

RESEARCH ARTICLE

Video Relationship Detection Using Mixture of Experts

ALA SHAABANA, ZAHRA GHARAE¹, AND PAUL FIEGUTH, (Senior Member, IEEE)

Vision and Image Processing Laboratory (VIP), Department of Systems Design Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada

Corresponding author: Zahra Gharaee (zahra.gharaee@uwaterloo.ca)

This work was supported by the Microsoft Office Media Group.

ABSTRACT Machine comprehension of visual information from images and videos by neural networks suffers from two limitations: (1) the computational and inference gap in vision and language to accurately determine which object a given agent acts on and then to represent it by language, and (2) the shortcoming in stability and generalization of the classifier trained by a single, monolithic neural network. To address these limitations, we propose **MoE-VRD**, a novel approach to visual relationship detection via a mixture of experts. MoE-VRD recognizes language triplets in the form of a \langle subject, predicate, object \rangle tuple to extract the relationship between subject, predicate, and object from visual processing. Since detecting a relationship between a subject (acting) and the object(s) (being acted upon) requires that the action be recognized, we base our network on recent work in visual relationship detection. To address the limitations associated with single monolithic networks, our mixture of experts is based on multiple small models, whose outputs are aggregated. That is, each expert in MoE-VRD is a visual relationship learner capable of detecting and tagging objects. MoE-VRD employs an ensemble of networks while preserving the complexity and computational cost of the original underlying visual relationship model by applying a sparsely-gated mixture of experts, which allows for conditional computation and a significant gain in neural network capacity. We show that the conditional computation capabilities and massive ability to scale the mixture-of-experts leads to an approach to the visual relationship detection problem which outperforms the state-of-the-art.

INDEX TERMS Computer vision, video analysis, visual relationship detection, mixture-of-experts, deep learning.

I. INTRODUCTION

In the last decade, there has been a surge in research on the machine comprehension of visual information from images and video sequences. In particular, the application of large neural networks has allowed problems to be tackled such as video object segmentation [1], [2], [3], [4], object recognition and classification [5], [6], [7], [8], [9], and action recognition [10], [11], [12], [13], [14].

This unprecedented progress in the comprehension of visual information, however, suffers from the computational and inference gap between vision and language [15] to accurately determine which object a given agent acts on and how it might represent it by language.

The associate editor coordinating the review of this manuscript and approving it for publication was Amin Zehtabian¹.

We can begin by drawing inspiration from the two-streams hypothesis of the brain and how it processes visual information, such that the brain distinguishes between a ventral stream (the “what” pathway) and a dorsal stream (the “where” or “how” pathway) [16]. In parallel with this distinction, natural languages contain two classes of verbs describing actions: manner verbs, describing how an action is performed by expressing cause, such as waving arms (implying cheering) or nodding head (for assent); as opposed to result verbs, that describe the result of an action by expressing their effect, such as move, heat, clean, enter etc. [17].

Computationally, then, several approaches to visual information processing focus on the manner in which an action is performed (how) [18], [19], [20], [21], versus the result of the action, whether an object moves (where) or changes in appearance [22], [23].

When an action results in nearby changes, the visual information processing problem consists of detecting the three inter-related entities of subject, predicate (action), and the object(s) involved; that is, to recognize language triplets in the form of a \langle subject, predicate, object \rangle tuple. To detect the relationship between a subject (acting) and the object(s) (acted upon), the action must be recognized. Some recent approaches to visual relationship detection have focused on static images [24], however static relationship detection clearly has limitations in understanding temporal constraints inherent in video sequences, which offer significant richness regarding relationships [15], [24], [25]. Therefore, there has been a significant research emphasis on detecting visual relationships in video sequences [24], [25], [26], [27], [15], [28], [29].

The main challenges associated with this problem stem from the very large datasets, high ambiguity, and huge amount of background clutter. Moreover, the objects involved may only barely be recognizable due to pose, motion blur, occlusion and lighting. On the other hand, large variance in predicate representations [25] also makes it difficult to learn latent patterns, thus it is essential to consider visual and spatial features, and language ambiguity with synonyms. There is also a combinatorial effect, in that the number of unique tuple classes can be exceptionally large (the product of the vocabulary of subjects, objects, and predicates).

In this article, we propose an approach to video visual relationship detection (VidVRD), implemented by a multi-expert framework, where each expert is trained using the same model, and where the outputs of all experts are gated based on a separate neural network. Our performance results show that our novel architecture substantially outperforms known state-of-the-art methods. The contributions of this work are thus two-fold:

- 1) We construct a novel multi-expert detection framework.
- 2) We capture recent developments in video visual relationship detection as experts in our proposed multi-expert architecture. Our proposed approach is not tuned to a particular choice of expert, and other choices of expert should be equally valid and applicable.

The rest of this paper is organized as follows: Section II further develops the VidVRD problem and overviews past work, Section III describes the encapsulation of an existing VidVRD approach into an expert [15], Section IV describes the experiments and results which are discussed in Section V.

II. RELATED WORK

Cascading failures caused by low-level misclassifications make a monolithic solution, such as a single very large neural network, not an ideal strategy to resolve large visual relationship detection problems. To address such limitations, earlier approaches proposed solutions such as dividing a given problem based on pre-processing and post-processing heuristics [15], [24], [25], [30], to attempt to correctly classify predicates.

The capacity of a neural network to learn is limited by its degrees of freedom (number of parameters) and further limited by the available data. When datasets are, in fact, large enough, then increasing the number of parameters can lead to significant improvements in performance. However for a typical deep learning model, where the entire network is activated on each training sample, the computational cost is roughly quadratic in the number of parameters, as both the model size and the number of training samples increase together [31], [32], [33]. This phenomenon is particularly and uniquely exacerbated in visual information processing since the input layer is very large. The statistical relationships of pixels and objects detected in images and videos are subtle, and the networks are often expected to perform multiple tasks like object segmentation and action recognition, all of which lead to significant increases in network size.

Limitations in computing power will eventually fall short of meeting the training demand, and the networks trained in this fashion tend to be brittle and sensitive to slight changes in the data distribution [34] and task specification [35]. That is, current systems are better characterized as narrow experts rather than as resilient generalists [36].

With the above context in mind, we make the following observation about the VidVRD task: A training approach based on a single, monolithic network is not sufficient to have stable and generalizable classifiers, at least in certain problem contexts [34], [36]. Therefore, applying divide-and-conquer by developing multiple small models and aggregating their outputs could be a promising solution [37], [38] to create more compact and/or resilient networks. As a consequence, the mixture-of-experts architecture suggests a strategy to achieve the larger network capacity needed to solve large numbers of sub-problems, which is typical of the VidVRD domain.

A. VIDEO RELATIONSHIP DETECTION

Video visual relationship detection (VidVRD) is made up of multiple problems that must be resolved simultaneously: object recognition, subject recognition, and action recognition. Hence, the detection often takes the form of classifying a triplet in the form of a \langle subject, predicate, object \rangle tuple for each detected subject-object pair in a video.

When dealing with static images, this problem is comparatively straightforward as it involves detecting such a triplet only once, using language and spatial features [39], [40], [41], [42], [43]. In contrast, in video the problem becomes significantly more complicated as spatio-temporal features come into play, and relationships can *change* over time, cascading the difficulty of the problem [15], [25], [44].

In 2017, Shang et al. [25] proposed the first approach for VidVRD, decomposing a given video into segments and stitching relationship predictions in preceding and succeeding segments through a greedy association algorithm. They also introduced the first fully annotated video visual relationship dataset [25]. Later, Qian et al. [26] proposed tackling the problem using a fully connected spatio-temporal

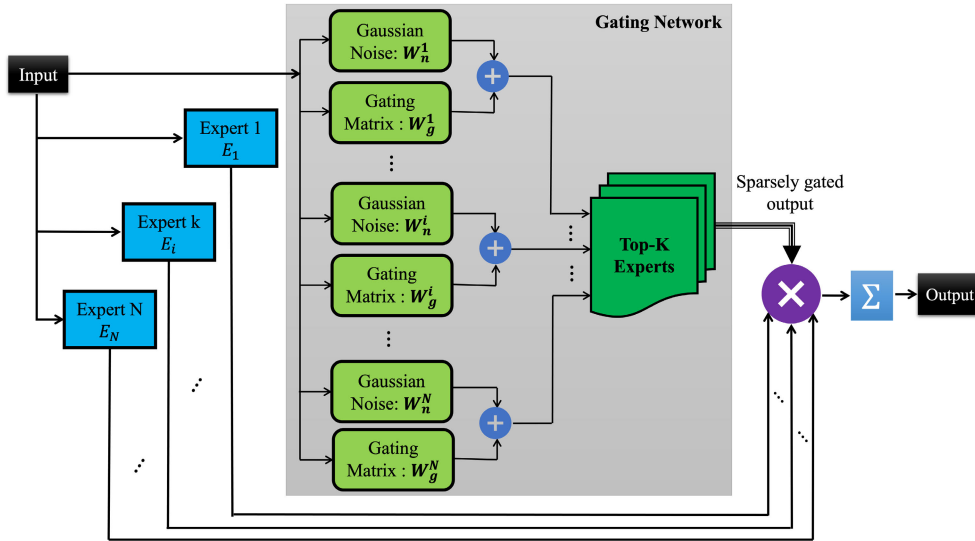


FIGURE 1. A Mixture of Experts (MoE) layer as described by Shazeer et al. [33].

graph. Tsai et al. [24] also constructed a spatio-temporal graph, but utilized Conditional Random Fields to exploit the statistical dependency between objects. Liu et al. [27] proposed a sliding-window scheme to simultaneously predict short-term and long-term relationships using different kernel sizes on object tracklets to generate sub-tracklet proposals with different durations.

Xiao et al. [29] proposed to spatio-temporally localize the relations (predicates) in a video sequence. Their proposed approach, the Visual Relation Grounding in Videos (vRGV) first produces region proposals from each frame of a video and then learns to ground a pre-defined relation from two trajectories. A trajectory is created by connecting the consecutive bounding boxes linked to a visual entity (subject or object) across a video segment.

Wu et al. [45] proposed using two graph-based networks to predict the spatial-temporal relations (actions) between subjects and objects in videos. They applied a gated graph network together with a long short-term graph network to, respectively, extract spatial relations within video frames and multi-scale temporal relations between consecutive frames.

Gao et al. [46] proposed a tracklet-based visual transformer composed of a temporal-aware decoder, which performs feature interactions between tracklets and predicate embeddings for relationship detection. Zheng et al. [47] also proposed the VRDFormer, in which a first module encodes a video into a sequential frame-level feature map, and a second one processes the sequential feature map in order to generate the relation instances.

Li et al. [48] proposed a method to address the long-tailed bias in VidVRD datasets, which results in poor generalization. Their approach, the Interventional Video Relation Detection (IVRD) applies causality-inspired intervention on the model input to decrease the effect of the spurious correlation in the training data, and therefore to enhance the robustness of the output prediction.

Cao et al. [49] proposed using comprehensive semantic representations that are useful for knowledge transfer across relationships to solve the VidVRD problem. Their approach, the Concept-Enhanced Relation Network (CKERN) produces conceptually richer semantic representations of the detected object pairs, and then predicts the relationship based on the integration of multi-modal features.

Gao et al. [50] proposed a classification-then-grounding approach based on the temporal bipartite graphs of the videos, where the nodes are entities and predicates, and the edges denote different semantic roles between the nodes. Their proposed approach, the Bipartite Graph model (BIG) first classifies all of the nodes and edges of the graph (classification), and then localizes the temporal location of each relation instance (grounding).

Chen et al. [51] introduced a compositional encoding for VidVRD. Their proposed approach, the Social Fabric Encoding (SFE) encodes a pair of object “tubelets” as a composition of interaction primitives. Learning these primitives, the resulting representation is used to localize and classify relationships from co-occurring objects.

More recently, Shang et al. [15] modified their earlier approach [25], whereby they proposed an iterative relation inference that exploits the inter-dependency of relation components (subject/objects and predicates) for better visual relation recognition. To achieve this, they created three preferential predictors with learnable tensors alongside the normal visual predictors to model the inter-dependency relationship between subjects/objects and predicate classes [15]. Hence, each relational component has three classifiers, and each of them consists of a visual predictor and a preferential predictor. The visual predictor is a deep neural network for recognizing the visual patterns of subject/object and predicate, whereas the preferential predictor refines the prediction of one variable (subject, object, or predicate) conditioned on the values of the other two [15]. Following a

similar architecture to their initial work [15], the authors take a sliding time window and generate object tracklet proposals as the detected entities, and then predict associated relation triplets.

Due to its lightweight network architecture, modularity, and state-of-the-art performance, we have chosen the work by Shang et al. [15] as the basis for the expert in the architecture proposed in this paper.

B. MIXTURE OF EXPERTS

Proposed more than three decades ago by Jacobs et al. [52], the mixture of experts (MoE) architecture has been applied to problems including the modeling of task relationships [53], increasing network breadth and depth [54], multi-modal generative models [55], and volunteer computing [56].

In 2017, Shazeer et al. proposed a new general purpose neural network component: the Sparsely-Gated Mixture-of-Experts (MoE) Layer [33], consisting of a number of experts, each a simple feed-forward neural network, together with a trainable gating network which selects a sparse subset of the experts to be trained on each given input [33]. The gating network essentially determines which of the experts are best suited to a given type of input. All parts of the network, both gating and experts, are trained jointly by back-propagation [33]. In their paper, Shazeer et al. applied their technique, illustrated in Figure 1, to language modeling and machine translation.

Riquelme et al. also proposed a Vision Mixture of Expert (V-MoE) [57] for image classification. V-MoE replaces a subset of feedforward layers in a vision transformer with sparse MoE layers, where each image patch is “routed” to a subset of “experts”. This halves the computation consumption at inference while performing equally well as the state-of-the-art.

It has been observed [33], [54], [58] that a gating network is inclined to converge to a state where it produces large weights for the same few experts regardless of input, a phenomenon very much analogous to the problems encountered with self-organized maps [59] (essentially a very flat single-layer network) from pattern recognition. This imbalance becomes self-reinforcing / self-perpetuating, as the favored experts are trained more frequently / more rapidly and thus are even more likely to be selected by the gating network [33]. To address this problem, Shazeer et al. [33] defined the importance of an expert relative to a batch of training samples to be the batch-wise sum of the gate values for that expert.

In the following section we will present our proposed architecture, which applies a sparsely gated mixture of experts [33] to video visual relationship detection [15].

III. ARCHITECTURE

Figure 3 illustrates our proposed architecture and its main components.

A. SPARSELY-GATED MIXTURE OF EXPERTS

The MoE consists of a set of N expert networks E_1, \dots, E_N and one gating network, G , whose output is a sparse binary N -dimensional vector. The experts are themselves identical feed-forward neural networks, each with their own parameters. Since our interest in this paper is the MoE concept, for the individual experts we have adopted the baseline state-of-the-art approach of [15] shown in Figure 1.

Given an input x , the output of the i th expert’s function is denoted as $E_i(x)$. These N outputs are combined in the MoE layer as

$$y = \sum_{i=1}^N G(x)_i \cdot E_i(x), \quad (1)$$

where $G(x)_i$ represents the output of the gating network. The sparsity in computation, one of the key strengths of the MoE approach, is realized by the explicit sparsity of the gating output. That is,

$$G(x)_i = 0 \quad \text{for most } i, \quad (2)$$

such that whenever $G(x)_i = 0$ the corresponding expert E_i is not invoked (the particular input x is not fed-forward into the network representing expert E_i).

There is significant flexibility in the choice of gating function. In this paper we adopt the absolute simplest case — a single-layer gating function, with more capable multi-layer generalizations to be considered as future work. The gating output $G(x)$ is computed as

$$G(x) = \text{Softmax}(\text{top}_K(W_g^i \cdot x + N_g(W_n^i \cdot x))), \quad (3)$$

where top_K selects the K largest values (the best experts), and W_g^i and W_n^i are trainable gating and noise weight matrices, respectively, which are parametrized for each expert i .

The number of samples sent to the gating layer is discrete, and therefore not applicable to back-propagation, however the inclusion of the noise term $N_g(x)$ allows for a smooth estimate of the number of samples used for each expert in each batch, thus allowing for the back-propagation of gradients. The noise function is defined as

$$N_g(x) = \text{StandardNormal}() \cdot \text{Softplus}(x), \quad (4)$$

where $\text{Softplus}(x) = \frac{1}{\beta} \log(1 + \beta x)$ is a smooth approximation of the ReLU function to constrain the output to be positive.

An importance term is considered in the overall loss to address imbalances resulting from the *self-reinforcing effect* [60], which occurs when certain favoured experts are trained more rapidly and thus are selected even more by the gating network. The importance loss is

$$L_{\text{importance}}(x) = \alpha(\text{CV}(g) + \text{CV}(l)), \quad (5)$$

where α is a hand-tuned scaling factor, g is the batch-wise sum of gate values

$$g = \sum_{x \in B} G(x) \quad (6)$$

over batch B , and

$$l = \sum_{x \in B, G(x) > 0} G(x) \quad (7)$$

represents the load, summed over the positive gate values.

We applied $CV(\cdot)$, the coefficient of variation,

$$CV(x) = \frac{\text{var}(x)}{(\text{mean}(x))^2 + \epsilon} \quad (8)$$

as additional loss terms in (5) which encourage experts to have a more balanced (equal) importance [33].

B. VISUAL RELATIONSHIP DETECTION

To address the problem of visual relationship detection in video sequences, we take inspiration from VidVRD-II [15], which learns relationship detection using iterative inference shown in Figure 2.

VidVRD-II assumes a set of three entities $\mathbb{E} = [e_1, e_2, e_3]$ representing subject e_1 , predicate e_2 and object e_3 , and their corresponding features $f_{e_1}, f_{e_2}, f_{e_3}$, which builds the language triplet $\langle \text{subject}, \text{predicate}, \text{object} \rangle$, and models the problem of video visual relationship detection as the joint probability

$$P(\langle e_1, e_2, e_3 \rangle | f_{e_1}, f_{e_2}, f_{e_3}). \quad (9)$$

This joint probability can be factorized as three conditional probabilities

$$P(e_1 | f_{e_1}, e_2, e_3) P(e_2 | f_{e_2}, e_1, e_3) P(e_3 | f_{e_3}, e_1, e_2), \quad (10)$$

which aids in inference when there is ambiguous visual information, since the classes of any two components imply a preference over the class of the third.

Each conditional probability of (10) is modelled by a classifier consisting of a visual predictor and a preferential predictor also shown in Figure 2. The visual predictor simply learns visual patterns of the subject e_1 , predicate e_2 , and object e_3 , using a deep neural network. The preferential predictor applies learnable dependency tensors to refine the prediction of one variable conditioned on the values of the other two:

$$e_{pr} = \begin{cases} P(e_1 | f_{e_1}, e_2, e_3) = \Phi(V_{e_1} \cdot f_{e_1}) + p_{e_2} \cdot W_{e_1} \cdot p_{e_3} \\ P(e_2 | f_{e_2}, e_1, e_3) = \Phi(V_{e_2} \cdot f_{e_2}) + p_{e_1} \cdot W_{e_2} \cdot p_{e_3} \\ P(e_3 | f_{e_3}, e_1, e_2) = \Phi(V_{e_3} \cdot f_{e_3}) + p_{e_1} \cdot W_{e_3} \cdot p_{e_2} \end{cases} \quad (11)$$

Here e_{pr} is the conditional probability vector of three entities, and $V = [V_{e_1}, V_{e_2}, V_{e_3}]$ are the learnable weights of the visual predictors. In our case study, the weights of the subject and object classifiers are shared, thus $V_{e_1} = V_{e_3}$. The paper of [15] applied the nonlinearity Φ to the entire expression of (11), as in

$$P(e_i | f_{e_i}, e_j, e_k) = \Phi(V_{e_i} \cdot f_{e_i} + p_{e_j} \cdot W_{e_i} \cdot p_{e_k}) \quad (12)$$

However actual implementations applied Φ to the first term only, as in (11), a convention which we have preserved for consistency. The weights $[W_{e_1}, W_{e_2}, W_{e_3}]$ model the

dependency of one class over the other two, separately parametrized for each classifier. Φ represents the nonlinear activation, here implemented by a Softmax function for the subject and object classes, and a Sigmoid function for the predicate class.

To design our architecture, we apply the MoE concept of Section III-A by incorporating VidVRD-II into the MoE framework, giving rise to the combined MoE-VRD shown in Figure 3. The object tracklet proposals are extracted and fed into the relation prediction experts. Based on the gating layer, the top K experts are chosen. Upon the completion of the forward pass, back-propagation is applied so the gradients back-propagate through the gating network and the selected experts.

1) OBJECT TRACKLET PROPOSALS

We use Seq-NMS [7] to generate object tracklet proposals as a pre-processing step to use as inputs to the relational classifier experts. For frame-level object detection, a Faster-RCNN with an Inception-ResNet foundation [61] is pretrained on the Open Images dataset [43].

The model serves as a suitably generic object detector [15], [61]. The bounding boxes and corresponding region features are extracted, after which Seq-NMS generates a compact set of object tracklets, which form the inputs to the expert neural networks.

2) FEATURE EXTRACTION

Applying the object tracklet proposals, we generate two types of features: *Visual Features* and *Relative Positional Features*, shown in Figure 3.

To generate the visual features f of (9) – (10), the bounding boxes are applied to extract the pretrained deep visual features of the subject and object entities, and the predicate’s visual feature is computed through a concatenation of the subject and object visual feature vectors.

In addition to the visual features, we extract a relative positional feature to represent the spatio-temporal relationship between the entities. For each pair of object tracklets, the algorithm computes the relative distance between the subject and object by encoding the spatial and temporal relative positional feature f_r^p :

$$f_r^p = \left[\frac{x_{e_1}^p - x_{e_3}^p}{x_{e_3}^p}, \frac{y_{e_1}^p - y_{e_3}^p}{y_{e_3}^p}, \log \frac{w_{e_1}^p}{w_{e_3}^p}, \log \frac{h_{e_1}^p}{h_{e_3}^p}, \log \frac{w_{e_1}^p h_{e_1}^p}{w_{e_3}^p h_{e_3}^p}, \frac{t_{e_1}^p - t_{e_3}^p}{30} \right], \quad (13)$$

where $p \in [b, e]$ represents the beginning or ending bounding box, characterized by coordinates (x, y) , width w , height h , and time t for subject e_1 and object e_3 . A feed-forward network is used to fuse the subject’s and object’s visual features f_{e_1}, f_{e_3} with the relative positional features of the beginning and ending bounding boxes f_r^b, f_r^e , where the relative positional feature f_r^p provides the expert with additional information to recognize visual relationships.

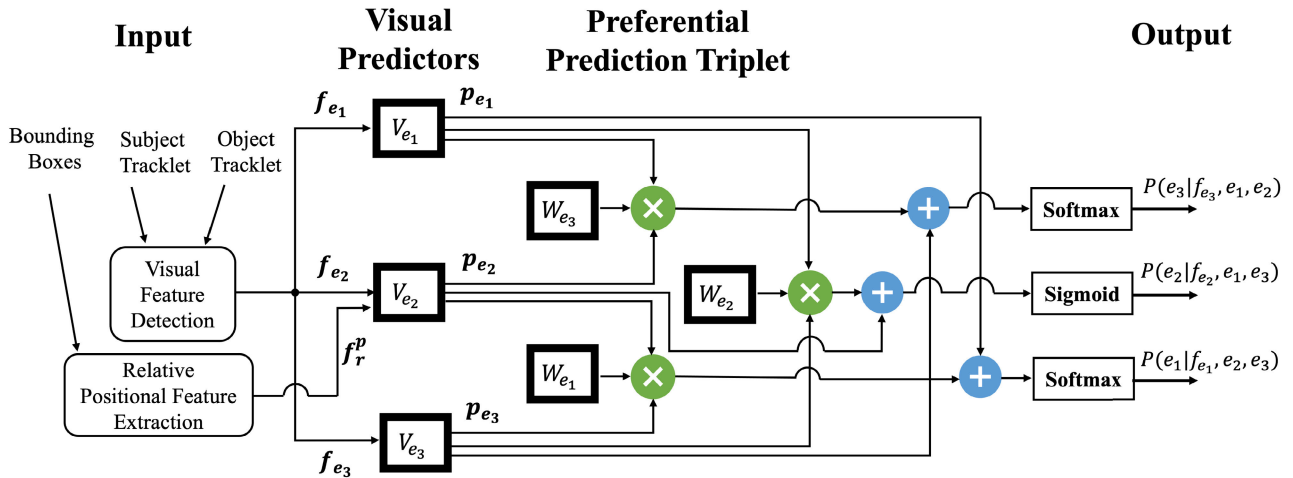


FIGURE 2. Visual relationship detection framework proposed by [15], which is used as the basis of our expert.

In summary, each encapsulated expert consists of an object predictor, a subject predictor, and a predicate predictor — each of which is a basic feed-forward network, allowing for a set of modestly-sized, nimble experts to speed up training and inference, when compared to an equivalent single monolithic network.

IV. EXPERIMENTS AND RESULTS

In order to properly assess the improvements offered by our proposed framework and to make a fair comparison with the state-of-the-art, we conduct our experiments using a similar experimental setup and the same datasets as used by the iterative inference approach proposed by Shang et al. [15].

A. DATASETS

Our experiments are conducted using two VidVRD benchmark datasets: ImageNet-VidVRD [25] and VidOR [62], [63]. The ImageNet-VidVRD is the first dataset for video visual relation detection, created by Shang et al. [25]. It consists of 1,000 videos collected from ILSVRC2016-VID [64]. These videos are manually annotated with video relation instances [15], [25]. VidOR is a recently-released large-scale benchmark also compiled by Shang et al. [62], which contains 10,000 social media videos from YFCC-100M [63].

B. EVALUATION METRICS

In object detection there are two related problems to be solved: localization and classification. Localization determines the location of an object (e.g., its bounding box), whereas classification infers the object's identity.

For object detection tasks, it is standard to calculate precision and recall based on a given threshold on the IoU (Intersection over Union), which measures the fractional overlap between predicted and true bounding boxes. If the IoU result for a predicted bounding box exceeds the threshold, then the prediction is classified as a true positive, otherwise it is a false positive.

We can thus define metrics, such as Recall@50, implying an evaluation of the recall metric based on an IoU threshold of 50% (i.e., allowing bounding boxes to be only 50% overlapping). Other metrics, such as Recall@100, Precision@10 etc. are equivalently defined.

We designed and implemented our evaluation metrics consistent with those in previous works [15], [25], [62]. To this end, we calculate how many ground truth relation instances are detected by the mixture of experts in each testing video. Metrics are divided into two categories: relation tagging and relation detection.

Relation tagging focuses on the precision of the relation triplet without considering the precision of its spatio-temporal location in the video. In other words, relation tagging simply checks whether the relationship was detected properly, but not whether it was detected with any accuracy in time or space. The tagging performance is evaluated by Precision@1 (P@1), Precision@5 (P@5), and Precision@10 (P@10).

In contrast, relation detection measures the precision of both the relation triplet and the corresponding subject/object trajectories for every detected relation instance [25]. Relation detection of instances are considered to be correct only if they match the ground truth relation instance, and the voluminal Intersection-over-Union (vIoU) of the trajectories of subject and object are both larger than a pre-defined threshold. The vIoU threshold is set to 50% in order to address the objectives of our experiments. The detection performance is measured using the Mean Average Precision (mAP), Recall@50 (R@50), and Recall@100 (R@100).

All experiments are run ten times with different random seeds for each expert, and the mean plus/minus standard-deviation scores are reported.

C. SINGLE EXPERT PERFORMANCE

To ensure the proper functioning of the proposed mixture of experts, clearly we must first validate the performance

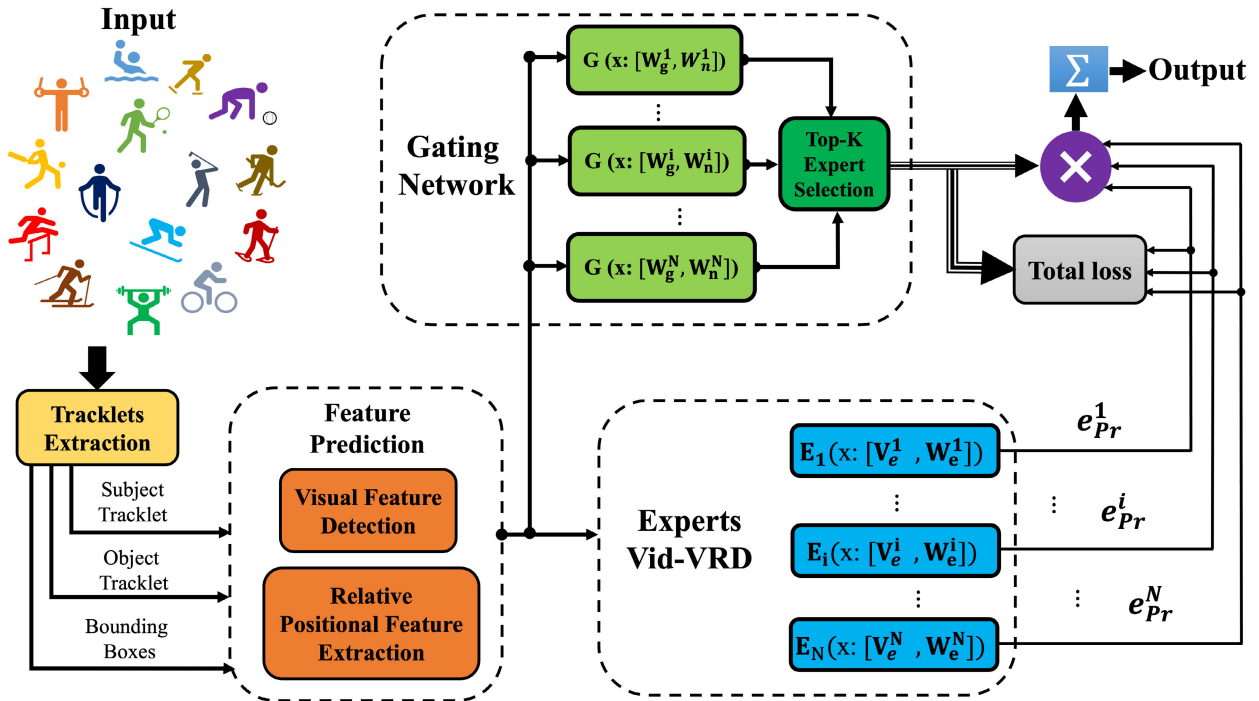


FIGURE 3. An illustration of the MoE-VRD architecture proposed in this article. Raw RGB images are taken as input; for each given image frame the subject and object tracklets are extracted and given to the feature extraction network together with bounding box information in order to generate visual and relative positional features representing all three entities: subject, predicate and object. The visual and positional features are applied as the input to our experts and gating networks. Every expert outputs a score corresponding to each entity, which represents both visual and preferential predictions. The gating network outputs a sparsely gated vector, which evaluates each expert’s learning. Selecting the top K experts, the sum-product of the sparsely gated expert scores is calculated and represented as the output of our MoE-VRD architecture.

TABLE 1. Comparisons of a single expert with the method of by Shang et al. [15] on the ImageNet-VidVRD Dataset [25] (top) and on the VidOR Dataset [62] (bottom). In both cases, the expert performs essentially equivalently to that of [15].

ImageNet-VidVRD Dataset	Relation detection			Relation tagging		
	mAP	R@50	R@100	P@1	P@5	P@10
VidVRD-II [15]	29.37 ± 0.40	19.63 ± 0.19	22.92 ± 0.48	70.40 ± 1.53	53.88 ± 0.31	40.16 ± 0.70
MoE-VRD ($N = 1$)	29.49 ± 0.32	18.65 ± 0.2	22.71 ± 0.33	68.50 ± 1.81	54.70 ± 0.23	39.50 ± 0.86

VidOR Dataset	Relation detection			Relation tagging		
	mAP	R@50	R@100	P@1	P@5	P@10
VidVRD-II [15]	8.65 ± 0.11	8.59 ± 0.11	10.69 ± 0.08	57.40 ± 0.57	44.54 ± 0.68	33.30 ± 0.31
MoE-VRD ($N = 1$)	7.98 ± 0.16	8.23 ± 0.28	11.02 ± 0.14	57.53 ± 0.43	43.12 ± 0.54	32.89 ± 0.24

of a single expert. The performance of the MoE framework consisting of only a single expert ($N = 1$, and thus necessarily $K = 1$) during training and testing should be essentially unchanged from the published performance of the underlying expert [15], [25].

In the multi-expert architecture, at every iteration the gating layer outputs a sparse vector selecting K experts. If there is only a single expert, then the gating layer is essentially irrelevant, simply selecting the same one expert every time, and the resulting performance should be the same as if we had simply run the model built by Shang et al. without any adjustments [15].

We have shown the results of the state-of-the-art approaches in Table 1, illustrating a comparison between the relation detection and relation tagging results of Shang et al.’s VidVRD-II [15] and our single-expert ($N = 1$) MoE-VRD architecture over the ImageNet-VRD and VidOR datasets.

Both approaches perform quite similarly, validating that the MoE-VRF framework is not interfering with the operation of the underlying expert, allowing us to generalize to multiple experts, next.

D. MULTI-EXPERT PERFORMANCE

We now wish to evaluate the performance of our proposed MoE-VRD architecture when more than one expert ($N > 1$) is at play. We evaluate our approach on the VidOR dataset [62] and ImageNet-VRD dataset [25].

We evaluate our work against recent representative approaches:

- **VidVRD-II** [15], which builds upon the same authors’ work in [25]. It uses an iterative inference approach to video relationship detection.
- **GSTEG** [24], which constructs a fully-connected spatio-temporal graph for relation inference.

TABLE 2. Performance of our proposed MoE-VRD with $K = 2$ and a total of $N = 10$ experts, in comparison with state-of-the-art approaches on the ImageNet-VidVRD dataset [25]. For every criterion the proposed MoE-VRD outperforms all other approaches. The substantial increase in performance stems unambiguously from the mixture-of-experts approach, since our expert on its own is no better than the method in VidVRD-II, as was shown in Table 1. “–” indicates that no corresponding results were reported.

ImageNet-VidVRD Dataset	Relation detection			Relation tagging		
	mAP	R@50	R@100	P@1	P@5	P@10
VidVRD [25]	8.58	5.54	6.37	43.00	28.90	20.80
GSTEG [24]	9.52	7.05	7.67	51.50	39.50	28.23
VRD-GCN [26]	14.23	7.43	8.75	59.50	40.50	27.85
3DRN [65]	14.68	5.53	6.39	57.89	41.80	29.15
VRD-STGC [44]	18.38	11.21	13.69	60.00	43.10	32.24
SFE [51]	20.08	13.73	16.88	62.50	49.20	38.45
IVRD [48]	22.97	12.40	14.46	68.83	49.87	35.57
BIG-C [50]	26.08	14.10	16.25	73.00	55.10	40.00
CKERN [49]	–	–	–	74.50	55.59	41.34
VidVRD-II [15]	29.37 ± 0.40	19.63 ± 0.19	22.92 ± 0.48	70.40 ± 1.53	53.88 ± 0.31	40.16 ± 0.70
MoE-VRD (K = 2)	33.02 ± 0.23	22.77 ± 0.28	24.20 ± 0.22	74.12 ± 1.44	56.47 ± 0.17	42.05 ± 0.92

TABLE 3. A comparison with the state-of-the-art, as in Table 2, but here on the VidOR dataset [62].

VidOR Dataset	Relation detection			Relation tagging		
	mAP	R@50	R@100	P@1	P@5	P@10
3DRN [65]	2.47	1.58	1.85	33.05	35.27	–
VRD-STGC [44]	6.85	8.21	9.90	48.92	36.78	–
IVRD [48]	7.42	7.36	9.41	53.40	42.70	–
CKERN [49]	–	–	–	58.80	46.07	34.29
BIG [50]	8.54	8.03	10.04	64.42	51.80	40.96
Ens-5 [46]	9.48	8.56	10.43	63.46	54.07	41.94
SFE [51]	11.21	9.99	11.94	68.86	55.16	–
VidVRD-II [15]	8.65 ± 0.11	8.59 ± 0.11	10.69 ± 0.08	57.40 ± 0.57	44.54 ± 0.68	33.30 ± 0.31
MoE-VRD(K = 2)	9.44 ± 0.21	9.54 ± 0.13	11.51 ± 0.31	58.92 ± 0.67	45.11 ± 0.19	34.85 ± 0.22

- **VRD-GCN [26]**, which builds a model that can take advantage of spatial-temporal contextual cues to make better predictions on objects as well as their dynamic relationships.
- **VRD-STGC [27]**, which proposes a novel sliding-window scheme to simultaneously predict short-term and long-term relationships [27], and extracts spatio-temporal features.
- **3DRN [65]**, which develops a 3-D CNN to learn the visual features for relation recognition in an end-to-end manner.
- **IVRD [48]**, which proposes a causality-inspired intervention on the model input to improve prediction robustness.
- **CKERN [49]**, which generates comprehensive semantic representations by incorporating retrieved concepts with local semantics.
- **BIG [50]**, which proposes a classification-then-grounding approach based on temporal bipartite graphs.
- **Ens-5 [46]**, which proposes a tracklet-based visual Transformer composed of a temporal-aware decoder.
- **SFE [51]**, which proposes encoding the representation of a pair of objects as a composition of interaction primitives.

For the ImageNet-VRD dataset [25] we compare to all ten of these methods; for the VidOR dataset [62] we compare to eight of the preceding methods, due to the choice of results reported in the respective papers.

Table 2 shows the results, comparing our proposed MoE architecture with all ten other methods on the ImageNet-VRD dataset [25]. The proposed MoE-VRD performs significantly better, in every metric, than any method tested, including the most recent state of the art. The large margin of improvement stems from having a gating function that allows experts to be trained quite separately on different sorts of inputs, leading to a degree of robustness due to heterogeneity, which is very difficult for single large monolithic networks.

Similar to Table 2, Table 3 now shows the comparative results on the VidOR dataset [62]. Our proposed MoE-VRD still exhibits superior performance in every metric when compared to most of the state of the art approaches, although by a lesser margin than in Table 2, likely due to the increased diversity of the VidOR dataset [62], and the related naivety or limitation of the MoE-VRD in using a set of identical experts. The creation of heterogeneous or differently-specialized experts is a subject for future research.

BIG [50], Ens-5 [46], and SFE [51] do outperform the proposed MoE-VRD for one or more metrics in Table 3, although for the Relation Detection assessment the MoE-VRD is highly competitive, outperforming BIG and Ens-5. In any event, a universal improvement on every possible dataset and/or metric is not to be expected, and the impressive results of SFE in Table 3 are, for example, significantly less impressive in its rather lackluster performance, relative to MoE-VRD, in Table 2.




FIGURE 4. mAP of the MoE-VRD approach having $N = 10$ experts, as a function of K during training. Note that performance drops after $K = 2$; due to the averaging nature of the architecture before the final output, such that well-performing experts may become drowned out by more poorly performing peers if K is set too large.

E. ABLATION STUDY

Really the only aspect of the proposed architecture which can be removed, via ablation, is the collection of experts. That is, our ablation study assesses performance as a function of the number of “top experts” K chosen by the gating function for each input. The total number of experts was fixed to $N = 10$, since $N > K$ needs to be large enough to test a meaningful range of K , at the same time increasing N far past 10 leads to challenges in network memory requirements and training reliability. The results of the multi-expert experiments are presented in Figure 4, plotting mAP K , ablating K from 6 down to 1.

The best MoE results are achieved when we select the two top experts ($K = 2$) for each input, such that the performance drops with increased K for both datasets. Note that a low optimum value of K is an asset, not a liability, in that a small K implies a modest computational complexity, since only K experts are actually engaged for any given input.

V. DISCUSSION AND CONCLUSION

The problem of video-based visual relationship detection (Vid-VRD) is relatively new compared to static image-based visual relationship detection. The spatio-temporal dimensions in the video domain cascade the difficulty of the problem, given the far greater data volumes and the ability for relationships to change over time. There have been a few approaches to address this problem, however they uniformly rely on monolithic neural networks [24], [26], [44].

In this work, we developed a new framework, the MoE-VRD, based on a mixture of experts approach. MoE-VRD is developed by encapsulating a Vid-VRD framework [15] into an expert within a sparsely gated mixture of experts architecture.

Our proposed approach to video visual relationship detection also addresses limitations in computing power and distributed computation, which arises from the limited capacity of neural networks to absorb information due to the limitations in network size (number of parameters) in comparatively blunt architectures based on a single, monolithic network.

We have observed that the performance of the network drops when we select more than two experts ($K > 2$). We believe that this stems from the averaging operation, which acts prior to the final output, resulting in well-performing experts being increasingly drowned out by those experts having inferior performance. Studying this effect more carefully, and in other settings, is one subject for future work.

We achieved highly promising results from MoE-VRD based on experiments on two different datasets, ImageNet-VRD [25] and VidOR [62]. Our MoE-VRD outperforms nearly all state-of-the-art approaches in most metrics, and outperforms every tested approach on the ImageNet-VRD dataset.

The proposed approach in this paper is perhaps still naive, as all of the experts are tackling the same problem. In principle one could imagine dividing the problem into

smaller subproblems (with certain experts only aimed at subject/object recognition, for example), to address multimodal datasets or to explore hierarchical MoE architecture, in which a primary gating network chooses a combination of experts — each of which itself is a secondary or tertiary MoE with its own respective set of experts and gating network [33].

Finally, almost certainly the gating network itself would benefit from further scrutiny. The gating network of this paper is the simplest possible choice, a single-layer feed-forward network, taking as input the same spatio-temporal object-tracklet features as are being provided to the experts. In a sense, it would seem that too much is being asked of the gating function, to go all the way from low-level input to expert-assessment output, such that the gating function tacitly must emulate or reproduce certain elements of expert behaviour. It would seem preferable to have the gating function operate at a higher / more abstracted level, and having certain aspects of expert-assessment made the responsibility of the expert networks themselves.

ACKNOWLEDGMENT

The authors acknowledge advice from Chong Luo (Microsoft Research Asia) and the Alliance Program of the Natural Science and Engineering Research Council of Canada.

REFERENCES

- [1] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool, "One-shot video object segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 221–230.
- [2] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung, "Learning video object segmentation from static images," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2663–2672.
- [3] P. Tokmakov, K. Alahari, and C. Schmid, "Learning video object segmentation with visual memory," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 4481–4490.
- [4] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool, "The 2017 Davis challenge on video object segmentation," 2017, *arXiv:1704.00675*.
- [5] R. Yan, J. Zhang, J. Yang, and A. G. Hauptmann, "A discriminative learning framework with pairwise constraints for video object classification," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 4, pp. 578–593, Apr. 2006.
- [6] L. Zhang, S. Z. Li, X. Yuan, and S. Xiang, "Real-time object classification in video surveillance based on appearance learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2007, pp. 1–8.
- [7] W. Han, P. Khorrani, T. Le Paine, P. Ramachandran, M. Babaeizadeh, H. Shi, J. Li, S. Yan, and T. S. Huang, "Seq-NMS for video object detection," 2016, *arXiv:1602.08465*.
- [8] K. Kang, W. Ouyang, H. Li, and X. Wang, "Object detection from video tubelets with convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 817–825.
- [9] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Large-scale video classification with convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 1725–1732.
- [10] C. Feichtenhofer, A. Pinz, and R. P. Wildes, "Spatiotemporal multiplier networks for video action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4768–4777.
- [11] J. Carreira and A. Zisserman, "Quo vadis, action recognition? A new model and the kinetics dataset," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 6299–6308.
- [12] Z. Liu and H. Hu, "Spatiotemporal relation networks for video action recognition," *IEEE Access*, vol. 7, pp. 14969–14976, 2019.
- [13] J. Zhang and H. Hu, "Deep spatiotemporal relation learning with 3D multi-level dense fusion for video action recognition," *IEEE Access*, vol. 7, pp. 15222–15229, 2019.
- [14] A. Ullah, J. Ahmad, K. Muhammad, M. Sajjad, and S. W. Baik, "Action recognition in video sequences using deep bi-directional LSTM with CNN features," *IEEE Access*, vol. 6, pp. 1155–1166, 2018.
- [15] X. Shang, Y. Li, J. Xiao, W. Ji, and T.-S. Chua, "Video visual relation detection via iterative inference," in *Proc. 29th ACM Int. Conf. Multimedia*. New York, NY, USA: Association for Computing Machinery, Oct. 2021, pp. 3654–3663, doi: 10.1145/3474085.3475263.
- [16] M. A. Goodale and A. D. Milner, "Separate visual pathways for perception and action," *Trends Neurosci.*, vol. 15, no. 1, pp. 20–25, 1992.
- [17] M. Warglien, P. Gärdenfors, and M. Westera, "Event structure, conceptual spaces and the semantics of verbs," *Theor. Linguistics*, vol. 38, nos. 3–4, pp. 159–193, Jan. 2012.
- [18] J. Arunehru, G. Chamundeeswari, and S. P. Bharathi, "Human action recognition using 3D convolutional neural networks with 3D motion cuboids in surveillance videos," *Proc. Comput. Sci.*, vol. 133, pp. 471–477, Jan. 2018.
- [19] Z. Gharraee, P. Gärdenfors, and M. Johnsson, "First and second order dynamics in a hierarchical SOM system for action recognition," *Appl. Soft Comput.*, vol. 59, pp. 574–585, Oct. 2017.
- [20] K. Muhammad, M. Khan, A. Ullah, A. S. Imran, M. Sajjad, M. S. Kiran, G. Sannino, and V. H. C. de Albuquerque, "Human action recognition using attention based LSTM network with dilated CNN features," *Future Gener. Comput. Syst.*, vol. 125, pp. 820–830, Dec. 2021.
- [21] Z. Gharraee, "Hierarchical growing grid networks for skeleton based action recognition," *Cogn. Syst. Res.*, vol. 63, pp. 11–29, Oct. 2020.
- [22] Z. Gharraee, P. Gärdenfors, and M. Johnsson, "Online recognition of actions involving objects," *Biol. Inspired Cogn. Archit.*, vol. 22, pp. 10–19, Oct. 2017.
- [23] I. Giorgi, A. Cangelosi, and G. L. Masala, "Learning actions from natural language instructions using an ON-world embodied cognitive architecture," *Frontiers Neurobot.*, vol. 15, p. 48, May 2021.
- [24] Y.-H.-H. Tsai, S. Divvala, L.-P. Morency, R. Salakhutdinov, and A. Farhadi, "Video relationship reasoning using gated spatio-temporal energy graph," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 10424–10433.
- [25] X. Shang, T. Ren, J. Guo, H. Zhang, and T.-S. Chua, "Video visual relation detection," in *Proc. 25th ACM Int. Conf. Multimedia*, Mountain View, CA USA, Oct. 2017, pp. 1300–1308.
- [26] X. Qian, Y. Zhuang, Y. Li, S. Xiao, S. Pu, and J. Xiao, "Video relation detection with spatio-temporal graph," in *Proc. 27th ACM Int. Conf. Multimedia*, Oct. 2019, pp. 84–93.
- [27] C. Liu, Y. Jin, K. Xu, G. Gong, and Y. Mu, "Beyond short-term snippet: Video relation detection with spatio-temporal global context," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10840–10849.
- [28] Z. Su, X. Shang, J. Chen, Y.-G. Jiang, Z. Qiu, and T.-S. Chua, "Video relation detection via multiple hypothesis association," in *Proc. 28th ACM Int. Conf. Multimedia*. New York, NY, USA: Association for Computing Machinery, Oct. 2020, pp. 3127–3135, doi: 10.1145/3394171.3413764.
- [29] J. Xiao, X. Shang, X. Yang, S. Tang, and T. Chua, "Visual relation grounding in videos," in *Computer Vision—ECCV 2020 (Lecture Notes in Computer Science)*, vol. 12351. Glasgow, U.K.: Springer, Aug. 2020, pp. 447–464.
- [30] Z. Gharraee, K. Holmquist, L. He, and M. Felsberg, "A Bayesian approach to reinforcement learning of vision-based vehicular control," in *Proc. 25th Int. Conf. Pattern Recognit. (ICPR)*, Jan. 2021, pp. 3947–3954.
- [31] A. S. Davis and I. Arel, "Low-rank approximations for conditional feedforward computation in deep neural networks," in *Proc. 2nd Int. Conf. Learn. Represent. (ICLR)*, 2014, pp. 1–14.
- [32] E. Bengio, P. Bacon, J. Pineau, and D. Precup, "Conditional computation in neural networks for faster models," *CoRR*, vol. abs/1511.06297, pp. 1–12, Nov. 2015.
- [33] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. V. Le, G. E. Hinton, and J. Dean, "Outrageously large neural networks: The sparsely-gated mixture-of-experts layer," in *Proc. 5th Int. Conf. Learn. Represent. (ICLR)*, 2017, pp. 1–19. [Online]. Available: <https://openreview.net/forum?id=B1ckMDqIq>
- [34] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar, "Do CIFAR-10 classifiers generalize to CIFAR-10?" 2018, *arXiv:1806.00451*.
- [35] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, "Overcoming catastrophic forgetting in neural networks," *Proc. Nat. Acad. Sci. USA*, vol. 114, no. 13, pp. 3521–3526, Mar. 2017.

- [36] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," *OpenAI Blog*, vol. 1, no. 8, p. 9, 2019.
- [37] X. Dong, Z. Yu, W. Cao, Y. Shi, and Q. Ma, "A survey on ensemble learning," *Frontiers Comput. Sci.*, vol. 14, no. 2, pp. 241–258, 2020.
- [38] Z. Yin, M. Zhao, Y. Wang, J. Yang, and J. Zhang, "Recognition of emotions using multimodal physiological signals and an ensemble deep learning model," *Comput. Methods Programs Biomed.*, vol. 140, pp. 93–110, Mar. 2017.
- [39] C. Lu, R. Krishna, M. Bernstein, and L. Fei-Fei, "Visual relationship detection with language priors," in *Proc. 14th Eur. Conf. Comput. Vis. (ECCV)*, Amsterdam, The Netherlands, Cham, Switzerland: Springer, Oct. 2016, pp. 852–869.
- [40] Y. Zhu, S. Jiang, and X. Li, "Visual relationship detection with object spatial distribution," in *Proc. IEEE Int. Conf. Multimedia Expo (ICME)*, Jul. 2017, pp. 379–384.
- [41] T. Yao, Y. Pan, Y. Li, and T. Mei, "Exploring visual relationship for image captioning," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 684–699.
- [42] K. Liang, Y. Guo, H. Chang, and X. Chen, "Visual relationship detection with deep structural ranking," in *Proc. 32nd AAAI Conf. Artif. Intell.* Palo Alto, CA, USA: AAAI Press, 2018, pp. 7098–7105. [Online]. Available: <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16491>
- [43] A. Kuznetsova, H. Rom, N. Alldrin, J. Uijlings, I. Krasin, J. Pont-Tuset, S. Kamali, S. Popov, M. Mallocci, A. Kolesnikov, T. Duerig, and V. Ferrari, "The open images dataset V4," *Int. J. Comput. Vis.*, vol. 128, no. 7, pp. 1956–1981, 2020.
- [44] Y. Liu, K. Chen, C. Liu, Z. Qin, Z. Luo, and J. Wang, "Structured knowledge distillation for semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2604–2613.
- [45] X. Wu, R. Wang, J. Hou, H. Lin, and J. Luo, "Spatial-temporal relation reasoning for action prediction in videos," *Int. J. Comput. Vis.*, vol. 129, no. 5, pp. 1484–1505, May 2021.
- [46] K. Gao, L. Chen, Y. Huang, and J. Xiao, "Video relation detection via tracklet based visual transformer," in *Proc. 29th ACM Int. Conf. Multimedia*, Oct. 2021, pp. 4833–4837.
- [47] S. Zheng, S. Chen, and Q. Jin, "VRDFormer: End-to-end video visual relation detection with transformers," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 18836–18846.
- [48] Y. Li, X. Yang, X. Shang, and T.-S. Chua, "Interventional video relation detection," in *Proc. 29th ACM Int. Conf. Multimedia*, Oct. 2021, pp. 4091–4099.
- [49] Q. Cao, H. Huang, M. Ren, and C. Yuan, "Concept-enhanced relation network for video visual relation inference," *IEEE Trans. Circuits Syst. Video Technol.*, early access, Nov. 9, 2022, doi: [10.1109/TCSVT.2022.3220426](https://doi.org/10.1109/TCSVT.2022.3220426).
- [50] K. Gao, L. Chen, Y. Niu, J. Shao, and J. Xiao, "Classification-then-grounding: Reformulating video scene graphs as temporal bipartite graphs," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 19497–19506.
- [51] S. Chen, Z. Shi, P. Mettes, and C. G. M. Snoek, "Social fabric: Tubelet compositions for video relation detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 13485–13494.
- [52] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton, "Adaptive mixtures of local experts," *Neural Comput.*, vol. 3, no. 1, pp. 79–87, Mar. 1991.
- [53] J. Ma, Z. Zhao, X. Yi, J. Chen, L. Hong, and E. H. Chi, "Modeling task relationships in multi-task learning with multi-gate mixture-of-experts," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2018, pp. 1930–1939.
- [54] X. Wang, F. Yu, L. Dunlap, Y.-A. Ma, R. Wang, A. Mirhoseini, T. Darrell, and J. E. Gonzalez, "Deep mixture of experts via shallow embedding," in *Proc. Uncertainty Artif. Intell.*, 2020, pp. 552–562.
- [55] Y. Shi, N. Siddharth, B. Paige, and P. Torr, "Variational mixture-of-experts autoencoders for multi-modal deep generative models," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 15718–15729.
- [56] M. Ryabinin and A. Gusev, "Towards crowdsourced training of large neural networks using decentralized mixture-of-experts," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 3659–3672.
- [57] C. Riquelme, J. Puigcerver, B. Mustafa, M. Neumann, R. Jenatton, A. S. Pinto, D. Keysers, and N. Houlsby, "Scaling vision with sparse mixture of experts," in *Proc. Adv. Neural Inf. Process. Syst., Annu. Conf. Neural Inf. Process. Syst. (NeurIPS)*, vol. 34, Dec. 2021, pp. 8583–8595.
- [58] D. Eigen, M. Ranzato, and I. Sutskever, "Learning factored representations in a deep mixture of experts," in *Proc. 2nd Int. Conf. Learn. Represent. (ICLR)*, 2014, pp. 1–8.
- [59] T. Kohonen, "The self-organizing map," *Proc. IEEE*, vol. 78, no. 9, pp. 1464–1480, Sep. 1990.
- [60] J. Lu, L. Li, and C. Zhang, "Self-reinforcing unsupervised matching," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 8, pp. 4404–4418, Aug. 2022.
- [61] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-ResNet and the impact of residual connections on learning," in *Proc. 31st AAAI Conf. Artif. Intell.* Palo Alto, CA, USA: AAAI Press, 2017, pp. 4278–4284. [Online]. Available: <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14806>
- [62] X. Shang, D. Di, J. Xiao, Y. Cao, X. Yang, and T.-S. Chua, "Annotating objects and relations in user-generated videos," in *Proc. Int. Conf. Multimedia Retr.*, Jun. 2019, pp. 279–287.
- [63] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li, "YFCC100M: The new data in multimedia research," *Commun. ACM*, vol. 59, no. 2, pp. 64–73, 2016.
- [64] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, Dec. 2015.
- [65] Q. Cao, H. Huang, X. Shang, B. Krause, and T.-S. Chua, "3-D relation network for visual relation recognition in videos," *Neurocomputing*, vol. 432, pp. 91–100, Apr. 2021.



ALA SHAABANA received the degree from McMaster University, in 2017. He joined VMware Inc., to perform research and development on distributed computing, where he researched unstructured data analysis, natural language processing, and predictive scalability of large-scale systems. In 2020, he joined the University of Waterloo as a Postdoctoral Researcher looking into video visual relationship detection. He is an expert in machine learning, data analytics, and distributed and scalable systems. Currently, he is with the Opentensor Foundation researching incentivized artificial intelligence.



ZAHRA GHARAEI received the B.Sc. degree in electrical control engineering and the M.Sc. degree in mechatronics from K. N. Toosi University of Technology, Iran, in 2009 and 2012, respectively, and the Ph.D. degree in cognitive science from the Department of Philosophy and Cognitive Science, Lund University, Sweden, in 2018. She was a Postdoctoral Researcher with the Computer Vision Laboratory (CVL), Department of Electrical Engineering, Linköping University, Sweden, from 2018 to 2022. She is currently a Postdoctoral Research Fellow with the Vision and Image Processing Laboratory (VIP Lab), Department of Systems Design Engineering, University of Waterloo, Canada. She has been involved in different research projects with WASP, EU-FET WYSIWYD, Microsoft, Miovision, and BioScan. Her research interests include artificial intelligence, machine learning, and computational cognitive science.



PAUL FIEGUTH (Senior Member, IEEE) received the B.A.Sc. degree from the University of Waterloo, Canada, in 1991, and the Ph.D. degree from Massachusetts Institute of Technology, in 1995, all in electrical engineering. He joined the University of Waterloo as a Faculty Member, in 1996, where he is currently a Professor of systems design engineering. He is the Co-Director of the Vision and Image Processing Research Group, where his research interests broadly involve machine learning for computer vision and statistical image processing. He is the author of a 2010 textbook on statistical image processing, a 2021 textbook on complex systems, and an upcoming textbook on pattern recognition and machine learning. His specific research interests include hierarchical algorithms for large problems, particularly in simplifying modeling and interpretation.

...