



UvA-DARE (Digital Academic Repository)

Multi-Label Meta Weighting for Long-Tailed Dynamic Scene Graph Generation

Chen, S.; Du, Y.; Mettes, P.; Snoek, C.G.M.

DOI

[10.48550/arXiv.2306.10122](https://doi.org/10.48550/arXiv.2306.10122)

[10.1145/3591106.3592267](https://doi.org/10.1145/3591106.3592267)

Publication date

2023

Document Version

Final published version

Published in

ICMR'23

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment>)

[Link to publication](#)

Citation for published version (APA):

Chen, S., Du, Y., Mettes, P., & Snoek, C. G. M. (2023). Multi-Label Meta Weighting for Long-Tailed Dynamic Scene Graph Generation. In *ICMR'23: proceedings of the 2023 ACM International Conference on Multimedia Retrieval : Thessaloniki, Greece, June 12-15, 2023* (pp. 39-47). Association for Computing Machinery. <https://doi.org/10.48550/arXiv.2306.10122>, <https://doi.org/10.1145/3591106.3592267>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)



Multi-Label Meta Weighting for Long-Tailed Dynamic Scene Graph Generation

Shuo Chen, Yingjun Du, Pascal Mettes, and Cees G. M. Snoek
University of Amsterdam

ABSTRACT

This paper investigates the problem of scene graph generation in videos with the aim of capturing semantic relations between subjects and objects in the form of (subject, predicate, object) triplets. Recognizing the predicate between subject and object pairs is imbalanced and multi-label in nature, ranging from ubiquitous interactions such as spatial relationships (*e.g. in front of*) to rare interactions such as *twisting*. In widely-used benchmarks such as Action Genome and VidOR, the imbalance ratio between the most and least frequent predicates reaches 3,218 and 3,408, respectively, surpassing even benchmarks specifically designed for long-tailed recognition. Due to the long-tailed distributions and label co-occurrences, recent state-of-the-art methods predominantly focus on the most frequently occurring predicate classes, ignoring those in the long tail. In this paper, we analyze the limitations of current approaches for scene graph generation in videos and identify a one-to-one correspondence between predicate frequency and recall performance. To make the step towards unbiased scene graph generation in videos, we introduce a multi-label meta-learning framework to deal with the biased predicate distribution. Our meta-learning framework learns a meta-weight network for each training sample over all possible label losses. We evaluate our approach on the Action Genome and VidOR benchmarks by building upon two current state-of-the-art methods for each benchmark. The experiments demonstrate that the multi-label meta-weight network improves the performance for predicates in the long tail without compromising performance for head classes, resulting in better overall performance and favorable generalizability. Code: <https://github.com/shanshuo/ML-MWN>.

CCS CONCEPTS

• **Computing methodologies** → **Scene understanding; Activity recognition and understanding.**

KEYWORDS

Scene Graph Generation, Long-Tailed Distribution, Multi-Label Meta-Learning, Imbalanced Data, Video Understanding, Semantic Relations, Action Genome, VidOR

ACM Reference Format:

Shuo Chen, Yingjun Du, Pascal Mettes, and Cees G. M. Snoek. 2023. Multi-Label Meta Weighting for Long-Tailed Dynamic Scene Graph Generation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICMR '23, June 12–15, 2023, Thessaloniki, Greece

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0178-8/23/06...\$15.00
<https://doi.org/10.1145/3591106.3592267>

In *International Conference on Multimedia Retrieval (ICMR '23)*, June 12–15, 2023, Thessaloniki, Greece. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3591106.3592267>

1 INTRODUCTION

Scene graph generation in videos focuses on detecting and recognizing relationships between pairs of subjects and objects. The resulting dynamic scene graph is a directed graph whose nodes are objects with their relationships as edges in a video. Extracting such graphs from videos constitutes a highly challenging research problem [13], with broad applicability in multimedia and computer vision. Effectively capturing such structural-semantic information boosts downstream tasks such as captioning [40], video retrieval [29], visual question answering [1], and numerous other visual-language tasks.

Current methods place a heavy emphasis on recognizing subject-to-object relationship categories. A leading approach to date involves extracting multi-modal features for relation instances, followed by either pooling the multi-modal features [23, 30, 38] or learning a feature representation [6] to feed into the predicate classifier network. Despite the strong focus on relation recognition, existing methods overlook the extremely long-tailed distribution of predicate classes. Figure 1 displays the recall per predicate class from STTran [7] and its corresponding occurrences on the Action Genome dataset. This trend is even more pronounced on the VidOR dataset. Figure 2 illustrates the occurrence distribution vs. Recall@50 from Social Fabric [6] for the video relation detection task on the VidOR dataset, where a few head predicates dominate all other classes. This phenomenon has not been actively investigated, as the evaluation metrics do not penalize lower scores for predicates in the long tail. In light of these observations, this paper advocates for the development for scene graph generation methods in videos that effectively handle both common and rare predicates.

We introduce a meta-learning framework to address the long-tailed dynamic scene graph generation problem. Drawing inspiration from the concept of meta weighting [28], we propose a Multi-Label Meta Weight Network (ML-MWN) to learn meta weights across both examples and classes explicitly. These meta weights are, in turn, used to steer the downstream loss to optimize the parameters of the predicate classifier. We adopt a meta-learning framework to optimize the ML-MWN parameters, where we compute each instance's per-class loss in a training batch and obtain a loss matrix. The loss matrix is fed into our ML-MWN, which outputs a weight matrix, with each row representing the weight vector for an instance's loss vector. We sample a meta-validation batch and use an unbiased meta-loss to guide the training of ML-MWN. We adopt the inverse frequency binary cross-entropy loss as the meta-loss. Finally, we integrate our framework with existing methods to guide the predicate classification.

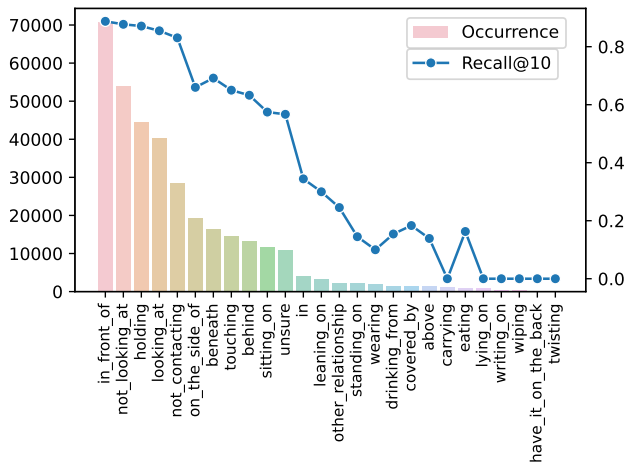


Figure 1: Long-tailed predicate occurrences vs. class-wise recall from STTran [7] on the Action Genome dataset [12]. The class-wise performance drops drastically, highlighting the importance of long-tailed dynamic scene graph generation.

To evaluate our meta-learning framework, we employ two recent state-of-the-art methods [6, 7], one for the scene graph generation task on the Action Genome dataset and one for video relation detection on the VidOR dataset. We empirically demonstrate that our approach enhances predicate predictions for these recent methods across various evaluation metrics. Furthermore, we show that our framework improves the performance of long-tailed predicates without hampering the performance of more common classes. Our approach is generic and works on top of any scene graph generation method, ensuring broad applicability. We make the code available on <https://github.com/shanshuo/ML-MWN>.

In summary, our contributions are three-fold:

1. We investigate the long-tail issue in dynamic scene graph generation and analyze the limitations of existing methods.
2. We introduce a multi-label meta-learning framework to address the biased predicate class distribution.
3. We propose a Multi-Label Meta Weight Network (ML-MWN) to explicitly learn a weighting function, which demonstrates generalization ability performance on two benchmarks when plugged into two existing approaches,

2 RELATED WORKS

Dynamic scene graph generation. Scene graph generation was first pioneered in [13] for image retrieval, and the task quickly gained further traction, as seen in *e.g.* [21, 33, 39, 42, 43]. Recently, a number of papers have identified the long-tailed distribution in image scene graphs and focused on generating unbiased scene graphs [8, 9, 15–17, 41]. We seek to bring the same problem to light in the video domain. Ji *et al.* [12] firstly extended scene graph generation to videos and introduced the Action Genome dataset. A wide range of works have since proposed solutions to the problem [3, 6, 10, 14, 20, 30–32, 38, 45]. Recently, Li *et al.* [17] proposed an anticipatory pre-training paradigm based on Transformer to

model the temporal correlation of visual relationships. Similarly, the VidOR dataset collected by Shang *et al.* [26] is another popular benchmark. Leading approaches generate proposals [4] for individual objects on short video snippets, encode the proposals, predict a relation, and associate the relations over the entire video, *e.g.* [23, 30, 38]. Liu *et al.* [20] generate the proposals using the sliding window way. More recently, Gao *et al.* [10] proposed a classification-then-grounding framework, which can avoid the high influence of proposal quality on performance. Chen *et al.* [5] performed a series of analyses on video relation detection. In this paper, we use STTran [7] and Social Fabric [6] to capture the relation feature and insert our multi-label meta-weight network on top. Cong *et al.* [7] proposed a spatial-temporal Transformer to capture the spatial context and temporal dependencies for a dynamic scene graph. Moreover, Chen *et al.* [6] proposed an encoding that represents a pair of object tubelets as a composition of interaction primitives. Both approaches provide competitive results and form a fruitful testbed for our meta-learning framework.

Multi-label long-tailed classification. Multi-label long-tailed recognition is a challenging problem that deals with sampling differences and biased label co-occurrences [44]. A few works have studied this topic, with most solutions based on new loss formulations. Specifically, Wu *et al.* [37] proposed a distribution-based loss for multi-label long-tailed image recognition. More recently, Tian *et al.* [35] proposed a hard-class mining loss for the semantic segmentation task by dynamically weighting the loss for each class based on instantaneous recall performance. Inspired by these loss-based works, we utilize inverse frequency cross-entropy loss during our meta-learning process.

Meta learning for sample weighting. Ren *et al.* [24] pioneered the adoption of a meta learning framework to re-weight samples for imbalanced datasets. Based on [24], Shu *et al.* [28] utilize an MLP to explicitly learn the weighting function. Recently, Bohdal *et al.* [2] presented EvoGrad to compute gradients more efficiently by preventing the computation of second-order derivatives in [28]. However, these methods are targeted for multi-class single-label classification. Therefore, we present the multi-label meta weight net for predicate classification, with an MLP that output a weight for each class loss.

3 MULTI-LABEL META WEIGHT NETWORK

Dynamic scene graph generation [7] takes a video as the input and generates directed graphs whose objects of interest are represented as nodes, and their relationships are represented as edges. Each relationship edge, along with its connected two object nodes, form a ⟨subject, predicate, object⟩ semantic triplet. These directed graphs are structural representations of the video’s semantic information. Highly related to dynamic scene graph generation, video relation detection [27] also outputs ⟨subject, predicate⟩ object triplets, aiming to classify and detect the relationship between object tubelets occurring within a video. Due to the high similarity between the two tasks, we consider them both in the experiments. For brevity, in this paper, we use the term dynamic scene graph generation to denote both tasks throughout this paper.




Figure 2: On the VidOR dataset [26], the long-tailed distribution is even worse. We can observe that Social Fabric [6] ignores most predicates with limited samples.

Action Genome [12] and VidOR [26] are two popular benchmark datasets for dynamic scene graph generation. However, both datasets suffer from a long-tailed distribution in predicate occurrences, as shown in Figure 1. The evaluation metrics forgo the class-wise differences and count all classes during inference, resulting in a trained predicate classifier with a strong bias toward head classes such as `in_front_of` and `next_to`. Although these predicate classes are often spatial-oriented and object-agnostic, tail classes like carrying, twisting, and driving are of more interest to us. In addition to the long-tailed distribution, predicate classification faces another challenge. Since multiple relationships can occur between a subject-object pair simultaneously, predicate classification is a multi-label classification problem. The co-occurrence of labels leads to head-class predicate labels frequently appearing alongside tail-class predicate labels, further exacerbating the imbalance problem.

In this paper, we propose a meta-learning framework that addresses the long-tailed multi-label predicate classification task. We introduce a Multi-Label Meta Weight Net (ML-MWN) to learn a weight vector for each training instance’s multi-label loss. The gradient of the sum of weighted loss is then calculated to optimize the classifier network’s parameters during backward propagation. Our model-agnostic approach can be incorporated into existing dynamic scene graph generation methods. In particular, the framework includes two stages: (1) Relation feature extraction, where we use existing dynamic scene graph generation methods to obtain the feature representation of the relation instances, and (2) multi-label meta-weighting learning. We adopt a meta-learning framework to re-weight each instance’s multi-label loss and propose learning an explicit weighting function that maps from training loss to weight vector. We learn a weight vector for each training instance to re-weight its multi-label loss, *i.e.*, multi-label binary cross-entropy loss. We achieve this by using an MLP, which takes the multi-label training loss as input and outputs the weight vector. We sample a

meta-validation set to guide the training of MLP. Ideally, the meta-validation set should be clean and free from the long-tailed issue, as in [28]. However, we cannot sample such a clean meta-validation set due to the label-occurrence issue. To deal with the issue, we adopt the inverse frequency binary cross-entropy loss on meta-validation set. In the following sections, we describe the ML-MWN and the meta-learning framework in detail.

3.1 Learning weights for multi-label losses

Let x_i denote the feature representation of i -th relation instance from the training set \mathcal{D} and $y_i \in \mathbb{R}^C$ represent the corresponding multi-label one-hot vector, where $\mathcal{D} = \{x_i, y_i\}_{i=1}^N$. The multi-label predicate classifier network is represented by f_θ with θ as its parameters. To enhance the robustness of training in the presence of long-tailed multi-label training instances, we impose weights $w_{i,c}$ on the i -th instance’s c -th class loss $l_{i,c}$. Instead of pre-specifying the weights based on class size [18, 35], we learn an explicit weighting function directly from the data. Specifically, we propose the ML-MWN (Multi-Label Meta Weight Net) denoted by g_ϕ , with ϕ as its parameters, to obtain the weighting vector for each relation instance’s multi-label loss. We use the loss from f_θ as the input.

A small meta-validation set $\hat{\mathcal{D}} = \{x_j, y_j\}_{j=1}^M$, where M is the number of meta-validation instances and $M \ll N$, is sampled to guide the training of ML-MWN. The meta-validation set does not overlap with the training set. The weighted losses are then calculated to guarantee that the learned multi-label predicate classifier is unbiased toward dominant classes.

During training, the optimal classifier parameter θ^* can be extracted by minimizing the training loss:

$$L^{train}(\theta) = \frac{1}{n} \frac{1}{C} \sum_{i=1}^n \sum_{c=1}^C w_{i,c} \cdot l_{i,c}, \quad (1)$$

where n is the number of training instances in a batch, and C is the number of classes. During inference, we only use the optimal classifier network f_{θ^*} for evaluation.

3.2 The meta-learning process

We adopt a meta-learning framework to update the classifier and ML-MWN. The meta-validation set represents the unbiased relation instances following a balanced predicate class distribution. Due to the multi-label classification label-occurrence issue [44], we employ an inverse frequency BCE loss on the meta-validation set to simulate a balanced label distribution. As illustrated in Figure 3, the process comprises three main steps to optimize θ and ϕ within a batch.

Suppose we are at t -th iteration during training. First, for a batch of n training instances with corresponding feature representations and multi-labels $\{x_i, y_i\}, 1 \leq i \leq n$, we feed x_i into the classifier and obtain $\hat{y}_i = f_{\theta^t}(x_i) \in \mathbb{R}^C$. The unweighted BCE training loss is calculated as

$$l_{i,c}(\theta^t) = -y_{i,c} \cdot \log(\hat{y}_{i,c}(\theta^t)) + (1 - y_{i,c}) \cdot \log(1 - \hat{y}_{i,c}(\theta^t)). \quad (2)$$

Then $l_{i,c}$ is fed into the ML-MWN to obtain the weight $\hat{w}_{i,c} = g_{\phi^t}(l_{i,c}(\theta^t))$. After calculating the weighted loss as $\hat{w}_{i,c} \cdot l_{i,c}$, we update θ^t :

$$\hat{\theta}^t = \theta^t - \alpha \frac{1}{n} \frac{1}{C} \sum_{i=1}^n \sum_{c=1}^C g'_{\phi^t}(l_{i,c}(\theta^t)) \nabla_{\theta^t} l_{i,c}(\theta^t) \Bigg|_{\theta^t}, \quad (3)$$

where α is the step size. We call the updated $\hat{\theta}^t$ the pseudo classifier parameters since they are not used for the next batch.

In the second step, we update the ML-MWN parameters based on the meta-validation loss. We feed the meta-validation relation instance into the pseudo classifier and obtain $\hat{y}_j = f_{\hat{\theta}^t}(x_j) \in \mathbb{R}^C$. Let M_c denote the total number of relation instances belonging to predicate class $c \in \{1, \dots, C\}$. The frequency of a predicate class is calculated as $freq(c) = M_c/M$. By using inverse frequency weighting, the meta-validation loss is re-balanced to mimic a balanced predicate label distribution. We then update the ML-MWN parameters ϕ on the meta-validation data:

$$\begin{aligned} \phi^{t+1} &= \phi^t - \beta \frac{1}{M} \sum_{j=1}^M \sum_{c=1}^C \frac{1}{freq(c)} \nabla_{\phi^t} l_{j,c}(\hat{\theta}^t) \Bigg|_{\phi^t} \\ &= \phi^t - \beta \sum_{c=1}^C \frac{M}{M_c} \nabla_{\phi^t} l_{j,c}(\hat{\theta}^t) \Bigg|_{\phi^t}, \end{aligned} \quad (4)$$

where β is the step size.

Lastly, the updated ϕ^{t+1} is employed to output the new weights $w_{i,c}$. The new weighted losses are used to improve the parameters θ of the classifier network:

$$\theta^{t+1} = \theta^t - \alpha \frac{1}{n} \frac{1}{C} \sum_{i=1}^n \sum_{c=1}^C g'_{\phi^{t+1}}(l_{i,c}(\theta^t)) \nabla_{\theta^t} l_{i,c}(\theta^t) \Bigg|_{\theta^t}. \quad (5)$$

The ultimate goal is to guide the classifier network to achieve a balanced performance on the unbiased meta-validation set. The sequences of steps are shown in Algorithm 1. By alternating between standard and meta-learning, we can learn unbiased dynamic scene graphs by specifically increasing the focus on those examples and predicate classes that do not often occur in a dataset.

Algorithm 1 The ML-MWN learning algorithm

Require: Training data set \mathcal{D} , meta-validation set $\hat{\mathcal{D}}$, max epochs N_{Epoch}

Ensure: Predicate multi-label classifier network parameter θ^*

```

1: for  $t = 1$  to  $N_{Epoch}$  do
2:   for each mini batch  $\{x_i, y_i\}, 1 \leq i \leq n$  do
3:     Calculate the prediction  $\hat{y}_i$ .
4:     Calculate the unweighted loss using Eq. 2.
5:     Formulate the pseudo predicate classifier  $\hat{\theta}^t$  by Eq. 3.
6:     Get meta-validation instances  $\{x_j, y_j\} \in \hat{\mathcal{D}}$ .
7:     Update  $\phi^{t+1}$  by Eq. 4.
8:     Update  $\theta^{t+1}$  by Eq. 5.
9:   end for
10: end for

```

4 EXPERIMENTS

4.1 Datasets

4.1.1 Action Genome. [12] is a dataset which provides frame-level scene graph labels. It contains 234,253 annotated frames with 476,229 bounding boxes of 35 object classes (without person) and 1,715,568 instances of 25 relationship classes. For the 25 relationships, there are three different types: (1) attention relationships indicating if a person is looking at an object or not, (2) spatial relationships describing where objects are relative to one another, and (3) contact relationships denoting the different ways the person is contacting an object. In AG, there are 135,484 subject-object pairs. Each pair is labeled with multiple spatial relationships (e.g. ⟨phone-in front of-person⟩ and ⟨phone-on the side of-person⟩) or contact relationships (e.g. ⟨person-eating-food⟩ and ⟨person-holding-food⟩). There are three strategies to generate a scene graph with the inferred relation distribution [7]: (a) *with constraint* allows each subject-object pair to have one predicate at most. (b) *semi constraint* allows a subject-object pair has multiple predicates. The predicate is regarded as positive only if the corresponding confidence is higher than the threshold (0.9 in the experiments). (c) *no constraint* allows a subject-object pair to have multiple relationships guesses without constraint.

Evaluation metrics. We have three tasks for evaluation following [7]: (1) predicate classification (PREDCLS): with the subject and object's ground truth labels and bounding boxes, only predict predicate labels of the subject-object pair. (2) scene graph classification (SGCLS): with the subject and object's ground truth bounding boxes given, predict the subject, object's label and their corresponding predicate. (3) scene graph detection (SGDET): detect the subject and object's bounding boxes and predict the subject, object, and predicate's labels. The object detection is regarded as positive if the IoU between the predicted and ground-truth box is at least 0.5. Since traditional metrics Recall@K (R@K) are not able to reflect the impact of long-tailed data, we use the mean Recall@K (mR@K), which evaluates the R@K ($K = [10, 20, 50]$) of each relationship class and averages them. We follow the same selection of K as [7].

Implementation details. We randomly sample 10% samples from the training set as the meta-validation set. In line with [7], we adopt

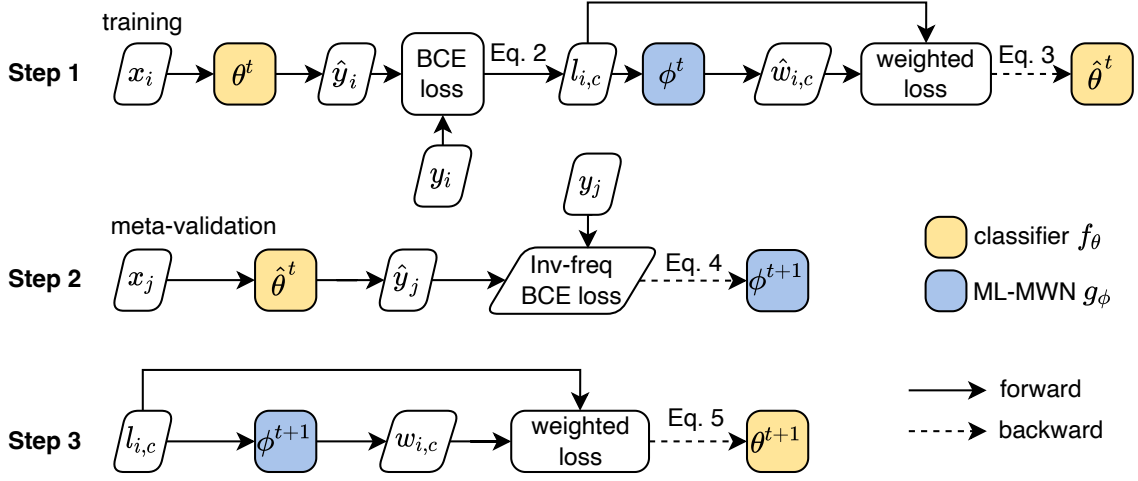


Figure 3: An overview of our proposed meta-learning process. We omit the relation feature extraction parts for simplification. x_i and x_j represent the features of the training instance and meta-validation instance, respectively. During a training batch, there are three steps: 1. Calculate the weighted loss and obtain a pseudo classifier; 2. Evaluate the pseudo classifier on the meta-validation set and update the ML-MWN; 3. Compute the new weighted loss with updated ML-MWN then update the classifier. The Binary Cross-Entropy (BCE) loss is adopted for multi-label classification. The inverse-frequency BCE loss is used to simulate an unbiased meta-validation set. During inference, only the predicate multi-label classifier network is utilized.

the Faster-RCNN [25] based on the ResNet101 [11] as the object detection backbone. The Faster-RCNN model is trained on AG and provided by Cong *et al.* [7]. We use an AdamW [22] optimizer with an initial learning rate $1e^{-4}$ and batch size 1 to train our relation feature model STTran part. We train ML-MWN using SGD with a momentum of 0.9, weight decay of 0.01, and an initial learning rate of 0.01. We train for 10 epochs. Other hyperparameter settings are identical to Cong *et al.* [7]. If not specified, the ML-MWN is an MLP of 1-100-1.

4.1.2 VidOR. [26] is a dataset that includes 10,000 user-generated videos selected from YFCC-100M [34], totaling approximately 84 hours of footage. It contains 80 object categories and 50 predicate categories. Besides providing annotated relation triplets, the dataset also provides bounding boxes of objects. VidOR is split into a training set with 7,000 videos, a validation set with 835 videos, and a testing set with 2,165 videos. Since the ground truth of the test set is unavailable, we follow [20, 23, 30, 38] and use the training set for training and the validation set for testing. We report the analysis of method performance on the VidOR validation set.

Evaluation metrics. We use the relation detection task for evaluation. The output requires a ⟨subject, predicate, object⟩ triplet prediction, along with the subject and object boxes. We adopt mR@K (K = [50, 100]) as the evaluation metric. We disregard the mAP used in Chen *et al.* [6] because we are more concerned with covering ground truth relationships belonging to tail classes during predictions. **Calculating mR@K.** For annotated video I_v , its G_v ground truth relationship triplets contain $G_{v,c}$ ground truth triplets with

relationship class c . With C relationship classes, the model successfully predicts $T_{v,c}^K$ triplets. In the V videos of validation/test dataset, for relationship c , there are V_c videos containing at least one ground truth triplet with this relationship. The R@K of relationship c can be calculated:

$$R@K_c = \frac{1}{V_c} \sum_{v=1, G_{v,c} \neq 0}^{V_c} \frac{T_{v,c}^K}{G_{v,c}} \quad (6)$$

Then we can calculate

$$mR@K = \frac{1}{C} \sum_{c=1}^C R@K_c. \quad (7)$$

Implementation details. We randomly sample 10% samples from the training set as the meta-validation set. Our experiments are conducted using 1 NVIDIA V100 GPU. We adopt the same training strategy of Chen *et al.* [6] for the relation feature extraction model. First, we detect all objects in each video frame using Faster R-CNN [25] with a ResNet-101 [11] backbone trained on MS-COCO [19]. The detected bounding boxes are linked with the Deep SORT tracker [36] to obtain individual object tubelets. Then, each tubelet is paired with any other tubelet to generate the tubelet pairs. We extract spatial location features [31], language features, I3D features, and location mask features for each pair. Then the multi-modal features are used as the representation of the relation instance. For the classifier and ML-MWN, we use an SGD optimizer with an initial learning rate of 0.01 and train 10 epochs.

	PredCLS			SGCLS			SGDET		
	mR@10	mR@20	mR@50	mR@10	mR@20	mR@50	mR@10	mR@20	mR@50
STTran [7]	37.96	39.65	39.66	27.61	28.14	28.14	17.89	21.76	22.89
STTran + MW-Net [28]	40.29	42.21	42.24	30.21	30.90	30.90	20.06	23.66	24.99
STTran + ML-MWN	43.23	44.43	44.64	32.13	32.70	32.72	23.46	27.13	28.52

Table 1: Evaluating the effect of meta learning on Action Genome in the *with constraint* setting. Enriching the recent STTran approach with meta learning improves recall across all metrics, with the best results achieved using the proposed multi-label meta weighting.

	PredCLS			SGCLS			SGDET		
	mR@10	mR@20	mR@50	mR@10	mR@20	mR@50	mR@10	mR@20	mR@50
STTran [7]	49.94	59.07	59.77	40.17	44.27	44.51	21.63	31.36	40.96
STTran + MW-Net [28]	52.61	62.32	63.1	43.12	47.11	47.77	24.19	34.83	43.85
STTran + ML-MWN	55.95	65.79	68.01	46.20	50.60	50.83	26.21	40.12	49.96

Table 2: Evaluating the effect of meta learning on Action Genome in the *semi constraint* setting. Similar to the *with constraint* setting, the proposed multi-label meta weighting obtains the best results across all metrics.

	PredCLS			SGCLS			SGDET		
	mR@10	mR@20	mR@50	mR@10	mR@20	mR@50	mR@10	mR@20	mR@50
STTran [7]	52.61	68.30	82.90	42.52	51.14	64.77	21.64	30.64	35.53
STTran + MW-Net [28]	55.12	70.42	85.45	45.55	54.46	67.24	24.24	34.25	37.98
STTran + ML-MWN	57.13	74.22	89.24	48.48	57.65	70.15	27.59	36.67	40.57

Table 3: Evaluating the effect of meta learning on Action Genome in the *no constraint* setting. Also in this challenging setting, our approach works best over all metrics.

4.2 Multi-label meta weighting on top of the state-of-the-art

Video scene graph generation. First, we investigate the effect of incorporating our meta-learning approach on top of existing state-of-the-art methods for scene graph generation in videos and video relation detection. We build upon the recent STTran approach of Cong *et al.* [7] for video scene graph generation. We compare STTran as is and as a baseline that uses conventional meta-learning without considering the multi-label nature of scene graphs, namely MW-Net [28]. Table 1 shows the results for the *with constraints* setting. Across the PredCLS, SGCLS, and SGDET tasks, incorporating our meta-learning approach improves the results. For PredCLS, our proposed STTran + ML-MWN enhances mR@10 by **5.27**, compared to the STTran baseline. On mean recall @ 50, we improve the scores by **4.98**, from 39.66 to 44.64. On SGDET, the mean recall @ 50 increases from 22.89 to 28.52. The MW-Net baseline already improves the STTran results, emphasizing the overall potential of meta-learning to address the long-tailed nature of scene graphs. However, our proposed multi-label meta-learning framework performs best across all tasks and recall thresholds. This improvement is a direct result of increasing the weight of classes in the long tail when optimizing the classifier network.

The results are consistent for the *semi constraint* and *no constraint* settings, as shown in Table 2 and Table 3. In Table 2, the mean recall is higher than in the *with constraint* setting since more predicted results are involved. For the SGCLS task, our framework achieves 50.60% on mR@20, which is 6.33% better than STTran and 3.49% better than STTran + MW-Net. Our framework outperforms all metrics in the *no constraint* setting. In particular, for SGDET, our method reaches 27.59% at mR@10, 5.95% better than STTran, and 3.35% higher than STTran + MV-Net. We conclude that our meta learning framework is effective for video scene graph generation and can be adopted by any existing work. In Table 3, the mean recall is the highest among the three settings. Unlimited predictions contribute to enhanced recall performance. Under this setting, STTran + ML-MWN still achieves the best on all metrics across all tasks. The results prove our method’s generality on various tasks with different settings.

Video relation detection. For video relation detection, we begin with the recent Social Fabric approach by Chen *et al.* [6]. Table 4 demonstrates the effect of incorporating our proposed meta learning framework for relation detection. The Social Fabric baseline, which is the current state-of-the-art in this setting, struggles to achieve good results for relation detection using mean recall as metrics. This underlines the problem’s difficulty. This holds similarly




Figure 4: Class-wise R@10 comparison of PredCLS on AG. Our framework outperforms STTran on all predicate classes.

Method	Relation detection	
	mR@50	mR@100
Sun <i>et al.</i> [31]	1.48	2.78
Social Fabric (SF) [6]	2.37	3.79
SF + MW-Net [28]	4.45	5.35
SF + ML-MWN	6.35	7.54

Table 4: Comparison on the VidOR dataset. Our meta learning framework provides clear improvements for relation detection on top of Social Fabric [6].

Architecture	PredCLS		
	mR@10	mR@50	mR@100
C-50-C	41.34	42.24	42.56
C-100-C	43.23	44.43	44.64
C-200-C	42.16	42.85	42.97
C-100-100-C	43.01	43.96	44.03
C-10-10-C	42.18	42.74	42.96
C-10-10-10-C	42.85	44.01	44.28

Table 5: Performance on AG with constraint for different MLP architecture. The 1-100-1 architecture is the best.

for the baseline by Sun et al [31]. When incorporating MW-Net [28], the results noticeably improve and further enhance with multi-label meta weighting. For mR@50, adding our meta-learning on top of Social Fabric boosts the results from 2.37 to 6.35. We conclude that multi-label meta-learning is crucial in video relation detection to achieve meaningful relation detection recalls across all classes.

4.3 Analyses, ablations, and qualitative examples

Predicate-level analysis. We present the class-wise R@10 of the predicate classification task on Action Genome in Figure 4. Observing Figure 4, we see that our method surpasses STTran [7] in all predicate categories. The improvement is much more significant for tail classes with limited training samples compared to head classes.

The superior performance demonstrates that the meta-validation set effectively guides the classifier to balance the tail classes without compromising the performance of head predicate classes.

Ablating the MLP architecture. We conduct an ablation study on the MLP architecture for the PredCLS task on Action Genome. Table 5 shows the results for six structures with varying depths and widths. We find that maximum width and depth are not necessary, with the best results achieved by the 1-100-1 variant, which we use as default in all experiments.

Qualitative examples. We provide the qualitative results in Figure 5 and Figure 6. In Figure 5, we compare our method with STTran [7] on the Action Genome dataset. Our method demonstrates better recognition of tail predicates in Action Genome. In the top row, STTran incorrectly classifies the tail class *beneath* as the head class *in front of*, and *sit on* as *touch*. In the bottom row, STTran misses *drink from* amongst others, while our method classifies them all correctly. In Figure 6, we compare our method with Social Fabric [6] on the VidOR dataset. Social Fabric fails to detect the tail class *lean on* in all frames, while our method successfully predicts it.

5 CONCLUSION

Predicate recognition plays a crucial role in contemporary dynamic scene graph generation methods, but the long-tailed and multi-label nature of the predicate distribution is commonly ignored. We observe that rare predicates on popular benchmarks are inadequately recovered or even disregarded by recent methods. To move toward unbiased scene graph generation in videos, we propose a multi-label meta-learning framework that learns to weight samples and classes to optimize any predicate classifier effectively. Our approach is versatile and can be incorporated into any existing methods. Experiments demonstrate the potential of our multi-label meta-learning framework, with superior overall performance and an improved focus on rare predicates. We believe our method could be extended to other multi-label long-tailed recognition tasks and may offer inspiration for future research.




Figure 5: Qualitative comparison on Action Genome predicate classification task. The gray box is the wrongly recognized predicates. Our method performs better than STTran [7] on recognizing the tail and the head both.




Figure 6: Qualitative comparison on VidOR predicate classification. The Social Fabric baseline [6] misses the *lean_on* predicate, while our method detects it correctly.

REFERENCES

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual Question Answering. In *ICCV*. IEEE Computer Society, Santiago, Chile, 2425–2433.
- [2] Ondrej Bohdal, Yongxin Yang, and Timothy M. Hospedales. 2021. EvoGrad: Efficient Gradient-Based Meta-Learning and Hyperparameter Optimization. In *NeurIPS*. Neural Information Processing Systems Foundation, Virtual, 22234–22246.
- [3] Qianwen Cao, Heyan Huang, Xindi Shang, Boran Wang, and Tat-Seng Chua. 2021. 3-D Relation Network for visual relation recognition in videos. *Neurocomputing* 432 (2021), 91–100.
- [4] Shuo Chen, Pascal Mettes, Tao Hu, and Cees GM Snoek. 2020. Interactivity Proposals for Surveillance Videos. In *ICMR*. ACM, Dublin, Ireland, 108–116.
- [5] Shuo Chen, Pascal Mettes, and Cees GM Snoek. 2021. Diagnosing Errors in Video Relation Detectors. In *BMVC*. BMVA Press, Online, 241.
- [6] Shuo Chen, Zenglin Shi, Pascal Mettes, and Cees GM Snoek. 2021. Social Fabric: Tubelet Compositions for Video Relation Detection. In *ICCV*. IEEE, Montreal, QC, Canada, 13465–13474.
- [7] Yuren Cong, Wentong Liao, Hanno Ackermann, Bodo Rosenhahn, and Michael Ying Yang. 2021. Spatial-temporal transformer for dynamic scene graph generation. In *ICCV*. IEEE, Montreal, QC, Canada, 16352–16362.
- [8] Alakh Desai, Tz-Ying Wu, Subarna Tripathi, and Nuno Vasconcelos. 2021. Learning of visual relations: The devil is in the tails. In *ICCV*. IEEE, Montreal, QC, Canada, 15384–15393.
- [9] Xingning Dong, Tian Gan, Xueming Song, Jianlong Wu, Yuan Cheng, and Liqiang Nie. 2022. Stacked Hybrid-Attention and Group Collaborative Learning for Unbiased Scene Graph Generation. In *CVPR*. IEEE, New Orleans, LA, USA, 19405–19414.
- [10] Kaifeng Gao, Long Chen, Yulei Niu, Jian Shao, and Jun Xiao. 2022. Classification-then-grounding: Reformulating video scene graphs as temporal bipartite graphs. In *CVPR*. IEEE, New Orleans, LA, USA, 19475–19484.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*. IEEE Computer Society, Las Vegas, NV, USA, 770–778.
- [12] Jingwei Ji, Ranjay Krishna, Li Fei-Fei, and Juan Carlos Niebles. 2020. Action genome: Actions as compositions of spatio-temporal scene graphs. In *CVPR*. Computer Vision Foundation / IEEE, Seattle, WA, USA, 10233–10244.
- [13] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. 2015. Image retrieval using scene graphs. In *CVPR*. IEEE Computer Society, Boston, MA, USA, 3668–3678.
- [14] Anna Kukleva, Makarand Tapaswi, and Ivan Laptev. 2020. Learning Interactions and Relationships Between Movie Characters. In *CVPR*. Computer Vision Foundation / IEEE, Seattle, WA, USA, 9846–9855.
- [15] Rongjie Li, Songyang Zhang, Bo Wan, and Xuming He. 2021. Bipartite graph network with adaptive message passing for unbiased scene graph generation. In *CVPR*. Computer Vision Foundation / IEEE, Virtual, 11109–11119.
- [16] Wei Li, Haiwei Zhang, Qijie Bai, Guoqing Zhao, Ning Jiang, and Xiaojie Yuan. 2022. PDDL: Predicate Probability Distribution Based Loss for Unbiased Scene Graph Generation. In *CVPR*. IEEE, New Orleans, LA, USA, 19425–19434.
- [17] Yiming Li, Xiaoshan Yang, and Changsheng Xu. 2022. Dynamic Scene Graph Generation via Anticipatory Pre-Training. In *CVPR*. IEEE, New Orleans, LA, USA, 13864–13873.
- [18] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. 2017. Focal Loss for Dense Object Detection. In *ICCV*. IEEE Computer Society, Venice, Italy, 2999–3007.
- [19] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *ECCV*. Springer, Zurich, Switzerland, 740–755.
- [20] Chenchen Liu, Yang Jin, Kehan Xu, Guoqiang Gong, and Yadong Mu. 2020. Beyond Short-Term Snippet: Video Relation Detection with Spatio-Temporal Global Context. In *CVPR*. Computer Vision Foundation / IEEE, Seattle, WA, USA, 10837–10846.
- [21] Hengyue Liu, Ning Yan, Masood Mortazavi, and Bir Bhanu. 2021. Fully convolutional scene graph generation. In *CVPR*. Computer Vision Foundation / IEEE, Virtual, 11546–11556.
- [22] Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *ICLR*. OpenReview.net, New Orleans, LA, USA.
- [23] Xufeng Qian, Yueting Zhuang, Yimeng Li, Shaoning Xiao, Shiliang Pu, and Jun Xiao. 2019. Video relation detection with spatio-temporal graph. In *ACM MM*. ACM, Nice, France, 84–93.
- [24] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. 2018. Learning to reweight examples for robust deep learning. In *ICML*. PMLR, Stockholm, Sweden, 4331–4340.
- [25] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *NeurIPS*. Neural Information Processing Systems Foundation, Montreal, Quebec, Canada, 91–99.
- [26] Xindi Shang, Donglin Di, Junbin Xiao, Yu Cao, Xun Yang, and Tat-Seng Chua. 2019. Annotating objects and relations in user-generated videos. In *ICMR*. ACM, Ottawa, ON, Canada, 279–287.
- [27] Xindi Shang, Tongwei Ren, Jingfan Guo, Hanwang Zhang, and Tat-Seng Chua. 2017. Video visual relation detection. In *ACM MM*. ACM, Mountain View, CA, USA, 1300–1308.
- [28] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. 2019. Meta-weight-net: Learning an explicit mapping for sample weighting. In *NeurIPS*. Neural Information Processing Systems Foundation, Vancouver, BC, Canada, 1917–1928.
- [29] Cees G. M. Snoek and Marcel Worring. 2009. Concept-Based Video Retrieval. *Found. Trends Inf. Retr.* 2, 4 (2009), 215–322.
- [30] Zixuan Su, Xindi Shang, Jingjing Chen, Yu-Gang Jiang, Zhiyong Qiu, and Tat-Seng Chua. 2020. Video Relation Detection via Multiple Hypothesis Association. In *ACM MM*. ACM, Virtual Event / Seattle, WA, USA, 3127–3135.
- [31] Xu Sun, Tongwei Ren, Yuan Zi, and Gangshan Wu. 2019. Video visual relation detection via multi-modal feature fusion. In *ACM MM*. ACM, Nice, France, 2657–2661.
- [32] Sai Praneeth Reddy Sunkesula, Rishabh Dabral, and Ganesh Ramakrishnan. 2020. LIGHTEN: Learning Interactions with Graph and Hierarchical Temporal Networks for HOI in Videos. In *ACM MM*. ACM, Virtual Event / Seattle, WA, USA, 691–699.
- [33] Leitian Tao, Li Mi, Nannan Li, Xianhang Cheng, Yaosi Hu, and Zhenzhong Chen. 2022. Predicate Correlation Learning for Scene Graph Generation. *TIP* 31 (2022), 4173–4185.
- [34] Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. 2016. YFCC100M: the new data in multimedia research. *Commun. ACM* 59, 2 (2016), 64–73.
- [35] Junjiao Tian, Niluthpol Chowdhury Mithun, Zachary Seymour, Han-Pang Chiu, and Zsolt Kira. 2022. Striking the Right Balance: Recall Loss for Semantic Segmentation. In *ICRA*. IEEE, Philadelphia, PA, USA, 5063–5069.
- [36] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. 2017. Simple online and realtime tracking with a deep association metric. In *ICIP*. IEEE, Beijing, China, 3645–3649.
- [37] Tong Wu, Qingqiu Huang, Ziwei Liu, Yu Wang, and Dahua Lin. 2020. Distribution-balanced loss for multi-label classification in long-tailed datasets. In *ECCV*. Springer, Glasgow, UK, 162–178.
- [38] Wentao Xie, Guanghui Ren, and Si Liu. 2020. Video Relation Detection with Trajectory-aware Multi-modal Features. In *ACM MM*. ACM, Virtual Event / Seattle, WA, USA, 4590–4594.
- [39] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. 2017. Scene graph generation by iterative message passing. In *CVPR*. IEEE Computer Society, Honolulu, HI, USA, 3097–3106.
- [40] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard Zemel, and Yoshua Bengio. 2015. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *ICML*. JMLR.org, Lille, France, 2048–2057.
- [41] Shaotian Yan, Chen Shen, Zhongming Jin, Jianqiang Huang, Rongxin Jiang, Yaowu Chen, and Xian-Sheng Hua. 2020. PCPL: Predicate-Correlation Perception Learning for Unbiased Scene Graph Generation. In *ACMMM*. ACM, Virtual Event / Seattle, WA, USA, 265–273.
- [42] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. 2018. Graph r-cnn for scene graph generation. In *ECCV*. Springer, Munich, Germany, 690–706.
- [43] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. 2018. Neural motifs: Scene graph parsing with global context. In *CVPR*. Computer Vision Foundation / IEEE Computer Society, Salt Lake City, UT, USA, 5831–5840.
- [44] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. 2021. Deep long-tailed learning: A survey. *arXiv:2110.04596*
- [45] Sipeng Zheng, Xiangyu Chen, Shizhe Chen, and Qin Jin. 2019. Relation understanding in videos. In *ACM MM*. ACM, Nice, France, 2662–2666.