

Spatially Conditioned Graphs for Detecting Human–Object Interactions

Frederic Z. Zhang^{1,3} Dylan Campbell^{2,3} Stephen Gould^{1,3}
¹The Australian National University ²University of Oxford
³Australian Centre for Robotic Vision

{firstname.lastname}@anu.edu.au dylan@robots.ox.ac.uk

<https://github.com/fredzhang/spatially-conditioned-graphs>


Abstract

We address the problem of detecting human–object interactions in images using graphical neural networks. Unlike conventional methods, where nodes send scaled but otherwise identical messages to each of their neighbours, we propose to condition messages between pairs of nodes on their spatial relationships, resulting in different messages going to neighbours of the same node. To this end, we explore various ways of applying spatial conditioning under a multi-branch structure. Through extensive experimentation we demonstrate the advantages of spatial conditioning for the computation of the adjacency structure, messages and the refined graph features. In particular, we empirically show that as the quality of the bounding boxes increases, their coarse appearance features contribute relatively less to the disambiguation of interactions compared to the spatial information. Our method achieves an mAP of 31.33% on HICO-DET and 54.2% on V-COCO, significantly outperforming state-of-the-art on fine-tuned detections.


1. Introduction

The task of detecting human–object interactions (HOIs) requires localising and describing pairs of interacting humans and objects. In particular, an HOI is defined as a (subject, predicate, object) triplet, following the definition of visual relations from Lu et al. [23], where the subject and object are typically represented as labelled bounding boxes. For HOI triplets, the subject is always a human, so the interactions of interest simplify to pairs of predicates and objects, e.g., *riding a horse* or *sitting on a bench*.


Since the output representations are inherently similar, HOI detection is most often approached as a downstream task of object detection. Given a set of object detections from an image, one may construct candidate human–object pairs by exhaustively matching between the detected human and object instances. Indeed, the vast majority of previous works [3, 6, 10, 17, 25, 24, 28, 5, 11] use an off-the-shelf



(a) An image with detected human and object instances



(b) Adjacency matrices computed with appearance features, normalised by rows (left) and columns (right)



(c) Adjacency matrices computed with spatial conditioning, normalised by rows (left) and columns (right)

Figure 1. Many images contain far more non-interactive human–object pairs than interactive ones (a). Correct inference of the interaction type and the correspondences requires a combination of appearance and spatial information. When using appearance features only, the adjacency matrix for a graphical neural network tends to be dominated by a few salient objects (b). Since messages from each node to its neighbours are identical apart from an adjacency scaling, this leads to the node features being dominated by those of the most salient objects, confusing the classifier. With spatial conditioning, the adjacency matrix is able to reflect the inherent interactive pairs without explicit supervision (c).

object detector [26] as a preprocessing stage. We take the same approach, leveraging the success of modern object detectors. While this converts the HOI detection task into the

Table 1. The use of appearance (A) and spatial (S) modalities at different stages of the graphical model, in recent HOI works. Refinement refers to late-stage fusion that takes place after message passing and fuses the graph features with other modalities.

Methods	Adjacency (early fusion)	Message (mid fusion)	Refinement (late fusion)
GPNN [25]	A	A	–
Wang et al. [11]	A, S	A	A, S
DRG [5]	S	S	–
VSGNet [28]	A	A	A, S
Ours	A, S	A, S	A, S

simpler HOI recognition task on a set of candidate human–object pairs, it is still far from being solved.

Recognising HOIs is extremely challenging. While image recognition discriminates between scene types [31] or prominent object types [27], focusing on the holistic understanding of an image, HOI recognition requires an understanding of the interactions between specific humans and objects at a much finer level. This requires reasoning about the subtle relationships between the instances as well as their contexts. This is particularly necessary when there are multiple human–object pairs with the same interaction type, where the model needs to correctly infer the interaction type and the correspondences between the individual instances. In addition, many interactions do not have strong visual cues and can be quite abstract, such as *buying an apple* or *inspecting a boat*. This poses a big challenge for standard CNNs, which excel at recognising physical qualities such as texture and shape. HOI detection demands a more sophisticated architecture capable of performing logical reasoning, not merely recognising the visual cues of the humans and objects of interest. The complexity and ambiguity of the problem is such that even humans can fail to correctly recognise HOIs in images, despite our ability to reason about visual cues and spatial relationships. Following prior work, we make use of graphical models to model these interrelationships and perform structured prediction.

Since humans and objects in an image play different roles in the interactions, we build a bipartite graph to characterise these interrelationships, wherein each human node is connected to each object node. As is intuitive, we use the appearance features of a detected instance as the node encoding, be it a person or an object. Edge encodings, however, have been under-explored in the HOI detection problem. Previous works [25, 11] take the appearance feature extracted from the minimum covering rectangle of the human and object boxes as the edge encoding. This representation does not necessarily encode the spatial relationships between a human–object pair, and there could be additional objects in the tight box other than the intended pair. Instead, we use explicitly learned spatial representations as the edge encodings. To shed some light on their significance, let us

consider the example shown in Figure 1a. Graphical models allow the propagation of contextual information between nodes. In this instance, each human node will receive information suggesting the presence of bikes. However, conventional algorithms send identical messages from a node to its neighbours, with the sole variable being a learnable weight that characterises the connectivity. And Figure 1b shows that this connectivity matrix fails to identify correct human–bike pairs with only appearance information, which causes confusion when distinguishing between all putative human–bike pairs. As such, we contend that it is crucial to incorporate spatial information to regulate the message passing procedure. Our intuition is that, with spatial conditioning, each human node receives information of the presence of a bike *and* its relative location. Therefore, the interaction *riding a bike* could potentially be suppressed for a human instance if all bikes in the image are, say, to its left, as opposed to being directly under it.

Our primary contribution is a spatially conditioned message passing algorithm that renders outgoing messages which are dependent on the receiving nodes. For our bipartite graph, the algorithm also passes anisotropic messages across the bipartition. Furthermore, we extend the spatial conditioning mechanism to other parts of the graph—the computation of the adjacency structure and the refinement of the graph features—through a proposed multi-branch fusion module. While previous works have also combined appearance and spatial modalities at these two stages of the network as shown in Table 1, our approach is consistent at each fusion stage and, in particular, gains significant performance improvements from using both modalities during message passing. Our secondary contribution is an analysis of the relative significance of the different modalities. We empirically show that as detection quality improves, the importance of the coarse appearance features decreases compared to that of the spatial information. We obtain state-of-the-art performance on the HICO-DET [3] and V-COCO [9] datasets, establishing a new benchmark for detecting human–object interactions.

2. Related Work

The HOI detection pipeline has significant overlap with that of object detection. Analogous to two-stage object detectors, a common approach is to first generate human–object pair proposals and then classify their interactions. Specifically, Faster R-CNN [26] has been used in many preceding works [3, 6, 10, 17, 25, 24, 28, 5, 11, 14, 16] to generate objects, each of which is associated with a predicted class and a confidence score. Afterwards, with appropriate filtering, human–object pairs are constructed exhaustively from the remaining detections. That is, each human instance will be paired up with each object instance. The rest of the pipeline varies, but typically employs a network with

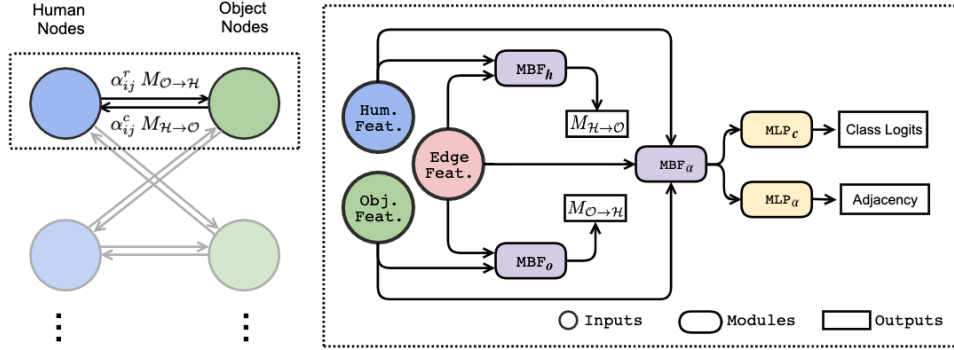


Figure 2. Diagram of proposed bipartite graph structure (left) and message passing algorithm (right). The graph structure and its connectivity is shown on the left, specifically highlighting the directed edges and anisotropic message passing. On the right, we zoom in on a particular pair of nodes and illustrate the computation of adjacency (Eq. 5), messages (Eq. 3, 4) and class logits. For better clarity, we intentionally leave the update functions out of the diagram and refer the readers to the equations (Eq. 1, 2).

multiple streams to exploit different modalities of information. For instance, Chao et al. [3] proposed a three-branch architecture to process the human and object appearance features and their pairwise spatial relationships. Different to many previous works, Liao et al. [18] presented a proposal-free HOI detection pipeline, where interactions are directly detected as keypoints. Such a keypoint represents the centre of the minimum covering rectangle for a human–object pair engaged in the predicted interaction. Positions of the human and object instances are obtained by regressing the displacements with respect to the detected interaction keypoint, similar to CornerNet [15], a keypoint-based object detector. Instead, we adopt the ubiquitous approach of using an off-the-shelf detector, due to their high performance and stability, and focus on improving the classification performance given a set of detections.

The choice of features has undergone significant development in recent research on HOI detection. Chao et al. [3] used RoIPool [7] to extract human and object appearance features and handcrafted a two-channel binary mask to encode the pairwise spatial relationships. While RoIAlign [12] is now used in preference to RoIPool for appearance feature extraction, the binary mask is still widely used [6, 5, 17, 28, 11]. However, Gupta et al. [10] argued that a handcrafted spatial feature is a more effective way to encode the spatial relationships, explicitly exposing the coordinates of the bounding box pairs, the intersection over union, the aspect ratios, etc. They and others [17, 32, 29] also proposed the use of human pose as additional information, which leads to some success in a few previous methods. We observe similar benefits to using handcrafted spatial encodings, but do not make use of human pose information in this work. Instead, we focus on showing how structured architectures can best exploit appearance and spatial information to disambiguate human–object interactions.

Graphical models were introduced to HOI detection by

Qi et al. [25]. They proposed a fully-connected graph with detected human and object instances as nodes. The node features are initialised with box appearance features and iteratively updated with a message passing algorithm. Wang et al. [11] argued that the graph should take into consideration the fact that there are two sets of heterogeneous nodes, that is, the human nodes and object nodes. Thus, message passing between homogeneous nodes (intra-class messages) should be modelled differently from that between heterogeneous nodes (inter-class messages). Gao et al. [5] also took advantage of the heterogeneity in nodes by constructing separate human-centric and object-centric graphs. They modelled human–object pairs as nodes, and employed the pairwise spatial relationships as node encodings. Last, Ulutan et al. [28] proposed a bipartite graph in addition to a visual branch, which makes use of the appearance features of human–object pairs and the global scene. Most of the previous methods use both appearance and spatial modalities in graphical models as shown in Table 1. However, the messages in all of their graphical models contain only one of the two modalities. Furthermore, the messages sent from a node to its neighbours are identical except weighted by adjacency values, which is what makes this work distinct.

3. Spatially Conditioned Graphs

To reason jointly about the appearance and spatial information of an image, we propose a graph neural network for detecting human–object interactions. The structure of the graph is shown in Figure 2. To obtain an initial set of detections $\{d_i\}_{i=1}^n$ for each image, we run an off-the-shelf object detector and apply appropriate filtering. We use Faster R-CNN [26], although our model is detector agnostic. The detections are given by the tuple $d_i = (\mathbf{b}_i, s_i, c_i)$, with bounding box coordinates $\mathbf{b}_i \in \mathbb{R}^4$, confidence score $s_i \in [0, 1]$ and predicted object class $c_i \in \mathcal{K}$, where \mathcal{K} is the set of object categories dependent on the dataset.

3.1. A Bipartite Graph Structure

We denote the bipartite graph as $\mathcal{G} = (\mathcal{H}, \mathcal{O}, \mathcal{E})$, where $\mathcal{H} = \{d_i \mid c_i = \text{“person”}\}$, $\mathcal{O} = \{d_i \mid c_i \neq \text{“person”}\}$, and \mathcal{E} is the set of edges, such that all vertices on one side of the bipartition are densely connected to those on the other. The node encodings are initialised with appearance features extracted using RoIAlign [12], and the edge encodings are computed as handcrafted feature vectors. We start by encoding the rudimentary spatial information: centre coordinates of the bounding boxes, widths, heights, aspect ratios and areas, all normalised by the corresponding dimension of the image. To characterise the pairwise relationships, we also include the intersection over union, the area of the human box normalised by that of the object box, and a directional encoding given by $[\text{ReLU}(d_x) \text{ReLU}(-d_x) \text{ReLU}(d_y) \text{ReLU}(-d_y)]$, where d_x and d_y are the differences between centre coordinates of the human and object boxes normalised by the dimensions of the human box. This gives us the pairwise spatial encoding $\mathbf{p} \in \mathbb{R}_+^{18}$. Following the practice of Gupta et al. [10], we concatenate the spatial encoding with its logarithm, allowing the network to learn second and higher order combinations of different terms. For numerical stability, a small constant $\epsilon > 0$ is added before taking the logarithm, which gives $\mathbf{p} \oplus \log(\mathbf{p} + \epsilon)$ as the pairwise spatial features.

To initialise the human and object nodes, the respective appearance features are mapped to a lower dimension with a multilayer perceptron (MLP) to get the node encodings $\mathbf{x}_i^0, \mathbf{y}_j^0 \in \mathbb{R}^n$ for indices $i \in \{1, \dots, |\mathcal{H}|\}$, $j \in \{1, \dots, |\mathcal{O}|\}$ and time step $t = 0$. Similarly, the edge encoding $\mathbf{z}_{ij} \in \mathbb{R}^n$ is obtained by mapping the pairwise spatial features to the same dimension using another MLP. The edge encodings are constant during message passing. We define our bidirectional message passing updates as

$$\mathbf{x}_i^{t+1} = \text{LN} \left(\mathbf{x}_i^t + \sigma \left(\sum_{j=1}^{|\mathcal{O}|} \alpha_{ij}^r M_{\mathcal{O} \rightarrow \mathcal{H}}(\mathbf{y}_j^t, \mathbf{z}_{ij}) \right) \right) \quad (1)$$

$$\mathbf{y}_j^{t+1} = \text{LN} \left(\mathbf{y}_j^t + \sigma \left(\sum_{i=1}^{|\mathcal{H}|} \alpha_{ij}^c M_{\mathcal{H} \rightarrow \mathcal{O}}(\mathbf{x}_i^t, \mathbf{z}_{ij}) \right) \right), \quad (2)$$

where LN denotes the LayerNorm operation [1], σ is the activation function (ReLU) and α is an adjacency weight between nodes. Notably, the message function M is, by design, anisotropic in that it has different parametrisations for different directions. This design allows nodes to send different message tailored to the type of receiving nodes.

3.2. Spatial Conditioning

Appearance and spatial features constitute the two most important sources of information in the disambiguation of complex interactions. However, in all previous works [25,

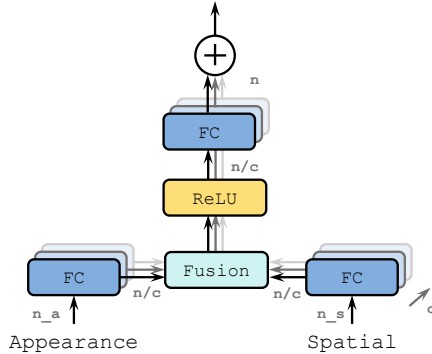


Figure 3. Structure of the multi-branch fusion module. The appearance and spatial features are mapped to c subspaces, fused and mapped to an intermediate representation size. The outputs of different branches are aggregated by taking the sum. The input and output dimensions of each FC layer are marked in the diagram.

[11, 5, 28], messages between nodes contain only one of the two modalities, and each node sends identical messages to its neighbours, modulo an adjacency scaling. We believe that this limits the representation power of the graphical model significantly. To this end, we propose to condition the messages between nodes on their spatial relationships, which allows messages to express the relative location of the human or object, not just their presence. To do so, we fuse the edge encoding and the node encoding (of the sender) by taking the elementwise product. We justify this design choice in the ablation analysis in Section 4.5.

We extend this strategy to two other parts of the graph. First, we apply spatial conditioning to compute the adjacency matrix. This allows the learned graph connectivity to also take into account the spatial relationship. As a result, it is able to infer the interactive human–object pairs as shown in Figure 1c. Second, we apply spatial conditioning to obtain the representations for human–object pairs. That is, after message passing is finished, we concatenate the graph features of each human–object pair, conditioned on their edge encoding. Our model therefore consistently applies spatial conditioning to compute the adjacency matrix, the messages, and the final pairwise features, which corresponds to early, mid and late fusion between the modalities.

3.3. Multi-Branch Fusion

To increase the expressive power of the spatial conditioning, we use a multi-branch structure for modality fusion. We map the modalities to c subspaces with reduced dimension, fuse the projections in each subspace, and then aggregate the outputs, as shown in Figure 3. We refer to the proposed module as *multi-branch fusion* (MBF). Following the nomenclature of Xie et al. [30], we refer to the number of homogeneous branches as the *cardinality*. Importantly, the number of parameters is independent of the cardinal-

ity by design, due to the subspace dimensionality reduction. We define the message functions as

$$M_{\mathcal{O} \rightarrow \mathcal{H}}(\mathbf{y}_j^t, \mathbf{z}_{ij}) = \text{MBF}_o(\mathbf{y}_j^t, \mathbf{z}_{ij}) \quad (3)$$

$$M_{\mathcal{H} \rightarrow \mathcal{O}}(\mathbf{x}_i^t, \mathbf{z}_{ij}) = \text{MBF}_h(\mathbf{x}_i^t, \mathbf{z}_{ij}). \quad (4)$$

The two fusion modules have independent weights, allowing for anisotropic messages.

MBFs are also used to compute the adjacency with spatial conditioning, with an additional linear layer to map the output to a scalar. The pre-normalised adjacency is

$$\tilde{\alpha}_k = \mathbf{w}_k^\top \sigma(\text{MBF}_\alpha(\mathbf{x}_i^t \oplus \mathbf{y}_j^t, \mathbf{z}_{ij})) + b_k \quad (5)$$

where $\mathbf{w}_k \in \mathbb{R}^n$, $b_k \in \mathbb{R}$ and k is a linear index corresponding to a pair of (i, j) , that is, $k \in \{1, \dots, |\mathcal{H} \times \mathcal{O}|\}$. During message passing, the adjacency value α_{ij}^r is obtained by applying softmax to the entries sharing the same index i (row normalisation). Similarly, α_{ij}^c is obtained via softmax while fixing j (column normalisation).

After all iterations of message passing, we fuse the spatial features and the graph features prior to binary classification for each target class. The computation of classification scores has the same form as that of the adjacency matrix in (Eq. 5), except with an additional sigmoid layer and that the output dimension is equal to the number of target classes. In fact, the adjacency can be interpreted as general interactiveness while the class probabilities are further conditioned on action types. For this reason, we use the same MBF module to compute the adjacency matrix and class probabilities.

3.4. Contextual Cues

As with most RoIPool-based feature extraction methods, the pooled information is local to a region. While this is reasonable for object detection, longer-range information about the context or even the global scene can be crucial for understanding human–object interactions. While Qi et al. [25] used appearance features extracted from the minimum covering rectangle of the human and object boxes as edge features, our model uses spatial information as edge features. To compensate for the loss of contextual cues, we employ another MBF module to fuse the global features and the spatial features for each pair as $\text{MBF}_g(\mathbf{g}, \mathbf{z}_{ij})$, where \mathbf{g} represents the global features. These features are concatenated with the spatially conditioned graph features to give $\text{MBF}_\alpha(\mathbf{x}_i^T \oplus \mathbf{y}_j^T, \mathbf{z}_{ij}) \oplus \text{MBF}_g(\mathbf{g}, \mathbf{z}_{ij})$ for classification.

3.5. Training and Inference

For each image during training, we append the ground-truth boxes to the set of detections and assign them a score of one. We then remove detected boxes below a threshold score and apply non-maximum suppression. The m highest scoring human and object boxes are then selected to initialise the bipartite graph. After message passing, we generate a set of human–object pairs from the graph, denoted

by $\{q_k\}_{k=1}^{|\mathcal{H} \times \mathcal{O}|}$, where $q_k = (\mathbf{b}_i^h, s_i^h, \mathbf{b}_j^o, s_j^o, \tilde{\mathbf{s}}_k)$. The bounding boxes \mathbf{b} and object detection scores s are obtained from the corresponding human and object nodes connected by edge (i, j) . The classification scores for all actions $\tilde{\mathbf{s}}_k$ are multiplied by the object detections scores. In practice, however, because the object detection scores do not consider the interactiveness of object instances, they tend to be overconfident. As a result, we raise the object detection scores to the power of λ during inference to counter this effect. The purpose of such operation is the same as the *Low-grade Instance Suppression* function [17]. But we found raising the power works better for our model. The final scores are computed as

$$\mathbf{s}_k = (s_i^h)^\lambda \cdot (s_j^o)^\lambda \tilde{\mathbf{s}}_k. \quad (6)$$

To associate the detected human–object pairs with the ground truth, the intersection-over-union is computed between each detected pair and ground-truth pair. Following previous practice [3], the IoU is computed for human and object boxes separately and taken as the minimum of the two. Detected pairs are considered to be positive when the IoU is above a designated threshold.

Due to the nature of proposal generation, there are overwhelmingly more negative examples than positive ones. In particular, the majority of examples are easy negatives. This inhibits the model from further improving on examples that are not well classified. To alleviate this issue, we adopt the focal loss [19] as a binary classification loss, given by

$$\text{FL}(\hat{y}, y) = \begin{cases} -\beta(1 - \hat{y})^\gamma \log(\hat{y}), & y = 1 \\ -(1 - \beta) \hat{y}^\gamma \log(1 - \hat{y}), & y = 0 \end{cases} \quad (7)$$

where $\hat{y} \in [0, 1]$ is the final score of an example for a certain class, $y \in \{0, 1\}$ is the binary label, and $\beta \in [0, 1]$ and $\gamma \in \mathbb{R}_+$ are hyper-parameters. In particular, β is a balancing factor between positive and negative examples. With $\beta > 0.5$, positive examples are assigned higher weights and vice versa. The parameter γ attenuates the loss incurred on well-classified examples. This prevents the large number of easy negatives from dominating the gradient. However, suppressing easy negatives reduces the focal loss’ magnitude [19], so normalisation is required. We extend Lin et al.’s [19] proposal to binary classification by normalising the loss by the number of positive logits.

It is also important to restrict the output space to meaningful interactions. Denote the set of actions by \mathcal{A} and the subset of valid actions for a specific object type $o \in \mathcal{K}$ by \mathcal{A}_o . Then the interactions of interest are in the set $\mathcal{I} = \cup_{o \in \mathcal{K}} \mathcal{A}_o \times \{o\}$, with $\mathcal{I} \subseteq \mathcal{A} \times \mathcal{K}$. Following the practice of Gupta et al. [10], we only compute the loss on the subset \mathcal{A}_o for each human–object pair, given the object type o . This removes predictions for non-existent interaction types, such as *eating a car*, allowing the network to dedicate its parameters to learning meaningful interactions.

In the HICO-DET dataset [3], interactions of interest include those between two humans (i.e., a human may be an object and a subject in an HOI triplet). To capture such interactions, we construct bipartite graphs such that object nodes subsume human nodes, that is, object nodes are identical to the set of all detections. Human nodes representing the same instance across the bipartition are initialised to be the same, yet will diverge as message passing proceeds.

4. Experiments

4.1. Dataset and Metric

We evaluated our model on the HICO-DET [3] and V-COCO [9] datasets. HICO-DET contains 37 633 training and 9 546 test images with bounding box annotations, 80 object classes (identical to those in the MS COCO dataset [21]), 117 action types, and 600 interaction types. There are 117 871 annotated human-object pairs in the training set and 33 405 in the test set. The distribution of pairs per interaction class is highly uneven, following a long tail distribution. In particular, there are 47 interaction categories with only one training example.

The evaluation metric is mean average precision (mAP). Detected human-object pairs are considered as positive when the IoU with any ground-truth pair is higher than 0.5. For multiple detected pairs associated with the same ground-truth instance, only the highest scoring pair is considered as positive. The computation of mAP follows the 11-point interpolation algorithm used in the Pascal VOC challenge [4]. To capture the effectiveness of our model across interactions with different numbers of annotations, we follow previous practice [3] and report results in three categories: full (all 600 interactions), rare (138 interactions with fewer than 10 training examples), and non-rare (462 interactions with 10 or more training examples).

V-COCO is a much smaller dataset with 2 533 images in the training set, 2 867 in the validation set and 4 946 in the test set. The dataset contains 26 different actions. We report our performance on this dataset for legacy reasons.

4.2. Implementation Details

We use Faster R-CNN [26] with ResNet50-FPN [13, 20] pretrained on MS COCO [21] to generate detections. For each image, we first filter out detections with scores lower than 0.2 and perform non-maximum suppression (NMS) with a threshold of 0.5. Afterwards, we extract the $m = 15$ highest scoring human boxes, and the $m = 15$ highest scoring object boxes. This gives us at most $15(30 - 1) = 435$ box pairs, after removing pairs involving the same person twice. Inference follows the same setup, except that ground-truth detections are not used.

We use ResNet50-FPN [13, 20] as the backbone for feature extraction. To utilise the feature pyramid, boxes are as-

signed to different pyramid levels based on their sizes [20]. The pooled box features are mapped to 1024-dimensional vectors with a two-layer MLP. Similarly, the spatial features are mapped to the same dimension (1024) with a three-layer MLP. For the MBF module, we use $c = 16$ and $n = 1024$. We use $T = 2$ iterations of message passing for all models unless otherwise specified. To counter the over-confidence in object scores, we set $\lambda = 2.8$ during inference while keeping $\lambda = 1.0$ during training. Lastly, for the focal loss, we set $\beta = 0.5$ and $\gamma = 0.2$. All hyper-parameters are selected using cross-validation.

We adopt an image-centric training strategy [7] with slight modifications. Input images are normalised and resized such that the shorter edge is 800 pixels. Bounding boxes are then resized accordingly. Afterwards, images are batched with zero padding. To train the model, we use AdamW [22] as the optimiser, with a momentum of 0.9 and weight decay of 10^{-4} . We use an initial learning rate of 10^{-5} for the backbone and 10^{-4} for the rest of the network. The learning rates are dropped by a magnitude at the sixth epoch. All models are trained for 10 epochs on 8 GeForce GTX TITAN X devices, with an effective batch size of 32.

4.3. Comparison with State-of-the-Art

Quantitative results on the HICO-DET [3] test set are shown in Table 2. We report the performance of our model with three different detectors: one pre-trained on the MS COCO dataset [21], one fine-tuned on the HICO-DET dataset as provided by Gao et al. [5], and an oracle supplying the ground truth detections. We achieve competitive performance when using the COCO pre-trained detector, but significantly outperform state-of-the-art when using the higher-quality fine-tuned detections, a 20% relative improvement. In particular, we outperform the next best method IDN [16] by 5 mAP, despite slightly underperforming that method when using the pre-trained detections. This suggests that our graph neural network can better exploit the high-quality detections. This is supported by the results for the oracle detector, where we outperform the next best method by 7.5 mAP. We show an example of the different detector outputs in Figure 4. Less salient people and objects are suppressed in the fine-tuned detector, making the spatial information more discriminative. We also report on V-COCO [9] test set, as shown in Table 3. Our model achieves competitive performance using a pre-trained detector and receives consistent gains from a fine-tuned detector.

4.4. Contribution of Different Modalities

Notably, our model is able to gain nearly 9 mAP by using a fine-tuned detector and a further 20 mAP by using an oracle detector that supplies ground-truth detections, which is much higher than what previous methods gain. Due to the use of spatial conditioning in our model, we hypothesise

Table 2. HOI detection performance (mAP×100) on the HICO-DET [3] test set under the default setting. See appendix for the known object setting. The most competitive method in each category is in bold, while the second best is underlined.

Method	Backbone	Full	Rare	Non-rare
DETECTOR PRE-TRAINED ON MS COCO				
HO-RCNN [3]	CaffeNet	7.81	5.37	8.54
InteractNet [8]	ResNet-50-FPN	9.94	7.16	10.77
GPNN [25]	ResNet-101	13.11	9.34	14.23
iCAN [6]	ResNet-50	14.84	10.45	16.15
Bansal et al. [2]	ResNet-101	16.96	11.73	18.52
TIN [17]	ResNet-50	17.03	13.42	18.11
Gupta et al. [10]	ResNet-152	17.18	12.17	18.68
RPNN [32]	ResNet-50	17.35	12.78	18.71
Wang et al. [11]	ResNet-50-FPN	17.57	16.85	17.78
DRG [5]	ResNet-50-FPN	19.26	17.74	19.71
Peyre et al. [24]	ResNet-50-FPN	19.40	14.63	20.87
VCL [14]	ResNet50	19.43	16.55	20.29
VSGNet [28]	ResNet-152	19.80	16.05	20.91
IDN [16]	ResNet50	23.36	22.47	23.63
Ours	ResNet-50-FPN	<u>21.85</u>	<u>18.11</u>	<u>22.97</u>
DETECTOR FINE-TUNED ON HICO-DET				
PPDM [18]	Hourglass-104	21.73	13.78	24.10
Bansal et al. [2]	ResNet-101	21.96	16.43	23.63
VCL [14]	ResNet50	23.63	17.21	25.55
DRG [5]	ResNet-50-FPN	24.53	19.47	26.04
IDN [16]	ResNet50	<u>26.29</u>	<u>22.61</u>	<u>27.39</u>
Ours	ResNet-50-FPN	31.33	24.72	33.31
ORACLE DETECTOR				
iCAN [6]	ResNet-50	33.38	21.43	36.95
TIN [17]	ResNet50	34.26	22.90	37.65
Peyre et al. [24]	ResNet-50-FPN	34.35	27.57	36.38
IDN [16]	ResNet50	<u>43.98</u>	<u>40.27</u>	<u>45.09</u>
Ours	ResNet-50-FPN	51.53	41.01	54.67

Table 3. Performance (mAP×100) on the V-COCO [9] test set. The most competitive method in each category is in bold, while the second best is underlined. *Using a fine-tuned detector.

Method	Backbone	Scenario 1	Scenario 2
InteractNet [8]	ResNet-50-FPN	40.0	–
GPNN [25]	ResNet-101	44.0	–
iCAN [6]	ResNet-50	45.3	52.4
TIN [17]	ResNet-50	47.8	54.2
DRG [5]	ResNet-50-FPN	51.0	–
VSGNet [28]	ResNet-152	51.8	57.0
Wang et al. [11]	ResNet-50-FPN	52.7	–
IDN [16]	ResNet50	53.3	60.3
Ours	ResNet-50-FPN	<u>53.0</u>	<u>58.2</u>
Ours*	ResNet-50-FPN	54.2	60.9

that as detection quality improves, spatial information plays a more significant role in the disambiguation of interactions, while coarse appearance features contribute relatively less. This is supported by evidence in Table 4, where we show




Figure 4. Object detections from the pre-trained MS COCO model (left) compared to the fine-tuned HICO-DET model (right). Boxes with scores higher than 0.5 are displayed. The fine-tuned detector suppresses objects that are less likely to be engaged in interactions.

Table 4. Difference in performance between models with appearance and spatial features (Ours) and with only appearance features (baseline), as detection quality increases to the right.

Detector	COCO	HICO-DET	Oracle
Performance Δ	+1.93	+2.90	+4.36

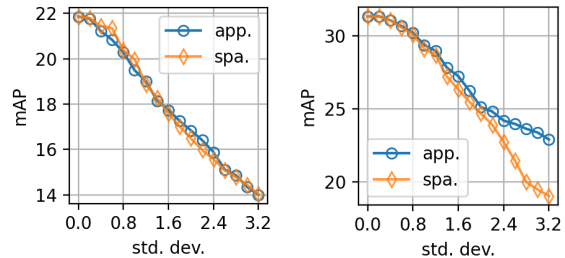


Figure 5. Model performance under different levels of corruption in appearance and spatial modalities, using the pre-trained detector (left) and the fine-tuned detector (right).

that the performance difference between the baseline model and our full model increases as detection quality improves. To investigate this hypothesis, we add Gaussian noise with zero mean and variable standard deviation to the appearance and spatial features separately, and observe how corruption in different modalities damages the performance. As shown in Figure 5, when using the pre-trained detector, noise in the appearance and spatial features has an approximately equal effect on performance. However, with the fine-tuned detector, noisy spatial features have a much larger impact. We conclude that spatial information contributes relatively more to performance as the detection quality improves.

4.5. Ablation Studies

We conducted a series of ablation studies to validate our design choices. Our baseline is a bipartite graph with appearance features only. Specifically, the message sent from a node is computed from its appearance encodings using a linear layer. The adjacency and class probabilities are computed from the concatenated node encodings of a human-object pair using an MLP, and the computation of classification scores shares weights with that of adjacency until the logistic layer. We first investigate the importance of spatial conditioning at different stages in our model: for comput-

Table 5. Ablating the addition of spatial conditioning at different stages of the model on the HICO-DET dataset (mAP \times 100).

Stage	COCO Detector	HICO-DET Detector
None	19.92	28.43
Adjacency	20.56	29.48
Messages	20.79	30.06
Global features	20.44	29.51
Refinement	21.03	30.11
All (Ours)	21.85	31.33

Table 6. Ablating the multi-branch fusion design choices, including the binary operation and the cardinality (c).

Design Choice	COCO Detector	HICO-DET Detector
Product ($c = 1$)	21.18	30.75
Sum ($c = 1$)	21.35	30.87
Concat. ($c = 1$)	21.02	30.66
Product ($c = 16$)	21.85	31.33
Sum ($c = 16$)	21.81	31.07
Concat. ($c = 16$)	21.67	31.65


Table 7. Varying the number of message passing iterations (T).

Model	COCO Detector	HICO-DET Detector
Ours ($T = 0$)	20.05	28.86
Ours ($T = 1$)	20.70	30.99
Ours ($T = 2$)	21.85	31.33
Ours ($T = 3$)	21.72	31.78


ing the adjacencies, messages, global features, and refined graph features. As shown in Table 5, every stage improves over the baseline, and they combine together to achieve the best performance. We next demonstrate the impact of different design choices for multi-branch fusion, including the choice of fusion methods and the number of branches (cardinality). As shown in Table 6, the performance improves with higher cardinality. We also show that the performance is insensitive to the choice of binary fusion operation, with our choice (elementwise product) being comparable to the elementwise sum and concatenation operations. Lastly, we show how the number of message passing iterations at test time affects the results. As shown in Table 7, message passing is clearly helpful for this problem, while an additional iteration further improves the results significantly.

4.6. Qualitative Results

We show qualitative results of our model in Figure 6. In Figure 6a, the ground-truth interaction is *riding a bike*. As shown in Table 8, positive human–bike pairs (1, 2), (3, 4) and (5, 6) have the highest scores. However, the network also assigns the negative pair (2, 3) a relatively high score. This is due to the spatial proximity and visual similarity between bike (3) and the correct bike instance (1), which can confuse our model. Another example is given in Figure 6b,



(a) Interaction: *riding a bike*



(b) Interaction: *sitting on a bench*

Figure 6. Qualitative results with success and failure cases for our model. The scores corresponding to (a) are in Table 8, and the scores for (b) are in Table 9.

where the true interaction is *sitting on a bench*. As shown in Table 9, our model assigns high scores to all correct human–bench pairs and suppresses the non-interactive pair (1, 5).

Table 8. Scores for the interaction *riding a bike* in Figure 6a.

Instance index	2	4	6
1	0.5742	0.0027	0.0000
3	0.4617	0.4735	0.0002
5	0.0006	0.0008	0.7899

Table 9. Scores for the interaction *sitting on a bench* in Figure 6b.

Instance index	1	3	4	6
5	0.0011	0.3890	0.4049	0.6056

5. Conclusion

In this paper, we have proposed a spatially conditioned graph neural network for detecting human–object interactions. To perform spatial conditioning, we applied a multi-branch fusion mechanism that modulates the appearance features with the spatial configuration of the human–object pairs. We use this mechanism consistently for computing adjacency, messages and refined graph features, and show that our model outperforms the state-of-the-art by a considerable margin with fine-tuned detections. We also show that the margin of improvement increases with the detection quality, allowing our model to most effectively exploit advances in object detector research.

Acknowledgements

This research is funded in part by the ARC Centre of Excellence for Robotic Vision (CE140100016) and Continental AG (D.C.).

References

- [1] Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *Adv. Neural Inform. Process. Syst.*, 2016. 4
- [2] Ankan Bansal, Sai Saketh Rambhatla, Abhinav Shrivastava, and Rama Chellappa. Detecting human-object interactions via functional generalization. *AAAI*, 2020. 7
- [3] Yu-Wei Chao, Yunfan Liu, Xieyang Liu, Huayi Zeng, and Jia Deng. Learning to detect human-object interactions. *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, 2018. 1, 2, 3, 5, 6, 7
- [4] Mark Everingham, S. M. Ali Eslami, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The Pascal visual object classes challenge: A retrospective. *Int. J. Comput. Vis.*, 111(1):98–136, 2014. 6
- [5] Chen Gao, Jiarui Xu, Yuliang Zou, and Jia-Bin Huang. DRG: Dual relation graph for human-object interaction detection. *Eur. Conf. Comput. Vis.*, 2020. 1, 2, 3, 4, 6, 7
- [6] Chen Gao, Yuliang Zou, and Jia-Bin Huang. iCAN: Instance-centric attention network for human-object interaction detection. *Brit. Mach. Vis. Conf.*, 2018. 1, 2, 3, 7
- [7] Ross Girshick. Fast R-CNN. *Int. Conf. Comput. Vis.*, (9):1440–1448, 2015. 3, 6
- [8] Georgia Gkioxari, Ross Girshick, Piotr Dollár, and Kaiming He. Detecting and recognizing human-object interactions. *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 7
- [9] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015. 2, 6, 7
- [10] Tanmay Gupta, Alexander Schwing, and Derek Hoiem. No-frills human-object interaction detection: Factorization, layout encodings, and training techniques. *Int. Conf. Comput. Vis.*, 2019. 1, 2, 3, 4, 5, 7
- [11] Wang Hai, Zheng Weishi, and Yingbiao Ling. Contextual heterogeneous graph network for human-object interaction detection. *Eur. Conf. Comput. Vis.*, 2020. 1, 2, 3, 4, 7
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. *Int. Conf. Comput. Vis.*, pages 2980–2988, 2017. 3, 4
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. 6
- [14] Zhi Hou, Xiaojiang Peng, Yu Qiao, and Dacheng Tao. Visual compositional learning for human-object interaction detection. In *Eur. Conf. Comput. Vis.*, 2020. 2, 7
- [15] Hei Law and Jia Deng. CornerNet: Detecting objects as paired keypoints. *Eur. Conf. Comput. Vis.*, 2018. 3
- [16] Yong-Lu Li, Xinpeng Liu, Xiaoqian Wu, Yizhuo Li, and Cewu Lu. Hoi analysis: Integrating and decomposing human-object interaction. In *Adv. Neural Inform. Process. Syst.*, 2020. 2, 6, 7
- [17] Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yanfeng Wang, and Cewu Lu. Transferable interactiveness knowledge for human-object interaction detection. *Int. Conf. Comput. Vis.*, 2019. 1, 2, 3, 5, 7
- [18] Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, and Jiashi Feng. PPDM: Parallel point detection and matching for real-time human-object interaction detection. *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 3, 7
- [19] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *Int. Conf. Comput. Vis.*, 2017. 5
- [20] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017. 6
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. *Eur. Conf. Comput. Vis.*, 2014. 6
- [22] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *Int. Conf. Learn. Represent.*, 2018. 6
- [23] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. *Eur. Conf. Comput. Vis.*, 2016. 1
- [24] Julia Peyre, Ivan Laptev, Cordelia Schmid, and Josef Sivic. Detecting unseen visual relations using analogies. *Int. Conf. Comput. Vis.*, 2019. 1, 2, 7
- [25] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu. Learning human-object interactions by graph parsing neural networks. *Eur. Conf. Comput. Vis.*, 2018. 1, 2, 3, 4, 5, 7
- [26] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *Adv. Neural Inform. Process. Syst.*, pages 91–99, 2015. 1, 2, 3, 6
- [27] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, 2015. 2
- [28] Oytun Ulutan, A S M Iftexhar, and B. S. Manjunath. VS-Net: Spatial attention network for detecting human object interactions using graph convolutions. *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 1, 2, 3, 4, 7
- [29] Bo Wan, Desen Zhou, Yongfei Liu, Rongjie Li, and Xuming He. Pose-aware multi-level feature network for human object interaction detection. *Int. Conf. Comput. Vis.*, 2019. 3
- [30] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5987–5995, 2017. 4
- [31] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2017. 2
- [32] Penghao Zhou and Mingmin Chi. Relation parsing neural network for human-object interaction detection. *Int. Conf. Comput. Vis.*, 2019. 3, 7