

Object-Centric World Models for Causality-Aware Reinforcement Learning

Yosuke Nishimoto¹, Takashi Matsubara²

¹The University of Osaka

²Hokkaido University

u835711c@ecs.osaka-u.ac.jp, matsubara@ist.hokudai.ac.jp

Abstract

World models have been developed to support sample-efficient deep reinforcement learning agents. However, it remains challenging for world models to accurately replicate environments that are high-dimensional, non-stationary, and composed of multiple objects with rich interactions since most world models learn holistic representations of all environmental components. By contrast, humans perceive the environment by decomposing it into discrete objects, facilitating efficient decision-making. Motivated by this insight, we propose *Slot Transformer Imagination with CAusality-aware reinforcement learning* (STICA), a unified framework in which object-centric Transformers serve as the world model and causality-aware policy and value networks. STICA represents each observation as a set of object-centric tokens, together with tokens for the agent action and the resulting reward, enabling the world model to predict token-level dynamics and interactions. The policy and value networks then estimate token-level cause–effect relations and use them in the attention layers, yielding causality-guided decision-making. Experiments on object-rich benchmarks demonstrate that STICA consistently outperforms state-of-the-art agents in both sample efficiency and final performance.

1 Introduction

Deep reinforcement learning (RL) has achieved success in a variety of fields (Mnih et al. 2013; Lillicrap et al. 2015; Mnih et al. 2016; Schulman et al. 2017; Haarnoja et al. 2018), including robotics (Levine et al. 2016; Andrychowicz et al. 2019; Kalashnikov et al. 2018) and autonomous driving (Sallab et al. 2017; Isele et al. 2017; Kendall et al. 2019). However, achieving high performance requires extensive interactions with the environment, which is costly and inefficient in the real-world tasks when considering real-time operation and physical device failures (Moerland et al. 2023).

Humans understand and predict real-world dynamics through interaction with the environment (Sutton and Barto 1981; Friston 2010). Inspired by this mechanism, world models were proposed for Model-Based RL (MBRL) (Sutton 1990; Ha and Schmidhuber 2018). In this setting, agents train world models to replicate their environments (observations and actions) and optimize their policies within “imag-

ined” environments generated by the world models, making their learning sample-efficient. Earlier MBRL agents have adopted Recurrent Neural Networks (RNNs) as the dynamics model of the world model (Hafner et al. 2020, 2021, 2025; Kaiser et al. 2020; Deng, Jang, and Ahn 2022), whereas recent works have begun to explore the use of Transformers (Vaswani et al. 2017; Robine et al. 2023; Micheli, Alonso, and Fleuret 2023, 2024; Zhang et al. 2023; Burchi and Timofte 2025). Compared with RNNs, Transformers provide superior learning efficiency, generalization performance, and long-term prediction accuracy.

However, even with Transformers, world models still struggle to replicate environments that are high-dimensional, non-stationary, and composed of multiple objects with their interactions, while such environments are common in real-world applications, including service robots and autonomous driving. This is because world models learn holistic representations of the environments, which may fail to capture the important relationships and interactions between individual objects (Santoro et al. 2017). When placed in such environments, humans perceive the environment by decomposing it into discrete concepts such as objects and events, enabling more efficient and causality-aware decision-making (Spelke and Kinzler 2007). Incorporating these cognitive mechanisms into world models potentially allow RL agents to operate more effectively even in complex settings.

Motivated by this insight, we propose *Slot Transformer Imagination with CAusality-aware reinforcement learning* (STICA), depicted in Figure 1. This is an RL agent built upon a unified framework in which the world model, policy network, and value network are all implemented using object-centric Transformers, and both the policy and value networks explicitly leverage causal information for more structured and effective decision-making. The slot-based autoencoder extracts object-centric representations (z_t^1, \dots, z_t^n) from observations o_t while excluding static background information z_{BG} , as shown in Figure 1 (a) and (b). The Transformer-based dynamics model accepts these representations (z_t^1, \dots, z_t^n), along with agent actions a_t and obtained rewards r_{t-1} , as input tokens and predicts token-level dynamics and interactions. The policy and value networks estimate causal relationships G among the input tokens, thereby facilitating policy learning for causality-

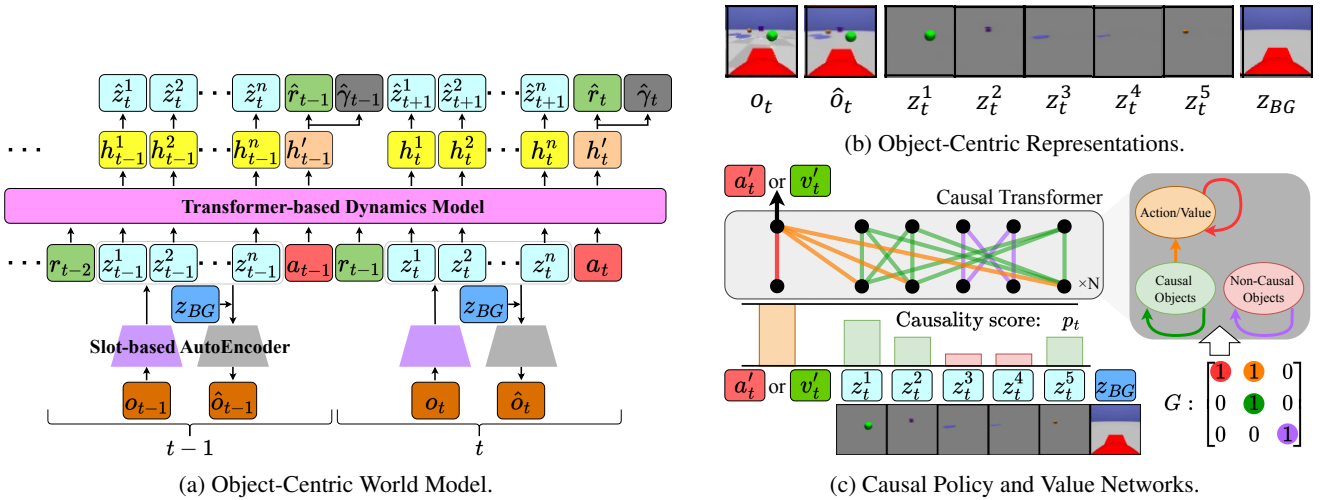


Figure 1: The architecture and object-centric representations of STICA. (a) Object-Centric World Model. Slot-based autoencoder extracts object-centric latent states (z_t^1, \dots, z_t^n) from observation o_t , excluding static background information z_{BG} at time t ($1 \leq t \leq T$). Transformer-based dynamics model computes hidden states ($h_{1:t}^1, \dots, h_{1:t}^n$) and $h'_{1:t}$ from latent states ($z_{1:t}^1, \dots, z_{1:t}^n$), actions $a_{1:t}$, and rewards $r_{1:t-1}$, followed by the multilayer perceptrons (MLPs) that predict the next latent states ($\hat{z}_{t+1}^1, \dots, \hat{z}_{t+1}^n$), the reward \hat{r}_t , and the discount factor $\hat{\gamma}_t$. (b) Examples of object-centric representations for Safety Gym benchmark task; the observation o_t , its reconstruction \hat{o}_t , the reconstructions from the extracted object-centric latent states (z_t^1, \dots, z_t^5), and that from the static background information z_{BG} . (c) Causal policy and value networks. They estimate causal relationships from the latent states to the action or value, based on a causal graph G and causality scores p_t^k , and adjust the attention weights within the Transformers accordingly, enabling the causality-aware decision-making. The latent states (z_t^1, z_t^2 , and z_t^5) of goal-related objects or obstacles are expected to have stronger causal influence on the target token (a_t^1 or v_t^1), while latent states (z_t^3 and z_t^4) of objects irrelevant to task completion have weaker causal influence.

aware decision-making, which is aware of ‘‘causality’’ in the sense of token-level dependency (not in the context of causal inference) (see Figure 1 (c)). The main contributions of this paper can be summarized as follows:

High-performance object-centric world model: We introduce STICA, a novel framework that significantly outperforms state-of-the-art agents on the Safety-Gym benchmark (Achiam and Amodei 2019), a first-person-view, non-stationary, object-rich, 3D environment, as well as on the Object-Centric Visual RL benchmark (Yoon et al. 2023), a fixed-view, 2D or 3D environment.

Causality-aware decision-making: We introduce the policy and value networks that explicitly leverage the token-level dependencies as the information of causal dependencies, enabling decision-making that accurately captures object properties. By visualizing the attention weights within these networks, we obtain the human-interpretable causal influences that STICA identifies and exploits to guide its actions.

End-to-end learning of object-centric representations: By removing static background information from object-centric representation, we enhance the feature-extraction capabilities of autoencoders, which have depended on supervised learning and pretraining in prior studies. Consequently, this modification enables precise, end-to-end learning of object-centric representations.

2 Related Work

2.1 Model-Based Reinforcement Learning

Model-based reinforcement learning (MBRL) agents build a world model that replicates the environmental dynamics from experience. By using this world model for planning, agents improve their generalization capabilities, sample efficiency, and decision-making performance. Dyna is a classic method that introduced this idea (Sutton 1990). It learns a model of the environment and uses it to refine rewards and value functions. Ha and Schmidhuber (2018) reintroduced the concept of a world model, employing neural networks to emulate the environment from high-dimensional visual inputs (i.e., images) successfully.

The Dreamer series is a set of MBRL agents that employ the Recurrent State-Space Model (RSSM) as world models, which learns deterministic and stochastic state transitions and allows robust prediction from observations (Hafner et al. 2019, 2020, 2021, 2025). This series trains world models by reconstructing observations and learns policies using an actor-critic method. DreamerV2 stabilizes representation learning by adopting categorical latent states and extends to complex, high-dimensional environments with discrete action spaces, such as Atari games (Hafner et al. 2021). DreamerV3 introduces several tricks to further improve the performance across diverse domains (Hafner et al. 2025). TD-MPC omitted the reconstructions but focused on the prediction of rewards (Hansen, Su, and Wang 2022). TD-MPC2 redesigned the architecture to stabilize the learning

procedure, improving its performance on challenging visual and continuous-control tasks under a fixed hyperparameter set (Hansen, Su, and Wang 2024).

Several recent works have employed Transformers in place of RSSMs for world models to overcome the long-term memory limitations (Vaswani et al. 2017). This approach was pioneered by TransDreamer (Chen et al. 2021) and advanced by subsequent works: TWM (Robine et al. 2023), STORM (Zhang et al. 2023), IRIS (Micheli, Alonso, and Fleuret 2023), and Δ -IRIS (Micheli, Alonso, and Fleuret 2024). TWISTER (Burchi and Timofte 2025) further uses Contrastive Predictive Coding (van den Oord, Li, and Vinyals 2018) to promote representation learning suited for long-term prediction.

These MBRL agents learn holistic representations of the environment, which may fail to capture the important relationships and interactions between individual objects (San-toro et al. 2017).

2.2 Object-Centric Representation Learning

Object-centric representation learning is an unsupervised method for decomposing a high-dimensional visual scene into individual objects and learning their individual representations. Most existing methods employed an autoencoder to extract multiple latent states, referred to as “slots”. These slot-based methods have been successfully applied to object detection and tracking in images and videos (Burgess et al. 2019; Locatello et al. 2020; Singh, Deng, and Ahn 2022; Seitzer et al. 2023; Jiang et al. 2023; Wu et al. 2023b; Singh, Schaub-Meyer, and Roth 2024; Kipf et al. 2022; Elsayed et al. 2022; Singh, Wu, and Ahn 2022; Wu et al. 2023b). This paradigm is also employed for video prediction, which learns the dynamics of individual objects to forecast their future movements and interactions. For instance, Transformer-based autoregressive models, such as SlotFormer (Wu et al. 2023a) and RSM (Nguyen et al. 2023), have demonstrated their robust generalization and impressive long-horizon prediction performance. SlotSSM (Jiang et al. 2024) employed a state-space model and showed further performance improvements.

2.3 Reinforcement Learning with Object-Centric Representation Learning

Some studies have proposed to combine model-free RL agents with object-centric representation learning (Zadaianchuk, Seitzer, and Martius 2021; Carvalho et al. 2021; Yi et al. 2022; Heravi et al. 2023; Yoon et al. 2023; Haramati, Daniel, and Tamar 2024). SMORL (Zadaianchuk, Seitzer, and Martius 2021) and OCARL (Yi et al. 2022) employ a Transformer to the policy network and operate it on object-centric representations. OCRL (Yoon et al. 2023) and EIT (Haramati, Daniel, and Tamar 2024) further employ a Transformer to the value network. While this enables decision-making with a focus on individual objects, due to the lack of world model, it does not explicitly leverage their dynamics and interactions.

Object-centric MBRL agents remain less developed, as they require additional resources to learn object-centric rep-

resentations. OODP (Zhu and Zhang 2018), COBRA (Watters et al. 2019a), OP3 (Veerapaneni et al. 2020), and SWM (Singh et al. 2021) require numerous episodes generated through random exploration (that is, random action selection without a policy or an estimated value). FOCUS (Ferraro et al. 2023) and OC-STORM (Zhang et al. 2025) are based on supervised learning. SOLD (Mosbach et al. 2025) is pre-trained on external datasets. Furthermore, most of these methods are not designed to handle non-stationary environments, where new objects appear or existing ones disappear. This limitation also implies a limited ability to operate in partially observable settings, including first-person views. Therefore, these agents face significant limitations in terms of generality and sample efficiency.

3 Method: STICA

We introduce STICA, an RL agent equipped with an object-centric world model and designed to make more structured and effective decisions by explicitly leveraging causal information about object interactions (see Figure 1). STICA is the first MBRL agent that extracts object-centric representations directly from observations without random exploration, supervision, or pretraining. Its architecture is composed of three components: an object-centric world model, a causal policy network, and a causal value network. The object-centric world model learns state transitions of latent states of individual objects extracted from visual observations, and simulates the environment to generate imaginary trajectories. The causal policy and value networks perform on these imaginary trajectories to maximize expected returns. The training process of STICA follows the standard one that most existing MBRL agents employed.

3.1 Object-Centric World Model

As illustrated in Figure 1 (a), our object-centric world model is composed of a slot-based autoencoder ϕ and a Transformer-based dynamics model ψ . This model is trained on batches of B sequences of n latent states ($z_{1:T}^1, \dots, z_{1:T}^n$), actions $a_{1:T}$, rewards $r_{1:T}$, and discount factors $\gamma_{1:T}$, where the subscripts $1 : T$ denote the sequences from $t = 1$ to $t = T$.

Slot-based AutoEncoder: The slot-based autoencoder ϕ is trained to extract object-centric representations and a background representation separately from observations through the reconstruction. Namely, it is composed of an encoder and a decoder:

$$\begin{aligned} \text{Encoder: } & p_{\phi}^z((z_t^1, \dots, z_t^n) | o_t), \\ \text{Decoder: } & p_{\phi}^{\hat{o}}(\hat{o}_t | (z_t^1, \dots, z_t^n), z_{BG}). \end{aligned} \quad (1)$$

Given an observation o_t , the encoder obtains n deterministic slots (s_t^1, \dots, s_t^n) using an attention mechanism called Slot Attention (Locatello et al. 2020). Each slot s_t^i serves as 128-dimensional logits to define 16 distinct categorical distributions of 8 categories, from which each latent state z_t^i is sampled as 16 sets of 8-dimensional one-hot vectors (Hafner et al. 2021). An additional learnable time-independent latent state z_{BG} is defined. This is expected to represent static

background information that is independent on time t , such as the environment layout and the agent’s body. This trick allows the encoder to extract the object-centric representations of individual objects separately from the static background information, which is crucial for the world model to capture the dynamics of individual objects. The decoder is implemented as a spatial broadcast decoder (Watters et al. 2019b), which reconstructs the individual RGB images ($\hat{o}_t^1, \dots, \hat{o}_t^n$) and the individual unnormalized masks (m_t^1, \dots, m_t^n) from the latent states (z_t^1, \dots, z_t^n). It also reconstructs the background \hat{o}_{BG} from the static background information z_{BG} , while the corresponding mask m_{BG} is filled with 0. A softmax function is applied to the masks ($m_t^1, \dots, m_t^n, m_{BG}$) for obtaining the normalized masks ($M_t^1, \dots, M_t^n, M_{BG}$).

Then, we mix the individual RGB images ($\hat{o}_t^1, \dots, \hat{o}_t^n$, \hat{o}_{BG}) with the normalized masks ($M_t^1, \dots, M_t^n, M_{BG}$) as weights, obtaining a single reconstruction \hat{o}_t ;

$$\hat{o}_t = \sum_{k=1}^n M_t^k \odot \hat{o}_t^k + M_{BG} \odot \hat{o}_{BG}. \quad (2)$$

Figure 1(b) visualizes examples of the reconstructed individual RGB images ($\hat{o}_t^1, \dots, \hat{o}_t^n, \hat{o}_{BG}$) for latent states ($z_t^1, \dots, z_t^n, z_{BG}$).

Loss for Slot-based AutoEncoder: The objective of the slot-based autoencoder ϕ is to minimize the following loss function \mathcal{L}_ϕ^{ae} :

$$\begin{aligned} \mathcal{L}_\phi^{ae} &= \mathbb{E}_B[\frac{1}{T} \sum_{t=1}^T (\mathcal{J}_{rec}^t + \alpha_1 \mathcal{J}_{ent}^t + \alpha_2 \mathcal{J}_{cross}^t)], \\ \text{where } \mathcal{J}_{rec}^t &= -\ln p_\phi(o_t | (z_t^1, \dots, z_t^n), z_{BG}), \\ \mathcal{J}_{ent}^t &= -\sum_{k=1}^n H(p_\phi^z(z_t^k | o_t)), \\ \mathcal{J}_{cross}^t &= \sum_{k=1}^n H(p_\phi^z(z_t^k | o_t), p_\psi^z(\hat{z}_t^k | h_{t-1}^k)), \end{aligned} \quad (3)$$

with hyperparameters $\alpha_1, \alpha_2 \geq 0$. The reconstruction error \mathcal{J}_{rec}^t ensures that the latent states (z_t^1, \dots, z_t^n) and z_{BG} contain sufficient information to reconstruct the observation o_t . The entropy term \mathcal{J}_{ent}^t is a regularization term that prevents the distribution of each latent state z_t^k from being deterministic. The cross-entropy term \mathcal{J}_{cross}^t is another regularization term that aligns each extracted latent state z_t^k with the latent state \hat{z}_t^k predicted by the Transformer-based dynamics model ψ . Here, as the slot encoder extracts the latent states (z_t^1, \dots, z_t^n) in a random order, the cross-entropy term \mathcal{J}_{cross}^t is computed after rearranging the indices $1, \dots, n$ to minimize the L1 distance between the logits of the categorical distributions for z_t^k and \hat{z}_t^k .

Transformer-based Dynamics Model: The Transformer-based dynamics model ψ predicts the future latent states ($z_{t+1}^1, \dots, z_{t+1}^n$), the reward r_t , and the discount factor γ_t from the past reward r_{s-1} , latent states (z_s^1, \dots, z_s^n), and action a_s for $s = t, t-1, t-2, \dots$. This model ψ is composed of the following modules:

Aggregation model:

$$(h_t^1, \dots, h_t^n, h'_t) = f_\psi^h(r_{1:t-1}, (z_{1:t}^1, \dots, z_{1:t}^n), a_{1:t}),$$

Latent state predictor: $p_\psi^z(\hat{z}_{t+1}^k | h_t^k)$ for $k = 1, \dots, n$,

Reward predictor: $p_\psi^r(\hat{r}_t | h'_t)$,

Discount predictor: $p_\psi^\gamma(\hat{\gamma}_t | h'_t)$,

(4)

where the aggregation model f_ψ^h is implemented using Transformer-XL (Dai et al. 2019), and the others p_ψ using multilayer perceptrons (MLPs).

The aggregation model f_ψ^h feeds the history of the reward $r_{1:t-1}$, latent states ($z_{1:t}^1, \dots, z_{1:t}^n$) and action $a_{1:t}$ as tokens and predicts the hidden states (h_t^1, \dots, h_t^n) and h'_t as the outputs of the current latent states (z_t^1, \dots, z_t^n) and action a_t in a deterministic manner. In this predictor, causal masking is applied to ensure self-attention layers not to access future time steps. The positional encoding depends only on time t and not on the indices $1, \dots, n$ of the latent states to ensure the equivariance to the order of the latent states.

The latent state predictor $p_\psi^z(\hat{z}_{t+1}^k | h_t^k)$ and the reward predictor $p_\psi^r(\hat{r}_t | h'_t)$ output a categorical distribution and a normal distribution through the reparameterization trick, respectively (Kingma and Welling 2014). The discount predictor $p_\psi^\gamma(\hat{\gamma}_t | h'_t)$ outputs a value in the range of $[0, 1]$.

Loss for Transformer-based Dynamics Model: The objective of the Transformer-based dynamics model ψ learns the transitions of the latent states (z_t^1, \dots, z_t^n) by minimizing the following loss function:

$$\begin{aligned} \mathcal{L}_\psi^{dyn} &= \mathbb{E}_B[\frac{1}{T} \sum_{t=1}^T (\mathcal{J}_{cross}^{t+1} + \beta_1 \mathcal{J}_{rew}^t + \beta_2 \mathcal{J}_{dis}^t)], \\ \text{where } \mathcal{J}_{cross}^{t+1} &= \sum_{k=1}^n H(p_\phi^z(z_{t+1}^k | o_{t+1}), p_\psi^z(\hat{z}_{t+1}^k | h_t^k)), \\ \mathcal{J}_{rew}^t &= -\ln(p_\psi^r(\hat{r}_t | h'_t)), \\ \mathcal{J}_{dis}^t &= -\ln(p_\psi^\gamma(\hat{\gamma}_t | h'_t)), \end{aligned} \quad (5)$$

with hyperparameters $\beta_1, \beta_2 \geq 0$, and $\gamma_t = 0$ for episode ends and $\gamma_t = \gamma$ otherwise. The cross-entropy term $\mathcal{J}_{cross}^{t+1}$ ensures that each predicted latent state \hat{z}_{t+1}^k is close to the extracted one z_{t+1}^k . This term $\mathcal{J}_{cross}^{t+1}$ is also computed after rearranging the indices of the latent states. The negative log-likelihoods \mathcal{J}_{rew}^t and \mathcal{J}_{dis}^t optimizes the predictions of the reward \hat{r}_t and discount factor $\hat{\gamma}_t$, respectively, where the discount factor $\hat{\gamma}_t$ is considered as a Bernoulli variable.

3.2 Policy and Value Networks with Causal Attention

Transformer-based Policy and Value Networks: STICA learns the policy on latent states ($\hat{z}_t^1, \dots, \hat{z}_t^n$) generated by our object-centric world model using an advantage actor-critic (A2C) method (Mnih et al. 2016). The advantage function is estimated using the generalized advantage estimator (GAE) (Schulman et al. 2016). The policy network $\pi_\theta(a_t | (z_t^1, \dots, z_t^n))$ and value network $V_\xi(v_t | (z_t^1, \dots, z_t^n))$ are implemented using two separate Transformers (see Figure 1 (c)). Each Transformer takes as input tokens the latent states (z_t^1, \dots, z_t^n) and an extra learnable token, whose corresponding output is used for the estimate of the target (action a_t or value v_t). All latent states (z_t^1, \dots, z_t^n) share the same positional encoding, ensuring the Transformer’s output to be independent of their order (in other words, making the output equivariant to the order of the latent states).

Causal Attention: An environment contains both causal objects, which are relevant for action selection and value es-




Figure 2: (Left) Eight tasks from the Safety Gym benchmark. Top panels show first-person views for training. Bottom panels show fixed-view images for reference and not for training. (Right) Three tasks from the OCVRL benchmark.

timation, and non-causal objects, which are not. Selecting an action requires attending to reward-relevant objects like goals and obstacles, while paying limited attention to others, such as the floor. Similarly, value estimation can ignore non-causal objects entirely. Attention mechanisms within the policy and value networks can leverage these causal relationships. We explicitly represent these causal relationships using a causal graph matrix G as follows:

$$G = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{cases} \text{policy or value} \\ \text{causal object} \\ \text{non-causal object} \end{cases} \quad (6)$$

Each element $G_{i,j}$ indicates a causal relationship from object j to object i , where the index takes 1 for the target (policy or value), 2 for a causal object, and 3 for a non-causal object. For example, $G_{1,2} = 1$ indicates the existence of a causal relationship from the causal objects to the target.

To leverage these causal relationships for the policy and value networks, we estimate the causality score $p_t^k \in [0, 1]$ for each latent state z_t^k at time step t by using an MLP, as shown in Figure 1 (c). The causality score p_t^k is a probability that the latent state z_t^k represents a causal object. Based on these estimates, we adjust the attention weights in the Transformers. Specifically, we define the weight matrix W_t as follows:

$$W_t = \begin{bmatrix} 1 & 0 & 0 \\ 0 & p_t^1 & 1 - p_t^1 \\ \vdots & \vdots & \vdots \\ 0 & p_t^n & 1 - p_t^n \end{bmatrix} \begin{cases} \text{policy or value} \\ \text{latent states} \end{cases}, \quad (7)$$

where the first row corresponds to the target token (policy or value), and the remaining rows correspond to the tokens for the latent states (z_t^1, \dots, z_t^n). Then, the matrix $W_t G W_t^\top$ represents the causal relationships among the tokens, where the (i, j) element of W_t represents the influence of the j -th token on the i -th token. Given this, we define the causal attention CA_t as follows:

$$CA_t = \text{Norm} \left(\text{softmax} \left(\frac{Q_t K_t^\top}{\sqrt{d}} \right) \odot W_t G W_t^\top \right) V_t. \quad (8)$$

Here, Q_t , K_t , and V_t denote the queries, keys, and values at time step t , respectively, and d is the dimensionality of the key vectors. \odot denotes element-wise multiplication, and $\text{Norm}(\cdot)$ denotes row-wise L1 normalization of the input

matrix. The matrix $W_t G W_t^\top$ scales the original attention weights $\text{softmax} \left(\frac{Q_t K_t^\top}{\sqrt{d}} \right)$ to reflect the causal relationships. Each attention layer in the Transformers of the policy and value networks uses this causal attention CA_t instead of the standard attention mechanism.

4 Experiments and Results

4.1 Benchmarks and Baselines

Safety Gym Benchmark: For validating the performance of STICA, we used eight object-rich 3D tasks from the Safety Gym benchmark (Achiam and Amodei 2019), as summarized in Figure 2 (left). The agent receives first-person images as observations o_t , exemplified in the top panels. For an overview of the environment layout, please refer to the fixed-view images in the bottom panels, while these are not used for experiments.

The tasks are classified into two types: Goal and Button. The Goal task is to reach a target (green cylinder) area among many static obstacles. After reaching the target, its location is randomly reset, while the rest of the layout is preserved. The Button task is to press a target (green) circular button among many other (yellow) circular buttons and dynamic cubic obstacles. Once the target button is pressed, another button is selected as a new target, with all other objects remaining fixed. The agent receives rewards for moving towards the current targets densely and for completing the tasks sparsely. We have two types of agents: Point and Car. The Point agent has separate actuators for rotation and forward/backward movement, while the Car agent is equipped with two independently-driven parallel wheels and a free-rolling rear wheel. Each task has two difficulty levels: Level 1 (less objects) and Level 2 (more objects). We measured the average undiscounted reward return, $\hat{J}(\pi) = \frac{1}{E} \sum_{i=1}^E \sum_{t=0}^{T_{\text{ep}}} r_t$, over $E = 10$ episodes of $T_{\text{ep}} = 1000$ steps, as described in Achiam and Amodei (2019).

While the Safety Gym benchmark was originally designed to evaluate the safeness of RL agents via both rewards and costs (i.e., failures), our goal is to develop an object-centric world model and an RL agent capable of causality-aware decision-making. Consequently, we use only the average undiscounted return as our performance metric. Nevertheless, the Safety Gym benchmark provides a suitable evaluation setting because its environments feature the key challenges: high-dimensional, partially observable

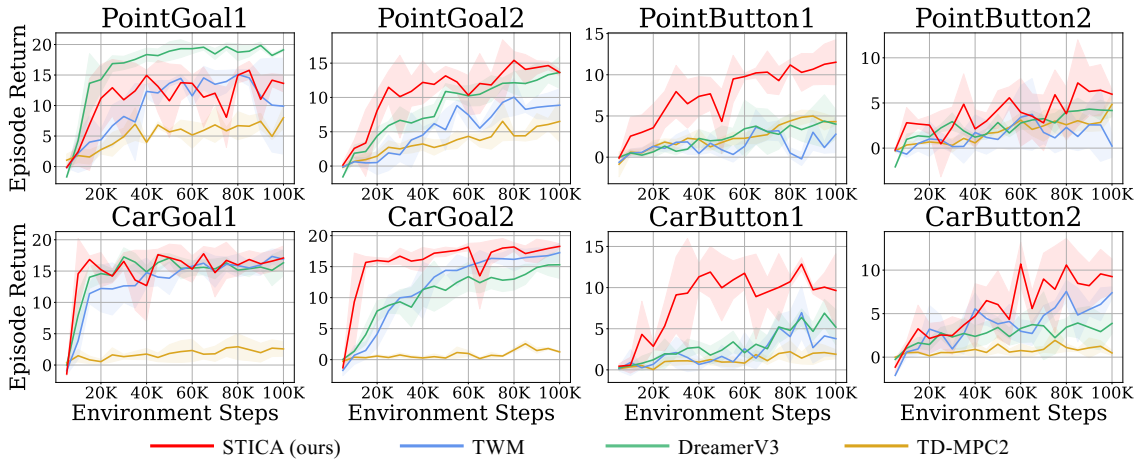


Figure 3: Training curves for Safety Gym benchmark.

Method	Point				Car				Mean	Normalized Mean
	Goal1	Goal2	Button1	Button2	Goal1	Goal2	Button1	Button2		
PPO	5.00	4.43	4.94	3.78	3.15	1.46	0.95	1.51	3.15	1.00
TWM	9.89	8.87	2.79	0.23	<u>16.89</u>	<u>17.25</u>	3.79	<u>7.40</u>	8.39	3.83
DreamerV3	19.13	13.64	4.01	4.16	16.32	15.30	<u>5.20</u>	3.87	<u>10.20</u>	<u>4.06</u>
TD-MPC2	8.00	6.50	4.31	4.84	2.57	1.24	1.90	0.48	3.73	1.15
STICA(ours)	<u>13.63</u>	13.64	11.52	5.97	17.09	18.27	9.65	9.25	11.90	5.49

Table 1: Agent Scores on Safety Gym benchmark.

first-person views; non-stationary targets; and rich interactions with many dynamic obstacles.

See Appendices A.1 for more details about benchmarks and B for the experimental settings.

Object-Centric Visual RL Benchmark: We also evaluated STICA on two 2D and one 3D tasks from the Object-Centric Visual RL (OCVRL) benchmark (Yoon et al. 2023). The agent receives fixed-view images as observations o_t , as exemplified in Figure 2 (right). Object Goal Task imposes the agent (red circle) to move towards a target object (blue square). If the agent touches other objects, it receives no reward, and the episode is terminated. Object Interaction Task requires the agent (red circle) to push a specific object (blue square) to the fixed goal position (blue square at the bottom left corner). In these tasks, the actions are to move the agent in four discrete directions. Object Reaching task is a 3D version of Object Goal Task, where the agent has a three-fingered robot arm and is tasked to touch the target object with one finger. The action adjusts the positions of three joints of the finger. In any tasks, the agent receives a sparse reward when it reaches the target object or pushes the target object to the goal position.

Although these environments are fully observable and stationary (i.e., no objects appear or disappear), we use this benchmark to demonstrate the generality of STICA.

Baselines: We also evaluated three MBRL agents as baselines. TWM is a Transformer-based world model that learns holistic state transitions (Robine et al. 2023). DreamerV3 (Hafner et al. 2025) and TD-MPC2 (Hansen, Su, and

Wang 2024) are state-of-the-art MBRL agents that also learn holistic representations of the environment. Because TD-MPC2 is specifically designed for tasks with continuous action spaces, we omitted it from the OCVRL benchmark’s 2D tasks. See Appendix A.2 for more details about baselines.

4.2 Results

Safety Gym Benchmark: We summarized the time courses of the reward returns for the Safety Gym benchmark in Figure 3. Solid lines denote the averages over three runs, while shaded areas denote the standard deviations. STICA achieved the best performance in all tasks except for PointGoal1 in terms of both the final performance and the learning speed. See also the final average returns in Table 1, where bold and underlined numbers indicate the best and the second best performances, respectively. Normalized mean is the average after being normalized against the performance by Proximal Policy Optimization (PPO), as suggested by the benchmark (Achiam and Amodei 2019). We attribute STICA’s superior performance to the object-centric and causality-aware modules as verified later. PointGoal1 is a simpler task with fewer and less diverse objects than others, where holistic representation may be sufficient, and STICA did not show a significant advantage.

Please refer to Appendix C for other visualizations.

Object-Centric Visual RL Benchmark: Figure 4 shows the training curves for the OCVRL benchmark, and Table 2 summarizes the final success rates averaged over three runs. STICA achieves superior performance on all tasks. In particular, STICA achieved a drastically better result on Object

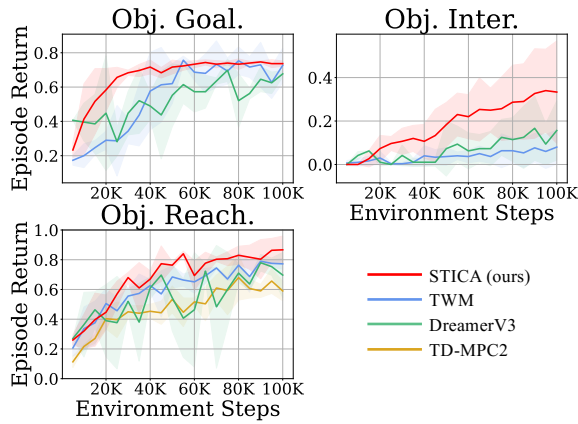


Figure 4: Training curves for the OCVRL benchmark.

Method	Obj. Goal.	Obj. Inter.	Obj. Reach.
TWM	0.727	0.080	0.772
DreamerV3	0.677	0.156	0.697
TD-MPC2	-	-	0.590
STICA(ours)	0.737	0.333	0.867

Table 2: Final success rates on the OCVRL benchmark.

Interaction Task, where the agent directly interacts with objects. In such case, the object-centric world model and the object-centric causal value network act synergistically.

4.3 Ablation Study

To assess the contributions of the modules in STICA, we conducted an ablation study on two tasks (PointButton1 and CarButton1) from the Safety Gym benchmark, and summarize the results in Figure 5. We refer to the use of the static background information z_{BG} as the background removal, as it removes the background information from the object-centric representations. We consider the following variants: STICA w/o CA, which omits the causal attention mechanism; STICA w/o BR, which omits the background removal; STICA w/o CA+BR, which omits both; and STICA w/o CA+BR+TP+TV, which further replaces the Transformer-based policy and value networks with their original architectures. A variant that omits object-centric representations effectively corresponds to TWM.

We observe that adopting an object-centric world model alone yields only a marginal performance improvement (STICA w/o CA+BR+TP+TV vs. TWM). The object-centric Transformers for the policy and value networks produce a noticeable gain in performance (STICA w/o CA+BR vs. STICA w/o CA+BR+TP+TV). The background removal leads to a modest improvement (STICA vs. STICA w/o BR, or STICA w/o CA vs. STICA w/o CA+BR). The introduction of causal attention yields a substantial performance increase (STICA vs. STICA w/o CA).

4.4 Visualization

We also examined the contributions of causal attention. From the policy and value networks, we obtained the av-




Figure 5: Ablation studies on the PointButton1 and CarButton1 from the Safety Gym benchmark.




Figure 6: Visualization of the attention weights in the policy and value networks.

erage attention weight from the token for each latent state z_t^k to the target token (action a_t^k or value v_t^k) in the attention layers, denoted by A_t^k . For visibility, we visualized the average attention weights A_t^k in the regions of the corresponding normalized masks M_t^k , that is, $\sum_{k=1}^n M_t^k A_t^k$, in Figure 6.

We observed that STICA allocates strong and selective attention to task-relevant (that is, causal) objects in both networks, compared to STICA w/o CA or w/o BR. In particular, the value network in STICA attends exclusively to the reward-yielding target objects (i.e., the green button or green cylinder area), with no appreciable attention paid to other objects. The policy network also focuses on the target object, but it pays some attention to the other objects, which is necessary for action selection. In contrast, the policy and value networks in STICA w/o CA fail to focus on the target object, and instead, they pay much attention to various objects in the environment. Moreover, the policy and value networks in STICA w/o BR also pay much attention to the background, which is completely irrelevant to the task. Overall, STICA is able to allocate selective attention to objects related to tasks by using causal attention and background removal.

5 Conclusion

In this paper, we introduced STICA, a novel RL agent with an object-centric world model and causality-aware policy and value networks. By decomposing observations into individual objects while removing static background information, STICA learns object-centric representations and their individual dynamics from scratch. Its causal attention enables the policy and value networks to focus on task-relevant objects, leading to more structured and effective decision-making. Through these modules, STICA significantly outperforms state-of-the-art RL agents on object-rich benchmarks in both sample efficiency and final performance.

Acknowledgments

This study was partly supported by JST PRESTO (JP-MJPR24TB, JPMJPR21C7), CREST (JPMJCR1914), ASPIRE (JPMJAP2329), Moonshot R&D (JPMJMS2033-14), and JSPS KAKENHI (24K15105), Japan, and achieved through the use of large-scale computer systems at the D3Center, Osaka University.

References

- Achiam, J.; and Amodei, D. 2019. Benchmarking Safe Exploration in Deep Reinforcement Learning.
- Andrychowicz, O. M.; Baker, B.; Chociej, M.; Józefowicz, R.; McGrew, B.; Pachocki, J.; Petron, A.; Plappert, M.; Powell, G.; Ray, A.; Schneider, J.; Sidor, S.; Tobin, J.; Welinder, P.; Weng, L.; and Zaremba, W. 2019. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1): 3–20.
- As, Y.; Usmanova, I.; Curi, S.; and Krause, A. 2022. Constrained Policy Optimization via Bayesian World Models. In *ICLR*.
- Burchi, M.; and Timofte, R. 2025. Learning Transformer-based World Models with Contrastive Predictive Coding. In *ICLR*.
- Burgess, C. P.; Matthey, L.; Watters, N.; Kabra, R.; Higgins, I.; Botvinick, M. M.; and Lerchner, A. 2019. MONet: Unsupervised Scene Decomposition and Representation. *ArXiv*, abs/1901.11390.
- Carvalho, W.; Liang, A.; Lee, K.; Sohn, S.; Lee, H.; Lewis, R.; and Singh, S. 2021. Reinforcement Learning for Sparse-Reward Object-Interaction Tasks in a First-person Simulated 3D Environment. In *IJCAI*.
- Chen, C.; Yoon, J.; Wu, Y.-F.; and Ahn, S. 2021. TransDreamer: Reinforcement Learning with Transformer World Models. In *Deep RL Workshop NeurIPS 2021*.
- Dai, Z.; Yang, Z.; Yang, Y.; Carbonell, J.; Le, Q.; and Salakhutdinov, R. 2019. Transformer-XL: Attentive Language Models beyond a Fixed-Length Context. In *ACL*.
- Deng, F.; Jang, I.; and Ahn, S. 2022. DreamerPro: Reconstruction-Free Model-Based Reinforcement Learning with Prototypical Representations.
- Elsayed, G. F.; Mahendran, A.; van Steenkiste, S.; Greff, K.; Mozer, M. C.; and Kipf, T. 2022. SAVi++: Towards End-to-End Object-Centric Learning from Real-World Videos. In *NeurIPS*.
- Ferraro, S.; Mazzaglia, P.; Verbelen, T.; and Dhoedt, B. 2023. FOCUS: Object-Centric World Models for Robotics Manipulation. *arXiv:2307.02427*.
- Friston, K. J. 2010. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11: 127–138.
- Ha, D.; and Schmidhuber, J. 2018. Recurrent World Models Facilitate Policy Evolution. In *NeurIPS*.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic Algorithms and Applications. *ArXiv*, abs/1812.05905.
- Hafner, D.; Lillicrap, T.; Ba, J.; and Norouzi, M. 2020. Dream to Control: Learning Behaviors by Latent Imagination. In *ICLR*.
- Hafner, D.; Lillicrap, T.; Fischer, I.; Villegas, R.; Ha, D.; Lee, H.; and Davidson, J. 2019. Learning Latent Dynamics for Planning from Pixels. In *ICML*.
- Hafner, D.; Lillicrap, T. P.; Norouzi, M.; and Ba, J. 2021. Mastering Atari with Discrete World Models. In *ICLR*.
- Hafner, D.; Pasukonis, J.; Ba, J.; and Lillicrap, T. 2025. Mastering diverse control tasks through world models. *Nature*, 640(8059): 647–653.
- Hansen, N.; Su, H.; and Wang, X. 2024. TD-MPC2: Scalable, Robust World Models for Continuous Control. In *ICLR*.
- Hansen, N. A.; Su, H.; and Wang, X. 2022. Temporal Difference Learning for Model Predictive Control. In *ICML*.
- Haramati, D.; Daniel, T.; and Tamar, A. 2024. Entity-Centric Reinforcement Learning for Object Manipulation from Pixels. In *ICLR*.
- Heravi, N.; Wahid, A.; Lynch, C.; Florence, P.; Armstrong, T.; Tompson, J.; Sermanet, P.; Bohg, J.; and Dwibedi, D. 2023. Visuomotor Control in Multi-Object Scenes Using Object-Aware Representations. In *ICRA*.
- Huang, W.; Ji, J.; Xia, C.; Zhang, B.; and Yang, Y. 2024. SafeDreamer: Safe Reinforcement Learning with World Models. In *ICLR*.
- Isele, D.; Rahimi, R.; Cosgun, A.; Subramanian, K.; and Fujimura, K. 2017. Navigating Occluded Intersections with Autonomous Vehicles using Deep Reinforcement Learning. *arXiv:1705.01196*.
- Jiang, J.; Deng, F.; Singh, G.; and Ahn, S. 2023. Object-Centric Slot Diffusion. In *NeurIPS*.
- Jiang, J.; Deng, F.; Singh, G.; Lee, M.; and Ahn, S. 2024. Slot State Space Models. In *NeurIPS*.
- Kaiser, Ł.; Babaeizadeh, M.; Miłos, P.; Osiński, B.; Campbell, R. H.; Czechowski, K.; Erhan, D.; Finn, C.; Kozakowski, P.; Levine, S.; Mohiuddin, A.; Sepassi, R.; Tucker, G.; and Michalewski, H. 2020. Model Based Reinforcement Learning for Atari. In *ICLR*.
- Kalashnikov, D.; Irpan, A.; Pastor, P.; Ibarz, J.; Herzog, A.; Jang, E.; Quillen, D.; Holly, E.; Kalakrishnan, M.; Vanhoucke, V.; and Levine, S. 2018. QT-Opt: Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation. *arXiv:1806.10293*.
- Kendall, A.; Hawke, J.; Janz, D.; Mazur, P.; Reda, D.; Allen, J.-M.; Lam, V.-D.; Bewley, A.; and Shah, A. 2019. Learning to Drive in a Day. In *2019 ICRA*.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- Kipf, T.; Elsayed, G. F.; Mahendran, A.; Stone, A.; Sabour, S.; Heigold, G.; Jonschkowski, R.; Dosovitskiy, A.; and Greff, K. 2022. Conditional Object-Centric Learning from Video. In *ICLR*.
- Levine, S.; Finn, C.; Darrell, T.; and Abbeel, P. 2016. End-to-end training of deep visuomotor policies. *J. Mach. Learn. Res.*, 17(1): 1334–1373.

- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N. M. O.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *CoRR*, abs/1509.02971.
- Locatello, F.; Weissenborn, D.; Unterthiner, T.; Mahendran, A.; Heigold, G.; Uszkoreit, J.; Dosovitskiy, A.; and Kipf, T. 2020. Object-Centric Learning with Slot Attention. In *NeurIPS*.
- Micheli, V.; Alonso, E.; and Fleuret, F. 2023. Transformers are Sample-Efficient World Models. In *The Eleventh ICLR*.
- Micheli, V.; Alonso, E.; and Fleuret, F. 2024. Efficient world models with context-aware tokenization. In *ICML*.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T. P.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous Methods for Deep Reinforcement Learning. In *ICML*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing Atari with Deep Reinforcement Learning.
- Moerland, T. M.; Broekens, J.; Plaat, A.; and Jonker, C. M. 2023.
- Mosbach, M.; Ewertz, J. N.; Villar-Corrales, A.; and Behnke, S. 2025. SOLD: Slot Object-Centric Latent Dynamics Models for Relational Manipulation Learning from Pixels. arXiv:2410.08822.
- Nguyen, T.; Mansouri, A.; Madan, K.; Khuong, N. D.; Ahuja, K.; Liu, D.; and Bengio, Y. 2023. Reusable Slotwise Mechanisms. In *NeurIPS*.
- Robine, J.; Höftmann, M.; Uelwer, T.; and Harmeling, S. 2023. Transformer-based World Models Are Happy With 100k Interactions. In *ICLR*.
- Sallab, A. E.; Abdou, M.; Perot, E.; and Yogamani, S. 2017. Deep Reinforcement Learning framework for Autonomous Driving. *Electronic Imaging*, 29(19): 70–76.
- Santoro, A.; Raposo, D.; Barrett, D. G.; Malinowski, M.; Pascanu, R.; Battaglia, P.; and Lillicrap, T. 2017. A simple neural network module for relational reasoning. In *NeurIPS*.
- Schulman, J.; Moritz, P.; Levine, S.; Jordan, M.; and Abbeel, P. 2016. High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *ICLR*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *CoRR*, abs/1707.06347.
- Seitzer, M.; Horn, M.; Zadaianchuk, A.; Zietlow, D.; Xiao, T.; Simon-Gabriel, C.-J.; He, T.; Zhang, Z.; Schölkopf, B.; Brox, T.; and Locatello, F. 2023. Bridging the Gap to Real-World Object-Centric Learning. In *ICLR*.
- Singh, G.; Deng, F.; and Ahn, S. 2022. Illiterate DALL-E Learns to Compose. In *ICLR*.
- Singh, G.; Peri, S.; Kim, J.; Kim, H.; and Ahn, S. 2021. Structured World Belief for Reinforcement Learning in POMDP. In *ICML*.
- Singh, G.; Wu, Y.-F.; and Ahn, S. 2022. Simple Unsupervised Object-Centric Learning for Complex and Naturalistic Videos. In *NeurIPS*.
- Singh, K.; Schaub-Meyer, S.; and Roth, S. 2024. Guided Latent Slot Diffusion for Object-Centric Learning. *ArXiv*, abs/2407.17929.
- Spelke, E. S.; and Kinzler, K. D. 2007. Core knowledge. *Developmental Science*, 10(1): 89–96.
- Sutton, R. S. 1990. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *ICML*.
- Sutton, R. S.; and Barto, A. G. 1981. An Adaptive Network That Constructs and Uses an Internal Model of its World. *Cognition and Brain Theory*, 4: 217–246.
- van den Oord, A.; Li, Y.; and Vinyals, O. 2018. Representation Learning with Contrastive Predictive Coding. *CoRR*, abs/1807.03748.
- Vaswani, A.; Shazeer, N. M.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *NeurIPS*.
- Veerapaneni, R.; Co-Reyes, J. D.; Chang, M.; Janner, M.; Finn, C.; Wu, J.; Tenenbaum, J.; and Levine, S. 2020. Entity Abstraction in Visual Model-Based Reinforcement Learning. In *CoRL*.
- Watters, N.; Matthey, L.; Bosnjak, M.; Burgess, C. P.; and Lerchner, A. 2019a. COBRA: Data-Efficient Model-Based RL through Unsupervised Object Discovery and Curiosity-Driven Exploration. arXiv:1905.09275.
- Watters, N.; Matthey, L.; Burgess, C. P.; and Lerchner, A. 2019b. Spatial Broadcast Decoder: A Simple Architecture for Learning Disentangled Representations in VAEs. *ArXiv*, abs/1901.07017.
- Wu, Z.; Dvornik, N.; Greff, K.; Kipf, T.; and Garg, A. 2023a. SlotFormer: Unsupervised Visual Dynamics Simulation with Object-Centric Models. In *ICLR*.
- Wu, Z.; Hu, J.; Lu, W.; Gilitshenski, I.; and Garg, A. 2023b. SlotDiffusion: Object-Centric Generative Modeling with Diffusion Models. In *NeurIPS*.
- Yi, Q.; Zhang, R.; Peng, S.; Guo, J.; Hu, X.; Du, Z.; Zhang, X.; Guo, Q.; and Chen, Y. 2022. Object-Category Aware Reinforcement Learning. In *NeurIPS*.
- Yoon, J.; Wu, Y.-F.; Bae, H.; and Ahn, S. 2023. An Investigation into Pre-Training Object-Centric Representations for Reinforcement Learning. In *ICML*.
- Zadaianchuk, A.; Seitzer, M.; and Martius, G. 2021. Self-supervised Visual Reinforcement Learning with Object-centric Representations. In *ICLR*.
- Zhang, W.; Jolley, A.; McInroe, T.; and Storkey, A. 2025. Objects matter: object-centric world models improve reinforcement learning in visually complex environments. arXiv:2501.16443.
- Zhang, W.; Wang, G.; Sun, J.; Yuan, Y.; and Huang, G. 2023. STORM: Efficient Stochastic Transformer based World Models for Reinforcement Learning. In *NeurIPS*.
- Zhu, G.; and Zhang, C. 2018. Object-Oriented Dynamics Predictor. In *NeurIPS*.