

Learning Object-Centric Representations of Multi-Object Scenes from Multiple Views

Li Nanbo
School of Informatics
University of Edinburgh
nanbo.li@ed.ac.uk

Cian Eastwood
School of Informatics
University of Edinburgh
c.eastwood@ed.ac.uk

Robert B Fisher
School of Informatics
University of Edinburgh
rbf@inf.ed.ac.uk

Abstract

Learning object-centric representations of multi-object scenes is a promising approach towards machine intelligence, facilitating high-level reasoning and control from visual sensory data. However, current approaches for *unsupervised object-centric scene representation* are incapable of aggregating information from multiple observations of a scene. As a result, these “single-view” methods form their representations of a 3D scene based only on a single 2D observation (view). Naturally, this leads to several inaccuracies, with these methods falling victim to single-view spatial ambiguities. To address this, we propose *The Multi-View and Multi-Object Network (MulMON)*—a method for learning accurate, object-centric representations of multi-object scenes by leveraging multiple views. In order to sidestep the main technical difficulty of the *multi-object-multi-view* scenario—maintaining object correspondences across views—MulMON iteratively updates the latent object representations for a scene over multiple views. To ensure that these iterative updates do indeed aggregate spatial information to form a complete 3D scene understanding, MulMON is asked to predict the appearance of the scene from novel viewpoints during training. Through experiments we show that MulMON better-resolves spatial ambiguities than single-view methods—learning more accurate and disentangled object representations—and also achieves new functionality in predicting object segmentations for novel viewpoints. Our implementation and pretrained models are given on GitHub ¹.

1 Introduction

Traditional VAEs[15] use “single-object” or “flat” vector representations that fail to capture the compositional structure of natural scenes, i.e. the existence of interchangeable objects with common properties. As a result, “multi-object” or object-centric representations have emerged as a promising approach to scene understanding, improving sample-efficiency and generalization for many downstream applications like relational reasoning and control [5, 20, 3]. However, recent progress in unsupervised object-centric scene representation has been limited to “single-view” methods which form their representations of 3D scenes based only on a single 2D observation (view). As a result, these methods form inaccurate representations that fall victim to single-view spatial ambiguities (e.g. occluded or partially occluded objects) and fail to capture 3D spatial structures.

¹<https://github.com/NanboLi/MulMON>

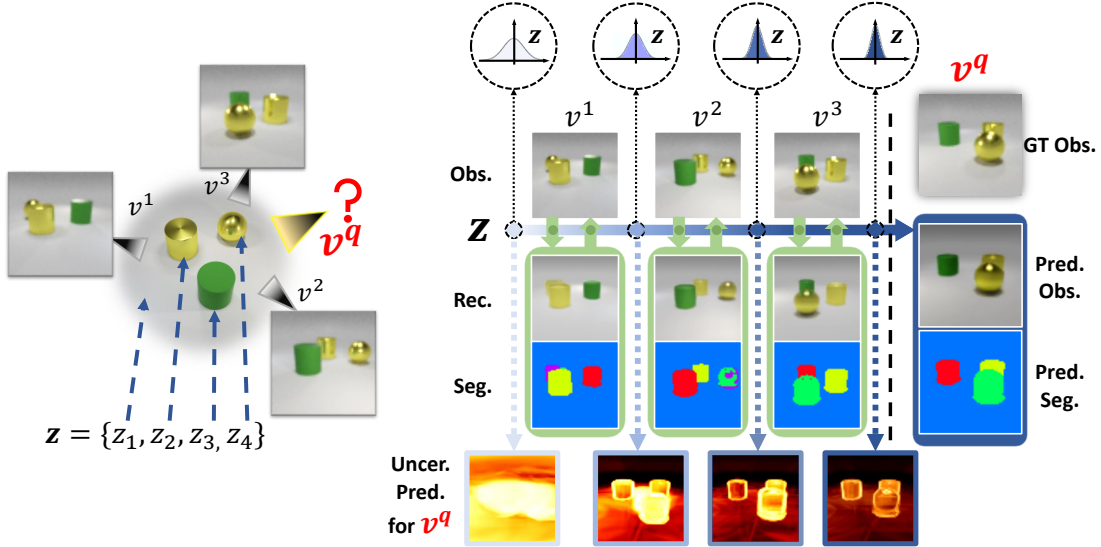


Figure 1: **Left:** *Multi-object-multi-view setup*. v^q denotes the query viewpoint, while z_k denotes “slot” k , i.e. the latent object representation of a scene object. **Right:** *MulMON overview*. Starting with a standard normal prior, MulMON iteratively refines \mathbf{z} over multiple views, each time reducing its uncertainty about the scene—as illustrated by the darkening, white-to-blue arrow. Within-view “inner loop” iterations are depicted by the green arrows and boxes. Cross-view “outer-loop” iterations are depicted by the white-to-blue arrows and boxes. At the bottom, we have visualised MulMON’s reduction in uncertainty about \mathbf{z} in image space, where each image shows the per-pixel variance of MulMON’s predicted observation from query viewpoint v^q . MulMON’s final predictions for v^q (observation and segmentation) are shown to the right of the vertical dotted line.

To address this, we present MulMON (Multi-View and Multi-Object Network)—an unsupervised method for learning object-centric scene representations from multiple views. Using a spatial mixture model [10] and iterative amortized inference [19], MulMON sidesteps the main technical difficulty of the multi-object-multi-view scenario—maintaining object correspondence across views—by iteratively updating the latent object representations for a scene over multiple views, each time using the previous iterations posterior as the new prior. To ensure that these iterative updates do indeed aggregate spatial information, rather than simply overwrite, MulMON is asked to predict the appearance of the scene from novel viewpoints during training. Given images of a *static* scene from several viewpoints, MulMON forms an object-centric representation, then uses this representation to predict the appearance and object segmentations of that scene from unobserved viewpoints. Through experiments we demonstrate that:

- MulMON better-resolves spatial ambiguities than single-view methods like IODINE [9], while providing all the benefits of object-based representations that “single-object” methods like GQN [8] lack, e.g. object segmentations and manipulations (see Section 5).
- MulMON accurately captures 3D scene information (rotation along the vertical axis) by integrating spatial information from multiple views (see Section 5.3).
- MulMON achieves both inter- and intra-object disentanglement—enabling both single-object and single-object-property scene manipulations (see Section 5.3).
- MulMON represents the first feasible solution to the *multi-object-multi-view problem*, permitting new functionality like viewpoint-queried object-segmentation (see Section 5.2).

2 Background

2.1 Multi-object representations

Assuming a single observation or view x^1 for now, we can formally describe the goal of “multi-object” (i.e. object-centric) scene learning as that of computing the posterior $p(z_1, z_2, \dots, z_K | x^1)$, where K is the number of 3D objects in the scene including the background “object”, $z_k \in \mathbb{R}^D$ is a D -dimensional latent representation of object k , and $x^1 \in \mathbb{R}^M$ is a 2D image with M pixels. As K is unknown, recent “multi-object” works [4, 9] have fixed K globally to be a number that is sufficiently large (greater than actual number of objects) to capture all the objects in a scene and allowing for empty slots. Thus, we will use K to represent the number of object *slots* hereafter.

2.2 Multi-object representations from multiple views

As 2D views of 3D scenes are inherently under-specified, they often contain spatial ambiguities (e.g. occlusions or partial occlusions). As a result, “single-view” methods often learn inaccurate representations of the underlying 3D scene. To solve this, it is desirable to aggregate information from multiple views into a single accurate representation of the (static) 3D scene. Recently, this was achieved with some success by GQN for single-object (i.e. single-slot) representations. However, doing so for multi-object representations is much more challenging due to the difficulty of the object matching problem, i.e. maintaining object-representation correspondences across views. As a result, this *multi-object-multi-view* (MOMV) problem remains unsolved.

With this in mind, we can define a more general object-centric scene representation learning problem as that of learning a representation of a K -object scene based on T uncalibrated observations from random viewpoints, where the scenes are static and assumed to be a spatial configuration that is independent of the observer. Formally, this involves computing the posterior $p(z_1, z_2, \dots, z_K | x^1, x^2, \dots, x^T)$. Since a 2D observation of a 3D scene must be associated with a viewpoint, we can better specify the problem as that of computing $p(z_1, z_2, \dots, z_K | x^1, x^2, \dots, x^T, v^1, v^2, \dots, v^T)$, where $v^t \in \mathbb{R}^J$ is the viewpoint associated with image x^t . This can be written more compactly as $p(\mathbf{z} = \{z_k\} | \{(x^t, v^t)\})$, where both the latent object set $\mathbf{z} = \{z_k\}$ and observation set $\{(x^t, v^t)\}$ are permutation-invariant.

3 Method

Our goal is to learn structured, object-centric scene representations that accurately capture the spatial structure of 3D multi-object scenes, and to do this by leveraging multiple 2D views. Key to achieving this is 1) an outer loop that iterates over views, aggregating information while avoiding the object matching problem, and 2) a training procedure that ensures that these outer loops are indeed used to form a complete 3D understanding of the scene, rather than just overwriting each other. We detail 1) in Section 3.1, and 2) in Section 3.4. Additionally, we describe the viewpoint-conditioned generative model and iterative inference procedure in Sections 3.2 and 3.3 respectively.

3.1 Iterating over views

Cross-view iterations (outer loop). For a static scene, we consider that the latent scene representation $\mathbf{z} = \{z_k\}$ is updated sequentially in T steps as the T observations are obtained one-by-one from $t = 1$ to $t = T$, where t denotes the updating step. This suggests that \mathbf{z}^t is obtained by updating \mathbf{z}^{t-1} using a new observation x^t , taken at viewpoint v^t (see the green box

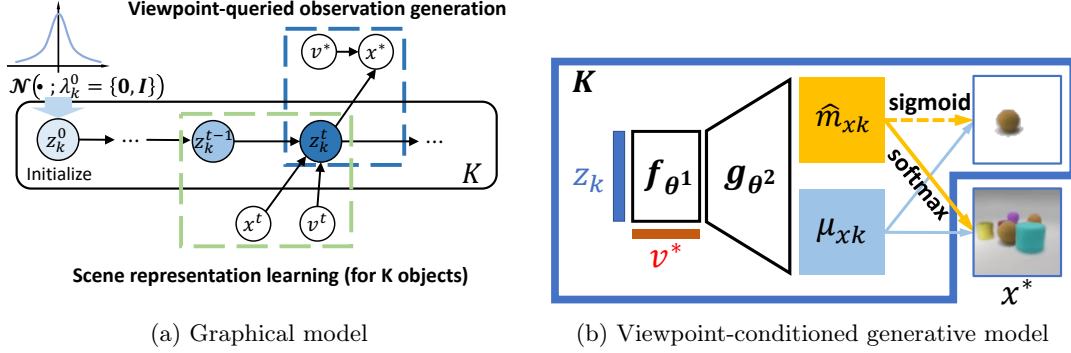


Figure 2: (a) Graphical view of MulMON’s cross-view iterations. The two core components, a viewpoint-conditioned generative model (Section 3.2) and an inference model (Section 3.3), are shown in blue and green boxes respectively. (b) For viewpoint-conditioned generation, each of the K latent object representations \mathbf{z}_k are transformed w.r.t. a viewpoint v^* using the function f_{θ^1} , before being passed through a decoder g_{θ^2} to render a viewpoint-queried observation x^* . As shown, g_{θ^2} actually outputs the pixel-wise means μ_{xk} and predicted object masks $\text{softmax}(\hat{m}_{xk})$ of the spatial Gaussian mixture model, which are combined to get x^* .

in Figure 2a). Therefore, by making an assumption that $\mathbf{z} = \mathbf{z}^t$ for any integer $t \in [1, T]$, we can compute the target multi-view posterior in a recursive form as:

$$p(\mathbf{z} = \{z_k\} | x^{1:T}, v^{1:T}) = p(\mathbf{z}^0) \prod_{t=1}^T p(\mathbf{z}^t | x^t, v^t, \mathbf{z}^{t-1}), \quad (1)$$

where z_k^{t-1} is the latent representation of a scene object (indexed by k) before observing x^t at viewpoint v^t , z_k^t the representation afterwards. and $p(\mathbf{z}^0)$ the initial guess which we assume to be a standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. The formulation in equation 1 turns the multi-view problem into a recursive single-view problem and, in theory, enables online learning of scenes from an infinitely large number of observations without causing memory overflow.

Within-view iterations (inner loop). As shown in Figure 2, MulMON consists of a scene-representation inference model and a viewpoint-conditioned generative model. In each iteration, the inference model starts with a prior assumption about the K objects in the latent space, i.e. $\mathbf{z} = \mathbf{z}^{t-1} = \{z_k^{t-1}\}$, and approximates the target posterior $p(\mathbf{z}^t = \{z_k^t\} | x^t, v^t, \mathbf{z}^{t-1})$ after observing x^t at viewpoint v^t . The approximation, as mentioned in Section 1, is handled by iterative amortized inference[19] and the approximate posterior is passed to the next iteration as the prior assumption. Therefore, a single iteration is a single-view process that takes a latent prior about $\mathbf{z}^{t-1} = \{z_k^{t-1}\}$ and an image observation x^t (taken at a viewpoint v^t) as inputs. We call the single-view iterative process the *inner loop*, and the cross-view Bayesian updating process (see equation 1) the *outer loop*.

3.2 Generative Model

We model an image x^t with a spatial Gaussian mixture model[24, 10], similar to MONet[4] and IODINE[9], and additionally we take as input (condition on) the viewpoint v^t . We can then write the generative likelihood as:

$$p_{\theta}(x^t | \mathbf{z}^t, v^t) = \prod_{i=1}^M \sum_{k=1}^K p_{\theta}(C_i^t = k | z_k^t, v^t) \cdot p_{\theta}(x_{ik}^t | z_k^t, v^t), \quad (2)$$

where x_{ik}^t are the RGB values in image t at a pixel location i that pertain to object k , $p_{\theta}(x_{ik}^t | z_k^t, v^t)$ is the Gaussian density function parametrized by a neural network θ , and m_{ik}

is the mixing coefficient for object k and pixel i , i.e. the probability that pixel i is assigned to the k -th object. More formally, $m_{ik} = p_\theta(C_i^t = k | z_k^t, v^t)$, where C_i^t is a categorical random variable and $C_i^t = k$ represents the event that pixel i is assigned to the k -th object. This is an important property for object segmentation, as it implies that every pixel in x^t must be explained by one and only one object. Together, the M mixing coefficients for object k (one per pixel) form a soft object segmentation mask $m_k = p_\theta(C^t = k | z_k^t, v^t)$. We assume all pixel values x_{ik}^t are independent given the corresponding latent object representation z_k^t and viewpoint v^t , and simplify computations by using a fixed variance $\sigma^2 = 0.01$ for all pixels. In practice, we split the parameters θ into two pieces, θ^1 and θ^2 , in order to handle the viewpoint-queried neural transformation and observation-generation separately in two consecutive stages. That is, we first transform the K latent object representations \mathbf{z}^t w.r.t. a viewpoint v^q using the function f_{θ^1} , then we pass the output through a decoder g_{θ^2} in order to render a viewpoint-queried observation x^q . We illustrate this process in Figure 2b and Algorithm 1.

3.3 Inference

Equation 1 simplifies the inference problem by breaking the computation of the multi-object-multi-view posterior $p(\mathbf{z} = \{z_k\} | \{(x^t, v^t)\})$, into a recursive computation of multi-object-single-view posteriors, $p(\mathbf{z}^t | x^t, v^t, \mathbf{z}^{t-1})$. However, exact inference of $p(\mathbf{z}^t = \{z_k^t\} | x^t, v^t, \mathbf{z}^{t-1})$ is still intractable. Similar to IODINE[9], we apply iterative amortized inference[19] to approximate the intractable target posterior. However, unlike IODINE which always initializes the prior from a standard Gaussian, the inference model of MulMON takes an approximate posterior from last iteration as the prior. Hence, we approximate the intractable posterior with $q_\lambda(\mathbf{z}^t | x^t, v^t, \mathbf{z}^{t-1})$, where $\lambda = \{\lambda_k\} = \{(\mu_k, \sigma_k)\}$ parametrizes a set of object-specific Gaussian distributions in the latent space. We denote the number of iterations for the inner loop with L , and each iteration is indexed by l . The parameter update in the iterative inference is thus:

$$\begin{aligned} z_k^{t(l)} &\stackrel{k}{\sim} q_{\lambda_k^{(l)}}(z_k^t | x^t, v^t, z_k^{t-1}) \\ \lambda_k^{(l+1)} &\stackrel{k}{\leftarrow} \lambda_k^{(l)} + f_\Phi(z_k^{t(l)}, x^t, v^t, \mathbf{a}(\cdot)), \end{aligned} \quad (3)$$

where the refinement function f_Φ , with trainable parameter Φ , is modeled by a recurrent neural network. The $\stackrel{k}{\sim}$ and $\stackrel{k}{\leftarrow}$ operators denote parallel operations over K independent object slots. The same auxiliary inputs such as mask gradients $\nabla_{m_k} \mathcal{L}$ and posterior gradient $\nabla_{\lambda_k} \mathcal{L}$, where \mathcal{L} is the objective function of MulMON (will be discussed in Section 3.4), as that of IODINE are also adopted to refine the posterior. These auxiliary inputs are computed by a function $\mathbf{a}(\cdot)$, namely the ‘‘auxiliary function’’, which takes in the refinement function’s inputs along with the posterior parameter $\lambda_k^{(l)}$.

The variational approximate posterior of MulMON is thus:

$$q_\lambda(\mathbf{z} = \{z_k\} | x^{1:T}, v^{1:T}) = q(\mathbf{z}^0) \prod_{t=1}^T q_\lambda(\mathbf{z}^t | x^t, v^t, \mathbf{z}^{t-1}), \quad (4)$$

where the initial guess $q(\mathbf{z}^0)$ is a standard Gaussian and this is the same as $p(\mathbf{z}^0)$ in equation 1. We refer to Algorithm 1 for more details about MulMON’s inference process and its behaviors at test time.

3.4 Training

Similar to IODINE, MulMON learns the decoder parameters θ and the refinement network parameters Φ by minimizing $\mathcal{D}_{\text{KL}}[q_\lambda(\mathbf{z} | x^{1:T}, v^{1:T}) || p_\theta(\mathbf{z} | x^{1:T}, v^{1:T})]$, which is equivalent to maximizing the evidence lower bound (i.e. the ELBO, denoted as \mathcal{L}) in a generative configuration. However, instead of directly maximizing the ELBO like IODINE, we simulate novel viewpoint-queried

Algorithm 1: MulMON at Test Time: Online Scene Learning

Input: Trained parameters Φ, θ
Hyperparameters $K, \sigma^2 = 0.01, L$ **Initialize** $\lambda^0 = \{\lambda_k^0\} \leftarrow \{(\mu_k = \mathbf{0}, \sigma_k = \mathbf{I})\};$
/* The outer loop for scene learning */
for $t = 1$ **to** T **do**
 Access a scene observation $(x^t, v^t);$
 $\lambda^{prior} = \lambda^{t(0)} \leftarrow \lambda^{t-1};$
 /* The inner loop for observation aggregation */
 for $l = 0$ **to** $L - 1$ **do**
 $\mathbf{z}^{t(l)} \sim \mathcal{N}(\mathbf{z}^{t(l)}; \lambda^{t(l)});$
 $\{\mu_{x_k}^{(l)}, \hat{m}_{x_k}^{(l)}\} \leftarrow g_{\theta^2}(f_{\theta^1}(\mathbf{z}^{t(l)}, v^t));$
 $\{m_k^{(l)}\} \leftarrow \mathbf{softmax}(\{\hat{m}_{x_k}^{(l)}\});$
 /* The spatial Gaussian mixture */
 $p_{\theta}(x^t | \mathbf{z}^{t(l)}, v^t) \leftarrow \sum_k m_k^{(l)} \mathcal{N}(x_k^t; \mu_{x_k}^{(l)}, \sigma^2 \mathbf{I});$
 if $l == 0$ **then**
 $\mathcal{L}_{\mathcal{T}}^{(l)} \leftarrow -\log p_{\theta}(x^t | \mathbf{z}^{t(l)}, v^t);$
 else
 $\mathcal{L}_{\mathcal{T}}^{(l)} \leftarrow \mathcal{D}_{\text{KL}}[\mathcal{N}(\mathbf{z}^{t(l)}; \lambda^{t(l)}) || \mathcal{N}(\mathbf{z}^{t(l)}; \lambda^{prior})] - \log p_{\theta}(x^t | \mathbf{z}^{t(l)}, v^t);$
 $\lambda^{t(l+1)} \leftarrow \lambda^{t(l)} + f_{\Phi}(z_k^{t(l)}, x^t, v^t, \mathbf{a}(\cdot));$
 $\lambda^t \leftarrow \lambda^{t(l+1)};$

generation in the training process (similar to GQN [8]). By asking MulMON to predict the appearance of a scene from unobserved viewpoints during training, we ensure that the iterative updates are indeed used to aggregate spatial information across views, as a complete 3D scene understanding is required to perform well. More formally, we randomly partition the set of T scene observations $\{(x^t, v^t)\}$ into two subsets \mathcal{T} and \mathcal{Q} , with $n \sim \mathcal{U}(1, 5)$ observations in \mathcal{T} and the remaining $T - n$ observations in \mathcal{Q} . We perform scene learning on \mathcal{T} and novel viewpoint-queried generation on \mathcal{Q} . We thus derive the MulMON ELBO (for one scene sample) as:

$$\begin{aligned} \mathcal{L} = & \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \mathbf{E}_{q_{\lambda}(\mathbf{z}^t | \cdot)} [\log p_{\theta}(x^t | \mathbf{z}^t, v^t)] + \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}, t \sim \mathcal{T}} \mathbf{E}_{q_{\lambda}(\mathbf{z}^t | \cdot)} [\log p_{\theta}(x^q | \mathbf{z}^t, v^q)] \\ & - \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \mathbf{IG}(\mathbf{z}^t, x^t; v^t, \mathbf{z}^{t-1}) \end{aligned} \quad (5)$$

where \mathbf{IG} is the information gain (aka. Bayesian surprise), the operation $|\cdot|$ measures the size of a discrete set, and $q_{\lambda}(\mathbf{z}^t | \cdot)$ is an abbreviation of the variational posterior $q_{\lambda}(\mathbf{z}^t | x^t, v^t, \mathbf{z}^{t-1})$, from which we sample \mathbf{z}^t by applying ancestral sampling. In practice, we use an efficient approximation of the information gain, i.e. an approximate $\mathbf{IG} \mathcal{D}_{\text{KL}}[q_{\lambda}(\mathbf{z}^t | x^t, v^t, \mathbf{z}^{t-1}) || q_{\lambda}(\mathbf{z}^t | v^t, \mathbf{z}^{t-1})]$ or $\mathcal{D}_{\text{KL}}[q_{\lambda}(\mathbf{z} = \mathbf{z}^t | x^t, v^t, x^{1:t-1}, v^{1:t-1}) || q_{\lambda}(\mathbf{z} = \mathbf{z}^{t-1} | x^{1:t-1}, v^{1:t-1})]$. Note that using a fixed number of observations could harm the model’s robustness at test time, hence why we randomly partition the observations into size-varying sets \mathcal{T} and \mathcal{Q} during training, i.e. we train the model with varying number of observations. See Appendix A for full details of the training algorithm of MulMON.

4 Related Work

Single-object-single-view (SOSV). Many recent breakthroughs in unsupervised representation learning of have come in the form of “disentanglement” models [11, 6, 14] that seek feature-attribute-level understanding by encouraging e.g. independence among latent dimensions. However, most of these models focus on a single view of a single object that has been placed

in front of some background (e.g. dSprites, CelebA, 3D Chairs). As a result, they fail to i) generalize to more realistic, multi-object scenes, and ii) accurately capture 3D scene information (e.g. resolve single-view spatial ambiguities and estimate e.g. rotation along the vertical axis).

Multi-object-single-view (MOSV). To avoid the additional computational complexities of factorising or segmenting the objects in a scene into explicit *multi-object* representations, many works have used pre-segmented images [26, 27]. However, this comes at the cost of decreased representational power (good object representation requires good object segmentation [9]) and a reliance on annotated data. In addition, these works struggle in a multi-view scenario where pre-segmented images require consistent multi-frame object registration and tracking, since the segmentation and representation models work independently. More recently, several works [4, 9, 18, 17, 1] have succeeded in approximating the factorized posterior $p(z_1, z_2, \dots, z_K|x)$ within the VAE framework, achieving impressive unsupervised object-level scene factorization. However, being single-view models, they fall victim to single-view spatial ambiguities. As a result, they fail to accurately capture the scene’s 3D spatial structure, causing problems for object-level segmentation. To overcome this and learn object-based representations that accurately capture 3D spatial structures, MulMON essentially extends these models to the multi-view scenario.

Single-object-multi-view (SOMV). Recently, unsupervised models like GQN [8] and EGQN [22] have been quite successful in aggregating multiple observations of a scene into a single-slot representation that accurately captures the spatial layout of the 3D scene, as shown by their ability to predict the appearance of a scene from unobserved viewpoints. However, being single-slot or “single-object” models, they fail to achieve object-level scene understanding in multi-object scenes, and as a result, miss out on the aforementioned benefits of object-centric scene representations. To overcome this, MulMON essentially extends these models to the case of multi-object representations. In addition, several works have sought explicit 3D representations either in the latent space [21] or output space [25, 2]. However, due to the complexity of 1) working with explicit 3D object representations and 2) maintaining object correspondences across views, these works have been limited to single-object scenes (often quite simple, with “floating” objects placed in front of a plain background).

Multi-object scenes in videos. While some works in multi-object discovery and tracking in videos appear to be MOMV models [16, 12], they in fact work with one view per scene (abiding strictly by our definition of a scene in Section 2) and are only capable of dealing with binarizable MNIST-like images.

5 Experiments

Our experiments are designed to demonstrate that MulMON is a feasible solution to the MOMV problem, and to demonstrate that MulMON learns better representations than the MOSV and SOMV models by resolving spatial ambiguity. To do so, we compare the performance of MulMON against two baseline models, IODINE[9] (MOSV) and GQN[8] (SOMV), in terms of segmentation, viewpoint-queried prediction (appearance and segmentation) and disentanglement (inter- and intra-object). To best facilitate these comparisons, we created two new datasets called CLEVR6-MultiView (abbr. CLE-MV) and CLEVR6-Augmented (abbr. CLE-Aug) which contain ground-truth segmentation masks and shape descriptions (e.g. colors, materials, etc.). The CLE-MV dataset is a multi-view, observation-enabled variant (10 views per scene) of the CLEVR6 dataset[13, 9]. The CLE-Aug adds more complex shapes (e.g. horses, ducks, and teapots etc.) to the CLE-MV environment. In addition, we compare the models on the GQN-Jaco dataset[8] and use the GQN-Shepard-Metzler7 dataset[8] (abbr. Shep7) for a specific ablation study. We train all models using an Adam optimizer with an initial learning rate 0.0003 for 300k

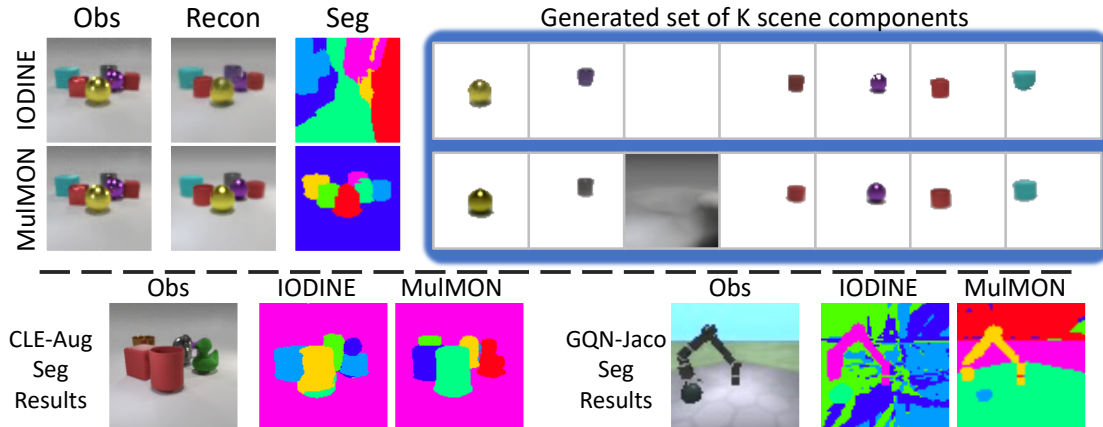


Figure 3: Qualitative comparison of MulMON vs. IODINE in terms of scene segmentation performance. **Top:** L) Reconstruction and segmentation comparison on a CLE-MV data sample. R) Individual masked-object generation using each object’s representation independently (see Appendix C for the computation details). **Bottom:** Segmentation performance on CLE-Aug and GQN-Jaco data samples (specific colors arbitrary).

| Models | CLE-MV | CLE-Aug |
|--------|---------------------------------------|---------------------------------------|
| GQN | N/A | N/A |
| IODINE | 0.1891 ± 0.0000 | 0.5137 ± 0.0007 |
| MulMON | 0.7852 ± 0.0008 | 0.7076 ± 0.0004 |

(a) Object Segmentation mIoU Scores

| Models | CLE-MV | CLE-Aug |
|--------|---------------------------------------|---------------------------------------|
| GQN | N/A | N/A |
| IODINE | N/A | N/A |
| MulMON | 0.7845 ± 0.0011 | 0.6860 ± 0.0006 |

(b) Predicted Object Segmentation mIoU Scores

| Models | CLE-MV | CLE-Aug | GQN-Jaco |
|--------|---------------------------------------|---------------------------------------|---------------------------------------|
| GQN | 0.1426 ± 0.0002 | 0.1482 ± 0.0001 | 0.1675 ± 0.0013 |
| IODINE | N/A | N/A | N/A |
| MulMON | 0.0464 ± 0.0004 | 0.0733 ± 0.0003 | 0.1607 ± 0.0018 |

(c) Predicted Observation RMSE (pixel avg.)

| Models | Disent. | Compl. | Inform. |
|--------|-----------------------------------|-----------------------------------|-----------------------------------|
| GQN | N/A | N/A | N/A |
| IODINE | 0.47 ± 0.00 | 0.60 ± 0.01 | 0.67 ± 0.01 |
| MulMON | 0.65 ± 0.01 | 0.73 ± 0.01 | 0.78 ± 0.00 |

(d) Disentanglement Analysis (CLE-MV)

Table 1: Quantitative comparisons of MulMON, IODINE and GQN. “N/A” denotes cases where a model is *unable* to perform a task. In tables (a), (b) and (d), higher is better and 1 is best. For table (c), lower is better and 0 is best.

gradient steps. In addition, all experiments were run across five different random seeds to simulate scenarios of different observation orders and view selections. For more details about the four datasets and model implementations see Appendix B and C respectively in the supplementary materials.

5.1 Scene Factorization

The ability of MulMON to perform scene object decomposition in the scene learning phase is crucial for learning object-centric scene representations. We evaluate its segmentation ability by computing mean-intersection-over-union (mIoU) scores between the output and the GT masks. However, since the segments produced by IODINE and MulMON are unordered, GT masks and object segmentation masks need to first be one-to-one registered for each scene. We solve this matching problem by first computing every possible object pairs of GT object masks and outputs, then, for each GT object mask, we find the output object mask that gives the highest IoU score. Table 1a shows that MulMON outperforms IODINE in object segmentation. The qualitative comparison in Figure 3 shows that IODINE captures each object well independently but fails to understand the spatial structure along depth directions (3D) – as described by the Categorical distribution (see Section 3.2). Note that IODINE’s poor segmentation performance is mostly due to its poor handling of the background, i.e. its tendency to split up the background.




Figure 4: Qualitative results on novel-viewpoint prediction. **Left:** Working example of MulMON, including uncertainty reduction across views, for a GQN-Jaco sample. **Right:** Qualitative comparison between MulMON and GQN.

Although the background is often considered a less-important “object”, correct handling of the background demonstrates better spatial-reasoning ability. Together, all of these results suggest that MulMON learns better single-object representations and spatial structures by overcoming spatial ambiguities.

5.2 Novel-viewpoint Prediction

MulMON can predict both observations and segmentation for novel viewpoints. This is the major advantage of our model (a MOMV model) over the MOSV and SOMV models in scene understanding. For our evaluation of online scene learning, each model is provided with 5 observations of each scene and then asked to predict both the observation and segmentation for randomly-selected novel viewpoints. We compute the root-mean-square error (RMSE) and mIoU as quality measures of the predicted observation and segmentation respectively. Table 1c shows that MulMON outperforms GQN on novel-view observation prediction. Table 1b shows that MulMON is the only model that can predict the object segmentation for novel viewpoints – and it does so with a similar quality to the original object segmentation (compare with Table 1a). However, as shown in Figure 4, GQN tends to capture more pixel details than MulMON, albeit at the risk of predicting wrong spatial configurations.

5.3 Disentanglement Analysis

To evaluate how well MulMON performs disentanglement at both the inter-object level and the intra-object level, we run disentanglement analyses on the representations learned by MulMON. For our qualitative analysis, we pick one of K objects in a scene, and traverse one dimension of the learned object-representation at a time. Figure 5 (left) shows i) MulMON’s intra-object disentanglement, encoding interpretable features in different latent dimensions; and ii) MulMON’s inter-object disentanglement, allowing single-object manipulation without affecting other objects in the scene. Figure 5 (right) shows that MulMON captures 3D information (vertical-axis rotation) and broadcasts consistent manipulations of this 3D information to different views. For our quantitative analysis, we employ the method of Eastwood and Williams[7] to compare the representations learned by each model on the CLE-MV and CLE-Aug datasets. As shown in Table 1d, MulMON learns object representations that are more disentangled, complete (compact) and informative (about ground-truth object properties). See Appendix D for further details.




Figure 5: Single-object manipulations via latent traversals. **Left** Traversing two dimensions of the duck’s latent representation (one per row). Top row cropped for visual clarity. **Right** For two different views (one per row), we manipulate the dimension of the learned representation that appears to capture vertical-axis rotation.

5.4 Ablation Study

We consider the number of observations T to be the most important hyperparameter of MulMON as the key insight of MulMON is to reduce multi-object spatial uncertainty by aggregating information across multiple observations. To visualize the effect of T on MulMON’s performance, we plot MulMON’s uncertainty about the scene as a function of T . More specifically, for a given scene and ordering of the observations, we: 1) draw 10 samples from the approximate variational posterior $q_{\lambda}(\mathbf{z} = \{z_k\} | x^{1:T}, v^{1:T})$, 2) obtain the corresponding viewpoint-queried observation predictions using the 10 latent samples (see Section 3.2 and Figure 2b), 3) compute the pixel-wise empirical variance over these observation predictions. Averaging over all scenes in the dataset and sampling 5 random view orderings (5 different random seeds), we can then create Figure 6 which shows that MulMON effectively reduces the spatial uncertainty/ambiguity by leveraging multiple views. In particular, MulMON’s uncertainty is rapidly reduced after only a small number of observations T . We also study the effects of two other important hyperparameters, namely the globally-fixed number of object slots K and the coefficient of information gain \mathbf{IG} (in the MulMON ELBO). For details on these further ablation studies, we refer the reader to Appendix D.

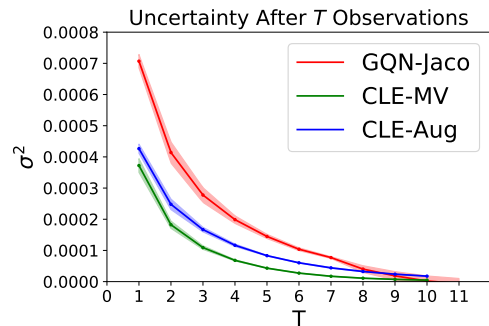


Figure 6: Uncertainty vs. T .

6 Conclusion

We have presented MulMON—a method for learning accurate, object-centric representations of multi-object scenes by leveraging multiple views. We have shown that MulMON’s ability to aggregate information across multiple views does indeed allow it to better-resolve spatial ambiguity (or uncertainty) and better-capture 3D spatial structures, and as a result, outperform state-of-the-art models for unsupervised object segmentation. We have also shown that, by virtue of addressing the more complicated multi-object-multi-view scenario, MulMON achieves new functionality—the prediction of both appearance and object segmentations for novel viewpoints. As all scenes in this paper are static, future work may look to extend MulMON to dynamic multi-object scenes.

Acknowledgements

This research is partly supported by the *Trimbot2020* project, which is funded by the European Union Horizon 2020 programme. The authors would like to thank Prof. C. K. I. Williams for his valuable advice that helps to improve this work and acknowledge the GPU computing support from Dr. Zhibin Li’s Advanced Intelligence Robotics Lab at the University of Edinburgh.

Broader Impact

In this paper, we presented a new method to learn object-centric representations of multi-object scenes. Object-centric scene representations can support many downstream tasks, such as autonomous scene exploration, object segmentation (tracking) and scene synthesizing.

Autonomous scene exploration has real-world applications in exploring hazardous environments, mines, potential bomb threats, nuclear waste zones. This could have societal impacts through increased worker safety or potential military (mis)uses.

Object detection and tracking has real-world applications in tracking people in CCTV footage, detecting buildings from aerial footage, and spotting potential hazards for autonomous vehicles. Potential societal impacts include safer autonomous vehicles and unwanted/increased surveillance.

Finally, scene synthesizing has applications in automated scene modelling for computer games. This further transitions society away from labor-intensive tasks to higher-level cognitive tasks. This could have both positive (more time for cognitive tasks) and negative (less employment) impacts on society.

References

- [1] T. Anciukevičius, C. Lampert, and P. Henderson. Object-centric image generation with factored depths, locations, and appearances. *arXiv:2004.00642*, 2020.
- [2] A. Arsalan Soltani, H. Huang, J. Wu, T. D. Kulkarni, and J. B. Tenenbaum. Synthesizing 3d shapes via modeling multi-view depth maps and silhouettes with deep generative networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1511–1519, 2017.
- [3] V. Bapst, A. Sanchez-Gonzalez, C. Doersch, K. L. Stachenfeld, P. Kohli, P. W. Battaglia, and J. B. Hamrick. Structured agents for physical construction. In *International Conference on Machine Learning*, pages 464–474, 2019.
- [4] C. P. Burgess, L. Matthey, N. Watters, R. Kabra, I. Higgins, M. Botvinick, and A. Lerchner. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019.
- [5] D. Carlos, A. Cohen, and M. L. Littman. An object-oriented representation for efficient reinforcement learning. In *International Conference on Machine Learning*, pages 240–247, 2008.
- [6] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2172–2180, 2016.
- [7] C. Eastwood and C. K. I. Williams. A framework for the quantitative evaluation of disentangled representations. In *International Conference on Learning Representations*, 2018.
- [8] S. A. Eslami, D. J. Rezende, F. Besse, F. Viola, A. S. Morcos, M. Garnelo, A. Ruderman, A. A. Rusu, I. Danihelka, K. Gregor, et al. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- [9] K. Greff, R. L. Kaufman, R. Kabra, N. Watters, C. Burgess, D. Zoran, L. Matthey, M. Botvinick, and A. Lerchner. Multi-object representation learning with iterative variational inference. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2424–2433, 2019.
- [10] K. Greff, S. Van Steenkiste, and J. Schmidhuber. Neural expectation maximization. In *Advances in Neural Information Processing Systems*, pages 6691–6701, 2017.

- [11] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.
- [12] J.-T. Hsieh, B. Liu, D.-A. Huang, L. F. Fei-Fei, and J. C. Niebles. Learning to decompose and disentangle representations for video prediction. In *Advances in Neural Information Processing Systems*, pages 517–526, 2018.
- [13] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. Lawrence Zitnick, and R. Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910, 2017.
- [14] H. Kim and A. Mnih. Disentangling by factorising. In *International Conference on Machine Learning*, pages 2649–2658, 2018.
- [15] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [16] A. Kosiosek, H. Kim, Y. W. Teh, and I. Posner. Sequential attend, infer, repeat: Generative modelling of moving objects. In *Advances in Neural Information Processing Systems*, pages 8606–8616, 2018.
- [17] Z. Lin, Y.-F. Wu, S. V. Peri, W. Sun, G. Singh, F. Deng, J. Jiang, and S. Ahn. Space: Unsupervised object-oriented scene representation via spatial attention and decomposition. In *International Conference on Learning Representations*, 2020.
- [18] F. Locatello, D. Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, A. Dosovitskiy, and T. Kipf. Object-centric learning with slot attention. In *Advances in Neural Information Processing Systems*, 2020.
- [19] J. Marino, Y. Yue, and S. Mandt. Iterative amortized inference. In *International Conference on Machine Learning*, pages 3403–3412, 2018.
- [20] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [21] D. J. Rezende, S. A. Eslami, S. Mohamed, P. Battaglia, M. Jaderberg, and N. Heess. Unsupervised learning of 3d structure from images. In *Advances in neural information processing systems*, pages 4996–5004, 2016.
- [22] J. Tobin, W. Zaremba, and P. Abbeel. Geometry-aware neural rendering. In *Advances in Neural Information Processing Systems*, pages 11555–11565, 2019.
- [23] N. Watters, L. Matthey, C. P. Burgess, and A. Lerchner. Spatial broadcast decoder: A simple architecture for learning disentangled representations in vaes. *arXiv preprint arXiv:1901.07017*, 2019.
- [24] C. K. I. Williams and M. K. Titsias. Greedy learning of multiple objects in images using robust statistics and factorial learning. *Neural Computation*, 16(5):1039–1062, 2004.
- [25] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. 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.
- [26] S. Yao, T. M. Hsu, J.-Y. Zhu, J. Wu, A. Torralba, B. Freeman, and J. Tenenbaum. 3d-aware scene manipulation via inverse graphics. In *Advances in Neural Information Processing Systems*, pages 1887–1898, 2018.
- [27] I. Yildirim, T. D. Kulkarni, W. A. Freiwald, and J. B. Tenenbaum. Efficient and robust analysis-by-synthesis in vision: A computational framework, behavioral tests, and modeling neuronal representations. In *CogSci*, volume 1, 2015.

Note: we use the same notations in the Appendix as that in the main paper.

A. Training Algorithm of MulMON

We refer to Algorithm 2 and 3 for the training algorithm of MulMON.

Algorithm 2: MulMON Training Algorithm

```

Data a set of  $N$  scenes as  $\{(images\ x^{1:T},\ viewpoints\ v^{1:T})\}_N$ 
begin
  Initialize trainable parameters  $\Phi^{(0)}, \theta^{(0)}$ , step count  $s = 0$ ;
  repeat
    Sample mini batch  $\{(x^{1:T}, v^{1:T})\}_M \sim \{(x^{1:T}, v^{1:T})\}_N$ , where  $M \leq N$ ;
    /* The below loop can go parallel as tensor operations */
    for  $(x^{1:T}, v^{1:T})$  in  $\{(x^{1:T}, v^{1:T})\}_M$  do
       $\mathcal{L}_m \leftarrow \text{SingleSampleELBO}((x^{1:T}, v^{1:T}), \Phi^{(s)}, \theta^{(s)});$ 
       $\mathcal{L} = \frac{1}{M} \sum_{m=1}^M \mathcal{L}_m;$ 
      /* gradient update */
       $\Phi^{(s+1)} \leftarrow \text{optimizer}(\mathcal{L}, \Phi^{(s)});$ 
       $\theta^{(s+1)} \leftarrow \text{optimizer}(\mathcal{L}, \theta^{(s)});$ 
       $s \leftarrow s + 1;$ 
  until  $\Phi, \theta$  converge;

```

B. Data Configurations

We show samples of the used datasets in Figure 14. CLEVR-MultiView & CLEVR-Augmented We




Figure 7: Examples of the four dataset used in this work.

adapt the Blender environment of the original CLEVR datasets[13] to render both datasets. We make a scene by randomly sampling 3 ~ 6 rigid shapes as well as their properties like poses, materials, colors etc.. For the CLEVR-MultiView (CLE-MV) dataset, we sample shapes from three categories: cubes, spheres, and cylinders, which are the same as the original CLEVR dataset. For the CLEVR-Augmented (CLE-Aug), we add more shape categories into the pool: mugs, teapots, ducks, owls, and horses. We render 10

Algorithm 3: SingleSampleELBO

Input: A single scene sample (images $x^{1:T}$, viewpoints $v^{1:T}$), trainable parameters Φ , θ
Hyperparameters K , σ^2 , L

begin

```
 $\mathcal{T} = \{(x^t, v^t)\}, \mathcal{Q} = \{(x^q, v^q)\} \xleftarrow{\text{random split } T} (x^{1:T}, v^{1:T});$   
Initialize  $\lambda^0 = \{\lambda_k^0\} \leftarrow \{(\mu_k = \mathbf{0}, \sigma_k = \mathbf{I})\};$   
/* The outer loop for scene learning */  
for  $t = 1$  to  $|\mathcal{T}|$  do  
  Access a scene observation  $(x^t, v^t);$   
   $\lambda^{prior} = \lambda^{t(0)} \leftarrow \lambda^{t-1};$   
  /* The inner loop for observation aggregation */  
  for  $l = 0$  to  $L - 1$  do  
     $\mathbf{z}^{t(l)} \sim \mathcal{N}(\mathbf{z}^{t(l)}; \lambda^{t(l)});$   
     $\{\mu_{x_k}^{(l)}, \hat{m}_k^{(l)}\} \leftarrow g_{\theta^2}(f_{\theta^1}(\mathbf{z}^{t(l)}, v^t));$   
     $\{m_k^{(l)}\} \leftarrow \text{softmax}(\{\hat{m}_k^{(l)}\});$   
     $p_{\theta}(x^t | \mathbf{z}^{t(l)}, v^t) \leftarrow \sum_k m_k^{(l)} \mathcal{N}(x_k^t; \mu_{x_k}^{(l)}, \sigma^2 \mathbf{I});$   
    if  $l == 0$  then  
       $\mathcal{L}_{\mathcal{T}}^{(l)} \leftarrow -\log p_{\theta}(x^t | \mathbf{z}^{t(l)}, v^t);$   
    else  
       $\mathcal{L}_{\mathcal{T}}^{(l)} \leftarrow \mathcal{D}_{\text{KL}}[\mathcal{N}(\mathbf{z}^{t(l)}; \lambda^{t(l)}) || \mathcal{N}(\mathbf{z}^{t(l)}; \lambda^{prior})] - \log p_{\theta}(x^t | \mathbf{z}^{t(l)}, v^t);$   
       $\lambda^{t(l+1)} \leftarrow \lambda^{t(l)} + f_{\Phi}(z_k^{t(l)}, x^t, v^t, \mathbf{a}(\cdot));$   
     $\lambda^t \leftarrow \lambda^{t(l+1)};$   
   $\mathcal{L}_{\mathcal{T}}^t \leftarrow \frac{2l+2}{L^2+L} \sum_l \mathcal{L}_{\mathcal{T}}^{(l)};$   
  /* Viewpoint-queried prediction */  
  for  $(x^q, v^q)$  in  $\mathcal{Q}$  do  
     $\mathbf{z}^t \sim \mathcal{N}(\mathbf{z}^t; \lambda^t);$   
     $\{\mu_{x_k}^q, \hat{m}_k^q\} \leftarrow g_{\theta^2}(f_{\theta^1}(\mathbf{z}^t, v^q));$   
     $\{m_k^q\} \leftarrow \text{softmax}(\{\hat{m}_k^q\});$   
     $p_{\theta}(x^q | \mathbf{z}^t, v^q) \leftarrow \sum_k m_k^q \mathcal{N}(x_k^q; \mu_{x_k}^q, \sigma^2 \mathbf{I});$   
     $\mathcal{L}_{\mathcal{Q}}^q \leftarrow -\log p_{\theta}(x^q | \mathbf{z}^t, v^q);$   
  /* Compute the MulMON ELBO */  
   $\mathcal{L} = \frac{1}{|\mathcal{T}|} \sum_t \mathcal{L}_{\mathcal{T}} + \frac{1}{|\mathcal{Q}|} \sum_q \mathcal{L}_{\mathcal{Q}}$ 
```

Output: \mathcal{L}

image observations for each scene and save the 10 camera poses as 10 viewpoint vectors. We use resolution 64×64 for the CLE-MV images and 128×128 for the CLE-Aug images. All viewpoints are at the same horizontal level but different azimuth with their focuses locked at the scene center. We thus parametrize a viewpoint 3-D viewpoint vector as $(\cos \alpha, \sin \alpha, r)$, where α is the azimuth angle and r is the distance to the scene center. In addition, we save the object properties (e.g. shape categories, materials, and colors etc.) and generate objects' segmentation masks for quantitative evaluations. CLEVR-MultiView (CLE-MV) contains 1500 training scenes, 200 testing images. CLEVR-Augmented (CLE-Aug) contains 2434 training scenes and 500 testing scenes.

GQN-Jaco We use a mini subset of the original GQN-Jaco dataset[8] in our paper. The original GQN-Jaco contains 4 million scenes, each of them contains 20 image observations (resolution: 64×64) and 20 corresponding viewpoint vectors (7D). To reduce the storage memory and accelerate the training, we randomly sample 2,000 scenes for training and 500 scenes for testing. Also, for each scene, we use only 11 observations (viewpoints) that are randomly sampled from the 20 observations of the original dataset.

GQN-Shepard-Metzler7 Same as the GQN-Jaco dataset, we make a mini GQN-Shepard-Metzler7 dataset[8] (Shep7) by randomly selecting 3000 scenes for training and 200 for testing. Each scene contains 15 images observations (resolution: 64×64) with 15 corresponding viewpoint vectors (7D). We use Shep7 to study the effect of K on our model.

C. Implementation Details

Training configurations See Table 2 for our training configurations.

Table 2: Training Configurations

| TYPE | MULMON, IODINE, GQN |
|--------------------------------|---|
| OPTIMIZER | ADAM |
| INITIAL LEARNING RATE η_0 | $3e^{-4}$ |
| LEARNING RATE AT STEP s | $\max\{0.1\eta_0 + 0.9\eta_0 \cdot (1.0 - s/6e^5), 0.1\eta_0\}$ |
| TOTAL GRADIENT STEPS | $300k$ |
| BATCH SIZE | 8 FOR CLE-MV, CLE-AUG, 16 FOR GQN-JACO, 12 FOR SHEP7 |

* GQN SCHEDULER WITH A FASTER ATTENUATION RATE

Table 3: Model State Space Specifications

| TYPE | CLE-MV | CLE-AUG | GQN-JACO | SHEP7 |
|--|--------|---------|----------|-------|
| Z_DIMS | 16 | 16 | 32 | 16 |
| V_DIMS | 3 | 3 | 7 | 7 |
| Z_DIMS: THE DIMENSION OF A LATENT REPRESENTATION | | | | |
| V_DIMS: THE DIMENSION OF A VIEWPOINT VECTOR | | | | |

Model architecture We show our model configurations in Table 3, 4, and 5.

Table 4: MulMON Refinement Network with Trainable Parameters Φ

| Parameters | Type | Channels (out) | Activations. | Descriptions |
|------------|-------------------|-----------------|--------------|--------------------------------------|
| Φ | Input | 17 | | * Auxiliary inputs $\mathbf{a}(x^t)$ |
| | Conv 3×3 | 32 | Relu | |
| | Conv 3×3 | 32 | Relu | |
| | Conv 3×3 | 64 | Relu | |
| | Conv 3×3 | 64 | Relu | |
| | Flatten | | | |
| | Linear | 256 | Relu | |
| | Linear | 128 | Linear | |
| | Concat | $128+4*z_dims$ | | |
| | LSTMCell | 128 | | |
| | Linear | 128 | Linear | output $\Delta\lambda$ |

z_dims : the dimension of a latent representation

v_dims : the dimension of a viewpoint vector

Stride= 1 set for all Convs.

* see IODINE[9] for details

LSTMCell channels: the dimensions of the hidden states

Decoder-output processing For a single view of a scene, our decoder g_θ outputs $K \times H \times W$ RGB values (i.e. $\{x_k\}$ as in equation 2 of the main paper) along with $K \times H \times W$ mask logits (denoted as $\{\hat{m}_k\}$). H and W are the same as the image sizes, i.e. height and width. In this section, we detail the computation of rendering K individual scene components’ images, segmentation masks, and reconstructed scene images. We compute the individual scene objects’ images as:

$$x_k \stackrel{k}{\leftarrow} \text{sigmoid}(\hat{m}_k) \cdot x_k.$$

As shown in Figure 8, this overcomes mutual occlusions of the objects since the **sigmoid** functions do not impose any dependence on K objects. We compute the segmentation masks as:

$$m_k \stackrel{k}{\leftarrow} \text{softmax}_k(\hat{m}_k).$$

Table 5: MulMON Decoder with Trainable Parameters θ

| Parameters | Type | Channels (out) | Activations. | Descriptions |
|-------------------------------|-------------------|---------------------|--------------|---|
| θ^1 (view transformer) | Input | $z_dims + v_dims$ | | $z_k \sim \mathcal{N}(z_k; \lambda_k), v$ |
| | Linear | 512 | Relu | |
| | Linear | z_dims | Linear | $\tilde{z}_k = f_{\theta_1}(z_k, v)$ |
| θ^2 (Generator) | Input | z_dims | | $\tilde{z}_k = f_{\theta_1}(z_k, v)$ |
| | Broadcast | $z_dims + 2$ | | * Broadcast to grid |
| | Conv 3×3 | 32 | Relu | |
| | Conv 3×3 | 32 | Relu | |
| | Conv 3×3 | 32 | Relu | |
| | Conv 3×3 | 4 | Linear | rgb mean (μ_{xk}) + mask logits (\hat{m}_k) |

z_dims : the dimension of a latent representation
 v_dims : the dimension of a viewpoint vector
*: see spatial broadcast decoder [23]
Stride= 1 set for all Convs.




Figure 8: Example of images of post-processed decoder outputs: (left to right) (predicted) scene reconstruction, generated K individual scene components (white background for visual clarity), segmentation map. The generated scene components overcome/impute occlusions (e.g. the purple glossy sphere).

To generate binary segmentation masks, we take **argmax** operation over the K \hat{m}_k at every pixel location and encode the maximum indicator (indices) using one-hot codes. We render a scene image using a composition of all scene objects as:

$$\begin{aligned}
 x &= \sum_k \mathbf{softmax}_k(\hat{m}_k) \cdot x_k \\
 &= \sum_k m_k \cdot x_k.
 \end{aligned}$$

D. Additional Results

| Models | Disent. | Compl. | Inform. |
|--------|-------------|-------------|-------------|
| GQN | N/A | N/A | N/A |
| IODINE | 0.54 | 0.48 | 0.21 |
| MulMON | 0.63 | 0.54 | 0.58 |

Table 6: Disentanglement Analysis (CLE-Aug)

D.1 Disentanglement Analysis

To compare quantitatively the intra-object disentanglement achieved by MulMON and IODINE, we employ the framework and metrics (DCI) of Eastwood and Williams[7]. Specifically, let \mathbf{y} be the ground-truth generative factors for a single-view of a single object in a single scene, and let \mathbf{z} be the corresponding learned representation. Following [7], we learn a mapping from $Z = (\mathbf{z}_1, \mathbf{z}_2, \dots)$ to $Y = (\mathbf{y}_1, \mathbf{y}_2, \dots)$ with random forests in order to quantify the disentanglement, completeness and informativeness of the learned object representations. Section 5.3 presented the results on the CLE-MV dataset, and here we present the results on the CLE-Aug dataset. As shown in Table 6, MulMON again outperforms IODINE, learning representations that are more disentangled, complete and informative (about ground-truth factor

values). It is worth noting the significant gap in informativeness ($1 - NRMSE$) in Table 6. This strongly indicates that the object representations learned by MulMON are more accurate, i.e. they better-capture object properties.

D.2 Generalization Results

To evaluate MulMON’s generalization ability, we trained MulMON, IODINE and GQN on CLE-Aug. Then, we compared their performance on CLE-MV and 2 new datasets—Black-Aug and UnseenShape (see Figure 9). Black-Aug contains the CLE-Aug objects but only in single, unseen colour (black). This tests the models’ ability perform segmentation without colour differences/cues. UnseenShape contains only novel objects that are not in the CLE-Aug dataset—cups, cars, spheres, and diamonds. This directly tests generalization capabilities. Both datasets contain 30 scenes, each with 10 views. Table 1 shows that 1) all models generalize well to novel scenes, 2) MulMON still performs best for all tasks but observation prediction—where GQN does slightly better due to its more direct prediction procedure (features \rightarrow layout vs. features \rightarrow objects \rightarrow layout), 3) MulMON can indeed understand the composition of novel objects in novel scenes—impressive novel-view predictions (observations and segmentations) and disentanglement. Despite the excellent quantitative performance achieved by MulMON in generalization, we discovered that MulMON tended to decompose some objects, e.g. cars, into pieces (see Figure 10). Future investigations are thus needed in order to enable MulMON to generalize to more complex objects.




Figure 9: Samples from novel-scene datasets.

Table 7: MulMON’s generalization performance.

| Tasks | Models | CLE-Aug (train) | CLE-MV | Black-Aug | UnseenShape |
|--------------------|--------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|
| Seg. (mIoU) | IODINE | 0.51 ± 0.001 | 0.61 ± 0.002 | 0.50 ± 0.006 | 0.51 ± 0.004 |
| | MulMON | 0.71 ± 0.000 | 0.71 ± 0.004 | 0.67 ± 0.002 | 0.64 ± 0.004 |
| Pred.Obs (RMSE) | GQN | 0.15 ± 0.000 | 0.15 ± 0.001 | 0.24 ± 0.003 | 0.17 ± 0.002 |
| | MulMON | 0.07 ± 0.000 | 0.16 ± 0.002 | 0.26 ± 0.002 | 0.21 ± 0.006 |
| Disent. (D,C,I) | IODINE | 0.54, 0.48, 0.21 | 0.14, 0.12, 0.26 | 0.2, 0.26, 0.27 | 0.13, 0.12, 0.26 |
| | MulMON | 0.63, 0.54, 0.68 | 0.52, 0.48, 0.63 | 0.55, 0.55, 0.66 | 0.5, 0.47, 0.67 |
| Pred.Seg (mIoU) | MulMON | 0.69 ± 0.001 | 0.71 ± 0.004 | 0.68 ± 0.005 | 0.60 ± 0.005 |

D.3 Ablation Study

Prediction performance vs. the number of observations T In Section 5.4 of the main paper, we show that the spatial uncertainty MulMON learns decrease as more observations (T) are acquired. Here we study the effect of T on MulMON’s task performance, i.e. novel-viewpoint prediction. We employ mIoU (mean intersection-over-union) and RMSE (root-mean-square error) to measure MulMON’s performance on observation prediction and segmentation prediction respectively. In Figure 11, we show that the spatial uncertainty reduction (Left) suggests boosts of task performance (Right). This means MulMON does leverage multi-view exploration to learn more accurate scene representations (than a single-view configuration), this also explains the performance gap (see Section 5.1 in the main paper) between IODINE and MulMON. To further demonstrate the advantage that MulMON has over both IODINE and GQN, we compare their performance in terms of both segmentation and novel-view appearance prediction, as a function of the number of observations given to the models. Figure 12 shows that; 1) MulMON significantly outperforms IODINE even with a single view, likely due to a superior 3D scene understanding gained during training (figures on the left), 2) Despite the more difficult task of achieving object-level segmentation, MulMON closely mirrors the performance of GQN in predicting the appearance of the scene from unobserved viewpoints (figures on the right), 3) MulMON achieves similar performance in scene segmentation from observed and unobserved viewpoints, with any difference diminishes as the




Figure 10: Failure cases of MulMON in generalization: splitting an car into pieces. Here the MulMON is trained on the CLE-MV dataset and test on the UnseenShape data.

number of views increase (see dashed lines vs. solid lines in the left-hand figures).

Effect the number of object slots K Although explicit assumption about the number of objects in a scene is not required for MulMON, selecting a appropriate K (i.e. the number of object slots) is crucial to have MulMON work correctly. In the main paper, we discussed that “ K needs to be sufficiently larger than the number of scene objects” and we show the experimental support here. We train our model on CLE-MV, where each scene contains 4 to 7 objects including the background, using $K = 9$ and run tests on novel-viewpoint prediction tasks using various K . Figure 13 shows that, for both observation prediction and segmentation prediction tasks, the model’s performance improves as K increases until reaching a critical point at $K = 7$, which is the maximum quantity of scene objects in the dataset. Therefore, one should select a K that is always greater or equal to the maximum number of objects in a scene. When this condition is satisfied, further increase K will mostly not affect MulMON’s performance.

Subtle cases in terms of K ’s selection do exist. As shown in Figure, instead of treating the Shep7 scene a combination of a single object and the background, MulMON performs as a part segmenter that discovers the small cubes and their spatial composition. This is because, in the training phase of MulMON, the amortized parameters Φ and θ are trained to capture the object features (possibly disentangled) shared across all the objects in the whole dataset instead of each scene with specific objects. These shared object features are what drives the segmentation process of MulMON. In Shep7, what is being shared are the cubes, the object itself is a spatial composition of the cubes. The results on Shep7 along with the results shown in Figure 10 illustrate the subjectiveness of perceptual grouping and leaves much space for us to study in the future.

Effect of the IG coefficient in scene learning In the MulMON ELBO, we fix the coefficient of the information gain at 1. In testing, we consider this coefficient controls the scene learning rate (uncertainty reduction rate). We denote the coefficient as α_{IG} hereafter. According to the MulMON ELBO (to maximize), the negative sign of the **IG** term suggests that greater value of the coefficient leads to less information gain (spatial exploration). To verify this, we try four different α_{IG} (0.1, 1.0, 10.0 and 100.0) and track the prediction uncertainty as observations are acquired (same as our ablation study of T). The results in Figure 15 verifies our assumption the scene learning rate: larger α_{IG} leads to slower scene learning and vice versa.

D.4 Random Scene Generation

As a generative model, MulMON can generate random scenes by composing independently-sampled objects. However, to focus on forming accurate, disentangled representations of multi-object scenes, we must assume objects are i.i.d. and thus ignore inter-object correlations—e.g. two objects can appear at the same location. Figure 16 shows some random scene examples generated by MulMON (trained on

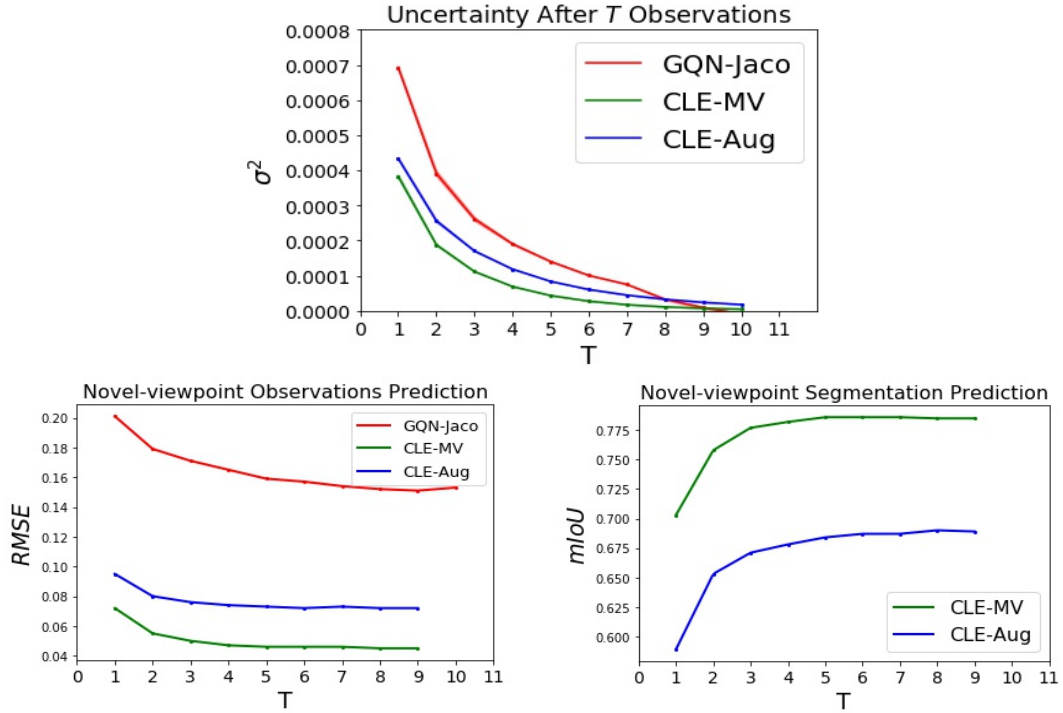


Figure 11: Effect of T : as more observations are acquired, (Top) the spatial uncertainty reduces and the performance of novel-viewpoint prediction on observation (Bottom left) and segmentation (Bottom right) prediction boosts.

the CLE-MV dataset). We can see that MulMON generates mostly good object samples by randomly composing different features but does not take into account the objects' spatial correlations.




Figure 12: Performance comparison w.r.t. a different number of observations T . (Top left) Segmentation performance vs. number of observations T on CLE-MV dataset. Note that “obs” means that MulMON reconstructs the observed images (scene appearances) and “unobs” means that MulMON predicts the appearance of the scene from unobserved viewpoints. (Top right) Novel-viewpoint observation prediction performance vs. number of observations given to the models on CLE-MV dataset. (Bottom left) Segmentation performance vs. number of observations T on CLE-MV dataset. (Bottom right) RMSE of appearance predictions for unobserved viewpoints vs. number of observations on CLE-Aug dataset.




Figure 13: Effect of K : (Left) the spatial uncertainty decreases and (Middle & Right) the task performance (novel-viewpoint prediction) boosts.




Figure 14: MulMON on Shep7. MulMON treats an Shep7 object as composition of parts (cubes) instead of one object.




Figure 15: Scene learning rate vs. **IG** coefficient (denoted as α_{IG}). (Left) Uncertainty reduction gets slower when increase increase α_{IG} . (Right) The computed uncertainty change rate or scene learning rate (lower means slower) shows larger α_{IG} slows the scene learning.




Figure 16: Random scene generation samples.