

GTNet: Generative Transfer Network for Zero-Shot Object Detection

Shizhen Zhao,¹ Changxin Gao,^{1*} Yuanjie Shao,¹ Lerenhan Li,¹
Changqian Yu,¹ Zhong Ji,² Nong Sang¹

¹Key Laboratory of Ministry of Education for Image Processing and Intelligent Control,
School of Artificial Intelligence and Automation, Huazhong University of Science and Technology

²School of Electrical and Information Engineering, Tianjin University

Email: {zhaosz, cgao}@hust.edu.cn

Abstract

We propose a Generative Transfer Network (GTNet) for zero-shot object detection (ZSD). GTNet consists of an Object Detection Module and a Knowledge Transfer Module. The Object Detection Module can learn large-scale seen domain knowledge. The Knowledge Transfer Module leverages a feature synthesizer to generate unseen class features, which are applied to train a new classification layer for the Object Detection Module. In order to synthesize features for each unseen class with both the intra-class variance and the IoU variance, we design an IoU-Aware Generative Adversarial Network (IoUGAN) as the feature synthesizer, which can be easily integrated into GTNet. Specifically, IoUGAN consists of three unit models: Class Feature Generating Unit (CFU), Foreground Feature Generating Unit (FFU), and Background Feature Generating Unit (BFU). CFU generates unseen features with the intra-class variance conditioned on the class semantic embeddings. FFU and BFU add the IoU variance to the results of CFU, yielding class-specific foreground and background features, respectively. We evaluate our method on three public datasets and the results demonstrate that our method performs favorably against the state-of-the-art ZSD approaches.

1. Introduction

In recent years, many deep learning methods have achieved desirable performance in object detection (Girshick et al. 2014; Girshick 2015; Ren et al. 2015; Redmon et al. 2016; Liu et al. 2016; He et al. 2017). However, the performance of the detectors relies on the large-scale detection datasets with fully-annotated bounding boxes. It is impractical to collect enough labeled data since the real world is overwhelmed with an enormous amount of categories. In this case, these methods are challenged by the task of zero-shot object detection (ZSD), which aims to simultaneously classify and localize new classes in the absence of any training instances.

Zero-shot object detection can be addressed in two spaces: the semantic embedding space and the visual feature space. (1) Existing methods (Bansal et al. 2018; Rahman et al. 2018; Li et al. 2019b) generally map the visual features from predicted bounding boxes to the semantic embedding space. During the inference stage, a class label is predicted by finding the nearest class based on the similarity scores with all




Figure 1: Illustration of our feature generating process. Visual features are first synthesized conditioned on the semantic embeddings. Then we further add the IoU variance to the features.

class embeddings. However, mapping high-dimensional visual feature to low-dimensional semantic space tends to cause the hubness problem due to the heterogeneity gap between these two spaces (Zhang and Peng 2018). (2) Directly classifying an object in the visual feature space can address the hubness problem. A number of zero-shot classification methods (Xian et al. 2018; Verma et al. 2018; Li et al. 2019a; Huang et al. 2019) have proved the effectiveness of the solution in the visual space. However, visual features contain not only the intra-class variance but also the IoU variance, which is a critical cue for object detection.

To address these problems, we propose a Generative Transfer Network (GTNet) for ZSD. Specifically, we introduce a generative model to synthesize the visual features to solve the hubness problem. Meanwhile, with the consideration of the IoU variance, we design an IoU-Aware Generative Adversarial Network (IoUGAN) to generate visual features with both the intra-class variance and the IoU variance.

The proposed GTNet consists of an Object Detection Module and a Knowledge Transfer Module. More specifically, the Object Detection Module contains a feature extractor, a bounding box regressor, and a seen category classifier. In addition, the feature extractor is used to extract features of Region-of-Interest (RoI) from an image. The Knowledge Transfer Module consists of a feature synthesizer and an unseen category classifier. The feature synthesizer is utilized to generate visual features for training the unseen category classifier. The trained unseen category classifier can be integrated with the feature extractor and the bounding box regressor to

*Corresponding author.

achieve ZSD.

As a detection-specific synthesizer, IoUGAN consists of three unit models: Class Feature Generating Unit (CFU), Foreground Feature Generating Unit (FFU), and Background Feature Generating Unit (BFU). Each unit contains a generator and a discriminator. The feature generating process of IoUGAN is shown as Figure 1. Specifically, CFU focuses on synthesizing features for each unseen class with intra-class variance conditioned on the class semantic embeddings. FFU aims to add the IoU variance to the CFU results, generating the foreground features. Additionally, in order to reduce the confusion between background and unseen categories, BFU synthesizes the class-specific background features conditioned on the CFU results. We conduct extensive experiments on three public datasets, and those experimental results show that the proposed approach performs favorably against the state-of-the-art ZSD algorithms.

The main contributions of this work can be summarized as follows:

- We propose a novel deep architecture GTNet for the ZSD problem. GTNet utilizes a feature synthesizer to generate unseen class features. It is only required to train a new classification layer for a pre-trained detector using the synthesized features. To the best of our knowledge, we are the first to propose a generative approach for ZSD.
- We propose a novel conditional generative model IoUGAN to synthesize unseen class features with both the intra-class variance and the IoU variance. Moreover, IoUGAN can be integrated into GTNet as the feature synthesizer.
- We conduct extensive experiments and component studies to demonstrate the superiority of the proposed approach for ZSD.

2. Related Work

Fully-Supervised Object Detection: In the past years, object detection is driven by deep convolutional neural networks (CNN). The most popular models can be categorized into single-stage networks like SSD (Liu et al. 2016), YOLO (Redmon and Farhadi 2017) and double-stage networks like Faster R-CNN (Ren et al. 2015), R-FCN (Dai et al. 2016). Specifically, single-stage networks like SSD (Liu et al. 2016), YOLO (Redmon and Farhadi 2017) take classification and bounding box regression in one step. While double-stage networks like Faster R-CNN (Ren et al. 2015), R-FCN (Dai et al. 2016) predict offsets of predefined anchor boxes in the first stage by Region Proposal Network (RPN) and then classify and finetune the bounding box prediction in the second stage. These object detection models can only detect categories that have appeared in the training dataset. However, they can not be directly applied to predict classes which are unseen during training.

Zero-Shot Classification: Most existing zero-shot classification methods project visual features to the semantic space, such as (Frome et al. 2013; Xian et al. 2016; Akata et al. 2015; Kodirov, Xiang, and Gong 2017). Instead of learning a visual-semantic mapping, some previous works (Shigeto et al. 2015; Dinu, Lazaridou, and Baroni 2014; Zhang, Xiang, and Gong 2017) propose to learn a semantic-visual mapping. Additionally, there are also some works to

learn an intermediate space, which is shared by the visual features and the semantic embeddings (Changpinyo et al. 2016; Zhang and Saligrama 2016). In contrast, some conditional generative models (Xian et al. 2018; Verma et al. 2018; Li et al. 2019a; Huang et al. 2019) transform the zero-shot classification problem to a general fully-supervised problem by leveraging GANs to synthesizing visual features of unseen categories. We also employ a generative model to generate pseudo features for converting the ZSD into a fully supervised object detection problem. Different the cGANs used in zero-shot classification, IoUGAN can add the IoU variance to the synthesized features.

Zero-Shot Object Detection: There are five contemporary studies on ZSD (Bansal et al. 2018; Rahman et al. 2018; Demirel, Cinbis, and Ikizler-Cinbis 2018; Li et al. 2019b; Rahman, Khan, and Barnes 2018). Specifically, Bansal et al. proposes a background-aware model to reduce the confusion between background and unseen classes. Rahman et al. propose a classification loss composing of a max-margin loss and a meta-class clustering loss. Demirel, Cinbis, and Ikizler-Cinbis uses a convex combination of embeddings to solve the ZSD problem. Li et al. detects unseen classes by exploring their natural language description. Rahman, Khan, and Barnes proposes a polarity loss function for better aligning visual features and semantic embeddings. All of these methods focus on mapping the visual features from predicted bounding boxes to the semantic embedding space. In contrast, we propose a generative approach to tackle the ZSD problem.

3. Generative Approach for Zero-Shot Object Detection

In this section, we first present GTNet, which consists of two modules: an Object Detection Module and a Knowledge Transfer Module. Then we design a novel conditional generative model (*i.e.* IoUGAN) as the feature synthesizer, which is embedded in the Knowledge Transfer Module.

3.1. Generative Transfer Network

Network Overview: As shown in Figure 2, GTNet consists of an Object Detection Module and a Knowledge Transfer Module. Specifically, the Object Detection Module is composed of a seen categories classifier, a bounding box regressor, and a feature extractor. The bounding box regressor is shared among all categories instead of being specific for each category, in order to reuse the regression parameters for detecting unseen classes. The feature extractor is used to extract RoI features from an image. Specifically, the feature extractor is in the fashion of Faster R-CNN (Ren et al. 2015), which has superior performance among competitive end-to-end detection models. The feature synthesizer (*i.e.* IoUGAN) is a conditional generative model, which can learn to generate visual features conditioned on the class semantic embeddings. IoUGAN will be detailed in Section 3.2.

Building Process of the Zero-Shot Detector: The Object Detection Module is pre-trained in the large-scale seen class dataset. Then the feature synthesizer can be trained using the

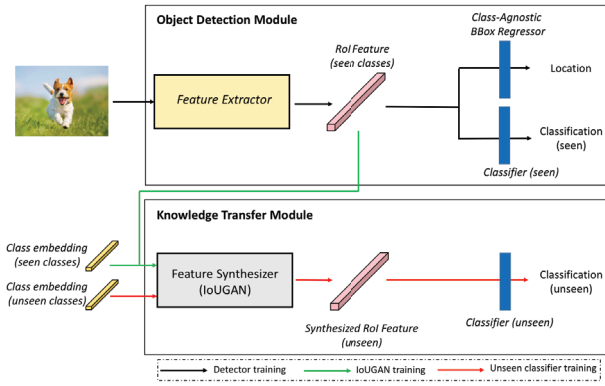


Figure 2: Illustration of the Generative Transfer Network (GTNet). It consists of an Object Detection Module and a Knowledge Transfer Module. The Object Detection Module contains a bounding box regressor, a seen category classifier, and a feature extractor. The Knowledge Transfer Module is composed of a feature synthesizer (*i.e.* IoUGAN) and an unseen category classifier. The synthesizer can be trained using the features from the feature extractor and the class semantic embeddings. A trained synthesizer can sample the unseen class features that can be applied to train an unseen category classifier. The trained unseen classifier can further be integrated with the feature extractor and the regressor to achieve zero-shot object detection.

real features from the feature extractor and the corresponding class embeddings. The trained synthesizer can generate the visual features of unseen classes by inputting the corresponding class semantic embeddings. To the end, we use the generated features to train an unseen categories classifier. In the inference stage, we integrate the trained unseen category classifier with the feature extractor and the regressor to detect unseen classes.

3.2. IoU-Aware Generative Adversarial Network

To generate features of different classes with the IoU variance, we propose a simple yet effective IoUGAN, which mainly is composed of three unit models (*i.e.* CFU, FFU, and BFU), as shown in Figure 3. Each unit contains a generator and a discriminator. In addition, we use WGAN (Martin Arjovsky and Bottou 2017) as the basic component for the proposed generative model due to its more stable training.

- **CFU**: it focuses on generating features with the intra-class variance conditioned on the given semantic embeddings and the noise vectors.
- **FFU**: it adds the IoU variance to the results of CFU and outputs the class-specific foreground features.
- **BFU**: it takes the results of CFU as the inputs and outputs the class-specific background features.

Definitions and Notations: We use the feature extractor to sample the visual features of seen categories. The training set for IoUGAN is denoted as $S = \{(v_i^c, \{v_{ij}^f, v_{ij}^b\}_{j=1}^{n^s}, y_i, e(y)_i)\}_{i=1}^{N^s}$, where $v_i^c \in \mathcal{V}^c$ denotes

the feature extracted from a ground-true box, $v_{ij}^f \in \mathcal{V}^f$ is the class-specific foreground feature from a positive bounding box whose IoU with corresponding the ground-truth box is larger than a threshold t_f , $v_{ij}^b \in \mathcal{V}^b$ is the class-specific background feature from a negative bounding box whose IoU with the corresponding ground-truth box is smaller than a threshold t_b , $y \in \mathcal{Y}^s$ denotes the class label in $\mathcal{Y}^s = \{y_1, \dots, y_s\}$ with S seen classes and $e(y) \in \mathcal{E}$ denotes the category-specific semantic embedding. Note that the positive boxes and the negative boxes, from which v_{ij}^f and v_{ij}^b are extracted, are predicted by RPN. During training, we randomly pick v_{ij}^f and v_{ij}^b among their n^s samples with the corresponding v_i^c to IoUGAN. Additionally, we have the semantic embeddings of unseen classes $\mathcal{U} = \{(u, e(u)) | u \in \mathcal{Y}^u, e(u) \in \mathcal{E}\}$, where u is a class from a disjoint label set $\mathcal{Y}^u = \{u_1, \dots, u_u\}$ of U labels and $e(u)$ is the semantic embedding of a corresponding unseen class.

CFU: Instead of directly generating features of different class with the IoU variance, we simplify the task to first generate features with the intra-class variance using our CFU. We use the feature v^c extracted from the ground-truth bounding box as the real feature to guide the generator to capture the overall characteristic of an object. Given the train data \mathcal{S} of seen classes, we aim to learn a conditional generator $G^c : \mathcal{Z} \times \mathcal{E} \rightarrow \mathcal{V}^c$ in CFU, which takes random Gaussian noise $z \in \mathcal{Z} \subset \mathbb{R}^{d_z}$ and class embedding $e(y) \in \mathcal{E}$ as its inputs, and outputs a visual feature $\tilde{v}^c \in \mathcal{V}^c$ of class y . The discriminator $D^c : \mathcal{V}^c \times \mathcal{E} \rightarrow \mathbb{R}$ is a multi-layer perceptron which outputs a real value. G^c tries to minimize it, while D^c seeks to maximize the loss. Once the generator G^c learns to synthesize visual features of seen categories, *i.e.* v^c , conditioned on the seen class embedding $e(y) \in \mathcal{Y}^s$, it can also generate \tilde{v}^c of any unseen class u via its class embedding $e(u)$. The conditional WGAN loss in CFU is given by

$$\mathcal{L}_{WGAN}^c = \mathbb{E}[D^c(v^c, e(y))] - \mathbb{E}[D^c(G^c(z, e(y)), e(y))] - \alpha^c \mathbb{E}[(\|\nabla_{\hat{v}^c} D^c(\hat{v}^c, e(y))\|_2 - 1)^2] \quad (1)$$

where \hat{v}^c is the convex combination of v^c and \tilde{v}^c : $\hat{v}^c = \eta^c v^c + (1 - \eta^c) \tilde{v}^c$ with $\eta^c \sim U(0, 1)$, α^c is the penalty coefficient and \mathbb{E} is the expectation. The Wasserstein distance is approximated by the first two terms, and the third term constrains the gradient of the discriminator to have unit norm along with the convex combination of real (*i.e.* v^c) and generated (*i.e.* \tilde{v}^c) pairs.

FFU: The bounding boxes predicted by RPN always can not overlap entirely with the ground-truth bounding boxes. Even a bounding box is a positive one that has large IoU with a ground-truth box, the feature extracted from the box still lacks some information compared to that extracted from a ground-true box. The detector should be robust to the information loss of a foreground feature. In this case, we use FFU to randomly add the IoU variance to the features outputted by CFU. Therefore, the foreground feature v^f is used as the real feature to train FFU.

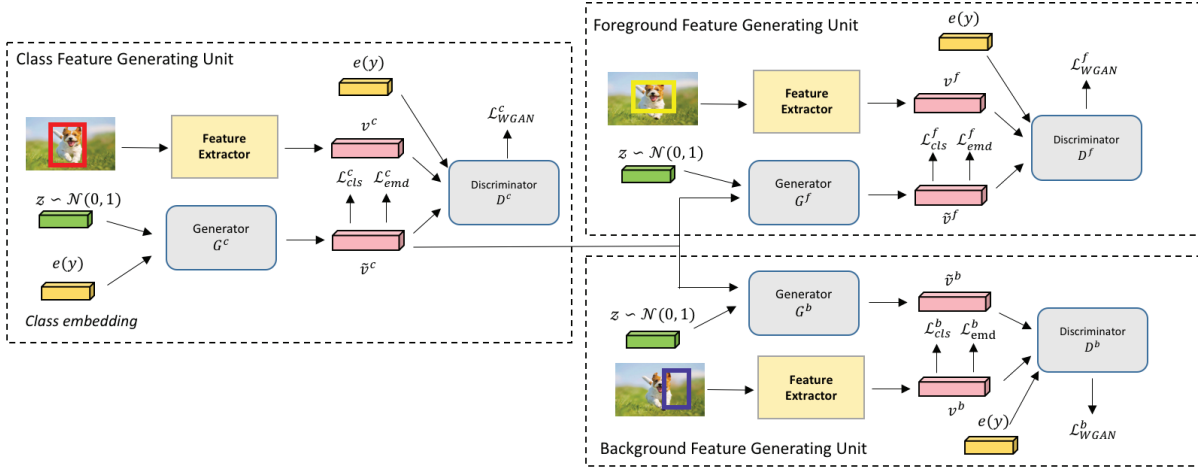


Figure 3: Illustration of the IoU-Aware Generative Adversarial Network (IoUGAN). The Class Feature Generating Unit (CFU) takes the class embeddings and the random noise vectors as input and outputs the features with the intra-class variance. Then the Foreground Feature Generating Unit (FFU) and the Background Feature Generating Unit (BFU) add the IoU variance to the results of CFU and output the class-specific foreground and background features, respectively.

With the output feature from CFU: $\tilde{v}^c = G^c(z, e(y))$ and Gaussian latent variable z , the conditional WGAN loss of FFU is given by

$$\mathcal{L}_{WGAN}^f = \mathbb{E}[D^f(v^f, e(y))] - \mathbb{E}[D^f(G^f(z, \tilde{v}^c), e(y))] - \alpha^f \mathbb{E}[(\|\nabla_{\hat{v}^f} D^f(\hat{v}^f, e(y))\|_2 - 1)^2] \quad (2)$$

where \hat{v}^f is a convex combination of v^f and \tilde{v}^f : $\hat{v}^f = \eta^f v^f + (1 - \eta^f) \tilde{v}^f$ with $\eta^f \sim U(0, 1)$, α^f is the penalty coefficient. Different from the generator in CFU, the class semantic embedding is not used as an input to the generator of FFU with the assumption that the semantic information has already been preserved by \tilde{v}^c .

BFU: Confusion between background and unseen classes limits the performance of a zero-shot detector (Rahman, Khan, and Barnes 2018; Bansal et al. 2018). To enhance the discrimination of detector between background and unseen classes, we use BFU to generate the class-specific background features. We use the background feature v^b as the real feature to train BFU.

With the output feature from CFU: $\tilde{v}^c = G^c(z, e(y))$ and Gaussian latent variable z , the conditional WGAN loss of BFU is given by

$$\mathcal{L}_{WGAN}^b = \mathbb{E}[D_b(v^b, e(y))] - \mathbb{E}[D_b(G_b(z, \tilde{v}^c), e(y))] - \alpha^b \mathbb{E}[(\|\nabla_{\hat{v}^b} D_b(\hat{v}^b, e(y))\|_2 - 1)^2] \quad (3)$$

where \hat{v}^b is a convex combination of v^b and \tilde{v}^b : $\hat{v}^b = \eta^b v^b + (1 - \eta^b) \tilde{v}^b$ with $\eta^b \sim U(0, 1)$, α^b is the penalty coefficient.

Overall Objective Function: Additionally, the synthesized features should be well suited for training a discriminative classification layer of a detector. Similar to (Xian et al. 2018), a discriminative classification layer trained on the

input data is used to classify the synthesized features \tilde{v} so that we can minimize the classification loss over the generated features. Note that for simplicity, we use the \tilde{v} to denote all the generated features (*i.e.* \tilde{v}^c , \tilde{v}^f and \tilde{v}^b). To this end, We use the negative log likelihood,

$$\mathcal{L}_{cls} = -\mathbb{E}_{\tilde{v} \sim p_{\tilde{v}}} [\log P(y|\tilde{v}; \theta)], \quad (4)$$

where y denotes the class label of \tilde{v} , $P(y|\tilde{v}; \theta)$ is the predicted probability of \tilde{v} belonging to its true class label y . In addition, we further use the \mathcal{L}_{cls}^c , \mathcal{L}_{cls}^f , \mathcal{L}_{cls}^b to denote the classification loss functions for the three units, respectively. The conditional probability is computed by a single-layer FC network parameterized by θ , which is pretrained on the real features of seen classes.

Moreover, we expect generated features of a class y_i to be close to the real features of the same class and far from the features of other classes y_j (for $j \neq i$). We generate matched (same classes) and unmatched (different classes) pairs by pairing the real and generated features in a batch. To the end, the distance between the matched and unmatched features can be minimized and maximized, respectively, by the cosine embedding loss, which is given by,

$$\mathcal{L}_{emb} = \mathbb{E}_m [1 - \cos(v, \tilde{v})] + \mathbb{E}_{um} [\max(0, \cos(v, \tilde{v}))] \quad (5)$$

where \mathbb{E}_m and \mathbb{E}_{um} are the expectations over the matched (m) and unmatched (um) pair distributions, respectively. Similarly, we use v to denote all the real features (*i.e.* v^c , v^f and v^b). We further use \mathcal{L}_{emb}^c , \mathcal{L}_{emb}^f , \mathcal{L}_{emb}^b to denote the embedding loss functions for the three units, respectively. While the other losses (\mathcal{L}_{WGAN} and \mathcal{L}_{cls}) focus on the similarity between the generated features and the real ones of the same classes, the embedding loss \mathcal{L}_{emb} also emphasizes the dissimilarity between the generated features and the other class ones. The final objective for training CFU, FFU and BFU, using $\beta_1, \beta_2, \beta_3, \gamma_1, \gamma_2$ and γ_3 as hyper-parameters for

weighting the respective losses, is given by

$$\min_{G^c} \max_{D^c} \mathcal{L}_{WGAN}^c + \beta_1 \mathcal{L}_{cls}^c + \gamma_1 \mathcal{L}_{emb}^c, \quad (6)$$

$$\min_{G^f} \max_{D^f} \mathcal{L}_{WGAN}^f + \beta_2 \mathcal{L}_{cls}^f + \gamma_2 \mathcal{L}_{emb}^f, \quad (7)$$

$$\min_{G^b} \max_{D^b} \mathcal{L}_{WGAN}^b + \beta_3 \mathcal{L}_{cls}^b + \gamma_3 \mathcal{L}_{emb}^b. \quad (8)$$

4. Experimental Results

In this section, we first describe the datasets and settings for evaluation. Then we compare our approach with the state-of-the-arts and conduct the ablation studies. Finally, we show the qualitative results of our method.

4.1. Dataset Description

Our proposed framework is evaluated on three challenging datasets.

ILSVRC-2017 detection dataset (Russakovsky et al. 2015) contains 200 object categories. The categories were carefully selected with variations of object scale, level of image clutterness, average number of object instance, *etc.* Following (Rahman et al. 2018), we select 23 categories as unseen classes and the rest are as seen classes.

MSCOCO (Lin et al. 2014) was introduced for object detection and semantic segmentation tasks, with each image containing multiple object instances with different factors such as clutter, views and many others. Following (Bansal et al. 2018), training images are selected from the 2014 training set and testing images are randomly sampled from the validation set. For seen/unseen category split, we use 48 seen classes for training and 17 unseen classes for testing.

VisualGenome (VG) (Krishna et al. 2017) was collected with a focus on visual relationship understanding. This dataset constitutes of 500 object categories with multiple bounding boxes per image. Following (Bansal et al. 2018), images from part-1 of the dataset are used for training and testing images are randomly sampled from part-2. For seen/unseen category split, we use 478 classes for training and 130 classes for testing.

Train/Test Split: In a zero-shot learning task, no visual instances of unseen classes are allowed during training. In terms of the ILSVRC-2017 detection dataset, following (Rahman et al. 2018), images that contain any unseen classes are removed from the training set. The test set is composed of the rest of images from ILSVRC training dataset and images from validation dataset that have at least one unseen class bounding box. For MSCOCO, following (Bansal et al. 2018), we also remove all the images from the training dataset which contain any unseen class bounding boxes. For VG, however, due to a large number of test classes and dense labeling, most images will be eliminated from the training set. Therefore, following (Bansal et al. 2018), we do not conduct the same procedure for VG.

Evaluation Metric: For ILSVRC-2017 detection dataset, we use mean average precision (mAP) to evaluate the performance. Following (Bansal et al. 2018), for MSCOCO and VG, Recall@100 is used as the evaluation metric which defined as the recall when only the top 100 proposals (sorted by

prediction score) are selected from an image. This is because, for large-scale crowd-annotated datasets such as MSCOCO and VG, it is difficult to annotate all bounding boxes for all instances of a category. mAP is sensitive to missing annotations and it will count such detections as false positive.

Implementation Details: The feature extractor is in the fashion of Faster R-CNN (Ren et al. 2015). In addition, we use Resnet-101 as the backbone of the feature extractor. The seen category classifier and the class-agnostic regressor are both single Fully-Connected (FC) layer. On the other hand, for all three units (*i.e.* CFU, FFU, and BFU) in IoUGAN, the generators G are three-layer FC networks with an output layer dimension equal to the size of the feature. The hidden layers are of size 4096. The discriminators D are a two-layer FC network with the output size equal to 1 and a hidden size equal to 4096. The new classification layer is a single-layer FC with an input size equal to the feature size and the output size equal to the number of unseen classes plus one (*i.e.* background class). We use $\beta_1 = \beta_2 = \beta_3 = 0.01$ and $\gamma_1 = \gamma_2 = \gamma_3 = 0.1$ across all datasets. Following (Bansal et al. 2018), the parameters t_f and t_b are set to 0.5 and 0.2. All the modules are trained using the Adam optimizer with a 10^{-4} learning rate.

Compared Methods: We compare the effect of our proposed approach to three recent state-of-the-art methods (Rahman et al. 2018; Bansal et al. 2018; Li et al. 2019b). In order to fairly compare with the method (Li et al. 2019b), we use the textual description embedding which is trained by fastText (Edouard Grave and Bojanowski 2017) as the class semantic embedding. The study (Li et al. 2019b) extends other the other two ZSD methods (Rahman et al. 2018; Bansal et al. 2018) using the textual description embedding and report the results. We just copy and paste the results from the paper (Li et al. 2019b) in Table 1 and Table 2.

4.2. Comparisons with the State-of-the-Arts

We compare the proposed approach with the state-of-the-arts on three challenging benchmark datasets. We report mAP for all the compared models on the ILSVRC-2017 detection dataset. As can be seen from Table 1, the proposed approach performs better than the other compared models on most of the categories, achieving 1.9% improvement against ZSDTD (Li et al. 2019b). The performance on the categories (*i.e.* hamster, tiger, scorpion) which have similar concepts are much better than those (*i.e.* ray, harmonica, maraca) which have less similar concepts. Moreover, our model has significant improvement in the syringe category from 3.1% to 30.4%, but the performance on the ray class is 25.5% less than the first place. We note that the syringe has twenty-three similar categories in training dataset while the ray only has five. This phenomenon indicates that our approach can effectively transfer the knowledge to the unseen class domain when we have enough similar categories in seen class domain. While there are few similar classes in the training dataset, it is difficult for our generative model to synthesize discriminative visual features.

For the MSCOCO and VG datasets, We report the results in terms of Recall@100 for all the compared methods in Table 2, with three different IoU overlap thresholds (*i.e.* 0.4, 0.5,

Table 1: Comparisons with the states-of-the-arts on the ILSVRC-2017 dataset. We use mean average precision (mAP) as the evaluation metric and present the results in percentages.

	p.box	syringe	harmonica	maraca	burrito	pineapple	electric-fan	ipod	dishwater	can-opener	plate-track	bench	bow-tie	s.trunks	scorpion	snail	hamster	tiger	ray	train	unicycle	g.ball	h.bar	mAP
SAN (Rahman et al. 2018)	5.9	1.5	0.3	0.2	40.6	2.9	7.7	28.5	13.3	5.1	7.8	5.2	2.6	4.6	68.9	6.3	53.8	77.6	21.9	55.2	21.5	31.2	5.3	20.3
SB (Bansal et al. 2018)	6.8	1.8	0.8	0.5	43.7	3.8	8.3	30.9	15.2	6.3	8.4	6.8	3.7	6.1	71.2	7.2	58.4	79.4	23.2	58.3	23.9	34.8	6.5	22.0
DSES (Bansal et al. 2018)	7.4	2.3	1.1	0.6	46.2	4.3	8.7	32.7	14.6	6.9	9.1	7.4	4.9	6.9	73.4	7.8	56.8	80.8	24.5	59.9	25.4	33.1	7.6	22.7
LAB (Bansal et al. 2018)	6.5	1.6	0.7	0.5	44.1	3.6	8.2	30.1	14.9	6.4	8.8	6.4	4.1	4.8	69.9	6.9	57.1	80.2	23.6	58.2	25.1	35.6	7.2	21.9
ZSDTD (Li et al. 2019b)	7.8	3.1	1.9	1.1	49.4	4.0	9.4	35.2	14.2	8.1	10.6	9.0	5.5	8.1	73.5	8.6	57.9	82.3	26.9	61.5	24.9	38.2	8.9	24.1
GTNet	4.4	30.4	2.3	1.2	51.1	5.2	18.5	40.6	18.0	13.1	4.7	13.7	4.6	19.2	69.7	10.2	74.7	72.7	1.4	65.7	27.1	40.4	9.1	26.0

Table 2: Comparisons with the states-of-the-arts on the MSCOCO and Visual Genome (VG) datasets. We use Recall@100 as the evaluation metric and present the results in percentages.

IoU	MSCOCO			Visual Genome		
	0.4	0.5	0.6	0.4	0.5	0.6
SAN (Rahman et al. 2018)	35.7	26.3	14.5	6.8	5.9	3.1
SB (Bansal et al. 2018)	37.8	28.6	15.4	7.2	5.6	3.4
DSES (Bansal et al. 2018)	42.6	31.2	16.3	8.4	6.3	3.3
LAB (Bansal et al. 2018)	35.2	22.4	12.1	8.6	6.1	3.3
ZSDTD (Li et al. 2019b)	45.5	34.3	18.1	9.7	7.2	4.2
GTNet	47.3	44.6	35.5	14.3	11.3	8.9

0.6). In addition, for convenient discussion, IoU 0.5 is used as an example in this paper unless specified otherwise. From the experimental results, we observe that the proposed model generally performs much better than the other compared baselines. Specifically, the proposed model increases Recall@100 to 44.6% from 34.3% achieved by the method (Li et al. 2019b) on MSCOCO, and increases Recall@100 to 11.3% from 7.2% on VG.

4.3. Ablation Studies

We study the effects of different components in our model and report the experimental results. We still use mAP as the evaluation metric for ILSVRC and Recall@100 for MSCOCO and VG. We design a baseline which uses the traditional visual-semantic embedding approach. Note that in the inference stage, the network architectures of the baseline and our approach are the same except that our approach directly maps the RoI features into label space while the baseline maps the features into semantic space followed by a nearest neighbor search.

Effectiveness of IoUGAN: To study the effect of the IoU information and whether the synthesized background features are useful or not, We first report the results which use the features generated by only CFU with the background features sampling from the detection training dataset. Note that these generated features do not contains the IoU variance. Then we use FFU to add the IoU variance to the features from CFU. In the end, we use the full architecture of IoUGAN with the ability to synthesize the class-specific background features. From the experimental results, as shown in Table 3, we can see that only using CFU, our approach can outperform the baseline by 8.9%, 0.6% and 3.8% on three datasets, respectively. Moreover, adding the IoU variance by FFU will dramatically improve the performance of ZSD from 27.0%, 3.6% and 20.1% to 39.2%, 7.5% and 24.5%, which indicates the great

importance of the IoU variance. The ZSD performance is further improved when using the synthesized background features which make the detector more discriminatory.

Effectiveness of the Loss Functions: To study the effect of the loss functions, we first report results only using the Wasserstein loss. Then we add the classification loss \mathcal{L}_{cls} and the embedding loss \mathcal{L}_{emb} , respectively. To the end, we report the results using all the loss functions. As shown in Table 4, both the classification loss and the embedding loss can further improve the performance.

4.4. Generalized Zero-Shot Object Detection

We also show the results of the Generalized Zero-Shot Object Detection (GZSD) in Table 5. We compare to the baseline method presented in the ablation study and the Polarity Loss (Rahman, Khan, and Barnes 2018). The results in Table 5 shows proposed GTNet performs better than the baseline and the Polarity Loss.

4.5. Qualitative Results

Figure 4 shows the output detection by the proposed approach with the MSCOCO dataset being used as an example here. These examples confirm that the proposed method can detect classes that are unseen during the training process. The failure detection tends to happen when objects are small and have similar classes in the set of unseen classes.

5. Conclusions

In this paper, we propose GTNet, which consists of an Object Detection Module and a Knowledge Transfer Module, to tackle the ZSD problem. GTNet naturally embeds a feature synthesizer in the Knowledge Transfer Module. The feature synthesizer learns to generate unseen class features, which will be further applied to train an unseen category classifier for a pre-trained detector. In order to generate unseen class

Table 3: Experimental results of different components on the three challenging benchmark datasets.

IoU	MSCOCO			Visual Genome			ILSVRC
	0.4	0.5	0.6	0.4	0.5	0.6	0.5
Baseline	27.4	18.1	14.5	5.2	3.0	2.3	16.3
CFU	29.2	27.0	24.7	4.5	3.6	2.7	20.1
CFU + FFU	43.7	39.2	33.0	9.3	7.5	5.8	24.5
CFU + FFU + BFU	46.2	43.4	34.9	13.6	10.6	7.9	25.5

Table 4: Experimental results of different loss functions on the three challenging benchmark datasets.

IoU	MSCOCO			Visual Genome			ILSVRC
	0.4	0.5	0.6	0.4	0.5	0.6	0.5
IoUGAN	46.2	43.4	34.9	13.6	10.6	7.9	25.5
IoUGAN (+ \mathcal{L}_{cls})	47.1	44.3	35.3	14.0	10.9	8.4	25.7
IoUGAN (+ \mathcal{L}_{emb})	47.5	44.9	35.2	13.7	10.8	8.3	25.8
IoUGAN (+ \mathcal{L}_{cls} + \mathcal{L}_{emb})	47.3	44.6	35.5	14.3	11.3	8.9	26.0



Figure 4: Some examples of zero-shot object detection using the proposed approach.

Table 5: Experimental results of Generalized Zero-Shot Object Detection on the MSCOCO Dataset.

	MSCOCO	
	Seen	Unseen
Baseline	20.3	8.9
Polarity Loss	38.2	26.3
GTNet	42.5	30.4

features with the IoU variance, we design a novel generative model IoUGAN as the feature synthesizer. IoUGAN is composed of three unit models (*i.e.* CFU, FFU, and BFU). The CFU aims to generate features with the intra-class variance for each unseen class. FFU and BFU add the IoU variance to the results of CFU and output the class-specific foreground

and background features, respectively. We conducted extensive experiments in order to demonstrate the superiority of our model and investigated the effectiveness of different components. In the future, we will investigate how to generate more discriminative features for the unseen classes that have few similar categories in the training dataset.

Acknowledgment

This work was supported by National Natural Science Foundation of China (No.61901184, No.61771329 and No. 61876210).

References

Akata, Z.; Reed, S.; Walter, D.; Lee, H.; and Schiele, B. 2015. Evaluation of output embeddings for fine-grained

- image classification. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Bansal, A.; Sikka, K.; Sharma, G.; Chellappa, R.; and Divakaran, A. 2018. Zero-shot object detection. In *European Conference on Computer Vision*.
- Changpinyo, S.; Chao, W.-L.; Gong, B.; and Sha, F. 2016. Synthesized classifiers for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Dai, J.; Li, Y.; He, K.; and Sun, J. 2016. R-fcn: Object detection via region-based fully convolutional networks. In *arXiv:1605.06409*.
- Demirel, B.; Cinbis, R. G.; and Ikizler-Cinbis, N. 2018. Zero-shot object detection by hybrid region embedding. In *British Machine Vision Conference*.
- Dinu, G.; Lazaridou, A.; and Baroni, M. 2014. Improving zero-shot learning by mitigating the hubness problem. In *arXiv:1412.6568*.
- Edouard Grave, Tomas Mikolov, A. J., and Bojanowski, P. 2017. Bag of tricks for efficient text classification. In *European Chapter of the Association for Computational Linguistics*.
- Frome, A.; Corrado, G. S.; Shlens, J.; Bengio, S.; Dean, J.; Ranzato, M.; and Mikolov, T. 2013. Devise: A deep visual-semantic embedding model. In *Neural Information Processing Systems*.
- Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Girshick, R. 2015. Fast r-cnn. In *IEEE International Conference on Computer Vision*.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask R-CNN. In *IEEE International Conference on Computer Vision*.
- Huang, H.; Wang, C.; Yu, P. S.; and Wang, C.-D. 2019. Generative dual adversarial network for generalized zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Kodirov, E.; Xiang, T.; and Gong, S. 2017. Semantic autoencoder for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; and et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*.
- Li, J.; Jin, M.; Lu, K.; Ding, Z.; Zhu, L.; and Huang, Z. 2019a. Leveraging the invariant side of generative zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Li, Z.; Yao, L.; Zhang, X.; Wang, X.; Kanhere, S.; and Zhang, H. 2019b. Zero-shot object detection with textual description. In *AAAI Conference on Artificial Intelligence*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *European Conference on Computer Vision*.
- Martin Arjovsky, S. C., and Bottou, L. 2017. Wasserstein generative adversarial networks. In *ICML*.
- Rahman, S.; Khan, S.; Porikli, F.; and Porikli, F. 2018. Zero-shot object detection: Learning to simultaneously recognize and localize novel concepts. In *Asian Conference on Computer Vision*.
- Rahman, S.; Khan, S.; and Barnes, N. 2018. Polarity loss for zero-shot object detection. In *arXiv:1811.08982*.
- Redmon, J., and Farhadi, A. 2017. Yolo9000: Better, faster, stronger. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Neural Information Processing Systems*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; and et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*.
- Shigeto, Y.; Suzuki, I.; Hara, K.; Shimbo, M.; and Matsumoto, Y. 2015. Ridge regression, hubness, and zero-shot learning. *Lecture Notes in Computer Science*.
- Verma, V. K.; Arora, G.; Mishra, A.; and Rai, P. 2018. Generalized zero-shot learning via synthesized examples. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Xian, Y.; Akata, Z.; Sharma, G.; Nguyen, Q.; Hein, M.; and Schiele, B. 2016. Latent embeddings for zero-shot classification. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Xian, Y.; Lorenz, T.; Schiele, B.; and Akata, Z. 2018. Feature generating networks for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Zhang, C., and Peng, Y. 2018. Visual data synthesis via gan for zero-shot video classification. In *International Joint Conference on Artificial Intelligence*.
- Zhang, Z., and Saligrama, V. 2016. Zero-shot learning via joint latent similarity embedding. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Zhang, L.; Xiang, T.; and Gong, S. 2017. Learning a deep embedding model for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*.