

The following publication Y. Fan, W. Xu, H. Wang, J. Wang and S. Guo, "PMR: Prototypical Modal Rebalance for Multimodal Learning," 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 2023, pp. 20029-20038 is available at <https://doi.org/10.1109/CVPR52729.2023.01918>.

## PMR: Prototypical Modal Rebalance for Multimodal Learning

Yunfeng Fan<sup>1</sup>, Wenchao Xu<sup>1</sup>, Haozhao Wang<sup>2,\*</sup>, Junxiao Wang<sup>3,4</sup>, and Song Guo<sup>1</sup>

<sup>1</sup>The Hong Kong Polytechnic University, <sup>2</sup>Huazhong University of Science and Technology

<sup>3</sup>KAUST, <sup>4</sup>SDAIA-KAUST AI Center

{yunfeng.fan, wenchao.xu}@polyu.edu.hk, hz\_wang@hust.edu.cn,  
junxiao.wang@kaust.edu.sa, song.guo@polyu.edu.hk

### Abstract

Multimodal learning (MML) aims to jointly exploit the common priors of different modalities to compensate for their inherent limitations. However, existing MML methods often optimize a uniform objective for different modalities, leading to the notorious "modality imbalance" problem and counterproductive MML performance. To address the problem, some existing methods modulate the learning pace based on the fused modality, which is dominated by the better modality and eventually results in a limited improvement on the worse modal. To better exploit the features of multimodal, we propose Prototypical Modality Rebalance (PMR) to perform stimulation on the particular slow-learning modality without interference from other modalities. Specifically, we introduce the prototypes that represent general features for each class, to build the non-parametric classifiers for uni-modal performance evaluation. Then, we try to accelerate the slow-learning modality by enhancing its clustering toward prototypes. Furthermore, to alleviate the suppression from the dominant modality, we introduce a prototype-based entropy regularization term during the early training stage to prevent premature convergence. Besides, our method only relies on the representations of each modality and without restrictions from model structures and fusion methods, making it with great application potential for various scenarios. The source code is available [here](#)<sup>1</sup>.

### 1. Introduction

Multimodal learning (MML) [24, 34, 36] emerges to mimic the way humans perceive the world, i.e., from multiple sense channels toward a common phenomenon for better understanding the external environment, which has attracted extensive attention in various scenarios, e.g. video classification [13, 28, 37], event localization [38, 43], ac-

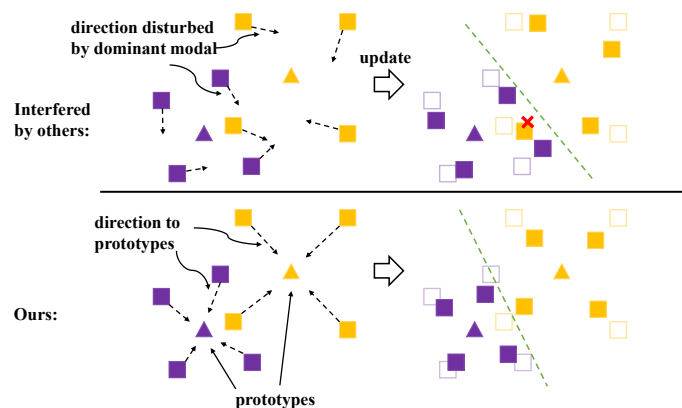


Figure 1. The slow-learning modal’s updating direction is severely disturbed by the dominant one, making it hard to exploit its features. We propose to use the prototypes, the centroids of each class in representation space, to adjust updating direction for better uni-modal performance. Other modalities will not interfere with the new direction, which ensures improvement.

tion recognition [10, 33], audiovisual speech recognition [23, 25]. By employing a complementary manner of multimodal training, it is expected that MML can achieve better performance than using a single modality. However, the heterogeneity of multimodal data poses challenges on how to learn multimodal correlations and complementarities.

According to recent research [29, 42], although the overall performance of multimodal learning exceeds that of single-modal learning, the performance of each modality tends to be far from their upper bound. The reason behind this phenomenon is the “modality imbalance” problem, in which the dominant modality will hinder the full utilization of multimodal. The researchers [40] also claimed that different modalities overfit and converge at different rates, meaning that optimizing the same objective for different modalities leads to inconsistent learning efficiency.

Several methods [29, 40, 44] have been proposed to address the problem. Some of them [29, 44] try to modu-

\*Corresponding author

<sup>1</sup><https://github.com/fanyunfeng-bit/Modal-Imbalance-PMR>

late the learning paces of different modalities based on the fusion modal. However, we find out through experiments that the dominant modality not only suppresses the learning rates of other modalities [29] but also interferes with their update direction, which makes it hard to improve the performance of slow-learning modalities. Moreover, existing methods either inevitably bring additional model structures [5, 40] or are limited by specific fusion methods [29, 42], which limit their application scenarios.

To tackle their limitations, we propose the **Prototypical Modal Rebalance** (PMR) strategy to stimulate the slow-learning modality via promoting the exploitation of features and alleviate the suppression from the dominant modality by slowing down itself in the early training stage.

Concretely, we introduce the prototypes for each modality, which are defined as “representative embeddings of instances of a class”. We utilize the prototypes to construct non-parametric classifiers by comparing the distances between each sample with all the prototypes to evaluate the performance of each modality and design a new prototype-based metric inspired by [29] to monitor the modality imbalance degree during the training process. Then, we propose the prototypical cross-entropy (PCE) loss to accelerate the slow-learning modality by enhancing its clustering process, as illustrated in Fig. 1. The PCE loss can achieve comparable performance to the cross-entropy (CE) loss [1] in the classification task, and more importantly, it is not affected by the dominant modality and gives internal impetus for full feature exploitation instead of going on in the disturbed direction. In addition, we introduce a prototypical entropy regularization (PER) term, which can be seen as a penalty on the dominant modality to prevent premature convergence for suppression effect alleviation. Our method only relies on the representations of each modality and without restrictions from model structures and fusion methods. Therefore, the PMR strategy has great generality potential. To summarize, our contributions in this paper are as follows:

- We analyze the modal imbalance problem and find that during the training process, the deviation of the gradient update direction of the uni-modal became larger, indicating that we should not regulate along the original gradient.
- We propose PMR to address the modal imbalance problem by actively accelerating slow-learning modalities with PCE loss and simultaneously alleviating the suppression of the dominant modality via PER.
- We conduct comprehensive experiments and demonstrate that 1) PMR can achieve considerable improvements over existing methods; 2) PMR is independent of the fusion method or model structure and has strong advantages in generality.

## 2. Related Works

### 2.1. Imbalanced multimodal learning

Several recent studies [5, 40, 41] have shown that many multimodal DNNs cannot achieve better performance compared to the best single-modal DNNs. Wang *et al.* [40] found that different modalities overfit and generalize at different rates and thus obtain suboptimal solutions when jointly training them using a unified optimization strategy. Peng *et al.* [29] proposed that the better-performing modality will dominate the gradient update while suppressing the learning process of the other modality. Furthermore, it has been reported that multimodal DNNs can exploit modal bias in the data, which is inconsistent with the expectation of exploiting cross-modal interactions in VQA [8, 12, 41].

Several approaches have been proposed recently to deal with the modal imbalance problem [5, 29, 40, 44]. Wang *et al.* [40] used additional classifiers for each modality and its fusion modality and then optimized the gradient mixing problem they introduced to obtain better weights for each branch. Du *et al.* [5] tried to improve uni-modal performance by distilling knowledge from well-trained models. However, these approaches introduce more model structure and computational effort, which makes the training process more complex and expensive. Xiao *et al.* [44] proposed DropPathway, which randomly drops the audio pathway during training as a regularization technique to adjust the learning paces between visual and audio pathways. Peng *et al.* [29] chose to slow down the learning rate of the mighty modality by online modulation to lessen the inhibitory effect on the other modality. Although a certain degree of improvement is achieved, such approaches do not impose the intrinsic motivation of improvement on the slow-learning modality, making the improvement of this modality a passive rather than an active behavior. Besides, the interference from other modalities will hinder the improvement by modulation based on the fused modality data. Furthermore, the application scenarios of these methods are limited by fusion methods or model structures. In this paper, we aim to power the slow-learning modality to reduce the gap with the dominant one and allow it to be used in conjunction with various fusion methods and model architectures.

### 2.2. Prototypical network

Prototypical networks were originally proposed to solve few-shot or zero-shot classification problems [4, 7, 32, 35], based on the idea that there is an embedding, defined as a prototype, which is surrounded by points from the same class. Recently, this approach has been widely used to address long-tail recognition [45], domain adaptation [3, 27], and facilitate unsupervised learning [20], since prototypes can be used to represent general features of a class. In [20, 35], prototypes were interpreted as non-parametric pro-

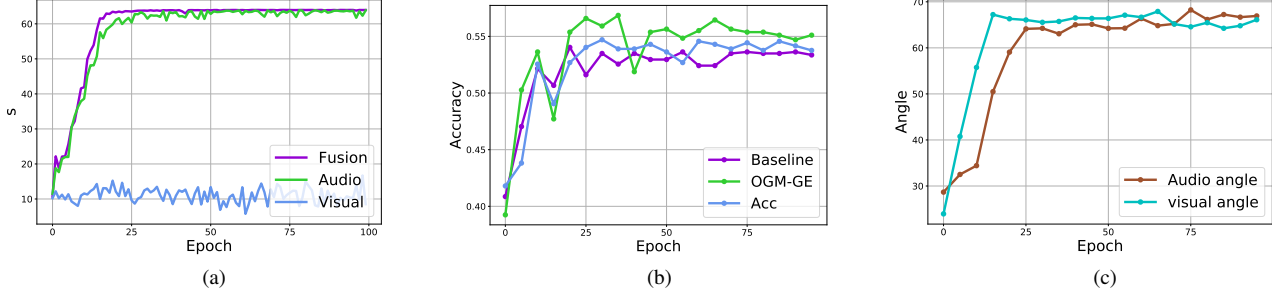


Figure 2. (a) Performance of each modality and their fusion. (b) Training accuracy of multimodal with different modulations. Baseline means no extra modulation. Acc increases the gradient magnitude of the slow-learning modality. OGM-GE [29] reduces the gradient magnitude of better modality. (c) The gradient direction discrepancy (angle) between uni-modal and multimodal on the baseline. The results were acquired from CREMA-D.

totypical classifiers that perform on par or even better than parametric linear classifiers. In this paper, we leverage prototypes to build non-parametric classifiers to evaluate the performance of each modality feature.

### 3. Modality Imbalance Analysis

Without loss of generality, we consider two input modalities as  $m_0$  and  $m_1$ . The dataset is denoted as  $\mathcal{D}$ , which consists of instances and their corresponding labels  $(\mathbf{x}, y)$ , where  $\mathbf{x} = (\mathbf{x}^{m_0}, \mathbf{x}^{m_1}, y) = \{x_i^{m_0}, x_i^{m_1}, y_i\}_{i=1,2,\dots,N}$ ,  $y = \{1, 2, \dots, M\}$ , and  $M$  is the number of categories. The goal is to train a model with this data to predict  $y$  from  $\mathbf{x}$ .

We use a multimodal DNN with two uni-modal branches for prediction. Each branch has an encoder, denoted as  $\phi^0, \phi^1$ , to extract features of respective modal data  $\mathbf{x}^{m_0}$  and  $\mathbf{x}^{m_1}$ . The representation outputs of encoders are denoted as  $\mathbf{z}^0 = \phi^0(\theta^0, \mathbf{x}^{m_0})$  and  $\mathbf{z}^1 = \phi^1(\theta^1, \mathbf{x}^{m_1})$ , where  $\theta^0$  and  $\theta^1$  are the parameters of encoders. The two uni-modal encoders are connected through the representations by some kind of fusion, which is prevalent in multimodal learning [19, 22, 46]. In this work, we have tried some different fusion methods. For simplicity, we use  $[\cdot, \cdot]$  to denote fusion operation. Let  $W \in \mathbb{R}^{M \times (d_{z^0} + d_{z^1})}$  and  $b \in \mathbb{R}^M$  denote the parameters of the linear classifier to produce the logits output:

$$f(\mathbf{x}) = W [\phi^0(\theta^0, \mathbf{x}^{m_0}); \phi^1(\theta^1, \mathbf{x}^{m_1})] + b \quad (1)$$

The cross-entropy loss of the multimodal model is

$$\mathcal{L}_{CE} = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{f(x_i)_{y_i}}}{\sum_{k=1}^M e^{f(x_i)_k}} \quad (2)$$

The gradient of the softmax logits output with true label  $y^i$  should be:

$$\frac{\partial \mathcal{L}_{CE}}{\partial f(x_i)_{y_i}} = \frac{e^{(W[\phi^0; \phi^1] + b)_{y_i}}}{\sum_{k=1}^M e^{(W[\phi^0; \phi^1] + b)_k}} - 1 \quad (3)$$

For convenience, we simplify  $\phi^0(\theta^0, \mathbf{x}^{m_0})$  and  $\phi^1(\theta^1, \mathbf{x}^{m_1})$  as  $\phi^0$  and  $\phi^1$ . According to Eq. (3), the final gradient is influenced by the performance of the fused modality. However, we cannot directly judge the contribution of the two modalities. To do it, we take a simple fusion method, summation, as the example here:

$$f(\mathbf{x}) = W^0 \cdot \phi^0(\theta^0, \mathbf{x}^{m_0}) + b^0 + W^1 \cdot \phi^1(\theta^1, \mathbf{x}^{m_1}) + b^1 \quad (4)$$

where  $W^0 \in \mathbb{R}^{M \times d_{z^0}}$ ,  $W^1 \in \mathbb{R}^{M \times d_{z^1}}$  and  $b^0, b^1 \in \mathbb{R}^M$  are the parameters for individual modal classifier. Therefore, we use the logits output of the ground truth as the performance for each modality and their summation fusion:

$$\begin{aligned} s^0 &= \text{softmax}(W^0 \cdot \phi^0 + b^0)_y \\ s^1 &= \text{softmax}(W^1 \cdot \phi^1 + b^1)_y \\ s^{fu} &= \text{softmax}(W^0 \cdot \phi^0 + b^0 + W^1 \cdot \phi^1 + b^1)_y \end{aligned} \quad (5)$$

As shown in Fig. 2a, the performance of the audio modal is very similar to the multimodal during training and the visual modal is much worse in CREMA-D [2], which means better modality contributes more to the gradient because of higher performance similarity with their fusion. Moreover, the visual modal is severely suppressed in the multimodal learning process because of the excessive dominance of gradient updates by the audio modal. Therefore, we have to mitigate the inhibition on visual modal to fully exploit visual features. A straightforward approach can be to increase the magnitude of the gradient. We test a similar way with OGM-GE [29], named Acc, to increase the gradient magnitude of the slow-learning modality instead of lowering the gradient magnitude of the better modality in OGM-GE. The results are shown in Figs. 2b and 2c.

As demonstrated in Fig. 2b, increasing the gradient magnitude of the slow-learning modality (visual) could improve the validation accuracy a little bit but not as obviously as

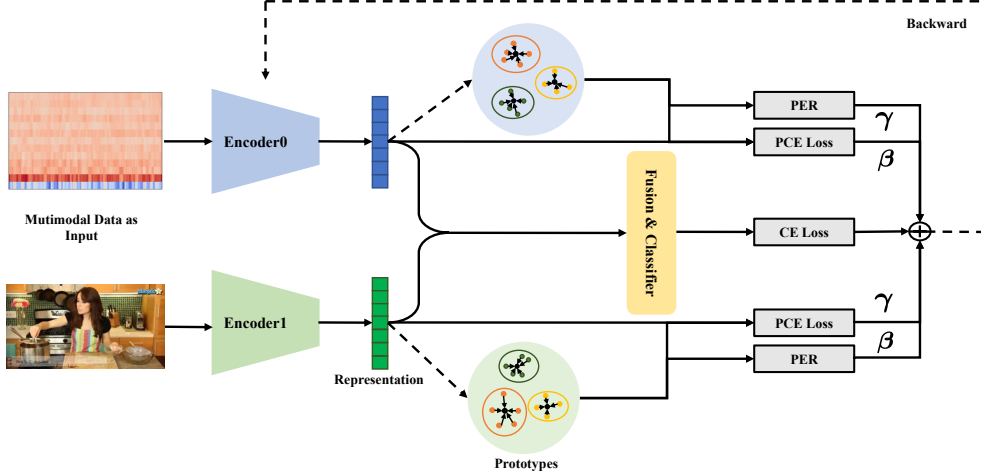


Figure 3. The pipeline of modality modulation with prototypical modal rebalance strategy.

OGM-GE does. To find the reason for the phenomenon, we use the uni-modal output  $s^0$  and  $s^1$  to calculate the gradient for each modal branch additionally with CE loss and illustrate the direction discrepancy of gradients between each uni-modal and multimodal  $s^{fu}$ , as shown in Fig. 2c. The angle between the actual gradient update direction (from the multimodal output) and each modal's guidance direction (from the uni-modal output) increases dramatically during training, in the meantime, remaining acute. Therefore, the two modalities influence each other, resulting in a larger gap between the gradient update direction obtained by the fused modality and the expected direction of each modality. That means the slow-learning modality cannot fully exploit its feature with the disturbance from other modalities, ultimately making the method of modulating the gradient amplitude limited, as illustrated in Fig. 1.

## 4. Prototypical Modal Rebalance

### 4.1. Prototypical CE loss for modal acceleration

As discussed above, due to inconsistency in performance between modalities, they will affect each other on the exploitation of self-feature in multimodal learning, and the modality with worse performance is particularly suppressed, which limits multimodal performance. In order to solve the problem, we aim to facilitate the exploitation of the features of the slow-learning modality via prototypes, as shown in Fig. 3.

The performance estimation in Eq. (5) needs the logits output after the classifier and of course, the logits have to be decomposable into separate components of the two modalities. However, this constraint is too strong, limiting the estimation's application in most scenarios. Aiming to implement a universal estimation method, we introduce the prototypes of the categories. With the data

$x = \{x_i^{m_0}, x_i^{m_1}, y_i\}_{i=1,2,\dots,N}$ , we could produce the representations  $z = \{z_i^0, z_i^1\}_{i=1,2,\dots,N}$  in the training process. We denote  $z_k^0 = \{z_{k_i}^0\}_{i=1,2,\dots,N_k}$ ,  $z_k^1 = \{z_{k_i}^1\}_{i=1,2,\dots,N_k}$  as the subset data of each category, where  $N_k$  is the number of class  $k$ , and  $k = 1, 2, \dots, M$ . The prototype is the centroid of each category data:

$$\begin{aligned} c_k^0 &= \frac{1}{N_k} \sum_{i=1}^{N_k} z_{k_i}^0 \\ c_k^1 &= \frac{1}{N_k} \sum_{i=1}^{N_k} z_{k_i}^1 \end{aligned} \quad (6)$$

Then, we use prototypes to produce a distribution over classes for the input data  $x$  based on a softmax over distances to the prototypes in the embedding space for each modality as in [35]:

$$\begin{aligned} p_i^{\phi^0}(y = k | x_i^{m_0}) &= \frac{\exp(-d(\phi^0(x_i^{m_0}), c_k^0))}{\sum_{k'} \exp(-d(\phi^0(x_i^{m_0}), c_{k'}^0))} \\ p_i^{\phi^1}(y = k | x_i^{m_1}) &= \frac{\exp(-d(\phi^1(x_i^{m_1}), c_k^1))}{\sum_{k'} \exp(-d(\phi^1(x_i^{m_1}), c_{k'}^1))} \end{aligned} \quad (7)$$

where  $d(\cdot, \cdot)$  is the distance function, which is the Euclidean distance in this paper. Here, we design the imbalanced ratio inspired by [29]:

$$\rho_t = \frac{\sum_{i \in B_t^0} p_i^{\phi^0}}{\sum_{i \in B_t^1} p_i^{\phi^1}} \quad (8)$$

$B_t^0$  and  $B_t^1$  are the batch data at training step time  $t$ , therefore, we could use it to evaluate the degree of imbalance between two modalities in real-time. The metric is calculated using only representations and prototypes, which are both computationally independent of the fusion method and classifier structure.

According to Sec. 3, in order to exploit more information of the slow-learning modality, increasing the gradient magnitude of the suppressed modal is not an ideal solution due to the perturbation by another modality. We leverage the prototypes to introduce the PCE loss that is independent of another modality to promote the slow-learning modal’s performance:

$$\begin{aligned}\mathcal{L}_{PCE}^0(f) &= \mathbb{E}_{p(x_i^{m_0}, y)} \left[ -\log \frac{\exp(-d(z_i^0, c_y^0))}{\sum_{k=1}^M \exp(-d(z_i^0, c_k^0))} \right] \\ \mathcal{L}_{PCE}^1(f) &= \mathbb{E}_{p(x_i^{m_1}, y)} \left[ -\log \frac{\exp(-d(z_i^1, c_y^1))}{\sum_{k=1}^M \exp(-d(z_i^1, c_k^1))} \right]\end{aligned}\quad (9)$$

The acceleration loss is the weighted combination of CE loss and PCE loss:

$$\mathcal{L}_{acc} = \mathcal{L}_{CE} + \alpha \cdot \beta \mathcal{L}_{PCE}^0 + \alpha \cdot \gamma \mathcal{L}_{PCE}^1 \quad (10)$$

where  $\alpha$  is a hyper-parameter to control the degree of modulation. With  $\rho_t$  to evaluate the imbalance degree dynamically, we are able to regulate the learning speed of each modality by adjusting the coefficients  $\beta$  and  $\gamma$  in a simple way:

$$\begin{cases} \beta = clip\left(0, \frac{1}{\rho_t} - 1, 1\right), \gamma = 0 & \rho_t < 1 \\ \beta = 0, \gamma = clip(0, \rho_t - 1, 1) & \rho_t \geq 1 \end{cases} \quad (11)$$

where  $clip(a, b, c)$  is the truncate function which constrains  $b$  to be between  $a$  and  $c$ . In this way, the slower-learning modality can be facilitated to exploit its features, while the better modality maintains the original learning strategy, which mitigates the modal imbalance phenomenon. Because the PCE loss only uses the representation of each modal encoder, our method can be applied to any modal fusion scenario, as long as the model itself has encoder(s) to extract features for two modalities respectively. Besides, to stabilize the learning process and reduce the computation cost, we compute the prototypes based on a subset of training data in a momentum fashion between each training epoch:

$$\begin{aligned}c_k^0|_{old} &= \varepsilon c_k^0|_{old} + (1 - \varepsilon) c_k^0|_{new} \\ c_k^1|_{old} &= \varepsilon c_k^1|_{old} + (1 - \varepsilon) c_k^1|_{new}\end{aligned}\quad (12)$$

where  $c_k|_{old}$  is the previous prototype in last epoch and  $c_k|_{new}$  is the prototype calculated in the current epoch.

## 4.2. Prototypical entropy regularization for inhibition reduction

With the help of PCE loss, we could accelerate the slow-learning modality. Nevertheless, the fusion output is still a hindrance that restrains the improvement of the slow-learning modality. As demonstrated in Figs. 2a and 2c, the distraction from other modalities increases dramatically

---

### Algorithm 1: multimodal learning with PMR.

---

**Input:** Input data  $\mathcal{D} = \{x_i^{m_0}, x_i^{m_1}, y_i\}_{i=1,2,\dots,N}$ , subset  $\mathcal{D}_f$ , initialized model parameters  $\theta^0$ ,  $\theta^1$ , hyper-parameters  $\alpha, \varepsilon$ , epoch number  $E$ , regularization epoch number  $E_r$ .

```
int e=0;
while e < E do
    Obtain the subset representations  $z_s^0, z_s^1$  by
    feeding-forward  $\mathcal{D}_f$  to the model;
    Calculate the detached prototypes  $c^0, c^1$  using
    Eq. (12);
    foreach mini-batch data  $B_t$  in  $\mathcal{D}$  at step  $t$  do
        Obtain the representations  $z_t^0, z_t^1$  by
        feeding-forward  $B_t$  to the model;
        Calculate  $\rho_t$  using Eq. (8);
        if  $e < E_r$  then
            Calculate the final loss  $\mathcal{L}_{final}$  using
            Eqs. (10), (11) and (13);
        else
            Calculate the final loss using Eq. (10)
        Update the model based on  $\mathcal{L}_{final}$ ;
    e=e+1;
```

---

when the performance gap between modalities is quite distinct. In order to reduce the inhibition, we propose the prototypical entropy regularization (PER) terms to slow down the convergence speed of the dominant modality:

$$\begin{aligned}\mathcal{L}_{final} &= \mathcal{L}_{acc} - \mu \cdot \gamma H^0(\pi(-d(z^0, c_y^0))) \\ &\quad - \mu \cdot \beta H^1(\pi(-d(z^1, c_y^1)))\end{aligned}\quad (13)$$

where  $\pi$  is the *softmax* function which produces the probabilities and  $H$  is the entropy.  $\mu$  is a hyper-parameter. The coefficients  $\beta$  and  $\gamma$  are multiplied on the opposite modality compared with Eq. (10), which means accelerating slow-learning modality and preventing modality premature convergence happens at the same time. One point worth emphasizing is that we only add the regularization term in the first training few epochs to reduce the inhibitory effect in the early training stage to avoid performance damage of this modal. Overall, the pseudo-code of PMR is provided in Algorithm 1.

## 5. Evaluation

### 5.1. Datasets

CREMA-D [2] is an audio-visual dataset for the study of emotion recognition, which consists of facial and vocal emotional expressions. The emotional states can be divided into 6 categories: *happy, sad, anger, fear, disgust* and *neutral*. There is a total of 7442 clips in the dataset, which are

randomly divided into 6698 samples as the training set and 744 samples as the testing set.

**AVE** [38] is an audio-visual video dataset for audio-visual event localization, in which there are 28 event classes and 4,143 10-second videos with both auditory and visual tracks as well as second-level annotations. All the videos are collected from YouTube. In our experiments, we extract the frames from event-localized video segments and capture the audio clips within the same segments, constructing a labeled multimodal classification dataset. The training and validation split of the dataset follows [38].

**Colored-and-gray MNIST** [16] is a synthetic dataset based on MNIST [18], which is denoted as CG-MNIST. Each instance contains two kinds of images, a gray-scale and a monochromatic image. In the training set, there are 60,000 instances and the monochromatic images are strongly color-correlated with their digit labels. In the validation set, the number of instances is 10,000, while the monochromatic images are weakly color-correlated with their labels. In this work, we consider the monochromatic image as the first modality and the gray-scale image as the second modality, as in [42]. This dataset is used to prove the method’s effectiveness beyond audio-visual.

## 5.2. Experimental settings

For datasets CREMA-D and AVE, we adopt the ResNet18 [9] as our encoder backbones and map the input data as 512-dimensional vectors. For audio modality, the data is converted to a spectrogram of size 257×1,004 for AVE and 257×299 for CREMA-D. For visual modality, we extract 10 frames from the video clips and randomly select 3 frames for CREMA-D (4 for AVE) to build the training dataset. For Colored-and-gray MNIST, we build a neural network with 4 convolution layers and 1 average pool layer as the encoder. We train all the models with mini-batch size 64 and an SGD optimizer [31] with a momentum of 0.9 and a weight decay of 1e-4. The learning rate is initialized as 1e-3 and gradually decays to 1e-4. The subset of data for prototype calculation is one-tenth the size of training data and is also extracted from the training data.  $\alpha$  is set to 1 or 2, for different datasets.  $\mu$  is set to a small value 1e-2 or 1e-3 for different datasets. All of our experiments were performed on one NVIDIA GeForce RTX 3090 GPU.

## 5.3. Effectiveness on the multimodal task

**Comparison on conventional fusion methods.** In this experiment, we apply the PMR strategy on 4 kinds of basic fusion methods: concatenation [26], summation, film [30] and gated [15]. Among these, summation is the type of late fusion and the other three belong to intermediate fusion [11, 21] method. The logit output of summation and concatenation can be split into two individual parts for each modality combined with a linear classifier, as discussed

Dataset	CREMA-D	AVE	CG-MNIST
Method	Acc	Acc	Acc
Uni-modal1	54.4	62.1	99.3
Uni-modal2	58.0	31.0	60.4
Concatenation	53.2	65.4	58.4
Summation	52.9	65.1	59.1
Film	57.2	64.4	60.0
Gated	58.4	63.5	59.8
Concatenation†	61.1	67.1	77.2
Summation†	59.4	<b>68.1</b>	<b>78.5</b>
Film†	<b>61.8</b>	66.4	76.3
Gated†	59.9	62.7	68.7

Table 1. Performance on CREMA-D, AVE and Colored-and-gray MNIST dataset with various fusion methods. † indicates PMR strategy is applied. PMR gets great performance improvement on nearly all scenarios

in [29]. In film and gated, features between the modalities are fused in more complicated ways therefore the logit output cannot be fully split. The results are shown in Tab. 1, which also includes the performance trained on every single modality. Modal1 is audio and modal2 is visual for CREMA-D and AVE, while modal1 is gray-modality and modal2 is colored-modality for CG-MNIST. According to the results, we can find that the performance of each uni-modal model in each dataset is inconsistent, as the audio performance is worse than the visual in CREMA-D, and on the contrary, the audio performance is better than the visual in AVE. In addition, uni-modal performance may outperform the vanilla fusion methods. For example, uni-visual performance is obviously better than the concatenation and summation for CREMA-D, so the uni-gray performance is in CG-MNIST, indicating the inhibitory relationship between modalities. We get a significant improvement on the three datasets compared with each vanilla fusion method when our PMR strategy is applied, except for a slight decrease in gated on the AVE dataset.

### Improved performance compared with other baselines.

We compare our PMR strategy with other three modulation strategies for modality imbalance: Modality-Drop [44], Gradient-Blending [40] and OGM-GE [29]. We compare them with concatenation and film fusion methods. All results are shown in Tab. 2. We can notice that all the modulation methods achieve better performance compared with the baseline and our PMR gets the best among them. The main improvement contribution comes from PCE loss to accelerate the slow-learning modality while combining with PER will still get almost a 1% increase or at least stay the same. Apart from these, we also note that Gradient-Blending has to train additional uni-modal classifiers and requires the validation results with extra computation, while OGM-GE can only be directly used in concatenation but not film, therefore, we modulate the training process every few epochs as

Dataset	CREMA-D		AVE	
	Concat	Film	Concat	Film
Naive	53.2	57.2	65.4	64.4
Modality-drop	55.8	58.3	66.4	63.9
Grad-Blending	56.7	58.7	65.5	65.2
OGM-GE	57.7	58.0	65.9	65.1
PMR w/o PER	60.3	61.3	<b>67.2</b>	65.6
PMR	<b>61.1</b>	<b>61.8</b>	67.1	<b>66.4</b>

Table 2. Comparison with various modulation strategies on CREMA-D and AVE dataset with concatenation and film fusion methods. PMR achieves the best performance among them.

Dataset	CREMA-D	CG-MNIST
Method	Acc	Acc
MMTM	51.6	71.4
CentralNet	50.2	69.3
MMTM†	<b>55.1</b>	<b>74.2</b>
CentralNet†	52.9	72.1

Table 3. Performance on CREMA-D and CG-MNIST with two kinds of intermediate fusion methods. † indicates PMR strategy is applied, which achieves better performance.

done in [29]. Compared to them, our method doesn’t need additional modules and is independent of the fusion methods and classifier structure, which makes it applicable to more scenarios.

**Performance on different architectures.** The fusion stage of the four fusion methods used above is after the encoder or the classifier. To validate the applicability of PMR in more scenarios, we combine it with two intermediate fusion methods MMTM [14] and CentralNet [39] with and without PMR on CREMA-D and CG-MNIST. MMTM could recalibrate the channel-wise features of different CNN streams by squeeze and multimodal excitation steps. We use ResNet18 as the backbone and apply MMTM in the three final residual blocks. CentralNet uses both unimodal hidden representations and a central joint representation at each layer, which are fused by a learned weighted summation. We use only one frame for each video on CREMA-D for convenience. According to Tab. 3, our proposed PMR achieves prominent improvement even if fusion is operated during the processing of encoders, indicating the applicability of PMR in complex scenarios.

**Application on another task.** We process the AVE dataset as a classification dataset in the above experiments. Here we use the original AVE dataset to complete the audio-visual event localization task. We apply PMR on the AVEL [38] with concatenation and DMRN [38] fusion methods, which utilize LSTM to model temporal dependencies in the two modalities respectively. We evaluate the performance of supervised event localization in the late fusion style. In experiments, we only apply PMR on the ground truth video

Dataset	AVE
Fusion	Acc
concat	72.7
DMRN	73.1
concat†	<b>74.3</b>
DMRN†	74.2

Table 4. Event localization experiments based on the AVEL with two fusion methods. † indicates PMR strategy is applied.

segments without extra processing on the video segments which are not relevant to the label. We choose to use Adam optimizer with the same settings in [38]. The results are shown in Tab. 4. Applying PMR on the event localization task still achieves a certain promotion on accuracy. The utilization of two kinds of fusion methods on this task indicates that our method has great application potential in various task scenarios and various fusion methods.

## 5.4. Ablation study

**Uni-modal performance comparison.** As we discussed in Sec. 3, the interaction between modalities eventually leads to the failure of each modality to effectively exploit its own features. Therefore, we first compare the uni-modal performance among uni-modal trained models and the uni-modal branch in multimodal models with and without PMR strategy. As shown in Figs. 4a and 4b, the uni-audio model is slightly better than the audio branch in vanilla multimodal learning, while the visual branch has few effective learning resulting in a distinct gap with uni-visual modal. After applying PMR strategy, the audio branch doesn’t show a noticeable change since our strategy is mainly for facilitation for the slow-learning modality, while the visual branch improves considerably, indicating that our PMR strategy outperforms by indeed promoting the feature exploitation of the slow-learning modality. Fig. 4c illustrates the curves of imbalance ratio  $\rho$  in vanilla multimodal and multimodal with PMR. It can be seen that our PMR helps to alleviate the modality imbalance in multimodal learning. Although audio modality learns fast in the early training process and converges quickly, visual modality can still get improvement. Compare with the baseline without PMR, the imbalance ratio with PMR (green line) decreases gradually even if the dominant modality has already converged, indicating the intrinsic encouragement from PCE is less affected by other modalities.

### Analysis of subset data scales for prototype calculation.

In our proposed PMR, the computation of prototypes is the major additional cost. We analyze the performance of our method on different scales of data for prototype calculation. The results are shown in Tab. 5. When the scale of subset is greater than 10%, PMR tends to be stable in accuracy, which means data of this scale can represent the approx-

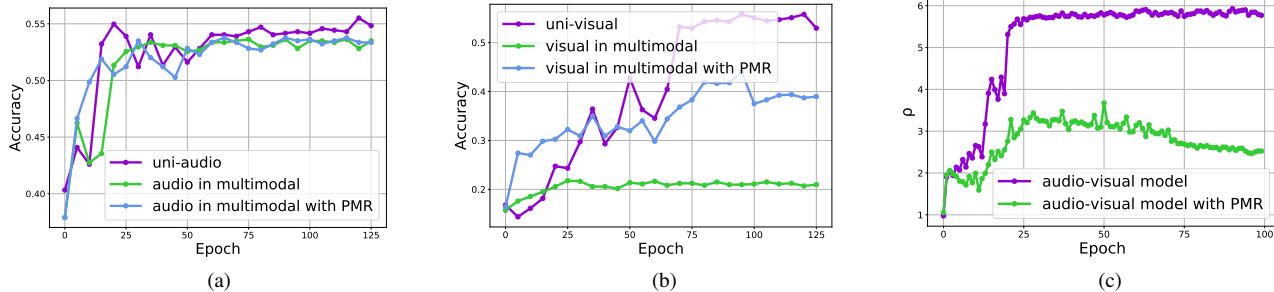


Figure 4. Performance of the uni-modal trained models, uni-modal branch in multimodal trained models, and uni-modal branch in multimodal with PMR on CREMA-D dataset. The fusion method in multimodal models is concatenation. (a) Performance of audio modality. (b) Performance of visual modality. (c) The change of imbalance ratio  $\rho$ .

Dataset	CREMA-D	AVE
Scale	Acc	Acc
baseline	53.2	65.4
1%	54.1	62.1
5%	58.4	65.8
10%	61.1	67.1
50%	61.2	<b>67.7</b>
100%	<b>61.5</b>	67.3

Table 5. Experiments on CREMA-D and AVE with different subset scales. n% indicates the proportion of the training dataset. Fusion method here is concatenation

Dataset	CREMA-D	AVE
Optimizer	Acc	Acc
SGD	53.2	65.4
AdaGrad	58.8	65.1
Adam	59.8	64.4
SGD†	61.1	<b>67.1</b>
AdaGrad†	61.5	66.2
Adam†	<b>65.3</b>	66.5

Table 6. Experiments with AdaGrad and Adam optimizers on CREMA-D and AVE. † indicates PMR strategy is applied. PMR consistently achieves better performance.

imate distribution of the overall data on CREMA-D and AVE. If the data scale is too small, the performance would drop a little bit, even worse than the vanilla baseline when the subset data scale is just 1% on AVE. The reason may be the calculated prototype is biased because of limited data, which further hinders the training of model. Although our method introduces a certain amount of extra computation, the required computation is not much and can be reduced with a reasonable selection of data.

**Adaptive optimizers.** We analyze the modality imbalance problem and propose the PMR strategy under the assumption: the same learning rate for each gradient calculation, which is true for popular SGD-based algorithms, but not rigorous for adaptive optimization algorithms. To validate the effectiveness of our method on more adaptive methods, we apply PMR on optimizers AdaGrad [6] and Adam [17], dynamically adjusting the learning rate for each parameter. As demonstrated in Tab. 6, we empirically show that PMR works with different optimizers. Different optimizers perform inconsistently on different datasets.

## 6. Discussion

Multimodal learning usually falls into a suboptimal solution because of the modality imbalance problem, which indicates that vanilla optimization with a single strategy for

different modalities is limited. We propose the prototypical modal rebalance (PMR) strategy to introduce different learning strategies for different modalities, i.e., accelerating the slow modality with prototypical cross entropy (PCE) loss and reducing the inhibition from dominant modality with prototypical entropy regularization (PER) term. This method achieves considerable performance improvement on the three multimodal datasets with different model structures and fusion methods. The non-parametric classifiers with prototypes can be applied in any scenario as long as we have the representations of instances for each modality.

## 7. Acknowledgement

Our work was supported by fundings from the Key-Area Research and Development Program of Guangdong Province (No. 2021B0101400003), Hong Kong RGC Research Impact Fund (No. R5060-19), Areas of Excellence Scheme (AoE/E-601/22-R), General Research Fund (No. 152203/20E, 152244/21E, 152169/22E, PolyU15222621), Shenzhen Science and Technology Innovation Commission (JCYJ20200109142008673) and the grant from Establishment of Distributed Artificial Intelligence Laboratory for Interdisciplinary Research (UGC/IDS(R)11/19).



## References

- [1] Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019. [2](#)
- [2] Houwei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini Verma. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE transactions on affective computing*, 5(4):377–390, 2014. [3](#), [5](#)
- [3] Yuhan Chai, Lei Du, Jing Qiu, Lihua Yin, and Zhihong Tian. Dynamic prototype network based on sample adaptation for few-shot malware detection. *IEEE Transactions on Knowledge and Data Engineering*, 2022. [2](#)
- [4] Kaize Ding, Jianling Wang, Jundong Li, Kai Shu, Chenghao Liu, and Huan Liu. Graph prototypical networks for few-shot learning on attributed networks. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 295–304, 2020. [2](#)
- [5] Chenzhuang Du, Tingle Li, Yichen Liu, Zixin Wen, Tianyu Hua, Yue Wang, and Hang Zhao. Improving multimodal learning with uni-modal teachers. *arXiv preprint arXiv:2106.11059*, 2021. [2](#)
- [6] John Duchi, Elad Hazan, and Yoram Singer. Adaptive sub-gradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011. [8](#)
- [7] Tianyu Gao, Xu Han, Zhiyuan Liu, and Maosong Sun. Hybrid attention-based prototypical networks for noisy few-shot relation classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019. [2](#)
- [8] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. [2](#)
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [6](#)
- [10] Javed Imran and Balasubramanian Raman. Evaluating fusion of rgb-d and inertial sensors for multimodal human action recognition. *Journal of Ambient Intelligence and Humanized Computing*, 11(1):189–208, 2020. [1](#)
- [11] Giridharan Iyengar and Harriet J Nock. Discriminative model fusion for semantic concept detection and annotation in video. In *Proceedings of the eleventh ACM international conference on Multimedia*, pages 255–258, 2003. [6](#)
- [12] Allam Jabri, Armand Joulin, and Laurens van der Maaten. Revisiting visual question answering baselines. In *European conference on computer vision*, pages 727–739. Springer, 2016. [2](#)
- [13] Yu-Gang Jiang, Zuxuan Wu, Jinhui Tang, Zechao Li, Xiangyang Xue, and Shih-Fu Chang. Modeling multimodal clues in a hybrid deep learning framework for video classification. *IEEE Transactions on Multimedia*, 20(11):3137–3147, 2018. [1](#)
- [14] Hamid Reza Vaezi Joze, Amirreza Shaban, Michael L Iuzzolino, and Kazuhito Koishida. Mmtm: Multimodal transfer module for cnn fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13289–13299, 2020. [7](#)
- [15] Douwe Kiela, Edouard Grave, Armand Joulin, and Tomas Mikolov. Efficient large-scale multi-modal classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. [6](#)
- [16] Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn: Training deep neural networks with biased data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9012–9020, 2019. [6](#)
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [8](#)
- [18] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. [6](#)
- [19] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. [3](#)
- [20] Junnan Li, Pan Zhou, Caiming Xiong, and Steven CH Hoi. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*, 2020. [2](#)
- [21] Dong Liu, Kuan-Ting Lai, Guangnan Ye, Ming-Syan Chen, and Shih-Fu Chang. Sample-specific late fusion for visual category recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 803–810, 2013. [6](#)
- [22] Xinwang Liu, Xinzhong Zhu, Miaomiao Li, Lei Wang, Chang Tang, Jianping Yin, Dinggang Shen, Huaimin Wang, and Wen Gao. Late fusion incomplete multi-view clustering. *IEEE transactions on pattern analysis and machine intelligence*, 41(10):2410–2423, 2018. [3](#)
- [23] Youssef Mroueh, Etienne Marcheret, and Vaibhava Goel. Deep multimodal learning for audio-visual speech recognition. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2130–2134. IEEE, 2015. [1](#)
- [24] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *ICML*, 2011. [1](#)
- [25] Dan Oneață and Horia Cucu. Improving multimodal speech recognition by data augmentation and speech representations. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 4578–4587. IEEE, 2022. [1](#)
- [26] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648, 2018. [6](#)
- [27] Yingwei Pan, Ting Yao, Yehao Li, Yu Wang, Chong-Wah Ngo, and Tao Mei. Transferrable prototypical networks

- for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2239–2247, 2019. [2](#)
- [28] Yagya Raj Pandeya, Bhuwan Bhattarai, and Joonwhoan Lee. Deep-learning-based multimodal emotion classification for music videos. *Sensors*, 21(14):4927, 2021. [1](#)
- [29] Xiaokang Peng, Yake Wei, Andong Deng, Dong Wang, and Di Hu. Balanced multimodal learning via on-the-fly gradient modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8238–8247, 2022. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [30] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. [6](#)
- [31] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951. [6](#)
- [32] Vivek Roy, Yan Xu, Yu-Xiong Wang, Kris Kitani, Ruslan Salakhutdinov, and Martial Hebert. Few-shot learning with intra-class knowledge transfer. *arXiv preprint arXiv:2008.09892*, 2020. [2](#)
- [33] Amir Shahroudy, Tian-Tsong Ng, Yihong Gong, and Gang Wang. Deep multimodal feature analysis for action recognition in rgb+ d videos. *IEEE transactions on pattern analysis and machine intelligence*, 40(5):1045–1058, 2017. [1](#)
- [34] Shashi Kant Shankar, Luis P Prieto, María Jesús Rodríguez-Triana, and Adolfo Ruiz-Calleja. A review of multimodal learning analytics architectures. In *2018 IEEE 18th international conference on advanced learning technologies (ICALT)*, pages 212–214. IEEE, 2018. [1](#)
- [35] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017. [2](#), [4](#)
- [36] Nitish Srivastava and Russ R Salakhutdinov. Multimodal learning with deep boltzmann machines. *Advances in neural information processing systems*, 25, 2012. [1](#)
- [37] Haiman Tian, Yudong Tao, Samira Pouyanfar, Shu-Ching Chen, and Mei-Ling Shyu. Multimodal deep representation learning for video classification. *World Wide Web*, 22(3):1325–1341, 2019. [1](#)
- [38] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event localization in unconstrained videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 247–263, 2018. [1](#), [6](#), [7](#)
- [39] Valentin Vielzeuf, Alexis Lechervy, Stéphane Pateux, and Frédéric Jurie. Centralnet: a multilayer approach for multimodal fusion. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. [7](#)
- [40] Weiyao Wang, Du Tran, and Matt Feiszli. What makes training multi-modal classification networks hard? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12695–12705, 2020. [1](#), [2](#), [6](#)
- [41] Thomas Winterbottom, Sarah Xiao, Alistair McLean, and Noura Al Moubayed. On modality bias in the tvqa dataset. *arXiv preprint arXiv:2012.10210*, 2020. [2](#)
- [42] Nan Wu, Stanislaw Jastrzebski, Kyunghyun Cho, and Krzysztof J Geras. Characterizing and overcoming the greedy nature of learning in multi-modal deep neural networks. In *International Conference on Machine Learning*, pages 24043–24055. PMLR, 2022. [1](#), [2](#), [6](#)
- [43] Yan Xia and Zhou Zhao. Cross-modal background suppression for audio-visual event localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19989–19998, 2022. [1](#)
- [44] Fanyi Xiao, Yong Jae Lee, Kristen Grauman, Jitendra Malik, and Christoph Feichtenhofer. Audiovisual slowfast networks for video recognition. *arXiv preprint arXiv:2001.08740*, 2020. [1](#), [2](#), [6](#)
- [45] Yuzhe Yang, Hao Wang, and Dina Katabi. On multi-domain long-tailed recognition, generalization and beyond. *arXiv preprint arXiv:2203.09513*, 2022. [2](#)
- [46] Litian Zhang, Xiaoming Zhang, and Junshu Pan. Hierarchical cross-modality semantic correlation learning model for multimodal summarization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022. [3](#)