

Visual Relationship Detection Using Part-and-Sum Transformers with Composite Queries

Qi Dong Zhuowen Tu Haofu Liao Yuting Zhang Vijay Mahadevan Stefano Soatto
Amazon Web Services

{qdon, ztu, liahaofu, yutingzh, vmahad, soattos}@amazon.com

Abstract

Computer vision applications such as visual relationship detection and human object interaction can be formulated as a composite (structured) set detection problem in which both the parts (subject, object, and predicate) and the sum (triplet as a whole) are to be detected in a hierarchical fashion. In this paper, we present a new approach, denoted Part-and-Sum detection Transformer (PST), to perform end-to-end visual composite set detection. Different from existing Transformers in which queries are at a single level, we simultaneously model the joint part and sum hypotheses/interactions with composite queries and attention modules. We explicitly incorporate sum queries to enable better modeling of the part-and-sum relations that are absent in the standard Transformers. Our approach also uses novel tensor-based part queries and vector-based sum queries, and models their joint interaction. We report experiments on two vision tasks, visual relationship detection and human object interaction and demonstrate that PST achieves state of the art results among single-stage models, while nearly matching the results of custom designed two-stage models.

1. Introduction

In this paper, we study problems such as *visual relationship detection* (VRD) [29, 21] and *human object interaction* (HOI) [11, 35, 4] where a composite set of a two-level (part-and-sum) hierarchy is to be detected and localized in an image. In both VRD and HOI, the output consists of a set of entities. Each entity, referred to as a “sum”, represents a triplet structure composed of parts: the parts are (*subject, object, predicate*) in VRD and (*human, interaction, object*) in HOI. The sum-and-parts structure naturally forms a two-level hierarchical output - with the sum at the root level and the parts at the leaf level. In the general setting for composite set detection, the hierarchy consists of two levels, but the number of parts can be arbitrary.

Many existing approaches in VRD [29, 16, 30, 54, 54]

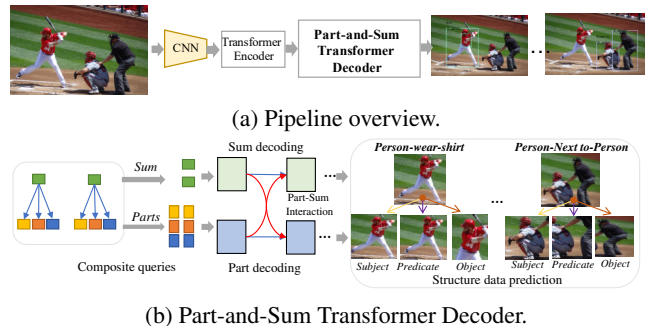


Figure 1: Overview of visual composite set detection by Part-and-Sum Transformer (PST).

and HOI [36, 4, 24, 14, 10] are based on two stage processes in which some parts (e.g. the subject and the object in VRD) are detected first, followed by detection of the association (sum) and the additional part (the predicate). Single stage approaches [50, 49, 15, 19] with end-to-end learning for VRD and HOI also exist. In practice, two-stage approaches produce better performance while single-stage methods are easier to train and use.

The task for an object detector is to detect and localize all valid objects in an input image, making the output a set. Though object detectors such as FasterRCNN [33] are considered end-to-end trainable, they perform instance level predictions, and require post-processing using non-maximum suppression to recover the entire set of objects in an image. Recent developments in Transformers [44] and their extensions to object detection [3] enable set-level end-to-end learning by eliminating anchor proposals and non-maximum suppression.

In this paper, we formulate the visual relationship detection (VRD) [29, 21] and human object interaction (HOI) [11, 35, 4] as composite set (two-level hierarchy) detection problems and propose a new approach, part-and-sum Transformers (PST) to solve them. PST is different from existing detection transformers where each object is represented by a vector-based query in either a one-level set [3]. We show the importance of establishing an explicit “sum” representation for the triplet as a whole to be simultaneously

ture/relationships. For accurate modeling composite data, we propose composite query based Part-and-Sum transformer decoder, to learn each composite/structure data on both entity and relationship levels, as shown in Figure 1 (b). In the following sections, we detail the PST model and the corresponding processes of training and inference.

3.2. Part-and-Sum Transformer (PST)

To construct the PST model, we first describe the vector based query, tensor based query, and composite queries that are used for composite set prediction. We then formulate the composite transformer decoder layer based on the composite queries.

3.2.1 Query Design for Composite Data

Vector based query A standard decoder used in DETR takes a set of vector based queries as input, as shown in Figure 2 (a). When applying this formulation to the relationship detection task, one can use feedforward networks (FFNs) to directly predict a subject-predicate-object triplet from the output of each query. This straightforward extension of DETR, while serving as a reasonable baseline, is sub-optimal as each query mixes the parts and their interaction altogether inside a vector. This makes the parts and their interaction implicitly modeled, limiting the expressiveness and representation capability of the visual relationship model.

Tensor based query To explicitly model the parts and their relationships (e.g. Subject, Predicate, and Object), we propose a tensor-based query representation using disjoint sub-vectors as sub-queries. Specifically for VRD, in a tensor based query representation, three sub-queries represent Subject, Predicate, and Object. All queries together form a $M \times P \times D$ matrix, where M is the number of queries, P is the number of entities in a relationship ($P = 3$), and D is the feature dimension of sub-queries. This formulation enables part-wise decoding in the transformer decoder, as shown in Figure 2 (b). Technically, the vector based query represents each relationship as a whole/Sum, whereas tensor based query models the parts disjointly. The difference in the query design leads to a difference in the learned contexts: self-attention layers among vector based queries mine inter-relation context, while self-attention layers among part queries mine the inter-part context.

Composite query On the one hand, vector based query is able to capture the relationship as a sum/whole, but there exists an intrinsic ambiguity in the parts. On the other hand, tensor based query models each part explicitly, but it lacks the knowledge of the relationships as a sum, which is important for the subject-object association. Based on the observation above, we propose a *composite* query representation. Formally, each composite query Q_i is composed of

part queries Q_i^P (a tensor query) as well as a sum/whole query q_i^G (a vector based query). In VRD, each composite query Q_i is composed of three sub-queries to represent Subject, Predicate, and Object; and one sum query to represent the relationship. $Q_i = \{Q_i^P, q_i^G\}$, and $Q_i^P = \{q_i^s, q_i^p, q_i^o\}$, where q_i^s , q_i^p and q_i^o denote subject, predicate and object sub-queries. Assuming M composite queries in the decoder, the overall query is a $M \times D \times 4$ tensor, where D is the dimension of sub-queries.

3.2.2 Part-and-Sum Transformer Decoder

As the composite query includes both part and sum queries, we separately decode the part queries Q^P and the sum query q^G . To enable mutual benefits of part and sum learning, we also set up part-sum interaction. Additionally, we propose a factorized self-attention layer for further enhancing part level learning. The architecture of Part-and-Sum Transformer Decoder is illustrated in Figure 2 (c).

Part-and-Sum separate decoding. The PST decoder has a two-stream architecture, for part and sum queries decoding, respectively. Each decoding stream contains a self-attention layers (SA), cross-attention layers (CA), and feed-forward neural networks (FFN). Let f and φ denote respectively SA and CA layers for part queries. Decoding the part queries is written as:

$$\begin{aligned} f_{\text{Part}}(Q^P) &= \text{SA}(Q_1^P, \dots, Q_M^P) \\ &= \text{SA}(q_1^s, q_1^p, q_1^o, \dots, q_M^s, q_M^p, q_M^o) \quad (1) \\ \varphi_{\text{Part}}(Q^P, I) &= \text{CA}([Q_1^P, \dots, Q_M^P], I), \end{aligned}$$

where I denotes the tokenized image features from the Transformer Encoder. Similarly, decoding the sum queries can be written as:

$$\begin{aligned} f_{\text{Sum}}(Q^G) &= \text{SA}(q_1^G, \dots, q_M^G) \\ \varphi_{\text{Sum}}(Q^G, I) &= \text{CA}([q_1^G, \dots, q_M^G], I) \quad (2) \end{aligned}$$

Each entity has both part and global embeddings via two separate sequential modules $\text{FFN}(\varphi(f(Q), I))$. The self-attention exploits the context among all queries. Part-and-Sum separated decoding effectively models two different types of contexts: the self-attention in part queries explores the inter-component context, for example, when one part query predicts “person”, it reinforces related predicates such as “eat” and “hold”; while self-attention for global queries exploits the inter-relationship context, for example, a sum query that predicts “Person read book”, is a clue to infer “person sit” relationship. These contexts provide the interactions needed for the accurate inference of the structured output.

Factorized self-attention layer. To make the interactions within a group-wise part query more structured, we design a factorized self-attention layer, as shown in Figure 2

(b). Instead of doing self-attention among all part queries as in Eq. 1, a factorized self-attention layer first conducts intra-relation self-attention, and then conducts inter-relation self-attention. The intra-relation self attention layer leverages the contexts of the parts to benefit relationship prediction, for example, subject query and object query are “person” and “horse” helps predict predicate “Ride”. The inter-relation self-attention layer leverages the inter-relation context, to enhance the holistic relation prediction per image, which is particularly important for multiple interactions detection for the same subject entity. More details are in the supplementary materials.

Part-Sum interaction. Part query decoding embeds more accurate component information, while global embedding contains more accurate component association. These two aspects are both important to the structured output detection, and are mutually beneficial to each other [8]. Thus, we design interaction between the two decoding streams, enabling part-sum conditions. Specifically, after FFNs in the decoder, for each part embedding $q_i^k, k \in \{s, o, p\}$, we combine it with the sum query embedding, while for each sum query q_i^G , we fuse all three part query embeddings. The part-sum interaction is formulated as:

$$\begin{aligned} q_i^k &= \mathcal{N}(q_i^k + q_i^G), k \in \{s, o, p\} \\ q_i^G &= \mathcal{N}(q_i^G + \sum_{k \in \{s, o, p\}} q_i^k), \end{aligned} \quad (3)$$

where \mathcal{N} is layer normalisation [1].

3.3. Model Training and Inference

Composite prediction. For each composite query $Q_i = \{q_i^G, q_i^s, q_i^o, q_i^p\}$, we predict the classes of subject, object and predicate; and the bounding boxes for subject and object. Specifically, for each part query, we predict corresponding classes by using a one-layer linear layer, and predict boxes using a shallow MLP head. Besides, we can also construct the global representation from group-wise part queries, by concatenating all part queries, and denote this as $q_i^{\text{spo}} = [q_i^s, q_i^o, q_i^p]$. The part query predication is:

$$\begin{aligned} \hat{b}_i^k &= f_{\text{box}}^k(q_i^k), k \in \{s, o\} \\ \hat{p}_i^k &= f_{\text{cls}}^k(q_i^k), k \in \{s, o, p\} \\ \hat{p}_i^{\text{spo}} &= f_{\text{cls}}^{\text{spo}}(q_i^{\text{spo}}) \end{aligned} \quad (4)$$

where f_{cls}^* are the FFNs for subject, object and predicate classification; and f_{box}^* are the FFNs for predicting the boxes of subject and object; $f_{\text{cls}}^{\text{spo}}$ is an FFN for predicting relation triplet.

For Sum query prediction, we predict classes and boxes of all parts from a Sum query q_i^G , that is:

$$\begin{aligned} \hat{b}_i^G &= g_{\text{box}}^k(q_i^G), k \in \{s, o\} \\ \hat{p}_i^G &= g_{\text{cls}}^k(q_i^G), k \in \{s, o, p\} \end{aligned} \quad (5)$$

where g_{cls}^* are FFNs for subject, object and predicate classification; and g_{box}^* are the FFNs for predicting the boxes of subject and object in the global level. Note that the last layer in f_{cls} is a Softmax layer, while the last layer in f_{box} is a Sigmoid layer.

Composite bipartite matching. We conduct composite group-wise bipartite matching, i.e. considering all components belonging to a relation jointly in computing set-to-set similarity. Specifically, for a relation, there are three Part queries (subject, object and predicate), and a triplet embedding. The bipartite matching algorithm finds a permutation σ of the M predictions $\{\hat{y}_i\}_{i=1}^M$ so that the total matching cost is minimized

$$\hat{\sigma} = \arg \min_{\sigma \in \mathcal{P}} \sum_i^M C_{\text{match}}(y_i, \hat{y}_{\sigma(i)}), \quad (6)$$

where \mathcal{P} is a set of all possible permutation of M elements, and $\sigma(i)$ is the i th element of the permutation σ .

We define the following matching cost between the i th ground truth y_i and the corresponding i th prediction determined by permutation σ :

$$\begin{aligned} C_{\text{match}}(y_i, \hat{y}_{\sigma(i)}) &= C_{\text{match}}^{\text{Part}}(y_i, \hat{y}_{\sigma(i)}) + C_{\text{match}}^{\text{Sum}}(y_i, \hat{y}_{\sigma(i)}) \\ &= \sum_{t \in \{s, p, o, \text{spo}, G_s, G_p, G_o\}} -\mathbb{1}_{\{c_i^t \neq \emptyset\}} \hat{p}_{\sigma(i)}^t(c_i^t) \\ &+ \sum_{t \in \{s, p, o, G_s, G_o\}} \mathbb{1}_{\{c_i^t \neq \emptyset\}} \mathcal{L}_{\text{box}}(b_i^t, \hat{b}_{\sigma(i)}^t). \end{aligned} \quad (7)$$

where $\hat{p}_{\sigma(i)}^t(c_i^t)$ is the probability of classifying t as c_i^t computed by Eq. 4 and 5, and $\hat{b}_{\sigma(i)}^t$ is a predicted bounding box (G_s and G_o denote the subject and object embeddings from a Sum query branch). \mathcal{L}_{box} includes GIoU and L1 losses, same as [3]. Here we use a union box of the subject and object in a relation to represent the target box of a corresponding predicate.

Training loss. Given two-level Part and Sum outputs, we compute classification losses and box regression losses on both levels. Once we obtained the best permutation σ that minimizes the overall matching cost between y and \hat{y} , we can compute the total loss as

$$\begin{aligned} \mathcal{L}(y, \hat{y}) &= \sum_{i=1}^M (\mathcal{L}_{\text{Part}}(y, \hat{y}) + \mathcal{L}_{\text{Sum}}(y, \hat{y})) \\ \mathcal{L}_{\text{Part}}(y, \hat{y}) &= \sum_{t \in \{s, p, o, \text{spo}\}} -\log \hat{p}_{\sigma(i)}^t(c_i^t) + \mathcal{L}_{\text{box}}(b_i^t, \hat{b}_{\sigma(i)}^t) \\ \mathcal{L}_{\text{Sum}}(y, \hat{y}) &= \sum_{t \in \{G_s, G_p, G_o\}} -\log \hat{p}_{\sigma(i)}^t(c_i^t) + \mathcal{L}_{\text{box}}(b_i^t, \hat{b}_{\sigma(i)}^t) \end{aligned} \quad (8)$$

Note that Eq. 8 is very similar to Eq. 7, except that a negative log-likelihood loss is used to train classifiers, for more effective learning.

Table 1: Phrase and relationship detection result comparison (%) under various Transformer architectures on the VRD dataset. For a fair comparison, all transformer decoders are based on the same CNN backbone and Transformer encoder, using the same number of queries.

No.	Query Type			Transformer Decoder Design			Phrase Detection				Relationship Detection			
	Vanilla	Tensor	Composite	Vanilla	Part	Part-and-Sum	$k = 1$		$k = 70$		$k = 1$		$k = 70$	
							R@50	R@100	R@50	R@100	R@50	R@100	R@50	R@100
(a)	✓			✓			26.17	29.43	27.66	32.71	17.88	19.41	19.97	23.08
(b)		✓		✓			26.69	31.46	28.67	34.35	19.36	22.63	21.89	25.89
(c)		✓			✓		30.40	34.86	32.29	37.68	23.28	26.30	25.46	29.65
(d)			✓	✓			25.70	29.66	28.01	34.11	17.75	20.20	20.17	24.53
(e)			✓			✓	30.63	33.82	32.55	40.63	23.57	27.63	26.48	31.83

4. Experiments

We evaluate our method on two composite set detection applications: Visual Relationship Detection (**VRD**), and Human Object Interaction detection (**HOI**).

Datasets. (1) For the VRD task, we evaluate the proposed PST on VRD dataset [29], containing 5,000 images, and 100 entity categories and 70 predicate categories. Relationships are labeled as a set of $\langle \text{subject}, \text{predicate}, \text{object} \rangle$ triplets, and all subject and object entities in relationships are annotated with an entity category and a bounding box. We follow the data split of [29] and use 3,700/300/1,000 images for training/validation/test. There are 37,993 visual relationship instances that belong to 6,672 triplet types, and 1,169 relation types that only appear in the test set, which are used for zero-shot relationship detection. (2) For HOI task, we conduct an evaluation on HICO-DET dataset [4], including 38,118 training images and 9,658 testing images. In this dataset, there are the same 80 object categories as MS-COCO [27] and 117 verb categories, and objects and verbs construct 600 classes of HOI triplets. One person is able to interact with multiple objects in various ways at the same time in this dataset.

Task Settings. For the VRD task, we test PST on Phrase Detection and Relationship Detection [29, 54]. In Phrase Detection, the model detects one bounding box for each relationship, and recognizes the categories of subject, object and predicate in the relationship. In Relationship Detection, the model detects two individual bounding boxes for both subject and object entities in the relationship, and classifies the subject, predicate and object in the relationship. In both tasks, we consider two settings: single and multiple predicates between a pair of subject and object entities, with k representing the number of predicates between a pair.

Performance Metrics. (1) In VRD, we use relationship detection recall@K as the evaluation metric, as the true relationships annotations are incomplete. Following the evaluation in [29], for each detected relationship, we compute joint probability of subject, predicate and object category prediction as the score for that relationship, and then rank all detected relationships to compute the recall metric. For a relationship to be detected correctly, all three elements are required to be correctly classified and the IoU between the predicted bounding boxes and groundtruth bounding boxes are greater than 0.5. (2) In HOI, we use mean average pre-

cision (mAP) [11] as the evaluation metric. An HOI detection is considered correct only when the action and the object class are both correctly recognized and the corresponding human and object bounding boxes detection have higher than 0.5 IoU with the groundtruth boxes.

Implementation Details. On both VRD and HOI task, PST shares configurations. PST uses the standard ResNet-50 network as the backbone, followed by a Transformer encoder with six encoder layers, the same as Deformable DETR [58]. The proposed PST decoder contains six layers of the proposed two-stream Part-and-Sum decoder layers. All feed-forward networks are two-linear-layer shallow networks. We set up auxiliary losses after each decoder layer, and use 400 composite queries, with three part queries representing the subject, object and predicate in a visual relation or human, object and interaction in a HOI, respectively. Note that, in our experiment, we use the vanilla multi-head self-attention module [44] as a self-attention layer, and use a deformable multi-head cross-attention module as a cross-attention layer. More details are in the supplementary.

4.1. Part-and-Sum Transformer Analysis

Part-and-Sum Transformer Decoder We first compare and analyze different Transformer designs, with different query types on VRD. The various Transformer designs are compared in Figure 2. Vector based query is the most straightforward way to detect structured outputs using a Transformer, by formulating an individual structure entity as a vector query, and feeding queries into a vanilla Transformer decoder [58] to learn embeddings for each relation. Then, three one-linear-layer heads are used to predict the classes of subjects, objects and predicates, and two three-layer MLP heads to regress the boxes. The result comparison are shown in Table 1.

We can see that (1) with vanilla transformer decoder, tensor based query outperforms vector based query (in (a) vs (b)), with the margin 1.01/1.64% and 2.03/2.81% at R@50/100 on Phrase and Relationship detection ($k = 70$). It is because vector query models the structure entity as a whole, and embeds multiple parts in one query. This design increases the difficulty of Hungarian matching. (2) For Tensor based query, Part Transformer outperforms the Vanilla Transformer with a clear margin (in (b) vs (c)). This benefit mainly comes from the factorized design in the self-

attention layer, and the relation-level constraint (in Eq. 8). The former enhances the intra-relation context, reducing the ambiguity of entity recognition, for example, Subject “Person” and Object “horse” are important clues to infer Predicate “Ride”. The latter learns the relation as a whole to reduce the entity instance confusion [54]. (3) Despite the part-and-sum benefit inside the composite query, vanilla transformer with composite query degrades (in (d) vs (a)), compared with using the vanilla query. It shows that directly mixing Part and Sum queries cannot benefit structured output learning, because per Sum query contains multiple parts, and some relations may share the same entity instance, which could confuse the similarity computation in self-attention modules. To leverage two-level information and context effectively, PST decodes Part and Sum queries separately, with group-wise part-sum interaction. By comparing (d) vs (e), this design outperforms vanilla transformer, with the margin 4.54/6.52% and 6.31/7.30% at R@50/100 on Phrase and Relationship detection ($k = 70$).

Factorized self-attention layer We check the effectiveness of the factorization self-attention layer in the Part query decoding stream. We compare the performance of PST with factorization self-attention layer vs PST with vanilla self-attention layer on VRD, and the results are shown in Table 2. It shows that factorized self-attention design leads to 1.18/2.66% improvement at R@50/100 on Relationship detection ($k = 70$).

Part-Sum interaction We compare two Part-Sum interaction schemes: Vanilla Self-attention vs Summation operation. The results are shown in Table 2. From it, we can see that Part-Sum bidirectional summation works better than self-attention interaction, mainly due to the determined grouping configuration between part and sum queries in PST as shown in Figure 2(c)¹.

Table 2: Ablation study on Part-Sum Transformer designs. We report Relationship detection result comparison (%) on VRD.

Module		Relationship Detection			
		$k = 1$		$k = 70$	
		R@50	R@100	R@50	R@100
Factorization SA	✗	22.14	26.48	25.30	29.17
	✓	23.57	27.63	26.48	31.83
Part-Sum Interaction	Self-Attention	22.04	25.42	23.89	28.87
	Part↔Sum	23.57	27.63	26.48	31.83

Composite prediction Given the two-stream design of Part-and-Sum decoding, we obtain the predictions from both part and sum levels. Thus, we study the various inference schemes: predicting structure data from part query branch, or from sum query branch, or combining two branches. To combine predictions from part and sum queries, for classification probability, we average the pre-

¹In PST, the component order of a composite query is fixed. On VRD, for instance, the first part query is for Subject, the second for Object, and the third for Predicate. The grouping between part queries and sum query is also fixed by design, i.e. the first sum query and the first group of part queries represent the same relation instance.

diction probability of group-wise part and sum queries; for box prediction, we just average predicted positions of left-top and right-bottom points. The results comparison are shown in Table 3, and it shows that Part only inference slightly outperforms Sum only inference, and combining the predictions from two levels is able to bring a minor improvement for relationship detection².

Table 3: Result comparison of various inference schemes.

Inference	Relationship Detection			
	$k = 1$		$k = 70$	
	R@50	R@100	R@50	R@100
Part only	23.57	27.63	26.48	31.83
Sum Only	22.06	25.43	25.76	30.45
Part-Sum	24.34	27.01	27.03	31.90

4.2. Visual Relationship Detection

We compare PST with existing visual relationship detection solutions on the VRD datasets [29].

Competitors. Existing visual relationship detection solutions can be classified into two categories: (I) *Stage-wise methods*: These approaches first detect objects using pre-trained detector, and then use the outputs of the object detector as *fixed* inputs of a relationship detection module. Specifically, we compare with: (1) VRD-Full [29] which combines the visual appearance and the language features of candidates boxes to learn relationships. (2) NMP [16] which builds a relationship graph and optimizes it by node-to-edge and edge-to-node message passing mechanisms. (3) CDDN [5] which proposes a context guided visual-semantic feature fusion scheme for predicate detection. (4) LSVR [53] which learns a better representation by aligning the features on both entity and relationship levels. (5) ReIDN [54] which uses contrastive loss functions to learn fine-grained visual features. (6) BCVRD [17] which proposes a new box-wise fusion method to better combine visual, semantic and spatial features. (7) HGAT [30] which proposes to use object-level and triplet-level reasoning to improve relationship detection. (II) *End-to-End methods*: These approaches detect objects and relationships jointly. Specifically, we compare with: (1) CAI [59] - leverages the subject-object context to detect relationships. (2) KL distillation[50] - uses a linguistic model to regularize the visual model learning. (3) DR-Net [6] - designs a fully connected network to mine object-pair relationships. (4) Zoom-Net [49] - leverages multi-scaled relation contexts. (5) VTransE [52] - learns to map the visual features to the relationship space.

Results. The visual relation detection comparisons on VRD dataset are shown in Table 4. For clarity, the stage-wise methods are grouped in the *first* block, and end-to-end methods are in the *second* block. The proposed PST belongs to the *second* block, and particularly it is the first

²We report the results by part query only based inference in both VRD and HOI experiments for clarity.

Table 4: Phrase and relationship detection result comparison (%) on VRD dataset. - denotes that the results are not reported in the original paper. k is the number of predicates associated with each subject-object pair. Note on VRD dataset, the maximum number of predicates is $k = 70$. The *first* block is for the stage-wise detection methods, and the *second* block is for end-to-end detection methods. Our method belongs to the latter. †: the reported results of BC-VRD are based on Faster R-CNN for a fair comparison.

Method	Phrase Detection				Relationship Detection			
	$k = 1$		$k = 70$		$k = 1$		$k = 70$	
	R@50	R@100	R@50	R@100	R@50	R@100	R@50	R@100
VRD-Full [29]	16.17	17.03	20.04	24.90	13.86	14.70	17.35	21.51
LSVR[53]	18.32	19.78	21.39	25.65	16.08	17.07	18.89	22.35
BC-VRD [17]†	19.72	20.95	24.47	28.38	15.87	16.63	19.91	22.86
MLA-VRD [55]	23.36	28.12	-	-	20.54	24.91	-	-
NMP [16]	-	-	-	-	20.19	23.98	21.50	27.50
HGAT [30]	-	-	-	-	22.52	24.63	22.90	27.73
RelDN-IMG [54]	26.37	31.42	28.24	35.44	19.82	22.96	21.52	26.38
MF-URLN [51]	31.50	36.10	-	-	23.90	26.80	-	-
RelDN [54]	31.34	36.42	34.45	42.12	25.29	28.62	28.15	33.91
DR-Net [6]	-	-	19.93	23.45	-	-	17.73	20.88
VTransE [52]	19.42	22.42	-	-	14.07	15.20	-	-
CAI [59]	17.60	19.24	-	-	15.63	17.39	-	-
Vip [23]	22.80	27.90	-	-	17.30	20.00	-	-
KL distillation[50]	23.14	24.03	26.32	29.43	19.17	21.34	22.68	31.89
Zoom-Net [49]	24.82	28.09	29.05	37.34	18.92	21.41	21.37	27.30
PST (ours)	30.63	33.82	32.55	40.63	23.57	27.63	26.48	31.83

holistic end-to-end VRD solution (directly outputs all predicted relationships without any post-processing). It is evident that PST outperforms the existing end-to-end methods on both Phrase and Relationship Detection tasks, e.g. surpassing the second best end-to-end method Zoom-Net [51] with a margin of 5.81%/5.73% in Phrase Detection, and 4.65%/6.22% in Relationship detection at R@50/100 when $k = 1$. It shows that PST is able to learn the relationships between all entities effectively.

From the comparisons with the stage-wise VRD methods, PST outperforms the second best method HGAT [30] with a margin of 2.8% at R@50 in the Relationship detection, but lags behind the best method RelDN [54] with a margin of 0.71% and 1.32% on Phrase and relationship detection at R@50 with $k = 1$. We note that RelDN is a sophisticated two-stage method that:(1) leverages two CNNs for the entity and predicate visual feature learning; (2) tunes the thresholds of three margins in metric learning based losses functions; (3) combines multi-modality information (visual, semantic and spatial information) for relationship prediction. By contrast, PST predicts relations just based on visual features and detects relationships end-to-end and holistically without any post-processing. PST is simple, without any hand-designed components to represent the prior knowledge.

4.3. Human Object Interaction Detection

Competitors. We compare our model with two types of state-of-the-art HOI methods: the two-stage methods [36, 12, 32, 11, 45, 13, 25, 46, 31, 43, 28, 20, 2, 56, 24, 14, 10, 4] and the single-stage methods [19, 26, 60]. The two-stage methods aim to detect individual objects in the

first stage. Then, they associate the detected objects and infer the HOI predictions in the second stage. Two-stage methods rely on good object detections in the first stage and mostly focus on the second stage where language priors [13, 31, 28, 20, 2, 56, 24, 10] and human pose features [45, 25, 28, 20, 24] may be leveraged to facilitate the inference of HOI predictions from detected object. The single-stage methods aim to bypass the object detection step and directly output HOI predictions in one step. Previous single-stage methods [19, 26] are not end-to-end solutions. They employ multiple branches with each branch outputs complementary HOI-related predictions and rely on post-processing to decode the final HOI predictions. The most related to our approach is HoiT [60] which is an end-to-end single stage solution. HoiT employs DETR-like structure and predicts an HOI triplet from each vector query.

Results. Table 5 shows the results of our method and the other state-of-the-art HOI methods on HICO-DET dataset. We see that most of the two-stage models have mAP around 20 (default, full) on the HICO-DET test set. The best two-stage model is DRG [10] which achieves 24.5 mAP. However, it is a complex model and requires three-stage training. In comparison, as end-to-end single stage models, our model and the contemporary HoiT [60] model are able to achieve 20+ mAP without using a dedicated object detector or extra pose or language information. Our model with composite queries has mAP of 23.9 and achieves the state-of-the-art performance for single-stage HOI.

4.4. Ablation Study

We report more components analysis of our proposed part-and-sum transformer model (PST) here.

Shared-stream vs independent-stream decoding Our part-and-sum transformer (PST) consists an independent-stream decoder for part queries and sum queries, i.e. part queries and sum queries are feed into variant self-attention layers, cross-attention layers and FFNs, and decoded independently. We compare this design with a shared-stream design, where part and sum queries are decoded by the same layers. The results are shown in Table 6 where the independent-stream design is shown to outperform the shared-branch design on both relationship detection and phrase detection tasks. We hypothesize that part queries and sum queries represent different aspects of a relationship, and it is better to decode these two kinds of queries independently.

Varying the number and dimension of the queries We show the comparisons for the tensor-based query strategy vs the vector-based query strategy by varying their numbers and dimensions. Specifically, we use 500 tensor queries in PST, and each query contains three sub-vector queries of 256 dimension. For a fair comparison, we also employ 500 vector queries but of 256×3 dimension. Note that increas-

Table 5: Comparison with state-of-the-art HOI methods on HICO-DET dataset. For “Detector”, “COCO” refers to an off-the-shelf object detector trained on COCO. “HICO-DET” means the COCO pretrained object detector is further finetuned on HICO-DET. “Pose” refers to using human pose as additional features. “Language” refers to adopting the language priors.

Method	Backbone	Detector	Pose	Language	Default			Known Object		
					Full↑	Rare↑	NonRare↑	Full↑	Rare↑	NonRare↑
<i>Two-stage methods</i>										
Shen et al.[36]	VGG19	COCO			6.46	4.24	7.12	-	-	-
HO-RCNN[4]	CaffeNet	COCO			7.81	5.37	8.54	10.41	8.94	10.85
InteractNet[12]	ResNet-50-FPN	COCO			9.94	7.16	10.77	-	-	-
GPNN[32]	ResNet-101	COCO			13.11	9.34	14.23	-	-	-
iCAN[11]	ResNet-50	COCO			14.84	10.45	16.15	16.26	11.33	17.73
PMFNet-Base[45]	ResNet-50-FPN	COCO			14.92	11.42	15.96	18.83	15.30	19.89
PMFNet[45]	ResNet-50-FPN	COCO	✓		17.46	15.65	18.00	20.34	17.47	21.20
No-Frills[13]	ResNet-152	COCO		✓	17.18	12.17	18.68	-	-	-
TIN[25]	ResNet-50	COCO	✓		17.22	13.51	18.32	19.38	15.38	20.57
CHG[46]	ResNet-50	COCO			17.57	16.85	17.78	21.00	20.74	21.08
Peyre et al.[31]	ResNet-50-FPN	COCO		✓	19.40	14.63	20.87	-	-	-
IPNet[47]	Hourglass	COCO			19.56	12.79	21.58	22.05	15.77	23.92
VSGNet[43]	ResNet-152	COCO			19.80	16.05	20.91	-	-	-
FCMNet[28]	ResNet-50	COCO	✓	✓	20.41	17.34	21.56	22.04	18.97	23.12
ACP[20]	ResNet-152	COCO	✓	✓	20.59	15.92	21.98	-	-	-
Bansal et al.[2]	ResNet-50-FPN	HICO-DET	✓	✓	21.96	16.43	23.62	-	-	-
PD-Net[56]	ResNet-152	COCO		✓	20.81	15.90	22.28	24.78	18.88	26.54
PastaNet[24]	ResNet-50	COCO	✓	✓	22.65	21.17	23.09	24.53	23.00	24.99
VCL[14]	ResNet-101	HICO-DET			23.63	17.21	25.55	25.98	19.12	28.03
DRG[10]	ResNet-50-FPN	HICO-DET		✓	24.53	19.47	26.04	27.98	23.11	29.43
<i>One-stage methods</i>										
UnionDet[19]	ResNet-50-FPN	HICO-DET			17.58	11.52	19.33	19.76	14.68	21.27
PPDM[26]	Hourglass	HICO-DET			21.73	13.78	24.10	24.58	16.65	26.84
HoIT[60]	ResNet-50	-			23.46	16.91	25.41	26.15	19.24	28.22
PST (Ours)	ResNet-50	-			23.93	14.98	26.60	26.42	17.61	29.05

Table 6: Comparison(%) of a shared-stream vs independent-stream in the PST decoder. Shared-stream: Part and Sum queries share the same SA and CA layers; independent-stream: Part and Sum queries are independently decoded by different SA and CA layers (SA: Self-attention and CA: Cross-attention).

Decoder design	Phrase Detection				Relationship Detection			
	$k = 1$		$k = 70$		$k = 1$		$k = 70$	
	R@50	R@100	R@50	R@100	R@50	R@100	R@50	R@100
Shared-stream	27.32	32.71	31.59	37.82	20.04	23.10	24.89	29.87
Independent-stream	30.63	33.82	32.55	40.63	23.57	27.63	26.48	31.83

ing the dimension of vector-based queries to three times larger (256×3) requires to three time larger dimension of image memory. To this end, we repeat the features of the encoder three times. By doing it, each query in both models can have the equal embedding dimension to represent each relationship. The comparison of results is shown in Table 7. From it, we see that (1) increasing the number of vector queries from 500 to 1500 is not able to bring the clear benefit; (2) Increasing the dimension of each vector query does not provide more information; (3) Tensor based query outperforms variant vector based query, with a margin of 4.65%/5.16% increase on R@50/100 on Relationship detection when $k = 1$; (4) Compared to the tensor based query, composite queries further improve the performance due to part-and-sum two-level learning.

Part-and-Sum design in the HOI task The comparisons for the vector based Transformer (PST-Sum), tensor based Transformer (PST-Part), and part-and-sum Transformer (PST) on the HOI task are shown in Table 8. PST outperforms PST-Part and PST-Sum, same as the perfor-

Table 7: Comparison of variant query designs in the PST.

Query Design			Phrase Detection				Relationship Detection			
Formulation	Number	Dimension	$k = 1$		$k = 70$		$k = 1$		$k = 70$	
			R@50	R@100	R@50	R@100	R@50	R@100	R@50	R@100
Vector	1500	256	26.39	29.68	29.41	32.68	18.63	21.14	20.26	23.64
Vector	500	256	26.17	29.43	27.66	32.71	17.88	19.41	19.97	23.08
Vector	500	256×3	25.13	30.17	27.88	33.34	18.65	20.94	21.20	25.92
Tensor	500	256	30.40	34.86	32.29	37.68	23.28	26.30	25.46	29.65
Composite	500	256	30.63	33.82	32.55	40.63	23.57	27.63	26.48	31.83

mance comparison in the VRD task. Note that [60] can be regarded as a kind of PST-Sum model, but with different implementations such as different classification losses, and Transformer attention designs.

Table 8: Results of variant PST decoders on HOI.

Method	Default			Known Object		
	Full↑	Rare↑	NonRare↑	Full↑	Rare↑	NonRare↑
PST-Sum	21.37	13.85	23.62	23.28	15.24	25.69
PST-Part	22.24	14.15	24.65	24.15	15.61	26.70
PST	23.93	14.98	26.60	26.42	17.61	29.05

5. Conclusion

In this work, we have presented a Transformer-based detector, Part-and-Sum Transformers (PST), for visual relationship detection and human object interaction detection. PST maintains separate representations for the sum and parts while enhancing their interactions with composite queries. This design helps reduce instance ambiguity in structure data detection by learning rich intra-relationship and inter-relationship simultaneously. More qualitative results are provided in the supplementary document.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [2] Ankan Bansal, Sai Saketh Rambhatla, Abhinav Shrivastava, and Rama Chellappa. Detecting human-object interactions via functional generalization. In *AAAI*, 2020. 7, 8
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. *ECCV*, 2020. 1, 2, 4
- [4] Yu-Wei Chao, Yunfan Liu, Xieyang Liu, Huayi Zeng, and Jia Deng. Learning to detect human-object interactions. In *WACV*, 2018. 1, 2, 5, 7, 8
- [5] Zhen Cui, Chunyan Xu, Wenming Zheng, and Jian Yang. Context-dependent diffusion network for visual relationship detection. In *ACM MM*, 2018. 6
- [6] Bo Dai, Yuqi Zhang, and Dahua Lin. Detecting visual relationships with deep relational networks. In *CVPR*, 2017. 6, 7
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2
- [8] Pedro F Felzenszwalb, Ross B Girshick, and David McAllester. Cascade object detection with deformable part models. In *CVPR*, 2010. 4
- [9] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *TPAMI*, 2009. 2
- [10] Chen Gao, Jiarui Xu, Yuliang Zou, and Jia-Bin Huang. Drg: Dual relation graph for human-object interaction detection. In *ECCV*, 2020. 1, 2, 7, 8
- [11] Chen Gao, Yuliang Zou, and Jia-Bin Huang. ican: Instance-centric attention network for human-object interaction detection. In *BMVC*, 2018. 1, 2, 5, 7, 8
- [12] Georgia Gkioxari, Ross Girshick, Piotr Dollár, and Kaiming He. Detecting and recognizing human-object interactions. In *CVPR*, 2018. 7, 8
- [13] Tanmay Gupta, Alexander G. Schwing, and Derek Hoiem. No-frills human-object interaction detection: Factorization, appearance and layout encodings, and training techniques. *ICCV*, 2018. 7, 8
- [14] Zhi Hou, Xiaojiang Peng, Yu Qiao, and Dacheng Tao. Visual compositional learning for human-object interaction detection. In *ECCV*, 2020. 1, 2, 7, 8
- [15] Jian-Fang Hu, Wei-Shi Zheng, Jianhuang Lai, Shaogang Gong, and Tao Xiang. Recognising human-object interaction via exemplar based modelling. In *ICCV*, 2013. 1, 2
- [16] Yue Hu, Siheng Chen, Xu Chen, Ya Zhang, and Xiao Gu. Neural message passing for visual relationship detection. In *ICML Workshop*, 2019. 1, 2, 6, 7
- [17] Sho Inayoshi, Keita Otani, Antonio Tejero-de Pablos, and Tatsuya Harada. Bounding-box channels for visual relationship detection. In *ECCV*, 2020. 6, 7
- [18] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *arXiv preprint arXiv:2101.01169*, 2021. 2
- [19] Bumsoo Kim, Taeho Choi, Jaewoo Kang, and Hyunwoo J Kim. Uniondet: Union-level detector towards real-time human-object interaction detection. In *ECCV*, 2020. 1, 2, 7, 8
- [20] Dong-Jin Kim, Xiao Sun, Jinsoo Choi, Stephen Lin, and In So Kweon. Detecting human-object interactions with action co-occurrence priors. In *ECCV*, 2020. 7, 8
- [21] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017. 1, 2
- [22] John Lafferty, Andrew McCallum, and Fernando CN Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, 2001. 2
- [23] Yikang Li, Wanli Ouyang, Xiaogang Wang, and Xiao'ou Tang. Vip-cnn: Visual phrase guided convolutional neural network. In *CVPR*, 2017. 7
- [24] Yong-Lu Li, Liang Xu, Xinpeng Liu, Xijie Huang, Yue Xu, Shiyi Wang, Hao-Shu Fang, Ze Ma, Mingyang Chen, and Cewu Lu. Pastanet: Toward human activity knowledge engine. In *CVPR*, 2020. 1, 2, 7, 8
- [25] 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. In *CVPR*, 2019. 7, 8
- [26] 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. In *CVPR*, 2020. 7, 8
- [27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2, 5
- [28] Yang Liu, Qingchao Chen, and Andrew Zisserman. Amplifying key cues for human-object-interaction detection. In *ECCV*, 2020. 7, 8
- [29] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *ECCV*, 2016. 1, 2, 5, 6, 7
- [30] Li Mi and Zhenzhong Chen. Hierarchical graph attention network for visual relationship detection. In *CVPR*, 2020. 1, 2, 6, 7
- [31] Julia Peyre, Ivan Laptev, Cordelia Schmid, and Josef Sivic. Detecting unseen visual relations using analogies. In *ICCV*, 2019. 7, 8
- [32] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu. Learning human-object interactions by graph parsing neural networks. In *ECCV*, 2018. 7, 8
- [33] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Adv. Neural Inform. Process. Syst.*, 2015. 1, 2

- [34] Thomas Serre, Aude Oliva, and Tomaso Poggio. A feedforward architecture accounts for rapid categorization. *PNAS*, 2007. 2
- [35] Liyue Shen, Serena Yeung, Judy Hoffman, Greg Mori, and Li Fei-Fei. Scaling human-object interaction recognition through zero-shot learning. In *WACV*, 2018. 1, 2
- [36] L. Shen, S. Yeung, J. Hoffman, G. Mori, and L. Fei-Fei. Scaling human-object interaction recognition through zero-shot learning. In *WACV*, 2018. 1, 2, 7, 8
- [37] Jamie Shotton, John Winn, Carsten Rother, and Antonio Criminisi. Textonboost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation. In *CVPR*, 2006. 2
- [38] Zhangzhang Si and Song-Chun Zhu. Learning and-or templates for object recognition and detection. *TPAMI*, 2013. 2
- [39] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *EMLP*, 2019. 2
- [40] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877*, 2020. 2
- [41] Ioannis Tsochantaris, Thorsten Joachims, Thomas Hofmann, Yasemin Altun, and Yoram Singer. Large margin methods for structured and interdependent output variables. *JMLR*, 2005. 2
- [42] Zhuowen Tu. Auto-context and its application to high-level vision tasks. In *CVPR*, 2008. 2
- [43] Oytun Ulutan, ASM Iftekhar, and Bangalore S Manjunath. Vsgnet: Spatial attention network for detecting human object interactions using graph convolutions. In *CVPR*, 2020. 7, 8
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Adv. Neural Inform. Process. Syst.*, 2017. 1, 2, 5
- [45] Bo Wan, Desen Zhou, Yongfei Liu, Rongjie Li, and Xuming He. Pose-aware multi-level feature network for human object interaction detection. In *ICCV*, 2019. 7, 8
- [46] Hai Wang, Wei-shi Zheng, and Ling Yingbiao. Contextual heterogeneous graph network for human-object interaction detection. In *ECCV*, 2020. 7, 8
- [47] Tiancai Wang, Tong Yang, Martin Danelljan, Fahad Shahbaz Khan, Xiangyu Zhang, and Jian Sun. Learning human-object interaction detection using interaction points. In *CVPR*, 2020. 8
- [48] Yifan Xu, Weijian Xu, David Cheung, and Zhuowen Tu. Line segment detection using transformers without edges. In *CVPR*, 2021. 2
- [49] Guojun Yin, Lu Sheng, Bin Liu, Nenghai Yu, Xiaogang Wang, Jing Shao, and Chen Change Loy. Zoom-net: Mining deep feature interactions for visual relationship recognition. In *ECCV*, 2018. 1, 2, 6, 7
- [50] Ruichi Yu, Ang Li, Vlad I Morariu, and Larry S Davis. Visual relationship detection with internal and external linguistic knowledge distillation. In *ICCV*, 2017. 1, 2, 6, 7
- [51] Yibing Zhan, Jun Yu, Ting Yu, and Dacheng Tao. On exploring undetermined relationships for visual relationship detection. In *CVPR*, 2019. 7
- [52] Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. Visual translation embedding network for visual relation detection. In *CVPR*, 2017. 6, 7
- [53] Ji Zhang, Yannis Kalantidis, Marcus Rohrbach, Manohar Paluri, Ahmed Elgammal, and Mohamed Elhoseiny. Large-scale visual relationship understanding. In *AAAI*, 2019. 6, 7
- [54] Ji Zhang, Kevin J Shih, Ahmed Elgammal, Andrew Tao, and Bryan Catanzaro. Graphical contrastive losses for scene graph parsing. In *CVPR*, 2020. 1, 2, 5, 6, 7
- [55] Sipeng Zheng, Shizhe Chen, and Qin Jin. Visual relation detection with multi-level attention. In *ACM MM*, 2019. 7
- [56] Xubin Zhong, Changxing Ding, Xian Qu, and Dacheng Tao. Polysemy deciphering network for robust human-object interaction detection. In *IJCV*, 2021. 7, 8
- [57] Song-Chun Zhu, David Mumford, et al. A stochastic grammar of images. *Foundations and Trends® in Computer Graphics and Vision*, 2007. 2
- [58] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable DETR: Deformable transformers for end-to-end object detection. *ICLR*, 2021. 2, 5
- [59] Bohan Zhuang, Lingqiao Liu, Chunhua Shen, and Ian Reid. Towards context-aware interaction recognition for visual relationship detection. In *ICCV*, 2017. 6, 7
- [60] Cheng Zou, Bohan Wang, Yue Hu, Junqi Liu, Qian Wu, Yu Zhao, Boxun Li, Chenguang Zhang, Chi Zhang, Yichen Wei, and Jian Sun. End-to-end human object interaction detection with hoi transformer. In *CVPR*, 2021. 2, 7, 8