

Disentangled World Models: Learning to Transfer Semantic Knowledge from Distracting Videos for Reinforcement Learning

Qi Wang^{1,2,3*} Zhipeng Zhang^{4,5*} Baao Xie^{2,3*} Xin Jin^{2,3†} Yunbo Wang¹
 Shiyu Wang^{5,6} Liaomo Zheng^{5,6} Xiaokang Yang¹ Wenjun Zeng^{2,3}

¹ MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

² Ningbo Institute of Digital Twin, Eastern Institute of Technology, Ningbo, China

³ Ningbo Key Laboratory of Spatial Intelligence and Digital Derivative, Ningbo, China

⁴ University of Chinese Academy of Sciences

⁵ Shenyang Institute of Computing Technology, Chinese Academy of Sciences

⁶ Shenyang CASNC Technology Co., Ltd

<https://qiwang067.github.io/diswm>

Abstract

Training visual reinforcement learning (RL) in practical scenarios presents a significant challenge, i.e., RL agents suffer from low sample efficiency in environments with variations. While various approaches have attempted to alleviate this issue by disentangled representation learning, these methods usually start learning from scratch without prior knowledge of the world. This paper, in contrast, tries to learn and understand underlying semantic variations from distracting videos via offline-to-online latent distillation and flexible disentanglement constraints. To enable effective cross-domain semantic knowledge transfer, we introduce an interpretable model-based RL framework, dubbed Disentangled World Models (DisWM). Specifically, we pre-train the action-free video prediction model offline with disentanglement regularization to extract semantic knowledge from distracting videos. The disentanglement capability of the pretrained model is then transferred to the world model through latent distillation. For finetuning in the online environment, we exploit the knowledge from the pretrained model and introduce a disentanglement constraint to the world model. During the adaptation phase, the incorporation of actions and rewards from online environment interactions enriches the diversity of the data, which in turn strengthens the disentangled representation learning. Experimental results validate the superiority of our approach on various benchmarks.

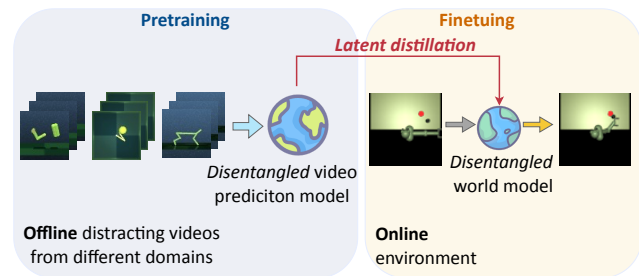


Figure 1. Overview of our proposed framework. The key idea is to leverage distracting videos for semantic knowledge transfer, enabling the downstream agent to improve sample efficiency on unseen tasks.

1. Introduction

Visual reinforcement learning (VRL) presents a promising approach for training agents within complex environments [11, 18, 20, 28, 45]. However, VRL frequently suffers from performance degradation in practical scenarios due to the complexity, volatility, and visual distractions in environments. Even minor environmental variations can result in significant pixel-level shifts, making the trained VRL policies ineffective or suboptimal [6, 7]. For instance, a slight change in lighting conditions can affect an object’s appearance (e.g., color, shadow, or other visual attributes). Therefore, it is crucial to enhance models with interpretability, enabling them to perceive, learn, and understand the semantic environmental variations.

Disentangled representation learning (DRL) presents a promising approach to addressing the interpretability challenges inherent in the “black-box” nature of deep learning algorithms. Fundamentally, DRL approaches mimic the cognitive processes of biological intelligence, where under-

*Equal contribution.

†Corresponding author: Xin Jin <jinxin@eitech.edu.cn>.

Table 1. MuJoCo (*downstream domain*) vs. DMC (*accessible distracting videos*).

	Video: <i>DMC</i>	Target: <i>MuJoCo</i>	Similarity / Difference
Task	Reacher Easy	Pusher	Relevant robotic control tasks
Dynamics	Two-link planar	Multi-jointed robot arm	Different
Action space	Box(-1, 1, (2.), float32)	Box(-2, 2, (7.), float32)	Different
Reward range	[0, 1]	[-4.49, 0]	Different

standing the world is facilitated by decomposing observations into distinct and independent factors [2, 14, 38, 40, 41]. In this form, when a factor of variation is changed (*e.g.*, color), only a small portion of features in the disentangled representation will be affected, enabling the agent to recover performance quickly. Several studies have explored the integration of DRL algorithms in the domain of VRL. For example, Higgins *et al.* [15] trained a β -VAE offline to obtain disentangled representations for reinforcement learning. TED [7] adopts a self-supervised auxiliary task to learn temporally disentangled representations for reinforcement learning. Additionally, Dunion *et al.* [6] introduced conditional mutual information to achieve a disentangled representation with correlated data. However, existing methods typically learn the representations from scratch, lacking any prior knowledge of the world. These approaches often require extensive interactions with the environment to acquire desired behaviors.

Towards this challenge, we introduce a model-based interpretable VRL framework, dubbed **Disentangled World Models (DisWM)**, which leverages prior knowledge extracted from distracting videos to facilitate the learning of unseen downstream tasks through latent distillation. It is crucial to note that **distracting videos refer to videos with visual distractions**, which are beneficial for learning disentangled representations. Specifically, as depicted in Figure 1, our framework consists of two phases: first, we pre-train a DRL encoder to learn disentangled latent representations from distracting videos. By doing so, the pretrained DRL encoder is “knowledgeable” in terms of representation disentanglement. Subsequently, we finetune an orthogonally designed world model with dual constraints of disentanglement and distillation, leveraging semantic knowledge transferred from the pretrained model via offline-to-online latent distillation. Another benefit of disentangled world model adaptation is that incorporating actions and rewards from online interactions with the environment enriches the diversity of the visual observations, which in turn strengthens the process of disentangled representation learning. It is worth mentioning that, as a cross-domain framework, DisWM does not require the pretraining videos to originate from the same domain as the downstream tasks.

Experimental results demonstrate the effectiveness of our proposed approach in improving the sample efficiency of VRL agents across our modified DeepMind Control and

MuJoCo *Pusher*. The contributions of this work can be summarized as follows:

- We frame the problem of learning interpretable VRL agents as a domain transfer learning problem. The key idea is to extract semantic knowledge from distracting videos and transfer this disentanglement capability to downstream control tasks.
- We present DisWM, an approach that follows the pretraining-finetuning paradigm using distracting videos, incorporating specific techniques of offline-to-online latent distillation and flexible disentanglement constraints.

2. Problem Setup

We formulate visual reinforcement learning as a partially observable Markov decision process (POMDP) that uses DMC and MuJoCo *Pusher* as the test bench. Specifically, we concentrate on scenarios where videos without actions and rewards are accessible, enabling world knowledge transfer. The goal is to maximize the cumulative reward of the target POMDP $\langle \mathcal{O}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma \rangle$ by transferring the shared world knowledge from the videos. These notations correspond to the visual observation space, the action space, the transition probabilities, the reward function, and the discount factor, respectively.

For instance, in one of the cross-domain experiments, we use MuJoCo as the downstream domain and the frames collected from DMC as the distracting video. Table 1 highlights the differences between the two domains in terms of visual appearances, physical dynamics, action spaces, and reward functions.

3. Method

3.1. Overview of DisWM

In this section, we present the details of DisWM, which involves three main stages (see Figure 2):

- Disentangled representation pretraining*: Pretrain a DRL-based video prediction model from distracting videos to extract disentangled features.
- Offline-to-online latent distillation*: Transfer the semantic knowledge from the pretrained model to the world model via cross-domain latent distillation.
- Disentangled world model adaptation*: Finetune the downstream agent with disentanglement constraints by incorporating the action and reward information.

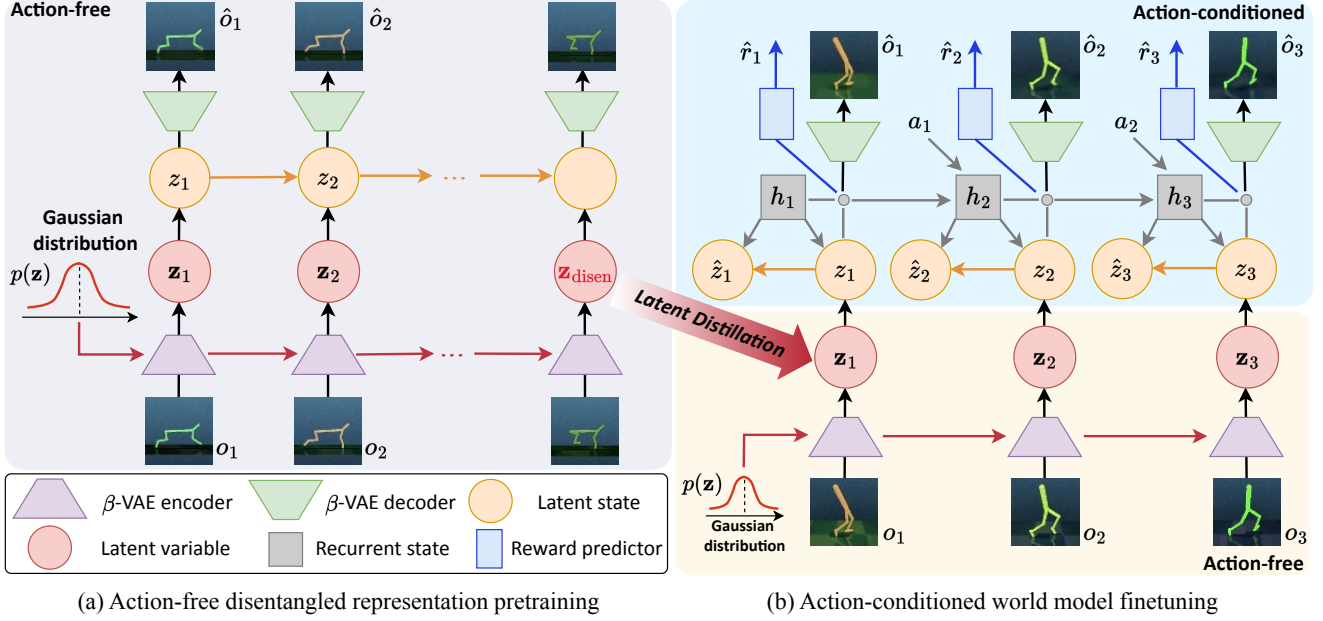


Figure 2. Architecture of Disentangled World Models. **(a)** The action-free video prediction model with disentanglement constraints is pre-trained on distracting videos offline for the well-disentangled latent variable $\mathbf{z}_{\text{disen}}$, which extracts semantic knowledge from the visual observations. The disentangled capability of $\mathbf{z}_{\text{disen}}$ is then transferred to the world model through latent distillation. **(b)** The action-conditioned world model is finetuned through disentanglement regularization online, which encourages the factorized representation. Moreover, the incorporation of actions and rewards enriches the diversity of the data, which in turn strengthens the disentangled representation learning.

3.2. Disentangled Representation Pretraining

To extract well-disentangled representations that can be transferred to the downstream world model, we first train a video prediction model on distracting videos without incorporating action information (**Lines 4-8** of Alg. 1). This model comprises three key components: (i) the posterior learner that encodes the observation o_t to latent state z_t via β -VAE encoder¹, which serves as a typical DRL framework to extract latent features \mathbf{z}_t from observations, (ii) the prior module that predicts future latent states based on historical states, without directly relying on the current observation o_t , and (iii) the β -VAE-based decoder that reconstructs \hat{o}_t from the latent state z_t . Concretely, the model can be formulated as follows:

$$\begin{aligned}
 \beta\text{-VAE encoder:} & \quad \mathbf{z}_t = e_{\phi'}(o_t) \\
 \text{Posterior state:} & \quad z_t \sim q_{\phi'}(z_t | z_{t-1}, \mathbf{z}_t) \\
 \text{Prior state:} & \quad \hat{z}_t \sim p_{\phi'}(\hat{z}_t | z_{t-1}) \\
 \text{Reconstruction:} & \quad \hat{o}_t \sim p_{\phi'}(\hat{o}_t | z_t) \\
 \text{Isotropic unit Gaussian:} & \quad p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, I).
 \end{aligned} \tag{1}$$

where ϕ' denotes the parameters of the model. The β -VAE-based video prediction model is trained to minimize the fol-

lowing loss function:

$$\begin{aligned}
 \mathcal{L}(\phi') = \mathbb{E}_{q_{\phi'}} \left[\underbrace{\sum_{t=1}^T -\ln p_{\phi'}(o_t | z_t)}_{\text{image reconstruction}} \right. \\
 \left. + \underbrace{\beta_1 \text{KL}[q_{\phi'}(z_t | z_{t-1}, o_t) \| p_{\phi'}(\hat{z}_t | z_{t-1})]}_{\text{action-free KL loss}} \right. \\
 \left. + \underbrace{\beta_2 \text{KL}[q_{\phi'}(\mathbf{z}_t | o_t) \| p(\mathbf{z}_t)]}_{\text{disentanglement loss}} \right]. \tag{2}
 \end{aligned}$$

The variational posterior distribution $q_{\phi'}(\mathbf{z}_t | o_t)$ is encouraged to be close to the standard multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}, I)$ to strengthen the orthogonality and disentanglement of the latent space. The importance of the disentanglement loss term is governed by β_2 .

3.3. Offline-to-Online Latent Distillation

After the offline pretraining with the distracting videos, the model is fine-tuned online to adapt to the downstream task by integrating actions and rewards (**Lines 13** of Alg. 1). A straightforward approach to transfer disentangled features involves initializing the action-conditioned world model with checkpoints obtained from the pretrained video prediction model. Nevertheless, it may experience a potential mismatch issue caused by the discrepancies between the two domains in visual appearances and physical dynamics. Directly applying the pretraining-finetuning paradigm

¹ β -VAE [14] is a variant of autoencoder that introduces a hyperparameter β to balance the reconstruction quality and disentanglement capability.

for downstream tasks tends to overwrite the disentangled information encoded in the pretrained latent features, leading to decreased performance when there are large domain discrepancies between the source and the target domains.

Through comprehensive pretraining on distracting videos that contain diverse visual variations, the video prediction model thus builds an interpretable and orthogonality latent space. In this space, the latent variable $\mathbf{z}_{\text{disen}}$ achieves a high degree of disentanglement. To exploit the prior semantic knowledge from the pretrained model and improve the sample efficiency of the downstream tasks, we introduce an offline-to-online latent distillation. This approach enables the disentangling capability of $\mathbf{z}_{\text{disen}}$ from the pretrained model to be effectively transferred to the latent variable \mathbf{z}_{task} of the world model. Specifically, this is achieved by minimizing the Kullback-Leibler (KL) divergence between the latent distributions of the two domains. The corresponding distillation loss $\mathcal{L}_{\text{distill}}$ can be formulated as follows:

$$\mathcal{L}_{\text{distill}} = \text{KL}(\mathbf{z}_{\text{disen}} \parallel \mathbf{z}_{\text{task}}) = \sum \mathbf{z}_{\text{disen}} \cdot \log \left(\frac{\mathbf{z}_{\text{disen}}}{\mathbf{z}_{\text{task}}} \right) \quad (3)$$

3.4. Disentangled World Model Adaptation

By obtaining well-disentangled representations of $\mathbf{z}_{\text{disen}}$ and employing the latent distillation for knowledge transfer, we then propose a DRL-based world model \mathcal{M}_ϕ , designed to harness these features to enhance interoperability and robustness against environmental variations (Lines 14-15 of Alg. 1). The components of \mathcal{M}_ϕ can be detailed as follows:

$$\begin{aligned} \text{Recurrent transition:} & \quad h_t = f_\phi(h_{t-1}, z_{t-1}, a_{t-1}) \\ \beta\text{-VAE encoder:} & \quad \mathbf{z}_t \sim e_\phi(o_t) \\ \text{Posterior state:} & \quad z_t \sim q_\phi(z_t \mid h_t, \mathbf{z}_t) \\ \text{Prior state:} & \quad \hat{z}_t \sim p_\phi(\hat{z}_t \mid h_t) \\ \text{Reconstruction:} & \quad \hat{o}_t \sim p_\phi(\hat{o}_t \mid h_t, z_t) \\ \text{Reward prediction:} & \quad \hat{r}_t \sim r_\phi(\hat{r}_t \mid h_t, z_t) \\ \text{Discount factor:} & \quad \hat{\gamma}_t \sim p_\phi(\hat{\gamma}_t \mid h_t, z_t) \\ \text{Isotropic unit Gaussian:} & \quad p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, I), \end{aligned} \quad (4)$$

where ϕ represents the combined parameters of the world model. We train \mathcal{M}_ϕ on the sampled data from the replay buffer \mathcal{B} with the following loss function:

$$\begin{aligned} \mathcal{L}(\phi) = \mathbb{E}_{q_\phi} \left[\underbrace{\sum_{t=1}^T -\ln p_\phi(o_t \mid h_t, z_t)}_{\text{image reconstruction}} - \underbrace{\ln r_\phi(r_t \mid h_t, z_t)}_{\text{reward prediction}} \right. \\ \left. - \underbrace{\ln p_\phi(\gamma_t \mid h_t, z_t)}_{\text{discount prediction}} + \alpha \underbrace{\text{KL}[q_\phi(z_t \mid h_t, o_t) \parallel p_\phi(\hat{z}_t \mid h_t)]}_{\text{KL divergence}} \right. \\ \left. + \beta \underbrace{\text{KL}[q_\phi(\mathbf{z}_t \mid o_t) \parallel p(\mathbf{z}_t)]}_{\text{disentanglement loss}} + \underbrace{\eta \mathcal{L}_{\text{distill}}}_{\text{distillation loss}} \right]. \end{aligned} \quad (5)$$

where β is a hyperparameter used to balance reconstruction quality and disentanglement capability. In this adaptation stage, η serves as a hyperparameter that gradually decreases from 0.1 to 0.01. Intuitively, η controls the progressive adaptation of the world model with the shared world knowledge transfer from the pretrained video prediction model.

Through the comprehensive training processes of this framework, we equip the DisWM with the capability to learn and understand the underlying semantic representations. This enhancement enables the model to be less sensitive to environmental variations, such as changes in object colors, positions, and backgrounds. Furthermore, by incorporating actions and rewards during the finetuning phase, the world model can generate data with more diverse representations, thereby improving the disentangled representation learning. For behavior learning, we utilize the actor-critic method that is in line with DreamerV2 [10] (Lines 16-18 of Alg. 1). For more details of behavior learning, please refer to [Supplementary Material B](#).

4. Experiments

4.1. Experimental Setups

Benchmark We evaluate DisWM on DeepMind Control Suite (DMC) [34], MuJoCo *Pusher* [35], and DrawerWorld [37]. DMC is a widely adopted benchmark with comprehensive and flexible robotic-control tasks. For DMC benchmark, we use 5 tasks, *i.e.*, *Walker Walk*, *Cheetah Run*, *Hopper Stand*, *Finger Spin*, *Cartpole Swingup*. In MuJoCo *Pusher*, a multi-jointed robotic arm is employed to manipulate a target cylinder (object). The goal is to move the object to a designated target position using the robot’s end effector (fingertip). The agent receives the negative reward, which is a combination of three components: the distance between the fingertip and the target, the distance between the object and the target goal position, and a penalty for large actions. DrawerWorld is a modified Metaworld [43] benchmark designed to evaluate texture adaptability in manipulation tasks. It includes five additional textures from realistic photos and *grid* texture. During training, we initially employ the *grid* texture and then change to the *wood* texture midway, while adopting the *metal* texture exclusively for evaluation. The corresponding results on DrawerWorld are reported in [Supplementary Material C.2](#). Furthermore, the introduction of compared baselines is given in [Supplementary Material A](#).

Implementation details. The visual observations of online finetuning stages are resized to 64×64 pixels. Inspired by APV [29], we build distracting video datasets with 1M frames using DreamerV2 [10] to interact with the environments with visual color distractors. This video datasets consist of the samples stored in the replay buffer throughout the

Algorithm 1 The training pipeline of DisWM.

```
1: Hyperparameters:  $H$ : Horizon of latent imagination.
2: Require: Distracting video dataset  $\mathcal{D}$ .
3: Initialize: Parameters of the model  $\{\phi, \psi, \xi\}$ .
4: for training step  $t = 1, 2, \dots, K_1$  do ▷ Disentangled representation pretraining
5:   Sample random minibatch  $\{o_t\}_{t=1}^T \sim \mathcal{D}$ .
6:   Obtain Gaussian prior  $\mathbf{z}_t$  from Isotropic unit Gaussian  $\mathcal{N}(\mathbf{0}, I)$ .
7:   Pretrain the action-free video prediction model with disentanglement regularization by minimizing Eq. (2).
8: end for
9: Train the random agent and collect a replay buffer  $\mathcal{B}$ .
10: while not converged do
11:   for training step  $t = 1, 2, \dots, K_2$  do
12:     Sample  $\{(o_t, a_t, r_t)\}_{t=1}^T \sim \mathcal{B}$ .
13:     Distill the disentangled features to the world model using Eq. (3). ▷ Offline-to-online latent distillation
14:     Obtain Gaussian prior  $\mathbf{z}_t$  from Isotropic unit Gaussian  $\mathcal{N}(\mathbf{0}, I)$ . ▷ Disentangled world model adaptation
15:     Train the world model  $\mathcal{M}_\phi$  with latent distillation and disentanglement constraints using Eq. (5).
16:     Generate  $\{(\hat{z}_i, \hat{a}_i)\}_{i=t}^{t+H}$  using  $\pi_\psi$  and  $\mathcal{M}_\phi$ . ▷ Behavior learning
17:     Train the critic  $v_\xi$  over  $\{(\hat{z}_i, \hat{a}_i)\}_{i=t}^{t+H}$ .
18:     Train the actor  $\pi_\psi$  over  $\{(\hat{z}_i, \hat{a}_i)\}_{i=t}^{t+H}$ .
19:   end for
20:    $o_1 \leftarrow \text{env.reset}()$  ▷ Environment interaction
21:   for time step  $t = 1, 2, \dots, T$  do
22:     Sample  $\hat{a}_t \sim \pi_\psi(\hat{a}_t | \hat{z}_t)$ .
23:      $r_t, o_{t+1} \leftarrow \text{env.step}(\hat{a}_t)$ .
24:   end for
25:   Append data to the replay buffer  $\mathcal{B}$ .
26: end while
```

training process until the agent achieves maximum score. For tasks in the DMC benchmark, the training steps of the agent are limited to 1×10^6 environment steps. Each run of DisWM requires roughly 5GB of VRAM and takes around 16 hours to train on a single RTX 3090 GPU. The dimensions of well-disentangled latent $\mathbf{z}_{\text{disen}}$ and downstream task latent \mathbf{z}_{task} are both set to 20 in our approach. In Figure 3, we showcase the example observations of various tasks with color distractors. We train the agent using a fixed set of colors, where the RGB values are varied within a restricted range around the original values. Furthermore, at the midpoint of the training process, we change to a different color scheme for varying distractors.

4.2. Main Comparison

We evaluate the sample efficiency and task performance for all the methods with the training curves of episode return. Figure 4 shows the performance of DisWM and all the baselines. Remarkably, it achieves better performance than TED [7], a method on top of RAD [17] and tailored for environments with distractors. For the offline-to-online finetuning models, *DV2 Finetune* achieves the second-best performance by transferring knowledge from the distracting videos. However, we observe a significant decline in sample efficiency, particularly in scenarios with large data distribution shifts between the source and target domains (*e.g.*,

DMC \rightarrow MuJoCo). These shifts can occur in various aspects, including visual observation, physical dynamics, reward definition, or the action space of the robots. Another crucial baseline is APV [29], which focuses on transferring knowledge obtained from videos with a stacked latent prediction model. Nevertheless, without environment-specific designs for visual distractors, directly training may eventually result in a decrease in performance in the downstream tasks. The CURL model struggles to learn effective behavior policies, especially in *Hopper Stand*. Additional results on the challenging DMC *Humanoid Walk* can be found in [Supplementary Material C.1](#).

Additionally, we present qualitative results in Figure 5 and Figure 6. Figure 5 shows the traversals of β -VAE during the pretraining phase. In each row of traversals, a distinct attribute varies, while other attributes remain constant, indicating that the pretrained model has successfully disentangled and learned this attribute, thereby improving the sample efficiency of the RL agent. Figure 6 displays the fine-grained disentanglement results on MuJoCo *Pusher* during the finetuning phase, demonstrating that the world model can effectively disentangle the variations.

4.3. Model Analyses

Ablation studies. We conduct ablation studies to validate the effect of the latent distillation and disentanglement con-




Figure 3. Example image observations of our modified DMC and MuJoCo *Pusher* with color distractors.

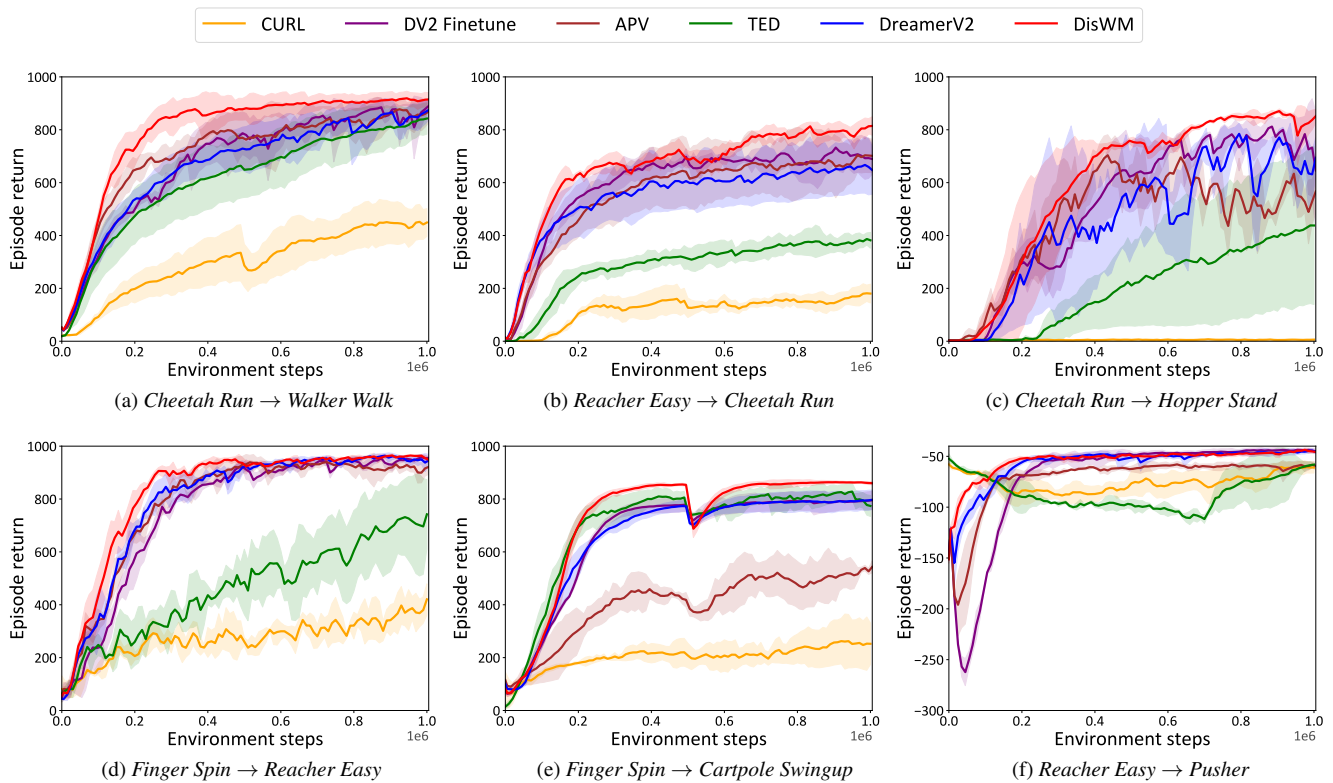


Figure 4. Comparison of DisWM against visual RL baselines, including *DreamerV2* [10], *APV* [29], *DV2 Finetune*, *TED* [7], *CURL* [18].

straints. Figure 7 (Left) shows corresponding results in the DMC *Walker Walk* \rightarrow *Cheetah Run*. The green curve shows that removing the latent distillation of DisWM results in a decreased performance, which indicates that the latent

distillation is essential during the early training stage. For the model represented by the blue curve, we do not adopt disentanglement constraints for both pretrain and finetune stages. It can be seen that the necessity of introducing DRL-




Figure 5. Visualization of traversals of β -VAE during the pretraining phase.




Figure 6. Fine-grained disentanglement results on MuJoCo *Pusher* during world model adaptation phase. Each row in *Pusher* images displays the traversal results on a specific attribute.

based training and disentangled representation significantly improves the learning efficiency of the agent.

Sensitivity analyses. We conduct sensitivity analyses on DMC (*Cheetah Run* \rightarrow *Walker Walk*). As shown in Figure 7 (Middle), we observe that when β for the representation disentanglement is too small, the model learns entangled latent representations. When β is too large, it will impede the reconstruction of the image, leading to a decline in performance. Latent distillation weight η controls the cross-domain transfer scale. Intuitively, setting this hyperparameter too low may result in the downstream agent not getting enough knowledge from the pretrained mod-

els. Conversely, excessively high η may result in the model overfitting to the pretrained model, which is not conducive to the learning of the downstream task. Additional sensitivity analyses on the latent space dimension are provided in [Supplementary Material C.3](#).

Effects of video domain. In Figure 8, we evaluate DisWM on DMC *Cartpole Swingup* by pretraining on alternative video datasets, including frames collected from *Finger Spin*, *Reacher Easy*, *Walker Walk*, and *Hopper Stand*. Interestingly, compared with the DreamerV2 agent without pretraining, DisWM can always benefit from pretraining via offline-to-online latent distillation. It obtains the semantic




Figure 7. These figures illustrate the ablation studies and sensitivity analyses of DisWM on DMC *Walker Walk* \rightarrow *Cheetah Run*. **Left:** Comparison with DisWM without cross-domain latent distillation or disentanglement constraints. **Middle:** The sensitivity analyses of cross-domain latent distillation weight. **Right:** The performance of DisWM with different disentanglement scale.

knowledge from the pretrained model and strengthens the disentanglement capability during finetuning phase.




Figure 8. Performance of DisWM on DMC *Cartpole Swingup* with different video datasets.

5. Related Work

Visual MBRL. Visual RL learns control policies from raw pixels, which has achieved remarkable performance in various tasks [3, 4, 31, 37], while prior RL studies focus on learning policies from low-dimensional states. Extant approaches can be divided into two main directions: model-free RL [18, 19, 24, 32, 42, 44, 47] and model-based RL [1, 8–10, 12, 20, 21, 23, 27, 30, 36, 46]. The following methods specifically address the variations of environments in visual MBRL. Pan [28] *et al.* decompose visual dynamics into controllable and uncontrollable states through the optimization of inverse dynamics. SeeX [16] proposed a bi-level optimization framework that adopts a separated world model and maximizes the task-relevant uncertainty. Orthogonal to these studies, our approach employs a DRL-based world model to alleviate the issue of visual variations.

Transfer RL To facilitate the learning of unseen tasks, transfer RL [22, 24, 26, 33, 36, 48] leverages the knowledge learned from past tasks. One promising way is to transfer world knowledge from accessible videos to improve the downstream control. APV [29] established a pretraining-finetuning framework with a stacked latent prediction model and video-based intrinsic bonus. IPV [39] introduced contextualized world models that pretrained on diverse in-the-wild videos. It incorporates a context encoder that works alongside the latent dynamics model into the image encoder to capture rich contextual information. PreLAR [45] pretrained the world model with the derived meaningful actions from the action-free video using an inverse dynamics encoder. Different from these approaches, we propose a new solution to transfer world knowledge from distracting videos to improve the learning efficiency of downstream tasks via offline-to-online latent distillation.

6. Conclusions and Limitations

In this paper, we present a transfer RL dubbed DisWM, which addresses the challenge of environment variations in practical scenarios. Our key insight is to leverage the accessible distracting videos to facilitate the sample efficiency of downstream tasks to offer flexible disentanglement constraints. Specifically, we introduce *disentangled representation pretraining*, *offline-to-online latent distillation*, and *disentangled world model adaptation* to improve the downstream control. DisWM demonstrates superior performance than existing visual RL baselines across various benchmarks.

One limitation of our approach is that disentangled representation learning encounters challenges in complex environments. Exploring the non-stationary environments with more intricate variations, such as time-varying background video distractions, could further highlight the potential of our approach for practical scenarios.

Acknowledgements. This work was supported by Grants of NSFC 62302246 & 62250062, ZJNSFC LQ23F010008, Ningbo 2023Z237 & 2024Z284 & 2024Z289 & 2023CX050011 & 2025Z038, the Smart Grid National Science and Technology Major Project (2024ZD0801200), the Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), the Fundamental Research Funds for the Central Universities, and the IDT Foundation of Youth Doctoral Innovation (S203.2.01.32.002). Additional support was provided by the project of Supporting Program for Young and Middle-aged Scientific and Technological Innovation Talents in Shenyang (Grant RC210488), the project of Provincial Doctoral Research Initiation Fund Program (Grant 2023-BS-214), the High Performance Computing Center at Eastern Institute of Technology, Ningbo, and Ningbo Institute of Digital Twin.

References

- [1] Eloi Alonso, Adam Jelley, Vincent Micheli, Anssi Kanervisto, Amos Storkey, Tim Pearce, and François Fleuret. Diffusion for world modeling: Visual details matter in atari. In *NeurIPS*, 2024. 8
- [2] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *TPAMI*, 35(8):1798–1828, 2013. 2
- [3] Hyesong Choi, Hunsang Lee, Seongwon Jeong, and Dongbo Min. Environment agnostic representation for visual reinforcement learning. In *ICCV*, pages 263–273, 2023. 8
- [4] Hyesong Choi, Hunsang Lee, Wonil Song, Sangryul Jeon, Kwanghoon Sohn, and Dongbo Min. Local-guided global: Paired similarity representation for visual reinforcement learning. In *CVPR*, pages 15072–15082, 2023. 8
- [5] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015. 1
- [6] Mhairi Dunion, Trevor McInroe, Kevin Sebastian Luck, Josiah Hanna, and Stefano Albrecht. Conditional mutual information for disentangled representations in reinforcement learning. In *NeurIPS*, 2023. 1, 2
- [7] Mhairi Dunion, Trevor McInroe, Kevin Sebastian Luck, Josiah P Hanna, and Stefano V Albrecht. Temporal disentanglement of representations for improved generalisation in reinforcement learning. In *ICLR*, 2023. 1, 2, 5, 6
- [8] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *ICML*, 2019. 8
- [9] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *ICLR*, 2020.
- [10] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. In *ICLR*, 2021. 4, 6, 8, 1
- [11] Danijar Hafner, Kuang-Huei Lee, Ian Fischer, and Pieter Abbeel. Deep hierarchical planning from pixels. *arXiv preprint arXiv:2206.04114*, 2022. 1
- [12] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *Nature*, 2025. 8, 1
- [13] Nicklas Hansen, Hao Su, and Xiaolong Wang. Td-mpc2: Scalable, robust world models for continuous control. In *ICLR*, 2024. 1
- [14] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017. 2, 3
- [15] 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. In *ICML*, 2017. 2
- [16] Kaichen Huang, Shenghua Wan, Minghao Shao, Hai-Hang Sun, Le Gan, Shuai Feng, and De-Chuan Zhan. Leveraging separated world model for exploration in visually distracted environments. *NeurIPS*, 2024. 8
- [17] Misha Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. In *NeurIPS*, 2020. 5
- [18] Michael Laskin, Aravind Srinivas, and Pieter Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. In *ICML*, pages 5639–5650, 2020. 1, 6, 8
- [19] Haoran Li, Zhennan Jiang, YUHUI CHEN, and Dongbin Zhao. Generalizing consistency policy to visual rl with prioritized proximal experience regularization. In *NeurIPS*, 2024. 8
- [20] Jiajian Li, Qi Wang, Yunbo Wang, Xin Jin, Yang Li, Wenjun Zeng, and Xiaokang Yang. Open-world reinforcement learning over long short-term imagination. In *ICLR*, 2025. 1, 8
- [21] Jessy Lin, Yuqing Du, Olivia Watkins, Danijar Hafner, Pieter Abbeel, Dan Klein, and Anca Dragan. Learning to model the world with language. In *ICML*, 2024. 8
- [22] Chris Lu, Yannick Schroecker, Albert Gu, Emilio Parisotto, Jakob Foerster, Satinder Singh, and Feryal Behbahani. Structured state space models for in-context reinforcement learning. *NeurIPS*, 2023. 8
- [23] Haoyu Ma, Jialong Wu, Ningya Feng, Chenjun Xiao, Dong Li, Jianye Hao, Jianmin Wang, and Mingsheng Long. Harmonydream: Task harmonization inside world models. In *ICML*, 2024. 8
- [24] Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training. In *ICLR*, 2023. 8
- [25] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 9(Nov):2579–2605, 2008. 2
- [26] Pietro Mazzaglia, Tim Verbelen, Bart Dhoedt, Alexandre Lacoste, and Sai Rajeswar. Choreographer: Learning and adapting skills in imagination. In *ICLR*, 2023. 8

- [27] Pietro Mazzaglia, Tim Verbelen, Bart Dhoedt, Aaron C Courville, and Sai Rajeswar Mudumba. Genrl: Multimodal-foundation world models for generalization in embodied agents. *NeurIPS*, 2024. 8
- [28] Minting Pan, Xiangming Zhu, Yunbo Wang, and Xiaokang Yang. Iso-dream: Isolating and leveraging noncontrollable visual dynamics in world models. In *NeurIPS*, 2022. 1, 8
- [29] Younggyo Seo, Kimin Lee, Stephen L James, and Pieter Abbeel. Reinforcement learning with action-free pre-training from videos. In *ICML*, 2022. 4, 5, 6, 8, 1
- [30] Younggyo Seo, Junsu Kim, Stephen James, Kimin Lee, Jinwoo Shin, and Pieter Abbeel. Multi-view masked world models for visual robotic manipulation. In *ICML*, 2023. 8
- [31] Wonil Song, Hyesong Choi, Kwanghoon Sohn, and Dongbo Min. A simple framework for generalization in visual rl under dynamic scene perturbations. *NeurIPS*, 37:121790–121826, 2024. 8
- [32] Adam Stooke, Kimin Lee, Pieter Abbeel, and Michael Laskin. Decoupling representation learning from reinforcement learning. In *ICML*, pages 9870–9879, 2021. 8
- [33] Yanchao Sun, Ruijie Zheng, Xiyao Wang, Andrew Cohen, and Furong Huang. Transfer rl across observation feature spaces via model-based regularization. In *ICLR*, 2022. 8
- [34] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018. 4
- [35] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *IROS*, 2012. 4
- [36] Qi Wang, Junming Yang, Yunbo Wang, Xin Jin, Wenjun Zeng, and Xiaokang Yang. Making offline rl online: Collaborative world models for offline visual reinforcement learning. In *NeurIPS*, 2024. 8
- [37] Xudong Wang, Long Lian, and Stella X Yu. Unsupervised visual attention and invariance for reinforcement learning. In *CVPR*, pages 6677–6687, 2021. 4, 8, 1
- [38] Xin Wang, Hong Chen, Zihao Wu, Wenwu Zhu, et al. Disentangled representation learning. *TPAMI*, 2024. 2
- [39] Jialong Wu, Haoyu Ma, Chaoyi Deng, and Mingsheng Long. Pre-training contextualized world models with in-the-wild videos for reinforcement learning. *NeurIPS*, 2023. 8, 1
- [40] Baao Xie, Bohan Li, Zequn Zhang, Junting Dong, Xin Jin, Jingyu Yang, and Wenjun Zeng. Navinerf: Nerf-based 3d representation disentanglement by latent semantic navigation. In *ICCV*, 2023. 2
- [41] Baao Xie, Qiuyu Chen, Yunnan Wang, Zequn Zhang, Xin Jin, and Wenjun Zeng. Graph-based unsupervised disentangled representation learning via multimodal large language models. In *NeurIPS*, 2024. 2
- [42] Denis Yarats, Rob Fergus, Alessandro Lazaric, and Lerrel Pinto. Mastering visual continuous control: Improved data-augmented reinforcement learning. In *ICLR*, 2022. 8
- [43] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *CoRL*, 2019. 4
- [44] Amy Zhang, Rowan McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning invariant representations for reinforcement learning without reconstruction. In *ICLR*, 2021. 8
- [45] Lixuan Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Prelar: World model pre-training with learnable action representation. In *ECCV*, 2024. 1, 8
- [46] Weipu Zhang, Gang Wang, Jian Sun, Yetian Yuan, and Gao Huang. Storm: Efficient stochastic transformer based world models for reinforcement learning. In *NeurIPS*, 2023. 8
- [47] Ruijie Zheng, Xiyao Wang, Yanchao Sun, Shuang Ma, Jieyu Zhao, Huazhe Xu, Hal Daumé III, and Furong Huang. Taco: Temporal latent action-driven contrastive loss for visual reinforcement learning. In *NeurIPS*, 2023. 8
- [48] Zhuangdi Zhu, Kaixiang Lin, Anil K Jain, and Jiayu Zhou. Transfer learning in deep reinforcement learning: A survey. *TPAMI*, 45(11):13344–13362, 2023. 8