

# 000 001 002 003 004 005 REVISITING CONFIDENCE CALIBRATION FOR MISCLAS- 006 SIFICATION DETECTION IN VLMS 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053

## ABSTRACT

Confidence calibration has been widely studied to improve the trustworthiness of predictions in vision-language models (VLMs). However, we theoretically reveal that standard confidence calibration inherently *impairs* the ability to distinguish between correct and incorrect predictions (*i.e.*, Misclassification Detection, MisD), which is crucial for reliable deployment of VLMs in high-risk applications. In this paper, we investigate MisD in VLMs and propose confidence recalibration to enhance MisD. Specifically, we design a new confidence calibration objective to replace the standard one. This modification theoretically achieves higher precision in the MisD task and reduces the mixing of correct and incorrect predictions at every confidence level, thereby overcoming the limitations of standard calibration for MisD. As the calibration objective is not differentiable, we introduce a differentiable surrogate loss to enable better optimization. Moreover, to preserve the predictions and zero-shot ability of the original VLM, we develop a post-hoc framework, which employs a lightweight meta network to predict sample-specific temperature factors, trained with the surrogate loss. Extensive experiments across multiple metrics validate the effectiveness of our approach on MisD.

## 1 INTRODUCTION

Pretrained Vision-language models (VLMs) (Radford et al., 2021; Zhou et al., 2022b;a; Khattak et al., 2023), such as CLIP (Radford et al., 2021), have demonstrated impressive zero-shot capabilities. Owing to their strong generalization ability and pretrained nature, [they have been applied to a wide range of downstream tasks](#), including autonomous driving (Cui et al., 2024), medical diagnosis (Zhao et al., 2023), and 3D scene understanding (Chen et al., 2023). While these models improve flexibility and accuracy, [ensuring their reliability remains essential](#), which is crucial for real-world deployment and safety-critical applications. Consequently, confidence calibration (Guo et al., 2017), which adjusts model confidence to better match true correctness, [is therefore an important component in developing reliable VLM-based systems](#).

Confidence calibration aims to align the model’s predicted confidence with the true likelihood of correctness. A classical calibration method is Temperature Scaling (Guo et al., 2017), which adjusts the sharpness of the output probabilities using a temperature coefficient to better align confidence with empirical accuracy. However, a single global temperature overlooks the instance-wise variation in confidence miscalibration. Therefore, many instance-wise calibration methods have been proposed (Huang et al., 2025; Krishnan & Tickoo, 2020). In VLMs, confidence calibration becomes more challenging due to modality-specific factors or the process of fine-tuning. [To address these issues, several works \(Wang et al., 2024; Lv et al., 2025\) show that distances in the text embedding space play a crucial role in calibration errors, and they incorporate this text-modality signal as an additional temperature adjustment.](#) In contrast, to mitigate miscalibration introduced by fine-tuning, DOR (Wang et al., 2025) uses a large vocabulary set to preserve the semantic structure of the pretrained CLIP, thereby reducing the shift in text features caused by prompt tuning.

Although previous methods can achieve great calibration for VLMs, many high-risk tasks rely more on the model’s ability to rank predictions by correctness than confidence calibration performance. Concretely, in the misclassification detection task, it is more desirable for the VLMs to consistently assign lower confidence scores to misclassified samples than to correctly classified ones, rather than merely aligning predicted probabilities with accuracy, which is illustrated in Figure 1(a).

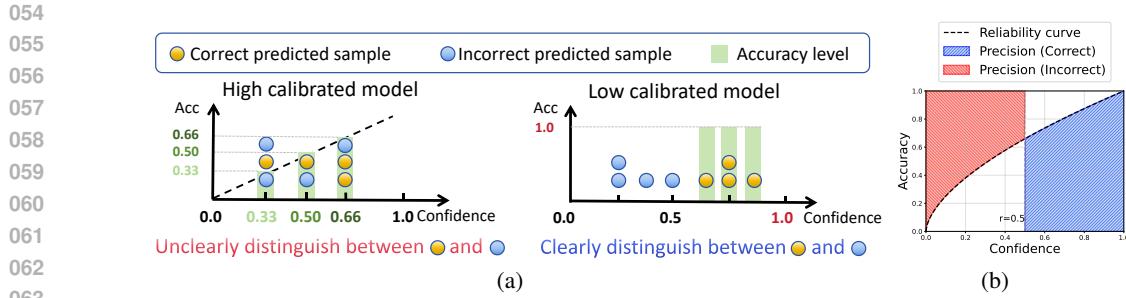


Figure 1: (a) Comparison between a calibrated model and one with improved MisD performance. Left: a calibrated model aligns predicted confidence with accuracy. Right: a model with higher MisD performance that better ranks correct predictions above incorrect ones. (b) Illustration of the relationship between the reliability curve and the MisD. Given a reliability curve and a confidence threshold  $r$ , samples with confidence above  $r$  are regarded as correct predictions. Their precision can be derived from the area under the curve within that interval. Conversely, samples with confidence below  $r$  are treated as incorrect predictions, and their precision can be derived from the area above the curve.

In this work, we first analyze the reliability diagram of confidence calibration for MisD, and find that the region under (above) the reliability curve is related to the precision of detecting correct (incorrect) prediction, as shown in Figure 1(b). Motivated by this insight, we reveal that standard calibration inherently limits the upper bound of this precision by analyzing its reliability curve. To remedy this limitation, we introduce a new target curve. It increases the area under the curve in the high-confidence region and enlarges the area above the curve in the low-confidence region, as illustrated in Figure 2. This design theoretically improves precision and reduces the mixing between correctly and incorrectly classified samples, thereby encouraging better separation between them. Although the proposed reliability curve is ideal for MisD, it cannot be directly expressed as a differentiable loss function. To bridge this gap, we design a surrogate loss that faithfully captures its effect, making the objective practically optimizable. In practice, to avoid interfering with the underlying VLM predictions, we adopt a post-hoc framework. Specifically, we introduce a lightweight meta network that predicts the temperature scaling factor to each sample, conditioned on the VLM logits, image embeddings, and predicted text embeddings. This meta network is trained using our surrogate loss. Compared to previous works<sup>1</sup>, our contributions are listed as follows:

- We make the first attempt to address MisD in VLMs from a calibration perspective. By analyzing the reliability diagram for MisD, we uncover key limitations of standard calibration and introduce a new perspective to systematically analyze MisD and guide its improvement.
- We propose a new calibration objective tailored to improving MisD performance. The theoretical analysis guarantees that the proposed objective can improve the precision in the MisD task and reduce mixing of correct and incorrect predictions on every confidence level.
- To ensure that the proposed loss does not interfere with the pretrained capabilities of VLMs, we develop a post-hoc calibration framework that learns individualized temperature coefficients for each sample. Empirical results show that the proposed method consistently improves MisD over the recent uncertainty estimation methods across diverse settings.

## 2 PRELIMINARIES

**Contrastive Language-Image Pretraining (CLIP).** CLIP is a powerful vision-language model that aligns image and text representations in a shared embedding space through contrastive learning (Radford et al., 2021). Owing to its contrastive training on large-scale image-text pairs, CLIP demonstrates strong zero-shot generalization and is readily deployable in various downstream scenarios. Let  $\xi : \mathbf{x} \rightarrow \mathbb{R}^d$  and  $\psi : \mathbf{t} \rightarrow \mathbb{R}^d$  denote the image and text encoders of CLIP, respectively. Given an image instance  $v$  and a text label  $c$ , the output logit of CLIP can be formulated as:

$$z_{v,c} = \tau_{\text{clip}} \cdot \text{sim}(\xi(\mathbf{x}_v), \psi(\mathbf{t}_c)), \quad (1)$$

<sup>1</sup>The details of related work are summarized in the Appendix A.

108 where  $\text{sim}(\cdot, \cdot)$  denotes cosine similarity,  $\mathbf{t}_c$  is a hand-crafted prompt, typically set to “a photo  
 109 of a  $\{\text{class}\}$ ” and  $\tau_{\text{clip}}$  is a fixed constant, usually set to 100. In multi-class classification, let  
 110  $C = \{c_0, c_1, \dots, c_{|C|-1}\}$  denote the set of candidate classes. The predicted label corresponds to the  
 111 class with the highest predicted probability: which can be formally expressed as follows:

$$112 \quad \hat{y} = \arg \max_{c \in C} (e^{z_c} / \sum_{i=0}^{|C|-1} e^{z_i}), \quad (2)$$

115 where  $\hat{y}$  is the predicted class and the associated probability (i.e.,  $s = \max(e^{z_c} / \sum_{i=0}^{|C|-1} e^{z_i})$ ) is  
 116 referred to as the confidence of the prediction.

117 **Confidence Calibration.** Confidence calibration is particularly important in high-risk and open-  
 118 world scenarios (Guo et al., 2017). It refers to alignment between a model’s predicted confidence and  
 119 its actual accuracy, making the predicted confidence more trustworthy. For example, if the average  
 120 confidence is 0.8, then approximately 80% of the predicted examples should be correct. Formally,  
 121 the objective of confidence calibration can be defined as follows:

$$122 \quad \mathbb{P}(\hat{y} = y | s = p) = p, \forall p \in [0, 1], \quad (3)$$

123 where  $y$  is the ground-truth label. The performance of confidence calibration can be evaluated using  
 124 a **reliability diagram**, where the vertical axis represents accuracy and the horizontal axis represents  
 125 confidence, as shown in Figure 1(b). Points closer to the diagonal indicate better calibration. The  
 126 **reliability curve** is obtained by smoothly connecting these points, with the diagonal line representing  
 127 the perfect-calibration curve (i.e.,  $f(x) = x$ ). Confidence calibration aims to make the model’s  
 128 reliability curve as close as possible to the perfect calibration curve.

129 **Misclassification detection (MisD).** Misclassification detection is a critical safeguard for deploying  
 130 models in real-world applications (Hendrycks & Gimpel, 2016), aiming to distinguish incorrect  
 131 predictions from correct predictions based on confidence ranking. Formally, given a confidence  
 132 threshold  $r \in [0, 1]$ , predictions with confidence above  $r$  are detected as correct predictions, while  
 133 those with confidence below  $r$  are detected as incorrect predictions.

### 135 3 METHOD

137 **Overview.** The overall research framework can be divided into three parts: ① We first analyze  
 138 the upper bound of calibration for MisD, demonstrating its inherent limitation. To overcome this  
 139 limitation, we redesign the calibration reliability curve tailored for MisD. In addition, we establish  
 140 several favorable properties of the proposed curve, highlighting its theoretical soundness and practical  
 141 advantages in improving MisD. ② As the reliability curve is not differentiable, we design a surrogate  
 142 loss (i.e.,  $\mathcal{L}_{\text{SUR}}$ ) to realize its optimization. ③ Finally, we present a post-hoc and lightweight  
 143 implementation, which adjusts confidence without modifying the VLMs’ parameters.

#### 145 3.1 REVISITING CALIBRATION WITH MISD

147 The reliability diagram is a key tool for evaluating the calibration of probabilistic classifiers, where the  
 148 diagonal ( $\text{acc} = s$ ) denotes perfect calibration. Thus, calibration quality is assessed by how closely  
 149 the reliability curve aligns with this line. So far, it has been used exclusively for this purpose. In this  
 150 work, we find that the reliability diagram also encodes information relevant to MisD, in particular that  
 151 the **regions under and above the reliability curve relate to the precision of detecting correct and**  
 152 **incorrect predictions.** This insight reveals a direct link between the reliability diagram and MisD.

153 Specifically, we first analyze the meaning of the region under the reliability curve. Since the vertical  
 154 axis of reliability diagram denotes accuracy, each curve value reflects the proportion of correct  
 155 predictions among all samples at that confidence. The enclosed region naturally accumulates these  
 156 proportions across the confidence spectrum. Building on this view, we formalize the following  
 157 relationship. The formally proof is listed in Appendix B.1:

158 **Lemma 3.1.** *Given a confidence interval  $[a, b]$ , let  $w(s)$  denote the density of the sample at confidence  
 159 level  $s$ . Then, the precision of correct predictions within  $[a, b]$  can be driven from region under  
 160 the reliability curve  $f(s)$  over  $[a, b]$ , i.e.,  $\text{Prec}^+ = (\int_a^b w(s)f(s)ds) / \int_a^b w(s)ds$ . Similarly, the  
 161 precision of incorrect predictions over  $[a, b]$  can be derived from the region above the curve via  
 $\text{Prec}^- = (\int_a^b w(s)(1 - f(s))ds) / \int_a^b w(s)ds$ .*

162 Lemma 3.1 establishes a strong connection between the region under (above) the reliability curve  
 163 and the task of detecting correct (incorrect) predictions. Since MisD can be quantified by the  
 164 joint precision of detecting high-confidence predictions (*i.e.*,  $s \geq r$ ) as correct and low-confidence  
 165 predictions (*i.e.*,  $s < r$ ) as incorrect, it can be characterized by the region below the reliability curve  
 166 in the high-confidence interval and the region above it in the low-confidence interval.

167 Based on the above observations, we can examine the effectiveness of confidence calibration for  
 168 MisD. Specifically, we analyze the precision of the perfect calibration curve in detecting correct and  
 169 incorrect predictions, as formalized in the following theorem (the proof is listed in Appendix B.3):

170 **Theorem 3.2.** *Let  $r \in [0, 1]$  be a confidence threshold. Under perfect calibration, the precision for  
 171 the correct prediction detection and incorrect prediction detection tasks is  $\text{Prec}^+ = \mathbb{E}_{s \sim w(s|s \in [r, 1])}[s]$   
 172 and  $\text{Prec}^- = \mathbb{E}_{s \sim w(s|s \in [0, r])}[1 - s]$ , respectively.*

173 Theorem 3.2 demonstrates that, even under the goal of confidence calibration (*i.e.*, perfect calibration),  
 174 the precision of correct prediction detection equals the conditional expectation of the confidence  
 175  $s$  in  $[r, 1]$ . In practice, this expectation is strictly less than 1 **unless all samples are concentrated**  
 176 **at confidence 1**, which rarely occurs. Similarly, the precision of incorrect prediction detection  
 177 corresponds to the conditional expectation of  $1 - s$  in  $[0, r]$ , and is strictly less than 1 **unless all**  
 178 **samples are concentrated at confidence 0**. In realistic scenarios, such as with pretrained CLIP  
 179 models, test samples are often drawn from diverse distributions, causing confidence values to be  
 180 widely spread across  $[0, 1]$  (see Appendix E.8 for empirical validation). Consequently, both the  
 181 precision of correct and incorrect prediction detection are far below 1. This observation explains why  
 182 strict calibration alone is insufficient to achieve high MisD performance.

183 Motivated by Lemma 3.1 and Theorem 3.2, we seek a reliability curve as the calibration objective  
 184 that explicitly guides optimization toward better MisD. Let  $f : [0, 1] \rightarrow [0, 1]$  denote the reliability  
 185 curve. Based on the above discussion, we therefore formulate the following MisD-oriented objective:

$$187 \max_{f \in \mathcal{F}} \left( \int_0^{0.5} w(s)[1 - f(x)]dx + \int_{0.5}^1 w(s)f(x)dx \right), \quad (4)$$

188 where  $\mathcal{F} = \{f : [0, 1] \rightarrow [0, 1], f(0) = 0, f(1) = 1, f \text{ nondecreasing}\}$  is a function family. This  
 189 objective maximizes the region above the curve in the low-confidence region  $[0, 0.5]$ , and the region  
 190 under the curve in the high-confidence region  $[0.5, 1]$ , thereby improving the precision of correct and  
 191 incorrect prediction detection. However, directly maximizing Eq.(4) over  $\mathcal{F}$  causes overly aggressive  
 192 probability adjustments and instability. To balance improvement in MisD with smoothness, we  
 193 impose an additional requirement that  $f$  should have a controllable and gradual transition. As a  
 194 practical and analytically tractable instantiation, we adopt the normalized sigmoid curve:  
 195

$$196 \Psi(s) \triangleq \frac{\sigma\left(\frac{s-0.5}{\lambda}\right) - \sigma\left(\frac{-0.5}{\lambda}\right)}{\sigma\left(\frac{0.5}{\lambda}\right) - \sigma\left(\frac{-0.5}{\lambda}\right)}, \quad \sigma(z) \triangleq \frac{1}{1 + e^{-z}}, \quad (5)$$

197 where  $\lambda \in \mathbb{R}^+$  controls the smoothness of the transition, enabling a trade-off between separation  
 198 strength and stability. Figure 2 illustrates the curve for different values of  $\lambda$ , highlighting its flexibility.  
 199 We summarize the key properties of the proposed curve that are effective for MisD as follows.

200 **Property 1a: Perfect calibration as a special case ( $\lambda \rightarrow \infty$ ).**

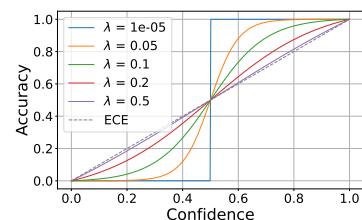
201 As  $\lambda \rightarrow \infty$ ,  $\Psi(x)$  converges uniformly to the diagonal  $\Psi(x) = x$  on  $[0, 1]$ , which corresponds to perfect calibration.

202 **Property 1b: Step-function limit that maximizes MisD effectiveness ( $\lambda \rightarrow 0$ ).** As  $\lambda \rightarrow 0$ ,  $\Psi(x)$  converges to the  
 203 closed-form solution of Eq.(4) (*i.e.*, step function), then the following corollary holds. The proof is listed in Appendix B.2.

204 **Corollary 3.3.** *When  $\lambda \rightarrow 0$ , the proposed normalized sigmoid curve assigns higher confidence to all correctly classified  
 205 samples than to any misclassified sample, *i.e.*,*

$$206 \mathbb{P}(\text{Conf}(N^+) > \text{Conf}(N^-)) = 1, \quad (6)$$

207 where  $N^+$  and  $N^-$  denote the sets of correctly classified and  
 208 misclassified samples, respectively. In this case, the reliability  
 209 curve approaches a step-like curve (cf. Figure 2 with  $\lambda = 1e-5$ ).



210 Figure 2: Visualization of the proposed normalized sigmoid curve with different  $\lambda$  values alongside  
 211 ECE in the reliability diagram.

216 **Properties 1a** and **1b** characterize the effect of the limit behavior as  $\lambda$  varies. A smaller  $\lambda$  leads  
 217 to stronger separation and hence better MisD performance, but pushing  $\lambda \rightarrow 0$  results in overly  
 218 aggressive probability updates, causing instability and potential over-fitting on the calibration set.  
 219 Therefore, in practice, an intermediate  $\lambda$  is chosen to balance separation strength and stability.

220 **Property 2: Precision dominance over perfect calibration.** The proposed reliability curve  
 221 guarantees strictly higher precision than perfect calibration for correctly classified samples in the high-  
 222 confidence region, and symmetrically outperforms perfect calibration for misclassification detection  
 223 in the low-confidence region. Consequently, given the distribution of samples over confidence values  
 224 (i.e.,  $w(s)$ ), we arrive at the following theorem (The proof is listed in the Appendix B.4):  
 225

226 **Theorem 3.4.** *Let  $r \in (0.5, 1)$  and let  $w(s) \geq 0$  be any weight function on  $[0, 1]$  with  $\int_r^1 w(s) ds > 0$ .  
 227 The precision for correct-prediction detection (i.e.,  $\text{Prec}^+$ ) satisfies the following inequality:*

$$\text{Prec}_{\Psi}^+(r; w) := \frac{\int_r^1 w(s) \Psi(s) ds}{\int_r^1 w(s) ds} \geq \frac{\int_r^1 w(s) s ds}{\int_r^1 w(s) ds} =: \text{Prec}_{\text{diag}}^+(r; w).$$

231 Moreover, for the incorrect prediction detection with  $r \in (0, 0.5)$ , the precision (i.e.,  $\text{Prec}^-$ ) satisfies  
 232 the following inequality:

$$\text{Prec}_{\Psi}^-(r; w) := \frac{\int_0^r w(s) (1 - \Psi(s)) ds}{\int_0^r w(s) ds} \geq \frac{\int_0^r w(s) (1 - s) ds}{\int_0^r w(s) ds} =: \text{Prec}_{\text{diag}}^-(r; w).$$

237 **Property 3: Less tolerance for prediction mixing.** MisD emphasizes separating correctly and  
 238 incorrectly classified samples; hence, less *mixing* between these two groups at the confidence level is  
 239 critical. Entropy is a natural metric that can be used to measure the degree of mixing. Therefore, we  
 240 can characterize the tolerance of the reliability curve for prediction mixing at every confidence level  
 241 by the entropy. We then have the following theorem (The proof is listed in Appendix B.6):

242 **Theorem 3.5.** *Given arbitrary confidence  $s$ , then the corresponding entropy for the proposed  
 243 reliability curve (i.e.,  $\Psi(s)$ ) and perfect-calibration curve satisfy the following inequality.*

$$\mathbf{M}_{\Psi}(s) \leq \mathbf{M}_{\text{diag}}(s), \quad \mathbf{M}(s) := -\mathbb{P}(N_s^+) \log(\mathbb{P}(N_s^+)) - \mathbb{P}(N_s^-) \log(\mathbb{P}(N_s^-)), \quad (7)$$

244 where  $\mathbb{P}(N_s^+)$  and  $\mathbb{P}(N_s^-)$  denote the probabilities of being correctly and incorrectly classified given  
 245 confidence  $s$ , respectively; i.e.,  $\mathbb{P}(N_s^+) = \mathbb{P}(\text{correct} | \text{conf} = s)$  and  $\mathbb{P}(N_s^-) = 1 - \mathbb{P}(N_s^+)$ .

246 This theorem indicates that the proposed reliability curve (i.e.,  $\Psi(s)$ ) has less tolerance for mixing of  
 247 correctly and incorrectly classified samples.

248 Taken together, **Properties 1–3** provide a coherent theoretical picture: the proposed curve continuously  
 249 interpolates between perfect calibration and ideal separation (**Property 1**), provably achieves  
 250 higher precision for both correct and incorrect prediction detection (**Property 2**), and admits lower  
 251 entropy at each confidence level, indicating less tolerance for prediction mixing (**Property 3**). These  
 252 results highlight that  $\Psi(s)$  aligns better with the goal of improving MisD. Building on these insights,  
 253 we next introduce a practical post-hoc method that operationalizes the proposed reliability curve.

### 257 3.2 SURROGATE LOSS

258 Despite the advantages of the proposed reliability curve (i.e.,  $\Psi(s)$ ), directly realizing it as a training  
 259 objective is challenging. In particular, the reliability curve cannot be straightforwardly expressed as  
 260 a differentiable loss, and an alternative would be to approximate it through binning on a held-out  
 261 calibration set, similar to ECE (Guo et al., 2017). However, calibration sets are usually small in  
 262 practice, which causes large variance in the empirical bin estimates and makes such optimization  
 263 unreliable. To tackle this issue, we propose a simple alternative to achieve it.

264 Specifically, we approximate the desired behavior of the reliability curve through a differentiable  
 265 surrogate penalty. Recall our proposed curve  $\Psi(s)$ , which smoothly maps the confidence score  
 266  $s \in [0, 1]$  to an “expected accuracy” curve that imposes stronger constraints at both ends of the  
 267 confidence spectrum. Instead of directly optimizing this non-differentiable target, we design penalties  
 268 that follow the trend of expected accuracy (i.e.,  $\Psi(s)$ ): *For incorrectly predicted samples, the penalty  
 269 strength follows expected accuracy  $\Psi(s)$* , meaning that incorrect predictions in regions of high

expected accuracy incur a larger penalty, *while correctly predicted samples follows*  $1 - \Psi(s)$ , so that correct predictions in regions of low expected accuracy receive a larger penalty. In short, the penalties are assigned in a way that directly mirrors the expected accuracy encoded by  $\Psi(s)$ , providing a differentiable surrogate for the proposed reliability curve. However,  $\Psi(s)$  only constrains the confidence, *i.e.*, the highest probability, while ignoring the distribution of the remaining probabilities. This may lead to optimization difficulties and suboptimal performance. Therefore, we further incorporate a constraint on the full probability distribution to make the constraint more comprehensive:

$$\begin{aligned} \mathcal{L}_{\text{SUR}} = & \min \beta \frac{1}{|N^+|} \sum_{i \in N^+} \phi\{1 - \Psi(s), -\mathbf{y}_i^\top \log f(\mathbf{x}_i)\} \\ & + (1 - \beta) \frac{1}{|N^-|} \sum_{j \in N^-} \phi\{\Psi(s), -[1/c, \dots, 1/c]^\top \log f(\mathbf{x}_j)\}, \end{aligned} \quad (8)$$

where  $\beta$  is a hyperparameter balancing the constraints on correctly and incorrectly predicted samples, and  $\phi\{\cdot, \cdot\}$  is a fusion function, which can be chosen as summation or multiplication. The first term of  $\phi\{\cdot, \cdot\}$  encourages the reliability curve to align with the  $\Psi(s)$ . The second term of  $\phi\{\cdot, \cdot\}$  provides a full probability constraint to avoid the limitation of only constraining the predicted class probability from the first term. Specifically, for correctly predicted samples, it is a standard Cross Entropy, which enforces a low-entropy probability distribution, while for incorrectly predicted samples, it imposes a constraint by comparing the predicted probability vector with the uniform distribution  $[1/c, \dots, 1/c]$ , thereby encouraging a high-entropy probability distribution. In summary, the first term drives the main alignment with the proposed reliability curve (*i.e.*,  $\Psi(s)$ ), and the second term complements it with full-probability regularization, jointly yielding a more stable and effective objective.

### 3.3 LIGHTWEIGHT META NETWORK

To avoid altering the model predictions and introducing significant overhead, we propose a *post-hoc* method. It employs a lightweight meta network (*i.e.*, LMN) to learn the temperature coefficients  $\tau$  for refining the VLM’s confidence. For instance, given an image  $v$ , its logits can be refined as:

$$\mathbf{z}'_v = [\tau_v \cdot z_{v,1}, \tau_v \cdot z_{v,2}, \dots, \tau_v \cdot z_{v,c}], \quad z_{v,i} = \tau_{\text{clip}} \cdot \text{sim}(\xi(\mathbf{x}_v), \psi(\mathbf{t}_i)). \quad (9)$$

To obtain the instance-specific temperature coefficient  $\tau_v$ , we design a lightweight network that takes as input the information from both modalities, without modifying the parameters of the pretrained or fine-tuned model. Concretely, the model’s output logits  $\mathbf{z}_v$ , the image embedding  $\xi(\mathbf{x}_v)$ , and the predicted text embeddings  $\psi(\mathbf{t}_p)$  are each passed through a separate fully connected (FC) layer to map them into a common latent space. These representations are then concatenated and fed into another FC layer, which projects them to a scalar value corresponding to  $\tau_v$ . Formally, we have:

$$\tau_v = \sigma^+(\text{FC}_\tau(\mathbf{h}_z || \mathbf{h}_x || \mathbf{h}_t)), \quad \mathbf{h}_z = \text{FC}_z(\mathbf{z}_v), \quad \mathbf{h}_x = \text{FC}_x(\xi(\mathbf{x}_v)), \quad \mathbf{h}_t = \text{FC}_t(\psi(\mathbf{t}_p)), \quad (10)$$

where  $[\cdot || \cdot]$  denotes concatenation and  $\sigma^+(x) = \log(1 + \exp(x))$  is an element-wise softplus activation (Dugas et al., 2000). Finally, by substituting Eq.(9) into the objective function in Eq.(8), we update only the parameters of these FC layers on the calibration set, making the overall training procedure lightweight and efficient. The flowchart of the proposed meta network can be found in Figure 5.

## 4 EXPERIMENT

In this section, we conduct experiments on six public datasets to evaluate the proposed method in terms of different settings. Details of experiments are shown in Appendix D and additional experiments are shown in Appendix E. The code is released at [Anonymous Code Link](#).

### 4.1 EXPERIMENTAL SETUP

**Datasets.** We conduct the analysis on six datasets covering specialized and fine-grained domains, which include DTD (Cimpoi et al., 2014), Flowers102 (Nilsback & Zisserman, 2008), EuroSAT (Helber et al., 2019), RESICS45 (Cheng et al., 2017), MNIST (Deng, 2012), and CUB (Wah et al., 2011). For each dataset, we adopt the official training and test splits provided in Li et al. (2022), and construct a few-shot calibration set by randomly sampling a small subset from the training set.

Table 1: Comparison results of recent confidence calibration OOD detection methods in the few-shot setting. Note that  $\uparrow$  indicates higher is better,  $\downarrow$  indicates lower is better.

Methods	DTD					Flowers102				
	AUROC↑	AUPR-S↑	AUPR-E↑	FPR90-S↓	FPR90-E↓	AUROC↑	AUPR-S↑	AUPR-E↑	FPR90-S↓	FPR90-E↓
Zero-shot CLIP	0.762	0.740	0.771	0.669	0.572	0.864	0.922	0.759	0.435	0.354
FeatureClipping	0.749	0.716	0.764	0.687	0.571	0.873	0.935	0.695	0.416	0.321
DOR	0.768	0.726	0.798	0.656	0.550	0.828	0.907	0.682	0.570	0.393
SCT	0.759	0.741	0.776	0.685	0.557	0.868	0.926	0.761	0.429	0.337
ViLU	<b>0.769</b>	<b>0.759</b>	<b>0.762</b>	<b>0.678</b>	<b>0.521</b>	<b>0.875</b>	<b>0.913</b>	<b>0.772</b>	<b>0.401</b>	<b>0.329</b>
LMN (Ours)	<b>0.802</b>	<b>0.800</b>	<b>0.804</b>	<b>0.636</b>	<b>0.457</b>	<b>0.886</b>	<b>0.937</b>	<b>0.799</b>	<b>0.378</b>	<b>0.305</b>
Methods	EuroSAT					RESICS45				
	AUROC↑	AUPR-S↑	AUPR-E↑	FPR90-S↓	FPR90-E↓	AUROC↑	AUPR-S↑	AUPR-E↑	FPR90-S↓	FPR90-E↓
Zero-shot CLIP	0.65	0.501	0.771	0.782	0.742	0.779	0.824	0.711	0.636	0.508
FeatureClipping	0.685	0.613	0.727	0.729	0.536	0.781	0.827	0.705	0.638	0.501
DOR	0.719	0.649	0.734	0.716	0.529	0.778	0.828	0.688	0.621	0.516
SCT	0.681	0.534	0.791	0.754	0.682	0.784	0.826	0.716	0.633	0.501
ViLU	<b>0.723</b>	<b>0.618</b>	<b>0.787</b>	<b>0.723</b>	<b>0.538</b>	<b>0.787</b>	<b>0.829</b>	<b>0.730</b>	<b>0.618</b>	<b>0.493</b>
LMN (Ours)	<b>0.788</b>	<b>0.698</b>	<b>0.855</b>	<b>0.655</b>	<b>0.468</b>	<b>0.808</b>	<b>0.845</b>	<b>0.741</b>	<b>0.597</b>	<b>0.445</b>
Methods	MNIST					CUB				
	AUROC↑	AUPR-S↑	AUPR-E↑	FPR90-S↓	FPR90-E↓	AUROC↑	AUPR-S↑	AUPR-E↑	FPR90-S↓	FPR90-E↓
Zero-shot CLIP	0.813	0.565	0.919	0.511	0.482	0.807	0.839	0.758	0.767	0.554
FeatureClipping	0.816	0.654	0.843	0.501	0.461	0.805	0.834	0.756	0.768	0.541
DOR	0.824	0.645	0.868	0.497	0.453	0.808	0.838	0.761	0.765	0.551
SCT	0.837	0.664	0.933	0.486	0.423	0.808	0.839	0.759	0.766	0.547
ViLU	<b>0.877</b>	<b>0.769</b>	<b>0.954</b>	<b>0.350</b>	<b>0.263</b>	<b>0.801</b>	<b>0.827</b>	<b>0.753</b>	<b>0.769</b>	<b>0.563</b>
LMN (Ours)	<b>0.915</b>	<b>0.779</b>	<b>0.965</b>	<b>0.200</b>	<b>0.205</b>	<b>0.812</b>	<b>0.846</b>	<b>0.764</b>	<b>0.756</b>	<b>0.532</b>

**Baselines.** We evaluate our method using pretrained CLIP as the base model and compare it with two recent calibration approaches (FeatureClipping (Tao et al., 2025) and DOR (Wang et al., 2025)), a strong VLM-based OOD detector (SCT (Yu et al., 2024)), **and the latest MisD-oriented method for VLMs (ViLU (Lafon et al., 2025))**. We also acknowledge FSMisD (Zeng et al., 2025), a recent MisD-oriented method for VLMs, but its code is not publicly available and cannot be included. In addition, we evaluate two commonly used prompt-tuning CLIP variants: CoOP (Zhou et al., 2022b) (textual) and VPT (Jia et al., 2022) (visual).

**Implementation details.** We use CLIP ViT-B/32 as the visual backbone. For both visual and textual prompt learning, we set the prefix size to 16 (Zhou et al., 2022b; Jia et al., 2022). We use SGD as the optimizer, with the number of training epochs selected from {100, 150, 200} and the learning rate selected from {0.001, 0.002, 0.005}. We adopt a 16-shot setting for the calibration set. When the base model is prompt-tuning CLIP, we split the original 16-shot calibration set into two parts: one is used for learning the prompts, and the other is reserved for training our post-hoc calibration network.

**Evaluation metrics.** To evaluate the MisD, we measure the ranking capability of confidence scores, *i.e.*, the ability to rank correctly classified samples ahead of misclassified ones. Following existing works (Corbière et al., 2019; Hendrycks et al., 2018), we adopt several widely used metrics: AUROC, AUPR-Success, AUPR-Error, FPR90%-Success-TPR, and FPR90%-Error-TPR. Since our MisD evaluation simultaneously considers both correct and incorrect predictions, the positive class in AUPR and FPR is redefined according to the detection target: for *Success*, correctly predicted samples are treated as positives, whereas for *Error*, incorrectly predicted samples are treated as positives.

## 4.2 MAIN RESULTS

We first evaluate the MisD performance of recent confidence calibration methods, the OOD detection method, and our proposed method applied to the pretrained CLIP. The results on five commonly used metrics (*i.e.*, AUROC, AUPR-Success, AUPR-Error, FPR90-Success, and FPR90-Error) are reported in Table 1, from which we make the following observations:

First, the confidence calibration methods can not help CLIP to distinguish between the correctly predicted samples and the incorrectly predicted samples. For example, on average across all datasets, FeatureClipping improves AUROC by only about 0.7% over the pretrained CLIP, while DOR achieves an average gain of merely 1.1%. Moreover, such marginal improvements are further undermined by their inconsistency, as both methods fail to enhance CLIP uniformly across all evaluation metrics. This aligns with Theorem 3.2, confirming that confidence calibration imposes a fundamental limit on achievable MisD performance.

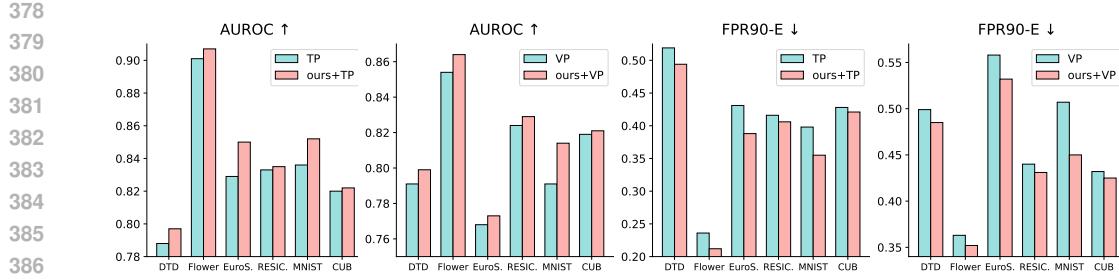


Figure 3: Bar plot comparison of textual-prompt (*i.e.*, TP in Figure) and visual-prompt (*i.e.*, VP in Figure) CLIP before and after applying our proposed method, evaluated in terms of AUROC and FPR90-Error across DTD, Flowers102, EuroSAT, RESICS45, MNIST, and CUB datasets.

Second, compared to the VLM-based calibration and OOD detection methods, *the proposed method always outperforms them by large margins across all evaluation metrics*. For example, the proposed method on average improves by 6.1%, 10.5%, 4.7%, 13.4%, and 22.9%, compared to the pretrained CLIP, in terms of AUROC, AUPR-Success, AUPR-Error, FPR90-Success, and FPR90-Error. Furthermore, *even when compared with the strongest baseline ViLU, our method still achieves substantial gains of 2.8%, 5.1%, 2.1%, 5.4%, and 18.7% on the same metrics*. These results demonstrate the superiority of the proposed method.

### 4.3 EVALUATION ON VARIOUS BASE MODELS

To evaluate the effectiveness of the proposed method on more base models, we adopt two types of prompt-based CLIP models (*i.e.*, textual prompt and visual prompt). Since prompt-based methods require a training set to learn the prompt embeddings, we leverage the few-shot calibration set, which is originally divided into training and validation portions. Specifically, following (Wang et al., 2021), the training portion is used to learn the prompt embeddings, while the validation portion serves both to validate the learned prompts and to train the proposed post-hoc calibration model.

The results are shown in Figure 3. We have the following observations: our method consistently improves the performance of both types of prompt-based CLIP across multiple datasets. For example, on average, our method brings an improvement of 1.2% in AUROC and a reduction of 6.2% in FPR90-Error for the textual prompt, and 1.1% and 4.4%, respectively, for the visual prompt. These results demonstrate the applicability of our model to different forms of base models.

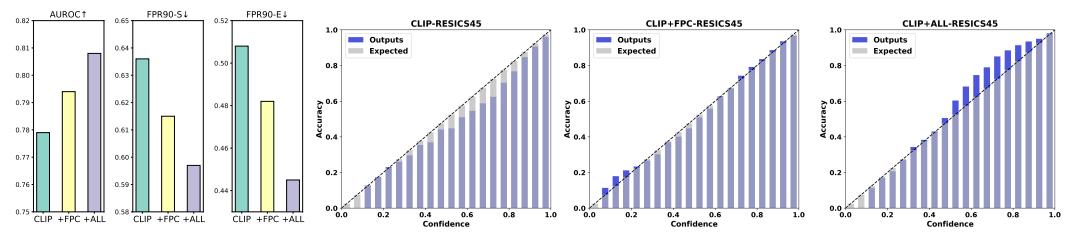
### 4.4 EVALUATION ON OPEN-VOCABULARY SETTING

A key advantage of pretrained CLIP lies in its strong zero-shot capability, enabling it to generalize well to unseen classes. However, fine-tuning operations, including prompt learning, may compromise this ability by overfitting to the training classes. To examine whether our proposed post-hoc method preserves the zero-shot capability, we design an open-vocabulary evaluation *and compare LMN with both zero-shot CLIP and the open-vocabulary calibration method DOR* (Wang et al., 2025) (*a CoOP-based calibration method*). Specifically, we randomly sample a subset of classes for calibration, while the remaining unseen classes are reserved for testing. The results are shown in Table 2.

The experimental results demonstrate that our method not only avoids degrading CLIP’s performance under the open-vocabulary setting, but also leads to consistent improvements across almost all datasets and evaluation metrics. For example, *on the DTD dataset, our method achieves a 7.1% and 5.9% relative improvement over CLIP and DOR in terms of the FPR90-S metric*. This indicates that our approach is capable of enhancing model performance while preserving CLIP’s inherent strengths. We attribute this to the post-hoc nature of our method, which does not modify the pretrained parameters of CLIP and thus maintains the intrinsic generalization ability of its image and text embeddings. By retaining this property, the proposed calibration method can leverage the strong representational power of CLIP while improving its reliability on unseen classes.

432  
 433 Table 2: Results of open-vocabulary evaluation on six datasets, measured by AUROC, FPR90-S, and  
 434 FPR90-E.  $\uparrow$  denotes that higher values are better, while  $\downarrow$  denotes that lower values are better.

435	436	DTD			Flowers102			EuroSAT		
		AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$	AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$	AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$
437	Zero-shot CLIP	0.760	0.642	0.604	0.853	0.407	0.397	0.608	0.747	0.882
438	<b>DOR</b>	<b>0.758</b>	<b>0.631</b>	<b>0.612</b>	<b>0.847</b>	<b>0.410</b>	<b>0.394</b>	<b>0.605</b>	<b>0.751</b>	<b>0.883</b>
439	LMN (Ours)	<b>0.770</b>	<b>0.596</b>	<b>0.598</b>	<b>0.858</b>	<b>0.396</b>	<b>0.387</b>	<b>0.614</b>	<b>0.743</b>	<b>0.875</b>
440	441	RESICS45			MNIST			CUB		
		AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$	AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$	AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$
442	Zero-shot CLIP	0.780	0.601	0.536	0.854	0.545	0.322	0.789	<b>0.641</b>	0.490
443	<b>DOR</b>	<b>0.764</b>	<b>0.621</b>	<b>0.543</b>	<b>0.842</b>	<b>0.573</b>	<b>0.331</b>	<b>0.779</b>	<b>0.667</b>	<b>0.512</b>
444	LMN (Ours)	<b>0.782</b>	<b>0.599</b>	<b>0.527</b>	<b>0.859</b>	<b>0.539</b>	<b>0.301</b>	<b>0.790</b>	0.653	<b>0.472</b>



452 Figure 4: Bar plot comparison of AUROC, FPR90-S, and FPR90-E metrics of CLIP, CLIP with full  
 453 probabilities constraint, and CLIP with our full method (i.e., first three figures) on the RESICS45  
 454 dataset. The corresponding reliability diagrams are shown in the last three figures.

## 455 4.5 ABLATION STUDIES

456 The key component of the proposed method is the surrogate loss  $\mathcal{L}_{\text{SUR}}$ , which consists of two parts:  
 457 (i) a confidence regularization term that encourages the predictions to align with the target reliability  
 458 curve (i.e.,  $\Phi(s)$ ), and (ii) a full-probability constraint (FPC) that complements the confidence  
 459 regularization by regularizing the entire probability distribution, preventing the training instability  
 460 cause by single probability (i.e., confidence) were optimized. Therefore, we do not report a separate  
 461 ablation with only the confidence regularization term, as it is intended to be used together with  
 462 FPC. To verify the effectiveness of these components, we visualize the reliability diagrams on the  
 463 RESICS45 dataset under three settings: the original CLIP, CLIP with FPC, and CLIP with the full  
 464 proposed method. The visualizations together with the corresponding quantitative results are reported  
 465 in Figure 4. Additional ablation studies on other datasets are provided in Appendix E.1.

466 From Figure 4, we draw the following observations. First, the complete objective achieves the  
 467 best overall performance, and using FPC alone also improves MisD performance compared to  
 468 the pretrained CLIP, indicating that both components of the surrogate loss contribute effectively.  
 469 Second, from the reliability diagrams, we see that the pretrained CLIP model exhibits overconfidence.  
 470 Incorporating FPC alleviates this issue to some extent; however, the resulting curve still deviates from  
 471 the target  $\Phi(s)$ . Finally, with the proposed confidence regularization term, the reliability curve aligns  
 472 much more closely with  $\Phi(s)$ , demonstrating that our surrogate loss indeed possesses the desired  
 473 ability to calibrate predictions toward the target normalized sigmoid reliability curve.

## 474 5 CONCLUSION

475 In this work, we revisited the reliability diagram of confidence calibration and established its connec-  
 476 tion with MisD. We showed that the standard calibration objective inherently limits MisD performance,  
 477 and proposed a new reliability curve as the calibration objective. This reliability curve theoretically  
 478 improves MisD performance and reduces the mixing of correct and incorrect predictions. As the  
 479 value of the reliability curve is a statistical result and thus non-differentiable, we address this issue by  
 480 designing a differentiable surrogate loss. Furthermore, to preserve the predictive power of VLMs, we  
 481 developed a lightweight post-hoc framework that employs a meta network to produce sample-specific  
 482 temperature factors. Both theoretical analysis and extensive experiments confirmed that our approach  
 483 consistently enhances MisD performance while maintaining model accuracy.

486 REPRODUCIBILITY STATEMENT  
487

488 We have made extensive efforts to ensure the reproducibility of our work. The model architecture,  
489 training procedure, and hyperparameter settings are described in Appendix D.2. Complete proofs of  
490 the theoretical results are presented in Appendix B. The datasets used in our experiments are publicly  
491 available, and the preprocessing steps are explained in the Appendix D E. Moreover, we provide an  
492 anonymous link to the source code in the supplementary materials to facilitate reproduction of our  
493 experiments ([Anonymous Code Link](#)).

494  
495 REFERENCES  
496

497 Runnan Chen, Youquan Liu, Lingdong Kong, Xinge Zhu, Yuexin Ma, Yikang Li, Yuenan Hou,  
498 Yu Qiao, and Wenping Wang. Clip2scene: Towards label-efficient 3d scene understanding by clip.  
499 In *CVPR*, pp. 7020–7030, 2023.

500 Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark  
501 and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017.

502 Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. De-  
503 scribing textures in the wild. In *CVPR*, pp. 3606–3613, 2014.

504 Charles Corbière, Nicolas Thome, Avner Bar-Hen, Matthieu Cord, and Patrick Pérez. Addressing  
505 failure prediction by learning model confidence. In *NeurIPS*, volume 32, 2019.

506 Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, Yang Zhou, Kaizhao Liang, Jintai Chen, Juanwu Lu,  
507 Zichong Yang, Kuei-Da Liao, et al. A survey on multimodal large language models for autonomous  
508 driving. In *CVPR*, pp. 958–979, 2024.

509 Li Deng. The mnist database of handwritten digit images for machine learning research [best of the  
510 web]. *IEEE signal processing magazine*, 29(6):141–142, 2012.

511 Charles Dugas, Yoshua Bengio, François Bélisle, Claude Nadeau, and René Garcia. Incorporating  
512 second-order functional knowledge for better option pricing. In *NeurIPS*, volume 13, 2000.

513 Arindam Ghosh, Thomas Schaaf, and Matthew Gormley. Ada focal: Calibration-aware adaptive focal  
514 loss. In *NeurIPS*, volume 35, pp. 1583–1595, 2022.

515 Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural  
516 networks. In *ICML*, pp. 1321–1330. PMLR, 2017.

517 Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset  
518 and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected  
519 Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019.

520 Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution  
521 examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.

522 Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier  
523 exposure. *arXiv preprint arXiv:1812.04606*, 2018.

524 Jincheng Huang, Jie Xu, Xiaoshuang Shi, Ping Hu, Lei Feng, and Xiaofeng Zhu. The final layer  
525 holds the key: A unified and efficient gnn calibration framework. *arXiv preprint arXiv:2505.11335*,  
526 2025.

527 Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and  
528 Ser-Nam Lim. Visual prompt tuning. In *ECCV*, pp. 709–727. Springer, 2022.

529 Ulf Johansson, Cecilia Sonstrod, Tuwe Lofstrom, and Henrik Bostrom. Confidence classifiers with  
530 guaranteed accuracy or precision. In *Conformal and Probabilistic Prediction with Applications*,  
531 pp. 513–533. PMLR, 2023.

532 Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz  
533 Khan. Maple: Multi-modal prompt learning. In *CVPR*, pp. 19113–19122, 2023.

540 Ranganath Krishnan and Omesh Tickoo. Improving model calibration with accuracy versus uncer-  
 541 tainty optimization. In *NeurIPS*, volume 33, pp. 18237–18248, 2020.

542

543 Marc Lafon, Yannis Karmim, Julio Silva-Rodríguez, Paul Couairon, Clément Rambour, Raphaël  
 544 Fournier-Sniedotta, Ismail Ben Ayed, Jose Dolz, and Nicolas Thome. Vilu: Learning vision-  
 545 language uncertainties for failure prediction. *arXiv preprint arXiv:2507.07620*, 2025.

546

547 Chunyuan Li, Haotian Liu, Liunian Li, Pengchuan Zhang, Jyoti Aneja, Jianwei Yang, Ping Jin,  
 548 Houdong Hu, Zicheng Liu, Yong Jae Lee, et al. Elevater: A benchmark and toolkit for evaluating  
 549 language-augmented visual models. In *NeurIPS*, volume 35, pp. 9287–9301, 2022.

550

551 Song-Lin Lv, Yu-Yang Chen, Zhi Zhou, Yu-Feng Li, and Lan-Zhe Guo. Contrast-aware calibration  
 552 for fine-tuned clip: Leveraging image-text alignment. *arXiv preprint arXiv:2501.19060*, 2025.

553

554 Jishnu Mukhoti, Viveka Kulharia, Amartya Sanyal, Stuart Golodetz, Philip Torr, and Puneet Dokania.  
 555 Calibrating deep neural networks using focal loss. In *NeurIPS*, volume 33, pp. 15288–15299,  
 556 2020.

557

558 Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? In  
 559 *NeurIPS*, volume 32, 2019.

560

561 Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number  
 562 of classes. In *ICCV*, pp. 722–729, 2008.

563

564 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
 565 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
 566 models from natural language supervision. In *ICML*, pp. 8748–8763. PMLR, 2021.

567

568 Linwei Tao, Minjing Dong, and Chang Xu. Feature clipping for uncertainty calibration. In *AAAI*,  
 569 volume 39, pp. 20841–20849, 2025.

570

571 Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd  
 572 birds-200-2011 dataset. 2011.

573

574 Shuoyuan Wang, Jindong Wang, Guoqing Wang, Bob Zhang, Kaiyang Zhou, and Hongxin Wei.  
 575 Open-vocabulary calibration for fine-tuned clip. In *ICML*, 2024.

576

577 Shuoyuan Wang, Yixuan Li, and Hongxin Wei. Understanding and mitigating miscalibration in  
 578 prompt tuning for vision-language models. In *ICML*, 2025.

579

580 Xiao Wang, Hongrui Liu, Chuan Shi, and Cheng Yang. Be confident! towards trustworthy graph  
 581 neural networks via confidence calibration. In *NeurIPS*, volume 34, pp. 23768–23779, 2021.

582

583 Miao Xiong, Ailin Deng, Pang Wei W Koh, Jiaying Wu, Shen Li, Jianqing Xu, and Bryan Hooi.  
 584 Proximity-informed calibration for deep neural networks. In *NeurIPS*, volume 36, pp. 68511–  
 585 68538, 2023.

586

587 Jia-Qi Yang, De-Chuan Zhan, and Le Gan. Beyond probability partitions: Calibrating neural networks  
 588 with semantic aware grouping. In *NeurIPS*, volume 36, pp. 58448–58460, 2023.

589

590 Geng Yu, Jianing Zhu, Jiangchao Yao, and Bo Han. Self-calibrated tuning of vision-language models  
 591 for out-of-distribution detection. In *NeurIPS*, volume 37, pp. 56322–56348, 2024.

592

593 Fanhu Zeng, Zhen Cheng, Fei Zhu, and Xu-Yao Zhang. Towards efficient and general-purpose  
 594 few-shot misclassification detection for vision-language models. *arXiv preprint arXiv:2503.20492*,  
 595 2025.

596

597 Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language  
 598 image pre-training. In *CVPR*, pp. 11975–11986, 2023.

599

Xu-Yao Zhang, Guo-Sen Xie, Xiuli Li, Tao Mei, and Cheng-Lin Liu. A survey on learning to reject.  
 600 *Proceedings of the IEEE*, 111(2):185–215, 2023.

594 Zihao Zhao, Yuxiao Liu, Han Wu, Mei Wang, Yonghao Li, Sheng Wang, Lin Teng, Disheng Liu,  
595 Zhiming Cui, Qian Wang, et al. Clip in medical imaging: A comprehensive survey. *arXiv preprint*  
596 *arXiv:2312.07353*, 2023.

597 Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for  
598 vision-language models. In *CVPR*, pp. 16816–16825, 2022a.

600 Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-  
601 language models. *ICCV*, 130(9):2337–2348, 2022b.

602 Fei Zhu, Zhen Cheng, Xu-Yao Zhang, and Cheng-Lin Liu. Rethinking confidence calibration for  
603 failure prediction. In *ECCV*, pp. 518–536. Springer, 2022.

605 Fei Zhu, Zhen Cheng, Xu-Yao Zhang, and Cheng-Lin Liu. Openmix: Exploring outlier samples for  
606 misclassification detection. In *CVPR*, pp. 12074–12083, 2023.

607

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 A RELATED WORK  
649  
650651 This section briefly reviews the topics related to this work, including confidence calibration and  
652 misclassification detection, and the recent works in VLMs.  
653654  
655 A.1 CONFIDENCE CALIBRATION  
656657 Confidence calibration is a crucial technique for trustworthy machine learning, whose goal is to  
658 calibrate the confidence, making it accurately reflect the actual correctness. The calibration methods  
659 can be divided into two categories: train-time calibration and post-hoc calibration. For the train-time  
660 calibration, a notable example is focal loss (Mukhoti et al., 2020), with subsequent works such as  
661 adaptive focal loss (Ghosh et al., 2022) modifying hyperparameters for different sample groups based  
662 on prior training knowledge. Similarly, label smoothing (Müller et al., 2019). For the post-hoc  
663 methods, the most commonly used method is temperature scaling (TS) (Guo et al., 2017), which  
664 proposes to use a hyperparameter (*i.e.*, temperature coefficient) as the denominator of logits, making  
665 the confidence adjustable. However, TS is not flexible enough, since TS uses a unified temperature  
666 coefficient for every sample. Therefore, Many subsequent methods (Wang et al., 2021; Xiong et al.,  
667 2023; Yang et al., 2023) aim to improve TS by applying adaptive temperature parameters, treating  
668 samples differently for a more effective maximum-entropy regularizer.  
669670 In VLMs, calibration becomes more challenging due to modality-specific factors or the process of  
671 fine-tuning (Wang et al., 2024; 2025). Consequently, DAC (Wang et al., 2024) utilizes the proximity of  
672 the text modality to adjust confidence accordingly. DOR (Wang et al., 2025) mitigating miscalibration  
673 introduced during fine-tuning by introducing an extra-large vocabulary set. Although many calibration  
674 methods are proposed, the fundamental goal of calibration is to enable a model to distinguish between  
675 correct predictions and incorrect predictions. However, in this work, we reveal that even a perfectly  
676 calibrated model remains fundamentally limited, and thus may still fail to identify the misclassified  
677 samples or correctly predicted samples that bring the potential risk in many scenarios.  
678  
679680 A.2 MISCLASSIFICATION DETECTION  
681682 Misclassification detection (Hendrycks & Gimpel, 2016), which is also called failure prediction, and  
683 the goal is to detect incorrect predictions from correct predictions. It is crucial for machine learning  
684 models deployed in high-risk scenarios. Calibration methods (Zhang et al., 2023) are often viewed as  
685 an effective strategy for misclassification detection, as aligning predicted confidence with the true  
686 likelihood of correctness allows the model to distinguish between reliable and unreliable predictions.  
687 However, empirical studies (Zhu et al., 2022) have shown that commonly used calibration methods  
688 only bring limited improvement for misclassification detection. Yet, no theoretical analysis has been  
689 conducted to uncover the underlying reasons, and consequently, the proposed method of this work  
690 has not been optimized from a calibration perspective. Therefore, to the best of our knowledge, no  
691 prior work has systematically investigated misclassification detection from the calibration perspective.  
692 Instead, existing approaches typically focus on confidence regression (Corbière et al., 2019), exposing  
693 outlier samples (Zhu et al., 2023), flat minima (Zhu et al., 2022), among other techniques.  
694695 In VLMs, a few recent efforts have been devoted to misclassification detection. For instance,  
696 FSMisD (Zeng et al., 2025) adopts a prompt-based strategy, which, however, overlooks the use of  
697 confidence information, limiting flexibility and making it difficult to integrate with existing prompt-  
698 tuning models. ViLU (Lafon et al., 2025), on the other hand, formulates uncertainty modeling as  
699 a binary classification problem, which essentially ignores the role of confidence in distinguishing  
700 low-confidence correct predictions from high-confidence misclassifications, and lacks theoretical  
701 guarantees. While these approaches provide useful heuristics, they mainly rely on task-specific  
702 designs or empirical observations, resulting in limited ability to capture the fundamental role of  
703 confidence in misclassification detection.

702 **B THEORETICAL PROOF**  
 703

704 **B.1 PROOF FOR LEMMA 3.1**  
 705

706 **Lemma B.1.** *Given a confidence interval  $[a, b]$ , let  $w(s)$  denote the density of the sample at confidence  
 707 level  $s$ . Then, the precision of correct predictions within  $[a, b]$  can be derived from the region under  
 708 the reliability curve  $f(s)$  over  $[a, b]$ , i.e.,  $\text{Prec}^+ = (\int_a^b w(s)f(s)ds) / \int_a^b w(s)ds$ . Similarly, the  
 709 precision of incorrect predictions over  $[a, b]$  can be derived from the region above the curve via  
 710  $\text{Prec}^- = (\int_a^b w(s)(1 - f(s))ds) / \int_a^b w(s)ds$ .*  
 711

712 *Proof.* The formula of precision is defined as:  
 713

714 
$$\text{Precision}_{[a,b]} = \frac{\text{correct predictions with } s \in [a, b]}{\text{all predictions with } s \in [a, b]}. \quad (11)$$
  
 715

716 The reliability curve can be represented as  $f(s) = \mathbb{P}(\text{correct} | \text{confidence} = s)$ , thus the number of  
 717 correct predictions in the interval  $[a, b]$  can be represented as:  
 718

719 
$$\text{correct predictions} = \int_a^b w(s)\text{Accuracy}(s)ds = \int_a^b w(s)\mathbb{P}(\text{correct} | \text{confidence} = s)ds. \quad (12)$$
  
 720

721 The number of the whole sample is:  
 722

723 
$$\text{all predictions} = \int_a^b w(s)ds. \quad (13)$$
  
 724

725 Combining the above three equations, we have the following equation:  
 726

727 
$$\text{Prec}^+_{[a,b]} = \frac{\int_a^b w(s)\mathbb{P}(\text{correct} | \text{confidence} = s)ds}{\int_a^b w(s)ds} = \frac{\int_a^b w(s)f(s)ds}{\int_a^b w(s)ds}, \quad (14)$$
  
 728

729 where  $\int_a^b \mathbb{P}(\text{correct} | \text{confidence} = s)ds$  is the area under the reliability curve over a given confidence  
 730 interval  $[a, b]$ . Then the precision equals that area divided by the length of the interval.  
 731

732 For the incorrect prediction detection, which just needs to replace the numerator of Eq. (11) with the  
 733 number of incorrect predictions, thus we have:  
 734

735 
$$\text{Prec}^-_{[a,b]} = \frac{\int_a^b w(s)\mathbb{P}(\text{incorrect} | \text{confidence} = s)ds}{\int_a^b w(s)ds} = \frac{\int_a^b w(s)(1 - f(s))ds}{\int_a^b w(s)ds}. \quad (15)$$
  
 736

737 The proof is completed.  
 738

739  $\square$

740 **B.2 PROOF FOR COROLLARY 3.3**  
 741

742 **Corollary B.2** (Ideal separation case). *When  $\lambda \rightarrow 0$ , the proposed normalized sigmoid curve assigns  
 743 higher confidence to all correctly classified samples than to any misclassified sample, i.e.,*

744 
$$\mathbb{P}(\text{Conf}(N^+) > \text{Conf}(N^-)) = 1, \quad (16)$$
  
 745

746 where  $N^+$  and  $N^-$  denote the sets of correctly classified and misclassified samples, respectively. In  
 747 this case, the reliability curve approaches a step-like curve (cf. Fig. 2 with  $\lambda = 1e - 5$ ).  
 748

749 *Proof.* As  $\lambda \rightarrow 0$ , Figure 2 shows that all samples with confidence greater than 0.5 are correctly  
 750 classified (accuracy 100%), while all samples with confidence below 0.5 are misclassified (accuracy  
 751 0%). This implies that the confidence of every correctly classified sample is strictly greater than 0.5,  
 752 and the confidence of every misclassified sample is strictly less than 0.5. Therefore, the confidence of  
 753 any correctly classified sample is always higher than that of any misclassified sample.  $\square$

756 B.3 PROOF FOR THEOREM 3.2  
757

758 **Theorem B.3.** Let  $r \in [0, 1]$  be a confidence threshold. Under perfect calibration, the precision  
759 for the correct prediction detection and incorrect prediction detection tasks is  $\text{Prec}_{[r,1]}^+ =$   
760  $\mathbb{E}_{s \sim w(s|s \in [r,1])}[s]$  and  $\text{Prec}_{[0,r]}^- = \mathbb{E}_{s \sim w(s|s \in [0,r])}[1 - s]$ , respectively.  
761

762 *Proof.* Given a perfect calibration model, which reliability curve is a diagonal line (i.e.,  $f(s) =$   
763  $s$ ). For the correct prediction detection task, predictions with confidence above the threshold  $r$   
764 (i.e.,  $s \in [r, 1]$ ) are considered as detected correct predictions. By Lemma 3.1, the precision of correct  
765 prediction in  $[r, 1]$  is

$$766 \quad \text{Prec}_{[r,1]}^+ = \frac{\int_r^1 w(s)f(s)ds}{\int_r^1 w(s)ds}. \quad (17)$$

769 Substituting  $f(s) = s$ , we have:

$$771 \quad \text{Prec}_{[r,1]}^+ = \frac{\int_r^1 w(s)f(s)ds}{\int_r^1 w(s)ds} = \mathbb{E}_{s \sim w(s|s \in [r,1])}[s]. \quad (18)$$

774 Similarly, the precision of incorrect predictions is:

$$776 \quad \text{Prec}_{[0,1-r]}^- = \frac{\int_0^r w(s)(1 - f(s))ds}{\int_0^r w(s)ds} = \mathbb{E}_{s \sim w(s|s \in [0,r])}[1 - s]. \quad (19)$$

779 The proof is completed.  $\square$

781 B.4 PROOF FOR THEOREM 3.4  
782

783 **Theorem B.4** (Theorem 3.4 (restated)). Let  $r \in (0.5, 1)$  and let  $w \geq 0$  be any weight function on  
784  $[0, 1]$  with  $\int_r^1 w(s) ds > 0$ . The precision for correct-prediction detection (i.e.,  $\text{Prec}^+$ ) satisfies the  
785 following inequality:

$$786 \quad \text{Prec}_{\Psi}^+(r; w) := \frac{\int_r^1 w(s) \Psi(s) ds}{\int_r^1 w(s) ds} \geq \frac{\int_r^1 w(s) s ds}{\int_r^1 w(s) ds} =: \text{Prec}_{\text{diag}}^+(r; w).$$

789 Moreover, for the incorrect prediction detection with  $r \in (0, 0.5)$ , the precision (i.e.,  $\text{Prec}^-$ ) satisfies  
790 the following inequality:

$$792 \quad \text{Prec}_{\Psi}^-(r; w) := \frac{\int_0^r w(s) (1 - \Psi(s)) ds}{\int_0^r w(s) ds} \geq \frac{\int_0^r w(s) (1 - s) ds}{\int_0^r w(s) ds} =: \text{Prec}_{\text{diag}}^-(r; w).$$

795 *Proof.* Before proving the Theorem, we state the following lemma (proof in Appendix B.5).

797 **Lemma B.5** (Above-diagonal on the high-confidence side). Given the proposed normalized sigmoid  
798 curve (i.e.,  $\Psi(x)$ ) and diagonal line (i.e.,  $y = x$ ), then  $\Psi$  is strictly concave on  $[0.5, 1]$ , satisfies  
799  $\Psi(0.5) = 0.5$  and  $\Psi(1) = 1$ , and hence:

$$800 \quad \Psi(x) \geq x \quad \text{for all } x \in [0.5, 1], \quad \Psi(x) > x \quad \text{for all } x \in (0.5, 1).$$

803 By Lemma B.5, let  $D(s) := \Psi(s) - s \geq 0$  for  $s \in [r, 1]$  and  $D(s) > 0$  on  $(r, 1)$ . Therefore

$$805 \quad \int_r^1 w(s) \Psi(s) ds - \int_r^1 w(s) s ds = \int_r^1 w(s) D(s) ds \geq 0,$$

807 and the inequality is strict if  $w$  places positive mass on some subset of  $(r, 1)$  where  $D > 0$ . Division  
808 by the common normalizer  $\int_r^1 w(s) ds > 0$  yields the claim, then the inequality  $\text{Prec}_{\Psi}^+(r; w) >$   
809  $\text{Prec}_{\text{diag}}^+$  holds.

810 For the incorrect prediction detection, since  $\Psi(s) \leq s$  on  $[0, 0.5]$  (strict on  $(0, 0.5)$  by symmetry of  
 811 the above argument), an analogous statement holds for incorrect prediction detection on  $[0, r]$  with  
 812  $r < 0.5$ : for any nonnegative  $w$  with  $\int_0^r w(s) ds > 0$ ,

$$\frac{\int_0^r w(s) (1 - \Psi(s)) ds}{\int_0^r w(s) ds} \geq \frac{\int_0^r w(s) (1 - s) ds}{\int_0^r w(s) ds}.$$

817 Then the inequality  $\text{Prec}_{\Psi}^-(r; w) > \text{Prec}_{\text{diag}}^-$  holds. □

## 821 B.5 PROOF FOR LEMMA B.5

823 **Lemma** (Lemma B.5 (restated)). *Given the proposed normalized sigmoid curve (i.e.,  $\Psi(x)$ ) and  
 824 diagonal line (i.e.,  $y = x$ ), then  $\Psi$  is strictly concave on  $[0.5, 1]$ , satisfies  $\Psi(0.5) = 0.5$  and  $\Psi(1) = 1$ ,  
 825 and hence:*

$$826 \quad \Psi(x) \geq x \quad \text{for all } x \in [0.5, 1], \quad \Psi(x) > x \quad \text{for all } x \in (0.5, 1).$$

828 *Proof.* Let  $\phi(x) := \sigma\left(\frac{x-0.5}{\lambda}\right)$ . Since  $\sigma'(z) = \sigma(z)(1 - \sigma(z))$  and  $\sigma''(z) = \sigma(z)(1 - \sigma(z))(1 - 2\sigma(z))$ , we have for  $x > 0.5$  that  $\frac{x-0.5}{\lambda} > 0$  and thus  $\sigma\left(\frac{x-0.5}{\lambda}\right) > \frac{1}{2}$ , implying  $\sigma''\left(\frac{x-0.5}{\lambda}\right) < 0$ .  
 829 Hence

$$830 \quad \phi''(x) = \frac{1}{\lambda^2} \sigma''\left(\frac{x-0.5}{\lambda}\right) < 0 \quad (x \in (0.5, 1)),$$

833 i.e.,  $\phi$  is strictly concave on  $[0.5, 1]$ . The normalization

$$835 \quad \Psi(x) = \frac{\phi(x) - \phi(0)}{\phi(1) - \phi(0)}, \quad \phi(0) = \sigma\left(-\frac{0.5}{\lambda}\right), \quad \phi(1) = \sigma\left(\frac{0.5}{\lambda}\right),$$

838 is an affine transform and therefore preserves strict concavity. Direct evaluation gives

$$839 \quad \Psi(1) = 1, \quad \Psi(0.5) = \frac{\sigma(0) - \sigma(-a)}{\sigma(a) - \sigma(-a)} = \frac{\frac{1}{2} - (1 - \sigma(a))}{\sigma(a) - (1 - \sigma(a))} = \frac{1}{2}, \quad a = \frac{0.5}{\lambda} > 0.$$

842 Define  $F(x) := \Psi(x) - x$ . Since  $\Psi$  is strictly concave and  $-x$  is linear (hence both concave and  
 843 convex),  $F$  is strictly concave on  $[0.5, 1]$ , and  $F(0.5) = F(1) = 0$ . For any  $x \in (0.5, 1)$ , write  
 844  $x = \lambda \cdot 0.5 + (1 - \lambda) \cdot 1$  with  $\lambda \in (0, 1)$ ; strict concavity yields

$$845 \quad F(x) > \lambda F(0.5) + (1 - \lambda) F(1) = 0.$$

847 Thus  $\Psi(x) > x$  on  $(0.5, 1)$  and  $\Psi(x) \geq x$  on  $[0.5, 1]$ . □

## 849 B.6 PROOF FOR THEOREM 3.5

851 **Theorem B.6** (Theorem 3.5 (restated)). *Given arbitrary confidence  $s$ , then the corresponding entropy  
 852 for the proposed reliability curve (i.e.,  $\Psi(s)$ ) and perfect-calibration curve satisfy the following  
 853 inequality.*

$$854 \quad \mathbf{M}_{\Psi}(s) \leq \mathbf{M}_{\text{diag}}(s), \quad \mathbf{M}(s) := -\mathbb{P}(N_s^+) \log(\mathbb{P}(N_s^+)) - \mathbb{P}(N_s^-) \log(\mathbb{P}(N_s^-)), \quad (20)$$

856 where  $N_s^+$  and  $N_s^-$  denote the conditional probabilities of being correctly and incorrectly classified  
 857 given confidence  $s$ , respectively; i.e.,  $N_s^+ = \mathbb{P}(\text{correct} | \text{conf} = s)$  and  $N_s^- = 1 - N_s^+$ .

859 *Proof.* Since the vertical axis of the reliability diagram is accuracy, and given an arbitrary confidence  
 860  $s$ , we have  $\mathbb{P}(N_s^+) = \text{Acc}(s)$  and  $\mathbb{P}(N_s^-) = 1 - \text{Acc}(s)$ . Due to  $\mathbb{P}(N_s^+) + \mathbb{P}(N_s^-) = 1$ , then the  
 861 entropy can be written as binary entropy form:

$$862 \quad \mathbf{M}(s) = -(p_s \ln p_s + (1 - p_s) \ln(1 - p_s)), \quad (21)$$

863 where  $p_s = \mathbb{P}(N_s^+)$ .

864 Then, we can obtain the first-order derivation of  $p_s$ :

$$866 \quad \frac{\partial \mathbf{M}(s)}{\partial p_s} = -(\ln p_s + 1 - (\ln(1-p) + 1)) = \ln \frac{1-p}{p}. \quad (22)$$

869 When  $p_s = 0.5$ ,  $\mathbf{M}(s)$  reaches its maximum value in the interval  $[0, 1]$  and  $\frac{\partial^2 \mathbf{M}(s)}{\partial p_s^2} < 0$ . Thus, the  
870 closer  $p_s$  is to 0.5, the larger the entropy becomes.

871 Let us recall the Lemma B.5, when  $s \in [0, 0.5]$  we have  $\Psi(s) \leq s \leq 0.5$ , then:

$$873 \quad |\Psi(s) - 0.5| = 0.5 - \Psi(s) \geq 0.5 - s = |s - 0.5|. \quad (23)$$

875 Similarly, we can easily obtain that when  $s \in [0.5, 1]$  we have  $\Psi(s) \geq s \geq 0.5$ , then:

$$876 \quad |\Psi(s) - 0.5| = \Psi(s) - 0.5 \geq s - 0.5 = |s - 0.5|. \quad (24)$$

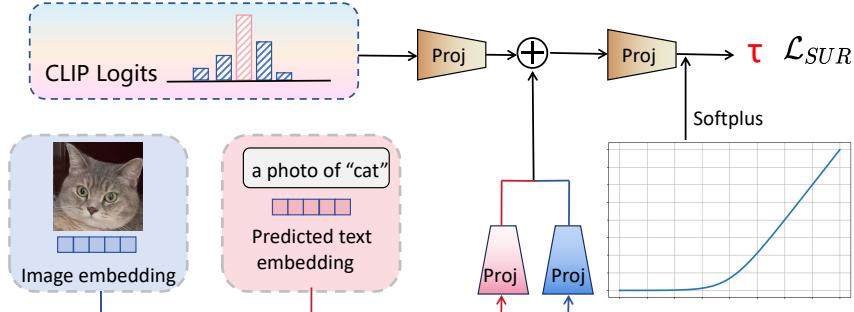
878 Based on Eq. (23) and Eq. (24),  $|\Psi(s) - 0.5| \geq |s - 0.5|$  holds for all  $s \in [0, 1]$ . Therefore, the  $p_s$   
879 of the perfect-calibration curve  $y = x$  lies closer to 0.5 than that of the proposed  $\Psi(s)$ . As a result,  
880 the perfect calibration curve exhibits higher entropy, thereby tolerating a greater degree of mixing  
881 between correctly and incorrectly predicted samples.

882  $\square$

## 885 C MODEL DETAILS

### 887 C.1 FLOWCHART VISUALIZATION

888 This section gives a flowchart (*i.e.*, Figure 5) of the proposed lightweight meta networks (LMN).



901 Figure 5: Flowchart of the proposed lightweight meta network (LMN). Given the output logits  
902 of CLIP, image embedding, and predicted text embedding, LMN first projects them into a shared  
903 subspace using distinct FC layers (indicated by different colors in the figure). The resulting vectors  
904 are concatenated and mapped to a scalar via another FC layer, followed by a Softplus activation to  
905 produce the sample-specific  $\tau$ . Finally, the surrogate loss  $\mathcal{L}_{SUR}$  updates the FC layers.

### 907 C.2 UNDERSTANDING THE CONTRIBUTION OF EACH MODALITY IN LMN

909 The meta-network indeed leverages three types of signals (*i.e.*, logits, image embeddings, and  
910 predicted text embeddings) and each contributes complementary information for predicting the  
911 effective temperature factor.

912 **Logits.** Logits as the primary calibration signal. Logits provide the most direct evidence for  
913 confidence misalignment, as they encode the inter-class margins and the overall shape of the predictive  
914 distribution. This is the core quantity used by traditional temperature scaling. Here we describe the  
915 complementary roles of each modality and why all of them contribute meaningfully to MisD:

917 **Image embeddings.** Image embeddings capture sample difficulty. Even when two samples share  
918 similar logits, their underlying visual characteristics may differ substantially. Image embeddings help

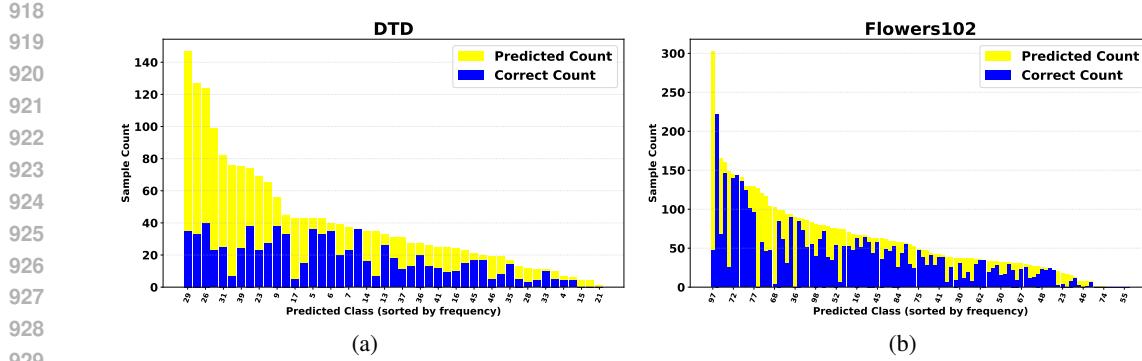


Figure 6: Predicted count (yellow) and correct count (blue) for each predicted class, sorted by prediction frequency.

LMN identify hard or visually atypical samples (e.g., unusual textures, crowded scenes, rare visual patterns). This allows LMN to incorporate sample-level visual cues.

**Predicted text embeddings.** Predicted text embeddings reveal systematic category-level patterns. To further illustrate the semantic confusion patterns captured by the predicted text embeddings, we visualize the distribution of predicted samples across classes in Figure 6. The figure shows that CLIP tends to cluster certain misclassified samples around a few semantically related text prototypes. Incorporating text embeddings enables LMN to capture these category-level tendencies and selectively increase correction strength for categories prone to systematic confusion.

Taken together, the three modalities provide complementary and non-redundant information: logits characterize confidence geometry, image embeddings characterize sample difficulty, and text embeddings capture semantic misalignment patterns.

### C.3 MOTIVATION FOR CHOOSING R

We set the midpoint parameter  $r = 0.5$  because it corresponds to the natural neutral threshold separating ‘‘uncertain’’ from ‘‘confident’’ predictions, which aligns with standard confidence-based calibration intuition (Johansson et al., 2023). Moreover, placing the inflection point at the center of the interval  $[0, 1]$  avoids introducing asymmetry or favoring either low- or high-confidence regions, providing a balanced and unbiased target curve.

## D EXPERIMENTAL SETTING

### D.1 DATASETS DETAILS

Table 3: The statistics of the used datasets.

	Num. classes ( $ \mathcal{Y} $ )	Size training data	Avg. labeled data per class	Size test
DTD	47	3760	64	1880
Flowers102	102	2040	16	6149
EuroSAT	10	27000	2200	5000
RESICS45	45	6300	110	25200
MNIST	10	60000	4696	10000
CUB	200	5594	26	5794

In the experiment section, we use six datasets. The statistics of the used datasets are reported in the Table 3. Here, we provide a description of each of them:

1. **DTD** (Cimpoi et al., 2014) The Describable Textures Dataset (DTD) is a continuously expanding collection of texture images captured in unconstrained environments. The

972 annotations are based on human-interpretable attributes, reflecting perceptual characteristics  
 973 of textures.  
 974

2. **Flowers102** (Nilsback & Zisserman, 2008) It is a benchmark dataset that collects images of  
 975 102 distinct flower categories, each representing species commonly occurring in the United  
 976 Kingdom.  
 977

3. **EuroSAT** (Helber et al., 2019) EuroSAT focuses on the classification of satellite scene  
 978 images, providing 10 predefined categories for evaluation.  
 979

4. **RESICS45** (Cheng et al., 2017) Designed for remote sensing image scene classification,  
 980 this benchmark dataset provides 45 categories of scenes and is publicly accessible.  
 981

5. **MNIST** (Deng, 2012) MNIST serves as a benchmark collection of handwritten digit images,  
 982 in which all samples are normalized in terms of size and aligned to the center of a uniform  
 983 image grid.  
 984

6. **CUB** (Wah et al., 2011) CUB is a fine-grained bird species classification dataset, containing  
 985 200 categories with large intra-class variation and subtle inter-class differences. It serves as  
 986 a benchmark for subordinate categorization tasks and provides comprehensive annotations  
 987 to support research in fine-grained visual recognition.  
 988

989

990  
 991 Table 4: Settings for the proposed method.

	learning rate	$\beta$	Num. hidden	$\lambda$	$\phi$
DTD	0.002	0.6	16	0.05	summation
Flowers102	0.001	0.8	16	0.05	summation
EuroSAT	0.002	0.6	16	0.1	summation
RESICS45	0.001	0.8	64	0.05	multiply
MNIST	0.005	0.4	16	0.05	summation
CUB	0.002	0.9	8	0.05	multiply

1000  
 1001 D.2 IMPLEMENTATION DETAILS

1002 We use CLIP ViT-B/32 as the visual backbone and report results averaged over 3 runs. For both  
 1003 visual and textual prompt learning, we set the prefix size to 16 (Zhou et al., 2022b; Jia et al., 2022).  
 1004 We use SGD as the optimizer, with the number of training epochs selected from {100, 150, 200}  
 1005 and the learning rate selected from {0.001, 0.002, 0.005}. We adopt 16 labeled samples per class  
 1006 (*i.e.*, 16-shot) as the calibration set. When the base model is prompt-tuning CLIP, we split the original  
 1007 16-shot calibration set into two parts: one is used for learning the prompts, and the other is reserved  
 1008 for training our post-hoc calibration network. Table 4 describes the detailed settings and architecture  
 1009 for the proposed lightweight meta network based on pretrained CLIP.  
 1010

1011 E ADDITIONAL EXPERIMENTS  
 10121013 E.1 ABLATION STUDY  
 1014

1015 The key component of the proposed method is the surrogate loss  $\mathcal{L}_{SUR}$ , which consists of two parts:  
 1016 (i) a confidence regularization term that encourages the predictions to align with the target reliability  
 1017 curve (*i.e.*,  $\Phi(s)$ ), and (ii) a full-probability constraint (FPC) that complements the confidence  
 1018 regularization by regularizing the entire probability distribution, preventing the training instability  
 1019 cause by single probability (*i.e.*, confidence) were optimized. Therefore, we do not report a separate  
 1020 ablation with only the confidence regularization term, as it is intended to be used together with FPC.  
 1021 To verify the effectiveness of these components, we investigate the performance of the following  
 1022 variants: the original CLIP, CLIP with FPC, and CLIP with the full proposed method. The results are  
 1023 reported in Table 5.  
 1024

1025 According to Table 5, the complete objective function achieves the best performance. Notably,  
 1026 employing only partial components (*i.e.*, +FPC) also leads to improvements. Taken together, these  
 1027 findings verify the effectiveness of all proposed modules.

1026

1027 Table 5: Results of the ablation study on six datasets, measured by AUROC, FPR90-S, and FPR90-E.  
1028  $\uparrow$  denotes that higher values are better, while  $\downarrow$  denotes that lower values are better.

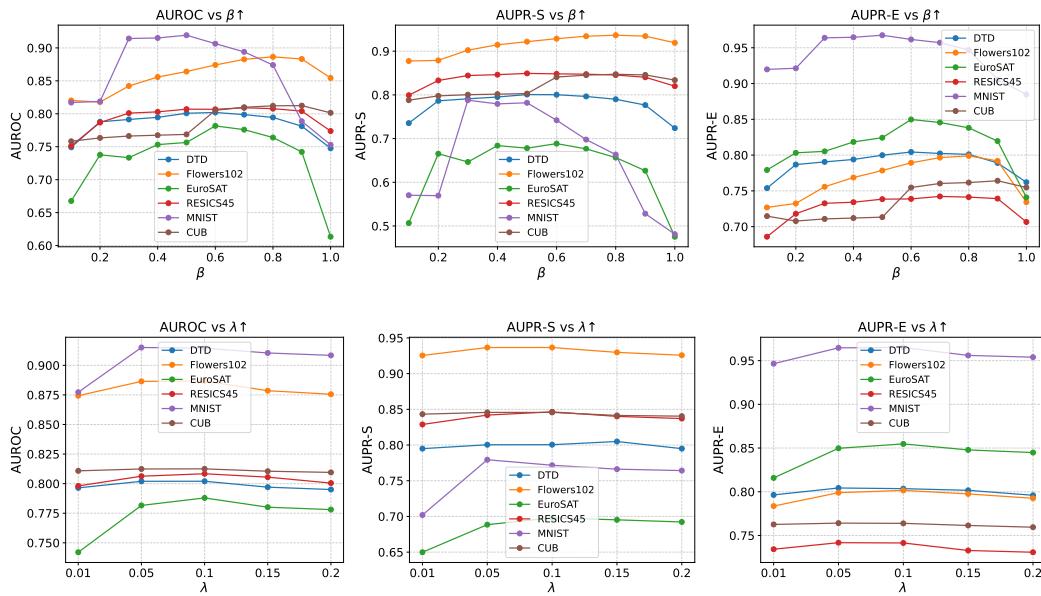
1029	Methods	DTD			Flowers102			EuroSAT		
		AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$	AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$	AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$
1030	Zero-shot CLIP	0.762	0.669	0.572	0.864	0.435	0.354	0.650	0.782	0.742
1031	+FPC	0.795	0.648	0.471	0.878	0.388	0.325	0.778	0.704	0.471
1032	+ALL	<b>0.802</b>	<b>0.636</b>	<b>0.457</b>	<b>0.886</b>	<b>0.378</b>	<b>0.305</b>	<b>0.788</b>	<b>0.665</b>	<b>0.468</b>
1033	RESICS45			MNIST			CUB			
1034	Methods	AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$	AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$	AUROC $\uparrow$	FPR90-S $\downarrow$	FPR90-E $\downarrow$
1035	Zero-shot CLIP	0.779	0.636	0.508	0.813	0.511	0.482	0.807	0.839	0.758
1036	+FPC	0.794	0.615	0.482	0.908	0.228	0.235	0.808	0.613	0.440
1037	+ALL	<b>0.808</b>	<b>0.597</b>	<b>0.445</b>	<b>0.915</b>	<b>0.200</b>	<b>0.205</b>	<b>0.812</b>	<b>0.602</b>	<b>0.438</b>

1038

## 1039 E.2 PARAMETER ANALYSIS

1040

1041



1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063 Figure 7: Line plots of model performance across six datasets (*i.e.*, DTD, Flowers102, EuroSAT,  
1064 RESICS45, MNIST, and CUB) under varying hyperparameters. The first row shows the effect of  
1065  $\beta$ , and the second row shows the effect of  $\lambda$ . From left to right, each subfigure corresponds to one  
1066 evaluation metric: AUROC, AUPR-S, and AUPR-E. The curves illustrate how performance changes  
1067 as the hyperparameter values increase.

1068

1069

1070

1071

1072

1073 In the proposed method, we employ the non-negative parameters (*i.e.*,  $\beta$  and  $\lambda$ ). For  $\beta$ , it is used to  
1074 achieve a trade-off between the two terms of the objective function (*i.e.*,  $\mathcal{L}_{SUR}$ ). For  $\lambda$ , which is used  
1075 to control the smoothness of the reliability curve. To investigate the impact of  $\beta$  and  $\lambda$  with different  
1076 settings, we conduct experiments on all six datasets by varying the value of  $\beta$  in the range of [0.1,  
1077 1.0] and the value of  $\lambda$  in the range of [0.01, 0.2]. The results are shown in Figure 7.

1078

1079

1080 From Figure 7, we make the following observations. First, for the hyperparameter  $\beta$ , the proposed  
1081 method achieves peak performance around  $\beta = 0.6$ . Performance deteriorates when  $\beta$  is too large  
1082 or too small, as the method then fails to balance focus between correctly and incorrectly predicted  
1083 samples. Second, for the hyperparameter  $\lambda$ , the method consistently performs well when  $\lambda$  is set  
1084 appropriately (e.g., [0.05, 0.1]). Similarly to  $\beta$ , extreme values of  $\lambda$  lead to inferior performance. This  
1085 is because  $\lambda$  controls the smoothness of  $\Psi(s)$ : if  $\Psi(s)$  is too sharp, the overly aggressive probability  
1086 updates can cause unstable optimization; if it is too smooth, the model loses the ability to effectively  
1087 adjust predictions. **In practice, we find that the parameter  $\lambda$  is highly stable across different datasets.**

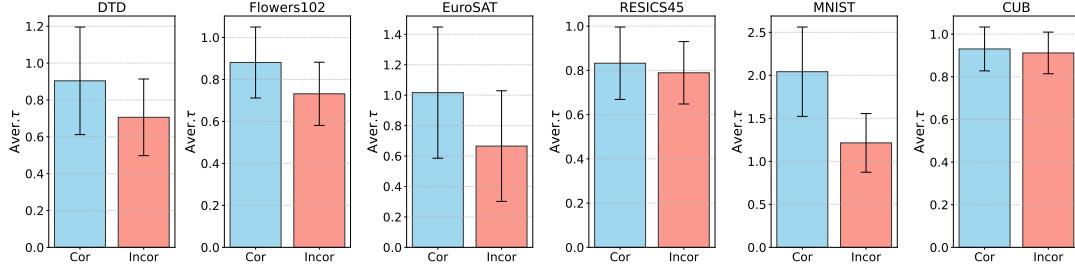
1080 As summarized in Table 4, five out of six datasets (*i.e.*, spanning textures, fine-grained objects,  
 1081 satellite imagery, handwriting, and birds) select the same value  $\lambda = 0.05$ . This consistency indicates  
 1082 that  $\lambda$  does not require dataset-specific tuning. Therefore, we recommend  $\lambda = 0.05$  as a reliable  
 1083 default, which already achieves near-optimal performance in our experiments, with adjustments  
 1084 needed only in rare cases.

1085

### 1086 E.3 ANALYSIS OF THE LEARNED TEMPERATURE COEFFICIENT

1087

1088



1099 Figure 8: Comparison of learned  $\tau$  values for correct versus incorrect predictions across six datasets.  
 1100 Each bar represents the average  $\tau$  for the corresponding sample type. The blue bars represent the  
 1101 correctly predicted samples, and the red bars represent the incorrectly predicted samples.

1102

1103 To verify that the learned instance-wise  $\tau$  indeed increases the confidence of correctly predicted  
 1104 samples compared with incorrectly predicted ones. Since the temperature coefficient  $\tau$  is multiplied  
 1105 by the logits, a larger value of  $\tau$  results in higher confidence. We compare the average learned  $\tau$   
 1106 values between correct and incorrect predictions on all six datasets. The results are visualized in  
 1107 Figure 8. In Figure 8, we can observe that the average value of the learned  $\tau$  for correct predictions is  
 1108 larger than that for incorrect predictions, with a larger gap on most datasets. For example, on MNIST,  
 1109 the gap is about 68.5% of the incorrect predictions. This demonstrates that the proposed method is  
 1110 able to capture the intrinsic differences between correctly and incorrectly classified samples, thereby  
 1111 enabling their discrimination.

1112

1113

### 1114 E.4 RESULTS ON SIGLIP BACKBONE

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1114 To further examine whether LMN generalizes beyond the CLIP-B/32, we additionally evaluate the  
 1115 method using CLIP-L/14 and SigLIP-B/16 Zhai et al. (2023) as the backbones. SigLIP represents  
 1116 a modern vision-language architecture trained with a different alignment objective and embedding  
 1117 structure, offering a stronger testbed for assessing the backbone-agnostic nature of LMN. Because  
 1118 most existing calibration and MisD-oriented baselines do not provide publicly available or repro-  
 1119 ducible implementations for SigLIP, SCT is the only method that can be feasibly adapted to this  
 1120 backbone, and we therefore adopt it as the comparison baseline in this setting. The results are shown  
 1121 in Table 6 and Table 7.

1122 From Table 6 and Table 7, we can observe that, across all datasets and all backbones, LMN continues  
 1123 to deliver consistent improvements over SCT under both AUROC and FPR90E, demonstrating that  
 1124 its effectiveness is not restricted to the CLIP-B/32 backbone and extends naturally to more advanced  
 1125 VLM architectures. These results confirm that LMN’s surrogate objective and meta-network design  
 1126 retain their advantages even when the underlying visual encoder and multimodal representation  
 1127 mechanism differ substantially from those of CLIP-B/32.

### 1128 E.5 CALIBRATION DATA SENSITIVITY ANALYSIS

1129

1130

1131

1132

1133

1128 To evaluate whether the effectiveness of LMN depends on the specific 16-shot calibration configura-  
 1129 tion used in the main experiments, we additionally conduct a sensitivity study by varying the number  
 1130 of calibration samples per class. Specifically, we consider 4-shot, 8-shot, 16-shot, 32-shot, and  
 1131 64-shot settings on datasets with sufficiently large training sets, including DTD, EuroSAT, RESICS45,  
 1132 and MNIST. For Flowers102 and CUB, the available training samples are too limited to support

1134  
 1135 Table 6: MisD Performance comparison under CLIP-L/14. For each dataset, AUROC ( $\uparrow$ ) and FPR90E  
 1136 ( $\downarrow$ ) are reported.

1137 1138 1139 1140 1141 1142 1143 1144 1145 1146 1147 1148 1149 1150 1151 1152 1153 1154 1155 1156 1157 1158 1159 1160 1161 1162 1163 1164 1165 1166 1167 1168 1169 1170 1171 1172 1173 1174 1175 1176 1177 1178 1179 1180 1181 1182 1183 1184 1185 1186 1187 1188 1189 1190 1191 1192 1193 1194 1195 1196 1197 1198 1199 1200 1201 1202 1203 1204 1205 1206 1207 1208 1209 1210 1211 1212 1213 1214 1215 1216 1217 1218 1219 1220 1221 1222 1223 1224 1225 1226 1227 1228 1229 1230 1231 1232 1233 1234 1235 1236 1237 1238 1239 1240 1241 1242 1243 1244 1245 1246 1247 1248 1249 1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275 1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292 1293 1294 1295 1296 1297 1298 1299 1299 1300 1301 1302 1303 1304 1305 1306 1307 1308 1309 1310 1311 1312 1313 1314 1315 1316 1317 1318 1319 1320 1321 1322 1323 1324 1325 1326 1327 1328 1329 1330 1331 1332 1333 1334 1335 1336 1337 1338 1339 1340 1341 1342 1343 1344 1345 1346 1347 1348 1349 1350 1351 1352 1353 1354 1355 1356 1357 1358 1359 1360 1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377 1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394 1395 1396 1397 1398 1399 1399 1400 1401 1402 1403 1404 1405 1406 1407 1408 1409 1410 1411 1412 1413 1414 1415 1416 1417 1418 1419 1420 1421 1422 1423 1424 1425 1426 1427 1428 1429 1430 1431 1432 1433 1434 1435 1436 1437 1438 1439 1440 1441 1442 1443 1444 1445 1446 1447 1448 1449 1450 1451 1452 1453 1454 1455 1456 1457 1458 1459 1460 1461 1462 1463 1464 1465 1466 1467 1468 1469 1470 1471 1472 1473 1474 1475 1476 1477 1478 1479 1480 1481 1482 1483 1484 1485 1486 1487 1488 1489 1490 1491 1492 1493 1494 1495 1496 1497 1498 1499 1499 1500 1501 1502 1503 1504 1505 1506 1507 1508 1509 1509 1510 1511 1512 1513 1514 1515 1516 1517 1518 1519 1519 1520 1521 1522 1523 1524 1525 1526 1527 1528 1529 1529 1530 1531 1532 1533 1534 1535 1536 1537 1538 1539 1539 1540 1541 1542 1543 1544 1545 1546 1547 1548 1549 1549 1550 1551 1552 1553 1554 1555 1556 1557 1558 1559 1560 1561 1562 1563 1564 1565 1566 1567 1568 1569 1570 1571 1572 1573 1574 1575 1576 1577 1578 1579 1580 1581 1582 1583 1584 1585 1586 1587 1588 1589 1590 1591 1592 1593 1594 1595 1596 1597 1598 1599 1599 1600 1601 1602 1603 1604 1605 1606 1607 1608 1609 1609 1610 1611 1612 1613 1614 1615 1616 1617 1618 1619 1619 1620 1621 1622 1623 1624 1625 1626 1627 1628 1629 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 1674 1675 1676 1677 1678 1679 1680 1681 1682 1683 1684 1685 1686 1687 1688 1689 1689 1690 1691 1692 1693 1694 1695 1696 1697 1698 1699 1699 1700 1701 1702 1703 1704 1705 1706 1707 1708 1709 1709 1710 1711 1712 1713 1714 1715 1716 1717 1718 1719 1719 1720 1721 1722 1723 1724 1725 1726 1727 1728 1729 1729 1730 1731 1732 1733 1734 1735 1736 1737 1738 1739 1739 1740 1741 1742 1743 1744 1745 1746 1747 1748 1749 1749 1750 1751 1752 1753 1754 1755 1756 1757 1758 1759 1759 1760 1761 1762 1763 1764 1765 1766 1767 1768 1769 1769 1770 1771 1772 1773 1774 1775 1776 1777 1778 1779 1780 1781 1782 1783 1784 1785 1786 1787 1788 1789 1789 1790 1791 1792 1793 1794 1795 1796 1797 1798 1799 1799 1800 1801 1802 1803 1804 1805 1806 1807 1808 1809 1809 1810 1811 1812 1813 1814 1815 1816 1817 1818 1819 1819 1820 1821 1822 1823 1824 1825 1826 1827 1828 1829 1829 1830 1831 1832 1833 1834 1835 1836 1837 1838 1839 1839 1840 1841 1842 1843 1844 1845 1846 1847 1848 1849 1849 1850 1851 1852 1853 1854 1855 1856 1857 1858 1859 1859 1860 1861 1862 1863 1864 1865 1866 1867 1868 1869 1869 1870 1871 1872 1873 1874 1875 1876 1877 1878 1879 1880 1881 1882 1883 1884 1885 1886 1887 1888 1889 1889 1890 1891 1892 1893 1894 1895 1896 1897 1898 1899 1899 1900 1901 1902 1903 1904 1905 1906 1907 1908 1909 1909 1910 1911 1912 1913 1914 1915 1916 1917 1918 1919 1919 1920 1921 1922 1923 1924 1925 1926 1927 1928 1929 1929 1930 1931 1932 1933 1934 1935 1936 1937 1938 1939 1939 1940 1941 1942 1943 1944 1945 1946 1947 1948 1949 1949 1950 1951 1952 1953 1954 1955 1956 1957 1958 1959 1959 1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1980 1981 1982 1983 1984 1985 1986 1987 1988 1989 1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019 2019 2020 2021 2022 2023 2024 2025 2026 2027 2028 2029 2029 2030 2031 2032 2033 2034 2035 2036 2037 2038 2039 2039 2040 2041 2042 2043 2044 2045 2046 2047 2048 2049 2049 2050 2051 2052 2053 2054 2055 2056 2057 2058 2059 2059 2060 2061 2062 2063 2064 2065 2066 2067 2068 2069 2069 2070 2071 2072 2073 2074 2075 2076 2077 2078 2079 2079 2080 2081 2082 2083 2084 2085 2086 2087 2088 2089 2089 2090 2091 2092 2093 2094 2095 2096 2097 2098 2099 2099 2100 2101 2102 2103 2104 2105 2106 2107 2108 2109 2109 2110 2111 2112 2113 2114 2115 2116 2117 2118 2119 2119 2120 2121 2122 2123 2124 2125 2126 2127 2128 2129 2129 2130 2131 2132 2133 2134 2135 2136 2137 2138 2139 2139 2140 2141 2142 2143 2144 2145 2146 2147 2148 2149 2149 2150 2151 2152 2153 2154 2155 2156 2157 2158 2159 2160 2161 2162 2163 2164 2165 2166 2167 2168 2169 2170 2171 2172 2173 2174 2175 2176 2177 2178 2179 2180 2181 2182 2183 2184 2185 2186 2187 2188 2189 2189 2190 2191 2192 2193 2194 2195 2196 2197 2198 2199 2199 2200 2201 2202 2203 2204 2205 2206 2207 2208 2209 2209 2210 2211 2212 2213 2214 2215 2216 2217 2218 2219 2219 2220 2221 2222 2223 2224 2225 2226 2227 2228 2229 2229 2230 2231 2232 2233 2234 2235 2236 2237 2238 2239 2239 2240 2241 2242 2243 2244 2245 2246 2247 2248 2249 2249 2250 2251 2252 2253 2254 2255 2256 2257 2258 2259 2259 2260 2261 2262 2263 2264 2265 2266 2267 2268 2269 2269 2270 2271 2272 2273 2274 2275 2276 2277 2278 2279 2279 2280 2281 2282 2283 2284 2285 2286 2287 2288 2289 2289 2290 2291 2292 2293 2294 2295 2296 2297 2298 2299 2299 2300 2301 2302 2303 2304 2305 2306 2307 2308 2309 2309 2310 2311 2312 2313 2314 2315 2316 2317 2318 2319 2319 2320 2321 2322 2323 2324 2325 2326 2327 2328 2329 2329 2330 2331 2332 2333 2334 2335 2336 2337 2338 2339 2339 2340 2341 2342 2343 2344 2345 2346 2347 2348 2349 2349 2350 2351 2352 2353 2354 2355 2356 2357 2358 2359 2359 2360 2361 2362 2363 2364 2365 2366 2367 2368 2369 2369 2370 2371 2372 2373 2374 2375 2376 2377 2378 2379 2379 2380 2381 2382 2383 2384 2385 2386 2387 2388 2389 2389 2390 2391 2392 2393 2394 2395 2396 2397 2398 2399 2399 2400 2401 2402 2403 2404 2405 2406 2407 2408 2409 2409 2410 2411 2412 2413 2414 2415 2416 2417 2418 2419 2419 2420 2421 2422 2423 2424 2425 2426 2427 2428 2429 2429 2430 2431 2432 2433 2434 2435 2436 2437 2438 2439 2439 2440 2441 2442 2443 2444 2445 2446 2447 2448 2449 2449 2450 2451 2452 2453 2454 2455 2456 2457 2458 2459 2459 2460 2461 2462 2463 2464 2465 2466 2467 2468 2469 2469 2470 2471 2472 2473 2474 2475 2476 2477 2478 2479 2479 2480 2481 2482 2483 2484 2485 2486 2487 2488 2489 2489 2490 2491 2492 2493 2494 2495 2496 2497 2498 2499 2499 2500 2501 2502 2503 2504 2505 2506 2507 2508 2509 2509 2510 2511 2512 2513 2514 2515 2516 2517 2518 2519 2519 2520 2521 2522 2523 2524 2525 2526 2527 2528 2529 2529 2530 2531 2532 2533 2534 2535 2536 2537 2538 2539 2539 2540 2541 2542 2543 2544 2545 2546 2547 2548 2549 2549 2550 2551 2552 2553 2554 2555 2556 2557 2558 2559 2559 2560 2561 2562 2563 2564 2565 2566 2567 2568 2569 2569 2570 2571 2572 2573 2574 2575 2576 2577 2578 2579 2579 2580 2581 2582 2583 2584 2585 2586 2587 2588 2589 2589 2590 2591 2592 2593 2594 2595 2596 2597 2598 2599 2599 2600 2601 2602 2603 2604 2605 2606 2607 2608 2609 2609 2610 2611 2612 2613 2614 2615 2616 2617 2618 2619 2619 2620 2621 2622 2623 2624 2625 2626 2627 2628 2629 2629 2630 2631 2632 2633 2634 2635 2636 2637 2638 2639 2639 2640 2641 2642 2643 2644 2645 2646 2647 2648 2649 2649 2650 2651 2652 2653 2654 2655 2656 2657 2658 2659 2659 2660 2661 2662 2663 2664 2665 2666 2667 2668 2669 2669 2670 2671 2672 2673 2674 2675 2676 2677 2678 2679 2679 2680 2681 2682 2683 2684 2685 2686 2687 2688 2689 2689 2690 2691 2692 2693 2694 2695 2696 2697 2698 2699 2699 2700 2701 2702 2703 2704 2705 2706 2707 2708 2709 2709 2710 2711 2712 2713 2714 2715 2716 2717 2718 2719 2719 2720 2721 2722 2723 2724 2725 2726 2727 2728 2729 2729 2730 2731 2732 2733 2734 2735 2736 2737 2738 2739 2739 2740 2741 2742 2743 2744 2745 2746 2747 2748 2749 2749 2750 2751 2752 2753 2754 2755 2756 2757 2758 2759 2759 2760 2761 2762 2763 2764 2765 2766 2767 2768 2769 2769 2770 2771 2772 2773 2774 2775 2776 2777 2778 2779 2779 2780 2781 2782 2783 2784 2785 2786 2787 2788 2789 2789 2790 2791 2792 2793 2794 2795 2796 2797 2798 2799 2799 2800 2801 2802 2803 2804 2805 2806 2807 2808 2809 2809 2810 2811 2812 2813 2814 2815 2816 2817 2818 2819 2819 2820 2821 2822 2823 2824 2825 2826 2827 2828 2829 2829 2830 2831 2832 2833 2834 2835 2836 2837 2838 2839 2839 2840 2841 2842 2843 2844 2845 2846 2847 2848 2849 2849 2850 2851 2852 2853 2854 2855 2856 2857 2858 2859 2859 2860 2861 2862 2863 2864 2865 2866 2867 2868 2869 2869 2870 2871 2872 2873 2874 2875 2876 2877

1188

1189 Table 8: AUROC ( $\uparrow$ ) under different numbers of calibration shots. ‘-’ denotes that the result can not  
1190 be obtained.

Dataset	CLIP	4-shot	8-shot	16-shot	32-shot	64-shot
DTD	0.762	0.787	0.779	0.792	0.804	0.807
Flowers102	0.864	0.873	0.884	0.886	–	–
EuroSAT	0.650	0.732	0.765	0.788	0.792	0.797
RESICS45	0.779	0.793	0.804	0.808	0.810	0.815
MNIST	0.813	0.883	0.901	0.915	0.936	0.941
CUB	0.807	0.810	0.810	0.812	–	–

1198

1199

1200 Table 9: FPR90E ( $\downarrow$ ) under different numbers of calibration shots. CLIP denotes the zero-shot  
1201 baseline.

Dataset	CLIP	4-shot	8-shot	16-shot	32-shot	64-shot
DTD	0.572	0.512	0.474	0.504	0.438	0.431
Flowers102	0.354	0.329	0.317	0.305	–	–
EuroSAT	0.742	0.619	0.522	0.468	0.450	0.439
RESICS45	0.508	0.477	0.462	0.445	0.446	0.421
MNIST	0.482	0.301	0.264	0.205	0.160	0.142
CUB	0.554	0.549	0.538	0.532	–	–

1209

1210

1211 For smaller datasets, inference finishes within tens of seconds (e.g., 15.77s on DTD and 53.83s on  
1212 EuroSAT). These results validate that LMN remains practical and efficient even when applied to  
1213 large-scale test sets.1214 Overall, the empirical measurements demonstrate that LMN maintains its lightweight property in  
1215 both training and inference, making it suitable for real-world and large-scale deployments.

1216

1217 

## E.7 EVALUATION UNDER DISTRIBUTION SHIFT

1218

1219 To further assess the robustness of the proposed LMN under distribution shift, we conduct an  
1220 additional evaluation using the ImageNet-Val to train the LMN and evaluating on two distribution-  
1221 shifted benchmarks (*i.e.*, ImageNet-A and ImageNet-Sketch). Specifically, ImageNet-Sketch, which  
1222 introduces substantial style and texture shift, and ImageNet-A, which contains adversarially curated  
1223 natural images designed to induce model failures. We report both misclassification-detection metrics  
1224 (AUROC and FPR90) and calibration metrics (Brier score) to obtain a comprehensive view of the  
1225 model’s behavior under shift. The results are provided in Table 11.

1226

1227 From the Table 11, we can observe that LMN consistently improves AUROC and reduces FPR90-E  
1228 across both distribution-shifted datasets, indicating stronger discrimination between correct and  
1229 incorrect predictions in the presence of distribution shifts. Meanwhile, LMN maintains competitive  
1230 calibration performance. These findings confirm that LMN remains effective even when the input  
1231 distribution deviates significantly from the calibration set, further demonstrating the practicality of  
1232 the proposed approach.

1233

1234

1235 Table 10: Training/inference time and number of LMN parameters across six datasets. LMN is highly  
1236 lightweight, with only 17K–20K parameters and second-level training time on all datasets.

Time (s)	DTD	Flowers102	EuroSAT	RESICS45	MNIST	CUB
Train	13.31	29.63	7.05	13.14	6.81	48.35
Test	15.77	113.58	53.83	339.01	96.07	95.48
Num. Parameters	DTD	Flowers102	EuroSAT	RESICS45	MNIST	CUB
LMN	17.2K	18.8K	17.3K	18.1K	17.3K	20.3K

1242

1243

Table 11: Performance under distribution shift on ImageNet-A and ImageNet-Sketch.

1244

1245

Method	ImageNet-A			ImageNet-Sketch		
	AUROC↑	FPR90-E↓	Brier↓	AUROC↑	FPR90-E↓	Brier↓
Zero-shot CLIP	0.657	0.740	0.814	0.805	0.495	0.732
LMN	<b>0.682</b>	<b>0.716</b>	<b>0.796</b>	<b>0.813</b>	<b>0.484</b>	<b>0.721</b>

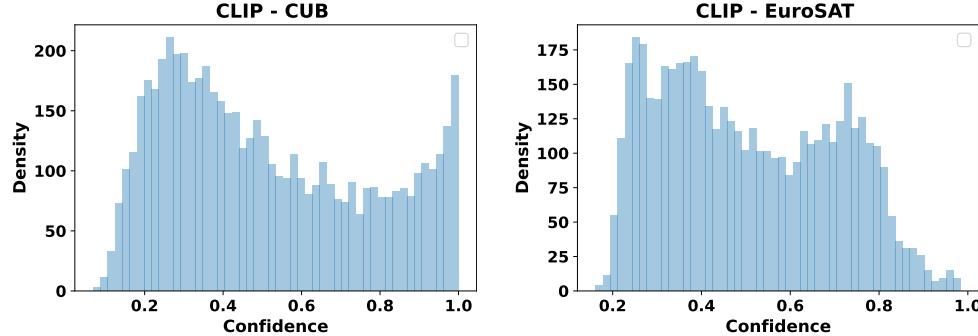
1246

1247

1248

1249

1250



1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

### E.8 ADDITIONAL EMPIRICALLY SUPPORT

1266

We empirically examine the confidence distribution on two benchmark datasets: CUB and EuroSAT, respectively. For the test set, we record the confidence assigned by the pretrained CLIP model to each sample. We then visualize the distribution using histograms to visualize the density of each confidence level. The results are shown in Figure 9.

1271

1272

1273

1274

1275

Empirically, as shown in Figure 9, CLIP confidence values on test samples are widely spread across the  $[0, 1]$  interval, with many samples falling in the middle range. As a result, the conditional expectations in  $[r, 1]$  for correct prediction detection and in  $[0, r]$  for incorrect prediction detection are far below 1, confirming that strict confidence calibration alone cannot achieve high MisD performance.

1276

1277

### F LLM USAGE STATEMENT

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295