




# Benchmarking and Analysis of Unsupervised Object Segmentation from Real-World Single Images

Yafei Yang<sup>1,2</sup> · Bo Yang<sup>1,2</sup> 

Received: 22 April 2023 / Accepted: 5 December 2023 / Published online: 6 January 2024  
© The Author(s) 2024

## Abstract

In this paper, we study the problem of unsupervised object segmentation from single images. We do not introduce a new algorithm, but systematically investigate the effectiveness of existing unsupervised models on challenging real-world images. We first introduce seven complexity factors to quantitatively measure the distributions of background and foreground object biases in appearance and geometry for datasets with human annotations. With the aid of these factors, we empirically find that, not surprisingly, existing unsupervised models fail to segment generic objects in real-world images, although they can easily achieve excellent performance on numerous simple synthetic datasets, due to the vast gap in objectness biases between synthetic and real images. By conducting extensive experiments on multiple groups of ablated real-world datasets, we ultimately find that the key factors underlying the failure of existing unsupervised models on real-world images are the challenging distributions of background and foreground object biases in appearance and geometry. Because of this, the inductive biases introduced in existing unsupervised models can hardly capture the diverse object distributions. Our research results suggest that future work should exploit more explicit objectness biases in the network design.

**Keywords** Unsupervised multi-object segmentation · Object-centric learning · Object discovering · Object-oriented scene representation

## 1 Introduction

The ability of automatically identifying individual objects from complex visual observations is a central aspect of human intelligence (Spelke et al., 1992). It serves as the key building block for higher-level cognition tasks such as planning and reasoning (Greff et al., 2020). In last years, a plethora of models have been proposed to segment objects from single static images in an unsupervised fashion: from the early AIR (Eslami et al., 2016), MONet (Burgess et al., 2019) to the recent SPACE (Lin et al., 2020), SlotAtt

(Locatello et al., 2020), GENESIS-V2 (Engelcke et al., 2021), etc. They jointly learn to represent and segment multiple objects from a single image, without needing any human annotations in training. This process is often called perceptual grouping/binding or object-centric learning. These methods and their variants have achieved impressive segmentation results on numerous synthetic datasets such as dSprites (Matthey et al., 2017) and CLEVR (Johnson et al., 2017). Such advances come with great expectations that the unsupervised techniques would likely close the gap with fully-supervised methods for real-world visual understanding. However, few work has systematically investigated the true potential of the emerging unsupervised object segmentation models on complex real-world images such as COCO dataset (Lin et al., 2014). This naturally raises an essential question:

*Is it promising or even possible to segment generic objects from real-world single images using (existing) unsupervised methods?*

*What is an object?* To answer the above question involves another fundamental question: what is an object? Exactly 100 years ago in Gestalt psychology, Wertheimer (Wertheimer,

Communicated by O. Veksler.

✉ Bo Yang  
bo.yang@polyu.edu.hk

Yafei Yang  
ya-fei.yang@connect.polyu.hk

<sup>1</sup> Shenzhen Research Institute, The Hong Kong Polytechnic University, Shenzhen, China

<sup>2</sup> vLAR Group, Department of Computing, The Hong Kong Polytechnic University, Hung Hom, Hong Kong

1923) first introduced a set of *perceptual grouping* and *figure-ground organization* principles to heuristically define visual data as individual objects and/or backgrounds. Perceptual grouping principles such as proximity and similarity specify how foreground visual elements are grouped into individual objects, whereas figure-ground organization principles investigate how visual data are just separated as the foreground against background, e.g. a relatively small area tends to be seen as foreground (Wagemans et al., 2012).

However, these principles are highly subjective, whilst the real-world scenes and objects are far more complex with extremely diverse appearances and shapes. Therefore, it is practically impossible to quantitatively define what is an object, i.e., the objectness, from visual inputs (e.g., a set of image pixels). Nevertheless, to thoroughly understand whether unsupervised methods can truly learn objectness akin to the psychological process of humans, it is vital to investigate the underlying factors that potentially facilitate or otherwise hinder the ability of unsupervised models. In this regard, by drawing on Gestalt principles, we instead define a series of new factors to quantitatively measure the complexity of individual foreground objects and the background in Sect. 2. By taking into account both appearance and geometry, our complexity factors explicitly assess the difficulty of segmenting single objects and the background. For example, it is harder to segment a chair with colorful textures from a cluttered background than a single-color ball from a clean background for unsupervised methods. With the aid of these factors, we extensively study whether and how existing unsupervised models can discover objects in Sect. 4.

*What is the problem of unsupervised object segmentation from single images?* A large number of models (Yuan et al., 2022) aim to tackle the problem of unsupervised object segmentation from single images. They share several key problem settings: (1) all training images do not have any human annotations; (2) every single image has multiple objects optionally with a textured background; (3) each image is treated as a static data point without any dynamic or temporal information; (4) all models are trained from scratch without requiring any pretrained networks on additional datasets. Ultimately, the goal of these models is to segment all individual objects as accurate as the ground truth human annotations. In this paper, we regard these settings as the basic and necessary part of unsupervised object segmentation from single images, and empirically evaluate how successfully the existing models can exhibit on real-world images.

*Contributions and findings* This paper addresses the essential question regarding the potential of unsupervised segmentation of generic objects from real-world single images. Our contributions are:

- We first introduce 4 complexity factors to quantitatively measure the difficulty of individual objects and 3 fac-



**Fig. 1** SlotAtt can perfectly segment simple objects on three synthetic images (left-hand side) but clearly fails on three real-world images (right-hand side)

tors to quantify backgrounds. These factors are key to investigating the true potential of existing unsupervised models.

- We extensively evaluate current unsupervised approaches in a large-scale experimental study. We implement 4 representative unsupervised methods and train more than 200 models on 4 groups of curated datasets from scratch. The datasets, code and pretrained models are available at <https://github.com/vLAR-group/UnsupObjSeg>.
- We analyze our experimental results and find that: (1) existing unsupervised object segmentation models cannot discover generic objects from single real-world images, although they can achieve outstanding performance on synthetic datasets, as qualitatively illustrated in Fig. 1; (2) the challenging distributions of both foreground object and the background biases in appearance and geometry from real-world images are the key factors incurring the failure of existing models; (3) the inductive biases introduced in existing unsupervised models are fundamentally not matched with the objectness biases exhibited in real-world images, and therefore fail to discover the real objectness.

*Related Work* Recently, ClevrTex (Karazija et al., 2021) and the concurrent work (Papa et al., 2022) also study unsupervised object segmentation on single images. Through evaluation on (complex) synthetic datasets only, both works focus on benchmarking the effectiveness of particular network designs of baselines. By comparison, our paper aims to explore what and how the objectness distribution gaps between synthetic and real-world images incur the failure of existing models. The recent work by Weis et al. (2021) which investigates video object discovery is orthogonal to our work as the motion signals do not exist in single images.

*Scope of this research* In addition to our core study on purely unsupervised object segmentation approaches from single images, we also include discussions about more recent methods such as Odin (Hénaff et al., 2022), DINOSAUR (Seitzer et al., 2022), FreeSOLO (Wang et al., 2022), and CutLER (Wang et al., 2023) that require pretrained models to obtain objectness biases on monolithic object images such as ImageNet (Russakovsky et al., 2015) in Sect. 4.7. How-

ever, this paper does not investigate object segmentation from saliency maps (Wang et al., 2021), static multi-views (Yuan et al., 2022), or dynamic videos (Singh et al., 2022a, b), because the input information of these methods is clearly different from that on single images.

A preliminary version of this work has been published in Yang and Yang (2022) and our new technical extensions include: (1) 3 additional complexity factors to quantitatively measure the difficulty of background in each image in Sect. 2.3, (2) new benchmarks with 6 more datasets, a new baseline, and 3 extra metrics in Sect. 3, (3) and more comprehensive experiments to analyze object segmentation of existing models.

## 2 Complexity Factors

As illustrated in the left three images of Fig. 2, an individual object, represented by a set of color pixels within a mask, can vary significantly given different types of appearance and geometric shape. A specific scene, represented by a set of objects placed on a clean canvas, can also differ vastly given different relative appearances and geometric layouts between objects, as illustrated in the middle images. If these objects are instead placed on diverse backgrounds, they tend to appear more differently as shown in the right images.

Unarguably, such variations and complexity of appearance and geometry in object level, scene level, and background level, directly affect human's ability to precisely separate all objects. Naturally, the performance of unsupervised segmentation models is also expected to be influenced by these variations. In this regard, we carefully define the following three groups of factors to quantitatively describe the complexity of visual scenes in datasets.

### 2.1 Object-Level Complexity Factors

As to an object, all its information can be described by appearance and geometry. Therefore we define the below two factors to measure the complexity of appearance and geometry respectively. Notably, both factors are nicely invariant to the object scale.

**Object Color Gradient:** This factor aims to calculate how frequently the appearance changes within the object mask. As shown in Fig. 3, given the RGB image and mask of an




Fig. 3 Illustrations for object color gradient




Fig. 4 Illustrations for object shape concavity

object, we firstly convert RGB into grayscale and then apply Sobel filter (Sobel & Feldman, 1973) to compute the gradients horizontally and vertically for each pixel within the mask. The final gradient value is obtained by averaging out all object pixels. Note that, the object boundary pixels are removed to avoid the interference of background. Numerically, the higher this factor is, the more complex texture and/or lighting effect the object has, and therefore it is likely harder to segment.

**Object Shape Concavity:** This factor is designed to evaluate how irregular the object boundary is. As shown in Fig. 4, given an object (binary) mask, denoted as  $M_{obj} \in \mathbb{R}^{H \times W}$ , we firstly find the smallest convex polygon mask ( $M_{cvx} \in \mathbb{R}^{H \times W}$ ) that surrounds the object mask using an existing algorithm (Eddins, 2011), and then the object shape concavity value is computed as:  $1 - \sum M_{obj} / \sum M_{cvx}$ . Clearly, the higher this factor is, the more irregular object shape is, and segmentation is trickier.

### 2.2 Scene-Level Complexity Factors

Given an image, in addition to the object-level complexity, the spatial and appearance relationships between all objects can also incur extra difficulty for segmentation. We define the following two factors to quantify the complexity of relative appearance and geometry between objects in an image.

**Inter-object Color Similarity:** This factor intends to assess the appearance similarity between all objects in the same image. As shown in Fig. 5, we first calculate the average color




Fig. 2 Complexity in appearance and geometry for objects and scenes



Fig. 5 Illustrations for inter-object color similarity




Fig. 6 Illustrations for inter-object shape variation

for each object, and then compute the pair-wise Euclidean distances of object colors, obtaining a  $K \times K$  matrix where  $K$  represents the object number. The *average color distance* is calculated by averaging the matrix excluding diagonal entries, and the final inter-object color similarity is computed as:  $1 - \text{average color distance} / (255 \times \sqrt{3})$ . Intuitively, the higher this factor is, the more similar all objects appear to be, the less distinctive each object is, and the harder it is to separate each object.

**Inter-object Shape Variation:** This factor aims to measure the relative geometry diversity between all  $K$  objects in an image. We first find the diagonal vector of each object bounding box. Then we calculate the differences between each pair of diagonal vectors, resulting in  $\frac{K \times (K-1)}{2}$  vectors (denoted as black vectors in Fig. 6). The final inter-object shape variation is the averaged norm of these  $\frac{K \times (K-1)}{2}$  vectors. The higher this factor, objects within an image are more diverse with imbalanced sizes, and therefore segmenting both gigantic and tiny objects is likely harder.

By capturing the appearance and geometry of objects in both object and scene levels, the four factors in Sects. 2.1 and 2.2 are designed to quantify the image complexity only




Fig. 8 Illustrations for background color gradient

with clean backgrounds. For illustration, Fig. 7 shows sample images for the four factors at different values. The higher the values, the more complex foreground objects are in both object and scene levels.

In fact, these factors are carefully selected from more than 10 candidates because they are empirically more suitable to differentiate the gaps between synthetic and real-world images, and they eventually serve as key indicators to diagnose existing unsupervised models in Sect. 4. Details of other candidates are in the appendix.

### 2.3 Background Complexity Factors

Since the diversity of image backgrounds also plays a critical role in potentially distinguishing or messing up foreground objects, we introduce the following three factors to measure the complexity of backgrounds. The background here is defined as a collection of all pixels which do not belong to any considered objects in a single image.

**Background Color Gradient:** This factor is designed to measure how frequently the appearance changes within background pixels. The calculation is the same Object Color Gradient as illustrated in Fig. 8. The background boundary is also removed to avoid the interference of foreground objects. The higher this factor, the more complex texture and/or lighting of the background, and therefore both the background and objects may be more difficult to be segmented.

**Background-Foreground Color Similarity:** This factor aims to measure the appearance similarity between background and foreground pixels in an image. Specifically, we first calculate the Euclidean distance between each background pixel RGB and each foreground pixel RGB, resulting

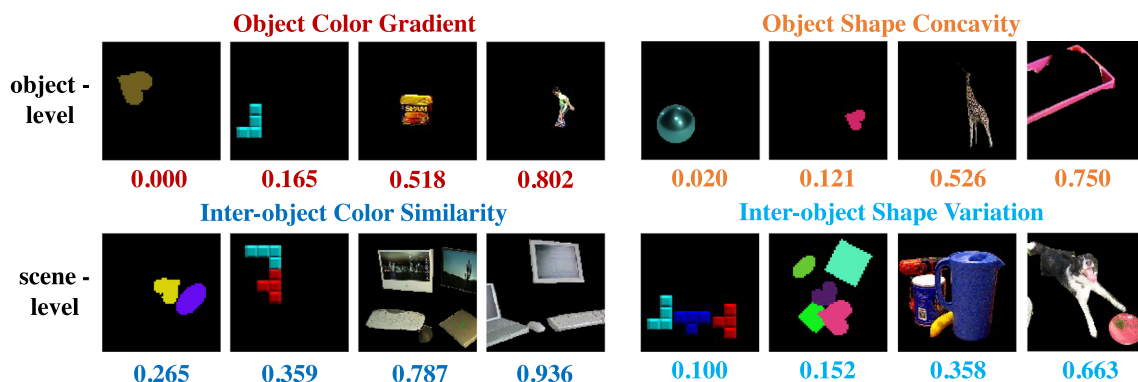



Fig. 7 Sample objects and scenes for the four factors at different complexity values





**Fig. 9** Illustrations for background shape irregularity

in a  $U \times V$  matrix where  $U$  and  $V$  represent the number of pixels in background and foreground respectively. We then treat it as a form of cost matrix where the Hungarian algorithm is applied to obtain an optimal *cost value* between all background and foreground pixels. The final background-foreground color similarity is computed as:  $1 - \text{cost value}$ . Intuitively, the higher this factor, the more similar background and foreground appear to be, and it is harder to separate each other. Details are in the appendix.

**Background Shape Irregularity:** This factor aims to measure the irregularity of the background shape, i.e., the full contour where the background intersects with all foreground objects. As illustrated in Fig. 9, given the (binary) background mask, we first divide its contour into subcontours, where each subcontour is fully self-connected and encloses a region. In fact, that region may originally contain one or more foreground objects, which means that the previously designed Inter-object Shape Variation factor for foreground objects cannot be simply reused to measure the shape complexity of background. For each enclosed region, we first compute its maximal inscribed convex set based on an existing algorithm (Borgefors & Strand, 2005) whose details are provided in the Appendix. The area of each enclosed region is denoted as  $A_i$  and the area of its inscribed convex set is denoted as  $C_i$ . The irregularity of each subcontour is then calculated as  $1 - \frac{C_i}{A_i}$ . The final Background Shape Irregularity score for each image is calculated as the average of irregularity score for all subcontours within the image as:  $\frac{1}{N} \sum_{i=1}^N \left(1 - \frac{C_i}{A_i}\right)$ , where  $N$  is the total number of enclosed regions within an image. Clearly, the higher this factor, the more irregular the background shape (contour), and therefore the harder to separate background pixels from foreground objects.

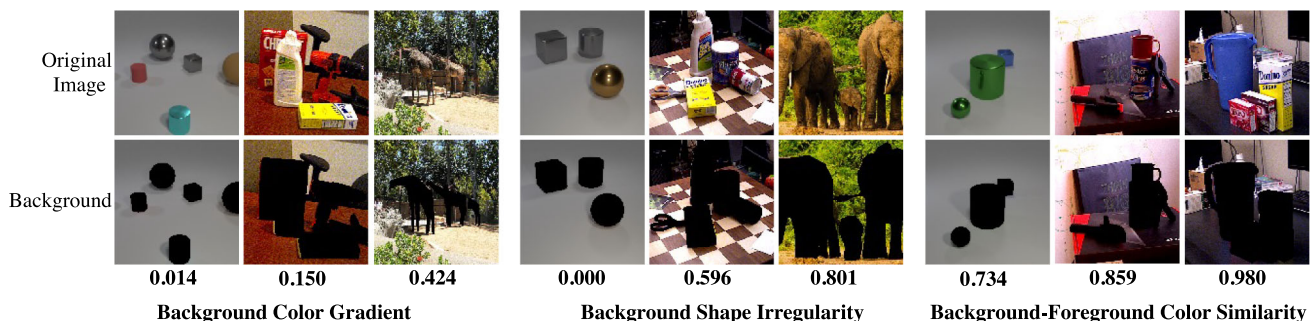
Above all, the three factors together are designed to quantify the complexity of image backgrounds in terms of appearance and shape. Figure 10 shows sample images for three background complexity factors at different values. The original full images are also included for reference.

### 3 Experimental Design

#### 3.1 Considered Methods

A range of works have explored unsupervised object segmentation in recent years. They are typically formulated as (variational) autoencoders (AE/VAE) (Kingma & Welling, 2014) or generative adversarial networks (GAN) (Goodfellow et al., 2014). GAN based models (Abdal et al., 2021; Arandjelovic & Zisserman, 2019; Azadi et al., 2020; Bielski & Favaro, 2019; Chen et al., 2019; van Steenkiste et al., 2020; Voynov et al., 2021) are usually limited to identifying a single foreground object and can hardly discover multiple objects due to the training instabilities, therefore not considered in this paper. There is another line of works that adopt Energy-based Models (EBMs) (Du et al., 2021; LeCun et al., 2006; Liu et al., 2022) on scene decomposition. However, they do not explicitly generate object masks but encode objects into energy functions, therefore not discussed in this paper. As shown in Table 1, the majority of existing models are based on AE/VAE and can be generally divided into two groups according to the object representation:

- *Factor-based models:* Each object is represented by explicit factors such as size, position, appearance, etc., and the whole image is a spatial organization of multiple objects. Basically, such representation explicitly enforces objects to be bounded within particular regions.
- *Layer-based models:* Each object is represented by an image layer, i.e., a binary mask, and the whole image is a spatial mixture of multiple object layers. Intuitively, this representation does not have strict spatial constraints, and instead is more flexible to cluster similar pixels as objects.



**Fig. 10** Sample image backgrounds for the three factors at different complexity values

**Table 1** Existing unsupervised models for object segmentation on single images

Factor-based models		Inductive Biases			Layer-based Models		Inductive Biases		
		VAE	Iter	Rel			VAE	Iter	Rel
CST-VAE (Huang & Murphy, 2016)	ICLRW'16	✓			Tagger (Greff et al., 2016)	NIPS'16		✓	
AIR (Eslami et al., 2016)	NIPS'16	✓			RC (Greff et al., 2016)	ICLRW'16		✓	
SPAIR (Crawford and Pineau, 2019)	AAAI'19	✓			NEM (Greff et al., 2017)	NIPS'17		✓	
SuPAIR (Stelzner et al., 2019)	ICML'19	✓			MONet (Burgess et al., 2019)	arXiv'19	✓		
GMIO (Yuan et al., 2019)	ICML'19	✓	✓		IODINE (Greff et al., 2019)	ICML'19	✓	✓	
ASR (Xu et al., 2019)	NeurIPS'19	✓			ECON (von Kügelgen et al., 2020)	ICLRW'20	✓		✓
SPACE (Lin et al., 2020)	ICLR'20	✓			GENESIS (Engelcke et al., 2020)	ICLR'20	✓		✓
GNM (Jiang & Ahn, 2020)	NeurIPS'20	✓		✓	SlotAtt (Locatello et al., 2020)	NeurIPS'20		✓	
SPLIT (Charakorn et al., 2020)	arXiv'20	✓			GENESIS-V2 (Engelcke et al., 2021)	NeurIPS'21	✓		✓
OCIC (Anciukevicius et al., 2020)	arXiv'20	✓		✓	CAE (Löwe et al., 2022)	TMLR'22			✓
GSGN (Deng et al., 2021)	ICLR'21	✓		✓	BO-QSA (Jia et al., 2023)	ICLR'23		✓	
ISA (Biza et al., 2023)	ICML'23	✓	✓		ISA (Biza et al., 2023)	ICML'23	✓	✓	

Each model includes different inductive biases, such as variational autoencoding (VAE), iterative inference (Iter), object relationship regularization (Rel), etc.

To decompose input images into objects against backgrounds, these approaches introduce different types of network architecture, losses, and regularization terms as inductive biases. These biases broadly include: (1) variational encoding which encourages the disentanglement of latent variables; (2) iterative inference which likely ends up with better scene representations over occlusions; (3) object relationship regularization such as depth estimation and autoregressive prior which aims at capturing the dependency of multiple objects; and many other biases. With different combinations of these biases, many methods have shown outstanding performance on synthetic datasets. Among them, we select 4 representative models for investigation: (1) AIR (Eslami et al., 2016), (2) MONet (Burgess et al., 2019), (3) IODINE (Greff et al., 2019), and (4) SlotAtt (Locatello et al., 2020). We also add the fully-supervised Mask R-CNN (He et al., 2017) as an additional baseline for a comprehensive comparison in Sects. 4.1–4.6.

More recently, a number of new methods such as Odin (Hénaff et al., 2022), DINOSAUR (Seitzer et al., 2022), FreeSOLO (Wang et al., 2022), and CutLER (Wang et al., 2023) make use of pretrained models on monolithic object images such as ImageNet to segment objects from real-world images. We additionally evaluate and discuss the representative method DINOSAUR (Seitzer et al., 2022) in Sect. 4.7.

### 3.2 Considered Datasets

We consider four groups of datasets for extensive benchmarking and analysis:

- *Group 1*: Three commonly-used synthetic datasets with blank backgrounds: dSprites (Matthey et al., 2017), Tetris (Kabra et al., 2019) and CLEVR (Johnson et al., 2017);
- *Group 2*: One semi-realistic and three real-world datasets with blank backgrounds: MOVIC (Greff et al., 2022), YCB (Calli et al., 2017), ScanNet (Dai et al., 2017), and COCO (Lin et al., 2014), representing semi-realistic, small-scale, indoor- and outdoor-level real scenes respectively;
- *Group 3*: One synthetic dataset with synthetic backgrounds: CLEVR<sub>bg</sub> (Johnson et al., 2017), whose foreground objects are the same as CLEVR in Group 1;
- *Group 4*: One semi-realistic and three real-world datasets with backgrounds: MOVIC<sub>bg</sub> (Greff et al., 2022), YCB<sub>bg</sub> (Calli et al., 2017), ScanNet<sub>bg</sub> (Dai et al., 2017), and COCO<sub>bg</sub> (Lin et al., 2014), whose foreground objects are the same as the three datasets in Group 2.

Naturally, objects and scenes in different datasets tend to have very different types of biases. For example, the objects

in dSprites tend to have a single-color bias, while COCO does not; the objects in ScanNet<sub>bg</sub> tend to be influenced by cluttered backgrounds, while ScanNet does not. Generally, the object-level biases can be divided as: (1) appearance biases including different textures and lighting effects for objects, and (2) geometry biases including the object shape and occlusions. Similarly, the scene-level biases include: (1) appearance biases such as the color similarity between all objects, and (2) geometry biases such as the diversity of all object shapes. The background biases can be divided as: (1) appearance biases including textures and lighting effects of backgrounds, and the color similarity between background and foreground, (2) geometry biases such as the irregularity of background shape. In fact, our complexity factors introduced in Sect. 2 are designed to well capture these biases. Table 3 in the appendix qualitatively summarizes the biases of datasets in Groups 1/2/3/4. We may hypothesize that the large gaps of biases between synthetic and real-world datasets would have a huge impact on the effectiveness of existing models.

To guarantee the fairness and consistency of all experiments, we carefully prepare the four groups of datasets using the following same protocols. Preparation details of each dataset are provided in the appendix.

- All images are rerendered or cropped with the same resolution of  $128 \times 128$ .
- Each image has about 2 to 6 solid objects, with a blank background in Groups 1&2, with a synthetic background in Group 3, and with a (semi)real background in Group 4.
- Each dataset has about 10,000 images for training, 2000 images for testing.

The datasets in Groups 1&2 are primarily used to evaluate how the four object- and scene-level complexity factors affect object segmentation performance of existing methods, while the datasets in Groups 3&4 are used to investigate how the diversity of backgrounds affects the segmentation performance.

### 3.3 Considered Metrics

Having the 12 representative datasets and existing unsupervised methods at hand, we choose the following metrics to evaluate the object segmentation performance: (1) AP score which is widely used for object detection and segmentation (Everingham et al., 2015), (2) PQ score which is used to measure non-overlap panoptic segmentation (Kirillov et al., 2019), and (3) Precision and Recall scores. A predicted mask is considered correct if its IoU against a ground truth mask is above 0.5. All objects are treated as a single class. The blank background is not taken into account for foreground object

segmentation. To compute AP, we simply treat the mean value of the soft object mask as the object confidence score. In order to quantitatively analyze over-/under-segmentation issues, we additionally calculate ARP/ARR scores (Zimmermann et al., 2023) and ARI score (Rand, 1971). Note that, the alternative segmentation covering (SC) (Arbeláez et al., 2011) is not considered as it can be easily saturated.

## 4 Key Experimental Results


### 4.1 Can Current Unsupervised Models Succeed on Real-world Datasets?

We first evaluate the performance of all five baselines on synthetic datasets (dSprites, Tetris, CLEVR) in Group 1, as well as on the datasets MOVIC, YCB, ScanNet, and COCO in Group 2. It is important to note that all images in these datasets only have blank backgrounds. Both quantitative and qualitative results are presented in Figs. 11 and 12. Detailed breakdown ARP/ARR scores for comparing over-/under-segmentation are provided in Fig. 13.

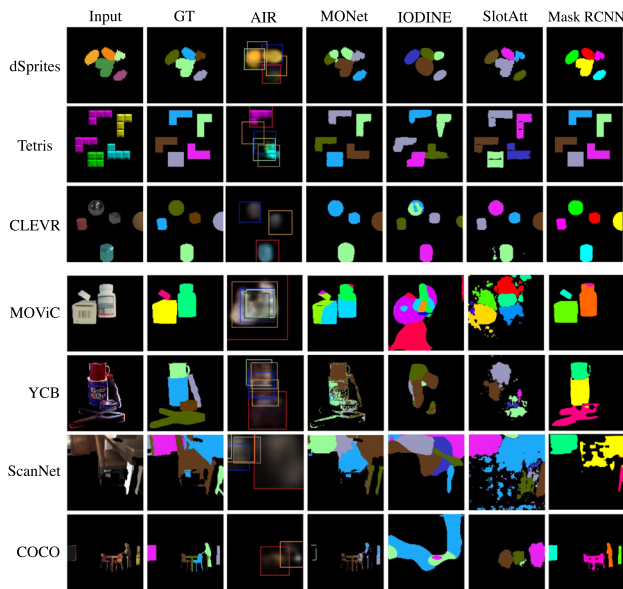
Notably, on synthetic datasets, all methods, especially the recent baselines IODINE and SlotAtt, demonstrate satisfactory segmentation outcomes. However, it is not surprising that all unsupervised methods encounter significant challenges when applied to real-world datasets, despite the blank backgrounds observed in all images. In addition, from ARP/ARR scores, the methods AIR, MONet and IODINE are more prone to over-segmentation, while SlotAtt tends to exhibit under-segmentation across all 7 datasets. Such an observation is aligned with results in (Zimmermann et al., 2023).

We then evaluate three baselines (IODINE/SlotAtt/MaskRCNN, because of their better abilities to model backgrounds) on datasets (CLEVR<sub>bg</sub>/MOVIC<sub>bg</sub>/YCB<sub>bg</sub>/ScanNet<sub>bg</sub>/COCO<sub>bg</sub>) from Groups 3&4, where all images have either synthetic or real-world backgrounds. Figs. 14 and 15 show the quantitative/qualitative results respectively. We can see that both unsupervised IODINE&SlotAtt achieve excellent performance on the synthetic dataset CLEVR<sub>bg</sub>, but again fail on real-world datasets with real backgrounds.


*Preliminary Diagnosis:* To diagnose the failure on real-world images, we hypothesize that it is because of the gaps in biases between synthetic and real-world datasets. In this regard, we quantitatively compute the distributions of all 7 complexity factors on the four groups of datasets. In particular, the two object-level factors, i.e., Object Color Gradient and Object Shape Concavity, and the two scene-level factors, i.e., Inter-object Color Similarity and Inter-object Shape Variation, are computed on the seven training splits of Groups 1&2 datasets (dSprites/Tetris/CLEVR, MOVIC/YCB/ScanNet/COCO). The three background-level



**Fig. 11** Quantitative results of object segmentation on 7 datasets from Groups 1&2 with blank backgrounds




**Fig. 12** Qualitative results of object segmentation from five methods on datasets of Groups 1&2 with blank backgrounds




**Fig. 15** Qualitative results of object segmentation on CLEVR<sub>bg</sub>/MOVIC<sub>bg</sub>/YCB<sub>bg</sub>/ScanNet<sub>bg</sub>/COCO<sub>bg</sub>

factors, i.e., Background Color Gradient, Background-Foreground Color Similarity and Background Shape Irregularity, are computed for each image of the four training splits of Groups 3&4 datasets (CLEVR<sub>bg</sub>, MOVIC<sub>bg</sub>, YCB<sub>bg</sub>/ScanNet<sub>bg</sub>/COCO<sub>bg</sub>).



**Fig. 13** Quantitative results for comparing over/under-segmentation of 4 unsupervised methods on 7 datasets of Groups 1&2



**Fig. 14** Quantitative results of object segmentation on 5 datasets of Groups 3&4 with synthetic or real-world backgrounds



As shown in the *top* row of Fig. 16, we can see that: (1) for the two object-level factors (Subfigs 1-a and 1-b), the complexity scores for synthetic datasets are substantially lower than those of real-world datasets, and the semi-realistic MOViC is in between. This implies that synthetic objects are more likely to have uniform colors and convex shapes; (2) for the two scene-level factors (Subfigs 1-c and 1-d), the images in synthetic datasets tend to include less similar objects in terms of color, which means that multiple objects in real-world scenes are less distinctive in appearance. In addition, the multiple objects in synthetic scenes tend to have similar sizes, whereas real-world scenes usually have diverse object scales in single images. As shown in the *bottom* row of Fig. 16, we can see that: (1) real-world image backgrounds are more likely to have non-uniform colors and irregular shapes (Subfigs 2-a and 2-b); (2) synthetic images tend to have more distinctive backgrounds against foreground objects than real-world images (Subfig 2-c).

To validate whether these distribution biases are the true reasons incurring the failure, we conduct extensive ablative experiments in Sects. 4.2, 4.3, 4.4, and 4.5.

## 4.2 How do Object-Level Factors Affect Current Models?

In this section, we aim to verify to what extent the distributions of object-level factors affect the segmentation performance by doing the following three ablative experiments. Examples of object-level ablations can be found in the appendix Fig. 36.

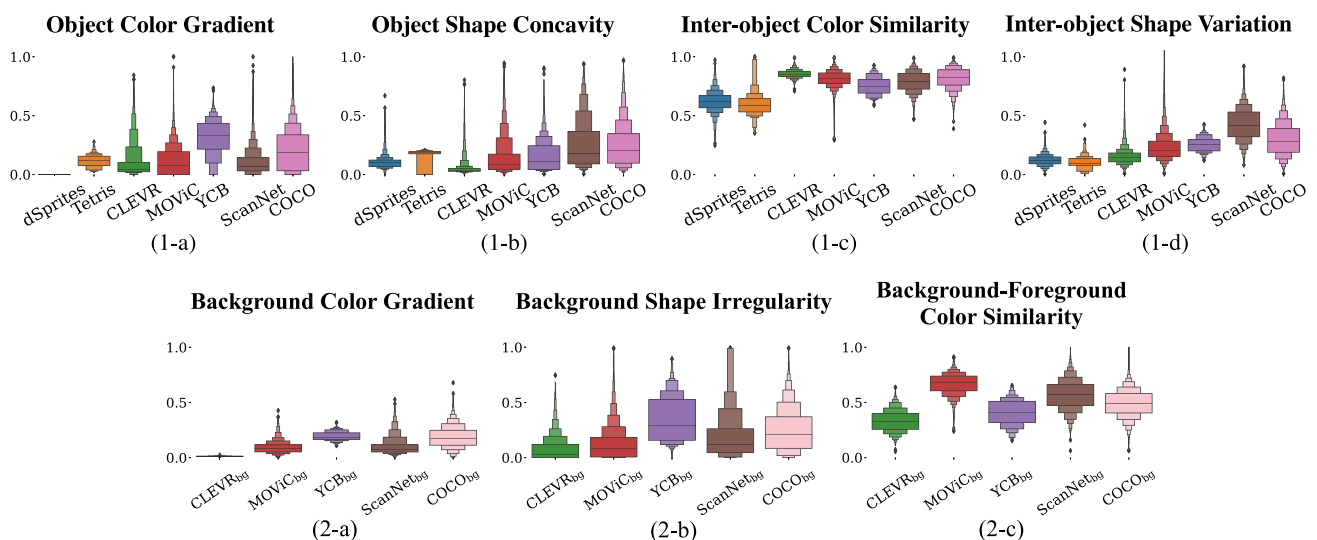
- *Ablation of Object Color Gradient*: For each object of Group 2 datasets (MOViC/YCB/ScanNet/COCO), we only replace all pixel colors by its average color *rgb* value

within each object mask, without touching the object shapes. In this way, the color gradients of each object are totally erased, thus removing the potential impact of Object Color Gradient. The three ablated datasets are: MOViC-C/YCB-C/ScanNet-C/COCO-C.

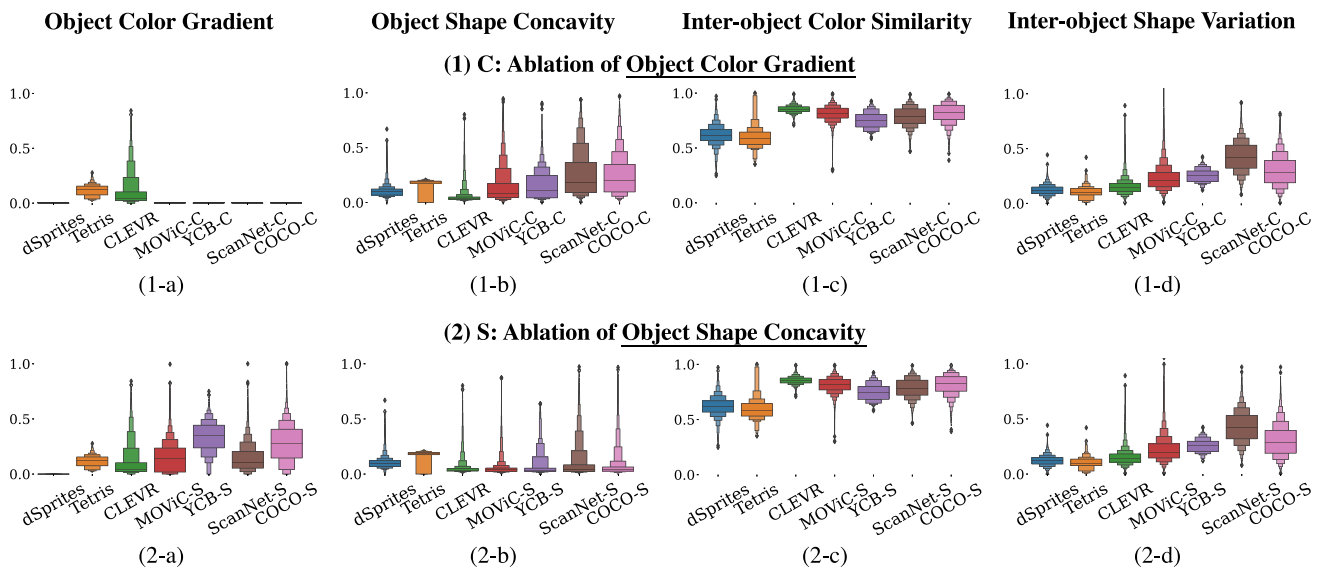
- *Ablation of Object Shape Concavity*: For each object in Group 2 datasets (MOViC/YCB/ScanNet/COCO), we find the smallest convex hull (Eddins, 2011) for its object mask and then fill the empty pixels by shifting original object pixels. Basically, this ablation aims to only reduce the irregularity of object shapes, yet retain the distributions of color gradients. The ablated datasets are: MOViC-S/YCB-S/ScanNet-S/COCO-S.
- *Ablation of Object Color Gradient and Shape Concavity*: We combine the above two ablations for each object, getting datasets: MOViC-C+S/YCB-C+S/ScanNet-C+S/COCO-C+S.

Figure 17(1a)–(1d) show new distributions of two object-level factors and two scene-level factors on the -C ablated datasets. We can see that the Object Color Gradient becomes all zeros in Fig. 17(1-a), even simpler than three synthetic datasets. Yet, the distributions of other three factors are almost the same as the original ones, i.e., Fig. 17(1-b)–(1-d) being similar to Fig. 16(1-b)–(1-d).

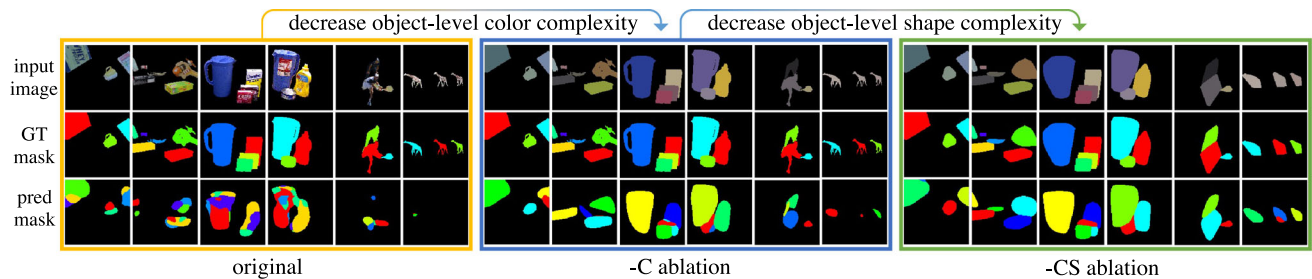
Figure 17(2-b) shows that the distributions of Object Shape Concavity of -S ablated datasets now become similar to the synthetic datasets, while the distributions of other three factors are unaffected, i.e., Fig. 17(2-a)–(2-d) being similar to Fig. 16(1-a), (1-c), (1-d). Note that, for the -C+S ablations, the distributions will be the same as shown in Fig. 17(1-a), (2-b), (1-c) or (2-c), (1-d) or (2-d). Having the three groups of object-level ablated real-world datasets, we then train and



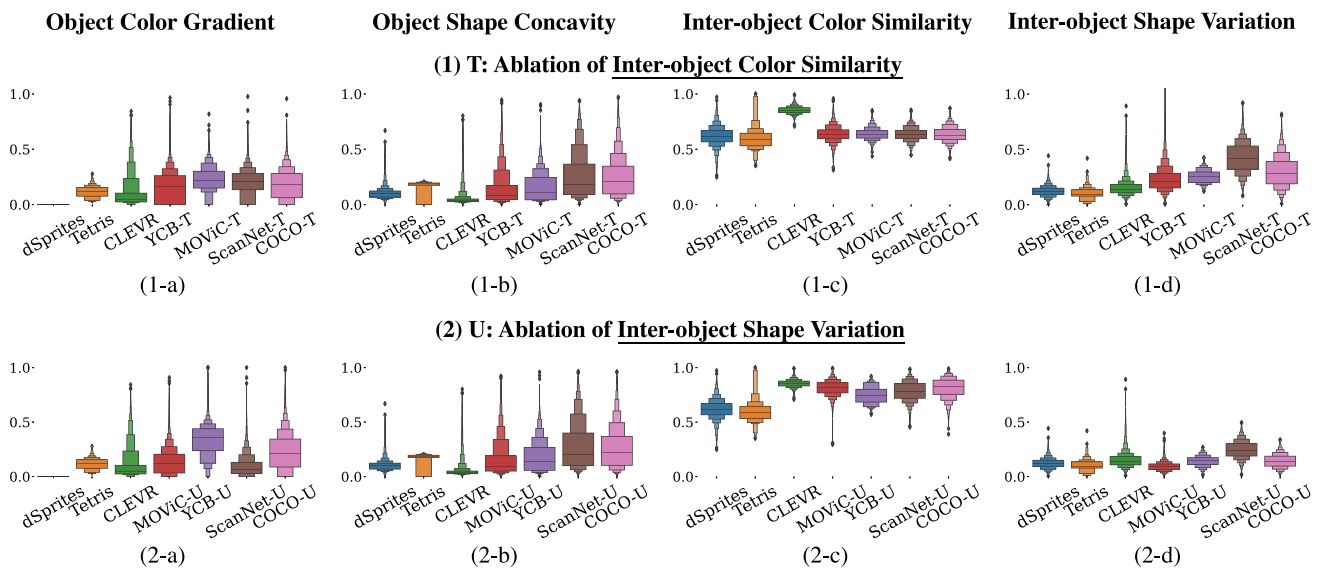
**Fig. 16** Distributions of the 7 complexity factors on four groups of datasets




**Fig. 17** Distributions of four foreground factors on datasets ablated in object-level



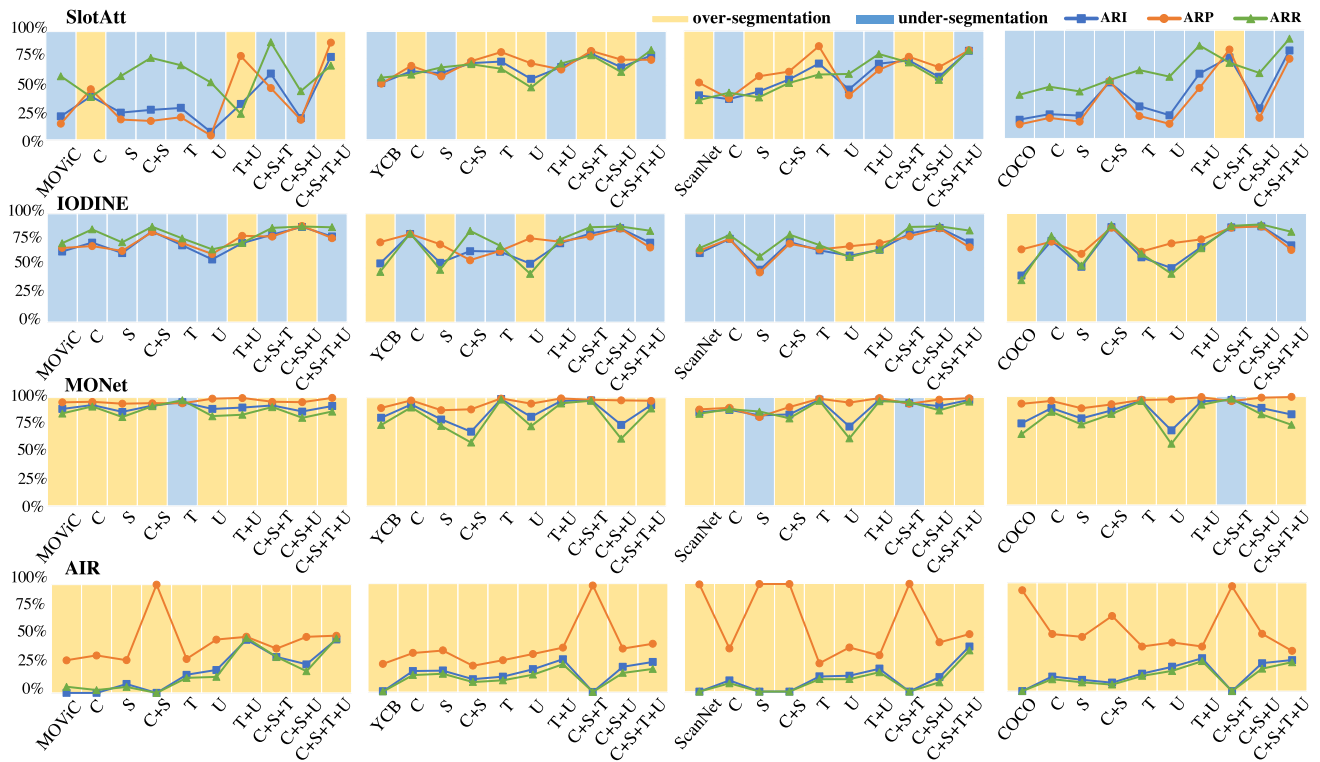
**Fig. 18** Qualitative results of object-level factors ablations from IODINE



**Fig. 19** Distributions of four foreground factors on datasets ablated in scene-level

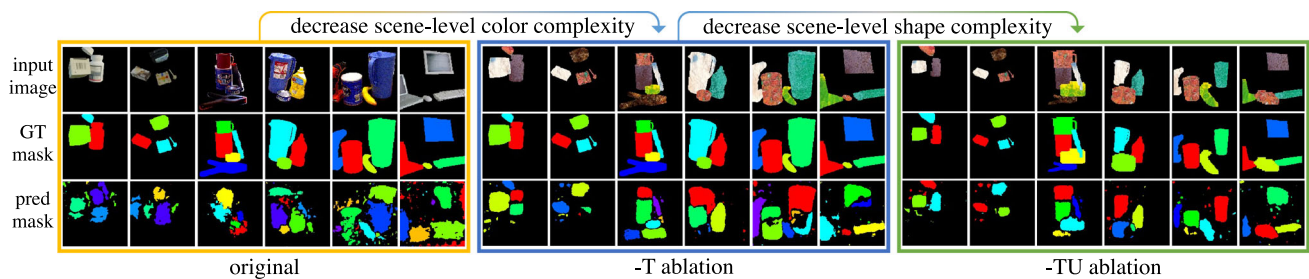


(a) Quantitative results of baselines on the group 2 datasets (with blank backgrounds) and their variants.

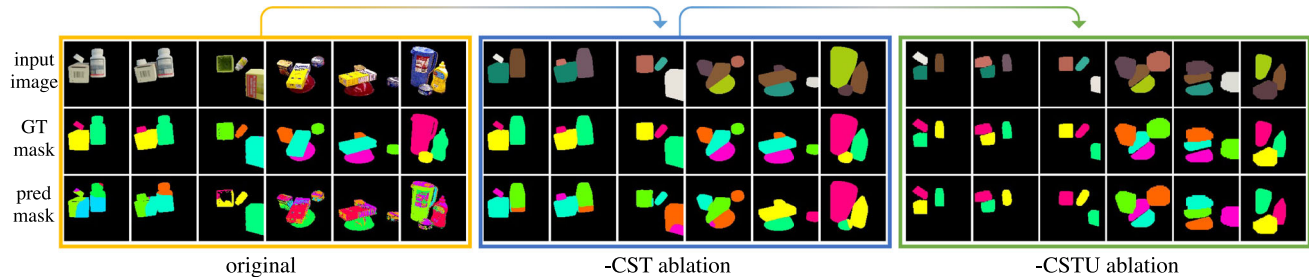


(b) The ARP/ARR/ARI scores of baselines on the group 2 datasets for analyzing over-/under-segmentation.

**Fig. 20** The letters C/S/C+S represent the three ablated datasets in Sect. 4.2; T/U/T+U represent the three ablated datasets in Sect. 4.3; C+S+T/C+S+U/C+S+T+U represent the three ablated datasets in Sect. 4.4



**Fig. 21** Qualitative results of scene-level factors ablations from SlotAtt




**Fig. 22** Qualitative results of both object-level and scene-level factors ablations from MONet

evaluate our four unsupervised baselines from scratch on each of the ablated datasets separately.

**Brief Analysis:** Figure 20 shows the quantitative segmentation results and the ARP/ARR scores for over/under-segmentation. Figure 18 qualitatively presents the evolution of segmentation performance from IODINE according to the adjustment in object-level complexity. More qualitative results for all approaches can be found in Fig. 40 in the appendix. We can see that: (1) Once the pixels of real-world objects are replaced by their mean colors, i.e., no color gradients, the object segmentation performance has been significantly improved for almost all methods. The ARP and ARR scores tend to be closer with  $-C$  ablation, indicating both over- and under-segmentation are alleviated. (2) Reducing the irregularity of real-world objects can also improve object segmentation, although not significantly. Simplifying object shapes does not effectively reduce the gap between ARP and ARR scores. In most cases, the over/under-segmentation issues remain the same as on the unablated datasets. (3) Overall, these results show that existing methods are more likely to learn the objectness represented by uniform colors and/or regular objects. However, compared with Fig. 11, the results of current ablated datasets in Fig. 20a still lag behind the synthetic datasets. This means that there should be some other factors that also potentially affect the object segmentation of existing models.

### 4.3 How do Scene-Level Factors Affect Current Models?

In this section, we investigate to what extent the distributions of scene-level factors affect the segmentation performance



**Fig. 23** Selected textures

with three ablative experiments. Image examples of scene-level ablations can be found in Fig. 36 in the appendix.

- *Ablation of Inter-object Color Similarity:* In each image of datasets of Group 2 (MOVIC/YCB/ScanNet/COCO), we replace all object textures with a set of new distinctive textures from DTD database (Cimpoi et al., 2014), as shown in Fig. 23. In this way, the multiple objects look more distinctive in appearance, while the per-object texture gradients are roughly preserved. The ablated datasets are denoted as: MOVIC-T/YCB-T/ScanNet-T/COCO-T.
- *Ablation of Inter-object Shape Variation:* In each image of Group 2 datasets, we normalize the scales of multiple objects by shrinking or expanding the diagonal length of their bounding boxes, such that the new object sizes tend to be uniform. For each object, its shape and texture are linearly scaled up or down. Basically, this aims to



remove the diversity of object sizes within single images. The ablated datasets are denoted as: MOVIC-U/YCB-U/ScanNet-U/COCO-U.

- *Ablation of both Inter-object Color Similarity and Shape Variation*: We simply combine the above two ablation strategies for images. The ablated datasets are denoted as: MOVIC-T+U/YCB-T+U/ScanNet-T+U/COCO-T+U.

Figure 19(1-a)–(1-d) shows new distributions of object-level and scene-level complexity factors for -T ablated datasets. We can see that the distributions of Inter-object Color Similarity in Fig. 19(1-c) becomes more similar to synthetic datasets, while the distributions of other three factors are still similar to original ones, i.e., Fig. 19(1-a), (1-b), (1-d) being similar to Fig. 16(1-a), (1-b), (1-d).

Figure 19(2-d) shows that the distributions of Inter-object Shape Variation for ablated datasets now become similar to the synthetic datasets, whereas the distributions of other three factors are still the same as the original ones, i.e., Fig. 19(2-a)–(2-c) being similar to Fig. 16(1-a)–(1-c). Note that, for the ablation of -T+U datasets, the distributions are the same as shown in Fig. 19(1-a), (2-b), (1-c), (2-d). Having these three groups of scene-level ablated real-world datasets, we then train and evaluate our four unsupervised baselines from scratch on each of the ablated datasets separately.

**Brief Analysis:** Figure 20 shows the quantitative segmentation results and ARP/ARR/ARI scores for over-/under-segmentation. Figure 21 presents qualitative performance of SlotAtt according to the decrease of scene-level complexity. More qualitative results for all methods can be found in Fig. 41 of the appendix. We can see that: (1) Once the textures of real-world images are replaced by more distinctive textures, i.e., with a lower similarity between object appearances, the segmentation performance has been surprisingly boosted remarkably for almost all methods. For MONet and IODINE, the gap between ARR and ARP scores is effectively reduced, indicating a mitigation of over-segmentation (for MONet) and under-segmentation (for IODINE). (2) Normalizing object sizes over images can also reasonably improve the segmentation performance for AIR through the alleviation of over-segmentation. Its effectiveness, however, is less obvious for other three approaches. (3) Overall, these results clearly show that existing unsupervised models significantly favor objectness with distinctive appearances in single images. However, compared with Fig. 11, the results on current scene-level ablated datasets are still inferior to synthetic datasets, meaning that the scene-level factors alone are not enough to explain the performance gap.

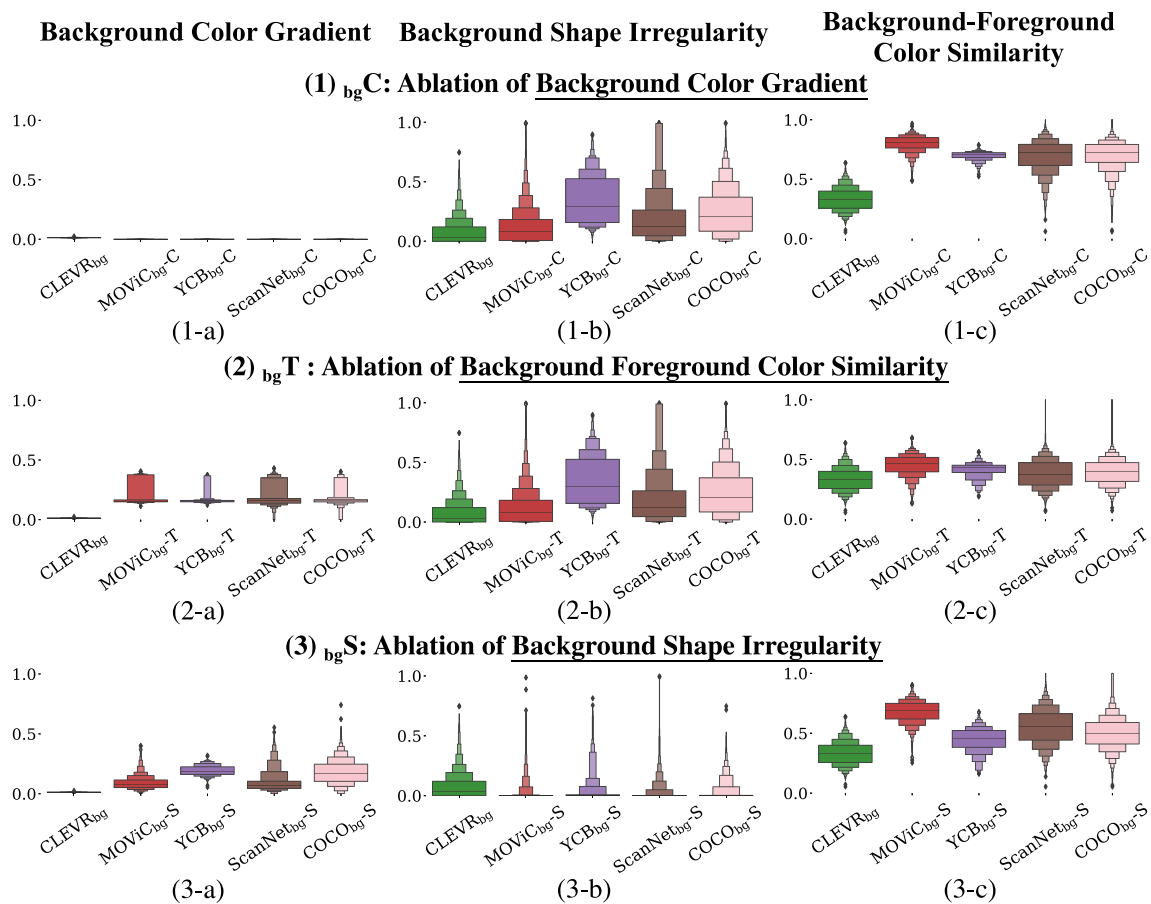
#### 4.4 How do Object- and Scene-Level Factors Jointly Affect Current Models?

In this section, we aim to study how the object- and scene-level factors jointly affect the segmentation performance. We conduct the following three ablative experiments. Image examples of joint ablations are in Fig. 37 in the appendix.

- *Ablation of Object Color Gradient, Object Shape Concavity, and Inter-object Color Similarity*: In each image of Group 2 datasets (MOVIC/YCB/ScanNet/COCO), we replace the object color by averaging all pixels of the distinctive texture, and also replace the object shape with a simple convex hull. The ablated datasets are denoted as: MOVIC-C+S+T/YCB-C+S+T/ScanNet-C+S+T/COCO-C+S+T.
- *Ablation of Object Color Gradient, Object Shape Concavity, and Inter-object Shape Variation*: In each image of the Group 2 datasets (MOVIC/YCB/ScanNet/COCO), we replace the object color by averaging its own texture, and modify the object shape as convex hull following by size normalization. The ablated datasets are denoted as: MOVIC-C+S+U/YCB-C+S+U/ScanNet-C+S+U/COCO-C+S+U.
- *Ablation of all four factors*: We aggressively combine all four ablations and get datasets: MOVIC-C+S+T+U/YCB-C+S+T+U/ScanNet-C+S+T+U/COCO-C+S+T+U.

Since these ablations are conducted independently, the new distributions of four complexity factors on current jointly ablated datasets are the same as Figs. 17(1-a), (2-b), and 19(1-c), (2-d). We train and evaluate the four baselines from scratch on each of the ablated datasets separately.

**Brief Analysis:** Figure 20 shows the quantitative segmentation results and the ARP/ARR/ARI scores for comparing over-/under-segmentation. Figure 22 presents qualitative performance of MONet according to the decrease of object-level and scene-level complexity. More qualitative results for all methods can be found in Fig. 42 in the appendix. We can see that: (1) Combining the two object-level factors with scene-level color ablation (-C+S+T) can significantly improve segmentation results, especially for SlotAtt as shown in Fig. 20a. Their ARP and ARR scores tend to be closer, suggesting an alleviation of over-/under-segmentation. Combining the two object-level factors with scene-level shape ablation (-C+S+U), however, can have a notable effect only on AIR, where its over-segmentation issue is relieved. (2) If the challenging real-world objects and images are ablated in



**Fig. 24** Distributions of three background-level factors on three types of ablation datasets in Sect. 4.5

both object- and scene-level, the segmentation performance of all unsupervised models achieves the same level with three synthetic datasets as shown in Fig. 11. (3) Overall, these three groups of experiments demonstrate that the failure of unsupervised models on real-world images involves both object- and scene-level dataset biases.

#### 4.5 How do Background Factors Affect Current Models?


In this section, we investigate to what extent the distributions of background-level factors affect the segmentation performance by conducting the following ablative experiments.

- *Ablation of Background Color Gradient*: For each image of datasets in Group 4, we replace the background pixels with its average color  $rgb$ . The four ablated datasets are denoted as:  $MOVIC_{bg-C}/YCB_{bg-C}/ScanNet_{bg-C}/COCO_{bg-C}$ .
- *Ablation of Background-Foreground Color Similarity*: For each image of Group 4 datasets, we replace the background texture with a new one from DTD database as

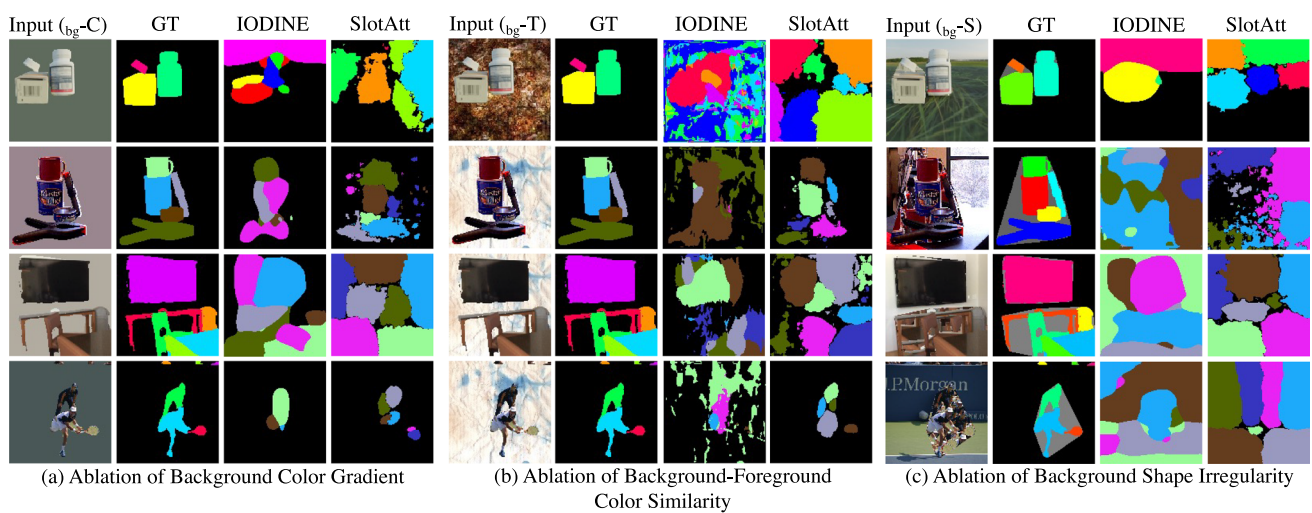
shown in Fig. 23. We select the most distinctive color against foreground pixels. In this way, the foreground objects are more distinguishable against the background, while the background gradients are roughly preserved. The three ablated datasets are denoted as:  $MOVIC_{bg-T}/YCB_{bg-T}/ScanNet_{bg-T}/COCO_{bg-T}$ .

- *Ablation of Background Shape Irregularity*: For each image, we first find all connected subcontours of its background. For each region enclosed by a subcontour, we find the smallest convex hull (Eddins, 2011) that surrounds it. The original enclosed regions are enlarged to their corresponding convex hulls and filled by shifting around the original region pixels. In this way, the background shape becomes more regular, while the background appearance keeps the same. The three ablated datasets are denoted as:  $MOVIC_{bg-S}/YCB_{bg-S}/ScanNet_{bg-S}/COCO_{bg-S}$ .

Figure 24 shows new distributions of the three background-level factors for the ablated datasets above. Particularly, subfigures (1-a)/(1-b)/(1-c) compares new distributions for the ablated datasets  $MOVIC_{bg-C}/YCB_{bg-C}/ScanNet_{bg-C}/$



**Fig. 25** Quantitative results of IODINE&SlotAtt on four datasets from Group 4 (with backgrounds) and their ablated datasets



**Fig. 26** Qualitative object segmentation results of SlotAtt and IODINE on four datasets of Group 4 ( $MOVIC_{bg}/YCB_{bg}/ScanNet_{bg}/COCO_{bg}$ ) with backgrounds

$COCO_{bg}-C$ . We can see that their Background Color Gradient decreases to zero (Subfig 1-a), and their Background Shape Irregularity (Subfig 1-b) remains the same as original datasets  $MOVIC_{bg}/YCB_{bg}/ScanNet_{bg}/COCO_{bg}$  (Subfig (2-b) in Fig. 16). Comparing Figs. 24(1-c) and 16(2-c), we observe the Background-Foreground Color Similarity is increased inevitably as a byproduct.

Figure 24(2-a)/(2-b)/(2-c) represents new distributions for the ablated datasets  $MOVIC_{bg}-T/YCB_{bg}-T/ScanNet_{bg}-T/COCO_{bg}-T$ . We can see that their Background-Foreground Color Similarity (Subfig 2-c) decreases, while Background Color Gradient (Subfig 2-a) and Background Shape Irregularity (Subfig 2-b) are similar to the original  $MOVIC_{bg}/YCB_{bg}/ScanNet_{bg}/COCO_{bg}$ . Figure 24(3-a)/(3-b)/(3-c) represents new distributions for the ablated datasets  $MOVIC_{bg}-S/YCB_{bg}-S/ScanNet_{bg}-S/COCO_{bg}-S$ . We can see that their Background Shape Irregularity (Subfig 3-b) decreases, while

the Background Color Gradient (Subfig 3-a) and Background-Foreground Color Similarity (Subfig 3-c) remain similar to the original  $MOVIC_{bg}/YCB_{bg}/ScanNet_{bg}/COCO_{bg}$ . We train and evaluate SlotAtt and IODINE from scratch on each of the ablated datasets. Note that, the performance of background segmentation is separately measured by a Recall score, denoted as BG Recall, where a predicted background mask is considered correct if its IoU against a ground truth background mask is above 0.5.

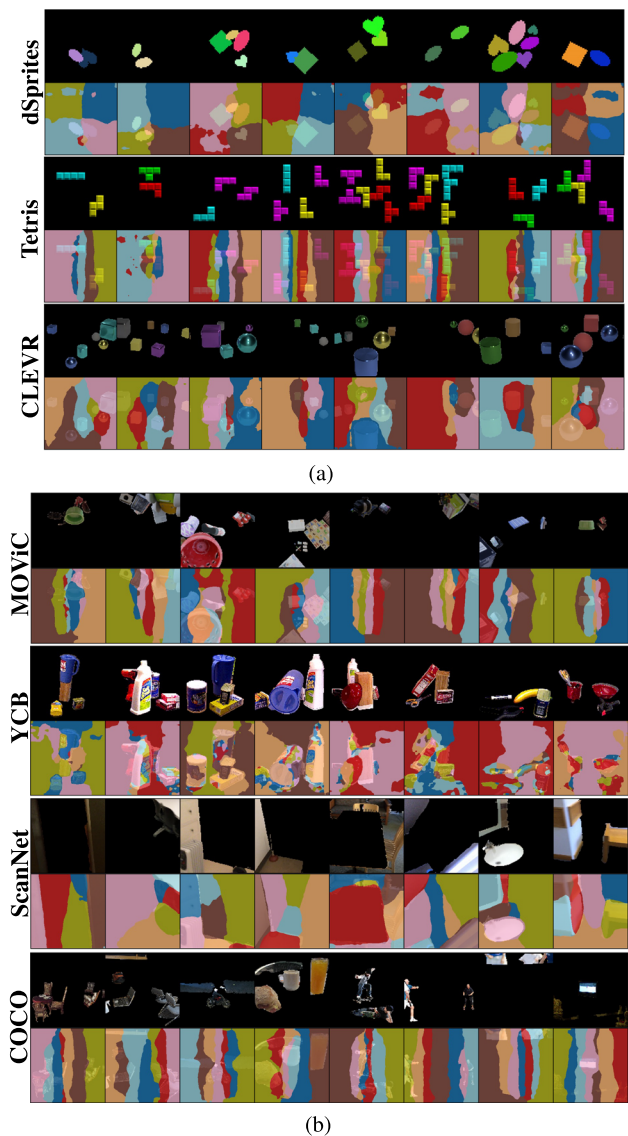
**Brief Analysis:** Figures 25 and 26 show the quantitative and qualitative results respectively. We can see that: (1) Removing the Background Color Gradient or replacing the background texture with a more discriminative one against foreground objects can largely increase BG Recall score. (2) Making the background contours to be more regular alone can hardly benefit the segmentation of objects and backgrounds. (3) Overall, the background can be easily rec-

ognized if it is of simple and discriminative color, which however cannot fundamentally alleviate the difficulty of segmenting individual foreground objects. This further confirms that the four object- and scene-level complexity factors introduced in Sects. 2.1 and 2.2 play an essential role in object segmentation of existing models.

#### 4.6 Why do Current Unsupervised Models Fail on Real-World Datasets?

As demonstrated in Sects. 4.2/4.3/4.4, once the object- and scene-level complexity factors are removed from the challenging real-world datasets, existing unsupervised models can perform as excellent as on the synthetic datasets, as qualitatively illustrated in Fig. 42. From this, we can safely conclude that the inductive biases designed in existing unsupervised models are far from able to match with and fully capture the true and complex objectness biases exhibited in real-world images. Nevertheless, from Fig. 20, each baseline tends to favor different objectness biases. In particular,

- AIR (Eslami et al., 2016): As a factor based model, AIR has a strong spatial-locality bias. Despite its poor segmentation performance across all datasets, there is a notable improvement when inter-object shape variation is ablated from real-world datasets (U/T+U/C+S+U/C+S+T+U). More convincingly, even when all other three factors are ablated (C+S+T), it can be hardly improved, showing that object shape variation is a key factor for AIR. Qualitatively (Fig. 12) and quantitatively (Fig. 13), AIR tends to use multiple bounding boxes for single real-world objects, which leads to over-segmentation. We can see from Fig. 20b that the uniform-scale ablation, especially –T+U datasets, can effectively decrease the gap between ARP and ARR scores, which implies an alleviated over-segmentation. Since AIR is designed to attend and infer objects as bounding boxes and does not explicitly model backgrounds, the background-level factors are not studied here.
- MONet (Burgess et al., 2019): MONet is more sensitive to color-related factors than shape-related factors. The ablations of object color gradient and inter-object color similarity significantly improve its performance, while ablations of object shape concavity and inter-object shape variation make little difference. For the two color-related factors, the scene-level one is more important than the object-level factor. From this, we can see that MONet has a strong dependency on color. Similar colors tend to be grouped together while different colors are separated apart. From Fig. 20b, we observe that the gaps between ARP and ARR scores are relatively small for

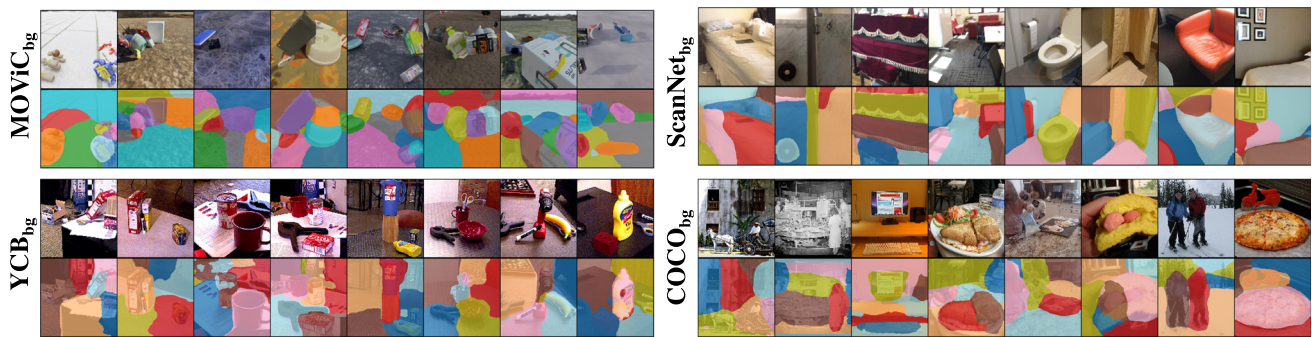


**Fig. 27** Qualitative object segmentation results of DINOSAUR. **a** On three datasets (dSprites/Tetris/CLEVR) from Group 1 with blank backgrounds; **b** On four datasets (MOVIC/UCB/ScanNet/COCO) from Group 2

all experiments on MONet, even when the AP scores are low. This suggests there are both over-segmentation and under-segmentation in the results. An object with various colors tends to be over-segmented, and different objects with similar colors are likely to be grouped together. Due to its heavy reliance on color clues, MONet cannot identify backgrounds without regular and distinctive colors. Thus, background-level factors are not studied for it.

- IODINE (Greff et al., 2019): IODINE also has a heavy dependency on object-/scene-level color-related factors. However, different from MONet, the ablation on object color gradient brings better performance





**Fig. 28** Qualitative object segmentation results of DINOSAUR on four datasets (MOVIC<sub>bg</sub>/YCB<sub>bg</sub>/ScanNet<sub>bg</sub>/COCO<sub>bg</sub>) from Group 4 with backgrounds

than inter-object color similarity. We speculate it is because the regularization on shape latent alleviates under-segmentation by biasing towards more regular shapes. From Fig. 20b, we observe both under-/over-segmentation, and the gaps between ARR and ARP scores are not obvious in most cases. Unsurprisingly, as shown in Fig. 25, IODINE cannot segment real backgrounds. The ablations of background-level color related factors do improve background-foreground separation, showing that it relies on colors to identify foregrounds from backgrounds. However, the successful background-foreground separation does little help in segmenting individual foreground objects, since the complex biases in real objects still stand.

- SlotAtt (Locatello et al., 2020): The ablations on all four factors increase the performance of SlotAtt at different levels, among which object- and scene-level color-related factors are more significant. We conjecture that it is because the feature embeddings used by Slot Attention module are learnt from both pixel colors and coordinates, which contributes to its sensitivity to both shape and color factors. By comparing ARP and ARR scores in Fig. 20, it is observed that MOVIC and COCO are more prone to under-segmentation while YCB and ScanNet are more prone to over-segmentation. From Fig. 16, it is observed that MOVIC and COCO have less distinguishable colors while YCB and ScanNet have more complex shapes. This also illustrates both color and shape are important factors for SlotAtt. As shown in Fig. 25, SlotAtt is not able to segment real-world backgrounds. The background-foreground separation in SlotAtt is particularly sensitive to the color contrast between the background and foreground. However, similar to IODINE, a distinguishable background cannot alleviate the burden of foreground object segmentation, as it does not fundamentally remove the complex object- and scene-level factors in real-world images.

**Table 2** Quantitative object segmentation results of DINOSAUR on four datasets (MOVIC<sub>bg</sub>/YCB<sub>bg</sub>/ScanNet<sub>bg</sub>/COCO<sub>bg</sub>) of Group 4 with backgrounds

	FG-ARI	mBO	AP	PQ	Pre	Rec
MOVIC <sub>bg</sub>	64.9	39.2	9.2	16.9	21.6	32.9
YCB <sub>bg</sub>	52.0	24.9	1.0	3.0	4.2	6.1
ScanNet <sub>bg</sub>	56.2	39.1	7.6	17.2	21.0	33.3
COCO <sub>bg</sub>	51.1	27.0	1.7	7.6	9.7	16.6

All scores are in percentage (%)

#### 4.7 Can Models Pretrained on Monolithic Object Images Segment Real-World Images?

The models pretrained on monolithic object images have recently achieved promising results on real-world datasets. FreeSOLO (Wang et al., 2022) and CutLER (Wang et al., 2023) generate pseudo labels from pretrained features and train detectors interactively with pseudo labels. Odin (Hénaff et al., 2022) and DINOSAUR (Seitzer et al., 2022) make use of features pretrained on ImageNet (Russakovsky et al., 2015) and apply different clustering strategies. Among this new line of works, we select the representative DINOSAUR to evaluate its segmentation performance on datasets of Groups 1/2/4.

**Brief Analysis:** As shown in Fig. 27, DINOSAUR exhibits clear limitations on datasets in Groups 1&2 where all images just have blank backgrounds. Primarily, this is because DINOSAUR relies heavily on pretrained ViT features, specifically those extracted from models like DINO (Caron et al., 2021) trained on the ImageNet dataset. Such features can hardly generalize to images in Groups 1&2 datasets where the synthetic objects and blank backgrounds are significantly different from those in ImageNet.

As shown in Fig. 28 and Table 2, as expected, DINOSAUR demonstrates better and reasonable object segmentation results on the 4 datasets in Group 4. In particular, it shows favorable FG-ARI and mean Best Overlap (mBO) scores, but tends to over-segment objects when the number of available

slots exceeds the number of objects present within images. Moreover, it often fails to identify backgrounds as complete entities, instead segmenting them into fragmented pieces. That is why the precision score tend to be lower than the recall score, and the AP score is further diminished.

Overall, by leveraging the large-scale pretrained real-world image features, DINOSAUR shows promising performance to segment objects in real-world images thanks to the feature grouping capability of slot attention based mechanisms (Yu et al., 2022; Xu et al., 2022). Nevertheless, it is still in its infancy to accurately identify generic objects and separate complex backgrounds. We conjecture that more explicit object biases need to be encoded into the pretrained models or the unsupervised learning process, although the existing pretrained features on monolithic object images may implicitly have the concepts of objectness. More advanced research works are expected in this direction in the future.

## 5 Conclusions

We systematically show that existing unsupervised methods are practically impossible to segment generic objects from single real-world images, and investigate the underlying factors that incur the failure. With the aid of our carefully designed seven object-, scene-, and background-level complexity factors, we conduct extensive experiments on multiple groups of ablated real-world objects and images, and safely conclude that the distributions of both object-, scene-, and background-level biases in appearance and geometry of real-world datasets are particularly diverse and indiscriminative, such that current unsupervised models cannot segment real objects or backgrounds. Based on this finding, we suggest two main directions for future study: (1) To exploit more discriminative objectness biases such as object motions which expressively describe the ownership of visual pixels as recently explored in Bear et al. (2020), Chen et al. (2022) and Tangemann et al. (2021), for 2D images and in Song and Yang (2022) for 3D point clouds. (2) To leverage pretrained features from single-object-dominant datasets which explicitly regard each image as an object as recently studied in Caron et al. (2021), Hénaff et al. (2022) and Seitzer et al. (2023), although such settings are no longer purely unsupervised.

**Acknowledgements** This work was supported in part by National Natural Science Foundation of China (62271431), In part by Shenzhen Science and Technology Innovation Commission (JCYJ20210324120603011), in part by Research Grants Council of Hong Kong (25207822).

**Funding** Open access funding provided by The Hong Kong Polytechnic University

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as

long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

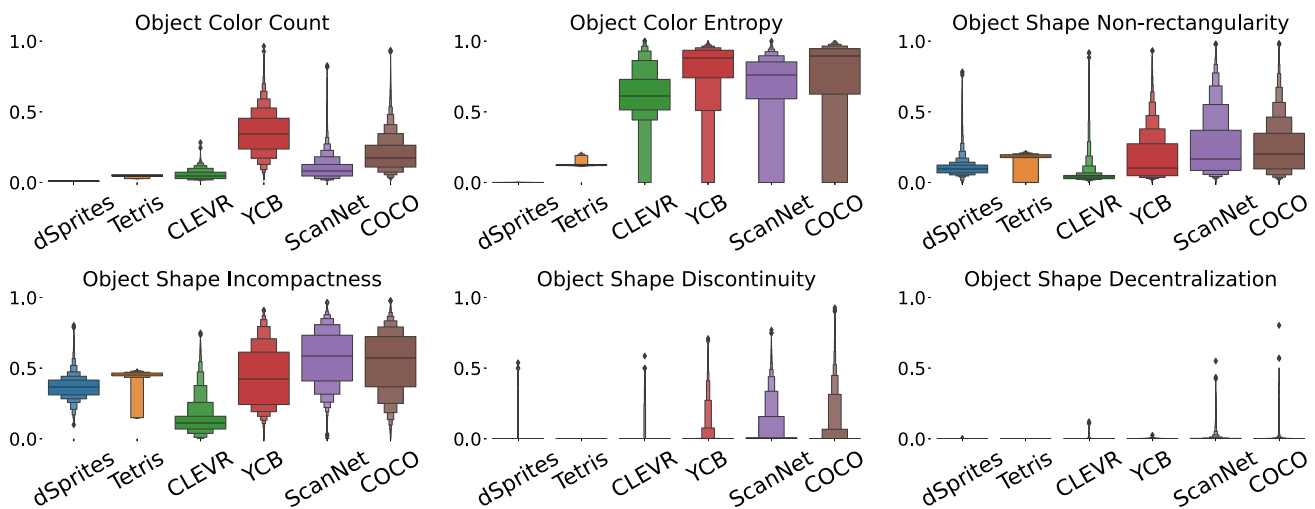
## Appendix

### Appendix A.1: Other Candidates of Object- and Scene-Level Complexity Factors

In addition to the primary four complexity factors in Sects. 2.1 and 2.2, we also explore other potential complexity factors to quantitatively measure the distributions of object- and scene-level biases in appearance and geometry. Basically, we aim to consider as many aspects as possible to investigate key factors underlying the distribution gaps between synthetic and real-world datasets. However, we empirically find that these candidate factors do not show significant discrepancies between synthetic and real-world datasets. Details are shown below.

#### Appendix A.1.1: Candidates of Object-Level Complexity Factors

- *Object Color Count*: This factor is defined as the total number of unique colors within an object mask. Basically, this is to simply measure the diversity of object colors.
- *Object Color Entropy*: Inspired by Shannon entropy (Shannon, 1948), we calculate the entropy value at each pixel by applying a  $3 \times 3$  filter on the grayscale image converted from RGB. In particular, for each pixel, its color value becomes a discrete value in  $[0, 255]$ . We compute its entropy score:  $H(x) = -\sum_{i=1}^n p(x_i) \log_2 p(x_i)$ , where  $p(x_i)$  denotes the probability of a specific color value  $x_i$  within the  $3 \times 3$  neighbourhood. Basically, this factor aims to measure the color diversity within  $3 \times 3$  image patches. The higher this factor, the more frequently the object color changes in small local areas.
- *Object Shape Non-rectangularity*: Given the binary mask of an object ( $\mathbf{M}_{obj} \in \mathbb{R}^{H \times W}$ ), we first calculate its axis-aligned bounding box ( $\mathbf{M}_{bbox} \in \mathbb{R}^{H \times W}$ ). Object shape non-rectangularity is calculated as  $1 - \sum \mathbf{M}_{obj} / \sum \mathbf{M}_{bbox}$ . Similar to object shape concavity, this factor is also designed to measure the complexity of object shapes. However, this factor is more likely to be affected by the object orientation since it takes axis-aligned bounding boxes as a reference.



**Fig. 29** Distributions of additional candidates of object-level complexity factors

- *Object Shape Incompactness*: There are two similar methods to quantify the compactness of object shapes. The first one is Polsby–Popper test (Polsby & Popper, 1991):  $PP(M_{obj}) = 4\pi A(M_{obj}) / P(M_{obj})^2$ . The other is Schwartzberg (1965) compactness score:  $S(M_{obj}) = (2\pi \sqrt{A(M_{obj}) / \pi}) / P(M_{obj})$ . In both formula,  $P(M_{obj})$  is the object perimeter and  $A(M_{obj})$  is the object area. For simplicity, we choose  $PP(M_{obj})$  to calculate the object shape incompactness score:  $1 - PP(M_{obj})$ .
- *Object Shape Discontinuity*: Given an object mask ( $M_{obj} \in \mathbb{R}^{H \times W}$ ), we first find the largest connected component ( $M_{lcc} \in \mathbb{R}^{H \times W}$ ) in its binary mask. The discontinuity of shape is calculated as:  $1 - \sum M_{lcc} / \sum M_{obj}$ . This factor is to evaluate how continuous an object shape is.
- *Object Shape Decentralization*: Given an object mask, we first calculate its centroid ( $\bar{x}, \bar{y}$ ) by averaging all pixel coordinates in the object. Then, the second moment of this object is calculated as:  $\sum_x \sum_y (x - \bar{x})^2 (y - \bar{y})^2$ , where  $(x, y)$  is the coordinates of pixels within the object. The higher this factor, the object shape is less likely to be centralized.


As shown in Fig. 29, we compare the distributions of the object-level factor candidates on both synthetic and real-world datasets. It can be seen that the majority of these factors do not show significant gaps between the simple synthetic and the challenging real-world datasets. Therefore, we do not conduct relevant ablation experiments.

#### Appendix A.1.2: Candidates of Scene-Level Complexity Factors


- *Inter-object Color Similarity with Chamfer Distance*: In the calculation of this factor, we first convert each pixel into a point in RGB space. In this way, each object can be

represented by a point set in the RGB space. This factor is calculated between each pair of objects by measuring the Chamfer distance of two point sets in the RGB space. Since Chamfer distance is an asymmetric measurement, we calculate and average out the bidirectional Chamfer distances. Compared with Euclidean distance, this measurement favors the most similar colors between two objects.

- *Inter-object Color Similarity with Hausdorff Distance*: This factor is similar to the previous one. The only difference is that we replace Chamfer distance with Hausdorff distance. Hausdorff distance is also a directed and asymmetric measurement, so the final score is the average of distance values in both directions.
- *Inter-object Shape Similarity over Boundaries*: For each object mask, we first find its boundary using the method in Cheng et al. (2021), and then crop it with its axis-aligned bounding box. Each bounding box is scaled and fits into a unit box with its original aspect ratio. Lastly, we calculate the IoU between the boundaries of two objects to measure their shape similarity.
- *Inter-object Shape Entropy between Boundaries*: We first combine all object masks into a single image by assigning different indices to different objects. Then we compute the entropy of each pixel with a  $3 \times 3$  filter. The final factor score is calculated by averaging all non-zero entropy values. Note that, the interior part of objects and background will not be considered because their entropy values will always be zeros. Basically, this factor is designed to evaluate how crowded an image is. The higher this factor, the more objects are spatially adjacent.
- *Inter-object Proximity between Centroids*: We first calculate the centroid ( $\bar{x}, \bar{y}$ ) of each object by averaging all pixel coordinates in the object mask. Euclidean distances between object centroids are then computed pair-wisely before they are averaged to be the final factor score. This



**Fig. 30** Distributions of additional candidates of scene-level complexity factors



**Fig. 31** Distributions of additional candidates of inter-Object area variation

factor is designed to measure the spatial proximity of multiple objects in a single image.

- *Inter-object Proximity with Chamfer Distance*: To measure the spatial proximity between objects, we also calculate the spatial Chamfer distance between objects. Specifically, each object is represented by a set of  $(h/w)$  coordinates, and the average of pair-wise bidirectional Chamfer distances is calculated as the proximity score for each image.
- *Inter-object Area Variation*: In order to measure the variation of objects in terms of their scale, We first calculate the area of each object, and then compute the pair-wise absolute difference for all object areas, obtaining a  $K \times K$  matrix. The final inter-object area variation is the average of the matrix excluding diagonal entries.

As shown in Fig. 30, we compare the distributions of the scene-level candidate factors on both synthetic and real-world datasets. We can see that both the inter-object color similarity with Chamfer and Hausdorff distances share similar distribution gaps with our primary inter-object color similarity factor defined in Sect. 2. The remaining four candidate factors relating to inter-object shape complexity do not show significant distribution gaps between the synthetic

and real-world datasets. In this regard, we choose not to conduct ablation experiments on these six candidate factors. The distribution for Inter-object Area Variation is presented in Fig. 31. Overall, the inter-object area variation for synthetic datasets is higher than real-world datasets. However, comparing Inter-object Area Variation in Fig. 31 and Inter-object Shape Variation in Fig. 2(1-d), we find that factor values for Inter-object Shape Variation have smaller variation within each dataset, indicating it better captures the scene-level property of datasets. Besides, Inter-object Shape Variation does not only take into account the scale of objects but also



the orientations. Thus, we select Inter-object Shape Variation over Inter-object Area Variation.

### Appendix A.1.3: Complexity Factors Measured Across Dataset

Object-level complexity factors are calculated for each object as shown in Fig. 32(1-a), (1-b). Scene-level complexity factors are calculated for each image as shown in Fig. 32(1-c), (1-d). We additionally measure the four complexity factors across the datasets as shown in Fig. 32(2-a)–(2-d). To measure object-level complexity factors across the datasets, we directly calculate the average of all object complexity values. To measure the scene-level complexity factors across the datasets, we make the color/shape comparison across the dataset and calculate the factor value. We find that complexity factors calculated across the dataset follow similar patterns as those calculated across the object/image.

### Appendix A.2: More Details of Background-Level Complexity Factors

- *Background-Foreground Color Similarity*: Given all background pixels and the foreground pixels in a single image, we calculate the Euclidean distance between each background pixel and each foreground pixel in the RGB space. This results in a  $U \times V$  matrix  $E$  where  $U$  and  $V$  represent the number of pixels in background and foreground respectively. We treat  $255\sqrt{3} - E$  as the cost matrix for a Hungarian algorithm and solve the optimal assignment between foreground pixels and background pixels. Each background pixel will be assigned to a distant foreground pixel in the RGB space. Given the assignment between foreground pixels and background pixels, the background-foreground color distance is calculated as the Euclidean distance between assigned pairs. The final factor value for background-foreground color similarity is computed as:  $1 - \text{background-foreground color distance}/255\sqrt{3}$ . The higher this factor, the more similar background and foreground appear to be, and it is harder to separate each other.
- *Background Shape Irregularity*: In order to compute the maximal inscribed convex set of a connected region, we follow the algorithm proposed by Borgefors and Strand (2005). The main idea is to iteratively find the deepest concavity and remove it with minimum area cost. Given a randomly shaped connected region  $R \in \mathbb{R}^{H \times W}$  (Fig. 33a), we first compute the smallest convex polygon mask (convex hull)  $C \in \mathbb{R}^{H \times W}$  as shown in Fig. 33b. The convex deficiency  $D$  is defined as the difference between the region  $R$  and its convex hull  $C$ :  $D = C - R$ ,  $D \in \mathbb{R}^{H \times W}$  as shown in the blue region in Fig. 33c. Then we calculate the constrained distance transform (Piper & Granum, 1987) for the convex deficiency  $D$  where

the region  $R$  is the constraint. On the calculated distance transform map, the deepest concavity  $dc \in \mathbb{R}^2$  is the element with the largest distance label as shown in Fig. 33d. In order to remove the deepest concavity  $dc$ , we cut the region with a straight line from  $dc$  with 8 directions ( $\frac{n \times \pi}{4}$ ,  $n = 0, 1, \dots, 7$ ) as shown in Fig. 34]. The cut that generates the smallest sub-region is selected and the smallest sub-region is removed. The above process is repeated until the remained sub-region is convex ( $\text{distance label}(dc) \leq 3$ ). The remained convex region is the maximal inscribed convex set.

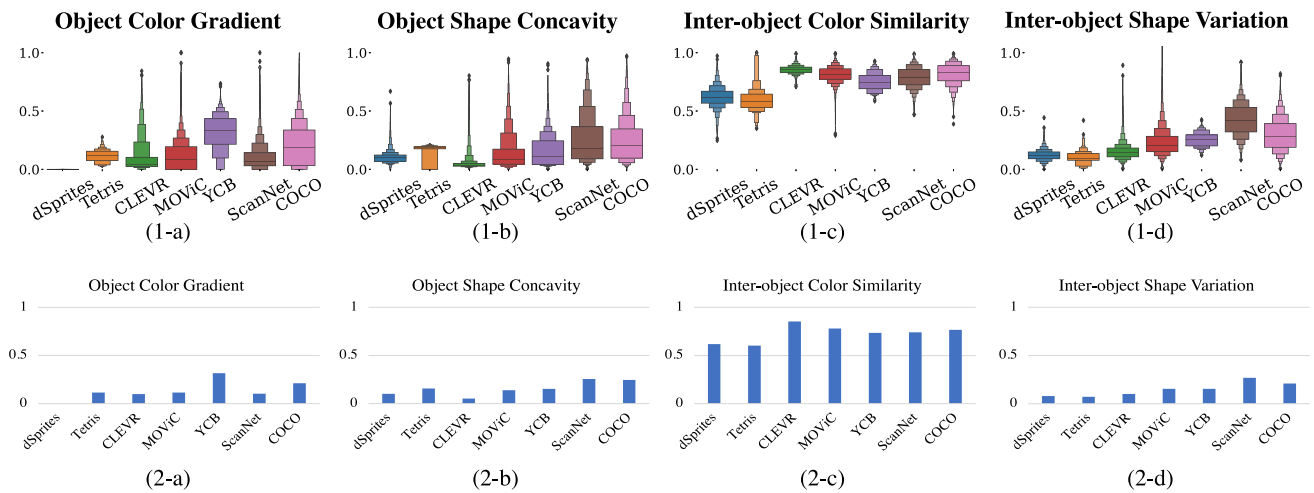
### Appendix A.3: Implementation Details of Baselines

*AIR* (Eslami et al., 2016)

- *Source Code*: We refer to <https://pyro.ai/examples/air.html> and <https://github.com/addtt/attend-infer-repeat-pytorch> for the implementation.
- *Important Adaptations*: We use an additional parameter to weight the KL divergence loss and reconstruction loss. For each experiment, we choose the weight for KL divergence from 1, 10, 25 and 50. The highest AP score is kept.
- *Training Details*: All experiments of AIR (Eslami et al., 2016) are conducted with a batch size of 64. The learning rate is set to  $1e-4$  for training inference networks and decoders,  $1e-3$  for baselines which are the same as the original paper. Since the number of objects in our datasets ranges between 2 and 6, we set the maximum number of steps at inference to be 6 for all experiments. All models are trained on a single GPU for 1000 epochs. We perform an evaluation every 50 epochs and select the highest AP score.
- *Implementation Details*: LSTMs have 256 cell units and object appearances are coded with 50 units. The hidden dimensions for appearance encoder and decoder are both 200. Images are normalized to hold values between 0 and 1 and the likelihood function is a Gaussian with a fixed standard deviation equal to 0.3. The presence prior  $p(n)$  is fixed to a geometric distribution that favors sparse reconstructions ( $p = 0.01$ ). The location prior is a normal distribution with mean  $[3, 0, 0]$  and standard deviation  $[0.2, 1, 1]$ , presenting the scale, x-position and y-position respectively. The appearance prior is a standard normal distribution with a dimension of 50.



*MONet* (Burgess et al., 2019)

- *Source Code*: We refer to (Engelcke et al., 2020)'s re-implementation at: <https://github.com/applied-ai-lab/genesis>.



**Fig. 32** Complexity calculated across object/image vs. complexity calculated across the dataset

**Fig. 33** Find the deepest concavity for a 2D shape



**Fig. 34** Find the cut from the deepest concavity that leads to minimum area reduction. The first row shows eight rays with pre-defined directions. The second row shows the sub-regions generated from cuts

- *Important Adaptations:* We train MONet with the GECO objective (Rezende & Viola, 2018) following the protocol mentioned in Engelcke et al. (2020).
- *Training Details:* All experiments on MONet are conducted with a batch size of 32 and a learning rate of  $1e-4$ . Since the maximum number of components is 7, including 1 background and 6 objects, we set the number of steps to be 7 for all experiments. All models are trained on a single GPU for 200 epochs with the training loss converged. We perform an evaluation every 10 epochs and select the highest AP score.
- *Implementation Details:* The VAE encoder is a standard CNN with  $3 \times 3$  kernels, stride of 2, and ReLU activa-

tions. It employs an attention network with  $[32, 32, 64, 64, 64]$  filters in the encoder and the reverse in the decoder. It receives the concatenation of the input image  $x$  and the attention mask in logarithmic units,  $\log m_k$  as input. The CNN layers output  $(32, 32, 64, 64)$  channels respectively. The CNN output is flattened and fed to a 2 layer MLP with output sizes of  $(256, 32)$ . The MLP output parameterises the  $\mu$  and  $\log \sigma$  of a 16-dim Gaussian latent posterior. The VAE uses a broadcast decoder (Watters et al., 2019) consisting of a four-layer CNN with no padding,  $3 \times 3$  kernels, stride 1, 32 output channels and ReLU activations. A final  $1 \times 1$  convolutional layer transforms the output into 4 channels. For all experiments, the output component

distribution is an independent pixel-wise Gaussian with fixed scales ( $\text{std} = 0.7$ ). The attention network is a standard UNet (Ronneberger et al., 2015) with five blocks. Each block consists of the following: a  $3 \times 3$  bias-free convolution with stride 1, followed by instance normalization with a learned bias term, followed by a ReLU activation, and finally downsampled or upsampled by a factor of 2 using nearest neighbour-resizing (no resizing occurs in the last block of each path).

#### IODINE (Greff et al., 2019)

- *Source Code*: We use the official implementation at: <https://github.com/deepmind/deepmind-research/tree/master/iodine>.
- *Important Adaptations*: The architecture is set the same as what is used for CLEVR dataset (Johnson et al., 2017) in the original paper (Greff et al., 2019).
- *Training Details*: Since we use a single GPU for the training of all models, the batch size is adjusted to be 4 and the learning rate  $0.0001 \times \sqrt{1/8}$ . The number of slots  $K$  is set as 7 and the inference iteration  $T$  as 5. We train each model for 500K iterations until the loss is fully converged.
- *Implementation Details*: The convolutional layers in the encoder and decoder have five layers with size-5 kernels, strides of [1, 2, 1, 2, 1], and filter sizes of [32, 32, 64, 64, 64] and [64, 32, 32, 32, 32], respectively. Fully-connected layers are used at the lowest resolution. The autoregressive prior is implemented as an LSTM with 256 units. The conditional distribution is parameterized by a multi-layer perceptron (MLP) with two hidden layers, 256 units per layer.

#### SlotAtt (Locatello et al., 2020)

- *Source Code*: We use the official implementation at: [https://github.com/google-research/google-research/tree/master/slot\\_attention](https://github.com/google-research/google-research/tree/master/slot_attention).
- *Important Adaptations*: The architecture is set the same as what is used for CLEVR dataset (Johnson et al., 2017) in the original paper (Locatello et al., 2020).
- *Training Details*: All experiments of SlotAtt are conducted with a batch size of 32 and learning rate selected from  $[4e-4, 4e-5]$ . The number of slots  $K$  is set as 7 and the number of iterations  $T$  is set as 3. All models are trained on a single GPU for 500K iterations until the loss is fully converged.
- *Implementation Details*: For all experiments, we use a slot feature dimension of 64. The GRU has a 64-dimensional hidden state and the feedforward block is an MLP with a single hidden layer of size 128 and ReLU activation followed by a linear layer. The CNN Encoder used in our experiments consists of 5 convolutional layers where each filter is  $5 \times 5$  with channel 64, and ReLU activations are used after each convolutional layer. After this backbone, we add position embeddings and then flatten the spatial dimensions. After applying a layer normalization, we finally add  $1 \times 1$  convolutions which we implement as a shared MLP applied at each spatial location with one hidden layer of 64 units. The spatial broadcast decoder (Watters et al., 2019) is applied independently on each slot representation with shared parameters between slots. We first copy the slot representation vector of dimension 64 onto a grid of shape  $8 \times 8 \times 64$ , after which we add a positional embedding. Finally, this representation is passed through several deconvolutional layers symmetrical to the encoder.

#### Mask R-CNN (He et al., 2017)

- *Source Code*: We use the implementation at: [https://github.com/matterport/Mask\\_RCNN](https://github.com/matterport/Mask_RCNN).
- *Important Adaptations*: We use the same settings as training for COCO in the above repository.
- *Training Details*: Training for all datasets starts from the pre-trained COCO weights (mask\_rcnn\_coco.h5) from [https://github.com/matterport/Mask\\_RCNN/releases](https://github.com/matterport/Mask_RCNN/releases). All models are trained on a single GPU for 30 epochs until the loss is fully converged.
- *Training Details*: The stride for the backbone is set to [4, 8, 16, 32, 64]. The maximum number of instances to detect in one image is 100. The minimum confidence is 0.5 and NMS ratio is 0.3. The scale of the RPN anchor is (32, 64, 128, 256, 512) and stride is 2.

### DINOSAUR (Seitzer et al., 2022)

- *Source Code*: We use the official implementation at: <https://github.com/amazon-science/object-centric-learning-framework>.
- *Important Adaptations*: We use a similar model architecture as for MOVIC dataset in the original paper. In light of the fact that our datasets only consist of 2–6 objects, we adapt the number of slots to be 7. And since we are using a single GPU for training, we adopt 64 as batch size.
- *Training Details*: We train DINOSAUR using the Adam optimizer with a learning rate of  $4e-4$ , a linear learning rate warm-up of 10 000 optimization steps and an exponentially decaying learning rate schedule. All models are trained on a single GPU with a batch size of 64.
- *Implementation Details*: We use a ViT with patch size 8 as the feature extractor and the MLP decoder for all datasets considered. The ViT-8 we used has the token dimensionality of 384 and 6 heads. All models use 12 Transformer blocks, linear patch embedding and additive positional encoding. The output of the last block (not applying the final layer norm) is passed on to the Slot Attention module and used in the feature reconstruction loss, after removing the entry corresponding to the CLS token. For the pre-trained weights, we use the timm library for DINO and the specific timm model name is `vit_small_patch8_224_dino`. We use a four-layer MLP with ReLU activations with hidden layer sizes of 1024.

2 criteria: (1) all objects shall be on the canvas with complete shapes; (2) all objects shall not overlap with each other.

*CLEVR* (Johnson et al., 2017) We first generate CLEVR images following <https://github.com/facebookresearch/clevr-dataset-gen>, where the number of objects per image is restricted between 3 to 6. Given generated images with a resolution  $640 \times 480$ , we perform center-cropping and then resize them to be  $128 \times 128$ . Then, we remove tiny objects which have less than 35 pixels from each image. Subsequently, the images with less than 2 objects are removed. Being consistent with the previous 2 synthetic datasets, all images have a black background.


## Appendix A.4: Details of Benchmark Datasets

In this section, we present the details of synthetic and real-world datasets. Figure 35 presents sample images from all four dataset groups. Table 3 qualitatively summarizes the biases of datasets in Groups 1/2/3/4.

*dSprites* (Matthey et al., 2017) To generate a specific image for this dataset, we first sample a random integer  $K$  from a uniform distribution with interval  $[2, 6]$  as the number of objects in that image. Then,  $K$  object shapes are selected from the binary dsprites dataset (Matthey et al., 2017) also in a uniformly random manner. Each object is assigned a random RGB color by sampling three random integers from a uniform distribution with interval  $[0, 255]$ . In total, we generate 10,000 images for training, and 2000 for testing.

*Tetris* (Kabra et al., 2019) For each image in this dataset, we first sample a random integer  $K$  from a uniform distribution with interval  $[2, 6]$  as the number of objects in this image. To render one tetris-like object onto the canvas, we randomly pick up a Tetris object from a randomly selected image from Kabra et al. (2019). Each object is resized to be  $88 \times 88$  and then placed onto the canvas. The position of each object is also sampled from a uniform distribution with





**Fig. 35** Samples from all datasets considered

**Table 3** The object-, scene-, and background-level biases in appearance and geometry of the datasets in Groups 1/2/3/4

	Synthetic datasets			Semi-realistic datasets	Real-world datasets		
	dSprites	Tetris	CLEVR/CLEVR <sub>bg</sub>	MOViC/MOViC <sub>bg</sub>	YCB/YCB <sub>bg</sub>	ScanNet/ScanNet <sub>bg</sub>	COCO/COCO <sub>bg</sub>
<i>Object-level Biases</i>							
<i>Appearance</i>							
Texture	Simple	Simple	Simple	Diverse	Diverse	Simple	Diverse
Lighting	No	No	Synthetic	Synthetic	Real	Real	Real
<i>Geometry</i>							
Shape	Simple	Simple	Simple	Simple	Simple	Diverse	Diverse
Occlusion	Minor	No	Minor	Severe	Severe	Severe	Severe
<i>Scene-level Biases</i>							
<i>Appearance</i>							
Similarity	Low	Low	High	High	High	High	High
<i>Geometry</i>							
Diversity	Low	Low	Low	Medium	High	High	High
<i>Background-level Biases</i>							
<i>Appearance</i>							
Texture	–	–	–/Simple	–/Simple	–/Diverse	–/Simple	–/Diverse
Lighting	–	–	–/Synthetic	–/Synthetic	–/Real	–/Real	–/Real
Similarity	–	–	–/Minor	–/High	–/High	–/High	–/High
<i>Geometry</i>							
Diversity	–	–	–/Simple	–/Simple	–/Complex	–/Complex	–/Complex

**MOViC** (Greff et al., 2022) We use the training and validation sets of the released MOVi dataset from [gs://kubric-public/tfds](https://kubric-public/tfds). Specifically, the downsampled variant at a resolution of  $128 \times 128$  of MOVi-C is used. In the original MOVi-C dataset, each image consists of  $3 \sim 10$  objects. Only images with  $2 \sim 6$  objects are kept. All background pixels are replaced by the black color.

**YCB** (Calli et al., 2017) We sample single frames from the YCB video dataset (Calli et al., 2017) every 20 images. Given the sampled frames with a resolution of  $640 \times 480$ , we first center crop and then resize them to be  $128 \times 128$ . Then, the images consisting of less than 2 or more than 6 objects

are removed. Similarly, all background pixels are replaced by a black color.

**ScanNet** (Dai et al., 2017) We sample single frames from the ScanNet dataset (Dai et al., 2017) every 20 images. Given the selected frames with a resolution of  $1296 \times 968$ , we first center crop the images with a size of  $800 \times 800$  and then resize them to be  $128 \times 128$ . For each resized image, we remove objects that contain more than  $128 \times 128 \times 0.2$  pixels or less than  $128 \times 128 \times 0.007$  pixels. The images with less than 2 or more than 6 are also dropped. All background pixels are replaced by the black color.

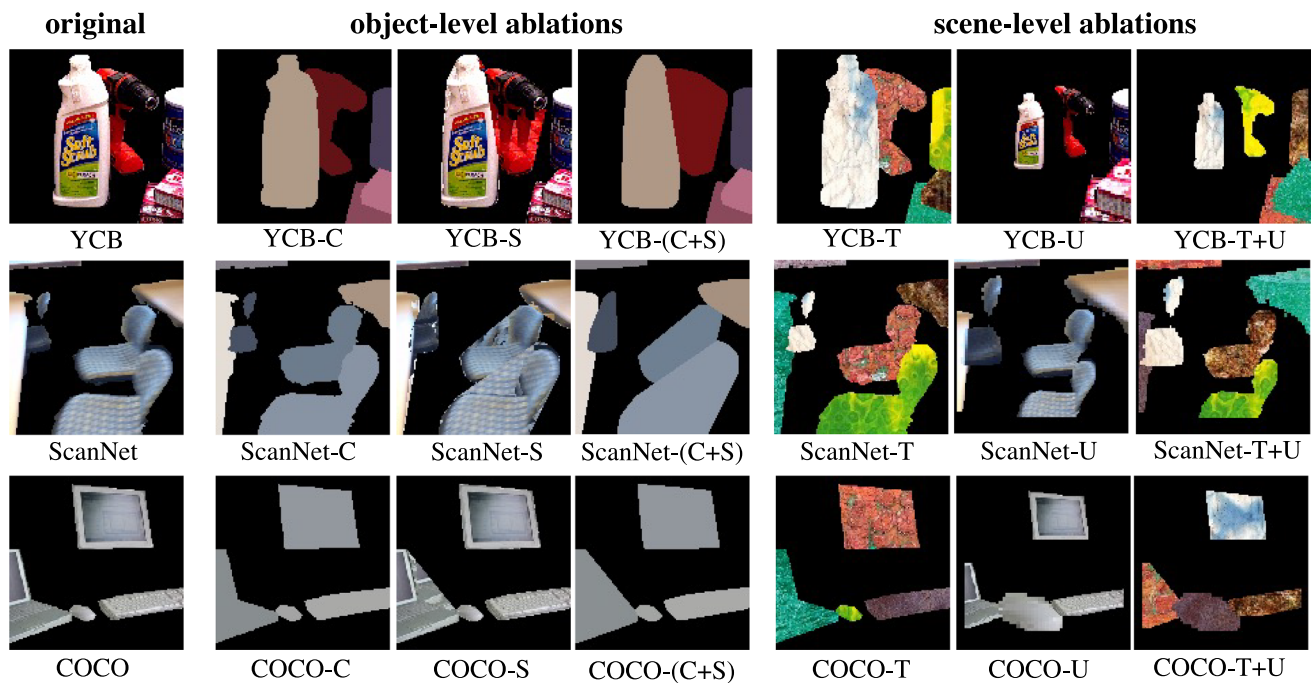
*COCO* (Lin et al., 2014) Given images in COCO-2017 (Lin et al., 2014) with various resolutions, we first center crop and then resize images to be  $128 \times 128$ . For each resized image, we use the same criteria applied to ScanNet to remove too-large and too-small objects. The images with  $2 \sim 6$  objects are kept. All background pixels are replaced by the

black color. Since the number of images that could meet all requirements is less than 2000 (only 1597) in the official validation split, we additionally select 403 different images from its official training split for testing. There is no overlap between our training and testing splits.

*CLEVR<sub>bg</sub>*, *MOViC<sub>bg</sub>*/*YCB<sub>bg</sub>*/*ScanNet<sub>bg</sub>*/*COCO<sub>bg</sub>* These five datasets are created in the same way as the above (CLEVR, MOViC/YCB/ScanNet/COCO) datasets, with synthetic or real-world backgrounds being added back to each image.

### Appendix A.5: Examples of the Ablated Datasets

Figure 36 presents before/after view for images ablated with object-level factors and scene-level factors respectively. Figure 37 shows examples of the joint ablations on both object-level and scene-level.



**Fig. 36** Example images of real-world datasets ablated with object-level and scene-level factors

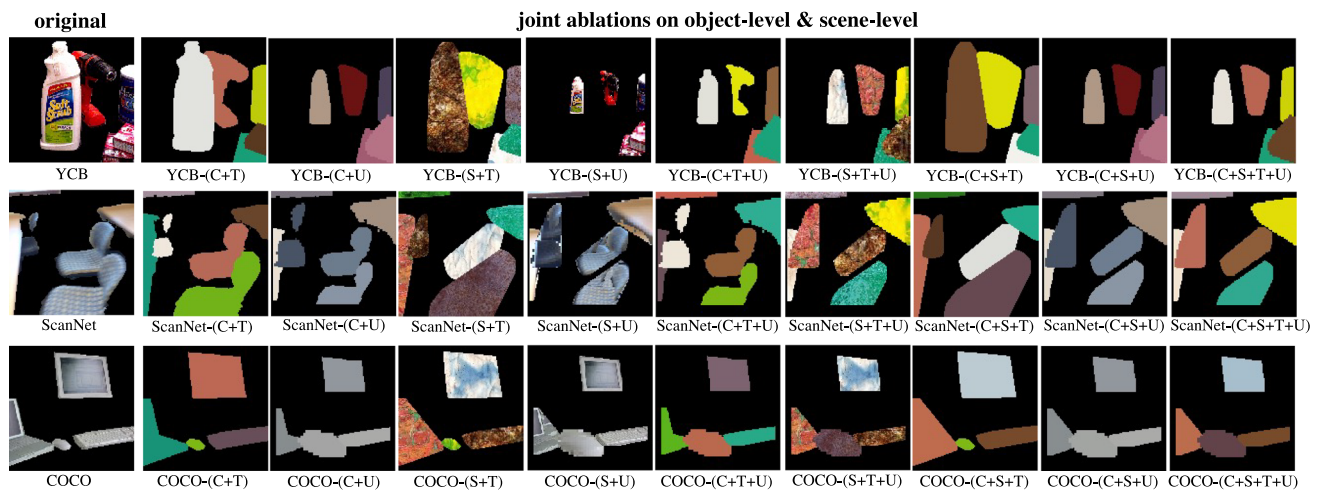


Fig. 37 Example images of real-world datasets ablated with object-level and scene-level factors (continued)

## Appendix A.6: Details of Experimental Results in Sect. 4.1

Tables 4 and 5 present the detailed results of experiments conducted in Sect. 4.1. Standard deviations for all scores are calculated over 3 independent runs.

**Table 4** Quantitative results of object segmentation from the five methods on datasets of Groups 1&2 with blank backgrounds

	dSprites AP/PQ/Pre/Rec	Tetris AP/PQ/Pre/Rec	CLEVR AP/PQ/Pre/Rec
AIR	45.4 <b>1.8</b> /38.2 <b>3.0</b> /57.6 <b>7.4</b> /58.1 <b>7.5</b>	25.2 <b>13.9</b> /23.4 <b>12.4</b> /36.8 <b>20.9</b> /39.9 <b>12.9</b>	46.4 <b>14.0</b> /44.3 <b>12.4</b> /67.4 <b>9.9</b> /52.5 <b>15.9</b>
MONet	69.7 <b>4.1</b> /61.6 <b>6.0</b> /70.4 <b>8.1</b> /73.9 <b>1.9</b>	85.9 <b>13.0</b> /75.8 <b>13.6</b> /85.1 <b>16.4</b> /89.7 <b>8.2</b>	39.0 <b>8.5</b> /37.3 <b>6.3</b> /65.6 <b>11.8</b> /42.8 <b>10.8</b>
IODINE	92.9 <b>4.3</b> /71.3 <b>6.1</b> /82.6 <b>2.3</b> /96.0 <b>5.2</b>	52.2 <b>2.3</b> /37.9 <b>4.6</b> /48.0 <b>2.3</b> /61.8 <b>1.7</b>	82.8 <b>2.8</b> /73.0 <b>5.7</b> /77.5 <b>3.1</b> /87.4 <b>2.0</b>
SlotAtt	92.8 <b>1.4</b> /82.8 <b>1.6</b> /88.8 <b>3.4</b> /92.9 <b>1.6</b>	94.3 <b>1.2</b> /79.9 <b>6.4</b> /90.5 <b>3.3</b> /94.4 <b>1.3</b>	91.7 <b>6.4</b> /82.9 <b>10.9</b> /90.8 <b>9.7</b> /92.7 <b>5.3</b>
MRCNN	98.4 <b>1.5</b> /90.2 <b>2.4</b> /99.6 <b>0.6</b> /98.4 <b>1.4</b>	99.8 <b>0.0</b> /90.3 <b>0.6</b> /99.8 <b>0.2</b> /99.8 <b>0.0</b>	98.2 <b>1.1</b> /90.0 <b>1.4</b> /97.8 <b>0.4</b> /99.5 <b>1.3</b>
	YCB AP/PQ/Pre/Rec	ScanNet AP/PQ/Pre/Rec	COCO AP/PQ/Pre/Rec
AIR	0.0 <b>0.1</b> /0.6 <b>0.3</b> /1.1 <b>0.4</b> /0.8 <b>0.2</b>	2.7 <b>1.4</b> /6.3 <b>1.7</b> /15.6 <b>2.8</b> /7.3 <b>1.6</b>	2.7 <b>0.1</b> /6.7 <b>0.5</b> /14.3 <b>2.6</b> /8.6 <b>0.8</b>
MONet	3.1 <b>1.6</b> /7.0 <b>2.6</b> /9.8 <b>3.6</b> /1.2 <b>0.8</b>	24.8 <b>1.6</b> /24.6 <b>1.6</b> /31.0 <b>1.6</b> /40.7 <b>1.8</b>	11.8 <b>2.0</b> /12.5 <b>1.1</b> /16.1 <b>0.9</b> /21.9 <b>1.7</b>
IODINE	1.8 <b>0.2</b> /3.9 <b>1.3</b> /6.2 <b>2.0</b> /7.3 <b>1.9</b>	10.1 <b>2.9</b> /13.7 <b>2.7</b> /18.6 <b>4.2</b> /24.4 <b>3.8</b>	4.0 <b>1.2</b> /6.3 <b>1.2</b> /9.9 <b>1.8</b> /10.8 <b>2.0</b>
SlotAtt	9.2 <b>0.4</b> /13.5 <b>0.9</b> /20.0 <b>1.3</b> /26.2 <b>6.8</b>	5.7 <b>0.3</b> /9.0 <b>1.5</b> /12.4 <b>2.5</b> /18.3 <b>2.7</b>	0.8 <b>0.3</b> /3.5 <b>1.2</b> /5.3 <b>1.7</b> /7.3 <b>2.2</b>
MRCNN	62.9 <b>4.0</b> /58.4 <b>5.9</b> /83.3 <b>7.2</b> /66.9 <b>3.3</b>	41.4 <b>4.3</b> /43.3 <b>6.2</b> /65.2 <b>3.5</b> /50.5 <b>1.8</b>	46.0 <b>0.8</b> /47.9 <b>2.1</b> /71.7 <b>4.0</b> /53.2 <b>0.4</b>

This table corresponds to Figs. 11 and 12. Standard deviations are calculated over 3 runs (marked with bold)

**Table 5** Quantitative results of object segmentation from three methods on datasets of Groups 3&4 with blank backgrounds

CLEVR <sub>bg</sub> AP/PQ/Pre/Rec/BG-Rec												
IODINE	89.0 <b>6.1</b> /67.5 <b>4.9</b> /86.6 <b>13.5</b> /90.5 <b>5.5</b> /99.9 <b>0.0</b>											
SlotAtt	92.6 <b>6.2</b> /81.5 <b>8.3</b> /90.9 <b>4.0</b> /93.6 <b>6.1</b> /100 <b>0.0</b>											
MRCNN	97.2 <b>1.0</b> /85.8 <b>4.2</b> /96.6 <b>1.1</b> /98.4 <b>1.0</b> /100.0 <b>0.0</b>											
	YCB <sub>bg</sub> AP/PQ/Pre/Rec/BG-Rec				ScanNet <sub>bg</sub> AP/PQ/Pre/Rec/BG-Rec				COCO <sub>bg</sub> AP/PQ/Pre/Rec/BG-Rec			
IODINE	0.1 <b>0.4</b> /1.2 <b>1.9</b> /1.8 <b>0.6</b> /2.6 <b>1.7</b> /8.2 <b>2.9</b>	0.8 <b>0.3</b> /3.9 <b>0.5</b> /5.4 <b>0.5</b> /7.8 <b>1.6</b> /26.5 <b>4.6</b>				0.1 <b>0.1</b> /1.0 <b>0.8</b> /1.3 <b>1.0</b> /2.2 <b>1.9</b> /10.1 <b>3.5</b>						
SlotAtt	0.0 <b>0.0</b> /0.9 <b>0.1</b> /1.4 <b>0.2</b> /1.9 <b>0.3</b> /0.3 <b>0.2</b>	1.8 <b>0.5</b> /6.2 <b>0.4</b> /8.1 <b>0.5</b> /12.9 <b>0.7</b> /0.0 <b>1.3</b>				0.0 <b>0.0</b> /0.3 <b>0.0</b> /0.4 <b>0.0</b> /0.7 <b>0.1</b> /0.1 <b>0.1</b>						
MRCNN	42.6 <b>2.1</b> /40.1 <b>0.2</b> /61.8 <b>4.5</b> /47.9 <b>2.5</b> /100 <b>0.2</b>	14.4 <b>0.9</b> /19.4 <b>0.3</b> /28.1 <b>0.2</b> /28.8 <b>0.7</b> /98.1 <b>0.0</b>				11.3 <b>0.2</b> /15.3 <b>0.0</b> /21.2 <b>1.0</b> /24.5 <b>0.6</b> /99.3 <b>0.3</b>						

This table corresponds to Figs. 14 and 15. Standard deviations are calculated over 3 runs (marked with bold)

## Appendix A.7: Details of Experimental Results in Sects. 4.2, 4.3, 4.4 and 4.5

Tables 6, 7, 8 and 9 present the detailed results of experiments conducted in Sects. 4.2, 4.3, 4.4 and 4.5. Standard deviations for all scores are calculated over 3 independent runs.

**Table 6** Quantitative results on datasets ablated with object-level factors

	YCB-C AP/PQ/Pre/Rec	ScanNet-C AP/PQ/Pre/Rec	COCO-C AP/PQ/Pre/Rec
AIR	4.4 <b>0.2</b> /1.0 <b>7.3</b> /20.7 <b>7.1</b> /12.5 <b>1.2</b>	2.7 <b>1.0</b> /7.6 <b>2.1</b> /29.1 <b>20.2</b> /7.4 <b>1.5</b>	2.9 <b>1.8</b> /7.9 <b>2.5</b> /32.0 <b>29.3</b> /7.7 <b>4.5</b>
MONet	55.1 <b>6.0</b> /49.5 <b>5.2</b> /64.8 <b>1.7</b> /58.8 <b>7.8</b>	31.1 <b>15.9</b> /33.1 <b>9.6</b> /41.6 <b>10.3</b> /45.6 <b>11.1</b>	34.5 <b>7.2</b> /34.6 <b>3.3</b> /45.8 <b>2.0</b> /42.6 <b>11.0</b>
IODINE	76.8 <b>0.3</b> /64.4 <b>0.1</b> /76.5 <b>0.1</b> /79.5 <b>0.0</b>	64.4 <b>6.4</b> /54.3 <b>5.9</b> /66.3 <b>4.6</b> /71.0 <b>5.9</b>	42.8 <b>14.7</b> /34.0 <b>11.8</b> /42.5 <b>14.4</b> /55.8 <b>11.9</b>
SlotAtt	39.2 <b>1.9</b> /30.2 <b>0.0</b> /40.2 <b>0.8</b> /50.4 <b>0.7</b>	5.7 <b>8.2</b> /9.0 <b>6.1</b> /12.4 <b>7.6</b> /18.3 <b>8.9</b>	7.2 <b>1.9</b> /9.8 <b>1.7</b> /13.9 <b>2.2</b> /18.7 <b>3.4</b>
	YCB-S AP/PQ/Pre/Rec	ScanNet-S AP/PQ/Pre/Rec	COCO-S AP/PQ/Pre/Rec
AIR	3.6 <b>3.0</b> /6.8 <b>3.2</b> /12.5 <b>6.7</b> /9.7 <b>2.7</b>	3.3 <b>2.5</b> /8.5 <b>1.1</b> /36.6 <b>23.0</b> /8.1 <b>8.5</b>	3.3 <b>0.0</b> /7.1 <b>1.1</b> /16.8 <b>5.6</b> /8.3 <b>1.0</b>
MONet	2.5 <b>1.1</b> /6.4 <b>1.8</b> /8.9 <b>2.6</b> /11.3 <b>3.2</b>	18.3 <b>10.4</b> /22.2 <b>3.9</b> /28.4 <b>4.4</b> /37.5 <b>5.3</b>	8.7 <b>3.2</b> /11.2 <b>1.7</b> /14.4 <b>3.0</b> /21.0 <b>0.7</b>
IODINE	2.5 <b>2.5</b> /5.0 <b>5.0</b> /7.9 <b>7.8</b> /1.1 <b>1.0</b>	8.3 <b>1.1</b> /13.7 <b>0.4</b> /18.4 <b>0.4</b> /24.8 <b>0.7</b>	4.2 <b>0.4</b> /8.3 <b>0.4</b> /12.1 <b>0.3</b> /14.8 <b>0.6</b>
SlotAtt	19.1 <b>0.7</b> /21.0 <b>1.7</b> /30.4 <b>2.8</b> /37.7 <b>0.8</b>	1.9 <b>8.5</b> /6.6 <b>7.5</b> /8.7 <b>9.8</b> /13.7 <b>11.9</b>	2.6 <b>1.0</b> /5.9 <b>0.9</b> /8.3 <b>1.3</b> /11.9 <b>10.9</b>
	YCB-C+S AP/PQ/Pre/Rec	ScanNet-C+S AP/PQ/Pre/Rec	COCO-C+S AP/PQ/Pre/Rec
AIR	2.0 <b>0.3</b> /6.4 <b>0.2</b> /15.1 <b>5.2</b> /7.3 <b>4.1</b>	3.5 <b>1.0</b> /9.1 <b>0.5</b> /38.3 <b>25.8</b> /8.5 <b>5.9</b>	4.8 <b>0.1</b> /10.8 <b>0.1</b> /41.5 <b>10.2</b> /10.1 <b>1.7</b>
MONet	9.1 <b>13.6</b> /12.7 <b>10.5</b> /34.1 <b>4.6</b> /12.6 <b>24.0</b>	48.1 <b>1.0</b> /47.9 <b>6.8</b> /70.4 <b>23.3</b> /51.5 <b>6.5</b>	10.5 <b>24.3</b> /16.5 <b>18.1</b> /56.5 <b>6.9</b> /14.9 <b>28.2</b>
IODINE	69.9 <b>30.9</b> /55.7 <b>29.6</b> /60.3 <b>28.3</b> /71.5 <b>26.4</b>	73.8 <b>11.4</b> /63.2 <b>10.1</b> /73.5 <b>6.1</b> /76.8 <b>10.8</b>	73.5 <b>2.5</b> /61.2 <b>2.4</b> /70.0 <b>2.9</b> /77.6 <b>1.2</b>
SlotAtt	39.1 <b>22.3</b> /30.2 <b>13.3</b> /42.9 <b>12.2</b> /54.5 <b>12.7</b>	22.5 <b>7.6</b> /21.5 <b>6.2</b> /28.8 <b>6.0</b> /39.6 <b>5.4</b>	20.3 <b>3.0</b> /19.9 <b>0.4</b> /26.1 <b>0.4</b> /34.8 <b>0.7</b>

Standard deviations are calculated over 3 runs (marked with bold). This table corresponds to Sect. 4.2



**Table 7** Quantitative results on datasets ablated with scene-level factors

	YCB-T AP/PQ/Pre/Rec	ScanNet-T AP/PQ/Pre/Rec	COCO-T AP/PQ/Pre/Rec
AIR	5.9 <b>3.7</b> /9.6 <b>3.3</b> /19.9 <b>10.3</b> /12.7 <b>1.3</b>	2.9 <b>1.2</b> /6.3 <b>0.1</b> /13.4 <b>3.7</b> /8.6 <b>3.2</b>	7.4 <b>3.1</b> /13.0 <b>4.0</b> /29.1 <b>9.8</b> /16.4 <b>4.3</b>
MONet	86.5 <b>12.4</b> /78.3 <b>13.3</b> /81.1 <b>11.6</b> /86.8 <b>12.2</b>	82.9 <b>12.3</b> /77.8 <b>7.6</b> /86.3 <b>1.6</b> /83.7 <b>11.9</b>	80.0 <b>9.4</b> /68.8 <b>13.5</b> /74.5 <b>15.0</b> /82.0 <b>5.6</b>
IODINE	32.4 <b>9.0</b> /27.3 <b>6.8</b> /35.3 <b>8.3</b> /43.6 <b>10.6</b>	33.2 <b>11.9</b> /27.3 <b>6.3</b> /34.5 <b>7.0</b> /44.7 <b>9.3</b>	40.8 <b>12.7</b> /33.7 <b>10.2</b> /41.8 <b>10.4</b> /55.9 <b>15.8</b>
SlotAtt	64.6 <b>5.3</b> /48.5 <b>3.2</b> /58.3 <b>3.5</b> /74.9 <b>4.6</b>	20.5 <b>22.3</b> /18.2 <b>20.1</b> /22.9 <b>23.1</b> /33.2 <b>31.4</b>	31.8 <b>17.8</b> /28.0 <b>11.2</b> /35.2 <b>12.4</b> /50.1 <b>19.8</b>
	YCB-U AP/PQ/Pre/Rec	ScanNet-U AP/PQ/Pre/Rec	COCO-U AP/PQ/Pre/Rec
AIR	9.8 <b>3.5</b> /12.2 <b>2.6</b> /18.6 <b>4.1</b> /19.9 <b>2.8</b>	7.1 <b>1.4</b> /10.4 <b>0.2</b> /19.9 <b>2.1</b> /14.0 <b>1.0</b>	12.9 <b>4.8</b> /16.3 <b>4.6</b> /28.6 <b>1.8</b> /24.3 <b>9.3</b>
MONet	3.7 <b>0.5</b> /8.3 <b>1.2</b> /12.2 <b>1.4</b> /13.2 <b>1.4</b>	25.4 <b>8.3</b> /23.6 <b>7.4</b> /29.2 <b>8.9</b> /37.2 <b>12.5</b>	14.5 <b>0.6</b> /16.5 <b>0.8</b> /27.9 <b>1.4</b> /20.9 <b>0.8</b>
IODINE	2.5 <b>1.4</b> /4.9 <b>2.0</b> /7.6 <b>2.9</b> /9.3 <b>3.6</b>	16.9 <b>0.8</b> /17.2 <b>0.0</b> /23.8 <b>0.8</b> /30.2 <b>0.1</b>	5.5 <b>1.0</b> /7.7 <b>0.0</b> /11.7 <b>0.6</b> /13.5 <b>0.3</b>
SlotAtt	28.9 <b>2.0</b> /21.9 <b>1.2</b> /30.7 <b>3.1</b> /37.3 <b>2.4</b>	9.4 <b>4.8</b> /10.4 <b>3.8</b> /15.4 <b>4.6</b> /18.1 <b>7.2</b>	4.4 <b>1.1</b> /7.9 <b>0.3</b> /11.9 <b>0.9</b> /14.7 <b>0.7</b>
	YCB-T+U AP/PQ/Pre/Rec	ScanNet-T+U AP/PQ/Pre/Rec	COCO-T+U AP/PQ/Pre/Rec
AIR	18.9 <b>13.7</b> /21.0 <b>9.9</b> /32.5 <b>15.7</b> /32.6 <b>13.2</b>	12.3 <b>5.4</b> /16.8 <b>5.6</b> /37.4 <b>20.8</b> /20.4 <b>1.1</b>	25.3 <b>13.0</b> /28.3 <b>9.0</b> /49.4 <b>13.2</b> /38.2 <b>12.4</b>
MONet	89.1 <b>0.8</b> /84.7 <b>1.5</b> /91.7 <b>4.1</b> /89.3 <b>0.8</b>	87.0 <b>8.9</b> /77.5 <b>10.3</b> /80.5 <b>11.1</b> /87.4 <b>8.7</b>	88.2 <b>2.9</b> /76.2 <b>0.1</b> /79.9 <b>3.7</b> /89.0 <b>2.9</b>
IODINE	35.3 <b>1.2</b> /26.4 <b>0.3</b> /34.9 <b>0.6</b> /42.8 <b>0.2</b>	27.4 <b>0.8</b> /25.0 <b>0.8</b> /32.1 <b>0.5</b> /41.8 <b>1.8</b>	53.3 <b>16.7</b> /35.6 <b>8.3</b> /45.0 <b>9.7</b> /59.8 <b>13.0</b>
SlotAtt	76.8 <b>1.2</b> /57.1 <b>3.8</b> /72.4 <b>3.4</b> /77.7 <b>1.8</b>	47.3 <b>11.2</b> /33.8 <b>9.2</b> /42.6 <b>9.1</b> /57.3 <b>8.4</b>	34.5 <b>17.8</b> /25.3 <b>7.6</b> /33.7 <b>8.3</b> /43.0 <b>14.1</b>

Standard deviations are calculated over 3 runs (marked with bold). This table corresponds to Sect. 4.3

**Table 8** Quantitative results on datasets ablated with both object- and scene-level factors

	YCB-C+S+T AP/PQ/Pre/Rec	ScanNet-C+S+T AP/PQ/Pre/Rec	COCO-C+S+T AP/PQ/Pre/Rec
AIR	6.3 <b>3.7</b> /10.7 <b>3.3</b> /24.0 <b>13.2</b> /13.1 <b>0.5</b>	1.9 <b>3.9</b> /6.4 <b>5.2</b> /27.6 <b>10.7</b> /6.3 <b>13.9</b>	10.8 <b>5.6</b> /18.6 <b>7.0</b> /45.1 <b>14.4</b> /20.3 <b>6.9</b>
MONet	67.7 <b>4.2</b> /55.0 <b>8.5</b> /62.7 <b>19.8</b> /75.4 <b>10.2</b>	76.1 <b>7.4</b> /61.7 <b>2.6</b> /64.4 <b>2.1</b> /77.1 <b>2.8</b>	67.1 <b>15.3</b> /56.4 <b>10.1</b> /62.2 <b>6.2</b> /73.9 <b>17.3</b>
IODINE	82.6 <b>2.6</b> /63.8 <b>0.9</b> /75.5 <b>5.1</b> /86.0 <b>1.6</b>	69.0 <b>14.9</b> /53.5 <b>14.1</b> /60.7 <b>14.5</b> /77.3 <b>9.4</b>	77.6 <b>7.8</b> /59.8 <b>6.8</b> /72.5 <b>1.8</b> /83.0 <b>7.5</b>
SlotAtt	65.9 <b>23.2</b> /51.6 <b>18.5</b> /62.2 <b>19.8</b> /72.9 <b>17.6</b>	60.2 <b>8.7</b> /45.5 <b>8.1</b> /52.9 <b>9.2</b> /68.3 <b>7.3</b>	51.8 <b>11.0</b> /41.5 <b>7.7</b> /48.6 <b>7.5</b> /62.9 <b>8.5</b>
	YCB-C+S+U AP/PQ/Pre/Rec	ScanNet-C+S+U AP/PQ/Pre/Rec	COCO-C+S+U AP/PQ/Pre/Rec
AIR	12.5 <b>6.4</b> /17.7 <b>7.0</b> /27.5 <b>11.2</b> /29.6 <b>10.5</b>	8.2 <b>1.5</b> /11.1 <b>0.5</b> /20.4 <b>0.8</b> /15.0 <b>1.3</b>	23.7 <b>9.8</b> /27.5 <b>8.0</b> /51.3 <b>9.6</b> /32.8 <b>8.3</b>
MONet	3.5 <b>2.5</b> /4.8 <b>3.2</b> /23.5 <b>2.6</b> /5.3 <b>6.5</b>	60.1 <b>9.5</b> /56.3 <b>13.3</b> /69.7 <b>20.7</b> /64.0 <b>6.2</b>	35.6 <b>1.4</b> /33.0 <b>3.8</b> /43.4 <b>16.6</b> /38.7 <b>0.1</b>
IODINE	65.3 <b>9.2</b> /55.9 <b>10.4</b> /73.9 <b>4.6</b> /67.7 <b>8.4</b>	64.6 <b>3.2</b> /52.1 <b>5.1</b> /61.3 <b>7.6</b> /68.2 <b>2.2</b>	64.8 <b>7.3</b> /53.4 <b>6.5</b> /65.8 <b>8.5</b> /68.2 <b>2.5</b>
SlotAtt	58.1 <b>3.7</b> /39.1 <b>2.8</b> /52.4 <b>2.3</b> /62.0 <b>3.6</b>	50.5 <b>2.9</b> /38.2 <b>3.6</b> /51.8 <b>6.9</b> /55.0 <b>1.3</b>	20.2 <b>5.2</b> /18.1 <b>0.3</b> /26.9 <b>0.1</b> /33.4 <b>0.6</b>
	YCB-C+S+T+U AP/PQ/Pre/Rec	ScanNet-C+S+T+U AP/PQ/Pre/Rec	COCO-C+S+T+U AP/PQ/Pre/Rec
AIR	24.8 <b>20.1</b> /23.2 <b>11.9</b> /33.7 <b>16.8</b> /39.1 <b>20.1</b>	10.0 <b>1.8</b> /13.5 <b>1.8</b> /24.0 <b>1.4</b> /18.2 <b>7.4</b>	36.4 <b>19.3</b> /39.4 <b>14.5</b> /63.1 <b>16.9</b> /49.2 <b>17.9</b>
MONet	70.8 <b>9.5</b> /79.2 <b>13.0</b> /96.8 <b>9.4</b> /70.8 <b>8.5</b>	79.2 <b>3.4</b> /65.0 <b>3.6</b> /69.1 <b>4.2</b> /82.4 <b>2.2</b>	76.2 <b>3.0</b> /75.9 <b>3.8</b> /90.9 <b>6.4</b> /76.5 <b>3.3</b>
IODINE	72.4 <b>4.8</b> /54.2 <b>6.7</b> /65.1 <b>5.8</b> /74.3 <b>5.8</b>	76.1 <b>8.7</b> /56.5 <b>9.2</b> /66.9 <b>7.0</b> /80.3 <b>6.2</b>	81.4 <b>3.4</b> /59.3 <b>5.1</b> /70.5 <b>7.7</b> /84.3 <b>6.7</b>
SlotAtt	92.0 <b>2.0</b> /65.5 <b>6.1</b> /84.4 <b>6.5</b> /92.5 <b>1.7</b>	62.7 <b>25.4</b> /42.6 <b>28.4</b> /50.5 <b>34.3</b> /69.4 <b>19.0</b>	83.7 <b>4.7</b> /60.7 <b>7.2</b> /76.4 <b>4.6</b> /84.0 <b>5.0</b>

Standard deviations are calculated over 3 runs (marked with bold). This table corresponds to Sect. 4.4

**Table 9** Quantitative results of IODINE and SlotAtt on datasets of Group 4 and their ablations

	YCB <sub>bg</sub> -C AP/PQ/Pre/Rec/BG-Rec	ScanNet <sub>bg</sub> -C AP/PQ/Pre/Rec/BG-Rec	COCO <sub>bg</sub> -C AP/PQ/Pre/Rec/BG-Rec
IODINE	3.4 <b>3.3</b> /6.2 <b>4.7</b> /9.8 <b>3.7</b> /10.8 <b>6.1</b> /99.9 <b>4.8</b>	13.6 <b>9.3</b> /14.3 <b>6.8</b> /21.0 <b>6.5</b> /22.9 <b>2.6</b> /98.2 <b>4.2</b>	1.7 <b>1.6</b> /3.5 <b>2.0</b> /9.2 <b>2.1</b> /4.3 <b>3.3</b> /98.3 <b>2.9</b>
SlotAtt	1.5 <b>0.9</b> /4.5 <b>1.3</b> /6.6 <b>1.8</b> /9.5 <b>3.0</b> /28.7 <b>1.2</b>	5.1 <b>2.0</b> /9.7 <b>1.2</b> /12.9 <b>1.7</b> /18.8 <b>1.9</b> /47.4 <b>14.2</b>	8.2 <b>4.4</b> /12.3 <b>4.2</b> /16.9 <b>5.4</b> /23.2 <b>7.2</b> /99.7 <b>3.5</b>
	YCB <sub>bg</sub> -T AP/PQ/Pre/Rec/BG-Rec	ScanNet <sub>bg</sub> -T AP/PQ/Pre/Rec/BG-Rec	COCO <sub>bg</sub> -T AP/PQ/Pre/Rec/BG-Rec
IODINE	0.7 <b>0.6</b> /3.4 <b>1.7</b> /5.3 <b>2.5</b> /6.6 <b>2.9</b> /91.6 <b>0.1</b>	10.6 <b>0.0</b> /11.0 <b>0.3</b> /15.1 <b>0.6</b> /20.3 <b>0.8</b> /84.5 <b>0.6</b>	1.2 <b>0.1</b> /4.4 <b>0.2</b> /6.5 <b>0.1</b> /8.4 <b>0.4</b> /92.7 <b>0.2</b>
SlotAtt	19.8 <b>2.9</b> /21.2 <b>1.5</b> /29.1 <b>2.1</b> /40.1 <b>2.5</b> /93.8 <b>0.2</b>	12.9 <b>3.2</b> /16.7 <b>2.5</b> /21.9 <b>2.9</b> /30.1 <b>3.1</b> /83.6 <b>10.3</b>	2.1 <b>3.2</b> /4.5 <b>4.7</b> /6.3 <b>6.3</b> /9.5 <b>7.5</b> /88.1 <b>3.5</b>
	YCB <sub>bg</sub> -S AP/PQ/Pre/Rec/BG-Rec	ScanNet <sub>bg</sub> -S AP/PQ/Pre/Rec/BG-Rec	COCO <sub>bg</sub> -S AP/PQ/Pre/Rec/BG-Rec
IODINE	0.20 <b>0.3</b> /2.4 <b>1.1</b> /3.4 <b>1.5</b> /4.9 <b>1.9</b> /12.5 <b>5.6</b>	2.5 <b>2.3</b> /7.6 <b>5.8</b> /10.2 <b>6.8</b> /14.5 <b>11.9</b> /25.9 <b>10.3</b>	0.2 <b>0.2</b> /1.9 <b>1.0</b> /2.5 <b>1.2</b> /4.1 <b>1.9</b> /10.4 <b>0.8</b>
SlotAtt	0.1 <b>0.1</b> /1.6 <b>0.6</b> /2.4 <b>0.9</b> /3.3 <b>1.5</b> /0.0 <b>0.0</b>	1.2 <b>0.1</b> /4.6 <b>0.9</b> /6.2 <b>1.0</b> /9.5 <b>2.1</b> /3.3 <b>2.1</b>	0.1 <b>0.1</b> /1.2 <b>0.0</b> /1.6 <b>0.0</b> /2.8 <b>0.0</b> /0.1 <b>0.1</b>
	YCB <sub>bg</sub> -C+T AP/PQ/Pre/Rec/BG-Rec	ScanNet <sub>bg</sub> -C+T AP/PQ/Pre/Rec/BG-Rec	COCO <sub>bg</sub> -C+T AP/PQ/Pre/Rec/BG-Rec
IODINE	2.4 <b>1.8</b> /5.2 <b>0.6</b> /8.1 <b>0.8</b> /9.9 <b>0.7</b> /95.1 <b>2.2</b>	11.9 <b>1.1</b> /17.1 <b>1.3</b> /23.0 <b>0.5</b> /29.3 <b>3.1</b> /90.7 <b>1.0</b>	2.5 <b>2.0</b> /6.3 <b>0.5</b> /10.7 <b>0.2</b> /9.9 <b>1.6</b> /94.5 <b>3.8</b>
SlotAtt	14.1 <b>4.1</b> /18.1 <b>2.1</b> /25.0 <b>3.1</b> /32.1 <b>8.9</b> /22.2 <b>1.1</b>	4.5 <b>1.1</b> /10.1 <b>2.0</b> /13.8 <b>2.6</b> /19.1 <b>3.3</b> /71.7 <b>12.2</b>	4.9 <b>2.9</b> /7.8 <b>1.6</b> /10.3 <b>1.6</b> /15.1 <b>3.0</b> /96.7 <b>8.7</b>
	YCB <sub>bg</sub> -C+S AP/PQ/Pre/Rec/BG-Rec	ScanNet <sub>bg</sub> -C+S AP/PQ/Pre/Rec/BG-Rec	COCO <sub>bg</sub> -C+S AP/PQ/Pre/Rec/BG-Rec
IODINE	8.3 <b>1.3</b> /10.0 <b>0.6</b> /16.7 <b>0.1</b> /15.6 <b>2.1</b> /99.1 <b>0.7</b>	13.6 <b>3.3</b> /17.7 <b>2.0</b> /27.7 <b>1.2</b> /26.1 <b>4.3</b> /94.1 <b>1.8</b>	3.7 <b>1.0</b> /8.4 <b>0.2</b> /17.9 <b>0.9</b> /10.7 <b>0.7</b> /95.7 <b>2.5</b>
SlotAtt	0.2 <b>0.5</b> /1.4 <b>0.5</b> /2.2 <b>1.5</b> /2.5 <b>1.4</b> /67.4 <b>1.2</b>	0.4 <b>1.1</b> /2.3 <b>0.7</b> /4.5 <b>2.1</b> /3.2 <b>1.6</b> /66.5 <b>2.4</b>	3.5 <b>1.0</b> /7.9 <b>1.4</b> /11.2 <b>1.8</b> /15.6 <b>3.4</b> /93.1 <b>3.1</b>
	YCB <sub>bg</sub> -T+S AP/PQ/Pre/Rec/BG-Rec	ScanNet <sub>bg</sub> -T+S AP/PQ/Pre/Rec/BG-Rec	COCO <sub>bg</sub> -T+S AP/PQ/Pre/Rec/BG-Rec
IODINE	2.0 <b>1.4</b> /7.3 <b>0.8</b> /10.6 <b>1.8</b> /13.2 <b>1.4</b> /92.8 <b>5.3</b>	12.2 <b>1.0</b> /12.2 <b>1.4</b> /16.4 <b>2.4</b> /22.6 <b>1.3</b> /86.0 <b>1.3</b>	2.3 <b>0.4</b> /7.1 <b>0.4</b> /10.4 <b>0.5</b> /13.0 <b>0.8</b> /88.8 <b>0.0</b>
SlotAtt	7.7 <b>0.0</b> /13.5 <b>0.0</b> /8.6 <b>0.2</b> /26.6 <b>0.2</b> /92.6 <b>0.0</b>	7.4 <b>1.0</b> /11.0 <b>0.3</b> /14.7 <b>0.2</b> /19.7 <b>0.2</b> /60.4 <b>2.6</b>	10.0 <b>4.0</b> /13.9 <b>7.7</b> /17.7 <b>9.0</b> /26.2 <b>7.0</b> /89.5 <b>1.5</b>
	YCB <sub>bg</sub> -C+T+S AP/PQ/Pre/Rec/BG-Rec	ScanNet <sub>bg</sub> -C+T+S AP/PQ/Pre/Rec/BG-Rec	COCO <sub>bg</sub> -C+T+S AP/PQ/Pre/Rec/BG-Rec
IODINE	3.0 <b>1.5</b> /5.9 <b>2.5</b> /9.3 <b>3.6</b> /11.3 <b>3.5</b> /99.4 <b>1.9</b>	12.0 <b>0.9</b> /15.8 <b>3.0</b> /21.3 <b>4.5</b> /28.0 <b>3.5</b> /93.7 <b>0.5</b>	6.4 <b>1.3</b> /10.9 <b>0.3</b> /16.7 <b>0.2</b> /17.7 <b>0.4</b> /97.9 <b>3.8</b>
SlotAtt	14.8 <b>4.2</b> /15.2 <b>1.5</b> /20.4 <b>3.1</b> /28.3 <b>4.1</b> /90.0 <b>4.2</b>	2.1 <b>0.3</b> /6.2 <b>0.4</b> /8.5 <b>0.4</b> /11.9 <b>0.5</b> /36.3 <b>9.5</b>	2.4 <b>2.1</b> /6.0 <b>2.8</b> /8.2 <b>3.7</b> /12.0 <b>4.8</b> /84.5 <b>1.1</b>

This table corresponds to Fig. 25 in Sect. 4.5. Standard deviations are calculated over 3 runs (marked with bold)

## Appendix A.8: Additional Joint Ablations on Object- and Scene-Level Factors

In addition to the experiments in Sect. 4.4, we generate additional 6 groups of datasets ablated with different combinations of object- and scene-level factors as follows and conduct corresponding experiments. Experiment results can be found in Figs. 38 and 39, and Table 10.

### Additional Ablated Datasets:

- *Ablation of Object Color Gradient and Inter-object Color Similarity*: In each image of the three real-world datasets, we replace the object color by averaging all pixels of a distinctive texture, and keep the original shape unchanged, getting three ablated datasets: YCB-C+T/ScanNet-C+T/COCO-C+T.
- *Ablation of Object Color Gradient and Inter-object Shape Variation*: In each image of the three real-world datasets, we replace the object color by averaging its own texture, and then apply size normalization on the original shape of objects, getting three ablated datasets: YCB-C+U/ScanNet-C+U/COCO-C+U.
- *Ablation of Object Color Gradient, Inter-object Color Similarity and Inter-object Shape Variation*: In each image of the three real-world datasets, we replace the object color by averaging all pixels of a distinctive tex-


ture, and then apply size normalization on the original shape of objects. The ablated datasets are denoted as: YCB-C+T+U/ScanNet-C+T+U/COCO-C+T+U.

- *Ablation of Object Shape Concavity and Inter-object Color Similarity*: In each image of the three real-world datasets, we replace the object color by a distinctive texture, and modify the object shape as a convex hull, getting three ablated datasets: YCB-S+T/ScanNet-S+T/COCO-S+T.
- *Ablation of Object Shape Concavity and Inter-object Shape Variation*: In each image of the three real-world datasets, we keep the texture of object unchanged, and modify the object shape as a convex hull followed by size normalization, getting three ablated datasets: YCB-S+U/ScanNet-S+U/COCO-S+U.
- *Ablation of Object Shape Concavity, Inter-object Color Similarity and Inter-object Shape Variation*: In each image of the three real-world datasets, we replace the object color with a distinctive texture, and modify the object shape as a convex hull followed by size normalization. The ablated datasets are denoted as: YCB-S+T+U/ScanNet-S+T+U/COCO-S+T+U.

**Qualitative and Quantitative Results:** As shown in Table 10, Figs. 38 and 39, all 6 additional combinations of object- and scene-level factors are explored, demonstrating consistent findings as our experiments in Sect. 4.4. Overall, all four



**Fig. 38** Quantitative results of baselines on Group 2 datasets and their variants in Sect. 1



**Fig. 39** Qualitative results of additional joint ablations on both object- and scene-level factors in Sect. 1

methods show a high sensitivity to both object- and scene-level factors relating to appearance. This can be seen from the fact that: for datasets without ablations in appearance, i.e., the (S+U) ablated datasets, the object segmentation performance is inferior. By contrast, the object segmentation accuracy can be greatly improved on the datasets only with appearance factors ablated, i.e., the (C+T) datasets. Meanwhile, more regular shapes and uniform scales of objects still have a significant positive influence on the success of object segmentation especially when the appearance factors are combined in ablated datasets (Figs. 40, 41 and 42).


To be specific, AIR (Eslami et al., 2016) is quite sensitive to the scale of objects apart from object color gradient and inter-object color similarity. MONet (Burgess et al., 2019) can obtain comparable performance to the simple synthetic datasets once object color gradient and inter-object color similarity are ablated. All four factors are closely relevant to the results of IODINE (Greff et al., 2019) and SlotAtt (Locatello et al., 2020).




**Table 10** Quantitative results on additional datasets ablated with both object- and scene-level factors

	YCB-S+U AP/PQ/Pre/Rec	ScanNet-S+U AP/PQ/Pre/Rec	COCO-S+U AP/PQ/Pre/Rec
AIR	8.0 <b>6.1</b> /13.2 <b>7.3</b> /26.6 <b>17.6</b> /17.9 <b>7.3</b>	2.8 <b>4.7</b> /6.7 <b>6.5</b> /14.8 <b>4.0</b> /8.1 <b>14.8</b>	12.4 <b>9.4</b> /30.2 <b>23.0</b> /40.1 <b>25.6</b> /24.2 <b>13.8</b>
MONet	82.9 <b>9.6</b> /66.7 <b>0.5</b> /73.0 <b>1.2</b> /87.8 <b>12.2</b>	36.5 <b>28.4</b> /33.5 <b>21.8</b> /40.4 <b>24.1</b> /50.3 <b>19.7</b>	66.1 <b>5.6</b> /53.2 <b>3.7</b> /56.2 <b>2.3</b> /74.4 <b>6.4</b>
IODINE	78.9 <b>2.6</b> /59.7 <b>2.1</b> /71.5 <b>0.6</b> /83.6 <b>2.7</b>	65.1 <b>2.9</b> /51.1 <b>1.0</b> /62.4 <b>2.9</b> /74.2 <b>0.5</b>	55.1 <b>10.4</b> /42.2 <b>7.5</b> /56.3 <b>6.4</b> /66.2 <b>9.2</b>
SlotAtt	58.7 <b>12.1</b> /43.2 <b>5.0</b> /57.8 <b>9.4</b> /71.0 <b>9.4</b>	29.2 <b>2.9</b> /27.6 <b>1.3</b> /34.4 <b>1.4</b> /49.2 <b>1.2</b>	22.1 <b>8.6</b> /19.6 <b>4.1</b> /25.2 <b>4.1</b> /35.8 <b>8.1</b>
	YCB-C+U AP/PQ/Pre/Rec	ScanNet-C+U AP/PQ/Pre/Rec	COCO-C+U AP/PQ/Pre/Rec
AIR	11.4 <b>4.9</b> /16.4 <b>5.4</b> /25.8 <b>8.7</b> /25.1 <b>6.5</b>	5.4 <b>0.5</b> /12.3 <b>2.6</b> /35.1 <b>15.3</b> /12.6 <b>0.4</b>	20.5 <b>10.8</b> /24.9 <b>7.1</b> /47.7 <b>10.2</b> /30.5 <b>8.8</b>
MONet	49.1 <b>3.5</b> /44.6 <b>4.1</b> /60.6 <b>3.2</b> /52.5 <b>3.3</b>	31.9 <b>5.7</b> /33.7 <b>5.7</b> /46.8 <b>5.4</b> /35.7 <b>8.7</b>	32.6 <b>11.1</b> /28.5 <b>12.1</b> /34.9 <b>14.2</b> /39.0 <b>18.5</b>
IODINE	66.3 <b>0.1</b> /54.6 <b>0.0</b> /71.3 <b>4.1</b> /70.0 <b>0.2</b>	34.8 <b>20.9</b> /30.0 <b>14.7</b> /36.4 <b>20.9</b> /47.2 <b>14.0</b>	44.4 <b>8.9</b> /32.6 <b>7.9</b> /41.7 <b>12.3</b> /51.1 <b>7.6</b>
SlotAtt	45.6 <b>7.9</b> /31.2 <b>4.2</b> /41.9 <b>7.7</b> /51.3 <b>6.4</b>	30.4 <b>8.9</b> /24.2 <b>6.2</b> /33.8 <b>9.5</b> /40.6 <b>11.0</b>	12.7 <b>7.7</b> /11.3 <b>3.7</b> /15.8 <b>5.0</b> /23.5 <b>6.8</b>
	YCB-S+T AP/PQ/Pre/Rec	ScanNet-S+T AP/PQ/Pre/Rec	COCO-S+T AP/PQ/Pre/Rec
AIR	9.3 <b>7.8</b> /13.4 <b>8.4</b> /21.9 <b>14.3</b> /21.2 <b>12.3</b>	3.2 <b>0.5</b> /7.7 <b>0.6</b> /18.5 <b>6.0</b> /8.9 <b>6.0</b>	10.5 <b>5.4</b> /18.2 <b>7.2</b> /4.3 <b>23.2</b> /20.3 <b>7.3</b>
MONet	86.7 <b>2.2</b> /78.9 <b>1.6</b> /81.9 <b>1.8</b> /87.0 <b>2.2</b>	83.2 <b>3.7</b> /77.8 <b>13.0</b> /86.0 <b>20.0</b> /83.8 <b>2.9</b>	80.2 <b>0.0</b> /68.5 <b>7.2</b> /73.2 <b>11.7</b> /82.3 <b>1.0</b>
IODINE	41.9 <b>17.9</b> /33.2 <b>11.4</b> /41.5 <b>12.0</b> /51.3 <b>14.7</b>	54.9 <b>23.1</b> /44.4 <b>16.2</b> /52.0 <b>16.7</b> /68.0 <b>21.6</b>	44.1 <b>1.1</b> /34.3 <b>0.7</b> /41.0 <b>0.3</b> /56.8 <b>0.6</b>
SlotAtt	77.3 <b>7.8</b> /60.6 <b>1.5</b> /75.0 <b>2.3</b> /69.2 <b>19.7</b>	24.9 <b>46.1</b> /21.2 <b>35.0</b> /25.5 <b>38.3</b> /37.9 <b>43.2</b>	67.4 <b>3.1</b> /50.2 <b>1.6</b> /59.2 <b>1.6</b> /76.4 <b>0.3</b>
AIR	0.8 <b>5.4</b> /2.9 <b>7.1</b> /5.0 <b>10.0</b> /5.2 <b>12.7</b>	6.3 <b>0.2</b> /13.1 <b>3.3</b> /32.4 <b>16.3</b> /14.3 <b>1.2</b>	20.0 <b>9.5</b> /25.8 <b>10.0</b> /48.3 <b>15.6</b> /31.6 <b>10.9</b>
MONet	5.5 <b>1.2</b> /9.8 <b>0.4</b> /14.1 <b>0.8</b> /17.0 <b>0.2</b>	36.9 <b>2.9</b> /31.4 <b>2.6</b> /38.2 <b>3.6</b> /50.1 <b>0.1</b>	26.1 <b>3.5</b> /23.8 <b>1.8</b> /29.6 <b>1.6</b> /41.2 <b>10.0</b>
IODINE	2.8 <b>2.4</b> /4.6 <b>2.0</b> /7.2 <b>2.9</b> /8.9 <b>3.6</b>	17.3 <b>2.4</b> /17.8 <b>1.4</b> /24.1 <b>1.8</b> /31.6 <b>2.1</b>	5.7 <b>0.5</b> /8.7 <b>1.7</b> /1.3 <b>13.7</b> /16.3 <b>2.4</b>
SlotAtt	36.2 <b>3.3</b> /23.8 <b>4.1</b> /33.6 <b>5.5</b> /45.6 <b>1.7</b>	21.1 <b>0.9</b> /18.9 <b>0.2</b> /26.2 <b>0.1</b> /32.5 <b>2.0</b>	12.7 <b>7.1</b> /12.1 <b>2.3</b> /16.4 <b>2.1</b> /24.7 <b>5.2</b>
	YCB-C+T+U AP/PQ/Pre/Rec	ScanNet-C+T+U AP/PQ/Pre/Rec	COCO-C+T+U AP/PQ/Pre/Rec
AIR	18.2 <b>12.0</b> /20.5 <b>7.9</b> /32.5 <b>13.5</b> /30.9 <b>9.0</b>	13.5 <b>6.6</b> /20.2 <b>9.1</b> /42.4 <b>25.5</b> /24.1 <b>5.6</b>	24.2 <b>11.1</b> /28.3 <b>8.5</b> /50.5 <b>13.1</b> /36.4 <b>9.7</b>
MONet	63.2 <b>6.9</b> /66.5 <b>11.8</b> /86.3 <b>13.0</b> /64.4 <b>5.7</b>	74.6 <b>5.6</b> /62.1 <b>6.6</b> /67.8 <b>7.2</b> /79.2 <b>3.2</b>	73.0 <b>3.3</b> /75.7 <b>10.3</b> /94.2 <b>23.1</b> /73.1 <b>7.5</b>
IODINE	54.9 <b>23.1</b> /38.7 <b>18.2</b> /58.0 <b>12.0</b> /59.4 <b>20.8</b>	65.5 <b>3.8</b> /48.0 <b>3.2</b> /59.8 <b>5.9</b> /70.7 <b>2.9</b>	47.9 <b>11.7</b> /35.5 <b>7.4</b> /53.8 <b>2.1</b> /53.4 <b>14.2</b>
SlotAtt	73.8 <b>4.6</b> /53.0 <b>2.9</b> /67.6 <b>7.8</b> /75.2 <b>4.5</b>	58.5 <b>8.1</b> /45.5 <b>10.2</b> /52.9 <b>8.9</b> /68.3 <b>11.9</b>	30.5 <b>27.6</b> /20.3 <b>18.0</b> /27.8 <b>23.0</b> /38.9 <b>23.3</b>
	YCB-S+T+U AP/PQ/Pre/Rec	ScanNet-S+T+U AP/PQ/Pre/Rec	COCO-S+T+U AP/PQ/Pre/Rec
AIR	24.6 <b>20.0</b> /24.8 <b>13.7</b> /37.4 <b>20.7</b> /37.6 <b>18.6</b>	16.7 <b>10.1</b> /23.9 <b>12.0</b> /47.8 <b>30.2</b> /28.5 <b>7.8</b>	29.4 <b>16.5</b> /34.6 <b>18.4</b> /58.3 <b>35.7</b> /42.2 <b>13.4</b>
MONet	88.4 <b>0.6</b> /84.6 <b>1.5</b> /92.0 <b>1.5</b> /88.6 <b>0.7</b>	97.8 <b>0.3</b> /87.3 <b>0.8</b> /85.4 <b>0.1</b> /98.1 <b>0.4</b>	86.9 <b>0.1</b> /83.8 <b>11.6</b> /89.9 <b>13.1</b> /87.0 <b>1.5</b>
IODINE	42.3 <b>15.0</b> /32.1 <b>9.7</b> /41.0 <b>11.5</b> /52.0 <b>15.7</b>	39.1 <b>4.7</b> /29.9 <b>1.8</b> /41.8 <b>7.2</b> /47.9 <b>2.4</b>	64.3 <b>9.7</b> /45.1 <b>6.9</b> /51.4 <b>5.6</b> /72.0 <b>11.7</b>
SlotAtt	88.6 <b>6.5</b> /67.0 <b>9.6</b> /80.9 <b>13.0</b> /89.6 <b>5.6</b>	82.6 <b>8.5</b> /61.5 <b>10.1</b> /73.2 <b>10.2</b> /83.9 <b>7.7</b>	84.2 <b>0.9</b> /59.6 <b>6.8</b> /74.7 <b>3.8</b> /85.3 <b>0.6</b>


Standard deviations are calculated over 3 runs (marked with bold)




(a) Ablation of Object Color Gradient.




(b) Ablation of Object Shape Concavity.




(c) Ablation of both Object Color Gradient and Object Shape Concavity.

**Fig. 40** Qualitative results of object-level factors ablations

(a) Ablation of Inter-object Color Similarity

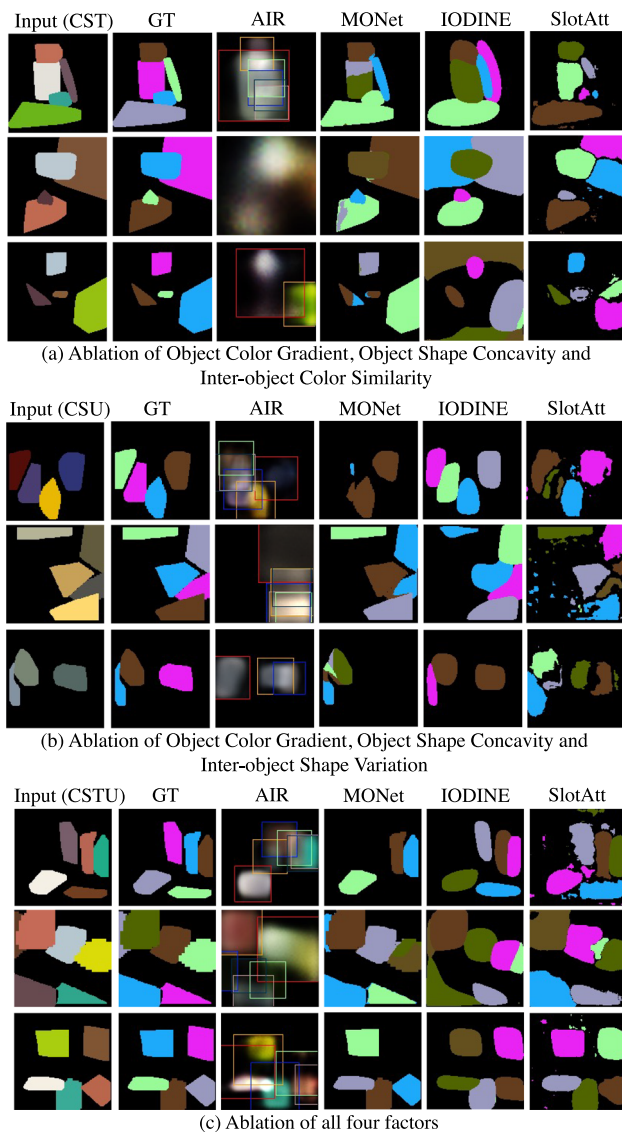


(b) Ablation of Inter-object Shape Variation



(c) Ablation of Inter-object Color Similarity and Shape Variation

**Fig. 41** Qualitative results of scene-level factors ablations



**Fig. 42** Qualitative results of joint ablation on object- and scene-level factors

## References

- Abdal, R., Zhu, P., Mitra, N., & Wonka, P. (2021). Labels4Free: Unsupervised segmentation using StyleGAN. In *ICCV*.
- Anciukevicius, T., Lampert, C. H., & Henderson, P. (2020). Object-centric image generation with factored depths, locations, and appearances. [arXiv:2004.00642](https://arxiv.org/abs/2004.00642)
- Arandjelovic, R., & Zisserman, A. (2019). Object discovery with a copy-pasting GAN. [arXiv:1905.11369](https://arxiv.org/abs/1905.11369)
- Arbeláez, P., Maire, M., Fowlkes, C., & Malik, J. (2011). Contour detection and hierarchical image segmentation. *TPAMI*, 33(5), 898–916.
- Azadi, S., Pathak, D., Ebrahimi, S., & Darrell, T. (2020). Compositional GAN: Learning image-conditional binary composition. In *IJCV*.
- Bear, D. M., Fan, C., Mrowca, D., Li, Y., Alter, S., Nayeibi, A., Schwartz, J., Fei-Fei, L., Wu, J., Tenenbaum, J. B., & Yamins, D. L. (2020). Learning physical graph representations from visual scenes. In *NeurIPS*.

- Bielski, A., & Favaro, P. (2019). Emergence of object segmentation in perturbed generative models. In *NeurIPS*.
- Biza, O., van Steenkiste, S., Sajjadi, M. S., Elsayed, G. F., Mahendran, A., & Kipf, T. (2023). Invariant slot attention: Object discovery with slot-centric reference frames. In *ICML*.
- Borgefors, G., & Strand, R. (2005). An approximation of the maximal inscribed convex set of a digital object. In *Image analysis and processing—ICIAP 2005: 13th international conference, Cagliari, Italy, September 6–8, 2005. Proceedings 13* (pp. 438–445). Springer.
- Burgess, C. P., Matthey, L., Watters, N., Kabra, R., Higgins, I., Botvinick, M., & Lerchner, A. (2019). MONet: Unsupervised scene decomposition and representation. [arXiv:1901.11390](https://arxiv.org/abs/1901.11390)
- Calli, B., Singh, A., Bruce, J., Walsman, A., Konolige, K., Srinivasa, S., Abbeel, P., & Dollar, A. M. (2017). Yale-CMU-Berkeley dataset for robotic manipulation research. *IJRR*, 36(3), 261–268.
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., & Joulin, A. (2021). Emerging properties in self-supervised vision transformers. In *ICCV*.
- Charakorn, R., Thawornwattana, Y., Itthipuripat, S., Pawlowski, N., & Dilokthanakul, N. (2020). An explicit local and global representation disentanglement framework with applications in deep clustering and unsupervised object detection. [arXiv:2001.08957](https://arxiv.org/abs/2001.08957)
- Chen, H., Venkatesh, R., Friedman, Y., Wu, J., Tenenbaum, J. B., Yamins, D. L. K., & Bear, D. M. (2022). Unsupervised segmentation in real-world images via Spelke object inference. [arXiv:2205.08515](https://arxiv.org/abs/2205.08515)
- Chen, M., Artières, T., & Denoyer, L. (2019). Unsupervised object segmentation by redrawing. In *NeurIPS*.
- Cheng, B., Girshick, R., Dollár, P., Berg, A. C., & Kirillov, A. (2021). Boundary IoU: Improving object-centric image segmentation evaluation. In *CVPR*.
- Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., & Vedaldi, A. (2014). Describing textures in the wild. In *CVPR*.
- Crawford, E., & Pineau, J. (2019). Spatially invariant unsupervised object detection with convolutional neural networks. In *AAAI*.
- Dai, A., Chang, A. X., Savva, M., Halber, M., Funkhouser, T., & Nießner, M. (2017). ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *CVPR*.
- Deng, F., Zhi, Z., Lee, D., & Ahn, S. (2021). Generative scene graph networks. In *ICLR*.
- Du, Y., Li, S., Sharma, Y., Tenenbaum, J., & Mordatch, I. (2021). Unsupervised learning of compositional energy concepts. In *NeurIPS*.
- Eddins, S. (2011). Binary image convex hull. <https://blogs.mathworks.com/steve/2011/10/04/binary-image-convex-hull-algorithm-notes/>
- Engelcke, M., Kosiorek, A. R., Jones, O. P., & Posner, I. (2020). GENE-SIS: Generative scene inference and sampling with object-centric latent representations. In *ICLR*.
- Engelcke, M., Jones, O. P., & Posner, I. (2021). GENESIS-V2: Inferring unordered object representations without iterative refinement. In *NeurIPS*.
- Eslami, S. M. A., Heess, N., Weber, T., Tassa, Y., Kavukcuoglu, K., & Hinton, G. E. (2016). Attend, infer, repeat: Fast scene understanding with generative models. In *NIPS*.
- Everingham, M., Eslami, S. M. A., Van Gool, L., Williams, C. K. I., Winn, J., & Zisserman, A. (2015). The PASCAL visual object classes challenge: A retrospective. In *IJCV*.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *NIPS*.
- Greff, K., Belletti, F., Beyer, L., Doersch, C., Du, Y., Duckworth, D., Fleet, D. J., Gnanaprasam, D., Golemo, F., Herrmann, C., Kipf, T., Kundu, A., Lagun, D., Laradji, I., Liu, H. T. D., Meyer, H., Miao, Y., Nowrouzezahrai, D., Oztireli, C., Pot, E., Radwan, N., Rebain, D., Sabour, S., Sajjadi, M. S. M., Sela, M., Sitzmann, V.,

- Stone, A., Sun, D., Vora, S., Wang, Z., Wu, T., Yi, K. M., Zhong, F., & Tagliasacchi, A. (2022). Kubric: A scalable dataset generator. In *CVPR 2022*.
- Greff, K., Rasmus, A., Berglund, M., Hao, T.H., Schmidhuber, J., & Valpola, H. (2016). Tagger: Deep unsupervised perceptual grouping. In *NIPS*.
- Greff, K., Kaufman, R. L., Kabra, R., Watters, N., Burgess, C., Zoran, D., Matthey, L., Botvinick, M., & Lerchner, A. (2019). Multi-object representation learning with iterative variational inference. In *ICML*.
- Greff, K., van Steenkiste, S., & Schmidhuber, J. (2020). On the binding problem in artificial neural networks. [arXiv:2012.05208](https://arxiv.org/abs/2012.05208)
- Greff, K., Van Steenkiste, S., & Schmidhuber, J. (2017). Neural expectation maximization. In *NIPS*.
- He, K., Gkioxari, G., Dollar, P., & Girshick, R. (2017). Mask R-CNN. In *ICCV*.
- Hénaff, O. J., Koppula, S., Shelhamer, E., Zoran, D., Jaegle, A., Zisserman, A., Carreira, J., & Arandjelović, R. (2022). Object discovery and representation networks. In *ECCV*.
- Huang, J., & Murphy, K. (2016). Efficient inference in occlusion-aware generative models of images. In *ICLR Workshops*.
- Jia, B., Liu, Y., & Huang, S. (2023). Improving object-centric learning with query optimization. In *ICLR*.
- Jiang, J., & Ahn, S. (2020). Generative neurosymbolic machines. In *NeurIPS*.
- Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Zitnick, C. L., & Girshick, R. (2017). CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*.
- Kabra, R., Burgess, C., Matthey, L., Kaufman, R. L., Greff, K., Reynolds, M., & Lerchner, A. (2019). Multi-object datasets. <https://github.com/deepmind/multi-object-datasets/>
- Karazija, L., Laina, I., & Rupprecht, C. (2021). ClevrTex: A texture-rich benchmark for unsupervised multi-object segmentation. In *NeurIPS*.
- Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. In *ICLR*.
- Kirillov, A., He, K., Girshick, R., Rother, C., & Dollár, P. (2019). Panoptic segmentation. In *CVPR*.
- LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., & Huang, F. (2006). A tutorial on energy-based learning. *Predicting Structured Data*, 1.
- Lin, T. Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., & Zitnick, C. L. (2014). Microsoft COCO: Common objects in context. In *ECCV*.
- Lin, Z., Wu, Y. F., Peri, S. V., Sun, W., Singh, G., Deng, F., Jiang, J., & Ahn, S. (2020). SPACE: Unsupervised object-oriented scene representation via spatial attention and decomposition. In *ICLR*.
- Liu, N., Li, S., Du, Y., Torralba, A., & Tenenbaum, J. B. (2022). Compositional visual generation with composable diffusion models. In *ECCV*.
- Locatello F, Weissenborn, D., Unterthiner, T., Mahendran, A., Heigold, G., Uszkoreit, J., Dosovitskiy, A., & Kipf, T. (2020). Object-centric learning with slot attention. In *NeurIPS*.
- Löwe, S., Lippe, P., Rudolph, M., & Welling, M. (2022). Complex-valued autoencoders for object discovery. In *TMLR*.
- Matthey, L., Higgins, I., Hassabis, D., & Lerchner, A. (2017). dSprites—Disentanglement testing Sprites dataset. <https://github.com/deepmind/dsprites-dataset/>
- Papa, S., Winther, O., & Dittadi, A. (2022). Inductive biases for object-centric representations in the presence of complex textures. [arXiv:2204.08479](https://arxiv.org/abs/2204.08479)
- Piper, J., & Granum, E. (1987). Computing distance transformations in convex and non-convex domains. *Pattern Recognition*, 20(6), 599–615.
- Polsby, D. D., & Popper, R. D. (1991). The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale L & Pol'y Rev*, 9.
- Rand, W. M. (1971). Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66(336), 846–850.
- Rezende, D. J. & Viola, F. (2018). Taming VAEs. [arXiv:1810.00597](https://arxiv.org/abs/1810.00597)
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18* (pp. 234–241). Springer.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., & Fei-Fei, L. (2015). ImageNet large scale visual recognition challenge. In *IJCV*.
- Schwartzberg, J. E. (1965). Reapportionment, gerrymanders, and the notion of compactness. *Minn. L. Rev.*, 50, 443.
- Seitzer, M., Horn, M., Zadaianchuk, A., Zietlow, D., Xiao, T., Simon-Gabriel, C. J., He, T., Zhang, Z., Schölkopf, B., Brox, T., & Locatello, F. (2022). Bridging the gap to real-world object-centric learning. In *ICLR 2023*.
- Seitzer, M., Horn, M., Zadaianchuk, A., Zietlow, D., Xiao, T., Simon-Gabriel, C. J., He, T., Zhang, Z., Schölkopf, B., Brox, T., & Locatello, F. (2023). Bridging the gap to real-world object-centric learning. In *ICLR*.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3), 379–423.
- Singh, G., Deng, F., & Ahn, S. (2022a). Illiterate DALL-E learns to compose. In *ICLR*.
- Singh, G., Wu, Y. F., & Ahn, S. (2022b). Simple unsupervised object-centric learning for complex and naturalistic videos. In *NeurIPS*.
- Sobel, I., & Feldman, G. (1973). A  $3 \times 3$  isotropic gradient operator for image processing. *Pattern Classification and Scene Analysis*, 271–272.
- Song, Z., & Yang, B. (2022). OGC: Unsupervised 3D object segmentation from rigid dynamics of point clouds. In *NeurIPS*.
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, 99(4), 605–632.
- Stelzner, K., Peharz, R., & Kersting, K. (2019). Faster attend-infer-repeat with tractable probabilistic models. In *ICML*.
- Tangemann, M., Schneider, S., von Kügelgen, J., Locatello, F., Gehler, P., Brox, T., Kümmerer, M., Bethge, M., & Schölkopf, B. (2021). Unsupervised object learning via common fate. [arXiv:2110.06562](https://arxiv.org/abs/2110.06562)
- van Steenkiste, S., Kurach, K., Schmidhuber, J., & Gelly, S. (2020). Investigating object compositionality in generative adversarial networks. *Neural Networks*.
- von Kügelgen, J., Ustyuzhaninov, I., Gehler, P., Bethge, M., & Schölkopf, B. (2020). Towards causal generative scene models via competition of experts. In *ICLR Workshops*.
- Voynov, A., Morozov, S., & Babenko, A. (2021). Object segmentation without labels with large-scale generative models. In *ICML*.
- Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A century of gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychological Bulletin*.
- Wang, W., Lai, Q., Fu, H., Shen, J., Ling, H., & Yang, R. (2021). Salient object detection in the deep learning era: An in-depth survey. In *TPAMI*.
- Wang, X., Girdhar, R., Yu, S. X., & Misra, I. (2023). Cut and learn for unsupervised object detection and instance segmentation. In *CVPR 2023* (pp. 3124–3134).
- Wang, X., Yu, Z., De Mello, S., Kautz, J., Anandkumar, A., Shen, C., & Alvarez, J. M. (2022). FreeSOLO: Learning to segment objects without annotations. In *CVPR 2022*.



- Watters, N., Matthey, L., Burgess, C. P., & Lerchner, A. (2019). Spatial broadcast decoder: A simple architecture for learning disentangled representations in VAEs. arXiv preprint [arXiv:1901.07017](https://arxiv.org/abs/1901.07017)
- Weis, M. A., Chitta, K., Sharma, Y., Brendel, W., Bethge, M., Geiger, A., & Ecker, A. S. (2021). Benchmarking unsupervised object representations for video sequences. In *JMLR*.
- Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt. *Psychologische Forschung*.
- Xu, J., De Mello S, Liu, S., Byeon, W., Breuel, T., Kautz, J., & Wang, X. (2022). GroupViT: Semantic segmentation emerges from text supervision. In *CVPR*.
- Xu, K., Li, C., Zhu, J., & Zhang, B. (2019). Multi-object generation with amortized structural regularization. In *NeurIPS*.
- Yang, Y., & Yang, B. (2022). Promising or elusive? Unsupervised object segmentation from real-world single images. In *NeurIPS*.
- Yu, Q., Wang, H., Qiao, S., Collins, M., Zhu, Y., Adam, H., Yuille, A., & Chen, L. C. (2022). k-Means mask transformer. In *ECCV*.
- Yuan, J., Li, B., & Xue, X. (2019). Generative modeling of infinite occluded objects for compositional scene representation. In *ICML*.
- Yuan, J., Chen, T., Li, B., & Xue, X. (2022). Compositional scene representation learning via reconstruction: A survey. [arXiv:2202.07135](https://arxiv.org/abs/2202.07135)
- Zimmermann, RS., van Steenkiste, S., Sajjadi, MS., Kipf, T., & Greff, K. (2023). Sensitivity of slot-based object-centric models to their number of slots. arXiv preprint [arXiv:2305.18890](https://arxiv.org/abs/2305.18890)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.