0 = X^0 \cup \{x_{t_0}^0\}$ 
  Sample  $o_{t_{00}}^e \in D$  such that  $t_{00} \neq f; t; t+1; g$ , and get representation  $z_{t_{00}}^e = f(o_{t_{00}}^e)$ 
  Create non-temporal same episode sample  $x_{t_{00}}^0 = (z_t^e; z_{t_{00}}^e), X^{00} = X^{00} \cup \{x_{t_{00}}^0\}$ 
end for
for each sample  $x \in X; X^0; X^{00}$  do
  Classifier prediction  $y = g(x)$  (see Equation 1)
  Calculate binary cross-entropy loss  $\mathcal{L}_{\text{TED}}(x; l)$  (see Equation 2)
end for
Calculate average loss for the batch  $\bar{\mathcal{L}}_{\text{TED}} = \text{mean}(\mathcal{L}_{\text{TED}}(X))$ 
Backpropagate loss to update encoder parameters and classifier parameters
Update target encoder parameters  $\tau = \tau + (1 - \beta) \bar{\mathcal{L}}_{\text{TED}}$ 
Output: Loss  $\bar{\mathcal{L}}_{\text{TED}}$  and updated parameters  $\tau, \theta$ 

```

samples. The non-temporal sample $x_{t_0}^0 = (o_t^e; o_{t_0}^e)$ consists of non-consecutive timesteps from different episodes, where $t_0 \in \mathcal{B}$ such that $t_0 \in e$. These samples encourage the encoder to learn a representation (o_t^e) , that disentangles episodic features that are chosen randomly at the start of an episode, such as colours, as this is sufficient to distinguish $x_{t_0}^0$ from the temporal sample x_t . The second non-temporal sample $x_{t_{00}}^0 = (o_t^e; o_{t_{00}}^e)$ consists of non-consecutive timesteps from the same episode where $t_{00} \in D$. These same episode non-temporal samples encourage the encoder to disentangle features that change during the episode, such as agent and object positions, because episodic features will be the same for both x_t and $x_{t_{00}}^0$.

Classification objective. Given a batch of samples $\{x_t \in X; X^0; X^{00}\}$, a logistic regression classifier is trained to discriminate between the corresponding representation of temporal samples $x_t \in X$, and non-temporal samples $x_{t_0}^0 \in X^0$ and $x_{t_{00}}^0 \in X^{00}$. To learn a disentangled representation, we use a regression function of the form proposed by Hyen & Morioka (2017):

$$y(x_t) = g(x_t^1; x_t^2) = \sum_{i=1}^N [k_1^i z^i(x_t^1) + k_2^i z^i(x_t^2) + b^i] / (k^i z^i(x_t^1) + b^i)^2 + c \quad (1)$$

for all $x_t \in X; X^0; X^{00}$, where N is the dimensionality of the representation $z(o_t^e)$, $x_t = (x_t^1; x_t^2)$, and $c = f(k_1; k_2; b; k; b; c; g)$ are the classifier parameters to be trained simultaneously with the encoder parameters.

Temporal samples $x_t \in X$, are given a classification label $l = 1$, and non-temporal samples $x_{t_0}^0 \in X^0$ and $x_{t_{00}}^0 \in X^{00}$, are given a classification label $l = 0$. The classifier is trained using the cross-entropy loss for binary classification for all $x_t \in X; X^0; X^{00}$:

$$\mathcal{L}_{\text{TED}}(x_t; l) = -\lambda (2l \log(y(x_t)) + (1-l) \log(1-y(x_t))) \quad (2)$$

where λ is the sigmoid function. Since there are two non-temporal samples $x_{t_0}^0 \in X^0$ and $x_{t_{00}}^0 \in X^{00}$, for each transition in the batch \mathcal{B} but only one temporal $x_t \in X$, the positive temporal samples are weighted by 2. The coefficient λ is a hyperparameter to be tuned to the task. It is used to scale up the classification auxiliary loss to ensure it is not dwarfed by the RL loss and to prioritise representation learning over policy learning at the start of training while the factors of variation, are independent. The structure of the classifier $g(x_t^1; x_t^2)$ (Equation 1) ensures each feature z^i is considered separately to the other features. This encourages the encoder to not only uncover the temporal structure necessary for classification, but to do so by separating the factors of variation into distinct features in the representation so that the temporal structure of each feature can be determined independently of the other features. The second term of Equation 1 approximates the marginal log-pdf. Due to this structure, we expect the encoder to learn a disentangled representation and the classifier to approximate the distribution of the samples (Hyen & Morioka, 2017). The pseudocode for an update step is provided in Algorithm 1, which is performed for every update of the

base RL algorithm. We train the encoder with both \mathcal{L}_{TED} and the RL loss. It is important to note that the TED auxiliary loss requires only creating temporal and non-temporal pairs of representations to train the classifier. The classifier is not required for execution so it can be discarded after training.

5 EXPERIMENTAL RESULTS

Our experiments are designed to evaluate whether TED allows zero-shot generalisation to unseen values of an environment variable, and whether TED promotes faster adaptation of the base RL algorithm after a change in environment variables if generalisation is not instant. We evaluate our approach on settings where the agent must generalise to unseen values of task relevant and irrelevant distractor variables with continued learning on the test environment. We use a training environment with a subset of some of the environment variables (such as colours), and evaluate generalisation on a test environment with unseen values of the variables. Images of the train and test environment for each experiment are provided in Appendix C. We show results on a variety of different tasks with RAD (Laskin et al., 2020a), SVEA (Hansen et al., 2021) and PPO (Schulman et al., 2017) as different base algorithms utilising the TED loss, which covers on-policy and off-policy, continuous and discrete control, and with and without image augmentations, demonstrating the flexibility of our approach. Our results show that TED consistently improves the generalisation of the base RL algorithm in all tasks. TED also shows lower variance across seeds compared to baselines, making it more reliable.

5.1 GENERALISATION TO TASK-IRRELEVANT VARIABLES

To demonstrate generalisation to unseen values of task-irrelevant environment variables, we use continuous control tasks from the DeepMind Control Suite (DMC) (Tunyasuvunakool et al., 2020) and Panda Gym (Galloudec et al., 2021) as simulations of robotics tasks. We adapt DMC wrappers from the Distracting Control Suite (Stone et al., 2021) to add colour distractors to the observations as irrelevant factors of variation.

Experiment setup. We show results with RAD and SVEA as two different base algorithms for continuous control. We use the cartpole swingup and walkerwalk tasks from DMC with SVEA as the base algorithm. For RAD, we use the easier nspin task instead of walkerwalk since RAD was unable to learn an optimal policy on the walkerwalk task due to the difficulty of the task with the colour distractors. We also use the Reach task with dense rewards from Panda Gym where the agent receives a reward of the negative of its distance to the goal. We train each algorithm on a fixed set of colours by varying the RGB colour values within a small bounded region of the original value. We test on a set of colours of the same size that were not observed during training. Following the original setup of the base RL algorithms, the encoder for SVEA has 11 convolutional layers and RAD has 4 convolutional layers. Full implementation details are described in Appendix B.

Baselines. We compare to the chosen base algorithm for each task to demonstrate the performance improvement achieved by TED. For further comparison, we also compare our method to representative baselines from each category of RL generalisation method discussed in Section 2.1. Our data augmentation baseline is DrQ (Yarats et al., 2021). We use DBC (Zhang et al., 2021) as a baseline method that learns invariant representations. We also compare with CURL (Laskin et al., 2020b) as a state-of-the-art contrastive auxiliary task. Finally, we include the base algorithm with domain randomisation, shown as base algorithm-DR in the figures, as a privileged baseline that is trained on both the 'train' and 'test' colours together, demonstrating that when the test variations are known a priori, it is often less sample efficient to use domain randomisation and sometimes unable to learn the optimal policy. All baselines use the same size encoder as the TED base algorithm on a given task, so results for the same baseline differ slightly depending on the base algorithm it is being compared to. For all baselines, we tuned hyperparameters by grid search and report the best performing ones.

Disentanglement metric. To evaluate the learned representations, we measure the disentanglement using the disentanglement metric proposed by Higgins et al. (2017a). The metric measures disentanglement using pairs of images with one factor of variation fixed to the same value in both images and the other factors randomised. For example, the background colour could be held fixed while all other colours and variables are randomly assigned. We use pairs of frame stacks

(a) cartpoleswingup (b) nger_spin (c) pandareach

Figure 2: Generalisation to unseen colours at the vertical dotted line with RAD base algorithm. RAD-TED (ours) recovers more quickly than baselines and achieves higher returns than domain randomisation (RAD-DR). Returns are the average of 10 evaluation episodes, averaged over 5 seeds, shaded region shows standard deviation.

(a) cartpoleswingup (b) walker_walk (c) pandareach

Figure 3: Generalisation to unseen colours at the vertical dotted line. SVEA-TED (ours) recovers more quickly than baselines and achieves similar performance to domain randomisation but with fewer assumptions. Returns are the average of 10 evaluation episodes, averaged over 5 seeds, shaded region shows standard deviation.

corresponding to ω_t instead of individual images to allow the policy to extract velocity information. One sample for the classifier is calculated for a batch of observation pairs, given by:

$$z_{diff} = \frac{1}{B} \sum_{b=1}^B |z(o_t^{(b)}) - z(o_{t'}^{(b)})| \tag{3}$$

where $o_t^{(b)}$ and $o_{t'}^{(b)}$ have the same fixed factor. We then train a linear classifier that takes as one input and predicts which factor was held fixed across a batch of inputs. The accuracy of this classifier is the disentanglement score. The intuition is that there will be low variance in the features corresponding to the fixed factor in a disentangled representation so the classifier will be highly accurate. Although independence of the factors does not strictly hold in practice in many RL environments, it does hold when generating the observations for the metric because the value of each factor is assigned randomly, so we can fairly assess the disentanglement. A score of 1.0 is the maximum for a fully disentangled representation.

Results. The results for RAD as the base algorithm are shown in Figure 2, and the results for SVEA are shown in Figure 3. The figures show that TED improves the generalisation of both RAD and SVEA across all tasks. In many tasks, TED is able to generalise instantly without any reduction in performance at all. In other tasks, e.g. see Figure 2a, TED experiences some reduction in performance but is able to recover more quickly than the base algorithm and other baselines. Surprisingly, the baselines that achieve equally good training performance as TED fail to adapt and, in some cases, are unable to recover after overfitting to the training colours. Even though TED increases the disentanglement score of the base algorithm, shown in Figure 1, there can still be an initial drop in performance on the training environment because the RL policy may still unnecessarily

	cartpole swingup	nger spin	panda reach		cartpole swingup	walker walk	panda reach
RAD-TED	0.99	0.79	0.95	SVEA-TED	0.64	0.79	0.90
RAD	0.88	0.56	0.83	SVEA	0.53	0.71	0.83
RAD-DR	0.67	0.53	0.49	SVEA-DR	0.58	0.62	0.72
CURL	0.87	0.89	0.91	CURL	0.92	0.77	0.65
DBC	0.65	0.46	0.58	DBC	0.63	0.61	0.34
DrQ	0.73	0.59	0.91	DrQ	0.86	0.66	0.60

(a) Tasks with RAD base algorithm (shown in Figure 2)

(b) Tasks with SVEA base algorithm (shown in Figure 3)

Table 1: Disentanglement scores at the end of training before changing to the test environment at the vertical dotted line.

put some weight on the colour features, however the disentangled structure allows faster recovery. TED achieves higher returns than the privileged domain randomisation baseline in many tasks (e.g. Figure 2), which assumes the test colours are known a priori, because it can be difficult to learn an optimal policy with such a larger set of colour distractors. The poor training performance of some of the baselines, particularly DBC, is due to the difficulty of learning an optimal policy with the colour distractors as it increases the size of the state space. DBC aims to maintain performance on difficult distractors resulting in reduced performance on 'easy' distractors.

5.2 GENERALISATION TO BOTH TASK-RELEVANT AND IRRELEVANT VARIABLES

We use the Procgen generalisation benchmark (Cobbe et al., 2020) to demonstrate generalisation to unseen values of both task-relevant and irrelevant variables together. Procgen is a set of discrete-control tasks that uses procedural generation to determine the layout, objects, entities, background, colours and other game details for each level of a game. This produces many factors of variation across the game levels, some of which are relevant to solving the game and others are irrelevant distractors. Due to the nature of procedural task generation with Procgen, we are unable to extract the factors of variation to calculate the disentanglement scores for these tasks.

Experiment setup. We use the coinrun and jumper Procgen environments. We train on 100 levels of the hard difficulty, and test generalisation on 100 unseen hard levels. Following the setup of Cobbe et al. (2020), we use PPO as the base algorithm for the Procgen tasks.

Results. The results are shown in Figure 4. The results show that PPO-TED (ours) recovers more quickly than PPO on the unseen levels. PPO is unable to recover optimal performance on both tasks after overfitting to the training levels, but PPO-TED converges to higher performance on the unseen levels, similar to that achieved in training. In the jumper environment, Figure 4b, PPO-TED also demonstrates better zero-generalisation performance (at the vertical dotted line) as well as faster adaptation to the unseen levels.

(a) coinrun

(b) jumper

Figure 4: Generalisation to unseen levels at the vertical dotted line. PPO-TED (ours) recovers more quickly than PPO on the unseen levels. Returns are the average of 10 evaluation episodes, averaged over 5 seeds, shaded region shows standard deviation. The graphs show the 10-point rolling average for readability.

(a) Comparison of different TED loss coefficients. (b) Comparison with modified TED using same episode samples or different episode samples only. (c) Comparison with modified TED using a standard linear classifier.

Figure 5: Ablations for RAD-TED on the cartpole/wingup task with colour distractors, averaged over 5 seeds.

5.3 ABLATION STUDIES

How does the loss coefficient affect performance? Figure 5a shows how the choice of the loss coefficient affects performance. We found $\lambda = 100$ to be optimal for this task, and the results show some robustness as $\lambda = 200$ still achieves good performance. A lower coefficient $\lambda = 50$ does not prioritise disentanglement enough reducing generalisation performance. A higher coefficient reduces performance because the agent is prioritising disentanglement too much over the optimal policy.

Is one type of non-temporal sample sufficient to improve generalisation? TED requires two types of non-temporal samples: from the same episode, and from different episodes. We explored two simplified versions of TED using non-temporal samples from the same episode only and non-temporal samples from different episodes only. The results in Figure 5b show that using a single type of non-temporal sample improves the generalisation of RAD to an extent, but RAD-TED further improves the generalisation performance.

How important is the disentangled structure of the TED classifier? We compared TED to a simplified version that uses a standard linear classifier instead of the disentangled structure of the TED classifier. The results in Figure 5c show that while the linear classifier improves the generalisation performance of RAD, it does not generalise as well as RAD-TED.

6 LIMITATIONS AND FUTURE WORK

Our approach has some limitations that could be addressed in future work. Guarantees of disentanglement usually assume the factors of variation are independent and do not generalise to correlated factors (Toube et al., 2021). While we have shown that TED improves disentanglement in practice, RL observations generated by a learning agent do not have independent factors of variation in general. We leave further exploration of how to relax this assumption for future work.

We introduced two types of non-temporal samples to disentangle features controlled by the agent (which form a non-stationary time-series) and episodic features (which form a stationary time-series). While there exists a theoretical proof that disentanglement is possible with either non-stationary (Hyvärinen & Morioka, 2016) or stationary (Hyvärinen & Morioka, 2017) factors, it is still an open problem to identify an approach that guarantees disentanglement in a time-series that contains both stationary and non-stationary factors.

Our approach introduces the TED loss coefficient as a new hyperparameter which must be tuned to the task. Future work could consider automatically tuning λ based on the current disentanglement score. Finally, our temporal samples are constructed from the current timestep t and next timestep $t + 1$ to allow TED to be easily added to existing algorithms. Future work could explore extending the horizon of the temporal sample to k with $k > 1$.

7 CONCLUSION

In this work, we introduced TED, a disentangled representation learning auxiliary task for RL. Our approach is the first to consider an auxiliary task based on disentangled representation learning

for online RL. TED can be used with existing algorithms and does not require a decoder. We demonstrated experimentally that TED improves generalisation of three different base algorithms by adapting with continued learning to previously unseen values of environment variables that are both task-relevant and irrelevant. We also adapted the notion of factors of variation to frame stacking in RL and used a disentanglement metric to show that TED improves the disentanglement of learned representations. TED is a step toward making RL algorithms more robust for real-world deployment and life-long learning as the agent is able to quickly recover when presented with previously unseen values of environment variables and continue learning **while reducing** catastrophic forgetting.

REPRODUCIBILITY STATEMENT

A public and open-source implementation of TED can be found at github.com/usr/repo [currently anonymised for double blind review]. Full details of the architecture and hyperparameter settings for each of our experiments, including implementation details of the disentanglement metric, can be found in Appendix B.

REFERENCES

- Rishabh Agarwal, Marlos C. Machado, Pablo Samuel Castro, and Marc G Bellemare. Contrastive behavioral similarity embeddings for generalization in reinforcement learning. *International Conference on Learning Representations*, 2021.
- Y. Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* volume 35, pp. 1798–1828, 2013.
- Christopher P. Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in *vae*. In *31st Conference on Neural Information Processing Systems (NIPS 2017)*, 2017.
- Yevgen Chebotar, Ankur Handa, Viktor Makoviychuk, Miles Macklin, Jan Issac, Nathan Ratliff, and Dieter Fox. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8973–8979. IEEE, 2019.
- Karl Cobbe, Oleg Klimov, Chris Hesse, Taehoon Kim, and John Schulman. Quantifying generalization in reinforcement learning. In *Proceedings of the 36th International Conference on Machine Learning (ICML 2019)* volume 97 of *Proceedings of Machine Learning Research*, pp. 1282–1289. PMLR, 2019.
- Karl Cobbe, Christopher Hesse, Jacob Hilton, and John Schulman. Leveraging procedural generation to benchmark reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)* volume 119 of *Proceedings of Machine Learning Research*, pp. 2048–2056. PMLR, 2020.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Vlad Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. IMPALA: scalable distributed deep-rl with importance weighted actor-learner architecture. *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)* volume 80 of *Proceedings of Machine Learning Research*, pp. 1407–1416. PMLR, 2018.
- Quentin Galouédec, Nicolas Cazin, Emmanuel Dellaëra, and Liming Chen. panda-gym: Open-Source Goal-Conditioned Environments for Robotic Learning. In *Robot Learning Workshop: Self-Supervised and Lifelong Learning at NeurIPS*, 2021.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Offpolicy maximum entropy deep reinforcement learning with a stochastic actor. *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)* volume 80 of *Proceedings of Machine Learning Research*, pp. 1861–1870. PMLR, 2018.
- Aleksi Härmäläinen, Karol Arndt, Ali Ghadirzadeh, and Ville Kyrki. Affordance learning for end-to-end visuomotor robot control. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1781–1788. IEEE, 2019.
- Nicklas Hansen and Xiaolong Wang. Generalization in reinforcement learning by soft data augmentation. In *International Conference on Robotics and Automation*, 2021.
- Nicklas Hansen, Hao Su, and Xiaolong Wang. Stabilizing deep q-learning with convnets and vision transformers under data augmentation. In *Conference on Neural Information Processing Systems*, 2021.

- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9726–9735, 2020.
- Irina Higgins, Lőc Matthey, Arka Pal, Christopher P. Burgess, Xavier Glorot, Matthew M. Botvinick, Shakir Mohamed, and Alexander Lerchner.-vae: Learning basic visual concepts with a constrained variational framework. *International Conference on Learning Representations 2017a*.
- Irina Higgins, Arka Pal, Andrei Rusu, Loic Matthey, Christopher Burgess, Alexander Pritzel, Matthew Botvinick, Charles Blundell, and Alexander Lerchner. DARLA: Improving zero-shot transfer in reinforcement learning. *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1480–1490. PMLR, 2017b.
- Aapo Hyvärinen and Hiroshi Morioka. Unsupervised feature extraction by time-contrastive learning and nonlinear ica. *Advances in Neural Information Processing Systems (NIPS 2006)*.
- Aapo Hyvärinen and Hiroshi Morioka. Nonlinear ica of temporally dependent stationary sources. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)* volume 54, 2017.
- Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. In *International Conference on Learning Representations 2017*.
- Hyunjik Kim and Andriy Mnih. Disentangling by factorising. *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2649–2658. PMLR, 2018.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *International Conference on Learning Representations 2014*.
- Robert Kirk, Amy Zhang, Edward Grefenstette, and Tim Raschel. A survey of generalisation in deep reinforcement learning. *arXiv:2111.09794* 2022.
- Charline Le Lan, Stephen Tu, Adam Oberman, Rishabh Agarwal, and Marc G. Bellemare. On the generalization of representations in reinforcement learning. *International Conference on Artificial Intelligence and Statistics (AISTATS 2022)*.
- Michael Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. *34th Conference on Neural Information Processing Systems (NeurIPS 2020)* 2020a.
- Michael Laskin, Aravind Srinivas, and Pieter Abbeel. CURL: Contrastive unsupervised representations for reinforcement learning. *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*, volume 119 of *Proceedings of Machine Learning Research* pp. 5639–5650. PMLR, 2020b.
- Bonnie Li, Vincent François-Lavet, Thang Doan, and Joelle Pineau. Domain adversarial reinforcement learning. *arXiv:2102.07097* 2021.
- Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. *Proceedings of the 36th International Conference on Machine Learning (ICML 2019)* *Proceedings of Machine Learning Research*. PMLR, 2019.
- Francesco Locatello, Ben Poole, Gunnar Raetsch, Bernhard Schölkopf, Olivier Bachem, and Michael Tschannen. Weakly-supervised disentanglement without compromise. *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6348–6359. PMLR, 2020.

- Bogdan Mazoure, Remi Tachet des Combes, Thang Long Doan, Philip Bachman, and R Devon Hjelm. Deep reinforcement and infomax learning. *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, 2020.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithm. *arXiv:1707.06347*, 2017.
- Max Schwarzer, Nitarshan Rajkumar, Michael Noukhovitch, Ankesh Anand, Laurent Charlin, R Devon Hjelm, Philip Bachman, and Aaron C Courville. Pretraining representations for data-efficient reinforcement learning. *15th Conference on Neural Information Processing Systems (NeurIPS 2021)*, 2021.
- Rui Shu, Yining Chen, Abhishek Kumar, Stefano Ermon, and Ben Poole. Weakly supervised disentanglement with guarantees. *International Conference on Learning Representations*, 2020.
- Austin Stone, Oscar Ramirez, Kurt Konolige, and Rico Jonschkowski. The distracting control suite – a challenging benchmark for reinforcement learning from pixels. *arXiv:2101.02722*, 2021.
- Frederik Täubele, Elliot Creager, Niki Kilbertus, Anirudh Goyal, Francesco Locatello, Bernhard Schölkopf, and Stefan Bauer. Is independence all you need? on the generalization of representations learned from correlated data. *Proceedings of the 38th International Conference on Machine Learning volume 139 of Proceedings of Machine Learning Research*, PMLR, 2021.
- Saran Tunyasuvunakool, Alistair Muldal, Yotam Doron, Siqi Liu, Steven Bohez, Josh Merel, Tom Erez, Timothy Lillicrap, Nicolas Heess, and Yuval Tassa. *control: Software and tasks for continuous control*. *Software Impacts*, 6:100022, 2020.
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018.
- Mel Vecerik, Jean-Baptiste Regli, Oleg Sushkov, David Barker, Rugile Pevcevičute, Thomas Roth, Raia Hadsell, Lourdes Agapito, and Jonathan Scholz. S3k: Self-supervised semantic keypoints for robotic manipulation via multi-view consistency. *Conference on Robot Learning*, pp. 449–460. PMLR, 2021.
- Denis Yarats, Ilya Kostrikov, and Rob Fergus. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. *International Conference on Learning Representations*, 2021.
- Amy Zhang, Clare Lyle, Shagun Sodhani, Angelos Filos, Marta Kwiatkowska, Joelle Pineau, Yarin Gal, and Doina Precup. Invariant causal prediction for block mdp. *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*, volume 119 of *Proceedings of Machine Learning Research*, pp. 11214–11224. PMLR, 2020.
- Amy Zhang, Rowan McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning invariant representations for reinforcement learning without reconstruction. *International Conference on Learning Representations*, 2021.
- Chiyuan Zhang, Oriol Vinyals, Remi Munos, and Samy Bengio. A study on overfitting in deep reinforcement learning. *arXiv: 1804.06893*, 2018.

A EXTENDED BACKGROUND

In this section, we provide details of the RL algorithms that we use as the base algorithms for TED in our experiments. We use RAD (Laskin et al., 2020a) and SVEA (Hansen et al., 2021) for the experiments in Section 5.1. Both base algorithms are extensions of the Soft Actor-Critic (SAC) algorithm (Haarnoja et al., 2018). We use PPO (Schulman et al., 2017) as the base algorithm in Section 5.2.

SAC. Soft-Actor Critic (SAC) is an off-policy, actor-critic RL algorithm that learns a policy π to maximise the expected discounted future rewards and the entropy of the policy. SAC uses transitions from the replay buffer D to train the critic $Q: \mathcal{O} \times \mathcal{A} \rightarrow \mathbb{R}$ by minimising the loss:

$$L_Q = \mathbb{E}_{(\mathbf{o}_t, \mathbf{a}_t, \mathbf{o}_{t+1}, r_t) \sim D} \left[\left(Q(\mathbf{o}_t; \mathbf{a}_t) - r_t - \gamma V(\mathbf{o}_{t+1}) \right)^2 \right] \quad (4)$$

The target value function V is estimated by:

$$V(\mathbf{o}_{t+1}) = \mathbb{E}_{\mathbf{a}_{t+1}} \left[\min_{i=1,2} Q_i(\mathbf{o}_{t+1}; \mathbf{a}_{t+1}) \right] - \text{SAC} \log \pi(\mathbf{a}_{t+1} | \mathbf{o}_{t+1}) \quad (5)$$

where Q is the target Q network whose parameters are an exponential moving average of the corresponding Q network parameters. We maintain two Q networks, Q_1 and Q_2 , and use the minimum of the two networks for the updates. The actor π is trained by minimising the loss:

$$L = \mathbb{E}_{\mathbf{o}_t \sim D} \mathbb{E}_{\mathbf{a}_t \sim \pi} \left[\text{SAC} \log \pi(\mathbf{a}_t | \mathbf{o}_t) - \min_{i=1,2} Q_i(\mathbf{o}_t; \mathbf{a}_t) \right] \quad (6)$$

RAD. RAD adds data augmentations to the observations before the SAC network updates. We use image padding and random crop augmentations for the observations in each transition sampled from the replay buffer $(\mathbf{o}_t; \mathbf{o}_{t+1}) \sim D$.

SVEA. SVEA stabilises Q-learning using a combination of augmented and unaugmented images with the updated loss:

$$L_Q^{\text{SVEA}} = \text{SVEA} L_Q(o_t; a_t; o_{t+1}) + \text{SVEA} L_Q(o_t^{\text{aug}}; a_t; o_{t+1}) \quad (7)$$

The policy π is trained using only unaugmented images, using the standard loss in Equation 6.

PPO. Proximal Policy Optimisation (PPO) is an on-policy, actor-critic RL algorithm for continuous or discrete action spaces. We use a discrete policy π for the Progen environments. PPO learns a policy by minimising a clipped loss over minibatches transitions:

$$L^{\text{CLIP}} = \mathbb{E}_{(\mathbf{o}_t, \mathbf{a}_t, \mathbf{o}_{t+1}, r_t)} \left[\min(\lambda A_t; \text{clip}(\lambda \frac{\pi(\mathbf{a}_t | \mathbf{o}_t)}{\pi(\mathbf{a}_t | \mathbf{o}_t)}; 1 - \epsilon, 1 + \epsilon)) A_t \right] \quad (8)$$

where λ is the ratio of the action probability under the new and old policies, and A_t is the action advantage given by: $A_t = Q(\mathbf{o}_t; \mathbf{a}_t) - V(\mathbf{o}_t)$.

B IMPLEMENTATION DETAILS

In this section, we provide the implementation details for TED. Our codebase is built on top of the publicly released DrQ PyTorch implementation by Yarats et al. (2021). We adapt the codebase to implement the base RL algorithms as well as the TED auxiliary task. A public and open-source implementation of TED can be found in the supplementary material and at github.com/usc-repo [currently anonymised for double blind review].

B.1 RAD AND SVEA IMPLEMENTATION DETAILS

Encoder architecture. We use the same encoder architecture as Yarats et al. (2021). The encoder weights are shared between the actor π and critic Q . For the RAD experiments, the encoder consists of 4 convolutional layers following the original RAD implementation (Laskin et al., 2020a). For SVEA experiments, we use 11 convolution layers following the original SVEA paper (Hansen et al.,

2021). Baselines follow the same encoder size as the base RL algorithm for TED in each experiment. Each convolutional layer has a 3×3 kernel size and 32 channels. The first layer has a stride of 2, all other layers have a stride of 1. There is a ReLU activation between each convolutional layer. The convolutional layers are followed by a trunk network with a linear layer, layer normalisation, and finally a tanh activation.

Both RAD and SVEA use the critic loss L_Q to update the encoder parameters. So our implementation of TED requires both the critic loss L_Q and TED loss L_{TED} to update the encoder together. In practice, this can be done by adding the two losses $L_{ENC} = L_Q + L_{TED}$ and backpropagating the encoder loss L_{ENC} to update the encoder parameters.

Actor and critic architecture. The actor and critic Q networks are both 2-layer MLPs with a hidden dimension of 1024. We apply ReLU activations after each layer except the last layer.

TED architecture. The TED classifier is implemented with the following parameters: k_1 , k_2 , k , b , and \mathbf{b} are vectors of length 50; and c is a scalar. The output of the classifier is defined in Equation 1.

Hyperparameters. Table 2 shows the value of the TED loss coefficient for each task. All other hyperparameters are provided in Table 3. Unless specified otherwise, we use the same hyperparameter settings for both RAD-TED and SVEA-TED algorithms.

Environment	Base algorithm	Value of
cartpole_swingup	RAD	100
cartpole_swingup	SVEA	0.5
finger_spin	RAD	25
walker_walk	SVEA	1
panda_reach	RAD	200
panda_reach	SVEA	0.5

Table 2: TED loss coefficient

Hyperparameter name	Value
Replay buffer capacity	100000
Initial steps	1000
Stacked frames	3
Action repeat	2 for finger_spin, 4 otherwise
Batch size	128
Discount factor	0.99
Optimizer	Adam
Learning rate (actor, critic and encoder)	1e-3
Target soft-update rate	0.01
Actor update frequency	2
Actor log stddev bounds	[10; 2]
Latent representation dimension	50 for RAD, 100 for SVEA
Image size	(84, 84)
Image pad	4
Initial temperature	0.1

Table 3: Hyperparameter values for both RAD-TED and SVEA-TED.

Disentanglement metric. To calculate the disentanglement metric described in Section 5.1, we collect a batch of 10,000 samples. To create each sample, we use 32 pairs of observations with the same fixed factor to evaluate Equation 3 with $B = 32$. The batch of samples is split into 8,000 training samples to train the classifier, and 2,000 testing samples to calculate the accuracy of the classifier as the disentanglement score. We use a Scikit-learn logistic regression classifier with L1 regularisation and the saga solver.

B.2 PPO IMPLEMENTATION DETAILS

We follow the PPO architecture and hyperparameters used in the Procgen benchmark (Cobbe et al., 2020). The encoder uses the IMPALA (Espenholt et al., 2018) architecture. PPO is augmented with TED by adding the loss terms $\mathcal{L} = \mathcal{L}_{\text{PPO}} + \mathcal{L}_{\text{TED}}$ and backpropagating the total loss \mathcal{L} with a shared optimiser. The hyperparameters are shown in Table 4.

Hyperparameter name	Value
Image size	(64, 64)
Discount factor	0.999
GAE	0.95
# Timesteps per rollout	250
Epochs per rollout	3
# Minibatches per rollout	8
Entropy bonus	0.01
PPO clip range	0.2
Learning rate	5e-4
# Workers	1
# Environments per worker	64
LSTM?	No
Frame stack?	No
TED coefficient	1 for coinrun, 0.25 for jumper

Table 4: Hyperparameter values for PPO-TED

C ENVIRONMENT IMAGES

In Figure 6, we provide images of example observations for each of the DMC and Panda Gym training and testing environments used in our experiments to visualise the generalisation challenge. A video is also provided in the supplementary materials.

In Figure 7, we provide images of example observations for each of the Procgen environments used in our experiments.

