Proceedings of Machine Learning Research vol 140:1–16, 2022 1st Conference on Causal Learning and Reasoning

VIM: Variational Independent Modules for Video Prediction

Author names withheld

Editor: Under Review for CLeaR 2022

Abstract

We introduce a variational inference model called VIM, for Variational Independent Modules, for sequential data that learns and infers latent representations as a set of objects and discovers modular causal mechanisms over these objects. These mechanisms - which we call modules - are independently parametrized, define the stochastic transitions of entities and are shared across entities. At each time step, our model infers from a low-level input sequence a high-level sequence of categorical latent variables to select which transition modules to apply to which high-level object. We evaluate this model in video prediction tasks where the goal is to predict multi-modal future events given previous observations. We demonstrate empirically that VIM can model 2D visual sequences in an interpretable way and is able to identify the underlying dynamically instantiated mechanisms of the generation process. We additionally show that the learnt modules can be composed at test time to generalize to out-of-distribution observations.

Keywords: Objects, Modularity, Unsupervised Representation Learning, Videos, Interpretability, Compositionality

1. Introduction

Predicting future events is believed to be a fundamental function of the human brain (Clark, 2013; Mullally and Maguire, 2014) having implications for representation learning, planning or counterfactual reasoning. Humans have the ability to decompose a visual scene into abstract objects and predict their changes far into the future (Kahneman et al., 1992; Spelke et al., 1993). Most interestingly, humans can adapt to a new situation quickly by re-using relevant prior knowledge about objects and their dynamics. In particular, they are good at generalizing in a compositional way because they represent knowledge as re-usable components that can be further composed to explain new observations.

Recent work (Eslami et al., 2016; Burgess et al., 2019; Greff et al., 2019; Crawford and Pineau, 2019; Locatello et al., 2020) have made significant progress in learning unsupervised entity-centric representations of images and a few of them (Lin et al., 2020a; Kossen et al., 2019; Jiang et al., 2019; Crawford and Pineau, 2020; Kosiorek et al., 2018) have extended those models to generate both deterministic and stochastic videos. However, except for SCOFF (Goyal et al., 2020), to the best of our knowledge none of them have considered architectures that disentangle the underlying dynamic rules of the generation process and rather model the transition part of their models with a shared monolithic module that is applied at each time step. We extend the architectural inductive biases from SCOFF and incorporate them in a probabilistic generative model of videos.

In this work we propose a variational framework to learn both entity-centric representations and entity-centric transition modules. We show that in a simple 2D stochastic environment we are able to identify the exact generating factors of variation both in terms of declarative (e.g. constitutive objects of the visual scene) and procedural knowledge (e.g. transitions rules that govern the stochastic dynamics of the objects). We argue that having those 2 layers of structure in a generative model

© 2022 .

is an essential first step towards generalizing out-of distribution, particularly when the new unseen data results from a composition of seen dynamic rules.

Our model called Variational Independent Modules (VIM), is a probabilistic generative model where both latent states and transitions functions over these latents have an entity-centric inductive bias. It thus learns a latent state composed of a set of abstract entities, or slots (Locatello et al., 2020) and a set of stochastic transition functions over the entities. These transitions functions, which we call modules, are independently parametrized and are shared across entities, following the principle of reusable independent causal mechanisms (Peters et al., 2017; Goyal et al., 2019, 2020). At each time step t, our model infers a set of categorical latent variables (called *selection variables* $\mathbf{r}^t$) to select which transition modules to apply to each represented entity in the latent set $\mathbf{z}^t$. We train VIM using variational inference over both the entity-centric set of K latents $\{\mathbf{z}_k\}_{k=1,,K}$ and their respective selection variables $\{\mathbf{r}_k\}_{k=1,,K}$. In particular, the distribution of the categorical selection variables is implemented with a key-query attention mechanism (Bahdanau et al., 2014; Vaswani et al., 2017) in which all the possible modules compete (Goyal et al., 2019; Goyal et al., 2021; Locatello et al., 2020) to explain future states and they are sampled according to their attention importance weights.

One of the key assumptions behind our framework is that the abstract entities are evolving mostly independently and only interact sparsely with each other (Pearl, 2009; Goyal et al., 2019). Consequently, the causal graph is sparse (Bengio, 2017; Goyal and Bengio, 2020) and a module can only consider a handful of slots as input argument. However in this work we will mostly consider unary modules and leave the exploration of n-ary modules to model interactions between multiple entities for future work.

We evaluate VIM in a 2D stochastic setting where the goal is to predict multi-modal future events given previous observations and we make the following contributions:

- We propose a simple stochastic environment in which we know the rules of the dynamics, corresponding to the different modes in the transition distribution, to enable checking for correctly identified solutions (Locatello et al., 2018).
- We show via simulations that in this environment, our framework is able to identify the ground truth dynamics rules that govern the data generation process.
- We show that VIM is able to generalize in a compositional and interpretable way in a simple out-of-distribution (OOD) tracking task.

2. Background : Recurrent State Space Models

Given a set of D observed frames $\mathbf{c} = (c_1, ..., c_D)$ and the T following future frames $\mathbf{x} = (x_1, ..., x_T)$, our goal is to learn a generative model that maximizes the probability $p(\mathbf{x}|\mathbf{c})$. To solve this task, our proposed model builds upon the Recurrent State-Space Model (RSSM) (Hafner et al., 2019) that we first review in this section.

RSSMs define a variational framework to model sequential data. They introduce a sequence of latent variables $\mathbf{z} = (z_1, z_2, ..., z_T)$ to capture the stochasticity of the observation at each time step such that the joint distribution is factorized as:

$$p(\mathbf{x}, \mathbf{z}|\mathbf{c}) = \prod_{t=1}^{T} p(x_t|\mathbf{z}_{\leq \mathbf{t}})p(z_t|\mathbf{z}_{<\mathbf{t}})p(z_0|\mathbf{c}). \tag{1}$$

where $p(x_t|\mathbf{z}_{\leq \mathbf{t}})$ is the likelihood model, $p(z_t|\mathbf{z}_{<\mathbf{t}})$ is the prior transition model, c is some given context and $p(z_0|\mathbf{c})$ the discovery model that maps the observed context to a distribution over the initial latent state. The main advantage of RSSMs over previous autoregressive models is computational: it can make multi-step future predictions without having to render/encode observed frames at each time step.

RSSMs are trained using variational inference (Jordan et al., 1999), using an amortized approximate posterior $q(\mathbf{z}|\mathbf{x},\mathbf{c}) = \prod_{t=1}^{T} q(z_t|\mathbf{z}_{<\mathbf{t}},\mathbf{x}_\mathbf{t})p(z_0|\mathbf{c})$ where $q(z_t|\mathbf{z}_{<\mathbf{t}},\mathbf{x}_\mathbf{t})$ is called the posterior transition model and $q(\mathbf{z}|\mathbf{x},\mathbf{c})$ approximates the true posterior distribution $p(\mathbf{z}|\mathbf{x},\mathbf{c})$. Training is done end-to-end by maximizing the Evidence Lower Bound (ELBO) (Kingma and Welling, 2013; Rezende et al., 2014).

3. VIM: Variational Independent Modules

Similar to (Lin et al., 2020a; Kossen et al., 2019; Jiang et al., 2019; Crawford and Pineau, 2020; Kosiorek et al., 2018), VIM is a probabilistic RSSM where the latent space is structured as a set of K latent vectors $\mathbf{z} = \{\mathbf{z}_k\}_{k=1..K}$ (e.g. slots (Locatello et al., 2020)) where each slot is supposed to represent an *abstract* entity of the input (e.g. visual objects). In VIM we add an additional layer of structure to account for the fact that only a few rules (Goyal et al., 2020) govern the dynamics of objects in the world and that these rules (that we call *modules*) are shared amongst entities. Our transition distribution $p_\theta(\mathbf{z}_t|\mathbf{z}_{<t})$ is thus parametrized with M independent modules that compete against each other to explain future observations. The modules operate over entities and can have a predefined number of arguments. For our experiments, unless specified otherwise, we use unary modules that operate independently on a single entity. The high-level computation steps of our model's transition function are the following:

- For each entity k we first compute all the M possible next distribution given its current latent state $\mathbf{z}_t$ defined by the M independant transition modules.
- We then score each candidate with a learned key-query attention mechanism to define a categorical distribution over these possible futures represented by an entity-centric selection variable $\mathbf{r}_t^k$.
- Finally we sample the module index according to $\mathbf{r}_t^k$ and select the corresponding future candidate as the next step entity state $\mathbf{z}_{t+1}^k$.

VIM can thus be interpreted as a RSSM with 2 latent variables, $\mathbf{r} = \{\mathbf{r}_t\}_{t=1..T}$ and $\mathbf{z} = \{\mathbf{z}_t\}_{t=1..T}$ where $\mathbf{r}$ variables specify which modules to apply at which time step for all the slots in $\mathbf{z}$. Its generative model is factorized as

$$p_\theta(\mathbf{x},\mathbf{r},\mathbf{z}|\mathbf{c}) = \underbrace{p_\theta(\mathbf{z}_0^k|\mathbf{c})}_{\text{discovery model}} \prod_{t=1}^{T} \underbrace{p_\theta(\mathbf{x}_t|\mathbf{z}_t)}_{\text{observation model}} \prod_{k=1}^{K} \underbrace{\underbrace{p_\theta(\mathbf{z}_t^k|f_\theta(\mathbf{z}_{t-1}^k,\mathbf{r}_t^k))}_{\text{latent update}} \underbrace{p_\theta(\mathbf{r}_t^k|\mathbf{z}_{t-1})}_{\text{module selection}}}_{\text{transition model}} \tag{2}$$

where the observation model $p_\theta(\mathbf{x}_t|\mathbf{z}_t)$ and the discovery model $p_\theta(\mathbf{z}_0|\mathbf{c})$ are both parametrized with a slot-attention network (Locatello et al., 2020). In the following we describe in more details the parametrization of both the inference and generation model of the transition part of VIM.

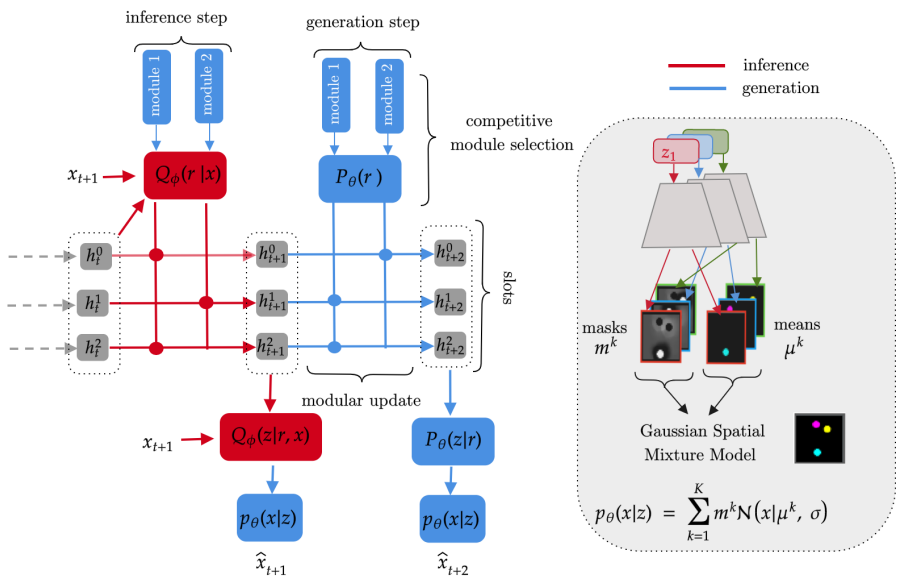

Figure 1: **Overview of VIM** VIM disentangles both Objects and Dynamics with slots and transition modules. *Left*: structure of the transitions, controlled by the selection of an inferred rule r and its application to the latent state z, and by an observation model of x. *Right*: generative decoder architecture mapping a set of object descriptions z to an image where the different objects are rendered and composed using a Spatial Guassian Mixture Model.

3.1. Generation

For both generation and inference, at time step t we start by computing all the M possible future candidates for each entity k. We denote this set of candidates $\mathbf{h}_t^k = \{f_{\theta_m}(\mathbf{z}_t^k) = f_\theta(\mathbf{z}_t^k, \mathbf{r}_t^k = m)\}_{m=1..M}$ where the M modules have independent parameters $\{\theta_{m=1..M}\}$.

Module Selection Prior In this step we describe the parametrization of the module selection prior $p(\mathbf{r}_{t+1}^k|\mathbf{z}_t)$ where $\mathbf{r}_{t+1}^k$ defines a categorical variable that indexes which module to apply to entity k at time step t. The probability of the categorical distribution from which $\mathbf{r}_{t+1}^k$ is obtained are computed with a key-query attention mechanism where the keys are extracted from the M candidates $\mathbf{h}_t^k$ and the queries from a function of the current state $g_\theta(\mathbf{z}_t^k, \mathbf{z}_t)$. We then sample the module m to apply with a Gumbel softmax (Jang et al., 2016). We denote $f_\theta(\mathbf{z}_t^k, \mathbf{r}_{t+1}^k)$ the application of the module indexed by $\mathbf{r}_{t+1}^k$ but there really are M different modules, each with their separate parameters.

Latent Update Prior Now we describe the parametrization of the entity-centric latent update $p_\theta(\mathbf{z}_{t+1}^k|f_\theta(\mathbf{z}_t^k, \mathbf{r}_{t+1}^k))$. Since the module selection is discrete, it is necessary to give some flexibility to the updates to account for non-discrete variations of the data (e.g., rendering uncertainty) or to simply ease the training by giving some slack to the optimization process. To do so we first sample from a Gaussian prior such that: $\mathbf{\Delta}_{t+1}^k \sim \mathcal{N}(\mu_\theta(f(\mathbf{z}_t^k, \mathbf{r}_{t+1}^k)), \sigma_\theta(f(\mathbf{z}_t^k, \mathbf{r}_{t+1}^k)))$ where μ_θ, and σ_θ are small MLPs . We then compute the latents with a simple additive update: $\mathbf{z}_{t+1}^k = \text{LayerNorm}(\mathbf{z}_t^k + f(\mathbf{z}_t^k, \mathbf{r}_{t+1}^k) + \mathbf{\Delta}_{t+1}^k)$

3.2. Inference

In this section we describe the parametrization of the inference machinery, which is factorized as:

$$q(\mathbf{z}, \mathbf{r}|\mathbf{x}) = \prod_{t=1}^{T} \prod_{k=1}^{K} q_\phi(\mathbf{z}_t^k | f_\theta(\mathbf{z}_{t-1}^k, \mathbf{r}_t^k), \mathbf{x}_t) q_\phi(\mathbf{r}_t^k | \mathbf{z}_{t-1}, \mathbf{x}_t). \tag{3}$$

During inference we need to encode information about the target frame to infer the updates to apply to each entity slot k. However not all the visual information is needed for all the abstracted entities. Each entity needs to attend to a specific perceptual grouping of the target frame. To that end, we propose to use a slot-attention mechanism where the grouping is learned through a key-query attention mechanism. We denote $\mathbf{x}_t^k$ the perceptual grouping for entity k obtained at time step t such that :

$$\hat{\mathbf{x}}_k^t = \sum_i \beta_{k,i}^t V_i^t \text{ with } \beta^t = \text{softmax}(\frac{KQ^T}{\sqrt{D}}, \text{dim="slots"}) \tag{4}$$

where the keys K are extracted from the current states $\{\mathbf{z}_t^k\}_{k=1..K}$ and the queries and values are extracted from the target input $\mathbf{x}_t$ that is first encoded with a size-preserving CNN backbone and then processed as a set of vectors to which we append learned positional encodings.

Module Selection Posterior We describe the parametrization of the rule selection posterior $q_\phi(\mathbf{r}_t^k|\mathbf{z}_{t-1}, \mathbf{x})$ for each entity k. Like with the module selection prior in generation, the idea is to score the same future update candidates $\{f_{\theta_m}(\mathbf{z}_t^k)\}_{m=1..M}$ with a key-query attention mechanism, but this time the query for entity k comes from the entity-centric target encoding $\mathbf{x}_t^k$. The intuition behind this design choice is to select the update that best explains the current observation. We sample the module m to apply with a Gumbel softmax trick and denote $f_\theta(\mathbf{z}_{t-1}^k, \mathbf{r}_t^k)$ the application of the module indexed by $\mathbf{r}_t^k$.

Latent Update Posterior The latent update posterior is similar to prior, where this time $\mathbf{\Delta}_t^k \sim \mathcal{N}(\mu_\phi([f(\mathbf{z}_{t-1}^k, \mathbf{r}_t^k); \mathbf{x}_t^k]), \sigma_\phi([f(\mathbf{z}_{t-1}^k, \mathbf{r}_t^k); \mathbf{x}_t^k]))$ where μ_ϕ, and σ_ϕ are small MLPs. Each latent follows the same additive update: $\mathbf{z}_{t+1}^k = \text{LayerNorm}(\mathbf{z}_t^k + f(\mathbf{z}_t^k, \mathbf{r}_{t+1}^k) + \mathbf{\Delta}_{t+1}^k)$.

3.3. Training

Training is done using variational inference maximizing the following evidence lower bound:

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x},\mathbf{r}) q_\phi(\mathbf{r}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \underbrace{KL(q_\phi(\mathbf{z}|\mathbf{r}, \mathbf{x}) || p_\theta(\mathbf{z}|\mathbf{r}))}_{\text{gaussian KL}} - \underbrace{KL(q_\phi(\mathbf{r}|\mathbf{x}) || p_\theta(\mathbf{r}))}_{\text{module KL}} \tag{5}$$

where there is a tension between explaining the multi-modality of the future outcomes either using the Gaussian distribution of a single module (but this would mean trying to fit a unimodal Gaussian distribution to multi-modal observations and would result in a high noise KL) or using several modules indexed by $\mathbf{r}$ to explain the different modes. The module KL also prevents the model to learn duplicate transition modules and we verify this experimentally by training the model with more modules than actual modes in the distribution. In that case, additional modules are simply rarely selected during inference.

4. Related Work

Video Prediction Since the initial models inspired by language modelling (Ranzato et al., 2014; Srivastava et al., 2015), video prediction has seen great progress leveraging advances in generative models and deep learing architectures. The temporal and spatial dependencies between pixels is typically captured via autoregressive models (Larochelle and Murray, 2011; Dinh et al., 2016; Kalchbrenner et al., 2017; Reed et al., 2017; Weissenborn et al., 2020) or latent variables models such as the VAE (Kingma and Welling, 2013; Chung et al., 2015; Denton and Fergus, 2018; Lee et al., 2018; Castrejon et al., 2019; Villegas et al., 2017a,b) or GAN (Goodfellow et al., 2014; Vondrick et al., 2016; Mathieu et al., 2015). Our approach uses variational inference for training, like other VAE-like models. Most of those previous models, however, rely on fixed-size unstructured vectorial representations of the state and monolithic prediction models while we explore the use of structured latent space (as a set of slots) and modular architectures (for the mechanisms) to explain multimodality in the transitions.

Unsupervised Object Discovery A recent research direction explores unsupervised object-centric representation learning from visual inputs. The main motivation behind this line of work is to disentangle a latent representation in terms of objects composing the visual scene. They can be divided into two types of models: on one hand, the spatial mixture models (Locatello et al., 2020; Burgess et al., 2019; Greff et al., 2019) learn a set of unstructured latents that are decoded into a pixel-wise mean and mask such that each pixel location defines a Gaussian mixture model weighted by the slot masks. On the other hand, spatial-transformer based models propose to further disentangle each slot latent representation into several variables (e.g. content, location, presence, depth) where the location variable parametrizes the input to a spatial transformer that fills a canvas, as in (Eslami et al., 2016; Crawford and Pineau, 2019; Stelzner et al., 2019; Lin et al., 2020a). The contribution of our work is orthogonal and we want to show the benefits of disentangling the mechanisms that process the slots in the transition module of a sequential generative model when the target distribution has multimodal uncertainty. Lin et al. (2020a) argues that a careful design of the prior is needed to account for multimodal future trajectories and propose to do so with a slot-wise hierarchical Gaussian prior model. In this work we propose an *interpretable* alternative to account for several modes in the target distribution and show that a modular approach is capable of capturing the underlying factors (e.g. modes in our case) of the generating process. In terms of architectural components, we extend Locatello et al. (2020)'s slot attention module to a sequential setting where the slot attention module is used to encode a meaningful part of the target frame in a slot-wise manner. We show that by using independently parametrized modules in the recurrent transition function, our approach is able to discover object-centric dynamical rules in an unsupervised manner and that the slot attention inference machinery is able to select compositions of those dynamical rules at test time.

Independent Mechanisms Recent approaches have explored architectures composed of a set of independently parametrized modules which compete with each other to communicate and attend or process an input (Goyal et al., 2019; Goyal et al., 2021). Those architectures are inspired by the notion of independent mechanisms (Pearl, 2009; Bengio et al., 2019; Goyal et al., 2019), which suggests that a set of independently parametrized modules capturing causal mechanisms should remain robust in case of distribution shift due to an intervention, as adapting one should not require adapting the other modules. The recurrent independent mechanisms architectures (Goyal et al., 2019; Goyal et al., 2021) however are not probabilistic models; they cannot capture the uncertainty inher-

ent with future predictions. In this work we formulate the same intuition of mechanisms separation in a variational inference framework where the selection of the the mechanisms is expressed as a categorical random variable whose posterior distribution the model must infer. We further showcase an interpretability advantage in the case of a simple 2D stochastic environment.

5. Experiments

In this section we describe experiments designed to showcase three main properties of VIM: i) object-centric *multi-modality* of the predictions in stochastic environments, ii) *interpretability* of the learned modules, and iii) *generalization* to compositional out-of-distribution settings in an interpretable way.

5.1. Interpretability

In this section we show that the modular architecture of VIM is able to capture the underlying generating factors of the target distribution in an interpretable way. Through an ablation study on the number of available modules we also show that with enough modules to capture the different distribution modes precisely, VIM performs better than models that rely on a unimodal prior.

Dataset We evaluate VIM on a dataset built to showcase interpretability of the modules. This dataset, called the **Random Walk Dataset** consists of 2D image sequences of multiple colored-balls that evolve on a black background. At each time step each ball can move randomly in one of the 4 cardinal directions. The resolution of the videos is 32×32 and the model is trained on up to 15 consecutive frames.

Setup We train VIM with N modules and $K = 4$ slots on trajectories that contain 3 balls. During training, the selection variable $\mathbf{r}$ is sampled using a Gumbel-softmax (Jang et al., 2016) with a fixed temperature of 1 and we use the hard version during testing so that the variables are categorical. For evaluation we match each ground truth ball of the visual scene to a latent slot explaining it. Note that Spatial-Transformer based models (Jang et al., 2016; Crawford and Pineau, 2019) directly exploit the slot bounding box center coordinate to compute this matching. Our model does not consider object bounding boxes. Instead, we use the slot-wise masks outputted by the observation model to match them to objects. More specifically, to determine the location (x_k, y_k) of the object explained by the k-th slot we sample coordinates uniformly over the 2D space covering the whole frame and then we average these coordinates weighted by the value of the slot mask value at each location $m_k^{x,y}$, such that:

$$x_k, y_k = \frac{\sum_{(x,y)\in[1,H]\times[1,W]} (x, y) * m_k^{x,y}}{m_k^{x,y}}.$$

Once we have those slot-wise coordinates we compare them against ground truth balls positions and match a slot with the closest ball position in the first frame and keep the same matching for the rest of the sequence. We then use these matched pairings to compute ball position errors for tracking purposes or to compare the slot-wise module selection variable to the ground-truth action transition.

Results Our objective is to verify that the learnt modules can directly be interpreted as causal dynamics rules underlying the generation process for the Random Walk dataset. We show that each module correspond to a particular mode of the generating distribution (e.g. moving in one of the cardinal directions). To do so, we evaluate our inference model on some test sequences of

length 20 and store the selection variables of all the slots $((\mathbf{r}_k^t)_{k=1..K})_{t=1..20}$. Using the first set of slots of the sequence we match each ground truth ball with its explaining slot (whose location has been obtained with its decoded mask) by choosing the closest one. We denote the slot associated with the d-th ball by i_d. We compare at each time step t the ground truth transitions of each ball d with the index of the module that slot i_d has selected and report the proportion of the correct correspondences. The percentage in Figure 2 shows that each cardinal transition can be explained by one of the modules. We further visualize the repeated effect of each module on a single ball to confirm that correspondence. Figure 2 shows that four of the modules move a ball in one of the cardinal direction exclusively. When the model is trained with more modules than actually needed (e.g. more modules than modes in the distribution), then additional modules are not used. This behaviour is in part encouraged by the module KL term of the loss which will be larger if 2 modules are identical: given the way we compute the module selection weights with a key-query attention machinery, modules that explain the same mechanisms would have the same selection probability. We show in Figure 2 the correspondence between selected modules and ground truth actions for VIM trained with 5 modules. In this case, only 4 modules appear to be selected by the slots that contain balls.

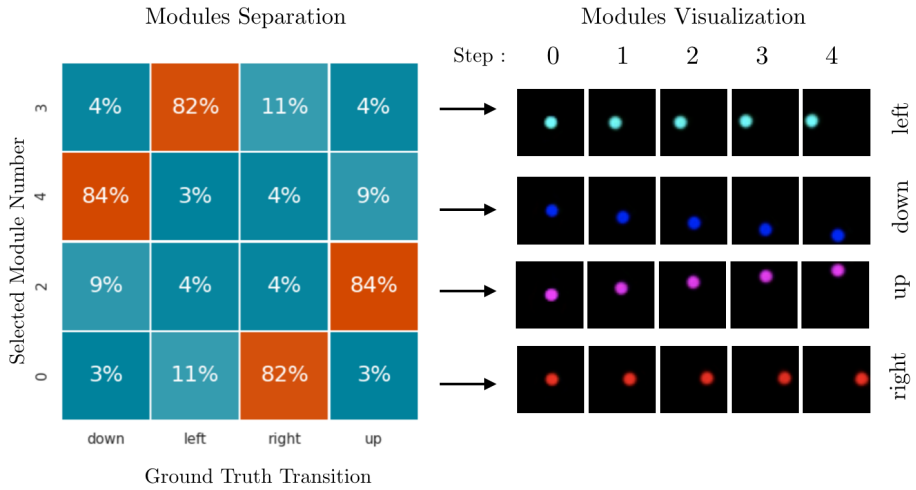

Figure 2: **Dynamics Rules Discovery** The proposed VIM model learns to match modules to ground truth entity transitions. *Left:* For each rule, we show the distribution over selected modules for a 5-module model. Modules specialize to implement one rule, with additional modules not being used. *Right:* We show the effects of repeatedly applying a module to a slot. Module transitions are interpretable and correspond to ground-truth transitions.

Effect of the number of Modules In this section we study the effect of the number of available modules to model the sequences in terms of generation abilities. The model is trained on the Random Walk dataset with a varying number of modules. Larger per-module capacity is given to models with less modules to factor out the effect of total number of parameters in the modeling ability. In Figure 3 we identify three regimes: monolithic transition module (e.g. 1 Module does the work), transition function with less modules than actual modes in the distribution (e.g. 2 modules share the work), and a transition function with more modules than actual modes in the distribution (5 and

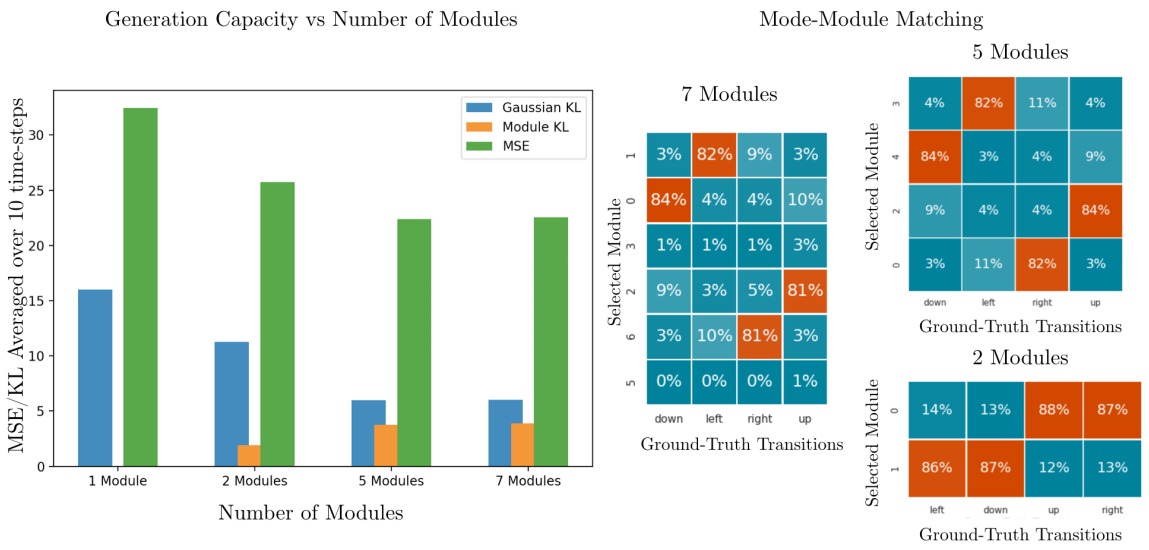

Figure 3: **Modules number ablation** *Left:* MSE/KL of the inference model averaged over 10 time steps for VIM with a varying number of modules *Right:* Modules specialize to each explain a mode in the distribution when enough modules are available.

7 modules). The monolithic transition model can only rely on the unimodal Gaussian distribution to capture the multimodal stochasticity of the transitions, which results in lesser ability to capture multimodality. In the low modules regime, the model separates the modules to explain the modes in a hierarchical way, but some modes are shared between modules which results in lower modeling capacity than the model with enough modules to explain all the modes. In the case of 5 modules, module number 1 is never selected and thus doesn't appear in the color map. When VIM has more modules than actual modes in the distribution, additional (and non useful) modules are almost never selected. We show the resulting modules specialization and mode-module matching in Figure 3.

5.2. Compositionality

Dataset In this section we introduce a variant of the Random Walk dataset where the coloured balls dynamics at each step is a composition (e.g. left-up-up-left) of the previous atomic cardinal transitions. Between two consecutive frames, we sample the sequence of modules that composes the transition of each ball at random. When trained on this dataset, the model never sees the the atomic transitions that compose the random walk of each ball. We call it the Compositional Random Walk.

Setup We introduce a variant of VIM where at each time step we allow the model to perform multiple selections and update steps before rendering. The rationale behind this variant is that for sequences whose transition dynamics rules are compositional we want to check whether VIM is able to compose some atomic building blocks (e.g. the modules) to explain an observation both at training and testing time. When the model is allowed k iterations before rendering we simply denote the variant VIM-k.

Results The insights we derived from the results of these experiments are two-fold:

- When trained on the atomic setting, VIM generalizes in an interpretable manner to compositions of transitions seen during training.
- When trained on the compositional setting, the individual modules that are learned in the multi-step version of VIM are still interpretable in terms of atomic transitions that were never seen during training.

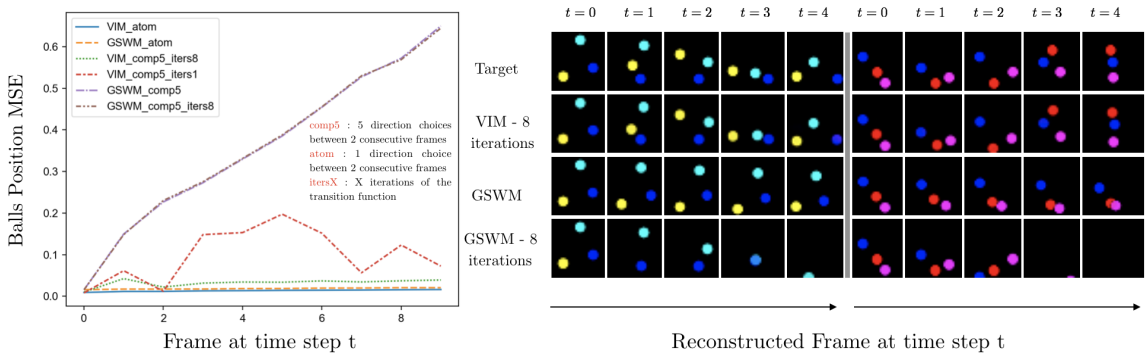

Figure 4: **OOD Tracking**. VIM's ability to track to out-of-distribution composition of transitions seen during training. Compositional dynamics tracking comparison of VIM and G-SWM and their respective multi-step versions. Both trained on the atomic setting and are compared qualitatively *Right*: and quantitatively *Left* reporting their respective ball position errors.

OOD Compositional Generalization We train VIM on the atomic setting and test it on out-of-distribution compositional sequences where up to 5 choices of atomic transitions are allowed for each ball between two consecutive frames.(e.g. *left-down-right-down-right*). At test time we consider a multi-step version of VIM where 8 iterations of the selection-update step are allowed in between frames. We show that the inference key-query selection part of VIM acts at each iteration as a kind of planner in the space of modules and selects the module that best explains the current observation. In Figure 5 we render the resulting frame at each iteration and notice that VIM composes modules in an iterative manner such that each ball progressively gets closer its target position before oscillating around it. Additionally, in Figure 4 we compare the tracking ability of a state-of-the art object-centric video generation model, G-SWM (Lin et al., 2020b), with this multi-step version of VIM . We show that both the key-query selection mechanism and the composable modules enables VIM to better track objects that have a different amplitude than the one seen during training. For a fair comparison we also consider a multi-step version of G-SWM where several iterations of the same monolithic propagation module is allowed between 2 consecutive frames.

Multi-step Interpretability We train a 2-step version of VIM on the compositional setting and restrict the number of modules to 5 so that we still have more modules than atomic transitions (that are never seen) but less modules than possible compositional transitions seen during training (e.g. *left-right, up-left, down-down*, and so on). We test this model on sequences of atomic transitions to evaluate the proportion of selected modules corresponding to each ground truth transition. We show in Figure 6 that the same modularisation emerges as when VIM is trained on the atomic setting. In this compositional setting, when evaluated on atomic transitions we observe the same module specialization as when it was trained on atomic transitions. This shows that when trained

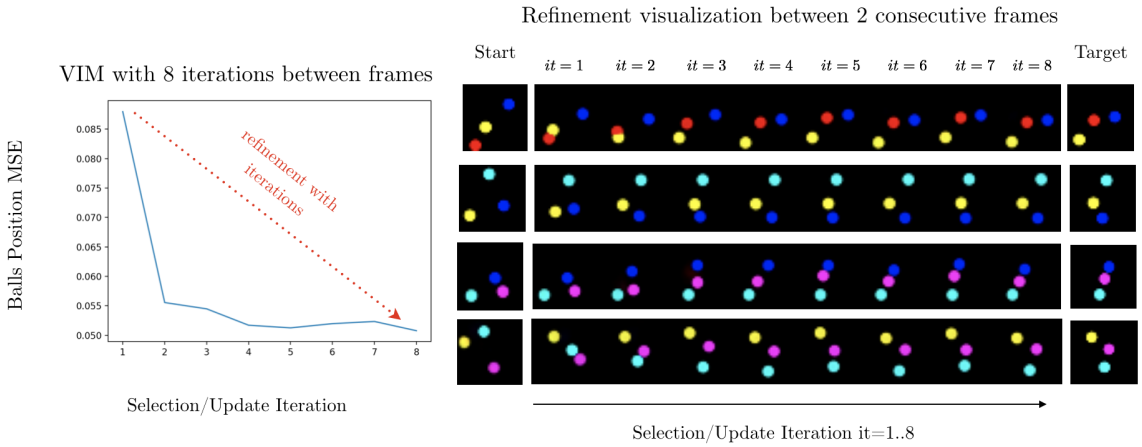

Figure 5: **Iterative Refinement**. We show that VIM can model out-of-distribution compositions of transitions. *Left:* VIM trained on the Random Walk Dataset and tested on the Compositional Dataset allowing 8 iterations between frames. Between 2 consecutive frames, the ball position error decreases with the number of iterations. *Right*: Visualization of the module iterations. VIM iteratively selects modules to modify the starting observation and reconstruct the target frame.

Figure 6: **Rule-Module matching for compositional transitions** We show a rule-module matching histogram for a VIM model with 5 modules and 2 module iterations per frame trained on the Compositional Random Walk dataset. We observe that, despite only observing compositions of actions, the modules implement and match ground-truth atomic transitions.

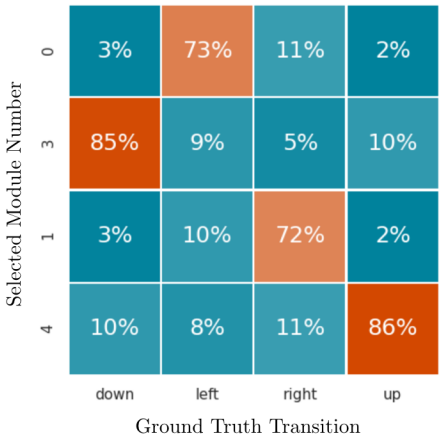

with 2 selection/update iterations between 2 consecutive frames VIM successfully identifies the underlying atomic transformations that compose the transitions VIM has been trained on without directly observing them.

5.3. Dynamics Transfer

In this section we are interested in showing that VIM learns a factorized representation formed by entities and transition rules over these entities. In particular we would like to evaluate whether VIM is able to recognize known dynamics applied to unseen objects and that the inference machinery proposes to transfer those dynamics to known objects using the selection variables for these dynamics.

Dataset In this section we test VIM on a dataset where objects have out-of-distribution attributes (e.g. shape, color, size), beyond those seen during training, but following the same Random Walk dynamics seen during training: they can go in one of the four cardinal directions between two consecutive frames and with the same amplitude.

Results We train VIM on the Random Walk Dataset with coloured balls of the same size and test it on a OOD version where objects have different shape, colour and size than the balls seen during training. During testing, we run the inference model on these OOD target sequences $(\mathbf{x}_t^{target})_{t=0..T}$ and extract the selection variables at each time step $(\mathbf{r}_t^{target})_{t=1..T}$. We then consider an image c where balls with known attributes are placed at the exact same location as the first frame of the target sequence and transfer the extracted target dynamics such that the rendered sequence is sampled following :

$$\{\mathbf{x}_t^{transfer}\}_{0:T} \sim \prod_{t=1}^{T} p(\mathbf{x}_t^{transfer}|\mathbf{z}_t) \prod_{k=1}^{K} p(\mathbf{z}_t^k | f(\mathbf{z}_{t-1}^k, \mathbf{r}_t^{target,k})) p(\mathbf{z}_0^k|\mathbf{c}).$$

In Figure 7 we show that VIM is able to recognize known dynamics of objects that have OOD attributes. VIM has successfully factorized individual objects identity and the shared dynamic rules that are applied in the same way irrespective of the objects identity.

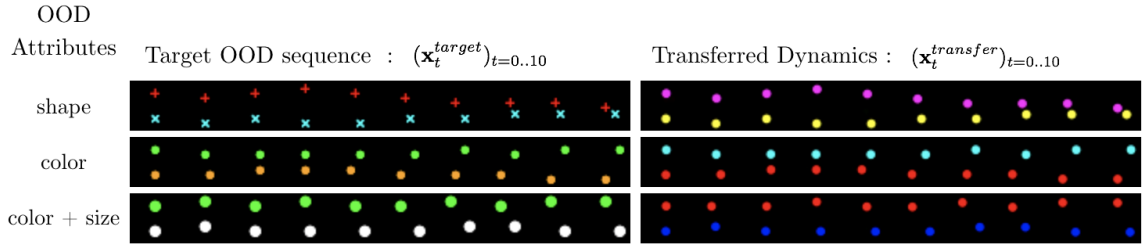

Figure 7: **Dynamics Transfer with OOD Objects**. VIM is able to recognize and transfer known dynamics of novel objects with OOD attributes. *Left*: Target trajectory from OOD objects from which we extract the module selection variables. *Right*: Transferred dynamics on known objects by applying the same selected modules as for the target sequences.

6. Conclusion

VIM introduces a variational framework to disentangle both the constituent entities (slots) and the generating factors (rules) of object dynamics in stochastic scenes. We show that the modular architecture of VIM captures different dynamics rules in an interpretable manner. Moreover, VIM modules can be composed at test time to explain OOD dynamics not seen during training. Additionally, we show that VIM is able to recognize the dynamics of objects with OOD attributes and transfer them to known objects by re-using the learned modules on different slots. The work presented here could be extended in different ways. For example, we could add interactions between slots through the use of n-ary modules. We could additionally scale up the model architecture and capacity to more complex datasets with in-the-wild 3D scenes and test its performance on downstream tasks. We plan to explore these directions in follow-up work.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.

Yoshua Bengio. The consciousness prior. *arXiv preprint arXiv:1709.08568*, 2017.

Yoshua Bengio, Tristan Deleu, Nasim Rahaman, Rosemary Ke, Sébastien Lachapelle, Olexa Bilaniuk, Anirudh Goyal, and Christopher Pal. A meta-transfer objective for learning to disentangle causal mechanisms. *arXiv:1901.10912*, 2019.

Christopher P. Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. Monet: Unsupervised scene decomposition and representation, 2019.

Lluis Castrejon, Nicolas Ballas, and Aaron Courville. Improved conditional vrnns for video prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7608–7617, 2019.

Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, pages 2980–2988, 2015.

Andy Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, 36(3):181–204, 2013.

Eric Crawford and Joelle Pineau. Spatially invariant unsupervised object detection with convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3412–3420, 2019.

Eric Crawford and Joelle Pineau. Exploiting spatial invariance for scalable unsupervised object tracking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3684–3692, 2020.

Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. In *International Conference on Machine Learning*, pages 1182–1191, 2018.

Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.

SM Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, Koray Kavukcuoglu, and Geoffrey E Hinton. Attend, infer, repeat: Fast scene understanding with generative models. *arXiv preprint arXiv:1603.08575*, 2016.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

Anirudh Goyal and Yoshua Bengio. Inductive biases for deep learning of higher-level cognition. *arXiv preprint arXiv:2011.15091*, 2020.

Anirudh Goyal, Alex Lamb, Phanideep Gampa, Philippe Beaudoin, Sergey Levine, Charles Blundell, Yoshua Bengio, and Michael Mozer. Object files and schemata: Factorizing declarative and procedural knowledge in dynamical systems.

Anirudh Goyal, Alex Lamb, Jordan Hoffmann, Shagun Sodhani, Sergey Levine, Yoshua Bengio, and Bernhard Schölkopf. Recurrent independent mechanisms. *arXiv preprint arXiv:1909.10893*, 2019.

Anirudh Goyal, Alex Lamb, Phanideep Gampa, Philippe Beaudoin, Sergey Levine, Charles Blundell, Yoshua Bengio, and Michael Mozer. Object files and schemata: Factorizing declarative and procedural knowledge in dynamical systems. *arXiv preprint arXiv:2006.16225*, 2020.

Anirudh Goyal, Aniket Didolkar, Nan Rosemary Ke, Charles Blundell, Philippe Beaudoin, Nicolas Heess, Michael Mozer, and Yoshua Bengio. Neural production systems. *arXiv preprint arXiv:2103.01937*, 2021.

Klaus Greff, Raphaël Lopez Kaufman, Rishabh Kabra, Nick Watters, Chris Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference, 2019.

Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *International Conference on Machine Learning*, pages 2555–2565. PMLR, 2019.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax, 2016.

Jindong Jiang, Sepehr Janghorbani, Gerard de Melo, and Sungjin Ahn. Scalable object-oriented sequential generative models. *Unknown Journal*, 2019.

Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.

Daniel Kahneman, Anne Treifsman, and Brian J Gibbs. The reviewing of object files: Object-specific integration of information. *Cognitive psychology*, 24(2):175–219, 1992.

Nal Kalchbrenner, Aäron van den Oord, Karen Simonyan, Ivo Danihelka, Oriol Vinyals, Alex Graves, and Koray Kavukcuoglu. Video pixel networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1771–1779. JMLR. org, 2017.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Adam R Kosiorek, Hyunjik Kim, Ingmar Posner, and Yee Whye Teh. Sequential attend, infer, repeat: Generative modelling of moving objects. *arXiv preprint arXiv:1806.01794*, 2018.

Jannik Kossen, Karl Stelzner, Marcel Hussing, Claas Voelcker, and Kristian Kersting. Structured object-aware physics prediction for video modeling and planning. *arXiv preprint arXiv:1910.02425*, 2019.

Hugo Larochelle and Iain Murray. The neural autoregressive distribution estimator. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 29–37, 2011.

Alex X Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. Stochastic adversarial video prediction. *arXiv preprint arXiv:1804.01523*, 2018.

Zhixuan Lin, Yi-Fu Wu, Skand Peri, Bofeng Fu, Jindong Jiang, and Sungjin Ahn. Improving generative imagination in object-centric world models. In *International Conference on Machine Learning*, pages 6140–6149. PMLR, 2020a.

Zhixuan Lin, Yi-Fu Wu, Skand Vishwanath Peri, Bofeng Fu, Jindong Jiang, and Sungjin Ahn. Improving generative imagination in object-centric world models. *CoRR*, abs/2010.02054, 2020b. URL https://arxiv.org/abs/2010.02054.

Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. 2018.

Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *arXiv preprint arXiv:2006.15055*, 2020.

Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015.

Sinéad L Mullally and Eleanor A Maguire. Memory, imagination, and predicting the future: a common brain mechanism? *The Neuroscientist*, 20(3):220–234, 2014.

Judea Pearl. Causality: Models, reasoning, and inference. 2nd edition. 2009.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.

MarcAurelio Ranzato, Arthur Szlam, Joan Bruna, Michael Mathieu, Ronan Collobert, and Sumit Chopra. Video (language) modeling: a baseline for generative models of natural videos. *arXiv preprint arXiv:1412.6604*, 2014.

Scott Reed, Aäron van den Oord, Nal Kalchbrenner, Sergio Gómez Colmenarejo, Ziyu Wang, Yutian Chen, Dan Belov, and Nando de Freitas. Parallel multiscale autoregressive density estimation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2912–2921. JMLR. org, 2017.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.

Elizabeth S Spelke, Karen Breinlinger, Kristen Jacobson, and Ann Phillips. Gestalt relations and object perception: A developmental study. *Perception*, 22(12):1483–1501, 1993.

Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852, 2015.

Karl Stelzner, Robert Peharz, and Kristian Kersting. Faster attend-infer-repeat with tractable probabilistic models. In *International Conference on Machine Learning*, pages 5966–5975. PMLR, 2019.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.

Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing motion and content for natural video sequence prediction. *arXiv preprint arXiv:1706.08033*, 2017a.

Ruben Villegas, Jimei Yang, Yuliang Zou, Sungryull Sohn, Xunyu Lin, and Honglak Lee. Learning to generate long-term future via hierarchical prediction. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3560–3569. JMLR. org, 2017b.

Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating visual representations from unlabeled video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 98–106, 2016.

Dirk Weissenborn, Oscar Täckström, and Jakob Uszkoreit. Scaling autoregressive video models. *International Conference on Learning Representations*, 2020.