

ROOTS: Object-Centric Representation and Rendering of 3D Scenes

Chang Chen*

*Department of Computer Science
Rutgers University
Piscataway, NJ 08854, USA*

CHANG.CHEN@RUTGERS.EDU

Fei Deng*

*Department of Computer Science
Rutgers University
Piscataway, NJ 08854, USA*

FEI.DENG@RUTGERS.EDU

Sungjin Ahn

*Department of Computer Science and Center for Cognitive Science
Rutgers University
Piscataway, NJ 08854, USA*

SUNGJIN.AHN@RUTGERS.EDU

Editor: Sebastian Nowozin

Abstract

A crucial ability of human intelligence is to build up models of individual 3D objects from partial scene observations. Recent works either achieve object-centric generation but without the ability to infer the representation, or achieve 3D scene representation learning but without object-centric compositionality. Therefore, learning to both represent and render 3D scenes with object-centric compositionality remains elusive. In this paper, we propose a probabilistic generative model for learning to build modular and compositional 3D object models from partial observations of a multi-object scene. The proposed model can (i) infer the 3D object representations by learning to search and group object areas, and also (ii) render from an arbitrary viewpoint not only individual objects but also the full scene by compositing the objects. The entire learning process is unsupervised and end-to-end. In experiments, in addition to generation quality, we also demonstrate that the learned representation permits object-wise manipulation and novel scene generation, and generalizes to various settings. Results can be found on our project website: <https://sites.google.com/view/roots3d>.

Keywords: object-centric representations, latent variable models, 3D scene generation, variational inference, 3D-aware representations

1. Introduction

At the core of human learning is the ability to build up mental models of the world along with the growing experience of our life. In building such models, a particularly important aspect is to factorize underlying structures of the world such as objects and their relationships. This ability is believed to be crucial in enabling various advanced cognitive functions in human-like AI systems (Lake et al., 2017) such as systematic generalization (Bahdanau et al., 2019; van

*. Both authors contributed equally.

Steenkiste et al., 2019), reasoning (Bottou, 2014), and causal learning (Schölkopf, 2019; Peters et al., 2017). While humans seem to learn such object-centric representations (Kahneman et al., 1992; Rolls et al., 2005; Hood and Santos, 2009; von Hofsten and Spelke, 1985; Martin, 2007; Høydal et al., 2019) in a 3D-aware fashion through partial observations of scenes without supervision, in machine learning this problem has only been tackled for simple 2D fully observable images (Eslami et al., 2016; Lin et al., 2020; Crawford and Pineau, 2019; Greff et al., 2017, 2019; Locatello et al., 2020; Burgess et al., 2019; Engelcke et al., 2020). Therefore, the more challenging yet realistic setting of learning 3D-aware object-centric representation of 3D space from partial observations has remained elusive.

Regarding this, there have been a number of recent approaches that can (only) *generate* 3D scene images via object-centric compositional rendering (Nguyen-Phuoc et al., 2020; van Steenkiste et al., 2020; Ehrhardt et al., 2020). However, none of the existing models provide the crucial ability of the reverse that we seek in this paper: the object-centric inverse graphics, i.e., learning object-centric 3D representations from partial observations. In learning representations and rendering of 3D scenes, GQN (Eslami et al., 2018) and its variants (Kumar et al., 2018; Tobin et al., 2019; Singh et al., 2019; Yoon et al., 2020) are the closest to our work. However, the 3D representations inferred by these models provide only scene-level representation without explicit object-centric decomposition.

In this paper, we tackle the problem of learning to build modular and compositional 3D object models from partial scene images. Our proposed model, ROOTS (Representation and Rendering of Object-Oriented Three-D ScenS), is able to decompose partial observations into objects, group them object-wise, and build a modular compositional 3D representation at the level of individual objects. Such representation also enables compositional rendering. As our object model provides object-wise 3D rendering from arbitrary viewpoints, we can also render the entire scene from arbitrary viewpoints by first rendering individual objects and then compositing them according to the scene layout. In particular, this enables a novel *nested autoencoder architecture* in which we can reuse the GQN model as an internal autoencoder module for object modeling, making the model simpler. The entire process is unsupervised and end-to-end trainable. We demonstrate the above capabilities of our model on simulated 3D scenes with multiple objects. We evaluate our model in terms of generation quality, structure accuracy, generalization ability, and downstream task performance. We also showcase that by manipulating the scene layout, we can generate scenes with many more objects than typical of the training regime.

2. Preliminary: Generative Query Networks

The Generative Query Network (GQN) is a latent variable model for learning to represent and render 3D scenes. Given a set of context images and viewpoints, it learns a *3D-viewpoint-steerable representation* (in short, *3D representation* throughout this paper) in the sense that any target image viewed from an arbitrary viewpoint can be generated from the representation. We note that such 3D representations are different from and more challenging to learn than *2D representations* that only model the scene from a single viewpoint. Recent advances in unsupervised object-centric representation learning (Eslami et al., 2016; Lin et al., 2020; Crawford and Pineau, 2019; Greff et al., 2017, 2019; Locatello et al., 2020; Burgess et al., 2019; Engelcke et al., 2020) mostly require the representation to model only

a single 2D image. Therefore, these methods can only learn 2D representations, even if the 2D image is a view of an underlying 3D scene.

More formally, consider an agent navigating a 3D environment (called a scene) and collecting K pairs of image \mathbf{x}_c and the corresponding viewpoint \mathbf{v}_c for $c = 1, 2, \dots, K$. This collection is called *context* $\mathcal{C} = \{(\mathbf{x}_c, \mathbf{v}_c)\}_{c=1}^K$. GQN learns a scene-level 3D representation \mathbf{z} by encoding \mathcal{C} , such that the target image $\hat{\mathbf{x}}_q$ from an arbitrary query viewpoint \mathbf{v}_q can be generated by the decoder $\hat{\mathbf{x}}_q = \text{GQN}_{\text{dec}}(\mathbf{z}, \mathbf{v}_q)$. The generative process can be written as:

$$p(\mathbf{x}_q | \mathbf{v}_q, \mathcal{C}) = \int p(\mathbf{x}_q | \mathbf{z}, \mathbf{v}_q) p(\mathbf{z} | \mathcal{C}) d\mathbf{z} .$$

The prior encoder $p(\mathbf{z} | \mathcal{C}) = \text{GQN}_{\text{enc}}(\mathcal{C})$ first obtains an order-invariant encoding (e.g., a sum encoding) $\mathbf{r}_{\mathcal{C}}$ of context \mathcal{C} , and then uses ConvDRAW (Gregor et al., 2016) to autoregressively sample \mathbf{z} from $\mathbf{r}_{\mathcal{C}}$. The decoder $\text{GQN}_{\text{dec}}(\mathbf{z}, \mathbf{v}_q)$ uses a deterministic version of ConvDRAW to render the target image $\hat{\mathbf{x}}_q$ from \mathbf{z} , and $p(\mathbf{x}_q | \mathbf{z}, \mathbf{v}_q)$ is often modeled as a Gaussian distribution $\mathcal{N}(\hat{\mathbf{x}}_q, \sigma^2 \mathbf{1})$ with σ being a hyperparameter. Since computing the posterior distribution $p(\mathbf{z} | \mathbf{x}_q, \mathbf{v}_q, \mathcal{C})$ is intractable, GQN uses variational inference for posterior approximation and is trained by maximizing its evidence lower bound. Backpropagation through random variables is done by the reparameterization trick (Kingma and Welling, 2014; Rezende et al., 2014).

Note that the model described above is actually a more consistent version of the GQN named CGQN (Kumar et al., 2018). In the original GQN (Eslami et al., 2018), the latent \mathbf{z} is also conditioned on \mathbf{v}_q , i.e., $p(\mathbf{z} | \mathbf{v}_q, \mathcal{C})$, and rendering is query-agnostic, i.e., $p(\mathbf{x}_q | \mathbf{z})$, leading to potential inconsistency across multiple query viewpoints. Throughout the paper, we use the abbreviation GQN to refer to the general GQN framework embracing both GQN and CGQN.

3. ROOTS

GQN represents a multi-object 3D scene as a single vector without learning explicit object-wise decomposition. Hence, it cannot entertain the potential and various advantages of object-centric representations. To resolve this limitation, we propose ROOTS, a probabilistic generative model that learns to represent and render 3D scenes via composition of object-centric 3D representations in a fully unsupervised and end-to-end trainable way. This problem has never been tackled, and it is highly challenging because not only can an object be unobservable from certain viewpoints, but also the appearance, position, pose, size, and occlusion of an object can vary significantly across the context images. The premise of our approach to tackling this challenge is that: if we can collect the local regions corresponding to a specific object across the context images, then we can reuse GQN on those filtered local observations to learn the 3D representation for that object.

To this end, we propose the following approaches. First, ROOTS has a *nested autoencoder architecture*, one autoencoder at scene level and the other at object level. Further, the scene-level autoencoder is constructed by the composition of the object-level autoencoders. For the scene-level encoding, the model encodes the context set to a 3D spatial structure of the scene and infers the 3D position of each object in the 3D space. Given the inferred 3D position of objects, we then propose a method, called *Attention-by-Perspective-Projection*

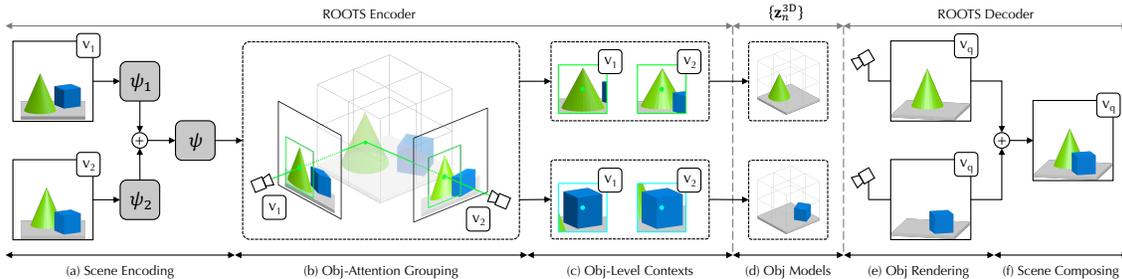


Figure 1: Overview of ROOTS pipeline. *ROOTS encoder* (a - c): (a) Context observations are encoded and aggregated into a scene-level representation ψ . (b) ψ is reorganized into a feature map of the 3D space, from which 3D center positions are inferred for each object. By applying perspective projection to the inferred 3D center positions, we identify image regions for each object across viewpoints. (c) Object regions are cropped and grouped into object-level contexts. *Object models* (d): The object-level contexts allow us to obtain the 3D appearance representation of each object through an object-level GQN. *ROOTS decoder* (e - f): To render the full scene for a given query viewpoint, we composite the rendering results of individual objects.

to efficiently find and attend the local regions, corresponding to a specific object, across all the context images. This grouping allows us to construct a new object-level context set containing only a specific object and thus to reuse the standard GQN encoder as an in-network module for object-level 3D-aware encoding of the object appearance. Scene-level decoding is also composed by object-level decoding and background decoding. We decode the appearance representation of each object using the object-level GQN decoder, and place the decoded images in the target image by mapping the 3D positions to the 2D positions in the target image. Together with background rendering, we can complete the rendering of a scene image. See Figure 1 for an overview of ROOTS pipeline.

3.1 ROOTS Encoder

The goal of ROOTS encoder is to infer the *3D object models* from scene-level context observations $\mathcal{C} = \{(\mathbf{x}_c, \mathbf{v}_c)\}_{c=1}^K$ but without any object-level supervision. Each object model consists of the 3D representation of an object, fully disentangled into its 3D position and 3D appearance. The modularity and compositionality of these object models allow them to be collected from multiple scenes, and then reconfigured to generate novel scenes that are out of the training distribution.

To infer the 3D object models, it is imperative that the encoder should be properly structured. In particular, we find in our experiments that directly inferring object models from an order-invariant encoding of \mathcal{C} would fail, potentially because the lack of proper structure prohibits learning and optimization. To solve this problem, we extract object

regions from each scene image \mathbf{x}_c and group them into object-level contexts, which provide more relevant information for inferring the 3D appearance of each object. We call this grouping process object-attention grouping.

For 2D fully observable images, extracting object regions can be solved by recent 2D scene decomposition methods (Eslami et al., 2016; Crawford and Pineau, 2019; Lin et al., 2020). However, in our 3D and partially observed setting, it remains a challenge to efficiently group together the regions that correspond to the same object across viewpoints. One naive approach is to find the best among all possible groupings, but its time complexity is exponential in the number of viewpoints. Another possible way is to treat the extracted object regions in one viewpoint as anchors, and match regions from other viewpoints to one of these anchors, by computing pairwise matching scores. The time complexity is quadratic in the number of objects. By contrast, our proposed object-attention grouping scales linearly in both the number of viewpoints and the number of objects. The key idea is to first infer the center position of each object in 3D coordinates. This allows us to use perspective projection (Hartley and Zisserman, 2003) from 3D to 2D to efficiently locate the same object across different context images.

We develop a scene encoder (Section 3.1.1) to infer the object positions, describe in more detail the object-attention grouping in Section 3.1.2, and use an object-level GQN encoder (Section 3.1.3) to infer the object appearance.

3.1.1 SCENE ENCODER

The goal of the scene encoder is to infer the 3D center position of each object in world coordinates—the same coordinate system where the camera viewpoints \mathbf{v}_c are measured. We assume that the objects resides in a bounded 3D space. The scene encoder partitions this bounded 3D space into a rectangular cuboid of $N_{\max} = N_x \times N_y \times N_z$ cells. For each cell (i, j, k) , we infer a Bernoulli variable $z_{ijk}^{\text{pres}} \in \{0, 1\}$ that is 1 if and only if the cell contains the center of an object (note that the full appearance volume of an object need not be contained within the cell). We also infer a continuous variable $\mathbf{z}_{ijk}^{\text{where}} \in \mathbb{R}^3$ that, when $z_{ijk}^{\text{pres}} = 1$, specifies the center position of the object in the 3D world coordinates. Here $\mathbf{z}_{ijk}^{\text{where}}$ is constrained to be within the boundary of cell (i, j, k) . This prior on cell-wise object preference helps efficient training and obviates the need for expensive autoregressive processing (Lin et al., 2020).

In the above scene encoding, each cell handles one or no object. In actual implementation, however, the partition is soft, meaning neighboring cells can have some overlap. Hence, when a cell does contain more than one object, the scene encoder can learn to distribute them to adjacent cells. A similar 2D version of this approach has been used in SPAIR (Crawford and Pineau, 2019) and SPACE (Lin et al., 2020), showing impressive decomposition of 2D scenes into dozens of objects.

Specifically, to infer $\{z_{ijk}^{\text{pres}}, \mathbf{z}_{ijk}^{\text{where}}\}$ from the context observations \mathcal{C} , we encode \mathcal{C} into a Geometric Volume Feature Map (GVFM) $\mathbf{r} \in \mathbb{R}^{N_x \times N_y \times N_z \times d}$, yielding a d -dimensional feature vector \mathbf{r}_{ijk} for each cell. Then, $\{z_{ijk}^{\text{pres}}, \mathbf{z}_{ijk}^{\text{where}}\}$ can be computed in parallel for all cells by a neural network $f_{\text{pres,where}}$:

$$p(z_{ijk}^{\text{pres}}, \mathbf{z}_{ijk}^{\text{where}} \mid \mathcal{C}) = f_{\text{pres,where}}(\mathbf{r}_{ijk}^{\text{neighbor}}),$$

where $\mathbf{r}_{ijk}^{\text{neighbor}}$ includes the feature vectors of cell (i, j, k) and its neighboring cells, allowing inter-object relations to be taken into consideration.

While GVFM may seem similar to the grid cells used in SPAIR and SPACE, there are fundamental differences. As a feature map of the 3D space, GVFM must *aggregate* information from multiple partial 2D observations and *reorganize* it in an object-wise fashion. This is in contrast to the 2D feature map learned by grid cells which have a natural alignment with the single fully observed 2D image. Therefore, we obtain GVFM in two steps. First, we compute an order-invariant summary ψ of \mathcal{C} as the summation over encodings of individual context observations:

$$\psi = \sum_{c=1}^K \psi_c = \sum_{c=1}^K f_\psi(\mathbf{x}_c, \mathbf{v}_c) ,$$

where f_ψ is a learned encoding network. Second, we apply a 3D transposed convolution over ψ to turn the sum of 2D image representations ψ into 3D spatial representation \mathbf{r} where individual \mathbf{r}_{ijk} slots contains object-specific information:

$$\mathbf{r} = \text{ConvTranspose3D}(\psi) .$$

3.1.2 OBJECT-ATTENTION GROUPING

Object-attention grouping aims to identify image regions that correspond to the same object across different observation images. This is crucial in obtaining object-wise 3D appearance representations. More precisely, for each object n present in the scene and each context image \mathbf{x}_c , we seek a 2D bounding box capturing object n in \mathbf{x}_c . The bounding box is parameterized by its center position and scale (width and height), denoted $(\mathbf{o}_{n,c}^{\text{center}}, \mathbf{o}_{n,c}^{\text{scale}})$. Notice that here each object index n corresponds to a distinct cell index (i, j, k) with $z_{ijk}^{\text{pres}} = 1$.

Our key observation is that inferring the 3D object center positions in the first step allows us to solve object-attention grouping by using perspective projection. We call this Attention-by-Perspective-Projection (APP). Assuming that the projection operation takes constant time, the time complexity of APP is linear in both the number of objects and the number of viewpoints.

Attention-by-Perspective-Projection (APP). Let us focus on object n and find its 2D bounding box in \mathbf{x}_c . We first analytically compute its 2D center position $\mathbf{o}_{n,c}^{\text{center}} \in \mathbb{R}^2$ in \mathbf{x}_c and its distance from the camera, denoted $o_{n,c}^{\text{depth}} \in \mathbb{R}$, by applying perspective projection to its 3D center position $\mathbf{z}_n^{\text{where}}$:

$$[\mathbf{o}_{n,c}^{\text{center}}, o_{n,c}^{\text{depth}}]^\top = \text{APP}_{\text{pos}}(\mathbf{z}_n^{\text{where}}, \mathbf{v}_c) = \text{normalize}(\mathbf{T}_{\text{World} \rightarrow \text{Camera}}(\mathbf{v}_c)[\mathbf{z}_n^{\text{where}}, 1]^\top) .$$

Here, $\mathbf{z}_n^{\text{where}}$ is first converted to camera coordinates by the viewpoint-dependent transformation matrix $\mathbf{T}_{\text{World} \rightarrow \text{Camera}}(\mathbf{v}_c) \in \mathbb{R}^{3 \times 4}$, and then normalized. See Appendix E for more details.

To compute the 2D bounding box scale $\mathbf{o}_{n,c}^{\text{scale}} \in \mathbb{R}^2$, one option is to learn a 3D bounding box for object n , project its eight vertices onto the image plane, and find the smallest rectangle that covers all eight vertices. Unfortunately, the resulting 2D bounding box will only be tight under specific viewpoints, and we will likely encounter optimization difficulties. Hence, to allow better gradient flow and provide the model with the opportunity to predict tighter 2D bounding boxes, we design $\text{APP}_{\text{scale}}$ that implicitly learns the projection:

$$p(\mathbf{o}_{n,c}^{\text{scale}} | \mathbf{z}_n^{\text{where}}, \mathcal{C}) = \text{APP}_{\text{scale}}(\mathbf{o}_{n,c}^{\text{center}}, o_{n,c}^{\text{depth}}, \mathbf{r}_n, \mathbf{v}_c) = \text{MLP}(\text{concat}[\mathbf{o}_{n,c}^{\text{center}}, o_{n,c}^{\text{depth}}, \mathbf{r}_n, \mathbf{v}_c]) .$$

To work properly, $\text{APP}_{\text{scale}}$ should learn to perform the following operations implicitly: (i) extract 3D scale information from \mathbf{r}_n , (ii) make a projection from viewpoint \mathbf{v}_c , and (iii) refine the projection using $\mathbf{o}_{n,c}^{\text{center}}$ and $\mathbf{o}_{n,c}^{\text{depth}}$.

3.1.3 OBJECT ENCODER

With object-attention grouping, we can decompose the scene-level context \mathcal{C} into object-level context \mathcal{C}_n for each object $n = 1, 2, \dots, N$, where

$$N = \sum_{ijk} z_{ijk}^{\text{pres}} \leq N_{\text{max}}$$

is the total number of objects present in the scene. Specifically, we first use a spatial transformer \mathcal{ST} (Jaderberg et al., 2015) to differentially crop object patch $\mathbf{x}_{n,c}^{\text{att}}$ from scene image \mathbf{x}_c using $\mathbf{o}_{n,c}^{\text{center}}$ and $\mathbf{o}_{n,c}^{\text{scale}}$:

$$\mathbf{x}_{n,c}^{\text{att}} = \mathcal{ST}(\mathbf{x}_c, \mathbf{o}_{n,c}^{\text{center}}, \mathbf{o}_{n,c}^{\text{scale}}).$$

After collecting these patches from all viewpoints, we group them based on the object index n to obtain object-level context

$$\mathcal{C}_n = \{(\mathbf{x}_{n,c}^{\text{att}}, \mathbf{v}_c, \mathbf{o}_{n,c}^{\text{where}})\}_{c=1}^K,$$

where we include $\mathbf{o}_{n,c}^{\text{where}} = (\mathbf{o}_{n,c}^{\text{center}}, \mathbf{o}_{n,c}^{\text{scale}}, \mathbf{o}_{n,c}^{\text{depth}})$ to provide information complementary to $\mathbf{x}_{n,c}^{\text{att}}$. The object-level context allows us to use an object-level GQN encoder

$$p(\mathbf{z}_n^{\text{what}} | \mathcal{C}_n) = \text{GQN}_{\text{enc}}(\mathcal{C}_n)$$

to obtain independent and modular object-level 3D appearance $\mathbf{z}_n^{\text{what}}$ for each object n . A summary of ROOTS encoder is provided in Appendix D.

3.2 ROOTS Decoder

Given partial observations of a multi-object 3D scene, ROOTS not only learns to infer the 3D object models, but also learns to render them independently and individually from arbitrary viewpoints. The full scene is also rendered from arbitrary query viewpoints by compositing object rendering results. By collecting and re-configuring the inferred object models, ROOTS can easily generate novel scenes that are out of the training distribution.

Object Renderer. For each object n , given its 3D appearance representation $\mathbf{z}_n^{\text{what}}$ and a query viewpoint \mathbf{v}_q , ROOTS is able to generate a 4-channel (RGB+mask) image $\mathbf{o}_{n,q}^{\text{what}}$ depicting the object’s 2D appearance when viewed from \mathbf{v}_q . This is achieved by an object-level GQN decoder:

$$\mathbf{o}_{n,q}^{\text{what}} = \text{GQN}_{\text{dec}}(\text{concat}[\mathbf{z}_n^{\text{what}}, \mathbf{r}_n^{\text{att}}], \mathbf{v}_q).$$

Here, $\mathbf{r}_n^{\text{att}}$ is an order-invariant summary of object-level context \mathcal{C}_n :

$$\mathbf{r}_n^{\text{att}} = \sum_{c=1}^K f_{\text{att}}(\mathbf{x}_{n,c}^{\text{att}}, \mathbf{v}_c, \mathbf{o}_{n,c}^{\text{where}}),$$

where f_{att} is a learnable encoding network.

Scene Composer. The final scene image $\hat{\mathbf{x}}_q$ corresponding to query \mathbf{v}_q is obtained by superimposing layers of object-wise images with proper masking. For this, we first use an inverse spatial transformer \mathcal{ST}^{-1} (Jaderberg et al., 2015) to differentiably place each object at the right position in the scene canvas with proper scaling:

$$[\hat{\mathbf{x}}_{n,q}, \hat{\mathbf{m}}_{n,q}] = \mathcal{ST}^{-1}(\mathbf{o}_{n,q}^{\text{what}}, \mathbf{o}_{n,q}^{\text{center}}, \mathbf{o}_{n,q}^{\text{scale}}) .$$

Here, the 3-channel image $\hat{\mathbf{x}}_{n,q}$ can be regarded as an object-specific image layer containing only object n , and the single-channel $\hat{\mathbf{m}}_{n,q}$ is the mask for object n . The position and scaling parameters are computed by simply reusing the APP module for query viewpoint \mathbf{v}_q :

$$\begin{aligned} [\mathbf{o}_{n,q}^{\text{center}}, \mathbf{o}_{n,q}^{\text{depth}}]^\top &= \text{APP}_{\text{pos}}(\mathbf{z}_n^{\text{where}}, \mathbf{v}_q) , \\ p(\mathbf{o}_{n,q}^{\text{scale}} \mid \mathbf{z}_n^{\text{where}}, \mathbf{v}_q, \mathcal{C}) &= \text{APP}_{\text{scale}}(\mathbf{o}_{n,q}^{\text{center}}, \mathbf{o}_{n,q}^{\text{depth}}, \mathbf{r}_n, \mathbf{v}_q) . \end{aligned}$$

We then composite these N image layers into a single image, ensuring that occlusion among objects is properly handled. Similar to previous works (van Steenkiste et al., 2020; Crawford and Pineau, 2019; Burgess et al., 2019; Greff et al., 2019; Engelcke et al., 2020; Lin et al., 2020), for each layer n , we compute a transparency map

$$\alpha_{n,q} = \mathbf{w}_{n,q} \odot \hat{\mathbf{m}}_{n,q} ,$$

where \odot is pixel-wise multiplication. This masks out occluded pixels of object n . To obtain the values of $\{\mathbf{w}_{n,q}\}_{n=1}^N$ at each pixel, we first use $\{\hat{\mathbf{m}}_{n,q}\}_{n=1}^N$ to find the objects that contain the pixel, and then assign the values based on their relative depth $\{\mathbf{o}_{n,q}^{\text{depth}}\}_{n=1}^N$. See Appendix F for more details. The final rendered scene $\hat{\mathbf{x}}_q$ is composited as:

$$\hat{\mathbf{x}}_q = \sum_{n=1}^N \alpha_{n,q} \odot \hat{\mathbf{x}}_{n,q} .$$

A summary of ROOTS decoder is provided in Appendix D.

3.3 Probabilistic Model

We now piece things together and formulate ROOTS as a conditional generative model. Given a collection of context observations $\mathcal{C} = \{(\mathbf{x}_c, \mathbf{v}_c)\}_{c=1}^K$ of a multi-object scene, ROOTS learns to infer the number of objects, denoted N , the 3D object model $\mathbf{z}_n^{\text{3D}} = (\mathbf{z}_n^{\text{where}}, \mathbf{z}_n^{\text{what}})$ for each object $n = 1, 2, \dots, N$, and the 2D representation $\mathbf{o}_{n,\mathcal{C}}^{\text{2D}} = \{\mathbf{o}_{n,c}^{\text{where}}\}_{c=1}^K$ collected for each object n from all context viewpoints. In addition, ROOTS also learns a background representation \mathbf{z}^{bg} through a scene-level GQN encoder:

$$p(\mathbf{z}^{\text{bg}} \mid \mathcal{C}) = \text{GQN}_{\text{enc}}^{\text{bg}}(\mathcal{C}) .$$

Using these representations, ROOTS can then generate the target image \mathbf{x}_q from an arbitrary query viewpoint \mathbf{v}_q of the same scene. During generation, ROOTS also infers the 2D object representation $\mathbf{o}_{n,q}^{\text{2D}} = \mathbf{o}_{n,q}^{\text{where}}$ for the query viewpoint. We do not include $\mathbf{o}_{n,q}^{\text{what}}$ here because it is a deterministic variable.

Let $\mathcal{Q} = \{(\mathbf{x}_q, \mathbf{v}_q)\}_{q=1}^M$ be the collection of queries for the same scene, $\mathbf{x}_{\mathcal{Q}} = \{\mathbf{x}_q\}_{q=1}^M$ and $\mathbf{v}_{\mathcal{Q}} = \{\mathbf{v}_q\}_{q=1}^M$ be the target images and query viewpoints respectively, and $\mathcal{D} = \mathcal{C} \cup \mathcal{Q}$ be the

union of contexts and queries. To simplify notations, we collect all viewpoint-independent 3D representations into a single variable \mathbf{z}^{3D} , including the number of objects, the 3D object models, and the background representation:

$$\mathbf{z}^{3D} = (N, \{\mathbf{z}_n^{3D}\}_{n=1}^N, \mathbf{z}^{bg}) .$$

We also collect the viewpoint-dependent 2D representations for all objects into a single variable $\mathbf{o}_{\mathcal{S}}^{2D}$, where the subscript \mathcal{S} denotes the set of viewpoints. For example,

$$\mathbf{o}_{\mathcal{C}}^{2D} = \{\mathbf{o}_{n,\mathcal{C}}^{2D}\}_{n=1}^N , \quad \mathbf{o}_q^{2D} = \{\mathbf{o}_{n,q}^{2D}\}_{n=1}^N .$$

The generative process can then be written as:

$$p(\mathbf{x}_{\mathcal{Q}} | \mathbf{v}_{\mathcal{Q}}, \mathcal{C}) = \iint \underbrace{p(\mathbf{z}^{3D}, \mathbf{o}_{\mathcal{C}}^{2D} | \mathcal{C})}_{\text{Encoder}} \prod_{q=1}^M \underbrace{p(\mathbf{o}_q^{2D} | \mathbf{z}^{3D}, \mathbf{v}_q, \mathcal{C})}_{\text{Object Renderer}} \underbrace{p(\mathbf{x}_q | \mathbf{z}^{3D}, \mathbf{o}_{\mathcal{C} \cup q}^{2D}, \mathbf{v}_q, \mathcal{C})}_{\text{Scene Composer}} d\mathbf{z}^{3D} d\mathbf{o}_{\mathcal{D}}^{2D} .$$

The encoder can be further factorized in an object-wise fashion:

$$p(\mathbf{z}^{3D}, \mathbf{o}_{\mathcal{C}}^{2D} | \mathcal{C}) = \underbrace{p(\mathbf{z}^{bg} | \mathcal{C})}_{\text{Background}} \underbrace{p(N | \mathcal{C})}_{\text{Density}} \prod_{n=1}^N \underbrace{p(\mathbf{z}_n^{\text{where}} | \mathcal{C})}_{\text{Scene Encoder}} \underbrace{p(\mathbf{o}_{n,\mathcal{C}}^{2D} | \mathbf{z}_n^{\text{where}}, \mathcal{C})}_{\text{APP}} \underbrace{p(\mathbf{z}_n^{\text{what}} | \mathcal{C}_n)}_{\text{Object Encoder}} ,$$

where the object-level context \mathcal{C}_n is obtained as a deterministic function of $\mathbf{z}_n^{\text{where}}$, $\mathbf{o}_{n,\mathcal{C}}^{2D}$, and \mathcal{C} . The object renderer can be factorized similarly:

$$p(\mathbf{o}_q^{2D} | \mathbf{z}^{3D}, \mathbf{v}_q, \mathcal{C}) = \prod_{n=1}^N \underbrace{p(\mathbf{o}_{n,q}^{2D} | \mathbf{z}_n^{\text{where}}, \mathbf{v}_q, \mathcal{C})}_{\text{APP}} .$$

The scene composer obtains the full scene from the foreground image $\hat{\mathbf{x}}_q$ and the background image $\hat{\mathbf{x}}_q^{\text{bg}}$ through alpha compositing:

$$p(\mathbf{x}_q | \mathbf{z}^{3D}, \mathbf{o}_{\mathcal{C} \cup q}^{2D}, \mathbf{v}_q, \mathcal{C}) = \mathcal{N}\left(\hat{\mathbf{x}}_q + (\mathbf{1} - \sum_{n=1}^N \alpha_{n,q}) \odot \hat{\mathbf{x}}_q^{\text{bg}}, \sigma^2 \mathbf{1}\right) ,$$

where $\hat{\mathbf{x}}_q^{\text{bg}}$ is rendered by a scene-level GQN decoder:

$$\hat{\mathbf{x}}_q^{\text{bg}} = \text{GQN}_{\text{dec}}^{\text{bg}}(\mathbf{z}^{\text{bg}}, \mathbf{v}_q) ,$$

and σ^2 is a hyperparameter called pixel-variance.

3.4 Inference and Learning

Due to the intractability of the log-likelihood $\log p(\mathbf{x}_{\mathcal{Q}} | \mathbf{v}_{\mathcal{Q}}, \mathcal{C})$, we train ROOTS using variational inference with the following approximate posterior:

$$q(\mathbf{z}^{3D}, \mathbf{o}_{\mathcal{D}}^{2D} | \mathcal{D}) = q(\mathbf{z}^{\text{bg}} | \mathcal{D}) q(N | \mathcal{D}) \prod_{n=1}^N q(\mathbf{z}_n^{\text{where}} | \mathcal{D}) q(\mathbf{o}_{n,\mathcal{D}}^{2D} | \mathbf{z}_n^{\text{where}}, \mathcal{D}) q(\mathbf{z}_n^{\text{what}} | \mathcal{D}_n) ,$$

where \mathcal{D}_n is the object-level context deterministically obtained from \mathcal{D} using the inferred $\mathbf{z}_n^{\text{where}}$ and $\mathbf{o}_{n,\mathcal{D}}^{2D}$. The implementation of the approximate posterior is almost the same as ROOTS encoder described in Section 3.1, except that the summary vector ψ should now encode the entire \mathcal{D} instead of only \mathcal{C} . We treat all continuous variables as Gaussian variables, and use reparameterization trick (Kingma and Welling, 2014) to sample from the approximate posterior. For discrete variables, we use Gumbel-Softmax trick (Jang et al., 2017; Maddison et al., 2017). The entire model can be trained end-to-end by maximizing the Evidence Lower Bound (ELBO):

$$\begin{aligned} \mathcal{L} = & \mathbb{E}_{q(\mathbf{z}^{3D}, \mathbf{o}_{\mathcal{D}}^{2D})} \left[\sum_{q=1}^M \log p(\mathbf{x}_q | \mathbf{z}^{3D}, \mathbf{o}_{\mathcal{C} \cup q}^{2D}, \mathbf{v}_q, \mathcal{C}) \right] - D_{\text{KL}}[q(\mathbf{z}^{\text{bg}} | \mathcal{D}) \| p(\mathbf{z}^{\text{bg}} | \mathcal{C})] \\ & - D_{\text{KL}}[q(N | \mathcal{D}) \| p(N | \mathcal{C})] - \mathbb{E}_{q(N|\mathcal{D})} \left[\sum_{n=1}^N D_{\text{KL}}[q(\mathbf{z}_n^{3D}, \mathbf{o}_{n,\mathcal{D}}^{2D} | \mathcal{D}) \| p(\mathbf{z}_n^{3D}, \mathbf{o}_{n,\mathcal{D}}^{2D} | \mathbf{v}_{\mathcal{Q}}, \mathcal{C})] \right]. \end{aligned}$$

Combining with Unconditioned Prior. One difficulty in using the conditional prior is that it may not reflect the inductive bias we have for the latent variables. In our experiments, it turns out that introducing inductive bias to the posterior of some variables helps stabilize the model. We achieve this by adding KL terms between the posterior and *unconditioned* prior (like in VAEs, Kingma and Welling 2014; Higgins et al. 2017) to the ELBO. Specifically, the model is trained by maximizing:

$$\begin{aligned} \tilde{\mathcal{L}} = & \mathcal{L} - \gamma D_{\text{KL}}[q(N | \mathcal{D}) \| \text{Geom}(\rho)] - \mathbb{E}_{q(N|\mathcal{D})} \left[\sum_{n=1}^N D_{\text{KL}}[q(\mathbf{z}_n^{\text{where}} | \mathcal{D}) \| \mathcal{N}(\mathbf{0}, \mathbf{1})] \right] \\ & - \mathbb{E}_{q(N|\mathcal{D})} \left[\sum_{n=1}^N \mathbb{E}_{q(\mathbf{z}_n^{\text{where}}|\mathcal{D})} [D_{\text{KL}}[q(\mathbf{o}_{n,\mathcal{D}}^{2D} | \mathbf{z}_n^{\text{where}}, \mathcal{D}) \| \mathcal{N}(\mathbf{0}, \mathbf{1})]] \right]. \end{aligned}$$

Here, γ is a weighting hyperparameter, and $\text{Geom}(\rho)$ is a truncated Geometric distribution with support $\{0, 1, \dots, N_{\text{max}}\}$ and success probability ρ . We set $\gamma = 7$ and $\rho = 0.999$ during training, thereby encouraging the model to decompose the scenes into as few objects as possible.

4. Related Work

ROOTS is broadly related to recent advances in learning representations for the appearance and geometry of 3D scenes, and more closely related to those that do not require 3D supervision. ROOTS is also inspired by recent works that learn to decompose 2D scenes into object-wise representations.

Geometric Deep Learning. Learning representations that capture the geometry of 3D scenes has been of growing interest. Recent works have explored integrating voxels (Maturana and Scherer, 2015; Kar et al., 2017; Tulsiani et al., 2017; Wu et al., 2016; Choy et al., 2016), meshes (Kato et al., 2018; Kanazawa et al., 2018), point clouds (Qi et al., 2017; Achlioptas et al., 2018), and many other classical representations into deep learning models to achieve better 3D scene understanding. However, they often require 3D supervision (Huang et al., 2018; Tulsiani et al., 2018; Cheng et al., 2018; Shin et al., 2019; Du et al., 2018) and work on single-object scenes (Wu et al., 2016; Yan et al., 2016; Choy et al., 2016; Kar et al., 2017; Nguyen-Phuoc et al., 2019). By contrast, ROOTS learns to decompose a multi-object scene into object-wise representations without being given the groundtruth scene-level or object-level 3D representations.

Neural Representation of 3D Scenes. Recent works (Eslami et al., 2018; Kumar et al., 2018; Tobin et al., 2019; Tung et al., 2019; Sitzmann et al., 2019a,b; Singh et al., 2019; Mildenhall et al., 2020; Dupont et al., 2020) have explored learning 3D scene representations from 2D images without 3D supervision. While the rendering quality (Tobin et al., 2019; Sitzmann et al., 2019b; Mildenhall et al., 2020) and efficiency (Dupont et al., 2020) have been improved, these methods are not able to decompose the full scene into objects without object-level supervision, and cannot learn object-wise representation and rendering models. We believe these works are complementary to ROOTS and may allow object models to be learned from more realistic scenes. Crawford and Pineau (2020) recently proposed to learn 3D object-centric representations from unlabeled videos. Although their model can infer the 3D position of each object, the object appearance is modeled in 2D. Another line of work (Nguyen-Phuoc et al., 2020; Liao et al., 2020) learns object-aware 3D scene representations for generative adversarial networks (Goodfellow et al., 2014). They only support rendering and are unable to infer the object models for a given scene.

Object-Oriented Representation of 2D Images. There have been prolific advances in unsupervised object-oriented representation learning from fully observed 2D images. They mainly fall into two categories: detection-based and mixture-based. The detection-based approaches (Eslami et al., 2016; Crawford and Pineau, 2019; Lin et al., 2020) first identify object regions and then learn object representations from object patches cropped by the spatial transformer (Jaderberg et al., 2015). The mixture-based approaches (Greff et al., 2017; Burgess et al., 2019; Greff et al., 2019; Engelcke et al., 2020; Locatello et al., 2020) model the observed image as a pixel-level Gaussian mixture where each component is expected to capture a single object. None of these approaches consider the 3D structure of the scene, let alone the 3D appearance of objects.

5. Experiments

In this section, we evaluate the quality of object models learned by ROOTS and demonstrate the benefits they bring in terms of generation quality, generalization ability, and downstream task performance. We also showcase the built-in compositionality and disentanglement properties of ROOTS. We first introduce the data sets and baselines we use, and then show both qualitative and quantitative results.

Data Sets. Existing data sets in previous work on unsupervised 3D scene representation learning (Eslami et al., 2018; Tobin et al., 2019) either do not contain multi-object scenes or cannot provide object-wise groundtruth information, and thus cannot serve our purpose. Hence, we created two data sets: the Shapes data set and the Multi-Shepard-Metzler (MSM) data set, using MuJoCo (Todorov et al., 2012) and Blender (Blender Online Community, 2017) respectively. Both data sets contain 60K multi-object scenes (50K for training, 5K for validation, and 5K for testing) with complete groundtruth scene specifications including object shapes, colors, positions, and sizes. Each scene is rendered as 128×128 color images from 30 random viewpoints. During training, we sample 10-20 viewpoints uniformly at random as contexts and use the rest as queries. For evaluation and visualization, we use 15 viewpoints as contexts and the rest as queries. Notice that the scene specifications are for evaluation only and are not used during training.

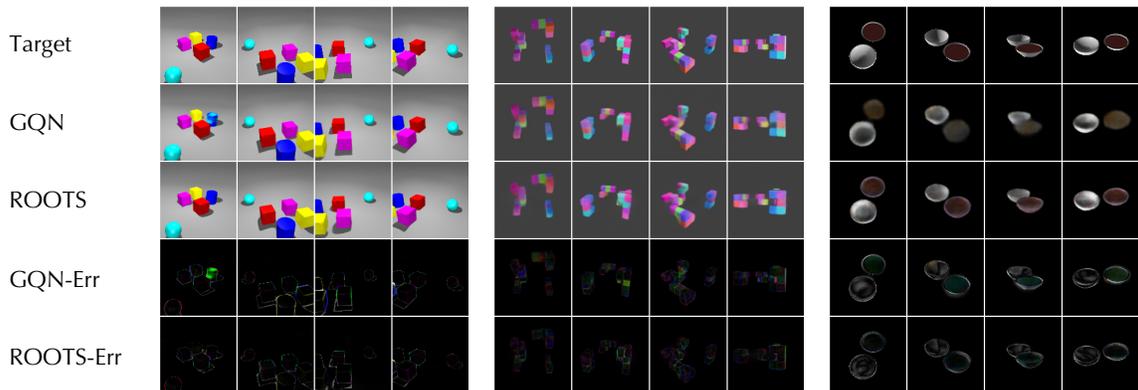


Figure 2: Sample generations from three scenes. Columns correspond to query viewpoints. ROOTS gives better generations in regions of occlusion while GQN sometimes misses occluded objects and predicts wrong colors. GQN-Err and ROOTS-Err are difference maps (pixel-wise absolute difference) between targets and generations of GQN and ROOTS, respectively.

We generated three versions of the Shapes data set, containing scenes with 1-3, 2-4, and 3-5 objects respectively. The position, size, shape, and color of the objects are randomized. The MSM data set contains scenes with 2-4 randomly positioned Shepard-Metzler objects. Each object consists of 5 cubes whose positions are generated by a self-avoiding random walk. The color of each cube is independently sampled from a continuous color space, as described in GQN (Eslami et al., 2018). Since these objects have complex shapes randomly generated per scene, they span a large combinatorial space, and it is unlikely that two different scenes will share a same object. Also, the objects can have severe occlusion with each other, making this data set significantly more challenging than the single-object version considered in GQN.

For evaluation on realistic objects, we also included a publicly available ShapeNet arrangement data set (Tung et al., 2019; Cheng et al., 2018). Each scene of this data set consists of 2 ShapeNet (Chang et al., 2015) objects placed on a table surface, and is rendered from 54 fixed cameras positioned on the upper hemisphere. Following prior work, we split the data set into a training set of 300 scenes and a test set of 32 scenes containing unseen objects. During training, we sample 2-4 viewpoints uniformly at random as contexts, and also sample the queries so that the number of contexts and queries combined for each scene is 6. For evaluation and visualization, we use 3 contexts and 3 queries. The contexts and queries are guaranteed to be distinct. Because object-wise annotations are not available, we did not perform quantitative evaluation of object-level decomposition on this data set.

Baselines. Because there is no previous work that can build 3D object models from multi-object scene images, we use separate baselines to evaluate scene-level representation and object-level decomposition respectively. For scene-level representation and generation quality, we use CGQN (Kumar et al., 2018) as the baseline model, and refer to it as GQN in the rest of this section to indicate the general GQN framework. For object-level

decomposition, we compare the image segmentation ability embedded in ROOTS with that of IODINE (Greff et al., 2019), which focuses on this ability without learning 3D representations.

5.1 Qualitative Evaluation

In this section, we qualitatively evaluate the learned object models by showing scene generation and decomposition results and object model visualizations. We also demonstrate the built-in compositionality and disentanglement properties by compositing novel scenes out of the training distribution and visualizing latent traversals, respectively.

Scene Generation. Like GQN, ROOTS is able to generate target observations for a given scene from arbitrary query viewpoints. Figure 2 shows a comparison of scene generations using 15 contexts. ROOTS gives better generations in regions of occlusion (especially on the MSM data set), and correctly infers partially observable objects (e.g., the yellow cube in the 4th column). In contrast, GQN tends to miss heavily occluded and partially observable objects, and sometimes predicts wrong colors. As highlighted in the difference maps in Figure 2, on the Shapes data set, GQN sometimes generates inconsistent colors within an object. On the MSM data set, GQN samples may look slightly clearer than those of ROOTS as GQN generates sharper boundaries between the unit cubes. However, the difference map reveals that GQN more frequently draws the objects with wrong colors. On the ShapeNet arrangement data set, GQN samples are more blurry and also with wrong colors. We believe that the object models learned by ROOTS and the object-level modular rendering provide ROOTS with a stronger capacity to represent the appearance of individual objects, leading to its better generation quality.

Object Models. We further visualize the learned object models in Figure 4A, by applying the object renderer to $\mathbf{z}_n^{\text{what}}$ and a set of query viewpoints. We also show the scene rendering process in Figure 3, where object rendering results are composited to generate the full scene. As can be seen, from images containing multiple objects with occlusion, ROOTS is able to learn the complete 3D appearance of each object, predict accurate object positions, and correctly handle occlusion. Such object models are not available from GQN because it only learns scene-level representations.

Compositionality. Once object models are learned, they can be reconfigured to form novel scenes that are out of the training distribution. As an example, in Figure 4B, we first provide ROOTS with context images from three scenes (top three rows) with 3 objects each, and collect the learned object representations $\{(\mathbf{r}_n^{\text{att}}, \mathbf{z}_n^{\text{where}}, \mathbf{z}_n^{\text{what}})\}$. A new scene with 9 objects can then be composed and rendered from arbitrary query viewpoints. Rendering results are shown in the bottom row of Figure 4B. We would like to emphasize that the model is trained on scenes with 1-3 objects. Thus, a scene with 9 objects has never been seen during training.

Disentanglement. Since object position and appearance are disentangled in the learned object models, by manipulating the position latent, we are able to move objects around without changing other factors like object appearance. In Figure 5, we visualize traversals of $\mathbf{z}_n^{\text{where},x}$ and $\mathbf{z}_n^{\text{where},y}$ of the yellow ball through generations from 5 query viewpoints. It can be seen that the change of one coordinate does not affect the other. In addition, the appearance of the yellow ball remains complete and clean during the traversal. Other

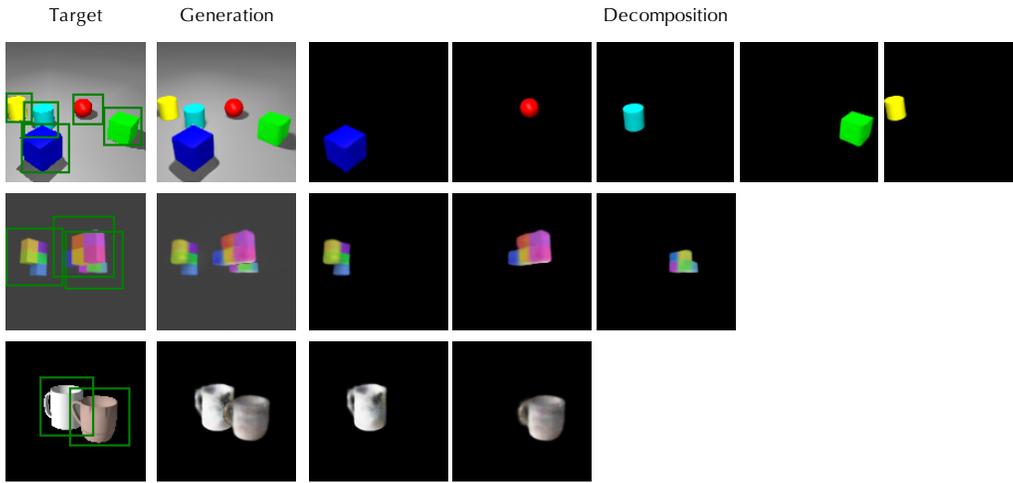


Figure 3: The full scene is composited from individual object rendering results. Predicted bounding boxes are drawn on target images.

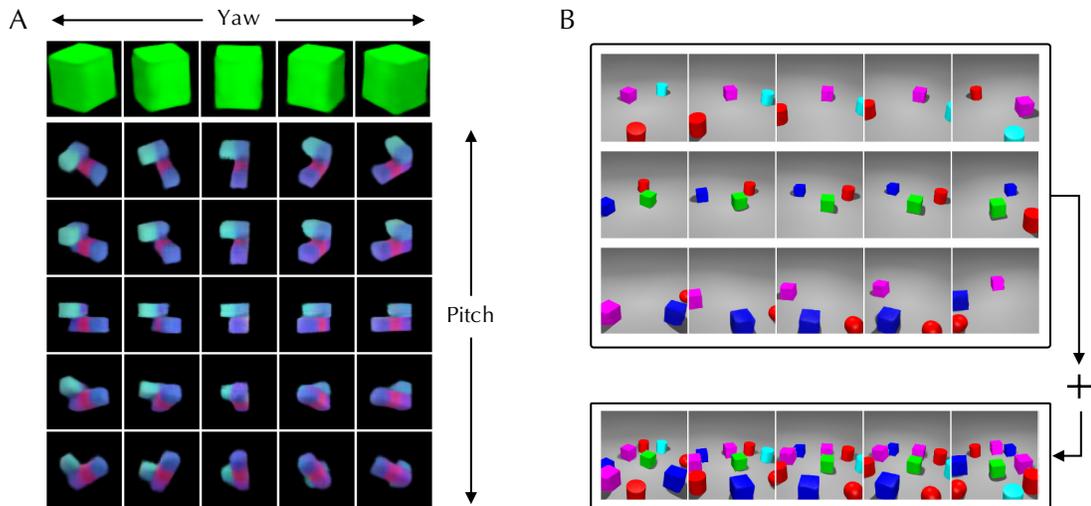


Figure 4: (A) Visualization of learned object models from a set of query viewpoints. (B) Learned object models are reconfigured into a novel scene. Columns correspond to query viewpoints.

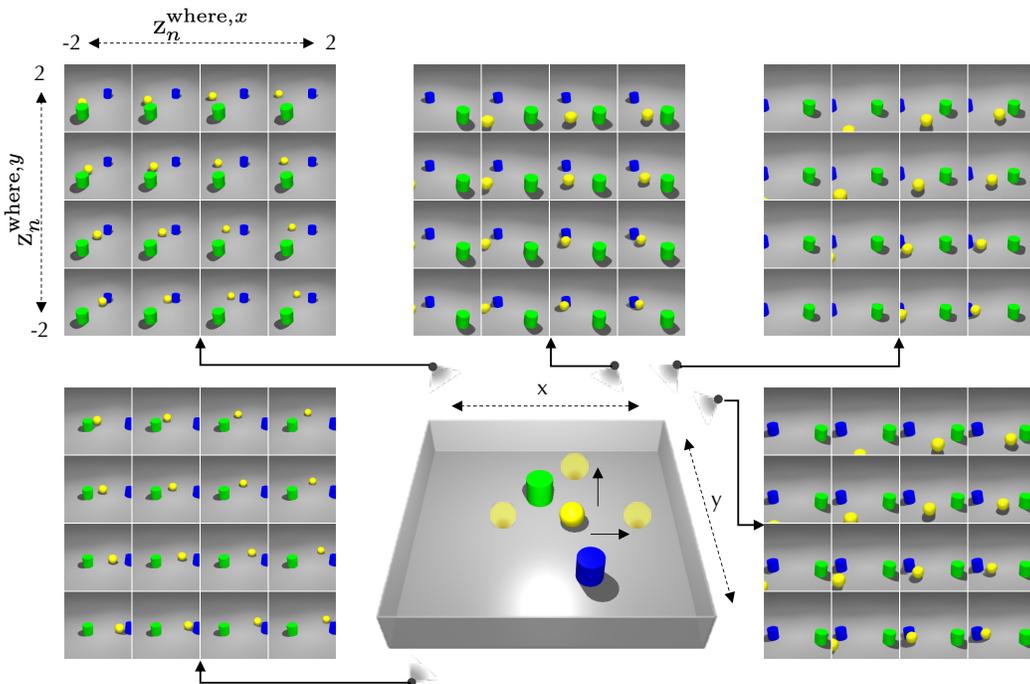


Figure 5: Traversal of the position latent $z_n^{\text{where},x}$ and $z_n^{\text{where},y}$ of the yellow ball in the scene. We show generations from five query viewpoints after manipulating the position latent.

untouched components (the green cylinder, the blue cylinder, and the background) remain unchanged. Moreover, we also notice some desired rendering effects. For example, the size of the yellow ball becomes smaller as it moves further away from the camera.

5.2 Quantitative Evaluation

In this section, we report quantitative results on scene generation and decomposition, which reflect the quality of the learned object models. We also highlight the benefit of learning object models in two downstream tasks.

Scene Generation. To compare the generation quality of ROOTS and GQN, in Table 1 and Table 2, we report negative log-likelihood (NLL) and mean squared error (MSE) on the test sets. We provide 15 context observations for both models. Similar to previous works (Kumar et al., 2018; Babaeizadeh et al., 2018), we estimate the NLL and report the minimum MSE both using 100 samples from the learned conditional prior. This measures the ability of a conditional generative model to capture the true outcome within its conditional prior of all possible outcomes. ROOTS outperforms GQN on both metrics, showing that learning object models also contributes to better generation quality.

Object Models. To evaluate the quality of learned object models, we report object counting accuracy and an adapted version of average precision (AP, Everingham et al. 2010) in Figure 6. AP measures the object localization ability. To compute AP, we set some

Data Set	1-3 Shapes		2-4 Shapes		3-5 Shapes	
Metrics	NLL↓	MSE↓	NLL↓	MSE↓	NLL↓	MSE↓
ROOTS	-207595.81	30.60	-206611.07	42.41	-205608.07	54.45
GQN	-206760.87	40.62	-205604.74	54.49	-204918.39	62.73

Table 1: Quantitative evaluation of scene generation on the Shapes data sets.

Data Set	Multi-Shepard-Metzler		ShapeNet Arrangement	
Metrics	NLL↓	MSE↓	NLL↓	MSE↓
ROOTS	-206627.56	42.22	-192414.85	212.77
GQN	-206294.22	46.22	-185010.31	301.62

Table 2: Quantitative evaluation of scene generation on the Multi-Shepard-Metzler data set and the ShapeNet arrangement data set.

thresholds t_i on the 3D distance between the predicted $\mathbf{z}_n^{\text{where}}$ and the groundtruth object center position. If the distance is within the threshold, the prediction is considered a true positive. Clearly, a smaller threshold requires the model to locate objects more accurately. We set three thresholds: 1/4, 2/4, and 3/4 of the average object size. For each threshold t_i , we obtain the area under the precision-recall curve as $\text{AP}(t_i)$. The final AP is averaged over the three thresholds: $\text{AP} = \sum_{i=1}^3 \text{AP}(t_i)/3$. We vary the number of contexts provided, and compute counting accuracy and AP using the predicted N and $\mathbf{z}_n^{\text{where}}$ that achieve the minimum MSE over 10 samples from the conditional prior. As shown in Figure 6, both counting accuracy and AP increase as the number of context observations becomes larger. This indicates that ROOTS can effectively accumulate information from the given contexts.

Segmentation of 2D Observations. The rendering process of ROOTS implicitly segments 2D observations under query viewpoints. The segmentation performance reflects the quality of learned 3D object appearance. Since GQN cannot provide such segmentation, we compare ROOTS with IODINE (Greff et al., 2019) in terms of the adjusted Rand index (ARI, Rand 1971; Hubert and Arabie 1985) on the Shapes data sets (IODINE completely failed on the MSM data set—it tends to split one object into multiple slots based on color similarity, as we show in Appendix K). We train IODINE on all the images available in the training set, using the official implementation. At test time, ROOTS is given 15 random contexts for each scene and performs segmentation for an unseen query viewpoint. ROOTS does not have access to the target image under the query viewpoint. In contrast, IODINE directly takes the target image as input. Results in Table 3 show that ROOTS outperforms IODINE on both foreground segmentation (ARI-NoBg) and full image segmentation (ARI). We would like to emphasize that IODINE specializes in 2D scene segmentation, whereas ROOTS obtains its 2D segmentation ability as a by-product of learning 3D object models.

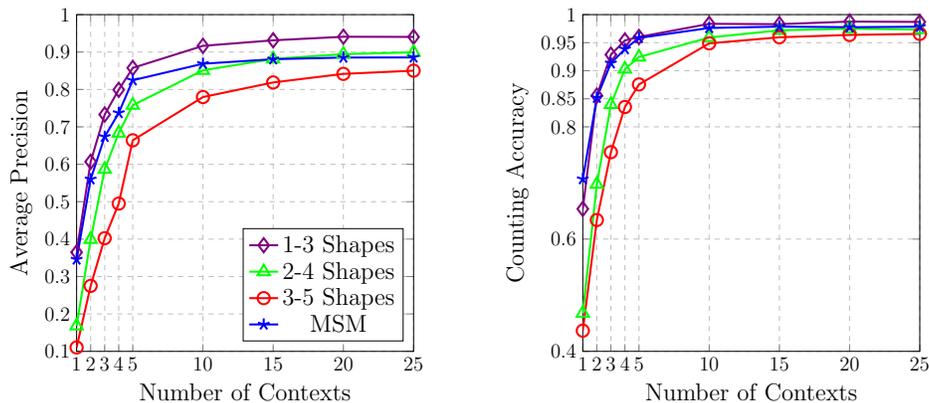


Figure 6: Average precision and counting accuracy.

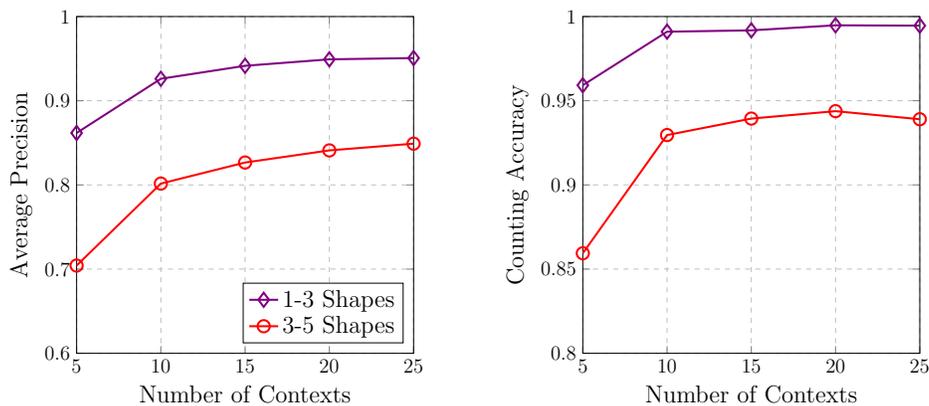


Figure 7: Generalization performance of average precision and counting accuracy. ROOTS is trained on the 2-4 Shapes data set.

Generalization. To evaluate the generalization ability, we first train ROOTS and GQN on the Shapes data set with 2-4 objects, and then test on the Shapes data sets with 1-3 objects and 3-5 objects respectively. As shown in Table 4, ROOTS achieves better NLL and MSE in both interpolation and extrapolation settings. We further report AP and counting accuracy for ROOTS when generalizing to the above two data sets. As shown in Figure 7, ROOTS generalizes well to scenes with 1-3 objects, and performs reasonably when given more context observations on scenes with 3-5 objects.

Downstream 3D Reasoning Tasks. The 3D object models can facilitate object-wise 3D reasoning. We demonstrate this in two downstream tasks on the Shapes data set with 3-5 objects. *Retrieve Object.* The goal of this task is to retrieve the object that lies closest to a given position \mathbf{p} . We consider both 3D and 2D versions of the task. In 3D version, we set \mathbf{p} as the origin of the 3D space, whereas in 2D version, \mathbf{p} is the center point of the target image from viewpoint \mathbf{v}_q . We treat this task as a classification problem, where the input is the learned representation (along with \mathbf{v}_q in 2D version), and the output is the label of

Data Set	1-3 Shapes		2-4 Shapes		3-5 Shapes		Multi-Shepard-Metzler	
Metrics	ARI \uparrow	ARI-NoBg \uparrow	ARI \uparrow	ARI-NoBg \uparrow	ARI \uparrow	ARI-NoBg \uparrow	ARI \uparrow	ARI-NoBg \uparrow
ROOTS	0.9477	0.9942	0.9482	0.9947	0.9490	0.9930	0.9303	0.9608
IODINE	0.8217	0.8685	0.8348	0.9854	0.8422	0.9580	Failed	Failed

Table 3: Quantitative evaluation of 2D segmentation.

Training Set	2-4 Shapes			
Test Set	1-3 Shapes		3-5 Shapes	
Metrics	NLL \downarrow	MSE \downarrow	NLL \downarrow	MSE \downarrow
ROOTS	-208122.58	24.27	-204480.37	67.98
GQN	-207616.49	30.35	-202922.03	86.68

Table 4: Quantitative evaluation of generalization ability.

the desired object. Here, the label is an integer assigned to each object based on its shape and color. We compare ROOTS with the GQN baseline, and report testing accuracies in Table 5. ROOTS outperforms GQN, demonstrating the effectiveness of the learned object models in spatial reasoning. **Find Pair.** In this task, the goal is to find two objects that have the smallest pair-wise distance in 3D space. Again, we treat this as a classification task, where the target label is the sum of labels of the two desired objects. The testing accuracies are reported in Table 5. Clearly, this task requires pair-wise relational reasoning. The object models learned by ROOTS naturally allows extraction of pair-wise relations. In contrast, the scene-level representation of GQN without object-wise factorization leads to incompetence in relational reasoning.

5.3 Ablation Study

Our ablation study shows that the components of ROOTS are necessary for obtaining object models. In particular, we tried the following alternative design choices.

ROOTS Encoder. One may think that $\mathbf{z}_n^{\text{what}}$ can be directly inferred from scene-level contexts without object-attention grouping. Thus, instead of using the object-level GQN encoder, we tried inferring $\mathbf{z}_n^{\text{what}}$ from GVFM along with $\mathbf{z}_n^{\text{where}}$:

$$p(\mathbf{z}_{ijk}^{\text{pres}}, \mathbf{z}_{ijk}^{\text{where}}, \mathbf{z}_{ijk}^{\text{what}} \mid \mathcal{C}) = f_{\text{pres,where,what}}(\mathbf{r}_{ijk}^{\text{neighbor}}).$$

The model, however, failed to decompose scenes into objects and hence was not trainable.

ROOTS Decoder. One may also think that the object-specific image layer $\hat{\mathbf{x}}_{n,q}$ can be directly generated from the 3D object model \mathbf{z}_n^{3D} without having the intermediate 2D representation $\mathbf{o}_{n,q}^{2D}$. To implement this, instead of using the inverse spatial transformer, we directly obtain $\hat{\mathbf{x}}_{n,q}$ and $\hat{\mathbf{m}}_{n,q}$ from the object-level GQN decoder:

$$[\hat{\mathbf{x}}_{n,q}, \hat{\mathbf{m}}_{n,q}] = \text{GQN}_{\text{dec}}(\text{concat}[\mathbf{z}_n^{\text{what}}, \mathbf{r}_n^{\text{att}}], \mathbf{v}_q).$$

This model was also not trainable as it could not use the object positions effectively.

Tasks	Retrieve Object		Find Pair
	3D Version	2D Version	
ROOTS	90.38%	93.71%	84.70%
GQN	81.31%	84.18%	12.48%

Table 5: Testing accuracies on downstream tasks.

6. Conclusion

We proposed ROOTS, a probabilistic generative model for unsupervised learning of 3D object models from partial observations of multi-object 3D scenes. The learned object models capture the complete 3D appearance of individual objects, yielding better generation quality of the full scene. They also improve generalization ability and allow out-of-distribution scenes to be easily generated. Moreover, in downstream 3D reasoning tasks, ROOTS shows superior performance compared to the baseline model.

ROOTS is built upon SPAIR (Crawford and Pineau, 2019) and SPACE (Lin et al., 2020), and similarly has difficulty dealing with complex natural images. However, recent work (Engelcke et al., 2021) in unsupervised object-centric representation learning from 2D scenes is starting to show promising results on complex real-world data sets. Future work can seek to incorporate those advancements into ROOTS. Another interesting future direction is to learn the knowledge of the 3D world in an online manner similarly as we humans keep updating our knowledge of the world.

Acknowledgments

We would like to acknowledge support for this project from Kakao Brain and Center for Super Intelligence (CSI). We would like to thank Jindong Jiang, Skand Vishwanath Peri, and Yi-Fu Wu for helpful discussion, and anonymous reviewers for valuable feedback.

Appendix A. Generation Samples

We provide more generation samples in this section. For each scene in Figure 8, we show 8 sampled context images in the top row, superimposed with predicted bounding boxes. We also show generations from three query viewpoints, together with the decomposed object-wise rendering results. Similar visualizations for two scenes from the 3-5 Shapes data set are provided in Figure 9.

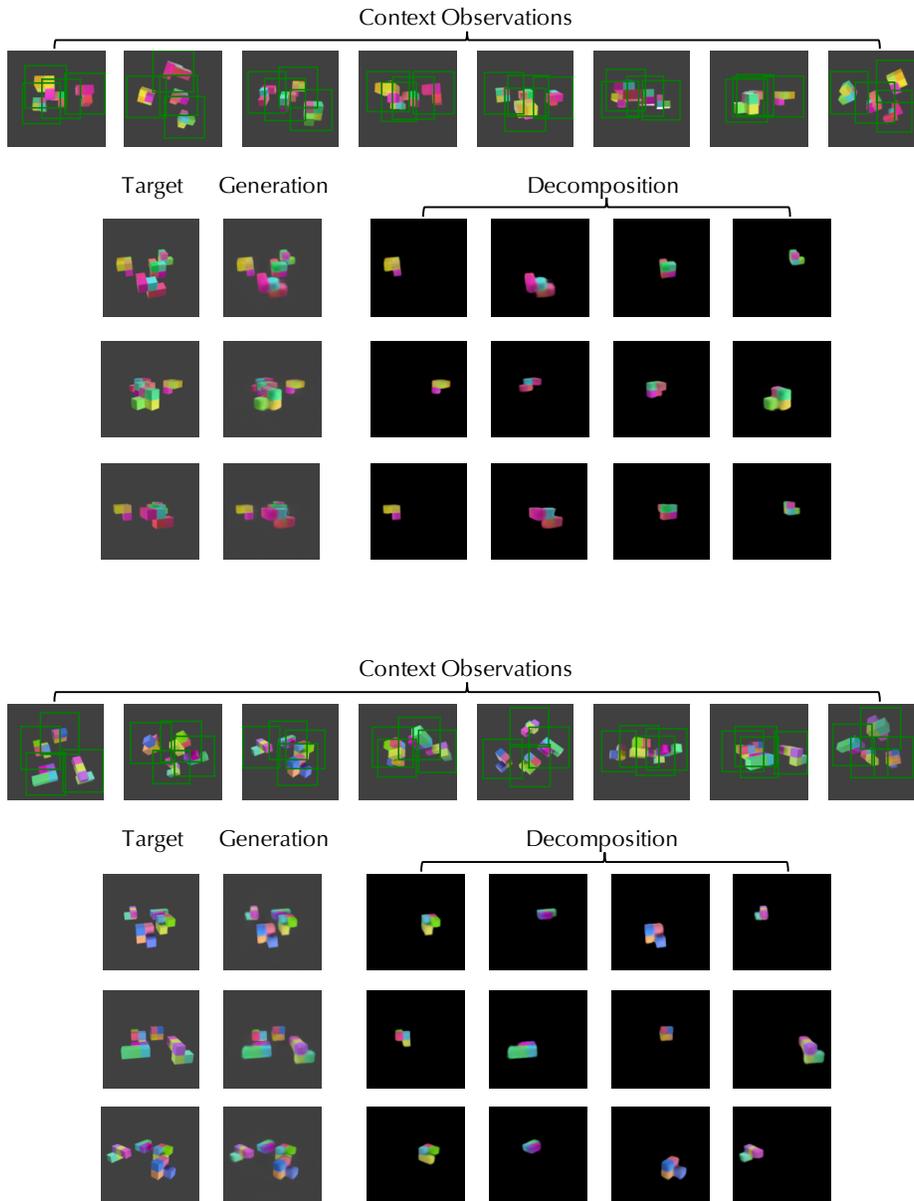


Figure 8: Generation samples from the Multi-Shepard-Metzler data set.

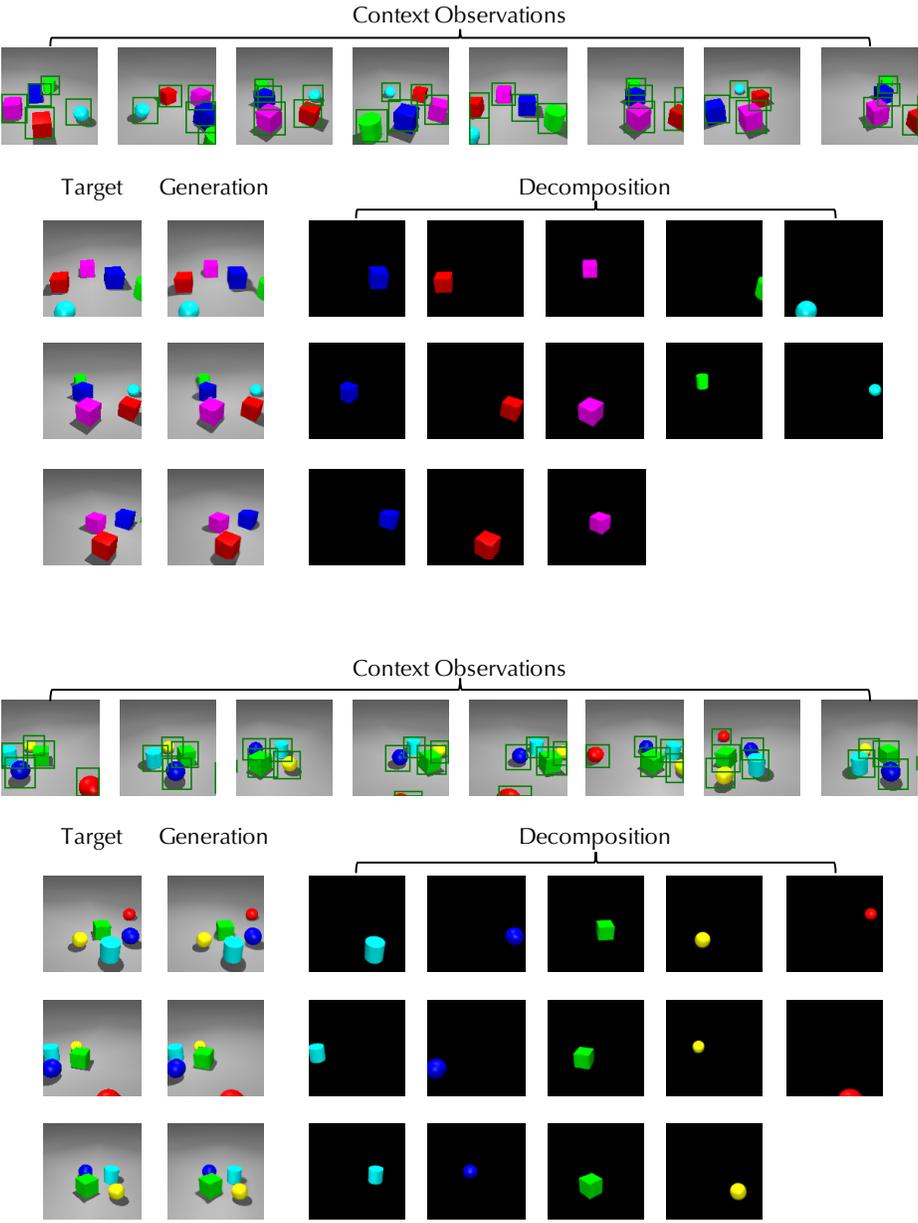


Figure 9: Generation samples from the 3-5 Shapes data set.

Appendix B. Generation with Limited Contexts

In this section, we demonstrate how ROOTS copes with uncertainty when given a limited set of contexts. As shown in Figure 10 left, we incrementally provide ROOTS with one to four contexts. Each time a new context is observed, we recompute the conditional prior and take two samples (corresponding to two plausible explanations) to generate the objects and the scene from four query viewpoints, as shown on the right. The targets for the query viewpoints are shown on the top.

In general, we find that when an object is observed in only one or two contexts, ROOTS can be quite uncertain about its size and color. Because the inference of object size is conditioned on its 3D position, the uncertainty in size probably comes from the uncertainty in depth estimation. The color uncertainty mostly appears in query view 2, which is quite different from the first three context views. As more contexts become available, ROOTS is able to correct the inferred object models.

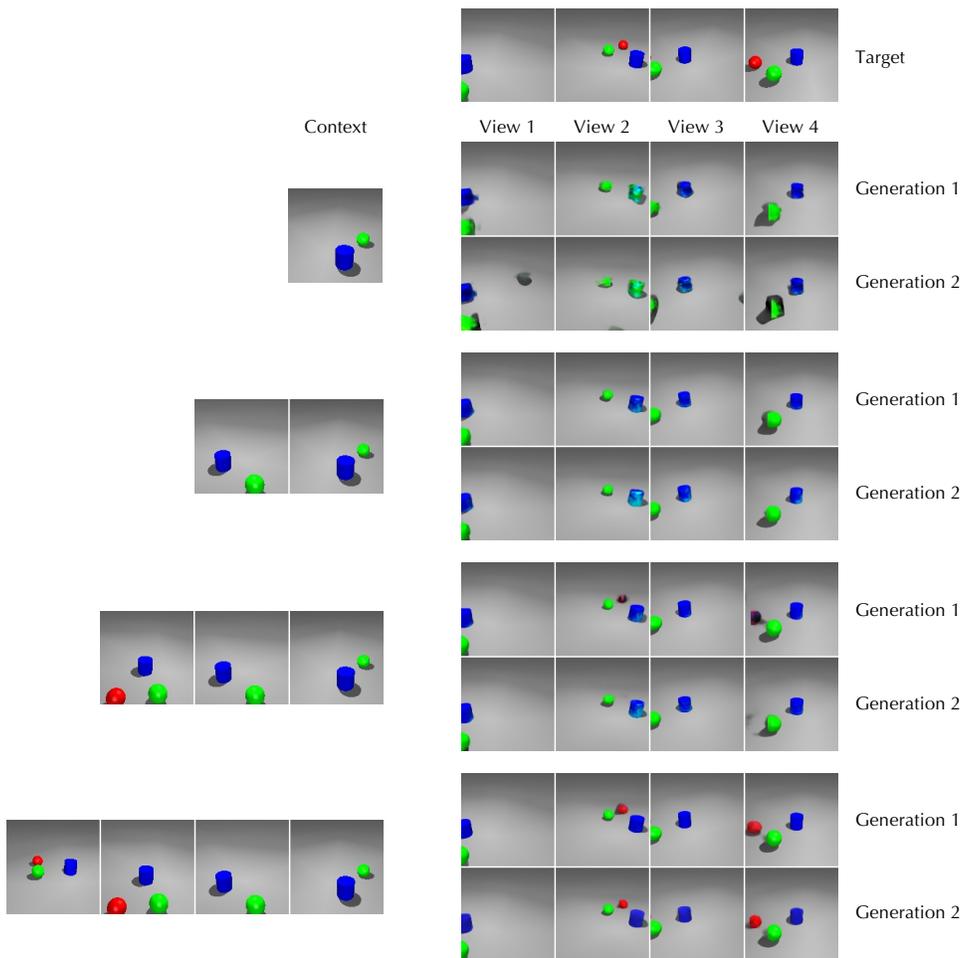


Figure 10: Generation samples from four contexts provided incrementally.

Appendix C. Object Models

In this section, we provide two more samples of the learned object models. As shown in Figure 11, each object model inferred from a multi-object scene can generate complete object appearance given different query viewpoints.

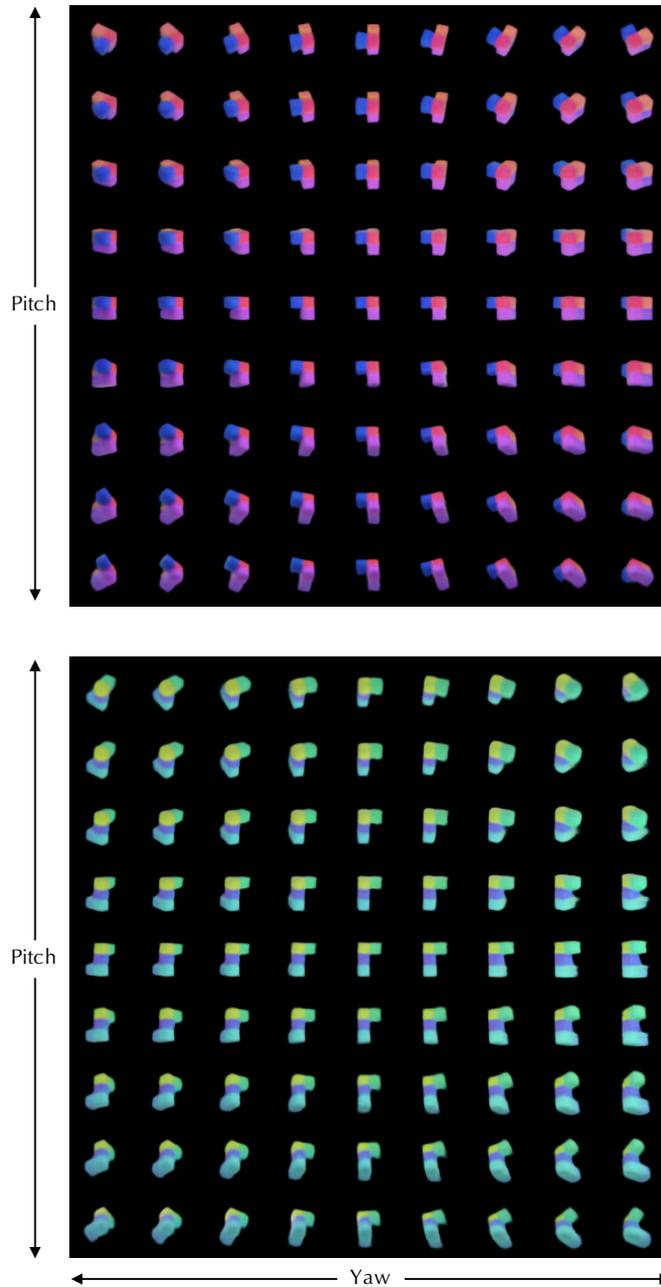


Figure 11: Visualization of learned object models from a set of query viewpoints.

Appendix D. Summary of ROOTS Encoder and Decoder

Algorithm 1 ROOTS Encoder

Input: contexts $\mathcal{C} = \{(\mathbf{x}_c, \mathbf{v}_c)\}_{c=1}^K$, partition resolutions N_x, N_y, N_z

Define: $[J] = \{1, 2, \dots, J\}$ for any natural number J

- 1: Obtain Geometric Volume Feature Map $\mathbf{r} = f_{\text{ctx_enc}}(\mathcal{C})$
 - 2: **for each** $(i, j, k) \in [N_x] \times [N_y] \times [N_z]$ **parallel do**
 - 3: Infer object presence and position in 3D world coordinates: $z_{ijk}^{\text{pres}}, \mathbf{z}_{ijk}^{\text{where}} \sim f_{\text{pres,where}}(\mathbf{r})$
 - 4: **end for**
 - 5: Obtain the number of objects $N = \sum_{ijk} z_{ijk}^{\text{pres}}$
 - 6: Map each (i, j, k) with $z_{ijk}^{\text{pres}} = 1$ to an object index $n \in [N]$
 - 7: **for each** object $n \in [N]$ **parallel do**
 - 8: **for each** context $(\mathbf{x}_c, \mathbf{v}_c) \in \mathcal{C}$ **parallel do**
 - 9: Infer 2D object location $\mathbf{o}_{n,c}^{\text{where}}$ using Attention-by-Perspective-Projection
 - 10: Crop 2D object patch $\mathbf{x}_{n,c}^{\text{att}}$ from \mathbf{x}_c using $\mathbf{o}_{n,c}^{\text{where}}$
 - 11: **end for**
 - 12: Obtain object context $\mathcal{C}_n = \{(\mathbf{x}_{n,c}^{\text{att}}, \mathbf{v}_c, \mathbf{o}_{n,c}^{\text{where}})\}_{c=1}^K$
 - 13: Infer 3D object appearance $\mathbf{z}_n^{\text{what}} \sim \text{GQN}_{\text{enc}}(\mathcal{C}_n)$
 - 14: Build object model $\mathbf{z}_n^{\text{3D}} = (\mathbf{z}_n^{\text{where}}, \mathbf{z}_n^{\text{what}})$
 - 15: **end for**
 - 16: **return** object models $\{\mathbf{z}_n^{\text{3D}}\}_{n=1}^N$, object contexts $\{\mathcal{C}_n\}_{n=1}^N$
-

Algorithm 2 ROOTS Decoder

Input: object models $\{\mathbf{z}_n^{\text{3D}}\}_{n=1}^N$, object contexts $\{\mathcal{C}_n\}_{n=1}^N$, query viewpoints $\mathbf{v}_Q = \{\mathbf{v}_q\}_{q=1}^M$

Define: $[J] = \{1, 2, \dots, J\}$ for any natural number J

- 1: **for each** query viewpoint $\mathbf{v}_q \in \mathbf{v}_Q$ **parallel do**
 - 2: **for each** object $n \in [N]$ **parallel do**
 - 3: Obtain object context encoding $\mathbf{r}_n^{\text{att}} = f_{\text{obj_ctx_enc}}(\mathcal{C}_n)$
 - 4: Decode 3D appearance into 2D image patch:

$$\mathbf{o}_{n,q}^{\text{what}} = \text{GQN}_{\text{dec}}(\text{concat}[\mathbf{z}_n^{\text{what}}, \mathbf{r}_n^{\text{att}}], \mathbf{v}_q)$$
 - 5: Infer 2D object location $\mathbf{o}_{n,q}^{\text{where}}$ using Attention-by-Perspective-Projection
 - 6: Obtain image layer $\hat{\mathbf{x}}_{n,q}$ and transparency map $\alpha_{n,q}$ from $\mathbf{o}_{n,q}^{\text{what}}$ and $\mathbf{o}_{n,q}^{\text{where}}$
 - 7: **end for**
 - 8: Composite the full image $\hat{\mathbf{x}}_q = \sum_{n=1}^N \alpha_{n,q} \odot \hat{\mathbf{x}}_{n,q}$
 - 9: **end for**
 - 10: **return** generations $\{\hat{\mathbf{x}}_q\}_{q=1}^M$
-

Appendix E. Perspective Projection

Following GQN (Eslami et al., 2018), we parameterize the viewpoint \mathbf{v} as a tuple (\mathbf{w}, y, p) , where $\mathbf{w} \in \mathbb{R}^3$ is the position of the camera in world coordinates, and $y \in \mathbb{R}$ and $p \in \mathbb{R}$ are its yaw and pitch respectively. We also assume access to the intrinsic camera parameters, including focal length $f \in \mathbb{R}$ and sensor size, that are the same across all scenes. APP_{pos} converts the center position of an object n from world coordinates $\mathbf{z}_n^{\text{where}} \in \mathbb{R}^3$ to image coordinates $[\mathbf{o}_n^{\text{center}}, o_n^{\text{depth}}]^\top \in \mathbb{R}^3$ as follows:

$$[a, b, d]^\top = \mathbf{R}_{y,p}(\mathbf{z}_n^{\text{where}} - \mathbf{w}) , \quad \mathbf{o}_n^{\text{center}} = \text{normalize}([fa/d, fb/d]^\top) , \quad o_n^{\text{depth}} = d .$$

Here, $\mathbf{R}_{y,p}$ is a 3×3 rotation matrix computed from the camera yaw and pitch, $[a, b, d]^\top$ represents the center position of the object in camera coordinates, and $[a, b]^\top$ is further normalized into image coordinates $\mathbf{o}_n^{\text{center}}$, using the focal length and sensor size, so that the upper-left corner of the image corresponds to $[-1, -1]^\top$ and the lower-right corner corresponds to $[1, 1]^\top$.

Appendix F. Transparency Map

The transparency map $\alpha_{n,q}$ ensures that occlusion among objects is properly handled for a query viewpoint \mathbf{v}_q . Ideally, $\alpha_{n,q}(i, j) = 1$ if, when viewed from \mathbf{v}_q , the pixel (i, j) is contained in object n and is not occluded by any other object, and $\alpha_{n,q}(i, j) = 0$ otherwise. On the other hand, the object mask $\hat{\mathbf{m}}_{n,q}$ is expected to capture the non-occluded full object, that is, $\hat{\mathbf{m}}_{n,q}(i, j) = 1$ if the pixel (i, j) is contained in object n when viewed from \mathbf{v}_q , regardless of whether it is occluded or not. Therefore, we compute $\alpha_{n,q}$ by masking out occluded pixels from $\hat{\mathbf{m}}_{n,q}$:

$$\alpha_{n,q} = \mathbf{w}_{n,q} \odot \hat{\mathbf{m}}_{n,q} ,$$

where \odot is pixel-wise multiplication, and $\mathbf{w}_{n,q}(i, j) = 1$ if object n is the closest one to the camera among all objects that contain the pixel (i, j) . In actual implementation, $\alpha_{n,q}$, $\hat{\mathbf{m}}_{n,q}$, and $\mathbf{w}_{n,q}$ are not strictly binary, and we obtain the value of $\mathbf{w}_{n,q}$ at each pixel (i, j) by the masked softmax over negative depth values:

$$\mathbf{w}_{n,q}(i, j) = \frac{\hat{\mathbf{m}}_{n,q}(i, j) \exp(-o_{n,q}^{\text{depth}})}{\sum_{n=1}^N \hat{\mathbf{m}}_{n,q}(i, j) \exp(-o_{n,q}^{\text{depth}})} .$$

Appendix G. Data Set Details

In this section, we provide details of the two data sets we created.

Shapes. There are 3 types of objects: cube, sphere, and cylinder, with 6 possible colors to choose from. Object sizes are sampled uniformly between $[0.56, 0.66]$ units in the MuJoCo (Todorov et al., 2012) physics world. All objects are placed on the $z = 0$ plane, with a range of $[-2, 2]$ along both x -axis and y -axis. We randomly sample 30 cameras for each scene. They are placed at a distance of 3 from the origin, but do not necessarily point to the origin. The camera pitch is sampled between $[\pi/7, \pi/6]$ so that the camera is always above the $z = 0$ plane. The camera yaw is sampled between $[-\pi, \pi]$.

Multi-Shepard-Metzler. We generate the Shepard-Metzler objects as described in GQN (Eslami et al., 2018). Each object consists of 5 cubes with edge length 0.8. Each cube

is randomly colored, with hue between $[0, 1]$, saturation between $[0.75, 1]$, and value equal to 1. Like the Shapes data set, all objects are placed on the $z = 0$ plane, with a range of $[-3, 3]$ along both x -axis and y -axis. We randomly sample 30 cameras for each scene and place them at a distance of 12 from the origin. They all point to the origin. The camera pitch is sampled between $[-\frac{5}{12}\pi, \frac{5}{12}\pi]$, and the yaw is sampled between $[-\pi, \pi]$.

Appendix H. ROOTS Implementation Details

In this section, we introduce the key building blocks for implementing ROOTS.

Context Encoder. The context encoder is modified based on the ‘tower’ representation architecture in GQN (Eslami et al., 2018). It encodes each pair of context image and the corresponding viewpoint into a vector. Summation is applied over the context encodings to obtain the order-invariant representation ψ .

Object-Level Context Encoder. The object-level context encoder is also an adaptation of the ‘tower’ representation architecture, but takes the extracted object-level context \mathcal{C}_n as input.

ConvDRAW. We use ConvDRAW (Gregor et al., 2016) to infer the prior and posterior distributions of the latent variables. To render the objects and the background, we use a deterministic version of ConvDRAW (i.e., without sampling). In the following, we describe one rollout step (denoted l) of ConvDRAW used in ROOTS generative and inference processes, respectively. We provide detailed configurations of each ConvDRAW module in Table 6.

- Generative Process:

$$\begin{aligned} (\mathbf{h}_p^{(l+1)}, \mathbf{c}_p^{(l+1)}) &\leftarrow \text{ConvLSTM}_\theta(\psi_C, \mathbf{z}^{(l)}, \mathbf{h}_p^{(l)}, \mathbf{c}_p^{(l)}) \\ \mathbf{z}^{(l+1)} &\sim \text{StatisNet}_\theta(\mathbf{h}_p^{(l+1)}) \end{aligned}$$

- Inference Process:

$$\begin{aligned} (\mathbf{h}_p^{(l+1)}, \mathbf{c}_p^{(l+1)}) &\leftarrow \text{ConvLSTM}_\theta(\psi_C, \mathbf{z}^{(l)}, \mathbf{h}_p^{(l)}, \mathbf{c}_p^{(l)}) \\ (\mathbf{h}_q^{(l+1)}, \mathbf{c}_q^{(l+1)}) &\leftarrow \text{ConvLSTM}_\phi(\psi_C, \psi_Q, \mathbf{h}_p^{(l)}, \mathbf{h}_q^{(l)}, \mathbf{c}_q^{(l)}) \\ \mathbf{z}^{(l+1)} &\sim \text{StatisNet}_\theta(\mathbf{h}_q^{(l+1)}) \end{aligned}$$

Here, ψ_C and ψ_Q are order-invariant encodings of contexts and queries respectively, and $\mathbf{z}^{(l+1)}$ is the sampled latent at the $(l+1)$ -th step. The prior module is denoted by subscript p , with learnable parameters θ , and the posterior module is denoted by subscript q , with learnable parameters ϕ . StatisNet maps hidden states to distribution parameters, and will be explained in the following. ConvLSTM replaces the fully-connected layers in LSTM (Hochreiter and Schmidhuber, 1997) by convolutional layers.

Sufficient Statistics Network. The sufficient statistics network (StatisNet) outputs sufficient statistics for pre-defined distributions, e.g., μ and σ for Gaussian distributions, and ρ for Bernoulli distributions. We list the configuration of all sufficient statistics networks in Table 7. For $\mathbf{z}_n^{\text{where}}$, $\mathbf{z}_n^{\text{what}}$, and \mathbf{z}^{bg} , we use ConvDraw to learn the sufficient statistics. For $\mathbf{z}_n^{\text{pres}}$, two ConvBlocks are first used to extract features, and then a third ConvBlock

Module Name	Rollout Steps	Hidden Size
\mathbf{z}^{bg}	2	128
$\mathbf{z}^{\text{where}}$	2	128
\mathbf{z}^{what}	4	128
Object Renderer	4	128
Background Renderer	2	128

Table 6: Configuration of ConvDRAW modules.

Object-Level Latent Variables		
$\mathbf{z}^{\text{where}}$	\mathbf{z}^{what}	
Conv3D(128, 3, 1, GN, CELU)	Conv3D(128, 1, 1, GN, CELU)	
Conv3D(64, 1, 1, GN, CELU)	Conv3D(64, 1, 1, GN, CELU)	
Conv3D(32, 1, 1, GN, CELU)	Conv3D(32, 1, 1, GN, CELU)	
Conv3D(3×2 , 1, 1, GN, CELU)	Conv3D(4×2 , 1, 1, GN, CELU)	
\mathbf{z}^{pres}		
ConvBlock 1	ConvBlock 2	ConvBlock 3
Conv3D(256, 3, 1, GN, CELU)	Conv3D(256, 1, 1, GN, CELU)	Conv3D(128, 3, 1, GN, CELU)
	Conv3D(256, 1, 1, GN, CELU)	Conv3D(64, 1, 1, GN, CELU)
		Conv3D(1, 1, 1, GN, CELU)
Scene-Level Latent Variables		
\mathbf{z}^{bg}		
Conv3D(1×2 , 1, 1, GN, CELU)		

Table 7: Configuration of sufficient statistics networks.

combines the features and outputs the parameter of the Bernoulli distribution. GN denotes group normalization (Wu and He, 2018), and CELU denotes continuously differentiable exponential linear units (Barron, 2017).

Training Details. ROOTS is trained for 200 epochs with a batch size of 12, using RMSprop (Hinton et al., 2012) with learning rates chosen from $\{1 \times 10^{-3}, 3 \times 10^{-4}, 1 \times 10^{-4}, 3 \times 10^{-5}\}$. We find the learning rate of 3×10^{-5} works best, and larger learning rates tend to cause inaccurate position predictions.

Appendix I. Downstream Task Details

We use the 3-5 Shapes data set for the downstream tasks. To generate the groundtruth labels, we assign a label to each object based on its type and color. There are 3 different types and 6 different colors in total, thus the label value for one object lies in the range from 0 to 17. We split the data set into training set, validation set, and test set of size 50K, 5K, and 5K, respectively. During training, 10 to 20 randomly sampled context observations

Scene Representation Encoder		GQN-GVFM
GQN	ROOTS	
Conv(256, 2, 2, GN, CELU)	Conv(256, 1, 1, GN, CELU)	Conv(256, 5, 3)
Conv(256, 2, 2, GN, CELU)	Conv(256, 3, 1, GN, CELU)	Conv(128, 3, 2)
Conv(256, 4, 4, GN, CELU)	Conv(256, 3, 1, GN, CELU)	Conv(256, 3, 1)
		ConvTrans3D(256, [3,1,3], 1)
Graph Network		Classifier Network
Edge Encoder	Graph Encoder	
MLP(256, 256, 256+1)	MLP(256, 256)	MLP(128, 64, N_{cls})

Table 8: Configuration of downstream task networks.

are provided for both GQN and ROOTS to learn representations of a 3D scene. All latent variables are sampled from the learned priors.

Retrieve Object. To predict the correct class of the object that lies closest to a given point in the 3D space, the classifier first encodes the scene representation $\hat{\mathbf{r}}$ into a vector, and then uses MLP to predict class probabilities. For GQN, we concatenate the scene embedding and scene latent representation together, that is, $\hat{\mathbf{r}} = [\mathbf{r}_{\mathcal{C}}, \mathbf{z}]$. For ROOTS, we use the object-level representations, that is, $\hat{\mathbf{r}} = \{\mathbf{r}_n^{\text{att}}, \mathbf{z}_n^{\text{3D}}\}_{n=1}^N$, where $\mathbf{r}_n^{\text{att}}$ is an order-invariant encoding of object-level context \mathcal{C}_n . For 2D version of this task, we provide \mathbf{v}_q as an additional input to the classifier. The structure of the scene representation encoder is specified in Table 8. We use three linear layers for the classifier network, as listed in Table 8. For both versions of this task, the number of classes $N_{\text{cls}} = 18$.

Find Pair. The object-level representation provided by ROOTS naturally allows us to extract pair-wise relationships by using a graph net (Battaglia et al., 2018; Veličković et al., 2018). Specifically, we use the object-level representation as the node feature. We then extract edge features for each pair of objects using shared MLPs (Edge Encoder). The edge features are pooled into a single vector using attention, and fed into another MLP (Graph Encoder) to produce the final classification result. For GQN, we encode the scene representation provided by GQN into a GVFM, so that object-specific features can be split into individual cells. We then treat each cell as a node feature, and apply a similar graph net for classification. The Edge Encoder and Graph Encoder of the graph net is specified in Table 8. We use the same classifier structure as the Retrieve Object task. For this task, the number of classes $N_{\text{cls}} = 35$.

Appendix J. CGQN Baseline

We use the ‘tower’ representation architecture in GQN (Eslami et al., 2018) to encode the context \mathcal{C} into a scene-level representation $\mathbf{r}_{\mathcal{C}}$ of size $16 \times 16 \times 256$. We then follow the CGQN paper (Kumar et al., 2018) and use ConvDraw (Gregor et al., 2016) to sample \mathbf{z} from $\mathbf{r}_{\mathcal{C}}$ and decode the target image \mathbf{x}_q from \mathbf{z} and query viewpoint \mathbf{v}_q . The main hyperparameters are ConvDraw steps (denoted L) and the number of \mathbf{z} channels (denoted c) that are sampled at

Data Set	1-3 Shapes	2-4 Shapes	3-5 Shapes
σ	0.3	0.3	0.25
K	4	7	8

Table 9: Hyperparameters for IODINE.

each ConvDraw step. Generally, larger values of L and c give stronger model capacity, but take more computational resource. We choose $L = 12$ and $c = 4$ for the Shapes data set and the ShapeNet arrangement data set, and $L = 16$ and $c = 8$ for the Multi-Shepard-Metzler data set.

We trained the CGQN baseline using Adam (Kingma and Ba, 2015) with learning rates chosen from $\{1 \times 10^{-3}, 3 \times 10^{-4}, 1 \times 10^{-4}, 3 \times 10^{-5}\}$. We found the learning rate of 3×10^{-4} worked best for the Shapes data set and the Multi-Shepard-Metzler data set, while for the ShapeNet arrangement data set, the learning rate of 1×10^{-4} worked best. For fair comparison of NLL, we used a fixed pixel-variance σ^2 of 0.09 during training for both ROOTS and CGQN. To be consistent with the pixel-variance annealing strategy used in CGQN, we multiplied the KL divergence by a value β that is linearly annealed from 20 to 5 at the start of training.

Appendix K. IODINE Baseline

We use the implementation from DeepMind for IODINE (Greff et al., 2019). We adjusted the output standard deviation σ and the slot number K for best performance. We tried values for σ in the range of $[0.1, 0.3]$, and found that the model tends to be unstable with smaller σ values. We varied K from 4 to 8. The final values used for the Shapes data set are listed in Table 9. For the Multi-Shepard-Metzler data set, we found that IODINE tends to segment objects based on color similarity, thereby splitting a single object into multiple slots. We show two examples in Figure 12 and Figure 13, where the spatial broadcast decoder (Watters et al., 2019) and deconvolution-based decoder are used respectively. We use $K = 5$ and $\sigma = 0.25$ in both settings.

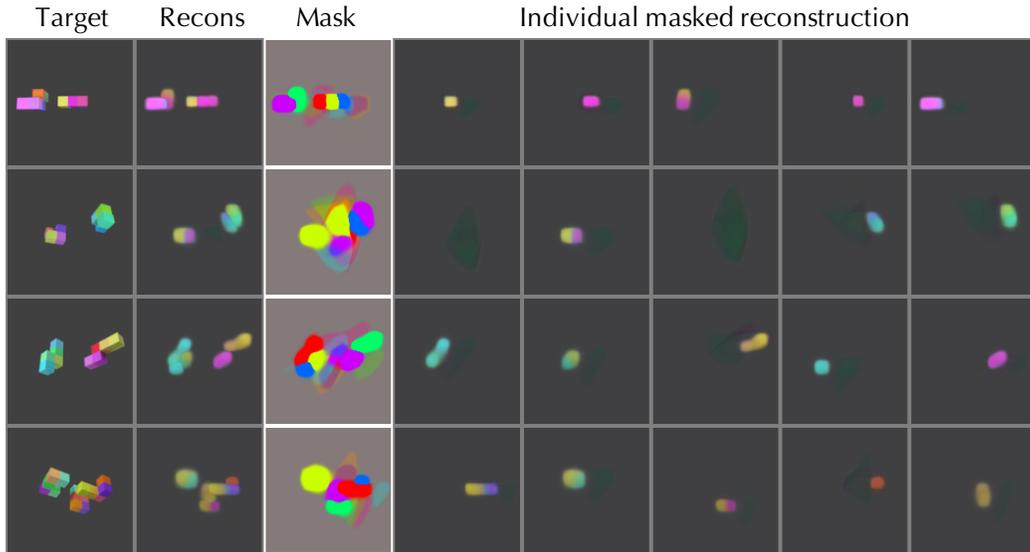


Figure 12: Scene decomposition results of IODINE with spatial broadcast decoder on the Multi-Shepard-Metzler data set.

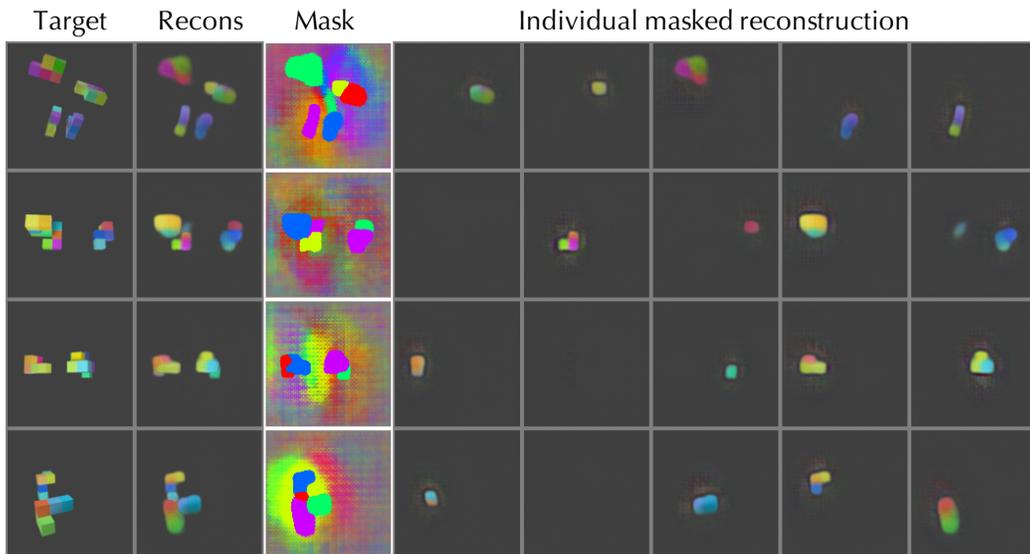


Figure 13: Scene decomposition results of IODINE with deconvolution-based decoder on the Multi-Shepard-Metzler data set.

References

- Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3D point clouds. In *International Conference on Machine Learning*, 2018.
- Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, Roy H. Campbell, and Sergey Levine. Stochastic variational video prediction. In *International Conference on Learning Representations*, 2018.
- Dzmitry Bahdanau, Shikhar Murty, Michael Noukhovitch, Thien Huu Nguyen, Harm de Vries, and Aaron Courville. Systematic generalization: What is required and can it be learned? In *International Conference on Learning Representations*, 2019.
- Jonathan T. Barron. Continuously differentiable exponential linear units. *arXiv preprint arXiv:1704.07483*, 2017.
- Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Caglar Gulcehre, Francis Song, Andrew Ballard, Justin Gilmer, George Dahl, Ashish Vaswani, Kelsey Allen, Charles Nash, Victoria Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matt Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Blender Online Community. *Blender—a 3D modelling and rendering package*. Blender Foundation, Blender Institute, Amsterdam, 2017.
- Léon Bottou. From machine learning to machine reasoning. *Machine Learning*, 94(2): 133–149, 2014.
- Christopher P. Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. MONet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019.
- Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An information-rich 3D model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- Ricson Cheng, Ziyang Wang, and Katerina Fragkiadaki. Geometry-aware recurrent neural networks for active visual recognition. In *Advances in Neural Information Processing Systems*, pages 5081–5091, 2018.
- Christopher B. Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3D-R2N2: A unified approach for single and multi-view 3D object reconstruction. In *Proceedings of the European Conference on Computer Vision*, 2016.

- Eric Crawford and Joelle Pineau. Spatially invariant unsupervised object detection with convolutional neural networks. In *Thirty-Third AAAI Conference on Artificial Intelligence*, 2019.
- Eric Crawford and Joelle Pineau. Learning 3D object-oriented world models from unlabeled videos. *Workshop on Object-Oriented Learning at ICML*, 2020.
- Yilun Du, Zhijian Liu, Hector Basevi, Ales Leonardis, Bill Freeman, Josh Tenenbaum, and Jiajun Wu. Learning to exploit stability for 3D scene parsing. In *Advances in Neural Information Processing Systems*, pages 1726–1736, 2018.
- Emilien Dupont, Miguel Angel Bautista, Alex Colburn, Aditya Sankar, Carlos Guestrin, Josh Susskind, and Qi Shan. Equivariant neural rendering. In *International Conference on Machine Learning*, 2020.
- Sebastien Ehrhardt, Oliver Groth, Aron Monszpart, Martin Engelcke, Ingmar Posner, Niloy Mitra, and Andrea Vedaldi. RELATE: Physically plausible multi-object scene synthesis using structured latent spaces. *arXiv preprint arXiv:2007.01272*, 2020.
- Martin Engelcke, Adam R. Kosior, Oivi Parker Jones, and Ingmar Posner. GENESIS: Generative scene inference and sampling of object-centric latent representations. In *International Conference on Learning Representations*, 2020.
- Martin Engelcke, Oivi Parker Jones, and Ingmar Posner. Genesis-v2: Inferring unordered object representations without iterative refinement. *arXiv preprint arXiv:2104.09958*, 2021.
- S.M. Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, and Geoffrey E. Hinton. Attend, Infer, Repeat: Fast scene understanding with generative models. In *Advances in Neural Information Processing Systems*, pages 3225–3233, 2016.
- S.M. Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S. Morcos, Marta Garnelo, Avraham Ruderman, Andrei A. Rusu, Ivo Danihelka, Karol Gregor, David P. Reichert, Lars Buesing, Theophane Weber, Oriol Vinyals, Dan Rosenbaum, Neil Rabinowitz, Helen King, Chloe Hillier, Matt Botvinick, Daan Wierstra, Koray Kavukcuoglu, and Demis Hassabis. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- Mark Everingham, Luc Van Gool, Christopher K.I. Williams, John Winn, and Andrew Zisserman. The PASCAL Visual Object Classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010.
- Ian J. 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*, page 2672–2680, 2014.
- Klaus Greff, Sjoerd van Steenkiste, and Jürgen Schmidhuber. Neural expectation maximization. In *Advances in Neural Information Processing Systems*, pages 6691–6701, 2017.

- Klaus Greff, Raphaël Lopez Kaufmann, Rishab Kabra, Nick Watters, Chris Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference. In *International Conference on Machine Learning*, 2019.
- Karol Gregor, Frederic Besse, Danilo Jimenez Rezende, Ivo Danihelka, and Daan Wierstra. Towards conceptual compression. In *Advances In Neural Information Processing Systems*, pages 3549–3557, 2016.
- Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003.
- Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew M. Botvinick, Shakir Mohamed, and Alexander Lerchner. β -VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.
- Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. Neural networks for machine learning - Lecture 6a - Overview of mini-batch gradient descent. 2012.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Bruce Hood and Laurie Santos. *The Origins of Object Knowledge*. Oxford University Press, 2009.
- Øyvind Arne Høydal, Emilie Ranheim Skytøen, Sebastian Ola Andersson, May-Britt Moser, and Edvard I. Moser. Object-vector coding in the medial entorhinal cortex. *Nature*, 568(7752):400–404, 2019.
- Siyuan Huang, Siyuan Qi, Yinxue Xiao, Yixin Zhu, Ying Nian Wu, and Song-Chun Zhu. Cooperative holistic scene understanding: Unifying 3D object, layout, and camera pose estimation. In *Advances in Neural Information Processing Systems*, pages 207–218, 2018.
- Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of Classification*, 2(1):193–218, 1985.
- Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In *Advances in Neural Information Processing Systems*, pages 2017–2025, 2015.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparametrization with Gumbel-Softmax. In *International Conference on Learning Representations*, 2017.
- Daniel Kahneman, Anne Treisman, and Brian J. Gibbs. The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24(2):175–219, 1992.
- Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proceedings of the European Conference on Computer Vision*, 2018.

- Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *Advances in Neural Information Processing Systems*, pages 365–376, 2017.
- Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3D mesh renderer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.
- Ananya Kumar, S.M. Ali Eslami, Danilo Jimenez Rezende, Marta Garnelo, Fabio Viola, Edward Lockhart, and Murray Shanahan. Consistent generative query networks. *arXiv preprint arXiv:1807.02033*, 2018.
- Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017.
- Yiyi Liao, Katja Schwarz, Lars Mescheder, and Andreas Geiger. Towards unsupervised learning of generative models for 3D controllable image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- Zhixuan Lin, Yi-Fu Wu, Skand Vishwanath Peri, Weihao Sun, Gautam Singh, Fei Deng, Jindong Jiang, and Sungjin Ahn. SPACE: Unsupervised object-oriented scene representation via spatial attention and decomposition. In *International Conference on Learning Representations*, 2020.
- Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. In *Advances in Neural Information Processing Systems*, pages 11525–11538, 2020.
- Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The Concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*, 2017.
- Alex Martin. The representation of object concepts in the brain. *Annual Review of Psychology*, 58:25–45, 2007.
- Daniel Maturana and Sebastian Scherer. VoxNet: A 3D convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 922–928. IEEE, 2015.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Proceedings of the European Conference on Computer Vision*, 2020.

- Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. HoloGAN: Unsupervised learning of 3D representations from natural images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- Thu Nguyen-Phuoc, Christian Richardt, Long Mai, Yongliang Yang, and Niloy Mitra. BlockGAN: Learning 3D object-aware scene representations from unlabelled images. In *Advances in Neural Information Processing Systems*, pages 6767–6778, 2020.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017.
- Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- William M. Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66(336):846–850, 1971.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, 2014.
- Edmund T. Rolls, Jianzhong Xiang, and Leonardo Franco. Object, space, and object-space representations in the primate hippocampus. *Journal of Neurophysiology*, 94(1):833–844, 2005.
- Bernhard Schölkopf. Causality for machine learning. *arXiv preprint arXiv:1911.10500*, 2019.
- Daeyun Shin, Zhile Ren, Erik B. Sudderth, and Charless C. Fowlkes. 3D scene reconstruction with multi-layer depth and epipolar transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- Gautam Singh, Jaesik Yoon, Youngsung Son, and Sungjin Ahn. Sequential neural processes. In *Advances in Neural Information Processing Systems*, pages 10254–10264, 2019.
- Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. DeepVoxels: Learning persistent 3D feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019a.
- Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3D-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, pages 1119–1130, 2019b.
- Joshua Tobin, Wojciech Zaremba, and Pieter Abbeel. Geometry-aware neural rendering. In *Advances in Neural Information Processing Systems*, pages 11555–11565, 2019.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.

- Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Shubham Tulsiani, Saurabh Gupta, David Fouhey, Alexei A. Efros, and Jitendra Malik. Factoring shape, pose, and layout from the 2D image of a 3D scene. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- Hsiao-Yu Fish Tung, Ricson Cheng, and Katerina Fragkiadaki. Learning spatial common sense with geometry-aware recurrent networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- Sjoerd van Steenkiste, Klaus Greff, and Jürgen Schmidhuber. A perspective on objects and systematic generalization in model-based RL. *Workshop on Generative Modeling and Model-Based Reasoning for Robotics and AI at ICML*, 2019.
- Sjoerd van Steenkiste, Karol Kurach, Jürgen Schmidhuber, and Sylvain Gelly. Investigating object compositionality in generative adversarial networks. *Neural Networks*, 130:309–325, 2020.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- Claes von Hofsten and Elizabeth S. Spelke. Object perception and object-directed reaching in infancy. *Journal of Experimental Psychology: General*, 114(2):198, 1985.
- Nicholas Watters, Loic Matthey, Christopher P. Burgess, and Alexander Lerchner. Spatial Broadcast decoder: A simple architecture for learning disentangled representations in VAEs. *arXiv preprint arXiv:1901.07017*, 2019.
- Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In *Advances in Neural Information Processing Systems*, pages 82–90, 2016.
- Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision*, 2018.
- Xinchen Yan, Jimei Yang, Ersin Yumer, Yijie Guo, and Honglak Lee. Perspective transformer nets: Learning single-view 3D object reconstruction without 3D supervision. In *Advances in Neural Information Processing Systems*, pages 1696–1704, 2016.
- Jaesik Yoon, Gautam Singh, and Sungjin Ahn. Robustifying sequential neural processes. In *International Conference on Machine Learning*, 2020.