

# Modality-specific Distillation

Woojeong Jin<sup>1\*</sup> Maziar Sanjabi<sup>2</sup> Shaoliang Nie<sup>2</sup> Liang Tan<sup>2</sup> Xiang Ren<sup>1</sup> Hamed Firooz<sup>2</sup>

<sup>1</sup>University of Southern California <sup>2</sup>Facebook AI  
{woojeong.jin, xiangren}@usc.edu  
{maziars, snie, liangtan, mhfirooz}@fb.com

## Abstract

Large neural networks are impractical to deploy on mobile devices due to their heavy computational cost and slow inference. Knowledge distillation (KD) is a technique to reduce the model size while retaining performance by transferring knowledge from a large “teacher” model to a smaller “student” model. However, KD on multimodal datasets such as vision-language datasets is relatively unexplored and digesting such multimodal information is challenging since different modalities present different types of information. In this paper, we propose modality-specific distillation (MSD) to effectively transfer knowledge from a teacher on multimodal datasets. Existing KD approaches can be applied to multimodal setup, but a student doesn’t have access to modality-specific predictions. Our idea aims at mimicking a teacher’s modality-specific predictions by introducing an auxiliary loss term for each modality. Because each modality has different importance for predictions, we also propose weighting approaches for the auxiliary losses; a meta-learning approach to learn the optimal weights on these loss terms. In our experiments, we demonstrate the effectiveness of our MSD and the weighting scheme and show that it achieves better performance than KD.

## 1 Introduction

Recent advances in computer vision and natural language processing are attributed to deep neural networks with large number of layers. Current state-of-the-art architectures are getting wider and deeper with billions of parameters, e.g., BERT (Devlin et al., 2019) and GPT-3 (Brown et al., 2020). In addition to their huge sizes, such wide and deep models suffer from high computational costs and latencies at inference. These shortcomings greatly

<sup>\*</sup>The work in progress was mainly done during internship at Facebook AI.

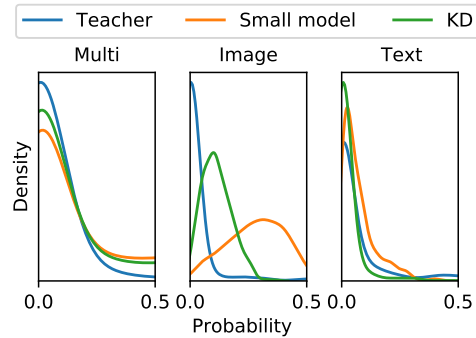


Figure 1: **Density of model outputs on Hateful-Memes:** given multimodality samples as input (Multi), given only image modality as input (Image), and given only text modality as input (Text). KD denotes conventional knowledge distillation and the small model is a model with distillation. We observe that there is still a prediction gap between the teacher and the student trained by KD. To minimize the gap, we propose modality-specific distillation (MSD).

limit these models practicality and make them unsuitable for many mobile applications.

To mitigate the heavy computational cost and the memory requirement, there have been several attempts to compress a larger model (a teacher) into a smaller model (a student) (Ba and Caruana, 2014; Hinton et al., 2015; Romero et al., 2014; Park et al., 2019; Müller et al., 2020). Among them, *knowledge distillation* (KD) (Hinton et al., 2015) assumes the knowledge in the teacher as a learned mapping from inputs to outputs, and transfers the knowledge by training the student model with the teacher’s outputs (of the last or a hidden layer) as targets. Recently, KD has been explored in various studies such as improving a student model (Hinton et al., 2015; Park et al., 2019; Romero et al., 2014; Tian et al., 2019; Müller et al., 2020) and improving a teacher model itself by self-distillation (Xie et al., 2020; Kim et al., 2020; Furlanello et al., 2018).

There has been a surge of interest in distillation in a multimodal setup such as cross-modal

distillation (Gupta et al., 2016; Tian et al., 2019). Multimodal problems involve relating information from multiple sources. For example, visual question answering (VQA) requires answering questions about an image (Antol et al., 2015; Goyal et al., 2017; Gurari et al., 2018; Singh et al., 2019) and models should incorporate information from the text and image sources to answer the questions. Multimodal problems are important because many real-world problems requires understanding signals from different modalities to make accurate predictions; information on the web and social media is often represented as textual and visual description. Digesting such multimodal information in an effective manner is challenging due to their different natures, e.g., visual and textual sources present different types of information. Also, they don't have comparable amounts of information in each modality; usually the textual modality tends to dominate and have more information.

While KD approaches can be applied to multimodal applications, student models in these approaches are directly trained to mimic a teacher's outputs without access to teacher's modality-specific behaviors. As a result, the student and teacher models may significantly differ in their outputs using each modality as input. We illustrate the point in Fig 1. The gap verifies that the student's and the teacher's modality-specific behaviors are not well matched. We hypothesize that it may lead to inefficient distillation, because the student does not carefully mimic the teacher's modality-specific predictions.

Thus, we propose *modality-specific distillation* (MSD) which is to mimic the teacher's modality-specific behavior to minimize the gaps. We improve the transfer by splitting the multimodality into separate modalities, using them as additional inputs, and thus distilling modality-specific behavior of the teacher. Our MSD introduces auxiliary losses per modality to encourage each modality to be distilled effectively; we transfer the modality-specific knowledge from the teacher. Furthermore, we propose weighting approaches for weighting the auxiliary losses to take importance of each modality into account; one modality might have more important information. There are two main strategies to weight these auxiliary losses in the objective; population-based and instance-wise weighting schemes. In the population based, the weight of each loss term is fixed for the whole popula-

tion. But in many cases the samples' modalities might carry different amount of information; one of modalities has more important information for predictions. Thus, we explore an intuitive instance-wise weighting scheme. In the end, we propose a meta-learning approach to find optimal weights.

As we will see in our empirical study on multimodal datasets, MSD significantly improves the performance of student models over KD. Also, our extensive experiments verify that MSD with weighting functions learned by our method shows the best performance among other weighting schemes. In our analysis, we show that datasets are different in the need of population-based or sample-specific weighting; the MM-IMDB dataset, for example, shows less improvement on instance-wise weighting compared to population-based weighting.

## 2 Background

In this section, we first define notations and revisit conventional knowledge distillation (KD).

### 2.1 Problem Definition and Notations

Given a trained and frozen teacher model  $T$  and a student model  $S$ , the output of our task is a trained student model. Our goal is to transfer knowledge from the teacher to the student on multimodal datasets. We let  $f_T$  and  $f_S$  be functions of the teacher and the student, respectively.  $t$  and  $s$  refer to softmax output of the teacher and the student. Typically the models are deep neural networks and the teacher is deeper than the student. The function  $f$  can be defined using output of the last layer of the network (e.g., logits).  $X$  is a multimodal (language-vision) dataset,  $X^t$  refers to only the text modality of  $X$ ,  $X^v$  is refers to only the image modality of  $X$ , and  $x_i$  is a dataset instance. In this work, we focus on one text and one image modalities, but it is easy to extend the work to more/other modalities.

### 2.2 Conventional Knowledge Distillation

In knowledge distillation (Hinton et al., 2015), a student is trained to minimize a weighted sum of two different losses: (a) cross entropy with hard labels (one-hot encodings on correct labels) using a standard softmax function, (b) cross entropy with soft labels (probability distribution of labels) produced by a teacher with a temperature higher than 1 in the softmax of both models. The temperature controls the softness of the probability distributions.

Thus, the loss for the student is defined as:

$$\mathcal{L}_{\text{student}} = \lambda \mathcal{L}_{\text{CE}} + (1 - \lambda) \mathcal{L}_{\text{distill}}, \quad (1)$$

where  $\mathcal{L}_{\text{CE}}$  is a standard cross-entropy loss on hard labels,  $\mathcal{L}_{\text{distill}}$  is a distillation loss, which is a cross-entropy loss on soft labels, and  $\lambda \in [0, 1]$  controls the balance between hard and soft targets.

To be specific, knowledge distillation (Hinton et al., 2015) minimizes Kullback-Leibler divergence between soft targets from a teacher and probabilities from a student. The soft targets (or soft labels) are defined as softmax on outputs of  $f_T$  with temperature  $\tau$ . The distillation loss is defined as follows:

$$\mathcal{L}_{\text{distill}} = \tau^2 \frac{1}{|X|} \sum_{x_i \in X} \text{KL}(t(x_i; \tau), s(x_i; \tau)), \quad (2)$$

where

$$t(x_i; \tau) = \sigma\left(\frac{f_T(x_i)}{\tau}\right), \quad s(x_i; \tau) = \sigma\left(\frac{f_S(x_i)}{\tau}\right), \quad (3)$$

$\sigma$  is a softmax function. The temperature parameter  $\tau$  controls the entropy of the output distribution (higher temperature  $\tau$  means higher entropy in the soft labels). Following (Hinton et al., 2015), we scale the loss by  $\tau^2$  in order to keep gradient magnitudes approximately constant when changing the temperature. We omit  $\tau$  for brevity.

**Limitations.** This KD can be applied to multimodal setups and student models in this distillation are directly trained to mimic a teacher’s outputs without access to teacher’s modality-specific behaviors. As a result, the student and teacher models may significantly differ in their modality-specific outputs, which leads to inefficient distillation. To better mimic the teacher’s behaviors, we propose modality-specific distillation in the next section.

### 3 Proposed Method

In this section, we introduce our proposed approach, modality-specific distillation (MSD) for multimodal datasets.

#### 3.1 Modality-specific Distillation

Samples in multimodal datasets are constructed from multiple modalities such as text modality and image modality. In this work, we focus on vision-language datasets. The core idea of MSD is to feed each modality as a separate input into a teacher and a student, and transfer the modality-specific

knowledge of the teacher to the student. This will minimize the gap between a teacher and its student with regard to individual modalities predictions and thus the student learns more effectively from a teacher. Fig. 2 illustrates comparison between KD and MSD. From this perspective, MSD serves as a data augmentation strategy (Xie et al., 2019b, 2020), where the augmented data is naturally generated from the modalities of the input. Our approach can be viewed as an extension of Cutout (DeVries and Taylor, 2017) that masks out random sections of input images during training while our approach masks out one of the modalities during distillation. Unlike some other data augmentation techniques such as Mixup (Zhang et al., 2017) where the labels for augmented data is generated through simple interpolation, we use the teacher to guide us for setting soft-labels in MSD.

To be specific, we introduce two loss terms,  $\mathcal{L}_{\text{textKD}}$  and  $\mathcal{L}_{\text{imageKD}}$  to minimize difference between probability distributions between the teacher and the student given each modality (assuming text and image as the only two modalities).

$$\mathcal{L}_{\text{textKD}} = \tau^2 \frac{1}{|X_t|} \sum_{x_i \in X_t} \text{KL}(t(x_i), s(x_i)). \quad (4)$$

$\mathcal{L}_{\text{imageKD}}$  is similarly defined; the input is image modality instead.

With above two auxiliary losses, the MSD loss for the student is defined as follows:

$$\begin{aligned} \mathcal{L}_{\text{distill}} = & \sum_{x_i \in X} w_i \text{KL}(t(x_i), s(x_i)) \\ & + \sum_{x_i \in X^v} w_i^v \text{KL}(t(x_i), s(x_i)) + \sum_{x_i \in X^t} w_i^t \text{KL}(t(x_i), s(x_i)), \end{aligned} \quad (5)$$

where we omit the scaling factor  $\tau^2 \frac{1}{|X|}$  for brevity.  $w_i, w_i^t, w_i^v \in [0, 1]$  control the balance between three distillation losses. These weights determine importance of each modality and they affect the student’s performance on multimodal datasets.

**Weighting on Each Modality.** Samples from multimodal datasets have different information on each modality. Fig. 3 shows the teacher model predictions for samples in Hateful-Memes and MM-IMDB test sets. For each sample, three probabilities are calculated: 1) predictions of samples with both of its modalities, 2) predictions of samples with just its text modality, and 3) predictions of samples with just its image modality. As one can

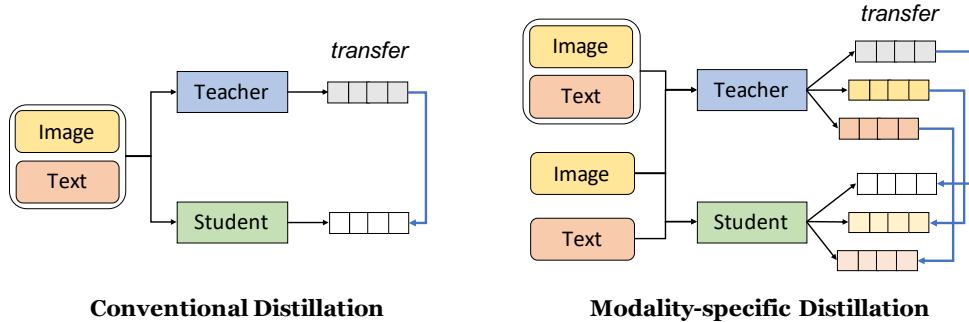


Figure 2: **Comparison between KD and MSD.** In KD, multimodal datasets are taken as a teacher and a student’s inputs to compute the distillation loss. However, we factorize the multimodal input into each modality, and use it as a separate input to a teacher and student.

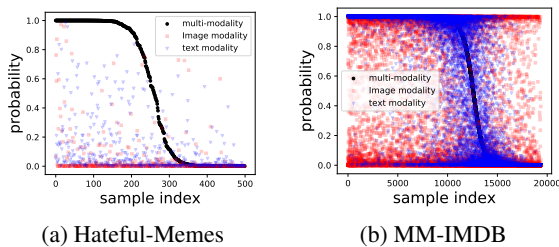


Figure 3: **Prediction probabilities of test samples for different modalities.** Black points correspond to the predictions of samples with both modalities (original input), red points do with image modality, and blue points do with text modality. The samples are ordered based on their multi modal output probabilities.

see for MM-IMDB there is a strong correlation between multimodal predictions and predictions from text modality, indicating the fact that in MM-IMDB text is a dominant modality. On the other hand, for Hateful-Memes dataset there is no such global pattern but one can observe some correlations for individual instances. This behavior is actually expected based on the way Hateful-Memes is built to include unimodal confounders (Kielar et al., 2020).

Following these observations we propose three weighting schemes for distillation losses, presented in the order of complexity: 1) population-based (Section 3.2), 2) instance-wise (Section 3.3) weighting approaches for the losses, and 3) meta-learning approach (Section 3.4) to find the optimal weights on meta data. We will discuss each of these in the following sections.

### 3.2 Population-based Weighting

Population-based weighting is to assign weights depending on modality; we give constant weights ( $w_i, w_i^v, w_i^t$ ) for each loss term in equation (5). This

weighting approach assumes the weights are determined by the types of modality. Best weights or coefficients for each loss term are obtained by grid search on the validation set. However, population-based weighting is limited because it does not assign finer-grained weights to each data instance; each data instance might have different optimal weights for the loss terms. This is what we pursue next in the instance-wise weighting.

### 3.3 Instance-wise Weighting

Instance-wise weighting is to give different weights to each loss term depending on a data sample. The assumption is that each data point has different optimal weights for knowledge distillation. By assigning instance-level weights, we expect a better learning for the student to mimic teacher’s modality-specific behavior. In this sense, population-based weighting can be regarded as one version of instance-wise weighting that assigns weights depending on modality. As it is not possible to tune sample-weights as separate hyper-parameters, we instead propose to use simple/intuitive fixed weighting functions, described as follows. We exploit teacher’s output as the input to these fixed weighting schemes. Obviously, the next step to this approach would be to learn this weighting function alongside the rest of the model, i.e. meta-learning, which we discuss further in the Section 3.4.

**Importance-based weighting.** The idea is to weight each loss term based on the importance of its corresponding modality. To measure the importance of each modality, we compute the change in the output of teacher after dropping the other modality:

$$I_{i,t} = \delta(t(x_i), t(x_i^v)), I_{i,v} = \delta(t(x_i), t(x_i^t)), \quad (6)$$

where  $t(x_i), t(x_i^v), t(x_i^t)$  is teacher’s probabilities (i.e. softmax output), given the multimodal, image alone and text alone inputs, respectively. We use Kullback-Leibler Divergence to measure the difference denoted by  $\delta$ . Thus weights for loss terms are defined as  $w_i^v = g(I_{i,t})$  and  $w_i^t = g(I_{i,v})$ , where  $g = \tanh(\cdot)$  to ensure the weights are in the range  $[0, 1]$ . In this strategy, we assign  $w_i = 1$  for the loss term for multimodality. Note that in this strategy we do not explicitly use the true labels to decide the distillation weights, and we use the teacher’s predictions instead.

**Correctness-based weighting.** Another idea of instance-wise weighting is to weight terms depending on how accurate predictions of the teacher on each modality are. This is to measure the correctness between ground truth and predictions with each modality. If the prediction with one modality is close to the ground truth, then we assign a larger weight to that. To measure the correctness, we adopt cross entropy loss on each instance. We choose the weights proportionally according to the following rule:

$$w_i : w_i^v : w_i^t = 1/h(t(x_i)) : 1/h(t(x_i^v)) : 1/h(t(x_i^t)) \quad (7)$$

where  $h(x) = -\sum_{j=1}^c y_{i,j} \log x$  and  $y_{i,j} \in \{0, 1\}$  are the correct targets for the  $j$ -th class of the  $i$ -th example.  $h(x)$  measures the distance between ground-truth labels and predictions and thus the inverse of  $h(x)$  can represent the correctness of the predictions. In order to choose the actual weights, we add a normalization constraint such that,  $w_i + w_i^v + w_i^t = 1$ . It is worth noting that in this weighting scheme, the actual labels are directly used in deciding the weights unlike the previous one.

### 3.4 Meta Learning for Weights

Although, the aforementioned weighting schemes are intuitive, there is no reason to believe they are the optimal way of getting value out of modality-specific distillation. Moreover, it is not trivial to get optimal weight functions since it can depend on a dataset. Thus, we propose a meta-learning approach to find optimal weight functions. Inspired by (Shu et al., 2019), we design meta learners to find the optimal coefficients.  $(w_i, w_i^v, w_i^t)$  is defined as follows:

$$(w_i, w_i^v, w_i^t) = f(t(x_i), t(x_i^v), t(x_i^t); \Theta) \quad (8)$$

---

### Algorithm 1: Meta-Learning Algorithm

---

**Input:** Training data  $\mathcal{D}$ , Meta-data set  $\hat{\mathcal{D}}$ , batch size  $n, m$ , learning rates  $\alpha, \beta$ , max iterations  $T$ .

- 1 **for**  $t \leftarrow 0$  **to**  $T - 1$  **do**
- 2      $\{x, y\} \leftarrow \text{SampleMiniBatch}(D, n)$ .
- 3      $\{x^{(\text{meta})}, y^{(\text{meta})}\} \leftarrow \text{SampleMiniBatch}(\hat{D}, m)$ .
- 4      $\hat{\mathbf{w}}^{(t)}(\Theta^{(t)}) \leftarrow$   
        $\mathbf{w}^{(t)} - \alpha \frac{1}{n} \sum_{i=1}^n \nabla_{\mathbf{w}} \mathcal{L}_{\text{student}}(\mathbf{w}^{(t)}, \Theta^{(t)})$
- 5      $\Theta^{(t+1)} \leftarrow$   
        $\Theta^{(t)} - \beta \frac{1}{m} \sum_{i=1}^m \nabla_{\Theta} \mathcal{L}_{\text{meta}}(\hat{\mathbf{w}}^{(t)}(\Theta^{(t)}))$
- 6      $\mathbf{w}^{(t+1)} \leftarrow$   
        $\mathbf{w}^{(t)} - \alpha \frac{1}{n} \sum_{i=1}^n \nabla_{\mathbf{w}} \mathcal{L}_{\text{student}}(\mathbf{w}^{(t)}, \Theta^{(t+1)})$
- 7 **return** Network parameters  $\mathbf{w}^{(T)}, \Theta^{(T)}$

---

where  $\Theta$  defines the parameters for the meta learner network, an Multi-Layer Perceptron (MLP) with a sigmoid layer, which approximates a wide range of functions (Csáji et al., 2001). In general, the meta function for defining weights can depend on any input from the sample; but here we limit ourselves to the teacher’s predictions.

**Meta-Learning Objective.** We assume that we have a small amount of unbiased meta-data set  $\{x_i^{(\text{meta})}, y_i^{(\text{meta})}\}_{i=1}^M$ , representing the meta knowledge of ground-truth sample-label distribution, where  $M$  is the number of meta samples and  $M \ll N$ . In our setup, we use the validation set as the meta-data set. The optimal parameter  $\Theta^*$  can be obtained by minimizing the following cross-entropy loss:

$$\begin{aligned} \mathcal{L}_{\text{meta}}(\mathbf{w}^*(\Theta)) \\ = -\frac{1}{M} \sum_{i=1}^M \sum_{j=1}^c y_{i,j} \log(s(x_i; \mathbf{w}^*(\Theta))), \end{aligned} \quad (9)$$

where  $\mathbf{w}^*$  is an optimal student’s parameter, which is defined as follows:

$$\mathbf{w}^*(\Theta) = \arg \min_{\mathbf{w}} \mathcal{L}_{\text{student}}(\mathbf{w}, \Theta). \quad (10)$$

$w^*$  is parameterized by  $\Theta$ , a meta learner’s parameter.

The meta learner is optimized for generating instance weights that minimize the average error of the student over the meta-data set, while the student is trained on the training set with the generated instance weights from the meta learner.

**Meta-Learning Algorithm.** Finding the optimal  $\Theta^*$  and  $w^*$  requires two nested loops; one gradient update of a meta learner requires a trained student on the training set. Thus, we adopt an online strategy following (Shu et al., 2019), which updates

the meta learner with only one gradient update of the student. Algorithm 1 illustrates its learning process. First, we sample mini batches from the training and meta-data sets, respectively (lines 2 and 3). Then, we update the student’s parameter along the descent direction of the student’s loss on a mini-batch training data (line 4). Note that the student’s parameter is parameterized by the meta learner’s parameter. With the updated parameter, the meta learner can be updated by moving the current parameter  $\Theta(t)$  along the objective gradient of equation (9) on a mini-batch meta data (line 5). After updating the meta-learner, the student’s parameter can be updated on a mini-batch training data (line 6).

## 4 Experiments

In this section, we empirically show the effectiveness of our proposed approaches.

### 4.1 Experimental Setup

We use VisualBERT (Li et al., 2019), a pre-trained multimodal model, as the teacher model. For a student model, we use TinyBERT (Jiao et al., 2019). VisualBERT consists of 12 layers and a hidden size of 768, and has 109 million number of parameters, while TinyBERT consists of 4 layers and a hidden size of 312, and has 14.5 million number of parameters. We use the region features from images for both the teacher and the student and fine-tune the student on each dataset. For training the meta learner we use the datasets’ validation set as meta data. We find the best hyperparameters on the validation set. For comparison, we include various knowledge distillation approaches: Conventional KD (Hinton et al., 2015), FitNet (Romero et al., 2014), RKD (Park et al., 2019), and SP (Tung and Mori, 2019). We empirically show that our MSD approaches, i.e. population-based weighting, instance-wise weighting based on importance of each modality and correctness of predictions of each modality, and meta learning, can improve the performance of the small model compared to other KD approaches. Moreover, meta-learning approach provides the closest performance to the teacher model in all three multimodal datasets by finding the optimal weights per sample for MSD.

### 4.2 Datasets

To examine our proposed approaches, we use three multimodal datasets: Hateful-Memes (Kiela et al.,

2020) MM-IMDB (Arevalo et al., 2017), and Visual Entailment (SNLI-VE) (Xie et al., 2019a; Young et al., 2014).

The Hateful-Memes dataset (Kiela et al., 2020) consists of 10K multimodal memes. The task is a binary classification problem, which is to detect hate speech in multimodal memes. We use Accuracy (ACC), and AUC as evaluation metrics of choice for hateful memes (Kiela et al., 2020).

The MM-IMDB (Multimodal IMDB) dataset consists of 26K movie plot outlines and movie posters. The task involves assigning genres to each movie from a list of 23 genres. This is a multi-label prediction problem, i.e., one movie can have multiple genres and we use Macro F1 and Micro F1 as evaluation metrics following (Arevalo et al., 2017).

The goal of Visual Entailment is to predict whether a given image semantically entails an input sentence. Classification accuracy over three classes (“Entailment”, “Neutral” and “Contradiction”) is used to measure model performance. We use accuracy as an evaluation metric following (Xie et al., 2019a).

### 4.3 Results

Table 1 shows our main results on Hateful-Memes, MM-IMDB, and SNLI-VE datasets. The small model refers to the student model without distillation from the teacher. Existing KD approaches improve the student model on all datasets. However, our MSD approaches improve the small model substantially. We observe that among weighting strategies, MSD with meta learning shows the best performance, indicating it finds effective weights for each dataset. We note that population-based weighting shows good improvement, which means weighting based on modality only is still very effective on multimodal datasets. Also, population-based weighting outperforms instance-wise weighting on the MM-IMDB dataset, suggesting all samples are likely to have the same preference or dependency on each modality of the dataset.

In addition, we present improvements over KD approaches with/without our MSD (meta-learning) in Table 2. Here, we use MSD on top of each KD approach. Note that our MSD approach is orthogonal to existing KD approach. The results show the benefits of our MSD method on top of other approaches; MSD improves these KD methods on multimodal datasets.

Table 1: **Main Results.** Mean results ( $\pm$ std) over 5 repetitions are reported. Our MSD outperforms all the KD approaches. Here, we use our MSD on top of conventional KD (Hinton et al., 2015). Also our meta learning for weights shows the best performance.

Method	Hateful-Memes		MM-IMDB		SNLI-VE
	ACC	AUC	Macro F1	Micro F1	ACC
Teacher	65.28	71.82	59.92	66.53	77.57
Small model	60.83 ( $\pm$ 0.20)	65.54 ( $\pm$ 0.25)	38.78 ( $\pm$ 4.03)	58.10 ( $\pm$ 1.23)	72.30 ( $\pm$ 0.35)
Conventional KD (Hinton et al., 2015)	60.84 ( $\pm$ 1.50)	66.53 ( $\pm$ 0.27)	41.76 ( $\pm$ 4.72)	58.96 ( $\pm$ 1.62)	72.61 ( $\pm$ 0.55)
FitNet (Romero et al., 2014)	62.00 ( $\pm$ 0.26)	67.13 ( $\pm$ 0.51)	46.21 ( $\pm$ 2.12)	60.46 ( $\pm$ 0.30)	73.06 ( $\pm$ 0.50)
RKD (Park et al., 2019)	61.43 ( $\pm$ 0.40)	67.03 ( $\pm$ 0.21)	51.16 ( $\pm$ 1.64)	62.52 ( $\pm$ 0.70)	73.09 ( $\pm$ 0.53)
SP (Tung and Mori, 2019)	61.70 ( $\pm$ 1.10)	66.11 ( $\pm$ 0.45)	49.07 ( $\pm$ 0.82)	61.41 ( $\pm$ 0.34)	73.00 ( $\pm$ 0.98)
MSD (Population)	62.15 ( $\pm$ 1.71)	68.16 ( $\pm$ 1.60)	51.85 ( $\pm$ 0.34)	62.13 ( $\pm$ 0.19)	73.64 ( $\pm$ 0.54)
MSD (Instance, Importance)	62.78 ( $\pm$ 1.00)	67.94 ( $\pm$ 0.52)	49.20 ( $\pm$ 1.27)	61.84 ( $\pm$ 0.49)	73.34 ( $\pm$ 0.48)
MSD (Instance, Correctness)	63.27 ( $\pm$ 0.45)	67.72 ( $\pm$ 0.82)	51.02 ( $\pm$ 0.70)	62.05 ( $\pm$ 0.45)	73.52 ( $\pm$ 0.54)
MSD (Meta learning)	<b>63.86 (<math>\pm</math>1.28)</b>	<b>68.30 (<math>\pm</math>0.62)</b>	<b>53.12 (<math>\pm</math>0.08)</b>	<b>63.00 (<math>\pm</math>0.09)</b>	<b>74.04 (<math>\pm</math>0.15)</b>

Table 2: **Improvement over KD approaches with MSD.** Our MSD substantially improves existing KD approaches.

Method	Hateful-Memes		MM-IMDB	
	ACC	AUC	Macro F1	Micro F1
KD (Hinton et al., 2015)	60.84	66.53	41.76	58.96
+MSD	<b>63.86</b>	<b>68.30</b>	<b>53.12</b>	<b>63.00</b>
FitNet (Romero et al., 2014)	62.00	67.13	46.21	60.46
+MSD	<b>62.50</b>	<b>68.77</b>	<b>51.75</b>	<b>62.13</b>
RKD (Park et al., 2019)	61.43	67.03	51.16	62.52
+MSD	<b>63.10</b>	<b>67.58</b>	<b>52.36</b>	<b>63.24</b>
SP (Tung and Mori, 2019)	61.70	66.11	49.07	61.41
+MSD	<b>62.30</b>	<b>67.92</b>	<b>52.83</b>	<b>62.70</b>

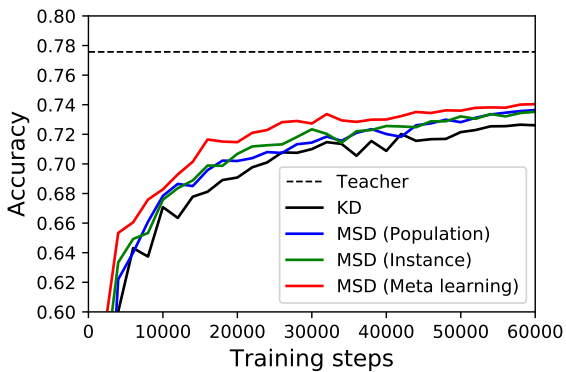


Figure 4: Test accuracy of a student on SNLI-VE during training and comparison between knowledge distillation (KD) and modality-specific distillation (MSD) with population-based weighting, instance-wise weighting, and meta learning for weights.

#### 4.4 Learning Curve

Our proposed MSD approaches can also help with training speed, measured by test metrics over training steps. Fig 4 shows the evolution of accuracy on the *test set* during training on the SNLI-VE dataset. When we train the student with MSD, training progresses faster than KD. Since the teacher provides two additional probabilities with each modality, the

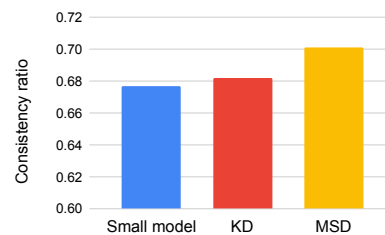


Figure 5: **Teacher-Student consistency ratio.** We investigate the student model’s sensitiveness to changes in modalities. Higher ratio indicates its sensitiveness is closer to the teacher’s.

student learns faster and the final performance is better than KD. We observe a large performance increase early in training with the meta-learning approach, thus leading to the best accuracy. In this case, the meta learning for sample weighting finds the optimal weights for each data instance, so the student quickly learns from more important modality that is vital for the predictions.

#### 4.5 Analysis

In this section, we empirically investigate the benefits of our approach by analyzing MSD.

##### 4.5.1 Teacher-Student Consistency

To showcase that our approach helps the student model to be more sensitive to important changes in modalities, we take examples from the Hateful Memes test set and randomly replace one of the modalities with a modality from another sample. Hateful Memes is a multimodal dataset and changing the modalities might or might not change the final label. In this case, we do not have the ground truth, but we use the teacher’s predicted label on the new generated sample as a proxy for ground truth and count the times that the student/small

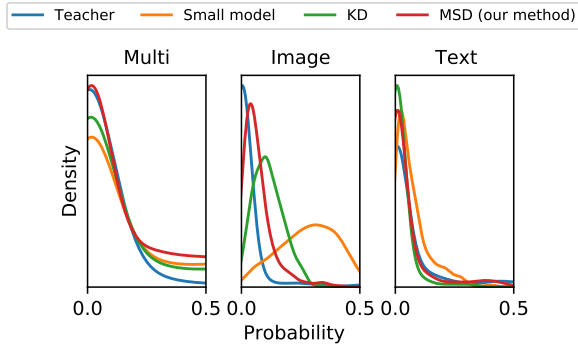


Figure 6: **Density of model outputs on samples of label 0 on the test set of Hateful-Memes:** given multimodal samples as input (Multi), given only image modality as input (Image), and given only text modality as input (Text). Our proposed approach, MSD with the meta-learning approach, minimizes the gap between the teacher and the student trained by KD.

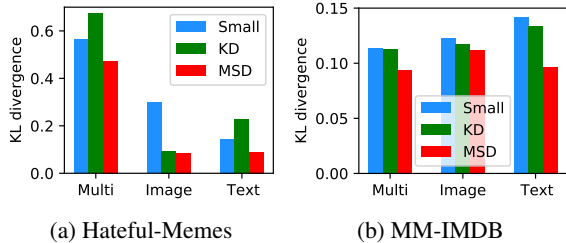


Figure 7: **Kullback-Leibler divergence on the test set between the teacher's outputs and other models' outputs.** This is a measure of how teacher's probability distribution is different from other models'. The lower divergence is, the closer a model is to the teacher.

model is consistent with the teacher on these generated samples. We define the ratio of such consistent predictions over the total generated samples as “**Teacher-Student consistency ratio**”. Note that none of the models have seen these samples during the training. As it can be seen from Fig. 5, our MSD approach has a larger “Teacher-Student consistency ratio” than small model with and without KD. This indicates that MSD not only improves the accuracy but also improves the sensitivity of the student model to better match the teacher on the changes in modalities on unseen data.

#### 4.5.2 Probability Distribution of Model Outputs

There is a performance gap between the teacher model and student model in predicting true labels given a multimodal sample and each of its individual modalities. Fig 6 shows this gap for Hateful-Memes dataset. For example, given only image

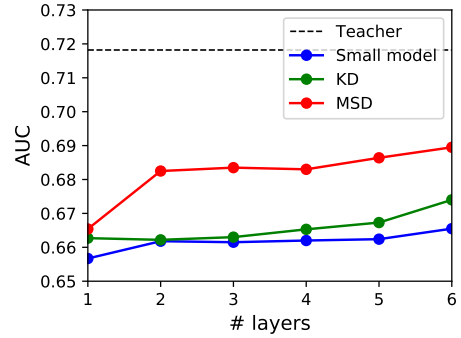


Figure 8: **Test AUC on Hateful-Memes with varying the number of layers in the student model.**

modality as input (the middle plot in Fig 6), there is a considerable difference between the teacher and the small model for predicting benign samples. KD minimizes the gap and our MSD with the meta-learning approach shows the similar density curve to the teacher's.

In addition, we measure Kullback-Leibler (KL) divergence between the teacher's outputs and other models' outputs on the test set as shown in Fig 7. This is to measure the difference between teacher's probability distribution and others'. As is shown, our MSD approach shows the smallest KL divergence from the teacher which means the student learned with MSD outputs probability distribution close to the teacher's.

#### 4.5.3 Student Model Size

To examine how the size of student model affects the performance, we evaluate the baselines and our MSD method on the Hateful-Memes dataset, with varying number of layers in the student model. The result is depicted in Fig 8. In this case, the number of layers is proportional to the number of parameters, i.e. student model size. We use the meta-learning weighting as our MSD method of choice here. As is shown, we observe that the AUC score improves as the model size is getting larger. Also the improvement of KD over the small model is marginal and MSD significantly outperforms KD in any number of layers in the student.

## 5 Related Work

**Knowledge Distillation.** There have been several studies of transferring knowledge from one model to another (Ba and Caruana, 2014; Hinton et al., 2015; Romero et al., 2014; Park et al., 2019; Müller et al., 2020; Tian et al., 2019; Furlanello et al., 2018; Kim et al., 2020). Ba and Caruana (Ba and



Caruana, 2014) improve the accuracy of a shallow neural network by training it to mimic a deep neural network with penalizing the difference of logits between the two networks. Hinton et al. (Hinton et al., 2015) introduced knowledge distillation (KD) that trains a student model with the objective of matching the softmax distribution of a teacher model at the output layer. Romero et al. (Romero et al., 2014) distill a teacher using additional linear projection layers and minimize  $L_2$  loss at the earlier layers to train a students. Park et al. (Park et al., 2019) focused on mutual relations of data examples instead and they proposed relational knowledge distillation. The transfer works best when there are many possible classes because more information can be transferred, but in cases where there are only a few possible classes the transfer is less effective. To deal with the problem, Müller et al. (Müller et al., 2020) improved the transfer by forcing the teacher to divide each class into many subclasses. Tian et al. (Tian et al., 2019) proposed to distill from the penultimate layer using a contrastive loss for cross-modal transfer. A few recent papers (Furlanello et al., 2018; Kim et al., 2020) have shown that distilling a teacher model into a student model of identical architecture, i.e., self-distillation, can improve the student over the teacher.

**Meta Learning for Sample Weighting.** Recently, some methods were proposed to learn an adaptive weighting scheme from data to make the learning more automatic and reliable including Meta-Weight-Net (Shu et al., 2019), learning to reweight (Ren et al., 2018), FWL (Dehghani et al., 2017), MentorNet (Jiang et al., 2018), and learning to teach (Fan et al., 2018; Wu et al., 2018; Fan et al., 2020). These approaches were proposed to deal with noisy and corrupted labels and learn optimal functions from clean datasets. They are different in that they adopt different weight functions such as a multilayer perceptron (Shu et al., 2019), Bayesian function approximator (Dehghani et al., 2017), and a bidirectional LSTM (Jiang et al., 2018); and they take different inputs such as loss values and sample features. In our case, we adopt these ideas of meta-learning, and specifically Meta-Weight0Net, and utilize it in a different context, i.e. multimodal knowledge distillation.

**Bias in Multimodal Datasets.** Different multimodal datasets were proposed to study whether a model uses a single modality’s features and the im-

plications for its generalization properties (Agrawal et al., 2018). Different approaches were proposed to deal with such problems where the model overfits to a single modality. Wang et al. (Wang et al., 2020) suggest to regularize the overfitting behavior to different modalities. REPAIR (Li and Vasconcelos, 2019) prevents a model from dataset biases by re-sampling the training data. Cadene et al. (Cadene et al., 2019) proposed RUBi that uses a bias-only branch in addition to a base model during training to overcome language priors. In our study, although we do not directly deal with the overfitting phenomena, we use different weighting schemes to better transfer the modality specific information from the teacher to the student.

## 6 Conclusion

We studied knowledge distillation on multimodal datasets; we observed that conventional KD may lead to inefficient distillation since a student model does not fully mimic a teacher’s modality-specific predictions. To better transfer knowledge from a teacher on the multimodal datasets, we proposed modality-specific distillation; the student mimics the teacher’s outputs on each modality. Furthermore, we proposed weighting approaches, population-based and instance-wise weighting schemes, and a meta-learning approach for weighting the auxiliary losses to take importance of each modality into consideration. We empirically showed that we can improve the student’s performance with our modality-specific distillation compared to conventional distillation. Our MSD approach results on modality specific outputs that better resemble the teacher’s outputs. We showed that the results hold across different student sizes. Moreover, our meta-learning approach is flexible enough to find different effective weighting schemes, depending on the dataset.

## References

- Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Anirudha Kembhavi. 2018. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4971–4980.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- John Arevalo, Tamar Solorio, Manuel Montes-y Gómez,

- and Fabio A González. 2017. Gated multimodal units for information fusion. *arXiv preprint arXiv:1702.01992*.
- Jimmy Ba and Rich Caruana. 2014. Do deep nets really need to be deep? In *Advances in neural information processing systems*, pages 2654–2662.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. 2019. Rubi: Reducing unimodal biases for visual question answering. In *Advances in neural information processing systems*, pages 841–852.
- Balázs Csanád Csáji et al. 2001. Approximation with artificial neural networks. *Faculty of Sciences, Eötvös Loránd University, Hungary*, 24(48):7.
- Mostafa Dehghani, Arash Mehrjou, Stephan Gouws, Jaap Kamps, and Bernhard Schölkopf. 2017. Fidelity-weighted learning. *arXiv preprint arXiv:1711.02799*.
- J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Terrance DeVries and Graham W Taylor. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*.
- Yang Fan, Fei Tian, Tao Qin, Xiang-Yang Li, and Tie-Yan Liu. 2018. Learning to teach. *arXiv preprint arXiv:1805.03643*.
- Yang Fan, Yingce Xia, Lijun Wu, Shufang Xie, Weiqing Liu, Jiang Bian, Tao Qin, Xiang-Yang Li, and Tie-Yan Liu. 2020. Learning to teach with deep interactions. *arXiv preprint arXiv:2007.04649*.
- T. Furlanello, Zachary Chase Lipton, Michael Tschannen, L. Itti, and Anima Anandkumar. 2018. Born again neural networks. In *ICML*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. 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.
- Saurabh Gupta, Judy Hoffman, and Jitendra Malik. 2016. Cross modal distillation for supervision transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2827–2836.
- Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3608–3617.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. 2018. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, pages 2304–2313.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2019. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*.
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. The hateful memes challenge: Detecting hate speech in multimodal memes. *arXiv preprint arXiv:2005.04790*.
- Kyungyul Kim, ByeongMoon Ji, Doyoung Yoon, and Sangheum Hwang. 2020. Self-knowledge distillation: A simple way for better generalization. *arXiv preprint arXiv:2006.12000*.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Yi Li and Nuno Vasconcelos. 2019. Repair: Removing representation bias by dataset resampling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9572–9581.
- Rafael Müller, Simon Kornblith, and Geoffrey Hinton. 2020. Subclass distillation. *arXiv preprint arXiv:2002.03936*.
- Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. 2019. Relational knowledge distillation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3967–3976.
- Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. 2018. Learning to reweight examples for robust deep learning. *arXiv preprint arXiv:1803.09050*.
- Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*.
- Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. 2019. Meta-weight-net: Learning an explicit mapping for sample weighting. In *Advances in Neural Information Processing Systems*, pages 1919–1930.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards vqa models that can read. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8317–8326.
- Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2019. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*.
- Frederick Tung and Greg Mori. 2019. Similarity-preserving knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1365–1374.
- Weiyao Wang, Du Tran, and Matt Feiszli. 2020. What makes training multi-modal classification networks hard? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12695–12705.

- Lijun Wu, Fei Tian, Yingce Xia, Yang Fan, Tao Qin, Lai Jian-Huang, and Tie-Yan Liu. 2018. Learning to teach with dynamic loss functions. In *Advances in Neural Information Processing Systems*, pages 6466–6477.
- Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019a. Visual entailment: A novel task for fine-grained image understanding. *arXiv preprint arXiv:1901.06706*.
- Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. 2019b. Unsupervised data augmentation for consistency training. *arXiv preprint arXiv:1904.12848*.
- Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. 2020. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10687–10698.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2017. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*.



Figure 9: A multimodal violating sample (Left). We further replaced its image modality with a background picture that makes it benign and examined models on both examples (Right).

## A Case Study

We demonstrate the motivation behind our work through an example. Fig. 9 shows an example of a multimodal sample from Hateful Memes test dataset. The sample is violating based on both modalities together, and all models correctly predict that. To further probe the models, we replace the background image of the sample with a picture that makes the label benign. On this artificially generated sample we notice that only the teacher and MSD model correctly predict benign, while the other two models make wrong predictions (presumably by just looking at the text only).

## B Hyperparameters

The teacher model is a VisualBERT (Li et al., 2019), and the student model is TinyBERT (Jiao et al., 2019). We used the MMF library and pre-trained checkpoints from it for VisualBERT<sup>1</sup> and used a pre-trained checkpoint in TinyBERT<sup>2</sup>. VisualBERT consists of 12 layers and a hidden size of 768, and has 109 million number of parameters, while TinyBERT consists of 4 layers and a hidden size of 312, and has 14.5 million number of parameters. For all experiments, we performed a grid search to find the best hyperparameters. We adopt the AdamW optimizer to train networks. We use a linear learning rate schedule that drops to 0 at the end of training with warmup steps of 10% maximum iterations.

**Hateful-Memes.** We performed a grid search over learning rates (1e-5, 3e-5, **5e-5**, 1e-4), and temperatures (1, 2, **4**, 8), and, batch sizes (**10**, 20, 30, 40, 50, 60), and the meta learner’s learning rates (1e-1, 1e-2, **1e-3**, 1e-4). We set the maximum number

<sup>1</sup><https://mmf.sh>

<sup>2</sup><https://github.com/huawei-noah/Pretrained-Language-Model/tree/master/TinyBERT>

Table 3: Dataset Statistics.

Stat. \ Data	Hateful-Memes	MM-IMDB	SNLI-VE
Type	Binary	Multilabel	Multi-class
# Classes	2	23	3
# Examples	10,000	25,959	565,286
# Training	8,500	15,552	529,527
# Validation	500	2,608	17,858
# Test	1,000	7,799	17,901

of iterations to 5000. The balance parameter  $\lambda$  between cross entropy and distillation is set among (0.2, 0.4, **0.5**, 0.6, 0.8).

**MM-IMDB.** For MM-IMDB experiments, we follow a similar procedure, a grid search, to the Hateful-Memes. The batch size is 20, temperature is 1, and the meta learner’s learning rate is 1e-4. We set the maximum number of iterations to 10000. The balance parameter  $\lambda$  is set to 0.5.

**SNLI-VE.** For Visual Entailment (SNLI-VE), the batch size is 64, temperature is 4, and the meta learner’s learning rate is 1e-4. We set the maximum number of iterations to 60000. The balance parameter  $\lambda$  is set to 0.6.