

Enhanced Multimodal Representation Learning with Cross-modal KD

Mengxi Chen¹, Linyu Xing¹, Yu Wang^{1,2}✉, Ya Zhang^{1,2}✉

¹Shanghai Jiao Tong University, ²Shanghai AI Laboratory

{mxchen_mc, xly8991, yuwanhsjtu, ya-zhang}@sjtu.edu.cn

Abstract

This paper explores the tasks of leveraging auxiliary modalities which are only available at training to enhance multimodal representation learning through cross-modal Knowledge Distillation (KD). The widely adopted mutual information maximization-based objective leads to a shortcut solution of the weak teacher, i.e., achieving the maximum mutual information by simply making the teacher model as weak as the student model. To prevent such a weak solution, we introduce an additional objective term, i.e., the mutual information between the teacher and the auxiliary modality model. Besides, to narrow down the information gap between the student and teacher, we further propose to minimize the conditional entropy of the teacher given the student. Novel training schemes based on contrastive learning and adversarial learning are designed to optimize the mutual information and the conditional entropy, respectively. Experimental results on three popular multimodal benchmark datasets have shown that the proposed method outperforms a range of state-of-the-art approaches for video recognition, video retrieval and emotion classification.

1. Introduction

Multimodal learning has shown much promise in a wide spectrum of applications, such as video understanding [20, 39, 44], sentiment analysis [1, 19, 47] and medical image segmentation [7, 49]. It is generally recognized that including more modalities helps the prediction accuracy. For many real-world applications, while one has to make predictions based on limited modalities due to efficiency or cost concerns, it is usually possible to collect additional modalities at training. To enhance the representation learned, several studies have thus considered to leverage such modalities as a form of auxiliary information at training [11, 23, 29, 50]. Several attempts have explored to fill in the auxiliary modalities at test time, by employing approaches such as Generative Adversarial Network

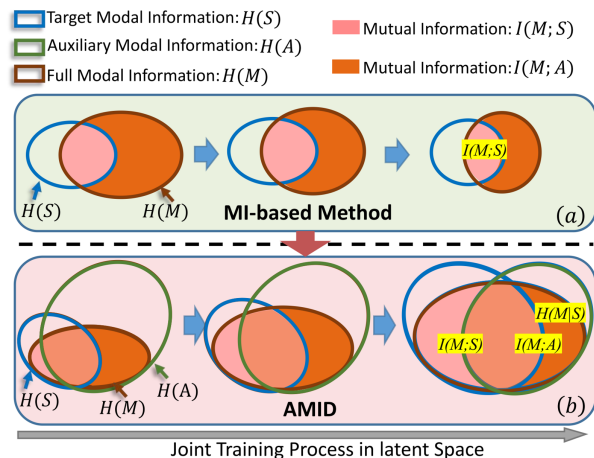


Figure 1. Illustration of the difference between the MI-based method and the proposed AMID method. (a) MI-based methods maximize $I(M; S)$. (b) AMID maximizes $I(M; S)$ and $I(M; A)$, while minimizing $H(M|S)$.

(GAN) [4, 22]. However, such data generation-based approaches introduce extra computation costs at inference.

As an alternative, cross-modal knowledge distillation is proposed, where a teacher model trained with both the target modalities and the auxiliary modalities (denoted as *full modalities* hereafter), is employed to guide the training of a student model with target modalities only [10, 11, 16, 23, 45, 46]. In such a way, the student model directly learns to embed information about auxiliary modalities and involves no additional computational cost at inference. A typical practice is to pre-train the teacher model from the full modality data, and fix it during the training. While such a setting enables the teacher to provide information-rich representation [23, 46], it may not be favorable for the alignment of the teacher and student, due to the large information gap between them, especially at the early stage of training [32].

To better facilitate the knowledge transfer, online KD is proposed to jointly train the teacher and the student. One common solution is to maximize the mutual information (MI) between the teacher and the student in a shared

space [42]. However, such an objective only focuses on maximizing the shared knowledge, in terms of percentage, from the teacher to the student, which is necessary but not sufficient to guarantee a complete knowledge transfer. As illustrated in Fig. 1(a), a short-cut solution is to simply make the teacher model capture less information, *i.e.*, shrinking the entropy of the teacher.

Under the framework online KD, this paper explores to simultaneously achieve: 1) maintaining the representation capacity of the full modality teacher as much as possible, and 2) narrowing down the information gap between the teacher and the student. As Fig. 1(b) shows, in addition to maximizing the MI between the teacher and the student, maximizing the MI between the teacher and an additional auxiliary model may contribute to the former goal, while minimizing the conditional entropy of the teacher given the student may benefit the latter goal. We thus propose a new objective function consisting of all the above three terms. To achieve the optimization of the MI, we further derive a new form of its lower bound by taking into account the supervision of the label, and propose a contrastive learning-based approach to implement it. To minimize the conditional entropy, an adversarial learning-based approach is introduced, where additional discriminators are used to distinguish representations between the teacher and the student. We name the proposed method Adversarial Mutual Information Distillation (AMID) hereafter.

To validate the effectiveness of AMID, extensive experiments on three popular multimodal tasks including video recognition, video retrieval and emotion classification tasks are conducted. The performance of these tasks are conducted on UCF51 [40], ActivityNet [20] and IEMOCAP [6] benchmark datasets, respectively. The results show that AMID outperforms a range of state-of-the-art approaches on all three tasks. To summarize, the main contribution of this paper is threefold:

- We propose a novel cross-modal KD method, Adversarial Mutual Information Distillation (AMID), to prevent the teacher from losing information by maximizing the MI between it and an auxiliary modality model.
- AMID maximally transfers information from the teacher to the student by optimizing the MI between them and minimizing the conditional entropy of the teacher given the student.
- To implement the joint optimization of the three objective terms, a novel training approach that leverages the combined advantage of both contrastive learning and adversarial learning is proposed.

2. Related Work

Knowledge Distillation. Knowledge distillation (KD) is

first proposed in [21], where soft labels in predictive space produced by the teacher through temperature hyperparameter transfer the dark knowledge to the student. Later, many works studied and demonstrated the benefits of distilling knowledge in the latent space [25, 35, 38, 48]. There have been some works that explore KD-based methods for cross-modal representation learning. CMC [41] aligns different views of the same samples by contrastive learning to extract shared information. CRD [42] leverages instance-level contrastive objectives to perform KD. GMC [36] learns geometrically aligned representations by aligning modality-specific representations with that of the corresponding complete observation. However, these methods can not transfer the correlation among instances, which are also valuable in many tasks. To transfer the inter-sample relation, [9] propose CCL to close the cross-modal semantic gap between the multimodality teacher and the student by contrasting different modal embeddings in the same classes in a common latent space. CRCD [51] distills the information of the teacher space by maximizing two complementary MI criteria. In contrast to these recent KD works which only consider the alignment between the teacher and student but neglect the objectives' impact on the teacher, we proposed a new KD objective that helps the teacher capture as much auxiliary information as possible and constraints the alignment simultaneously.

Mutual Information Maximization. MI is first introduced to deep learning in [2], and since then, many methods have proposed to combine variational bounds with parameterized neural networks to enable differentiable and tractable estimation of MI in high-dimensional spaces [15, 34]. The character that MI can be used to represent the relationship between two variables is well-suited for correlated multimodal data. [19] hierarchically optimizes MI in unimodal pairs and between multimodal and unimodal representations for maintaining useful information while filtering noise in unimodal input. [28] proposed TupleInfoNCE to optimize MI in multimodal input tuples which enables the model to capture both the modal-shared information and the specific information. However, these methods rarely consider the correlation among samples with the supervision of the label, which are still useful. In this paper, we derive a new lower bound on MI in teacher-student pairs which takes into account the inter-sample correlation and propose a contrastive-learning method for its maximization.

Adversarial Learning. Since [18] developed adversarial learning in deep learning, many researchers attempt to improve its performance by stabilizing the training process. In addition to modifying the structure of discriminator or generator [8, 31, 37], several works still considered incorporating multiple discriminators to improve learning [13, 14, 33]. Similar to [13], in this paper, we design two additional discriminators and adjust their weights adaptively.

3. Method

Suppose we have a multimodal dataset with n modalities, consisting of k target modalities $\{x_1, \dots, x_k\}$ and $n-k$ auxiliary modalities $\{x_{k+1}, \dots, x_n\}$. Denote the model trained with all n modalities, k target modalities only, and $n-k$ auxiliary modalities only as full modality model $M(x_1, \dots, x_n)$, target modality model $S(x_1, \dots, x_k)$, and auxiliary modality model $A(x_{k+1}, \dots, x_n)$, respectively. This paper aims to transfer knowledge from the full modality model M to the target modality model S with the assistance of auxiliary modalities at training.

A widely adopted objective for cross-modal knowledge transfer is to maximize the MI between the representation of the full modality teacher M and the representation of the target modality student S , *i.e.* $I(M; S)$. But there exists a short-cut solution as shown in Fig. 1(a). In order to prevent this problem, we propose AMID, which maximizes the MI between M and A , *i.e.* $I(M; A)$, so as to ensure the full modality model captures information from both target modalities and auxiliary modalities, and maximizes $I(M; S)$ while minimizing the conditional entropy of the teacher given the student $H(M|S)$, so as to narrow down the information gap between the teacher and the student. For optimizing the MI, we derive a new form of its lower bound and maximize it through a contrastive learning-based approach. For minimizing $H(M|S)$, we introduce an adversarial learning-based approach.

To sum up, the overall objective is to optimize:

$$\mathcal{O} = \alpha_1 \text{lb}(I(M; S)) + \alpha_2 \text{lb}(I(M; A)) - \alpha_3 H(M|S), \quad (1)$$

where $\text{lb}(\cdot)$ means the lower bound on MI, $\alpha_1, \alpha_2, \alpha_3$ are the different weights for the corresponding terms.

3.1. Mutual Information Optimization

In this section, we will focus on how to optimize $I(M; S)$ and $I(M; A)$ through maximizing their lower bounds for capturing as much information as possible. Unlike [42] which only considers the correlation between the teacher's representation and that of the student in the same sample, we derive a new form of lower bound on MI by taking into account the correlation of samples within a class. Notice that as the deduction procedures for $I(M; S)$ and $I(M; A)$ are the same, we will construct an approximate loss for optimizing $I(M; S)$.

Firstly, we define a distribution q which models the joint distribution of M and S . A latent variable, C , is used to indicate whether a tuple (M, S) comes from the same class ($C = 1$) or not ($C = 0$):

$$\begin{aligned} q(M, S|C = 1) &= p(M, S); \\ q(M, S|C = 0) &= p(M)p(S). \end{aligned} \quad (2)$$

We sample (M, S) by considering the relationship between full modality and target modality representations in the

same classes, which is more effective for transferring information. The pair (M, S) is from the joint distribution when M, S are in the same class or from the product of the marginal distributions when M, S are in different classes.

Suppose in a mini-batch, given a specific M_i , there are $k-1$ positive pairs $(M_i, S_i^{\eta_j})$ drawn from the joint distribution, where $\{\eta_j\}_{j=1}^{k-1}$ is an index set and every element indicates the position of other S with the same class as M_i , 1 positive pair $(M_i, S_i^{\eta_0})$ from the same sample, and N negative pairs $(M_i, S_i^{\xi_l})$ from the product of marginals, where $\{S_i^{\xi_l}\}_{l=1}^N$ are N target modality representations from different classes than M_i . Then the priors of C can be calculated as $q(C = 1) = \frac{k}{N+k}$ and $q(C = 0) = \frac{N}{N+k}$. Using Bayes' rule, the logarithmic posterior for $C = 1$ is given by:

$$\begin{aligned} \log q(C = 1|M, S) &= \log \frac{q(M, S|C = 1)q(C = 1)}{\sum_{c=0,1} q(M, S|C = c)q(C = c)} \\ &= -\log\left(1 + \frac{N}{k} \frac{p(M)p(S)}{p(M, S)}\right) \\ &\leq -\log\left(\frac{N}{k}\right) + \log \frac{p(M, S)}{p(M)p(S)}. \end{aligned} \quad (3)$$

With Eq. (3), we can get a lower bound on $I(M; S)$ as:

$$\begin{aligned} I(M; S) &= \mathbb{E}_{(M, S) \sim p(M, S)} \log \frac{p(M, S)}{p(M)p(S)} \\ &\geq \log\left(\frac{N}{k}\right) + \mathbb{E}_{q(M, S|C=1)} \log q(C = 1|M, S). \end{aligned} \quad (4)$$

Thus maximizing $\mathbb{E}_{q(M, S|C=1)} \log q(C = 1|M, S)$ is to maximize the lower bound on $I(M; S)$. However, the true distribution $q(C = 1|M, S)$ cannot be calculated directly. A classification model $h : \{M, S\} \rightarrow [0, 1]$ is fitted to approximate this distribution. The log-likelihood of the sampled tuples with model h is defined as:

$$\begin{aligned} \mathcal{I}(h) &= k \mathbb{E}_{q(M, S|C=1)} [\log h(M, S)] \\ &\quad + N \mathbb{E}_{q(M, S|C=0)} [\log (1 - h(M, S))]. \end{aligned} \quad (5)$$

This likelihood is maximized to make h a good approximation to $q(C = 1|M, S)$. Thus, we can get the optimal $h^*(M, S) = q(C = 1|M, S)$. Because that both $\log h(M, S)$ and $\log (1 - h(M, S))$ are non-positive for any h , we can rewrite Eq. (4) according to Eq. (5):

$$\begin{aligned} I(M; S) &\geq \log\left(\frac{N}{k}\right) + \mathbb{E}_{q(M, S|C=1)} [\log h^*(M, S)] \\ &\geq \log\left(\frac{N}{k}\right) + k \mathbb{E}_{q(M, S|C=1)} [\log h^*(M, S)] \\ &\quad + N \mathbb{E}_{q(M, S|C=0)} [\log (1 - h^*(M, S))] \\ &= \log\left(\frac{N}{k}\right) + \max_h \mathcal{I}(h) \geq \log\left(\frac{N}{k}\right) + \mathcal{I}(h) \end{aligned} \quad (6)$$

Thus we can optimize $I(M; S)$ indirectly by maximizing this lower bound $\log\left(\frac{N}{k}\right) + \mathcal{I}(h)$. Inspired by [9], we




Figure 2. An overview of the proposed AMID. There are three different models, M , S and A , in AMID for the full modality, target modality and auxiliary modality, respectively. AMID distills multimodal information to the target modality by maximizing the lower bound on $I(M; S)$ and minimizing $H(M|S)$, and prevents the shout-cut solution by maximizing the lower bound on $I(M; A)$. MI optimization is a contrastive-based approach, and conditional entropy minimization is an adversarial learning-based approach that contains two discriminators with different inputs.

choose h as:

$$h(M_i, S_u) = \frac{\exp(\Phi(M_i, S_u)/\tau)}{\sum_{j=1}^{N+k} \exp(\Phi(M_i, S_j)/\tau)}, \quad (7)$$

where Φ is a cosine similarity function, τ is a temperature adjusting the correlation of different pairs. Substitute Eq. (7) into Eq. (6), optimizing $I(M; S)$ is equivalent to minimizing the following loss:

$$\begin{aligned} \mathcal{L}_{MI_s} = & -\frac{1}{N+k} \sum_{i=1}^{N+k} \left(\sum_{j=0}^{k-1} \log h(M_i, S_i^{\eta_j}) \right. \\ & \left. + \sum_{l=1}^N \log \left(1 - h(M_i, S_i^{\xi_l}) \right) \right). \end{aligned} \quad (8)$$

Similarly, we can optimize the $I(M; A)$ by minimizing the following loss:

$$\begin{aligned} \mathcal{L}_{MI_a} = & -\frac{1}{N+k} \sum_{i=1}^{N+k} \left(\sum_{j=0}^{k-1} \log h(M_i, A_i^{\eta_j}) \right. \\ & \left. + \sum_{l=1}^N \log \left(1 - h(M_i, A_i^{\xi_l}) \right) \right). \end{aligned} \quad (9)$$

The form of Eq. (8) and Eq. (9) are similar to the CCL loss in [9]. However, our formulation is associated with MI and derives a tighter lower bound on it by taking account into the supervision of the label, which in our experiments is found to be more effective.

3.2. Conditional Entropy Minimization

Maximizing the lower bound on $I(M; S)$ is insufficient for narrowing the information gap between the teacher and the student as the optimal lower bound does not mean that MI is maximum. This subsection focuses on further transferring knowledge by minimizing the conditional entropy of the teacher given the student $H(M|S)$.

The posterior $p(M|S)$ is difficult to calculate, thus we define a new distribution $q(M|S)$ to approximate it. Then $H(M|S)$ can be calculated as in [8]:

$$\begin{aligned} H(M|S) &= -\mathbb{E}_{S \sim p(S)} [\mathbb{E}_{M' \sim p(M|S)} [\log p(M'|S)]] \\ &= -\mathbb{E}_{S \sim p(S)} [\mathbb{E}_{M' \sim p(M|S)} [\log q(M'|S)]] \\ &\quad - D_{KL}(p(M|S) || q(M|S)) \\ &\leq -\mathbb{E}_{S \sim p(S)} [\mathbb{E}_{M' \sim p(M|S)} [\log q(M'|S)]] . \end{aligned} \quad (10)$$

Thus we can minimize $H(M|S)$ by maximizing the negative of its upper bound. We achieve this by introducing an adversarial learning-based approach, and using it to maximize $\mathbb{E}_{S \sim p(S)} [\mathbb{E}_{M' \sim p(M|S)} [\log q(M'|S)]]$ w.r.t. the parameters of the student network during the generation and minimize $\mathbb{E}_{S \sim p(S)} [\mathbb{E}_{M' \sim p(M|S)} [\log q(M'|S)]] - \log p(M|S)$ w.r.t. the parameters of the discriminators. In other words, we use the discriminator $D(\cdot)$ to recognize whether features come from multimodal distribution or not and feedback informative gradient to the student model through minimizing:

$$\mathcal{L}_{adv} = -\mathbb{E}_{S \sim p(S)} [\mathbb{E}_{M' \sim p(M|S)} [\log D(S)]] . \quad (11)$$

Minimizing Eq. (11) w.r.t. the parameters of the target modality student enables the student to imitate representations coming from the multimodal distribution. It is worth noting that optimizing $H(M|S)$ through Eq. (11) does not affect the learning of the teacher since no gradient back propagates to the teacher.

Nevertheless, for the samples of the same class, it is more challenging to transfer the multimodal information by adversarial learning. This is because the generated target modality representation tends to mimic that of multimodality in the same instance with the constraint of MI, easier but can not capture the correlation among samples with the supervision of the label. As a result, we introduce two discriminators to learn more informative target modal representations. One discriminator called fundamental discriminator (D_1) is designed to distinguish the target modality and full modality representations from the entire training data. The other discriminator, which is called enhanced dis-

criminator (D_2), is designed to further distinguish the two kinds of representations coming from the classes with more than one sample in a batch since we think paying attention to these representations can help the student further aggregate representations in the same class. Thus, minimizing Eq. (11) w.r.t. the parameters of the student becomes minimizing:

$$\mathcal{L}_{adv} = -\frac{\lambda_1}{N+k} \sum_{i=1}^{N+k} \log D_1(S_i^{\eta_0}) - \frac{\lambda_2}{d} \sum_{v=1}^w \log D_2(S_v^{\eta_1}). \quad (12)$$

And the parameters of discriminators are trained by minimizing:

$$\begin{aligned} \mathcal{L}_D = & -\frac{\lambda_1}{N+k} \sum_{i=1}^{N+k} (\log D_1(M_i) + \log(1 - D_1(S_i^{\eta_0}))) \\ & - \frac{\lambda_2}{d} \sum_{v=1}^w (\log D_2(M_v) + \log(1 - D_2(S_v^{\eta_1}))), \end{aligned} \quad (13)$$

where $\{M_v\}_{v=1}^w$ is a set where every element comes from the classes with more than one samples and $\{S_v^{\eta_1}\}_{v=1}^w$ is its corresponding set of target modal representations. d is used for more stabilized training and we set it by the frequency of the data with samples of the same classes. λ_1 and λ_2 are weighting parameters for D_1 and D_2 respectively, and $\lambda_1 + \lambda_2 = 1$. To further improve the efficiency of the transfer, we adjust λ_1 and λ_2 adaptively according to the alignment between S and M for each epoch t . More specifically, we use the averaged cosine similarity in the tuples $\Phi(M_i, S_i^{\eta_0})$ and $\Phi(M_v, S_v^{\eta_1})$, $j \neq 0$ in Eq. (7) to provide meaningful feedback for the cooperation of discriminators. Large cosine similarity indicates that M and S are well aligned, thus the weight of the corresponding discriminator can be relatively small. Moving average [30] is also applied to improve the training stability. The formula of λ is given as:

$$\begin{aligned} \tilde{\lambda}(t+1) &= \beta \tilde{\lambda}(t) + (1-\beta)(1 - \text{Avg}_t(\Phi(M, S))) \\ \lambda(t+1) &= \sigma(\tilde{\lambda}(t+1)), \end{aligned} \quad (14)$$

where $\sigma(\cdot)$ is the softmax function, $\text{Avg}_t(\cdot)$ represents the average of cosine similarity on epoch t , and $\beta \in (0, 1)$ is the smoothing parameter, and we set $\beta = 0.9$ empirically.

Adversarial learning is known to often generate random images if not constrained [31]. This also occurs when target modality representations are generated without any constraints. Our method implicitly overcomes this problem by maximizing the lower bound on $I(M; S)$ and its advantages come twofold: First, the maximization of the lower bound on MI can not only transfer the multimodal information but also place a constraint on the adversarial distillation. Second, adversarial distillation further reduces the information gap that still exists after maximizing the lower bound on MI. Thus, the maximization of the lower bound on MI and adversarial distillation are complementary to each other in AMID to generate informative and robust representations.

3.3. Overall Loss

Assuming that after the optimization above, the target modality and full modality representations are already aligned, then we can use two classifiers with shared parameters to perform the classification task for these representations respectively. The overall classification loss is:

$$\mathcal{L}_{cls} = \mathcal{L}_{ce}^S(S_i, y_i) + \mathcal{L}_{ce}^M(M_i, y_i), \quad (15)$$

where \mathcal{L}_{ce} is the cross-entropy loss, y_i is the ground-truth. Classifiers with the same parameters can also help facilitate the generation of representations from the target modality network to imitate those of multimodality. In addition, the target modality can be aligned with the full modality in the predictive space. We use the Jensen–Shannon divergence (JSD) [27] to optimize the MI between the predictive target modality distribution (P_S) and the multimodality distribution (P_M):

$$\mathcal{L}_{JSD} = \frac{1}{N+k} \sum_{i=1}^{N+k} \text{JSD}(P_{S_i} || P_{M_i}). \quad (16)$$

The final loss function for distilling knowledge from auxiliary modalities is given as:

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{JSD} + \alpha_1 \mathcal{L}_{MI_s} + \alpha_2 \mathcal{L}_{MI_a} + \alpha_3 \mathcal{L}_{adv}, \quad (17)$$

The whole procedure is shown in the supplementary and the overall model structure is shown in Fig. 2.

4. Experiments

4.1. Datasets

Video understanding. We experiment on two popular multimodal datasets for video recognition (Top-1 accuracy, %) and kNN video retrieval (R@K, %) [5]. Cross-modal knowledge distillation is performed from the acoustic to the visual modality.

- **UCF51**, as a subset of UCF101 [40], includes 6,845 videos with audio belonging to 51 categories. The public split1 is adopted.
- **ActivityNet** [20] is a large-scale video benchmark of 14,950 videos with audio, covering 200 complex human activities, where the training set has 10,024 videos while 4,926 videos in the validation set.

Emotion classification. We conduct experiments on IEMOCAP [6], a popular multimodal dataset for 4-class emotion classification, which contains totaling 5531 videos in 5 dyadic conversation sessions. Each video consists of acoustic, textual and visual modalities, and the target model is based on visual modality only. A 5-fold cross-validation is adopted for the evaluation, taking one session as the test

Table 1. The Top-1 accuracy (%) of video recognition on UCF51 and ActivityNet. Notice that we change all the activation functions to Leaky ReLU to ensure a fair comparison.

	UCF51	ActivityNet
Baseline	66.7	48.8
CRD [42]	67.1	50.4
GMC [36]	64.1	53.1
CRCD [51]	69.6	52.4
CCL [9]	68.9	51.9
AMID (Ours)	73.8	53.6

set for each fold. Following [50], two evaluation metrics, weighted accuracy (WA, %) and unweighted accuracy (UA, %), are adopted.

4.2. Experimental Setup

Because different modalities such as audio, video and text are heterogeneous, we use individual extractors for each modality, then the features are concatenated and passed through a 2-layer MLP as the full modality teacher. The teacher shares the same extractors as that of the student and the auxiliary model for reducing the number of parameters. **Video recognition and retrieval.** Following [9], we use a 1D-CNN14 [26] pre-trained on AudioSet [17] as the audio extractor and use R(2+1)D-18 [43] as the student video network. The two discriminators are 5-layer MLP and 3-layer MLP respectively. Besides, we change all activation functions to Leaky ReLU for adversarial learning because it gives better training stability compared to the ReLU. The dimension of the representations is set to 512. The batch size and temperature τ are set to 16 and 0.5, respectively. $\alpha_1, \alpha_2, \alpha_3$ are set to 4, 1.6, 1. More details are shown in the supplementary.

Emotion classification. We extract raw acoustic, textual and visual embeddings by the pre-trained wav2vec 2.0 [3], BERT [12] and DenseNet [24] respectively. RNN is adopted as the student network with a final feature dimension of 768. The batch size and temperature τ are set to 128 and 0.6, respectively.

A warm-up stage is employed for all experiments. At the first T_{start} epochs, the student model and the teacher model are trained by minimizing $\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{JSD}$ to alleviate the converge difficulties of adversarial learning. This helps a preliminary alignment between the teacher and student and can also prevent discriminators from recognizing readily at the beginning of the training. T_{start} is set to 20 for the video understanding tasks and 5 for the emotion classification task. And notice that without special mention, the leveraged pre-trained auxiliary modality extractors are fixed since the generated representations already contain rich auxiliary information.

Table 2. R@K (K = 1, 5) of video retrieval on UCF51 and ActivityNet.

Metric	UCF51		ActivityNet	
	R1	R5	R1	R5
Baseline	64.0	71.9	41.5	63.9
CRD [42]	66.0	72.3	43.6	65.5
GMC [36]	61.8	71.0	46.1	67.7
CRCD [51]	68.2	75.2	45.7	67.0
CCL [9]	67.0	71.7	46.2	64.8
AMID (Ours)	73.6	77.1	46.4	67.2




Figure 3. The per-class Top-1 accuracy (%) of video recognition on UCF51 comparing AMID with CCL.

4.3. Comparison with State-of-the-arts

Video recognition. We compare AMID with two types of SOTA cross-modal KD methods, *i.e.* intra-sample methods including CRD [42], GMC [36], and inter-sample methods including CCL [9], CRCD [51]. The results are summarized in Tab. 1. Here "Baseline" represents the model directly trained with the visual modality only, representing a lower bound for the tasks. It can be seen that AMID significantly outperforms all the compared methods, suggesting that it is more effective and consistent in learning information from auxiliary modalities. AMID outperforms the second-best method by 4.2% and 0.5% on the two datasets, respectively.

To further investigate the breakdown of performance, we compare AMID with CCL on UCF51 and give the results of class-wise Top-1 accuracy in Fig. 3. It can be seen that AMID achieves better performance than CCL for most of the classes, although for a few of them it is comparable to or slightly worse than CCL. The classes where AMID is of inferior performance are mainly those audio events that are not or weakly correlated with the videos, such as 'FrisbeeCatch' and 'SkyDiving'. In these classes, the auxiliary modal information is less helpful. Thus, AMID, which is proposed for better cross-modal knowledge distillation, is not very effective for these classes.

Video retrieval. Tab. 2 presents the results of video retrieval on the UCF51 and ActivityNet datasets, which is a

Table 3. Performance comparison under different setups on IEMOCAP, where A-V, L-V, AL-V mean distillation from acoustic to visual, textual to visual, and acoustic and textual to visual.

Setup	Method	WA	UA
V	Baseline	51.2	50.8
A-V	GMC [36]	51.7	50.9
	CRCD [51]	51.8	50.2
	CCL [9]	51.5	51.4
	AMID (Ours)	52.5	51.7
L-V	GMC [36]	50.9	50.5
	CRCD [51]	50.2	49.9
	CCL [9]	52.7	51.8
	AMID (Ours)	53.0	52.4
AL-V	GMC [36]	51.9	50.5
	CRCD [51]	52.0	49.4
	CCL [9]	53.2	51.4
	AMID (Ours)	53.8	52.7

more fine-grained task to evaluate the discriminability of video representations. The results show that AMID consistently outperforms almost all competing approaches in terms of R@1 and R@5.

Emotion classification. Tab. 3 compares AMID with three state-of-the-art methods, CCL, CRCD and GMC, for emotion classification. We conduct experiences on IEMOCAP, a three-modality dataset, to further evaluate the performance of AMID in different auxiliary modalities. It is worth noting that the "Baseline" model is trained directly with only the visual modality. The results demonstrate that AMID outperforms the Baseline model for all setups, where the UA metrics increase by 0.9%, 1.6% and 1.9% respectively, while the other three methods are not always effective. Besides, AMID outperforms CCL, CRCD and GMC on almost all the setups.

4.4. Ablation Study

4.4.1 Effectiveness of MI and conditional entropy

In order to verify the effectiveness of our proposed method for capturing as much multimodal information as possible, the impact of using different training configurations on UCF51 is studied in Tab. 4. Here "student alone" represents the result of a model trained with the visual modality only and "teacher alone" represents that of the teacher model with the full modalities. These two cases show the initial performance of the student and the teacher without distillation, respectively. When optimizing \mathcal{L}_{MI_s} alone for the distillation, although the student model can achieve performance gains, from 66.7% to 69.8%, it degrades the teacher's performance, from 74.3% to 71.2% and thereby reducing the upper bound of the student's performance. Alternatively, more informative multimodal representations can be obtained by maximizing the lower bound on $I(M; A)$ and

Table 4. Ablation study for the proposed loss functions.

Configuration	Top-1 accuracy	
	Student	Teacher
student alone (w/o distillation)	66.7	\
teacher alone (w/o distillation)	\	74.3
\mathcal{L}_{MI_s}	69.8	71.2
$\mathcal{L}_{MI_s} + \mathcal{L}_{MI_a}$	72.7	76.0
$\mathcal{L}_{MI_s} + \mathcal{L}_{MI_a} + \mathcal{L}_{adv}$ (AMID)	73.8	75.7
AMID with \mathcal{L}_{nce}	72.3	74.3

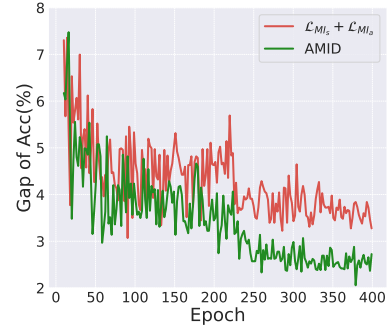


Figure 4. The accuracy gap between the teacher and student in using $\mathcal{L}_{MI_s} + \mathcal{L}_{MI_a}$ and AMID loss.

$I(M; S)$ simultaneously ($\mathcal{L}_{MI_s} + \mathcal{L}_{MI_a}$), since a better teacher can improve the distillation performance. Tab. 4 shows that this improves the accuracy of the student from 66.7% to 72.7% and also improves that of the teacher from 74.3% to 76.0%. Besides, adding \mathcal{L}_{adv} (*i.e.*, AMID) to $\mathcal{L}_{MI_s} + \mathcal{L}_{MI_a}$ does improve the performance from 72.7% to 73.8%, and the performance gap of AMID is smaller than $\mathcal{L}_{MI_s} + \mathcal{L}_{MI_a}$, which indicates that AMID achieves better alignment between the teacher and the student due to the further constraint on the gap by minimizing $H(M|S)$ without degrading the teacher's performance obviously (the teacher's performance decreases from 76.0% to 75.7%).

To further understand the effectiveness of the conditional entropy minimization approach in reducing the information gap between the teacher and the student, we explore the gap of AMID and $\mathcal{L}_{MI_s} + \mathcal{L}_{MI_a}$ during training. This information gap is represented by the gap of accuracy between the teacher and student on the validation set over the training epochs. The results are shown in Fig. 4. It can be seen that the accuracy gap of using AMID loss is consistently smaller than that of $\mathcal{L}_{MI_s} + \mathcal{L}_{MI_a}$, which validates the proposed adversarial learning-based approach indeed helps to narrow down the information gap.

AMID is also compared to an ablative setup using the contrastive loss \mathcal{L}_{nce} [9] to verify the effectiveness of the proposed loss function in maximizing the lower bound on MI. The result in Tab. 4 shows that AMID performs much better than the baseline "AMID with \mathcal{L}_{nce} ", improving the accuracy by 1.5%. This validates that our proposed ap-

Table 5. Effectiveness of proposed training modifications.

Modification	Top-1 accuracy
AMID	73.8
with fixed teacher	66.0
w/o warm-up	65.8
w/o D_2	72.8
w/o dynamic weight	72.7




Figure 5. (Left) The cosine similarity between the teacher’s and student’s representations only in the same class during training. (Right) The cosine similarity during testing.

proach, which takes into account the correlation of samples within a class, achieves better optimization for MI.

4.4.2 Effectiveness of training modifications

To evaluate the effect of each modification applied in AMID, we conduct a set of experiments on UCF51 and give the results in Tab. 5. As mentioned before, using a fixed pre-trained teacher to guide the student’s learning leads to an undesired gap between the teacher and the student at the beginning of training, which is unfavorable for distillation. The result affirms this, where the performance of AMID degrades from 73.8% to 66.0% without joint training. In addition, it also shows that without employing the warm-up setup described before, the recognition performance degrades by 8.0%. The reason may be that the target modality and multimodality representations are noisy and uncorrelated at the start of the training, which results in that the discriminators can recognize them readily and thus affecting the feedback of the informative gradient of adversarial learning.

We also study the effectiveness of the two discriminators of our proposed AMID by comparing AMID with two ablative cases: w/o D_2 and w/o dynamic weight. One constraint is removed at a time. The results of these ablation studies are shown in Tab. 5. The accuracy decreases by 1.0% when only the fundamental discriminator D_1 is used, suggesting that the enhanced discriminator D_2 indeed helps to distill the multimodal information of samples of the same class. Moreover, it also demonstrates that using uniform weight leads to a performance drop, indicating the effectiveness of the proposed dynamic weighting scheme.

To further investigate the effectiveness of the proposed adversarial learning-based approach in transferring the

Table 6. Top-1 accuracy (%) on UCF51 with different hyperparameters.

α_1	value	3.7	3.8	3.9	4.0	4.1	4.2
	accuracy	72.6	73.7	73.7	73.8	73.8	73.4
α_2	value	1.3	1.4	1.5	1.6	1.7	1.8
	accuracy	71.2	73.7	73.8	73.8	73.6	72.4
α_3	value	0.7	0.8	0.9	1.0	1.1	1.2
	accuracy	73.3	73.4	73.7	73.8	73.4	73.1

inter-sample multimodal information, we compare AMID with two cases: w/o D_2 and w/o \mathcal{L}_{adv} . Fig. 5 displays the average of cosine similarity between the teacher and the student only in the same class, *i.e.* the average of all $\Phi(M_v, S_v^{n_j}), j \neq 0$ in Eq. (7). We can see that the cosine similarity of AMID is consistently larger than that of the other two cases and the cosine similarity of w/o D_2 is larger than that of w/o \mathcal{L}_{adv} , meaning that the proposed approach indeed helps the student to learn more compact representations for each class.

4.5. Hyperparameters Analysis

Sensitivity analysis studies are conducted for the hyperparameters, α_1 , α_2 and α_3 , referenced in the loss function Eq. (17). Similar to the previous study, we experiment on UCF51 and show the results in Tab. 6. The other two hyperparameters are kept the same when one changes. We test different values for α_1 from 3.7 to 4.2, α_2 from 1.3 to 4.8 and α_3 from 0.7 to 1.2, all with an increment of 0.1. The results show that the hyperparameters α_1 , α_2 and α_3 are not sensitive within a certain range. In general, α_1 between 3.8 to 4.1 and α_2 between 1.4 to 1.7 works well, and we set $\alpha_3 = 1.0$ for all other experiments.

5. Conclusion

In this paper, we propose a generic cross-modal distillation method, AMID, to enhance the target modality representations. The proposed approach simultaneously maximizes the lower bounds on MI between the teacher and the student as well as between the teacher and an auxiliary model while minimizing the conditional entropy of the teacher given the student. For optimizing the lower bound on MI, we derive a new form and solve it by a contrastive learning-based approach. For minimizing the conditional entropy, an adversarial learning-based approach is used. Our extensive experimental evaluation demonstrates the effectiveness of AMID on three popular datasets.

Acknowledgement

This work is supported by the National Key R&D Program of China (No. 2022ZD0160702), STCSM (No. 22511106101, No. 18DZ2270700, No. 21DZ1100100, No. 21511101100), 111 plan (No. BP0719010), and State Key Laboratory of UHD Video and Audio Production and Presentation.

References

- [1] Sarah A. Abdu, Ahmed H. Yousef, and Ashraf Salem. Multimodal video sentiment analysis using deep learning approaches, a survey. *Information Fusion*, pages 204–226, 2021. **1**
- [2] Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. Deep variational information bottleneck. In *International Conference on Learning Representations, ICLR*, 2017. **2**
- [3] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *CoRR*, 2020. **6**
- [4] Avi Ben-Cohen, Eyal Klang, Stephen P. Raskin, Shelly Soffer, Simona Ben-Haim, Eli Konen, Michal Marianne Amitai, and Hayit Greenspan. Cross-modality synthesis from CT to PET using FCN and GAN networks for improved automated lesion detection. *Eng. Appl. Artif. Intell.*, pages 186–194, 2019. **1**
- [5] Uta Buchler, Biagio Brattoli, and Bjorn Ommer. Improving spatiotemporal self-supervision by deep reinforcement learning. In *Proceedings of the European Conference on Computer Vision, ECCV*, 2018. **5**
- [6] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan. IEMO-CAP: interactive emotional dyadic motion capture database. *Lang. Resour. Evaluation*, pages 335–359, 2008. **2, 5**
- [7] Cheng Chen, Qi Dou, Yueming Jin, Hao Chen, Jing Qin, and Pheng-Ann Heng. Robust multimodal brain tumor segmentation via feature disentanglement and gated fusion. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 447–456, 2019. **1**
- [8] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2172–2180, 2016. **2, 4**
- [9] Yanbei Chen, Yongqin Xian, A. Sophia Koepke, Ying Shan, and Zeynep Akata. Distilling audio-visual knowledge by compositional contrastive learning. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 7016–7025, 2021. **2, 3, 4, 6, 7**
- [10] Nieves Crasto, Philippe Weinzaepfel, Karteek Alahari, and Cordelia Schmid. Mars: Motion-augmented rgb stream for action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7882–7891, 2019. **1**
- [11] Rui Dai, Srijan Das, and François Bremond. Learning an augmented rgb representation with cross-modal knowledge distillation for action detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13053–13064, 2021. **1**
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*, pages 4171–4186, 2019. **6**
- [13] Thang Doan, João Monteiro, Isabela Albuquerque, Bogdan Mazoure, Audrey Durand, Joelle Pineau, and R. Devon Hjelm. On-line adaptative curriculum learning for gans. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3470–3477, 2019. **2**
- [14] Ishan Durugkar, Ian Gemp, and Sridhar Mahadevan. Generative multi-adversarial networks. *arXiv preprint arXiv:1611.01673*, 2016. **2**
- [15] Adam Foster, Martin Jankowiak, Eli Bingham, Paul Horsfall, Yee Whye Teh, Tom Rainforth, and Noah D. Goodman. Variational bayesian optimal experimental design. In *Advances in Neural Information Processing Systems*, pages 14036–14047, 2019. **2**
- [16] Nuno C Garcia, Pietro Morerio, and Vittorio Murino. Modality distillation with multiple stream networks for action recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 103–118, 2018. **1**
- [17] Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, pages 776–780, 2017. **6**
- [18] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014. **2**
- [19] Wei Han, Hui Chen, and Soujanya Poria. Improving multimodal fusion with hierarchical mutual information maximization for multimodal sentiment analysis. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pages 9180–9192, 2021. **1, 2**
- [20] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 961–970, 2015. **1, 2, 5**
- [21] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. **2**
- [22] Judy Hoffman, Saurabh Gupta, and Trevor Darrell. Learning with side information through modality hallucination. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 826–834, 2016. **1**
- [23] Minhao Hu, Matthis Maillard, Ya Zhang, Tommaso Ciceri, Giammarco La Barbera, Isabelle Bloch, and Pietro Gori. Knowledge distillation from multi-modal to mono-modal segmentation networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention, MICCAI*, pages 772–781, 2020. **1**
- [24] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 2261–2269, 2017. **6**

- [25] Jangho Kim, Seonguk Park, and Nojun Kwak. Paraphrasing complex network: Network compression via factor transfer. In *Advances in Neural Information Processing Systems*, pages 2765–2774, 2018. 2
- [26] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D. Plumbley. Panns: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE ACM Trans. Audio Speech Lang. Process.*, pages 2880–2894, 2020. 6
- [27] J. Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information Theory*, pages 145–151, 1991. 5
- [28] Yunze Liu, Qingnan Fan, Shanghang Zhang, Hao Dong, Thomas Funkhouser, and Li Yi. Contrastive multimodal fusion with tupleinfonce. In *Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV*, pages 754–763, 2021. 2
- [29] Mengmeng Ma, Jian Ren, Long Zhao, Sergey Tulyakov, Cathy Wu, and Xi Peng. SMIL: multimodal learning with severely missing modality. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2302–2310, 2021. 1
- [30] Tabet Matiisen, Avital Oliver, Taco Cohen, and John Schulman. Teacher–student curriculum learning. *IEEE Transactions on Neural Networks and Learning Systems*, pages 3732–3740, 2020. 5
- [31] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 2, 5
- [32] Seyed-Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5191–5198, 2020. 1
- [33] Behnam Neyshabur, Srinadh Bhojanapalli, and Ayan Chakrabarti. Stabilizing gan training with multiple random projections. *arXiv preprint arXiv:1705.07831*, 2017. 2
- [34] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2019. 2
- [35] Peyman Passban, Yimeng Wu, Mehdi Rezagholizadeh, and Qun Liu. ALP-KD: attention-based layer projection for knowledge distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 13657–13665, 2021. 2
- [36] Petra Poklukar, Miguel Vasco, Hang Yin, Francisco S. Melo, Ana Paiva, and Danica Kragic. Geometric multimodal contrastive representation learning. In *International Conference on Machine Learning, ICML*, pages 17782–17800, 2022. 2, 6, 7
- [37] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *International Conference on Learning Representations, ICLR*, 2016. 2
- [38] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *International Conference on Learning Representations, ICLR*, 2015. 2
- [39] Zhensheng Shi, Ju Liang, Qianqian Li, Haiyong Zheng, Zhaorui Gu, Junyu Dong, and Bing Zheng. Multi-modal multi-action video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV*, pages 13678–13687, October 2021. 1
- [40] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. A dataset of 101 human action classes from videos in the wild. *Center for Research in Computer Vision*, 2(11), 2012. 2, 5
- [41] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Proceedings of the European Conference on Computer Vision, ECCV*, pages 776–794, 2020. 2
- [42] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *International Conference on Learning Representations, ICLR*, 2020. 2, 3, 6
- [43] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 6450–6459, 2018. 6
- [44] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *Proceedings of the European Conference on Computer Vision, ECCV*, pages 20–36, 2016. 1
- [45] Chen Xu, Quan Li, Junfeng Ge, Jinyang Gao, Xiaoyong Yang, Changhua Pei, Fei Sun, Jian Wu, Hanxiao Sun, and Wenwu Ou. Privileged features distillation at taobao recommendations. In *Conference on Knowledge Discovery and Data Mining, KDD*, pages 2590–2598, 2020. 1
- [46] Chen Xu, Quan Li, Junfeng Ge, Jinyang Gao, Xiaoyong Yang, Changhua Pei, Fei Sun, Jian Wu, Hanxiao Sun, and Wenwu Ou. Privileged features distillation at taobao recommendations. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2590–2598, 2020. 1
- [47] Amir Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Multimodal language analysis in the wild: CMU-MOSEI dataset and interpretable dynamic fusion graph. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics, ACL*, pages 2236–2246, 2018. 1
- [48] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *International Conference on Learning Representations, ICLR*, 2017. 2
- [49] Wen Zhang, Kai Shu, Suhang Wang, Huan Liu, and Yalin Wang. Multimodal fusion of brain networks with longitudinal couplings. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 3–11, 2018. 1
- [50] Jinming Zhao, Ruichen Li, and Qin Jin. Missing modality imagination network for emotion recognition with uncertain missing modalities. In *ACL/IJCNLP*, pages 2608–2618, 2021. 1, 6
- [51] Jinguo Zhu, Shixiang Tang, Dapeng Chen, Shijie Yu, Yakun Liu, Mingzhe Rong, Aijun Yang, and Xiaohua Wang. Complementary relation contrastive distillation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 9260–9269, 2021. 2, 6, 7