



HAL
open science

Review of Recent Deep Learning Based Methods for Image-Text Retrieval

Jianan Chen, Lu Zhang, Cong Bai, Kidiyo Kpalma

► **To cite this version:**

Jianan Chen, Lu Zhang, Cong Bai, Kidiyo Kpalma. Review of Recent Deep Learning Based Methods for Image-Text Retrieval. IEEE 3rd International Conference on Multimedia Information Processing and Retrieval, Aug 2020, Shenzhen, China. 10.1109/MIPR49039.2020.00042 . hal-02480975

HAL Id: hal-02480975

<https://hal.science/hal-02480975>

Submitted on 3 Sep 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Review of Recent Deep Learning Based Methods for Image-Text Retrieval

Jianan Chen¹, Lu Zhang¹, Cong Bai², Kidiyo Kpalma¹

Univ Rennes, INSA Rennes, CNRS, IETR - UMR 6164, F-35000 Rennes, France¹
College of Computer Science, Zhejiang University of Technology, Hangzhou, China²
{Jianan.Chen, Lu.Ge, Kidiyo.Kpalma}@insa-rennes.fr¹, congbai@zjut.edu.cn²

Abstract

Cross-modal retrieval has drawn much attention in recent years due to the diversity and the quantity of information data that exploded with the popularity of mobile devices and social media. Extracting relevant information efficiently from large-scale multi-modal data is becoming a crucial problem of information retrieval. Cross-modal retrieval aims to retrieve relevant information across different modalities. In this paper, we highlight key points of recent cross-modal retrieval approaches based on deep-learning, especially in the image-text retrieval context, and classify them into four categories according to different embedding methods. Evaluations of state-of-the-art cross-modal retrieval methods on two benchmark datasets are shown at the end of this paper.

1. Introduction

Over the last decade, cross-modal retrieval has made significant progress. The goal of cross-modal retrieval is to retrieve relevant information across heterogeneous modalities. It is widely used in many fields, such as visual question and answering [28], image or video caption [2, 6], phrase localization [38], knowledge transfer [10] and text-to-image generation [43, 31, 17]. Benefiting from Content-Based Image Retrieval (CBIR) and Natural Language Processing (NLP) techniques, the computer could almost bridge the semantic gap between high-level human perception and low-level features in single mode. As deep learning achieves remarkable results in both vision and language domain, researchers begin to explore the semantic gap between image and text. In bidirectional image-text cross-modal retrieval, taking image as the query to retrieve relevant information in text data is called image-to-text retrieval, and vice versa.

The first cross-modal retrieval review is written by Liu *et al.* [22], which focus on summarizing traditional methods. The overviews of multimedia information retrieval in [1] and [29] are not only for image-text but also for

video and audio modalities. The most recent cross-modal overview [27] focuses on music and sound data retrieval. In this paper, we focus on cross-modal retrieval methods based on deep-learning, only for image-text context, and proposed in the last two years as some new methods based on deep learning have been proposed, which significantly improve the performance. We give an analysis of this relatively narrow topic in the image-text cross-modal retrieval domain and propose to classify these algorithms into four categories according to their embedding methods: 1) pairwise learning embedding methods; 2) adversarial learning methods; 3) interaction learning methods; 4) attributes learning methods.

The rest of the paper is organized as follows: we classify the most recent image-text cross-modal retrieval algorithms by their embedding methods and highlight their pros and cons in Section 2; we show the performance comparison results of the representative algorithms in each category using two most popular datasets in this domain (Flickr30K and MSCOCO) in Section 3; then the paper concludes with the recent image-text retrieval works and gives some perspectives in Section 4.

2. Image-Text Retrieval Methods Classification

In general, the cross-modal retrieval architecture could be divided into three parts: “Image Features Branch”, “Text Features Branch” and “Latent Space”, as shown in Fig. 1. The first two branches extract image features and text features, separately. Long Short-Term Memory (LSTM [9]) or Recurrent Neural Network (RNN [4]) is used to extract stylistic features from texts, while Convolutional Neural Network (CNN [35]) is used to extract image features. Then the “Latent Space” part projects the features corresponding to different modalities to one common space, and measure the similarity between the projected text features and the projected image features. Here, the methods following the above architecture are classified as the *pairwise learning*, cf. the blue regions in Fig. 1. If adversarial machine learning methods are adopted in the “Latent Space” part, we classify these methods into the *adversarial learning*, cf. yellow

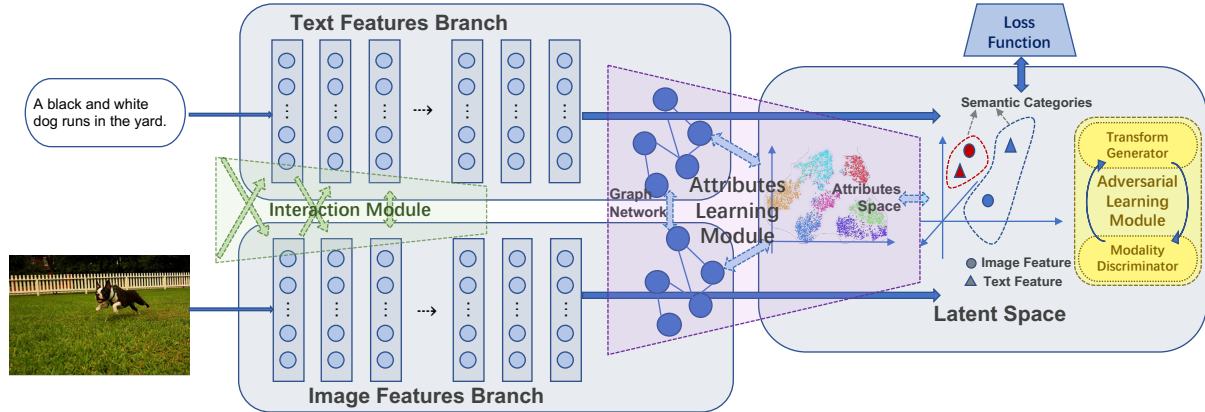


Figure 1: General architecture of Image-Text cross-modal retrieval methods. The blue regions are the basic three-part structure and denotes the *pairwise learning* methods. With an additional dotted green region, it is the structure of the *interaction learning* methods. With the dotted yellow square in “Latent Space”, it indicates the *adversarial learning* methods. With an additional dotted violet region, it becomes the structure of the *attributes learning* methods.

region in Fig. 1. If there are some interaction flows between the “Image Features Branch” and the “Text Features Branch” in addition to the general architecture, we classify these methods into the *interaction learning*, cf. green region in Fig. 1. If high-level semantic attributes are exploited, instead of the direct use of the basic image features and text features, we classify these methods into the *attributes learning*, cf. violet region in Fig. 1. In the following, we detail these four types of methods.

2.1 Pairwise Learning Methods

Pairwise learning methods attempt to find a cross-modal loss function that can calculate the distance between corresponding feature pairs directly in a common space. By learning this loss function, the distance between associated images and texts reduces, and the distance between independent samples increases. There are some different forms of pairwise learning, but all of them represent two different features in the same common space directly. Pairwise learning methods differ in the factors of the loss function, such as corresponding label relation, feature distance space, similarity measure evaluation.

Zhang and Lu introduce a new matching loss called Cross-Modal Projection Matching (CMPM [44]). The idea behind this is to increase the correlation of matching pairs and to reduce it for unmatching pairs by minimizing the Kullback-Leibler divergence between the probability of matching image features to text features and the normalized true matching probability. All the positive and negative samples are thus considered in the CMPM. The disadvantage may lie in the lack of intrinsic association of words in the text since bidirectional LSTM only integrally words sequence information without sufficient se-

semantic context logic. After CMPM, Deep Pairwise Ranking model with multi-label information for Cross-Modal retrieval (DPRCM [14]) is proposed, which employs a bi-triplet loss to reduce the distance between positive samples in the same identity and increase the distance between negative independent samples. DPRCM also combines cross-entropy loss with bi-triplet loss in their retrieval network so that multi-label information can be learned in common space under supervision. DPRCM extracts image features and text features only by two-layer neural networks separately, which is a more straightforward way than other cross-modal retrieval methods. Unlike DPRCM, Deep Supervised Cross-Modal Retrieval (DSCMR [45]) uses fully connected layers to build common representation space. A linear classifier is used to predict the category of each sample in the common representation space. Simultaneously, another discrimination loss is minimized in label space. Both DPRCM and DSCMR belong to supervised learning methods. They use label information to enhance the learning progress when they deal with multi-modal pairs. Finally, Liu *et al.* propose neighbor-aware network (NAN [20]), which calculates the neighbor-aware ranking loss in common semantic space under the influence of the intra-attention module. The neighbor-aware ranking loss can be divide into inter-modal and intra-modal parts. Inter-modal neighbor-aware ranking loss focuses on semantic relation inside a homogeneous modality while intra-modal neighbor-aware ranking loss points on semantic relation between heterogeneous modalities. NAN adds an attention module to re-weight feature map since different semantics are distinguished in intra-modal and inter-modal neighbor-aware networks. As attention features map could be associated with the semantic relation during the neighbor-aware

ranking loss learning, the intra-attention module plays an important role in image-text matching.

Unlike CMPM, DPRCM and DSCMR rely more heavily upon label distance information. There are some other pairwise loss functions belong to supervised learning, such as kNN-margin loss [21], hard negatives loss [5]. The key point of pairwise learning is to design an efficient loss function that could reduce features distance of the same semantic category in common space.

2.2 Adversarial Learning Methods

Adversarial learning methods are enlightened by Generative Adversarial Nets (GAN [7]). Wang *et al.* [36] introduced adversarial learning firstly into cross-modal retrieval domain. In latent space, a two-player minimax value game has been played between a discriminator and a generator in adversarial network learning. The expectation value $V_{D,G}$ is defined as:

$$V_{D,G} = \mathbb{E}_{I_i \sim I} [\log D((I_i))] + \mathbb{E}_{T_i \sim T} [\log (1 - D(G(T_i)))] \quad (1)$$

where I and T indicate the image and text modalities. Adversarial learning method uses minimax game to bridge image and text features that played between generator and discriminator. That is a bright method for cross-modal retrieval.

After that, Sarafianos *et al.* propose Text-Image Modality Adversarial Matching (TIMAM [34]), which adopts an Adversarial Representation Learning (ARL) framework to learn modality-invariant representations for more effective image-text matching. In the ARL framework, a two-layer fully-connected network adversarial discriminator is optimized in the common space. The better discriminator pain, the better cross-modal retrieval gain. TIMAM also adds Bidirectional Encoder Representations from Transformers (BERT [3]) in front of LSTM branch to optimize text features. At the same time, Liu *et al.* propose a new deep adversarial graph attention convolution network (A-GANet [23]). A-GANet extracts image features not only from the CNN branch but also from a graph attention convolution layer based on a visual scene graph. The visual scene graph carries information about object regions and relationships according to human visual perception characteristics. High-level structured semantic visual features are learned from this designed graph attention convolution layers. Particular joint embedding layers connect the image and text features through the adversarial learning module. Furthermore, Wang *et al.* [37] and Zhu *et al.* [46] use adversarial learning in food images and recipes matching.

Adversarial learning methods have not been around for a long time in the field of cross-modal retrieval. It has also been used in other areas such as image synthesis and style transfer that require more inference.

2.3 Interaction Learning Methods

In this section, we define interaction learning methods as those having a large amount of information transfer between the two branches before the image and text features enter the common space.

Lou *et al.* propose a Multitask learning approach for Cross-modal Image-Text Retrieval (MCITR [24]) to take into account the common features extracted from image-text cross-modal data. MCITR employs relation-enhanced correspondence cross-modal autoencoder [8] to correlate the hidden representations, before text and image are projected into an embedding space. Simultaneously, Cross-Modal Adaptive Message Passing (CAMP [41]) adopts a cross-modal message-passing aggregation at the beginning of the network. CAMP explores the interactions between images and text before calculation in common space. Other methods add attention module to transfer the information between image and text branch, such as Wang *et al.* [40] and Wu *et al.* [42].

Due to the information transfer between image and text branches in the initial and low-level processing, more corresponding connections could be represented in latent space. Nevertheless, this kind of algorithms is more complicated, and the amount of calculation increases exponentially.

2.4 Attributes Learning Methods

In deep learning, a vast number of parameters are trained by large-scale calculations to obtain excellent results, which means that a massive amount of data is needed for training deep neural networks. However, human beings can learn the properties of things from a few examples. Attributes learning imitates human thought processes and learns the characteristics of objects. The essence of attributes learning is “learning to learn”. There are also some attempts to apply attributes learning in cross-modal retrieval.

Ji *et al.* propose an Attribute-Guided Network (Ag-Net [13]) for cross-modal retrieval, which combines with zero-shot learning and hashing retrieval. Objective functions are designed to transform image and text feature vectors into object attribute vectors in attributes space. Then a three-layer neural network transforms attribute vectors to hash codes. Without the supervision information, instances clusters themselves in attribute space. Hamming distance is selected to calculate the similarity between different modalities. Although hash coding is an efficient representation, we cannot determine whether there is a linear relationship between the hash code length and the number of attributes. From then on, Huang and Wang propose Aligned Cross-Modal Memory (ACMM [12]) for few-shot image and sentence matching. ACMM includes two key steps: aligned memory controller network and mem-

| Method | | Flickr30K | | | | | | MSCOCO (1K test images) | | | | | | | |
|----------------------|-----------------|---------------|-------------|-------------|---------------|-------------|-------------|-------------------------|---------------|-------------|-------------|---------------|-------------|-------------|-------------|
| | | Text-to-Image | | | Image-to-Text | | | mR | Text-to-Image | | | Image-to-Text | | | mR |
| | | R@1 | R@5 | R@10 | R@1 | R@5 | R@10 | | R@1 | R@5 | R@10 | R@1 | R@5 | R@10 | |
| Pairwise Learning | kNN-margin [21] | 36.0 | 64.4 | 72.5 | 26.7 | 54.3 | 65.7 | 53.3 | 65.4 | 91.9 | 97.1 | 49.6 | 82.7 | 91.2 | 79.7 |
| | CMPM [44] | 48.3 | 75.6 | 84.5 | 35.7 | 63.6 | 74.1 | 63.6 | 56.1 | 86.3 | 92.9 | 44.6 | 78.7 | 89.0 | 74.6 |
| | CMPM+CMPC [44] | 49.6 | 76.8 | 86.1 | 37.3 | 65.7 | 75.5 | 65.2 | 52.9 | 83.8 | 92.1 | 41.3 | 74.6 | 85.9 | 71.8 |
| | VSE++ [5] | 52.9 | 80.5 | 87.2 | 39.6 | 70.1 | 79.5 | 68.3 | 64.6 | 90.0 | 95.7 | 52.0 | 84.3 | 92.0 | 79.7 |
| | NAN [20] | 55.1 | 80.3 | 89.6 | 39.4 | 68.8 | 79.9 | 68.9 | 61.3 | 87.9 | 95.4 | 47.0 | 80.8 | 90.1 | 77.1 |
| Adversarial Learning | A-GANet [23] | - | - | - | 39.5 | 69.9 | 80.9 | - | - | - | - | - | - | - | - |
| | TIMAM [34] | 53.1 | 78.8 | 87.6 | 42.6 | 71.6 | 81.9 | 69.3 | - | - | - | - | - | - | - |
| Interaction Learning | MFM* [25] | 50.2 | 78.1 | 86.7 | 38.2 | 70.1 | 80.2 | 67.3 | 58.9 | 86.3 | 92.4 | 47.7 | 80.1 | 90.9 | 76.2 |
| | SCAN* [15] | 67.4 | 90.3 | 95.8 | 48.6 | 77.7 | 85.2 | 77.5 | 72.7 | 94.8 | 98.4 | 58.8 | 88.4 | 94.8 | 84.7 |
| | MTFN-RR [39] | 65.3 | 88.3 | 93.3 | 52.0 | 80.1 | 86.1 | 77.5 | 74.3 | 94.9 | 97.9 | 60.1 | 89.1 | 95.0 | 85.2 |
| | BFAN* [19] | 68.1 | 91.3 | - | 50.8 | 78.4 | - | - | 74.9 | 95.2 | - | 59.4 | 88.4 | - | - |
| | CAMP [41] | 68.1 | 89.7 | 95.2 | 51.5 | 77.1 | 85.3 | 77.8 | 72.3 | 94.8 | 98.3 | 58.5 | 87.9 | 95.0 | 84.5 |
| | PFAN* [40] | 70.0 | 91.8 | 95.0 | 50.4 | 78.7 | 86.1 | 78.7 | 76.5 | 96.3 | 99.0 | 61.6 | 89.6 | 95.2 | 86.4 |
| | SAEM [42] | 69.1 | 91.0 | 95.1 | 52.4 | 81.1 | 88.1 | 79.5 | 71.2 | 94.1 | 97.7 | 57.8 | 88.6 | 94.9 | 84.1 |
| Attributes Learning | GVSE* [11] | 68.5 | 90.9 | 95.5 | 50.6 | 79.8 | 87.6 | 78.8 | 72.2 | 94.1 | 98.1 | 60.5 | 89.4 | 95.8 | 85.0 |
| | ACMM [12] | 80.0 | 95.5 | 98.2 | 50.2 | 76.8 | 84.7 | 80.9 | 81.9 | 98.0 | 99.3 | 58.2 | 87.3 | 93.9 | 86.4 |
| | ACMM* [12] | 85.2 | 96.7 | 98.4 | 53.8 | 79.8 | 86.8 | 83.5 | 84.1 | 97.8 | 99.4 | 60.7 | 88.7 | 94.9 | 87.6 |

Table 1: The performance of state-of-the-art methods on Flickr30K and MSCOCO. Red, green, and blue represent the best, second, and third performance respectively. * indicates ensemble methods.

ory read and update. A cross-modal graph convolutional network based on aligned memory controller network aims to generate modality-specific interface vectors to connect with shared memory item. Memory read and update module aims to score the similarity between sentence and image semantically and to update the memory when few-shot content learning is used.

Attributes learning can learn the characteristics of objects from a few examples. Comparing with direct pairwise learning by the deep neural network, the attributes learning owns stronger cross-modal correlation features extraction. Thus attributes learning methods show the highest performances in Table 1.

3. Databases and Evaluation

There are many established databases in the cross-modal retrieval field, especially for image-text retrieval tasks. For example, CUHK-PEDES dataset [16] focuses on pedestrians on the road; Wikipedia dataset [32] has more text information which could mine NLP capabilities; Recipes1M dataset [26, 33] owns large-scale food images and recipes, etc. Since various image-text retrieval methods have reported their performances on the two most common database Flickr30K and MSCOCO, we sum them up here for comparison.

3.1 Databases

Flickr30K [30] is a standard dataset for image-text retrieval, containing more than 31K images and 155K sentences in total, each image has five corresponding sentences. Flickr30K has 44,518 categories in total. It is usually split

into 29K images for training, 1K images for validation and 1K images for test.

MSCOCO [18] contains 123,287 images and each one is described by five sentences. It has 91 objects categories. Generally experiments use 5K images for validation, 1K or 5K images for test. Table 1 shows performances on 5 folds of 1K test images as the most commonly setting.

3.2 Evaluation

We collect the state-of-the-art approaches $Recall@K$ results shown on papers, which measures the number of correct items are found among the top K retrieval results. For convenience, we also give a general evaluation indicator mR , which means the mean of $Recall@K$. For all the algorithms, we show the best results in the database. However, these performances may get from module ensemble method, we use * indicates that in Table 1. The best results of all retrieval methods are in red, second in green, and third in blue.

3.3 Discussion

The comparison of the results from these two databases can only show part of the performances of the algorithms. Comparing the results, we can see that some algorithms have better results in retrieving text from images, and others are better in retrieving images from text. If a method gets better results in image-to-text retrieval direction than text-to-image direction, it means that the image feature representations in latent space are more precise, and vice versa. But no algorithm can achieve the best results in both directions currently. In other words, no algorithm gets the best balance

point between two directions. Moreover, the same method performs differently on different databases, which may be caused by the number of object categories in the data. As we know, the number of object categories in Flickr30k is hundred times that of MSCOCO. More categories means higher learning costs, which is a challenge for retrieval algorithms. Some categories have a large sample size, but some have a small sample size. A large number of samples may cause overfitting, while a small number is not enough to train appropriate network parameters. For this reason, attention models and attributes learning methods are used to reduce the impact of small number of samples on retrieval results. Therefore, the interaction learning methods and attributes learning methods archive as high performance on the Flick30K as on the MSCOCO database.

4. Conclusion and Perspective

Deep learning based methods in image-text cross-modal retrieval have achieved significant progress in recent years. Pairwise learning proposes two-branch architecture; adversarial learning and interaction learning methods are based on it; attributes learning may become a popular trend for cross-modal retrieval tasks due to the usage of relations between attributes and semantics. In the future, more image-text free pairs databases for cross-modal retrieval should be explored. More evaluation metrics should be used to compare different methods.

5. Acknowledgement

This work is supported in part by the China Scholarship Council (CSC) No. 201801810074 and Natural Science Foundation of China under Grant No. U1908210 and 61976192.

References

[1] R. Aygun and W. Benesova. Multimedia retrieval that works. In *2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, pages 63–68, April 2018.

[2] A. F. Biten, L. Gomez, M. Rusinol, and D. Karatzas. Good news, everyone! context driven entity-aware captioning for news images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[3] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018.

[4] J. L. Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990.

[5] F. Faghri, D. J. Fleet, J. R. Kiros, and S. Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. 2018.

[6] Z. Fan, Z. Wei, S. Wang, and X.-J. Huang. Bridging by word: Image grounded vocabulary construction for visual captioning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6514–6524, 2019.

[7] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.

[8] D. R. Hardoon, S. Szedmak, and J. Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural computation*, 16(12):2639–2664, 2004.

[9] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[10] X. Huang and Y. Peng. Deep cross-media knowledge transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8837–8846, 2018.

[11] Y. Huang, Y. Long, and L. Wang. Few-shot image and sentence matching via gated visual-semantic embedding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8489–8496, 2019.

[12] Y. Huang and L. Wang. Acmm: Aligned cross-modal memory for few-shot image and sentence matching. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.

[13] Z. Ji, Y. Sun, Y. Yu, Y. Pang, and J. Han. Attribute-guided network for cross-modal zero-shot hashing. *IEEE transactions on neural networks and learning systems*, 2019.

[14] Y. Jian, J. Xiao, Y. Cao, A. Khan, and J. Zhu. Deep pairwise ranking with multi-label information for cross-modal retrieval. In *2019 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1810–1815, July 2019.

[15] K.-H. Lee, X. Chen, G. Hua, H. Hu, and X. He. Stacked cross attention for image-text matching. In *The European Conference on Computer Vision (ECCV)*, September 2018.

[16] S. Li, T. Xiao, H. Li, B. Zhou, D. Yue, and X. Wang. Person search with natural language description. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

[17] W. Li, P. Zhang, L. Zhang, Q. Huang, X. He, S. Lyu, and J. Gao. Object-driven text-to-image synthesis via adversarial training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12174–12182, 2019.

[18] T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014.

[19] C. Liu, Z. Mao, A.-A. Liu, T. Zhang, B. Wang, and Y. Zhang. Focus your attention: A bidirectional focal attention network for image-text matching. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, pages 3–11, New York, NY, USA, 2019. ACM.

[20] C. Liu, Z. Mao, W. Zang, and B. Wang. A neighbor-aware approach for image-text matching. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3970–3974, May 2019.

- [21] F. Liu and R. Ye. A strong and robust baseline for text-image matching. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 169–176, Florence, Italy, July 2019. Association for Computational Linguistics.
- [22] J. Liu, C. Xu, and H. Lu. Cross-media retrieval: State-of-the-art and open issues. *Int. J. of Multimedia Intelligence and Security*, 1:33 – 52, 01 2010.
- [23] J. Liu, Z.-J. Zha, R. Hong, M. Wang, and Y. Zhang. Deep adversarial graph attention convolution network for text-based person search. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, pages 665–673, New York, NY, USA, 2019. ACM.
- [24] J. Luo, Y. Shen, X. Ao, Z. Zhao, and M. Yang. Cross-modal image-text retrieval with multitask learning. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, CIKM '19, pages 2309–2312, New York, NY, USA, 2019. ACM.
- [25] L. Ma, W. Jiang, Z. Jie, Y. Jiang, and W. Liu. Matching image and sentence with multi-faceted representations. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1, 2019.
- [26] J. Marin, A. Biswas, F. Ofli, N. Hynes, A. Salvador, Y. Aytar, I. Weber, and A. Torralba. Recipe1m+: A dataset for learning cross-modal embeddings for cooking recipes and food images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
- [27] M. Mueller, A. Arzt, S. Balke, M. Dorfer, and G. Widmer. Cross-modal music retrieval and applications: An overview of key methodologies. *IEEE Signal Processing Magazine*, 36(1):52–62, Jan 2019.
- [28] H. Nam, J. Ha, and J. Kim. Dual attention networks for multimodal reasoning and matching. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2156–2164, July 2017.
- [29] Y. Peng, X. Huang, and Y. Zhao. An overview of cross-media retrieval: Concepts, methodologies, benchmarks, and challenges. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(9):2372–2385, Sep. 2018.
- [30] B. A. Plummer, L. Wang, C. M. Cervantes, J. C. Caicedo, J. Hockenmaier, and S. Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *CoRR*, abs/1505.04870, 2015.
- [31] T. Qiao, J. Zhang, D. Xu, and D. Tao. Mirrorgan: Learning text-to-image generation by redescription. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1505–1514, 2019.
- [32] N. Rasiwasia, J. Costa Pereira, E. Coviello, G. Doyle, G. R. Lanckriet, R. Levy, and N. Vasconcelos. A new approach to cross-modal multimedia retrieval. In *Proceedings of the 18th ACM International Conference on Multimedia*, MM '10, pages 251–260, New York, NY, USA, 2010. ACM.
- [33] A. Salvador, N. Hynes, Y. Aytar, J. Marin, F. Ofli, I. Weber, and A. Torralba. Learning cross-modal embeddings for cooking recipes and food images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [34] N. Sarafianos, X. Xu, and I. A. Kakadiaris. Adversarial representation learning for text-to-image matching. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5814–5824, 2019.
- [35] A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, and K. J. Lang. Phoneme recognition using time-delay neural networks. *IEEE transactions on acoustics, speech, and signal processing*, 37(3):328–339, 1989.
- [36] B. Wang, Y. Yang, X. Xu, A. Hanjalic, and H. T. Shen. Adversarial cross-modal retrieval. In *Proceedings of the 25th ACM International Conference on Multimedia*, MM '17, pages 154–162, New York, NY, USA, 2017. ACM.
- [37] H. Wang, D. Sahoo, C. Liu, E.-p. Lim, and S. C. H. Hoi. Learning cross-modal embeddings with adversarial networks for cooking recipes and food images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [38] L. Wang, Y. Li, J. Huang, and S. Lazebnik. Learning two-branch neural networks for image-text matching tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):394–407, Feb 2019.
- [39] T. Wang, X. Xu, Y. Yang, A. Hanjalic, H. T. Shen, and J. Song. Matching images and text with multi-modal tensor fusion and re-ranking. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, pages 12–20, New York, NY, USA, 2019. ACM.
- [40] Y. Wang, H. Yang, X. Qian, L. Ma, J. Lu, B. Li, and X. Fan. Position focused attention network for image-text matching. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, IJCAI'19, pages 3792–3798. AAAI Press, 2019.
- [41] Z. Wang, X. Liu, H. Li, L. Sheng, J. Yan, X. Wang, and J. Shao. Camp: Cross-modal adaptive message passing for text-image retrieval. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [42] Y. Wu, S. Wang, G. Song, and Q. Huang. Learning fragment self-attention embeddings for image-text matching. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, pages 2088–2096, New York, NY, USA, 2019. ACM.
- [43] T. Xu, P. Zhang, Q. Huang, H. Zhang, Z. Gan, X. Huang, and X. He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1316–1324, 2018.
- [44] Y. Zhang and H. Lu. Deep cross-modal projection learning for image-text matching. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [45] L. Zhen, P. Hu, X. Wang, and D. Peng. Deep supervised cross-modal retrieval. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [46] B. Zhu, C.-W. Ngo, J. Chen, and Y. Hao. R2gan: Cross-modal recipe retrieval with generative adversarial network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.