

# FloCoDe: Unbiased Dynamic Scene Graph Generation with Temporal Consistency and Correlation Debiasing

Anant Khandelwal  
Glance AI

anant.iitd.2085@gmail.com

## Abstract

*Dynamic scene graph generation (SGG) from videos requires not only a comprehensive understanding of objects across scenes but also a method to capture the temporal motions and interactions with different objects. Moreover, the long-tailed distribution of visual relationships is a crucial bottleneck for most dynamic SGG methods. This is because many of them focus on capturing spatio-temporal context using complex architectures, leading to the generation of biased scene graphs. To address these challenges, we propose FLOCODE: **Flow**-aware Temporal Consistency and **Correlation Debiasing** with uncertainty attenuation for unbiased dynamic scene graphs. FLOCODE employs feature warping using flow to detect temporally consistent objects across frames. To address the long-tail issue of visual relationships, we propose correlation debiasing and a label correlation-based loss to learn unbiased relation representations for long-tailed classes. Specifically, we propose to incorporate label correlations using contrastive loss to capture commonly co-occurring relations, which aids in learning robust representations for long-tailed classes. Further, we adopt the uncertainty attenuation-based classifier framework to handle noisy annotations in the SGG data. Extensive experimental evaluation shows a performance gain as high as 4.1%, demonstrating the superiority of generating more unbiased scene graphs.*

## 1. Introduction

Scene graph generation for videos (VidSGG) aims to represent the video in the form of a dynamic graph that is able to capture the temporal evolution of the relationships between pairs of objects. VidSGG has direct applications for various downstream tasks, such as visual question answering [1, 43, 52], video captioning [53], and video retrieval [9, 39, 50], etc. VidSGG is considered more challenging compared to its image-based counterpart, as the relations between identified object pairs are dynamic along the temporal dimension, making it a multi-label problem. At the current stage, VidSGG is still in its nascent phase compared

to SGG (scene graph generation) for static images [8, 22–24, 27, 30, 41, 56, 57].

Several works [7, 12, 18, 26, 28, 35, 45] have proposed solutions for VidSGG, mostly employing spatio-temporal sequence processing with transformers [3, 4, 13, 19, 32, 40, 48]. However, many of these methods focus on building complex models to effectively aggregate spatio-temporal information in a video, without adequately addressing data imbalance in the relation/predicate classes. While their performance is often good in terms of the Recall@K metric, which is biased towards frequent classes, a more comprehensive metric, mean-Recall@K, has been proposed [5, 41] to assess performance in the presence of low-frequency classes, providing an overall view of SGG models rather than considering only high-frequency classes.

Although recent methods [31, 34] have proposed addressing class imbalance using memory-based debiasing and uncertainty attenuation for classification, the memory debiasing approach is based on learnable attention weights, which carries a risk of bias towards high-frequency classes. A case of this issue is illustrated in Figure 2 during a qualitative comparison with our method. To overcome these limitations, some works [25, 54] have attempted to address biased relation predictions. [25] proposed weakening the false correlation between input data and predicate labels, while [54] considered biases in a meta-learning paradigm. Although these approaches mitigate the long-tail problem to some extent, the performance is still not satisfactory.

Our analysis (see Table 4) pinpoints a significant issue in existing dynamic Scene Graph Generation (SGG): the inaccurate detection of objects across video frames. Owing to this issue, we propose leveraging flow-warped features in the temporal dimension to handle dynamic fluctuations in videos. Unlike previous methods [31] that employ learnable read from memory, we propose debiasing predicate embeddings during the generation stage. This is achieved by ensuring an unbiased correlation between predicates and objects. Additionally, existing methods [31] overlook label correlations, which can provide valuable clues during relation classification, especially for long-tailed predi-

cate classes. We propose supervised multi-label contrastive loss to consider label correlations. This aims to bring together predicate representations with overlapping classes and push apart negative samples without shared classes. Further, we introduce a Mixture of Logit Networks (MLNs) to calculate aleatoric (data) and epistemic (model) uncertainty<sup>1</sup>. We propose an uncertainty-aware regularization method aimed at improving the capacity to comprehend inconsistent output patterns, spanning both clean and noisy distributions. This method not only enables robust learning of the fundamental target distribution but also skillfully captures the injected noise patterns inherent in the data generation process. Combining these elements, our framework, FLOCODE (**F**low-aware Temporal Consistency and **C**orrelation **D**ebiasing with Uncertainty Attenuation), makes significant contributions: 1) We propose modeling both aleatoric and epistemic uncertainty and label correlations for more unbiased scene graphs. 2) Introducing a novel correlation-guided debiased learning of predicate embeddings. 3) Incorporating flow-aware, temporally consistent object detection for precise node classification in scene graphs. 4) Achieving substantial improvements in mR@K [41] and R@K, showcasing its superiority in generating unbiased scene graphs.

## 2. Related Work

**Image Scene Graph Generation:** Image-based scene graph generation (ImgSGG) refers to the task of generating a structured graph summary that represents objects as nodes and their relationships (formally known as predicates) as edges within an image. Extensive efforts have been dedicated to ImgSGG, particularly in comparison to the Visual Genome (VG) benchmark [22]. Previous research has focused on developing efficient approaches for aggregating spatial context [24, 27, 30, 56, 57]. Concurrently, recent studies [41, 42, 51] have concentrated on addressing foundational challenges, such as mitigating biased scene graphs resulting from long-tailed predicate distributions and dealing with noisy annotations in datasets.

**Video Scene Graph Generation (VidSGG):** Researchers have made significant strides in understanding spatial context within images, prompting an exploration of spatial context and temporal correlations between objects in videos. In the realm of video scene graph generation (VidSGG), challenges akin to those in image-based scene graph generation persist, including long-tailed predicates and noisy annotations. Action Genome [17], a prominent VidSGG benchmark, introduces an additional hurdle of addressing

<sup>1</sup>Aleatoric uncertainty pertains to the intrinsic and unavoidable aspects of uncertainty within the data generation process, such as measurement noise. Conversely, epistemic uncertainty encompasses the uncertainty associated with the model itself, and this type of uncertainty may diminish with an increase in the amount of training data available.

temporal fluctuations across frames. Various approaches [28, 35, 45, 47, 58] have employed object-tracking mechanisms to manage temporal fluctuations, but these models often come with high computational costs, memory consumption, and performance issues due to information accumulation from irrelevant frames. STTran [7] offers a robust baseline, using a spatial encoder and a temporal decoder to implicitly extract spatial-temporal contexts. Other works [26, 49] focus on extracting temporal correlations, ensuring temporal continuity through co-occurrence patterns or proposing pre-training paradigms. Building on the success of transformer-based models [3, 32, 40, 48], which excel in sequence processing for spatio-temporal context, there remains a challenge of bias towards high-frequency classes. TEMPURA [31] addresses long-tail bias with an uncertainty-guided loss function. In our contribution, we extend this by exploring label correlation to further enhance predictive accuracy in VidSGG.

## 3. Method

**Problem Statement:** Given a video  $\mathcal{V} = \{I_1, I_2, I_3, \dots, I_T\}$ , the goal of dynamic SGG is to generate scene graphs denoted as  $\mathcal{G} = \{G_t\}_{t=1}^T$  of video  $\mathcal{V}$  consisting of  $T$  frames.  $G_t = \{V_t, E_t\}$  is the scene graph of frame  $I_t$ , where  $V_t$  is the set of nodes and  $E_t$  is the set of edges representing relations between nodes in  $V_t$ . Nodes in  $V_t$  are connected to each other using predicates in  $E_t$ , forming multiple  $\langle \text{subject-predicate-object} \rangle$  triplets. The sets of object and predicate classes are referred to as  $\mathcal{Y}_o = \{y_{o_1}, y_{o_2}, y_{o_3}, \dots, y_{o_{c_o}}\}$  and  $\mathcal{Y}_r = \{y_{r_1}, y_{r_2}, y_{r_3}, \dots, y_{r_{c_r}}\}$  respectively.

**Object Detection and Relation Representation:** With the use of an off-the-shelf object detector (Faster-RCNN [36]), we obtain the set of objects  $O_t = \{o_i^t\}_{i=1}^{N(t)}$ , where  $N(t)$  is the number of objects detected in frame  $I_t$ . Each object in a  $t^{\text{th}}$  frame is denoted as  $o_i^t = \{b_i^t, v_i^t, c_{o_i}^t\}$  where  $b_i^t \in \mathbb{R}^4$  being the bounding box,  $v_i^t \in \mathbb{R}^{2048}$  the RoIAligned [15] proposal feature of  $o_i^t$  and  $c_{o_i}^t$  is its predicted class. However, the object class  $c_{o_i}^t$  fluctuates across the frames and is not coherent even for the same object. Existing works [45] address this by incorporating object tracking algorithms; opposed to this, our strategy compensates for these dynamic fluctuations using flow-warped features and ensures temporal coherence. Specifically, we extracted the base features  $f_t, t \in [1, T]$  and ROIs from Faster R-CNN using ResNet-101 [14]. We warp these base features using the temporal flow and compute the RoIAligned warped object features  $v_i^{t \rightarrow t'}$  ( $t'$  represents the immediate previous frame that contains the  $i^{\text{th}}$  object) using the predicted ROIs.

### 3.1. Temporal Flow-Aware Object Detection

Object detectors trained on static images are prone to misclassify the same object in different frames. Existing meth-

ods [7, 18, 26, 49] either use  $c_{o_i}^t$  obtained from object detection in each frame or use object features to classify objects. However, these methods do not compensate for temporal fluctuations in the videos. Inspired by FGFA[59], which uses flow-guided feature aggregation for object detection in videos, we propose to leverage flow-warped features and temporal processing for consistent object detection across frames. We introduce *Temporal Flow-Aware Object Detection (TFoD)*, which utilises transformer encoder [48] with masked self-attention (*TEnc*) to process the set of temporal object sequences  $\mathcal{O}_{\mathcal{V}}$ , which is constructed as follows:

$$\mathcal{O}_{\mathcal{V}} = \{\mathcal{O}_{t_1, k_1}^1, \mathcal{O}_{t_2, k_2}^2, \dots, \mathcal{O}_{t_{\hat{C}_o}, k_{\hat{C}_o}}^{\hat{C}_o}\}, \text{ where} \\ \mathcal{O}_{t_j, k_j}^j = \{v_i^{t_j}, v_i^{t_j+1}, \dots, v_i^{k_j}\} \quad (1)$$

Each entry of  $\mathcal{O}_{t_j, k_j}^j$  corresponds to an object of the same detected class  $c_{o_j}$ , here  $1 \leq t_j, k_j \leq T$  and  $\hat{C}_o \leq C_o$  denoting all the detected classes in a video  $\mathcal{V}$ . However, the detected class labels can be noisy since they are based only on frame-level predictions; hence, we use the flow-warped feature  $v_i^{t \rightarrow t'}$  instead of  $v_i^t$  before feeding it to the transformer encoder. The flow-warped feature is computed as:

$$f_{t \rightarrow t'} = \mathcal{W}(f_t, \mathcal{F}(I_{t'}, I_t)) \quad (2)$$

where  $\mathcal{W}$  is a bilinear warping function [59, 60] applied on all the locations for each channel in the feature maps and flow field  $\mathcal{F}(I_{t'}, I_t)$  is computed from the pre-trained Flow-Net [10], where  $t'$  is the index of the immediate frame previous to the  $t^{\text{th}}$  frame having the same object as depicted in the object sequences  $\mathcal{O}_{\mathcal{V}}$ . Using  $f_{t \rightarrow t'}$  and predicted ROIs, we obtained warped RoIAligned feature,  $\mathcal{O}_{t_j, k_j}^j$  is given as:

$$\mathcal{O}_{t_j, k_j}^j = \{v_i^{t_j}, v_i^{t_j+1 \rightarrow t_j}, \dots, v_i^{k_j}\} \quad (3)$$

Each of  $\mathcal{O}_{t_j, k_j}^j$  is zero-padded to prepare functional input tensor. *Tenc* uses masked multi-head self-attention instead of multi-head self-attention in transformer encoder [48]. Mask is introduced to learn the temporal dependencies in a unidirectional manner so that the object at frame index  $t$  can only attend to objects in previous frames. For any input  $\mathbf{X}$ , the single head masked attention  $\mathbb{A}$  is given as:

$$\mathbb{A}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax} \left( \frac{\text{mask}(\mathbf{Q}\mathbf{K}^T)}{\sqrt{D_K}} \right) \mathbf{V} \quad (4)$$

where  $D_K$  is the dimension of  $\mathbf{K}$ , and  $\mathbf{Q}, \mathbf{K}, \mathbf{V}$  are the query, key, and value vectors, respectively. Here,  $\mathbf{Q} = \mathbf{K} = \mathbf{V} = \mathbf{X}$ , the multi-head attention is  $\text{concat}(a_1, a_2, \dots, a_H)W_H$ , where  $a_i = \mathbb{A}(\mathbf{X}W_{Q_i}, \mathbf{X}W_{K_i}, \mathbf{X}W_{V_i})$  where  $W_{Q_i}, W_{K_i}, W_{V_i}$  and  $W_H$  are the learnable weight matrices. The rest of the components, like residual connection, normalisation, and FFN (feed-forward network),

remain the same as in the transformer encoder [48]. The output of  $n$ -layered *TEnc* is given as:

$$X_{out}^{(n)} = \text{TEnc}(X_{out}^{(n-1)}), X_{out}^{(0)} = \hat{\mathcal{O}}_{\mathcal{V}} \quad (5)$$

where  $\hat{\mathcal{O}}_{\mathcal{V}} = \mathcal{O}_{\mathcal{V}} + P_o^T$ , where  $P_o^T$  are the fixed positional embeddings injecting the temporal position of objects. Inspired by the properties of neural collapse [33] we prefixed the classifier weights (forming Equiangular Tight Frame (ETF)) for each object class to induce the maximal separable classifier even under the class imbalance setting. The pre-fixed classifier weights  $\mathbf{W}_{ETF}$  are given as:

$$\mathbf{W}_{ETF} = \sqrt{\frac{C_o}{C_o - 1}} \mathbf{U} \left( \mathbf{I}_{C_o} - \frac{1}{C_o} \mathbf{1}_{C_o} \mathbf{1}_{C_o}^T \right) \quad (6)$$

where  $\mathbf{W}_{ETF} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{C_o}] \in R^{d \times C_o}$ ,  $\mathbf{U} \in R^{d \times C_o}$ , satisfies  $\mathbf{U}^T \mathbf{U} = \mathbf{I}_{C_o}$ ,  $\mathbf{I}_{C_o}$  is the identity matrix, and  $\mathbf{1}_{C_o}$  is an all-ones vector. The object classification loss:

$$\mathcal{L}_o(x_{o_i}, \mathbf{W}_{ETF}) = \frac{1}{2} \left( \mathbf{w}_{c_{o_i}}^T \hat{x}_{o_i} - 1 \right) \quad (7)$$

where  $\hat{x}_{o_i} = x_{o_i} / \|x_{o_i}\|$  and  $x_{o_i} \in X_{out}^{(n)}$  and  $\mathbf{w}_{c_{o_i}}$  is the fixed prototype in  $\mathbf{W}_{ETF}$  for object class  $c_{o_i}$  and we have  $\|\mathbf{w}_{c_{o_i}}\| = 1$ . Finally, the converged features will be aligned with  $\mathbf{W}_{ETF}$ , and thus the ETF structure instructed by neural collapse is attained. The theoretical construction of the loss has been proved in [55].

### 3.2. Correlation-Aware Predicate Embedding

The relationship between objects is governed by three types of correlations: a) *spatial correlation between predicates*, b) *temporal correlation between predicates*, and c) *predicate-object correlation across the video frames*. We propose to model these correlations using the Vanilla Transformer [48]. For each object  $o_i$  of predicted class  $c_{o_i}$  obtained from object detection (Section 3.1), the input to the transformer encoder is the set of features describing the relation with each object  $o_j$  detected in all the frames where  $o_i$  is detected. The input is constructed as follows:

$$r_{i,j}^t = \text{concat}(\mathbf{x}_{c_{o_i}}, f_u(\mathbf{u}_{i,j}^t + f_{box}(\mathbf{b}_i^t, \mathbf{b}_j^t)), f_I(t)) \quad (8)$$

where  $\mathbf{x}_{c_{o_i}} \in X_{out}^{(n)}$  is the feature representation of object  $o_i$  belonging to class  $c_{o_i} \in [1, C_o]$ ,  $u_{i,j}^t \in \mathbb{R}^{256 \times 7 \times 7}$  is the feature map of the union box computed by RoIAlign [15].  $f_u, f_I$  is the FFN based on non-linear projections, and  $f_{box}$  is the bounding box to feature map projection of [56]. Both are configured to produce  $d$ -dimensional relation features.  $f_I$  serves as positional embeddings denoting the frame index. The single encoder input consists of both spatial and temporal relation features between object  $o_i$  and all other objects  $o_1, o_2, \dots, o_j$ . Specifically, it is denoted as  $R_i = r_{i,1}^{t_1}, r_{i,2}^{t_2}, \dots, r_{i,j}^{t_j}$ , where  $t_j$  are the frame indices in which

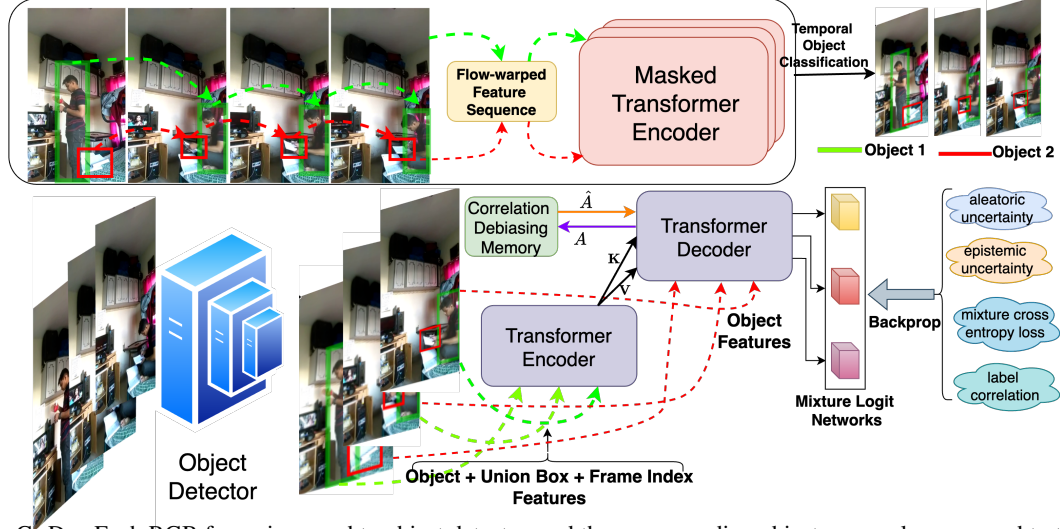


Figure 1. FLOCODE: Each RGB frame is passed to object detector, and the corresponding object proposals are passed to temporal flow-aware object classification. Finally, we predict relations (with correlation debiasing, Mixture Logit Network for uncertainty estimation, mixture loss and label correlation loss) from relation representation at the output of transformer decoder to generate unbiased scene graphs

$o_i$  and  $o_j$  are detected simultaneously. The transformer decoder leverages masked self-attention, and its input is a set of object representations,  $\{\mathbf{x}_{c_{o_1}}, \mathbf{x}_{c_{o_2}}, \dots, \mathbf{x}_{c_{o_j}}\}$ , corresponding to objects  $\{o_1, o_2, \dots, o_j\}$ , detected in all the frames where  $o_i$  is detected. Since the input to the transformer encoder contains all predicate (relation) features across the frames, and hence, with multi-head self-attention, it models both spatial and temporal correlations between predicates. Similarly, the cross-attention between the encoder and decoder models the predicate-object correlation. The predicate embeddings at the output of the transformer decoder are denoted as  $\hat{r}_{tem}^k = \{\hat{r}_{i,j}^t\} \forall k \in [1, N(t)], t \in [1, T]$ . At transformer decoder, we use the sliding window of size 10 for predicate representation with relating objects.

### 3.3. Debaised Predicate Embedding

Since the relations between objects are highly imbalanced, the representation becomes biased towards popular ones. Therefore, to produce unbiased relation embeddings, we propose updating correlation matrices as a weighted average of the correlation matrices across the epochs, where the weight is determined by the decay factor. For each triplet  $\{o_i, r_{i,j}, o_j\}$  the softmax attention scores between the transformer encoder and decoder across all the layers are stored. As some predicates are rare and thus more prone to bias, updating their correlation using a running average will generate debaised embeddings. Let's denote the running average as the stored correlation matrix between every object pair for all relations as  $\mathcal{M}_{e-1}$  at the end of the previous epoch  $e - 1$ , and the attention matrix (from Transformer cross-attention) at the current ongoing epoch  $e$  as  $A_e$ . During training, we update the attention matrix (denoted as  $\hat{A}_e$ ) using decaying factor  $\eta$  as the training progresses, given as:

$$\hat{A}_e(o_i, r_{i,j}, o_j) = \eta * A_e(o_i, r_{i,j}, o_j) + (1 - \eta) * \mathcal{M}_{e-1}(o_i, r_{i,j}, o_j) \quad (9)$$

where  $\mathcal{M}_0 = A_0$ ,  $\mathcal{M}_e(o_i, r_{i,j}, o_j) = \hat{A}_e(o_i, r_{i,j}, o_j)$  at the end of each epoch. We then use the attention value from  $\hat{A}_e$  in place of the attention value calculated from  $QK^T \in A_e$ . This avoids biasing attention weights towards popular predicates. Once the debaised weights are learned during training, we expect that they will generate debaised embeddings during the inference phase. Hence, there is no modification of attention matrices during inference time.

### 3.4. Predicate Classification

Following [31], we adopt classifier framework to handle noisy annotations. Different from [31], we propose: 1) an uncertainty-aware mixture of attenuated loss 2) supervised contrastive learning, which incorporates label correlation to improve predicate classification. Specifically, we propose the classifier head as a mixture-of-experts model named mixture logit networks (MLN) and a noise pattern estimation method utilizing the outputs of the MLN. The number of mixtures is denoted as  $\mathcal{K}$ .

**Uncertainty-Aware Mixture of Attenuated loss:** For a predicate embedding  $\mathbf{z}_i$ , the aleatoric ( $\sigma_a$ ) and epistemic uncertainty ( $\sigma_e$ ) are computed as:

$$\sigma_e^2 = \sum_{p=1}^{C_r} \sum_{k=1}^{\mathcal{K}} \pi_{i,p}^k \|\mu_{i,p}^k\|_2^2 - \sum_{j=1}^{\mathcal{K}} \pi_{i,p}^j \mu_{i,p}^j \|\mu_{i,p}^j\|_2^2 \quad (10)$$

$$\sigma_a^2 = \sum_{p=1}^{C_r} \sum_{k=1}^{\mathcal{K}} \pi_{i,p}^k \Sigma_{i,p}^k \quad (11)$$

where the mean  $\mu_{i,p}$ , variance  $\Sigma_{i,p}^k$ , and mixture weights  $\pi_{i,p}^k$  are the logits of label  $p$  in  $k$ th mixture. These are estimated as follows:

$$\mu_i^k = f_\mu^k(\mathbf{z}_i), \Sigma_i^k = \sigma(f_\Sigma^k(\mathbf{z}_i)), \pi_i^k = \frac{e^{f_\pi^k(\mathbf{z}_i)}}{\sum_{k=1}^{\mathcal{K}} e^{f_\pi^k(\mathbf{z}_i)}} \quad (12)$$

where  $f_\mu, f_\Sigma, f_\pi$  are the FFN projection functions and  $\sigma$  is the sigmoid non-linearity which ensures  $\Sigma_{i,p}^k \geq 0$  for the  $p^{\text{th}}$  predicate class. During training,  $\mathbf{z}_i = \hat{r}_{tem}^i$ , the mixture of attenuated loss ( $\mathcal{L}_{MAL}$ ) is given as:

$$\mathcal{L}_{MAL} = \frac{1}{N} \sum_{i=1}^N \sum_{p=1}^{\mathcal{C}_r} \sum_{k=1}^{\mathcal{K}} \pi_{i,p}^k \frac{\mathcal{L}(\mu_{i,p}^k, y_{r_p}^i)}{\Sigma_{i,p}^k} \quad (13)$$

where  $\mathcal{L}(\mu_{i,p}^k, y_{r_p}^i)$  is the sigmoidal cross-entropy loss,  $y_{r_p}^i$  is the ground-truth predicate class mapped to  $\mathbf{z}_i$ ,  $\mu_{i,p}^k$  is the logit of label  $p$  in the  $k^{\text{th}}$  mixture. For the corrupted input, it is more likely to make a false prediction; hence,  $\Sigma_{i,p}^k$  will increase to reduce the overall loss function for such instances. This, in turn, prevents over-fitting to the corrupted instances, making the model more robust.

#### Uncertainty-aware Supervised Contrastive Learning:

The MAL loss function, which independently classifies labels, presents challenges in capturing correlations among co-occurring semantic labels. To address this limitation, we introduce the multi-label contrastive loss, denoted as  $\mathcal{L}_{MCL}$ . The primary objective of this loss function is to minimize the distance between representations of predicates sharing at least one class with the anchor representation  $\hat{r}_{tem}^n$  while maximizing the distance from negative samples that do not share any classes. Let  $\mathcal{A}(n) = \{m \in \{N \setminus n\} : \mathcal{Y}_r^n \cdot \mathcal{Y}_r^m \neq 0, \text{ represent the positive set, where } \cdot \text{ denotes the dot product}\}$ . This set comprises samples that share at least one label with the anchor  $\hat{r}_{tem}^n$ . Additionally, let  $\mathcal{Y}_r(n, m) = \{y_{r_p}^{\{n,m\}} \in \mathcal{Y}_r^{\{n,m\}} \text{ s.t. } y_{r_p}^n = y_{r_p}^m = 1\}$  indicate the indices of samples  $m$  that possess at least one shared label with  $\mathcal{Y}_r^n$ . The loss  $\mathcal{L}_{MCL}$  is formulated as:

$$\mathcal{L}_{MCL} = \frac{1}{N} \sum_{n=1}^{n=N} \frac{-1}{|\mathcal{A}(n)|} \sum_{m \in \mathcal{A}(n)} J(n, m) \sum_{y_{r_p} \in \mathcal{Y}_r(n, m)} \left( \log \frac{\exp(\rho_{y_{r_p}}^{n,m} / \tau)}{\sum_{i \in N \setminus n} \exp(\rho_{y_{r_p}}^{n,i} / \tau)} \right) \quad (14)$$

where  $J(n, m)$  is the Jaccard similarity between labels  $y_{r_p}^n$  and  $y_{r_p}^m$ , the similarity  $\rho_{y_{r_p}}^{n,i}$  is given as:

$$\rho_{y_{r_p}}^{n,i} = \left( \prod_{k=1}^{\mathcal{K}} \left( \frac{(\Sigma_{n,p}^k)^2 + (\Sigma_{i,p}^k)^2}{2(\Sigma_{n,p}^k)(\Sigma_{i,p}^k)} \right)^{-\frac{1}{2}} \right) \exp \left( -\frac{1}{4} \sum_{k=1}^{\mathcal{K}} \frac{(\mu_{n,p}^k - \mu_{i,p}^k)^2}{(\Sigma_{n,p}^k)^2 + (\Sigma_{i,p}^k)^2} \right) \quad (15)$$

Specifically,  $\rho_{y_{r_p}}^{n,i}$  represents the Bhattacharyya coefficient, a commonly used metric for evaluating the similarity between probability distributions in various domains, including computer vision, pattern recognition, and statistical analysis[6, 38].

**EMA Teacher:** During training, we adopt the EMA weight update [2, 20, 44, 46] for transformers in Section 3.2. Let's say  $\phi_T, \theta_T$  are the weights of transformers for teacher and student, respectively. The weight update is then given as:

$$\phi_{T,e} = \alpha * \phi_{T,e-1} + (1 - \alpha) * \theta_{T,e} \quad (16)$$

where  $e$  is the training epoch. The EMA teacher effectively an ensemble of student models at different training steps, which is a most widely used learning strategy in semi-supervised setting[11, 16].

### 3.5. Training and Testing

**Training:** The entire framework is trained end-to-end minimizing the loss equation:

$$\mathcal{L} = \mathcal{L}_o + \mathcal{L}_{MAL} + \mathcal{L}_{MCL} - \lambda_1 \sigma_e + \lambda_2 \sigma_a \quad (17)$$

where  $\lambda_1, \lambda_2$  determines the amount of regularization for aleatoric and epistemic uncertainties respectively. This prevents  $\Sigma_{i,p}^k$  to grow indefinitely to minimise  $\mathcal{L}_{MAL}$  loss. Further we regularize with  $\sigma_e$  to increase epistemic uncertainty, thereby encouraging the utilization of a more mixtures.

**Testing:** During testing we utilize the EMA teacher  $\phi_T$  to generate the predicate embeddings  $\hat{r}_{tem}^i$ . The debiasing of predicate embeddings is only limited to training. These predicate embeddings are then passed to MLN which outputs the predicate confidence scores from all mixture components. The combined predicate confidence score  $\hat{y}_{r_p}^i$  from  $\mathcal{K}$  mixtures is calculated as:

$$\hat{y}_{r_p}^i = \sum_{k=1}^{\mathcal{K}} \pi_{i,p}^k \frac{\mu_{i,p}^k}{\Sigma_{i,p}^k} \quad (18)$$

## 4. Experiments

**Dataset:** Following previous works [7, 31], we also evaluated our method on the most widely used benchmark, Action Genome [17]. Action Genome is the largest benchmark for video SGG; it is built on top of Charades[37]. It contains 476,229 bounding boxes for 35 object classes (without person) and 1,715,568 instances of 26 predicate classes annotated for 234,253 frames. For all experiments, we use the same training and test split following [7, 26, 31].

**Metrics and Evaluation Setup:** We evaluated the performance of FLoCoDe with popular metrics, namely, recall@K (i.e., R@K) and mean-recall@K (i.e., mR@K), for  $K = [10, 20, 50]$ . R@K measures the ratio of correct instances among the top-K predicted instances with the highest confidence, but this is biased towards frequent predicate

classes [41], whereas mR@K averages out the R@K over all relationships. Hence, mR@K is a more reliable metric for balanced evaluation across predicates [41].

**Tasks:** Following previous works [7, 17, 22, 45], we also evaluated our method on three different experimental tasks: 1) **Predicate Classification (PREDCLS)**: predict the predicate class of object pairs, given the ground-truth bounding boxes and labels of objects. 2) **Scene graph classification (SGCLS)**: predict both predicate labels and the category labels of objects given the bounding boxes of objects. 3) **Scene graph detection (SGDET)**: simultaneously detects objects appearing in the frame and the predicate labels of each object pair in a frame. Following, we also evaluated our method using two setups: a) **With Constraint** and b) **No Constraints**. Later one allows each object pair to have more than one predicates simultaneously while the former one restricts to only one predicate.

**Implementation details:** Following previous works [7, 17, 22, 31, 45], we adopted Faster R-CNN [36] with a ResNet-101 [14] backbone as the object detector. The object detector was trained on the training set of Action Genome [17], resulting in a 24.6 mAP at 0.5 IoU with CoCo metrics. To ensure a fair comparison, we utilized this detector across all the baselines. Per-class non-maximal suppression at 0.4 IoU (intersection over union) was applied, following previous works [7, 17, 22, 31, 45], to reduce region proposals provided by the Region Proposal Network (RPN). The parameters of the object detector (excluding the object classifier) remained fixed during training when training scene graph generation models. For correlation-aware predicate embedding, it is necessary to match object pairs across frames. In cases with multiple objects in the same category, we use the IoU between the two objects across different images to match the subject-object pair. The IoU between the bounding box of the object in the previous frame and the object of the same category in the next frame is calculated. If the IoU is higher than 0.8, they are considered the same object. In the event of multiple candidates, the one with the highest IoU is chosen. An AdamW optimizer [29] is employed with a batch size of 1 and an initial learning rate of  $2e^{-5}$ . The number of mixture components  $\mathcal{K}$  is set to 4 for SGCLS and 6 for PREDCLS and SGDET. The self-attention and cross-attention layers in our framework have 8 heads with  $d = 1536$ , and dropout is set to 0.1. Regularizer hyper-parameters are set as  $\lambda_1 = 1, \lambda_2 = 1$ . For debiased predicate embedding, initial  $\eta$  is small i.e. 0.1 and we further reduce it with patience of 3. For EMA teacher update,  $\alpha = 0.999$  is used. All experiments are conducted on a single NVIDIA RTX-3090.

#### 4.1. Comparison with state-of-the-art

We compared our method FLOCODE with several state-of-the-art methods for dynamic SGG, namely TEMPURA

[31], STTran [7], TRACE [45], STTran-TPI [49], APT [26], and ISGG [21]. Additionally, we compared our method with ReLDN [57], which is a static method. Performance comparisons in terms of mR@K and R@K for  $K = [10, 20, 50]$  are reported in Tables 1, 2, and 3. These tables contain comparisons with two experimental setups: a) **With Constraint** and b) **No Constraints**. For the tasks, i.e., *PREDCLS + SGCLS*, we presented results for these experimental setups, specifically reporting mR@K and R@K in Tables 2 and 3, respectively. Table 1 compares results for the task *SGDET*. From the tables, it has been observed that our method consistently outperforms other methods across all tasks and for both experimental setups. Specifically, in comparison to the best baselines, we observe improvements of 4.1% on *SGDET*-mR@10, 3.4% on *SGCLS*-mR@10, and 1.9% on *PREDCLS*-mR@10 under the "With Constraint" setup. Under the "No Constraints" setup, we observe more significant improvements, with 3.9% on *SGDET*-mR@10, 1.4% on *SGCLS*-mR@10, and 1.7% on *PREDCLS*-mR@10. This demonstrates the capability of FLOCODE in generating more unbiased SGG for videos incorporating dynamic fluctuations and long-tailed relations. We further verified this in Figure 3 for **With Constraint** and **No Constraints**. In these figures, we compared our method on HEAD, BODY, and TAIL classes with mR@10 values. We split the classes into HEAD, BODY, and TAIL with the same definition as mentioned in [31]. Clearly, FLOCODE improved the performance across all the classes, but the improvement for TAIL classes is more confirming the unbiased predictions. Per-class performance is shown in Fig. 4, comparing with other methods STTran and TRACE, showing improvement at the class level. Additionally, our method outperforms in terms of R@K values, as shown in Table 3, demonstrating improvements overall compared to existing methods. This shows that our method has better generalization since it performs better on both mR@K (long-tail) and R@K (overall). Qualitative visualizations are illustrated in Fig. 2.

#### 4.2. Ablation Studies

We have conducted extensive ablation studies on SGCLS and SGDET tasks. Specifically, we studied the impact of *MCL* (uncertainty-aware contrastive learning), *Debiasing* (correlation-aware debiasing), *TFoD* (flow-aware temporal consistency), *Regularizer* (aleatoric and epistemic regularizer), and *EMA Teacher*. When all these components have been removed, FLOCODE boils down to the baseline STTran [7], where the object proposals and predicate embeddings are fed to the FFN layers before finally predicting the predicate class using a classification layer. The results for these ablation studies are presented in Table 4. **Uncertainty Attenuation and Debiasing:** We first discuss the impact of uncertainty-aware contrastive learning

Method	With Constraint						No Constraint					
	mR@10	mR@20	mR@50	R@10	R@20	R@50	mR@10	mR@20	mR@50	R@10	R@20	R@50
ReIDN	3.3	3.3	3.3	9.1	9.1	9.1	7.5	18.8	33.7	13.6	23.0	36.6
HCRD supervised	-	8.3	9.1	-	27.9	30.4	-	-	-	-	-	-
TRACE	8.2	8.2	8.2	13.9	14.5	14.5	22.8	31.3	41.8	26.5	35.6	45.3
ISGG	-	19.7	22.9	-	29.2	35.3	-	-	-	-	-	-
STTran	16.6	20.8	22.2	25.2	34.1	37.0	20.9	29.7	39.2	24.6	36.2	48.8
STTran-TPI	15.6	20.2	21.8	26.2	34.6	37.4	-	-	-	-	-	-
APT	-	-	-	26.3	36.1	38.3	-	-	-	25.7	37.9	50.1
TEMPURA	18.5	22.6	23.7	28.1	33.4	34.9	24.7	33.9	43.7	29.8	38.1	46.4
FloCoDe	<b>22.6</b>	<b>24.2</b>	<b>27.9</b>	<b>31.5</b>	<b>38.4</b>	<b>42.4</b>	<b>28.6</b>	<b>35.4</b>	<b>47.2</b>	<b>32.6</b>	<b>43.9</b>	<b>51.6</b>

Table 1. Comparative results for SGDET task, on AG[17] in terms of mean-Recall@K and Recall@K, best results are in bold.

Method	With Constraint						No Constraints					
	PredCLS			SGCLS			PredCLS			SGCLS		
	mR@10	mR@20	mR@50	mR@10	mR@20	mR@50	mR@10	mR@20	mR@50	mR@10	mR@20	mR@50
ReIDN	6.2	6.2	6.2	3.4	3.4	3.4	31.2	63.1	75.5	18.6	36.9	42.6
TRACE	15.2	15.2	15.2	8.9	8.9	8.9	50.9	73.6	82.7	31.9	42.7	46.3
STTran	37.8	40.1	40.2	27.2	28.0	28.0	51.4	67.7	82.7	40.7	50.1	58.8
STTran-TPI	37.3	40.6	40.6	28.3	29.3	29.3	-	-	-	-	-	-
TEMPURA	42.9	46.3	46.3	34.0	35.2	35.2	61.5	85.1	98.0	48.3	61.1	66.4
FloCoDe	<b>44.8</b>	<b>49.2</b>	<b>49.3</b>	<b>37.4</b>	<b>39.2</b>	<b>39.4</b>	<b>63.2</b>	<b>86.9</b>	<b>98.6</b>	<b>49.7</b>	<b>63.8</b>	<b>69.2</b>

Table 2. Comparative results for PREDCLS and SGCLS task, on AG[17] in terms of mean-Recall@K, best results are in bold.

Method	With Constraint						No Constraints					
	PredCLS			SGCLS			PredCLS			SGCLS		
	R@10	R@20	R@50	R@10	R@20	R@50	R@10	R@20	R@50	R@10	R@20	R@50
ReIDN	20.3	20.3	20.3	11.0	11.0	11.0	44.2	75.4	89.2	25.0	41.9	47.9
TRACE	27.5	27.5	27.5	14.8	14.8	14.8	72.6	91.6	96.4	37.1	46.7	50.5
STTran	68.6	71.8	71.8	46.4	47.5	47.5	77.9	94.2	99.1	54.0	63.7	66.4
STTran-TPI	69.7	72.6	72.6	47.2	48.3	48.3	-	-	-	-	-	-
APT	69.4	73.8	73.8	47.2	48.9	48.9	78.5	95.1	99.2	55.1	65.1	68.7
TEMPURA	68.8	71.5	71.5	47.2	48.3	48.3	80.4	94.2	99.4	56.3	64.7	67.9
FloCoDe	<b>70.1</b>	<b>74.2</b>	<b>74.2</b>	<b>48.4</b>	<b>51.2</b>	<b>51.2</b>	<b>82.8</b>	<b>97.2</b>	<b>99.9</b>	<b>57.4</b>	<b>66.2</b>	<b>68.8</b>

Table 3. Comparative results for PREDCLS and SGCLS task, on AG[17] in terms of Recall@K, best results are in bold.

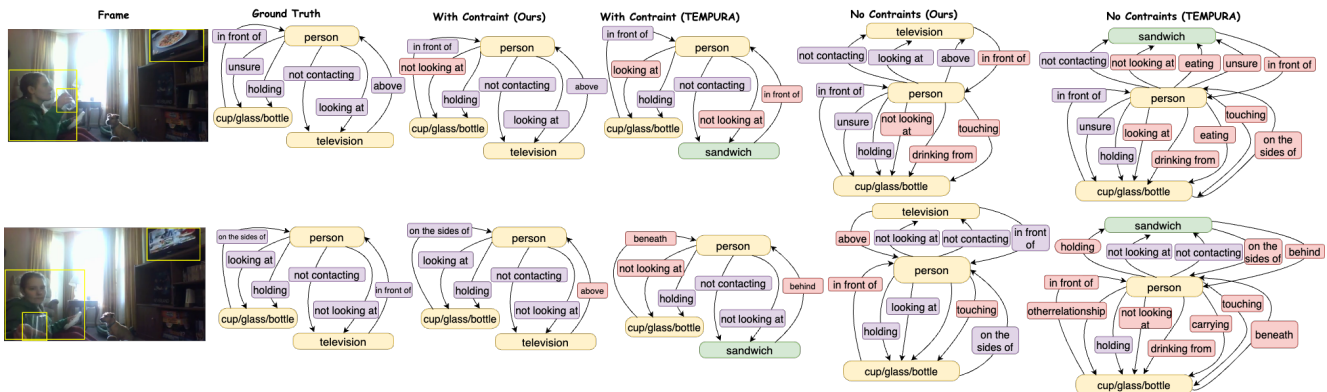


Figure 2. **Qualitative Comparison** with TEMPURA[31] for both **With Constraint** and **No Constraints** setup. From left to right: input video frames, ground truth graphs, graphs generated by FLOCoDe, graphs generated by TEMPURA[31]. Incorrect object and predicate predictions are shown in green and red, respectively.

Uncertainty-aware Contrastive Learning	Correlation-aware Debiasing	Flow-aware Temporal Consistency	Regularizer	EMA Teacher	With Constraint				No Constraints			
					SGCLS		SGDET		SGCLS		SGDET	
					mR@10	mR@20	mR@10	mR@20	mR@10	mR@20	mR@10	mR@20
-	-	-	-	-	27.2	28.0	16.5	20.8	40.7	50.1	20.9	29.7
✓	-	✓	✓	✓	34.1	33.8	19.6	22.1	46.9	61.1	26.6	32.7
-	✓	✓	✓	✓	33.6	34.3	19.4	21.8	46.2	60.6	25.8	32.5
✓	✓	-	✓	✓	32.2	33.4	18.1	19.8	45.9	59.1	21.8	31.6
✓	✓	✓	-	✓	35.8	36.6	21.2	22.7	48.3	61.4	27.5	34.4
✓	✓	✓	✓	-	36.7	38.8	22.1	23.8	49.2	62.9	28.3	35.2
✓	✓	✓	✓	✓	<b>37.4</b>	<b>39.2</b>	<b>22.6</b>	<b>24.2</b>	<b>49.7</b>	<b>63.8</b>	<b>28.6</b>	<b>35.4</b>

Table 4. **Ablation Studies:** Importance of *MCL*, *Debiasing*, *TfOD*, *Regularizer* & *EMA Teacher* for SGCLS and SGDET.




Figure 3. Comparison of mR@10 for the HEAD, BODY and TAIL classes for "with constraint"(top) and "no constraints"(bottom)




Figure 4. Comparative per class performance for PREDCLS task in R@10 for "with constraint" setup

Task	$\mathcal{K}$				
	1	2	4	6	8
PREDCLS	39.8	41.2	43.4	<b>44.8</b>	44.2
SGCLS	30.2	35.5	<b>37.4</b>	36.2	35.8
SGDET	16.1	18.1	21.9	<b>22.6</b>	22.1

Table 5. Results (in mR@10) with varying number of mixtures  $\mathcal{K}$  for **With Constraint** setup

and correlation-aware debiasing. First, we remove the loss  $\mathcal{L}_{MCL}$  to study the improvement on top of  $\mathcal{L}_{MAL}$ . In the second case, we remove the correlation-aware debiasing during training. The results for these are shown in Table 4, rows 1 and 2. Comparing the resulting models with FLOCODE shows a significant drop in mR@10 values, indicating the value addition from each of these in generating unbiased SGG. This also shows that both can address the bias, specifically, contrastive learning deals with label correlation especially for long tailed classes and hence focus to generate unbiased predicate embeddings. **Temporally Consistent Object Detection:** Table 4 (row 3) clearly illustrates the impact of flow-aware detection on ensuring temporal consistency in object identification. The absence of *TfOD* leads to a noticeable performance decline when contrasted with FLOCODE, underscoring the pivotal role of accurate object detection as a key bottleneck in SGG methods. For the *PREDCLS* task, leveraging only ground-truth boxes and labels results in significantly higher mR@k and R@K values compared to other tasks. Regarding the **Uncertainty Regularizer and EMA Teacher** components (rows 4 and 5), their ablation underscores their crucial role in reducing noise, especially for TAIL classes. Aleatoric and epistemic regularization helps better prediction of the noise distribution i.e. mixture variances. EMA teachers contribute to balanced predicate embeddings, enhancing performance specific to their respective classes. **Number of Mixtures  $\mathcal{K}$ :** The performance of FLOCODE with varying numbers of mixtures in MLN is shown in Table 5. The number of mixtures between 4 to 6 is optimal.

## 5. Conclusion

In summary, this paper proposes FLOCODE, a method for generating unbiased dynamic scene graphs from videos. By addressing issues such as biased scene graph generation and the long-tailed distribution of visual relationships, FLOCODE achieves significant performance gains. Through features like flow-aware temporal consistency, correlation debiasing, label correlation and uncertainty attenuation it offers a robust solution for capturing accurate scene representations in dynamic environments.



## References

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015. **1**
- [2] Nikita Araslanov and Stefan Roth. Self-supervised augmentation consistency for adapting semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15384–15394, 2021. **5**
- [3] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6836–6846, 2021. **1, 2**
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. **1**
- [5] Tianshui Chen, Weihao Yu, Riquan Chen, and Liang Lin. Knowledge-embedded routing network for scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6163–6171, 2019. **1**
- [6] Marc Combalia, Ferran Hueto, Susana Puig, Josep Malvehy, and Veronica Vilaplana. Uncertainty estimation in deep neural networks for dermoscopic image classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 744–745, 2020. **5**
- [7] Yuren Cong, Wentong Liao, Hanno Ackermann, Bodo Rosenhahn, and Michael Ying Yang. Spatial-temporal transformer for dynamic scene graph generation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16372–16382, 2021. **1, 2, 3, 5, 6**
- [8] Alakh Desai, Tz-Ying Wu, Subarna Tripathi, and Nuno Vasconcelos. Learning of visual relations: The devil is in the tails. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15404–15413, 2021. **1**
- [9] Jianfeng Dong, Xirong Li, Chaoxi Xu, Xun Yang, Gang Yang, Xun Wang, and Meng Wang. Dual encoding for video retrieval by text. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4065–4080, 2021. **1**
- [10] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2758–2766, 2015. **3**
- [11] Geoff French, Samuli Laine, Timo Aila, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, varied perturbations. *arXiv preprint arXiv:1906.01916*, 2019. **5**
- [12] Raghav Goyal<sup>12</sup> and Leonid Sigal<sup>123</sup>. A simple baseline for weakly-supervised human-centric relation detection. 2021. **1**
- [13] Kai Han, Yunhe Wang, Hanting Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chun-jing Xu, Yixing Xu, et al. A survey on vision transformer. *IEEE transactions on pattern analysis and machine intelligence*, 45(1):87–110, 2022. **1**
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **2, 6**
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. **2, 3**
- [16] Lukas Hoyer, Dengxin Dai, Yuhua Chen, Adrian Koring, Suman Saha, and Luc Van Gool. Three ways to improve semantic segmentation with self-supervised depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11130–11140, 2021. **5**
- [17] Jingwei Ji, Ranjay Krishna, Li Fei-Fei, and Juan Carlos Niebles. Action genome: Actions as compositions of spatio-temporal scene graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10236–10247, 2020. **2, 5, 6, 7**
- [18] Jingwei Ji, Rishi Desai, and Juan Carlos Niebles. Detecting human-object relationships in videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8106–8116, 2021. **1, 3**
- [19] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *European Conference on Computer Vision*, pages 105–124. Springer, 2022. **1**
- [20] Anant Khandelwal. Segda: Maximum separable segment mask with pseudo labels for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2158–2168, 2023. **5**
- [21] Siddhesh Khandelwal and Leonid Sigal. Iterative scene graph generation. *Advances in Neural Information Processing Systems*, 35:24295–24308, 2022. **6**
- [22] 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. *International journal of computer vision*, 123:32–73, 2017. **1, 2, 6**
- [23] Rongjie Li, Songyang Zhang, and Xuming He. Sgtr: End-to-end scene graph generation with transformer. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19486–19496, 2022.
- [24] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *Proceedings of the IEEE international conference on computer vision*, pages 1261–1270, 2017. **1, 2**
- [25] Yicong Li, Xun Yang, Xindi Shang, and Tat-Seng Chua. Interventional video relation detection. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4091–4099, 2021. **1**
- [26] Yiming Li, Xiaoshan Yang, and Changsheng Xu. Dynamic scene graph generation via anticipatory pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision*

- and *Pattern Recognition*, pages 13874–13883, 2022. 1, 2, 3, 5, 6
- [27] Xin Lin, Changxing Ding, Jinqun Zeng, and Dacheng Tao. Gps-net: Graph property sensing network for scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3746–3753, 2020. 1, 2
- [28] Chenchen Liu, Yang Jin, Kehan Xu, Guoqiang Gong, and Yadong Mu. Beyond short-term snippet: Video relation detection with spatio-temporal global context. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10840–10849, 2020. 1, 2
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [30] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 852–869. Springer, 2016. 1, 2
- [31] Sayak Nag, Kyle Min, Subarna Tripathi, and Amit K Roy-Chowdhury. Unbiased scene graph generation in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22803–22813, 2023. 1, 2, 4, 5, 6, 7
- [32] Megha Nawhal and Greg Mori. Activity graph transformer for temporal action localization. *arXiv preprint arXiv:2101.08540*, 2021. 1, 2
- [33] Vardan Papyan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020. 3
- [34] Tao Pu, Tianshui Chen, Hefeng Wu, Yongyi Lu, and Liang Lin. Spatial-temporal knowledge-embedded transformer for video scene graph generation. *arXiv preprint arXiv:2309.13237*, 2023. 1
- [35] Xufeng Qian, Yueting Zhuang, Yimeng Li, Shaoning Xiao, Shiliang Pu, and Jun Xiao. Video relation detection with spatio-temporal graph. In *Proceedings of the 27th ACM international conference on multimedia*, pages 84–93, 2019. 1, 2
- [36] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 2, 6
- [37] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 510–526. Springer, 2016. 5
- [38] Aman Sinha, Matthew O’Kelly, Russ Tedrake, and John C Duchi. Neural bridge sampling for evaluating safety-critical autonomous systems. *Advances in Neural Information Processing Systems*, 33:6402–6416, 2020. 5
- [39] Cees GM Snoek, Marcel Worring, et al. Concept-based video retrieval. *Foundations and Trends® in Information Retrieval*, 2(4):215–322, 2009. 1
- [40] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7464–7473, 2019. 1, 2
- [41] Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. Learning to compose dynamic tree structures for visual contexts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6619–6628, 2019. 1, 2, 6
- [42] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiabin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3716–3725, 2020. 2
- [43] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhofen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4631–4640, 2016. 1
- [44] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. 5
- [45] Yao Teng, Limin Wang, Zhifeng Li, and Gangshan Wu. Target adaptive context aggregation for video scene graph generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13688–13697, 2021. 1, 2, 6
- [46] Wilhelm Truhedden, Viktor Olsson, Juliano Pinto, and Lennart Svensson. Dacs: Domain adaptation via cross-domain mixed sampling. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1379–1389, 2021. 5
- [47] Yao-Hung Hubert Tsai, Santosh Divvala, Louis-Philippe Morency, Ruslan Salakhutdinov, and Ali Farhadi. Video relationship reasoning using gated spatio-temporal energy graph. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10424–10433, 2019. 2
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 1, 2, 3
- [49] Shuang Wang, Lianli Gao, Xinyu Lyu, Yuyu Guo, Pengpeng Zeng, and Jingkuan Song. Dynamic scene graph generation via temporal prior inference. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5793–5801, 2022. 2, 3, 6
- [50] Yinwei Wei, Xiang Wang, Weili Guan, Liqiang Nie, Zhouchen Lin, and Baoquan Chen. Neural multimodal cooperative learning toward micro-video understanding. *IEEE Transactions on Image Processing*, 29:1–14, 2019. 1
- [51] Bin Wen, Jie Luo, Xianglong Liu, and Lei Huang. Unbiased scene graph generation via rich and fair semantic extraction. *arXiv preprint arXiv:2002.00176*, 2020. 2

- [52] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9777–9786, 2021. [1](#)
- [53] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015. [1](#)
- [54] Li Xu, Haoxuan Qu, Jason Kuen, Jiuxiang Gu, and Jun Liu. Meta spatio-temporal debiasing for video scene graph generation. In *European Conference on Computer Vision*, pages 374–390. Springer, 2022. [1](#)
- [55] Yibo Yang, Haobo Yuan, Xiangtai Li, Zhouchen Lin, Philip Torr, and Dacheng Tao. Neural collapse inspired feature-classifier alignment for few-shot class incremental learning. *arXiv preprint arXiv:2302.03004*, 2023. [3](#)
- [56] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5831–5840, 2018. [1](#), [2](#), [3](#)
- [57] Ji Zhang, Kevin J Shih, Ahmed Elgammal, Andrew Tao, and Bryan Catanzaro. Graphical contrastive losses for scene graph parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11535–11543, 2019. [1](#), [2](#), [6](#)
- [58] Sipeng Zheng, Shizhe Chen, and Qin Jin. Vrdformer: End-to-end video visual relation detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18836–18846, 2022. [2](#)
- [59] Xizhou Zhu, Yujie Wang, Jifeng Dai, Lu Yuan, and Yichen Wei. Flow-guided feature aggregation for video object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 408–417, 2017. [3](#)
- [60] Xizhou Zhu, Yuwen Xiong, Jifeng Dai, Lu Yuan, and Yichen Wei. Deep feature flow for video recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2349–2358, 2017. [3](#)