

# Shepherding Slots to Objects: Towards Stable and Robust Object-Centric Learning

Jinwoo Kim<sup>1\*</sup> Janghyuk Choi<sup>2\*</sup> Ho-Jin Choi<sup>2</sup> Seon Joo Kim<sup>1</sup>  
<sup>1</sup>Yonsei University <sup>2</sup>KAIST

{jinwoo-kim, seonjookim}@yonsei.ac.kr  
 {janghyuk.choi, hojinc}@kaist.ac.kr

## Abstract

*Object-centric learning (OCL) aspires general and compositional understanding of scenes by representing a scene as a collection of object-centric representations. OCL has also been extended to multi-view image and video datasets to apply various data-driven inductive biases by utilizing geometric or temporal information in the multi-image data. Single-view images carry less information about how to disentangle a given scene than videos or multi-view images do. Hence, owing to the difficulty of applying inductive biases, OCL for single-view images remains challenging, resulting in inconsistent learning of object-centric representation. To this end, we introduce a novel OCL framework for single-view images, SLoT Attention via SHepherding (SLASH), which consists of two simple-yet-effective modules on top of Slot Attention. The new modules, Attention Refining Kernel (ARK) and Intermediate Point Predictor and Encoder (IPPE), respectively, prevent slots from being distracted by the background noise and indicate locations for slots to focus on to facilitate learning of object-centric representation. We also propose a weak semi-supervision approach for OCL, whilst our proposed framework can be used without any assistant annotation during the inference. Experiments show that our proposed method enables consistent learning of object-centric representation and achieves strong performance across four datasets. Code is available at <https://github.com/object-understanding/SLASH>.*

## 1. Introduction

*Object-centric learning (OCL) decomposes an image into a set of vectors corresponding to each distinct object to acquire object-wise representations [16]. Learning object-centric representation enables machines to perceive the visual world in a manner similar to humans. We recog-*

\*Equal contribution, a coin is flipped.




Figure 1. Results of training Slot Attention [35] with different seeds, which show inconsistent learning results. In the first trial, object-centric representations fail to grasp each distinct object due to the background noise. In the second, the model succeeds in distinguishing each different object from the background.

nize the world as a composition of *objects* [27] and extend the object-related knowledge to various environments [48]. Therefore, OCL enables a compositional understanding of an image and generalization for downstream tasks, such as visual reasoning [36] and object localization [6].

Mainstream OCL has adopted an autoencoding-based compositional generative model [10, 15, 35]. Slot Attention [35] is the most prominent technique for OCL, which

uses *slots* as the intermediate representation bottlenecks. In the Slot Attention, randomly initialized slots compete with each other to occupy their attention regions in terms of pixels. Eventually, each slot attains object-centric representation by aggregating visual features according to the attention map between the slot and pixels.

Recently, OCL has been extended to multi-view images [4, 42] and videos [9, 30, 46]. Multi-view image [43] or video [13, 14, 51] datasets allow models to learn spatial geometry or temporal dynamics of objects through supplementary objective tasks such as novel view synthesis [42] and optical flow inference [30]. Consequently, these datasets provide additional information that enables the adoption of data-driven inductive biases, facilitating the learning of better object-centric representations.

In contrast, it is challenging to obtain data-driven inductive biases, such as geometric or temporal information, for single-view images. To address this problem, novel architectures, such as auto-regressive generative models [3, 10, 11, 15] and Transformer [53] for encoders [44] and decoders [45], have been proposed. However, owing to the absence of additional inductive biases, OCL for complex single-view images suffers from unstable training results.

This stability issue implies inconsistent learning of object-centric representation, that is, not all trials of training a model with the same architecture consistently succeed in distinguishing objects from the background (Fig. 1). The attention-leaking problem, or bleeding issue, can mislead a model to yield object-centric representations based on distorted attention maps. The bleeding issue is fatal for OCL because it is difficult to predict the behavior of a model, that is, whether a slot will seize a distinct object or an object entangled with a background.

To solve this bleeding issue, we propose a novel OCL framework, SLASH (SLot Attention via SHepherding). SLASH resolves the bleeding by guiding the randomly initialized slots to successfully grasp objects 1) without being distracted by the background and 2) by keeping informed of the destination. These are accomplished by adding two simple-yet-effective modules, Attention Refining Kernel (ARK) and Intermediate Point Predictor and Encoder (IPPE), to the Slot Attention framework.

ARK is a single-channel single-layer convolutional kernel, designed to prevent slots from focusing on a noisy background. We adopt the Weights-Normalized Convolutional (WNConv) kernel, a learnable low-pass filter, as the kernel for ARK. This simple kernel refines the attention map between slots and pixels by reducing noise and solidifying object-like patterns.

IPPE serves as an indicator to nudge a slot to focus on the proper location. Thus, the slots can consistently update their representations without being confused by the background. IPPE consists of two submodules with simple MLPs. The

first submodule predicts the position of an object in two-dimensional coordinates, and the second encodes the predicted coordinates into a high-dimensional vector.

Since IPPE needs to be trained to provide locational cues to slots, it is necessary to introduce positional labels. However, using fully annotated ground-truths is costly, particularly for densely-annotated labels such as object masks. Hence, we adopt a weak semi-supervision approach in which only a small subset of the dataset includes weak annotations, such as the centers of bounding boxes. We show that IPPE can be successfully trained with weakly semi-supervised learning and can be deployed under circumstances where no assistant ground-truth exists.

For a comprehensive study, we validate our method on numerous datasets, including CLEVR, CLEVRTEX, PTR, and MOVi. Moreover, we conduct 10 trials of training for each method, including the baselines and ours, to thoroughly evaluate the results. We estimate the performance of the models using three metrics: mean Intersection over Union (mIoU), Adjusted Rand Index (ARI), and foreground-ARI (fg-ARI). In particular, mIoU and ARI investigate whether the bleeding issue occurs by considering the background separation. A model is defined as being stable over the metrics when deviations are lower, and as being robust when averages are higher across all datasets. Experimental results demonstrate that our method achieves stable and robust OCL that prevents the bleeding issue.

Our main contributions are as follows:

- We observe OCL for single-view images suffers from the stability issue with inconsistent training results. To resolve this issue, we propose a novel framework, SLASH (SLot Attention via SHepherding) consisting of two simple-yet-strong modules: ARK and IPPE.
- ARK is a learnable low-pass filter designed to prevent the bleeding issue where the attention of a slot leaks into a background.
- IPPE is introduced to inform slots of the regions to be focused. By leveraging weak semi-supervision, IPPE can inject positional information into a slot.
- We empirically prove SLASH achieves stable and robust OCL against four distinctive datasets. SLASH shows the best stability while outperforming the previous methods for all datasets over multiple metrics.

## 2. Related Works

### 2.1. Object-Centric Representation Learning

A line of works for OCL adopts the scene reconstruction, where a model learns to decompose an image into several components without using any human-annotated ground-truths [7, 12, 16, 17, 23, 34, 40, 50]. MONet [3] and IODINE

[15] proposed unsupervised auto-regressive approaches to sequentially disentangle object-centric representations from a scene. GENESIS [10, 11] improved object-centric learning by enabling interactions between slots while using an auto-regressive approach. Slot Attention [35] introduced an attention-based mechanism between slots and pixels, where slots parallelly and iteratively compete with each other to occupy their own territory in the pixel space. Slot Attention improved training speed and memory efficiency by enabling the parallel update of slots. Recently, SLATE [45] and DINOSAUR [44] adopted Transformer [53] as an encoder and decoder for Slot Attention, respectively, to learn object-centric representations over real-world images.

Several studies have adapted *novel view synthesis* (NVS) to OCL [4, 20, 42, 49]. ROOTS [4] proposed an approach to infer 3D disentangled object representation using 3D-to-2D perspective projection [19] with multi-view images. Other studies [42, 49] directly applied Slot Attention for multi-view images and demonstrated that using multi-view images with NVS significantly improves OCL performance.

OCL for videos has been actively studied [9, 23, 26, 30, 46, 52]. SAVi [30], SAVi++ [9] and STEVE [46] extended Slot Attention to videos, in which a model iteratively infers object-centric representations across a sequence of images. With a sequence of images, models can learn to distinguish objects from backgrounds by referring to the temporal consistency and dynamics of the objects. In this study, we focus on a more challenging case, OCL for single images, where less information about an object and its background is provided than in OCL for multi-view images and videos.

## 2.2. Weakly Supervised OCL

In weakly supervised learning, training is conducted with human-annotated labels that provide insufficient or indirect information but are pertinent to obtaining the target outputs. In OCL, GFS-Net [40] viewed the learning of object representations as a combination of *what and where* problem. GFS-Net is first trained with images containing only a single object and then fine-tuned with images containing multiple objects to resolve the what and where problem. PriSMONet [8] used shape priors in weakly supervised learning for multi-object 3D scene decomposition over RGB-D images. Furthermore, in OCL for videos, SAVi [30] and SAVi++ [9] used position information, such as the center of a mass, bounding box, or object mask, for each object in the first frame of the given video to deal with the video-level OCL. In this work, motivated by [9, 30], we utilize the point information of objects in a scene to iron out the subjectiveness problem in the image-level OCL.

## 2.3. Semi-Supervised Learning

In deep learning studies, any type of human-annotated labels, even with coarse or sparse information, can help en-

hance the performance of models. However, it is unfeasible to place ground-truths everywhere during the training and testing phases. Several semi-supervision studies have investigated how to leverage the lack of labels to solve image classification [2, 18, 33, 38], object detection [5, 22, 47, 56], and semantic/instance segmentation [1, 29, 32, 39] problems.

In this study, we adopt a novel approach for OCL, where models can only use weak supervision labels for a fraction of a given dataset. The most comparable study for different tasks is Point DETR [5]. Point DETR focused on weakly semi-supervised object detection using a dataset consisting of a few fully-annotated images with bounding boxes and object category labels, and a rich amount of weakly-annotated images with center points and object category labels. However, instead of fully-annotated images with semantic labels, our method only uses a few amounts of weakly-annotated images with point-level labels and a significant amount of non-annotated images.

## 3. Method

### 3.1. Preliminary: Slot Attention

In Slot Attention [35], the object-centric representation is implemented with the concept of `slots`  $\in \mathbb{R}^{K \times D_{slot}}$ , which is a set of  $K$  vectors of dimension  $D_{slot}$ . The slots are initialized by a Gaussian distribution with a learnable mean  $\mu$  and sigma  $\sigma$ , and updated over  $T$  iterations by the Slot Attention module. The slots are then decoded into the target reconstruction image.

We first describe the overall procedure of how Slot Attention is trained for the completeness of this study. Given an image, a convolutional neural network (CNN) encoder produces a visual feature map of dimension  $HW \times D_{enc}$ , where  $H$  and  $W$  are the height and width of an input image. The Slot Attention module takes `slots` and the visual feature map, called `inputs`, then projects them to dimension  $D$  by a linear transformation  $k$  for slots and  $q, v$  for `inputs`  $\in \mathbb{R}^{HW \times D_{enc}}$ . Dot-product attention is applied to generate an attention map, `attn`, with query-wise normalized coefficients where slots compete with each other to occupy the more relevant pixels of the visual feature map (Eq. (1) [35]).

$$\text{attn}_{i,j} := \frac{\exp(M_{i,j})}{\sum_l \exp(M_{i,l})}, \quad \text{where} \quad (1)$$

$$M := \frac{1}{\sqrt{D}} k(\text{inputs}) \cdot q(\text{slots})^T \in \mathbb{R}^{HW \times K}.$$

The projected visual feature map weighted by the attention map coefficients (Eq. (2) [35]) is aggregated to produce the updated slots, `updates`. As the Slot Attention module runs iteratively, `slots` can gradually update their representation. Each updated slot is then decoded into an RGB-

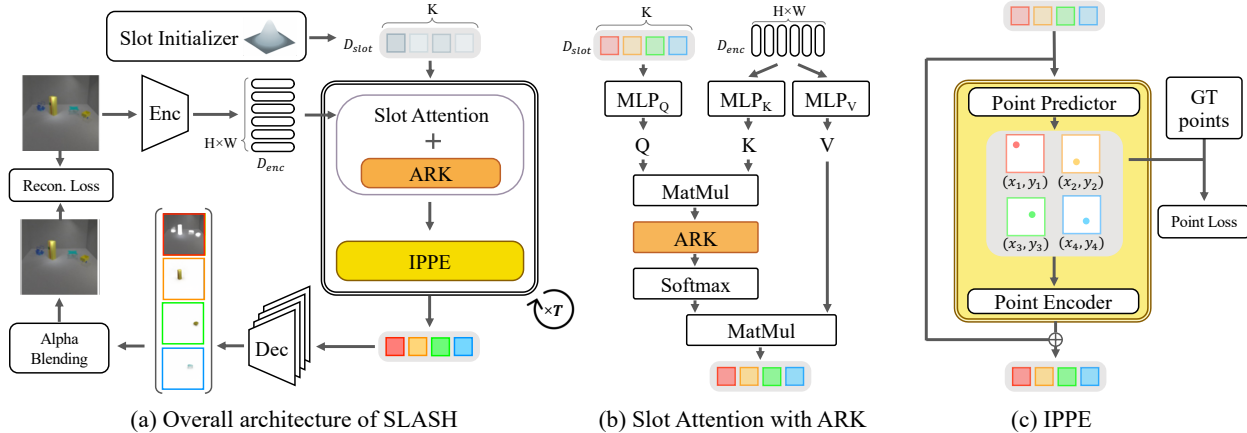


Figure 2. (a) Overall architecture of the proposed framework. Upon Slot Attention (modules without color fillings), we add Attention Refining Kernel (ARK, filled with orange color) and Intermediate Point Predictor and Encoder (IPPE, filled with yellow color). (b) Within the Slot Attention module, we insert ARK before the softmax function. (c) IPPE predicts 2D point coordinates and encode the coordinates into vectors of dimension  $D_{slots}$ . The point-encoded vectors are then added to the slots so that the slots can incorporate position information.

An image using a spatial broadcast decoder [55], where the weights are shared across slots. The decoded images are blended into a single image using alpha masks to yield the final reconstructed image. The mean squared error (MSE) between the original input image and the predicted reconstruction image is chosen for the objective function so that the overall training follows unsupervised learning.

$$\text{updates} := W^T \cdot v(\text{inputs}) \in \mathbb{R}^{K \times D},$$

$$\text{where } W_{i,j} := \frac{\text{attn}_{i,j}}{\sum_{i=1}^N \text{attn}_{i,j}}. \quad (2)$$

### 3.2. SLoT Attention via SHepherding (SLASH)

In this work, our goal is to achieve stable and robust OCL in single-view images by preventing the bleeding issue incurred when slots are distracted by background noise. To achieve this goal, the model needs to provide guidance to the slots about where to focus or not. To this end, we propose a novel OCL framework, SLASH (SLoT Attention via SHepherding), which steers slots to correctly seize objects using two newly introduced modules: Attention Refining Kernel (ARK) and Intermediate Point Predictor and Encoder (IPPE). ARK guards and stabilizes slots against background noise by reducing the noise and solidifying object-like patterns in the attention map between the slots and pixels. IPPE guides slots towards the area where an object is likely to exist by providing positional indications to the slots. Using these two simple-yet-effective modules that shepherd slots to the desired region, SLASH accomplishes stable and robust OCL. The overall architecture of SLASH is shown in Fig. 2.

#### 3.2.1 Attention Refining Kernel

Attention Refining Kernel (ARK) is designed to prevent slots from being distracted by background noise by refining the attention map between slots and visual features. As depicted in the upper part of Fig. 3, we can observe that Slot Attention [35] generates attention maps with salt-and-pepper-like noise around the objects. Noisy attention maps are likely to provoke unstable learning of object-centric representations. We address this issue by introducing an inductive bias for *local density* of objects. The bias for local density assumes that the density of the attention values should be higher near an object and lower outside the object. Thus, the inductive bias is materialized using the Weights-Normalized Convolutional (WNConv) kernel which aims to refine an attention map by reducing noise and solidifying object-like patterns around objects. WNConv kernel is a single-channel single-layer convolutional network trained under the constraints that the sum of the kernel weights equals 1 while maintaining every weight greater than or equal to 0. With these constraints, the WNConv kernel serves as a low-pass filter, smoothing the attention map as shown in the lower part of Fig. 3. As depicted in Fig. 2 (b), ARK is applied to the logit values of the attention map.

#### 3.2.2 Intermediate Point Predictor and Encoder

Intermediate Point Predictor and Encoder (IPPE) expedites learning “where objects exist”. In order for IPPE to understand the location of objects, it is necessary to introduce external supervision related to the object positions. To train our model practically, we utilize a weak semi-supervision approach. We use the low-cost information, center points




Figure 3. Visualization of the results by Attention Refining Kernel (ARK). Each colored box shows an attention map between a slot and pixels. The upper part of a rectangular box is a visualization of the attention map on the input image. The lower part represents the attention map in grayscale. The top and bottom row, split by the dotted line, corresponds to the attention map before and after applying ARK, respectively. One can observe that ARK refines the scattered attention around objects so that slots can escape from the background noise.

of bounding boxes, as the weak supervision among the possible positional cues. Furthermore, instead of using a fully annotated dataset, we assume that only a fraction (10%) of the dataset and not all objects in a given image (75%) have labels. The following describes how IPPE leverages weak semi-supervision.

IPPE consists of two modules, a point predictor and a point encoder, as shown in Fig. 2 (c). The point predictor is a 3-layer MLP that predicts 2D point coordinates of objects from slots. The point encoder, also a 3-layer MLP, encodes the point coordinates into  $D_{slot}$  dimensional vectors, which are added to the original slots. The updated slots can now contain information about the location of objects and become less likely to wander around the background.

The point predictor is trained by weak semi-supervision with an auxiliary point loss which is MSE between the predicted and ground-truth coordinates. Hungarian algorithm [31] is used to match the predictions and ground-truths. Fig. 4 shows the results of the point predictor, where the predictions get closer to objects through slot updates.

The point encoder is trained using the image reconstruction loss in a self-supervised manner [35]. As the reconstruction loss is shared with the Slot Attention module, the point encoder generates position-encoded vectors that are well-aligned with the Slot Attention module. It is worth noting that the point encoder can take either ground-truths as inputs, if available, or the predicted coordinates from the point predictor, otherwise.

Our method differs from the previous weakly supervised OCL method for videos [30] in that we use weak semi-

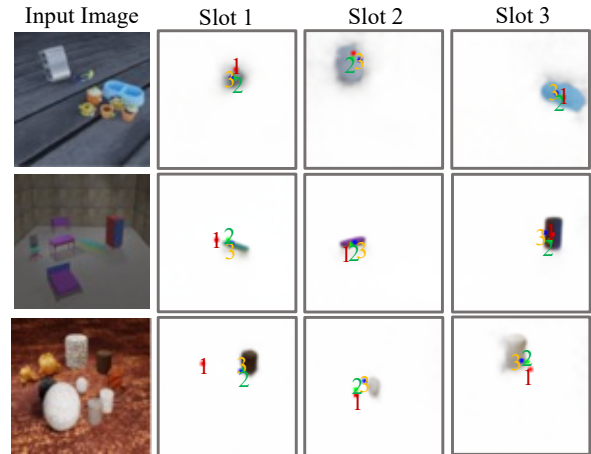


Figure 4. Visualization of point predictions by Intermediate Point Predictor and Encoder (IPPE). The leftmost column shows the input images, and the right three columns show the prediction results by IPPE for each slot. Each number stands for the order of iteration  $T$ . Best viewed in color.

annotations as the ground-truth labels for the module (Point Predictor) as well as the input for the module (Point Encoder). Conversely, SAVi exploits weak supervision to initialize slots using an MLP with the ground-truth position information as its input. That weakly supervised slot initialization shows outstanding performance for video OCL; however, it is limited in the sense that the model requires labels for all samples, even during inference time. This limitation arises from the lack of preparation for the cases where the ground-truths are not provided or are partially provided in both the training and inference phases. By virtue of the design of IPPE, our method can be trained in a weakly semi-supervised manner and can be deployed under circumstances where no ground-truth exists. In the following section, we validate the proposed method against a SAVi-like OCL method for images.

## 4. Experiments

### 4.1. Experimental Setup

**Task & Metrics** To validate the effectiveness of our method, we conduct experiments on the object discovery task following the previous OCL works [9–11, 30, 35, 40]. In the object discovery task, a model is required to cluster pixels into object segments. Though the task seems similar to the instance segmentation, the object discovery differs from the image segmentation in that it does not require semantic classes or captions for each segmentation.

To evaluate the models, we use mean Intersection over Union (mIoU) and Adjusted Rand Index (ARI) [41]. Similar to [10, 37], we avoid focusing on foreground-ARI (fg-ARI) where the annotation for the background is excluded

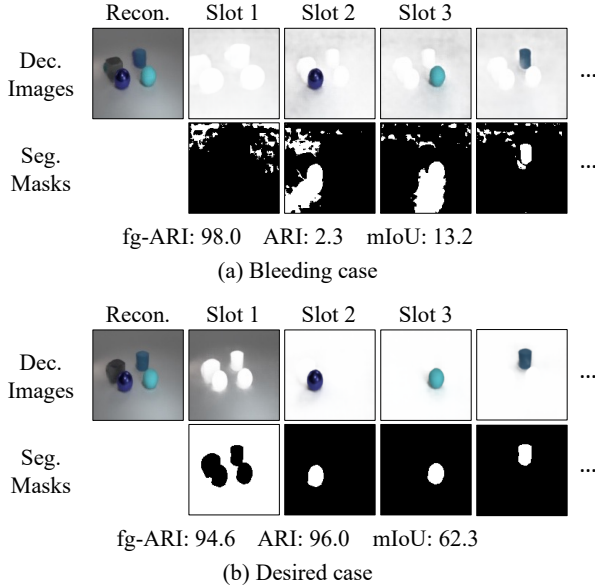


Figure 5. Demonstration of the bleeding issue. The top row of each case shows decoded images by a model [35] and the bottom row shows segmentation masks corresponding to the decoded images. At the bottom of each case, evaluation results by three different metrics are reported in %. One can observe that fg-ARI cannot represent the bleeding case in contrast to ARI and mIoU.

from the evaluation. Fig. 5 demonstrates that fg-ARI cannot describe the stability issue, such as the bleeding. On the other hand, the stability issue can be demonstrated using mIoU and ARI since the annotation of backgrounds is concerned with those metrics.

**Baselines** We compare SLASH with Slot Attention (SA) [35], GENESIS-v2 (GenV2) [11], and weakly supervised Slot Attention (WS-SA). GenV2 is a recent study on OCL in single-view images, derived from GENESIS [10]. The official GenV2<sup>1</sup> is used in our datasets. Additionally, we compare SLASH with WS-SA, a simple variant of SA, equipped with a weakly supervised slot initializer. The WS-SA initializes each slot using an MLP, which takes the point coordinates of an object as input by following SAVi [30]. Unlike SAVi, we assume the datasets do not contain labels for the precise number of objects in an image. Therefore, we initialize the surplus slots with randomly sampled values from the Gaussian distribution as opposed to SAVi which initializes surplus slots that receive no ground-truth point coordinates with  $(-1, -1)$  to deactivate the slots.

**Datasets** The experiments cover four multi-object datasets: CLEVR6 [24], CLEVRTEX [28], PTR [21] and MOVi-C [14]. CLEVR was designed to assess the models’ comprehension of compositional elements, such as visual reasoning. CLEVR6 contains 35K train and 7.5K validation samples consisting of scenes with three to six ob-

<sup>1</sup><https://github.com/applied-ai-lab/genesis>

	mIoU	ARI	fg-ARI
CLEVR6			
SA [35]	49.5 ± 20.5	63.0 ± 42.1	97.1 ± 1.6
SA <sup>†</sup>	46.4 ± 24.6	59.2 ± 48.2	98.3 ± 0.9
WS-SA*	61.8 ± 4.3	90.7 ± 2.0	93.3 ± 1.6
<b>SLASH*</b>	63.6 ± 4.3	90.3 ± 4.3	94.2 ± 1.3
CLEVRTEX			
SA	22.2 ± 4.3	38.1 ± 12.5	52.1 ± 5.9
MONet [3] <sup>‡</sup>	19.8 ± 1.0	—	36.7 ± 0.9
IODINE [15] <sup>‡</sup>	29.2 ± 0.8	—	59.2 ± 2.2
GenV2 [11] <sup>‡</sup>	7.9 ± 1.5	—	31.2 ± 12.4
WS-SA*	22.4 ± 4.5	36.0 ± 7.2	52.3 ± 7.3
<b>SLASH*</b>	34.7 ± 5.3	59.4 ± 11.5	61.9 ± 6.4
PTR			
SA	17.6 ± 14.7	19.6 ± 29.8	44.5 ± 18.8
GenV2	28.5 ± 11.3	41.0 ± 25.4	56.8 ± 13.0
WS-SA*	23.8 ± 15.5	21.4 ± 33.6	52.9 ± 11.6
<b>SLASH*</b>	44.1 ± 9.6	67.9 ± 22.6	59.0 ± 3.2
MOVi			
SA	23.0 ± 9.8	25.9 ± 20.3	48.7 ± 7.0
GenV2	10.8 ± 1.1	3.6 ± 0.2	47.1 ± 5.8
WS-SA*	21.6 ± 11.5	22.8 ± 21.3	46.2 ± 8.5
<b>SLASH*</b>	27.7 ± 5.9	34.6 ± 13.5	51.9 ± 4.0

Table 1. Results over the object discovery task (mean ± std for 10 trials, reported in %). \* indicates that the model is trained by weakly semi-supervised learning. All models performed inference with no assistant label. † is for the results by [35] which uses a center crop. ‡ is for the results from [28] which conducted three trials of training for each method with a center crop.

jects [25, 35]. CLEVRTEX is a variant of CLEVR, having complicated shapes, textures, materials, and backgrounds. CLEVRTEX contains 50K samples, which we split into 40K train and 10K validation set. PTR, which contains 52K train and 9K validation samples, is a visual reasoning dataset in which objects have part-whole hierarchies. MOVi-C is a synthetic video dataset comprising realistic and textured daily objects and backgrounds. We collected the first frames of the randomly rendered videos that have scenes with at most six objects. Our MOVi-C dataset contains 39K train and 9K validation samples. The supplementary material contains details of the data collection process. **Training** All models are trained by the MSE reconstruction loss in an autoencoding fashion. The training environments for WS-SA and SLASH are the same as those of SA [35], while those of GenV2 follow the official paper [11]. The number of slots,  $K$ , is set to 7 for CLEVR6, PTR, and MOVi-C, and 11 for CLEVRTEX.

## 4.2. Object Discovery

The quantitative results on the object discovery task are summarized in Tab. 1. The bleeding case causes significant degradation of mIoU and ARI, that is, the metrics have higher deviations and lower averages. We argue that

	mIoU	ARI	fg-ARI
CLEVR6			
SA [35]	49.5 ± 20.5	63.0 ± 42.1	97.1 ± 1.6
+ ARK	64.1 ± 3.1	89.9 ± 2.2	95.3 ± 1.3
+ IPPE	57.8 ± 14.6	84.9 ± 27.9	95.7 ± 1.0
+ ARK + IPPE	63.6 ± 4.3	90.3 ± 4.3	94.2 ± 1.3
CLEVRTEX			
SA	22.2 ± 4.3	38.1 ± 12.5	52.1 ± 5.9
+ ARK	31.4 ± 6.6	55.6 ± 13.2	57.8 ± 7.7
+ IPPE	25.1 ± 7.4	40.4 ± 15.6	54.9 ± 7.3
+ ARK + IPPE	34.7 ± 5.3	59.4 ± 11.5	61.9 ± 6.4
PTR			
SA	17.6 ± 14.7	19.6 ± 29.8	44.5 ± 18.8
+ ARK	43.8 ± 3.0	62.3 ± 19.4	60.4 ± 3.2
+ IPPE	38.4 ± 12.8	58.4 ± 31.3	58.5 ± 3.1
+ ARK + IPPE	44.1 ± 9.6	67.9 ± 22.6	59.0 ± 3.2
MOVi			
SA	23.0 ± 9.8	25.9 ± 20.3	48.7 ± 7.0
+ ARK	26.2 ± 6.1	33.2 ± 13.7	51.0 ± 3.7
+ IPPE	27.2 ± 7.9	36.2 ± 16.8	50.8 ± 5.7
+ ARK + IPPE	27.7 ± 5.9	34.6 ± 13.5	51.9 ± 4.0

Table 2. Results of the ablation studies on the modules of SLASH (mean ± std for 10 trials, reported in %).

a model is stable when it has lower deviations and robust when it has higher averages across all datasets. Thus it is crucial to prevent the bleeding case for stable and robust OCL. SLASH records the highest average value of mIoU and ARI for almost all datasets except for ARI on CLEVR6 with a minimal difference. In addition, SLASH demonstrates lower standard deviation values of mIoU and ARI across overall datasets. To sum up, SLASH scores the highest and the most consistent performance across all datasets, achieving stable and robust OCL. We provide abundant qualitative results in the supplementary material due to spatial constraints.

### 4.3. Ablation Studies

#### 4.3.1 ARK and IPPE

To prove the effectiveness of ARK and IPPE, we conduct an ablation study on those modules by training SA [35] with or without each module. Tab. 2 demonstrates that SLASH benefits from both ARK and IPPE.

We observe that ARK apparently stabilizes the model, resulting in low standard deviation values for overall datasets. IPPE boosts the performance of both SA and ‘+ARK’. Although the standard deviation values tend to be high due to the absence of ARK, we find that IPPE aids in learning against a complicated dataset, i.e. MOVi, where slots struggle to grasp the visual patterns of objects. We argue that the positional information given by IPPE can aid the slots in binding appropriate objects more effectively.

	mIoU	ARI	fg-ARI
CLEVR6			
SA ( $\tau = 1$ ) [35]	49.5 ± 20.5	63.0 ± 42.1	97.1 ± 1.6
SA ( $\tau = 2$ )	52.7 ± 24.6	66.4 ± 40.6	95.8 ± 2.8
Gaussian	50.1 ± 14.9	75.6 ± 29.4	93.2 ± 2.4
Conv	52.7 ± 19.6	76.8 ± 32.0	92.5 ± 2.2
WNConv	64.1 ± 3.1	89.9 ± 2.2	95.3 ± 1.6
CLEVRTEX			
SA ( $\tau = 1$ )	22.2 ± 4.3	38.1 ± 12.5	52.1 ± 5.9
SA ( $\tau = 2$ )	25.6 ± 2.0	39.6 ± 3.7	54.9 ± 2.7
Gaussian	26.0 ± 8.5	43.5 ± 14.6	55.9 ± 11.0
Conv	24.8 ± 6.0	42.5 ± 9.7	54.3 ± 11.1
WNConv	31.4 ± 6.6	55.6 ± 13.2	57.8 ± 7.7
PTR			
SA ( $\tau = 1$ )	17.6 ± 14.7	19.6 ± 29.8	44.5 ± 18.8
SA ( $\tau = 2$ )	34.3 ± 12.0	56.6 ± 26.5	50.0 ± 7.9
Gaussian	20.6 ± 15.1	20.0 ± 29.9	53.8 ± 10.2
Conv	12.4 ± 9.7	11.6 ± 13.1	32.1 ± 26.0
WNConv	43.8 ± 3.0	62.3 ± 19.4	60.4 ± 3.2
MOVi			
SA ( $\tau = 1$ )	23.0 ± 9.8	25.9 ± 20.3	48.7 ± 7.0
SA ( $\tau = 2$ )	27.1 ± 5.5	28.7 ± 12.5	54.6 ± 1.8
Gaussian	25.5 ± 10.8	33.7 ± 21.5	48.5 ± 7.1
Conv	25.5 ± 8.8	28.2 ± 18.0	53.0 ± 2.4
WNConv	27.2 ± 6.1	33.2 ± 13.7	57.0 ± 3.7

Table 3. Results of ablation studies on the alternatives of the kernel for ARK (mean ± std for 10 trials, reported in %). SA stands for the baseline Slot Attention [35].  $\tau$  is a temperature coefficient in the attention mechanism.

#### 4.3.2 Kernels in ARK

In our method, ARK is applied as a low-pass filter with a WNConv network to eliminate the noise and strengthen the object-like patterns in the attention maps. In this study, we look into the alternatives for the WNConv. Firstly, we compare our WNConv with the global smoothing scheme where we increase the temperature  $\tau$ . The logit values of an attention map are divided by  $\tau$  so that the larger  $\tau$  yield a more smoothed attention map. Secondly, we apply a representative smoothing technique, Gaussian filter [54]. Lastly, we conduct experiments on a Convolutional (Conv) kernel without any constraints like weight normalization.

Tab. 3 demonstrates the results of the possible kernels for ARK. We observe that the SA models with high temperature and the Gaussian smoothing outperform the original SA. This result implies that simple global and local smoothing can help a model boost its performance by erasing noises in the attention maps. For the Conv kernel, we observe that the overall performance is worse than the other kernels. We argue that the poor performance of the Conv kernel is incurred due to the high degree of freedom for the single-channel single-layer network. In contrast, owing to the reduced degree of freedom, WNConv consistently performs well by recording high average values than the alternatives.

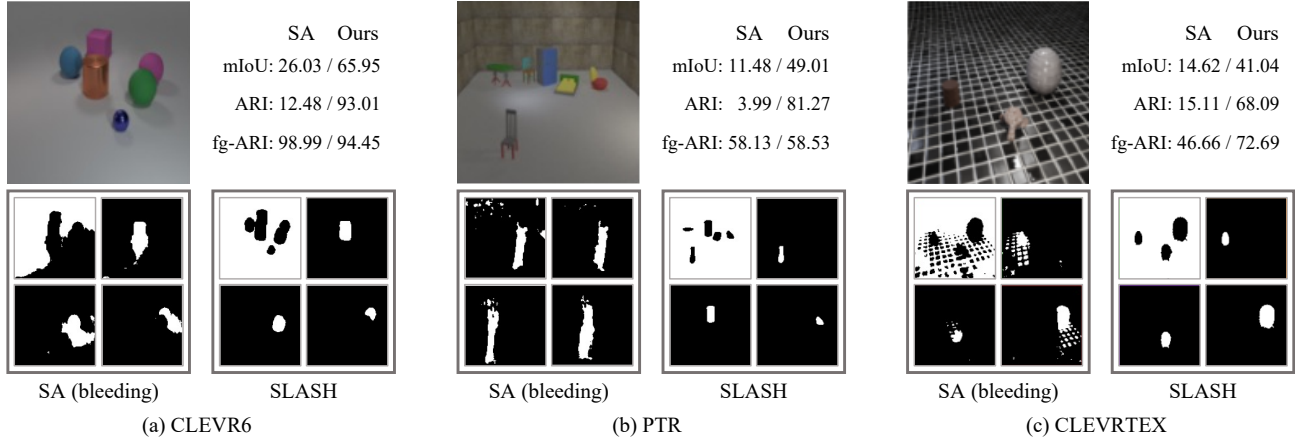


Figure 6. Qualitative evaluation of SLASH compared to the baseline, Slot Attention (SA) [35]. The images in the top left of each dataset section are randomly selected inputs fed to the models. The bottom row shows the segmentation masks generated from attention maps of slots. The numbers in the top right of each section are for the qualitative results of each model over each dataset (reported in %).

#### 4.4. Analysis of Bleeding Issues

In this section, we investigate the cases where the baseline, Slot Attention [35], fails to prevent the bleeding issue and SLASH succeeds in that. Fig. 6 shows the results by the baseline and SLASH. Here, we present our analysis of the failure cases for each dataset.

As depicted in Fig. 6 (a), for CLEVR6, SA encounters the bleeding issue due to the simplicity of the background. As shown in the top-left image, CLEVR6 only contains simple white backgrounds without any complicated pattern. Since the background has almost no information, a model is likely to get into the trivial solution that every slot binds to the piece of the background.

As shown in Fig. 6 (b), stripping frequently occurs for PTR. The striping issue is a phenomenon where each slot is trapped into a simple and meaningless stripe pattern of an image. We assume the striping issue occurs as a model tends to focus on positional embedding rather than object-related patterns that are difficult for the model to figure out.

CLEVRTEX contains a variety of complex backgrounds as shown in Fig. 6 (c). In SA, slots tend to be attracted by the explicit eye-catching patterns on backgrounds. We argue that this phenomenon is attributed to the design of SA which focuses on the versatility towards domain- and task-agnostic models. This design principle results in the lack of inductive biases and locational information for discovering objects rather than backgrounds.

As SLASH is designed to not only eliminate background noise and solidify object-like patterns in the attention map, but also encode the positional information into the slots, we observe that SLASH is robust against the aforementioned failure cases in various datasets.

## 5. Conclusion and Limitation

In this paper, we observed that OCL for single-view images has a stability issue that some training trials end up with having the bleeding issue. We attributed this problem to the lack of inductive bias about the appearance of objects and additional cues like positional information. We presented a new OCL framework for single-view images, called SLASH, acting as a shepherd, guiding the slots to the correct destination without being distracted by the background noise. To accomplish this, we proposed two simple modules: ARK for smoothing the noise in the attention and IPPE for inducing positional information through a weak semi-supervision. Experimental results show the effectiveness of our method, which achieved strong and consistent results for stable and robust OCL.

Although our model shows impressive results on various challenging synthetic datasets, extending our method to real-world image datasets remains a problem and promising path for future work. There are several potential challenges to achieving real-world OCL: understanding backgrounds, controlling a large number of objects, handling the representation of an object having intricate shape and texture, and designing an efficient model that can process high-resolution images. We expound on additional limitations of our study in the supplementary material.

## 6. Acknowledgement

This research was supported and funded by Artificial Intelligence Graduate School Program under Grant 2020-0-01361, Artificial Intelligence Innovation Hub under Grant 2021-0-02068, and the Korean National Police Agency. [Project Name: XR Counter-Terrorism Education and Training Test Bed Establishment/Project Number: PR08-04-000-21].



## References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, 2018. 3
- [2] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019. 3
- [3] Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019. 2, 6
- [4] Chang Chen, Fei Deng, and Sungjin Ahn. Roots: Object-centric representation and rendering of 3d scenes. *The Journal of Machine Learning Research*, 2021. 2, 3
- [5] Liangyu Chen, Tong Yang, Xiangyu Zhang, Wei Zhang, and Jian Sun. Points as queries: Weakly semi-supervised object detection by points. In *CVPR*, 2021. 3
- [6] Minsu Cho, Suha Kwak, Cordelia Schmid, and Jean Ponce. Unsupervised object discovery and localization in the wild: Part-based matching with bottom-up region proposals. In *CVPR*, 2015. 1
- [7] Eric Crawford and Joelle Pineau. Spatially invariant unsupervised object detection with convolutional neural networks. In *AAAI*, number 01, 2019. 2
- [8] Cathrin Elich, Martin R Oswald, Marc Pollefeys, and Joerg Stueckler. Weakly supervised learning of multi-object 3d scene decompositions using deep shape priors. *CVIU*, 2022. 3
- [9] Gamaleldin F Elsayed, Aravindh Mahendran, Sjoerd van Steenkiste, Klaus Greff, Michael C Mozer, and Thomas Kipf. Savi++: Towards end-to-end object-centric learning from real-world videos. In *NeurIPS*, 2022. 2, 3, 5
- [10] Martin Engelcke, Adam R Kosiorek, Oiwi Parker Jones, and Ingmar Posner. Genesis: Generative scene inference and sampling with object-centric latent representations. In *ICLR*, 2020. 1, 2, 3, 5, 6
- [11] Martin Engelcke, Oiwi Parker Jones, and Ingmar Posner. Genesis-v2: Inferring unordered object representations without iterative refinement. In *NeurIPS*, 2021. 2, 3, 5, 6
- [12] SM Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, Geoffrey E Hinton, et al. Attend, infer, repeat: Fast scene understanding with generative models. In *NeurIPS*, 2016. 2
- [13] Rohit Girdhar and Deva Ramanan. Cater: A diagnostic dataset for compositional actions and temporal reasoning. In *ICLR*, 2020. 2
- [14] Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J Fleet, Dan Gnanaprasam, Florian Golemo, Charles Herrmann, Thomas Kipf, Abhijit Kundu, Dmitry Lagun, Issam Laradji, Hsueh-Ti (Derek) Liu, Henning Meyer, Yishu Miao, Derek Nowrouzezahrai, Cengiz Oztireli, Etienne Pot, Noha Radwan, Daniel Rebain, Sara Sabour, Mehdi S. M. Sajjadi, Matan Sela, Vincent Sitzmann, Austin Stone, Deqing Sun, Suhani Vora, Ziyu Wang, Tianhao Wu, Kwang Moo Yi, Fangcheng Zhong, and Andrea Tagliasacchi. Kubric: a scalable dataset generator. In *CVPR*, 2022. 2, 6
- [15] Klaus Greff, Raphaël Lopez Kaufman, Rishabh Kabra, Nick Watters, Christopher Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference. In *ICML*. PMLR, 2019. 1, 2, 3, 6
- [16] Klaus Greff, Antti Rasmus, Mathias Berglund, Tele Hao, Harri Valpola, and Jürgen Schmidhuber. Tagger: Deep unsupervised perceptual grouping. In *NeurIPS*, 2016. 1, 2
- [17] Klaus Greff, Sjoerd Van Steenkiste, and Jürgen Schmidhuber. Neural expectation maximization. In *NeurIPS*, 2017. 2
- [18] Philip Haeusser, Alexander Mordvintsev, and Daniel Cremers. Learning by association—a versatile semi-supervised training method for neural networks. In *CVPR*, 2017. 3
- [19] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 3
- [20] Paul Henderson and Christoph H. Lampert. Unsupervised object-centric video generation and decomposition in 3D. In *NeurIPS*, 2020. 3
- [21] Yining Hong, Li Yi, Joshua B Tenenbaum, Antonio Torralba, and Chuang Gan. Ptr: A benchmark for part-based conceptual, relational, and physical reasoning. In *NeurIPS*, 2021. 6
- [22] Jisoo Jeong, Seungeui Lee, Jeesoo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. In *NeurIPS*, 2019. 3
- [23] Jindong Jiang, Sepehr Janghorbani, Gerard De Melo, and Sungjin Ahn. Scalar: Generative world models with scalable object representations. In *ICLR*, 2020. 2, 3
- [24] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2017. 6
- [25] Rishabh Kabra, Chris Burgess, Loic Matthey, Raphael Lopez Kaufman, Klaus Greff, Malcolm Reynolds, and Alexander Lerchner. Multi-object datasets. <https://github.com/deepmind/multi-object-datasets/>, 2019. 6
- [26] Rishabh Kabra, Daniel Zoran, Goker Erdogan, Loic Matthey, Antonia Creswell, Matt Botvinick, Alexander Lerchner, and Chris Burgess. Simone: View-invariant, temporally-abstracted object representations via unsupervised video decomposition. In *NeurIPS*, 2021. 3
- [27] Daniel Kahneman, Anne Treisman, and Brian J Gibbs. The reviewing of object files: Object-specific integration of information. *Cognitive psychology*, (2), 1992. 1
- [28] Laurynas Karazija, Iro Laina, and Christian Rupprecht. Clevrtex: A texture-rich benchmark for unsupervised multi-object segmentation. In *NeurIPS*, 2021. 6
- [29] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *CVPR*, 2017. 3
- [30] Thomas Kipf, Gamaleldin F Elsayed, Aravindh Mahendran, Austin Stone, Sara Sabour, Georg Heigold, Rico Jonschkowski, Alexey Dosovitskiy, and Klaus Greff. Condi-

- tional object-centric learning from video. In *ICLR*, 2022. [2](#), [3](#), [5](#), [6](#)
- [31] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, (1-2), 1955. [5](#)
- [32] Xin Lai, Zhuotao Tian, Li Jiang, Shu Liu, Hengshuang Zhao, Liwei Wang, and Jiaya Jia. Semi-supervised semantic segmentation with directional context-aware consistency. In *CVPR*, 2021. [3](#)
- [33] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. [3](#)
- [34] Zhixuan Lin, Yi-Fu Wu, Skand Vishwanath Peri, Weihao Sun, Gautam Singh, Fei Deng, Jindong Jiang, and Sungjin Ahn. Space: Unsupervised object-oriented scene representation via spatial attention and decomposition. In *ICLR*, 2020. [2](#)
- [35] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. In *NeurIPS*, 2020. [1](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [36] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B Tenenbaum, and Jiajun Wu. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In *ICLR*, 2019. [1](#)
- [37] Tom Monnier, Elliot Vincent, Jean Ponce, and Mathieu Aubry. Unsupervised layered image decomposition into object prototypes. In *ICCV*, 2021. [5](#)
- [38] Youngtaek Oh, Dong-Jin Kim, and In So Kweon. Distribution-aware semantics-oriented pseudo-label for imbalanced semi-supervised learning. *arXiv preprint arXiv:2106.05682*, 2021. [3](#)
- [39] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *ICCV*, 2015. [3](#)
- [40] Mihir Prabhudesai, Anirudh Goyal, Deepak Pathak, and Katerina Fragkiadaki. Generating fast and slow: Scene decomposition via reconstruction. *arXiv preprint arXiv:2203.11194*, 2022. [2](#), [3](#), [5](#)
- [41] William M Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical association*, (336), 1971. [5](#)
- [42] Mehdi SM Sajjadi, Daniel Duckworth, Aravindh Mahendran, Sjoerd van Steenkiste, Filip Pavetić, Mario Lučić, Leonidas J Guibas, Klaus Greff, and Thomas Kipf. Object scene representation transformer. In *NeurIPS*, 2022. [2](#), [3](#)
- [43] Mehdi SM Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lučić, Daniel Duckworth, Alexey Dosovitskiy, et al. Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. In *CVPR*, 2022. [2](#)
- [44] Maximilian Seitzer, Max Horn, Andrii Zadaianchuk, Dominik Zietlow, Tianjun Xiao, Carl-Johann Simon-Gabriel, Tong He, Zheng Zhang, Bernhard Schölkopf, Thomas Brox, et al. Bridging the gap to real-world object-centric learning. *arXiv preprint arXiv:2209.14860*, 2022. [2](#), [3](#)
- [45] Gautam Singh, Fei Deng, and Sungjin Ahn. Illiterate dall-e learns to compose. In *ICLR*, 2021. [2](#), [3](#)
- [46] Gautam Singh, Yi-Fu Wu, and Sungjin Ahn. Simple unsupervised object-centric learning for complex and naturalistic videos. *arXiv preprint arXiv:2205.14065*, 2022. [2](#), [3](#)
- [47] Kihyuk Sohn, Zizhao Zhang, Chun-Liang Li, Han Zhang, Chen-Yu Lee, and Tomas Pfister. A simple semi-supervised learning framework for object detection. *arXiv preprint arXiv:2005.04757*, 2020. [3](#)
- [48] Elizabeth S Spelke and Katherine D Kinzler. Core knowledge. *Developmental science*, (1), 2007. [1](#)
- [49] Karl Stelzner, Kristian Kersting, and Adam R Kosiorek. Decomposing 3d scenes into objects via unsupervised volume segmentation. *arXiv preprint arXiv:2104.01148*, 2021. [3](#)
- [50] Karl Stelzner, Robert Peharz, and Kristian Kersting. Faster attend-infer-repeat with tractable probabilistic models. In *ICML*. PMLR, 2019. [2](#)
- [51] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2020. [2](#)
- [52] Sjoerd Van Steenkiste, Michael Chang, Klaus Greff, and Jürgen Schmidhuber. Relational neural expectation maximization: Unsupervised discovery of objects and their interactions. In *ICLR*, 2018. [3](#)
- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. [2](#), [3](#)
- [54] David Vernon. Machine vision: Automated visual inspection and robot vision. 1991. [7](#)
- [55] Nicholas Watters, Loic Matthey, Christopher P Burgess, and Alexander Lerchner. Spatial broadcast decoder: A simple architecture for learning disentangled representations in vaes. *arXiv preprint arXiv:1901.07017*, 2019. [4](#)
- [56] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyun Wei, Xiang Bai, and Zicheng Liu. End-to-end semi-supervised object detection with soft teacher. In *ICCV*, 2021. [3](#)