

Detecting Human-Object Interactions with Action Co-occurrence Priors

Dong-Jin Kim¹ Xiao Sun² Jinsoo Choi¹ Stephen Lin² In So Kweon¹

¹ KAIST, South Korea. {dijnjusa, jinsc37, iskweon77}@kaist.ac.kr

² Microsoft Research, China. {xias, stevelin}@microsoft.com

<https://github.com/Dong-JinKim/ActionCooccurrencePriors/>

Abstract. A common problem in human-object interaction (HOI) detection task is that numerous HOI classes have only a small number of labeled examples, resulting in training sets with a long-tailed distribution. The lack of positive labels can lead to low classification accuracy for these classes. Towards addressing this issue, we observe that there exist natural correlations and anti-correlations among human-object interactions. In this paper, we model the correlations as *action co-occurrence matrices* and present techniques to learn these priors and leverage them for more effective training, especially on rare classes. The utility of our approach is demonstrated experimentally, where the performance of our approach exceeds the state-of-the-art methods on both of the two leading HOI detection benchmark datasets, HICO-Det and V-COCO.

1 Introduction

Human-object interaction (HOI) detection aims to localize humans and objects in an image and infer the relationships between them. An HOI is typically represented as a human-action-object triplet with the corresponding bounding boxes and classes. Detecting these interactions is a fundamental challenge in visual recognition that requires both an understanding of object information and high-level knowledge of interactions.

A major issue that exists in HOI detection is that its datasets suffer from long-tailed distributions, in which many HOI triplets have few labeled instances. Similar to datasets for general visual relationship detection (VRD) [37], a reason for this is missing labels, where the annotation covers only a subset of the interactions present in an image. For the widely-used HICO-Det dataset [3], 462 out of the 600 HOI classes have fewer than 10 training samples. For such classes, the lack of positive labels can lead to inadequate training and low classification performance. How to alleviate the performance degradation on rare classes is thus a key issue in HOI detection.

To address the problem of long-tailed distributions, we propose to take advantage of natural *co-occurrences* in human actions. For example, the HOI of ‘operate-hair dryer’ is rarely labeled and consequently hard to detect in the left image of Fig. 1. However, ‘operate-hair dryer’ often occurs when the more commonly labeled HOI of ‘hold-hair dryer’ is present. As a result, detection of



Fig. 1. Marginal/conditional probability values are computed from the distribution of the training label. Intuitively, detection of HOIs (operate-hair dryer) can be facilitated by detection of commonly co-occurring HOIs (hold-hair dryer). Also, non-detection of HOIs (blow-cake) can be aided by detection of incompatible HOIs (cut-cake). We leverage this intuition as a prior to learn an HOI detector on long-tailed datasets.

‘operate-hair dryer’ can be facilitated by detection of ‘hold-hair dryer’ in an image. On the other hand, the detection of an HOI may preclude other incompatible HOIs, such as for ‘cut-cake’ and ‘blow-cake’ in the right image of Fig. 1.

In this paper, we introduce the new concept of utilizing co-occurring actions as prior knowledge, termed as action co-occurrence priors (ACPs), to train an HOI detector. In contrast to the language-based prior knowledge which requires external data sources [27,37,61], the co-occurrence priors can be easily obtained from the statistics of the target dataset. We also propose two novel ways to exploit them. First, we design a neural network with hierarchical structure where the classification is initially performed with respect to *action groups*. Each action group is defined by one anchor action, where the anchor actions are mutually exclusive according to the co-occurrence prior. Then, our model predicts the fine-grained HOI class within the action group. Second, we present a technique that employs knowledge distillation [20] to expand HOI labels so they can have more positive labels for potentially co-occurring actions. During training, the predictions are regularized by the refined objectives to improve robustness, especially for classes in the data distribution tail. To the best of our knowledge, we are the first to leverage the label co-occurrence in HOI detection to alleviate long-tailed distribution.

The main contributions of this work can be summarized as: (1) The novel concept of explicitly leveraging correlations among HOI labels to address the problem of long-tailed distributions in HOI detection; (2) Two orthogonal ways to leverage action co-occurrence priors, namely through a proposed hierarchical architecture and HOI label expansion via knowledge distillation. The resulting model is shown to be consistently advantageous in relation to state-of-the-art techniques on both the HICO-Det [3] and V-COCO [16] benchmark datasets.

2 Related Work

Human-Object Interaction Human-Object Interaction was originally studied in the context of recognizing the function or ‘affordance’ of objects [6,12,14,15,46].

Early works focus on learning more discriminative features combined with variants of SVM classifiers [7,8,57], and leverage the relationship with human poses for better representation [7,8,59] or mutual context modeling [58].

Recently, a completely data-driven approach based on convolutional neural networks (CNNs) has brought dramatic progress to HOI. Many of the pioneering works created large scale image datasets [3,4,16,68] to set new benchmarks in this field. Henceforth, significant progress have been seen in using CNNs for this problem [1,3,11,13,18,25,32,33,35,43,45,49,50,52,53,54]. Most of these works follow a two-step scheme of CNN feature extraction and multi-information fusion, where the multiple information may include human and object appearance [3,11,13,43]; box relation (either box configuration or spatial map) [1,3,13,18,54]; object category [1,18,41]; human pose [18,32,33]; and particularly, linguistic prior knowledge [25,41]. More recent works tend to combine these various cues [18,33,50,53]. These works differ from one another mainly in their techniques for exploiting external knowledge priors. Kato *et al.* [25] incorporate information from WordNet [39] using a Graph Convolutional Network (GCN) [29] and learn to compose new HOIs. Xu *et al.* [54] also use a GCN to model the general dependencies among actions and object categories by leveraging a VRD dataset [37]. Li *et al.* [33] utilize interactiveness knowledge learned across multiple HOI datasets. Peyre *et al.* [41] transfer knowledge from triplets seen at training to new unseen triplets at test time by analogy reasoning.

Different from the previous works, our approach is to reformulate the target action label space and corresponding loss function in a manner that leverages *co-occurrence* relationships among action classes for HOI detection. In principle, the proposed technique is complementary to all of the previous works and can be combined with any of them. For our experiments, we implemented our approach on a baseline presented in [18], with details given in Sec. 3.3.

Visual Relationship Detection The closest problem to HOI detection is Visual Relationship Detection (VRD) [5,31,42,56,60,61,62,64,65,67], which deals with general visual relationships between two arbitrary objects. In the VRD datasets [30,37], the types of visual relationships that are modeled include verb (action), preposition, spatial and comparative phrase. The two tasks share common challenges such as long-tail distributions or even zero-shot problems [37]. Our work focuses on HOI detection, as co-occurrences of human-object interactions are often strong, but the proposed technique could be extended to model the general co-occurrences that exist in visual relationships.

Label Hierarchy in Multi-label Learning The hierarchical structure of label categories has long been exploited for multi-label learning in various vision tasks, *e.g.*, image/object classification [9,55], detection [23,38], and human pose estimation [47,48]. In contrast, label hierarchy has rarely been considered in HOI detection. Inspired by previous work [9] that uses Hierarchy and Exclusion (HEX) graphs to encode flexible relations between object labels, we present the first method to take advantage of an action label hierarchy for HOI recognition. While label hierarchies have commonly been used, our method is different in that it is defined by co-occurrences (rather than semantics or a taxonomy [9]). This

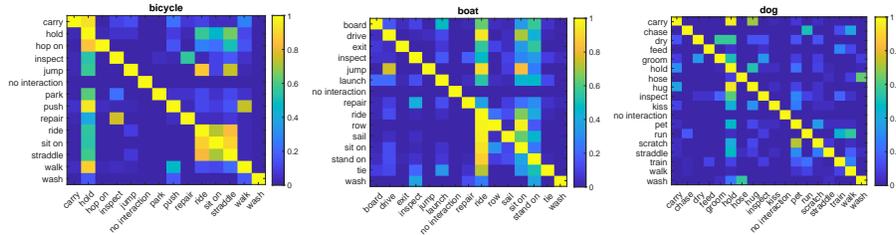


Fig. 2. Examples of co-occurrence matrices constructed for several objects (bicycle, boat, dog). Along the Y-axis is the given action, and the X-axis enumerates conditional actions. Each element represents the conditional probability that an action occurs when another action is happening.

co-occurrence based hierarchy can be determined statistically, without direct human supervision.

3 Proposed Method

Our method for utilizing the co-occurrence information of HOI labels consists of three key components: (1) establishing action co-occurrence priors (Sec. 3.1), (2) hierarchical learning including anchor action selection (Sec. 3.2) and devising the hierarchical architecture (Sec. 3.3), and (3) ACP projection for knowledge distillation (Sec. 3.4).

3.1 Action Co-occurrence Priors

Here, we formalize the action co-occurrence priors. The priors for the actions are modeled by a co-occurrence matrix $C \in \mathbb{R}^{N \times N}$ where an entry c_{ij} in C represents the conditional probability that action j occurs when action i is happening:

$$c_{ij} = p(j|i), i, j \in [0, N), \quad (1)$$

where N denotes the total number of actions classes and i, j are indices of two actions. C is constructed from the target HOI detection dataset by counting the image-level statistics of its training labels. Examples of co-occurrence matrices constructed for single object are visualized in Fig. 2.

Meanwhile, we also consider the complementary event of action i (*i.e.*, where the i -th action does not occur) and denote it as i' , such that $p(i') + p(i) = 1$. The complementary action co-occurrence matrix $C' \in \mathbb{R}^{N \times N}$ can thus be defined by entries c'_{ij} in C' that represent the conditional probability that an action j occurs when another action i does not occur:

$$c'_{ij} = p(j|i'), i, j \in [0, N). \quad (2)$$

It can be seen from Fig. 2, that different types of relationships can exist between actions. The types can be divided into three types. The first type is the

prerequisite relationship, where the given action is highly likely to co-occur with the conditional action. For example, the HOI ‘sit on-bicycle’ is a prerequisite of the HOI ‘ride-bicycle’. In this case, $p(\text{sit on-bicycle}|\text{ride-bicycle})$ is close to 1. Next, **exclusion**, where the given action is highly unlikely to co-occur with the conditional action. An example is that the HOI ‘wash-bicycle’ and HOI ‘ride-bicycle’ are unlikely to happen together. As a result, $p(\text{wash-bicycle}|\text{ride bicycle})$ is close to 0. Finally, **overlapping**, where the given action and conditional action may possibly co-occur, for example HOI ‘hold-bicycle’ and HOI ‘inspect-bicycle’, such that $p(\text{hold-bicycle}|\text{inspect-bicycle})$ is in between 0 and 1.

The strong relationships that may exist between action labels can provide strong priors on the presence or absence of an HOI in an image. In contrast to previous works where models may implicitly learn label co-occurrence via relational architectures [2,63], we explicitly exploit these relationships between action labels as priors, to effectively train our model especially for rare HOIs.

3.2 Anchor Action Selection via Non-Exclusive Action Suppression

From a co-occurrence matrix for an object, it can be seen that some actions are close in semantics or commonly co-occur while others are not. Intuitively, closely related actions (*e.g.*, ‘sit on-bicycle’ and ‘straddle-bicycle’) tend to be harder to distinguish from each other. If the positive labels for these actions are rare, then they become even more difficult to distinguish. Such cases require fine-grained recognition [10] and demand more dedicated classifiers. This motivates us to learn HOI classes in a coarse-to-fine manner. Specifically, we first identify a set of mutually exclusive action classes, called *anchor actions*, which tend to be distinguishable from one another. The anchor actions will be used to partition the entire action label space into fine-grained sub-spaces. The other action classes will be attributed to one or more sub-spaces and recognized in the context of a specific anchor action. In summary, unlike previous HOI detection works which predict action probabilities independently of one another, we divide the whole action label set into two sets, one for anchor actions and one for regular actions, which are modeled in different ways as explained in detail in Sec. 3.3.

In selecting anchor actions, we seek a set of action classes that are exclusive of one another. Toward this end, we define the exclusiveness of an action class as counting the number of actions that will never occur if action i is happening ($e_i = \sum_j (1 \text{ if } (c_{ij} = 0), \text{ else } 0)$). e_i will have a high value if few other actions can occur when i does. Based on exclusiveness values, the anchor action label set \mathcal{D} is generated through *non-exclusive suppression (NES)*. It iteratively finds the most exclusive action class as an anchor action and removes remaining action classes that are not exclusive to the selected anchor actions. The anchors in the list are action classes that never occur together in the training labels. For example, if an action such as ‘toast’ is on the anchor list, then actions like ‘stand’ and ‘sit’ cannot be on the list because they may co-occur with ‘toast’, while actions such as ‘hunt’ or ‘hop on’ can potentially be on the list. While there may exist other ways the anchor action selection could be done, we empirically

found this approach to be simple, effective (detection accuracy), and efficient (less than 0.01 second).

The anchor action label set acts as a finite partition of the action label space (a set of pairwise disjoint events whose union is the entire action label space). To form a complete action label space, we add an ‘*other*’ anchor action, denoted as \mathcal{O} , for when an action class does not belong to \mathcal{D} . Finally, we get $|\mathcal{D}| + 1$ anchor actions including \mathcal{D} and the ‘other’ action class \mathcal{O} .

There are several benefits to having this anchor action label set. First, only one anchor action can happen at one time between a given human and object. Thus, we can use the relative (one hot) probability representation with *softmax* activation, whose features were shown to compare well against distance metric learning-based features [21]. Second, anchor actions tend to be easier to distinguish from one another since they generally have prominent differences in an image. Third, it decomposes the action label space into several sub-spaces, which facilitates a coarse-to-fine solution. Each sub-task will have a much smaller solution space, which can improve learning. Finally, each sub-task will use a standalone sub-network which focuses on image features specific to the sub-task, which is an effective strategy for fine-grained recognition [10].

After selecting anchor actions, the entire action label set \mathcal{A} is divided into the anchor action label set \mathcal{D} and the remaining set of ‘regular’ action classes \mathcal{R} , so that $\mathcal{A} = \{\mathcal{D}, \mathcal{R}\}$. Each of the regular action classes is then associated with one or more anchor actions to form $|\mathcal{D}| + 1$ action groups $\mathbf{G} = \{\mathcal{G}_i; i \in \mathcal{D} \cup \mathcal{O}\}$, one for each anchor action. A regular action class $j \in \mathcal{R}$ will be assigned to the group of anchor action i (\mathcal{G}_i) if action j is able to co-occur with the anchor action i ,

$$j \in \mathcal{G}_i, \text{ if } c_{ij} > 0 \quad (i \in \mathcal{D} \cup \mathcal{O}, j \in \mathcal{R}). \quad (3)$$

Note that the anchor actions themselves are not included in the action groups and a regular action j can be assigned to multiple action groups since it may co-occur with multiple anchor actions.

3.3 Hierarchical Architecture

We implemented our hierarchical approach upon the ‘No-Frills’ (NFs) baseline presented in [18] on account of its simplicity, effectiveness, and code availability [17]. Here, we give a brief review of the NFs architecture.

Baseline Network NFs follows the common scheme of CNN feature extraction followed by multi-information fusion. It uses the off-the-shelf Faster R-CNN [44] object detector with ResNet152 [19] backbone network to detect human and object bounding boxes. As illustrated in Fig. 3, the multiple information used in [18] (denoted as X) are fed through four separate network streams to generate fixed dimension features. Then, all the features are added together and sent through a *sigmoid* activation to get the action probability prediction \hat{A} :

$$\hat{A} = \text{sigmoid}(F(X)) \in \mathbb{R}^N, \quad (4)$$

where $\hat{A}(a) = p(a|X)$ represents the probability prediction for action class a .

To eliminate training-inference mismatch, NFs directly optimizes the HOI class probabilities instead of separating the detection and interaction losses as done in [11,13]. The final HOI prediction is a joint probability distribution over M number of HOI classes computed from the probabilities \hat{H} , \hat{O} , and \hat{A} for human, object, and action, respectively:

$$\hat{Y} = \text{joint}(\hat{H}, \hat{O}, \hat{A}) \in \mathbb{R}^M. \quad (5)$$

Specifically, for a HOI class (h, o, a) ,

$$\hat{Y}(h, o, a) = \hat{H}(h) * \hat{O}(o) * \hat{A}(a) = p(h|I) * p(o|I) * p(a|X), \quad (6)$$

where $\hat{H}(h) = p(h|I)$ and $\hat{O}(o) = p(o|I)$ are the probability of a candidate box pair being a human h and object o , provided by the object detector [44]. Finally, the binary cross-entropy loss $\mathcal{L}(\hat{Y}, Y^{gt})$ is directly computed from the HOI prediction. This ‘No-Frills’ baseline network is referred to as **Baseline**.

Modified Baseline Network For a stronger baseline comparison, we make two simple but very effective modifications on the baseline network. (1) Replace the one-hot representation with the Glove word2vec [40] representation for the object category. (2) Instead of directly adding up the multiple information, we average them and forward this through another *action prediction module* to obtain the final action probability prediction. For a naive approach (the **Modified Baseline**), we simply use a sub-network f_{sub} of a few FC layers as the *action prediction module*. Then Eq. (4) is modified to

$$\hat{A} = \text{sigmoid}(f_{sub}(F(X))). \quad (7)$$

Our hierarchical architecture further modifies the *action prediction module* by explicitly exploiting ACP information which is described in the next paragraph. **Proposed Hierarchical Architecture** Now we introduce the *action prediction module* for our hierarchical architecture (illustrated in Fig. 3) that better exploits the inherent co-occurrence among actions. While the baseline network predicts all the action probabilities directly from $F(\cdot)$ with a single feed-forward sub-network f_{sub} , we instead use $|\mathcal{D}| + 2$ sub-networks where one ($f_{anchor}(\cdot)$) is first applied to predict the anchor action set and then one of the $|\mathcal{D}| + 1$ other sub-networks ($f_{\mathcal{G}_i}(\cdot)$) which corresponds to the predicted anchor action is used to estimate the specific action within the action group. Because of the mutually exclusive property of anchor actions, we use the *softmax* activation for anchor action predictions, while employing the *sigmoid* activation for regular action predictions conditional to the action groups:

$$\hat{A}_{anchor} = \text{softmax}(f_{anchor}(F(X))) \in \mathbb{R}^{|\mathcal{D}|+1} \quad (8)$$

$$\hat{A}_{\mathcal{G}_i} = \text{sigmoid}(f_{\mathcal{G}_i}(F(X))) \in \mathbb{R}^{N-|\mathcal{D}|}, \text{ where } i \in \mathcal{D} \cup \mathcal{O}, \quad (9)$$

where $\hat{A}_{anchor}(i)$ is directly used as the final probability predictions for the anchor actions ($p(i|X) = \hat{A}_{anchor}(i), i \in \mathcal{D}$). We let $\hat{A}_{\mathcal{G}_i}(j)$ represent the learned conditional probability that action j occurs when action i is happening ($p(j|i, X) =$

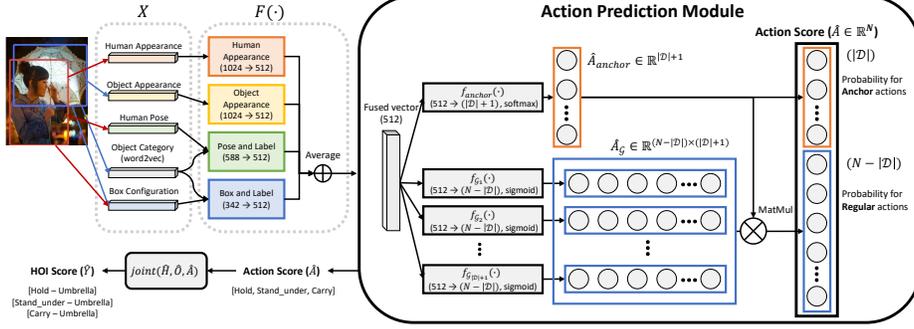


Fig. 3. Illustration of our overall network architecture. Our work differs from the baseline [18] by the addition of a hierarchical *action prediction module*. For our hierarchical architecture, anchor action probability is directly generated by a *softmax* sub-network. Regular action probability is generated by a matrix multiplication of the anchor probability and the output from a few *sigmoid* based conditional sub-networks.

$\hat{A}_{G_i}(j)$). Since the anchor action set is a finite partition of the entire action label space, the probability of a regular action j can be predicted according to the law of total probability:

$$\hat{A}_{regular}(j) = p(j|X) = \sum_{i \in \mathcal{D} \cup \mathcal{O}} p(i|X) * p(j|i, X) = \sum_{i \in \mathcal{D} \cup \mathcal{O}} \hat{A}_{anchor}(i) * \hat{A}_{G_i}(j), \quad (10)$$

where $j \in \mathcal{R}$. Thus, instead of Eq. (7), we obtain the final action probability predictions for our hierarchical architecture $\hat{A}(a) = p(a|X)$ as

$$\hat{A}(a) = \begin{cases} \hat{A}_{anchor}(a), & \text{if } a \in \mathcal{D} \\ \sum_{i \in \mathcal{D} \cup \mathcal{O}} \hat{A}_{anchor}(i) * \hat{A}_{G_i}(a), & \text{otherwise.} \end{cases} \quad (11)$$

We use the same method as in Eq. (6) and cross-entropy loss to compute the final HOI probability prediction \hat{Y} and the corresponding loss \mathcal{L} .

To demonstrate the effectiveness of the hierarchical learning, we introduce another two baselines, *MultiTask* and *TwoStream*, that lie between the *Modified Baseline* and our hierarchical learning. *MultiTask* only uses the anchor action classification as an additional multi-task element to the *Modified Baseline*. *TwoStream* separately predicts the anchor and the regular actions but without using the hierarchical modeling between anchor and regular actions.

3.4 ACP Projection for Knowledge Distillation

Knowledge distillation [20] was originally proposed to transfer knowledge from a large network to a smaller one. Recently, knowledge distillation has been utilized for various purposes such as life-long learning [34] or multi-task learning [28]. Hu *et al.* [22] extended this concept to distill prior knowledge in the form of logic

rules into a deep neural network. Specifically, they propose a teacher-student framework to project the network prediction (student) to a rule-regularized subspace (teacher), where the process is termed distillation. The network is then updated to balance between emulating the teacher’s output and predicting the true labels.

Our work fits this setting because the ACPs can act as a prior to distill. We first introduce an *ACP Projection* to map the action distribution to the ACP constraints. Then, we use the teacher-student framework [22] to distill knowledge from ACPs.

ACP Projection In ACP Projection, an arbitrary action distribution $A = \{p(i), i \in [0, N]\} \in \mathbb{R}^N$ is projected into the ACP-constrained probability space:

$$A^* = \text{project}(A, C, C') \in \mathbb{R}^N, \quad (12)$$

where A^* is the projected action prediction. The projected probability for the j -th action $A^*(j) = p(j^*)$ is generated using the law of total probability:

$$p(j^*) = \frac{1}{N} \sum_{i=1}^N (p(i) * p(j|i) + p(i') * p(j|i')) = \frac{1}{N} \left(\sum_{i=1}^N p(i) * c_{ij} + \sum_{i=1}^N (1-p(i)) * c'_{ij} \right). \quad (13)$$

In matrix form, the ACP projection is expressed as

$$\text{project}(A, C, C') = \frac{AC + (1 - A)C'}{N}. \quad (14)$$

In practice, we use the object-based action co-occurrence matrices $C_o \in \mathbb{R}^{N \times N}$ and $C'_o \in \mathbb{R}^{N \times N}$ which only count actions related to a specific object o . Fig. 2 shows examples of C_o with respect to object classes. Also, we give different weights α and β as hyper-parameters to the action co-occurrence matrix C_o and its complementary matrix C'_o , with the weights subject to $\alpha + \beta = 2$, $\alpha > \beta$. The projection function is then modified as

$$\text{project}(A, C_o, C'_o) = \frac{\alpha AC_o + \beta(1 - A)C'_o}{N}. \quad (15)$$

This is done because we empirically found the co-occurrence relationships in C_o to generally be much stronger than the complementary actions in C'_o .

Teacher-Student Framework Now we can distill knowledge from the ACPs using ACP Projection in both the training and inference phases. There are three ways ACP Projection can be used: (1) Directly project the action prediction \hat{A} into the ACP-constrained probability space at the testing phase to obtain the final action output (denoted as *PostProcess*). (2) Project the action prediction \hat{A} in the training phase and use the projected action as an additional learning target [22,61]. (3) Project the ground truth label H^{gt} , O^{gt} , and A^{gt3} to the ACP

³ The triplet ground truth labels H^{gt} , O^{gt} , and A^{gt} are straightforward to determine from the HOI ground truth label Y^{gt} .

space in the training phase and use the projected action $project(A^{gt}, C_{O^{gt}}, C'_{O^{gt}})$ as an additional learning target. The second and third items are incorporated into the teacher-student framework as terms in a new objective function (denoted as **Distillation**):

$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}(\hat{Y}, Y^{gt}) + \lambda_2 \mathcal{L}(\hat{Y}, \hat{Y}_{projO}) + \lambda_3 \mathcal{L}(\hat{Y}, Y_{projO}^{gt}), \quad (16)$$

where

$$\hat{Y}_{projO} = joint(\hat{H}, \hat{O}, project(\hat{A}, C_{\hat{O}}, C'_{\hat{O}})) \in \mathbb{R}^M, \quad (17)$$

$$Y_{projO}^{gt} = joint(H^{gt}, O^{gt}, project(A^{gt}, C_{O^{gt}}, C'_{O^{gt}})) \in \mathbb{R}^M. \quad (18)$$

$\lambda_1, \lambda_2, \lambda_3$ are balancing weights among the ground truth HOI term and the teacher objectives. The object type can be easily determined from the object probability predictions \hat{O} or the ground truth label O^{gt} .

4 Experiments

The goal of the experiments is to show the effectiveness and generalizability of our method. In particular, we show that our method can consistently alleviate the long-tailed distribution problem in various setups by improving performance especially for rare HOI classes. In this section, we describe the experimental setups, competing methods and provide performance evaluations of HOI detection.

4.1 Datasets and Metrics

We evaluate the performance of our model on the two popular HOI detection benchmark datasets, **HICO-Det** [3] and **V-COCO** [16]. HICO-Det [3] extends the HICO (Humans Interacting with Common Objects) dataset [4] which contains 600 human-object interaction categories for 80 objects. Different from the HICO dataset, HICO-Det additionally contains bounding box annotations for humans and objects of each HOI category. The vocabulary of objects matches the 80 categories of MS COCO [36], and there are 117 different verb (action) categories. The number of all possible triplets is 117×80 , but the dataset contains positive examples for only 600 triplets. The training set of HICO-Det contains 38,118 images and 117,871 HOI annotations for 600 HOI classes. The test set has 9,658 images and 33,405 HOI instances.

For evaluation, HICO-Det uses the mean average precision (mAP) metric. Here, an HOI detection is counted as a true positive if the minimum of the human overlap IOU and object overlap IOU with the ground truth is greater than 0.5. Following [3], HOI detection performance is reported for three different HOI category sets: (1) all 600 HOI categories (Full), (2) 138 categories with fewer than 10 training samples (Rare), and (3) the remaining 462 categories with more than 10 training samples (Non-rare).

V-COCO (Verbs in COCO) is a subset of MS-COCO [36], which consists of 10,346 images (2,533, 2,867, 4,946 for training, validation and test, respectively)

Table 1. Ablation study on the HICO-Det dataset. Our final model that includes both hierarchical architecture and distillation followed by post processing (Ours, ACP) shows the best performance among the baselines.

	Full	Rare	Non-rare
Baseline	17.56	13.23	18.85
Modified Baseline	19.09	13.09	20.89
+Hierarchical only	20.03	14.52	21.67
+Distillation only	19.98	13.67	21.86
+Hierarchical+Distillation	20.25	15.33	21.72
+Hierarchical+Distillation+Post (Ours, ACP)	20.59	15.92	21.98

Table 2. Performance of our models with different architectures for action prediction module. Our model ((D) Hierarchical) shows the best performance among the design choices.

	Full	Rare	Non-rare
(A) Modified Baseline	19.09	13.09	20.89
(B) MultiTask	19.54	13.93	21.22
(C) TwoStream	19.63	13.67	21.41
(D) Hierarchical	20.03	14.52	21.67

and 16,199 human instances. Each person is annotated with binary labels for 26 action classes. For the evaluation metric, same as for evaluation on HICO-Det, we use the AP score.

4.2 Quantitative Results

Ablation study In the ablations, the ‘No-Frills’ baseline network [18] we used is denoted as the *Baseline*. We first evaluate the effectiveness of the core design components in the proposed method including (1) our simple modification to *Baseline* in Sec. 3.3, denoted as *Modified Baseline*; (2) the hierarchical learning technique introduced in Sec. 3.3, denoted as *Hierarchical*; and (3) the knowledge distillation technique presented in Eq. (16) of Sec. 3.4, denoted as *Distillation*. Table 1 gives a comprehensive evaluation for each component. We draw conclusions from it one-by-one.

First, our baseline network is strong. Our *Modified Baseline* achieves 19.09 mAP and surpasses the ‘No-Frills’ *Baseline* by 1.51 mAP (a relative 8.7% improvement), which is already competitive to the state-of-the-art result [41] and serves as a strong baseline.

Second, both hierarchical learning and knowledge distillation are effective. This is concluded by adding *Hierarchical* and *Distillation* to the *Modified Baseline*, respectively. Specifically, *+Hierarchical* improves the modified baseline by 0.94 mAP (a relative 4.9% improvement), and *+Distillation* (training with Eq. (16)) improves the modified baseline by 0.89 mAP (a relative 4.7% improvement). Including both obtains 1.16 mAP improvement (relatively better by 6.1%).

Third, the proposed ACP method achieves a new state-of-the-art. Our final result is generated by further using the *PostProcess* step (introduced in Sec. 3.4) that projects the final action prediction into the ACP constrained space. Our method achieves 20.59 mAP (relative 7.9% improvement) for Full HOI categories, 15.92 mAP (relative 21.6% improvement) for Rare HOI categories, and 21.98 mAP (relative 5.2% improvement) for Non-rare HOI categories. Note that our method made especially significant improvements for Rare classes, which supports the claim that the proposed method can alleviate the long-tailed distribution problem of HOI detection datasets. This result sets the new state-

Table 3. Results on the HICO-Det dataset compared to the previous state-of-the-art methods. Our model shows favorable performance against the current state-of-the-art models on all the metrics.

	Full	Rare	Non-rare
Shen <i>et al.</i> [45]	6.46	4.24	7.12
HO-RCNN [3]	7.81	5.37	8.54
Gupta <i>et al.</i> [16] impl. by [13]	9.09	7.02	9.71
InteractNet [13]	9.94	7.16	10.77
GPNN [43]	13.11	9.34	14.23
iCAN [11]	14.84	10.45	16.15
iHOI [53]	13.39	9.51	14.55
with Knowledge [54]	14.70	13.26	15.13
Interactiveness Prior [33]	17.22	13.51	18.32
Contextual Attention [51]	16.24	11.16	17.75
No-Frills [18]	17.18	12.17	18.68
RPNN [66]	17.35	12.78	18.71
PMFNet [50]	17.46	15.65	18.00
Peyre <i>et al.</i> [41]	19.40	15.40	20.75
Our baseline	17.56	13.23	18.85
ACP (Ours)	20.59	15.92	21.98

Table 4. Results on the V-COCO dataset. For our method, we show results both for constructing the ACP from V-COCO and for using the ACP constructed from HICO-Det instead. Both of these models show favorable performance against the current state-of-the-art models.

	$AP_{r\text{ole}}$
Gupta <i>et al.</i> [16] impl. by [13]	31.8
InteractNet [13]	40.0
GPNN [43]	44.0
iCAN [11]	45.3
iHOI [53]	45.79
with Knowledge [54]	45.9
Interactiveness Prior [33]	48.7
Contextual Attention [51]	47.3
RPNN [66]	47.53
PMFNet [50]	52.0
Our baseline	48.91
ACP (Ours, V-COCO)	52.98
ACP (Ours, HICO-Det)	53.23

of-the-art on both the HICO-Det and V-COCO datasets as shown in Table 3 and Table 4.

In addition, the *MultiTask*, *TwoStream*, and our *Hierarchical* architecture are compared in Table 2. From *MultiTask*, it can be seen that the *softmax* based anchor action classification already brings benefits to the *Modified Baseline* when used only in a multi-task learning manner. From *TwoStream*, separately modeling the anchor and the regular classes leads to a slight more improvements compared to *MultiTask*. Moreover, our *Hierarchical* architecture improves upon *TwoStream* by explicitly modeling the hierarchy between anchor and regular action predictions.

Comparison with the state-of-the-art We compare our method with the previous state-of-the-art techniques in Table 3. Among the methods included in this comparison are the benchmark model of the HICO-Det dataset [3], the baseline model that we modified from [18], and the current published state-of-the-art method [41]. As shown in Table 3, our final model (Ours) shows significant improvements over our baseline model on all metrics, and shows favorable performance against the current state-of-the-art model in terms of all the metrics. In particular, our model surpasses the current state-of-the-art model [41] by 1.19 mAP.

Results on V-COCO dataset To show the generalizability of our method, we also evaluate our method on the V-COCO dataset. Note that the exact same method is directly applied to both HICO-Det and V-COCO, including the co-occurrence matrix, anchor action selection, and the architecture design. We also constructed a co-occurrence matrix from V-COCO, but the matrix was sparse. Thus, to better take advantage of our idea, we instead use the co-occurrence

Table 5. Results on the zero-shot triplets of the HICO-Det dataset. Our final model shows better performance than Peyre *et al.* by a large margin. Note that our ACP model under the zero-shot setting even outperforms the supervised setting of our baseline.

	mAP
Peyre <i>et al.</i> [41] Supervised	33.7
Peyre <i>et al.</i> [41] Zero-Shot	24.1
Peyre <i>et al.</i> [41] Zero-Shot with Aggregation	28.6
Ours Supervised (Modified Baseline)	33.27
Ours Zero-Shot (Modified Baseline)	20.34
Ours Zero-Shot (ACP)	34.95

matrix collected from the HICO-Det dataset to train on the V-COCO. Table 4 shows the performance of our model (Ours, HICO-Det) compared to the recent state-of-the-art HOI detectors on the V-COCO dataset. In addition, we show results of our model with the co-occurrence matrix constructed from the V-COCO dataset (Ours, V-COCO). Both of these models show favorable performance on the V-COCO dataset against the previous state-of-the-art model [50].

Results of the zero-shot setup on the HICO-Det dataset The *zero-shot setting* on the HICO-Det dataset is defined by Peyre *et al.* [41]. Specifically, we select a set of 25 HOI classes that we treat as unseen classes and exclude them and their labels in the training phase. However, we still let the model predict those 25 unseen classes in the test phase, which is known as the zero-shot problem. These HOI classes are randomly selected among the set of non-rare HOI classes. Since Peyre *et al.* did not provide which specific HOI classes they selected, we select the unseen HOI classes such that the performance (mAP) for these classes in our *Modified Baseline* model (introduced in Sec. 3.3) is similar to the corresponding *Supervised* baseline in [41]. In Table 5, we show results of our final model (ACP) and our modified baseline model compared to the corresponding setting reported in [41]. Our final model shows better performance (35.0 mAP) than Peyre *et al.* (28.6 mAP) by a large margin (relative 22.4% improvement). This result is remarkable in that our ACP model under the zero-shot setting even *outperforms* the supervised setting of our baseline model, indicating the power of the proposed ACP method to effectively leverage prior knowledge on action co-occurrences. Furthermore, the analogy transfer method proposed by Peyre *et al.* (denoted as aggregation) requires large-scale linguistic knowledge to train a word representation, whereas our model only requires the co-occurrence information of the labels in the dataset, which is much easier to obtain. We conclude that the proposed method is effective for the zero-shot problem while being easy to implement.

4.3 Additional Analysis

Score Improvement after ACP Projection We also show the HOI probability change from before to after applying the projection function $project(\cdot)$ on our model’s HOI prediction (*i.e.*, the effect of *PostProcess* introduced in Sec. 3.4) in Fig. 4. Leveraging co-occurrence matrix C can not only increase the

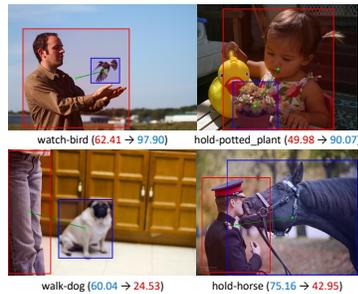


Fig. 4. The HOI probability before and after applying the projection function $project(\cdot)$ on our model’s HOI prediction (*PostProcess*). Note that *PostProcess* can be done without any optimization.

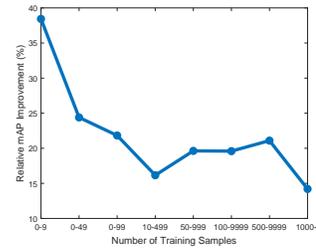


Fig. 5. The relative mAP score improvements for various HOI sets with the different number of training samples. Our method is able to improve the performance especially when the number of training samples is small (38.24% improvement for 0-9 samples).

score for true classes (top) but also reduce the score for false classes (bottom). Note that this change can be achieved *without any optimization* process.

Performance on various sets with different number of training samples Finally, in Fig. 5, we show the relative mAP score improvements of our model compared to the baseline model by computing mAP on various sets of HOI classes that have different number of training samples. Our method shows positive performance improvements for all numbers of training samples. Also, there is a trend that HOI classes with a small number of training samples mostly show larger performance improvements. In particular, for HOI classes with the number of training samples between 0 and 9, our model achieves 38.24% improvement compared to the baseline model. These results indicate that the proposed method is able to improve the performance of an HOI detector, especially for classes with few training samples.

5 Conclusion

We introduced a novel method to effectively train an HOI detector by leveraging prior knowledge on action co-occurrences in two different ways, via the architecture and via the loss function. Our proposed method consistently achieves favorable performance compared to the current state-of-the-art methods in various setups. Co-occurrence information not only is helpful for alleviating the long-tailed distribution problem but also can be easily obtained. A direction for future work is to construct and utilize co-occurrence priors for other relationship-based vision tasks [24,26,37].

Acknowledgements. This work was supported by the Institute for Information & Communications Technology Promotion (2017-0-01772) grant funded by the Korea government.

References

1. Bansal, A., Rambhatla, S.S., Shrivastava, A., Chellappa, R.: Detecting human-object interactions via functional generalization. In: AAAI Conference on Artificial Intelligence (AAAI) (2020)
2. Baradel, F., Neverova, N., Wolf, C., Mille, J., Mori, G.: Object level visual reasoning in videos. In: European Conference on Computer Vision (ECCV) (2018)
3. Chao, Y.W., Liu, Y., Liu, X., Zeng, H., Deng, J.: Learning to detect human-object interactions. In: IEEE Winter Conference on Applications of Computer Vision (WACV) (2018)
4. Chao, Y.W., Wang, Z., He, Y., Wang, J., Deng, J.: Hico: A benchmark for recognizing human-object interactions in images. In: IEEE International Conference on Computer Vision (ICCV) (2015)
5. Dai, B., Zhang, Y., Lin, D.: Detecting visual relationships with deep relational networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
6. Delaitre, V., Fouhey, D.F., Laptev, I., Sivic, J., Gupta, A., Efros, A.A.: Scene semantics from long-term observation of people. In: European Conference on Computer Vision (ECCV) (2012)
7. Delaitre, V., Laptev, I., Sivic, J.: Recognizing human actions in still images: a study of bag-of-features and part-based representations. In: British Machine Vision Conference (BMVC) (2010)
8. Delaitre, V., Sivic, J., Laptev, I.: Learning person-object interactions for action recognition in still images. In: Advances in Neural Information Processing Systems (NIPS) (2011)
9. Deng, J., Ding, N., Jia, Y., Frome, A., Murphy, K., Bengio, S., Li, Y., Neven, H., Adam, H.: Large-scale object classification using label relation graphs. In: European Conference on Computer Vision (ECCV) (2014)
10. Fu, J., Zheng, H., Mei, T.: Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
11. Gao, C., Zou, Y., Huang, J.B.: ican: Instance-centric attention network for human-object interaction detection. In: British Machine Vision Conference (BMVC) (2018)
12. Gibson, J.J.: The ecological approach to visual perception: classic edition. Psychology Press (2014)
13. Gkioxari, G., Girshick, R., Dollár, P., He, K.: Detecting and recognizing human-object interactions. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
14. Grabner, H., Gall, J., Van Gool, L.: What makes a chair a chair? In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2011)
15. Gupta, A., Davis, L.S.: Objects in action: An approach for combining action understanding and object perception. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2007)
16. Gupta, S., Malik, J.: Visual semantic role labeling. arXiv preprint arXiv:1505.04474 (2015)
17. Gupta, T., Schwing, A., Hoiem, D.: No-Frills Pytorch Github. <https://github.com/BigRedT/no.frills.hoi.det>
18. Gupta, T., Schwing, A., Hoiem, D.: No-frills human-object interaction detection: Factorization, layout encodings, and training techniques. In: IEEE International Conference on Computer Vision (ICCV) (2019)

19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
20. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
21. Horiguchi, S., Ikami, D., Aizawa, K.: Significance of softmax-based features in comparison to distance metric learning-based features. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2019)
22. Hu, Z., Ma, X., Liu, Z., Hovy, E., Xing, E.: Harnessing deep neural networks with logic rules. In: Annual Meeting of the Association for Computational Linguistics (ACL) (2016)
23. Hwang, S.J., Sha, F., Grauman, K.: Sharing features between objects and their attributes. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2011)
24. Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., Girshick, R.: Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
25. Kato, K., Li, Y., Gupta, A.: Compositional learning for human object interaction. In: European Conference on Computer Vision (ECCV) (2018)
26. Kim, D.J., Choi, J., Oh, T.H., Kweon, I.S.: Dense relational captioning: Triple-stream networks for relationship-based captioning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
27. Kim, D.J., Choi, J., Oh, T.H., Kweon, I.S.: Image captioning with very scarce supervised data: Adversarial semi-supervised learning approach. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP) (2019)
28. Kim, D.J., Choi, J., Oh, T.H., Yoon, Y., Kweon, I.S.: Disjoint multi-task learning between heterogeneous human-centric tasks. In: IEEE Winter Conference on Applications of Computer Vision (WACV) (2018)
29. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: International Conference on Learning Representations (ICLR) (2017)
30. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision (IJCV)* **123**(1), 32–73 (2017)
31. Li, Y., Ouyang, W., Wang, X.: VIP-CNN: A visual phrase reasoning convolutional neural network for visual relationship detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
32. Li, Y.L., Liu, X., Lu, H., Wang, S., Liu, J., Li, J., Lu, C.: Detailed 2d-3d joint representation for human-object interaction. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
33. Li, Y.L., Zhou, S., Huang, X., Xu, L., Ma, Z., Fang, H.S., Wang, Y., Lu, C.: Transferable interactiveness knowledge for human-object interaction detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
34. Li, Z., Hoiem, D.: Learning without forgetting. In: European Conference on Computer Vision (ECCV) (2016)
35. Liao, Y., Liu, S., Wang, F., Chen, Y., Feng, J.: Ppdm: Parallel point detection and matching for real-time human-object interaction detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)

36. 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 (ECCV) (2014)
37. Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. In: European Conference on Computer Vision (ECCV) (2016)
38. Marszalek, M., Schmid, C.: Semantic hierarchies for visual object recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2007)
39. Miller, G.A.: Wordnet: a lexical database for english. *Communications of the ACM* **38**(11), 39–41 (1995)
40. Pennington, J., Socher, R., Manning, C.: Glove: Global vectors for word representation. In: Conference on Empirical Methods in Natural Language Processing (EMNLP) (2014)
41. Peyre, J., Laptev, I., Schmid, C., Sivic, J.: Detecting unseen visual relations using analogies. In: IEEE International Conference on Computer Vision (ICCV) (2019)
42. Plummer, B.A., Mallya, A., Cervantes, C.M., Hockenmaier, J., Lazebnik, S.: Phrase localization and visual relationship detection with comprehensive linguistic cues. *IEEE International Conference on Computer Vision (ICCV)* (2017)
43. Qi, S., Wang, W., Jia, B., Shen, J., Zhu, S.C.: Learning human-object interactions by graph parsing neural networks. In: European Conference on Computer Vision (ECCV) (2018)
44. 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 (NIPS)* (2015)
45. Shen, L., Yeung, S., Hoffman, J., Mori, G., Fei-Fei, L.: Scaling human-object interaction recognition through zero-shot learning. In: *IEEE Winter Conference on Applications of Computer Vision (WACV)* (2018)
46. Stark, L., Bowyer, K.: Achieving generalized object recognition through reasoning about association of function to structure. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **13**(10), 1097–1104 (1991)
47. Sun, X., Li, C., Lin, S.: Explicit spatiotemporal joint relation learning for tracking human pose. In: *Proceedings of the IEEE International Conference on Computer Vision Workshops* (2019)
48. Sun, X., Wei, Y., Liang, S., Tang, X., Sun, J.: Cascaded hand pose regression. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
49. Ulutan, O., Iftekhhar, A., Manjunath, B.: Vsgnet: Spatial attention network for detecting human object interactions using graph convolutions. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
50. Wan, B., Zhou, D., Liu, Y., Li, R., He, X.: Pose-aware multi-level feature network for human object interaction detection. In: *IEEE International Conference on Computer Vision (ICCV)* (2019)
51. 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. In: *IEEE International Conference on Computer Vision (ICCV)* (2019)
52. Wang, T., Yang, T., Danelljan, M., Khan, F.S., Zhang, X., Sun, J.: Learning human-object interaction detection using interaction points. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
53. Xu, B., Li, J., Wong, Y., Zhao, Q., Kankanhalli, M.S.: Interact as you intend: Intention-driven human-object interaction detection. *IEEE Transactions on Multimedia* (2019)

54. Xu, B., Wong, Y., Li, J., Zhao, Q., Kankanhalli, M.S.: Learning to detect human-object interactions with knowledge. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
55. Yan, Z., Zhang, H., Piramuthu, R., Jagadeesh, V., DeCoste, D., Di, W., Yu, Y.: Hd-cnn: hierarchical deep convolutional neural networks for large scale visual recognition. In: IEEE International Conference on Computer Vision (ICCV) (2015)
56. Yang, X., Zhang, H., Cai, J.: Shuffle-then-assemble: learning object-agnostic visual relationship features. In: European Conference on Computer Vision (ECCV) (2018)
57. Yao, B., Fei-Fei, L.: Grouplet: A structured image representation for recognizing human and object interactions. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2010)
58. Yao, B., Fei-Fei, L.: Modeling mutual context of object and human pose in human-object interaction activities. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2010)
59. Yao, B., Jiang, X., Khosla, A., Lin, A.L., Guibas, L., Fei-Fei, L.: Human action recognition by learning bases of action attributes and parts. In: IEEE International Conference on Computer Vision (ICCV) (2011)
60. Yin, G., Sheng, L., Liu, B., Yu, N., Wang, X., Shao, J., Change Loy, C.: Zoom-net: Mining deep feature interactions for visual relationship recognition. In: European Conference on Computer Vision (ECCV) (2018)
61. Yu, R., Li, A., Morariu, V.I., Davis, L.S.: Visual relationship detection with internal and external linguistic knowledge distillation. In: IEEE International Conference on Computer Vision (ICCV) (2017)
62. Zhan, Y., Yu, J., Yu, T., Tao, D.: On exploring undetermined relationships for visual relationship detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
63. Zhang, H., Zhang, H., Wang, C., Xie, J.: Co-occurrent features in semantic segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
64. Zhang, H., Kyaw, Z., Chang, S.F., Chua, T.S.: Visual translation embedding network for visual relation detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
65. Zhang, J., Elhoseiny, M., Cohen, S., Chang, W., Elgammal, A.: Relationship proposal networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
66. Zhou, P., Chi, M.: Relation parsing neural network for human-object interaction detection. In: IEEE International Conference on Computer Vision (ICCV) (2019)
67. Zhuang, B., Liu, L., Shen, C., Reid, I.: Towards context-aware interaction recognition for visual relationship detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
68. Zhuang, B., Wu, Q., Shen, C., Reid, I., van den Hengel, A.: Hcvrd: a benchmark for large-scale human-centered visual relationship detection. In: AAAI Conference on Artificial Intelligence (AAAI) (2018)