

Disentangled (Un)Controllable Features

Jacob E. Kooi
Vrije Universiteit Amsterdam
Amsterdam, Netherlands
j.e.kooi@vu.nl

Mark Hoogendoorn
Vrije Universiteit Amsterdam
Amsterdam, Netherlands
m.hoogendoorn@vu.nl

Vincent François-Lavet
Vrije Universiteit Amsterdam
Amsterdam, Netherlands
vincent.francoislavet@vu.nl

Abstract—In the context of MDPs with high-dimensional states, downstream tasks are predominantly applied on a compressed, low-dimensional representation of the original input space. A variety of learning objectives have therefore been used to attain useful representations. However, these representations usually lack interpretability of the different features. We present a novel approach that is able to disentangle latent features into a controllable and an uncontrollable partition. We illustrate that the resulting partitioned representations are easily interpretable on three types of environments and show that, in a distribution of procedurally generated maze environments, it is feasible to interpretably employ a planning algorithm in the isolated controllable latent partition.

Index Terms—representation learning, interpretability, reinforcement learning

I. INTRODUCTION

Learning from high-dimensional data remains a challenging task. Particularly for reinforcement learning (RL) [1], the complexity and high dimensionality of the Markov Decision Process (MDP) [2] states often leads to complex or intractable solutions. In order to facilitate learning from high-dimensional input data, an encoder architecture can be used to compress the inputs into a lower-dimensional latent representation. To this extent, a plethora of work has successfully focused on discovering a compressed encoded representation that accommodates the underlying features for the task at hand [3]–[5].

The resulting low-dimensional representations however seldom contain specific disentangled features, which leads to disorganized latent information. This means that the individual latent states can represent the information from the state in any arbitrary way. The result is a representation with poor interpretability, as the latent states cannot be connected to certain attributes of the original observation space (e.g, the x - y coordinates of the agent). Prior work in structuring a latent representation has shown notions and use of interpretability in MDP representations [6]. When expanding this notion of interpretability to be compatible with RL, it has been argued that the controllable features should be an important element of a latent representation, since it generally represents what is directly influenced by the policy. In this light, [7] have introduced the concept of isolating and disentangling controllable features in a low-dimensional maze environment, by means of a selectivity loss. Furthermore, [8] provides an object-centric approach to isolate distinct objects in MDPs and [9] shows theoretical foundations for this isolation in a weakly-supervised controllable setting. Controllable features however

only represent a fragment of an environment, where in many cases the uncontrollable features are of equal importance. For example, in the context of a distribution of mazes, for the prediction of the next controllable (agent) state following an action, the information about the wall structure is crucial (see Fig. 1). We therefore hypothesize that a thorough representation should incorporate controllable and uncontrollable features, ideally in a disentangled, interpretable arrangement; Interpretability is crucial for future real-world deployment [10], while an additional benefit would be that the separation of the controllable and uncontrollable features can be exploited in downstream algorithms such as planning.

Our contribution consists of an algorithm that, showcased in three different MDP settings, explicitly disentangles the latent representation into a controllable and an uncontrollable latent partition. This is highlighted on three types of environments, each with a varying class of controllable and uncontrollable elements. This allows for a precise and visible separation of the latent features, improving interpretability, representation quality and possibly moving towards a basis for building causal relationships between an agent and its environment. The unsupervised learning algorithm consists of both an action-conditioned and a state-only forward predictor, along with a contrastive and an adversarial loss, which isolate and disentangle the controllable versus the non-controllable features. Furthermore, we show an application of learning and planning on the human-interpretable disentangled latent representation, where the properties of disentanglement allow the planning algorithm to operate solely in the controllable partition of the latent representation.

II. RELATED WORK

1) *General Representation Learning*: Many works have focused on converting high-dimensional inputs to a compact, abstract latent representation [11]. Learning this representation can make use of auxiliary, unsupervised tasks in addition to the pure RL objectives [4]. One way to ensure a meaningful latent space is to implement architectures that require a pixel reconstruction loss such as a variational [12] or a deterministic [13] autoencoder. Other approaches combined reward reconstruction with latent prediction [14], pixel reconstruction with planning [15] or used latent predictive losses without pixel reconstruction [5].

2) *Representing controllable features*: In representation learning for RL, a focus on controllable features can be

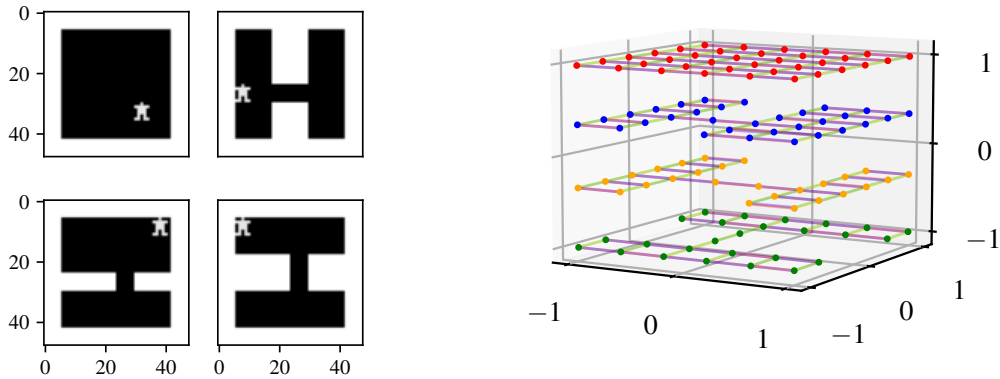


Fig. 1: Visualization in a maze environment of four random pixel observations $s \in \mathbb{R}^{48 \times 48}$ (left) and the encoded observations $z = f(s; \theta_{enc}) \forall s \in \mathcal{S}$ (right). On the right, we can see the disentanglement of the controllable latent $z^c \in \mathbb{R}^2$ on the horizontal axes, and the uncontrollable latent $z^u \in \mathbb{R}^1$ on the vertical axis. The encoder is trained on high-dimensional tuples (s_t, a_t, r_t, s_{t+1}) , sampled from a replay buffer \mathcal{B} , gathered from random trajectories in the four maze environments shown on the left. All possible states in all four mazes are encoded and plotted with the transition prediction for each possible action, revealing a clear disentanglement between the controllable latents (agent x-y position) and the uncontrollable latent (wall architecture). Note that all samples are taken from the same buffer, filled with samples from all four mazes.

beneficial as these features are strongly influenced by the policy [7]. This can be done using generative methods [16], but is most commonly pursued using an auxiliary inverse-prediction loss; predicting the action that was taken in the MDP [3]. The work in [17], [18] builds a latent representation with an emphasis on the controllable features of an environment with inverse-prediction losses, and uses these features to guide exploratory behavior. Furthermore, [19] employs multi-step inverse prediction to successfully encompass controllable features in their representation. However, these works have not retained the uncontrollable features in their representation, which is a key aspect in our work.

3) *Partitioning a latent representation*: Sharing similarity in terms of the separation of the latent representation, [20] disentangles the latent representation in the domain adaptation setting into a task-relevant and a context partition, by means of adversarial predictions with gradient reversals and cyclic reconstruction. In [21], a reconstruction-based adversarial architecture is used that divides a latent representation into reward-relevant and irrelevant features. Related work in [22] attempts to divide the latent representation of the model-based RL algorithm Dreamer [15], using action-conditioned and state-only forward predictors, into controllable, uncontrollable and their respective reward relevant and irrelevant features. As compared to [22], we focus on reward-free unsupervised representation learning and empirically prove the separation in a low-dimensional, structured setting.

4) *Interpretable representations in MDPs*: More closely related to our research is [7], which connects individual latent dimensions to independently controllable states in a maze using a reconstruction loss and a selectivity loss. The work in [6] visualizes the representation of an agent and its transitions in a maze environment, but does not disentangle the agent state in its controllable and uncontrollable parts,

which limits the interpretability analysis and does not allow simplifications during planning. Furthermore, [8] shows an object-oriented approach to isolate different (controllable) features, using graph neural networks (GNN's) and a contrastive forward prediction loss, but does not discriminate between controllable and uncontrollable features. Continued work in this direction focuses on theoretical foundations for an encoder to structurally represent a distinct controllable object [9]. We aim to progress the aforementioned lines of research by using a representation learning architecture that disentangles an MDP's latent representation into interpretable, disentangled controllable and uncontrollable features. Finally, we show that having separate partitions of controllable and uncontrollable features can be exploited in a planning algorithm.

III. PRELIMINARIES

We consider an agent acting within an environment, where the environment is modeled as a discrete Markov Decision Process (MDP) defined as a tuple $(\mathcal{S}, \mathcal{A}, T, R, \gamma)$. Here, \mathcal{S} is the state space, \mathcal{A} is the action space, $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is the environment's transition function, $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$ is the environment's reward mapping and γ is the discount factor. We consider the setting where we have access to a replay buffer (\mathcal{B}) of visited states $s_t \in \mathcal{S}$ that were followed by actions $a_t \in \mathcal{A}$ and resulted in the rewards $r_t \in \mathcal{R}$ and the next states s_{t+1} . One entry in \mathcal{B} contains a tuple of past experience (s_t, a_t, r_t, s_{t+1}) . The agent's goal is to learn a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes the expectation of the discounted return $V^\pi(s) = \mathbb{E}_\tau[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) \mid s_t = s]$, where τ is a trajectory following the policy π .

Furthermore, we examine the setting where a high-dimensional state ($s_t \in \mathbb{R}^v$) is compressed into a lower-dimensional latent state $z_t \in \mathcal{Z} = \mathbb{R}^w$ where \mathcal{Z} represents the latent space with $w \leq v$. This is done by means of a

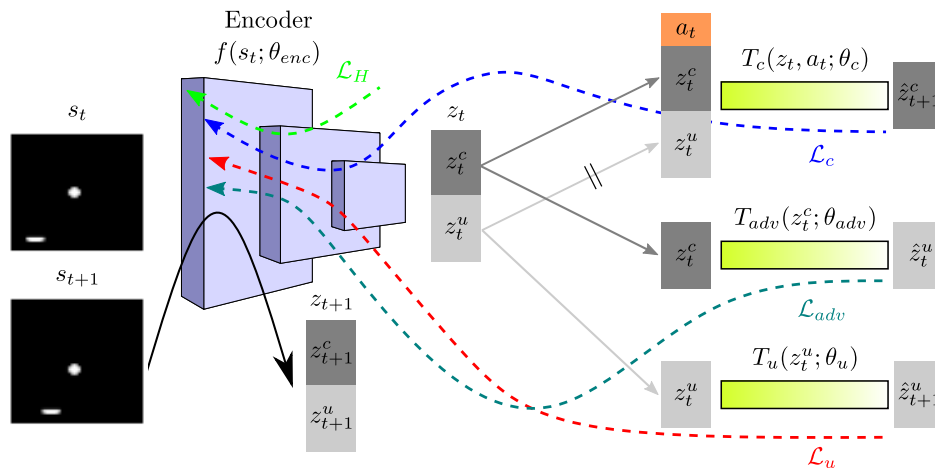


Fig. 2: Overview of the disentangling architecture, with dashed lines representing gradient propagation and green rectangles representing parameterized prediction functions. An observation s_t is encoded into a latent representation consisting of two parts; z_t^c and z_t^u , which represent controllable and uncontrollable features respectively. These separated representations are then independently used to make action-conditioned, state-only and adversarial predictions in order to provide gradients to the encoder that disentangle the latent representation z_t into controllable (z_t^c) and uncontrollable (z_t^u) partitions.

neural network encoding $f : \mathcal{S} \rightarrow \mathcal{Z}$ where f represents the encoder.

IV. ALGORITHM

We aim for an interpretable and disentangled representation of the controllable and uncontrollable latent features. We define controllable features as the characteristics of the MDP that are predominantly affected by any action $a \in \mathcal{A}$, such as the position of the agent in the context of a maze environment. The uncontrollable features are those attributes that are not or only marginally affected by the actions. We show that the proposed disentanglement is possible by designing losses and gradient propagation through two separate parts of the latent representation. Specifically, to assign controllable information to the controllable latent partition, the gradient from an action-conditioned forward predictor is propagated through it. To assign uncontrollable information to the uncontrollable latent partition, the gradient from a state-only forward predictor is propagated through it. The remaining details will be provided in the rest of this Section.

We consider environments with high-dimensional states, represented as pixel inputs. These pixel inputs are subsequently encoded into a latent representation $z_t = (z_t^c, z_t^u) \in \mathcal{Z} \in \mathbb{R}^{n_c} + \mathbb{R}^{n_u}$, with the superscripts c and u representing the controllable and uncontrollable features, and the superscripts n_c and n_u representing their respective dimensions. The compression into a latent representation $\mathcal{S} \rightarrow \mathcal{Z}$ is done by means of a convolutional encoder, parameterized by a set of learnable parameters θ_{enc} according to:

$$z_t = (z_t^c, z_t^u) = f(s_t; \theta_{enc}). \quad (1)$$

An overview of the proposed algorithm is illustrated in Fig. 2 and the details are provided hereafter. In this section, all losses

and transitions are given under the assumption of a continuous abstract representation and a deterministic transition function. The algorithm could be adapted by replacing the losses related to the internal transitions with generative approaches (in the context of continuous and stochastic transitions) or a log-likelihood loss (in the context of stochastic but discrete representations).

A. Controllable Features

To isolate controllable features in the latent representation, z_t^c is used to make an action-conditioned forward prediction in latent space. In the context of a continuous latent space and deterministic transitions, z_t^c is updated using a mean squared error (MSE) forward prediction loss $\mathcal{L}_c = |\hat{z}_{t+1}^c - z_{t+1}^c|^2$, where \hat{z}_{t+1}^c is the action-conditioned residual forward prediction of the parameterized function $T_c(z, a; \theta_c) : \mathcal{Z} \times \mathcal{A} \rightarrow \mathcal{Z}$:

$$\hat{z}_{t+1}^c = T_c(z_t, a_t; \theta_c) + z_t^c \quad (2)$$

and the prediction target z_{t+1}^c is part of the encoder output $f(s_{t+1}; \theta_{enc})$. Note that the full latent state z_t is necessary in order to predict \hat{z}_{t+1}^c (e.g. the uncontrollable features could represent a wall or other static structure that is necessary for the prediction of the controllable features). Furthermore, the uncontrollable latent partition input z_t^u is accompanied by a stop gradient to discourage the presence of controllable features in z_t^u . When minimizing \mathcal{L}_c , both the encoder (θ_{enc}) as well as the predictor (θ_c) are updated, which allows shaping the representation z_t^c as well as learning the internal dynamics.

B. Uncontrollable Features

To express uncontrollable features in the latent space, z_t^u is used to make a state-only (not conditioned on the action

a_t) forward prediction in latent space. This enforces uncontrollable features within the uncontrollable latent partition z^u , since features that are action-dependent cannot be accurately predicted with the preceding state only. Following a residual prediction, z^u is then updated using a MSE forward prediction loss $\mathcal{L}_u = |\hat{z}_{t+1}^u - z_{t+1}^u|^2$, with \hat{z}_{t+1}^u defined as:

$$\hat{z}_{t+1}^u = T_u(z_t^u; \theta_u) + z_t^u \quad (3)$$

and $T_u(z^u; \theta_u) : \mathcal{Z} \rightarrow \mathcal{Z}$ representing the parameterized prediction function. The target z_{t+1}^u is part of the output of the encoder $f(s_{t+1}; \theta_{enc})$. When minimizing \mathcal{L}_u , both θ_{enc} and θ_u are updated. In this way the loss \mathcal{L}_u drives the latent representation z^u , which is conditioned on θ_{enc} according to $(z_t^c, z_t^u) = f(s_t; \theta_{enc})$, to only represent the features of s_t that are not conditioned on the action a_t .

C. Avoiding Predictive Representation Collapse

Minimizing a forward prediction loss in latent space \mathcal{Z} is prone to collapse [6], [14], due to the convergence of \mathcal{L}_c and \mathcal{L}_u when $f(s_t; \theta_{enc})$ is a constant $\forall s_t \in \mathcal{S}$. To avoid representation collapse when using forward predictors, a contrastive loss is used to enforce sufficient diversity in the latent representation:

$$\mathcal{L}_{H_1} = \exp(-C_d \|z_t - \bar{z}_t\|_2) \quad (4)$$

where C_d represents a constant hyperparameter and \bar{z}_t is a ‘negative’ batch of latent states z_t , which is obtained by shifting each position of latent states in the batch by a random number between 0 and the batch size. In the random maze environment, an additional contrastive loss is added to further diversify the controllable representation:

$$\mathcal{L}_{H_2} = \exp(-C_d \|z_t^c - \bar{z}_t^c\|_2) \quad (5)$$

where z_t^c is obtained from randomly sampled trajectories. This additional regularizer proved necessary to avoid collapse of z^c when moving to a near infinite number of possible mazes. The resulting contrastive loss \mathcal{L}_H for the random maze environment then consists of $0.5\mathcal{L}_{H_1} + 0.5\mathcal{L}_{H_2}$. The total loss used to update the encoder’s parameters now consists of $\mathcal{L}_{enc} = \mathcal{L}_c + \mathcal{L}_u + \mathcal{L}_H$.

D. Guiding Feature Disentanglement with Adversarial Loss

When using a controllable latent space $z^c \in \mathbb{R}^x, x \in \mathbb{N}$, where $x > g$, with g representing the number of dimensions needed to portray the controllable features, some information about the uncontrollable features in the controllable latent representation might be present. This is due to the non-enforcing nature of \mathcal{L}_c , as the uncontrollable features are equally predictable with or without the action. To ensure that no information about the uncontrollable features is kept in the controllable latent representation, an adversarial component is added to the architecture in Fig. 2. This is done by updating the encoder with an adversarial loss \mathcal{L}_{adv} and reversing the gradient [23]. The adversarial loss is defined as

$$\mathcal{L}_{adv} = |\hat{z}_t^u - z_t^u|^2, \quad (6)$$

with $\hat{z}_t^u = T_{adv}(z_t^c; \theta_{adv})$, where \hat{z}_t^u is the uncontrollable prediction of the parameterized function $T_{adv}(z^c; \theta_{adv}) : \mathcal{Z} \rightarrow \mathcal{Z}$ and z_t^u is the target. Intuitively, since the parameters of $T_{adv}(z^c; \theta_{adv})$ are being updated with \mathcal{L}_{adv} and the parameters of $f(s; \theta_{enc})$ are being updated with $-\mathcal{L}_{adv}$, the prediction function can be seen as the discriminator and the encoder can be seen as the generator [24]. The discriminator tries to give an accurate prediction of the uncontrollable latent z^u given the controllable latent z^c , while the generator tries to counteract the discriminator by removing any uncontrollable features from the controllable representation. In our case, the predictor is a multi-layer perceptron (MLP), which means that minimizing \mathcal{L}_{adv} enforces that no nonlinear relation between z^c and z^u can be learned. We hypothesize that this is a deterministic approximation of minimizing the Mutual Information (MI) between z^u and z^c . When using the adversarial loss, the combined loss propagating through the encoder consists of $\mathcal{L}_{enc} = \mathcal{L}_c + \mathcal{L}_u + \mathcal{L}_H - \mathcal{L}_{adv}$. Here the minus term in $-\mathcal{L}_{adv}$ represents a gradient reversal to the encoder. Note that the losses are not scaled, as this did not prove to be necessary for the experiments conducted.

Algorithm 1 Interpretable (Un)Controllable Features

- 1: Initialize $\theta_{enc}, \theta_c, \theta_u, \theta_{adv}$
 - 2: **for** $iteration = 1, 2, \dots, N$ **do**
 - 3: Sample batch of tuples $\{s_t, a_t, s_{t+1}\}$
 - 4: Encode observations: $f(s; \theta_{enc}) = \{z^c, z^u\}$
 - 5: Predict $\hat{z}_{t+1}^c = T_c(z_t^c, z_t^u, a; \theta_c) + z_t^c$ // detach z_t^u
 - 6: Predict $\hat{z}_{t+1}^u = T_u(z_t^u; \theta_u) + z_t^u$
 - 7: Predict $\hat{z}_t^u = T_{adv}(z_t^c; \theta_{adv})$
 - 8: Compute losses $\mathcal{L}_c, \mathcal{L}_u, -\mathcal{L}_{adv}, \mathcal{L}_H$
 - 9: Update parameters $\theta_{enc}, \theta_c, \theta_u, \theta_{adv}$
 - 10: **end for**
-

E. Downstream Tasks

By disentangling a latent representation in a controllable and an uncontrollable part, one can more readily obtain human-interpretable features. While interpretability is generally an important aspect, it is also important to test how a notion of human interpretability affects downstream performance, as it is generally desired to strike a good balance between interpretability and performance. This is examined by training an RL agent on the learned and subsequently frozen latent representation. The action a_t is chosen following an ϵ -greedy policy, where a random action is taken with a probability ϵ , and with $(1 - \epsilon)$ probability the policy $\pi(z) = \arg \max_{a \in \mathcal{A}} Q(z, a; \theta)$ is evaluated, where $Q(z, a; \theta)$ is the Q-network trained by Deep Double Q-Learning (DDQN) [25], [26]. The Q-network is trained with respect to a target Y_t :

$$Y_t = r_t + \gamma Q(z_{t+1}, \arg \max_{a \in \mathcal{A}} Q(z_{t+1}, a; \theta); \theta^-). \quad (7)$$

With γ representing the discount factor and θ^- the target Q-network’s parameters. The target Q-network’s parameters are updated as an exponential moving average of the original

parameters θ according to: $\theta_{k+1}^- = (1-\tau)\theta_k^- + \tau\theta_k$, where subscript k represents a training iteration and τ represents a hyperparameter controlling the speed of the parameter update. The resulting DDQN loss is defined as $\mathcal{L}_Q = |Y_t - Q(z_t, a; \theta)|^2$. The full computation of all losses is shown in pseudocode in Algorithm 1.

V. EXPERIMENTS

In this section, we showcase the disentanglement of controllable and uncontrollable features on three different environments, the complexity of which is in line with prior work on structured representations [6]–[9], [27]: (i) a quadruple maze environment, (ii) the catcher environment and (iii) a randomly generated maze environment. The first environment yields a state space of 119 different observations, and is used to showcase the algorithm’s ability to disentangle a low-dimensional latent representation. The catcher environment examines a setting where the uncontrollable features are not static, and the random maze environment is used to showcase disentanglement in a more complex distribution of over 25 million possible environments, followed by the application of downstream tasks by applying reinforcement learning (DDQN) and a latent planning algorithm running in the controllable latent partition. The base of the encoder is derived from [28] and consists of two convolutional layers, followed by a fully connected layer for low-dimensional latent representations or an additional CNN for a higher-dimensional latent representation such as a feature map. In all environments, the encoder $f(s; \theta_{enc})$ is trained from a buffer \mathcal{B} filled with transition tuples (s_t, a_t, r_t, s_{t+1}) from random trajectories. Note that, in interpretability, there is generally not a specific metric to optimize for. In order to produce interpretable representations, finding the right hyperparameters required manual (human) inspection of the plotted latent representations.

A. Quadruple Maze Environment

The maze environment consists of an agent and a selection of four distinct, handpicked wall architectures. The environment’s state is provided as pixel observations $s_t \in \mathbb{R}^{48 \times 48}$, where an action moves the agent by 6 pixels in each direction (up, down, left, right) except if this direction is obstructed by a wall. We consider the context where there is no reward ($r_t = 0 \forall (s_t, a_t) \in (\mathcal{S}, \mathcal{A})$) and there is no terminal state.

We select a two-dimensional controllable representation ($z^c \in \mathbb{R}^2$) and a one-dimensional uncontrollable representation ($z^u \in \mathbb{R}^1$). The experiments are conducted using a buffer \mathcal{B} filled with random trajectories from the four different basic maze architectures. The encoder’s parameters are updated using \mathcal{L}_{enc} in Section IV-C with $\mathcal{L}_H = \mathcal{L}_{H_1}$. After 50k training iterations, a clear disentanglement between the controllable (z^c) and uncontrollable (z^u) latent representation can be seen in Fig. 1. One can observe that the encoder is updated so that the one-dimensional latent representation z^u learns different values that define the type of wall architecture.

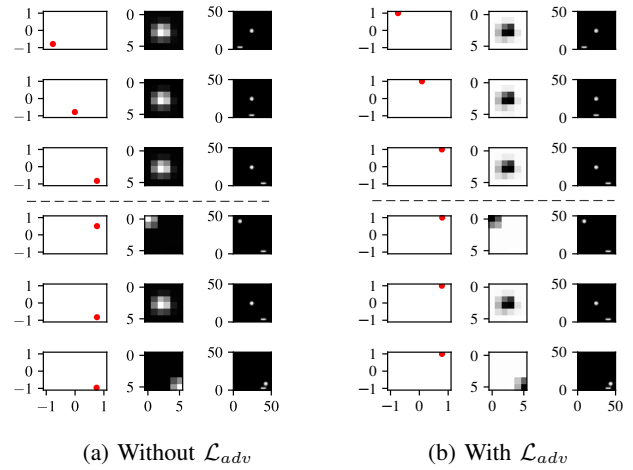


Fig. 3: Visualization of the latent feature disentanglement in the catcher environment after 200k training iterations, with $z_t = f(s_t; \theta_{enc}) \in \mathbb{R}^2 + \mathbb{R}^{6 \times 6}$. In (a) and (b), the left column shows z_t^c , the middle column is a feature map representing z_t^u and the right column is the pixel state s_t . The dashed lines separate observations where the ball position or the paddle position is kept fixed for illustration purposes. z^c tracks the agent position while z^u tracks the falling ball. In b), note that even when having a two-dimensional controllable state, the adversarial loss in b) makes sure that distinct ball positions have a negligible effect on z^c (left column), even when the high-level features of the agent and the ball might be hard to distinguish.

B. Catcher Environment

As opposed to the maze environment, the catcher environment encompasses uncontrollable features that are non-stationary. The ball is dropped randomly at the top of the environment and is falling irrespective of the actions, while the paddle position is directly modified by the actions. The environment’s states are defined as pixel observations s_t of size $\mathbb{R}^{51 \times 51}$. At each time step, the paddle moves left or right by 3 pixels. Since we are only doing unsupervised learning, we consider the context where there is no reward ($r_t = 0 \forall (s_t, a_t) \in (\mathcal{S}, \mathcal{A})$) and an episode ends whenever the ball reaches the paddle or the bottom.

We take $z^c \in \mathbb{R}^2$ and $z^u \in \mathbb{R}^{6 \times 6}$. To test disentanglement, z^c is of a higher dimension than needed since the paddle (agent) only moves on the x-axis and would therefore require only one feature. To show disentanglement, the redundant dimension of z^c should not or negligibly have information about z^u . The encoder’s parameters are updated according to Algorithm 1 with $\mathcal{L}_H = \mathcal{L}_{H_1}$. After training the encoder for 200k iterations, a selection of state observations s_t and their encoding into the latent representation $z = (z^c, z^u)$ can be seen in Fig. 3. A clear distinction between the ball and paddle representations can be observed, with the former residing in z^u and the latter in z^c .

C. Random Maze Environment

The random maze environment is similar to the maze environment from Section V-A, but consists of a large distribution of randomly generated mazes with complex wall structures. The environment’s state is provided as pixel observations $s_t \in \mathbb{R}^{48 \times 48}$, where an action moves the agent by 6 pixels in each direction. We consider $z^c \in \mathbb{R}^2$ and $z^u \in \mathbb{R}^{6 \times 6}$. This environment tests the generalization properties of a disentangled latent representation, as there are over 25 million possible maze architectures, corresponding to a probability of less than $4 \cdot 10^{-8}$ to sample the same maze twice. Note that because z^c is 2-dimensional, results with and without adversarial loss are in practice extremely close. After 50k training iterations, the latent representation $z = (z^c, z^u)$ shows an interpretable disentanglement between the controllable and the uncontrollable features (see Fig. 4a). A clear distinction between the agent and the wall structure can be found inside z^c and z^u . Note that Instead of using a single dimension to ‘describe’ the uncontrollable features z^u (see Fig. 1), using a feature map for z^u allows training an encoding that provides a more interpretable representation of the actual wall architecture.

1) *Reinforcement Learning*: In order to verify whether a human-interpretable disentangled latent encoding is informative enough for downstream tasks, we formalize the random maze environment into an MDP with rewards. The agent acquires a reward r_t of -0.1 at every time step, except when it finds the key in the top right part in which case it acquires a positive reward of 1. The episode ends whenever a positive reward is obtained or a total of 50 environment steps have been taken. For each new episode, a random wall structure is generated, and the agent starts over in the bottom left section of the maze (see Fig. 5). To see whether an interpretable disentangled latent representation is useful for RL, we compare different scenarios of (pre)training; (i) An encoder pretrained for 50k iterations to attain the representation in Fig. 4a and subsequently trained with DDQN for 500k iterations (ii) an encoder identical to the aforementioned but trained with DDQN and a planning algorithm (iii) an encoder pretrained for 50k iterations with \mathcal{L}_{inv} instead of \mathcal{L}_c and subsequently trained with DDQN for 500k iterations (iv) an encoder purely trained with DDQN gradients for 500k iterations. The resulting performances are compared in Fig. 5. We find that a disentangled structured representation is suitable for downstream tasks, as it achieves comparable performance to training an encoder end-to-end with DDQN for 500k iterations. Although performance is similar, Fig. 4b shows that an encoder updated solely with the DDQN gradient can lose any form of interpretability. Moreover, we show in Fig. 5 that a representation trained with an inverse prediction loss instead of a state-action forward prediction loss leads to poor downstream performance in the random maze environment.

2) *Planning*: As seen in Fig. 4a, after pre-training with the unsupervised losses, an interpretable disentangled representation with the corresponding agent transitions is obtained. Due

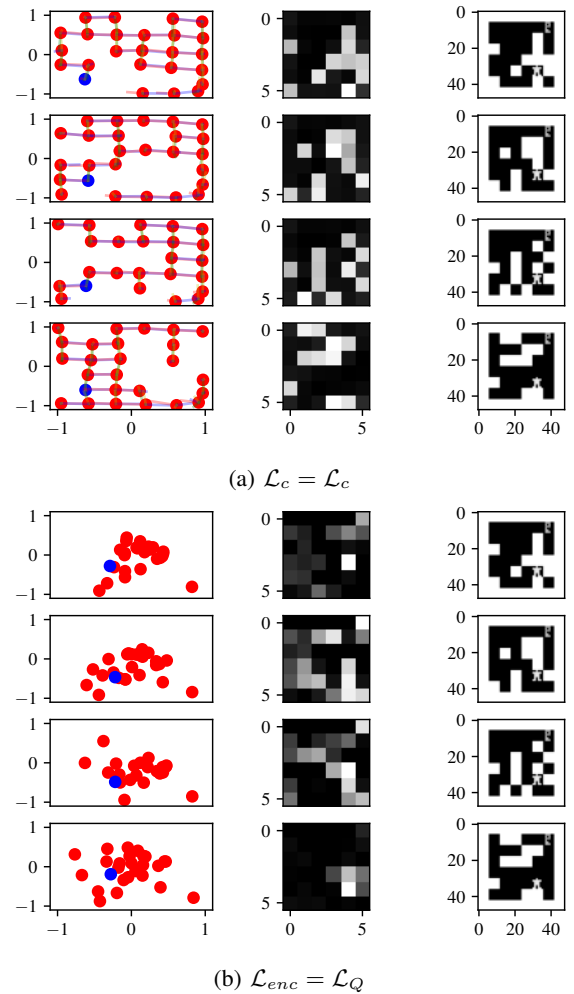


Fig. 4: A plot of the latent representation for all observations in a single randomly sampled maze when training with the aforementioned losses (a), and when end-to-end updating the encoder with only the Q-loss \mathcal{L}_Q from DDQN for 500k iterations (c). The left column shows the controllable latent $z_t^c \in \mathbb{R}^2$ with the current state in blue, the remaining states in red, and the predicted movement due to actions as different colored bars for each individual action. The middle column shows the uncontrollable latent $z_t^u \in \mathbb{R}^{6 \times 6}$ and the right column shows the original state $s_t \in \mathbb{R}^{48 \times 48}$. Evidently, the representation in (b) seems to have very little structure at all, showing that a representation that is optimized without prior structural incentives will often represent a black box.

to this disentanglement of the controllable and uncontrollable features, we can employ prior knowledge that the uncontrollable features are static, and employ latent planning in the controllable latent space only (see Fig. 6). The planning algorithm used is derived from [29], and is used to successfully plan only in the controllable partition of the latent representation z^c , while freezing the input for z^u regardless of planning depth. It can be observed that even when planning with a relatively small depth of 3, we achieve better performance than the pre-




Fig. 5: Performance of different (pre)trained representations on the random maze environment, measured as a mean (full line) and standard error (shaded area) over 5 seeds. The ‘Interpretable’ setting uses an encoder pre-trained with 50k iterations to acquire a representation as in Fig. 4a, after which the encoder is frozen and a Q-network is trained on top with DDQN for 500k iterations. The ‘Interpretable + Planning’ curve is similar to the ‘Interpretable’ setting but uses DDQN with a planning algorithm in the controllable partition of the latent space with a depth of 3. The ‘DDQN’ setting uses an encoder trained end-to-end with only DDQN for 500k iterations and the ‘Inverse Prediction’ setting is equal to the ‘Interpretable’ setting but has an encoder pre-trained with an ‘inverse’ action prediction loss instead of \mathcal{L}_c .

trained representation with an ϵ -greedy policy and than the purely DDQN-updated encoder.

VI. LIMITATIONS

While the work presented here provides a step towards a better understanding of disentangling controllable and uncontrollable features within an encoder architecture, there remain some limitations that we must acknowledge, and which can provide a basis for future research.

First, our method’s effectiveness was predominantly demonstrated on environments with relatively simple underlying dynamics. In these environments, the disentanglement process was easier to achieve due to the limited complexity of internal dynamics present. As we begin to transfer our approach to more complex environments characterized by more extensive internal dynamics, there can arise two problems; The first being that the separation of controllable from uncontrollable features may not be as clear-cut in more complex MDPs, but can be more on a spectrum, complicating the fundamental differences between a state-only and a state-action forward predictor. The second being that interpretability will be harder to enforce when there are a large number of underlying factors of variation. As distinct seeds can give different orderings and signs of the neurons in the final layer of the encoder, identifying a factor of variation can become exponentially harder for more complex environments.

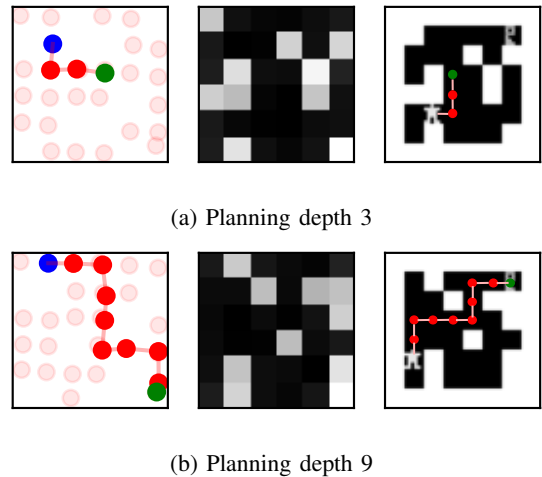


Fig. 6: Visualization of the latent representation through an actual planning iteration utilizing a planning depth of 3 (a) and a planning depth of 9 (b), with the controllable representation $z^c \in \mathbb{R}^2$ (left), the uncontrollable representation $z^u \in \mathbb{R}^{6 \times 6}$ (middle) that is kept static throughout planning depth and the original pixel input $s_t \in \mathbb{R}^{48 \times 48}$ (right). The translucent red dots represent every possible encoded state in the random maze environment, the full blue dot represents the current encoded state, the red dots represent intermediate encoded states estimated by planning and the green dot represents the final predicted state as chosen by the planning algorithm, consistent with its depth.

Lastly, while our work showed that an action-conditioned forward predictor could be preferred over an inverse predictor in some environments for isolating controllable features, it may not hold for all scenarios. The inherent properties of different environments might show a necessity of using different predictors. Consequently, there could very well be MDPs where our current approach might not provide the same level of disentanglement showed in the MDPs used in this paper.

Despite these limitations, we believe our work provides a strong foundation upon which future research can build and further extend the possibilities of achieving a highly interpretable latent representation through disentanglement of controllable and uncontrollable features.

VII. CONCLUSION AND FUTURE WORK

We have shown the possibility of disentangling controllable and uncontrollable features in an encoder architecture, strongly increasing the interpretability of the latent representation while also showing the potential use of this for downstream learning and planning, even in a single latent partition. This disentanglement of controllable and uncontrollable features in the latent representation of high-dimensional MDPs was achieved by propagating an action-conditioned forward prediction loss and a state-only forward prediction loss through distinct sections of the latent representation. Additionally, a contrastive loss and an adversarial loss were used to respectively avoid

collapse and further disentangle the latent representation. Finally, by employing forward prediction in latent space, we were able to successfully run a planning algorithm while leveraging the properties of the environment. In particular, the disentanglement of controllable and uncontrollable features allowed us to keep z^u frozen regardless of planning depth in the context of a distribution of randomly generated mazes, i.e. we only do forward prediction in z^c .

Future work could focus on gradually transferring our notion of disentanglement and interpretability to environments with more extensive underlying internal dynamics. Further work could also look at the ordering of the latent dimensions, as a latent representation is often arbitrarily ordered. This means that distinct seeds will lead to a different ordering and sign of the neurons in the final layer of the encoder. For example, if seed one would give agent position $+x$ and $+y$ for neurons 1 and 2 respectively, then seed two could give agent position $-y$ and $+x$ to the same neurons. As we are additionally using a contrastive loss while learning our representation, these results are compliant with the theory that a contrastive loss can recover the original latent information up to an orthogonal linear transformation [30].

Certain benefits can be obtained as well with a particular design of the encoder architecture, as we have done in this paper using estimates of the necessary dimensions of z^c and z^u for the different MDP environments. This can be seen as an inductive bias to aid disentanglement, as mentioned in [31]. Succeeding work could also focus on finding more algorithmic benefits of this disentanglement of controllable/uncontrollable features in more complex environments. For example, in the context of safety [32], a disentangled interpretable representation could allow incorporating latent state constraints in a planning algorithm. Lastly, as discussed in [10], [31], an interesting venue could be to further investigate the trade-off between interpretability and downstream performance. This is due to the fact that black-box representations such as Figure 4b still seem to have excellent downstream performance with DDQN, where for the task of maze navigation, a human would perform substantially better using the representation portrayed in Figure 4a as compared to using the representation in Figure 4b.

REFERENCES

- [1] R. Sutton and A. Barto, "Reinforcement learning: An introduction," *IEEE Transactions on Neural Networks*, 1998.
- [2] R. Bellman, "A markovian decision process," in *Journal of Mathematics and Mechanics*, vol. 6, no. 5, 1957.
- [3] R. Janschowski and O. Brock, "Learning state representations with robotic priors," *Autonomous Robots*, vol. 39, no. 3, 2015.
- [4] M. Jaderberg, V. Mnih, W. M. Czarnecki, T. Schaul, J. Z. Leibo, D. Silver, and K. Kavukcuoglu, "Reinforcement Learning with Unsupervised Auxiliary Tasks," in *ICLR*, 2017.
- [5] K. H. Lee, I. Fischer, A. Z. Liu, Y. Guo, H. Lee, J. Canny, and S. Guadarrama, "Predictive information accelerates learning in RL," in *NeurIPS*, 2020.
- [6] V. Francois-Lavet, Y. Bengio, D. Precup, and J. Pineau, "Combined Reinforcement Learning via Abstract Representations," in *AAAI*, 2019.
- [7] V. Thomas, J. Pondard, E. Bengio, M. Sarfati, P. Beaudoin, M.-J. Meurs, J. Pineau, D. Precup, and Y. Bengio, "Independently Controllable Factors," *arXiv preprint arXiv:1708.01289*, 2017.
- [8] T. Kipf, E. van der Pol, and M. Welling, "Contrastive Learning of Structured World Models," in *ICLR*, 2020.
- [9] K. Ahuja, J. Hartford, and Y. Bengio, "Weakly supervised representation learning with sparse perturbations," in *NeurIPS*, 2022.
- [10] C. Glanois, P. Weng, M. Zimmer, D. Li, T. Yang, J. Hao, and W. Liu, "A survey on interpretable reinforcement learning," *arXiv preprint arXiv:2112.13112*, 2021.
- [11] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2013.
- [12] I. Higgins, A. Pal, A. A. Rusu, L. Matthey, C. P. Burgess, A. Pritzel, M. Botvinick, C. Blundell, and A. Lerchner, "DARLA: Improving Zero-Shot Transfer in Reinforcement Learning," in *ICML*, 2017.
- [13] D. Yarats, A. Zhang, I. Kostrikov, B. Amos, J. Pineau, and R. Fergus, "Improving Sample Efficiency in Model-Free Reinforcement Learning from Images," in *AAAI*, 2021.
- [14] C. Gelada, S. Kumar, J. Buckman, O. Nachum, and M. G. Bellemare, "DeepMDP: Learning continuous latent space models for representation learning," in *ICML*, 2019.
- [15] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, "Learning latent dynamics for planning from pixels," in *ICML*, 2019.
- [16] A. Laversanne-Finot, A. Pere, and P.-Y. Oudeyer, "Curiosity driven exploration of learned disentangled goal spaces," in *PMLR*, 2018.
- [17] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, "Curiosity-driven Exploration by Self-supervised Prediction," in *CVPRW*, 2017.
- [18] A. P. Badia, P. Sprechmann, A. Vitvitskiy, D. Guo, B. Piot, S. Kapturowski, O. Tieleman, M. Arjovsky, A. Pritzel, A. Bolt, and C. Blundell, "Never Give Up: Learning Directed Exploration Strategies," in *ICLR*, 2020.
- [19] A. Lamb, R. Islam, Y. Efroni, A. Didolkar, D. Misra, D. Foster, L. Molu, R. Chari, A. Krishnamurthy, and J. Langford, "Guaranteed discovery of controllable latent states with multi-step inverse models," *arXiv preprint arXiv:2207.08229*, 2022.
- [20] D. Bertoin and E. Rachelson, "Disentanglement by cyclic reconstruction," in *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [21] X. Fu, G. Yang, P. Agrawal, and T. Jaakkola, "Learning task informed abstractions," in *ICML*, 2021.
- [22] T. Wang, S. Du, A. Torralba, P. Isola, A. Zhang, and Y. Tian, "Denoised MDPs: Learning world models better than the world itself," in *PMLR*, 2022.
- [23] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, V. Lempitsky, U. Dogan, M. Kloft, F. Orabona, T. Tommasi, and a. Ganin, "Domain-Adversarial Training of Neural Networks," in *Journal of Machine Learning Research*, vol. 17, 2016.
- [24] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Networks," in *NeurIPS*, 2014.
- [25] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, 2015.
- [26] H. van Hasselt, A. Guez, and D. Silver, "Deep Reinforcement Learning with Double Q-learning," in *AAAI*, 2016.
- [27] I. Higgins, D. Amos, D. Pfau, S. Racanière, L. Matthey, D. J. Rezende, and A. Lerchner, "Towards a definition of disentangled representations," *arXiv preprint arXiv:1812.02230*, 2018.
- [28] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. d. L. Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, T. Lillicrap, and M. Riedmiller, "DeepMind Control Suite," *arXiv preprint arXiv:1801.00690*, 2018.
- [29] J. Oh, S. Singh, and H. Lee, "Value Prediction Network," in *NeurIPS*, 2017.
- [30] R. S. Zimmermann, Y. Sharma, S. Schneider, M. Bethge, and W. Brendel, "Contrastive learning inverts the data generating process," in *PMLR*, 2021.
- [31] F. Locatello, S. Bauer, M. Lucic, G. Raetsch, S. Gelly, B. Schölkopf, and O. Bachem, "Challenging common assumptions in the unsupervised learning of disentangled representations," in *PMLR*, 2019.
- [32] O. Davoodi and M. Komeili, "Feature-based interpretable reinforcement learning based on state-transition models," in *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2021.