

# BiCro: Noisy Correspondence Rectification for Multi-modality Data via Bi-directional Cross-modal Similarity Consistency

Shuo Yang<sup>1</sup> Zhaopan Xu<sup>2</sup> Kai Wang<sup>3</sup> Yang You<sup>3</sup>  
 Hongxun Yao<sup>2</sup> Tongliang Liu<sup>4</sup> Min Xu<sup>1\*</sup>

<sup>1</sup> University of Technology Sydney <sup>2</sup> Harbin Institute of Technology  
<sup>3</sup> National University Singapore <sup>4</sup> The University of Sydney

## Abstract

As one of the most fundamental techniques in multi-modal learning, cross-modal matching aims to project various sensory modalities into a shared feature space. To achieve this, massive and correctly aligned data pairs are required for model training. However, unlike unimodal datasets, multimodal datasets are extremely harder to collect and annotate precisely. As an alternative, the co-occurred data pairs (e.g., image-text pairs) collected from the Internet have been widely exploited in the area. Unfortunately, the cheaply collected dataset unavoidably contains many mismatched data pairs, which have been proven to be harmful to the model's performance. To address this, we propose a general framework called BiCro (Bidirectional Cross-modal similarity consistency), which can be easily integrated into existing cross-modal matching models and improve their robustness against noisy data. Specifically, BiCro aims to estimate soft labels for noisy data pairs to reflect their true correspondence degree. The basic idea of BiCro is motivated by that – taking image-text matching as an example – similar images should have similar textual descriptions and vice versa. Then the consistency of these two similarities can be recast as the estimated soft labels to train the matching model. The experiments on three popular cross-modal matching datasets demonstrate that our method significantly improves the noise-robustness of various matching models, and surpass the state-of-the-art by a clear margin. The code is available at <https://github.com/xu5zhao/BiCro>.

## 1. Introduction

As a key step towards general intelligence, multimodal learning aims to endow an agent with the ability to extract and understand information from various sensory modalities. One of the most fundamental techniques in multimodal

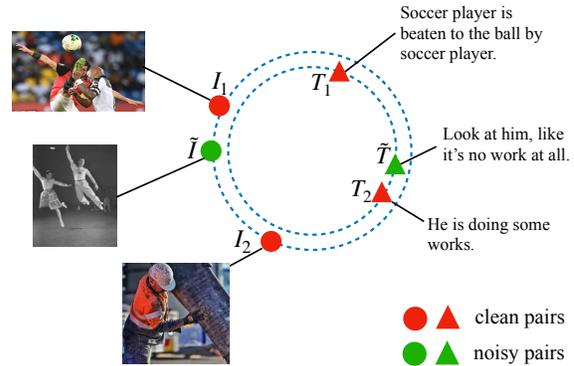


Figure 1. Illustration of the *Bidirectional Cross-modal similarity consistency (BiCro)*. In an image-text shared feature space, assume  $(I_1, T_1)$  and  $(I_2, T_2)$  are two clean positive data pairs. The image  $\tilde{I}$  is very close to the image  $I_1$ , but their corresponding texts  $\tilde{T}$  and  $T_1$  are far away from each other. Similarly, the text  $\tilde{T}$  is very close to the text  $T_2$  while their corresponding images  $\tilde{I}$  and  $I_2$  are irrelevant. Therefore the data pair  $(\tilde{I}, \tilde{T})$  is possibly noisy (mismatched).

learning is cross-modal matching. Cross-modal matching tries to project different modalities into a shared feature space, and it has been successfully applied to many areas, e.g., image/video captioning [1, 21, 23], cross-modal hashing [14, 18], and visual question answering [15, 58].

Like many other tasks in deep learning, cross-modal matching also requires massive and high-quality labeled training data. But even worse, the multimodal datasets (e.g., in image-text matching) are significantly harder to be manually annotated than those unimodal datasets (e.g., in image classification). This is because the textual description of an image is very subjective, while the categorical label of an image is easier to be determined. Alternatively, many works collect the co-occurred image-text pairs from the Internet to train the cross-modal matching model, e.g., CLIP [32] was trained over 400 million image-text pairs collected from the Internet. Furthermore, one of the largest public datasets in cross-modal matching, Conceptual Caption [36], was also

\*Corresponding Author(min.xu@uts.edu.au).

harvested automatically from the Internet. Such cheap data collection methods inevitably introduce noise into the collected data, *e.g.*, the Conceptual Caption dataset [36] contains 3% ~ 20% mismatched image-text pairs. The noisy data pairs undoubtedly hurt the generalization of the cross-modal matching models due to the memorization effect of deep neural networks [55].

Different from the noisy labels in classification, which refer to categorical annotation errors, the noisy labels in cross-modal matching refer to alignment errors in paired data. Therefore most existing noise-robust learning methods developed for classification cannot be directly applied to the cross-modal matching problem. Generally, all the data pairs in a noisily-collected cross-modal dataset can be categorized into the following three types based on their noise level: *well-matched*, *weakly-matched*, and *mismatched*. The *well-matched* data pairs can be treated as clean examples, and the *weakly-matched* and *mismatched* data pairs are noisy examples since they are all labeled as positive ( $y = 1$ ) in the dataset. We illustrate some noisy data pairs, including both *weakly-matched* and *mismatched* pairs in Fig. 2. The key challenge in noise-robust cross-modal matching is *how to estimate accurate soft correspondence labels for those noisy data pairs*. The estimated soft label is expected to be able to depict the *true correspondence degree* between different modalities in a data pair.

In this work, we propose a general framework for soft correspondence label estimation given only noisily-collected data pairs. Our framework is based on the *Bidirectional Cross-modal similarity consistency (BiCro)*, which inherently exists in paired data. An intuitive explanation of the *BiCro* is – taking image-text matching task as an example – similar images should have similar textual descriptions and vice versa. In other words, if two images are highly similar but their corresponding texts are irrelevant, there possibly exists image-text mismatch issue (see Fig. 1). Specifically, we first select a set of clean positive data pairs as *anchor points* out of the given noisy data. This step is achieved by modeling the per-sample loss distribution of the dataset using a Beta-Mixture-Model (BMM) and selecting those samples with a high probability of being clean. Then, with the help of the selected anchor points, we compute the *Bidirectional Cross-modal similarity consistency (BiCro)* scores for every noisy data pair as their soft correspondence labels. Finally, the estimated soft correspondence labels are recast as soft margins to train the cross-modal matching model in a co-teaching manner [11] to avoid the self-sample-selection error accumulation problem.

Before delving into details, we summarize our main contributions below:

- This paper tackles a widely-exist but rarely-explored problem in cross-modal matching, *i.e.*, the noisy cor-

respondence issue. We identify and explore the key challenge of noise-robust cross-modal learning: how to rectify the *binary noisy correspondence labels* into more accurate *soft correspondence labels* for noisy data pairs.

- We propose a general framework called *Bidirectional Cross-modal similarity consistency (BiCro)* for soft correspondence label estimation given only noisy data pairs. The BiCro is motivated by a rational assumption that – taking the image-text matching as an example – *similar images should have similar textual descriptions and vice versa*. The BiCro can be easily adapted to any cross-modal matching models to improve their noise robustness.
- To compute the BiCro score for noisy data pairs, we propose to first select some clean positive data pairs as *anchor points* out of the noisy dataset by modeling the per-sample loss distribution of the dataset through a *Beta-Mixture-Model (BMM)*. Then a set of examples with high clean possibilities can be selected. The BMM has a better modeling ability on the skewed clean data loss distribution compared to the *Gaussian-Mixture-Model (GMM)* used in the previous method [12].
- Extensive experiments on three cross-modal matching datasets, including both synthetic and real-world noisy labels, demonstrate that our method surpasses the state-of-the-art by a large margin. The visualization results also indicate that our estimated soft correspondence labels are highly-aligned with the correlation degree between different modalities in data pairs.

## 2. Related Works

### 2.1. Cross-modal Matching

Cross-modal Matching is an essential research area of computer vision and multimedia, which aims at retrieving similar items that are of different modals with respect to the query modal. Image and text information are the most common modalities thus have been widely developed. The main challenge of this task is to measure the similarity between different modalities. Most existing works map images and texts into a shared latent space where they can be aligned. According to the alignment, existing works could be divided into two categories: 1) Global alignment, the image and text are matched from a global feature computed by multiple neural networks. For example, [8, 35, 43, 57] propose a two-stream global feature learning network and compute the pairwise similarity between global features. 2) Local alignment, the regions within an image and words in a sentence correspond with each other such that finer-



Figure 2. Some noisy data pairs in the Flickr30K [53] and Conceptual Caption [36] datasets. The first row shows our estimated soft correspondence labels for the image-text pairs illustrated in the second and third rows. We highlight the matched words in green and the mismatched words in red.

grained detail is captured by mapping semantically similar items. For example, [5, 16] adopt attention mechanism to explore the semantic region-word correspondences. [6, 23] infer relation-aware similarities with both the local and global alignments by the graph convolutional network. However, all these alignment methods assume that the training data pairs are perfectly aligned, which is impossible to satisfy due to the high collection and annotation cost.

## 2.2. Learning with Noisy Labels

Supervised training of deep learning models requires precisely labeled datasets and noisy labels can significantly degrade the generalization of models. Therefore, various methods were proposed to improve robustness against noisy labels. The typical algorithms for combating noisy labels including adding regularization [9, 10, 22, 24, 40, 41, 44], designing robust loss functions [26, 27, 30, 45, 47–49], selecting possible clean samples [11, 13, 29, 34, 46, 52, 54] and correcting the labels [28, 38, 42, 50, 51]. We mainly introduce the label correction approaches which are most related to this work. Specifically, label correction algorithms aim to identify and correct suspicious labels in an iterative framework. Bootstrapping [33] combines the noisy label and the DNN prediction as a soft pseudo-label to replace the original label. Furthermore, Dynamic bootstrapping [2] use a Beta-Mixture-Model to combine label dynamically for every sample. Differently, Joint Optimization [39] trained their model with a large learning rate for several epochs firstly, and then generated pseudo-labels by averaging the predictions of that epochs. Recently, SELFIE [37] and Ada-Corr [59] selectively combine the labels of noisy data for corresponding to true labels with a high probability. However, almost all existing noisy label studies, focusing on the classification task, cannot be adapted to cross-modal matching directly in which the noisy label refers to the alignment errors between data pairs rather than the categorical labeling errors. The label of cross-modal matching can be consider as  $y_i \in \{0, 1\}$ , where data pair is positively correlated ( $y_i = 1$ ) or not ( $y_i = 0$ ). Therefore, label correction of noisy multi-modal data pairs can be translated into generating a continuous value between 0

and 1, called soft correspondence labels in this paper, to measure the true degree of correspondence. In this view, the methods proposed by NCR [12] and DECL [31] both improve their noise-robustness by correcting the soft correspondence label. Specifically, NCR [12] recasts the network’s prediction as the estimated soft correspondence label and DECL [31] models the uncertainty of cross-modal correspondence to predict correct correspondence of paired data. However, these methods rely on the network’s predictions, which would cause severe confirmation bias: Those confident but wrong predictions would be used to guide subsequent training, leading to a loop of self-reinforcing errors [4]. Differently, we rectify the noisy correspondence labels by the Bidirectional Cross-modal similarity consistency (BiCro) that is inherently contained in the paired data. Our soft correspondence labels are generated by considering the characteristics of multi-modal data itself, so as to avoid the confirmation bias problem of previous methods.

## 3. Methodology

### 3.1. Problem Definition

We first introduce the cross-modal matching task by taking image-text matching as an example. Then, we will present the noisy correspondence problem in cross-modal matching.

Traditionally, given a dataset  $\mathcal{D} = \{(I_i, T_i, y_i)\}_{i=1}^N$ , where  $N$  indicates the number of training samples,  $(I_i, T_i)$  is an image-text pair and  $y_i \in \{0, 1\}$  is the label. The binary label  $y_i$  is a hard-labeled correspondence score which indicates that the pair  $(I_i, T_i)$  is positively correlated ( $y_i = 1$ ) or not ( $y_i = 0$ ). The aim of cross-modal matching is to project the two modalities (visual and textual modalities in our case) into a shared feature space wherein positive data pairs have higher feature similarities and negative data pairs have lower feature similarities. Generally, the similarity of the given image-text pairs can be computed by  $S(f(I), g(T))$ , where  $S$  is a similarity measurement,  $f$  and  $g$  are two modal-specific feature extractors. In the following, we will denote the  $S(f(I), g(T))$  as  $S(I, T)$  for symbol simplicity. The feature extractors  $f$  and  $g$  can be learned

by minimizing the following triplet loss:

$$L_{hard}(I_i, T_i) = \left[ \alpha - S(I_i, T_i) + S(I_i, \hat{T}_h) \right]_+ + \left[ \alpha - S(I_i, T_i) + S(\hat{I}_h, T_i) \right]_+ \quad (1)$$

where  $(I_i, T_i)$  is a positive data pair,  $\alpha > 0$  denotes a given margin,  $[x]_+ = \max(x, 0)$ .  $\hat{I}_h$  and  $\hat{T}_h$  are hard negative samples [8] (the most similar negative samples in a mini-batch, i.e.,  $\hat{I}_h = \operatorname{argmax}_{I_j \neq I_i} S(I_j, T_i)$  and  $\hat{T}_h = \operatorname{argmax}_{T_j \neq T_i} S(I_i, T_j)$ ).

The solution in Eq. 1 is based on the assumption that all collected data pairs  $(I, T)$  are perfectly aligned. However, in practice, the multi-modal dataset is usually annotated by using cheap methods or harvested from the internet. Therefore, an unknown portion of data in the noisily-collected dataset  $\tilde{\mathcal{D}} = \{(I_i, T_i, \tilde{y}_i)\}_{i=1}^N$  is *mis-labeled*, which means that some data pairs  $(I, T)$  are *mis-matched* or *weakly-matched* but they are wrongly labeled as  $\tilde{y} = 1$ . Minimizing Eq. 1 on  $\tilde{\mathcal{D}}$  would result in a poorly-generalized cross-modal matching model since it would overfit to the noisy dataset and pulls close those negative samples.

### 3.2. Robust Matching Loss by Softening the Correspondence Label

As we have discussed in Section. 3.1, in a noisily-labeled dataset  $\tilde{\mathcal{D}}$ , the hard correspondence labels  $\tilde{y}$  are unreliable and cannot accurately reflect the degree of correlation between different modalities. Therefore, the noisy labels  $\tilde{y}$  need to be rectified to a more accurate estimation of the soft correspondence scores  $y^* \in [0, 1]$  between the modalities  $I$  and  $T$ . The rectified soft labels  $y^*$  are expected to be able to well-depict the correspondence degree between  $I$  and  $T$  (i.e.,  $y^*$  gradually grow from 0 to 1 as the correlation of  $I$  and  $T$  increases). Then, the rectified soft labels  $y^*$  can be recast as the soft margin to learn the shared feature space by:

$$L_{soft}(I_i, T_i) = \left[ \hat{\alpha}_i - S(I_i, T_i) + S(I_i, \hat{T}_h) \right]_+ + \left[ \hat{\alpha}_i - S(I_i, T_i) + S(\hat{I}_h, T_i) \right]_+ \quad (2)$$

where the  $\hat{\alpha}_i$  is the soft margin determined by  $y_i^*$ , i.e.,  $\hat{\alpha}_i = \frac{m y_i^* - 1}{m - 1} \alpha$  and  $m$  is a hyper-parameter [12],  $\hat{I}_h$  and  $\hat{T}_h$  are hard negative samples.

Now, the most challenging problem is: *how to estimate accurate soft correspondence labels  $y^*$  by only using the hard-labeled noisy dataset  $\tilde{\mathcal{D}} = \{(I_i, T_i, \tilde{y}_i)\}_{i=1}^N$* . NCR [12] proposes to leverage the network predictions to assign pseudo labels for every data pair. However, this would cause severe confirmation bias: Those confident but wrong predictions would be used to guide subsequent training, leading to a loop of self-reinforcing errors [4]. In this

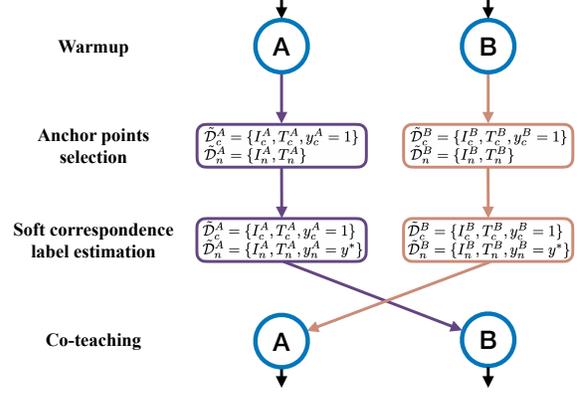


Figure 3. The overall training pipeline. We simultaneously train two models  $A$  and  $B$  in a co-teaching [11] manner to avoid the self-sample-selection error accumulation. At the inference stage, we average the predictions of  $A$  and  $B$ .

work, we propose to estimate the correspondence labels by leveraging the inherent similarity relationships between the two modalities, which gets rid of the confirmation bias problem.

### 3.3. Soft Correspondence Label Estimation

The key idea of our method is based on the assumption that *similar images should have similar textual descriptions and vice versa*. Assuming we have collected some clean positive data pairs (their soft correspondence labels  $y^*$  are 1), we can use these collected *anchor points* to infer the soft correspondence labels for the rest noisy data pairs, based on the aforementioned assumption. We will first show how to collect the anchor points out of a given noisy dataset, then we will present how to infer the soft correspondence labels for the noisy data pairs by using the anchor points.

#### 3.3.1 Anchor Points Selection by Modeling Per-sample Loss Value Distribution

We aim to identify those clean samples as anchor points in a noisy dataset  $\tilde{\mathcal{D}}$ . The memorization effect of deep neural networks [56] reveals that DNNs would first memorize training data of clean labels then those of noisy labels. This phenomenon indicates that noisy examples would have higher loss values while clean ones would have lower loss values during the early epochs of training. Therefore, given a matching model  $(f, g, S)$ , we can first compute the per-sample loss by:

$$\ell_{(f,g,S)} = \{\tilde{\ell}_i\}_{i=1}^N = \{L_{hard}(I_i, T_i)\}_{i=1}^N \quad (3)$$

Then, we can utilize the difference of the per-sample loss value distribution to identify clean data pairs. NCR [12] adopts a two-component Gaussian Mixture Model to fit the

per-sample loss value distribution:

$$p(\ell) = \sum_{k=1}^K \lambda_k p(\ell | k) \quad (4)$$

where  $\lambda_k$  and  $p(\ell|k)$  are the mixture coefficient and the probability density of the  $k$ -th component, respectively. However, we empirically found the clean set loss distribution is extremely skewed towards zero (see Fig. 4), making the Gaussian a poor approximation to the clean set loss distribution. As a more flexible distribution, the beta distribution has a better modeling ability for the symmetric and skewed distributions. We propose to fit a beta mixture model (BMM) to approximate the loss distribution for mixtures of clean and noisy samples (see Fig. 4). The beta distribution over a normalized loss  $\ell \in [0, 1]$  is defined to have pdf:

$$p(\ell | \gamma, \beta) = \frac{\Gamma(\gamma + \beta)}{\Gamma(\gamma)\Gamma(\beta)} \ell^{\gamma-1} (1 - \ell)^{\beta-1} \quad (5)$$

where  $\gamma, \beta > 0$  and  $\Gamma(\cdot)$  is the Gamma function. The mixture pdf can be formulated by substituting the Eq. 5 into Eq. 4. We use an Expectation Maximization (EM) procedure to fit the BMM to the observed per-sample loss value distribution. Then, we can compute the probability of the  $i$ -th data pair being clean or noisy as:

$$p(k | \ell_i) = \frac{p(k)p(\ell_i | k)}{p(\ell_i)} \quad (6)$$

where  $k = 0/1$  denotes the data pair  $(I_i, T_i, y_i)$  is clean/noisy. Now we can select a set of anchor points  $\tilde{D}_c = \{I_c, T_c, y_c = 1\}$  from the noisy dataset  $\tilde{D}$ :

$$\tilde{D}_c = \{(I_i, T_i, y_i = 1) | p(k = 0 | \ell_i) > \delta, \forall (I_i, T_i) \in \tilde{D}\} \quad (7)$$

where  $\delta$  is a threshold. Since the remained data  $\tilde{D}_n = \tilde{D} \setminus \tilde{D}_c$  is possibly noisy, we drop their labels as:

$$\tilde{D}_n = \{(I_i, T_i) | p(k = 0 | \ell_i) \leq \delta, \forall (I_i, T_i) \in \tilde{D}\} \quad (8)$$

In Section. 3.3.2, we will show how to estimate soft correspondence labels for  $\tilde{D}_n$  by using the anchor point samples in  $\tilde{D}_c$ . Then, the collected clean data and the noisy data with their estimated soft labels can be used for the model training.

However, training a model with high-confident (low loss) examples selected by the model itself would cause severe error accumulation problem, which is widely acknowledged in noisy label learning [11, 19, 54]. Similar to NCR [12], we also adopt the co-teaching [11] paradigm to alleviate the error accumulation problem. Specifically, we simultaneously train two networks  $A = \{f^A, g^A, S^A\}$  and  $B = \{f^B, g^B, S^B\}$  with the same architecture but different data sequences and initializations. At each training epoch,

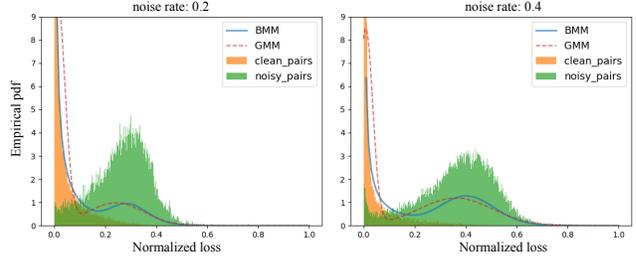


Figure 4. The per-sample loss distribution under 20% noisy labels and 40% noisy labels. Beta-Mixture-Model (BMM) performs better in modeling the skewed clean set loss distribution than Gaussian-Mixture-Model (GMM).

the network  $A$  will model the per-sample loss distribution on its own data batch sequence  $\tilde{D}^A$  and (1) select the anchor points  $\tilde{D}_c^A = \{I_c^A, T_c^A, y_c^A = 1\}$  according to Eq. 7, (2) estimate the soft correspondence labels for the noisy data  $\tilde{D}_n^A = \{I_n^A, T_n^A, y_n^A = y^*\}$ , then the data  $\tilde{D}_c^A \cup \tilde{D}_n^A$  is used for training the other network (*i.e.*,  $B$ ). Note that, before the co-teaching procedure, a warmup process is conducted to make the network  $A$  and  $B$  achieve an initial convergence. In order to reduce the impact of mismatched data pairs in the warmup process, we only select a certain ratio (*i.e.*, warmup selection ratio  $\alpha$ ) of the data pairs with smaller loss values for training in each batch. The training pipeline is illustrated in Fig. 3.

### 3.3.2 Soft Correspondence Label Estimation by Cross-modal Similarity Consistency

In this section, we proceed to show how to estimate accurate correspondence labels for  $\tilde{D}_n$  by using the anchor points  $\tilde{D}_c$  collected by Eq. 7. The key idea of our method is based on a rational assumption that *similar images should have similar textual descriptions and vice versa*. In other words, in an image-text shared feature space, if two images  $I_1$  and  $I_2$  are very close but their corresponding texts  $T_1$  and  $T_2$  are far away from each other, we can tell that at least one of these two data pairs (*i.e.*,  $(I_1, T_1)$  or  $(I_2, T_2)$ ) is mislabeled. Furthermore, if we assume the  $(I_1, T_1)$  is clean (well-matched), then the gap between  $D(f(I_2), g(I_1))$  and  $D(f(T_2), g(T_1))$  can reflect the degree of correlation between  $I_2$  and  $T_2$  to some extent, where  $D(\cdot, \cdot)$  is a distance function in feature space and we will write  $D(f(I), g(I))$  as  $D(I, I)$  in the following. That is to say, if the distances between  $(I_2, I_1)$  and  $(T_2, T_1)$  are more *consistent*, then  $I_2$  and  $T_2$  are more correlated.

Formally, for the  $i$ -th noisy data pair  $(I_n^i, T_n^i)$  in  $\tilde{D}_n$ , we first search its closest image  $I_c^\Delta$  in the collected anchor point set  $\tilde{D}_c$ , then we can compute its *image2text similarity consistency* by comparing the image feature distance  $D(I_n^i, I_c^\Delta)$

---

**Algorithm 1** The training pipeline of our robust cross-model matching framework.

---

**Input:** A noisily-labeled dataset  $\tilde{\mathcal{D}}$

- 1 **Required:** the clean probability threshold  $\delta$ , two individual matching models  $A$  and  $B$  with different initializations and data batch sequences
- 2 Warm up the model ( $A, B$ ) using  $L_{hard}$  Eq. 1.
- 3 **for**  $i = 1 : num\_epochs$  **do**
- 4     //Section. 3.3.1: modeling per-sample loss distribution using Eq. 6
- 5      $\mathcal{P}^A = \{p_i^A \mid p_i^A = p(k=0 \mid \ell_i)\}_{i=1}^N \leftarrow BetaMixtureModel(\tilde{\mathcal{D}}, B)$
- 6      $\mathcal{P}^B = \{p_i^B \mid p_i^B = p(k=0 \mid \ell_i)\}_{i=1}^N \leftarrow BetaMixtureModel(\tilde{\mathcal{D}}, A)$
- 7     **for**  $k \in \{A, B\}$  **do**
- 8          $\tilde{\mathcal{D}}_c^k = \{(I_i, T_i, y_i = 1) \mid p_i^k > \delta, \forall (I_i, T_i) \in \tilde{\mathcal{D}}\}$      //Section. 3.3.1: anchor points selection using Eq. 7
- 9          $\tilde{\mathcal{D}}_n^k = \{(I_i, T_i) \mid p_i^k \leq \delta, \forall (I_i, T_i) \in \tilde{\mathcal{D}}\}$      //Section. 3.3.1: noisy data selection using Eq. 8
- 10         **for**  $j = 1 : num\_steps$  **do**
- 11             Sample a mini-batch  $\{\mathcal{B}_j^c = (I_c, T_c, y_c = 1), \mathcal{B}_j^n = (I_n, T_n)\}$  from  $\{\tilde{\mathcal{D}}_c^k, \tilde{\mathcal{D}}_n^k\}$
- 12             //Section. 3.3.2: soft label estimation by bidirectional cross-modal similarity consistency
- 13             Keep the labels in  $\mathcal{B}_j^c$ , estimate soft correspondence labels  $y^*$  for the noisy data  $\mathcal{B}_j^n$  using Eq. 11
- 14             //Section. 3.2: optimize the soft matching loss on the estimated soft correspondence label
- 15             Train the network  $k$  on  $\{\mathcal{B}_j^c = (I_c, T_c, y_c = 1), \mathcal{B}_j^n = (I_n, T_n, y_n = y^*)\}$  by minimizing Eq. 2

**Output:** Matching models ( $A, B$ )

---

and their corresponding text feature distance  $D(T_n^i, T_c^\Delta)$ :

$$C_{i2t} = \frac{D(I_n^i, I_c^\Delta)}{D(T_n^i, T_c^\Delta)}, (I_n^i, T_n^i) \in \tilde{\mathcal{D}}_n, (I_c^\Delta, T_c^\Delta) \in \tilde{\mathcal{D}}_c \quad (9)$$

Similarly, we can compute its *text2image similarity consistency* by:

$$C_{t2i} = \frac{D(T_n^i, T_c^\Delta)}{D(I_n^i, I_c^\Delta)}, (I_n^i, T_n^i) \in \tilde{\mathcal{D}}_n, (I_c^\Delta, T_c^\Delta) \in \tilde{\mathcal{D}}_c \quad (10)$$

where  $T_c^\Delta$  is the closest text feature to  $T_n^i$  in the collected anchor point set  $\tilde{\mathcal{D}}_c$ . The estimated soft correspondence label  $y_i^*$  of  $(I_n^i, T_n^i)$  is finally formulated by the bidirectional cross-modal similarity consistency:

$$y_i^* = \left( \frac{D(I_n^i, I_c^\Delta)}{D(T_n^i, T_c^\Delta)} + \frac{D(T_n^i, T_c^\Delta)}{D(I_n^i, I_c^\Delta)} \right) / 2 \quad (11)$$

Then the noisy data with the estimated soft correspondence labels  $\tilde{\mathcal{D}}_n = \{I_n, T_n, y_n = y^*\}$  and the clean data  $\tilde{\mathcal{D}}_c = \{I_c, T_c, y_c = 1\}$  can be combined to train the matching model by minimizing the soft triplet loss in Eq. 2. The detailed training pipeline is illustrated in the Algorithm. 1.

## 4. Experiment

In this section, we evaluate the effectiveness of our proposed method *BiCro* on three image-text matching datasets, *i.e.*, Flickr30K [53], MS-COCO [25], and Conceptual Captions [36]. The Flickr30K [53] and MS-COCO [25] and two well-annotated datasets, we randomly corrupt their image-text data pairs for a specific percentage (*i.e.*, noise ratio) to

simulate the noisy correspondence issue. The Conceptual Captions [36] is with real noisy correspondence from the wild.

### 4.1. Datasets and Evaluation Metrics

The following three widely-used image-text matching datasets are used to evaluate our method and baselines:

**Flickr30K** [53] contains 31,000 images collected from the Flickr website and each image is associated with five captions. We use 1,000 images for model validation, 1,000 images for model testing, and 29,000 for model training.

**MS-COCO** [25] has 123,287 images with 5 captions each. Among them, 5,000 images are used for modal validation, 5,000 images are used for model testing, and 113,287 images are used for model training.

**Conceptual Captions** [36] is a large-scale image-text dataset with real-world noisy correspondence problem. It contains 3.3M images and each image is associated with one caption. All the data pairs in the Conceptual Captions dataset are automatically harvested from the Internet, therefore about 3%~20% image-text pairs in the dataset are mismatched or weakly-matched [36]. Following NCR [12], we use a subset of the Conceptual Captions dataset, *i.e.*, CC152K, in our experiments. 150,000 images in CC152K are used for model training, 1,000 images are used for model validation, and 1,000 images are used for model testing.

### 4.2. Implementation Details

As a general framework, *BiCro* can be applied to many existing cross-modal matching methods. Same as NCR [12]

Table 1. Image-Text Retrieval on Flickr30K and MS-COCO 1K.

Noise	Methods	Flickr30K							MS-COCO						
		Image→Text			Text→Image			Sum	Image→Text			Text→Image			Sum
		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	58.5	81.0	90.8	35.5	65.0	75.2	406.0	62.2	90.0	96.1	46.2	80.8	89.2	464.5
	VSRN	33.4	59.5	71.3	25.0	47.6	58.6	295.4	61.8	87.3	92.9	50.0	80.3	88.3	460.6
	IMRAM	22.7	54.0	67.8	16.6	41.8	54.1	257.0	69.9	93.6	97.4	55.9	84.4	89.6	490.8
	SAF	62.8	88.7	93.9	49.7	73.6	78.0	446.7	71.5	94.0	97.5	57.8	86.4	91.9	499.1
	SGR	55.9	81.5	88.9	40.2	66.8	75.3	408.6	25.7	58.8	75.1	23.5	58.9	75.1	317.1
	NCR	73.5	93.2	96.6	56.9	82.4	88.5	491.1	76.6	95.6	98.2	60.8	88.8	95.0	515.0
	DECL	77.5	93.8	97.0	56.1	81.8	88.5	494.7	77.5	95.9	98.4	61.7	89.3	95.4	518.2
	<b>BiCro</b>	<b>78.3</b>	94.1	97.3	60.0	83.7	89.5	502.9	78.2	95.9	98.4	62.5	89.8	95.5	520.3
	<b>BiCro*</b>	78.1	<b>94.4</b>	<b>97.5</b>	<b>60.4</b>	<b>84.4</b>	<b>89.9</b>	<b>504.7</b>	<b>78.8</b>	<b>96.1</b>	<b>98.6</b>	<b>63.7</b>	<b>90.3</b>	<b>95.7</b>	<b>523.2</b>
40%	SCAN	26.0	57.4	71.8	17.8	40.5	51.4	264.9	42.9	74.6	85.1	24.2	52.6	63.8	343.2
	VSRN	2.6	10.3	14.8	3.0	9.3	15.0	55.0	29.8	62.1	76.6	17.1	46.1	60.3	292.0
	IMRAM	5.3	25.4	37.6	5.0	13.5	19.6	106.4	51.8	82.4	90.9	38.4	70.3	78.9	412.7
	SAF	7.4	19.6	26.7	4.4	12.2	17.0	87.3	13.5	43.8	48.2	16.0	39.0	50.8	211.3
	SGR	4.1	16.6	24.1	4.1	13.2	19.7	81.8	1.3	3.7	6.3	0.5	2.5	4.1	18.4
	NCR	68.1	89.6	94.8	51.4	78.4	84.8	467.1	74.7	94.6	98.0	59.6	88.1	94.7	509.7
	DECL	72.7	92.3	95.4	53.4	79.4	86.4	479.6	75.6	95.5	98.3	59.5	88.3	94.8	512.0
	<b>BiCro</b>	73.6	<b>93.0</b>	<b>96.4</b>	<b>56.0</b>	80.8	87.4	487.2	76.4	95.2	<b>98.6</b>	61.5	<b>89.4</b>	<b>95.5</b>	516.6
	<b>BiCro*</b>	<b>74.6</b>	92.7	96.2	55.5	<b>81.1</b>	<b>87.4</b>	<b>487.5</b>	<b>77.0</b>	<b>95.9</b>	98.3	<b>61.8</b>	89.2	94.9	<b>517.1</b>
60%	SCAN	13.6	36.5	50.3	4.8	13.6	19.8	138.6	29.9	60.9	74.8	0.9	2.4	4.1	173.0
	VSRN	0.8	2.5	5.3	1.2	4.2	6.9	20.9	11.6	34.0	47.5	4.6	16.4	25.9	140.0
	IMRAM	1.5	8.9	17.4	1.9	5.0	7.8	42.5	18.2	51.6	68.0	17.9	43.6	54.6	253.9
	SAF	0.1	1.5	2.8	0.4	1.2	2.3	8.3	0.1	0.5	0.7	0.8	3.5	6.3	11.9
	SGR	1.5	6.6	9.6	0.3	2.3	4.2	24.5	0.1	0.6	1.0	0.1	0.5	1.1	3.4
	NCR	13.9	37.7	50.5	11.0	30.1	41.4	184.6	0.1	0.3	0.4	0.1	0.5	1.0	2.4
	DECL	65.2	88.4	94.0	46.8	74.0	82.2	450.6	73.0	94.2	<b>97.9</b>	57.0	86.6	93.8	502.5
	<b>BiCro</b>	<b>68.3</b>	90.4	93.8	<b>51.9</b>	76.9	84.4	465.7	<b>73.9</b>	<b>94.7</b>	97.7	<b>58.7</b>	87.0	93.8	<b>505.8</b>
	<b>BiCro*</b>	67.6	<b>90.8</b>	<b>94.4</b>	51.2	<b>77.6</b>	<b>84.7</b>	<b>466.3</b>	<b>73.9</b>	94.4	97.8	58.3	<b>87.2</b>	<b>93.9</b>	505.5

and DECL [31], we implement *BiCro* based on SGRAF [7], whose performance is the state-of-the-art in image-text matching. We follow the same training settings (*e.g.*, optimizer, network architecture, all the hyperparameters, etc.) as previous works [12,31] to make a fair comparison. Please refer to NCR [12] for more training details.

We first warmup the matching models *A* and *B* for 10 epochs to make them achieve an initial convergence. To reduce the effect of noisy data pairs, we select a small portion of data with small loss value in each batch to warmup the models based on a predefined warmup selection ratio  $\alpha$ . The choice of  $\alpha$  will be discussed in Section 4.5. At the training stage, the total number of iterations is 40 epochs, among them, the first 20 epochs are trained with clean samples screened by BMM, and the next 20 epochs were trained on all samples (with our estimated soft labels). At each training epoch, we select 10% data pairs with a higher probability of being clean in Eq. 7 as the selected anchor points. The remained data are regarded as noisy ones. At the inference stage, we average the similarities predicted by networks *A* and *B* for the retrieval evaluation. We propose two strategies to handle the noisy data pairs, (1) feed all of the data pairs with their estimated soft labels to the matching

model (denoted as *BiCro* in Table. 1 and Table. 2), (2) treat those data pairs with estimated soft labels under a threshold (mismatch threshold  $\theta$ ) as *mismatched data*, and set their correspondence labels as zero (denoted as *BiCro\** in Table. 1 and Table. 2). The effect of the mismatch threshold  $\theta$  will be discussed in Section 4.5.

### 4.3. Comparison with State-of-the-Arts

To demonstrate the effectiveness of the proposed *BiCro*, we evaluate the *BiCro* against several baseline methods, including general methods (SCAN [17], VSRN [20], IMRAM [3], SGRAF, SGR and SAF [7]) and robust learning method against noisy correspondence (NCR [12] and DECL [31]).

#### 4.3.1 Results on Simulated Noise

Table. 1 reports the experimental results on the 1K test images of Flickr30K dataset and over 5 folds of 1K test images of MS-COCO dataset. From the experimental results, we can find that the proposed *BiCro* performs significantly better than the state-of-the-art methods. In comparison to the best baseline DECL [31], *BiCro* improves the sum score for

Table 2. Image-Text Retrieval on CC152K.

Methods	Image→Text			Text→Image			Sum
	R@1	R@5	R@10	R@1	R@5	R@10	
SCAN	30.5	55.3	65.3	26.9	53.0	64.7	295.7
VSRN	32.6	61.3	70.5	32.5	59.4	70.4	326.7
IMRAM	33.1	57.6	68.1	29.0	56.8	67.4	312.0
SAF	31.7	59.3	68.2	31.9	59.0	67.9	318.0
SGR	11.3	29.7	39.6	13.1	30.1	41.6	165.4
NCR	39.5	64.5	73.5	40.3	64.6	73.2	355.6
DECL	39.0	66.1	75.5	40.7	66.3	76.7	364.3
<b>BiCro</b>	40.7	<b>67.3</b>	<b>76.7</b>	39.7	<b>67.6</b>	<b>76.9</b>	368.9
<b>BiCro*</b>	<b>40.8</b>	67.2	76.1	<b>42.1</b>	<b>67.6</b>	76.4	<b>370.2</b>

Table 3. Ablation studies on Flickr30K with 40% noise.

Methods				Image→Text			Text→Image		
Co-teaching	Soft label	Bmm	Warmup	R@1	R@5	R@10	R@1	R@5	R@10
✓	✓	✓	✓	<b>74.6</b>	<b>92.7</b>	<b>96.2</b>	<b>55.5</b>	<b>81.1</b>	<b>87.4</b>
	✓	✓	✓	69.9	91.9	95.7	52.0	79.0	85.5
✓		✓	✓	72.0	91.0	95.0	53.3	78.8	86.1
✓	✓		✓	72.0	92.3	95.5	55.2	79.8	86.6
✓	✓	✓		54.3	81.9	88.1	40.9	65.8	73.6

retrieving by 8.2%, 7.6%, 15.1%, 2.1%, 4.6% and 3.3% under different noise rates, respectively. In addition, BiCro\* further improves the overall performance by filtering out the mismatched pairs. The results at a noise rate of 60% show a strong robustness of our BiCro against the noisy correspondence with high noise rates. The failure of NCR to deal with high noise rates may be that it cannot divide the noisy and clean pairs well relying on GMM [12], while our method solves this drawback by using BMM.

### 4.3.2 Results on Real-world Noise

We evaluate the proposed method under the real noisy correspondence of CC152K. The experimental results are reported in Table 2. From the results, one could observe that our BiCro could achieve competitive performance under real noise. Specifically, BiCro is 4.6% higher than the best baseline in terms of sum in retrieval, respectively. Moreover, the performance gap between BiCro and BiCro\* shows that the filtering of data pairs according to soft correspondence labels can further reduce the impact of data mismatch issue on performance.

### 4.4. Hyperparameter Analysis

We now analyze the effect of the warmup selection ratio  $\epsilon$  and mismatch threshold  $\theta$ , which denote the ratio of data that participated in warmup and the threshold that sets the soft label to 0. We report the average performance of retrieving image and text with different  $\epsilon$  and  $\theta$  on Flickr30K with the noise ratio of 40%. As shown in Fig. 5, setting  $\epsilon = 0.3$  achieves the best overall accuracy for retrieving. Using all data ( $\epsilon = 1$ ) in warmup gives limited performance, which indicates it is necessary to select low loss data in warmup. On the other hand, with the increases of  $\theta$ , our BiCro\* improves the accuracy compared with using

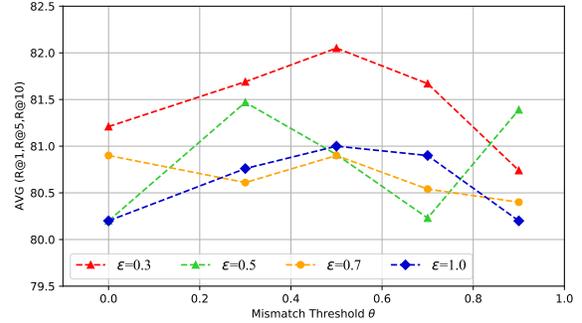


Figure 5. Variation of retrieval performance with different selection ratio  $\epsilon$  and mismatch threshold  $\theta$ . Note that  $\theta = 0$  denotes using the original generated soft labels for all data pairs.

the original generated soft labels ( $\theta = 0$ ). Moreover, further increasing the  $\theta$  leads to the performance drop since a too-large mismatch threshold  $\theta$  treats weakly labeled samples as mislabeled samples, resulting in the decline of generalization ability.

### 4.5. Ablation Study

In this section, we carry out the ablation study on the Flickr30K with a noise ratio of 40%. We show the effect of each component in Table 3. In order to prob the effectiveness of soft label estimation in BiCro, the labels of clean and noisy samples divided by co-teaching are set to 1 and 0 respectively as a comparison experiment to BiCro and the experimental results are shown in the third row. The fourth row shows the experiments of our method with GMM and BMM. From the results, we draw the following conclusions: 1) Since the quality of noise-robust cross-modal learning depends on the soft correspondence labels, the performance of adopting the soft labels estimated by BiCro surpasses the performance of using the original binary labels. 2) Although leveraging Gaussian-Mixture-Model achieves decent results, the performance can be further improved by Beta-Mixture-Model. 3) Co-teaching and Warmup, as important modules of the base network, are also crucial in our framework BiCro. 4) The model achieves the best test accuracy by utilizing all the components, which shows the complementarity of our proposed BiCro.

### 5. Conclusion

This paper focuses on the challenge of robust cross-modal matching on noisy data. To address this problem, We propose a general framework called *bidirectional cross-modal similarity consistency (BiCro)* for soft correspondence label estimation given only noisily-collected data pairs. The effectiveness of the proposed framework was verified on both synthetic noisy datasets and real noisy dataset. The visualization results also demonstrate that our estimated soft labels are the accurate estimation of the true correlation degree between different modality data.

## References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018. 1
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel E. O’Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *ICML*, pages 312–321, 2019. 3
- [3] 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. 7
- [4] Mingcai Chen, Hao Cheng, Yuntao Du, Ming Xu, Wenyu Jiang, and Chongjun Wang. Two wrongs don’t make a right: Combating confirmation bias in learning with label noise. *arXiv preprint arXiv:2112.02960*, 2021. 3, 4
- [5] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: universal image-text representation learning. In *ECCV*, pages 104–120, 2020. 3
- [6] 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. 3
- [7] Haiwen Diao, Ying Zhang, Lin Ma, and Huchuan Lu. Similarity reasoning and filtration for image-text matching. In *AAAI*, pages 1218–1226, 2021. 7
- [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, 4
- [9] Sheng Guo, Weilin Huang, Haozhi Zhang, Chenfan Zhuang, Dengke Dong, Matthew R Scott, and Dinglong Huang. Curriculumnet: Weakly supervised learning from large-scale web images. In *ECCV*, pages 135–150, 2018. 3
- [10] Bo Han, Jiangchao Yao, Gang Niu, Mingyuan Zhou, Ivor Tsang, Ya Zhang, and Masashi Sugiyama. Masking: A new perspective of noisy supervision. In *NeurIPS*, pages 5836–5846, 2018. 3
- [11] 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. 2, 3, 4, 5
- [12] 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. 2, 3, 4, 5, 6, 7, 8
- [13] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. MentorNet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, pages 2309–2318, 2018. 3
- [14] Qing-Yuan Jiang and Wu-Jun Li. Deep cross-modal hashing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3232–3240, 2017. 1
- [15] Kushal Kafle and Christopher Kanan. Visual question answering: Datasets, algorithms, and future challenges. *Computer Vision and Image Understanding*, 163:3–20, 2017. 1
- [16] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, pages 5583–5594, 2021. 3
- [17] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *ECCV*, 2018. 7
- [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. 1
- [19] Junnan Li, Richard Socher, and Steven C.H. Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *ICLR*, 2020. 5
- [20] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *ICCV*, pages 4653–4661, 2019. 7
- [21] Linghui Li, Sheng Tang, Lixi Deng, Yongdong Zhang, and Qi Tian. Image caption with global-local attention. In *Thirty-first AAAI conference on artificial intelligence*, 2017. 1
- [22] Mingchen Li, Mahdi Soltanolkotabi, and Samet Oymak. Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks. In *AISTATS*, 2020. 3
- [23] Sheng Li, Zhiqiang Tao, Kang Li, and Yun Fu. Visual to text: Survey of image and video captioning. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 3(4):297–312, 2019. 1, 3
- [24] Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Li-Jia Li. Learning from noisy labels with distillation. In *ICCV*, pages 1910–1918, 2017. 3
- [25] 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*, 2014. 6
- [26] Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*, 38(3):447–461, 2016. 3
- [27] Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah M. Erfani, and James Bailey. Normalized loss functions for deep learning with noisy labels. In *ICML*, 2020. 3
- [28] Xingjun Ma, Yisen Wang, Michael E Houle, Shuo Zhou, Sarah M Erfani, Shu-Tao Xia, Sudanthi Wijewickrema, and James Bailey. Dimensionality-driven learning with noisy labels. In *ICML*, pages 3361–3370, 2018. 3
- [29] Eran Malach and Shai Shalev-Shwartz. Decoupling” when to update” from” how to update”. In *NeurIPS*, pages 960–970, 2017. 3
- [30] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks

- robust to label noise: A loss correction approach. In *CVPR*, pages 1944–1952, 2017. 3
- [31] Yang Qin, Dezhong Peng, Xi Peng, Xu Wang, and Peng Hu. Deep evidential learning with noisy correspondence for cross-modal retrieval. In *ACMMM*, pages = 4948–4956, year = 2022. 3, 7
- [32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 1
- [33] Scott E. Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. In *ICLR, Workshop Track Proceedings*, 2015. 3
- [34] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, pages 4331–4340, 2018. 3
- [35] Nikolaos Sarafianos, Xiang Xu, and Ioannis A. Kakadiaris. Adversarial representation learning for text-to-image matching. In *ICCV*, pages 5813–5823, 2019. 2
- [36] 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*, 2018. 1, 2, 3, 6
- [37] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. SELFIE: refurbishing unclean samples for robust deep learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *ICML*, pages 5907–5915, 2019. 3
- [38] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *CVPR*, 2018. 3
- [39] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *CVPR*, pages 5552–5560, 2018. 3
- [40] Arash Vahdat. Toward robustness against label noise in training deep discriminative neural networks. In *NeurIPS*, pages 5596–5605, 2017. 3
- [41] Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, and Serge Belongie. Learning from noisy large-scale datasets with minimal supervision. In *CVPR*, pages 839–847, 2017. 3
- [42] Kai Wang, Xiangyu Peng, Shuo Yang, Jianfei Yang, Zheng Zhu, Xinchao Wang, and Yang You. Reliable label correction is a good booster when learning with extremely noisy labels. *arXiv preprint arXiv:2205.00186*, 2022. 3
- [43] Liwei Wang, Yin Li, Jing Huang, and Svetlana Lazebnik. Learning two-branch neural networks for image-text matching tasks. *TPAMI*, 41(2):394–407, 2019. 2
- [44] Songhua Wu, Xiaobo Xia, Tongliang Liu, Bo Han, Mingming Gong, Nannan Wang, Haifeng Liu, and Gang Niu. Class2simi: A new perspective on learning with label noise. *arXiv preprint arXiv:2006.07831*, 2020. 3
- [45] Xiaobo Xia, Tongliang Liu, Bo Han, Chen Gong, Nannan Wang, Zongyuan Ge, and Yi Chang. Robust early-learning: Hindering the memorization of noisy labels. In *ICLR*, 2021. 3
- [46] Xiaobo Xia, Tongliang Liu, Bo Han, Mingming Gong, Jun Yu, Gang Niu, and Masashi Sugiyama. Sample selection with uncertainty of losses for learning with noisy labels. *arXiv preprint arXiv:2106.00445*, 2021. 3
- [47] Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? In *NeurIPS*, 2019. 3
- [48] Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? In *NeurIPS*, pages 6838–6849, 2019. 3
- [49] Yilun Xu, Peng Cao, Yuqing Kong, and Yizhou Wang. L<sub>dmi</sub>: A novel information-theoretic loss function for training deep nets robust to label noise. In *NeurIPS*, pages 6222–6233, 2019. 3
- [50] Shuo Yang, Erkun Yang, Bo Han, Yang Liu, Min Xu, Gang Niu, and Tongliang Liu. Estimating instance-dependent label-noise transition matrix using dnns. *arXiv*, 2021. 3
- [51] Shuo Yang, Erkun Yang, Bo Han, Yang Liu, Min Xu, Gang Niu, and Tongliang Liu. Estimating instance-dependent bayes-label transition matrix using a deep neural network. In *International Conference on Machine Learning*, pages 25302–25312. PMLR, 2022. 3
- [52] Quanming Yao, Hansi Yang, Bo Han, Gang Niu, and James T Kwok. Searching to exploit memorization effect in learning with noisy labels. In *ICML*, 2020. 3
- [53] 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*, 2014. 3, 6
- [54] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor W Tsang, and Masashi Sugiyama. How does disagreement benefit co-teaching? In *ICML*, 2019. 3, 5
- [55] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2017. 2
- [56] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021. 4
- [57] Qi Zhang, Zhen Lei, Zhaoxiang Zhang, and Stan Z. Li. Context-aware attention network for image-text retrieval. In *CVPR*, pages 3533–3542, 2020. 2
- [58] Zhou Zhao, Qifan Yang, Deng Cai, Xiaofei He, Yueting Zhuang, Zhou Zhao, Qifan Yang, Deng Cai, Xiaofei He, and Yueting Zhuang. Video question answering via hierarchical spatio-temporal attention networks. In *IJCAI*, volume 2, page 8, 2017. 1
- [59] Songzhu Zheng, Pengxiang Wu, Aman Goswami, Mayank Goswami, Dimitris N. Metaxas, and Chao Chen. Error-bounded correction of noisy labels. In *ICML*, pages 11447–11457, 2020. 3