

APEX: Unsupervised, Object-Centric Scene Segmentation and Tracking for Robot Manipulation

Yizhe Wu¹, Oiwi Parker Jones¹, Martin Engelcke¹, and Ingmar Posner¹

Abstract—Recent advances in unsupervised learning for object detection, segmentation, and tracking hold significant promise for applications in robotics. A common approach is to frame these tasks as inference in probabilistic latent-variable models. In this paper, however, we show that the current state-of-the-art struggles with visually complex scenes such as typically encountered in robot manipulation tasks. We propose APEX, a new latent-variable model which is able to segment and track objects in more realistic scenes featuring objects that vary widely in size and texture, including the robot arm itself. This is achieved by a principled mask normalisation algorithm and a high-resolution scene encoder. To evaluate our approach, we present results on the real-world Sketchy dataset. This dataset, however, does not contain ground truth masks and object IDs for a quantitative evaluation. We thus introduce the Panda Pushing Dataset (P2D) which shows a Panda arm interacting with objects on a table in simulation and which includes ground-truth segmentation masks and object IDs for tracking. In both cases, APEX comprehensively outperforms the current state-of-the-art in unsupervised object segmentation and tracking. We demonstrate the efficacy of our segmentations for robot skill execution on an object arrangement task, where we also achieve the best or comparable performance among all the baselines.

I. INTRODUCTION

Scene segmentation and tracking are cornerstones of robotics (e.g. [1], [2]). A principal motivation is the ability to ground sensory observations in state-representations suitable for executing downstream tasks. While considerable advances have been made in using supervised methods for detecting and segmenting objects (e.g. [3], [4]), labelling data is resource intensive and quickly becomes intractable when trying to consider every possible object category for every deployment eventuality. Unsupervised learning – the discovery of representations suitable for task execution without the need for training labels – is therefore emerging as a promising alternative. In particular, recently developed *object-centric* scene representations (e.g. [5], [6], [7], [8], [9], [10]) have the potential of vastly improving data efficiency in robotics and other applications (e.g. [11], [12], [13]). Here we show, however, that current state-of-the-art methods fail on visually complex datasets that are representative of scenarios commonly encountered in robot manipulation.

Within the class of unsupervised, object-centric models, variational autoencoders (VAEs) ([14], [15]) are emerging as a popular choice. Object detection and segmentation are performed by using spatial transformer networks (STNs) [16] and spatial Gaussian mixture models (SGMMs) ([17], [18],

[19]) to separate objects, respectively. Some works do not impose any further structure on these latent representations (e.g. [5], [7], [11], [20]), while others further factorise the representations to explicitly disentangle foreground objects from the background, object location, appearance, and whether an object slot is used or not (e.g. [6], [9], [21], [22]). In contrast to, e.g., OP3 [11] which was also developed in the context of robotics, we argue the latter is of particular utility in robotics applications where such highly structured representations can be directly used as inputs to a planning module. A recently proposed model of this type called SCALOR [9] achieves particularly good results on segmenting and tracking objects in videos that contain a possibly large number of objects. We demonstrate however, that SCALOR struggles to learn object-centric representations on datasets with objects of widely varying sizes and textures as encountered in robot manipulation.

We therefore develop APEX (Amortised Parallel inference with miXture models), a novel object-centric generative model trained on videos. In contrast to prior work, APEX uses a principled mask normalisation procedure to parameterise an SGMM, which allows the explicit tuning of foreground and background standard deviations. APEX also features an improved scene encoder that outputs a high-resolution feature map that is particularly suitable for tracking [24]. To showcase the efficacy of APEX, we evaluate APEX qualitatively on Sketchy [23], an existing real-world robot manipulation dataset. We also introduce the Panda Pushing Dataset (P2D) as a quantitative benchmark against prior art. This contains videos of a Panda arm interacting with objects on a table in simulation and includes ground truth labels for segmentation masks and tracking objects between frames. It can be observed that APEX comprehensively outperforms recent state-of-the-art methods ([9], [11], [20], [25], [26]) by a large margin in terms of unsupervised segmentation and object tracking. Finally, inspired by the *Open Cloud Robot Table Organization Challenge* [27], we demonstrate the utility of APEX specifically in the context of a robot object re-arrangement task. Here, the superior segmentation performance of APEX, as illustrated in Figure 1, as well as the quality of the learned object representations lead to significantly better results compared to prior art.

II. RELATED WORK

This work builds on recent literature on object-centric generative models (OCGMs), which are typically formulated as VAEs (e.g. [21], [22]) or generative adversarial networks (GANs) (e.g. [28], [29], [30]). Unlike object-centric GANs,

¹Applied AI Lab, Oxford Robotics Institute, University of Oxford. Correspondence to: ywu@robots.ox.ac.uk.

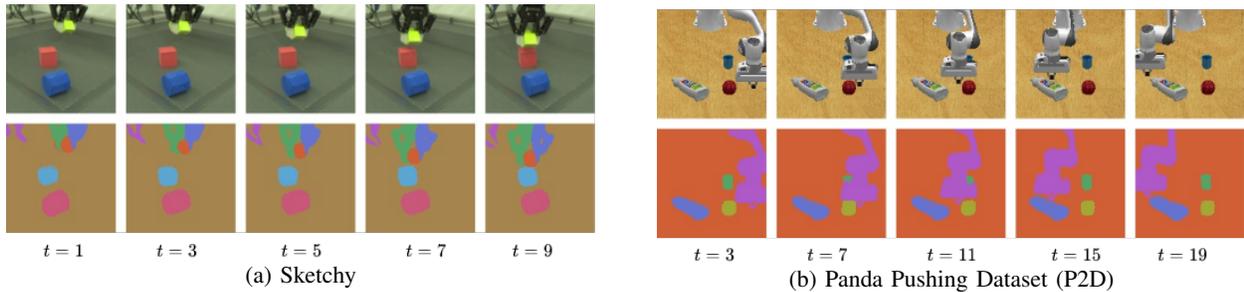


Fig. 1: APEX learns to segment and track objects in videos *without supervision* from (a) the established Sketchy dataset [23] and (b) our Panda Pushing Dataset (P2D). Segment colour indicates object ID. APEX accurately segments the diverse objects in the scenes and tracks the blue cup in (b) despite severe occlusion at $t = \{7, 11\}$.

VAE-based methods directly provide an amortised inference mechanism for extracting object-centric representations from input images. Early works that use STNs for separating objects do so by sequentially attending to different regions in an image, leading to a computational complexity that increases linearly with the number of objects ([21], [22]). More recent works parallelise the inference of object representations, which has been shown to be particularly useful for images with a large number of objects ([6], [9]). Segmentation-based models that parameterise an SGMM (e.g. [5], [7], [8], [20]) tend to be computationally more expensive than STN-based approaches where objects can be generated as smaller crops rather than image-sized components. A more informative pixel-level labelling, however, is required in applications such as the object arrangement task considered in this work. SLOT-ATTENTION [20] is a prominent recent model belonging to this category and is therefore selected as one of our baselines. The majority of related works learn object representations from individual images (e.g. [5], [7], [20]), but some also exploit temporal information in video sequences (e.g. [9], [10], [11]) to improve object separation. APEX also leverages a VAE, STNs, and an SGMM for the parallel inference of structured, object-centric latent representations from videos. The structure of the model is most closely related to SCALOR [9]. Instead of parameterising a Gaussian image likelihood, however, APEX parameterises an SGMM which allows direct tuning of loss magnitudes from background and foreground modules. Moreover, APEX uses a scene encoder that is particularly suitable for object tracking [31].

A small number of works also explore the use of object-centric representations in robotics ([11], [12]). OP3 [11] uses an object-centric model to predict goal images that contain a set of blocks in a desired configuration and the authors in [12] show that explicit object representations can accelerate the acquisition of robotic manipulation skills. Similarly, COBRA [13] leverages object-centric representations to improve data efficiency and policy robustness in several RL tasks in visually simple simulated environments. Inspired by [27], we benchmark a number of models on an object manipulation task, where we observe that segmentations and object representations learned by APEX lead to significantly better results compared to the recent state-of-the-art.

III. APEX: AMORTISED PARALLEL INFERENCE WITH MIXTURE MODELS

Let $\mathbf{x} \in [0, 1]^{H \times W \times C}$ be a frame from a video sequence $\mathbf{x}_{1:T}$, where H and W denote the image height and width, C represents the number of image channels (e.g. RGB), and T denotes the number of frames in the video sequence. Consider a scene \mathcal{S} to be formed of K object hypotheses – or *components* – which are each encoded by a set of latent variables such that $\mathcal{S} = \{\mathbf{z}_1, \dots, \mathbf{z}_K\}$. In particular, we consider two disjoint sets of object hypotheses comprised of *foreground* objects and a *background* component, such that $\mathcal{S} = \mathcal{O}^{\text{fg}} \cup \mathcal{O}^{\text{bg}}$, where $\mathcal{O}^{\text{fg}} = \{\mathbf{z}_1^{\text{fg}}, \dots, \mathbf{z}_{K-1}^{\text{fg}}\}$ and $\mathcal{O}^{\text{bg}} = \{\mathbf{z}_K^{\text{bg}}\}$, i.e., $\mathbf{z}_K^{\text{bg}} = \mathbf{z}_K$. Each foreground component is described by a set of latent variables $\mathbf{z}_k^{\text{fg}} = \{z_k^{\text{where}}, z_k^{\text{what}}, z_k^{\text{pres}}\}$ encoding its location, appearance, and existence in the scene (see [21]). \mathbf{z}_K^{bg} directly encodes the appearance of the background. For each frame \mathbf{x}_t at time t , foreground components can either be *propagated* from the previous time step or *discovered* in the current image (see [9], [10]).

APEX defines a generative model with learnable parameters θ that is formulated as an SGMM (e.g. [5], [7], [8]) via the image likelihood

$$p_\theta(\mathbf{x}_t | \mathbf{z}_{1:K,t}) = \sum_{k=1}^K \mathbf{m}_k(\mathbf{z}_{1:K,t}) \odot \mathcal{N}(\mu(\mathbf{z}_{k,t}), \sigma_k), \quad (1)$$

where $\mathbf{m}_{k,t}$ are the segmentation masks, $\mu_{k,t}$ are the means of the Gaussian components, and σ_k is a fixed component standard deviation. Separate standard deviations are used for the foreground and background, i.e. $\sigma_{1:K-1} = \sigma^{\text{fg}}$ and $\sigma_K = \sigma^{\text{bg}}$. A probabilistic prior $p_\theta(\mathbf{z}_{1:K,1:T})$ is defined to regularise the model.

To achieve segmentation and tracking, APEX provides an approximate inference model [32] with learnable parameters ϕ for an object-centric latent representation of how a scene evolves in a sequence of images, i.e.,

$$q_\phi(\mathbf{z}_{1:K,1:T} | \mathbf{x}_{1:T}) = \prod_{t=1}^T q_\phi(\mathbf{z}_{1:K,t} | \mathbf{z}_{1:K,<t}, \mathbf{x}_{\leq t}). \quad (2)$$

The inference and generative models are learned jointly as a VAE ([14], [15]). An overview of APEX is illustrated in

Figure 2. This section proceeds by first defining the underlying structure of the generative model, before defining the associated inference model that facilitates the segmentation and tracking of objects in video sequences.

A. Generative Model

To capture correlations between frames, we assume the latents $\mathbf{z}_{1:K,t}$ depend on the latents of the previous time steps

$$p_\theta(\mathbf{z}_{1:K,1:T}) = \prod_{t=1}^T p_\theta(\mathbf{z}_{1:K,t} | \mathbf{z}_{1:K,<t}). \quad (3)$$

$p_\theta(\mathbf{z}_{1:K,t})$ is further factorised into three terms to represent *propagated* objects \mathcal{O}_t^p , *discovered* objects \mathcal{O}_t^d , and the background $\mathcal{O}_t^{\text{bg}}$. Assuming there are P propagated objects and D newly discovered objects, we assume propagated objects depend on the previous latent variables, newly discovered objects in turn depend on objects that have already been propagated to avoid the *rediscovery* of propagated objects

$$p_\theta(\mathbf{z}_{1:K,t} | \mathbf{z}_{1:K,<t}) = p_\theta(\mathbf{z}_t^{\text{bg}} | \mathbf{z}_{<t}^{\text{bg}}) p_\theta(\mathbf{z}_{1:D,t}^d | \mathbf{z}_{1:P,t}^p) p_\theta(\mathbf{z}_{1:P,t}^p | \mathbf{z}_{1:K,<t}). \quad (4)$$

To facilitate application in robotics and similar to prior works (e.g. [9], [10]), the latents describing foreground objects in $\mathcal{O}_t^{\text{fg}}$ are factorised into $\{\mathbf{z}_k^{\text{where}}, \mathbf{z}_k^{\text{what}}, z_k^{\text{pres}}\}$ describing component location, appearance, and presence in the scene, respectively, such that

$$p_\theta(\mathbf{z}_{k,t}^p | \mathbf{z}_{k,<t}^{\text{fg}}) = p_\theta(z_{k,t}^{\text{pres}} | \mathbf{z}_{k,<t}^{\text{fg}}) p_\theta(\mathbf{z}_{k,t}^{\text{what}} | \mathbf{z}_{k,<t}^{\text{fg}})^{z_{k,t}^{\text{pres}}} p_\theta(\mathbf{z}_{k,t}^{\text{where}} | \mathbf{z}_{k,<t}^{\text{fg}})^{z_{k,t}^{\text{pres}}}, \quad (5)$$

$$p_\theta(\mathbf{z}_{k,t}^d | \mathbf{z}_{1:P,t}^p) = p_\theta(z_{k,t}^{\text{pres}} | \mathbf{z}_{1:P,t}^p) p_\theta(\mathbf{z}_{k,t}^{\text{what}} | \mathbf{z}_{k,t}^{\text{pres}})^{z_{k,t}^{\text{pres}}} p_\theta(\mathbf{z}_{k,t}^{\text{where}} | \mathbf{z}_{k,t}^{\text{pres}})^{z_{k,t}^{\text{pres}}}. \quad (6)$$

The segmentation masks \mathbf{m}_k and the means μ of the Gaussian components in Equation (1) are decoded from $\{\mathbf{z}_{k,t}^{\text{where}}, \mathbf{z}_{k,t}^{\text{what}}, z_{k,t}^{\text{pres}}\}$ and \mathbf{z}_t^{bg} . We consider $z_{k,t}^{\text{pres}}$ to be Bernoulli distributed and the remainder of the latents being Gaussian distributed (see [9], [25]). $\mathbf{z}_{k,t}^{\text{what}}$ encodes the appearance of a component and the logits $\alpha_{k,t}$ from which the mask $\mathbf{m}_{k,t}$ is later obtained. \mathbf{z}_t^{bg} only encodes the appearance of the background component. $\mathbf{z}_{k,t}^{\text{where}}$ encodes a transformation in an STN [16], describing the size and location of a bounding box that contains an object. A separate variable for encoding background location is not needed as every pixel that is not assigned to a foreground object is treated as background.

In contrast to SCALOR, where the foreground mask needs to be explicitly clamped to $[0, 1]$, we employ a principled approach to mask normalisation. In particular, we first introduce a foreground mask \mathbf{m}_t^{fg} to model the occupancy of the foreground as a whole, and then compute intermediate object masks $\hat{\mathbf{m}}_{1:K-1,t}$ which attribute specific occupancy responsibility to each individual object. Specifically, the mask logits $\alpha_{1:K-1,t}$ for the $K-1$ foreground objects are computed as

$$\alpha_{k,t} = c \tanh(\text{CNN}(\mathbf{z}_{k,t}^{\text{what}})), \quad (7)$$

where c is a fixed constant that constrains the mask logits to $[-c, c]$. The foreground mask \mathbf{m}_t^{fg} is computed as

$$\mathbf{m}_t^{\text{fg}} = \tanh\left(\sum_{k=1}^{K-1} \text{STN}\left(\text{softplus}(\alpha_{k,t}) z_{k,t}^{\text{pres}}, \mathbf{z}_{k,t}^{\text{where}}\right)\right), \quad (8)$$

where the softplus operation ensures that a possible foreground component makes a non-negative contribution to the foreground mask when $z_{k,t}^{\text{pres}} = 1$. The intermediate object masks $\hat{\mathbf{m}}_{1:K-1,t}$ are obtained such that

$$\hat{\alpha}_{k,t} = \text{STN}(\alpha_{k,t}, \mathbf{z}_{k,t}^{\text{where}}) + 2c z_{k,t}^{\text{pres}}, \quad (9)$$

$$\hat{\mathbf{m}}_{1:K-1,t} = \text{softmax}(\hat{\alpha}_{1,t}, \dots, \hat{\alpha}_{K-1,t}). \quad (10)$$

Equation (9) maps components with $z_{k,t}^{\text{pres}} = 1$ to an interval that does not overlap with components where $z_{k,t}^{\text{pres}} = 0$. This formalises the intuition that objects which are *not present* will not occupy any physical space. The masks for the $K-1$ foreground objects are then obtained by the element-wise multiplication of the foreground mask \mathbf{m}_t^{fg} with intermediate object masks $\hat{\mathbf{m}}_{1:K-1,t}$,

$$\mathbf{m}_{1:K-1,t} = \hat{\mathbf{m}}_{1:K-1,t} \odot \mathbf{m}_t^{\text{fg}}. \quad (11)$$

Finally, inspired by the stick-breaking process formulations in prior work ([5], [7]), the background mask \mathbf{m}_t^{bg} is computed as

$$\mathbf{m}_t^{\text{bg}} = 1 - \mathbf{m}_t^{\text{fg}}. \quad (12)$$

B. Inference Model

The true posterior over latent variables is generally intractable, so a variational approximation $q_\phi(\mathbf{z}_{1:K,1:T} | \mathbf{x}_{1:T})$ is introduced. Mirroring Equation (3) in the generative model, the approximate posterior is factorised as

$$q_\phi(\mathbf{z}_{1:K,1:T} | \mathbf{x}_{1:T}) = \prod_{t=1}^T q_\phi(\mathbf{z}_{1:K,t} | \mathbf{z}_{1:K,<t}, \mathbf{x}_{\leq t}). \quad (13)$$

The posterior at time step t exhibits the same dependencies as were assumed in Equation (4) and can thus be written as

$$q_\phi(\mathbf{z}_{1:K,t} | \mathbf{z}_{1:K,<t}, \mathbf{x}_{\leq t}) = q_\phi(\mathbf{z}_t^{\text{bg}} | \mathbf{z}_{1:D,t}^d, \mathbf{z}_{1:P,t}^p, \mathbf{x}_t)$$

$$q_\phi(\mathbf{z}_{1:D,t}^d | \mathbf{z}_{1:P,t}^p, \mathbf{x}_{\leq t}) q_\phi(\mathbf{z}_{1:P,t}^p | \mathbf{z}_{1:K,<t}, \mathbf{x}_{\leq t}). \quad (14)$$

We assume that the approximate posterior for the k^{th} propagated object in \mathcal{O}_t^p factorises such that

$$q_\phi(\mathbf{z}_{k,t}^p | \mathbf{z}_{k,<t}^{\text{fg}}, \mathbf{x}_{\leq t}) = q_\phi(z_{k,t}^{\text{pres}} | \mathbf{z}_{k,<t}^{\text{what}}, z_{k,<t}^{\text{pres}}, \mathbf{z}_{k,t-1}^{\text{where}}, \mathbf{x}_{\leq t})$$

$$q_\phi(\mathbf{z}_{k,t}^{\text{what}} | \mathbf{z}_{k,t}^{\text{where}}, \mathbf{z}_{k,<t}^{\text{what}}, \mathbf{x}_{\leq t})^{z_{k,t}^{\text{pres}}}$$

$$q_\phi(\mathbf{z}_{k,t}^{\text{where}} | \mathbf{z}_{k,<t}^{\text{where}}, \mathbf{z}_{k,<t}^{\text{what}}, \mathbf{x}_{\leq t})^{z_{k,t}^{\text{pres}}}$$

and that the posterior for a discovered object factorises as

$$q_\phi(\mathbf{z}_{k,t}^d | \mathbf{z}_{1:P,t}^p, \mathbf{x}_{\leq t}) = q_\phi(z_{k,t}^{\text{pres}} | \mathbf{z}_{k,t}^{\text{what}}, \mathbf{z}_{k,t}^{\text{where}}, \mathbf{z}_{1:P,t}^p, \mathbf{x}_{\leq t})$$

$$q_\phi(\mathbf{z}_{k,t}^{\text{what}} | \mathbf{z}_{k,t}^{\text{where}}, \mathbf{x}_{\leq t})^{z_{k,t}^{\text{pres}}} q_\phi(\mathbf{z}_{k,t}^{\text{where}} | \mathbf{x}_{\leq t})^{z_{k,t}^{\text{pres}}}. \quad (16)$$

Intuitively, $\mathbf{z}_{k,t}^{\text{where}}$ encodes where to look in an image in order to infer $\mathbf{z}_{k,t}^{\text{what}}$ which is used for determining the appearance of an object.

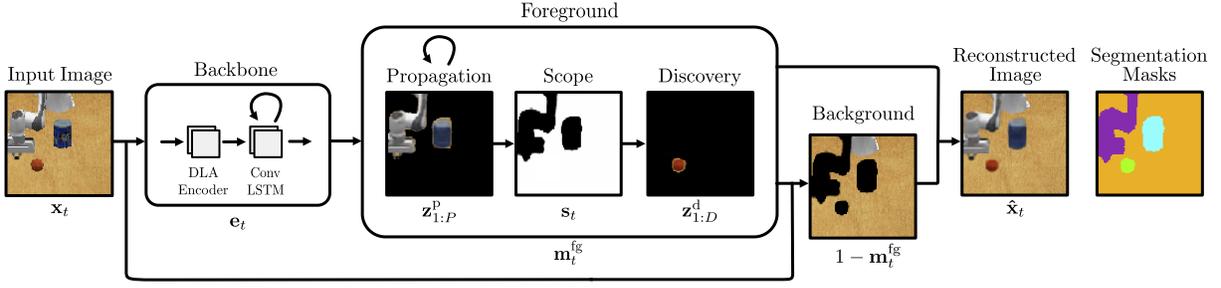


Fig. 2: Illustration of APEX. A backbone first extracts an encoding from an input image. This is used to propagate objects from the previous time step and to discover additional objects in the current image. The remainder of the image is treated as background. The final image is composed from both the foreground objects and the background.

C. Inference - Implementation

We now proceed with describing the implementation of how these posterior distributions are inferred for object *propagation* and *discovery* as well as the *background*.

Feature Extraction Features are extracted with a shared encoder. Observations $\mathbf{x}_{\leq t}$ are stacked and encoded into a feature map $\mathbf{e}_t \in \mathbb{R}^{H/4 \times W/4 \times F}$ with a deep layer aggregation (DLA) encoder [33] and a ConvLSTM [34] to capture spatio-temporal correlations. F is the number of the feature channels. The encoder outputs a high-resolution feature map that preserves spatial correspondences between the features and input images. This is beneficial for object tracking [24] and in contrast to SCALOR, where a low resolution feature map is used instead. *Propagation*: Using the previous object bounding box at $t-1$, a square feature map $\mathbf{f}_{k,t}^p$ is extracted from \mathbf{e}_t using an STN, whereby the square dimensions are set equal to the larger side of the previous bounding box. A CNN then maps $\mathbf{f}_{k,t}^p$ to a feature vector $\mathbf{c}_{k,t}^p$. *Discovery*: A grid of equally spaced features $\mathbf{c}_{1:D,t}^d$ is extracted from the feature map \mathbf{e}_t to discover new objects.

Object Location *Propagation*: The posterior over $\mathbf{z}_{k,t}^{\text{where}}$ is computed with an RNN whose inputs are $[\mathbf{c}_{k,t}^p, \mathbf{z}_{k,t-1}^{\text{what}}, \mathbf{z}_{k,t-1}^{\text{where}}]$. *Discovery*: The posterior over $\mathbf{z}_{k,t}^{\text{where}}$ is computed from $\mathbf{c}_{k,t}^d$ with a 1×1 convolution.

Object Appearance *Propagation*: A *glimpse* $\mathbf{G}_{k,t}$ is cropped from the feature map \mathbf{e}_t using a STN according to $\mathbf{z}_{k,t}^{\text{where}}$. A RNN takes the encoding of the glimpse $\mathbf{G}_{k,t}$ encoded by a CNN and $\mathbf{z}_{k,t-1}^{\text{what}}$ as inputs to infer the posterior over $\mathbf{z}_{k,t}^{\text{what}}$. *Discovery*: The posterior over $\mathbf{z}_{k,t}^{\text{what}}$ is inferred in same fashion as for propagated objects with weights being shared and $\mathbf{z}_0^{\text{what}}$ being initialised to a vector of zeros. *Background*: The posterior over \mathbf{z}_t^{bg} is obtained with a CNN using the background mask \mathbf{m}_t^{bg} and the current image \mathbf{x}_t .

Object Presence *Propagation*: The posterior over $z_{k,t}^{\text{pres}}$ is computed with an RNN whose inputs are $[\mathbf{c}_{k,t}^p, \mathbf{z}_{k,t}^{\text{what}}, z_{k,t-1}^{\text{pres}}]$. *Discovery*: The presence of objects in the *discovery* phase is determined by: 1. whether a new object is detected and 2. whether that object has been explained by the propagated objects. We use the scope \mathbf{s}_t to indicate which pixels have not been explained yet. This scope is defined as $\mathbf{s}_t = 1 - \mathbf{m}_t^p$ whereby \mathbf{m}_t^p is computed as in Equation (8) but only using propagated objects. We thus decompose the

posterior over $z_{k,t}^{\text{pres}}$ into two terms $\{p_{k,t}^{\text{proposal}}, p_{k,t}^{\text{context}}\} \in [0, 1]$ so that

$$q_\phi(z_{k,t}^{\text{pres}} | \mathbf{z}_{k,t}^{\text{what}}, \mathbf{z}_{k,t}^{\text{where}}, \mathbf{z}_{1:P,t}^p, \mathbf{x}_{\leq t}) = p_{k,t}^{\text{proposal}} p_{k,t}^{\text{context}}. \quad (17)$$

$p_{k,t}^{\text{proposal}}$ is obtained with a fully-connected layer from $[\mathbf{c}_{k,t}^d, \mathbf{z}_{k,t}^{\text{what}}, \mathbf{z}_{k,t}^{\text{where}}]$ and describes whether a cell might contain a new object or not. The purpose of $p_{k,t}^{\text{context}}$ is to avoid the rediscovery of objects via the scope \mathbf{s}_t . This removes discovered objects that overlap with propagated objects and it is computed as

$$p_{k,t}^{\text{context}} = \frac{\sum_{i,j} \mathbf{s}_{t,i,j} \tanh(\text{softplus}(\alpha_{k,t,i,j}))}{\sum_{i,j} \tanh(\text{softplus}(\alpha_{k,t,i,j}))}, \quad (18)$$

where (i, j) are all pixel coordinate tuples in an image. Intuitively, $p_{k,t}^{\text{context}}$ corresponds to the fraction of the proposed object mask that has not been explained by the propagated objects. *Object Filtering*: To reduce memory requirements, foreground objects are discarded at every time step when $z_{k,t}^{\text{pres}}$ is below a fixed, manually set threshold.

D. Learning

The inference and generative models can be jointly trained by maximising the evidence lower bound (ELBO). Omitting object subscripts, this is given by

$$\mathcal{L}(\theta, \phi) = \sum_{t=1}^T \mathbb{E}_{q_\phi(\mathbf{z}_t | \mathbf{z}_{<t}, \mathbf{x}_{\leq t})} [\log p_\theta(\mathbf{x}_t | \mathbf{z}_t)] + \mathbb{KL}[q_\phi(\mathbf{z}_t | \mathbf{z}_{<t}, \mathbf{x}_{\leq t}) \parallel p_\theta(\mathbf{z}_t | \mathbf{z}_{<t})]. \quad (19)$$

Prior distributions for continuous and discrete variables are assumed to be Gaussian and Bernoulli, respectively. Continuous variables are reparameterised (see [14], [15]) and discrete variables are obtained using the Gumbel-Softmax trick [35]. In practice, we minimise the inclusive KL divergence [36] for the \mathbf{z}^{what} and $\mathbf{z}^{\text{where}}$ as we empirically find this leads to better performance. We also include an entropy loss on the $K-1$ foreground object masks in the training objective,

$$\mathcal{L}_H = \sum_{i=1}^H \sum_{j=1}^W \sum_{t=1}^T \sum_{k=1}^{K-1} -m_{t,i,j}^{\text{fg}} m_{k,t,i,j} \log m_{k,t,i,j}. \quad (21)$$

This penalises pixels being explained by multiple components. The full training objective is the sum of the ELBO

and the mask entropy loss:

$$\mathcal{L} = \mathcal{L}(\theta, \phi) + \mathcal{L}_H. \quad (22)$$

IV. EXPERIMENTS

This section presents experiments on unsupervised scene segmentation, object tracking, and a simulated object manipulation task to showcase the capabilities of APEX. Our recent, state-of-the-art baselines consist of SCALOR [9], OP3 [11], SLOT-ATTENTION [20], SPACE [25], and G-SWM [26]. **Datasets** We perform a qualitative evaluation using the real-world Sketchy dataset [23], which contains demonstration trajectories of a robot arm performing different tasks involving a set of objects. The images are pre-processed as in [37], using a 128×128 resolution and sequences of length 10. Sketchy, however, does not contain object annotations, which prohibits the quantitative evaluation of object segmentation and tracking methods. We therefore introduce the Panda Pushing Dataset (P2D), which shows a Panda arm interacting with objects in simulation and includes pixel-level ground truth segmentations as well as object tracking IDs. Up to three objects are spawned in each episode and the robot-arm moves along a randomly selected straight line in the horizontal direction. Objects in the dataset are sampled from a set of 14 common objects (e.g. mugs, coffee cans, apples) with varying shapes, colours, and textures (see [38]). We collect a total of 2,400 trajectories (2,000 for training, 200 for validation, and 200 for testing) with each trajectory having a length of 20 frames with a resolution of 128×128 .

Metrics Similar to prior work (e.g. [7], [8]), the quality of object segmentations is evaluated using the Adjusted Rand Index (ARI) [39] and the Mean Segmentation Covering (MSC) on a held-out test set. To enable rigorous benchmarking, we report results on two variants of these metrics: one which only considers foreground objects (see [7]) as well as one where all pixels and ground truth masks are considered, including those belonging to the background. The latter is relevant in the context of this work as the background needs to be explicitly separated from the foreground for the manipulation task in Section IV-C. Multi-object tracking metrics are evaluated following the procedure in Weis et al. [40], which is based on the protocol from the established MOT16 tracking benchmark [41].

Implementation Details APEX is trained with ADAM [42] and a learning rate of 10^{-4} . The baselines are trained with the default learning rates and optimisers from the released implementations. The number of components in OP3 and SLOT-ATTENTION is set to $K = 5$ for P2D and to $K = 8$ for Sketchy. For SCALOR, the hard constraint on object size as found in the original implementation is removed as P2D and Sketchy contain objects of various sizes. APEX, SCALOR, and OP3 are trained with a batch size of 4 for 4×10^4 iterations. An additional 10^4 warm-up iterations are used for G-SWM. With the exception of OP3, the models are trained on the full image sequences. OP3 is trained on sub-sequences of length five due to the model capacity limitations on images that are resized to 64×64 as used in the original model. Training

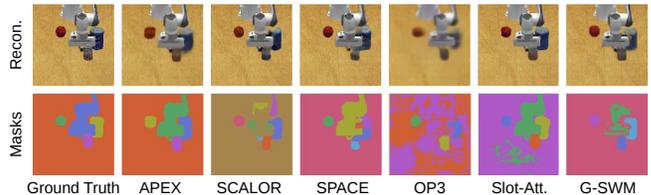


Fig. 3: Scene segmentation results on P2D. APEX achieves the qualitatively best results, cleanly segmenting all foreground objects as well as the Panda arm.

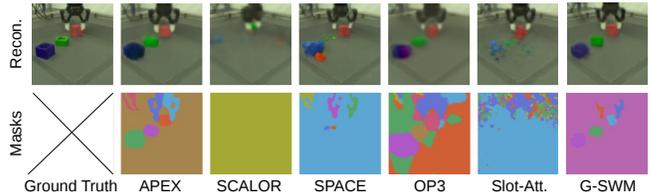


Fig. 4: Scene segmentation results on Sketchy. In contrast to the baselines, APEX manages to cleanly segment all objects even on this more challenging real-world dataset.

APEX, SCALOR, G-SWM and OP3 on a single NVIDIA Titan RTX GPU takes about 20 hours each. For SPACE, the batch size is increased to 16 and the number of training iterations are increased to 2×10^5 and 1×10^5 iterations for P2D and Sketchy, respectively, to account for the fact that SPACE is trained on individual images rather than image sequences. For SLOT-ATTENTION the number of training iterations is further increased to 4×10^5 and 2.5×10^5 iterations for P2D and Sketchy, respectively, as we found that the models take longer to learn reasonable segmentation masks. The resulting wall-clock times for SPACE and SLOT-ATTENTION are about 5 hours and 20 hours, respectively.

A. Unsupervised Segmentation and Tracking

Quantitative results for unsupervised segmentation and tracking on P2D are summarised in Table I and Table II, respectively. For both tasks, it can be seen that APEX outperforms the baselines by significant margins. We also observe a smaller standard deviation of the scores for APEX, indicating that APEX is more stable than the baselines. We attribute these improvements to the principled mask normalisation, the high-resolution feature maps, and the mask entropy loss. Qualitative tracking results for APEX are shown in Figure 1. It can be seen that APEX is clearly able to track individual objects even through occlusions, thus corroborating the quantitative results in Table II.

Qualitative segmentation results for APEX as well as for the baselines are shown in Figures 3 and 4 on P2D and Sketchy, respectively. In Figure 3, a fairly cluttered scene from P2D can be seen where the robot arm is interacting with two objects in close proximity to each other. While APEX clearly segments the foreground objects and the arm itself – a key prerequisite for task planning and control – the baselines struggle to do so. SCALOR is unable to accurately segment the arm with part of it being captured by the background

TABLE I: Mean and standard deviation of the segmentation metrics on P2D from four random seeds.

	Object tracking	ARI-FG	MSC-FG	ARI	MSC
OP3	✓	0.30 ± 0.03	0.30 ± 0.01	0.15 ± 0.05	0.39 ± 0.03
SCALOR	✓	0.31 ± 0.03	0.37 ± 0.01	0.43 ± 0.04	0.52 ± 0.01
G-SWM *	✓	0.38 ± 0.17	0.42 ± 0.05	0.33 ± 0.13	0.55 ± 0.04
SPACE	✗	0.33 ± 0.17	0.49 ± 0.12	0.72 ± 0.21	0.62 ± 0.09
SLOT-ATT.	✗	0.77 ± 0.24	0.46 ± 0.16	0.25 ± 0.22	0.50 ± 0.14
APEX	✓	0.89 ± 0.02	0.73 ± 0.01	0.93 ± 0.00	0.80 ± 0.01

* The G-SWM results are computed with one failed random seed being excluded.

TABLE II: Mean and standard deviation of the tracking metrics on P2D from four random seeds.

	MOTA ↑	MOTP ↑	Match ↑	ID S. ↓	Fps ↓	Miss ↓	MD ↑	MT ↑
OP3	-48.6 ± 17.0	66.6 ± 0.9	25.7 ± 4.4	0.3 ± 0.1	74.4 ± 13.4	73.9 ± 4.5	15.2 ± 4.7	15.6 ± 4.8
SCALOR	-125.6 ± 6.3	73.5 ± 0.5	31.6 ± 1.1	2.6 ± 0.0	157.2 ± 6.5	65.8 ± 1.1	19.0 ± 1.7	22.3 ± 2.3
G-SWM *	-41.8 ± 2.0	73.1 ± 0.5	36.7 ± 4.0	2.0 ± 1.0	78.4 ± 2.5	61.3 ± 3.5	25.3 ± 4.6	28.6 ± 3.2
APEX	50.5 ± 4.0	83.2 ± 0.6	79.7 ± 0.5	0.2 ± 0.1	29.2 ± 3.8	20.0 ± 0.5	70.8 ± 1.0	71.2 ± 0.7

* The G-SWM results are computed with one failed random seed being excluded.

module. We conjecture that this is due to the Gaussian image likelihood used for training SCALOR, as this does not untie the standard deviation of the foreground and the background likelihood as in APEX. A similar problem is observed with G-SWM which uses the same image likelihood modelling as SCALOR. For APEX and SPACE, in contrast, a smaller standard deviation is used for the background likelihood (0.04) than that of the foreground (0.1) to prevent the background module from also capturing the foreground objects. A Gaussian likelihood with a smaller standard deviation leads to a sharper distribution and thus results in a very low likelihood when the reconstructed colour deviates even slightly from the target. This will force the background module to focus on the more uniform background pixels which are easier to reconstruct compared to the foreground pixels. This is further examined in Section IV-B. While SPACE manages to segment the robot arm from the background, it fails to segment the small objects accurately. Unlike APEX, SPACE operates on static images and can thus not leverage temporal information. We argue that this helps APEX to learn to distinguish objects even when they are physically very close to each other in one frame, as they might move relative to each other in other frames. Both OP3 and SLOT-ATTENTION are able to roughly segment the objects but the segmentation results are noisy and inaccurate.

In Figure 4, APEX accurately segments all the objects in the scene – even the cables of the manipulator. We attribute the segmentation of the robot manipulator into left and right grippers to the relative motion (open/close) of the two gripper handles. In contrast, SCALOR fails to segment the objects and while SPACE is able to segment the arm, it omits the other foreground objects. This might be caused by the fact that the objects have a uniform colour and are therefore easily reconstructed by the background module. G-SWM successfully segments the objects, but part of the arm is treated as background. OP3 and SLOT-ATTENTION struggle to predict accurate masks.

B. Ablation Study

A set of ablation experiments is performed to validate the efficacy of the key design choices that set apart APEX from prior art: better scene encoding, principled mask normalisation, and a mask entropy loss. The results are summarised in Table III. Rather than re-using features extracted by the backbone with STNs when inferring object latents, it can be observed that using STNs to extract information from the input images, such as in SCALOR, performs consistently worse. Re-using the backbone features also facilitates a reduction in model parameters and computation. The foreground scores when training without the additional mask entropy loss are also smaller, indicating that the entropy loss is indeed beneficial for learning to disambiguate foreground objects. It can be seen that mimicking the normalisation scheme from SCALOR results in lower scores compared to the mask normalisation introduced in Section III. For the former, the STN outputs a mask $\alpha_k^s \in [0, 1]$ for each component and the masks are normalised according to $(\alpha_k^s)^2 / \sum_k \alpha_k^s$, which is numerically less stable than a softmax. This might also explain the increased standard deviation of the scores, especially for the foreground ARI. Finally, we compare the use of APEX’s SGMM image likelihood formulation to using a standard Gaussian likelihood. Both variants of APEX with a standard Gaussian likelihood fail to learn object-centric scene decompositions, highlighting the benefit of the SGMM formulation with separate standard deviations for foreground and background modules.

C. Object Arrangement Task

Generative models can be used to learn concise and informative representations that can be used in downstream tasks. In contrast to methods where a single latent vector encapsulates all object-relevant information (e.g. [11], [20]), further factorising the information into $[z_{k,t}^{\text{pres}}, z_{k,t}^{\text{where}}, z_{k,t}^{\text{what}}]$ as well as foreground and background latents allows us to directly use these representations as inputs to a controller.

We demonstrate this in an object arrangement task where a robot is required to pick and place objects on a table

TABLE III: Mean and standard deviation of the segmentation metrics on P2D from four random seeds.

	ARI-FG	MSC-FG	ARI	MSC
APEX	0.89 ± 0.02	0.73 ± 0.01	0.93 ± 0.00	0.80 ± 0.01
Image space STN	0.71 ± 0.27	0.62 ± 0.10	0.91 ± 0.01	0.71 ± 0.08
No entropy loss	0.63 ± 0.18	0.63 ± 0.07	0.92 ± 0.01	0.73 ± 0.05
SCALOR-norm	0.62 ± 0.30	0.60 ± 0.13	0.90 ± 0.03	0.70 ± 0.09
Gaussian likelihood *	0.10 ± 0.01	0.06 ± 0.00	0.00 ± 0.00	0.28 ± 0.00
Gaussian likelihood & SCALOR-norm. *	0.02 ± 0.00	0.06 ± 0.00	0.00 ± 0.00	0.28 ± 0.00

* These models consistently fail to meaningfully segment the images.

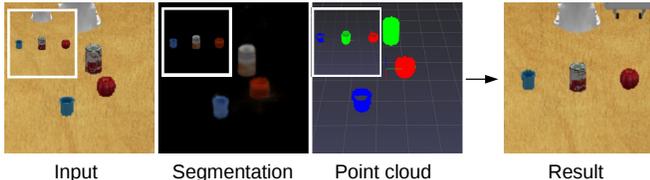


Fig. 5: Object arrangement task illustration. Given an input of the current scene and a goal image (a), objects are segmented (b) and matched (c). Depth information is used to obtain the 3D locations and shapes of all objects, which are used as inputs to a heuristic control policy that tries to move the objects to the desired locations specified by the target image.

according to a *goal image*. This image is first parsed into foreground objects and background. Assuming depth information is available, e.g. via a stereo camera, the 3D location and shape of an object are obtained by filtering a depth image according to the object mask. The current scene is processed in the same fashion, which allows the current objects to be matched to the associated objects in the goal image according to the smallest \mathcal{L}_2 distance in terms of the objects’ appearances encoded by z_k^{what} . This is facilitated by discarding *empty detections* according to z_k^{pres} and the explicit separation of foreground and background components. The current and desired location and shape of each object serve as inputs to a heuristic control policy that executes a sequence of sub-tasks to move the objects. This is illustrated in Figure 5. To avoid collisions, a check for whether the target location is occupied by other objects is conducted before executing a sub-task. If all target locations are occupied, one of the unsorted objects is moved to an edge of the table to make space for the other objects. This object will be moved to its target space at the end of the re-arrangement process.

Three object arrangement tasks are created involving two to four objects on a table. For each task, we create 150 test scenarios with the same group of objects spawned in different locations. The performance of APEX is compared to SCALOR, SPACE and G-SWM. Comparisons against OP3 and SLOT-ATTENTION are omitted here as neither explicitly models foreground and background components which are required for task execution. Performance is quantified using the mean distance of the final object locations relative to the desired object locations. We set a fixed penalty of 1m for outliers where an object falls off the table.

TABLE IV: Mean and standard deviation of the object distance to the desired goal positions in metres.

	Two objects	Three objects	Four objects
SPACE	0.07 ± 0.16	0.23 ± 0.25	0.23 ± 0.20
SCALOR	0.22 ± 0.17	0.15 ± 0.17	0.18 ± 0.21
G-SWM	0.02 ± 0.09	0.04 ± 0.12	0.25 ± 0.21
APEX	0.04 ± 0.10	0.06 ± 0.13	0.09 ± 0.18

The results are summarised in Table IV. APEX performs better than SPACE and SCALOR on all three tasks. Some failures are caused by missing detections or incorrect object matches. G-SWM performs better than APEX on tasks with two or three objects but fails to segment the goal image properly on task with four objects. Although G-SWM outperforms APEX when there are fewer objects, APEX appears to pull ahead as the number of objects increases. APEX also performs better on segmenting the robot arm (fig. 3, fig. 4), which is excluded from this task. SPACE performs worst on tasks with three objects, but better than SCALOR for two objects. We find that SPACE tends to oversegment the objects, i.e., a single object is divided into several components, which reduces the efficacy of both object matching and location estimation. We hypothesise that the improvement of SCALOR compared to SPACE is facilitated by the propagation module which is also incorporated into APEX and aids the learning of higher quality segmentations.

V. CONCLUSIONS

This paper proposes APEX, a novel, object-centric generative model designed to provide state-of-the-art unsupervised object segmentation and tracking on datasets commonly encountered in robotics. APEX is evaluated on the established Sketchy dataset [23] for qualitative results and on a custom Panda Pushing Dataset (P2D) for both quantitative and qualitative results. We show that APEX comprehensively outperforms prior art in terms of segmentation and tracking by leveraging improved feature encoding modules as well as a principled normalisation scheme for object and background masks. Finally, we demonstrate the efficacy of the unsupervised object representations learned by APEX on a robot manipulation task that involves the rearrangement of several objects on a table. APEX outperforms most of the baselines due to consistently providing segmentations of significantly higher quality, leading to improvements in object matching and 3D shape extraction.

ACKNOWLEDGMENT

This work was supported by EPSRC Programme Grant (EP/V000748/1), an Amazon Research Award and the China Scholarship Council. The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work. <http://dx.doi.org/10.5281/zenodo.22558>.

REFERENCES

- [1] A. Geiger, P. Lenz, and R. Urtasun, "Are We Ready for Autonomous Driving? The KITTI Vision Benchmark Suite," *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [2] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The Cityscapes Dataset for Semantic Urban Scene Understanding," *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [3] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," *arXiv preprint arXiv:1506.01497*, 2015.
- [4] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," *International Conference on Computer Vision (ICCV)*, 2017.
- [5] 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.
- [6] E. Crawford and J. Pineau, "Spatially Invariant Unsupervised Object Detection with Convolutional Neural Networks," *AAAI Conference on Artificial Intelligence*, 2019.
- [7] M. Engelcke, A. R. Kosiorek, O. Parker Jones, and I. Posner, "GENESIS: Generative Scene Inference and Sampling with Object-Centric Latent Representations," *International Conference on Learning Representations (ICLR)*, 2020.
- [8] 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," *International Conference on Machine Learning (ICML)*, 2019.
- [9] J. Jiang, S. Janghorbani, G. De Melo, and S. Ahn, "SCALOR: Generative World Models with Scalable Object Representations," *International Conference on Learning Representations (ICLR)*, 2020.
- [10] A. Kosiorek, H. Kim, Y. W. Teh, and I. Posner, "Sequential Attend, Infer, Repeat: Generative Modelling of Moving Objects," *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [11] R. Veerapaneni, J. D. Co-Reyes, M. Chang, M. Janner, C. Finn, J. Wu, J. Tenenbaum, and S. Levine, "Entity Abstraction in Visual Model-Based Reinforcement Learning," *Conference on Robot Learning (CoRL)*, 2020.
- [12] M. Wulfmeier, A. Byravan, T. Hertweck, I. Higgins, A. Gupta, T. Kulkarni, M. Reynolds, D. Teplyashin, R. Hafner, T. Lampe *et al.*, "Representation Matters: Improving Perception and Exploration for Robotics," *arXiv preprint arXiv:2011.01758*, 2020.
- [13] N. Watters, L. Matthey, M. Bosnjak, C. P. Burgess, and A. Lerchner, "COBRA: Data-Efficient Model-Based RL through Unsupervised Object Discovery and Curiosity-Driven Exploration," *arXiv preprint arXiv:1905.09275*, 2019.
- [14] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *International Conference on Learning Representations (ICLR)*, 2014.
- [15] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic Backpropagation and Approximate Inference in Deep Generative Models," *International Conference on Machine Learning (ICML)*, 2014.
- [16] M. Jaderberg, K. Simonyan, A. Zisserman *et al.*, "Spatial Transformer Networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- [17] K. Greff, A. Rasmus, M. Berglund, T. Hao, H. Valpola, and J. Schmidhuber, "Tagger: Deep Unsupervised Perceptual Grouping," *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [18] K. Greff, S. van Steenkiste, and J. Schmidhuber, "Neural Expectation Maximization," *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [19] S. van Steenkiste, M. Chang, K. Greff, and J. Schmidhuber, "Relational Neural Expectation Maximization: Unsupervised Discovery of Objects and their Interactions," *arXiv preprint arXiv:1802.10353*, 2018.
- [20] F. Locatello, D. Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, A. Dosovitskiy, and T. Kipf, "Object-Centric Learning with Slot Attention," *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [21] S. A. Eslami, N. Heess, T. Weber, Y. Tassa, D. Szepesvari, G. E. Hinton *et al.*, "Attend, Infer, Repeat: Fast Scene Understanding with Generative Models," *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [22] J. Huang and K. Murphy, "Efficient Inference in Occlusion-Aware Generative Models of Images," *arXiv preprint arXiv:1511.06362*, 2015.
- [23] S. Cabi, S. Gómez Colmenarejo, A. Novikov, K. Konyushkova, S. Reed, R. Jeong, K. Zolna, Y. Aytar, D. Budden, M. Vecerik *et al.*, "Scaling Data-Driven Robotics with Reward Sketching and Batch Reinforcement Learning," *arXiv preprint arXiv:1909.12200*, 2019.
- [24] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as Points," *arXiv preprint arXiv:1904.07850*, 2019.
- [25] 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," *International Conference on Learning Representations (ICLR)*, 2020.
- [26] Z. Lin, Y.-F. Wu, S. Peri, B. Fu, J. Jiang, and S. Ahn, "Improving generative imagination in object-centric world models," in *International Conference on Machine Learning*. PMLR, 2020, pp. 6140–6149.
- [27] "IROC 2020: Open Cloud Robot Table Organization Challenge (OCROTC)," 2020. [Online]. Available: <http://www.ocrotc.org/#/>
- [28] S. Ehrhardt, O. Groth, A. Monszpart, M. Engelcke, I. Posner, N. Mitra, and A. Vedaldi, "RELATE: Physically Plausible Multi-Object Scene Synthesis Using Structured Latent Spaces," *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [29] T. Nguyen-Phuoc, C. Richardt, L. Mai, Y.-L. Yang, and N. Mitra, "BlockGAN: Learning 3D Object-aware Scene Representations from Unlabelled Images," *arXiv preprint arXiv:2002.08988*, 2020.
- [30] M. Niemeyer and A. Geiger, "GIRAFFE: Representing Scenes as Compositional Generative Neural Feature Fields," *arXiv preprint arXiv:2011.12100*, 2020.
- [31] X. Zhou, V. Koltun, and P. Krähenbühl, "Tracking Objects as Points," in *European Conference on Computer Vision (ECCV)*, 2020.
- [32] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, "An Introduction to Variational Methods for Graphical Models," *Machine learning*, vol. 37, no. 2, pp. 183–233, 1999.
- [33] F. Yu, D. Wang, E. Shelhamer, and T. Darrell, "Deep Layer Aggregation," *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [34] X. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo, "Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2015, pp. 802–810.
- [35] E. Jang, S. Gu, and B. Poole, "Categorical Reparameterization with Gumbel-Softmax," *International Conference on Learning Representations (ICLR)*, 2017.
- [36] Y. Li and Y. Gal, "Dropout inference in bayesian neural networks with alpha-divergences," in *International conference on machine learning*. PMLR, 2017, pp. 2052–2061.
- [37] M. Engelcke, O. Parker Jones, and I. Posner, "GENESIS-V2: Inferring Unordered Object Representations without Iterative Refinement," *arXiv preprint arXiv:2104.09958*, 2021.
- [38] B. Calli, A. Singh, J. Bruce, A. Walsman, K. Konolige, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Yale-CMU-Berkeley Dataset for Robotic Manipulation Research," *International Journal of Robotics Research (IJRR)*, vol. 36, no. 3, pp. 261–268, 2017.
- [39] W. M. Rand, "Objective Criteria for the Evaluation of Clustering Methods," *Journal of the American Statistical association*, vol. 66, no. 336, pp. 846–850, 1971.
- [40] M. A. Weis, K. Chitta, Y. Sharma, W. Brendel, M. Bethge, A. Geiger, and A. S. Ecker, "Unmasking the Inductive Biases of Unsupervised Object Representations for Video Sequences," *arXiv preprint arXiv:2006.07034*, 2020.
- [41] A. Milan, L. Leal-Taixé, I. Reid, S. Roth, and K. Schindler, "MOT16: A Benchmark for Multi-Object Tracking," *arXiv preprint arXiv:1603.00831*, 2016.
- [42] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *International Conference on Learning Representations (ICLR)*, 2015.