

Symmetric Network with Spatial Relationship Modeling for Natural Language-based Vehicle Retrieval

Chuyang Zhao* Haobo Chen* Wenyuan Zhang Junru Chen
Sipeng Zhang Yadong Li Boxun Li
MEGVII Technology

{zhaochuyang, chenhaobo, zhangwenyuan, chenjunru, zhangsipeng, liyadong, liboxun}@megvii.com

Abstract

Natural language (NL) based vehicle retrieval aims to search specific vehicle given text description. Different from the image-based vehicle retrieval, NL-based vehicle retrieval requires considering not only vehicle appearance, but also surrounding environment and temporal relations. In this paper, we propose a Symmetric Network with Spatial Relationship Modeling (SSM) method for NL-based vehicle retrieval. Specifically, we design a symmetric network to learn the unified cross-modal representations between text descriptions and vehicle images, where vehicle appearance details and vehicle trajectory global information are preserved. Besides, to make better use of location information, we propose a spatial relationship modeling methods to take surrounding environment and mutual relationship between vehicles into consideration. The qualitative and quantitative experiments verify the effectiveness of the proposed method. We achieve 43.92% MRR accuracy on the test set of the 6th AI City Challenge on natural language-based vehicle retrieval track, yielding the 4th place on the public leaderboard. The code will be available at https://github.com/hbchen121/AICITY2022_Track2_SSM.

1. Introduction

Natural Language (NL) based vehicle retrieval aims to find specific vehicles given text descriptions, which is a vital part of the intelligent transportation and urban surveillance systems. Existing vehicle retrieval systems are generally based on image-to-image matching, *i.e.*, the vehicle re-identification. It requires providing a query image to retrieve the same vehicle in the image gallery. However, in practical applications, the query image may not be available and we only know the rough description of the target vehicle. This leads us to an urgent need for an NL-based

vehicle retrieval system.

Compared with image queries, natural language text descriptions are easier to obtain and modify, resulting in higher flexibility. However, because text belongs to a different modality from the image, it brings great challenges to retrieval. Specifically, natural language-based vehicle retrieval is a cross-modal matching that aims to learn a transferable model between visual and language. DUN [32] proposes a dual-path network to align vehicle features in video and text embeddings supervised by the circle loss [29]. But they ignore motion information in matching, which is conducive to distinguishing similar vehicles with different motion states. Different from DUN, CLV [1] utilize a symmetric InfoNCE loss [38] to learn cross-modal representation like CLIP [28]. They generate a global motion image for each track, retaining not only vehicle appearance information. In addition, the subject descriptions in the text are enhanced, improving the impact of vehicle appearance information. However, their motion maps suffer from a loss of vehicle appearance information, which degrades model performance.

Inspired by these methods, we propose a Symmetric Network with Spatial Relationship Modeling (SSM) approach to learn the visual and linguistic representations for NL-based vehicle retrieval. The symmetric network can learn both the vehicle internal characteristic (*e.g.* vehicle type, color, and shape) and the vehicle external characteristic (*e.g.* motion state and surrounding environment) simultaneously. More concretely, we adopt one visual encoder and text encoder to extract the vehicle appearance feature and the vehicle text embedding, respectively, which is optimized by the symmetric InfoNCE loss and pair-wise Circle loss. Symmetrically, another visual encoder and text encoder are employed to learn the representations of vehicle external characteristics, respectively. The visual encoder can be the EfficientNet B2 [38] and the text encoder can be RoBERTa [21]. Both internal and external features are fused to learn robust features of the vehicle track. We also propose a hard text sample mining method to distinguish similar texts, where

*Equal Contribution

texts from other views are utilized. In addition, by analyzing the ranking visualization, we found that the model can hardly learn the visual feature of the surrounding environment and relationship of multiple vehicles. To tackle this problem, we design a spatial relationship modeling approach to enhance tracklet information, improving the performance of our method. We have achieved 43.92% MRR accuracy on the test set of the 6th AI City Challenge on natural language-based vehicle retrieval track, yielding the 4th place on the public leaderboard.

2. Related Work

2.1. Natural Language Based Video Retrieval

Natural language based video retrieval aims to find the corresponding video given natural language description. There is increasing interest in this area. [11, 12, 18, 27] Most of the existing methods are based on representation learning, which try to make the feature representation of the corresponding text and video similar. These methods usually use a language model such as LSTM [16] or BERT, etc. [8] to extract the language feature and use a visual feature backbone such as ResNet [14] or C3D [2], etc. to extract the visual feature. Zhang, et al. [39] use hierarchical sequence embedding (HSE) to embed sequential data of different modalities into hierarchically semantic spaces with correspondence information. Antoine, et al. [24] propose a Mixture-of-Embedding-Experts (MEE) model with ability to handle missing input modalities during training. Dong, et al. [9] proposing a dual deep encoding network that encodes videos and queries into powerful dense representations of their own. Our method not only utilizes representation learning to learn transferable visual and text feature, but also utilize spatial relationship modeling to capture the surrounding environment and mutual relationship in vehicle retrieval task.

2.2. Vehicle Re-identification

In the last decade, vehicle ReID has achieved considerable progress, especially for deep learning based methods [3, 13, 20, 41]. Different from the person ReID methods [31, 40], it is still a challenging task due to the similar appearance and viewpoint variations problems. To deal with the two problems, some approaches [13, 23, 37] focus on subtle details to learn discriminative local features. For instance, OIFE [37] proposes 20 keypoints and raises an orientation invariant feature embedding module to emphasize regions with discriminative information. Part-regularized [13] trains a detector to focus on local regions and acquires local features to distinguish similar vehicles. VANet [5] adopts a view predictor and a modified triplet loss to generate viewpoint-aware features. In addition, many loss functions are applied to address the above challenges, including




Figure 1. Comparison of motion maps. We use cropped car image files to generate IOU-filtered motion maps that include more vehicle information.

representation learning loss functions and metric learning loss functions. representation learning losses [7, 19, 34–36] acquire feature representations by classification. Unlike them, metric learning losses [4, 15, 29] directly optimize the similarity score of image pairs and are termed as pair-wise loss, where margins are generally added between positive pairs and negative pairs to increase the distance. In this paper, we learn the unified cross-modal representation through the pair-wise losses.

3. Method

3.1. Overview

The pipeline of our method is illustrated in Fig. 3. The symmetric network consists of four branches, where the upper two branches are mainly used to acquire local visual and language representations of vehicle appearance, and the lower two branches are employed to learn the global information of the vehicle, including the motion state and the environment in which it is located. The two kinds of representations are fused together to generate the comprehensive feature representations of the vehicle track. We apply the Symmetric InfoNCE Loss [25] and Circle Loss [30] to connect the representations of the two modalities of the text and the image, ensuring that they are projected into a unified representation space. To further emphasize the motion and environment information, we construct triples to distinguish text descriptions from different views of the same vehicle. In addition, a spatial relationship modeling module is designed to use location information to constrain retrieval results in the retrieval process.

3.2. Data Augmentation

3.2.1 Image Augmentation

Similar to [1], we augment the motion image by pasting cropped vehicle images from the tracking frames to the background image. The background image is generated by computing the mean value of all frames from the same cam-

Original:	A white sedan turns left at an intersection.
Subject Aug:	A white sedan. A white sedan turns left at an intersection.
Motion Aug:	A white sedan turns left at an intersection. turns left.
Location Aug:	A white sedan turns left at an intersection. an intersection.

Figure 2. Examples of augmentation in text descriptions.

era, which can be formulated as:

$$B_c = \frac{1}{N_c} \sum_i^{N_c} F_i$$

where B_c is the background image of the c_{th} camera, N_c is the number of frames taken by the c_{th} camera, and F_i is the i_{th} frame.

The motion image is generated by pasting cropped vehicle images from the same track on the corresponding background image. As shown in the left of Fig. 1, too close consecutive frames may occlude adjacent frames on the motion image. To avoid this problem, we compute the Intersection Over Union (IOU) of adjacent frames and ignore frames whose IOU is larger than the threshold with already pasted frames. The threshold is set to 0.05 in our method.

3.2.2 Text Augmentation

Most text descriptions contain vehicle appearance information, motion information and location information. To make the model learn a specific information better, we adopt subject augmentation, motion augmentation and position augmentation respectively to emphasize different textual information, in which the respective text descriptions are repeated as shown in Fig. 2.

Subject augmentation. To enhance the appearance information, we need to repeat description about the vehicle appearance. Because the vehicle appearance description usually appears as the first noun phase of the sentence, we "spacy"¹ to extract all noun phrases and put the first one to the beginning of the sentence.

Motion augmentation. To enhance the motion information, we extract keywords which describe motion information and then repeat them. For motion information, We assume that there are only three motion types, *i.e.* turn left, turn right and go straight. If keywords "turn left" or "turn right" exist in any text description of one track's text description, we denote the motion type of this track is turning left or right. If none of these two keywords exist in the text

¹<https://spacy.io/>

description, we denote the motion type of this track is going straight. We append the motion type keyword "left", "right" or "straight" to the beginning of all text description of the track to emphasize the motion information.

Location augmentation. For location enhancement, We only distinguish whether a location described by a track is an intersection. So we look into the text descriptions of the track and check whether "intersection" exists in any of the text description to determine whether the described location is an intersection. We append "intersection" to the beginning of all text description of the track if the track is located in the intersection.

3.3. Visual and Linguistic Representation Learning

3.3.1 Symmetric Network

The task of natural language based vehicle retrieval refers to measuring the semantic similarity between a vehicle video and a language description. Generally speaking, humans pay more attention to unique information when describing vehicle videos, such as vehicle appearance, shape and other characteristics (termed as internal characteristics), as well as external characteristics such as motion state and surrounding environment (termed as external characteristics). We construct a dual-stream network to focus the internal and external characteristics, respectively.

Vehicle internal characteristic learning. As the upper two branches shown in Fig. 3, to enhance the internal characteristic in vehicle images, we crop the local vehicle image from the frame of the video and feed them into the image encoders. The encoders are EfficientNet B2 [38] or ibn-ResNet101-a [26] pretrained on ImageNet [6]. In addition, we employ a pretrained RoBERTa [21] as text encoder to learn the text embedding of vehicles. Its input text is augmented by subject emphasis to acquire vehicle embedding with more appearance and shape information. Following [1], we adopt projection heads to map visual feature and text embedding into a unified representation space, which is formulated as:

$$f_i = g_i(h_i) = W_2 \sigma(BN(W_1 h_i)), \quad (1)$$

$$f_t = g_t(h_t) = W_2 \sigma(LN(W_1 h_t)), \quad (2)$$

where h_i is the visual feature extracted by the image encoder and h_t is the linguistic embedding extracted by the image encoder. BN denotes the Batch Normalization (BN) layer. LN is a Layer Normalization (LN) layer. σ is a ReLU activation layer.

Vehicle external characteristic learning. Similar to the network structure in internal features learning, we adopt a

sine similarity is calculated by:

$$\cos(f_i, f_j) = \frac{f_i \cdot f_j}{\|f_i\| \cdot \|f_j\|}. \quad (5)$$

Then, the symmetric InfoNCE loss is as follows:

$$\mathcal{L}_{INCE} = \mathcal{L}_{i2t} + \mathcal{L}_{t2i}. \quad (6)$$

Different from the symmetric InfoNCE loss, Circle loss minimize the similarity of all negative pairs. Denote the positive pair and negative pairs as s_p and s_n , respectively. Circle loss is defined as follows:

$$\mathcal{L}_{circle} = \log \left[1 + \sum_{j=1}^L \exp(\gamma \alpha_n^j (s_n^j - \Delta_n)) \sum_{i=1}^K \exp(-\gamma \alpha_p^i (s_p^i - \Delta_p)) \right] \quad (7)$$

where the $K = 1$ and $L = 2(N - 1)$, Δ_p and Δ_n are the intra-class and inter-class margins, respectively. α_p and α_n are calculated as:

$$\begin{cases} \alpha_p^i = [O_p - s_p^i]_+ \\ \alpha_n^j = [s_n^j - O_n]_+ \end{cases}, \quad (8)$$

in which O_p and O_n are optimums of s_p and s_n , respectively. These parameters are set $O_p = 1 + m$, $\Delta_p = 1 - m$, $O_n = -m$ and $\Delta_n = m$, respectively.

The symmetric InfoNCE loss and Circle loss are enforced on all of three kinds of representations, which preserves information of different levels can be learned.

3.3.3 Hard Text Samples Mining

Although the non-appearance information is enhanced in external feature learning, our model still learns a lot of vehicle internal information and makes mistakes when discriminating difficult samples. For instance, when the same car appears in two different videos, the text descriptions of the two videos will have similar subjects as show in Fig. 3, which will degrade the retrieval performance.

To address this problem, we implement hard sample mining by composing triples with different view descriptions of the same vehicle and the current view motion image. More concretely, given N global motion images, N text descriptions and N text descriptions from other views, we have N triplets $\langle f_i^{img}, f_i^{text}, f_i^{text-v} \rangle$, where the anchor is f_i^{img} , the positive feature is f_i^{text} and the negative feature is f_i^{text-v} . Then we minimize the Triplet loss [17] to push away the f_i^{text-v} from f_i^{img} :

$$\mathcal{L}_{triple} = \frac{1}{N} \sum_{i=1}^N \left[\cos(f_i^{img}, f_i^{text-v}) - \cos(f_i^{img}, f_i^{text}) + m \right]_+ \quad (9)$$

in which m is the margin.

3.4. Spatial Relationship Modeling

3.4.1 Long-distance Relationship Modeling

Long-distance relationship modeling means that we build positional relationships between text and images, *e.g.*, intersection prediction. Then the text and images whose positions are both intersections are given greater similarity. Specifically, we distinguish whether a location is intersection. To increase the similarity of the text and image feature if they all describe the intersection, We calculate the location similarity between the text and the image and then add the similarity score to the final similarity matrix.

We found that the pictures in the dataset are all taken by cameras in fixed locations, so we can determine whether the location of an image is at the intersection by checking the location of its corresponding camera. We found if a camera is located in the intersection, some vehicles captured by this camera will stop for a while to wait the traffic light. We can calculate the location of the vehicle in each frame using the bounding box:

$$(x_i, y_i) = (left_i + \frac{1}{2}width_i, top_i + \frac{1}{2}height_i) \quad (10)$$

where $left_i$ and top_i are the left-top corner coordinate of the i_{th} frame's bounding box, $width_i$ and $height_i$ are the width and height of the i_{th} frame's bounding box.

We can infer the movement of the vehicle through the coordinates in each frames. If movement distance of the vehicle in consecutive n frames are all zero, we consider the car to be at the intersection and we label the corresponding camera to be located at the intersection. We can get the visual location vector by:

$$loc_v = \begin{cases} (0, 1)^T, & c_v \in \{c|c \text{ is at intersection}\} \\ (1, 0)^T, & \text{otherwise} \end{cases}, \quad (11)$$

where c_v is the camera of track v , $\{c|c \text{ is at intersection}\}$ is the collection of cameras which are located in the intersection.

We can determine whether a sentence is describing an intersection by checking whether keyword "intersection" exists in the sentence. We can get the text location vector by:

$$loc_s = \begin{cases} (0, 1)^T, & \text{"intersection"} \subseteq s \\ (1, 0)^T, & \text{otherwise} \end{cases}, \quad (12)$$

After we get the visual location embedding and text location embedding, we can calculate the similarity of them by the dot product. Assume we have n text queries and m visual tracks, we can get a matrix $S_l \in \mathbb{R}^{n \times m}$ representing the location similarity between each query and track.

3.4.2 Short-distance Relationship Modeling

In the queries, there are lots of sentences describe the relationship of more than one vehicles. However, Our model can not learn the relationship between multiple vehicles, which are very close in the video frame. To utilize this information, we perform relationship augmentation to make the proper text-visual pair more similar, termed short-distance relation modeling.

More concretely, we use ‘spacy’ to extract all noun-phases in the sentence and keep all phases describing vehicles. Then we use the vehicle description branch to extract the language embeddings of all the vehicle descriptions. We then use the provided detection file to extract the bounding boxes of all cars in the frame. We randomly select several frames in a track and then extract all detected cars in these frames. We use the cropped visual branch to extract the visual embedding for all cars in these frames. If a text query q describes relationship between vehicles v_1 and v_2 , we can calculate the similarity between v_2 and all detected vehicles in track t and take the maximum value as the relationship similarity between q and t . Assume we have n text queries and m visual tracks, we can get a matrix $S_r \in \mathbb{R}^{n \times m}$ representing the relationship similarity between each query and track.

The final similarity matrix S_{final} is formulated as the sum of the feature similarity matrix S , location similarity matrix S_l and relationship similarity matrix S_r :

$$S_{final} = S + \alpha S_l + \beta S_r \quad (13)$$

α and β are hyper-parameters. We set $\alpha = 1$ and $\beta = 0.2$ in our experiments.

4. Experiments

4.1. Datasets and Evaluation Settings.

Datasets. CityFlow-NL [10] dataset consist of 666 vehicles contains 3028 vehicle tracks collected from 40 cameras, of which 2155 vehicle tracks were used for training. Each track was annotated with three natural language descriptions. In addition, multiple natural language descriptions from other views of the vehicle in this track are collected in this track. The remaining data were adopted to evaluate the proposed method.

Evaluation. The vehicle retrieval by NL descriptions task generally utilize the Mean Reciprocal Rank (MRR) as the standard metrics, which is also used in [33]. It is formulated as:

$$MRR = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{rank_i}, \quad (14)$$

where $|Q|$ is the number of set of text descriptions. $rank_i$ refers to the rank position of the right track for the i_{th} text

Table 1. Ablation Study Results of our method.

Method	MRR
Baseline	0.3252
Baseline + Motion Aug	0.3362
Baseline + Location Aug	0.3305
Baseline + Other views	0.3405
Baseline + Circle loss	0.3548
Ensemble	0.4207
Ensemble + Spatial Modeling	0.4392

Table 2. Competition results of AI City Natural Language-Based Vehicle Retrieval Challenge.

Rank	Team Name	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	0.338
7	Terminus-AI	0.3320
8	MARS_WHU	0.3205
9	BUPT_MCPRL_T2	0.3012
10	folklore	0.2832

description.

4.2. Implement details

Both input local and global images are resized to 228×228 . We do not adopt any image augmentation methods to augment the data, including random horizontal flip and random clips, which can keep sports information such as motion unchanged. We set the batch size to 64 for each input, *i.e.*, vehicle local images, vehicle emphasis texts, global motion images, motion emphasis texts and text from other views. The total batch size is 320. Two text encoder is frozen in the training process. Our model is trained for 400 epochs employing the AdamW [22] optimizer with weight decay ($1e-2$). The learning rate starts from 0.01 and works as a warm-up strategy for 40 epochs. A step delay scheduler is adopted to decay the learning rate every 80 epochs. The γ and m in circle loss are set to 48 and 0.35, respectively. The m in triplet loss is set to 0. We train the model with different configures and ensemble them in final similarity calculating.

4.3. Performance on CityFlow-NL

We conduct ablation studies with different configures of our method. The results are illustrated in Table 1. The ‘‘Baseline’’ denotes the symmetric network with subject augmentation and symmetric InfoNCE loss. The ‘‘other view’’ refers to the hard samples mining approach. It can

be observed that each part learns different characteristics and brings a slight performance boost, respectively. In addition, the model ensemble integrates different information and achieves a great improvement.

We present the performances of top-10 team in the public leaderboard, as shown in Table 2. Our team (MegVideo) achieve a MRR score of 0.4392, taking 4-th highest among 15 total submissions on the AI City Challenge 2022 Track 2.

5. Conclusion

In this paper, we propose a Symmetric Network with Spatial Relationship Modeling method for NL-based vehicle retrieval. Firstly, we design a symmetric network to learn visual and linguistic representations. Then a spatial relationship modeling methods is proposed to make better use of location information. The qualitative experiments confirm the effectiveness of our method. It achieves 43.92% MRR accuracy on the test set of the 6th AI City Challenge on natural language-based vehicle retrieval track, yielding the 4th place on the public leaderboard.

Broader Impact. Our research can promote the practical application of vehicle retrieval technology, but the data used in the research may cause privacy violations. Therefore, we only use the data of the public dataset and mask the license plate to avoid increasing privacy leaks and other security issues.

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. 1, 2, 3
- [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
- [3] Junru Chen, Shiqing Geng, Yongluan Yan, Danyang Huang, Hao Liu, and Yadong Li. Vehicle re-identification method based on vehicle attribute and mutual exclusion between cameras. *arXiv preprint arXiv:2104.14882*, 2021. 2
- [4] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, pages 539–546, 2005. 2
- [5] Ruihang Chu, Yifan Sun, Yadong Li, Zheng Liu, Chi Zhang, and Yichen Wei. Vehicle re-identification with viewpoint-aware metric learning. In *ICCV*, pages 8281–8290, 2019. 2
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 3
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, pages 4690–4699, 2019. 2
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [9] Jianfeng Dong, Xirong Li, Chaoxi Xu, Shouling Ji, Yuan He, Gang Yang, and Xun Wang. Dual encoding for zero-example video retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9346–9355, 2019. 2
- [10] Qi Feng, Vitaly Ablavsky, and Stan Sclaroff. Cityflow-nl: Tracking and retrieval of vehicles at city scale by natural language descriptions. 2021. 6
- [11] Valentin Gabeur, Arsha Nagrani, Chen Sun, Kartek Alahari, and Cordelia Schmid. Masking modalities for cross-modal video retrieval. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1766–1775, 2022. 2
- [12] Ning Han, Jingjing Chen, Guangyi Xiao, Hao Zhang, Yawen Zeng, and Hao Chen. Fine-grained cross-modal alignment network for text-video retrieval. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3826–3834, 2021. 2
- [13] Bing He, Jia Li, Yifan Zhao, and Yonghong Tian. Part-regularized near-duplicate vehicle re-identification. In *CVPR*, pages 3997–4005, 2019. 2
- [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
- [15] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. 2
- [16] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [17] Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *International workshop on similarity-based pattern recognition*, pages 84–92. Springer, 2015. 5
- [18] Yehao Li, Yingwei Pan, Jingwen Chen, Ting Yao, and Tao Mei. X-modaler: A versatile and high-performance codebase for cross-modal analytics. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3799–3802, 2021. 2
- [19] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *ICML*, volume 48, pages 507–516, 2016. 2
- [20] Xinchun Liu, Wu Liu, Huadong Ma, and Huiyuan Fu. Large-scale vehicle re-identification in urban surveillance videos. In *ICME*, pages 1–6, 2016. 2
- [21] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 1, 3

- [22] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [23] Dechao Meng, Liang Li, Xuejing Liu, Yadong Li, Shijie Yang, Zheng-Jun Zha, Xingyu Gao, Shuhui Wang, and Qingming Huang. Parsing-based view-aware embedding network for vehicle re-identification. In *CVPR*, pages 7101–7110, 2020. 2
- [24] Antoine Miech, Ivan Laptev, and Josef Sivic. Learning a text-video embedding from incomplete and heterogeneous data. *arXiv preprint arXiv:1804.02516*, 2018. 2
- [25] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2
- [26] 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. 3
- [27] Mengshi Qi, Jie Qin, Yi Yang, Yunhong Wang, and Jiebo Luo. Semantics-aware spatial-temporal binaries for cross-modal video retrieval. *IEEE Transactions on Image Processing*, 30:2989–3004, 2021. 2
- [28] 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. *arXiv preprint arXiv:2103.00020*, 2021. 1, 4
- [29] Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. Circle loss: A unified perspective of pair similarity optimization. In *CVPR*, pages 6397–6406, 2020. 1, 2
- [30] Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6398–6407, 2020. 2, 4
- [31] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *Proceedings of the European conference on computer vision (ECCV)*, pages 480–496, 2018. 2
- [32] Ziruo Sun, Xinfang Liu, Xiaopeng Bi, Xiushan Nie, and Yilong Yin. Dun: Dual-path temporal matching network for natural language-based vehicle retrieval. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 4056–4062, 2021. 1
- [33] Ellen M Voorhees et al. The trec-8 question answering track report. In *Trec*, volume 99, pages 77–82. Citeseer, 1999. 6
- [34] Feng Wang, Jian Cheng, Weiyang Liu, and Haijun Liu. Additive margin softmax for face verification. *SPL*, 25(7):926–930, 2018. 2
- [35] Feng Wang, Xiang Xiang, Jian Cheng, and Alan Loddon Yuille. Normface: L_2 hypersphere embedding for face verification. In *ACM MM*, pages 1041–1049, 2017. 2
- [36] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *CVPR*, pages 5265–5274, 2018. 2
- [37] Zhongdao Wang, Luming Tang, Xihui Liu, Zhuliang Yao, Shuai Yi, Jing Shao, Junjie Yan, Shengjin Wang, Hongsheng Li, and Xiaogang Wang. Orientation invariant feature embedding and spatial temporal regularization for vehicle re-identification. In *ICCV*, pages 379–387, 2017. 2
- [38] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, pages 305–321, 2018. 1, 3, 4
- [39] Bowen Zhang, Hexiang Hu, and Fei Sha. Cross-modal and hierarchical modeling of video and text. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 374–390, 2018. 2
- [40] Yanwei Zheng, Zengrui Zhao, Xiaowei Yu, and Dongxiao Yu. Template-aware transformer for person reidentification. *Computational Intelligence and Neuroscience*, 2022, 2022. 2
- [41] Chaoran Zhuge, Yujie Peng, Yadong Li, Jiangbo Ai, and Junru Chen. Attribute-guided feature extraction and augmentation robust learning for vehicle re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 618–619, 2020. 2