

# Explore the Power of Synthetic Data on Few-shot Object Detection

Shaobo Lin, Kun Wang, Xingyu Zeng, Rui Zhao  
Sensetime Research

{linshaobo, wangkun, zengxingyu, zhaorui}@sensetime.com

## Abstract

*Few-shot object detection (FSOD) aims to expand an object detector for novel categories given only a few instances for training. The few training samples restrict the performance of FSOD model. Recent text-to-image generation models have shown promising results in generating high-quality images. How applicable these synthetic images are for FSOD tasks remains under-explored. This work extensively studies how synthetic images generated from state-of-the-art text-to-image generators benefit FSOD tasks. We focus on two perspectives: (1) How to use synthetic data for FSOD? (2) How to find representative samples from the large-scale synthetic dataset? We design a copy-paste-based pipeline for using synthetic data. Specifically, saliency object detection is applied to the original generated image, and the minimum enclosing box is used for cropping the main object based on the saliency map. After that, the cropped object is randomly pasted on the image, which comes from the base dataset. We also study the influence of the input text of text-to-image generator and the number of synthetic images used. To construct a representative synthetic training dataset, we maximize the diversity of the selected images via a sample-based and cluster-based method. However, the severe problem of high false positives (FP) ratio of novel categories in FSOD can not be solved by using synthetic data. We propose integrating CLIP, a zero-shot recognition model, into the FSOD pipeline, which can filter 90% of FP by defining a threshold for the similarity score between the detected object and the text of the predicted category. Extensive experiments on PASCAL VOC and MS COCO validate the effectiveness of our method, in which performance gain is up to 21.9% compared to the few-shot baseline.*

## 1. Introduction

In recent years, we have tremendous progress in object detection [2, 4, 24, 33, 34]. However, the impressive performance of these models relies on a large amount of annotated data. Therefore, the detectors cannot generalize well

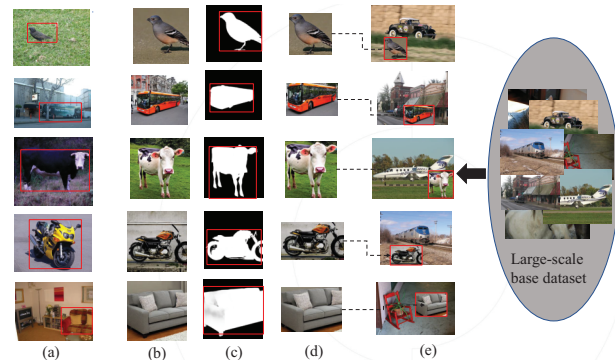


Figure 1. How to utilize synthetic data for FSOD? (a) The few-shot novel categories, including bird, bus, cow, motorbike, and sofa, are from the dataset of PASCAL VOC split 1. (b) The synthetic data from Stable Diffusion. (c) The saliency detection map of the synthetic data. (d) The cropped novel instances are based on the minimum enclosing box of the saliency map. (e) Combining synthetic instances and the randomly selected background from the base dataset.

to novel categories, especially when the annotated data are scarce. In contrast, humans can learn to recognize or detect a novel object with only a few labeled examples. Few-shot object detection (FSOD), which simulates this way, has attracted increasing attention. In FSOD, an object detector that is trained using base categories with sufficient data (base images) can learn to detect novel categories using only a few annotated images (novel images).

The few training samples restrict the performance of FSOD model. In FSOD, the number of objects for each novel category is  $K$  for  $K$ -shot detection. Existing literature always controls the number of novel data ( $K$ -shot), to carefully compare the performance of different methods. While this is helpful to drive research on sophisticated algorithms, it is unhelpful for the industry to enhance its few-shot based applications by incorporating external synthetic data. With the recent development of generative models, the text-to-image generator has made a great process. For example, DALL.E [32], Imagen [36], and Stable Diffusion [35] can generate high-quality images by simply using the input text description. These generators can pro-

duce diverse outcomes, implying a brighter future for industrial applications such as tackling many existing few-shot or long-tailed issues. This encourages us to explore the impact of synthetic data from text-to-image generators on FSOD tasks. Our study is carried out on the open-sourced Stable Diffusion. By using the generated images, we define a new setting of  $(K + G)$ -shot for the problem of using synthetic data for few-shot learning, consisting of  $K$  real novel instances and  $G$  generated novel instances. Two key questions must be answered: (1) How to use synthetic data for FSOD? (2) How to find representative samples from the large-scale synthetic dataset?

**How to use synthetic data for FSOD?** In Figure 1, the original FSOD dataset from PASCAL VOC split1 is annotated with the bounding boxes and categories. Since the synthetic images from Stable Diffusion lack annotations, we cannot utilize these data directly. Copy-Paste [9, 23, 28] is a simple and effective data augmentation strategy. Randomly pasting generated novel instances onto background images can generate a combinatorial number of training data for free. However, the original generated images can't be pasted onto the background image. The reason is that using the generated images, including the non-negligible background regions and the inaccurate bounding boxes, leads to performance degradation shown in Table 4. To solve this problem, saliency object detection is applied to the original generated images, and the minimum enclosing box is used for cropping the main object. After that, the cropped object is randomly down-sized and pasted on the image randomly selected from the base dataset. Besides, we design and compare several input text formats, finding that the performance of different text formats is related to the few-shot settings, shown in Figure 3. Intuitively, FSOD can benefit from more training data by increasing  $G$  in the  $(K + G)$ -shot setting. However, using more synthetic data can not continuously improve the FSOD model in our work. We find that using 20 synthetic images works best in most few-shot cases.

Since we can obtain a large amount of generated images from a text-to-image generator, **finding the representative samples** is necessary. Authenticity and diversity are two factors for evaluating the quality of a generated dataset. We focus on maximizing the diversity of training dataset for selecting the representative synthetic data because the authenticity of the images from Stable Diffusion is remarkable. We introduce the sample-based and cluster-based (k-means [13] / spectral clustering [26]) methods to build a high-diversity dataset. Specifically, the sample-based method applies uniform sampling to the data feature of each category. The cluster-based method uses a clustering algorithm to separate the data feature of each category into several sub-spaces and select one sample in each sub-space to maximize the diversity of training dataset. We se-

lect the sample nearest to the cluster centroid for each sub-space. CLIP [31] is used as the feature extractor for getting the data feature. However, the FP of novel categories is prominent, in which the base categories are often recognized as novel ones. The problem of high FP ratio of novel categories can not be solved by synthetic data, shown in Figure 4. CLIP is a zero-shot model that uses natural language to retrieve related images. We propose to integrate CLIP into the FSOD model for filtering the FP. By computing the similarity score between the detected object from the FSOD model and the name list of categories, we can remove the proposals whose similarity score with the predicted category is lower than the pre-defined threshold.

Our key contributions can be summarized as: (1) To the best of our knowledge, we are the first to leverage the external synthetic novel data from a text-to-image generator for FSOD, showing significant application potential. (2) We explore how to utilize the synthetic novel data, including the copy-paste pipeline, the design of input text, the number of used synthetic images, and the data selection strategy for discovering the representative synthetic samples. (3) Our method significantly improves multiple baselines and achieves state-of-the-art performance on PASCAL VOC and MS COCO.

## 2. Related Works

### 2.1. Few-shot Object Detection

FSOD is an important yet unsolved task in computer vision. Some works use meta-learning [8, 18, 21, 40, 44], where a meta-learner is introduced to acquire class-agnostic meta-knowledge which is transferred to novel classes. With the help of a meta learner that takes the support images as well as the bounding box annotations as inputs, the feature re-weighting modules are applied to a single-stage object detector (YOLOv2) [18] and a two-stage object detector (Faster R-CNN) [44]. [21] uses a class margin equilibrium (CME) approach, with the aim to optimize both feature space partition and novel class reconstruction in a systematic way. Transformation invariant principle [20] is proposed for various meta-learning models by introducing consistency regularization on predictions from the transformed images. TFA [39] is a simple two-stage fine-tuning approach, which significantly outperforms the earlier meta-learning methods. MPSR [42] adopts multi-scale positive sample refinement to handle scale variance problem. FSCE [38] aims to learn contrastive-aware object proposal encodings that facilitate the classification of detected objects. FADI [1] uses a two-step fine-tuning framework via association and discrimination, which builds up a discriminative feature space for each novel class with two integral steps. DeFRCN [29] extends Faster R-CNN by using gradient decoupled layer for multi-stage decoupling and

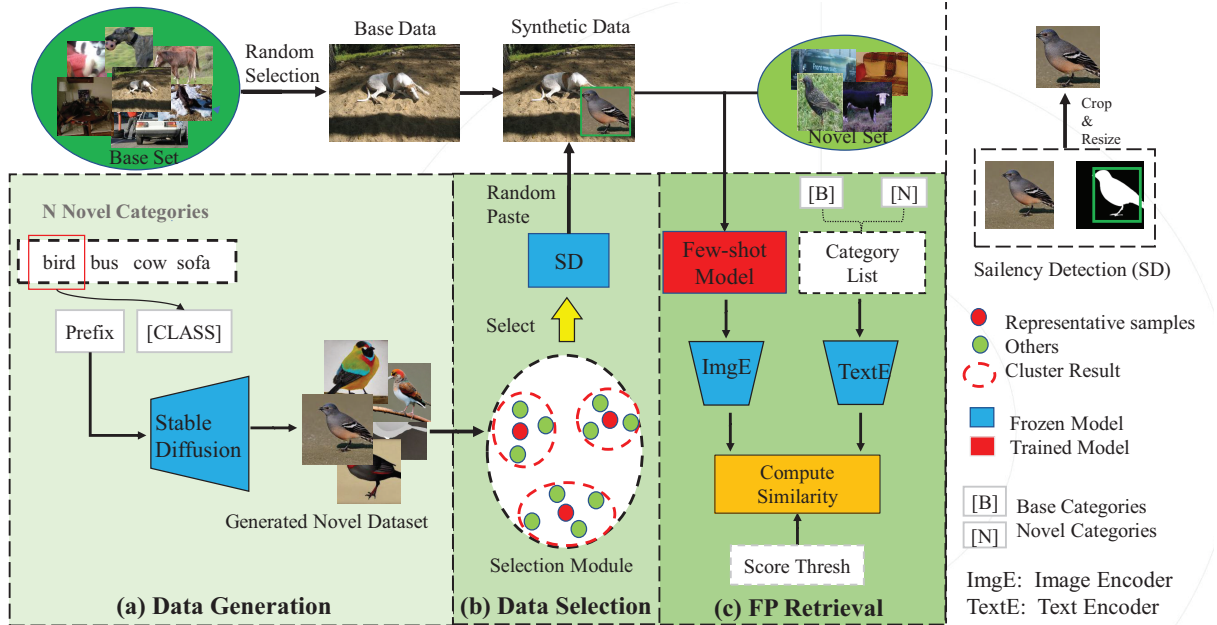


Figure 2. The pipeline of using synthetic data for FSOD. (a) During the data generation, the novel dataset is generated by Stable Diffusion. (b) In the data selection step, the generated data is processed by the selection module to find representative samples for training. We use the cluster-based method as an example. After that, a saliency detection is applied to the selected samples for cropping the novel instances. The synthetic images, added to the few-shot set for training the few-shot model, are obtained by randomly pasting novel instances onto the randomly selected background images. The Few-shot model can be any current few-shot object detector. (c) For FP retrieval, the similarity of the object feature from the image encoder of CLIP and the text feature from the text encoder of CLIP is computed. The score thresh filters the FP based on the similarity score.

prototypical calibration block for multi-task decoupling. Pseudo-Labeling [19] is proposed to find previously unlabelled instances to obtain high-quality pseudo-annotations for novel categories from the training dataset. MFDC [43] proposes to learn three types of class-agnostic commonalities between base and novel classes: recognition-related, localization-related and distribution commonalities, which are able to be integrated into most of existing fine-tuning based methods by performing distillation via a unified distillation framework.

## 2.2. Text-to-Image Generators

Recently, diffusion models [16, 37] have become a promising generative modeling framework, achieving state-of-the-art performance on image generation tasks [5, 17]. DALL.E [32] proposes a two-stage model: a prior that generates a CLIP [31] image embedding given a text caption and a decoder that generates an image based on the image embedding. Imagen [36] is a text-to-image diffusion model based on the power of language models and diffusion models. Using cross-attention layers, Stable Diffusion [35] uses diffusion models to synthesize high-resolution images for general conditioning inputs. GLIDE [27] studies diffusion models for the text-conditional image synthesis by comparing two guidance strategies: CLIP and classifier-free

guidance. The synthetic images generated from GLIDE are used for zero-shot, few-shot, and regular image recognition tasks [15], showing the power of synthetic data on image recognition. However, [15] does not consider the influence of the number of synthetic data defined by  $G$  in the  $(K+G)$ -shot setting. By comparing Stable Diffusion to GLIDE in Table 9, we notice that the quality of synthetic data from Stable Diffusion is superior to that of GLIDE. Therefore, we use Stable Diffusion in our experiments.

## 3. Method

### 3.1. Preliminary

In FSOD, given a labeled base dataset  $D_B$ , there are  $C_B$  base classes with sufficient images in each class. Novel dataset  $D_N$  with novel classes  $C_N$  consists of a few samples in each class.  $C_B$  and  $C_N$  do not have overlapping categories. The number of instances for each class in  $C_N$  is  $K$  for  $K$ -shot detection. Using the generated images, we define a new setting of  $(K+G)$ -shot for the problem of using synthetic data for FSOD, which consists of  $K$  real novel instances and  $G$  generated novel instances. There are two stages in FSOD methods [21, 38, 39]. In the pre-training stage, the model is trained on base classes to obtain a robust feature representation. In the fine-tuning stage, the pre-trained model is then fine-tuned on a balanced few-shot set

Table 1. The design of prefixes in the input text, including seven different settings: None, a, one, a5, one5, real, and adj.

Methods	Prefixs	Methods	Prefixs
None			
a	"a"	one	"one"
a5	"a"	one5	'one'
	'a photo of '		'a photo of one '
	'a photo of a '		'a picture of one '
	'a picture of '		'one photo of '
real	'a picture of a '	adj	'one picture of '
	'real '		'a photo of a good '
	'a real '		'a photo of a large '
	'one real '		'a photo of a nice '
	'a photo of a real '		'a photo of a cool '
	'a photo of one real '		'a photo of a clean '

(or adding the synthetic novel data in our setting), including both base and novel classes ( $C_B \cup C_N$ ).

### 3.2. Our Pipeline

In Figure 2, there are three steps in our overall pipeline of using synthetic data for FSOD: data generation, data selection, and FP retrieval.

**Data Generation** To study the influence of the input text of text-to-image generator, several representative text descriptions for Stable Diffusion are designed by following the format of prefix with the name of a category, shown in Table 1. Specifically, We can simply use the class name as the input text whose prefix is "None", or "a / one" with the class name. Besides, we can combine "a / one" and "photo / picture" or use some additional representative adjectives as the prefix of input text. The comparison among these text formats will be discussed in section 4.4.2. Using the input text, we can get many synthetic images according to the category name. Because the entity of the input text of Stable Diffusion is the name of a category, most of the generated images only have one target.

**Data Selection** The generated novel data from Stable Diffusion is processed by the selection module to find the representative synthetic samples, which should be high-quality and high-diversity. We focus on maximizing the diversity of the selected data because the quality of the synthetic images is high enough. The selection module is based on the CLIP similarity score or image feature from CLIP. We propose two types of selection methods, including sample-based and cluster-based methods.

$$G_n = \{G_n^i\}_{i=1}^N, R_n = \{R_n^i\}_{i=1}^M \quad (1)$$

where  $G_n$  includes the generated novel instances of category  $n$ .  $G_n^i$  is the  $i$ -th generated sample of category  $n$ .  $R_n$  is the real novel data of category  $n$  from the few-shot dataset.  $R_n^i$  is the  $i$ -th real sample of category  $n$ .  $N$  and  $M$  are the numbers of generated and real novel images, respectively.

The sample-based method is directly based on the CLIP score or image feature of CLIP. As illustrated in Equ 2,

the diversity of the generated dataset can be measured by the CLIP score of the generated instances or the distance between the CLIP features of generated and real novel instances. High-diversity means using uniform sampling to maximize the difference among the selected samples.

$$R_s = \begin{cases} Se(CLIP(G_n, T)) \\ Se(cos(CLIP(G_n), CLIP(R_n))) \end{cases} \quad (2)$$

where  $T$  is the input text. By inputting the text and image,  $CLIP$  outputs the similarity score. If  $T$  is not used,  $CLIP$  means the image encoder of CLIP, which processes the generated and real novel images for obtaining the image features.  $Se$  is the selection module for finding the representative samples using uniform sampling. The distance between the generated images and real images is computed by the cosine similarity.  $R_s$  is the selected samples based on the sample-based methods.

Using a clustering algorithm is the second way to find the representative samples.

$$C_n = Clu(CLIP(I_n), k_n) \quad (3)$$

where  $C_n$  is the class center of category  $n$ .  $Clu$  is the clustering method.  $k_n$  is the number of centers of category  $n$ . We use k-means or spectral clustering to compute the class center. We find that using the spectral clustering algorithm achieves the best performance.

$$R_c = \arg \min_i cos(C_n, G_n^i)_{i=1}^N \quad (4)$$

We select the sample nearest to the  $k_n$  cluster centroids  $C_n$  as the representative samples for category  $n$ . The distance is computed by the cosine similarity.  $R_c$  is the samples selected by the cluster-based methods.

After finding the representative samples, we apply a saliency detection method to these generated images. Then, the minimum enclosing box is used for cropping the main object based on the saliency map. Finally, the cropped object is down-sized and pasted on the randomly selected background from the base dataset. These synthetic novel images are introduced into the original few-shot set to obtain a new few-shot training dataset. The number of the selected synthetic images,  $G$  in the setting of  $(K + G)$ -shot, is a crucial hyper-parameter for FSOD, which will be discussed in section 4.4.2.

**FP Retrieval** In our work, the FP of novel categories is prominent. However, the problem of high FP ratio can not be solved by introducing synthetic data, which is shown in Figure 4 in section 4.4.4. With the increase of generated images, the FP ratio of the original few-shot model ( $w/o$  CLIP) does not decrease. We propose integrating the CLIP into the FSOD model for filtering the FP. Formally, given  $I_n$  which is the few-shot training dataset of the novel category

$n$ , we use the FSOD model  $F_s$  to get the bounding box  $bb_n$  and the corresponding confidence score  $conf_n$ .

$$R_n = I_n[bb_n]; conf_n, bb_n = F_s(I_n) \quad (5)$$

where  $R_n$  is the cropped region of  $I_n$  based on  $bb_n$ , which is predicted to category  $n$  by  $F_s$ .

$$score = \frac{e^{\cos(\text{ImgE}(R_n), \text{TextE}(t_n))}}{\sum_{i=1}^L e^{\cos(\text{ImgE}(R_n), \text{TextE}(t_i))}} \quad (6)$$

where  $\text{ImgE}$  and  $\text{TextE}$  are the image encoder and text encoder of CLIP.  $\text{ImgE}$  gets the image feature of the detected object  $R_n$ .  $\text{TextE}$  gets the text feature of the input text  $t_n$  of the novel category  $n$ . We compute the cosine similarity of the image and text features. The input text  $t_i \in (B \cup N)$  or  $N$  means the base and novel categories or only novel categories.  $L$  is the length of the category list. In Table 8, we find that the selection of the category list is related to the scale of the dataset, i.e., PASCAL VOC and MS COCO. Finally, we use the softmax operation to get the normalized similarity  $score$ . We remove the detection results of  $R_n$  (FP) whose  $score$  is lower than the pre-defined CLIP threshold, which is 0.1 in our experiment. The filtered results are used for final evaluation. With CLIP, we can reduce the FP ratio up to 90 % shown in Figure 4.

## 4. Experiments

### 4.1. Datasets and Evaluation Protocols

We evaluate our methods on PASCAL VOC [6, 7] and MS COCO [25]. In PASCAL VOC, we adopt the common strategy [33, 34] that using VOC 2007 test set for evaluating while VOC 2007 and 2012 train/val sets are used for training. Following [44], 5 out of its 20 object categories are selected as the novel classes, while the remaining as the base classes. We evaluate with three different novel/base splits from [44], named as split 1, split 2 and split 3. Each split contains 15 base categories with abundant data and 5 novel categories with K annotated instances for K = 1, 3, 5, 10. Following [38, 39, 44], we use the mean average precision (mAP) of novel categories at 0.5 IoU threshold as the evaluation metric and report the results on the official test set of VOC 2007. When using MS COCO, 20 out of 80 categories are reserved as novel classes, the rest 60 categories are used as base classes. The detection performance with COCO-style AP and AP<sub>75</sub> for K = 1, 3, 5, 10, 30 shots of novel categories are reported.

### 4.2. Implementation Details

Our baselines include DeFRCN [29] and MFDC [43], which are representative methods in FSOD. These methods use Faster R-CNN [34] with ResNet-101 [14], which are the same as almost all FSOD methods. The training strategies

of our methods follow the selected baselines. The difference is that we add the generated novel images to the few-shot set for training. The  $G$  in the setting of  $(K + G)$ -shot is 20. We adopt the ViT-B/32 [31] as the encoder of CLIP.

### 4.3. Comparison with State-of-the-art Methods

We compare our approaches to several competitive FSOD methods. The results are shown in Table 2 and Table 3. Following [1, 11, 19, 21, 42], we use a single run with the same training images to get the results of different shots.

#### 4.3.1 Results on PASCAL VOC and MS COCO.

Following [38, 39, 44], we provide the AP<sub>50</sub> of the novel classes on PASCAL VOC with three splits in Table 2. By using synthetic data, our methods can outperform the baselines in all few-shot settings and achieve state-of-the-art performance. The most obvious improvement is up to 21.9%.

We report the COCO-style AP and AP<sub>75</sub> of the 20 novel classes on MS COCO in Table 3. By using synthetic data, our methods can achieve state-of-the-art performance in almost all settings. The performance gain in the setting of 30-shot is low, since higher shots decline the influence of the synthetic images from another source which is different from the data from the current domain.

### 4.4. Ablation Study

Following [38, 44], we do ablation studies on PASCAL VOC split 1. We use DeFRCN as the baseline.

#### 4.4.1 The Components in Our Pipeline

There are four essential steps for applying synthetic data to FSOD: Add, Crop, Select, and CLIP. In Table 4, Add copies the original synthetic images and pasts them onto the background. Crop means using saliency detection for the main target before the pasting operation. Select is finding the top 20 representative samples from the generated dataset. CLIP uses the encoders of CLIP after the few-shot model to filter the FP. We discover that cropping is crucial, which can achieve 4% to 10% AP<sub>50</sub> improvement in PASCAL VOC, since directly using the original synthetic data hurt the performance due to the inaccurate bounding boxes. Our selection strategy can improve the precision on most settings (up to 2% in PASCAL VOC). CLIP helps the model to filter most of FP and improve the performance up to 4%. The best performance on MS COCO is also achieved by using all components.

#### 4.4.2 The Factors in Data Generation and Selection

There are three factors: the format of input text description, the number of used synthetic images ( $G$  in the setting

Table 2. Comparison with state-of-the-art few-shot object detection methods on VOC2007 test set for novel classes of the three splits. **Black** indicates state-of-the-art. **Red** is the improvement compared to the baseline. \* means using our methods: the synthetic images, the selection strategy, and the CLIP model for filtering the FP.

Methods / Shots		split 1				split 2				split 3			
		1	3	5	10	1	3	5	10	1	3	5	10
FRCN+ft [44]	ICCV2019	11.9	29	36.9	36.9	5.9	23.4	29.1	28.8	5.0	18.1	30.8	43.4
FRCN+ft-full [44]	ICCV2019	13.8	32.8	41.5	45.6	7.9	26.2	31.6	39.1	9.8	19.1	35	45.1
FR [18]	ICCV2019	14.8	26.7	33.9	47.2	15.7	22.7	30.1	40.5	21.3	28.4	42.8	45.9
MetaDet [40]	ICCV2019	18.9	30.2	36.8	49.6	21.8	27.8	31.7	43	20.6	29.4	43.9	44.1
Meta R-CNN [44]	ICCV2019	19.9	35	45.7	51.5	10.4	29.6	34.8	45.4	14.3	27.5	41.2	48.1
TFA [39]	ICML2020	39.8	44.7	55.7	56.0	23.5	34.1	35.1	39.1	30.8	42.8	49.5	49.8
MPSR [42]	ECCV2020	41.7	51.4	55.2	61.8	24.4	39.2	39.9	47.8	35.6	42.3	48	49.7
CME [21]	CVPR2021	41.5	50.4	58.2	60.9	27.2	41.4	42.5	46.8	34.3	45.1	48.3	51.5
FSCN [22]	CVPR2021	40.7	46.5	57.4	62.4	27.3	40.8	42.7	46.3	31.2	43.7	50.1	55.6
HallucFsDet [46]	CVPR2021	47	46.5	54.7	54.7	26.3	37.4	37.4	41.2	40.4	43.3	51.4	49.6
FSCE [38]	CVPR2021	44.2	51.4	61.9	63.4	27.3	43.5	44.2	50.2	37.2	47.5	54.6	58.5
UPE [41]	ICCV2021	43.8	50.3	55.4	61.7	31.2	41.2	42.2	48.3	35.5	43.9	50.6	53.5
QA-FewDet [10]	ICCV2021	42.4	55.7	62.6	63.4	25.9	46.6	48.9	51.1	35.2	47.8	54.8	53.5
Meta faster-rcnn [11]	AAAI2021	43	60.6	66.1	65.4	27.7	46.1	47.8	51.4	40.6	53.4	<b>59.9</b>	58.6
FADI [1]	NIPS2021	50.3	54.2	59.3	63.2	30.6	40.3	42.8	48	45.7	49.1	55	59.6
DeFRCN [29]	ICCV2021	53.6	61.5	64.1	60.8	30.1	47.0	53.3	47.9	48.4	52.3	54.9	57.4
FCT [12]	CVPR2022	49.9	57.9	63.2	67.1	27.6	43.7	49.2	51.2	39.5	52.3	57	58.7
KFSOD [45]	CVPR2022	44.6	54.4	60.9	65.8	37.8	43.1	48.1	50.4	34.8	44.1	52.7	53.9
Pseudo-Labeling [19]	CVPR2022	54.5	58.8	63.2	65.7	32.8	50.7	49.8	50.6	48.4	55.0	59.6	59.6
MFDC [43]	ECCV2022	63.4	67.7	69.4	68.1	42.1	53.4	55.3	53.8	56.1	59.0	62.2	63.7
DeFRCN*	Ours	67.5	<b>69.8</b>	<b>71.1</b>	71.5	<b>52</b>	54.3	57.5	<b>57.4</b>	55.9	58.6	59.6	63.9
	Improve	<b>+13.9</b>	<b>+8.3</b>	<b>+7</b>	<b>+10.7</b>	<b>+21.9</b>	<b>+7.3</b>	<b>+4.2</b>	<b>+9.5</b>	<b>+7.5</b>	<b>+6.3</b>	<b>+4.7</b>	<b>+6.5</b>
MFDC*	Ours	<b>68.9</b>	69.5	70.9	<b>73.3</b>	50.5	<b>56.8</b>	<b>58.9</b>	55.6	<b>59.9</b>	<b>62.8</b>	<b>65.3</b>	<b>66.9</b>
	Improve	<b>+5.4</b>	<b>+1.8</b>	<b>+1.5</b>	<b>+5.2</b>	<b>+8.4</b>	<b>+3.4</b>	<b>+3.6</b>	<b>+1.8</b>	<b>+3.8</b>	<b>+3.8</b>	<b>+3.1</b>	<b>+3.2</b>

Table 3. Few-shot object detection performance on MS COCO. **Black** indicates the state-of-the-art. **Red** is the improvement compared to the baseline. \* means using our methods: the synthetic images, the selection strategy, and the CLIP model for filtering the FP.

Methods / Shots	1		3		5		10		30	
	nAP	nAP <sub>75</sub>	nAP	nAP <sub>75</sub>	nAP	nAP <sub>75</sub>	nAP	nAP <sub>75</sub>	nAP	nAP <sub>75</sub>
FR [18]	-	-	-	-	-	-	5.6	4.6	9.1	7.6
Meta R-CNN [44]	-	-	-	-	-	-	8.7	6.6	12.4	10.8
CME [21]	-	-	-	-	-	-	15.1	16.4	16.9	17.8
FSCE [38]	-	-	-	-	-	-	11.1	9.8	15.3	14.2
UPE [41]	-	-	-	-	-	-	11	10.7	15.6	15.7
TFA [39]	3.4	3.8	6.6	6.5	8.3	8.0	10	9.3	13.7	13.4
MSPR [42]	2.3	2.3	5.2	5.1	6.7	6.4	9.8	9.7	14.1	14.2
QA-FewDet [10]	4.9	4.4	8.4	7.3	9.7	8.6	10.2	9.0	11.5	10.3
Meta faster-rcnn [11]	5.0	4.6	-	-	12.7	16.6	25.7	31.8	10.8	15.8
FADI [1]	5.7	6.0	8.6	8.3	10.1	9.7	12.2	11.9	16.1	15.8
DeFRCN [29]	6.5	6.9	13.4	13.6	15.3	14.6	18.5	17.6	22.6	22.3
FCT [12]	5.6	-	11.1	-	14.0	-	17.1	17	21.4	22.1
KFSOD [45]	-	-	-	-	-	-	18.5	18.7	-	-
Pseudo-Labeling [19]	-	-	-	-	-	-	17.8	17.8	<b>24.5</b>	<b>25.0</b>
MFDC [43]	10.8	11.6	15	15.5	16.4	17.3	19.4	20.2	22.7	23.2
DeFRCN*	<b>18.1</b>	<b>19.5</b>	<b>19.2</b>	<b>20.2</b>	<b>19.8</b>	<b>20.4</b>	<b>20.7</b>	<b>21.3</b>	23.1	23.9
Improve	<b>+11.6</b>	<b>+12.6</b>	<b>+5.8</b>	<b>+6.6</b>	<b>+4.5</b>	<b>+5.8</b>	<b>+2.2</b>	<b>+3.7</b>	<b>+0.5</b>	<b>+1.6</b>
MFDC*	<b>18.1</b>	19	18.6	19.4	19.5	20.1	20.3	20.9	22.5	23
Improve	<b>+7.3</b>	<b>+7.4</b>	<b>+3.6</b>	<b>+3.9</b>	<b>+3.1</b>	<b>+2.8</b>	<b>+0.9</b>	<b>+0.7</b>	-0.2	-0.2

of  $(K + G)$ -shot), and the type of pasting generated novel instances.

In Figure 3, we analyze the influence of the number of used synthetic images and the input text description for Stable Diffusion. As a result, we can get the following conclusions. First, almost all input text formats perform best with 20 synthetic images. Second, except for the format of "one", other formats have relatively high performance on average. In detail, "a5" and "one5" achieve the highest performance on average. "a5" has the most stable curve. Therefore, in other experiments, we use "a5" format with

20 synthetic images per class. In Table 5, we compare three different pasting methods: box, saliency map, and segmentation map, in which these types of novel instances are pasted onto the background. The saliency map uses the results of BASNet [30]. The segmentation map is from a supervised image segmentation method: deeplabv3+ [3]. Box uses the minimum enclosing rectangle of the saliency map. The results are averaged by using 10 to 200 generated images for training. In conclusion, the box type achieves the best performance in all settings.

Table 4. The components of our methods based on DeFRCN. The AP<sub>50</sub> of DeFRCN on 1 to 10-shot of PASCAL VOC split1 and MS-COCO is reported.

DeFRCN (Our impl.) [29]	Add	Crop	Select	CLIP	1-shot	3-shot	5-shot	10-shot
on PASCAL VOC					52.5	60.8	62.7	61.9
	✓				51.9	60	61.9	63.6
	✓	✓			62.2	65.5	66.5	67.8
	✓	✓	✓	✓	61.8	66.6	66.3	67.4
	✓	✓	✓	✓	63.5	66.3	67.5	69
	✓	✓	✓	✓	<b>67.5</b>	<b>69.8</b>	<b>71.1</b>	<b>71.5</b>
on MS-COCO					16.2	25.4	28.9	33.6
	✓	✓			25.4	28	30.1	33.8
	✓	✓	✓		27	30	32.3	34.7
	✓	✓	✓	✓	<b>29.2</b>	<b>32.2</b>	<b>33.1</b>	<b>35.1</b>

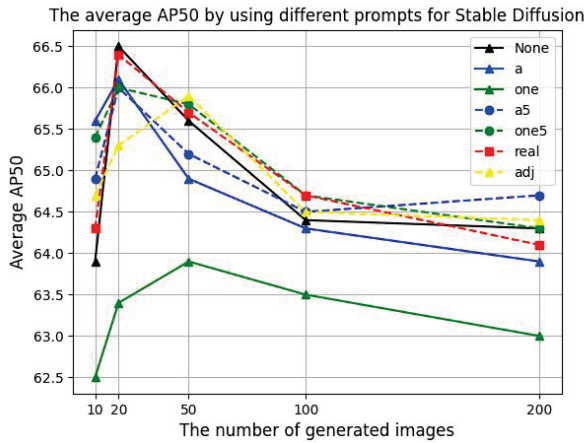


Figure 3. The influence of different input text for Stable Diffusion and the number of synthetic data on the average AP<sub>50</sub> of PASCAL VOC split1 1 to 10-shot. There are seven types of text descriptions: None, a, a5, one, one5, adj, and real. The X-axis is the number of generated images ( $G$ ) for each novel category.

Table 5. The type of several pasting methods. The results are averaged by using 10 to 200 generated images for training on 1 to 10-shot of PASCAL VOC split1. The saliency map is the result of saliency object detection. The segmentation map is the result of a supervised image segmentation method. Box uses the minimum enclosing rectangle of the saliency map.

Methods	1-shot	3-shot	5-shot	10-shot
Box	<b>62.7</b>	<b>65.2</b>	<b>65.4</b>	<b>66.8</b>
Saliency map	60.4	63.1	63.9	65.5
Segmentation map	61.5	63.9	64.2	66

#### 4.4.3 The Selection Module

Table 6 compares different synthetic data selection methods. All selection methods choose 20 samples out of 200 generated images per class. The features of samples (generated or original few-shot samples) are based on the image encoder of CLIP. We find that only considering the quality of synthetic data performs worse, such as CLIP\_max,

Table 6. The comparison of different data selection methods. All selection methods choose 20 samples out of 200 generated images per class. The features of samples are obtained from the image encoder of CLIP. The results of the random selection are averaged by five different selections. Syn\_max selects the top 20 samples which have the minimum distance to the mean features of all generated images. CLIP\_max selects the samples with the top-20 CLIP score. Instance\_max selects the nearest neighbor samples according to the mean feature of few-shot instances. CLIP\_uniform is uniform sampling based on the CLIP score. Instance\_uniform means uniform sampling based on the distance between the mean features of the few-shot instances and the features of generated images. Kmeans\_cluster and Spectral\_cluster get 20 clusters for each category to select the nearest neighbor sample for each cluster. **Black** is the best, and **Red** is the second best.

Types	Methods	1-shot	3-shot	5-shot	10-shot
Others	Random	62.6	65.9	66.8	68.2
	Syn_max	61.6	64.6	66.6	68.5
	CLIP_max	59.9	64.2	65.9	67.2
Quality-based	Instance_max	61.6	65.2	67.2	<b>68.6</b>
	CLIP_uniform	<b>63.4</b>	66.1	66.7	<b>68.6</b>
Sample-based	Instance_uniform	60.6	64.7	<b>67.4</b>	68.3
	Kmeans_cluster	62.6	<b>66.6</b>	66.6	67.9
Cluster-based	Spectral_cluster	<b>63.5</b>	<b>66.3</b>	<b>67.5</b>	<b>69</b>




Figure 4. The power of CLIP on reducing the number of FP. The generated images are added to the 1-shot setting of PASCAL VOC split1. w/o CLIP means the original few-shot model. w CLIP uses CLIP to filter the FP. The increase of the generated images used does not reduce the FP ratio. The FP ratio of using CLIP is much lower than that without CLIP.

since this operation reduces the diversity of data. Focusing on the high diversity of synthetic data can improve the few-shot model, such as using a clustering or uniform sampling method based on the CLIP score. The performance of Instance\_uniform is not satisfactory because there is a large deviation between the mean feature of few-shot samples and the actual prototype of a category. The best performance is achieved by spectral clustering.

Table 7. The filtering thresh for reducing the FP by using CLIP. The CLIP thresh is applied to CLIP score, which is lower than the CLIP thresh is defined as FP.

DeFRCN (Our impl.) [29]	Thresh	1-shot	3-shot	5-shot	10-shot
CLIP Thresh	0.5	65.2	68.4	68.1	69.6
	0.3	65.8	68.6	67.5	69.2
	0.2	66.3	<b>70.2</b>	70	71.2
	0.1	<b>67.5</b>	69.8	<b>71.1</b>	<b>71.5</b>
	0	63.5	66.3	67.5	69

Table 8. The category list for CLIP. The AP<sub>50</sub> of DeFRCN on 1 to 10-shot of PASCAL VOC split1 and MS-COCO is reported. The input of the text encoder of CLIP can be both the base and novel classes or the novel classes only. The experiments are conducted on PASCAL VOC and MS-COCO.

DeFRCN (Our impl.) [29]	Type	1-shot	3-shot	5-shot	10-shot
on PASCAL VOC	On Novel	62.2	66.3	65.7	68.1
	On Base+Novel	<b>66.3</b>	<b>70.2</b>	<b>70</b>	<b>71.2</b>
on MS-COCO	On Novel	<b>29.2</b>	<b>32.2</b>	<b>33.1</b>	<b>35.1</b>
	On Base+Novel	27.4	29.1	30.8	31.5

#### 4.4.4 How to Use CLIP

In Figure 4, increasing the generated images can not reduce the FP ratio. The reduction of FP ratio by using CLIP is up to 90%. In Table 8, we compare two category lists used as the input of the text-encoder of CLIP: novel and base categories or only the novel categories. Using novel and base categories means that CLIP calculates the similarity between the input image and the input text of all categories (20 in PASCAL VOC and 80 in MS COCO). When using novel categories, CLIP calculates the similarity between the input image and the input text of novel categories (5 in PASCAL VOC and 20 in MS COCO). The results show that the optimal length of the name list is 20, which is all categories of PASCAL VOC and the novel categories of MS COCO. Table 7 shows the effect of CLIP thresh on performance. The detected objects with the CLIP score of the predicted categories, which are lower than the CLIP threshold, are defined as false positives. These predictions are removed for final evaluation. The optimal CLIP threshold is 0.1.

#### 4.4.5 Other Analysis

In Figure 5, we study the upper bound of the performance of using synthetic data. With the increase of synthetic data, copy-paste-based methods, such as pasting real novel instances or synthetic instances, have a lower upper bound than using the original novel images. We can notice that the gap between pasting real and synthetic instances is marginal, which means the synthetic images are ready for the FSOD tasks. Using more synthetic data can boost accuracy when the number of data is lower than 50, but the gain does not continue with the increase of generated data. The performance of FSOD is mainly limited by the copy-paste




Figure 5. The influence of different types of additional images per category: the images with the pasted synthetic novel instance (copy-paste S), the images with the pasted real novel instance (copy-paste R), and the real novel images (R). The results are obtained by adding these images to the original 1-shot set of PASCAL VOC split1 for training the DeFRCN.

Table 9. The comparison between different text-to-image generators: Stable Diffusion and GLIDE. The results of both methods are averaged by training five times. Twenty samples are randomly selected for each training.

Methods	1-shot	3-shot	5-shot	10-shot
Stable Diffusion [35]	<b>62.6</b>	<b>65.9</b>	<b>66.8</b>	<b>68.2</b>
GLIDE [27]	60.3	64.7	64.1	66.1

itself, not the synthetic data.

We compare another text-to-image generator (GLIDE) to Stable Diffusion. The results in Table 9 show that the quality of the text-to-image generator affects the performance of FSOD.

## 5. Conclusion

In this work, we leverage the external synthetic novel data from a text-to-image generator for FSOD. We focus on utilizing the synthetic novel data, including the copy-paste pipeline, the design of input text for Stable Diffusion, the number of used synthetic images, and the data selection strategy for discovering the representative synthetic data. Furthermore, we propose integrating CLIP into the FSOD model to solve the problem of high FP ratio in FSOD. Extensive experiments on the few-shot object detection datasets, i.e., PASCAL VOC and MS COCO, validate the effectiveness of our method.

## References

- [1] Yuhang Cao, Jiaqi Wang, Ying Jin, Tong Wu, Kai Chen, Ziwei Liu, and Dahua Lin. Few-shot object detection via asso-



- ciation and discrimination. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 5, 6
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 1
- [3] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 6
- [4] Xiyang Dai, Yinpeng Chen, Bin Xiao, Dongdong Chen, Mengchen Liu, Lu Yuan, and Lei Zhang. Dynamic head: Unifying object detection heads with attentions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7373–7382, 2021. 1
- [5] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 3
- [6] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015. 5
- [7] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 5
- [8] Qi Fan, Wei Zhuo, Chi-Keung Tang, and Yu-Wing Tai. Few-shot object detection with attention-rpn and multi-relation detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4013–4022, 2020. 2
- [9] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2918–2928, 2021. 2
- [10] Guangxing Han, Yicheng He, Shiyuan Huang, Jiawei Ma, and Shih-Fu Chang. Query adaptive few-shot object detection with heterogeneous graph convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3263–3272, 2021. 6
- [11] Guangxing Han, Shiyuan Huang, Jiawei Ma, Yicheng He, and Shih-Fu Chang. Meta faster r-cnn: Towards accurate few-shot object detection with attentive feature alignment. *arXiv preprint arXiv:2104.07719*, 2021. 5, 6
- [12] Guangxing Han, Jiawei Ma, Shiyuan Huang, Long Chen, and Shih-Fu Chang. Few-shot object detection with fully cross-transformer. *arXiv preprint arXiv:2203.15021*, 2022. 6
- [13] John A Hartigan and Manchek A Wong. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108, 1979. 2
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [15] Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and Xiaojuan Qi. Is synthetic data from generative models ready for image recognition? *arXiv preprint arXiv:2210.07574*, 2022. 3
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 3
- [17] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23:47–1, 2022. 3
- [18] Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8420–8429, 2019. 2, 6
- [19] Prannay Kaul, Weidi Xie, and Andrew Zisserman. Label, verify, correct: A simple few shot object detection method. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14237–14247, 2022. 3, 5, 6
- [20] Aoxue Li and Zhenguo Li. Transformation invariant few-shot object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3094–3102, 2021. 2
- [21] Bohao Li, Boyu Yang, Chang Liu, Feng Liu, Rongrong Ji, and Qixiang Ye. Beyond max-margin: Class margin equilibrium for few-shot object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7363–7372, 2021. 2, 3, 5, 6
- [22] Yiting Li, Haiyue Zhu, Yu Cheng, Wenxin Wang, Chek Sing Teo, Cheng Xiang, Prahlad Vadakkepat, and Tong Heng Lee. Few-shot object detection via classification refinement and distractor retreatment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15395–15403, 2021. 6
- [23] Shaobo Lin, Kun Wang, Xingyu Zeng, and Rui Zhao. An effective crop-paste pipeline for few-shot object detection. *arXiv preprint arXiv:2302.14452*, 2023. 2
- [24] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 1
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 5
- [26] Andrew Ng, Michael Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems*, 14, 2001. 2
- [27] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 3, 8

- [28] Seulki Park, Youngkyu Hong, Byeongho Heo, Sangdoon Yun, and Jin Young Choi. The majority can help the minority: Context-rich minority oversampling for long-tailed classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6887–6896, 2022. [2](#)
- [29] Limeng Qiao, Yuxuan Zhao, Zhiyuan Li, Xi Qiu, Jianan Wu, and Chi Zhang. Defrcn: Decoupled faster r-cnn for few-shot object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8681–8690, 2021. [2](#), [5](#), [6](#), [7](#), [8](#)
- [30] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. Basnet: Boundary-aware salient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7479–7489, 2019. [6](#)
- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [2](#), [3](#), [5](#)
- [32] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. [1](#), [3](#)
- [33] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017. [1](#), [5](#)
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2016. [1](#), [5](#)
- [35] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [1](#), [3](#), [8](#)
- [36] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. [1](#), [3](#)
- [37] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. [3](#)
- [38] Bo Sun, Banghuai Li, Shengcai Cai, Ye Yuan, and Chi Zhang. Fsce: Few-shot object detection via contrastive proposal encoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7352–7362, 2021. [2](#), [3](#), [5](#), [6](#)
- [39] Xin Wang, Thomas E Huang, Trevor Darrell, Joseph E Gonzalez, and Fisher Yu. Frustratingly simple few-shot object detection. *arXiv preprint arXiv:2003.06957*, 2020. [2](#), [3](#), [5](#), [6](#)
- [40] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Meta-learning to detect rare objects. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9925–9934, 2019. [2](#), [6](#)
- [41] Aming Wu, Yahong Han, Linchao Zhu, and Yi Yang. Universal-prototype enhancing for few-shot object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9567–9576, 2021. [6](#)
- [42] Jiayi Wu, Songtao Liu, Di Huang, and Yunhong Wang. Multi-scale positive sample refinement for few-shot object detection. In *European Conference on Computer Vision*, pages 456–472. Springer, 2020. [2](#), [5](#), [6](#)
- [43] Shuang Wu, Wenjie Pei, Dianwen Mei, Fanglin Chen, Jiandong Tian, and Guangming Lu. Multi-faceted distillation of base-novel commonality for few-shot object detection. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, pages 578–594. Springer, 2022. [3](#), [5](#), [6](#)
- [44] Xiaopeng Yan, Ziliang Chen, Anni Xu, Xiaoxi Wang, Xiaodan Liang, and Liang Lin. Meta r-cnn: Towards general solver for instance-level low-shot learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9577–9586, 2019. [2](#), [5](#), [6](#)
- [45] Shan Zhang, Lei Wang, Naila Murray, and Piotr Koniusz. Kernelized few-shot object detection with efficient integral aggregation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19207–19216, 2022. [6](#)
- [46] Weilin Zhang and Yu-Xiong Wang. Hallucination improves few-shot object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13008–13017, 2021. [6](#)