

Visual Compositional Learning for Human-Object Interaction Detection

Zhi Hou^{1,2}, Xiaojiang Peng², Yu Qiao² *, and Dacheng Tao¹

¹ UBTECH Sydney AI Centre, School of Computer Science, Faculty of Engineering,
The University of Sydney, Darlington, NSW 2008, Australia
zhou9878@uni.sydney.edu.au, dachengtao@sydney.edu.au

² Shenzhen Key Lab of Computer Vision and Pattern Recognition, Shenzhen
Institutes of Advanced Technology, Chinese Academy of Sciences
{xj.peng, yu.qiao}@siat.ac.cn

Abstract. Human-Object interaction (HOI) detection aims to localize and infer relationships between human and objects in an image. It is challenging because an enormous number of possible combinations of objects and verbs types forms a long-tail distribution. We devise a deep Visual Compositional Learning (VCL) framework, which is a simple yet efficient framework to effectively address this problem. VCL first decomposes an HOI representation into object and verb specific features, and then composes new interaction samples in the feature space via stitching the decomposed features. The integration of decomposition and composition enables VCL to share object and verb features among different HOI samples and images, and to generate new interaction samples and new types of HOI, and thus largely alleviates the long-tail distribution problem and benefits low-shot or zero-shot HOI detection. Extensive experiments demonstrate that the proposed VCL can effectively improve the generalization of HOI detection on HICO-DET and V-COCO and outperforms the recent state-of-the-art methods on HICO-DET. Code is available at <https://github.com/zhihou7/VCL>.

Keywords: Human-Object Interaction, Compositional Learning

1 Introduction

Human-Object interaction (HOI) detection [7, 9, 20, 27] aims to localize and infer relationships (verb-object pairs) between human and objects in images. The main challenges of HOI come from the complexity of interaction and the long-tail distribution of possible verb-object combinations [7, 29, 36]. In practice, a few types of interactions dominate the HOI samples, while a large number of interactions are rare which are always difficult to obtain sufficient training samples.

The visual scenes are composed of basic elements, such as objects, parts, and other semantic regions. It is well-acknowledged that humans perceive world in

* corresponding author

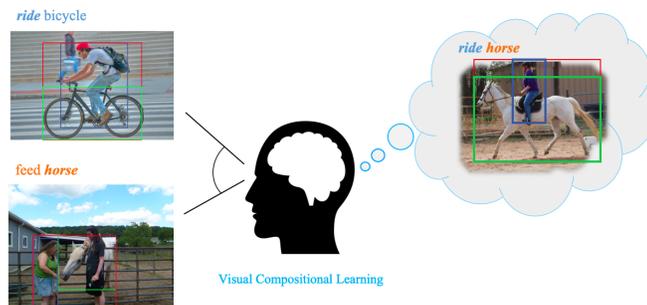


Fig. 1. An illustration of Visual Compositional Learning (VCL). VCL constructs the new concept of $\langle ride, horse \rangle$ from $\langle feed, horse \rangle$ and $\langle ride, bicycle \rangle$ via visual compositional learning

a compositional way in which visual scenes are treated as a layout of distinct semantic objects [16, 30]. We can understand HOIs by decomposing them into objects and human interaction (verb) types. This decomposition helps to solve the rare Human-Object Interactions with long-tailed distribution. For example, in HICO-DET dataset [7], $\langle hug, suitcase \rangle$ is a rare case with only one example, while we have more than 1000 HOI samples including object “suitcase”, and 500 samples including verb “hug”. Obviously, object representations can be shared among different HOIs. And samples with the same verb usually exhibit similar human pose or action characteristics [36]. By compositing the concepts of “suitcase” and “hug” learned from these large number samples, one can handle the rare case $\langle hug, suitcase \rangle$. This inspires to reduce the complexity of HOI detection and handle unseen/rare categories via learning compositional components, *i.e.* human verb and objects from visual scenes. Note this idea is near but different from disentangling representation learning [5] (e.g. factors method [29] in HOI detection) which aims to separate the distinct and informative factors from input examples. Similar to disentangling, compositional learning of HOI also includes the decomposing step. Unlike disentangling, compositional learning further composes novel HOI examples with decomposed factors, which is helpful to address low-shot and zero-shot HOI detection.

Inspired by the above analysis, this paper proposes a deep Visual Compositional Learning (VCL) frame work for Human-Object Interaction Detection, which performs compositional learning on visual verb and object representations. VCL simultaneously encourages shared verb and object representation across images and HOI types. As illustrated in Figure 1, with the semantic features of ‘horse’ and ‘ride’ in the images of $\langle feed, horse \rangle$ and $\langle ride, bicycle \rangle$, one may compose a new interaction feature $\langle ride, horse \rangle$ using off-the-shelf features.

To perform compositional learning for HOI detection, VCL faces three challenges. Firstly, verb features are usually highly correlated with human and object features. It is non-trivial to decouple the verb representations from those of hu-

man and objects in scenes. Unlike prior works [9, 12], which extract the verb representations from human boxes, we build verb representation from the union box of human and object. Meanwhile, we share weights between verb and human stream in the multi-branches structure to emphasize the verb representation on human box region. Our verb representation yields more discriminative cues for the final detection task. Secondly, the number of verbs and objects within a single image is limited for composition. We present a novel feature compositional learning approach by composing HOI samples with verb and object features from different images and different HOI types. In this way, our VCL encourages the model to learn the shared and distinctive verb and object representations that are insensitive to variations (*i.e.* the specific images and interactions). Thirdly, HOIs always exhibit long-tail distribution where a large number of categories have very few even zero training samples. VCL can compose new interactions and novel types of HOIs in the feature space (*i.e.* verb + object), e.g., the rare HOI $\langle wash, cat \rangle$ can be drawn from $\langle wash, dog \rangle$ and $\langle feed, cat \rangle$.

Overall, our main contributions can be summarized as follows,

- We creatively devise a deep Visual Compositional Learning framework to compose novel HOI samples from decomposed verbs and objects to relieve low-shot and zero-shot issues in HOI detection. Specifically, we propose to extract verb representations from union box of human and object and compose new HOI samples and new types of HOIs from pairwise images.
- Our VCL outperforms considerably previous state-of-the-art methods on the largest HOI Interaction detection dataset HICO-DET [7], particularly for rare categories and zero-shot.

2 Related Works

2.1 Human-Object Interaction Detection

Human-Object Interaction [7, 8] is essential for deeper scene understanding. Different from **Visual Relationship Detection** [24], **HOI** is a kind of human-centric relation detection. Several large-scale datasets (V-COCO [11] and HICO-DET [7]) were released for the exploration of HOI detection. Chao *et al.* [7] introduced a multi-stream model combining visual features and spatial location features to help tackle this problem. Following the multi-stream structure in [7], Gao *et al.* [9] further exploited an instance centric attention module and Li *et al.* [20] utilized interactiveness to explicitly discriminate non-interactive pairs to improve HOI detection. Recently, Wang *et al.* [32] proposed a contextual attention framework for Human-Object Interaction detection. GPNN [27] and RPNN [38] were introduced to model the relationships with graph neural network among parts or/and objects. Pose-aware Multi-level Feature Network [31] aimed to generate robust predictions on fine-grained human object interaction. Different from the previous works [7, 9, 20, 27, 31, 32, 38] who mainly focus on learning better HOI features, we address the long-tailed and zero-shot issues in HOI detection via Visual Compositional Learning.

2.2 Low-shot and Zero-shot Learning

Our work also ties with low-shot learning [33] within long-tailed distribution [22] and zero-shot learning recognition [35]. Shen *et al.* [29] introduced a factorized model for HOI detection that disentangles reasoning on verbs and objects to tackle the challenge of scaling HOI recognition to the long tail of categories, which is similar to our work. But we design a compositional learning approach to compose HOI examples. Visual-language joint embedding models [26, 36] enforce the representation of the same verb to be similar among different HOIs by the intrinsic semantic regularities. [26] further transferred HOIs embeddings from seen HOIs to unseen HOIs using analogies for zero-shot HOI detection. However, different from [26] who aims to zero-shot learning, VCL targets at Generalized Zero-Shot Learning [34]. In [3], a generic object detector was incorporated to generalize to interactions involving previously unseen objects. Also, Yang *et al.* [37] proposed to alleviate the predicate bias to objects for zero-shot visual relationship detection. Similar to previous approaches [3, 26, 29, 36], we also equally treat the same verb from different HOIs. However, all those works [3, 26, 29, 36] largely ignore the composition of verbs and objects. In contrast, we propose the Visual Compositional Learning framework to relieve the long-tail and zero-shot issues of HOI detection jointly, and we demonstrate the efficiency by massive experiments, especially for rare/unseen data in HOI detection.

2.3 Feature Disentangling and Composing

Disentangled representation learning has attracted increasing attention in various kinds of visual task [5, 6, 14, 15, 23] and the importance of Compositional Learning to build intelligent machines is acknowledged [4, 5, 10, 19]. Higgins *et al.* [15] proposed Symbol-Concept Association Network (SCAN) to learn hierarchical visual concepts. Recently, [6] proposed Multi-Object network (MONet) to decompose scenes by training a Variational Autoencoder together with a recurrent attention network. However, both SCAN [15] and MONet [6] only validate their methods on the virtual datasets or simple scenes.

Besides, Compositional GAN [2] was introduced to generate new images from a pair of objects. Recently, Label-Set Operations network (LaSO) [1] combined features of image pairs to synthesize feature vectors of new label sets according to certain set operations on the label sets of image pairs for multi-label few-shot learning. Both Compositional GAN [2] and LaSO [1], however, compose the features from two whole images and depend on generative network or reconstruct loss. In addition, Kato *et al.* [18] introduced a compositional learning method for HOI classification [8] that utilizes the visual-language joint embedding model to the feature of the whole of image. But [18] did not involve multiple objects detection in the scene. Our visual compositional learning framework differs from them in following aspects: i) it composes interaction features from regions of images, ii) it simultaneously encourages *discriminative* and *shared* verb and object representations.

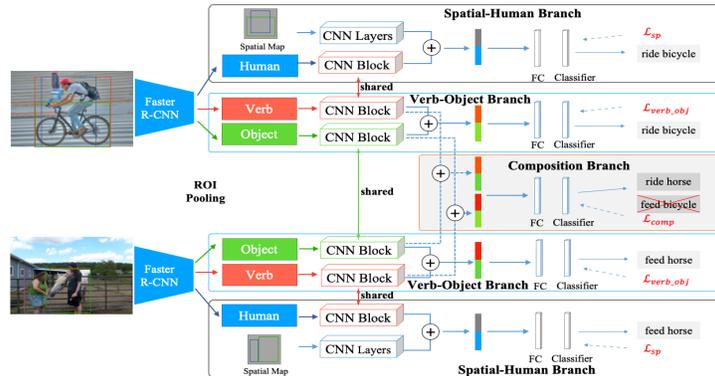


Fig. 2. Overview of the proposed Visual Compositional Learning (VCL) framework. Given two images, we first detect human and objects with Faster-RCNN [28]. Next, with ROI-Pooling and Residual CNN blocks, we extract human features, verb features (*i.e.* the union box of human and object), and object features. Then, these features are fed into the following branches: individual spatial-human branch, verb-object branch and composited branch. Finally, HOI representations from verb-object branch and composited branch are classified by a shared FC-Classifer, while HOI representations from spatial-human branch are classified by an individual FC-Classifer. *Note that all the parameters are shared across images and the newly composited HOI instances can be from a single image if the image includes multiple HOIs*

3 Visual Compositional Learning

In this section, we present our Visual Compositional Learning (VCL) framework for HOI detection. We first provide an overview of VCL and then detail the HOI branches of VCL. Last, we describe how we compose new HOIs and apply VCL to zero-shot detection.

3.1 Overview

To address the long-tail and zero-shot issues of HOI detection, we propose the Visual Compositional Learning (VCL) framework to learn shared object and verb features and compose novel HOI samples with these features. As shown in Figure 2, to perform compositional learning, our VCL takes as input a randomly selected image pair. Then we employ a Faster R-CNN [28] with backbone ResNet-50 [13] to detect human and objects in images. Subsequently, we use ROI-Pooling and Residual CNN blocks to obtain features of human, verbs, and objects individually. Then, to obtain HOI detection, these features together with a human-object spatial map are fed into a spatial-human branch and a verb-object branch. Particularly, composited features are fed into the composited branch for compositional learning. It is worth noting that all the parameters are shared across images.

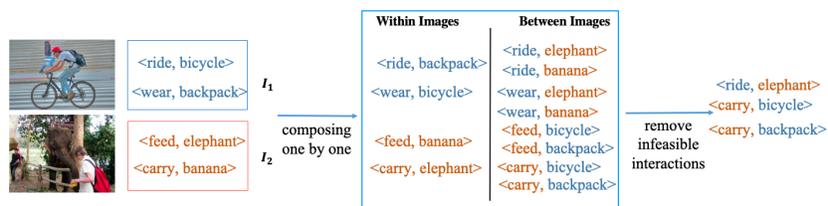


Fig. 3. Illustration of the process of composing new interactions. Given two images I_1 and I_2 , we compose new interaction samples within single image and between them by first considering all possible verb-object pairs and then removing infeasible interactions

3.2 Multi-branch Network

Multi-branch architectures are usually used in previous HOI detection works [7,9,20] where each branch processes a kind of input information. Similarly, VCL includes a spatial-human branch and a verb-object branch. But unlike previous works, VCL has a composition branch which helps the training of HOI detection.

Spatial-Human branch. The spatial-human branch processes human feature and the spatial map of human and object. Following [9], the input of spatial feature is a two-channel 64x64 tensor consisting of a person map and an object map. For the person map, the value of a position will be 1 if it is in the person box otherwise 0. The object map is similar. We concatenate the spatial map with the human feature.

Verb-Object branch. The verb-object branch in each image includes a verb stream and an object stream. Unlike prior works [9,12] which view the human features as verb representations, our newly introduced verb branch extracts *a verb representation from the union box of a human box and an object box*. Meanwhile, similar to [9,12], we share the weights between human stream in Spatial-Human branch and verb stream from the union box in Verb-Object branch. Our verb representation is more discriminative which contains more useful contextual information within the union box of human and object.

Composition branch. We compose new verb-object interaction samples in the feature space from the verbs and objects between and within images and then *these synthesized samples are trained jointly with the interactions annotated in the dataset*. It in turn improves the generalization of HOI detection, particularly for rare and unseen HOIs.

3.3 Composing Interactions

The key idea of our proposed VCL framework is to compose new interaction samples within and between images. This composition process encourages the network to learn shared object and verb features across HOIs by composing massive diverse HOI samples. As shown in Figure 3, the composition process mainly

contains two stages: generating all possible interactions (*i.e.* verb-object pairs) and removing infeasible interactions. Given two images I_1 and I_2 , we compose new interaction samples within single image and between images by first considering all possible verb-object pairs and then removing infeasible interactions in the HOI label space.

Existing HOI labels mainly contain one object and at least one verb, which set the HOI detection as a multi-label problem. To avoid frequently checking verb-object pairs, we design an efficient composing and removing strategy. First, we decouple the HOI label space into a verb-HOI matrix $\mathbf{A}_v \in R^{N_v \times C}$ and an object-HOI matrix $\mathbf{A}_o \in R^{N_o \times C}$, where N_v , N_o , and C denote the number of verbs, objects and HOI categories respectively. \mathbf{A}_v (\mathbf{A}_o) can be viewed as the cooccurrence matrix between verbs (objects) and HOIs. Then, given binary HOI label vectors $\mathbf{y} \in R^{N \times C}$, where N , C denote the number of interactions and HOI categories respectively. we can obtain the object label vector and verb label vector as follows,

$$\mathbf{l}_o = \mathbf{y}\mathbf{A}_o^\top, \mathbf{l}_v = \mathbf{y}\mathbf{A}_v^\top, \quad (1)$$

where $\mathbf{l}_o \in R^{N \times N_o}$ is usually one-hot vectors meaning one object of a HOI example, and $\mathbf{l}_v \in R^{N \times N_v}$ is possibly multi-hot vectors meaning multiple verbs. *e.g.* $\{\text{hold}, \text{sip}\}$, cup). Similarly, we can generate new interactions from arbitrary \mathbf{l}_o and \mathbf{l}_v as follows,

$$\hat{\mathbf{y}} = (\mathbf{l}_o\mathbf{A}_o) \& (\mathbf{l}_v\mathbf{A}_v), \quad (2)$$

where $\&$ denotes the ‘‘and’’ logical operation. The infeasible HOI labels that do not exist in the given label space are all-zero vectors after the logical operation. And then, we can filter out those infeasible HOIs. In implementation, we obtain verbs and objects from two images by ROI pooling and treat them within and between images as same. Therefore, we do not treat two levels of composition differently during composing HOIs.

Zero-Shot Detection. The composition process makes VCL handling zero-shot detection naturally. Specifically, with the above-mentioned operation, we can generate HOI samples for zero-shot (in the given HOI label space) between and within images which may not be annotated in the training set.

3.4 Training and Inference

Training. We train the proposed VCL in an end-to-end manner with Cross Entropy (CE) losses from multiple branches: \mathcal{L}_{sp} from the spatial-human branch for original HOI instances, \mathcal{L}_{verb_obj} from verb-object branch for original HOI instances, and \mathcal{L}_{comp} from verb-object branch for composited HOI instances. Formally, the total training loss is defined as follows,

$$\mathcal{L} = \mathcal{L}_{sp} + \lambda_1\mathcal{L}_{verb_obj} + \lambda_2\mathcal{L}_{comp}, \quad (3)$$

where λ_1 and λ_2 are two hyper-parameters. We employ the composition process (*i.e.* composing new HOI instances) in each minibatch at training stage.

Inference. At test stage, we remove the composition operation and use the spatial-human branch and the verb-object branch to recognize interaction (*i.e.* human-object pair) of an input image. We predict HOI scores in a similar manner to [9]. For each human-object bounding box pair (b_h, b_o) , we predict the score $S_{h,o}^c$ for each category $c \in 1, \dots, C$, where C denotes the total number of possible HOI categories. The score $S_{h,o}^c$ depends on the confidence for the individual object detection scores (s_h and s_o) and the interaction prediction based on verb-object branch $s_{verb.obj}^c$ and spatial-human branch s_{sp}^c . Specifically, Our final HOI score $S_{h,o}^c$ for the human-object bounding box pair (b_h, b_o) is :

$$S_{h,o}^c = s_h \cdot s_o \cdot s_{verb.obj}^c \cdot s_{sp}^c \quad (4)$$

4 Experiment

In this section, we first introduce datasets and metrics and then provide the details of the implementation of our method. Next, we report our experimental results compared with state-of-the-art approaches and zero-shot results. Finally, we conduct ablation studies to validate the components in our framework.

4.1 Datasets and Metrics

Datasets. we adopt two HOI datasets HICO-DET [7] and V-COCO [11]. HICO-DET [7] contains 47,776 images (38,118 in train set and 9,658 in test set), 600 HOI categories constructed by 80 object categories and 117 verb classes. HICO-DET provides more than 150k annotated human-object pairs. V-COCO [11] provides 10,346 images (2,533 for training, 2,867 for validating and 4,946 for testing) and 16,199 person instances. Each person has annotations for 29 action categories and there are no interaction labels including objects.

Metrics. We follow the settings in [9], *i.e.* a prediction is a true positive only when the human and object bounding boxes both have IoUs larger than 0.5 with reference to ground truth and the HOI classification result is accurate. We use the role mean average precision to measure the performance on V-COCO.

4.2 Implementation Details

For HICO-DET dataset, we utilize the provided interaction labels to decompose and compose interactions labels. For V-COCO dataset, which only provides verb (*i.e.* action) annotations for COCO images [21], we obtain object labels (only 69 classes in V-COCO images) from the COCO annotations. Therefore, we have 69 classes of objects and 29 classes of actions in V-COCO that construct 238 classes of Human-Object pairs to facilitate the detection of 29 classes of actions with VCL. In addition, following the released code of [20], we also apply the same reweighting strategy in HICO-DET and V-COCO (See supplementary materials for comparison), and we keep it for the composited interactions. Besides, to prevent composited interactions from dominating the training of the model, we

randomly select composited interactions in each minibatch to maintain the same number of composited interactions as we do of non-composited interactions.

For a fair comparison, we adopt the object detection results and pre-trained weights provided by authors of [9]. We apply two 1024-d fully-connected layers to classify the interaction feature concatenated by verb and object. We train our network for 1000k iterations on the HICO-DET dataset and 500k iterations on V-COCO dataset with an initial learning rate of 0.01, a weight decay of 0.0005, and a momentum of 0.9. We set λ_1 of 2 and λ_2 of 0.5 for HICO-DET dataset, while for V-COCO dataset, λ_1 is 0.5 and λ_2 is 0.1. In order to compose enough new interactions between images for training, we increase the number of interactions in each minibatch while reducing the number of augmentations for each interaction to keep the batch size unchanged. All experiments are conducted on a single Nvidia GeForce RTX 2080Ti GPU with Tensorflow.

4.3 Results and Comparisons

We compare our proposed Visual Compositional Learning with those state-of-the-art HOI detection approaches [3, 9, 12, 20, 26, 27, 31, 32, 36] on HICO-DET. The HOI detection result is evaluated with mean average precision (mAP) (%). We use the settings: Full (600 HOIs), Rare (138 HOIs), Non-Rare (462 HOIs) in Default mode and Known Object mode on HICO-DET.

Comparisons with state-of-the-art. From Table 1, we can find that we respectively improve the performance of Default and Know modes with Full setting to 19.43% and 22.00% without external knowledge. In comparison with [31] who achieves the best mAP in Rare category on HICO-DET among the previous works, we improve the mAP by 1.97% in Full category and by 0.9% in Rare category. Particularly, we do not use pose information like [31]. We improve dramatically by over **3%** in Rare category compared to other visual methods [9, 12, 20, 27, 38]. Besides, we achieve 1.92% better result than [26] in the rare category although we only use the visual information.

Particularly, [3] incorporates ResNet101 as their backbone and finetune the object detector on HICO-DET. When we the same backbone and finetune the object detector on HICO-DET, we can largely improve the performance to **23.63%**. This also illustrate the great effect of object detector on current two-stage HOI detection method. *in the two-stage HOI detection, we should not only focus on HOI recognition such as [9, 20, 38], but also improve the object detector performance.*

Visualization. In Figure 4, we qualitatively show that our proposed Visual Compositional Learning framework can detect those rare interactions correctly while the baseline model without VCL misclassifies on HICO-DET.

4.4 Generalized Zero-shot HOI Detection

We also evaluate our method in HOI zero-shot detection since our method can naturally be applied to zero-shot detection. Following [29], we select 120 HOIs from HOI label set and make sure that the remaining contains all objects and

Table 1. Comparisons with the state-of-the-art approaches on HICO-DET dataset [7]. Xu *et al.* [36], Peyre *et al.* [26] and Bansal *et al.* [3] utilize language knowledge. We include the results of [3] with the same COCO detector as ours. * means we use the res101 backbone and finetune the object detector on HICO-DET dataset like [3]

Method	Default			Known Object		
	Full	Rare	NonRare	Full	Rare	NonRare
GPNN [27]	13.11	9.34	14.23	-	-	-
iCAN [9]	14.84	10.45	16.15	16.26	11.33	17.73
Li <i>et al.</i> [20]	17.03	13.42	18.11	19.17	15.51	20.26
Wang <i>et al.</i> [32]	16.24	11.16	17.75	17.73	12.78	19.21
Gupta <i>et al.</i> [12]	17.18	12.17	18.68	-	-	-
Zhou <i>et al.</i> [38]	17.35	12.78	18.71	-	-	-
PMFNet [31]	17.46	15.65	18.00	20.34	17.47	21.20
Xu <i>et al.</i> [36]	14.70	13.26	15.13	-	-	-
Peyre <i>et al.</i> [26]	19.40	14.63	20.87	-	-	-
Bansal <i>et al.</i> [3]	16.96	11.73	18.52	-	-	-
Bansal* <i>et al.</i> [3]	21.96	16.43	23.62	-	-	-
Ours (VCL)	19.43	16.55	20.29	22.00	19.09	22.87
Ours* (VCL)	23.63	17.21	25.55	25.98	19.12	28.03

verbs. The annotations on HICO-DET, however, are long-tail distributed. The result of selecting preferentially rare labels is different from that of selecting non-rare labels and we do not know the specific partition of zero shot in [3, 29]. For a fair comparison, we split the HOI zero-shot experiment into two groups: rare first selection and non-rare first selection. Rare first selection means we pick out as many rare labels as possible for unseen classes according to the number of instances for zero-shot (remain 92,705 training instances), while non-rare first selection is that we select as many non-rare labels as possible for unseen classes (remain 25,729 training instances). The unseen HOI detection result is evaluated with mean average precision (mAP) (%). We report our result in the settings: Unseen (120 HOIs), Seen (480 HOIs), Full (600 HOIs) in Default mode and Known Object mode on HICO-DET.

We can find in Table 2, both selection strategies witness a consistent increase in all categories compared to the corresponding baseline. Noticeably, both kinds of selecting strategies witness a surge of performance by **over 4%** than baseline in unseen category with VCL. Particularly, our baseline without VCL still achieves low results (3.30% and 5.06%) because we predict 600 classes (including 120 unseen categories) in baseline. The model could learn the verb and object of unseen HOIs individually from seen HOIs. Meanwhile, non-rare first selection has more verbs and objects of unseen HOIs individually from seen HOIs than that of rare first selection. Therefore, non-rare first selection has better baseline than rare first selection. Besides, we improve [3, 29] largely nearly among all categories. In detail, for a fair comparison to [3], we also use Resnet-101 and finetune the object detector on HICO-DET dataset, which can largely improve the per-

Table 2. Comparison of Zero Shot Detection results of our proposed Visual Compositional Learning framework. * means we uses the res101 backbone and finetune object detector on HICO-DET. Full means all categories including Seen and Unseen

Method	Default			Known Object		
	Unseen	Seen	Full	Unseen	Seen	Full
Shen <i>et al.</i> [29]	5.62	-	6.26	-	-	-
Bansal* <i>et al.</i> [3]	10.93	12.60	12.26	-	-	-
w/o VCL (rare first)	3.30	18.63	15.56	5.53	21.30	18.15
w/o VCL (non-rare first)	5.06	12.77	11.23	8.81	15.37	14.06
VCL (rare first)	7.55	18.84	16.58	10.66	21.56	19.38
VCL (non-rare first)	9.13	13.67	12.76	12.97	16.31	15.64
VCL* (rare first)	10.06	24.28	21.43	12.12	26.71	23.79
VCL* (non-rare first)	16.22	18.52	18.06	20.93	21.02	20.90

formance. This demonstrates that our VCL effectively improves the detection of unseen interactions and maintains excellent HOI detection performance for seen interactions at the same time.

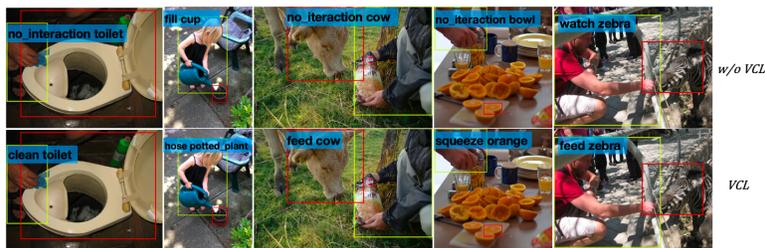


Fig. 4. Some rare HOI detections (Top 1 result) detected by the proposed Compositional Learning and the model without Compositional Learning. The first row is the results of baseline model without VCL. The second row is the results of VCL

4.5 Ablation Analysis

To evaluate the design of our VCL, we first conduct ablation studies about VCL and verb representation on HICO-DET dataset and V-COCO dataset. For V-COCO, we evaluate AP_{role} (24 actions with roles) following [7, 11]. Besides, we evaluate two branches and composing interactions between images and/or within images on HICO-DET. See supplementary material for more analysis.

Visual Compositional Learning is our core method. With VCL, we can respectively see an obvious increase by 1.00% in the Full category from Table 3 (row 6 vs row 7). Particularly, the improvement (2.41%) in Rare category is

Table 3. Ablation study of the proposed Visual Compositional Learning framework on HICO-DET and V-COCO test set. VCL means Visual Compositional Learning and Union Verb means we learn verb feature from union box of human and object. Sharing W means sharing weights between human stream and verb stream for verb representation. Re-weighting means we use re-weighting strategy in the code of [20]

VCL	Union Verb	Sharing W	Re-weighting	HICO			V-COCO
				Full	Rare	NonRare	AP_{role}
-	-	✓	-	16.87	10.07	18.90	46.9
✓	-	✓	-	17.35	12.10	18.91	47.2
✓	✓	-	✓	18.93	15.68	19.90	47.8
-	-	✓	✓	18.03	13.62	19.35	47.4
✓	-	✓	✓	18.57	14.83	19.69	47.7
-	✓	✓	✓	18.43	14.14	19.71	47.5
✓	✓	✓	✓	19.43	16.55	20.29	48.3

considerably better than that in Non-Rare category, which means the proposed VCL is more beneficial for rare data. Noticeably, the improvement of VCL with verb representation learning is better than that without verb representation, which implies that the more discriminative the verb representation is, the better improvement VCL obtains in HOI detection. We can see a similar trend of VCL on V-COCO in Table 3 where the performance decreases by 0.8% without VCL. Noticeably, V-COCO aims to recognize the actions (verbs) rather than the human-object pairs, which means the proposed VCL is also helpful to learn the discriminative action (verbs) representation. Besides, from row 1 and row 2 in Table 3, we can find the proposed VCL is pluggable to re-weighting strategy [17] (See details in supplementary materials).

Verb Representation, which we learn from the union box of human and object, is considerably helpful for the performance. We evaluate its efficiency by comparing human box and union box of human and object. From row 5 and row 7 in Table 3, we can find that with the proposed verb representation we improve the performance by 0.86% on HICO-DET and 0.6% in V-COCO respectively within VCL, which means learning verb representation from the union box of human and object is more discriminative and advantageous for HOI detection. Particularly, we *share the weights of resnet block between human and verb stream* after ROI pooling in the proposed verb representation method. Table 3 (row 3 vs row 7) shows sharing weights helps the model improve the performance largely from 18.93% to 19.43%. This may be explained by that the union region contains more noise and the model would emphasize the region of the human in the union box by sharing weights to obtain better verb representation.

Branches. There are two branches in our method and we evaluate their contributions in Table 4. Noticeably, we apply VCL to Verb-Object branch during training, while we do not apply VCL to Spatial-Human. By keeping one branch each time on HICO-DET dataset during inference, we can find the verb-object branch makes the larger contribution, particularly for rare category (**3%**).

Table 4. The branches ablation study of the model on HICO-DET test set. Verb-object branch only means we train the model without spatial-human branch.

Method	Full	Rare	NonRare
Two branches	19.43	16.55	20.29
Verb-Object branch only during inference	16.89	15.35	17.35
Spatial-Human branch only during inference	16.13	12.39	17.24
Verb-Object branch only	15.77	13.35	16.49
Verb-Object branch only (w/o VCL)	15.33	10.85	16.67

This efficiently illustrates the advantage of VCL for rare categories. But we can improve the performance dramatically from 16.89% to 19.43% with Spatial-Human Branch. Meanwhile, we can find the proposed VCL is orthogonal to spatial-human branch from the last two rows in Table 4. Noticeably, *by comparing verb-object branch only during inference and verb-object branch only from training, we can find the spatial-human branch can facilitate the optimization of verb-object branch (improving the mAP from 15.77% to 16.89%)*. See supplementary materials for more analysis of two branches in zero-shot detection where the performance of verb-object branch with VCL is over 3.5% better than spatial-human branch in unseen data.

Table 5. Composing Strategies study of VCL on HICO-DET test set

Method	Full (mAP %)	Rare (mAP %)	NonRare (mAP %)
Baseline (w/o VCL)	18.43	14.14	19.71
Within images	18.48	14.46	19.69
Between images	19.06	14.33	20.47
Between and within images	19.43	16.55	20.29

Composing interactions within and/or between images. In Table 5, we can find composing interaction samples between images is beneficial for HOI detection, whose performance in the Full category increases to 19.06% mAP, while composing interaction samples within images has similar results to baseline. It might be because the number of images including multiple interactions is few on HICO-DET dataset. Remarkably, composing interaction samples within and between images notably improves the performance up to 19.43% mAP in Full and **16.55%** mAP in Rare respectively. Those results mean composing interactions within images and between images is more beneficial for HOI detection.

4.6 Visualization of features

We also illustrate the verb and object features by t-SNE visualization [25]. Figure 5 illustrates that VCL overall improves the the discrimination of verb and

object features. There are many noisy points (see black circle region) in Figure 5 without VCL and verb presentation. Meanwhile, we can find the proposed verb representation learning is helpful for verb feature learning by comparing verb t-SNE graph between the left and the middle. Besides, the object features are more discriminative than verb. We think it is because the verb feature is more abstract and complex and the verb representation requires further exploration.

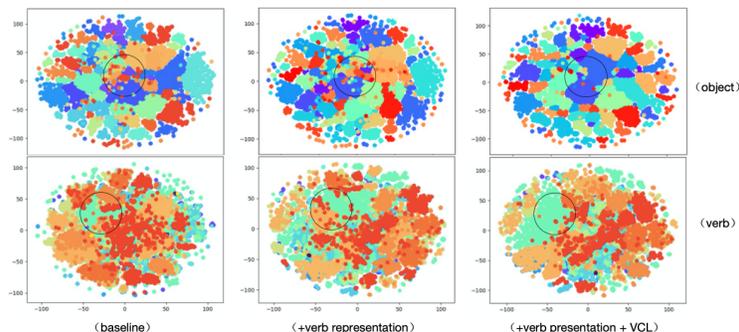


Fig. 5. Visual illustration of object features (80 classes) (up) and verb features (117 classes) (bottom) on HICO-DET dataset (20000 samples) via t-SNE visualization [25]. Left is the visual illustration of baseline, middle includes verb representation and the right uses both VCL and verb representation

5 Conclusion

In this paper, we propose a simple yet efficient deep Visual Compositional Learning framework, which composes the interactions from the shared verb and object latent representations between images and within images, to relieve the long-tail distribution issue and zero-shot learning in human-object interaction detection. Meanwhile, we extract a more discriminative verb representation from the union box of human and object rather than human box. Lastly, we evaluate the efficiency of our model on two HOI detection benchmarks, particularly for low-shot and zero-shot detection.

Acknowledge This work is partially supported by Science and Technology Service Network Initiative of Chinese Academy of Sciences (KFJ-ST-S-QYZX-092), Guangdong Special Support Program (2016TX03X276), National Natural Science Foundation of China (U1813218, U1713208), Shenzhen Basic Research Program (JCYJ20170818164704758, CXB201104220032A), the Joint Lab of CAS-HK, Australian Research Council Projects (FL-170100117).

References

1. Alfassy, A., Karlinsky, L., Aides, A., Shtok, J., Harary, S., Feris, R., Giryes, R., Bronstein, A.M.: Laso: Label-set operations networks for multi-label few-shot learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6548–6557 (2019)
2. Azadi, S., Pathak, D., Ebrahimi, S., Darrell, T.: Compositional gan: Learning conditional image composition. arXiv preprint arXiv:1807.07560 (2018)
3. Bansal, A., Rambhatla, S.S., Shrivastava, A., Chellappa, R.: Detecting human-object interactions via functional generalization. arXiv preprint arXiv:1904.03181 (2019)
4. Battaglia, P.W., Hamrick, J.B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al.: Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261 (2018)
5. Bengio, Y., Courville, A., Vincent, P.: Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* **35**(8), 1798–1828 (2013)
6. Burgess, C.P., Matthey, L., Watters, N., Kabra, R., Higgins, I., Botvinick, M., Lerchner, A.: Monet: Unsupervised scene decomposition and representation. arXiv preprint arXiv:1901.11390 (2019)
7. Chao, Y.W., Liu, Y., Liu, X., Zeng, H., Deng, J.: Learning to detect human-object interactions. In: 2018 IEEE winter conference on applications of computer vision (WACV). pp. 381–389. IEEE (2018)
8. Chao, Y.W., Wang, Z., He, Y., Wang, J., Deng, J.: Hico: A benchmark for recognizing human-object interactions in images. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1017–1025 (2015)
9. Gao, C., Zou, Y., Huang, J.B.: ican: Instance-centric attention network for human-object interaction detection. arXiv preprint arXiv:1808.10437 (2018)
10. Garnelo, M., Shanahan, M.: Reconciling deep learning with symbolic artificial intelligence: representing objects and relations. *Current Opinion in Behavioral Sciences* **29**, 17–23 (2019)
11. Gupta, S., Malik, J.: Visual semantic role labeling. arXiv preprint arXiv:1505.04474 (2015)
12. Gupta, T., Schwing, A., Hoiem, D.: No-frills human-object interaction detection: Factorization, appearance and layout encodings, and training techniques. arXiv preprint arXiv:1811.05967 (2018)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
14. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR* **2**(5), 6 (2017)
15. Higgins, I., Sonnerat, N., Matthey, L., Pal, A., Burgess, C.P., Bosnjak, M., Shanahan, M., Botvinick, M., Hassabis, D., Lerchner, A.: Scan: Learning hierarchical compositional visual concepts. arXiv preprint arXiv:1707.03389 (2017)
16. Hoffman, D.D., Richards, W.: Parts of recognition (1983)
17. Japkowicz, N., Stephen, S.: The class imbalance problem: A systematic study. *Intelligent data analysis* **6**(5), 429–449 (2002)

18. Kato, K., Li, Y., Gupta, A.: Compositional learning for human object interaction. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 234–251 (2018)
19. Lake, B.M., Ullman, T.D., Tenenbaum, J.B., Gershman, S.J.: Building machines that learn and think like people. *Behavioral and brain sciences* **40** (2017)
20. Li, Y.L., Zhou, S., Huang, X., Xu, L., Ma, Z., Fang, H.S., Wang, Y.F., Lu, C.: Transferable interactiveness prior for human-object interaction detection. arXiv preprint arXiv:1811.08264 (2018)
21. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
22. Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., Yu, S.X.: Large-scale long-tailed recognition in an open world. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2537–2546 (2019)
23. Locatello, F., Bauer, S., Lucic, M., Gelly, S., Schölkopf, B., Bachem, O.: Challenging common assumptions in the unsupervised learning of disentangled representations. arXiv preprint arXiv:1811.12359 (2018)
24. Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. In: European conference on computer vision. pp. 852–869. Springer (2016)
25. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov), 2579–2605 (2008)
26. Peyre, J., Laptev, I., Schmid, C., Sivic, J.: Detecting unseen visual relations using analogies. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
27. Qi, S., Wang, W., Jia, B., Shen, J., Zhu, S.C.: Learning human-object interactions by graph parsing neural networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 401–417 (2018)
28. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in neural information processing systems. pp. 91–99 (2015)
29. Shen, L., Yeung, S., Hoffman, J., Mori, G., Fei-Fei, L.: Scaling human-object interaction recognition through zero-shot learning. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1568–1576. IEEE (2018)
30. Spelke, E.S.: Principles of object perception. *Cognitive science* **14**(1), 29–56 (1990)
31. Wan, B., Zhou, D., Liu, Y., Li, R., He, X.: Pose-aware multi-level feature network for human object interaction detection. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9469–9478 (2019)
32. Wang, T., Anwer, R.M., Khan, M.H., Khan, F.S., Pang, Y., Shao, L., Laaksonen, J.: Deep contextual attention for human-object interaction detection. arXiv preprint arXiv:1910.07721 (2019)
33. Wang, Y.X., Girshick, R., Hebert, M., Hariharan, B.: Low-shot learning from imaginary data. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7278–7286 (2018)
34. Xian, Y., Lampert, C.H., Schiele, B., Akata, Z.: Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence* **41**(9), 2251–2265 (2018)
35. Xian, Y., Schiele, B., Akata, Z.: Zero-shot learning—the good, the bad and the ugly. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4582–4591 (2017)

36. Xu, B., Wong, Y., Li, J., Zhao, Q., Kankanhalli, M.S.: Learning to detect human-object interactions with knowledge. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019)
37. Yang, X., Zhang, H., Cai, J.: Shuffle-then-assemble: Learning object-agnostic visual relationship features. In: Proceedings of the European conference on computer vision (ECCV). pp. 36–52 (2018)
38. Zhou, P., Chi, M.: Relation parsing neural network for human-object interaction detection. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)