

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

OMG: Observe Multiple Granularities for Natural Language-Based Vehicle Retrieval

Yunhao Du¹*, Binyu Zhang¹*, Xiangning Ruan¹, Fei Su^{1,2}, Zhicheng Zhao^{1,2}[†], Hong Chen³ ¹Beijing University of Posts and Telecommunications ²Beijing Key Laboratory of Network System and Network Culture, China ³China Mobile Research Institute

{dyh_bupt, zhangbinyu, ruan, zhaozc, sufei}@bupt.edu.cn

chenhongyj@chinamobile.com

Abstract

Retrieving tracked-vehicles by natural language descriptions plays a critical role in smart city construction. It aims to find the best match for the given texts from a set of tracked vehicles in surveillance videos. Existing works generally solve it by a dual-stream framework, which consists of a text encoder, a visual encoder and a cross-modal loss function. Although some progress has been made, they failed to fully exploit the information at various levels of granularity. To tackle this issue, we propose a novel framework for the natural language-based vehicle retrieval task, OMG, which Observes Multiple Granularities with respect to visual representation, textual representation and objective functions. For the visual representation, target features, context features and motion features are encoded separately. For the textual representation, one global embedding, three local embeddings and a color-type prompt embedding are extracted to represent various granularities of semantic features. Finally, the overall framework is optimized by a cross-modal multi-granularity contrastive loss function. Experiments demonstrate the effectiveness of our method. Our OMG significantly outperforms all previous methods and ranks the 9th on the 6th AI City Challenge Track2. The codes are available at https://github.com/dyhBUPT/OMG.

1. Introduction

With the development of the artificial intelligence technology, smart transportation and city has elicited increasing attention in recent years. Vehicle retrieval, which aims to retrieve tracked vehicles given a query, occupies an important position. Most previous works solve this problem under a purely visual setting, which try to identify the images or



Figure 1. Visualization examples of the natural language-based vehicle retrieval task. The query generally describes the color, type, motion state and the relations (with neighbor objects) of the target vehicle.

tracklets of the same vehicle across different surveillance camera views [20, 23, 39]. However, image-based setting limits its application scenarios, bacause an image query is not always available. Comparatively speaking, Natural Language (NL) descriptions can be one of the most convenient way to give a query [7]. Therefore, it is not trivial to develop an effective system for natural language-based vehicle retrieval. Figure 1 gives example frames, bounding boxes and corresponding descriptions from the CityFlow-NL dataset [7], which is a large-scale benchmark for the text-based vehicle retrieval task.

Existing works tend to solve this task by using a dualstream framework [8] because of its flexibility and efficiency [22]. This framework consists of separate visual encoder and text encoder to extract modal-specific representations respectively and then computes their similarity. Then it is optimized by InfoNCE loss [42], Instance loss [46], etc. Though these methods have made some progress, they fail to mine adequate multi-granularity information. For instance, the relationships between the target vehicle and its

^{*}Equal contribution

[†]Corresponding authors

neighbors is ignored in [1, 14, 37, 38], and [30, 36] fail to fully utilze the color and type information. Moreover, they don't take good advantage of rich textual granular features.

In this paper, we propose OMG, which observes multiple granularities in an unified framework. It follows the dual-stream architecture, that is, one visual stream and one textual stream. The visual stream takes a cropped vehicle image, a context image and a foreground motion map as input, and then extracts their features by three individual image encoders. As for the textual stream, one global text, three local texts and one prompt are encoded separately by five tied text encoders. While training, a multi-granularity InfoNCE loss is applied to perform cross-modal contrastive learning between them. Furthermore, an ID loss is also used for the visual stream for more discriminative representation.

Extensive experiments verify the effectiveness of our proposed methods. To be specific, by taking "ResNet50 + BERT + InfoNCE" as baseline [10, 41, 42], a series of proposed optimization strategies boost its performance by **144%** on the validation set (MRR from 0.208 to 0.507). Entries based on the presented methodology ranks the 9th on the AI CIty Challenge 2022 Track2. The contributions of this work are summarized as follows:

1) We propose a novel framework for the natural language-based vehicle retrieval task, which fully exploits the multi-granularity information in a joint manner.

2) Various optimization tricks are employed to further improve the representation capability of the network, i.e., auxiliary ID loss, backtranslation-based text augmentation, stronger image and text encoders [24, 35].

3) Extensive experiments prove the effectiveness of our method.

2. Related Works

2.1. Text-based Video Retrieval

Text-based video retrieval aims to search the corresponding video via the given text description among a number of videos. Due to its great range of potential applications, it has received wide attention by the researchers. The most of existing methods adopt contrastive learning to train a network which can estimate the similarity between the video and the text descriptions. They take both video and text descriptions as input, and then extract the video features and text features via the modal encoders. They tend to update the networks by minimizing the distance between each pair of the video-text features. Due to the difference of the modality and limitation of the network, the early methods [18,28,29] usually use a dual-stream framework, which can reduce the computational complexity. These works encode the video by the visual feature extractor (C3D [40], I3D [2], etc.) and encode the text by the textual feature extractor (LSTM [11], GRU [4], etc.), and then estimate the video-text similarity in a potential semantic space.

Recently, with the successful migration of Transformer [43] from natural language processing to computer vision, (ViT [5], Swin-Transformer [24], etc.), the mainstream methods [6,9,16,17,19,22,26,35,44,45] on video-language retrieval tasks begin to utilize Transformer as encoders for both of the video and the natural language. HERO [17] and ALPRO [16] explore to boost the video-text alignment via the large-scale pre-training tasks. HiT [22] further proposed a hierarchical model with momentum contrast for videotext retrieval. Also, some single-stream networks come out, due to the consistency of the Transformer encoders structure. ViLBERT [26] firstly explore to utilize a single-stream Transformer network to retrieve the video based on natural language descriptions, which also attempts to align the corresponding video and text features in the potential semantic space. Further, due to the success of CLIP [35], which demonstrates the powerful performance of image-text contrastive learning, many researchers conduct experiments on the basis of CLIP. Besides, CLIP2Video [6] extends CLIP from image-text retrieval to video-text retrieval. Moreover, MFGATN [9], TACo [45], HANet [44] and FeatInter [19] also achieve good performance in the video-text retrieval tasks by designing different modules to enhance the comprehension of the video-text pair.

2.2. Text-based Vehicle Retrieval

AYCE [36] proposes a modular solution which applies BERT [41] to embed textual descriptions and a CNN [10] with a Transformer model [43] to embed visual information. SBNet [15] presents a substitution module that helps project features from different domains into the same space, and a future prediction module to learn temporal information by predicting the next frame. Pirazh et al. [14] and Tam et al. [30] adopts CLIP [35] to extract frame features and textual features. TIED [37] proposes an encoder-decoder based model in which the encoder embeds two modalities into the common space and the decoder jointly optimizes these embeddings by an input-token-reconstruction task. Tien-Phat et al. [31] adapts COOT [8] to model the cross-modal relationships with both appearance and motion attributes. Eun-Ju et al. [33] propose to perform color and type classification for both target and front-rear vehicles, and conduct movement analysis based on the Kalman filter algorithm [13]. DUN [38] uses pretrained CNN and GloVe [34] to extract modal-specific features and GRUs [3] to exploit temporal information. CLV [1] proposes the simple and effective global motion image, which is jointly encoded with the cropped image and language description, and ranks first on the AI City Challenge 2021 Track5.

Though these methods have achieved promising results on the CityFlow-NL benchmark, they don't utilze omnifeatures adequately. Differently, our method can mine rich



Figure 2. Framework of the proposed OMG, which consists of a visual stream and a textual stream. For the visual stream, it takes a cropped target image, a context image and a 4-channel foreground motion map as input, and encodes them into a target feature f_t^i , a context feature f_c^i and a motion feature f_m^i respectively. As to the textual stream, it embdes one global feature f_g^t , three local features f_l^t and one prompt feature f_p^t from different granularity of textual inputs. The overall framework is optimized with a multi-granularity InfoNCE loss and an ID loss in an end-to-end manner.

information from multiple granularities, which helps extract more discriminative features.

3. Method

In this section, we present the details of the OMG framework as shown in Figure 2. Specifically, we start with the visual stream in Section 3.1, and then introduce the textual stream in Section 3.2. Finally, Section 3.3 illustrates the multi-granularity InfoNCE loss and the auxiliary ID loss.

3.1. Multi-Granularity Visual Stream

For the visual stream, three granularities of images are inputed into image encoders to extract features respectively. Note that these three encoders don't share weights.

Target Features. Given a trajectory of the target vehicle, a frame is sampled and the target image is cropped according to the bounding box. Then it is taken as the input to generate the target feature f_t^i . By setting a high resolution, the network can focus on the details and implicitly encode the appearance information, e.g., color and type.

Context Features. The target feature only contains the target vehicle's own information by excluding the background, ignoring the relationship between neighbors. Instead, we extend the original bounding box [x, y, w, h] to [x - w, y - h, 3w, 3h] to crop the context image around the target vehicle. It has a larger receptive field and can usually

include nearby vehicles, which helps the encoder embed the relation information and predict the context feature f_c^i .

Motion Features. Another vital information is the motion pattern of the target vehicle. Inspired by the background modeling based motion image in [1], we propose a 4-channel foreground motion map, which consists of a 3-channel foreground cropped image and a single channel motion image. For the former one, the whole trajectory of the target vehicle is cropped and pasted into a black image. Note that we only keep the bounding boxes that don't overlap too much. For the latter one, the center positions of these bounding boxes are used to form a thick trajectory line. Finally, these two images are concatenated along the channel dimension and taken as the input of the image encoder. It's encoded into the motion feature f_m^i , which contains both coarse appearance and motion information simutaneously.

3.2. Multi-Granularity Textual Stream

Similar with the visual stream, the textual stream also contains three granularities of inputs, i.e., the global text, the local texts and the color-type prompt. The notable difference between these two streams is that the weights of all text encoders are shared and freezed except for the last fully connected layer.

Global Features. Every trajectory corresponds to three unique natural language sentences. The descriptions of



Figure 3. The illustration of the color-type prompt generation process, which consists of three steps, i.e., dependency parsing, voting and prompt generation. Note that we only visualize the dependency parsing results of one sentence for simplicity, i.e., "A blue wagon going straight down the street passing an intersection".

these sentences would have different emphases. For example, some may focus on the appearance and motion information (e.g., A black sedan makes a right turn at the intersection.), and others may contain the relation information (e.g., A white SUV switches to the left lane followed by another white vehicle.). In order to comprehensively consider them, an intuitive and simple way is to concatenate these three sentences into an unified one, termed the global text. The extracted global feature f_g^t from it includes rich semantic information from a global view.

Local Features. Though the global feature embeds the integrated semantics, it is too coarse to focus on details. To make the network pay attention to more fine-grained information, we also input the three sentences into the network respectively to predict three local features f_{l1}^t , f_{l2}^t , f_{l3}^t . Different from f_g^t , the f_{l*}^t places more emphasis on some specific features, e.g., appearance, motion and relation.

Prompt Features. The appearance would be the most vital information to distinguish different vehicles. Inspired by the prompt learning in Natural Language Processing (NLP) [21], we design a color-type prompt for better appearance representation. The prompt generation process is illustrated in the Figure 3. Specifically, the color and type of the vehicle are extracted from each sentence first by using a dependency parser [12]. Then the final color and type are determined by a simple voting mechanism ("gray" and "SUV" in the figure). Finally, the color-type prompt is generated by the template "*This is a [COLOR] [TYPE]*", e.g., "*This is a gray SUV*" in this example. By using this simple and intuitive prompt as input and predicting the prompt feature f_p^t , it is easier to focus on the appearance information for the text encoder.

3.3. Multi-Granularity Loss

InfoNCE loss [42] is a widely used objective function for contrastive learning, which is then extended to crossmodal tasks. In this work, we apply it to our cross-modal multi-granularity framework and name it *Multi-Granularity InfoNCE loss*. Given a batch with M text-vehicle pairs $\{(t_i, v_i)\}_{i=1}^M$, we first extract their multi-granularity features. Particularly, every t_i is embedded into N_t features $\{t_i^j\}_{j=1}^{N_t}$ and every v_i is embedded into N_v features $\{v_i^k\}_{k=1}^{N_v}$, where $N_t = 5$ and $N_v = 3$ in this work. Then the multi-granularity InfoNCE loss from text to image is:

$$L_{info}^{t2i} = \frac{1}{MN_tN_v} \sum_{i=1}^{M} \sum_{j=1}^{N_t} \sum_{k=1}^{N_v} \log(\frac{e^{s(t_i^j, v_i^k)/\tau}}{\sum_{n=1}^{M} e^{s(t_i^j, v_n^k)/\tau}}),$$
(1)

where τ denotes a learnable temperature parameter and $s(\cdot, \cdot)$ measures the cosine similarity as:

$$s(u,v) = \frac{u^T v}{||u|| \, ||v||}.$$
(2)

Similarly, the image-to-text one is:

$$L_{info}^{i2t} = \frac{1}{MN_t N_v} \sum_{i=1}^{M} \sum_{j=1}^{N_t} \sum_{k=1}^{N_v} log(\frac{e^{s(v_i^k, t_i^j)/\tau}}{\sum_{n=1}^{M} e^{s(v_i^k, t_n^j)/\tau}}).$$
(3)

Finally, the whole multi-granularity InfoNCE loss is:

$$L_{info} = (L_{info}^{t2i} + L_{info}^{i2t})/2.$$
 (4)

Though L_{info} can help improve the representational ability in a contrastive learning manner, it treats all samples as the negatives except that in the anchor pair. However, some vehicles from different text-vehicle pairs have

Table 1. Ablation study of OMG on the validation set. "↑" means higher is better.

Method	BoT	Context	Motion	Local	Prompt	CLIP	NLAug	IDLoss	Swin	MFT	MRR(†)	R@5 (↑)	R@10(†)
Baseline											0.208	0.269	0.363
	 ✓ 										0.218	0.285	0.382
	 ✓ 	\checkmark									0.228	0.315	0.422
	 ✓ 	\checkmark	\checkmark								0.204	0.263	0.435
	1			\checkmark							0.241	0.333	0.457
	1			\checkmark	 ✓ 						0.273	0.384	0.505
	 ✓ 	\checkmark		\checkmark	 ✓ 						0.285	0.411	0.556
	1	\checkmark		\checkmark	 ✓ 	\checkmark					0.316	0.435	0.586
	1	\checkmark		\checkmark	 ✓ 	\checkmark	\checkmark				0.401	0.567	0.723
	\checkmark	\checkmark		\checkmark	 ✓ 	\checkmark	\checkmark	\checkmark			0.438	0.591	0.750
		\checkmark		\checkmark	 ✓ 	\checkmark	\checkmark	\checkmark	\checkmark		0.490	0.672	0.790
OMG		\checkmark		\checkmark	 ✓ 	 ✓ 	\checkmark	\checkmark	\checkmark	 ✓ 	0.507	0.667	0.812

the same ID. To be specific, there are 2,155 pairs in the traning set of the CityFlow-NL dataset [7], but they are from C = 482 vehicles from different views. It would deteriorate the feature learning by treating these vehicles with the smae ID as negative samples, especially for the image encoders. For that reason, we introduce the cross-entropy loss on the multi-granularity visual features $\{v_i^k\}_{k=1}^{N_v}$ as an auxiliary ID loss. Given a vehicle v_i , we denote y_i as the truth ID label and $p_{i,c}^k$ as the ID prediction logits of class c for granularity k. The ID loss is computed as follows:

$$L_{id} = \frac{1}{MN_v} \sum_{i=1}^{M} \sum_{c=1}^{C} \sum_{k=1}^{N_v} -q_c log(p_{i,c}^k),$$
(5)

where $q_c = 1$ if $y_i = c$, otherwise $q_c = 0$.

Finally, the overall objective function is formulated as:

$$L_{OMG} = \lambda_1 L_{info} + \lambda_2 L_{id}, \tag{6}$$

where λ_1, λ_2 are the weights for the two losses. We set $\lambda_1 = 1, \lambda_2 = 1$ in our experiments.

4. Experiments

4.1. Dataset and Evaluation

We use the CityFlow-NL dataset [7] to train and evaluate our model, and it is extended from the CityFlow benchmark [39] by annotating natural language descriptions for vehicle targets. It is the first city-scale multi-target multi-camera tracking with natural language descriptions dataset that provides precise descriptions for multi-view ground truth vehicle tracks. The CityFlow-NL dataset consists of 666 targets vehicles in 3,028 single-view tracks and 5,289 unique NL descriptions. For the AI City Challenge, the dataset contains 2,498 tracks of vehicles with three descriptions. 530 unique vehicle tracks together with 530 query sets each with three descriptions are curated for this challenge.

The natural language-based vehicle retrieval task uses the Mean Reciprocal Rank (MRR) as the main evaluation metric. Specifically, the reciprocal rank of a query response is the multiplicative inverse of the rank of the first correct answer. The mean reciprocal rank is the average of the reciprocal ranks of the overall test set:

$$MRR = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{rank_i},\tag{7}$$

where $rank_i$ is the rank position of the first correct answer for the *i*-th query. Besides, Recall@5 and Recall@10 are also evaluated.

4.2. Implement Details

In the training phase, we train OMG with batch size 24 on three 16G Tesla T4 GPUs. For the visual stream, we first randomly sample a frame from the target trajectory. Then the input images are resized to 384×384 . *Random-Crop* and *RandomApply* are used to perform data augmentation. For the textual stream, the backbone is freezed except the last fully connected layer. The cosine annealing strategy is used with a base learning rate 6.7e-3 and a minimum learning rate 6.7e-6. The network is trained with the AdamW [25] optimizer for 600 epochs with 300 warm-up epochs.

While inference, the middle frame of the trajectory is sampled as the input of the visual stream. All $N_t N_v$ pairs of cross-modal multi-granularity of features are used to compute pairwise similarities. Then they are simply averaged for the final similarity prediction.

4.3. Ablation Study

In this section, we summarize the ablation study of our method. First, we conduct the ablation experiments for our OMG model, which has been introduced previously. Moreover, we also present a single-granularity model, named OSG (Observe Single Granularity) for the ensemble.

Ablation of OMG. Table 1 shows the ablation study of our OMG. The baseline is a dual-stream architecture with

Table 2. Ablation study of the OSG on the validation set. "↑" means higher is better.

Method	NLAug	Color	Туре	CLIP	D	Motion	Swin	MRR(†)	R@5(↑)	R@10(†)
Baseline								0.208	0.269	0.363
	\checkmark							0.212	0.293	0.401
	\checkmark	\checkmark						0.222	0.312	0.401
	\checkmark	\checkmark	 ✓ 					0.224	0.280	0.422
	\checkmark	\checkmark		\checkmark				0.326	0.433	0.566
	\checkmark	\checkmark		\checkmark	\checkmark			0.345	0.487	0.626
	\checkmark	\checkmark		\checkmark	\checkmark	\checkmark		0.380	0.505	0.621
OSG	\checkmark	\checkmark		✓	✓	\checkmark	\checkmark	0.441	0.567	0.685

ResNet50 [10], BERT [26] and InfoNCE loss [42]. Only the target image and the global text are used as input for the visual stream and textual stream respectively. Other optimization strategies are illustrated as follows:

- BoT: Pre-training a ReID model BoT [27] on Track1 with ResNet50-IBN [32] as the backbone to replace the original image encoder. Experiments shows that the stronger visual encoder brings 0.01 improvement for MRR.
- Context: Taking the context image as the input to extract context features. With the help of the context information, all metrics MRR, R@5 and R@10 improve.
- Motion: Taking the motion image as the input to extract motion features. The use of the foreground motion image doesn't boost the performance, perhaps beacuse of the too rough appearance information in it. Therefore, we discard it in the future version of our framework.
- Local: Adding three local texts as the input of the textual stream. Finer granularities of textual descriptions improve the performance by a large margin (from 0.218 to 0.241).
- Prompt: Adding the prompt text as the input of the textual stream. Though its simplicity, this simple colortype prompt boosts the performance significantly.
- CLIP: Replacing the text encoder BERT with the textual stream of CLIP [35]. The improvements can be attributed to its excellent ability of modal alignment.
- NLAug: Data augmentation with the backtranslation method for the natural language data. We can observe a huge gap between the effect of *NLAug* for OMG and OSG (in Table 2). Specifically, it improves the MMR of OMG by 0.085 (from 0.316 ro 0.401), but only 0.004 for OSG (from 0.208 to 0.212). The reason behind the gap is that the multi-granularity architecture makes better use of the augmented data.

- IDLoss: Applying the auxiliary ID loss. The ID loss helps the image encoder achieve more robust representation ability and improves the retrieval performance.
- Swin: Replacing the image encoder ResNet50 with the Swin-Transformer-B [24]. A stronger architecture of the image encoder further boosts the performance.
- MFT: Multi-frame testing, i.e., sampling uniformly 8 frames for the target trajectory to generate the mean features instead of only one frame. Experiments proves its effectiveness.

Ablation of OSG. We also conduct various experiments on a single granularity model (OSG) with different modules, as illustrated in Table 2. It has the same baseline with OMG. The "NLAug" denotes the same natural language augmentation in Table 1, which can also improve the performance of OSG slightly. The terms "Color", "Type" and "ID" indicate that we utilize additional classification heads to classify the color, type and ID of the vehicle. As shown in the Table 2, the color and ID supervision gains a relative MRR improvement which demonstrates the effectiveness of the supervision for the vehicle properties. However, the type of the vehicle can hardly help the OSG perform better. We infer this is due to the variety descriptions of vehicle type for the same target, such as "truck", "pickup", "hatchback", etc., which may lead to confusion. We also replace BERT with the text encoder of CLIP, and replace ResNet50 with Swin-Transformer, which gives a great boost to our model. To get the correct and representative motion information from the annotation, we first filter out the bounding box if its IoU is greater than 0.9 with the box in the previous frames. Then we uniformly sample 16 boxes from the remaining ones. At last, it is inputed into a GRU [4] to embed the motion features in the potential semantic space. The introduction of motion information enhances the model performance, which indicates the motion information plays an important role in this task.

4.4. Evaluation Results

Quantitative Results. We compare our OMG with previous state-of-the-art methods in Table 3. It is shown that our





A blue sedan drives through an intersection. A blue sedan moves straight passing cargo truck on the intersection. A blue sedan runs through an intersection.



A red sedan switches lane to left and passes a cargo truck. A red sedan keeps straight through an intersection. A dark-red car is going straight.





A brown sedan runs down the street followed by another vehicle. Convertible dark-colored car with a gray car behind it. A black sedan keeps straight followed by another vehicle.



Figure 4. Sample retrieval results visualization of our OMG on the test set of the CityFlow-NL dataset. Only the rank 1 3 results are listed and only 3 frames are sampled for each trajectory. The image is cropped for clarity. The target vehicle retrieved are marked with green bounding boxes. It's shown that our model has a certain semantic understanding ability.

Table 3. Comparision with previous methods.

Team	MRR
OMG(ours)	0.3012
Alibaba-UTS-ZJU [1]	0.1869
SDU-XidianU-SDJZU [38]	0.1613
SUNYKorea [33]	0.1594
Sun Asterisk [30]	0.1571
HCMUS [31]	0.1560
TUE [37]	0.1548
JHU-UMD [14]	0.1364
Modulabs-Naver-KookminU [15]	0.1195
Unimore [36]	0.1078

Table 4. The leaderboard of the 6th AI City Challenge Track2.

Rank	Team	MRR
1	Must Win	0.6606
2	Thursday	0.5251
3	HCMIU-CVIP	0.4773
4	MegVideo	0.4392
5	HCMUS	0.3611
6	P & L	0.3338
7	Terminus-AI	0.3320
8	MARS_WHU	0.3205
9	BUPT_MCPRL_T2(ours)	0.3012
10	folklore	0.2832
11	HYFL	0.2804
12	alpha	0.2802
13	SEEE-HUST	0.2333
14	ETRI_AIA	0.0389
15	Pair Lab	0,0216

method surpasses them by a large margin. Specifically, our method is 61.1% better than the last year's winning team, i.e., 0.3012 vs. 0.1869. Table 4 lists the top-15 team results on the AI City 2022 Challenge Track2 and we rank the 9th palce. Note that the final submission are the ensemble results from both OMG and OSG (in Section 4.3)

Qualitative Results. Figure 4 visualizes the retrieval results of OMG on the test set of the CityFlow-NL dataset. Only the top-3 results are listed here.

4.5. Limitations

Our OMG still has several limitations. First of all, though it can mine rich information from multiple granularities, we simply average the similarities between different cross-modal pairs. Improving the mechanism for fusing these multi-granularity features is expected to bring significant improvements, e.g., using the self-attention mechanism [43]. Besides, the leverage of the motion and relation information should be further improved. The temporal information is also underutilized. Data augmentation is vital for representation learning, which should be studied further. Last but not least, the cross-modal alignment problem is not trivial for the text-vision retrieval task, which is not given enough attention in this work. A more effective architecture or alignment mechanism would help boost the retrieval performance.

5. Conclusion

In this paper, we propose a novel framework for the natural language-based vehicle retrieval task, named OMG. It follows the dual-stream paradigm, which consists of a multi-granularity visual stream and a multi-granularity textual stream. By focusing on three-granularity of visual information (target, context, motion) and three-granularity of textual information (global, local, prompt), our model can mine rich semantics jointly with a multi-granularity InfoNCE loss and an auxiliary ID loss. With some optimization tricks, our method improves the performance of the baseline by 144% on the validation set. On the test set, OMG outperforms all previous methods by a large margin and ranks 9th on the 6th AI City Challenge Track2.

In the future, we will focus on discriminative representation learning and cross-modal alignment.

6. Acknowledgements

This work is supported by Chinese National Natural Science Foundation under Grants (62076033, U1931202) and MoE-CMCC "Artifical Intelligence" Project No.MCM20190701.

References

- [1] Shuai Bai, Zhedong Zheng, Xiaohan Wang, Junyang Lin, Zhu Zhang, Chang Zhou, Hongxia Yang, and Yi Yang. Connecting language and vision for natural language-based vehicle retrieval. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 4034– 4043, 2021. 2, 3, 8
- [2] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 6299–6308, 2017. 2
- [4] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. arXiv preprint arXiv:1406.1078, 2014. 2, 6

- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020. 2
- [6] Han Fang, Pengfei Xiong, Luhui Xu, and Yu Chen. Clip2video: Mastering video-text retrieval via image clip. arXiv preprint arXiv:2106.11097, 2021. 2
- [7] Qi Feng, Vitaly Ablavsky, and Stan Sclaroff. Cityflow-nl: Tracking and retrieval of vehicles at city scale by natural language descriptions. *arXiv preprint arXiv:2101.04741*, 2021.
 1, 5
- [8] Simon Ging, Mohammadreza Zolfaghari, Hamed Pirsiavash, and Thomas Brox. Coot: Cooperative hierarchical transformer for video-text representation learning. *Advances in neural information processing systems*, 33:22605–22618, 2020. 1, 2
- [9] Xiaoshuai Hao, Yucan Zhou, Dayan Wu, Wanqian Zhang, Bo Li, and Weiping Wang. Multi-feature graph attention network for cross-modal video-text retrieval. In *Proceedings of the 2021 International Conference on Multimedia Retrieval*, pages 135–143, 2021. 2
- [10] 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
- [11] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [12] Matthew Honnibal and Mark Johnson. An improved nonmonotonic transition system for dependency parsing. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 1373–1378, 2015. 4
- [13] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. 1960. 2
- [14] Pirazh Khorramshahi, Sai Saketh Rambhatla, and Rama Chellappa. Towards accurate visual and natural languagebased vehicle retrieval systems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4183–4192, 2021. 2, 8
- [15] Sangrok Lee, Taekang Woo, and Sang Hun Lee. Sbnet: Segmentation-based network for natural language-based vehicle search. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4054– 4060, 2021. 2, 8
- [16] Dongxu Li, Junnan Li, Hongdong Li, Juan Carlos Niebles, and Steven CH Hoi. Align and prompt: Video-andlanguage pre-training with entity prompts. arXiv preprint arXiv:2112.09583, 2021. 2
- [17] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. Hero: Hierarchical encoder for video+ language omni-representation pre-training. arXiv preprint arXiv:2005.00200, 2020. 2
- [18] Xirong Li, Jianfeng Dong, Chaoxi Xu, Jing Cao, Xun Wang, and Gang Yang. Renmin university of china and zhejiang gongshang university at trecvid 2018: Deep cross-modal embeddings for video-text retrieval. In *TRECVID*, 2018. 2

- [19] Baolong Liu, Qi Zheng, Yabing Wang, Minsong Zhang, Jianfeng Dong, and Xun Wang. Featinter: Exploring fine-grained object features for video-text retrieval. *Neurocomputing*, 2022. 2
- [20] Hongye Liu, Yonghong Tian, Yaowei Yang, Lu Pang, and Tiejun Huang. Deep relative distance learning: Tell the difference between similar vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2167–2175, 2016. 1
- [21] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. arXiv preprint arXiv:2107.13586, 2021. 4
- [22] Song Liu, Haoqi Fan, Shengsheng Qian, Yiru Chen, Wenkui Ding, and Zhongyuan Wang. Hit: Hierarchical transformer with momentum contrast for video-text retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11915–11925, 2021. 1, 2
- [23] Xinchen Liu, Wu Liu, Tao Mei, and Huadong Ma. A deep learning-based approach to progressive vehicle reidentification for urban surveillance. In *European conference* on computer vision, pages 869–884. Springer, 2016. 1
- [24] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 10012–10022, 2021. 2, 6
- [25] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101, 2017. 5
- [26] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. Advances in neural information processing systems, 32, 2019. 2, 6
- [27] Hao Luo, Wei Jiang, Youzhi Gu, Fuxu Liu, Xingyu Liao, Shenqi Lai, and Jianyang Gu. A strong baseline and batch normalization neck for deep person re-identification. *IEEE Transactions on Multimedia*, 22(10):2597–2609, 2019. 6
- [28] Niluthpol Chowdhury Mithun, Juncheng Li, Florian Metze, and Amit K Roy-Chowdhury. Learning joint embedding with multimodal cues for cross-modal video-text retrieval. In *Proceedings of the 2018 ACM on International Conference* on Multimedia Retrieval, pages 19–27, 2018. 2
- [29] Niluthpol Chowdhury Mithun, Juncheng B Li, Florian Metze, Amit K Roy-Chowdhury, and Samarjit Das. Cmuucr-bosch@ trecvid 2017: Video to text retrieval. In *TRECVID*, 2017. 2
- [30] Tam Minh Nguyen, Quang Huu Pham, Linh Bao Doan, Hoang Viet Trinh, Viet-Anh Nguyen, and Viet-Hoang Phan. Contrastive learning for natural language-based vehicle retrieval. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4245– 4252, 2021. 2, 8
- [31] Tien-Phat Nguyen, Ba-Thinh Tran-Le, Xuan-Dang Thai, Tam V Nguyen, Minh N Do, and Minh-Triet Tran. Traffic video event retrieval via text query using vehicle appearance and motion attributes. In *Proceedings of the IEEE/CVF Con-*

ference on Computer Vision and Pattern Recognition, pages 4165–4172, 2021. 2, 8

- [32] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 464–479, 2018. 6
- [33] Eun-Ju Park, Hoyoung Kim, Seonghwan Jeong, Byungkon Kang, and YoungMin Kwon. Keyword-based vehicle retrieval. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4220– 4227, 2021. 2, 8
- [34] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pages 1532–1543, 2014. 2
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2, 6
- [36] Carmelo Scribano, Davide Sapienza, Giorgia Franchini, Micaela Verucchi, and Marko Bertogna. All you can embed: Natural language based vehicle retrieval with spatiotemporal transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4253–4262, 2021. 2, 8
- [37] Clint Sebastian, Raffaele Imbriaco, Panagiotis Meletis, Gijs Dubbelman, Egor Bondarev, et al. Tied: A cycle consistent encoder-decoder model for text-to-image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4138–4146, 2021. 2, 8
- [38] Ziruo Sun, Xinfang Liu, Xiaopeng Bi, Xiushan Nie, and Yilong Yin. Dun: Dual-path temporal matching network for natural language-based vehicle retrieval. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4061–4067, 2021. 2, 8
- [39] Zheng Tang, Milind Naphade, Ming-Yu Liu, Xiaodong Yang, Stan Birchfield, Shuo Wang, Ratnesh Kumar, David Anastasiu, and Jenq-Neng Hwang. Cityflow: A city-scale benchmark for multi-target multi-camera vehicle tracking and re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8797–8806, 2019. 1, 5
- [40] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. 2
- [41] Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Well-read students learn better: On the importance of pre-training compact models. arXiv preprint arXiv:1908.08962, 2019. 2
- [42] Aaron Van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. arXiv e-prints, pages arXiv–1807, 2018. 1, 2, 4, 6

- [43] 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. 2, 8
- [44] Peng Wu, Xiangteng He, Mingqian Tang, Yiliang Lv, and Jing Liu. Hanet: Hierarchical alignment networks for videotext retrieval. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3518–3527, 2021. 2
- [45] Jianwei Yang, Yonatan Bisk, and Jianfeng Gao. Taco: Token-aware cascade contrastive learning for video-text alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11562–11572, 2021.
 2
- [46] Zhedong Zheng, Liang Zheng, Michael Garrett, Yi Yang, Mingliang Xu, and Yi-Dong Shen. Dual-path convolutional image-text embeddings with instance loss. ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM), 16(2):1–23, 2020. 1