
Improving Generative Imagination in Object-Centric World Models

Zhixuan Lin^{1,2*} Yi-Fu Wu¹ Skand Peri¹ Bofeng Fu^{1,3*} Jindong Jiang¹ Sungjin Ahn^{1,4}

Abstract

The remarkable recent advances in object-centric generative world models raise a few questions. First, while many of the recent achievements are indispensable for making a general and versatile world model, it is quite unclear how these ingredients can be integrated into a unified framework. Second, despite using generative objectives, abilities for object detection and tracking are mainly investigated, leaving the crucial ability of temporal imagination largely under question. Third, a few key abilities for more faithful temporal imagination such as multimodal uncertainty and situation-awareness are missing. In this paper, we introduce Generative Structured World Models (G-SWM). The G-SWM achieves the versatile world modeling not only by unifying the key properties of previous models in a principled framework but also by achieving two crucial new abilities, multimodal uncertainty and situation-awareness. Our thorough investigation on the temporal generation ability in comparison to the previous models demonstrates that G-SWM achieves the versatility with the best or comparable performance for all experiment settings including a few complex settings that have not been tested before. <https://sites.google.com/view/gswm>

1. Introduction

Endowing machines with the ability to learn world models without supervision is a grand challenge toward human-like AI (Johnson-Laird, 1983; Ha & Schmidhuber, 2018b). A promising approach to this is to develop structured generative models in which world states are represented by composition of abstract entities such as objects and agents, and their relationships. This abstraction grounded in the

*Work done while visiting Rutgers University ¹Rutgers University ²Zhejiang University ³Tianjin University ⁴Rutgers Center for Cognitive Science. Correspondence to: Zhixuan Lin <zxlin.cs@gmail.com>, Sungjin Ahn <sjn.ahn@gmail.com>.

Table 1. Summary of abilities addressed in the previous models. In the table, ‘Uncertainty’ and ‘Situating’ stand for multimodal uncertainty and situation-awareness, respectively.

	SQAIR	SILOT	SCALOR	STOVE
DISCOVERY	✓	✓	✓	×
INTERACTION	×	×	×	✓
OCCCLUSION	✓	✓	✓	×
SCALABILITY	×	✓	✓	×
BACKGROUND	×	×	✓	×
UNCERTAINTY	×	×	×	×
SITUATED	×	×	×	×

physical properties (Spelke et al., 1993) can help enable interpretability, reasoning, and optimal decision-making via efficient simulation (a.k.a., imagination) of the possible futures—a crucial ability for model-based decision making in both humans (Schacter et al., 2012; Addis & Schacter, 2012) and AI systems (Silver et al., 2017).

In this direction, there have been remarkable advances recently in approaches using spatial attention for unsupervised object detection, representation, and generation. These include the object propagation-discovery model in SQAIR (Kosiorok et al., 2018), the physical interaction models in STOVE (Kossen et al., 2019; van Steenkiste et al., 2018), the scalability models in SCALOR (Jiang et al., 2020) and SILOT (Crawford & Pineau, 2020), and the background context models in SCALOR. Although each of these abilities can be an indispensable ingredients toward the future, there are also a few key limitations and challenges that should be addressed.

A primary challenge is how we can integrate the benefits of such fragmental advances into a unified model, e.g. as shown in Rainbow (Hessel et al., 2018). While all of the above-mentioned properties are essential to learning a complex world model, the models are developed with their own goals in controlled and isolated settings where many abilities other than the main one are ignored (See Table 1). For instance, although it is usual for an object to occlude some objects while interacting with another, joint modeling of these two properties have not been investigated properly. It is, therefore, crucial to develop a unified model while maintaining a principled modeling framework without losing existing benefits.

Another significant problem is that, despite being trained with generative objectives, in many of the previous models, the generation abilities are not well investigated. Instead, object-centric representation (via unsupervised detection) and tracking were focused as the main task in these models (Kosiorrek et al., 2018; Crawford & Pineau, 2020; Jiang et al., 2020). Hence, it is largely under question how well these models can actually simulate or imagine the possible futures, which is the central purpose of world models.

Lastly, a few key abilities required to enable faithful simulation quality have been missing. Among such abilities are multimodal uncertainty and situated behavior. Multimodal uncertainty is crucial in model-based decision-making as the simulation should be capable of exploring various possibilities. Although previous models contain stochastic latent variables, being based on Gaussian modeling, their ability to explore diverse scenarios is significantly limited to stochastic unimodal search. In addition, agents should show situated behavior depending on, e.g., where it is placed in the environment. However, previous models do not support such situation-aware behavior, e.g. an agent navigating through the corridor of a maze.

In this paper, we propose a new generative world model, called Generative Structured World Models (G-SWM), for unsupervised learning of object-centric state representation and efficient future simulation. G-SWM not only unifies the key abilities of previous models in a principled framework but also achieves multimodal uncertainty and situated behavior. The contributions of the paper are as follows. First, we introduce a model with a versatile propagation module that can handle occlusion of objects as well as interactions both among objects and the environment–situation. Second, we incorporate a hierarchical object dynamics model, enabling the simulation of stochastic dynamics and multimodal behavior. Lastly, we thoroughly investigate the generation performance of G-SWM and previous models through a series of extensive experiments designed to highlight the abilities and limitations of the various models.

2. Object-Centric World Models

Our proposed model inherits its base architecture from SCALOR (Jiang et al., 2020) as it supports both foreground and background models (and more generally from object-centric world models based on spatial-attention (Kosiorrek et al., 2018; Jiang et al., 2020; Crawford & Pineau, 2020)). In this section, we introduce the problem setting, basic architecture shared in such models, and notations.

Video, Objects, and Contexts. Consider a video $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ of length T describing a dynamic scene containing objects interacting with each other. We consider a generative process in which each object k in the scene at

timestep t is generated from a latent variable $\dot{\mathbf{z}}_t^k$ and everything else (i.e. non-object related) is modeled by a context latent variable $\mathbf{z}_t^{\text{ctx}}$. We assume that both objects and context are dynamic. We use $\dot{\mathbf{z}}_t$ to denote the set of object latents $\{\dot{\mathbf{z}}_t^k\}_k$ at t . All latent variables at frame t are denoted by $\mathbf{z}_t = \dot{\mathbf{z}}_t \cup \mathbf{z}_t^{\text{ctx}}$.

Object Representation. The object latent variable $\dot{\mathbf{z}}_t^k$ usually consists of object *attribute* latent variables. With some abuse of notation, we denote this by $\mathbf{z}_t^{\text{att},k} = \{\mathbf{z}_t^{\alpha,k}\}_{\alpha \in \mathcal{A}}$ with \mathcal{A} the set of attribute index. Different models may use different attribute sets. In our model, we use the following attribute set $\mathbf{z}_t^{\text{att},k} = \{\mathbf{z}_t^{\text{what},k}, \mathbf{z}_t^{\text{where},k}, \mathbf{z}_t^{\text{pres},k}, \mathbf{z}_t^{\text{depth},k}\}$ to represent appearance, position, presence, and depth of an object, respectively. Unlike previous models, our model assumes that these attributes of an object are generated from the *state* latent variable $\mathbf{z}_t^{\text{state},k}$. Thus, $\dot{\mathbf{z}}_t^k = \{\mathbf{z}_t^{\text{att},k}, \mathbf{z}_t^{\text{state},k}\}$. The state latent $\mathbf{z}_t^{\text{state},k}$ along with an RNN connecting it across the timesteps is used to explicitly represent object dynamics and to model multimodal uncertainty.

The attribute latents $\mathbf{z}_t^{\text{att},k}$ represent randomness of its target object properties but may not be in a form useful for rendering. For example, it is more efficient to update the object position by learning the position *deviation* instead of directly predicting the absolute position. To distinguish between the attribute latent and the explicit attribute value (in an explicit form for rendering), we introduce the notation $\mathbf{o}_t^k = \{\mathbf{o}_t^{\text{pres},k}, \mathbf{o}_t^{\text{depth},k}, \mathbf{o}_t^{\text{where},k}, \mathbf{o}_t^{\text{what},k}\}$ to denote explicit attributes. The explicit attribute is a deterministic function of the attribute latent, e.g. $\mathbf{o}_t^{\text{where},k} = \mathbf{o}_{t-1}^{\text{where},k} + c \cdot \tanh(\mathbf{z}_t^{\text{where},k})$. These attributes have the same interpretation as previous related works (Jiang et al., 2020; Crawford & Pineau, 2020).

Propagation and Discovery. Using the propagation-discovery model (Kosiorrek et al., 2018; Jiang et al., 2020; Crawford & Pineau, 2020), we denote the object latents propagated from frame $t - 1$ to frame t by $\tilde{\mathbf{z}}_t$ and object latents discovered at frame t by $\bar{\mathbf{z}}_t$. Thus $\dot{\mathbf{z}}_t = \tilde{\mathbf{z}}_t \cup \bar{\mathbf{z}}_t$. To discover new objects, each frame \mathbf{x}_t is divided into a grid of $H \times W$ cells, and each cell is associated with a discovery object latent variable $\bar{\mathbf{z}}_t^k$. During inference, these latents are inferred in parallel with a fully convolutional network (Long et al., 2015). At the end of each timestep, we select a total of K objects \mathbf{o}_t with the highest presence from the union of the discovered objects and propagated objects $\tilde{\mathbf{o}}_t \cup \bar{\mathbf{o}}_t$. These objects are also the ones to be propagated to the next timestep. In this work, we use fixed prior distributions for discovered latents, so we only do discovery during inference.

Rendering. To render the foreground μ_t^{fg} , we first reconstruct individual objects using $\mathbf{o}_t^{\text{what}}$ and combine them with $\mathbf{o}_t^{\text{where}}$, $\mathbf{o}_t^{\text{pres}}$, and $\mathbf{o}_t^{\text{depth}}$. A foreground mask α_t combining individual object masks is also obtained. The background μ_t^{bg} is decoded from the context latent $\mathbf{z}_t^{\text{ctx}}$. The final recon-

struction is $\boldsymbol{\mu}_t = \boldsymbol{\mu}_t^{\text{fg}} + (1 - \alpha_t)\boldsymbol{\mu}_t^{\text{bg}}$, and the likelihood function $p_\theta(\mathbf{x}_t | \dot{\mathbf{z}}_{\leq t}, \mathbf{z}_{\leq t}^{\text{ctx}})$ is a pixel-wise Gaussian distribution $\mathcal{N}(\mathbf{x}_t | \boldsymbol{\mu}_t, \sigma^2 \mathbf{I})$ where σ is a hyperparameter.

3. Generative Structured World Models

In this section, we describe our proposed model, Generative Structured World Models (G-SWM). We first describe the probabilistic formulation of the model. Then, we describe the propagation module which contains the main contribution of the proposed model. Finally, we describe the inference and learning method.

3.1. Probabilistic Modeling

As a temporal generative model, G-SWM consists of the following four generation modules: (i) the *context* module generates a new context representation from its history, (ii) the *propagation* module updates the attributes of currently existing objects for the next timestep, (iii) the *discovery* module generates new objects, and (iv) the *rendering* module renders objects and the background into a canvas to generate the target image. The generative process of G-SWM is illustrated in Figure 1.

Specifically, the generative process of a video \mathbf{x} of length T can be described by the following joint distribution:

$$p_\theta(\mathbf{x}_{1:T}, \mathbf{z}_{1:T}) = \prod_{t=1}^T p_\theta(\mathbf{x}_t, \bar{\mathbf{z}}_t, \tilde{\mathbf{z}}_t, \mathbf{z}_t^{\text{ctx}} | \mathbf{z}_{<t}) \quad (1)$$

and the single timestep model $p_\theta(\mathbf{x}_t, \bar{\mathbf{z}}_t, \tilde{\mathbf{z}}_t, \mathbf{z}_t^{\text{ctx}} | \mathbf{z}_{<t})$ is further decomposed into:

$$\underbrace{p_\theta(\mathbf{x}_t | \dot{\mathbf{z}}_{\leq t}, \mathbf{z}_{\leq t}^{\text{ctx}})}_{\text{Rendering}} \underbrace{p(\bar{\mathbf{z}}_t)}_{\text{Discovery}} \underbrace{p_\theta(\tilde{\mathbf{z}}_t | \dot{\mathbf{z}}_{\leq t}, \mathbf{z}_{\leq t}^{\text{ctx}})}_{\text{Propagation}} \underbrace{p_\theta(\mathbf{z}_t^{\text{ctx}} | \mathbf{z}_{<t}^{\text{ctx}})}_{\text{Context}} .$$

The main contribution of the G-SWM model is implemented in the propagation and context module. Other modules are similar to the ones proposed in SCALOR and SILOT. Thus, in the following, our description of the model focuses the context and propagation modules. For the implementation of the discovery and rendering modules, refer to Appendix.

3.2. Context Modeling

The context latent variable $\mathbf{z}_t^{\text{ctx}}$ represents the state of global context, i.e., the background, that may affect the dynamics of the objects as well as the appearance of the background. For generality, we assume a dynamic context. The dynamics of the context is modeled by the conditional distribution $p_\theta(\mathbf{z}_t^{\text{ctx}} | \mathbf{z}_{<t}^{\text{ctx}})$ which is a Gaussian distribution. This is implemented with an RNN followed by an Multi-layer Perceptron (MLP) yielding the parameters of the Gaussian distribution.

3.3. Versatile Propagation

The core of our model is the propagation module which we call the versatile propagation (V-Prop) module. This module integrates diverse key abilities of existing object-centric world models into a unified framework as well as implementing new abilities such as situated behaviour and multimodal uncertainty. As a result, the V-Prop module can support complex object interaction-occlusion, situated behaviour, and multimodal uncertainty jointly. In this section, we describe how these abilities are implemented in our model.

Object-State RNN. The backbone of the V-Prop module is the object-state RNN (OS-RNN) associated with each object (Figure 1). The object-state RNN tracks dynamics of various states of its associated object. In particular, it takes as input (1) the interaction encoding of the object with other objects and (2) the context encoding modeling interaction with the background from the perspective of object k . Then, the state of the OS-RNN \mathbf{h}_t^k is updated as follows:

$$\mathbf{h}_t^k = \text{RNN}([\mathbf{o}_{t-1}^k, \mathbf{z}_{t-1}^{\text{state},k}, \mathbf{e}_{t-1}^{\text{ctx},k}, \mathbf{e}_{t-1}^{\text{rel},k}], \mathbf{h}_{t-1}^k) . \quad (2)$$

Here, $\mathbf{e}_{t-1}^{\text{ctx},k}$ is the *object-context* interaction encoding and $\mathbf{e}_{t-1}^{\text{rel},k}$ is the *object-object* interaction encoding. In the following, we describe in more detail how these representations are obtained.

3.3.1. INTERACTION & OCCLUSION

Objects and agents in a physical environment interact with complex relations. Collision and occlusion are among two most popular types of such object interactions. Occlusion may not be seen a direct interaction between objects when the objects are physically disconnected. However, from the perspective of an observer, occlusion can still be modeled as an implicit interaction between the objects where one object makes another one invisible fully or partially. Although in a realistic environment these are the most fundamental types of object relationship that can happen simultaneously in a scene, in previous works these are studied separately and thus its jointly modeling and working have not studied.

In our model, we model both collision and occlusion in a simple graph neural networks (GNN). This object interaction modeling is also used in STOVE. However, in our model, we take the depth of the object into consideration and thus make the GNN considers not only object position but also object depth to model both the collision and occlusion. In our experiments, we found that this interaction modeling can learn different behaviours depending on the depth of an object and situation.

Specifically, we compute the interaction encoding $\mathbf{e}_t^{\text{rel},k}$ of

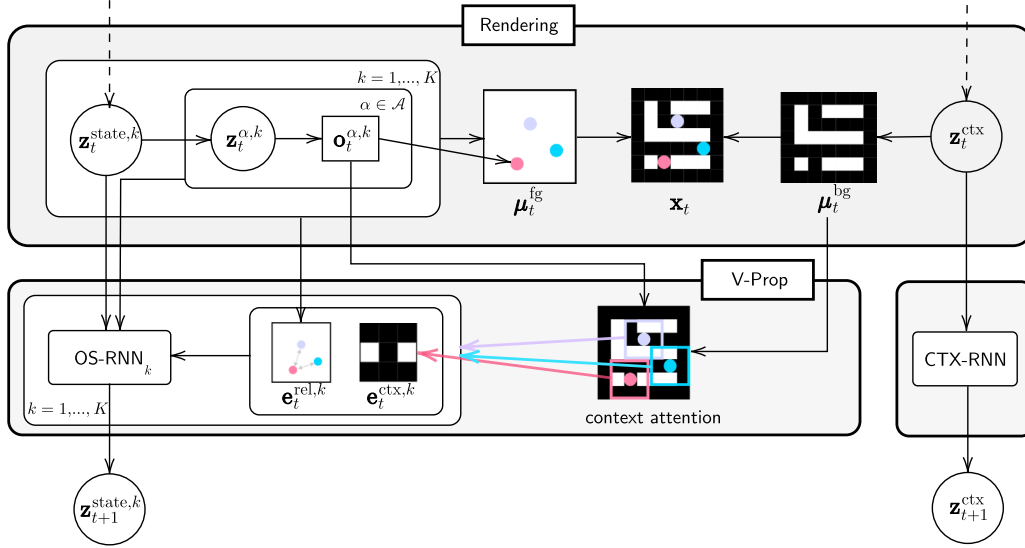


Figure 1. G-SWM Generative Process. For each timestep t , the generated image is a combination of a foreground image μ_t^{fg} and a background image μ_t^{bg} . To obtain the foreground image μ_t^{fg} , we determine the object attributes by decoding $z_t^{\text{state},k}$ and render them into a canvas. The background image is decoded from the context latent variable z_t^{ctx} . To propagate existing objects to the next timestep, the V-Prop module integrates information of object-object and object-context interaction in the object-state RNN, from which we compute the new object states $z_{t+1}^{\text{state},k}$. The context latent variable is also updated to z_{t+1}^{ctx} with a context RNN.

object k as follows:

$$\mathbf{e}_t^{\text{rel},k} = \mathbf{e}_t^{k,k} + \sum_{j \neq k} w_t^{k,j} \mathbf{e}_t^{k,j}. \quad (3)$$

To perform this, we first obtain the object representation \mathbf{u}_t^k by concatenating object attributes \mathbf{o}_t^k , the state latent $z_t^{\text{state},k}$, and the temporal object state \mathbf{h}_t^k . Then, we obtain the self-interaction encoding by $\mathbf{e}_t^{k,k} = \text{MLP}(\mathbf{u}_t^k)$, and the pairwise-interaction encoding $\mathbf{e}_t^{k,j}$ and the interaction weight (normalized with softmax) $w_t^{k,j}$ by $(\mathbf{e}_t^{k,j}, w_t^{k,j}) = \text{MLP}(\mathbf{u}_t^k, \mathbf{u}_t^j)$.

3.3.2. SITUATION AWARENESS

Objects not only behave interactively with other objects but also do so depending on their spatial situation. For example, an agent in PacMan should not only behave depending on the enemies but should also follow the corridor of a maze while not penetrating the wall.

We define such dependency of object states to the local spatial properties of the environment as *situation-awareness*. To implement this, we model the interaction between an object k and the environment by computing a context encoding $\mathbf{e}_t^{\text{ctx},k}$ using a simple attention mechanism, *attention on environment* (AOE).

The context encoding is obtained by first cropping the local area of the object using the Spatial Transformer (ST) (Jader-

berg et al., 2015) and then encoding the cropped patch:

$$\mathbf{g}_t^{\text{ctx},k} = \text{ST}(\mu_t^{\text{bg}}, \mathbf{o}_t^{\text{pos},k}, s^{\text{ctx}}), \quad (4)$$

$$\mathbf{e}_t^{\text{ctx},k} = \text{CNN}(\mathbf{g}_t^{\text{ctx},k}). \quad (5)$$

Here, μ_t^{bg} is the background image generated from $z_t^{\text{ctx},k}$ and $\mathbf{o}_t^{\text{pos},k}$ is the object position. A glimpse image $\mathbf{g}_t^{\text{ctx},k}$ extracted from μ_t^{bg} using attention by a Spatial Transformer is then encoded into the situation encoding $\mathbf{e}_t^{\text{ctx},k}$ using a CNN. The glimpse size is controlled by hyperparameter s^{ctx} . In other words, this is an encoding of *what happened nearby the object*. This situation encoding is then provided as input to the object-state RNN in Eqn. 2 to model temporal dynamics while being mixed with other encodings.

3.3.3. MULTIMODAL UNCERTAINTY

One of the main uses of a world model (Ha & Schmidhuber, 2018a) in model-based reinforcement learning or planning is to simulate diverse possible futures by learning a future distribution. Although previous models contain uncertainty modeling, their flexibility is limited to unimodal uncertainty because of the reliance on Gaussian distributions which provide only the unimodal uncertainty model. For instance, in predicting the future direction of an object, it would be hard to model the possibility of two very different directions, e.g. turning left or right, although it can model some randomness towards a direction, e.g. turning more or less to the left.

To this end, we adopt hierarchical latent modeling by introducing a high-level object-state latent variable $\tilde{z}_t^{\text{state},k}$ and

from the state latent we generate the attribute latents. This can be modeled by $p_\theta(\tilde{\mathbf{z}}_t^{\text{att},k}, \tilde{\mathbf{z}}_t^{\text{state},k} | \mathbf{z}_{<t}) =$

$$p_\theta(\tilde{\mathbf{z}}_t^{\text{att},k} | \tilde{\mathbf{z}}_t^{\text{state},k}) p_\theta(\tilde{\mathbf{z}}_t^{\text{state},k} | \mathbf{z}_{<t}). \quad (6)$$

That is, by learning the conditional behaviour distribution $p_\theta(\tilde{\mathbf{z}}_t^{\text{state},k} | \mathbf{z}_{<t})$, we can express the distribution of the expected behaviour of an object as a multivariate Gaussian distribution. Then, conditioning on a sampled behaviour, the object attribute $\tilde{\mathbf{z}}_t^{\text{att},k}$, such as position, can express highly complex multimodal uncertainty.

To model $p_\theta(\tilde{\mathbf{z}}_t^{\text{state},k} | \mathbf{z}_{<t})$, we use a multivariate Gaussian distribution whose parameters are estimated by an MLP taking \mathbf{h}_t^k of object-state RNN as input. Then, latent variables $\tilde{\mathbf{z}}_t^{\text{att},k}$ required to obtain object attributes $\tilde{\mathbf{o}}_t^k$ are sampled from Gaussian or Bernoulli distributions whose parameters are computed by MLPs taking $\tilde{\mathbf{z}}_t^{\text{state},k}$ as input. Refer to Appendix for more details of the implementation.

3.4. Inference and Learning

3.4.1. INFERENCE

For context inference, we use an approximation

$$q_\phi(\mathbf{z}_t^{\text{ctx}} | \mathbf{z}_{<t}^{\text{ctx}}, \mathbf{x}_t) = \mathcal{N}(\boldsymbol{\mu}^{\text{ctx}}, \boldsymbol{\sigma}^{\text{ctx}}) \quad (7)$$

which has a similar structure as the generation but with the encoding of \mathbf{x}_t as an additional input to the MLP generating $\boldsymbol{\mu}^{\text{ctx}}$ and $\boldsymbol{\sigma}^{\text{ctx}}$.

For propagation inference, given the factorization in Eqn. 6, it may be tempting to use the following full factorization of the approximate posterior $q_\phi(\tilde{\mathbf{z}}_t^k | \mathbf{x}_t, \mathbf{z}_{<t}) =$

$$q_\phi(\tilde{\mathbf{z}}_t^{\text{att},k} | \tilde{\mathbf{z}}_t^{\text{state},k}, \mathbf{x}_t, \mathbf{z}_{<t}) q_\phi(\tilde{\mathbf{z}}_t^{\text{state},k} | \mathbf{x}_t, \mathbf{z}_{<t}). \quad (8)$$

However, since \mathbf{x}_t provides all necessary information for accurately inferring the attributes $\tilde{\mathbf{z}}_t^{\text{att},k}$, this can easily make the model ignore $\tilde{\mathbf{z}}_t^{\text{state},k}$ during inference, consequently failing to learn $p_\theta(\tilde{\mathbf{z}}_t^{\text{att},k} | \tilde{\mathbf{z}}_t^{\text{state},k})$ as well. To address this issue, we use the following approximation replacing the factor for $\tilde{\mathbf{z}}_t^{\text{att},k}$ inference by the prior generation model. Thus, we have $q_\phi(\tilde{\mathbf{z}}_t^k | \mathbf{x}_t, \mathbf{z}_{<t}) =$

$$p_\theta(\tilde{\mathbf{z}}_t^{\text{att},k} | \tilde{\mathbf{z}}_t^{\text{state},k}) q_\phi(\tilde{\mathbf{z}}_t^{\text{state},k} | \mathbf{x}_t, \mathbf{z}_{<t}). \quad (9)$$

Implementation of the posterior $q_\phi(\tilde{\mathbf{z}}_t^k | \mathbf{x}_t, \mathbf{z}_{<t})$ follows a similar structure of the prior. In particular, a separate posterior object-state RNN is used to obtain the posterior hidden state \mathbf{h}_t^k as an encoding of the history $\mathbf{z}_{<t}$, and an encoding of a local proposal area from which $\tilde{\mathbf{z}}_t^{\text{state}}$ is inferred is extracted from \mathbf{x}_t .

3.4.2. LEARNING

Due to the posterior intractability, we use variational inference to train the model by maximizing the following

evidence lower bound:

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} \left[\frac{p_\theta(\mathbf{x}|\mathbf{z}) p_\theta(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{x})} \right]. \quad (10)$$

We adopt the standard ELBO in Eqn. (10) as the training objective. Besides, we found that including an auxiliary loss $\mathbb{KL}(q_\phi(\tilde{\mathbf{z}}^{\text{pres}}|\cdot) | \text{Bernoulli}(p))$ with a small p encourages the model to remove duplicate propagation during inference. This term also helps regularize the model and can be turned off after certain steps. We also apply the rejection mechanism proposed in SCALOR and the curriculum learning and discovery dropout used in SILOT. In our experiments, we found both these tricks help stabilize the training process. Related hyperparameters are provided in Appendix.

4. Related Work

Our work builds on top of a line of recent research in unsupervised object-centric representation learning. AIR (Eslami et al., 2016), SPAIR (Crawford & Pineau, 2019), SuPAIR (Stelzner et al., 2019), and SPACE (Lin et al., 2020) incorporate structured latent variables with a VAE (Kingma & Welling, 2014) in order to detect objects in a scene without any explicit supervision. SQAIR (Kosiorek et al., 2018) extends the AIR model by combining the object discovery module with a recurrent propagation module that can track changes in detected objects over time. SCALOR (Jiang et al., 2020) and SILOT (Crawford & Pineau, 2020) scale the SQAIR model to work on a large number of objects, using a parallel inference mechanism similar to SPACE. SCALOR and SPACE also introduce the dynamic background model and the background decomposition model, respectively.

STOVE (Kossen et al., 2019) introduces a state-space model for videos, combining an image model with a graph neural network (GNN) based dynamics model (Scarselli et al., 2009; Li et al., 2016; Kipf & Welling, 2017; Gilmer et al., 2017; Battaglia et al., 2018; Santoro et al., 2017; Sanchez-Gonzalez et al., 2018; Battaglia et al., 2016; Watters et al., 2017) that can model interactions between objects. STOVE leverages SuPAIR (Stelzner et al., 2019) as an object detection module and can achieve long-term video generation. R-NEM (van Steenkiste et al., 2018) uses a spatial mixture model learned via neural expectation maximization (Greff et al., 2017) to obtain disentangled representations for each entity in a scene and models the physical interactions between entities. Other works in this line of research include (Hsieh et al., 2018; He et al., 2019). Table 1 summarizes the abilities and limitations of a few of the most related models.

There have also been several recent works that utilize object-centric representations in the context of model-based reinforcement learning. OP3 (Veerapaneni et al., 2019) uses a spatial-mixture model similar to IODINE (Greff et al.,




Figure 2. Generated trajectories for the bouncing ball datasets. Conditioning on 10 ground truth frames, the first 20 generation steps are shown. To reduce visual clutter, results of the 2 LAYER-D setting are given in Appendix.

2019) to obtain a disentangled representation of entities in a scene. Similar to G-SWM, OP3 processes entities symmetrically and uses pairwise interaction modeling. However, it does not obtain explicit object locations and assumes the latent states follow the Markov property. COBRA (Waters et al., 2019) is another spatial-mixture model that uses MONet (Burgess et al., 2019) to obtain object representations, but it is not able to handle interactions between entities. Transporter (Kulkarni et al., 2019) uses KeyNet (Jakab et al., 2018) to discover keypoints and perform long term keypoint tracking. It makes the assumption that consecutive frames only differ in objects’ pose or appearance.

5. Experiments

We evaluate G-SWM on several datasets designed to illustrate generation quality with respect to the different abilities outlined in Table 1. We also test G-SWM in a 3D environment similar to the CLEVRER (Yi et al., 2020) dataset to show how our model performs in a more realistic setting.

5.1. Interaction, Occlusion, and Scalability

In these experiments, our dataset consists of a set of balls bouncing in a frame. Four different settings are tested. In both the OCCLUSION and INTERACTION settings, there are 3 balls each with a random color. In the OCCLUSION setting the balls pass through one another according to their respective depths. In the INTERACTION setting, the balls bounce off one another. The 2 LAYER setting combines these two settings with 3 red balls at one depth and 3 blue balls at a lower depth, so the red balls occlude the blue balls. When balls of the same color touch, they bounce off one another. The 2 LAYER-D (dense) setting is the same as the 2 LAYER setting, except there are 8 balls at each depth.

We compare our model with SCALOR and STOVE in these experiments. Each episode contains 100 timesteps and we train the models on random sequences of length 20. For




Figure 3. Generated trajectories for the the random single ball experiment. For each model, we show 100 samples.




Figure 4. Generated frames for the maze experiments. The Input column shows the 5 ground truth timesteps we provide to the model. The right-hand side of the figure shows 4 possible samples given the input. All 4 samples are combined in the Futures column, representing all the predicted states of the agents. Objects stay in corridor and randomly change their directions at intersections.

generation, each model is given the first 10 ground truth frames and asked to predict the remaining 90 frames. For measuring generation performance, we use the euclidean distance between the predicted and the ground truth center points of the balls at each timestep (Kossen et al., 2019).

Figure 2 shows qualitative results for the generated ball trajectories. Figure 5 plots the prediction error over 100 timesteps, and Table 2 shows the summed error over the first 10 prediction steps. STOVE performs reasonably well in both the INTERACTION and OCCLUSION settings where there are only 3 balls, but it cannot even detect the objects in the two 2 LAYER settings, so we leave these results empty. This is because it lacks a proper mechanism for handling the frequent occlusions in the 2 LAYER settings, and its LSTM-based detection is not scalable. SCALOR can handle occlusions with z^{depth} and is scalable to the 2 LAYER-D setting, but it fails to correctly predict the ball trajectories in the INTERACTION and the two 2 LAYER settings due to lack of interaction modeling. In many cases, the balls simply slow down and stop when they collide. G-SWM performs the best in all settings, showing a consistent ability to handle interaction and occlusion at the same time, while being scalable to videos with a large number of objects.




Figure 5. Plots of euclidean distance error of predicted ball positions per timestep. Each experiment is run with 5 different random seeds. Shaded areas indicate standard deviation.

Table 2. Euclidean distance error of predicted ball positions, summed over the first 10 timesteps of generation.

	G-SWM	SCALOR	STOVE
INTERACTION	0.242 ± 0.028	1.004 ± 0.019	0.411 ± 0.007
OCCLUSION	0.072 ± 0.015	0.272 ± 0.019	0.387 ± 0.021
2 LAYER	0.212 ± 0.007	1.012 ± 0.021	–
2 LAYER-D	0.535 ± 0.013	1.208 ± 0.011	–

5.2. Multimodal Uncertainty and Situation Awareness

As a simple demonstration of how different models work with multimodal uncertainty, we construct a dataset with a single ball moving down the center of the frame, and then randomly changes direction and moves towards either the bottom left corner or the bottom right corner. See Figure 3. At test time, each model is provided the first several frames before the ball changes direction. Even in this simple scenario, SCALOR, STOVE, and G-SWM-D (G-SWM without $\mathbf{z}^{\text{state}}$) fail to produce trajectories that match the dataset. In contrast, by introducing $\mathbf{z}^{\text{state}}$, G-SWM is able to model this multimodal behaviour and produce cleaner trajectories.

For a more complicated scenario, we construct a maze environment (see Figure 4) to demonstrate the ability of G-SWM to create realistic generations based on both interactions with the environment and multimodal uncertainty. Each episode in the dataset consists of a randomly created maze and several agents navigating through the maze. The agents only move within the corridors and continue in a straight path until they reach an intersection, at which point they choose a corridor to continue at random. For this experiment, we train on sequences of length 10 and at test time provide 5 ground truth steps to generate the following steps.

Figure 4 shows a few examples of the generated samples. More examples are given in Appendix. These examples show that G-SWM is able to generate predictions where the agents correctly stay within the corridors. Additionally, the random behaviour of the agents is also accurately modeled, as the predictions show the agents moving in different valid paths for different rollouts.

As an ablation study, we introduce several variants of G-SWM by respectively removing the multimodal uncertainty (MU) modeling (i.e., removing $\mathbf{z}^{\text{state}}$), the attention on environment (AOE) mechanism, and the modeling of situation awareness (SA). We count the number of agents that correctly stay within the corridors of the maze over time as an indicator of generation quality. Figure 6 shows the results over 100 timesteps and the convergence curve of this metric during training. Without $\mathbf{z}^{\text{state}}$ or situation awareness, the model cannot correctly predict the agent behaviour in the maze and thus clearly underperforming. Without AOE but with the attention-less encoding of the full context, the model eventually reaches a similar performance as the complete model, but it takes much longer to converge as shown in Figure 6. These results show that the proposed hierarchical modeling and the AOE mechanism can significantly improve generation performance when multimodal uncertainty and situation awareness are required.

5.3. 3D Interactions

Though the above experiments clearly exhibit the different capabilities of our model, we also want to demonstrate how our model performs in a more realistic setting. We use the CLEVR (Johnson et al., 2016) dataset as a starting point to create a dynamic 3D environment. In this dataset, several 3D objects are placed on a surface and a ball moves towards the objects, knocking them around. The shapes, positions, and colors of the objects are chosen randomly as well as the angle from which the ball enters the scene. For this dataset, we train on sequences of length 20. At test time, the model is given the first 10 ground truth frames, which are always

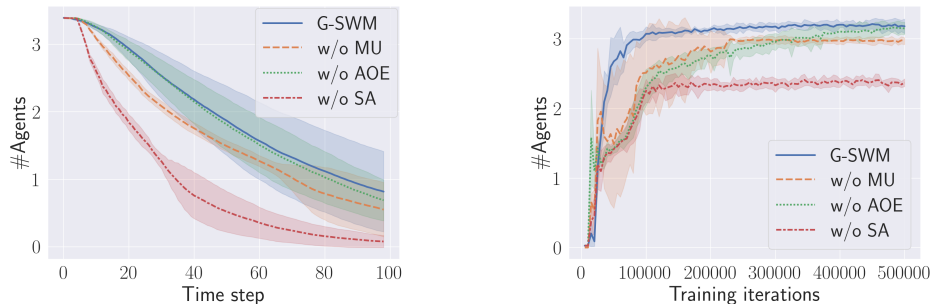


Figure 6. Quantitative results of the ablation study. The left plot shows the number of agents that stay within the maze corridors over 100 timesteps. The right plot shows the convergence curve of this metric during training, averaged over the first 10 generation steps. Each experiment is run with 5 random seeds.




Figure 7. Generation results for the 3D realistic physics experiments. Given the first 10 frames of the episode, 90 frames are predicted (15 out of the first 30 generated frames are shown).

before any interaction occurs.

Figure 7 shows the qualitative generation results on this dataset. More examples are given in Appendix. For comparison, we also test SCALOR and OP3 on this dataset. Since OP3 assumes the Markov property of the generation process, it cannot even model the ball movement. Besides, with implicit object representation, OP3 produces very blurry images. SCALOR produces clean reconstructions of objects, but due to lack of interaction modeling, it cannot model collisions and different objects overlap when they collide. G-SWM can produce realistic generations even in this complex setting. Occlusions are correctly captured when one object is in front of another and the generated collisions are also realistic. In particular, the model predicts a collision event only when two objects are close in the 3D space (e.g., at the same depth), which further demonstrates G-SWM’s ability to handle interactions based on the depth information. Although the generated object appearances are similar to the ground truth, we do notice that the model fails to capture the purple cube in this example, instead of generating a cylinder. This may be because, in the ground truth, the cube spins when it is hit, which is difficult to model correctly. Nevertheless, our model is able to predict other changes to appearance such as the change in the objects size (in the 2D

plane) as it moves further away from the camera.

6. Conclusion

In this paper, we proposed a new object-centric temporal generative model for world modeling. In experiments, we performed an extensive investigation on the generation performance of previously proposed models and ours, which has been missing in the literature. Through this comparison, we demonstrated that G-SWM implementing all important abilities jointly in a single model achieves superior or comparable performance in all tasks. We also demonstrated that G-SWM successively achieves two new important abilities, multimodal uncertainty and situated behavior. A future direction is to resolve some weakness of the spatial-attention-based object models and to support the temporal discovery model. This would help apply this model to improve model-based reinforcement learning and planning.

Acknowledgements

SA thanks Kakao Brain and Center for Super Intelligence (CSI) for their support.

References

- Addis, D. R. and Schacter, D. The hippocampus and imagining the future: where do we stand? *Frontiers in human neuroscience*, 5:173, 2012.
- Battaglia, P., Pascanu, R., Lai, M., Rezende, D. J., and Kavukcuoglu, K. Interaction networks for learning about objects, relations and physics. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pp. 4509–4517, Red Hook, NY, USA, 2016. Curran Associates Inc. ISBN 9781510838819.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gulcehre, C., Song, F., Ballard, A., Gilmer, J., Dahl, G., Vaswani, A., Allen, K., Nash, C., Langston, V., Dyer, C., Heess, N., Wierstra, D., Kohli, P., Botvinick, M., Vinyals, O., Li, Y., and Pascanu, R. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Burgess, C. P., Matthey, L., Watters, N., Kabra, R., Higgins, I., Botvinick, M., and Lerchner, A. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019.
- Crawford, E. and Pineau, J. Spatially invariant unsupervised object detection with convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3412–3420, 2019.
- Crawford, E. and Pineau, J. Exploiting spatial invariance for scalable unsupervised object tracking. In *AAAI*, pp. 3684–3692. AAAI Press, 2020.
- Eslami, S. A., Heess, N., Weber, T., Tassa, Y., Szepesvari, D., and Hinton, G. E. Attend, infer, repeat: Fast scene understanding with generative models. In *Advances in Neural Information Processing Systems*, pp. 3225–3233, 2016.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, pp. 1263–1272. JMLR.org, 2017.
- Greff, K., van Steenkiste, S., and Schmidhuber, J. Neural expectation maximization. In *Advances in Neural Information Processing Systems*, pp. 6691–6701, 2017.
- Greff, K., Kaufman, R. L., Kabra, R., Watters, N., Burgess, C., Zoran, D., Matthey, L., Botvinick, M., and Lerchner, A. Multi-object representation learning with iterative variational inference. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2424–2433. PMLR, 2019.
- Ha, D. and Schmidhuber, J. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems 31*, pp. 2451–2463. Curran Associates, Inc., 2018a.
- Ha, D. and Schmidhuber, J. World models. *arXiv preprint arXiv:1803.10122*, 2018b.
- He, Z., Li, J., Liu, D., He, H., and Barber, D. Tracking by animation: Unsupervised learning of multi-object attentive trackers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1318–1327, 2019.
- Hessel, M., Modayil, J., Van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M., and Silver, D. Rainbow: Combining improvements in deep reinforcement learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Hsieh, J.-T., Liu, B., Huang, D.-A., Fei-Fei, L. F., and Niebles, J. C. Learning to decompose and disentangle representations for video prediction. In *Advances in Neural Information Processing Systems*, pp. 517–526, 2018.
- Jaderberg, M., Simonyan, K., Zisserman, A., and Kavukcuoglu, K. Spatial transformer networks. In *Advances in neural information processing systems*, pp. 2017–2025, 2015.
- Jakab, T., Gupta, A., Bilen, H., and Vedaldi, A. Unsupervised learning of object landmarks through conditional image generation. In *Advances in Neural Information Processing Systems*, pp. 4016–4027, 2018.
- Jiang, J., Janghorbani, S., Melo, G. D., and Ahn, S. Scalor: Scalable object-oriented sequential generative models. In *International Conference on Learning Representations*, 2020.
- Johnson, J. E., Hariharan, B., van der Maaten, L., Fei-Fei, L., Zitnick, C. L., and Girshick, R. B. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1988–1997, 2016.
- Johnson-Laird, P. N. *Mental models: Towards a cognitive science of language, inference, and consciousness*. Number 6. Harvard University Press, 1983.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.

- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- Kosíorek, A., Kim, H., Teh, Y. W., and Posner, I. Sequential attend, infer, repeat: Generative modelling of moving objects. In *Advances in Neural Information Processing Systems*, pp. 8606–8616, 2018.
- Kossen, J., Stelzner, K., Hussing, M., Voelcker, C., and Kersting, K. Structured object-aware physics prediction for video modeling and planning. *arXiv preprint arXiv:1910.02425*, 2019.
- Kulkarni, T. D., Gupta, A., Ionescu, C., Borgeaud, S., Reynolds, M., Zisserman, A., and Mnih, V. Unsupervised learning of object keypoints for perception and control. In Wallach, H., Larochelle, H., Beygelzimer, A., dAlché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 10724–10734. Curran Associates, Inc., 2019.
- Li, Y., Tarlow, D., Brockschmidt, M., and Zemel, R. S. Gated graph sequence neural networks. In *International Conference on Learning Representations*, 2016.
- Lin, Z., Wu, Y.-F., Peri, S. V., Sun, W., Singh, G., Deng, F., Jiang, J., and Ahn, S. Space: Unsupervised object-oriented scene representation via spatial attention and decomposition. In *International Conference on Learning Representations*, 2020.
- Long, J., Shelhamer, E., and Darrell, T. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.
- Sanchez-Gonzalez, A., Heess, N. M. O., Springenberg, J. T., Merel, J., Riedmiller, M. A., Hadsell, R., and Battaglia, P. W. Graph networks as learnable physics engines for inference and control. In *ICML*, 2018.
- Santoro, A., Raposo, D., Barrett, D. G. T., Malinowski, M., Pascanu, R., Battaglia, P. W., and Lillicrap, T. A simple neural network module for relational reasoning. In *NIPS*, pp. 4967–4976, 2017.
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., and Szpunar, K. K. The future of memory: remembering, imagining, and the brain. *Neuron*, 76(4): 677–694, 2012.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- Spelke, E. S., Breinlinger, K., Jacobson, K., and Phillips, A. Gestalt relations and object perception: A developmental study. *Perception*, 22(12):1483–1501, 1993.
- Stelzner, K., Peharz, R., and Kersting, K. Faster attend-infer-repeat with tractable probabilistic models. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5966–5975. PMLR, 2019.
- van Steenkiste, S., Chang, M., Greff, K., and Schmidhuber, J. Relational neural expectation maximization: Unsupervised discovery of objects and their interactions. In *International Conference on Learning Representations*. OpenReview.net, 2018.
- Veerapaneni, R., Co-Reyes, J. D., Chang, M., Janner, M., Finn, C., Wu, J., Tenenbaum, J. B., and Levine, S. Entity abstraction in visual model-based reinforcement learning. In *CoRL*, volume 100 of *Proceedings of Machine Learning Research*, pp. 1439–1456. PMLR, 2019.
- Watters, N., Tacchetti, A., Weber, T., Pascanu, R., Battaglia, P., and Zoran, D. Visual interaction networks: Learning a physics simulator from video. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pp. 4542–4550, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Watters, N., Matthey, L., Bosnjak, M., Burgess, C. P., and Lerchner, A. COBRA: data-efficient model-based RL through unsupervised object discovery and curiosity-driven exploration. *CoRR*, abs/1905.09275, 2019. URL <http://arxiv.org/abs/1905.09275>.
- Yi, K., Gan, C., Li, Y., Kohli, P., Wu, J., Torralba, A., and Tenenbaum, J. B. Clevrer: Collision events for video representation and reasoning. In *International Conference on Learning Representations*, 2020.