

Noisy Correspondence Learning with Meta Similarity Correction

Haochen Han¹, Kaiyao Miao², Qinghua Zheng¹, Minnan Luo^{1*}

¹ School of Computer Science and Technology, Xi'an Jiaotong University, China

² School of Cyber Science and Engineering, Xi'an Jiaotong University, China

{hhc1997, miaoky814}@stu.xjtu.edu.cn, {qhzheng, minnluo}@xjtu.edu.cn

Abstract

Despite the success of multimodal learning in cross-modal retrieval task, the remarkable progress relies on the correct correspondence among multimedia data. However, collecting such ideal data is expensive and time-consuming. In practice, most widely used datasets are harvested from the Internet and inevitably contain mismatched pairs. Training on such noisy correspondence datasets causes performance degradation because the cross-modal retrieval methods can wrongly enforce the mismatched data to be similar. To tackle this problem, we propose a Meta Similarity Correction Network (MSCN) to provide reliable similarity scores. We view a binary classification task as the meta-process that encourages the MSCN to learn discrimination from positive and negative meta-data. To further alleviate the influence of noise, we design an effective data purification strategy using meta-data as prior knowledge to remove the noisy samples. Extensive experiments are conducted to demonstrate the strengths of our method in both synthetic and real-world noises, including Flickr30K, MS-COCO, and Conceptual Captions. Our code is publicly available.¹

1. Introduction

Recently, cross-modal retrieval has drawn much attention with the rapid growth of multimedia data. Given a query sample of specific modality, cross-modal retrieval aims to retrieve relevant samples across different modalities. Existing cross-modal retrieval works [1, 6, 11] usually learn a comparable common space to bridge different modalities, which achieved remarkable progress in many applications, including video-audio retrieval [1, 33], visual question answering [11, 26], and image-text matching [44].

Despite the promise, a core assumption in cross-modal retrieval is the correct correspondence among multiple




Figure 1. Illustration of noisy correspondence in image-text retrieval. A standard triplet loss is used to enforce the positive pairs to be closer than negatives in the common space. Noisy correspondence is the mismatched pairs but wrongly considered as positive ones, and thus results in model performance degradation.

modalities. However, collecting such ideal data is expensive and time-consuming. In practice, most widely used datasets are harvested from the Internet and inevitably contain noisy correspondence [31]. As illustrated in Fig. 1, the cross-modal retrieval method will wrongly enforce the mismatched data to be similar when learning with noisy correspondence, which may significantly affect the retrieval performance. To date, few effort has been made to address this. Huang [15] first researches this issue and proposes the NCR method to train from the noisy image-text pairs robustly. Inspired by the prior success for noisy labels [20], NCR divides the data into clean and noisy partitions and rectifies the correspondence with an adaptive model. However, NCR based on the memorization effect of DNNs [3], which leads to poor performance under high noise ratio.

To tackle the challenge, we propose a **Meta Similarity Correction Network (MSCN)** which aims to provide reliable similarity scores for the noisy features from main net. We view a binary classification task as the meta-process: given a multimodal sample, the MSCN will learn to determine whether the modalities correspond to each other or not, where the prediction of MSCN can be naturally regarded as the similarity score. Specifically, a small amount of clean pairs is used to construct positive and negative meta-data, both viewed as meta-knowledge that encourages MSCN to learn the discrimination. Meanwhile, the main net

*Minnan Luo is the corresponding author.

¹<https://github.com/hhc1997/MSCN>

trained by the corrected similarity score from MSCN with a self-adaptive margin to achieve robust learning. This interdependency makes the main net and MSCN benefits each other against noise. However, due to the property of triplet loss, it remains to produce positive loss for noisy pairs even if we employ the ideal similarity scores. To this end, we further propose a meta-knowledge guided data purification strategy to remove samples with potentially wrong correspondence. Extensive experiments are conducted to demonstrate the strengths of our method in both synthetic and real-world noises.

The main contributions of this work are summarized as follows:

- We pioneer the exploration of meta-learning for noisy correspondence problem, where a meta correction network is proposed to provide reliable similarity scores against noise.
- We present a novel meta-process that first considers both positive and negative data as meta-knowledge to encourage the MSCN to learn discrimination from them.
- We design an effective data purification strategy using meta-data as prior knowledge to remove the noisy samples.

2. Related Work

2.1. Cross-Modal Retrieval

Cross-modal retrieval works aim to take one modality of data as query to retrieve the related data in other modalities. Based on the utilization of annotation information, cross-modal retrieval methods can be roughly divided into two categories: 1) Unsupervised Methods [8, 22, 25, 36]. It only uses the co-occurrence information to learn the common representations among the multimodal data. 2) Supervised methods [14, 34, 38, 45]. The extra label information is employed to boost the discrimination of common representations. However, the lack of annotations limited the practicability in the real world.

Our proposed approach falls in the category of unsupervised methods. For example, Wang [36] uses a two-branch neural networks to learn the joint embeddings of multimodal data. Inspired by hard negative mining, VSE++ [8] uses hard negatives to improve the retrieval performance. To capture fine-grained interplay between modalities, SCAN [18] introduce stacked cross attention to enable attention with context from both image and sentence. Recently, motivated by the powerful learning ability of graph model, GSMN [25] and SGRAF [7] construct graph structure for multimodal data to benefit the learning of fine-grained correspondence. Although existing works have achieved remarkable results, they usually depend on the correct correspondence among multimodal data and cannot

tackle the noise issue. Thus, it is significant to explore how to learn cross-modal retrieval with noisy correspondence, but which is rarely touched in previous studies.

2.2. Noisy Correspondence Learning

As a newly proposed problem, noisy correspondence is the mismatched pairs but wrongly considered as positive ones. Huang [15] first research this issue and propose the NCR method to robustly train the image-text matching model with noisy correspondence pairs. Inspired by the prior success for noisy labels [20], NCR divides the data into clean and noisy part based on the memorization effect of DNNs, and then rectifies the correspondence with an adaptive model. Qin [29] proposes a uncertainty-based method to achieve efficient learning. Moreover, some works [40, 41] study the partially view-unaligned problem which can be viewed as a generalized noisy correspondence learning. Recently, some works [13, 39] explore a more challenging scenario where noisy labels and noisy correspondence may occur simultaneously. In contrast, our work tackles this challenge in cross-modal retrieval from a meta-learning perspective.

2.3. Meta-Learning

The objective of meta-learning is to learn at a more advanced level than conventional learning, such as learning the update rule [28], finding easily fine-tuned parameters [9], or adapting to new tasks [19]. Recently, researchers use meta-learning to find model’s parameters that robust against noisy labels [21, 32, 46]. For example, Li [21] optimizes a meta-objective before conventional training to enable the model not overfit to noise. MW-Net [32] use a meta-process to automatically assign weights to the training samples. MLC [46] presents a label correction network trained as a meta manner which generate reliable labels for the noisy training data.

The most relevant existing method is MLC [46], which first poses the noisy label problem as a meta label correction. However, we argue that our proposed method differs from MLC in two aspects. First, we focuses on the noisy correspondence problem in multimodal data instead of the unimodal classification scenarios. Second, we leverage both positive and negative meta-data to guide model learning the discrimination, which is unexplored in previous meta-learning methods.

3. The Proposed Method

3.1. Problem Formulation

Without losing generality, we use the image-text retrieval task as a proxy to investigate the noisy correspondence problem in cross-modal retrieval. Given a training set $\mathcal{D}_{train} = \{(I_i, T_i, y_i)\}_{i=1}^N$, where (I_i, T_i) is the i -th image-

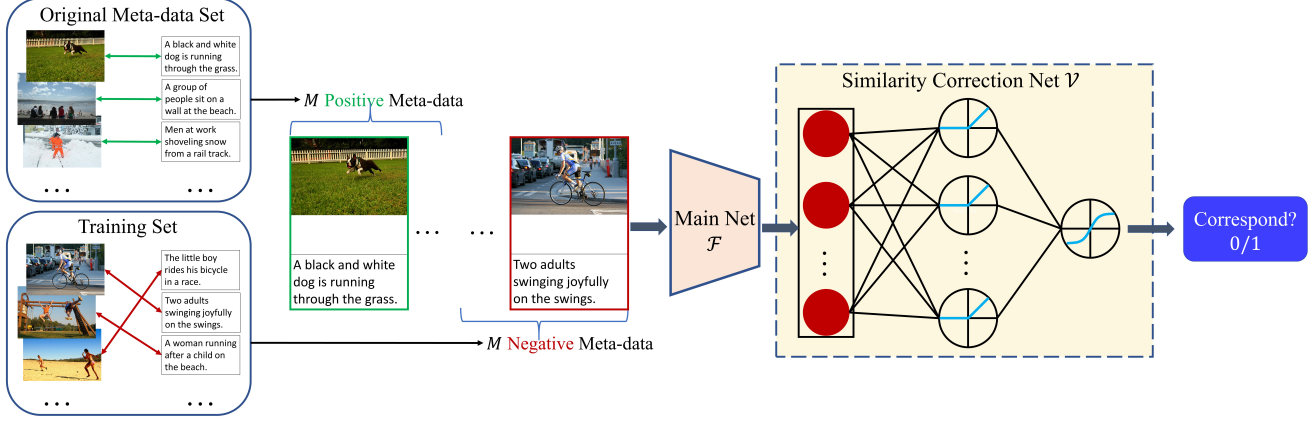


Figure 2. Overview of the proposed meta-process. The positive pairs are taken from the original meta-data which contain corresponding image and text, while negative pairs are extracting from different pairs in training set to construct mismatched pairs. Then the MSCN is learned to determine whether the modalities correspond to each other or not.

text pair, $y_i \in \{1, 0\}$ indicates the pair is positive (matched), *i.e.* (I_i, T_i) , or negative (mismatched), *i.e.* $(I_i, T_{j \neq i})$, and N is the number of the entire training data. The noisy correspondence in bimodal data is the negative pair but wrongly considered as $y_i = 1$. To address this issue, we propose a meta-learning based method to achieve robust training. The details are delineated next.

3.2. Meta-Learning Based Similarity Correction

The Training Objective. To begin with, we describe the standard projection module to enable cross-modal features to be comparable. Let $f(I; W_f)$ and $g(T; W_g)$ be the modal-specific networks to map the visual and textual modalities into the joint embedding space, respectively. Then we compute the similarity feature between bimodal features by the function

$$S(f(I), g(T); W_s) = \frac{W_s |f(I) - g(T)|^2}{\|W_s |f(I) - g(T)|^2\|_2}, \quad (1)$$

where W_s is a learnable parameter matrix to obtain low-dimensional similarity representation. For notation convenience, we collect $\{W_f, W_g, W_s\}$ into W as the parameters of main net $\mathcal{F} = (f, g, S)$ and denote similarity feature as $\mathcal{F}_W(I, T)$. With the presence of noisy correspondence, the model will inevitably overfit to noise if we use the similarity representation directly. To tackle this problem, we propose a meta similarity correction network (MSCN), regarding process of obtaining similarity score as a meta-process, which takes the noisy similarity representation as input and produces the corrected similarity score. Specifically, our MSCN is formulated as a MLP network $\mathcal{V}_\Theta(\cdot)$ with parameters Θ . The output $s_i = \mathcal{V}_\Theta(\mathcal{F}_W(I_i, T_i))$ is activated with Sigmoid function to represent the similarity score located in $[0, 1]$.

For the noisy training pairs, the main net learns noise-tolerant features guided by the corrected similarity scores produced from MSCN; while the MSCN predicts the reliable similarity score based on features from main net. This interdependency enables the main net and meta net reinforce each other against noise, and optimized via a bi-level optimization. Formally, the optimal parameters W^* is calculated by minimizing the cumulative loss over training data:

$$W^*(\Theta) = \arg \min_W \mathbb{E}_{(I_i, T_i) \in \mathcal{D}_{train}} l^{train}(I_i, T_i), \quad (2)$$

where $l^{train}(\cdot)$ is the triplet ranking loss:

$$l^{train}(I_i, T_i) = [\hat{\gamma} - \mathcal{V}_\Theta(\mathcal{F}_W(I_i, T_i)) + \mathcal{V}_\Theta(\mathcal{F}_W(I_i, T_i^-))]_+ + [\hat{\gamma} - \mathcal{V}_\Theta(\mathcal{F}_W(I_i, T_i)) + \mathcal{V}_\Theta(\mathcal{F}_W(I_i^-, T_i))]_+, \quad (3)$$

where $[x]_+ = \max(x, 0)$, I_i^- and T_i^- are the hardest negatives corresponding to given pair (I_i, T_i) similar to VSE++ [8]. To achieve robust cross-modal retrieval, we design a self-adaptive margin $\hat{\gamma}$ instead of a fixed value in VSE++. Ideally, the margin should put less value on the noisy pairs and more on the clean pairs. To this end, we adjust the margin according to the similarity score:

$$\hat{\gamma} = \frac{1}{1 + (\frac{s_i}{1-s_i})^{-\tau}} \gamma, \quad (4)$$

where γ is the original margin value and $\tau > 0$ is the parameter. This self-adaptive margin will enable the triplet loss put more attention on the samples with high similarity scores that possibly are clean pairs.

Meta Training Process. To ensure the reliability of the similarity scores produced by MSCN, we utilize a small

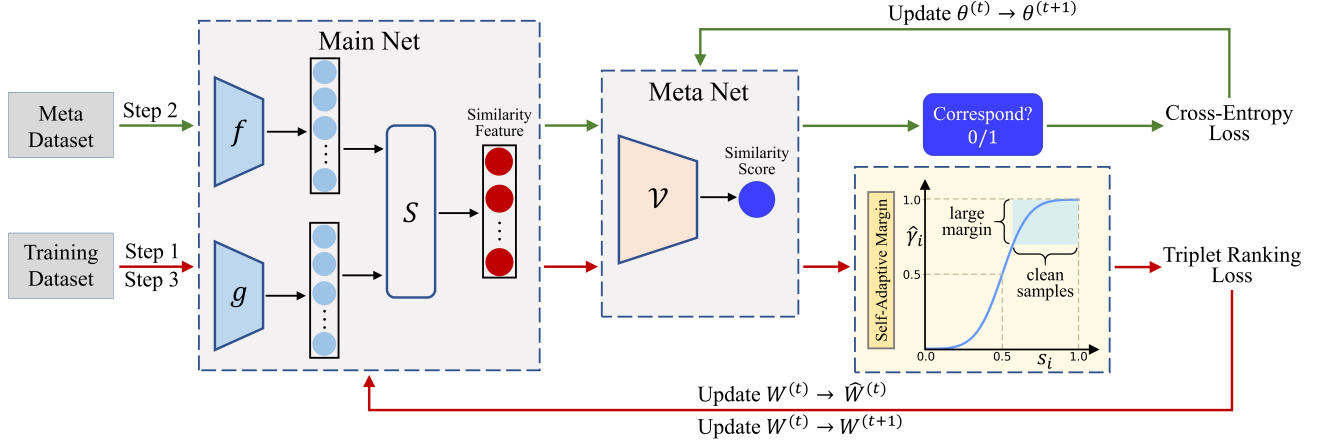


Figure 3. Flowchart of the proposed bi-level optimization. At each iteration, we alternately update the main net and the meta net (MSCN) following the steps in Eq. (7) - (9).

amount clean meta-data set $\mathcal{D}_{meta} = \{(I_i, T_i)\}_{i=1}^M$ to guide the training, where M is much smaller than N . While meta-learning has been extensively studied in recent literature [10, 32, 46], it remains challenging to construct the meta-process in our objective. On the one hand, the similarity features from main net can not used directly to produce loss. On the other hand, the meta-process should enable the output of MSCN to represent the similarity score. To achieve this, we view a binary classification task as the meta-process: given an image-text pair, the MSCN is learned to determine whether the modalities correspond to each other or not. Specifically, the positive pairs are taken from the meta-data which contain corresponding image and text, *i.e.* $(I_i, T_i, y_i = 1)$, while negative pairs are extracting from different pairs in training set to construct mismatched pairs, *i.e.* $(I_i, T_{j \neq i}, y_i = 0)$. The binary labels are produced from the data itself, and can be regarded as the ideal similarity score to guide the training of MSCN. The prediction of MSCN measures the probability of a pair being clean which can be naturally equivalent to the similarity score. Formally, in each iteration of training, we randomly construct M mismatched pairs from training set to extend the meta-data set as $\mathcal{D}'_{meta} = \{(I_i, T_i, y_i)\}_{i=1}^{2M}$. Then the optimal parameter Θ^* can be learned by minimizing the following meta-loss calculated on meta-data:

$$\Theta^* = \arg \min_{\Theta} \mathbb{E}_{(I_i, T_i) \in \mathcal{D}'_{meta}} l^{meta}(I_i, T_i, y_i), \quad (5)$$

where $l^{meta}(\cdot)$ is cross-entropy loss [35, 37]:

$$l^{meta}(I_i, T_i, y_i) = -y_i \cdot \log \mathcal{V}_{\Theta}(\mathcal{F}_{W^*(\Theta)}(I_i, T_i)). \quad (6)$$

As illustrated in Fig. 2, despite its simplicity, our proposed meta-process can not only enforce the MSCN to learn reliable similarity score but also has the following advantages:

- 1) previous works only guide the meta net with ground truth data; however, we leverage both positive and negative data as meta-knowledge that encourage the MSCN to learn discrimination from them.
- 2) We overcome one limitation of meta-learning methods that the lack of meta-data. We extract from different pairs among training set to construct massive negative meta-data. Moreover, it is more suitable to the rule of noisy correspondence, which the irregularity of noisy data causes in various patterns.

Bi-level Optimization. Motivated by recent works [9, 16, 30, 32], we use a bi-level optimization strategy to guarantee the efficiency for updating our main net and MSCN. As illustrated in Fig. 3, it is practically solved via two nested loops of optimization. Specifically, at each iteration t of training, we sample a mini-batch of training pairs $\{(I_i, T_i)\}_{i=1}^n$, where n is the size of mini-batch. We first update the main net parameters by taking a gradient descent step toward the direction of triplet loss, which can be formulated as follows:

$$\hat{W}^{(t)}(\Theta) = W^{(t)} - \alpha \sum_{i=1}^n \nabla_{W} l^{train}(I_i, T_i) \Big|_{W^{(t)}}, \quad (7)$$

where α means the learning rate. The updated $\hat{W}^{(t)}$ can be regard as the approximation of the optimal $W^*(\Theta)$ at t iteration, containing the gradient information of Θ . Then we update the parameters of MSCN with main net parameters fixed. Guided by the objective loss in Eq. (5) and a mini-batch of meta-data $\{(I_i, T_i, y_i)\}_{i=1}^m$, where m is the mini-batch size. The updated parameters of MSCN is calculated as:

$$\Theta^{(t+1)} = \Theta^{(t)} - \beta \frac{1}{m} \sum_{i=1}^m \nabla_{\Theta} l^{meta}(I_i, T_i, y_i) \Big|_{\Theta^{(t)}}, \quad (8)$$

where β is the learning rate for meta-process. Finally, we update the parameters of main net while fixing the MSCN parameters. With the corrected similarity scores produced by the newly updated $\Theta^{(t+1)}$, we can update the main net parameters $W^{(t+1)}$ by employing the gradient decent as:

$$W^{(t+1)} = W^{(t)} - \alpha \sum_{i=1}^n \nabla_W l^{train}(I_i, T_i)|_{W^{(t)}}. \quad (9)$$

The above procedure proceeds in an alternating manner to optimize the main net and MSCN until it exceeds the maximal iteration.

3.3. Meta-Knowledge Guided Data Purification

As mentioned, our MSCN can provide corrected similarity score to avoid noise. However, due to the property of triplet loss, it remains to produce positive loss for the noisy pairs even if we employ the ideal similarity score (*e.g.*, $\mathcal{V}(F(I_i, T_i)) \approx 0 < \mathcal{V}(F(I_i^-, T_i)) + \hat{\gamma}$), which results in the model fitting to the noise. To tackle this issue, we adopt sample selection strategy to remove the noisy data and only use clean samples as training data. Previous works [2, 42] motivated by the memorization effect of DNNs that clean sample has a lower loss during the beginning of training. Unfortunately, the triplet ranking loss is more complex, *i.e.*, affected by both positive and negative samples, which leads to sub-optimal selection results. Based on the observation shown in Fig. 4, the clean and noisy pairs are prone to distinguish from the similarity scores. Gaussian Mixture Model (GMM) is a widely used technique to model mixture distribution [15, 20]. However, we find the similarity distribution exhibits high skew toward 1 and causes a poor approximation for GMM. Therefore, we fit a two-component Beta Mixture Model (BMM) [27] of per-pair similarity score s_i to better distinguish clean and noisy samples:

$$p(s_i) = \sum_{k=1}^K \lambda_k \phi(s_i | \alpha_k, \beta_k), \quad (10)$$

where $K = 2$ and λ_k is the mixture coefficient, and $\phi(s | \alpha_k, \beta_k)$ denotes the probability density function with parameters $\alpha_k, \beta_k > 0$.

We use the Expectation-Maximization [5] algorithm to optimize the BMM, which poses a crucial challenge: the initialization of model parameters. Existing methods typically adopt the K -Means algorithm [23] for initializing; however, it is time-consuming (*i.e.*, the time complexity is $O(TKN)$, where T is the number of iterations) and underperforming. We thus propose a novel initialization approach which takes the meta-data as prior knowledge. Specifically, we fit the two Beta components with positive and negative meta-data to model the distribution of clean and noisy pairs, respectively. Given the similarity scores $\mathcal{S}_p = \{s_i\}_{i=1}^M$

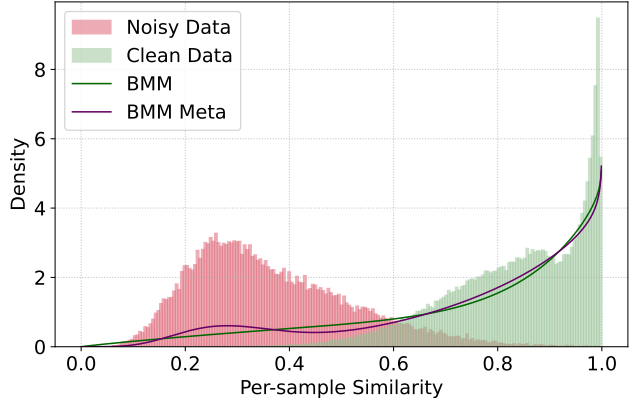


Figure 4. Per-sample similarity score distribution, estimated BMM and BMM with meta prior knowledge under 20% noisy correspondence in Flickr30K after 5 epochs warmup.

corresponding to positive meta-data, the parameters of the component modeling clean distribution can be initialised by

$$\begin{aligned} \alpha_k &= \frac{(1 - E(\mathcal{S}_p)) E(\mathcal{S}_p)^2}{V(\mathcal{S}_p)} - E(\mathcal{S}_p), \\ \beta_k &= \frac{\alpha_k (1 - E(\mathcal{S}_p))}{E(\mathcal{S}_p)}, \end{aligned} \quad (11)$$

where $E(\cdot)$ and $V(\cdot)$ are the mean and variance function, respectively. The initialization of component modeling noisy distribution is analogous but uses similarity scores $\mathcal{S}_n = \{s_i\}_{i=1}^M$ of negative meta-data. As illustrated in Fig. 4, our meta-knowledge guided initialization method achieves better approximation while with lower time complexity (*i.e.*, $O(KM)$).

Finally, we compute the posterior probability $p_i = p(k | s_i) = p(k) p(s_i | \alpha_k, \beta_k) / p(s_i)$ as the probability of i -th simple being clean. The purified training set is defined as:

$$\mathcal{D}'_{train} = \{(I_i, T_i) \in \mathcal{D}_{train} | p_i > 0.5\}. \quad (12)$$

At each iteration, we produce the purified dataset before the training of our meta-learning objective. Moreover, following [12, 15, 20], we maintain two networks $\{\mathcal{F}_W^1, \mathcal{V}_\Theta^1\}$ and $\{\mathcal{F}_W^2, \mathcal{V}_\Theta^2\}$ to avoid error accumulation that one network will produce the purified training set to train the other one. The full algorithm is shown in *supplementary material*.

4. Experiment

4.1. Datasets and Evaluation Protocol

Datasets. We evaluate our method on three standard image-text retrieval datasets. To be specific, the Flickr30K [43] collects 31,783 images with 5 corresponding captions each from Flickr website. Following the split in [15], we use 1,000 image-text pairs for validation, 1,000 image-text

pairs for testing, and the rest for training. The MS-COCO [24] contains 123,287 images with 5 corresponding captions each. We split the MS-COCO following [15] that 113,287 image-text pairs for training, 5,000 pairs for validation, and 5,000 pairs for testing. Conceptual Captions [31] is a large-scale dataset contains 3.3M images with one corresponding caption each. We employ the same subset following [15], *i.e.*, CC152K. Specifically, CC152K consists of 150,000 pairs from training set for training, 1,000 pairs from validation set for validation, and 1,000 pairs from validation set for testing. Moreover, we use an extra meta-data set from validation containing 3,000 pairs.

Evaluation Protocol. We evaluate the retrieval performance with the recall rate at K (R@K) metric. In a nutshell, R@K is the proportion for the relevant items retrieved in the closest K items to the query. In our experiments, we take image and text as queries, respectively, and report R@1, R@5, and R@10 results for a comprehensive evaluation. Following [15], we average the similarity scores from two networks at the inference phase.

4.2. Implementation Details

As a general approach, our method can be easily applied to almost all cross-modal retrieval methods to improve robustness. For fair comparison, we chose the same network backbones with NCR [15], *i.e.*, the projection modules and similarity function. For all datasets, the training processes contain 50 epochs after a 5 epochs warmup. We employ ADAM [17] as the optimizer for both main net and meta net with a batch size of 64. To fit the BMM, we set the stop threshold as 10^{-2} and maximum number of iterations as 10 for the EM procedure. Note that, we clamp the similarity scores into a range of $[10^{-4}, 1 - 10^{-4}]$ instead of $[0,1]$ for stable fitting. Moreover, we set the margin parameter as 0.2 and $\tau = 2$ to calculate the self-adaptive margin. For all datasets, the number of ground-truth meta-data is approximately 2% of training data, and the learning rate of main net and meta net are initialized with 2×10^{-4} and 1.7×10^{-5} , respectively. For CC152K containing real noise, we employ the validation data as the meta set due to the lack of ground-truth data. For all experiments, we decay the learning rate by 0.1 after 30 epochs.

4.3. Comparison with the State-of-the-Art

In this section, we conduct comparison experiments with 7 state-of-the-art methods that include SCAN [18], VSRN [22], IMRAM [4], SGR, SAF, SGRAF [7], and NCR [15]. As Flickr30K and MS-COCO are well annotated datasets, we generate synthetic noisy correspondence by randomly shuffling the training images and captions for a specific noise ratio, and we conduct experiments with three different

Methods	Image to Text			
	R@1	R@5	R@10	SUM
SCAN (ECCV'18)	30.5	55.3	65.3	151.1
VSRN(ICCV'19)	32.6	61.3	70.5	164.4
IMRAM(CVPR'20)	33.1	57.6	68.1	158.8
SAF (AAAI'21)	31.7	59.3	68.2	159.2
SGR (AAAI'21)	11.3	29.7	39.6	80.6
SGR* (AAAI'21)	35.0	63.4	73.3	171.7
NCR(NIPS'21)	39.5	64.5	73.5	177.5
MSCN*(Ours)	39.7	65.4	75.3	180.4
MSCN(Ours)	40.1	65.7	76.6	182.4
Methods	Text to Image			
	R@1	R@5	R@10	SUM
SCAN (ECCV'18)	26.9	53.0	64.7	144.6
VSRN(ICCV'19)	32.5	59.4	70.4	162.3
IMRAM(CVPR'20)	29.0	56.8	67.4	153.2
SAF (AAAI'21)	31.9	59.0	67.9	158.8
SGR (AAAI'21)	13.1	30.1	41.6	84.8
SGR* (AAAI'21)	34.9	63.0	72.8	170.7
NCR(NIPS'21)	40.3	64.6	73.2	178.1
MSCN*(Ours)	39.8	66.1	75.0	180.9
MSCN(Ours)	40.6	67.4	76.3	184.3

Table 1. Image-Text Retrieval Performance on CC152K with real noise, and the best results are highlighted in **bold**.

level of noise ratios, *i.e.*, 20%, 50%, and 70%. For all methods, we choose the best checkpoint on the validation set and report its performance on the test set. Following NCR, we also report two strong baselines based on SGR, *i.e.*, SGR-C and SGR*. In short, SGR-C uses only clean data for training, and SGR* employs a pre-training process while training without hard negatives to improve robustness. To evaluate our method on real-world noisy data, we conduct experiments on the CC152K, which is harvested from the Internet and contains about 3% ~ 20 % noisy correspondence [31].

Experiments on CC152K. Tab. 1 shows the results on CC152K with real noisy correspondence. On CC152K, our MSCN achieves the state-of-the-art performance in terms of all metrics. Specifically, the sum score of MSCN is 4.9 % and 6.2 % higher than the best baseline in text and image retrieval, respectively. In addition, we also present a more strict comparison (denoted by MSCN*) that trained our method with only 1,000 meta-data (approximately 0.67% of all data). We can see that our MSCN* also achieves a competitive performance. Specifically, all metrics except R@1 of image retrieval are superior to the baselines.

Experiments on Flickr30K and MS-COCO. Tab. 2 shows the results on Flickr30K and MS-COCO under a range of synthetic noise rates. Following [15], we report the

Noise	Methods	Flickr30K						MS-COCO					
		Image to Text			Text to Image			Image to Text			Text to Image		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
20%	SCAN	59.1	83.4	90.4	36.6	67.0	77.5	66.2	91.0	96.4	45.0	80.2	89.3
	VSRN	58.1	82.6	89.3	40.7	68.7	78.2	25.1	59.0	74.8	17.6	49.0	64.1
	IMRAM	63.0	86.0	91.3	41.4	71.2	80.5	68.6	92.8	97.6	55.7	85.0	91.0
	SAF	51.0	79.3	88.0	38.3	66.5	76.2	67.3	92.5	96.6	53.4	84.5	92.4
	SGR*	62.8	86.2	92.2	44.4	72.3	80.4	67.8	91.7	96.2	52.9	83.5	90.1
	SGR-C	72.8	90.8	95.4	56.4	82.1	88.6	75.4	95.2	97.9	60.1	88.5	94.8
	NCR	75.0	93.9	97.5	58.3	83.0	89.0	77.7	95.5	98.2	62.5	89.3	95.3
	MSCN	77.4	94.9	97.6	59.6	83.2	89.2	78.1	97.2	98.8	64.3	90.4	95.8
50%	SCAN	27.7	57.6	68.8	16.2	39.3	49.8	40.8	73.5	84.9	5.4	15.1	21.0
	VSRN	14.3	37.6	50.0	12.1	30.0	39.4	23.5	54.7	69.3	16.0	47.8	65.9
	IMRAM	9.1	26.6	38.2	2.7	8.4	12.7	21.3	60.2	75.9	22.3	52.8	64.3
	SAF	30.3	63.6	75.4	27.9	53.7	65.1	30.4	67.8	82.3	33.5	69.0	82.8
	SGR*	36.9	68.1	80.2	29.3	56.2	67.0	60.6	87.4	93.6	46.0	74.2	79.0
	SGR-C	69.8	90.3	94.8	50.1	77.5	85.2	71.7	94.1	97.7	57.0	86.6	93.7
	NCR	72.9	93.0	96.3	54.3	79.8	86.5	74.6	94.6	97.8	59.1	87.8	94.5
	MSCN	74.4	93.2	96.0	55.3	80.4	86.8	77.5	96.2	98.7	60.7	89.1	94.9
70%	SCAN	5.6	19.3	27.4	2.2	8.0	12.8	18.1	43.1	57.4	0.3	1.3	2.3
	VSRN	0.8	2.5	4.1	0.5	1.5	2.7	5.1	15.7	24.6	2.5	8.8	13.3
	IMRAM	1.3	3.1	3.9	0.3	1.2	2.8	7.1	20.0	33.4	5.3	15.2	22.0
	SAF	0.5	2.2	3.0	0.2	0.8	1.7	0.1	1.7	4.0	0.6	1.9	3.0
	SGR*	17.9	42.1	51.9	14.6	31.0	40.8	35.7	71.2	85.4	31.6	65.8	79.0
	SGR-C	65.0	89.3	94.7	48.1	74.5	81.1	69.8	93.6	97.5	56.5	86.0	93.4
	NCR	16.1	38.5	52.8	11.0	29.5	41.4	35.4	69.5	83.4	31.5	66.4	81.1
	MSCN	69.0	89.3	93.0	49.2	73.1	79.0	74.4	94.9	97.7	58.8	87.2	93.7

Table 2. Image-Text Retrieval Performance under synthetic noise rates of 20%, 50% and 70% on Flickr30K and MS-COCO 1K, and the best results are highlighted in **bold**.

results by averaging over 5 folds of 1K test images. From Tab. 2, we can draw the observation that noisy correspondence remarkably affect the performance of cross-modal retrieval methods. The retrieval accuracy will decrease fast with the noise rate increasing. Our MSCN is superior to the baselines in almost all metrics under different noise rates. In the low noise cases, *i.e.*, 20% and 50%, our MSCN improves R@1 (sum of image and text) by 3.7%, 2.2%, 2.5%, and 4.5% in the four valuations compared with NCR. In the high noise case, *i.e.*, 70%, all baselines can easily overfit to noisy correspondence and results in poor performance. Our MSCN improves R@1 by 52.9%, 38.2%, 39.4%, and 27.3% in these four valuations compared with NCR. Even compared with the SGR-C trained only on clean data, our MSCN also achieves competitive performance.

4.4. Ablation Study

To evaluate the performance of the proposed components (*i.e.*, self-adaptive margin and data purification), we conduct the ablation study on the Flickr30K with 20% noisy correspondence. Note that for MSCN without \mathcal{D}'_{train} , we only use a single model to perform the experiment. As shown in Tab. 3, we can see that the performance without $\hat{\gamma}$ or \mathcal{D}'_{train} are worse than the complete MSCN, which

indicates that all components are important to achieve advantageous results. We explain more on the data selection strategy in the *supplementary material*.

Method			Image to Text		
MSCN	w/o $\hat{\gamma}$	w/o \mathcal{D}'_{train}	R@1	R@5	R@10
✓			77.4	94.9	97.6
✓	✓		75.3	94.5	97.2
✓		✓	75.8	93.8	96.2
✓	✓	✓	74.1	91.5	94.7
			Text to Image		
✓			59.6	83.2	89.2
✓	✓		58.3	83.1	88.9
✓		✓	55.8	74.5	77.4
✓	✓	✓	53.4	71.2	72.2

Table 3. Ablation studies on Flickr30K with 20% noise rate.

4.5. Progressive Comparison

Fig. 6 plots the average of recalls on testing set of MSCN and NCR as training proceeds. We show a representative result using Flickr30K with 50% noisy correspondence. From the result, we can see that our MSCN achieves a more stable performance in overall training process. Although NCR



Figure 5. Real-world noisy examples detected by our MSCN.

does not need an extra meta-set to guide the training, it relies on a clean validation set to choose the best model parameters. The stableness also indicates that our method has better alleviated the interference of noisy correspondence.




Figure 6. Progressive performance comparison of MSCN and NCR on testing set as training proceeds.

4.6. Detected Noisy Samples

Fig. 5 shows some real-world noisy image-text pairs detected (with low similarity score) by MSCN in Conceptual Captions. We also present the similarity score predicted by our MSCN. For most noisy pairs, the image and text contains completely irrelevant semantic information. Moreover, our MSCN can even find the subtle mistake in the noisy image-text pair. In the last pair, the picture shows a man covering his ears but wrongly caption as using a mobile phone. Although this picture is misleading, our method successfully detects the subtle difference.

4.7. Visualization on Similarity Score

Fig. 7 plots the similarity score distribution predicted by MSCN for clean and noisy training samples. It can be seen that almost all large similarity scores belong to clean pairs, and the noisy pairs' scores are smaller than clean samples, which implies that the trained MSCN can successfully provide corrected similarity scores.

5. Conclusion

In this paper, we explore a meta-learning method to address the problem of learning with noisy correspon-

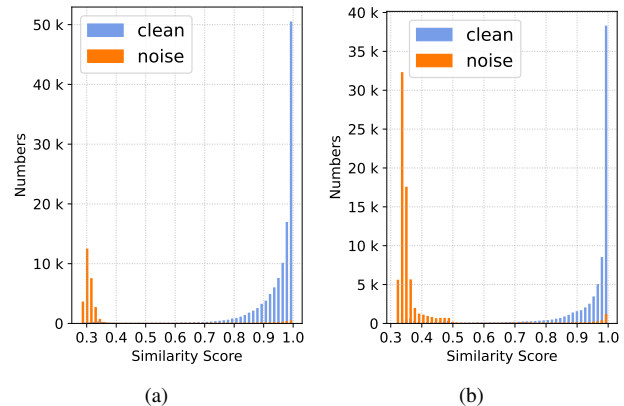


Figure 7. Per-sample similarity distribution on training data of Flickr30K. (a) 20% noisy correspondence. (b) 50% noisy correspondence.

dence. Specifically, we propose a meta correction network (MSCN) to provide reliable similarity scores. Our MSCN is trained by a novel meta-process that views both positive and negative data as meta-knowledge to encourage MSCN to learn discrimination from them. To further mitigate noise interference, we design a data purification strategy that uses meta-data as prior knowledge to purify the noisy data efficiently. We conduct comprehensive experiments on three widely used datasets. The results validate the effectiveness of our method in both synthetic and real noises.

Acknowledgements

This work is supported by the National Key Research and Development Program of China (2020AAA0108800), National Nature Science Foundation of China (62192781, 62272374, 62202367, 62250009, 62137002, 61937001), Innovative Research Group of the National Natural Science Foundation of China (61721002), Innovation Research Team of Ministry of Education (IRT_17R86), Project of China Knowledge Center for Engineering Science and Technology, and Project of Chinese academy of engineering “The Online and Offline Mixed Educational Service System for ‘The Belt and Road’ Training in MOOC China”. We would like to express our gratitude for the support of K. C. Wong Education Foundation.

References

- [1] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European conference on computer vision (ECCV)*, pages 435–451, 2018. 1
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel O’Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *International conference on machine learning*, pages 312–321. PMLR, 2019. 5
- [3] Devansh Arpit, Stanislaw K Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron C Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *ICML*, 2017. 1
- [4] Hui Chen, Guiguang Ding, Xudong Liu, Zijia Lin, Ji Liu, and Jungong Han. Imram: Iterative matching with recurrent attention memory for cross-modal image-text retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12655–12663, 2020. 6
- [5] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977. 5
- [6] Cheng Deng, Zhaojia Chen, Xianglong Liu, Xinbo Gao, and Dacheng Tao. Triplet-based deep hashing network for cross-modal retrieval. *IEEE Transactions on Image Processing*, 27(8):3893–3903, 2018. 1
- [7] Haiwen Diao, Ying Zhang, Lin Ma, and Huchuan Lu. Similarity reasoning and filtration for image-text matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1218–1226, 2021. 2, 6
- [8] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. *arXiv preprint arXiv:1707.05612*, 2017. 2, 3
- [9] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017. 2, 4
- [10] Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil. Bilevel programming for hyperparameter optimization and meta-learning. In *International Conference on Machine Learning*, pages 1568–1577. PMLR, 2018. 4
- [11] Dalu Guo, Chang Xu, and Dacheng Tao. Image-question-answer synergistic network for visual dialog. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10434–10443, 2019. 1
- [12] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in neural information processing systems*, 31, 2018. 5
- [13] Haochen Han, Qinghua Zheng, Minnan Luo, Kaiyao Miao, Feng Tian, and Yan Chen. Noise-tolerant learning for audio-visual action recognition. *arXiv preprint arXiv:2205.07611*, 2022. 2
- [14] Peng Hu, Xi Peng, Hongyuan Zhu, Liangli Zhen, and Jie Lin. Learning cross-modal retrieval with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5403–5413, 2021. 2
- [15] Zhenyu Huang, Guocheng Niu, Xiao Liu, Wenbiao Ding, Xinyan Xiao, Hua Wu, and Xi Peng. Learning with noisy correspondence for cross-modal matching. *Advances in Neural Information Processing Systems*, 34:29406–29419, 2021. 1, 2, 5, 6
- [16] Haoxuanye Ji, Le Wang, Sanping Zhou, Wei Tang, Nanning Zheng, and Gang Hua. Meta pairwise relationship distillation for unsupervised person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3661–3670, 2021. 4
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015. 6
- [18] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *Proceedings of the European conference on computer vision (ECCV)*, pages 201–216, 2018. 2, 6
- [19] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 2
- [20] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *International Conference on Learning Representations*, 2019. 1, 2, 5
- [21] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. Learning to learn from noisy labeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5051–5059, 2019. 2
- [22] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4654–4662, 2019. 2, 6
- [23] Aristidis Likas, Nikos Vlassis, and Jakob J Verbeek. The global k-means clustering algorithm. *Pattern recognition*, 36(2):451–461, 2003. 5
- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [25] Chunxiao Liu, Zhendong Mao, Tianzhu Zhang, Hongtao Xie, Bin Wang, and Yongdong Zhang. Graph structured network for image-text matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10921–10930, 2020. 2
- [26] Lin Ma, Zhengdong Lu, and Hang Li. Learning to answer questions from image using convolutional neural network. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016. 1
- [27] Zhanyu Ma and Arne Leijon. Bayesian estimation of beta mixture models with variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(11):2160–2173, 2011. 5
- [28] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. In *International Conference on Learning Representations*, 2018. 2

- [29] Yang Qin, Dezhong Peng, Xi Peng, Xu Wang, and Peng Hu. Deep evidential learning with noisy correspondence for cross-modal retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4948–4956, 2022. [2](#)
- [30] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International conference on machine learning*, pages 4334–4343. PMLR, 2018. [4](#)
- [31] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. [1](#), [6](#)
- [32] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. *Advances in neural information processing systems*, 32, 2019. [2](#), [4](#)
- [33] Didac Surís, Amanda Duarte, Amaia Salvador, Jordi Torres, and Xavier Giró-i Nieto. Cross-modal embeddings for video and audio retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. [1](#)
- [34] Jun Tang, Ke Wang, and Ling Shao. Supervised matrix factorization hashing for cross-modal retrieval. *IEEE Transactions on Image Processing*, 25(7):3157–3166, 2016. [2](#)
- [35] Bokun Wang, Yang Yang, Xing Xu, Alan Hanjalic, and Heng Tao Shen. Adversarial cross-modal retrieval. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 154–162, 2017. [4](#)
- [36] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5005–5013, 2016. [2](#)
- [37] Yunchao Wei, Yao Zhao, Canyi Lu, Shikui Wei, Luoqi Liu, Zhenfeng Zhu, and Shuicheng Yan. Cross-modal retrieval with cnn visual features: A new baseline. *IEEE transactions on cybernetics*, 47(2):449–460, 2016. [4](#)
- [38] Xing Xu, Li He, Huimin Lu, Lianli Gao, and Yanli Ji. Deep adversarial metric learning for cross-modal retrieval. *World Wide Web*, 22(2):657–672, 2019. [2](#)
- [39] Mouxing Yang, Zhenyu Huang, Peng Hu, Taihao Li, Jiancheng Lv, and Xi Peng. Learning with twin noisy labels for visible-infrared person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14308–14317, 2022. [2](#)
- [40] Mouxing Yang, Yunfan Li, Peng Hu, Jinfeng Bai, Jiancheng Lv, and Xi Peng. Robust multi-view clustering with incomplete information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):1055–1069, 2022. [2](#)
- [41] Mouxing Yang, Yunfan Li, Zhenyu Huang, Zitao Liu, Peng Hu, and Xi Peng. Partially view-aligned representation learning with noise-robust contrastive loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1134–1143, 2021. [2](#)
- [42] Quanming Yao, Hansi Yang, Bo Han, Gang Niu, and James Tin-Yau Kwok. Searching to exploit memorization effect in learning with noisy labels. In *International Conference on Machine Learning*, pages 10789–10798. PMLR, 2020. [5](#)
- [43] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014. [5](#)
- [44] Qi Zhang, Zhen Lei, Zhaoxiang Zhang, and Stan Z Li. Context-aware attention network for image-text retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3536–3545, 2020. [1](#)
- [45] Liangli Zhen, Peng Hu, Xu Wang, and Dezhong Peng. Deep supervised cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10394–10403, 2019. [2](#)
- [46] Guoqing Zheng, Ahmed Hassan Awadallah, and Susan Dumais. Meta label correction for noisy label learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11053–11061, 2021. [2](#), [4](#)