

Probabilistic Embeddings for Cross-Modal Retrieval

Sanghyuk Chun¹ Seong Joon Oh¹ Rafael Sampaio de Rezende² Yannis Kalantidis² Diane Larlus²

¹NAVER AI LAB ²NAVER LABS Europe

Abstract

Cross-modal retrieval methods build a common representation space for samples from multiple modalities, typically from the vision and the language domains. For images and their captions, the multiplicity of the correspondences makes the task particularly challenging. Given an image (respectively a caption), there are multiple captions (respectively images) that equally make sense. In this paper, we argue that deterministic functions are not sufficiently powerful to capture such one-to-many correspondences. Instead, we propose to use Probabilistic Cross-Modal Embedding (PCME), where samples from the different modalities are represented as probabilistic distributions in the common embedding space. Since common benchmarks such as COCO suffer from non-exhaustive annotations for cross-modal matches, we propose to additionally evaluate retrieval on the CUB dataset, a smaller yet clean database where all possible image-caption pairs are annotated. We extensively ablate PCME and demonstrate that it not only improves the retrieval performance over its deterministic counterpart, but also provides uncertainty estimates that render the embeddings more interpretable.

1. Introduction

Given a query and a database from different modalities, cross-modal retrieval is the task of retrieving the database items which are most relevant to the query. Most research on this topic has focused on the image and text modalities [6, 10, 27, 58]. Typically, methods estimate embedding functions that map visual and textual inputs into a common embedding space, such that the cross-modal retrieval task boils down to the familiar nearest neighbour retrieval task in a Euclidean space [10, 52].

Building a common representation space for multiple modalities is challenging. Consider an image with a group of people on a platform preparing to board a train (Figure 1). There are more than one possible caption describing this image. “People waiting to board a train in a train platform” and “The metro train has pulled into a large station” were




Figure 1. We propose to use **probabilistic embeddings** to represent images and their captions as probability distributions in a common embedding space suited for cross-modal retrieval. These distributions successfully model the uncertainty which results from the multiplicity of concepts appearing in a visual scene and implicitly perform many-to-many matching between those concepts.

two of the choices from the COCO [6] annotators. Thus, the common representation has to deal with the fact that an image potentially matches with a number of different captions. Conversely, given a caption, there may be multiple manifestations of the caption in visual forms. The multiplicity of correspondences across image-text pairs stems in part from the different natures of the modalities. All the different components of a visual scene are thoroughly and passively captured in a photograph, while language descriptions are the product of conscious choices of the key relevant concepts to report from a given scene. All in all, a common representation for image and text modalities is required to model the one-to-many mappings in both directions.

Standard approaches which rely on vanilla functions do not meet this necessary condition: they can only quantify one-to-one relationships [10, 52]. There have been attempts

to introduce multiplicity. For example, Song & Soleymani [47] have introduced the Polysemous Visual-Semantic Embeddings (PVSE) by letting an embedding function propose K candidate representations for a given input. PVSE has been shown to successfully capture the multiplicity in the task and improve over the baseline built upon one-to-one functions. Others [27] have computed region embeddings obtained with a pretrained object detector, such as Faster-RCNN [42], establishing multiple region-word matches. This strategy has led to significant performance gains at the expense of a significant increase in computational cost.

In this work, we propose **Probabilistic Cross-Modal Embedding (PCME)**. We argue that a probabilistic mapping is an effective representation tool that does not require an explicit many-to-many representation as is done by detection-based approaches, and further offers a number of advantages. First, PCME yields uncertainty estimates that lead to useful applications like estimating the difficulty or chance of failure for each query. Second, the probabilistic representation leads to a richer embedding space where set algebras make sense, whereas deterministic ones can only represent similarity relations. Third, PCME is complementary to the deterministic retrieval systems.

As harmful the assumption of one-to-one correspondence is for the method, the same assumption have introduced confusion in the evaluation benchmarks. For example, MS COCO [6] suffers from non-exhaustive annotations for cross-modal matches. The best solution would be to explicitly and manually annotate all image-caption pairs in a cross-modal dataset for evaluation. Unfortunately this process does not scale, especially for a large-scale dataset like COCO. Instead, we propose a smaller yet cleaner cross-modal retrieval benchmark based on CUB [56], as well as more sensible evaluation metrics.

Our contributions are as follows. (1) We propose Probabilistic Cross-Modal Embedding (PCME) to properly represent the one-to-many relationships in joint embedding spaces for cross-modal retrieval. (2) We identify shortcomings with existing cross-modal retrieval benchmarks and propose alternative solutions. (3) We analyse the joint embedding space using the uncertainty estimates provided by PCME and show how intuitive properties arise.

2. Related work

Cross-modal retrieval. In this work, we are interested in image-to-text and text-to-image cross-modal retrieval. Much research is dedicated to learn a metric space that jointly embeds images and sentences [9, 10, 11, 20, 27, 47, 49]. Early works [12, 25] relied on Canonical Correlation Analysis (CCA) [14] to build embedding spaces as linear projections between the two representation spaces. Frome *et al.* [11] use a hinge rank loss for triplets built from both modalities. Wang *et al.* [52] expand on this idea by also

training on uni-modal triplets to preserve the structure inherent for each modality in the joint space. Faghri *et al.* [10] propose to learn such space with a triplet loss, and only sample the hardest negative from a training batch with respect to a query-positive pair.

One of the drawbacks of relying on a single global representation is its inability to represent the diversity of semantic concepts present in an image or in a caption. Prior work [17, 55] observed a split between *one-to-one* and *many-to-many* matching in visual-semantic embedding spaces characterized by the use of one or several embedding representations per image or caption. Song & Soleymani [47] build many global representations for each image or sentence by using a multi-head self-attention on local descriptors to produce distinct attention maps. Other methods use region-level and word-level descriptors to build a global image-to-text similarity from many-to-many matching. Li *et al.* [27] employ a Graphical Convolutional Network (GCN) [24] for semantic reasoning of region proposals obtained from a Faster-RCNN [42].

Recently, the most successful way of addressing many-to-many image-to-sentence matching is through joint visual and textual reasoning modules appended on top of separate region-level encoders [26, 30, 32, 33, 36, 54, 55, 60]. Most such methods involve cross-modal attention networks and report state-of-the-art results on cross-modal retrieval. This, however, comes with a large increase in computational cost at test time: pairs formed by the query and every database entry need to go through the reasoning module. Focusing on the scalability, we choose to build on top of approaches that directly utilize the joint embedding space and are compatible with large-scale indexing.

Probabilistic embedding. Probabilistic representations of data have a long history in machine learning [34]; they were introduced in 2014 for word embeddings [51], as they gracefully handle the inherent hierarchies in language, and, since then, a line of research has explored different distribution families for word representations [28, 37, 38]. Recently, probabilistic embeddings have been introduced for computer vision tasks. Oh *et al.* [39] proposed the Hedged Instance Embedding (HIB) to handle the one-to-many correspondences for metric learning, while other works apply probabilistic embeddings to face understanding [45, 4], 2D-to-3D pose estimation [48], speaker diarization [46], and prototype embeddings [44]. Our work extends HIB to joint cross-modal embeddings between images and captions, in order to represent the different levels of granularities in the two domains and implicitly capture the resulting one-to-many associations. Recently Schönfeld *et al.* [43] utilized Variational Autoencoders [23] for zero-shot recognition. Their latent space is conceptually similar to ours, but is learned and used in very different ways: they simply use a 2-Wasserstein distance as their distribution alignment loss

and learn classifiers on top, while PCME uses a probabilistic *contrastive* loss that enables us to use the latent features directly for retrieval. To our knowledge, PCME is the first work that uses probabilistic embeddings for multi-modal retrieval.

3. Method

In this section, we present our **Probabilistic Cross-Modal Embedding (PCME)** framework and discuss its conceptual workings and advantages.

We first define the cross-modal retrieval task. Let $\mathcal{D} = (\mathcal{C}, \mathcal{I})$ denote a vision and language dataset, where \mathcal{I} is a set of images and \mathcal{C} a set of captions. The two sets are connected via ground-truth matches. For a caption $c \in \mathcal{C}$ (respectively an image $i \in \mathcal{I}$), the set of corresponding images (respectively captions) is given by $\tau(c) \subseteq \mathcal{I}$ (respectively $\tau(i) \subseteq \mathcal{C}$). Note that for every query q , there may be multiple cross-modal matches ($|\tau(q)| > 1$). Handling this multiplicity will be the central focus of our study.

Cross-modal retrieval methods typically learn an embedding space \mathbb{R}^D such that we can quantify the subjective notion of “similarity” into the distance between two vectors. For this, two embedding functions $f_{\mathcal{V}}, f_{\mathcal{T}}$ are learned to map image and text samples into the common space \mathbb{R}^D .

3.1. Building blocks for PCME

We introduce two key ingredients for PCME: joint visual-textual embeddings and probabilistic embedding.

3.1.1 Joint visual-textual embeddings

We describe how we learn visual and textual encoders. We then present a previous attempt at addressing the multiplicity of cross-modal associations.

Visual encoder $f_{\mathcal{V}}$. We use the ResNet image encoder [15]. Let $z_v = g_{\mathcal{V}}(i) : \mathcal{I} \rightarrow \mathbb{R}^{h \times w \times d_v}$ denote the output before the global average pooling (GAP) layer. Visual embedding is computed via $v = h_{\mathcal{V}}(z_v) \in \mathbb{R}^D$ where in the simplest case $h_{\mathcal{V}}$ is the GAP followed by a linear layer. We modify $h_{\mathcal{V}}$ to let it predict a distribution, rather than a point.

Textual encoder $f_{\mathcal{T}}$. Given a caption c , we build the array of word-level descriptors $z_t = g_{\mathcal{T}}(c) \in \mathbb{R}^{L(c) \times d_t}$, where $L(c)$ is the number of words in c . We use the pre-trained GloVe [40]. The sentence-level feature t is given by a bidirectional GRU [7]: $t = h_{\mathcal{T}}(z_t)$ on top of the GloVe features.

Losses used in prior work. The joint embeddings are often learned with a contrastive or triplet loss [10, 11].

Polysemous visual-semantic embeddings (PVSE) [47] are designed to model one-to-many matches for cross-modal retrieval. PVSE adopts a multi-head attention block on top of the visual and textual features to encode K possible embeddings per modality. For the visual case, each visual embedding $v^k \in \mathbb{R}^D$ for $k \in \{1, \dots, K\}$ is

given by: $v^k = \text{LN}(h_{\mathcal{V}}(z_v) + s(w^1 \text{att}_{\mathcal{V}}^k(z_v)z_v))$, where $w^1 \in \mathbb{R}^{d_v \times D}$ are the weights of fully connected layers, s is the sigmoid function and LN is the LayerNorm [1]. $\text{att}_{\mathcal{V}}^k$ denotes the k -th attention head of the visual self-attention $\text{att}_{\mathcal{V}}$. Textual embeddings t^k for $k \in \{1, \dots, K\}$ are given symmetrically by the multi-head attention: $t^k = \text{LN}(h_{\mathcal{T}}(z_t) + s(w^2 \text{att}_{\mathcal{T}}^k(z_t)z_t))$. PVSE learns the visual and textual encoders with the multiple instance learning (MIL) objective, where only the best best pair among the K^2 possible visual-textual embedding pairs is supervised.

3.1.2 Probabilistic embeddings for single modality

Our PCME models each sample as a distribution. It builds on the Hedged Instance Embeddings (HIB) [39], a single-modality methodology developed for representing instances as a distribution. HIB is the probabilistic analogue of the contrastive loss [13]. HIB trains a probabilistic mapping $p_{\theta}(z|x)$ that not only preserves the pairwise semantic similarities but also represents the inherent uncertainty in data. We describe the key components of HIB here.

Soft contrastive loss. To train $p_{\theta}(z|x)$ to capture pairwise similarities, HIB formulates a soft version of the contrastive loss [13] widely used for training deep metric embeddings. For a pair of samples (x_{α}, x_{β}) , the loss is defined as:

$$\mathcal{L}_{\alpha\beta}(\theta) = \begin{cases} -\log p_{\theta}(m|x_{\alpha}, x_{\beta}) & \text{if } \alpha, \beta \text{ is a match} \\ -\log(1 - p_{\theta}(m|x_{\alpha}, x_{\beta})) & \text{otherwise} \end{cases} \quad (1)$$

where $p_{\theta}(m|x_{\alpha}, x_{\beta})$ is the *match probability*.

Factorizing match probability. [39] has factorized $p_{\theta}(m|x_{\alpha}, x_{\beta})$ into the match probability based on the embeddings $p(m|z_{\alpha}, z_{\beta})$ and the encoders $p_{\theta}(z|x)$. This is done via Monte-Carlo estimation:

$$p_{\theta}(m|x_{\alpha}, x_{\beta}) \approx \frac{1}{J^2} \sum_j \sum_{j'} p(m|z_{\alpha}^j, z_{\beta}^{j'}) \quad (2)$$

where z^j are samples from the embedding distribution $p_{\theta}(z|x)$. For the gradient to flow, the embedding distribution should be reparametrization-trick-friendly [22].

Match probability from Euclidean distances. We compute the sample-wise match probability as follows:

$$p(m|z_{\alpha}, z_{\beta}) = s(-a\|z_{\alpha} - z_{\beta}\|_2 + b) \quad (3)$$

where (a, b) are learnable scalars and $s(\cdot)$ is sigmoid.

3.2. Probabilistic cross-modal embedding (PCME)

We describe how we learn a joint embedding space that allows for probabilistic representation with PCME.




Figure 2. **Method overview.** The visual and textual encoders for Probabilistic Cross-Modal Embedding (PCME) are shown. Each modality outputs mean and variance vectors in \mathbb{R}^D , which represent a normal distribution in \mathbb{R}^D .




Figure 3. **Head modules.** The visual and textual heads (h_V , h_T) share the same structure, except for the modality-specific modules. The mean and variance computations differ: variance module does not involve sigmoid $s(\cdot)$, LayerNorm (LN), and L2 projection.

3.2.1 Model architecture

An overview of PCME is shown in Figure 2. PCME represents an image i and caption c as normal distributions, $p(v|i)$ and $p(t|c)$ respectively, over the same embedding space \mathbb{R}^D . We parametrize the normal distributions with mean vectors and diagonal covariance matrices in \mathbb{R}^D :

$$\begin{aligned} p(v|i) &\sim N(h_V^\mu(z_v), \text{diag}(h_V^\sigma(z_v))) \\ p(t|c) &\sim N(h_T^\mu(z_t), \text{diag}(h_T^\sigma(z_t))) \end{aligned} \quad (4)$$

where $z_v = g_V(i)$ is the feature map and $z_t = g_T(c)$ is the feature sequence (§3.1.1). For each modality, two head modules, h^μ and h^σ , compute the mean and variance vectors, respectively. They are described soon.

Local attention branch. Inspired by the PVSE architecture (§3.1.1, [47]), we consider appending a *local attention branch* in the head modules (h^μ , h^σ) both for image and

caption encoders. See Figure 3 for the specifics. The local attention branch consists of a self-attention based aggregation of spatial features, followed by a linear layer with a sigmoid activation function. We will show with ablative studies that the additional branch helps aggregating spatial features more effectively, leading to improved performance. **Module for μ versus σ .** Figure 3 shows the head modules h^μ and h^σ , respectively. For h_V^μ and h_T^μ , we apply sigmoid in the local attention branch and add the residual output. In turn, LayerNorm (LN) [1] and L2 projection operations are applied [47, 50]. For h_V^σ and h_T^σ , we observe that the sigmoid and LN operations overly restrict the representation, resulting in poor uncertainty estimations (discussed in §D). We thus do not use sigmoid, LN, and L2 projection for the uncertainty modules.

Soft cross-modal contrastive loss. Learning the joint probabilistic embedding is to learn the parameters for the mappings $p(v|i) = p_{\theta_v}(v|i)$ and $p(t|c) = p_{\theta_t}(t|c)$. We adopt the probabilistic embedding loss in Equation (1), where the match probabilities are now based on the cross-modal pairs (i, c) : $\mathcal{L}_{\text{emb}}(\theta_v, \theta_t; i, c)$, where $\theta = (\theta_v, \theta_t)$ are parameters for visual and textual encoders, respectively. The match probability is now defined upon the visual and textual features: $p_\theta(m|i, c) \approx \frac{1}{\sqrt{2}} \sum_j^J \sum_{j'}^{J'} s(-a\|v^j - t^{j'}\|_2 + b)$ where v^j and $t^{j'}$ follow the distribution in Equation (4).

Additional regularization techniques. We consider two additional loss functions to regularize the learned uncertainty. Following [39], we prevent the learned variances from collapsing to zero by introducing the KL divergence loss between the learned distributions and the standard normal $\mathcal{N}(0, I)$. We also employ the *uniformity loss* that was recently introduced in [53], computed between all embeddings in the minibatch. See §A.1 for more details.

Sampling SGD mini-batch. We start by sampling B ground-truth image-caption matching pairs $(i, c) \in \mathcal{G}$. Within the sampled subset, we consider *every* positive and negative pair dictated by the ground truth matches. This would amount to B matching pairs and $B(B-1)$ non-matching pairs in our mini-batch.

Measuring instance-wise uncertainty. The covariance matrix predicted for each input represents the inherent uncertainty for the data. For a scalar uncertainty measure, we

take the determinant of the covariance matrix, or equivalently the geometric mean of the σ 's. Intuitively, this measures the volume of the distribution.

3.2.2 How does our loss handle multiplicity, really?

We perform a gradient analysis to study how our loss in Equation (1) handles multiplicity in cross-modal matches and learn uncertainties in data. In §A.2, we further make connections with the MIL loss used by PVSE (§3.1.1, [47]).

We first define the distance logit: $l_{jj'} := -a\|v^j - t^{j'}\|_2 + b$ and compare the amount of supervision with different (j, j') values. To see this, take the gradient on $l_{jj'}$.

$$\frac{\partial \mathcal{L}_{\text{emb}}}{\partial l_{jj'}} = \begin{cases} w_{jj'} \cdot (1 - s(l_{jj'})) & \text{for positive match} \\ -w_{jj'} \cdot s(l_{jj'}) & \text{for negative match} \end{cases} \quad (5)$$

$$w_{jj'} := \frac{e^{\pm l_{jj'}}}{\sum_{\alpha\alpha'} e^{\pm l_{\alpha\alpha'}}}$$


where \pm is the positivity of match.

We first observe that if $w_{jj'} = 1$, then Equation (5) is exactly the supervision from the soft contrastive loss (Equation (1)). Thus, it is the term $w_{jj'}$ that let the model learn multiplicity and represent associated uncertainty.

To study the behavior of $w_{jj'}$, first assume that (v, t) is a positive pair. Then, $w_{jj'}$ is the softmax over the pairwise logits $l_{jj'}$. Thus, pairs with smaller distances $\|v^j - t^{j'}\|_2$ have greater weights $w_{jj'}$ than distant ones. Similarly, if (v, t) is negative pair, then $w_{jj'}$ assigns greater weights on distant pairs than close ones. In other words, $w_{jj'}$ gives more weights on pair samples that correctly predicts the distance relationships on the embedding space. This results in the reward structure where wrong similarity predictions do not get penalized significantly, as long as there are at least one correct similarity prediction. Such a reward encourages the embeddings to produce more diverse samples and hedge the bets through non-zero values of σ predictions.

3.2.3 Test-time variants

Unlike methods that employ cross-modal reasoning modules [26, 30, 32, 33, 36, 54, 55, 60], computing match probabilities at test time for PCME reduces to computing a function over pairwise Euclidean distances. This means that the probabilistic embeddings of PCME can be used with various ways for computing the match probabilities at test time, with different variants having different computational complexities. The options are split into two groups. **(i) Sampling-based variants.** Similar to training, one can use Monte-Carlo sampling (Equation (2)) to approximate match probabilities. Assuming J samples, this requires $O(J^2)$ distance computations per match, as well as $O(J^2)$ space for every database entry. This implies that J plays an important role in terms of test time complexity. **(ii) Non-sampling**



- a) A baseball player swinging a bat at a ball.
- b) A baseball player is getting ready to hit a ball.
- c) A baseball player standing next to home plate holding a bat.
- d) A group of baseball players at the pitch.

Figure 4. Can you match the captions to the images? In the COCO annotations, each of the four captions corresponds to (only) one of the four images (Answer: P:D :a :c:B 'Q:V).

variants. One can simply use the distances based on μ to approximate match probabilities. In this case, both time and space complexities become $O(1)$. We ablate this variant (“ μ only”) in experiments, as it is directly comparable to deterministic approaches. We also may use any distributional distance measures with closed-form expressions for Gaussian distributions. Examples include the 2-Wasserstein distance, Jensen Shanon (JS) divergence, and Expected Likelihood Kernel (ELK). We ablate them as well. The details of each probabilistic distance can be found in §B.

4. Experiments

We present experimental results for PCME. We start with the experimental protocol and a discussion on the problems with current cross-modal retrieval benchmarks and evaluation metrics, followed by alternative solutions (§4.1). We then report experimental results on the CUB cross-modal retrieval task (§4.2) and COCO (§4.3). We present an analysis on the embedding space in §4.4.

4.1. Experimental protocol

We use an ImageNet-pretrained ResNet [15] and the pre-trained GloVe with 2.2M vocabulary [40] for initializing the visual and textual encoders. Training proceeds in two phases: a warm-up phase where only the head modules are trained, followed by end-to-end fine-tuning of all parameters. We use a ResNet-152 (resp. ResNet-50) backbone with embedding dimension $D = 1024$ (resp. $D = 512$) for MS-COCO (resp. CUB). For both datasets, models are always trained with Cutout [8] and random caption dropping [3] augmentation strategies with 0.2 and 0.1 erasing ratios, respectively. We use the AdamP optimizer [16] with the cosine learning rate scheduler [31] for a stable training. More implementation details are provided in §C.2. Hyper-

parameter details and ablations are presented in §D.

4.1.1 Metrics for cross-modal retrieval

Researchers have long been aware of many potentially positive matches in the cross-modal retrieval evaluation sets. Currently used metrics in part reflect such consideration.

Many works report the **Recall@k** ($R@k$) metrics with varying numbers for k . This evaluation policy, with larger values of k , becomes more lenient to plausible wrong predictions prevalent in COCO. However, it achieves leniency at the cost of failing to penalize obviously wrong retrieved samples. The lack of penalties for wrongly retrieved top- k samples may be complemented by the precision metrics.

Musgrave *et al.* [35] proposed the **R-Precision** (R-P) metric as an alternative; for every query q , we compute the ratio of positive items in the top- r retrieved items, where $r = |\tau(q)|$ is the number of ground-truth matches. This precision metric has a desirable property that a retrieval model achieves the perfect R-Precision score if and only if it retrieves all the positive items before the negatives.

For R-Precision to make sense, all the existing positive pairs in a dataset must be annotated. Hence, we expand the existing ground truth matches by seeking further plausible positive matches in a database through extra information (*e.g.* class labels for COCO). More concretely, a pair (i, c) is declared positive if the binary label vectors for the two instances, $y^i, y^c \in \{0, 1\}^{d_{label}}$, differ at most at ζ positions. In practice, we consider multiple criteria $\zeta \in \{0, 1, 2\}$ and average the results with those ζ values. We refer to metrics based on such class-based similarity as **Plausible Match (PM)** because we incentivize models to retrieve plausible items. We refer to the R-Precision metric based on the Plausible Match policy as **PMRP**. More details in §C.1.

4.1.2 Cross-modal retrieval benchmarks

COCO Captions [6] is a widely-used dataset for cross-modal retrieval models. It consists of 123,287 images from MS-COCO [29] with 5 human-annotated captions per image. We present experimental results on COCO. We follow the evaluation protocol of [19] where the COCO validation set is added to the training pool (referred to as rV or rVal in [9, 10]). Our training and validation splits contain 113,287 and 5,000 images, respectively. We report results on both 5K and (the average over 5-fold) 1K test sets.

The problem with COCO as a cross-modal retrieval benchmark is the binary relevance assignment of image-caption pairs (i, c) . As a result, the number of matching captions $\tau(i)$ for an image i is always 5. Conversely, the number of matching images $\tau(c)$ for a caption c is always 1. All other pairs are considered non-matching, independent of semantic similarity. This is far from representing the semantic richness of the dataset. See Figure 4 for an

illustration. While all 4×4 possible pairs are plausible positive pairs, 12 of them are assigned negative labels during training and evaluation. This results in a noisy training and, more seriously, unreliable evaluation results.

We re-purpose the CUB 200-2011 [56] as a more reliable surrogate for evaluating cross-modal retrieval models. We utilize the caption annotations by Reed *et al.* [41]; they consist of ten captions per image on CUB images (11,788 images of 200 fine-grained bird categories). False positives are suppressed by the fact that the captions and images are largely homogeneous within a class. False negatives are unlikely to happen because the images contain different types of birds across classes and the captions are generated under the instruction that the annotators should focus on class-distinguishing characteristics [41].

We follow the class splits proposed by Xian *et al.* [57], where 150 classes are used for training and validation, and the remaining 50 classes are used for the test. The hyperparameters are validated on the 150 training classes. We refer to this benchmark as *CUB Captions*.

4.2. Results on CUB

Similarity measures for retrieval at test time. We have discussed alternative similarity metrics that PCME may adopt at test time (§ 3.2.3). The “Mean only” metric only uses the h^μ features, as in deterministic retrieval scenarios. It only requires $O(N)$ space to store the database features. Probabilistic distance measures like ELK, JS-divergence, and 2-Wasserstein, require the storage for μ and σ features, resulting in the doubled storage requirement. Sampling-based distance computations, such as the average L2 distance and match probability, need J^2 times the storage required by the Mean-only baseline.

We compare the above variants in Table 1 and §E.1. First of all, we observe that PCME, with any test-time similarity measure, mostly beats the deterministically trained PCME (μ -only training). Even if the test-time similarity is computed as if the embeddings are deterministic (Mean only), PCME training improves the retrieval performances (24.7% to 26.1% for i2t and 25.6% to 26.7% for t2i). Other cheaper variants of probabilistic distances, such as 2-Wasserstein, also result in reasonable performances (26.2% and 26.7% i2t and t2i performances, respectively), while introducing only two times the original space consumption. The best performance is indeed attained by the similarity measure using the match probability, with 26.3% and 26.8% i2t and t2i performances, respectively. There exists a trade off between computational cost and performance and the deterministic test-time similarity measures. We use the match probability measure at test time for the rest of the paper.

Comparison against other methods. We compare PCME against VSE0 [10] and PVSE [47] in Table 2. As

PCME variant	Sampling	Test-time Similarity Metric	Space complexity	i2t R-P	t2i R-P
μ only	✗	Mean only	$O(N)$	24.70	25.64
PCME	✗	Mean only	$O(N)$	26.14	26.67
	✗	ELK	$O(2N)$	25.33	25.87
	✗	JS-divergence	$O(2N)$	25.06	25.55
	✗	2-Wasserstein	$O(2N)$	26.16	26.69
	✓	Average L2	$O(J^2N)$	26.11	26.64
	✓	Match prob	$O(J^2N)$	26.28	26.77

Table 1. **Pairwise distances for distributions.** There are many options for computing the distance between two distributions. What are the space complexity and retrieval performances for each option? R-P stands for the R-Precision.

Method	HNM	Image-to-text		Text-to-image	
		R-P	R@1	R-P	R@1
VSE0	✗	22.4	44.2	22.6	32.7
PVSE K=1	✓	22.3	40.9	20.5	31.7
PVSE K=2	✓	19.7	47.3	21.2	28.0
PVSE K=4	✓	18.4	47.8	19.9	34.4
PCME μ only	✗	24.7	46.4	25.6	35.5
PCME	✗	26.3	46.9	26.8	35.2

Table 2. **Comparison on CUB Caption test split.** R-P and R@1 stand for R-Precision and Recall@1, respectively. The usage of hardest negative mining (HNM) is indicated.

an important ingredient for PVSE, we consider the use of the hardest negative mining (HNM). We first observe that PVSE with HNM tends to obtain better performances than VSE0 under the R@1 metric, with 47.8% for $K=4$, compared to 44.2% for VSE0. However, under the R-Precision metric, we all PVSE models with HNM are worse than VSE0 (R-Precision drops from 22.4% for VSE0 to 18.4% for PVSE $K=4$). It seems that PVSE with HNM tends to retrieve items based on diversity, rather than precision. We conjecture that the HNM is designed to optimize the R@1 performances; more details in §E.2. Comparing PVSE with different values of K , we note that increasing K does not always bring about performance gains under the R-Precision metric (20.5%, 21.2% and 19.9% for $K=1,2,4$, respectively, for t2i), while the improvement is more pronounced under the R@1 metric. Finally, PCME provides the best performances on both R-Precision and R@1 metrics, except for the R@1 score for i2t. PCME also improves upon its deterministic version, PCME μ -only, with some margin: +1.6 pp and +1.2 pp on i2t and t2i R-Precision scores, respectively.

4.3. Results on COCO

As we have identified potential problems with measuring performance on COCO (§4.1.2), we report the results with our Plausible-Match R-Precision (PMRP) metrics (§4.1.1)

Method	1K Test Images				5K Test Images			
	i2t		t2i		i2t		t2i	
	PMRP	R@1	PMRP	R@1	PMRP	R@1	PMRP	R@1
VSE++ [10]	-	64.6	-	52.0	-	41.3	-	30.3
PVSE K=1 [47]	40.3*	66.7	41.8*	53.5	29.3*	41.7	30.1*	30.6
PVSE K=2 [47]	42.8*	69.2	43.6*	55.2	31.8*	45.2	32.0*	32.4
VSRN [27]	41.2*	76.2	42.4*	62.8	29.7*	53.0	29.9*	40.5
VSRN + AOQ [5]	44.7*	77.5	45.6*	63.5	33.0*	55.1	33.5*	41.1
PCME μ only	45.0	68.0	45.9	54.6	34.0	43.5	34.3	31.7
PCME	45.0	68.8	46.0	54.6	34.1	44.2	34.4	31.9

Table 3. **Comparison on MS-COCO.** PMRP stands for the Plausible Match R-Precision and R@1 for Recall@1. “*” denotes results produced by the published models.

that captures the model performances more accurately than the widely-used $R@k$ metrics. Table 3 shows the results with state-of-the-art COCO retrieval methods. We observe that the stochastic version of PCME performs better than the deterministic variant (μ only) across the board. In terms of the R@1 metric, PVSE $K=2$ [47], VSRN [27] and AOQ [5] work better than PCME (e.g. 45.2%, 53.0%, 55.1% versus 44.2% for the 5K, i2t task). However, on the more accurate PMRP metric, PCME outperforms previous methods with some margin (e.g. 31.8%, 29.7%, 33.0% versus 34.1% for the 5K, i2t task). The results on two metric imply that PCME retrieves the plausible matches much better than previous methods do. The full results can be found in §E.

4.4. Understanding the learned uncertainty

Having verified the retrieval performance of PCME, we now study the benefits of using probabilistic distributions for representing data. We show that the learned embeddings not only represent the inherent uncertainty of data, but also enable set algebras among samples that roughly correspond to their semantic meanings.

Measuring uncertainty with σ . In an automated decision process, it benefits a lot to be able to represent uncertainty. For example, the algorithm may refrain from making a decision based on the uncertainty estimates. We show that the learned cross-modal embeddings capture the inherent uncertainty in the instance. We measure the instance-wise uncertainty for all query instances by taking the geometric mean of over the $\sigma \in \mathbb{R}^D$ entries (§3.2.1). We then compute the average R@1 performances in each of the 10 uncertainty bins. Figure 6 plots the correlation between the uncertainty and R@1 on COCO test set. We observe performance drops with increasing uncertainty. In §F.2, we visualize which word affects more to uncertainty. Example uncertain instances and their retrieval results are in §F.3.

2D visualization of PCME. To visually analyze the behavior of PCME, we conduct a 2D toy experiment by using 9 classes of the CUB Captions (details in §C.3). Figure 5 visualizes the learned image and caption embeddings. We




Figure 5. **Visualization of the probabilistic embedding.** The learned image (left) and caption (right) embeddings on 9 subclass of CUB Captions. Classes are color-coded. Each ellipse shows the 50% confidence region for each embedding. The red ellipse corresponds to the generic CUB caption, “this bird has <unk> . . . <unk>” with 99% confidence region.




Figure 6. σ versus performance. Performance of PCME at different per-query uncertainty levels in COCO 1k test set.




Figure 7. σ captures ambiguity. Average σ values at different ratios of erased pixels (for images) and words (for captions).

also plot the embedding for the most generic caption for the CUB Captions dataset, “this bird has <unk> <unk> . . .”, where <unk> is a special token denoting the absence of a word. This generic caption covers most of caption variations in the embedding space (red ellipses).

Set algebras. To understand the relationship among distributions on the embedding space, we artificially introduce different types of uncertainties on the image data. In Figure 8, we start from two bird images and perform erasing and mixing transformations [59]. On the embedding space, we find that the mixing operation on the images results in embeddings that cover the *intersection* of the original embeddings. Occluding a small region in input images, on the other hand, amounts to slightly wider distributions, indicat-




Figure 8. **Set algebras.** For two images, we visualize the embeddings for either erased or mixed samples. Mixing (left) and erasing (right) operations roughly translate to the intersection and inclusion relations between the corresponding embeddings.

ing an *inclusion* relationship. We quantitatively verify that the sigma values positively correlate with the ratio of erased pixels in Figure 7. In COCO, we observe a similar behavior (shown in §F.1). We discover another positive correlation between the caption ambiguity induced by erasing words and the embedding uncertainty.

5. Conclusion

We introduce Probabilistic Cross-Modal Embedding (PCME) that learns probabilistic representations of multi-modal data in the embedding space. The probabilistic framework provides a powerful tool to model the widespread one-to-many associations in the image-caption pairs. To our knowledge, this is the first work that uses probabilistic embeddings for a multi-modal task. We extensively ablate our PCME and show that not only it improves the retrieval performance over its deterministic counterpart, but also provides uncertainty estimates that render the embeddings more interpretable.

Acknowledgement

We thank our NAVER AI LAB colleagues for valuable discussions. All experiments were conducted on NAVER Smart Machine Learning (NSML) [21] platform.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [3](#), [4](#)
- [2] Anil Bhattacharyya. On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math. Soc.*, 35:99–109, 1943. [11](#)
- [3] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21, Berlin, Germany, Aug. 2016. Association for Computational Linguistics. [5](#), [12](#)
- [4] Jie Chang, Zhonghao Lan, Changmao Cheng, and Yichen Wei. Data uncertainty learning in face recognition. In *Proc. CVPR*, pages 5710–5719, 2020. [2](#)
- [5] Tianlang Chen, Jiajun Deng, and Jiebo Luo. Adaptive offline quintuplet loss for image-text matching. In *Proc. ECCV*, 2020. [7](#), [15](#), [16](#)
- [6] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. [1](#), [2](#), [6](#)
- [7] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014. [3](#)
- [8] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. [5](#), [12](#)
- [9] Martin Engilberge, Louis Chevallier, Patrick Pérez, and Matthieu Cord. Finding beans in burgers: Deep semantic-visual embedding with localization. In *Proc. CVPR*, 2018. [2](#), [6](#), [12](#)
- [10] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. VSE++: Improving visual-semantic embeddings with hard negatives. In *Proc. BMVC*, 2018. [1](#), [2](#), [3](#), [6](#), [7](#), [12](#), [15](#), [16](#)
- [11] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Proc. NeurIPS*, pages 2121–2129, 2013. [2](#), [3](#)
- [12] Yunchao Gong, Liwei Wang, Micah Hodosh, Julia Hockenmaier, and Svetlana Lazebnik. Improving image-sentence embeddings using large weakly annotated photo collections. In *European Conference on Computer Vision*, pages 529–545. Springer, 2014. [2](#)
- [13] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *Proc. CVPR*, 2006. [3](#)
- [14] David R Hardoon, Sandor Szedmak, and John Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural computation*, 16(12):2639–2664, 2004. [2](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. CVPR*, 2016. [3](#), [5](#), [12](#)
- [16] Byeongho Heo, Sanghyuk Chun, Seong Joon Oh, Dongyoon Han, Sangdoon Yun, Youngjung Uh, and Jung-Woo Ha. Slowing down the weight norm increase in momentum-based optimizers. *arXiv preprint arXiv:2006.08217*, 2020. [5](#), [12](#)
- [17] Yan Huang, Wei Wang, and Liang Wang. Instance-aware image and sentence matching with selective multimodal lstm. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2310–2318, 2017. [2](#)
- [18] Tony Jebara, Risi Kondor, and Andrew Howard. Probability product kernels. *Journal of Machine Learning Research*, 5(Jul):819–844, 2004. [11](#)
- [19] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proc. CVPR*, pages 3128–3137, 2015. [6](#), [12](#)
- [20] Andrej Karpathy, Armand Joulin, and Li F Fei-Fei. Deep fragment embeddings for bidirectional image sentence mapping. In *Proc. NeurIPS*, pages 1889–1897, 2014. [2](#)
- [21] Hanjoo Kim, Minkyu Kim, Dongjoo Seo, Jinwoong Kim, Heungseok Park, Soeun Park, Hyunwoo Jo, KyungHyun Kim, Youngil Yang, Youngkwan Kim, et al. NSML: Meet the MLaaS platform with a real-world case study. *arXiv preprint arXiv:1810.09957*, 2018. [9](#)
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [3](#)
- [23] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *Proc. ICLR*, 2014. [2](#)
- [24] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. [2](#)
- [25] Benjamin Klein, Guy Lev, Gil Sadeh, and Lior Wolf. Fisher vectors derived from hybrid gaussian-laplacian mixture models for image annotation. *arXiv preprint arXiv:1411.7399*, 2014. [2](#)
- [26] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *Proc. ECCV*, 2018. [2](#), [5](#)
- [27] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *Proc. ICCV*, pages 4654–4662, 2019. [1](#), [2](#), [7](#), [15](#), [16](#)
- [28] Xiang Li, Luke Vilnis, Dongxu Zhang, Michael Boratko, and Andrew McCallum. Smoothing the geometry of probabilistic box embeddings. In *Proc. ICLR*, 2019. [2](#)
- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proc. ECCV*, 2014. [6](#)
- [30] Chunxiao Liu, Zhendong Mao, An-An Liu, Tianzhu Zhang, Bin Wang, and Yongdong Zhang. Focus your attention: A bidirectional focal attention network for image-text matching. In *Proceedings of the 27th ACM International Conference on Multimedia*, page 3–11, 2019. [2](#), [5](#)

- [31] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 5, 12
- [32] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Proc. NeurIPS*, pages 13–23, 2019. 2, 5
- [33] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10437–10446, 2020. 2, 5
- [34] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012. 2
- [35] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. A metric learning reality check. In *Proc. ECCV*, 2020. 6, 13
- [36] Hyeonseob Nam, Jung-Woo Ha, and Jeonghee Kim. Dual attention networks for multimodal reasoning and matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 299–307, 2017. 2, 5
- [37] Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. Efficient non-parametric estimation of multiple embeddings per word in vector space. In *Proc. EMNLP*, pages 1059–1069, 2014. 2
- [38] Dat Quoc Nguyen, Ashutosh Modi, Stefan Thater, Manfred Pinkal, et al. A mixture model for learning multi-sense word embeddings. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 121–127, 2017. 2
- [39] Seong Joon Oh, Kevin Murphy, Jiyan Pan, Joseph Roth, Florian Schroff, and Andrew Gallagher. Modeling uncertainty with hedged instance embedding. In *Proc. ICLR*, 2019. 2, 3, 4
- [40] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proc. EMNLP*, 2014. 3, 5, 12
- [41] Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. Learning deep representations of fine-grained visual descriptions. In *Proc. CVPR*, pages 49–58, 2016. 6
- [42] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proc. NeurIPS*, pages 91–99, 2015. 2
- [43] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *Proc. CVPR*, pages 8247–8255, 2019. 2
- [44] Tyler Scott, Karl Ridgeway, and Michael Mozer. Stochastic prototype embeddings. *ICML Workshop on Uncertainty and Robustness in Deep Learning*, 2019. 2
- [45] Yichun Shi and Anil K Jain. Probabilistic face embeddings. In *ICCV*, 2019. 2
- [46] Anna Silnova, Niko Brummer, Johan Rohdin, Themis Stafylakis, and Lukas Burget. Probabilistic embeddings for speaker diarization. In *Proc. Odyssey 2020 The Speaker and Language Recognition Workshop*, pages 24–31, 2020. 2
- [47] Yale Song and Mohammad Soleymani. Polysemous visual-semantic embedding for cross-modal retrieval. In *Proc. CVPR*, pages 1979–1988, 2019. 2, 3, 4, 5, 6, 7, 11, 15, 16
- [48] Jennifer J Sun, Jiaping Zhao, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, and Ting Liu. View-invariant probabilistic embedding for human pose. In *Proc. ECCV*, 2020. 2
- [49] Christopher Thomas and Adriana Kovashka. Preserving semantic neighborhoods for robust cross-modal retrieval. In *Proc. ECCV*, 2020. 2
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 4
- [51] Luke Vilnis and Andrew McCallum. Word representations via gaussian embedding. In *Proc. ICLR*, 2015. 2
- [52] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *Proc. CVPR*, pages 5005–5013, 2016. 1, 2
- [53] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, 2020. 4, 11
- [54] Zihao Wang, Xihui Liu, Hongsheng Li, Lu Sheng, Junjie Yan, Xiaogang Wang, and Jing Shao. CAMP: Cross-modal adaptive message passing for text-image retrieval. In *Proc. ICCV*, 2019. 2, 5
- [55] Xi Wei, Tianzhu Zhang, Yan Li, Yongdong Zhang, and Feng Wu. Multi-modality cross attention network for image and sentence matching. In *Proc. CVPR*, 2020. 2, 5
- [56] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010. 2, 6
- [57] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning—the good, the bad and the ugly. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4582–4591, 2017. 6
- [58] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *ACL*, 2:67–78, 2014. 1
- [59] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proc. ICCV*, 2019. 8
- [60] Qi Zhang, Zhen Lei, Zhaoxiang Zhang, and Stan Z. Li. Context-aware attention network for image-text retrieval. In *Proc. CVPR*, 2020. 2, 5

Supplementary Materials

We include additional materials in this document. We describe additional details on PCME to complement the main paper (§A). Various probabilistic distances are introduced (§B). We provide the experimental protocol details (§C), ablation studies (§D), and additional results (§E). Finally, more uncertainty analyses are shown (§F).

A. More details for PCME

In this section, we provide details for PCME.

A.1. The uniformity loss

Recently, Wang *et al.* [53] propose the uniformity loss which enforces the feature vectors to distribute uniformly on the unit hypersphere. In Wang *et al.* [53], the uniformity loss was shown to lead to better representations for L2 normalized features. Since our μ vectors are projected to the unit L2 hypersphere, we also employ the uniformity loss to learn better representations. We apply the uniformity loss on the joint embeddings $\mathcal{Z} = \{v_1^I, t_1^I, \dots, v_B^I, t_B^I\}$ in the mini-batch size of B as follows:

$$\mathcal{L}_{\text{Unif}} = \sum_{z, z' \in \mathcal{Z} \times \mathcal{Z}} e^{-2\|z-z'\|_2^2}. \quad (\text{A.1})$$

A.2. Connection of soft contrastive loss to the MIL objective of PVSE

In the main text we presented an analysis based on gradients to study how the loss function in Equation (1) handles plurality in cross-modal matches and learns uncertainties in data. Here we make connections with the MIL loss used by PVSE (§3.1.1, [47]); this section follows the corresponding section in the main paper.

To build connections with PVSE, consider a one-hot weight array $w_{jj'}$ where, given that (v, t) is a positive pair, the ‘‘one’’ value is taken only by the single pair (j, j') whose distance is smallest. Define $w_{jj'}$ for a negative pair (v, t) conversely. Then, we recover the MIL loss used in PVSE, where only the best match among J^2 predictions are utilized. As we see in the experiments, our *softmax* weight scheme provides a more interpretable and performant supervision for the uncertainty than the *argmax* version used by PVSE.

B. Probabilistic distances

We introduce probabilistic distance variants to measure the distance between two normal distributions $p = \mathcal{N}(\mu_1, \sigma_1^2)$ and $q = \mathcal{N}(\mu_2, \sigma_2^2)$. All distance functions are non-negative, and become zero if and only if two distributions are identical. Extension to multivariate Gaussian distributions with diagonal variance can be simply derived by taking the summation over the dimension-wise distances.

Kullback–Leibler (KL) divergence measures the difference between two distributions as follows:

$$\begin{aligned} KL(p, q) &= \int \log \frac{p}{q} dp \\ &= \frac{1}{2} \left[\log \frac{\sigma_2^2}{\sigma_1^2} + \frac{\sigma_1^2}{\sigma_2^2} + \frac{(\mu_1 - \mu_2)^2}{\sigma_2^2} \right]. \end{aligned} \quad (\text{B.1})$$

KL divergence is not a metric because it is asymmetric ($KL(p, q) \neq KL(q, p)$) and does not satisfy the triangular inequality. If q has very small variance, nearly zero, the KL divergence between p and q will be explored. In other words, if we have a very certain embedding, which has nearly zero variance, in our gallery set, then the certain embedding will be hardly retrieved by KL divergence measure. In the latter section, we will show that KL divergence leads to bad retrieval performances in the real-world scenario.

Jensen-Shannon (JS) divergence is the average of forward ($KL(p, q)$) and reverse ($KL(q, p)$) KL divergences. Unlike KL divergence, the square root of JS divergence is a metric function.

$$JS(p, q) = \frac{1}{2} [KL(p, q) + KL(q, p)]. \quad (\text{B.2})$$

Like KL divergence, JS divergence still has division term by variances σ_1, σ_2 , it can be numerically unstable when the variances are very small.

Probability product kernels [18] are generalized inner product for two distributions, that is:

$$PPK(p, q) = \int p(z)^\rho q(z)^\rho dz. \quad (\text{B.3})$$

When $\rho = 1$, it is called the expected likelihood kernel (ELK), and when $\rho = 1/2$, it is called Bhattacharyya’s affinity [2], or Bhattacharyya kernel.

Expected likelihood kernel (ELK) is a special case of PPK when $\rho = 1$ in Equation (B.3). In practice, we take log to compute ELK as follows:

$$ELK(p, q) = \frac{1}{2} \left[\frac{(\mu_1 - \mu_2)^2}{\sigma_1^2 + \sigma_2^2} + \log(\sigma_1^2 + \sigma_2^2) \right]. \quad (\text{B.4})$$

Bhattacharyya kernel (BK) is another special case of PPK when $\rho = 1/2$ in Equation (B.3). The log BK is defined as the follows:

$$BK(p, q) = \frac{1}{4} \left[\frac{(\mu_1 - \mu_2)^2}{\sigma_1^2 + \sigma_2^2} + 2 \log\left(\frac{\sigma_2}{\sigma_1} + \frac{\sigma_1}{\sigma_2}\right) \right]. \quad (\text{B.5})$$

Wasserstein distance is a metric function of two distributions on a given metric space M . The Wasserstein distance between two normal distributions on \mathbb{R}^1 , 2-Wasserstein distance, is defined as follows:

$$W(p, q)^2 = (\mu_1 - \mu_2)^2 + \sigma_1 - \sigma_2^2. \quad (\text{B.6})$$




Figure C.1. **Number of distinct categories in MS-COCO validation set.** Images which have more than 10 categories are omitted.

C. Experimental Protocol Details

We introduce the cross-modal retrieval benchmarks considered in this work. We discuss the issues with the current practice for evaluation and introduce new alternatives.

C.1. Plausible Match R-Precision (PMRP) details

In this work, we seek more reliable sources of pairwise similarity measurements through class and attribute labels on images. For example, on CUB caption dataset, we have established the positivity of pairs by the criterion that a pair (i, c) is positive if and only if both elements in the pair belong to the same bird class. Similarly, on COCO caption dataset, we judge the positivity through the multiple class labels (80 classes total) attached per image: a pair (i, c) is positive if and only if the binary class vectors for the two instances, $y^i, y^c \in \{0, 1\}^{80}$, differ at most at ζ positions (Hamming Distance). Note that because we use R-Precision, the ratio of positive items in top- r retrieved items where r is the number of the ground-truth matches, increasing ζ will make r larger, and will penalize methods more, which retrieve irrelevant items.

In Figure C.1, we visualize the number of distinct categories per image in MS-COCO validation set. In the figure, we can observe that about the half of images have more than two categories. To avoid penalty caused by almost neglectable objects (as shown in Figure C.2), we set $\zeta = 2$ for measuring PMRP score. For PMRP with different ζ rather than 2, results can be found in §E.

C.2. Implementation details

Common. As in Faghri *et al.* [10], we use the ImageNet-pretrained ResNet [15] and the pretrained GloVe with 2.2M vocabulary [40] for initializing the visual and textual encoders (f_V, f_T) . We first warm-up the models by training the head modules for each modality, with frozen feature extractors. Afterwards, the whole parameters are fine-tuned

in an end-to-end fashion. We use the ResNet-152 backbone with embedding dimension $D = 1024$ for MS-COCO and ResNet-50 with $D = 512$ for CUB. For all experiments, we set the number of samples $J = 7$ (the detailed study is in §E). We use AdamP optimizer [16] with the cosine learning rate scheduler [31] for a stable training.

MS-COCO. We follow the evaluation protocol of [19] where the validation set is added to the training pool (referred to as rV in [9, 10]). Our training and validation splits contain 113,287 and 5,000 images, respectively. We report results on both 5K and (the average over 5-fold) 1K test sets.

Hyperparameter search protocol. We validate the initial learning rate, number of epochs for the warm-up and fine-tuning, and other hyperparameters on the 150 CUB training classes and the MS-COCO caption validation split. For MS-COCO, we use the initial learning rate as 0.0002, 30 warm-up and 30 finetune epochs. Weights for regularizers \mathcal{L}_{KL} and \mathcal{L}_{Unif} are set to 0.00001 and 0, respectively. For CUB Caption, the initial learning rate is 0.0001, the number of warm-up epochs 10 and fine-tuning epochs 50. Weights for regularizers \mathcal{L}_{KL} and \mathcal{L}_{Unif} are set to 0.001 and 10, respectively. For both datasets, models are always trained with Cutout [8] and random caption dropping [3] augmentation strategies with 0.2 and 0.1 erasing ratios, respectively. The initial values for a, b in Equation (3) are set to -15 and 15 for COCO (-5 and 5 for CUB), respectively.

C.3. CUB 2D toy experiment details

We select nine bird classes from CUB caption; three swimming birds (“Western Grebe”, “Pied Billed Grebe”, “Pacific Loon”), three small birds (“Vermilion Flycatcher”, “Black And White Warbler”, “American Redstart”), and three woodpeckers (“Red Headed Woodpecker”, “Red Bellied Woodpecker”, “Downy Woodpecker”).

We slightly modify PCME to learn 2-dimensional embeddings. For the image encoder, we use the same structure as the other experiments, but omitting the attention modules from the μ and σ modules. For the caption encoder, we train 1024-dimensional bi-GRU on top of GloVe vectors, and apply two 2D projections to get the 1024 dimensional μ and σ embedding. The other training details are the same as the other CUB caption experiments.

D. Ablation studies

We provide ablation studies on PCME for regularization terms, σ module architectures, number of samples J during training, and embedding dimension D .

Regularizing uncertainty. PCME predicts probabilistic outputs. We have considered uncertainty-specific regular-




Figure C.2. **MS-COCO plausible match examples.** The plausible examples of the most left instance from $\zeta = 0$ to $\zeta = 2$. The contained instance classes, ζ , figure and captions are shown.

\mathcal{L}_{KL}	\mathcal{L}_{Unif}	i2t MAP@R	t2i MAP@R	Image $\mathbb{E}[\log \sigma]$	Caption $\mathbb{E}[\log \sigma]$
✗	✗	10.56	13.32	-13.01	-8.77
✓	✗	10.57	13.77	-3.84	-3.89
✗	✓	10.56	13.31	-11.26	-7.59
✓	✓	10.65	13.84	-3.63	-3.64

Table D.1. **Regularization for uncertainty.** Cross-validated MAP@R performances on CUB training set, with and without KL and uniformity loss terms. The scale estimate $\mathbb{E}[\log \sigma]$ is an averaged value over the σ dimensions as well as the validation samples.

Method	DoF(σ)	i2t	t2i
PCME μ only	0	24.7	25.6
PCME isotropic	1	25.7	26.0
PCME	512	26.3	26.8

Table D.2. **DoF for σ .** R-Precision on the CUB Caption test set.

ization strategy in the main paper, the information bottleneck loss \mathcal{L}_{KL} and the uniform loss \mathcal{L}_{Unif} . We study the benefits of those ingredients. Table D.1 shows our results. We report cross-validated MAP@R [35] on the 150 class training CUB caption datasets. The KL loss increases the sigma values to a meaningful range (from $e^{-13.01} \approx 2.2 \times 10^{-6}$ to $e^{-3.84} \approx 0.02$). The uniformity loss prevents the uncertainty from collapsing and slightly improves performances.

DoF for σ . Though by default we parametrize the full diagonal elements of the covariance matrix $\Sigma \in \mathbb{R}^{D \times D}$ with the vector $\sigma \in \mathbb{R}^D$, one may parametrize σ more cheaply via *e.g.* a scalar, by restricting the embedding distribution family to isotropic Gaussians. Table D.2 shows the trade-off between the degree of freedom (DoF) for σ and the




Figure D.1. **How isotropic are variances?** Sorted values of variance are compared against the trained values of isotropic PCME. Results on CUB test set.

R-Precision of PCME. Indeed, allowing greater degrees of freedom for σ brings better performance. Figure D.1 shows the average variance values for each dimension, which supports that the learned variances require high DoF.

Architecture study. Table D.3 shows the architecture design comparisons for PCME on CUB Caption test split. In the table, applying local attention to both μ and σ modules performs the best. Furthermore, we ablate sigmoid and LN parts of σ modules, which can restrict the representation of variances. As a result, limiting representations by sigmoid and layer norm harms the final performances.

Number of samples during training. In Figure D.2, we report the cross-validated mean R-Precision scores by varying the number of samples J during training. In the figure,

μ local attention	σ local attention	I-to-T R-Precision	T-to-I R-Precision
✗	✗	25.60	25.85
✗	✓	24.65	25.15
✓	✗	25.01	25.52
✓	✓	26.28	26.77

$s(\cdot)$ & LN in σ module	I-to-T R-Precision	T-to-I R-Precision
✓	23.81	24.58
✗	26.28	26.77

Table D.3. **Architectures for μ and σ .** Architecture design choices comparison on CUB caption test split.




Figure D.2. **Number of samples.** The cross-validated PCME performances against the number of samples J during training.




Figure D.3. **Embedding dimensions.** The PCME performance against the embedding dimensions D .

we observe that larger J leads to higher performances. In practice, we choose $J = 7$ for the computation budgets.

Embedding dimensions. Performances against different embedding space dimensions for PCME μ only and PCME are illustrated in Figure D.3. In all embedding dimensions, our stochastic approach (PCME) consistently outperforms the deterministic approach (PCME μ only).

PCME variant	Sampling	Test-time Similarity Metric	Space complexity	i2t R-P	t2i R-P
μ only	✗	Mean only	$O(N)$	24.70	25.64
PCME	✗	Mean only	$O(N)$	26.14	26.67
	✗	KL-divergence	$O(2N)$	21.99	20.92
	✗	JS-divergence	$O(2N)$	25.06	25.55
	✗	ELK	$O(2N)$	25.33	25.87
	✗	Bhattacharyya	$O(2N)$	24.93	25.27
	✗	2-Wasserstein	$O(2N)$	<u>26.16</u>	<u>26.69</u>
	✓	Average L2	$O(J^2N)$	26.11	26.64
✓	Match prob	$O(J^2N)$	26.28	26.77	

Table E.1. **Pairwise distances for distributions.** There are many options for computing the distance between two distributions. What are the space complexity and retrieval performances for each option? R-P stands for the R-Precision.




Figure E.1. **Comparison of different retrieval strategies.**

E. More results

In this section, we provide additional experimental results for PCME on CUB Caption and COCO Caption.

E.1. More results on similarity measures for retrieval at test time

In Table E.1, we report the full retrieval results obtained by the different distribution distances discussed in §B. As discussed in §B, KL-divergence even shows worse results than the “Mean only” baseline, a non-probabilistic distance. We also report the performances against the number of samples of matching probability in Figure E.1. In the figure, the matching probability strategy shows better results than non-sampling strategies from $J = 3$, and larger J leads to better performances. Due to the computation complexity, we use $J = 7$ in Table E.1.

E.2. Discussion on hardest negative mining

Since Recall@K is widely used for the evaluation of many cross-modal retrieval tasks, many recent cross-modal retrieval methods optimize Recall@1 directly by the hardest




Figure E.2. Hardest negative mining (HNM) vs. Non-HNM.

negative mining (HNM) strategy [10], that is:

$$\begin{aligned} & \max_{t'} [\alpha + \text{sim}(v, t') - \text{sim}(v, t)] \\ & + \max_{v'} [\alpha + \text{sim}(v', t) - \text{sim}(v, t)], \end{aligned} \quad (\text{E.1})$$

where sim is the cosine similarity. This strategy neglects all other possible positive candidates, but only considers the most similar positive and negative pairs. To reveal that HM strategy disadvantages to learn the global structure, we measure two metrics on CUB caption, R-Precision and Recall@1. For non-HM strategy, we replace \max to \sum in Equation (E.1). Figure E.2 shows R-Precision and recall@1 performances by different mining strategies. In the figure, PVSE with HNM strategy shows higher Recall@1 by increasing the number of embeddings K ($36.3 \rightarrow 37.6 \rightarrow 41.1$), but at the same time, it reduces the R-Precision scores ($21.4 \rightarrow 20.4 \rightarrow 19.2$). On the other hand, for all K , Non-HNM strategy PVSE results show worse R@1 than HNM results, but achieves higher R-Precision performances. In Table 3, we show that this phenomenon is also observed in MS-COCO by measuring PMRP scores.

E.3. Full results for CUB and COCO

CUB Caption. We report the full results on CUB Caption test data for unseen 50 classes and seen 150 classes in Table E.2 and Table E.3, respectively. In both splits, PCME shows the best R-Precision performances against baselines.

COCO Caption. We report the full results on MS-COCO Caption 1k test images and 5k test images in Table E.4 and Table E.5, respectively. In the tables, although PCME shows slightly worse R@1 results than PVSE K=2, PCME outperforms PVSE K=2 in PMRP scores.

Also, we report PMRP scores of four methods (PVSE [47], VSRN [27], VSRN + AOQ [5] and PCME) by varying ζ for PMRP in Figure E.3. In the figure, PMRP scores for VSRN and VSRN + AOQ are getting worse by

Method	HNM	Image-to-text		Text-to-image	
		R-P	R@1	R-P	R@1
VSE0	✗	22.35	44.19	22.57	32.71
PVSE K=1	✗	22.65	43.11	22.78	33.49
PVSE K=2	✗	21.62	44.05	21.49	31.31
PVSE K=4	✗	21.12	40.51	20.90	30.94
PVSE K=1	✓	22.34	40.88	20.51	31.71
PVSE K=2	✓	19.67	47.29	21.16	27.98
PVSE K=4	✓	18.38	47.76	19.94	34.39
PCME μ only	✗	24.70	46.38	25.64	35.50
PCME	✗	26.28	46.92	26.77	35.22

Table E.2. Comparison on CUB Caption unseen 50 class test set. R-P and R@1 stand for R-Precision and Recall@1, respectively. The usage of hardest negative mining (HNM) is indicated.

Method	HNM	Image-to-text		Text-to-image	
		R-P	R@1	R-P	R@1
VSE0	✗	19.85	40.88	18.72	25.51
PVSE K=1	✗	19.69	40.65	18.72	25.58
PVSE K=2	✗	18.84	41.45	17.72	24.99
PVSE K=4	✗	18.31	38.08	17.21	23.54
PVSE K=1	✓	18.98	38.77	18.23	23.49
PVSE K=2	✓	17.62	44.24	17.71	22.78
PVSE K=4	✓	17.47	44.98	17.44	26.19
PCME μ only	✗	20.65	42.70	20.16	26.94
PCME	✗	20.87	43.10	20.37	26.47

Table E.3. Comparison on CUB Caption seen 150 class test set. R-P and R@1 stand for R-Precision and Recall@1, respectively. The usage of hardest negative mining (HNM) is indicated.

increasing ζ , in other words, these method shows less coherence if we allow one missing or altering object class in the retrieved items. On the other hand, PCME shows even increased performance with $\zeta > 0$, in other words, PCME retrieves more plausible items than other methods.

F. More uncertainty analysis

Uncertainty estimation by PCME brings interesting insights for the cross-modal retrieval tasks. In this section, we show additional uncertainty analysis for PCME.

F.1. Corruption vs. uncertainty in MS-COCO

As Figure 7, we illustrate the uncertainty level by varying corruption levels on pixels and words in Figure F.1. The left figure shows the uncertainty levels against occluded pixels. As we expected, more occlusion leads to higher uncertainty. The right figure shows the uncertainty levels against the number of appended $\langle \text{unk} \rangle$ tokens.

Method	D	PMRP	Image-to-text			Text-to-image			
			R@1	R@5	R@10	PMRP	R@1	R@5	R@10
VSE++ BMVC'18 [10]	1024	-	64.6	90.0	95.7	-	52.0	84.3	92.0
PVSE K=1 CVPR'19 [47]	1024	40.3*	66.7	91.0	96.2	41.9*	53.5	85.1	92.7
PVSE K=2 CVPR'19 [47]	1024×2	42.8*	69.2	91.6	96.6	43.7*	55.2	86.5	93.7
VSRN ICCV'19 [27]	2048	41.2*	76.2	94.8	98.2	42.4*	62.8	89.7	95.1
VSRN + AOQ ECCV'20 [5]	2048×2	44.7*	77.5	95.5	98.6	45.6*	63.5	90.5	95.8
PCME $_{\mu}$ only	1024	45.0	68.0	92.0	96.2	45.9	54.6	86.3	93.8
PCME	1024×2	45.1	68.8	91.6	96.7	46.0	54.6	86.3	93.8

Table E.4. **1K MS-COCO results.** Plausible Match R-Precision (PMRP), Recall@K results on MS-COCO 1k test images. “*” denotes results produced by the published models.

Method	D	PMRP	Image-to-text			Text-to-image			
			R@1	R@5	R@10	PMRP	R@1	R@5	R@10
VSE++ BMVC'18 [10]	1024	-	41.3	71.1	81.2	-	30.3	59.4	72.4
PVSE K=1 CVPR'19 [47]	1024	29.3*	41.7	73.0	83.0	30.1*	30.6	61.4	73.6
PVSE K=2 CVPR'19 [47]	1024×2	31.8*	45.2	74.3	84.5	32.0*	32.4	63.0	75.0
VSRN ICCV'19 [27]	2048	29.7*	53.0	81.1	89.4	29.9*	40.5	70.6	81.1
VSRN + AOQ ECCV'20 [5]	2048×2	33.0*	55.1	83.3	90.8	33.5*	41.1	71.5	82.0
PCME $_{\mu}$ only	1024	34.0	43.5	73.1	84.2	34.3	31.7	62.2	74.9
PCME	1024×2	34.1	44.2	73.8	83.6	34.4	31.9	62.1	74.5

Table E.5. **Comparison on 5K MS-COCO.** Plausible Match R-Precision (PMRP), Recall@K results on MS-COCO 5k test images. “*” denotes results produced by the published models.




Figure E.3. **PMRP by varying ζ .** Plausible Match R-Precision scores for four methods with $\zeta = \{0, 1, 2\}$.

F.2. Frequent words for each uncertainty bin

Figure F.2 shows the frequent words per each uncertainty bin. We use term frequency–inverse document frequency (TF-IDF) as the frequent counter, defined as follows:

$$\text{TF-IDF}(i) = (1 + \log n_i) \log \frac{N}{n_i}, \quad (\text{F.1})$$

where N is the number of total captions, and n_i is the number of captions which contain word i . For the image word




Figure F.1. **σ captures ambiguity in COCO Caption.** Average $\log \sigma$ values at different ratios of erased pixels (for images) and appended `<unk>` tokens (for captions).

frequency, we use their ground truth captions for computing TF-IDF scores.

F.3. Example uncertain samples

We visualize the uncertain images and captions, and their corresponding retrieved items in Figure F.3 and Figure F.4. Interestingly, the retrieved captions and images are plausible results for the given query items. These qualitative results also show how Recall@1 measure is noisy, and the proposed Plausible Match R-Precision (PMRP) is a more plausible and reliable measure to compare different retrieval methods.




Figure F.2. **Frequent words in each uncertainty bin.** Term frequency–inverse document frequency (TF-IDF) sorted word frequencies are shown for each uncertainty bin (U-Bin, ascending order) for image (upper row) and caption (bottom row) modalities.

Query Image ($\sigma=0.0052$)



Retrieved captions

- A boy riding on ski's down a slope.
- A young boy is attempting to slide down a slope.
- A kid is riding down the street on a skateboard.
- a man with warm clothes skating on the snow
- a young person riding skis on a snowy field
- a person skating in very much snow with warm clothes

GT captions

- Two boys riding skateboards in the street, behind tree branches.
- two young people riding skate boards on a flat surface
- Two young men riding skateboards across a parking lot.
- two young men skateboarding in an open area during winter
- A couple of kids riding on top of skateboards.

Query image ($\sigma = 0.0054$)



Retrieved captions

- Two people in the midst of a tennis match on a grass court.
- Two men on grass court playing a game of tennis.
- Two men playing doubles tennis on a grass court.
- Two men playing tennis on a grass field.
- a couple of people play a game of tennis on a grass surface
- A male tennis players on the court with rackets.

GT captions

- A couple of men holding tennis racquets on a tennis court.
- Two men playing tennis at a somewhat large facility.
- Two men playing doubles tennis on a grass court.
- A couple of tennis players during a couples game about to deliver a hit.
- A male tennis players on the court with rackets.

Query image ($\sigma = 0.0054$)



Retrieved captions

- A surfer riding a wave in a blue ocean.
- A wet suited surfer riding the crest of an azure wave
- a male surfing a large ocean wave on a white surfboard
- The surfer is working on riding the big wave.
- A surfer is riding on a large wave.
- A surf boarder who is riding a wave.

GT captions

- A surfer is on his board in the middle of an ocean spraying wave.
- A man on a surfboard riding a wave
- A man is surfing a small wave in the ocean.
- A man riding on a wave on a surf board.
- a person riding a surf board on a wave

Figure F.3. **Uncertain image examples.** Highly uncertain images, retrieved captions by PCME, and their ground truth captions are shown.

Query caption ($\sigma = 0.0046$): A batter is swinging at a ball at the game.



Query caption ($\sigma = 0.0046$): a large clock tower is on top of a building.



Query caption ($\sigma = 0.0047$): A man playing tennis outside during a sunny day.



GT image

Retrieved images

Figure F.4. **Uncertain caption examples.** Highly uncertain captions, retrieved images by PCME, and their ground truth image are shown.