

# ProbVLM: Probabilistic Adapter for Frozen Vision-Language Models

Uddeshya Upadhyay<sup>\*,1</sup> Shyamgopal Karthik<sup>\*,1</sup> Massimiliano Mancini<sup>2</sup> Zeynep Akata<sup>1,3</sup>

<sup>1</sup>University of Tübingen <sup>2</sup>University of Trento <sup>3</sup>MPI for Intelligent Systems

## Abstract

Large-scale vision-language models (VLMs) like CLIP successfully find correspondences between images and text. Through the standard deterministic mapping process, an image or a text sample is mapped to a single vector in the embedding space. This is problematic: as multiple samples (images or text) can abstract the same concept in the physical world, deterministic embeddings do not reflect the inherent ambiguity in the embedding space. We propose *ProbVLM*, a probabilistic adapter that estimates probability distributions for the embeddings of pre-trained VLMs via inter/intra-modal alignment in a post-hoc manner without needing large-scale datasets or computing. On four challenging datasets, i.e., COCO, Flickr, CUB, and Oxford-flowers, we estimate the multi-modal embedding uncertainties for two VLMs, i.e., CLIP and BLIP, quantify the calibration of embedding uncertainties in retrieval tasks and show that *ProbVLM* outperforms other methods. Furthermore, we propose active learning and model selection as two real-world downstream tasks for VLMs and show that the estimated uncertainty aids both tasks. Lastly, we present a novel technique for visualizing the embedding distributions using a large-scale pre-trained latent diffusion model. Code is available at <https://github.com/ExplainableML/ProbVLM>

## 1. Introduction

Recently, large vision-language models (VLMs) [62, 51, 45, 74, 1, 35] have become exceedingly popular due to their ability to align images and text. These models such as CLIP [62] and BLIP [45] are trained on large-scale datasets such as LAION-400M [70] and YFCC-100M [79] and have shown strong performance when evaluated in a zero-shot fashion (i.e. without requiring fine-tuning on specific datasets) for a variety of downstream tasks. One of the most popular applications of VLMs is cross-modal retrieval [86, 88] i.e. retrieving images (text) for a queried text (images). However, image-to-text matching (and vice-




Figure 1: We provide probabilistic embeddings for deterministic pre-trained vision-language models that are *frozen*. By capturing the ambiguity inherently present in the inputs, we obtain well-calibrated uncertainty estimates.

versa) is fundamentally ill-posed due to the inherent ambiguity in either modality [97], i.e. the same caption (or image) can be valid for multiple images (or captions). Therefore, it becomes essential to model the ambiguity inherently present in the various modalities, and combinations thereof.

Instead of mapping inputs to embeddings, probabilistic embedding methods [57, 10] learn to map input samples to distributions. This is achieved by parameterizing the distributions of the embeddings and training a deep neural network to maximize its likelihood. Although they model ambiguities in the embedding space, such probabilistic models require training deep networks from scratch. This requires access to the large-scale datasets and the computational resources of the recent VLMs [62, 35, 51, 74, 45].

We propose *ProbVLM*, a post-hoc probabilistic adapter, the first method to convert the deterministic embeddings provided by a *frozen* large-scale vision-language models into probabilistic ones, as shown in Figure 1. This enables us to efficiently retain the benefits of large-scale pre-training while learning distributions that model the inherent ambiguities in the different modalities. Our *ProbVLM* models the embedding distribution as a heteroscedastic probability distribution and is trained using a combination of intra-modal and cross-modal alignment objectives and provides well-calibrated uncertainty estimates, useful for several tasks.

We demonstrate on two large vision-language datasets, i.e., COCO [46] and Flickr [60], and on two fine-grained image datasets, i.e., CUB [85] and Oxford-Flowers [55] with sentences from [66], that *ProbVLM* learns calibrated

<sup>\*</sup>Authors contributed equally.

uncertainties without requiring large-scale models to be trained from scratch. This sharply contrasts previous works on probabilistic embeddings [57, 10] that train new models from scratch. We perform a series of analyses to understand the impact of the training objective and to study the properties of the resulting uncertainties. Furthermore, we demonstrate that our uncertainty estimates can be used to select the optimal model from a set of finetuned vision-language models on an unlabeled target dataset. They can also be used to choose the most suitable samples for fine-tuning the model in an active learning setup. Finally, with the help of a pre-trained latent diffusion model [67], i.e., *Stable Diffusion*, we decode sampled embeddings from predicted distribution to visualize the predicted embedding distributions. We show that the predicted embedding distributions indeed capture meaningful modes of variation, that may be interpretable.

## 2. Related Work

**Vision-Language Models.** Such models [62, 51, 74, 1, 45, 47, 44, 100, 101, 90] have become ubiquitous in recent times due to their various applications in image classification [105, 21, 106, 50], cross-modal retrieval [4], as well as open-vocabulary semantic segmentation [24, 96]. The most notable among these is CLIP [62], which consists of an image and text encoder trained on 400M image-text pairs with a contrastive objective [28, 58]. As a result, the model is able to project images and text to a shared embedding space. In this paper, we focus on using the shared embedding space for the task of cross-modal retrieval [60, 46]. Recent works have predominantly relied on large-scale pre-training [62, 51, 74, 1, 104, 70, 69] to project images and text to the same metric space. However, it is essential to note that all of these vision-language models [62, 51, 45, 74, 1] provide deterministic mappings that do not model the inherent ambiguity in the inputs. In this work, we turn a deterministic model (i.e., CLIP) into a probabilistic one, without the need of a large-scale dataset.

**Probabilistic Embeddings.** These methods [57, 10, 43] provide an elegant solution to estimate the ambiguity present in the inputs [37]. The key idea here is to map inputs to probability distributions in the embedding space, as opposed to point estimates, thereby modeling the inherent ambiguity present in the input. In the context of cross-modal retrieval, this was done by optimizing a probabilistic analog of the contrastive objective to learn distributions for the image and text inputs [10]. Other works have further improved the performance [43, 59, 34], extended this formulation to achieve compositional retrieval [54], and have applied it to other tasks such as video retrieval [59, 17] and tasks like pose estimation [78]. However, most of these works focus on training a model from scratch, thereby not leveraging the power of the pre-trained models that are widely present. The notable exception to this is Probabilistic Face Embed-

ding (PFE) [73] that proposed to learn a probabilistic embedding while retaining a deterministic pre-trained model for the task of learning face embeddings. However, this was done in a unimodal setting using only images. In this work, we aim to utilize pre-trained vision-language models while providing probabilistic embeddings for both modalities. The probabilistic embeddings derived from our proposed `ProbVLM` are consistent with cross-modal learning at the core of pretrained vision-language models.

**Uncertainty Estimation.** These techniques have been widely explored for different tasks in computer vision [36, 7, 41, 42, 56, 102, 83, 53, 80, 27, 68, 103, 65, 81, 77, 82]. Uncertainties can be broadly categorized into aleatoric [36, 23, 3, 89, 12, 2, 87, 56, 95] and epistemic [25, 7, 41, 91, 20, 33, 19, 18] uncertainties. Uncertainty estimation has been used for a variety of tasks, such as identifying model failure [15, 5, 6, 92] and is extensively used in active learning to select the best samples to train the model [71, 38, 64, 72, 99, 98, 61, 52]. While several of these methods focus on training a new Bayesian model from scratch for quantifying the uncertainties in the prediction, some recent works like [83, 102, 29] have proposed methods to estimate the uncertainties for the pre-trained frozen models. However, these works tackle data from a single modality. This work efficiently estimates the uncertainty for the pre-trained frozen large-scale vision-language model.

## 3. Method

We first describe the problem formulation in Section 3.1. In Section 3.2, we describe our proposed method `ProbVLM` that estimates the complex probability distributions for the embeddings of the frozen deterministic vision-language encoders, quantifying the uncertainties for their predictions.

### 3.1. Problem Formulation

Let  $\mathcal{D} = (\mathcal{I}, \mathcal{C})$  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 where multiplicity is plausible. For a caption  $c \in \mathcal{C}$  (respectively an image  $i \in \mathcal{I}$ ), the set of corresponding images (respectively captions) is given by  $\kappa(c) \subseteq \mathcal{I}$  (respectively  $\kappa(i) \subseteq \mathcal{C}$ ). Recent advances in cross-modal vision-language models [62, 51, 74] often involve learning a shared embedding space,  $\mathcal{Z} \subseteq \mathbb{R}^D$  ( $D$ -dimensional space), for images and texts. This allows quantifying the similarity between cross-modal elements based on their distances in the shared embedding space. The shared embedding space is learned via a set of two encoders:  $\Phi_{\mathcal{V}}(\cdot; \theta_{\mathcal{V}}) : \mathcal{I} \rightarrow \mathcal{Z}$  for the images and  $\Phi_{\mathcal{T}}(\cdot; \theta_{\mathcal{T}}) : \mathcal{C} \rightarrow \mathcal{Z}$  for the texts, where  $\theta_{\mathcal{V}}$  and  $\theta_{\mathcal{T}}$  are the parameters for the respective mapping functions.

We consider a real-world scenario where the above set of encoders have already been trained on vast datasets using large models with high computational cost, e.g., CLIP [62],




Figure 2: Proposed framework (ProbVLM) takes an existing vision-language model and introduces a probabilistic adapter over the image and text encoders. These adapters predict the parameters of a parameterized distribution for a given embedding. Models are trained by minimizing an objective consisting of intra/cross-modal supervision as detailed in Section 3.

SLIP [51], Flava [74] and BLIP [45], are in *frozen state*, i.e., we have  $\Phi_V(\cdot; \theta_V^*)$  and  $\Phi_T(\cdot; \theta_T^*)$ , where  $\theta_V^*, \theta_T^*$  represents the parameters of the pretrained frozen encoders. These encoders are *deterministic* and map an image/text to vectors in the shared space, i.e., given a sample image  $x_V$  (and similarly sample text  $x_T$ ), the encoder provides an embedding  $z_V := \Phi_V(x_V; \theta_V^*)$  (and similarly,  $z_T := \Phi_T(x_T; \theta_T^*)$ ). However, the point estimates,  $z$ , do not capture the ambiguity inherent to these embeddings [57, 10, 17] that are better represented by the probability distribution  $P_{z|x}$ . Therefore, we propose to estimate  $P_{z|x}$  for the pretrained model efficiently, using ProbVLM, quantifying the uncertainties of the output without re-training the encoders.

### 3.2. Building ProbVLM

Despite being deterministic, large-scale *frozen* encoders already provide high-quality point estimates. Our proposed method leverages this fact, using the embeddings  $z$  as estimates for the mean of the desired distribution  $P_{z|x}$ , and estimating the remaining parameters.  $P_{z|x}$  can be modeled as a parametric distribution  $P_{z|x}(z|\{\hat{z}, \hat{\nu} \dots \hat{\rho}\})$  where the parameters can be estimated using a deep neural network [20, 36, 41]. Therefore, we introduce ProbVLM,

$$\Psi(\cdot; \zeta) := (\Psi_V(\cdot; \zeta_V), \Psi_T(\cdot; \zeta_T)) \quad (1)$$

where  $\Psi_V$  and  $\Psi_T$  represents the vision and text encoders parameterized by  $\zeta_V$  and  $\zeta_T$ , respectively. Also,  $\zeta := \zeta_V \cup \zeta_T$  represents the overall parameters for ProbVLM.

that learns to estimate the parameters  $\{\hat{z}, \hat{\nu} \dots \hat{\rho}\}$  with the help of frozen encoders  $\Phi_V(\cdot; \theta_V^*)$  and  $\Phi_T(\cdot; \theta_T^*)$ . The functions  $\Psi_V(\cdot; \zeta_V)$  and  $\Psi_T(\cdot; \zeta_T)$  operate on image and text embeddings respectively, but during training depend on both modalities, as discussed later. We design the learning scheme for  $\Psi(\cdot; \zeta)$  such that: (i) Estimated parameter  $\hat{z}$  should remain faithful to the original unimodal embedding  $z$  (*intra-modal alignment*), this makes the uncertainty of the ProbVLM serve as a good proxy for the uncertainty of frozen encoders. (ii) Estimated parameters  $\{\hat{\nu} \dots \hat{\rho}\}$  should capture the ambiguities and uncertainties present within and across modalities (*cross-modal alignment*). Figure 2 depicts ProbVLM in tandem with the frozen VLM.

**Intra-modal Alignment.** To ensure that the mean of the distribution estimated by  $\Psi(\cdot; \zeta)$  reflects the point estimates provided by the frozen encoders, we set up a probabilistic reconstruction problem for the embeddings within the modalities. That is, for a given sample  $x$  (either from image or text modality), we obtain the embedding from the frozen encoder  $z = \Phi(x; \theta)$  (using the appropriate encoder), then the modality-specific component of  $\Psi(\cdot; \zeta)$  learns to reconstruct the  $z$  (let the reconstruction be called  $\hat{z}$ ). The modality-specific component of  $\Psi(\cdot; \zeta)$  is designed to (i) relax the i.i.d constraints by assuming independent but *not* identically distributed residuals and (ii) learn the *heteroscedasticity* for the residuals at the time of reconstruction that may follow the heavy-tailed distributions [83, 84, 40, 39, 30]. The modality-specific component

is learned by maximizing the likelihood,  $\mathcal{L}(\zeta; \{\mathbf{z}_i\}_{i=1}^N)$  for the embeddings of  $N$  samples in the datasets. That is, the modality-specific optimal parameters are given by,

$$\zeta^* := \operatorname{argmax}_{\zeta} \mathcal{L}(\zeta; \{\mathbf{z}_i\}_{i=1}^N) = \prod_{i=1}^N \frac{\hat{\beta}_i e^{-(|\hat{\mathbf{z}}_i - \mathbf{z}_i|/\hat{\alpha}_i)^{\hat{\beta}_i}}}{2\hat{\alpha}_i \Gamma(1/\hat{\beta}_i)} \quad (2)$$

In the above equation,  $\frac{\hat{\beta}_i e^{-(|\hat{\mathbf{z}}_i - \mathbf{z}_i|/\hat{\alpha}_i)^{\hat{\beta}_i}}}{2\hat{\alpha}_i \Gamma(1/\hat{\beta}_i)}$  represents the *generalized Gaussian distribution* (GGD, represented by  $\mathcal{G}$ ) that is capable of modeling heavy-tailed distributions (note the Gaussian and Laplace are special cases of  $\mathcal{G}$  with  $\alpha = 1, \beta = 2$  and  $\alpha = 1, \beta = 1$ , respectively). The variables  $\hat{\mathbf{z}}_i, \hat{\alpha}_i, \hat{\beta}_i$  are the predicted mean, scale, and shape parameters of  $\mathcal{G}$  from our modality-specific components for the given input  $\mathbf{z}_i$ . We obtain modality-specific optimal parameters by minimizing negative log-likelihood (equivalent to Equation 2). Given  $\mathbf{z}$  and predicted  $\hat{\mathbf{z}}, \hat{\alpha}, \hat{\beta}$ , loss is given by,

$$L_{\text{rec}}(\zeta) := \left( \frac{|\hat{\mathbf{z}} - \mathbf{z}|}{\hat{\alpha}} \right)^{\hat{\beta}} - \log \frac{\hat{\beta}}{\hat{\alpha}} + \log \Gamma\left(\frac{1}{\hat{\beta}}\right) \quad (3)$$

Therefore, the vision-specific component of  $\text{PrObVLM}$ ,  $\Psi(\cdot; \zeta_V)$ , is trained by minimizing the Equation 3 using image embeddings, we denote this loss as  $L_{\text{rec}}^V(\zeta_V)$ . Similarly the text-specific component,  $\Psi(\cdot; \zeta_T)$ , is trained by minimizing  $L_{\text{rec}}^T(\zeta_T)$ . As discussed next, we also enforce cross-modal alignment so that the predicted distribution of  $\text{PrObVLM}$  captures the uncertainties across modalities from one-to-many correspondences for an embedding.

**Cross-modal Alignment.** While the intra-modal alignment seeks to match the means of the output distribution from  $\text{PrObVLM}$  to the embeddings derived from frozen vision-language encoders, we also enforce the image and text embedding output distribution (from  $\text{PrObVLM}$ ) belonging to similar concepts to remain close to each other. That is, given an image and text embedding pair  $(\mathbf{z}_V, \mathbf{z}_T)$  (from frozen model) representing similar concepts, the output distributions from  $\Psi(\cdot; \zeta)$ ,  $\mathcal{G}(\mathbf{z}; \hat{\mathbf{z}}_V, \hat{\alpha}_V, \hat{\beta}_V)$  and  $\mathcal{G}(\mathbf{z}; \hat{\mathbf{z}}_T, \hat{\alpha}_T, \hat{\beta}_T)$  (later referred to as  $\mathcal{G}_V(\mathbf{z})$  and  $\mathcal{G}_T(\mathbf{z})$ ) should match. This can be measured directly from the likelihood as,  $p(\mathbf{z}_V = \mathbf{z}_u)$ , where  $\mathbf{z}_V \sim \mathcal{G}_V(\mathbf{z})$  and  $\mathbf{z}_u \sim \mathcal{G}_T(\mathbf{z})$  as in [73], i.e.,

$$p(\mathbf{z}_V = \mathbf{z}_u) := \iint \mathcal{G}_V(\mathbf{z}_V) \mathcal{G}_T(\mathbf{z}_u) \delta(\mathbf{z}_V - \mathbf{z}_u) d\mathbf{z}_V d\mathbf{z}_u \quad (4)$$

where  $\delta(\cdot)$  refers to the *Dirac-delta distribution*. The above integral can be simplified further by defining  $\Delta \mathbf{z} = \mathbf{z}_V - \mathbf{z}_u$  and seeking  $p(\Delta \mathbf{z}) = 0$ . As both  $\mathbf{z}_V$  and  $\mathbf{z}_u$  are GGD random variables,  $\Delta \mathbf{z}$  follows the distribution based on the

*Bivariate Fox H-function* [76, 48, 49] given by,

$$\Delta \mathbf{z} \sim \frac{1}{2\Gamma(1/\hat{\beta}_V)\Gamma(1/\hat{\beta}_T)} \times \int \mathcal{H}_{1,2}^{1,1} \left[ \begin{matrix} At^2 \\ (0, 1) \end{matrix} \middle| \begin{matrix} (1 - \frac{1}{\hat{\beta}_V}, \frac{1}{\hat{\beta}_T}) \\ (\frac{1}{2}, 1) \end{matrix} \right] \mathcal{H}_{1,2}^{1,1} \left[ \begin{matrix} Bt^2 \\ (0, 1) \end{matrix} \middle| \begin{matrix} (1 - \frac{1}{\hat{\beta}_T}, \frac{1}{\hat{\beta}_V}) \\ (\frac{1}{2}, 1) \end{matrix} \right] \cos t(\mu - z) dt \quad (5)$$

Where  $A = \frac{\hat{\alpha}_V^2 \Gamma(1/\hat{\beta}_V)}{4\Gamma(3/\hat{\beta}_V)}$ ,  $B = \frac{\hat{\alpha}_T^2 \Gamma(1/\hat{\beta}_T)}{4\Gamma(3/\hat{\beta}_T)}$ ,  $\mu = \hat{\mathbf{z}}_V - \hat{\mathbf{z}}_u$ , and  $\mathcal{H}$  is the *Fox H function* [76, 48, 49]. Equation 5 does not provide a scalable objective function suitable for training deep neural networks. Hence, we propose an approximation that is easily scalable for deep-learning models given by,

$$p(\mathbf{z}_V = \mathbf{z}_u) = \iint \mathcal{G}_V(\mathbf{z}_V) \mathcal{G}_T(\mathbf{z}_u) \delta(\mathbf{z}_V - \mathbf{z}_u) d\mathbf{z}_V d\mathbf{z}_u \approx \int \frac{1}{2} (\mathcal{G}_V(\mathbf{z}) \delta(\mathbf{z} - \mathbf{z}_T) + \mathcal{G}_T(\mathbf{z}) \delta(\mathbf{z} - \mathbf{z}_V)) d\mathbf{z} \quad (6)$$

The appendix shows details of the above equation. The first term in the integral,  $\int \mathcal{G}_V(\mathbf{z}) \delta(\mathbf{z} - \mathbf{z}_T) d\mathbf{z}$ , is the likelihood of the text embedding  $\mathbf{z}_T$  under the predicted distribution,  $\mathcal{G}_V(\mathbf{z})$ , for the visual embedding. Similarly, the second term is the likelihood of the visual embedding  $\mathbf{z}_V$  under the predicted distribution,  $\mathcal{G}_T(\mathbf{z})$ , for the text embedding. Negative log of Equation 6 leads to a scalable objective function that can be used to learn the optimal parameters for vision and text components of  $\text{PrObVLM}$  ( $\Psi_V(\cdot; \zeta_V)$  and  $\Psi_T(\cdot; \zeta_T)$ ),

$$L_{\text{cross}}(\zeta_V, \zeta_T) := \underbrace{\left( \frac{|\hat{\mathbf{z}}_V - \mathbf{z}_T|}{\hat{\alpha}_V} \right)^{\hat{\beta}_V} - \log \frac{\hat{\beta}_V}{\hat{\alpha}_V} + \log \Gamma\left(\frac{1}{\hat{\beta}_V}\right)}_{\text{Cross-modal: vision} \rightarrow \text{text}} + \underbrace{\left( \frac{|\hat{\mathbf{z}}_T - \mathbf{z}_V|}{\hat{\alpha}_T} \right)^{\hat{\beta}_T} - \log \frac{\hat{\beta}_T}{\hat{\alpha}_T} + \log \Gamma\left(\frac{1}{\hat{\beta}_T}\right)}_{\text{Cross-modal: text} \rightarrow \text{vision}} \quad (7)$$

The overall objective used for  $\text{PrObVLM}$  is designed to be,

$$L_{\text{PrObVLM}}(\zeta_V, \zeta_T) = L_{\text{rec}}^V(\zeta_V) + L_{\text{rec}}^T(\zeta_T) + \lambda_{\text{cross}} L_{\text{cross}}(\zeta_V, \zeta_T) \quad (8)$$

where  $\lambda_{\text{cross}}$  is a hyperparameter controlling the relative contribution of inter-intra modality terms.

**Uncertainty Quantification.** Given embedding  $\mathbf{z}$  from a frozen encoder, predicted distributions from the trained  $\text{PrObVLM}$  (output from the appropriate component) allows aleatoric uncertainty estimation as  $\hat{\sigma}_{\text{aleatoric}}^2 = \frac{\hat{\alpha}^2 \Gamma(3/\hat{\beta})}{\Gamma(1/\hat{\beta})}$ . Moreover, we design both  $\Psi_V$  and  $\Psi_T$  to be simple 3-layer MLPs with dropout layers (with dropout probability set to 0.1 during training). Activating dropouts during inference, with multiple forward passes (say  $M$ ), allows estimating the epistemic uncertainty,  $\hat{\sigma}_{\text{epistemic}}^2 = \frac{1}{M} \sum_{m=1}^M (\hat{\mathbf{z}}_m -$

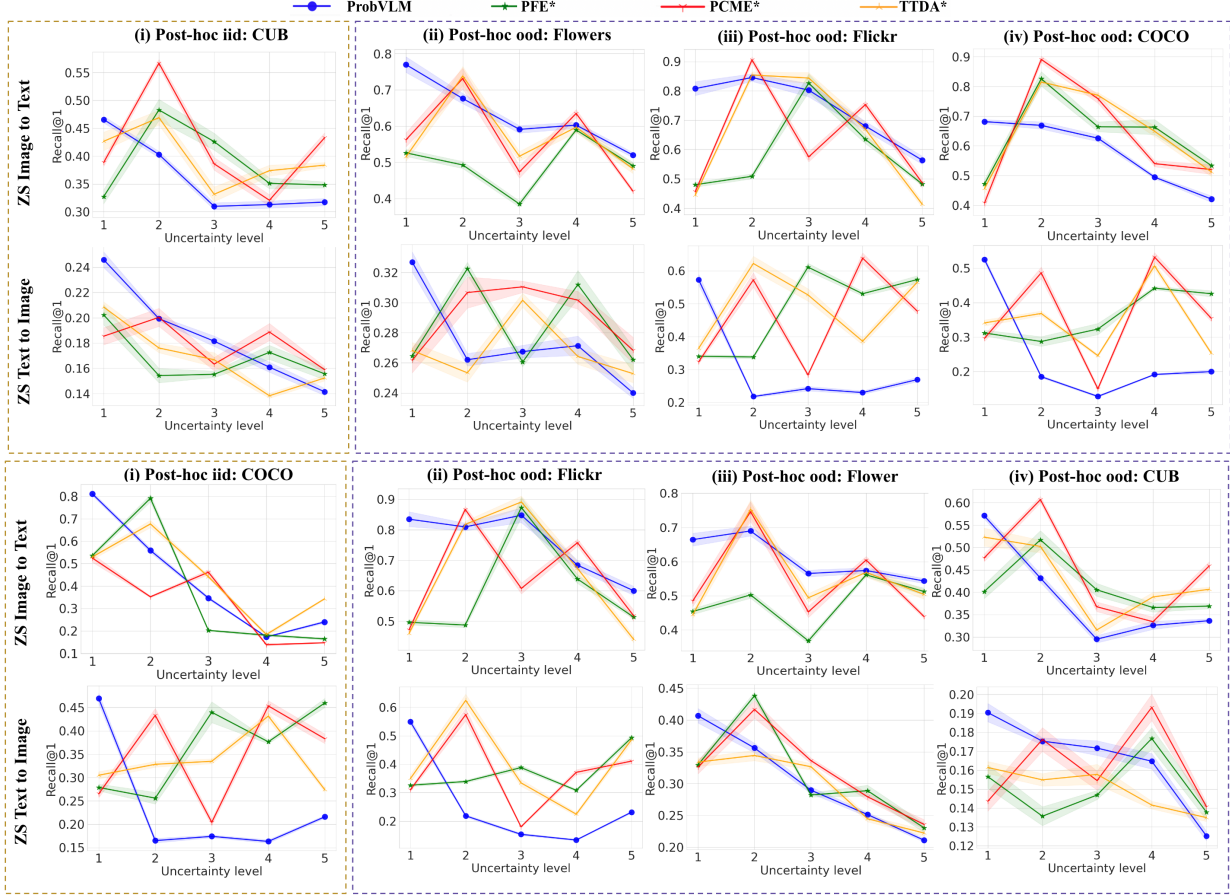


Figure 3: Measuring the calibration with various post-hoc method for Image-to-Text and Text-to-Image retrieval when trained on (top) CUB and (bottom) COCO, and evaluated on CUB, COCO, Flickr, FLO.

$\frac{1}{M} \sum_{j=1}^M \hat{\mathbf{z}}_j)^2$ . We estimate total uncertainty as,

$$\hat{\sigma}_{\text{total}}^2 = \hat{\sigma}_{\text{epistemic}}^2 + \hat{\sigma}_{\text{aleatoric}}^2 \quad (9)$$

### 3.3. Latent Diffusion for Probabilistic Embeddings

For a given text embedding  $\mathbf{z}_{\mathcal{T}}$ , the distribution estimated via `ProbVLM`,  $\mathcal{G}(\mathbf{z}; \hat{\mathbf{z}}_{\mathcal{T}}, \hat{\alpha}_{\mathcal{T}}, \hat{\beta}_{\mathcal{T}})$  can be visualized by drawing samples from the predicted distribution of vectors (say,  $\{\hat{\mathbf{z}}_{\mathcal{T},i}\}_{i=1}^Q$ ) and passing them through a latent diffusion model, e.g., *Stable Diffusion* (say,  $\Omega(\cdot; \theta_{\Omega}^*)$ ) using CLIP text encoder, to synthesize the set of samples (say,  $J$ ) from the corresponding distribution of images, i.e.,

$$J := \{\Omega(\hat{\mathbf{z}}_i; \theta_{\Omega})\}_{i=1}^Q \quad (10)$$

Section 4.4 uses this to visualize the predicted distributions.

## 4. Experiments and Results

We start by highlighting our tasks, datasets, and evaluation metrics. We also compare our model to various state-of-the-art methods quantitatively and qualitatively in Section 4.1. In Section 4.2, we provide an ablation analysis,

and Section 4.3 demonstrates some real-world applications of `ProbVLM` for model selection and active learning.

**Datasets, Metrics, and Baselines.** We use the MS-COCO [46], Flickr-30k [60], and the CUB [85] as they are widely used for cross-modal retrieval [10, 16, 75]. Furthermore, we adapt the Oxford-Flowers 102 (FLO) dataset [55] similar to [10] as an additional benchmark for cross-modal retrieval in a fine-grained setting. We evaluate the performance of both Image-to-Text retrieval and Text-to-Image Retrieval using the Recall@k (R@k) metric. To evaluate the calibration of the uncertainty estimates, we define uncertainty levels [10] and use the Spearman rank correlation (denoted by  $S$ ) between the uncertainty level and the Recall@k for retrieval. For an ideal model, performance would decrease monotonically with increasing uncertainty levels, leading to a score of -1. We also compute the  $R^2$  for the regression fit between the uncertainty levels and R@1 performances to measure if the drop in performance follows a linear trend. Finally, we also measure the product of these two scores (as a unified metric), i.e.,  $-SR^2$ , which should be 1.0 for an ideal model. Since there is *no prior work* to estimate probabilistic embeddings from a deterministic model

VL M	Metrics	i2t				t2i				
		COCO	Flickr	FLO	CUB	COCO	Flickr	FLO	CUB	
CLIP	ProbVLM	S ↓	<b>-0.99</b>	<b>-0.70</b>	<b>-0.90</b>	<b>-0.60</b>	<b>-0.30</b>	<b>-0.70</b>	<b>-0.99</b>	<b>-0.89</b>
		R <sup>2</sup> ↑	<b>0.93</b>	<b>0.71</b>	<b>0.62</b>	<b>0.67</b>	<b>0.35</b>	<b>0.50</b>	<b>0.99</b>	<b>0.70</b>
		-SR <sup>2</sup> ↑	<b>0.93</b>	<b>0.49</b>	<b>0.56</b>	<b>0.40</b>	<b>0.10</b>	<b>0.35</b>	<b>0.99</b>	<b>0.63</b>
	PFE*[73]	S ↓	-0.79	-0.19	0.60	<b>-0.60</b>	0.79	0.30	-0.89	-0.10
		R <sup>2</sup> ↑	0.59	0.01	0.30	0.28	0.74	0.44	0.52	0.00
		-SR <sup>2</sup> ↑	0.47	0.00	-0.18	0.17	-0.59	-0.13	0.47	-0.00
	PCME*[10]	S ↓	-0.89	-0.30	-0.30	<b>-0.60</b>	0.30	0.09	-0.70	0.30
		R <sup>2</sup> ↑	0.75	0.07	0.07	0.20	0.16	0.01	0.57	0.01
		-SR <sup>2</sup> ↑	0.68	0.02	0.02	0.12	-0.05	-0.00	0.40	-0.00
	TTDA[2]	S ↓	-0.79	-0.30	0.00	<b>-0.60</b>	-0.10	-0.19	-0.89	-0.50
		R <sup>2</sup> ↑	0.69	0.09	0.00	0.41	0.26	0.071	0.80	0.15
		-SR <sup>2</sup> ↑	0.55	0.03	0.00	0.24	0.00	0.01	0.73	0.07
BLIP	ProbVLM	S ↓	<b>-0.87</b>	<b>-0.79</b>	<b>-0.74</b>	<b>-0.66</b>	<b>-0.43</b>	<b>-0.38</b>	<b>-0.31</b>	<b>-0.22</b>
		R <sup>2</sup> ↑	<b>0.92</b>	<b>0.83</b>	<b>0.68</b>	<b>0.61</b>	<b>0.52</b>	<b>0.48</b>	<b>0.45</b>	<b>0.38</b>
		-SR <sup>2</sup> ↑	<b>0.80</b>	<b>0.66</b>	<b>0.50</b>	<b>0.40</b>	<b>0.22</b>	<b>0.18</b>	<b>0.14</b>	<b>0.08</b>
	PFE*[73]	S ↓	-0.82	-0.74	-0.63	-0.63	-0.39	-0.32	-0.28	-0.18
		R <sup>2</sup> ↑	0.72	0.76	0.62	0.44	0.48	0.38	0.39	0.37
		-SR <sup>2</sup> ↑	0.58	0.57	0.39	0.27	0.19	0.12	0.11	0.07
	PCME*[10]	S ↓	-0.76	-0.53	-0.60	-0.44	-0.28	-0.26	-0.28	-0.21
		R <sup>2</sup> ↑	0.81	0.56	0.60	0.53	0.50	0.34	0.44	0.36
		-SR <sup>2</sup> ↑	0.62	0.29	0.36	0.23	0.14	0.09	0.12	<b>0.08</b>
	TTDA[2]	S ↓	-0.44	-0.33	-0.74	-0.60	-0.19	-0.26	-0.21	-0.21
		R <sup>2</sup> ↑	0.66	0.56	0.42	0.55	0.49	0.23	0.35	0.36
		-SR <sup>2</sup> ↑	0.29	0.18	0.31	0.33	0.10	0.06	0.07	<b>0.08</b>

Table 1: Metrics to evaluate the calibration of the uncertainty estimates for both CLIP [62] and BLIP [45] Vision-Language models for all considered methods, trained on COCO and evaluated on COCO, Flickr, CUB, and FLO.

in a cross-modal setting, we adapt a few existing ideas for the task. The first baseline is adapted from PFE [73], where we learn the covariances for the heteroscedastic Gaussian distribution while keeping the mean fixed to the embeddings derived from the frozen encoders in each modality. The second is to use the soft-contrastive objective of PCME[10] to train a probabilistic adapter in a post-hoc fashion. Finally, we also have a baseline that performs perform Test-Time Data Augmentation (TTDA) on the inputs [2, 87]. This is done by perturbing the images and masking out words in the text. While TTDA does not require additional training, we train our ProbVLM and other baselines on datasets like COCO, Flickr, CUB, and FLO.

**Implementation Details.** Our ProbVLM consists of a Multi-Layer Perceptron (MLP) for both the image and text encoder, each consisting of an input layer going from embedding dimension to 256, a hidden layer of size 256, and

an output layer going from 256 to embedding dimensions. This is trained for 100 epochs with a learning rate of  $1e^{-4}$ . More details are available in the supplementary.

#### 4.1. Calibrated Uncertainty via ProbVLM

We investigate the calibration of the uncertainty derived from ProbVLM for the cross-modal retrieval task. All models trained on CUB and COCO were evaluated on all four datasets. Calibration plots are illustrated in Figure 3. We observe that the R@1 performance consistently drops for ProbVLM as we increase the uncertainty levels, whereas the baselines rarely see a monotonic drop in performance. We quantify this performance in Table 1. The highest score of 0.93 for  $-SR^2$  (i2t) on the COCO dataset indicates a decreasing performance trend with increasing uncertainty. Notably, the uncertainty estimates indicate the average performance in different bins even when ProbVLM is evaluated on datasets that are different from the train set. In some cases, we see that ProbVLM even achieves a nearly perfect score ( $-SR^2$  of 0.99, with CLIP VLM on FLO, after training on COCO for Image-to-Text Retrieval). On the contrary, we find that the baselines often achieve poor scores. It is important to note that all these models use the same underlying embeddings and achieve the same performance on the retrieval task. Of all the considered methods, ProbVLM provides the most calibrated uncertainty estimates. We see similar trends for ProbVLM with BLIP [45], where ProbVLM achieves a  $-SR^2$  of 0.80, when trained on COCO and evaluated on COCO, compared to other methods like PFE\* (0.58), PCME\* (0.62), and TTDA (0.29).

Figure 4-(Top) visualizes the ambiguities captured by ProbVLM on the visual embeddings. We take a bird image (source) from the CUB dataset and obtain the probability distribution for the visual embedding of that sample; we then compute the likelihood of the visual embeddings (i.e., point estimates derived from CLIP) for the other samples of CUB and COCO datasets, under the source distribution. We notice that within the CUB dataset, the bird images similar to the source image have a higher likelihood compared to other bird images. Also, the images from the COCO dataset tend to have a lower likelihood. However, some images from the COCO dataset have a likelihood similar to the samples from CUB. We visualize these samples and discover them to be bird images. Moreover, the overlapping region between CUB and COCO has samples from the COCO dataset that are ambiguous and related to bird images as they have similar backgrounds, etc. On the contrary, when a similar analysis is performed using the CLIP (by measuring the distance between the embeddings instead of likelihood, Figure 4-(Bottom)), we notice that the two datasets are well separated and ambiguities are not captured.

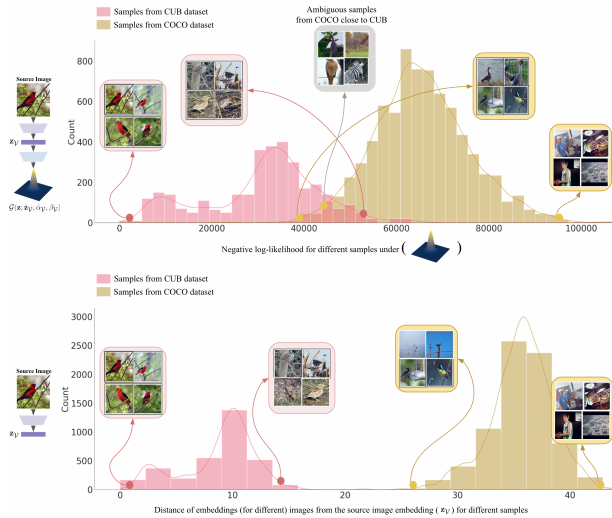


Figure 4: Visualizing the uncertainties of the vision encoder captured by  $\text{ProbVLM}$ . Fixing an image from CUB, we obtain the predicted embedding distribution and compute the likelihood of all other samples in CUB and COCO. We observe that the images in COCO are similar/ambiguous to CUB overlap (Top). However, deterministic embeddings lead to a separation between the two datasets (Bottom).




Figure 5: Plot indicating (left) necessity of the cross-modal alignment and (right) data required to train  $\text{ProbVLM}$ .




Figure 6: Uncertainty increases with increased masking of the input images (Left) and texts (Right). Results with three vision encoders and one language encoder from CLIP.




Figure 7: Results for active learning, with different vision encoders and varying training budgets. For a given encoder, uncertainty-based sampling outperforms random sampling.

## 4.2. Ablations

We ablate different components of our proposed  $\text{ProbVLM}$ , to provide a deeper understanding of its workings. First, we perform a sensitivity analysis on the cross-modal reconstruction objective, as shown in Figure 5-(Left), for  $\text{ProbVLM}$  on BLIP using the COCO dataset. We need a non-zero coefficient of the cross-modal loss to ensure that  $\text{ProbVLM}$  learns meaningful uncertainties that capture the ambiguities across modalities and are well-correlated with its performance on the downstream retrieval task. Similarly, having a large co-efficient for the cross-modal loss could hinder learning a faithful identity reconstruction, thereby hampering the performance of the downstream evaluation. Next, we investigate the amount of data that is required to train  $\text{ProbVLM}$  in Figure 5-(Right). We get satisfactory calibration of the uncertainty estimates while using only 50% of the dataset (shown for  $\text{ProbVLM}$  on BLIP using COCO), indicating that  $\text{ProbVLM}$  is highly data-efficient.

Further, we investigate the uncertainties by masking out increasing portions of the input image/text in Figure 6. We use three different CLIP backbones for the images, ViT-B/32, ViT-B/16, ResNet50, and GPT-based language encoder from CLIP [62, 63]. The mean uncertainty steadily increases as we mask increasing amounts of input.

## 4.3. Applications

We study the utility of the uncertainty estimates derived from  $\text{ProbVLM}$  on two critical applications not well reviewed for VLMs: active learning and model selection.

**Active Learning.** Here, we choose a small subset of the unlabeled dataset to fine-tune the model [11]. In this case, we wish to finetune the CLIP model on the FLO dataset while using a limited amount of labeled data. To achieve this, we estimate the uncertainty of the image embeddings us-

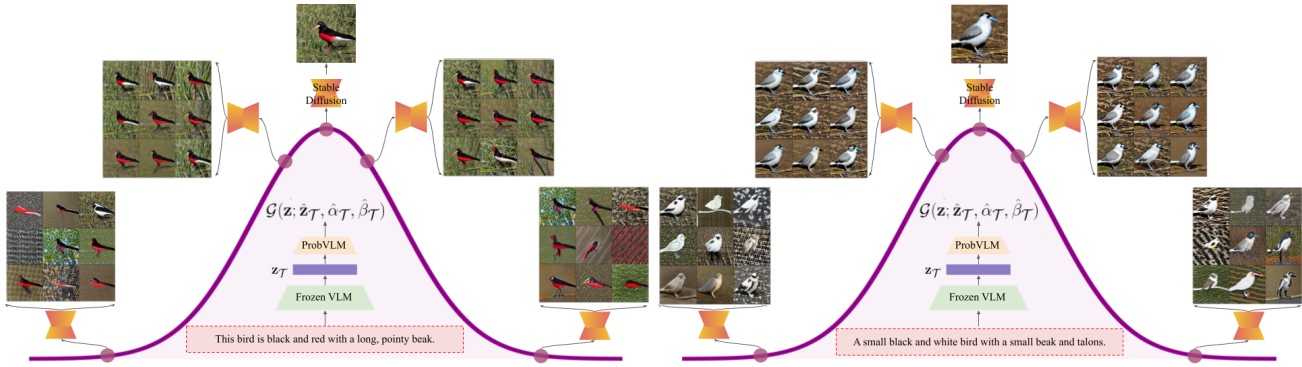


Figure 8: Visualizing the predicted embedding distributions from `ProbVLM` using a large-scale pre-trained diffusion model, i.e., *Stable Diffusion*. The example is shown for two different captions from CUB dataset, for which the point-estimate embedding vector is obtained via CLIP, and the distribution is obtained via `ProbVLM`.

D	Models	Metrics			
		Uncertainty score	R@1	R@5	R@10
CUB	CLIP-ViT32-COCO	11.92	31.5	61.0	75.8
	CLIP-ViT32-Flickr	<b>9.37</b>	<b>32.4</b>	<b>64.2</b>	<b>76.9</b>
	CLIP-ViT32-FLO	15.43	22.8	49.8	64.9
FLO	CLIP-ViT32-COCO	<b>11.83</b>	47.9	79.2	88.5
	CLIP-ViT32-Flickr	13.55	<b>49.5</b>	<b>84.6</b>	<b>93.9</b>
	CLIP-ViT32-CUB	18.39	37.7	69.4	82.8
Flickr	CLIP-ViT32-COCO	<b>9.61</b>	<b>88.8</b>	<b>97.8</b>	<b>99.8</b>
	CLIP-ViT32-CUB	16.49	25.8	47.4	55.6
	CLIP-ViT32-FLO	13.67	52.8	77.8	85.2
COCO	CLIP-ViT32-Flickr	<b>7.28</b>	<b>58.1</b>	<b>80.9</b>	<b>88.2</b>
	CLIP-ViT32-CUB	15.37	8.8	21.7	29.8
	CLIP-ViT32-FLO	12.44	23.9	46.6	58.8

Table 2: Results for the model selection experiment. `ProbVLM` accurately identifies the best performing source model using only unlabeled samples of the target dataset.

ing `ProbVLM` (trained using a diverse dataset like COCO). We then select the top-k samples sorted by their mean uncertainty in the visual embeddings and fine-tune the CLIP model using them with a contrastive objective [62]. Results with varying budgets are shown in Figure 7. Selecting samples based on uncertainty scores significantly outperforms random sampling for all considered visual backbones.

**Pretrained Model Selection.** We are given a set of models trained on different data distributions. We aim to select the best model for the target distribution for which we have unlabeled samples. This has been explored mostly in the context of classification previously [22, 26, 8, 9, 13, 14].

We consider the specific case of having the CLIP models fine-tuned on three datasets, and the fourth dataset is held out, for which we only have the images. We compute the mean uncertainty on these images using `ProbVLM` whose weights are interpolated from all the source datasets [93, 94, 32, 31]. This is to ensure that the uncertainties on all these models are comparable. The results for this experiment are

shown in Table 2. On CUB, Flickr, and COCO, the source model with the lowest uncertainty has the best performance on the target dataset, and on FLO dataset, the model with the least uncertainty has the 2nd best performance (R@1 of 47.9 vs 49.5 for the best model). This indicates that the uncertainties provided by `ProbVLM` can be used as a signal to predict the performance on unlabelled samples for retrieval.

#### 4.4. Latent Diffusion for Embedding Uncertainty

To further understand the semantics of the predicted embedding distributions from the `ProbVLM`, we visualize the text embedding distributions by sampling the embedding vectors from the predicted distribution for a caption (converted to embedding vector using CLIP) and passing it through the pre-trained latent diffusion model using CLIPs text encoder, *stable diffusion*, as shown in Figure 8 and described in details in Section 3.3. We observe from Figure 8 that the samples obtained closer to the mean (i.e., sampled embedding vector similar to the one generated by CLIP for the caption) lead to meaningful variations in the generated images, e.g., for the left caption, close to the mean of the distribution, the generated samples show variations in the shape and colour of the beak, wings, and feet. Whereas far away from the mean of the distributions, i.e., on the tails, we start seeing images with strong artifacts that no longer preserves the semantics of the caption. We observe this for another example as well shown in Figure 8-(Right). More results are available in the supplementary.

## 5. Conclusion

We introduce `ProbVLM`, a post-hoc method for estimating the embedding distribution for a frozen large-scale deterministic vision-language model. We efficiently estimate calibrated uncertainties using our framework and show that such calibrated estimates have a variety of applications in downstream tasks such as model selection and active learning. Furthermore, we perform experiments to interpret em-



bedding distribution predicted by `ProbVLM` using a large-scale pre-trained latent diffusion model (i.e., *Stable Diffusion*). We hope our work highlights and inspires future work on efficient methods for probabilistic embeddings.

**Acknowledgements.** This work was supported by DFG project number 276693517, by BMBF FKZ: 01IS18039A, by the ERC (853489 - DEXIM), by EXC number 2064/1 – project number 390727645, and by the MUR PNRR project FAIR - Future AI Research (PE00000013) funded by NextGenerationEU. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Uddeshya Upadhyay and Shyamgopal Karthik.

## References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022. 1, 2
- [2] Murat Seckin Ayhan and Philipp Berens. Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks. In *MIDL*, 2018. 2, 6
- [3] Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. Estimating and exploiting the aleatoric uncertainty in surface normal estimation. In *ICCV*, 2021. 2
- [4] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. A clip-hitchhiker’s guide to long video retrieval. *arXiv preprint arXiv:2205.08508*, 2022. 2
- [5] Victor Besnier, Andrei Bursuc, David Picard, and Alexandre Briot. Triggering failures: Out-of-distribution detection by learning from local adversarial attacks in semantic segmentation. In *ICCV*, 2021. 2
- [6] Victor Besnier, David Picard, and Alexandre Briot. Learning uncertainty for safety-oriented semantic segmentation in autonomous driving. In *ICIP*, 2021. 2
- [7] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *ICML*, 2015. 2
- [8] Mayee Chen, Karan Goel, Nimit S Sohoni, Fait Poms, Kayvon Fatahalian, and Christopher Ré. Mandoline: Model evaluation under distribution shift. In *ICML*, 2021. 8
- [9] Ching-Yao Chuang, Antonio Torralba, and Stefanie Jegelka. Estimating generalization under distribution shifts via domain-invariant representations. *arXiv preprint arXiv:2007.03511*, 2020. 8
- [10] Sanghyuk Chun, Seong Joon Oh, Rafael Sampaio De Rezende, Yannis Kalantidis, and Diane Larlus. Probabilistic embeddings for cross-modal retrieval. In *CVPR*, 2021. 1, 2, 3, 5, 6
- [11] Cody Coleman, Christopher Yeh, Stephen Mussmann, Baharan Mirzasoleiman, Peter Bailis, Percy Liang, Jure Leskovec, and Matei Zaharia. Selection via proxy: Efficient data selection for deep learning. In *ICLR*, 2020. 7
- [12] Yufei CUI, Ziquan Liu, Xiangyu Liu, Xue Liu, Cong Wang, Tei-Wei Kuo, Chun Jason Xue, and Antoni B. Chan. BayesMIL: A new probabilistic perspective on attention-based multiple instance learning for whole slide images. In *ICLR*, 2023. 2
- [13] Weijian Deng, Stephen Gould, and Liang Zheng. What does rotation prediction tell us about classifier accuracy under varying testing environments? In *ICML*, 2021. 8
- [14] Weijian Deng and Liang Zheng. Are labels always necessary for classifier accuracy evaluation? In *CVPR*, 2021. 8
- [15] Abbas Emami-Naeini, Muhammad M Akhter, and Stephen M Rock. Effect of model uncertainty on failure detection: the threshold selector. *IEEE Transactions on Automatic Control*, 1988. 2
- [16] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. In *BMVC*, 2018. 5
- [17] Bo Fang, Chang Liu, Yu Zhou, Min Yang, Yuxin Song, Fu Li, Weiping Wang, Xiangyang Ji, Wanli Ouyang, et al. Uatvr: Uncertainty-adaptive text-video retrieval. *arXiv preprint arXiv:2301.06309*, 2023. 2, 3
- [18] Gianni Franchi, Andrei Bursuc, Emanuel Aldea, Séverine Dubuisson, and Isabelle Bloch. Tradi: Tracking deep neural network weight distributions. In *ECCV*, 2020. 2
- [19] Gianni Franchi, Andrei Bursuc, Emanuel Aldea, Séverine Dubuisson, and Isabelle Bloch. One versus all for deep neural network for uncertainty (ovnni) quantification. *IEEE Access*, 10:7300–7312, 2021. 2
- [20] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016. 2, 3
- [21] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. 2
- [22] Saurabh Garg, Sivaraman Balakrishnan, Zachary C Lipton, Behnam Neyshabur, and Hanie Sedghi. Leveraging unlabeled data to predict out-of-distribution performance. In *ICLR*, 2022. 8
- [23] Jochen Gast and Stefan Roth. Lightweight probabilistic deep networks. In *CVPR*, 2018. 2
- [24] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Open-vocabulary image segmentation. *arXiv preprint arXiv:2112.12143*, 2021. 2
- [25] Alex Graves. Practical variational inference for neural networks. *NIPS*, 2011. 2
- [26] Devin Guillory, Vaishaal Shankar, Sayna Ebrahimi, Trevor Darrell, and Ludwig Schmidt. Predicting with confidence on unseen distributions. In *ICCV*, 2021. 8
- [27] Hongji Guo, Hanjing Wang, and Qiang Ji. Uncertainty-guided probabilistic transformer for complex action recognition. In *CVPR*, 2022. 2
- [28] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, 2010. 2

- [29] Julia Hornauer and Vasileios Belagiannis. Gradient-based uncertainty for monocular depth estimation. In *ECCV*, 2022. 2
- [30] Todd Huster, Jeremy Cohen, Zinan Lin, Kevin Chan, Charles Kamhoua, Nandi O. Leslie, Cho-Yu Jason Chiang, and Vyas Sekar. Pareto gan: Extending the representational power of gans to heavy-tailed distributions. In *ICML*, 2021. 3
- [31] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *arXiv preprint arXiv:2212.04089*, 2022. 8
- [32] Gabriel Ilharco, Mitchell Wortsman, Samir Yitzhak Gadre, Shuran Song, Hannaneh Hajishirzi, Simon Kornblith, Ali Farhadi, and Ludwig Schmidt. Patching open-vocabulary models by interpolating weights. *arXiv preprint arXiv:2208.05592*, 2022. 8
- [33] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018. 2
- [34] Yatai Ji, Junjie Wang, Yuan Gong, Lin Zhang, Yanru Zhu, Hongfa Wang, Jiaying Zhang, Tetsuya Sakai, and Yujie Yang. Map: Modality-agnostic uncertainty-aware vision-language pre-training model. *arXiv preprint arXiv:2210.05335*, 2022. 2
- [35] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. 1
- [36] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *NIPS*, 2017. 2, 3
- [37] Michael Kirchhof, Enkelejda Kasneci, and Seong Joon Oh. Probabilistic contrastive learning recovers the correct aleatoric uncertainty of ambiguous inputs. *arXiv preprint arXiv:2302.02865*, 2023. 2
- [38] Andreas Kirsch, Joost Van Amersfoort, and Yarin Gal. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *NeurIPS*, 2019. 2
- [39] Nitin Kumar, Ajit V Rajwade, Sharat Chandran, and Suyash P Awate. Kernel generalized gaussian and robust statistical learning for abnormality detection in medical images. In *ICIP*, 2017. 3
- [40] Nitin Kumar, Ajit V Rajwade, Sharat Chandran, and Suyash P Awate. Kernel generalized-gaussian mixture model for robust abnormality detection. In *MICCAI*, 2017. 3
- [41] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *NIPS*, 2017. 2, 3
- [42] Max-Heinrich Laves, Sontje Ihler, Jacob F Fast, Lüder A Kahrs, and Tobias Ortmaier. Well-calibrated regression uncertainty in medical imaging with deep learning. In *MIDL*, 2020. 2
- [43] Hao Li, Jingkuan Song, Lianli Gao, Pengpeng Zeng, Haonan Zhang, and Gongfu Li. A differentiable semantic metric approximation in probabilistic embedding for cross-modal retrieval. In *NeurIPS*, 2022. 2
- [44] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 2
- [45] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 1, 2, 3, 6
- [46] 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 *ECCV*, 2014. 1, 2, 5
- [47] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilt: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *NeurIPS*, 2019. 2
- [48] Francesco Mainardi, Gianni Pagnini, and R.K. Saxena. Fox h functions in fractional diffusion. *Journal of Computational and Applied Mathematics*, 2005. Proceedings of the Seventh International Symposium on Orthogonal Polynomials, Special Functions and Applications. 4
- [49] Arakaparampil M Mathai, Ram Kishore Saxena, and Hans J Haubold. *The H-function: theory and applications*. Springer Science & Business Media, 2009. 4
- [50] Yifei Ming, Ziyang Cai, Jiuxiang Gu, Yiyun Sun, Wei Li, and Yixuan Li. Delving into out-of-distribution detection with vision-language representations. In *NeurIPS*, 2022. 2
- [51] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *ECCV*, 2022. 1, 2, 3
- [52] Sri Aurobindo Munagala, Sidhant Subramanian, Shyamgopal Karthik, Ameya Prabhu, and Anoop Namboodiri. Clactive: Episodic memories for rapid active learning. In *CoLLAS*, 2022. 2
- [53] Jurijs Nazarovs, Zhichun Huang, Songwong Tasneeyapant, Rudrasis Chakraborty, and Vikas Singh. Understanding uncertainty maps in vision with statistical testing. In *CVPR*, 2022. 2
- [54] Andrei Neculai, Yanbei Chen, and Zeynep Akata. Probabilistic compositional embeddings for multimodal image retrieval. In *CVPR-W*, 2022. 2
- [55] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *ICVGIP*, 2008. 1, 5
- [56] David A Nix and Andreas S Weigend. Estimating the mean and variance of the target probability distribution. In *ICNN*, 1994. 2
- [57] Seong Joon Oh, Kevin Murphy, Jiyan Pan, Joseph Roth, Florian Schroff, and Andrew Gallagher. Modeling uncertainty with hedged instance embedding. *ICLR*, 2019. 1, 2, 3
- [58] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2
- [59] Jungin Park, Jiyoung Lee, Ig-Jae Kim, and Kwanghoon Sohn. Probabilistic representations for video contrastive learning. In *CVPR*, 2022. 2

- [60] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *CVPR*, 2015. [1](#), [2](#), [5](#)
- [61] Ameya Prabhu, Charles Dognin, and Maneesh Singh. Sampling bias in deep active classification: An empirical study. 2019. [2](#)
- [62] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. [1](#), [2](#), [6](#), [7](#), [8](#)
- [63] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. [7](#)
- [64] Anant Raj and Francis Bach. Convergence of uncertainty sampling for active learning. In *ICML*, 2022. [2](#)
- [65] Vikrant Rangnekar, Uddeshya Upadhyay, Zeynep Akata, and Biplab Banerjee. Usim-dal: Uncertainty-aware statistical image modeling-based dense active learning for super-resolution. 2023. [2](#)
- [66] Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. Learning deep representations of fine-grained visual descriptions. In *CVPR*, 2016. [1](#)
- [67] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022. [2](#)
- [68] Subhankar Roy, Martin Trapp, Andrea Pilzer, Juho Kannala, Nicu Sebe, Elisa Ricci, and Arno Solin. Uncertainty-guided source-free domain adaptation. In *ECCV*, 2022. [2](#)
- [69] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*, 2022. [2](#)
- [70] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. [1](#), [2](#)
- [71] Burr Settles. Active learning literature survey. 2009. [2](#)
- [72] Alexander Shapeev, Konstantin Gubaev, Evgenii Tsybalov, and Evgeny Podryabinkin. Active learning and uncertainty estimation. *Machine Learning Meets Quantum Physics*, 2020. [2](#)
- [73] Yichun Shi and Anil K Jain. Probabilistic face embeddings. In *ICCV*, 2019. [2](#), [4](#), [6](#)
- [74] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *CVPR*, 2022. [1](#), [2](#), [3](#)
- [75] Yale Song and Mohammad Soleymani. Polysemous visual-semantic embedding for cross-modal retrieval. In *CVPR*, 2019. [5](#)
- [76] Hamza Soury and Mohamed-Slim Alouini. New results on the sum of two generalized gaussian random variables. In *GlobalSIP*, 2015. [4](#)
- [77] Viswanath P Sudarshan, Uddeshya Upadhyay, Gary F Egan, Zhaolin Chen, and Suyash P Awate. Towards lower-dose pet using physics-based uncertainty-aware multimodal learning with robustness to out-of-distribution data. *Medical Image Analysis*, 2021. [2](#)
- [78] Jennifer J Sun, Jiaping Zhao, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, and Ting Liu. View-invariant probabilistic embedding for human pose. In *ECCV*, 2020. [2](#)
- [79] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 2016. [1](#)
- [80] Dustin Tran, Jeremiah Liu, Michael W Dusenberry, Du Phan, Mark Collier, Jie Ren, Kehang Han, Zi Wang, Zelda Mariet, Huiyi Hu, et al. Plex: Towards reliability using pretrained large model extensions. *arXiv preprint arXiv:2207.07411*, 2022. [2](#)
- [81] Uddeshya Upadhyay, Yanbei Chen, and Zeynep Akata. Robustness via uncertainty-aware cycle consistency. *NeurIPS*, 2021. [2](#)
- [82] Uddeshya Upadhyay, Yanbei Chen, Tobias Hepp, Sergios Gatidis, and Zeynep Akata. Uncertainty-guided progressive gans for medical image translation. In *MICCAI*. Springer, 2021. [2](#)
- [83] Uddeshya Upadhyay, Shyamgopal Karthik, Yanbei Chen, Massimiliano Mancini, and Zeynep Akata. Bayescap: Bayesian identity cap for calibrated uncertainty in frozen neural networks. In *ECCV*, 2022. [2](#), [3](#)
- [84] Uddeshya Upadhyay, Viswanath P Sudarshan, and Suyash P Awate. Uncertainty-aware gan with adaptive loss for robust mri image enhancement. In *ICCV-W*, 2021. [3](#)
- [85] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. [1](#), [5](#)
- [86] Bokun Wang, Yang Yang, Xing Xu, Alan Hanjalic, and Heng Tao Shen. Adversarial cross-modal retrieval. In *ACM-MM*, 2017. [1](#)
- [87] Guotai Wang, Wenqi Li, Michael Aertsen, Jan Deprest, Sébastien Ourselin, and Tom Vercauteren. Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing*, 338:34–45, 2019. [2](#), [6](#)
- [88] Kaiye Wang, Qiyue Yin, Wei Wang, Shu Wu, and Liang Wang. A comprehensive survey on cross-modal retrieval. *arXiv preprint arXiv:1607.06215*, 2016. [1](#)
- [89] Xi Wang and Laurence Aitchison. Bayesian ood detection with aleatoric uncertainty and outlier exposure. *AABI*, 2021. [2](#)
- [90] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*, 2021. [2](#)
- [91] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *ICML*, 2011. [2](#)

- [92] Spencer Whitehead, Suzanne Petryk, Vedaad Shakib, Joseph Gonzalez, Trevor Darrell, Anna Rohrbach, and Marcus Rohrbach. Reliable visual question answering: Abstain rather than answer incorrectly. In *ECCV*, 2022. 2
- [93] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *ICML*, 2022. 8
- [94] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *CVPR*, 2022. 8
- [95] Binhui Xie, Longhui Yuan, Shuang Li, Chi Harold Liu, and Xinjing Cheng. Towards fewer annotations: Active learning via region impurity and prediction uncertainty for domain adaptive semantic segmentation. In *CVPR*, 2022. 2
- [96] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for zero-shot semantic segmentation with pre-trained vision-language model. *arXiv preprint arXiv:2112.14757*, 2021. 2
- [97] Gengcong Yang, Jingyi Zhang, Yong Zhang, Baoyuan Wu, and Yujia Yang. Probabilistic modeling of semantic ambiguity for scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12527–12536, 2021. 1
- [98] Yazhou Yang and Marco Loog. Active learning using uncertainty information. In *ICPR*, 2016. 2
- [99] Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G Hauptmann. Multi-class active learning by uncertainty sampling with diversity maximization. *IJCV*, 2015. 2
- [100] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021. 2
- [101] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. 2
- [102] Xuanlong Yu, Gianni Franchi, and Emanuel Aldea. Slurp: Side learning uncertainty for regression problems. In *BMVC*, 2021. 2
- [103] Yu Yu, Hassan Sajjad, and Jia Xu. Learning uncertainty for unknown domains with zero-target-assumption. In *ICLR*, 2023. 2
- [104] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *CVPR*, 2021. 2
- [105] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, 2022. 2
- [106] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 2022. 2