

Risk Controlled Image Retrieval

Kaiwen Cai¹, Chris Xiaoxuan Lu², Xingyu Zhao³, Wei Huang⁴, Xiaowei Huang^{1*}

¹University of Liverpool

²University College London

³University of Warwick

⁴Purple Mountain Laboratories

Abstract

Most image retrieval research prioritizes improving predictive performance, often overlooking situations where the reliability of predictions is equally important. The gap between model performance and reliability requirements highlights the need for a systematic approach to analyze and address the risks associated with image retrieval. Uncertainty quantification technique can be applied to mitigate this issue by assessing uncertainty for retrieval sets, but it provides only a heuristic estimate of uncertainty rather than a *guarantee*. To address these limitations, we present Risk Controlled Image Retrieval (RCIR), which generates retrieval sets with coverage guarantee, i.e., retrieval sets that are guaranteed to contain the true nearest neighbors with a predefined probability. RCIR can be easily integrated with existing uncertainty-aware image retrieval systems, agnostic to data distribution and model selection. To the best of our knowledge, this is the first work that provides coverage guarantees to image retrieval. The validity and efficiency of RCIR are demonstrated on four real-world datasets: CAR-196, CUB-200, Pittsburgh, and ChestX-Det.

1 Introduction

Given a query image and a database, an image retrieval system selects the best matching candidates from the database. Image retrieval serves as the foundation for many computer vision applications, including face recognition (Cheng and Wang 2019), large-scale image classification (Warburg et al. 2021), and person re-identification (Yao et al. 2019).

An image retrieval system typically works by first representing each query image as a vector in a high-dimensional space and then finding the nearest neighbors of the query image in a database. It is not surprising that building discriminative image representations is crucial, and for this reason massive research has been devoted to improving the representations of image features, spanning from early hand-crafted features to nowadays deep learning-based ones (Dubey 2022). Nevertheless, like other deep learning-based applications, current image retrieval systems are black-box and data-driven, which renders them potentially unreliable due to limited training data and high model complexity.

*corresponding author: xiaowei.huang@liverpool.ac.uk
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

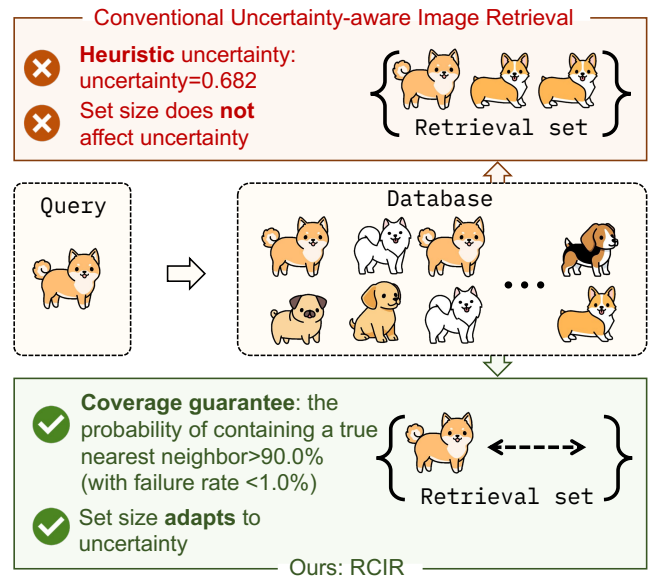


Figure 1: Illustration of our risk-controlled image retrieval RCIR and conventional uncertainty-aware image retrieval systems (e.g., (Warburg et al. 2021)): An uncertainty-aware image retrieval provides a heuristic uncertainty, and has a fixed retrieval set size. In contrast, RCIR provides a retrieval set that is guaranteed to cover a true nearest neighbor with a user-specified risk requirement (probability of containing, and failure rate), and moreover the retrieval size adapts to the uncertainty of the query sample and the risk requirement.

Meanwhile, reliability is essential for safety-critical applications ranging from autonomous driving to medical diagnosis, where it is vital to be aware of potential risks associated with the predictions. In the field of medical diagnosis, for instance, the user needs more than just a list of the best-matching candidates from disease databases. Instead, what they prefer or require is a cluster of retrieval results (or a retrieval set) with a specific level of risk, before concluding a diagnosis. The risk is defined as the probability of the retrieval set missing all true nearest neighbors of the query sample. See Sec. 3.3 for a formal definition. Conventional image retrieval systems are inadequate for these tasks as they

do not provide any measure for risk control.

Uncertainty estimation has emerged as a risk measure for image retrieval systems. It involves estimating the level of uncertainty for both query and database samples. Existing approaches such as (Warburg et al. 2021; Chang et al. 2020; Zhang, Wang, and Deng 2021) use high uncertainty as an (informal) signal to indicate the potentially inaccurate predictions. However, there are limitations. *First*, the estimated uncertainty is merely a heuristic measure (Angelopoulos et al. 2022a) rather than a guarantee. This leaves end users with only an *approximate* understanding about the prediction’s risk. *Second*, the estimated uncertainties do not consider the size of the retrieval set. That is, even if the retrieval set includes more candidates, which obviously increases its chance of containing the correct answer, the estimated uncertainty remains unaffected.

The above limitations motivate this work where we propose a novel image retrieval system called Risk Controlled Image Retrieval (RCIR). Fig. 1 highlights the advantages of RCIR over the conventional uncertainty-aware image retrieval systems, by adhering to a pre-specified risk requirement with provable guarantee for the uncertainty measure. We consider a risk requirement that involves two factors: a risk level α and an error rate δ , and require that, the probability that the assertion – the risk of the retrieval set not including at least one true nearest neighbor is less than or equal to α – occurs is greater than or equal to $1 - \delta$. Specifically, given such a risk requirement, RCIR can not only adapt the size of the retrieval set, but also ensure that the retrieval set satisfies the requirement, offering a provable statistical guarantee to the underlying image retrieval system.

Technically, RCIR employs two modules, a Retrieval Set Size Adapter (**Adapter**) and a Risk Controller (**Controller**). The **Controller** calculates a scale κ offline, which is then used to guide the **Adapter** to adjust the size of a retrieval set during inference. This collaboration enables the generation of a risk-controlled retrieval set.

We summarize the contributions of this paper as follows:

1. We for the first time utilize the common heuristic uncertainty estimation methods in a statistical manner to dynamically adjust the retrieval set size.
2. We for the first time allow an image retrieval system to generate retrieval sets that meet a user-specified risk requirement.
3. We demonstrate the effectiveness of RCIR through experimental results on four datasets: Stanford CAR-196, CUB-200, Pittsburgh, and ChestX-Det.
4. The source code of RCIR is available at <https://github.com/ramdrop/rcir>.

2 Related Work

2.1 Image Retrieval

Image retrieval often involves building a tagged database offline and searching it online, where images are represented by feature vectors. Therefore, image retrieval’s effectiveness relies on the feature extractor’s representation power. Initially, hand-crafted image features such as SIFT

(Lowe 1999) were primarily used. However, with the superior performance of deep learning-based features from pre-trained CNNs (Babenko et al. 2014), hand-crafted features have become obsolete. Recent research on image retrieval demonstrates that deep learning-based features can be effectively learned end-to-end with ranking loss functions (Hadsell, Chopra, and LeCun 2006; Schroff, Kalenichenko, and Philbin 2015; Sohn 2016). In this sense, image retrieval has evolved into a metric learning problem, with the aim of learning a mapping function that creates an embedding space where similar objects are positioned close together while dissimilar objects farther apart. Loss functions, including contrastive loss (Hadsell, Chopra, and LeCun 2006), triplet loss (Schroff, Kalenichenko, and Philbin 2015), and N-pairs loss (Sohn 2016), have been studied to enhance metric learning, which in turn benefits image retrieval.

2.2 Uncertainty Estimation

Deep learning has achieved tremendous success in numerous computer vision tasks, but its black-box working mechanism has raised concerns about its reliability. Uncertainty estimation is a way to quantify the confidence of the model in its prediction. Uncertainty of predictions can arise from either the data’s inherent uncertainty or the model’s uncertainty, known as aleatoric and epistemic uncertainty, respectively (Kendall and Gal 2017). To quantify epistemic uncertainty, variational inference based methods such as Monte Carlo Dropout are developed to approximate Bayesian Neural Networks, where the weights of the networks are modeled as distributions. Ensemble method (Fort, Hu, and Lakshminarayanan 2019) initiates multiple instances of the same model and then takes the variance of predictions as an uncertainty level. On the other hand, researchers typically estimate aleatoric uncertainty by learning it via an additional head that is parallel to the network.

Specific to the field of image retrieval, DUL (Chang et al. 2020) learns aleatoric uncertainty by constructing a stochastic embedding space where the uncertainty is regularized by a KL divergence loss. BTL (Warburg et al. 2021) utilizes a Bayesian loss function that enforces triplet constraints on stochastic embeddings. (Taha et al. 2019) shows MCD can be employed in image retrieval to quantify epistemic uncertainty. Current research on uncertainty estimation in image retrieval is primarily concerned with quantifying the likelihood of a pair of images being similar or dissimilar, albeit with only a heuristic estimate of uncertainty (Angelopoulos et al. 2022a). Our objective is to retrieve a set of images guaranteed to contain the true nearest neighbors of a given query sample with a user-specified risk requirement.

2.3 Conformal Prediction

The framework of conformal prediction (CP) provides a guarantee to the correctness of a prediction (Vovk, Gammerman, and Shafer 2005). Given a set of data points $(X_i, Y_i) \in \mathcal{Q} \times \mathcal{Y}$, for $i = 1, \dots, n$, on which a classifier f is trained, instead of returning a response $f(X_{n+1})$, the split conformal prediction, a popular CP method, returns a prediction set $\mathcal{C}(X_{n+1}) \subseteq \mathcal{Y}$ such that $P(Y_{n+1} \in \mathcal{C}(X_{n+1})) \geq 1 - \alpha$ for $\alpha \in [0, 1]$ an user-specified error level. It defines a score

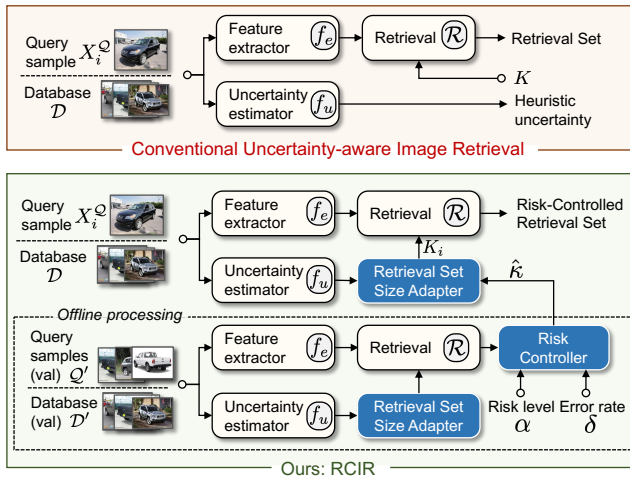


Figure 2: Diagrams of a conventional image retrieval system and our RCIR: Blue blocks (i.e., Retrieval Set Size Adapter and Risk Controller) highlight our contributions. Please see the text below for a detailed explanation.

function s , computes a quantile t for X_{n+1} over a set \mathcal{D}_c of calibration dataset, and then use t as a threshold to determine $\mathcal{C}(X_{n+1})$. The statistical guarantee depends on the exchangeability (a concept that is slightly relaxed from the i.i.d. assumption) in training and calibration datasets and query sample, as well as the hypothesis testing.

Beyond the initial study on classification, conformal quantile regression (Romano, Patterson, and Candes 2019) enhances CP to deliver prediction intervals with guaranteed accuracy. Building on this idea, DFRP (Bates et al. 2021) generates risk-controlled prediction set for image classification, and DFUQ (Angelopoulos et al. 2022a) proposes generating risk-guaranteed intervals in image regression. Different from these works on image classification or regression models, where the models are trained with one-to-one labels, our work focuses on image retrieval, where the model is trained using pairwise relations. Furthermore, the uncertainty estimation methods for regression tasks are distinct from those used in image retrieval since true labels are not available (Chang et al. 2020). These differences necessitate a thorough revisit to the formulation of the risk, as well as the formal proof on the guarantee, within this new paradigm.

3 Method

Fig. 2 presents the workflows of a conventional uncertainty-aware image retrieval system (Warburg et al. 2021), as well as our RCIR. The conventional uncertainty-aware image retrieval system first uses a feature extractor f_e to project the query and database samples into embeddings, and then applies a retrieval function \mathcal{R} to obtain the retrieval set, which forms of the K nearest neighbors of the query sample in the embedding space. Based on this, the uncertainty estimator f_u computes a heuristic uncertainty for each query sample. Our RCIR differs from the conventional one in that, the retrieval set is *risk-controlled*: during offline processing, given

the predefined risk requirements, i.e., the risk level α and the error rate δ , we first use a **Controller** to evaluate the risk of \mathcal{R} under various hypothesis on a scale κ in order to determine the best scale $\hat{\kappa}$. Then, during inference, given the i^{th} query sample X_i , an **Adapter** generates the retrieval set of size K_i , according to $\hat{\kappa}$, such that

$$P(\rho(\mathcal{R}) \leq \alpha) \geq 1 - \delta \quad (1)$$

where $\rho(\mathcal{R})$ is the risk of the retrieval set $\mathcal{R}(X)$ for a query set \mathcal{Q} where $X \in \mathcal{Q}$. The definitions of $\mathcal{R}(X)$ and $\rho(\mathcal{R})$ will be elicited below in Eqs. (3) and (7), respectively.

In the following, we first describe the conventional image retrieval pipeline, followed by the uncertainty-aware variant before moving into our risk controlled enhancement.

3.1 Image Retrieval

In image retrieval, a feature extractor f_e is trained to project high-dimensional images onto a lower-dimensional embedding space. In the embedding space, similar samples are positioned closely together, whereas dissimilar samples are positioned far apart. Assume that there is a dataset $\{\mathcal{Q}, \mathcal{D}\}$, where \mathcal{Q} denotes the query set and \mathcal{D} the database. Given the i^{th} query sample $X_i^{\mathcal{Q}} \in \mathcal{Q}$ (superscript denotes which set the sample belongs to), its most similar samples $\{Y_{i,j} | j = 1, 2, \dots, K\}$ are retrieved from \mathcal{D} as

$$\{Y_{i,j} | j = 1, 2, \dots, K\} = \mathcal{R}_{[K, f_e]}(X_i^{\mathcal{Q}}), \quad (2)$$

where $\mathcal{R}_{[K, f_e]}$ represents a retrieval function conditioned on K and f_e :

$$\begin{aligned} \mathcal{R}_{[K, f_e]}(X_i^{\mathcal{Q}}) = \{x | x \in \mathcal{P}, \mathcal{P} \subseteq \mathcal{D}, |\mathcal{P}| = K, \\ d[f_e(X_i^{\mathcal{Q}}), f_e(x)] \leq d[f_e(X_i^{\mathcal{Q}}), f_e(z)], \\ \forall z \in \mathcal{D} \setminus \mathcal{P}\}, \end{aligned} \quad (3)$$

where d denotes a distance metric function, and K denotes the number of the retrieved candidates, f_e denotes the feature extractor. The above image retrieval system has been the *de facto* standard in the image retrieval community. However, it only provides a list of best-matching candidates without indicating their reliability. The absence of reliability information can be problematic in situations where failures may have high-stakes consequences.

3.2 Uncertainty-aware Image Retrieval

Uncertainty estimation techniques mitigate this issue by providing a heuristic uncertainty for each query and database sample. Existing uncertainty-aware image retrieval systems include three methodologies: 1) estimate model uncertainty by BNNs (Gal and Ghahramani 2016), 2) estimate uncertainty by an ensemble of models, and 3) predict uncertainty by a deterministic model (Warburg et al. 2021). Since RCIR utilizes heuristic uncertainty, we briefly introduce them.

Uncertainty Estimated by a BNN Bayesian approaches treat the weights of a neural network as distributions instead of deterministic values. Since obtaining the analytical posterior distribution for the weights is intractable due

to the difficulty of acquiring evidence, researchers have utilized various methods to address this challenge, with variational inference being the most widely employed method, e.g., Monte Carlo Dropout (MCD) (Gal and Ghahramani 2016). MCD assumes a mixed Gaussian distribution on each weight, which makes the sampling of weights equivalent to applying dropout operations. In our image retrieval setting, we apply dropout to all conventional layers of the feature extractor with a dropout rate p . Feature μ and the heuristic uncertainty σ^2 are obtained by applying dropout at test time and T forward propagations:

$$\mu = \frac{1}{T} \sum_{t=1}^T f_{e_t \sim \theta}(X), \sigma^2 = \frac{1}{T} \sum_{t=1}^T (f_{e_t \sim \theta}(X) - \mu)^2. \quad (4)$$

where θ denotes the weight posterior distribution.

Uncertainty Estimated by an Ensemble Ensemble method (Fort, Hu, and Lakshminarayanan 2019) trains multiple instances of a same deterministic model, each with random initial weights. During inference, predictions of all instances are averaged, and the uncertainty is obtained as the variance of the predictions. Provided that there are N instances, the mean and variance of the predictions are:

$$\mu = \frac{1}{N} \sum_{i=1}^N f_{e_i}(X), \sigma^2 = \frac{1}{N} \sum_{i=1}^N (f_{e_i}(X) - \mu)^2. \quad (5)$$

Uncertainty Estimated by a Single Deterministic Model Multiple feed-forward propagations in MCD can cause overhead, and using multiple model instances can result in increased memory usage. In comparison, a single deterministic model is an attractive approach for estimating uncertainty. (Warburg et al. 2021) constructs a Gaussian distribution for each feature, i.e., $f_e(X) \sim \mathcal{N}(\mu, \sigma^2)$. A Bayesian triplet loss is introduced to enforce contrastive learning among these probabilistic features. The network is built upon a common image retrieval model by adding a variance head f_u parallel to the mean head f_e . Once trained, the model can output μ and σ^2 with one forward propagation:

$$[\mu, \sigma^2] = [f_e(X), f_u(X)]. \quad (6)$$

In summary, uncertainties of features are obtained by a BNN-based method, an ensemble or a single deterministic model. Nevertheless, the existing uncertainty-aware image retrieval methods are problematic in two senses:

1. They only provide a *heuristic* measure of uncertainty (Angelopoulos et al. 2022a) rather than a guarantee. That said, even with the estimated uncertainty, the retrieval set *cannot* be interpreted with a likelihood of encompassing the true nearest neighbors of a given query sample.
2. The retrieval set’s uncertainty is represented by the query sample’s uncertainty (Warburg et al. 2021), which does not consider the retrieval set size. As a result, the same uncertainty value is assigned to a retrieval set with one sample as to a retrieval set with ten samples. However, a larger retrieval set has a higher likelihood of capturing the true nearest neighbors compared to a smaller one.

We propose RCIR to overcome these limitations: RCIR takes as input the user-specified risk level α and error rate δ , and outputs a risk-controlled retrieval set guaranteed to contain at least one true nearest neighbor of the query sample with a probability $1 - \alpha$ and an error rate δ . We describe RCIR in what follows.

3.3 Risk Controlled Image Retrieval

Problem Setting We denote the risk of an image retrieval system \mathcal{R} by $\rho(\mathcal{R})$, defined as the expected value, over a set \mathcal{Q} of query samples, on the case that the retrieval set misses *all* true nearest neighbors:

$$\rho(\mathcal{R}) = \mathbb{E}_{X \in \mathcal{Q}}[\ell(\mathcal{R}(X), \mathcal{S}(X))], \quad (7)$$

where \mathcal{R} denotes a generic retrieval function as in Eq. (3), $\mathcal{S}(X)$ means retrieving all true nearest neighbors of X , and ℓ indicates if the retrieval set misses all true nearest neighbors:

$$\ell(\mathcal{R}(X), \mathcal{S}(X)) = \mathbb{1}(\mathcal{R}(X) \cap \mathcal{S}(X) = \emptyset). \quad (8)$$

The risk function $\rho(\cdot)$ evaluates the performance of a retrieval system \mathcal{R} on a query set \mathcal{Q} , and is bounded between 0 and 1: $\rho(\mathcal{R}) = 0$ occurs when the retrieval set $\mathcal{R}(X)$ covers at least one true nearest neighbor for all X in \mathcal{Q} , and $\rho(\mathcal{R}) = 1$ happens when the retrieval set $\mathcal{R}(X)$ never covers any true nearest neighbor for any X in \mathcal{Q} .

Based on the above definition, the goal of our RCIR is to show that the risk $\rho(\mathcal{R})$ can be provably guaranteed to be below a predefined level α with a probability of $1 - \delta$, in the form of Eq. (1). In RCIR, this is achieved by employing two components: an **Adapter** and a **Controller**, where the **Controller** determines a scale $\hat{\kappa}$ (Eq. (12)), whereas the **Adapter** uses the scale to adaptively determine the retrieval set (Eq. (9)). The guarantee is established formally in Theorem 1 (Eq. (13)).

Retrieval Set Size Adapter In practice, without compromising accuracy, end users prefer smaller retrieval sets for ‘easy’ query samples, compared to ‘hard’ ones, in order to save post-processing time. To achieve this, we propose adapting the retrieval set size based on the query’s *difficulty level*. While appealing, achieving this goal is not straightforward and has been barely explored in existing literature.

The first question we need to answer is: *what is the appropriate metric to measure a query’s difficulty level?* Our investigation shows that although existing heuristic uncertainties do not provide guarantees, they correlate with recall@1 (see Fig. 5). This discovery leads us to hypothesize that common heuristic uncertainty can serve as a coarse measure of a query’s difficulty level.

Moving on to the second question: *If such a metric exists, what would be the principled way of building the relation between this metric and the retrieval set size?* For this, we have identified a simple yet effective solution:

$$K_i = \lceil \kappa \cdot \Phi[f_u(X_i^{\mathcal{Q}})] \rceil, \quad (9)$$

where $\kappa \in \mathbb{R}^+$ denotes a scale, which will be computed by the **Controller** (details to explained below, with the best κ as $\hat{\kappa}$), f_u is a heuristic uncertainty estimator, Φ normalises the result into $[0, 1]$, and $\lceil \cdot \rceil$ rounds up to the nearest integer.

We denote the mapping function, Eq. (9), as the **Adapter**, which adjusts the retrieval set size based on the uncertainty of each query sample and a scale κ . Then, an image retrieval function with the **Adapter** is conditioned on κ , f_u and f_e :

$$\{Y_{i,j} | j = 1, 2, \dots, K_i\} = \mathcal{R}_{[\kappa, f_u, f_e]}(X_i^{\mathcal{Q}}), \quad (10)$$

and its risk is denoted as $\rho(\mathcal{R}_{[\kappa, f_u, f_e]})$. Since f_u and f_e are fixed once trained, $\rho(\mathcal{R}_{[\kappa, f_u, f_e]})$ only changes with respect to κ . Thereafter, we may also refer to the risk as $\rho(\kappa)$.

Risk Controller The **Adapter** alone is insufficient to provide provable guarantee on risk freeness for an image retrieval system. To complement it, we introduce **Controller**, which enhances the **Adapter** with risk awareness.

Given the given risk level is α and the error rate is δ , the **Controller** is responsible for determining the scale κ such that the risk of the retrieval system is below α with probability greater than $1 - \delta$. The general idea is as follows. Because $\rho(\kappa)$ is unknown, we consider instead its upper confidence bound (UCB) $\hat{\rho}^+(\kappa)$, which can be determined by applying Hoeffding's inequality (Eq. 16). Because the UCB $\hat{\rho}^+(\kappa)$ is greater than the risk $\rho(\kappa)$ with high probability (Eq. (11)), and the fact that the risk function is monotone nonincreasing (Lemma 1), we can prove in Theorem 1 that, as long as the scale $\hat{\kappa}$ is the smallest one such that $\hat{\rho}^+(\kappa) \leq \alpha$ for all $\kappa > \hat{\kappa}$, the guarantee holds on $\hat{\kappa}$ (Eq. (13)). It is not hard to see that Eq. (13) is an instantiation of Eq. (1).

Lemma 1. *Given a dataset $\{\mathcal{Q}, \mathcal{D}\}$, with a fixed feature extractor f_e and a fixed uncertainty estimator f_u , the risk function $\rho(\mathcal{R}_{[\kappa, f_u, f_e]})$ is a monotone nonincreasing function with respect to the scale $\kappa \in \mathbb{R}^+$.*

Proof. For any i^{th} query $X_i^{\mathcal{Q}}$, if $\kappa_1 > \kappa_2$, then $K_{i, \kappa_1} \geq K_{i, \kappa_2}$ by the mapping function Eq. (9) as f_u is fixed. This means $|\mathcal{R}_{[\kappa_1, f_u, f_e]}(X_i^{\mathcal{Q}})| \geq |\mathcal{R}_{[\kappa_2, f_u, f_e]}(X_i^{\mathcal{Q}})|$, according to Eq. (2). Note that, from Eq. (3), the risk function is monotonic with respect to K , the size of the retrieval set. Then, because in a retrieval system, the candidates are retrieved based on a consistent distance metric d (see Eq. (3)), we have

$$\mathcal{R}_{[\kappa_2, f_u, f_e]}(X_i^{\mathcal{Q}}) \subseteq \mathcal{R}_{[\kappa_1, f_u, f_e]}(X_i^{\mathcal{Q}}).$$

One step further, according to the loss definition in Eq. (8),

$$\ell(\mathcal{R}_{[\kappa_2, f_u, f_e]}(X_i^{\mathcal{Q}}), \mathcal{S}(X_i^{\mathcal{Q}})) \geq \ell(\mathcal{R}_{[\kappa_1, f_u, f_e]}(X_i^{\mathcal{Q}}), \mathcal{S}(X_i^{\mathcal{Q}})).$$

Finally, by the risk definition in Eq. (7), we have

$$\rho(\mathcal{R}_{[\kappa_2, f_u, f_e]}) \geq \rho(\mathcal{R}_{[\kappa_1, f_u, f_e]}),$$

which concludes the proof. \square

Theorem 1. *Assume the upper confidence bound of the risk, which we denote by $\hat{\rho}^+(\kappa)$, is accessible, and $\forall \kappa \in \mathbb{R}^+$,*

$$P(\rho(\kappa) \leq \hat{\rho}^+(\kappa)) \geq 1 - \delta. \quad (11)$$

Let $\hat{\kappa}$ denotes the smallest κ such that for any $\kappa > \hat{\kappa}$ we have $\hat{\rho}^+(\kappa) \leq \alpha$, i.e.,

$$\hat{\kappa} = \inf\{\kappa : \hat{\rho}^+(\kappa') < \alpha, \forall \kappa' > \kappa\}, \quad (12)$$

Then, we have that

$$P(\rho(\hat{\kappa}) \leq \alpha) \geq 1 - \delta. \quad (13)$$




Figure 3: The $\rho - \kappa$ curve on the CAR-196 calibration set: the risk $\rho(\kappa)$ is monotone nonincreasing with κ .

Algorithm 1: Compute $\hat{\kappa}$.

Input: Calibration set $\{\mathcal{Q}', \mathcal{D}'\}$, risk level α , error rate δ , trained feature extractor f_e , trained heuristic uncertainty estimator f_u , step $\Delta\kappa$

Output: $\hat{\kappa}$

```

1:  $\kappa \leftarrow 1, \hat{\rho}^+(\kappa) \leftarrow 1$ 
2: while  $\hat{\rho}^+(\kappa) > \alpha$  do
3:    $\kappa = \kappa + \Delta\kappa$ 
4:   for  $i$  in  $1, 2, \dots, |\mathcal{Q}'|$  do
5:      $\tilde{K}_i = \left\lceil \kappa \cdot \Phi[f_u(X_i^{\mathcal{Q}'})] \right\rceil$ 
6:      $\ell_i = \mathbb{1}[\mathcal{R}_{[\tilde{K}_i, f_e]}(X_i^{\mathcal{Q}'}) \cap \mathcal{S}(X_i^{\mathcal{Q}'}) = \emptyset]$ 
7:   end for
8:    $\hat{\rho}(\kappa) = \frac{1}{|\mathcal{Q}'|} \sum_i |\mathcal{Q}'| \ell_i$ 
9:    $\hat{\rho}^+(\kappa) = \hat{\rho}(\kappa) + \sqrt{\frac{1}{2n} \log \frac{1}{\delta}}$ 
10: end while
11:  $\hat{\kappa} = \kappa - \Delta\kappa$ 
12: return  $\hat{\kappa}$ 

```

Proof. Let κ^* be the smallest value of κ that ensures the risk is less than α , as given by the following equation:

$$\kappa^* = \inf\{\kappa : \rho(\kappa) \leq \alpha\}, \quad (14)$$

Then, due to the continuity and nonincreasing (Lemma 1) of $\rho(\kappa)$, it follows that $\rho(\kappa^*) = \alpha$. To make it easier to understand, we visualize in Fig. 3 the risks $\rho(\kappa^*)$, $\rho(\hat{\kappa})$ and α using the CAR-196 calibration set. Although it is an specific example, it does not lose the generality.

We define two events: the event A , which occurs when $\rho(\hat{\kappa}) > \alpha$, and the event B , which occurs when $\rho(\kappa^*) > \hat{\rho}^+(\kappa^*)$. We show that A is a sufficient condition for B . From $\rho(\hat{\kappa}) > \alpha$, we have that $\hat{\kappa} < \kappa^*$ (due to the monotonicity of $\rho(\kappa)$), which implies $\hat{\rho}^+(\kappa^*) < \alpha$ (according to Eq. (12)). Now, because $\rho(\kappa^*) = \alpha$ and $\hat{\rho}^+(\kappa^*) < \alpha$, we have that $\rho(\kappa^*) > \hat{\rho}^+(\kappa^*)$.

By the fact that A is a sufficient condition for B , we have that $P(A) \leq P(B)$. According to Eq. (11), we have $P(B) < \delta$. The combination of $P(A) \leq P(B)$ and $P(B) < \delta$ leads to $P(A) < \delta$. Consequently, $P(\bar{A}) \geq 1 - \delta$, i.e., $P(\rho(\hat{\kappa}) \leq \alpha) \geq 1 - \delta$. \square

With Eq. (12), we can proceed to determine an upper confidence bound $\hat{\rho}^+(\kappa)$. According to (Bates et al. 2021), Hoeffding's inequality applies to our risk function $\rho(\kappa)$, which is bounded by 1. We denote the risk on the calibration set by $\hat{\rho}(\kappa)$. Hoeffding's inequality indicates

$$P(\hat{\rho}(\kappa) - \rho(\kappa) \leq -x) \leq e^{-2nx^2}, \quad (15)$$

which implies an upper confidence bound

$$\hat{\rho}^+(\kappa) = \hat{\rho}(\kappa) + \sqrt{\frac{1}{2n} \log\left(\frac{1}{\delta}\right)}. \quad (16)$$

The **Controller** calculates κ by following the steps outlined by Alg. 1: Given a calibration set, α , δ , f_e , f_u and $\Delta\kappa$, the **Controller** computes $\hat{\kappa}$ by iteratively increasing κ until the upper confidence $\hat{\rho}^+(\kappa)$ bound is below α .

With all key steps explained, we summarize the pipeline of RCIR as follows: Given a predefined risk requirement, i.e., risk level α and error rate δ , the **Controller** uses Alg. 1 to compute $\hat{\kappa}$. Once $\hat{\kappa}$ is obtained, the **Adapter** generates the retrieval set size K_i , and the image retrieval function retrieves the top K_i candidates. The set of these candidates is the risk-controlled retrieval set as it is ensured to have at least one true nearest neighbor with a probability of $1 - \alpha$, and that the chance of the assert fail is limited to δ .

4 Experiments

4.1 Datasets

CUB-200 (Wah et al. 2011) contains 11,788 images of 200 classes. Each class has at least 50 images. Following (Warburg et al. 2021), we use the first 100 classes as a training set and the other 100 as a test set.

Stanford CAR-196 (Krause et al. 2013) contains 16,185 images of 196 classes. Each class has at least 80 images. Following (Warburg et al. 2021)(Musgrave, Belongie, and Lim 2020), we use the first 98 classes as the training set and the other 98 classes as the test set.

Pittsburgh (Torii et al. 2013) is a large image database from Google Street View. Following the split of NetVLAD (Arandjelovic et al. 2016), we adopt the subset that consists of 10k samples in the training, validation and test split.

ChestX-Det (Lian et al. 2021) is a subset of the public dataset NIH ChestX-ray14, and contains 3543 images with 14 classes. We choose samples of the six classes as the training set, and the rest as the test set.

For all datasets, we follow (Angelopoulos et al. 2022b) and randomly select 50% images from the test set to form the calibration set.

4.2 Implementations

RCIR is versatile and can be applied to any image retrieval system that provides heuristic uncertainty. To show the effectiveness of RCIR, we first build three conventional image retrieval systems, **BTL**, **MCD** and **Deep Ensemble**, based on the model used in (Warburg et al. 2021). Each system has a feature extractor and an uncertainty estimator as described in Sec. 3.2. BTL, MCD and Deep Ensemble retrieve fixed K database samples when given a query. For comparing purposes, we also include a **Deterministic** image retrieval system that only has a feature extractor.

All models are trained using the Adam optimizer with an initial learning rate of 10^{-5} and an exponential learning rate scheduler with gamma 0.99. The weight decay is set to 10^{-3} . The feature dimension is 2048 and the variance dimension is 1. We adopt a hard mining strategy, feeding




Figure 4: The Recall@1 results of different methods on various test sets. The error bars represent the standard deviation of the results of 10 trials.

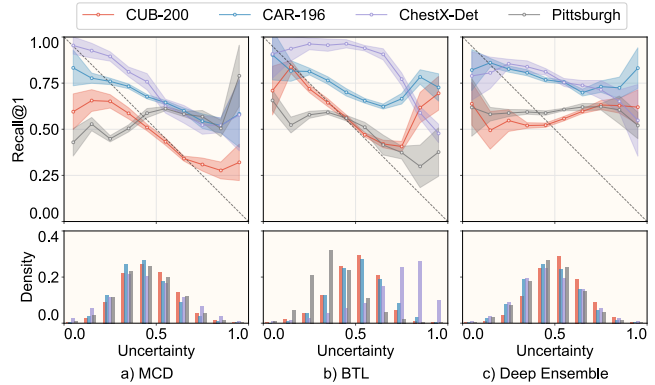


Figure 5: The reliability diagrams on different test sets. The dashed line denotes the ideal calibration line. (The shadows represent the standard deviation of the results of 10 trials, and the same goes for the other figures)

models with triplets that violate the triplet margin (Warburg et al. 2021). More information about the network and training can be found in the supplementary materials.

4.3 Metrics

We first follow the common evaluation steps (Warburg et al. 2021) for image retrieval systems, adopting **recall@1** to evaluate the retrieval performance of the feature extractor and **reliability diagram** (Guo et al. 2017) to assess the uncertainty quality of the uncertainty estimator. Then, we use **empirical risk**, which is the risk calculated on the test set, to evaluate the risk control performance. In addition, **retrieval set size** is used to verify that RCIR does not resort to large retrieval sets to lower the empirical risk. The retrieval set size should be minimized while the risk is controlled. Because large retrieval sets tend to include more negative samples and are inefficient for downstream processing. We also provide **qualitative visualization** to show more insights into the image retrieval system. Please see the supplementary materials for more information on these metrics.

4.4 Results

Recall@1 (Image Retrieval Performance) We begin by evaluating the retrieval performance of various methods. The Recall@1 of different image retrieval systems on various test sets are presented in Fig. 4. BTL achieves a higher recall@1 than Deterministic on most datasets, suggesting that a feature extractor can benefit from uncertainty-aware training. Nevertheless, MCD’s performance is inconsistent

when compared to Deterministic. The inconsistency in performance may stem from the varying impact of dropout layers on the representation power, depending on the dataset’s scale. Deep Ensemble achieves the best performance, which can be attributed to the fact that it utilizes multiple models to average out noises.

Reliability Diagram (Uncertainty Estimation Error)

Fig. 5 depicts the reliability diagrams of the uncertainty estimator of different image retrieval systems across datasets. It can be seen that none of these curves conform to the ideal calibration line, indicating that the heuristic uncertainty itself is not well-calibrated. Moreover, the gaps between the curves and the dashed line vary across datasets, implying that existing uncertainty estimation methods cannot consistently perform across different datasets.

Next, we build RCIR by plugging our **Adapter** and **Controller** modules into conventional image retrieval systems. We denote the three resulting RCIR systems as **BTL^R**, **MCD^R** and **Deep Ensemble^R**. Meanwhile, to evaluate the practical usefulness of heuristic uncertainty, we normalize heuristic uncertainties to the range $[0, 1]$ based on their statistics on the calibration set. Then each query-candidate pair’s uncertainty is calculated as the sum of their individual uncertainties (Warburg et al. 2021). This will result in three comparing methods: **BTL[⊕]**, **MCD[⊕]** and **Deep Ensemble[⊕]**. Fig. 6 depicts the data flow of the aforementioned methods.

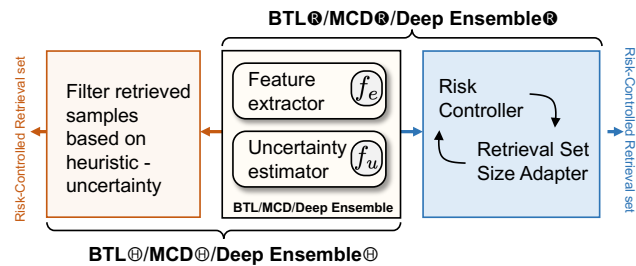


Figure 6: The data flow of our methods, **BTL^R**, **MCD^R** and **Deep Ensemble^R**, and the comparing methods, **BTL[⊕]**, **MCD[⊕]** and **Deep Ensemble[⊕]**.

For fairness, \times^{\oplus} will retrieve as many candidates as \times^{R} can achieve, then only pairs with uncertainties below $1 - \alpha$ will be kept. We adopt a step $\Delta\kappa = 0.2$ for \times^{R} .

Empirical Risk (Risk Control Effectiveness)

Fig. 7 shows the empirical risk of different image retrieval systems on different datasets. It is shown that the risks of \times^{R} are always below the predefined risk levels α . This means that before a retrieval we can say for sure that the retrievals of \times^{R} have a $1 - \alpha$ probability of covering the true nearest neighbors, with only a δ probability of being wrong. In comparison, the heuristic uncertainty-only methods, including **MCD[⊕]**, **BTL[⊕]** and **Deep Ensemble[⊕]**, exhibit risks higher than the predefined risk levels in the range $\alpha < 0.4$. The results show that the heuristic uncertainty alone is not reliable for risk control, while RCIR effectively controls the risk.




Figure 7: The risks on different test sets: risk is effectively controlled by RCIR (see \times^{R}) under different risk levels α (with $\delta = 0.1$).




Figure 8: The average K across different methods on different test sets (*solid line* — means \times^{R} (with $\delta = 0.1$), while *dashed line* — — means their conventional counterparts): When achieving the same recall, \times^{R} and their fixed-size counterparts have similar average retrieval set sizes, indicating that \times^{R} does not rely on simply increasing K to control the risk.

In addition, we investigate the effect of using different δ . Fig. 10 shows the empirical risk with different δ on the Pittsburgh dataset. With a smaller δ , RCIR would be more conservative (i.e., larger retrieval size) to ensure the risk is below the given α . (see supplementary for similar trends on the other datasets.)

Retrieval Set Size RCIR may always resort to large retrieval sets to lower the empirical risk, in which situation the risk is trivially controlled but the retrieval sets would be of less practical use. To examine this, we compare **MCD^R**, **BTL^R** and **Deep Ensemble^R** against their fixed-size counterparts, **MCD**, **BTL** and **Deep Ensemble**. Fig. 8 shows that when achieving the same recall@ K , \times^{R} and their fixed-size counterparts have similar average retrieval set sizes. This indicates that \times^{R} does not rely on simply increasing K to control the risk. But we also notice that **MCD^R** and **Deep Ensemble^R** on the Pittsburgh test set have a slightly larger average retrieval set size than their fixed-size counterparts. This is due to poor uncertainty estimation on the Pittsburgh test set, as shown in Fig. 5.

The impact of uncertainty estimator To study the impact of the uncertainty estimator on risk control performance, we modify the **MCD** image retrieval system by replacing the uncertainty estimator f_u with a random uncertainty estimator. This random uncertainty estimator assigns a randomly sam-




Figure 9: The reliability diagram, risk and average K of different models and pipelines on the CAR-196 dataset. Even with a poorly calibrated uncertainty (i.e., random guess), our $\times \textcircled{R}$ pipeline is still valid in controlling the risk. We also note that a poorly calibrated uncertainty results in a larger average retrieval set size.

pled value from the uniform distribution $[0, 1]$ as the uncertainty for each query. Such uncertainty is poor as it does not correlate with the recognition performance. We refer to this modified image retrieval pipeline as **Random** and the corresponding RCIR system as **Random \textcircled{R}** . Fig. 9a) presents the reliability diagram of the uncertainty estimator of MCD and Random, where the curve of Random diverges more from the ideal calibration line (dashed line), indicating a poorly calibrated uncertainty. Fig. 9b) shows the empirical risk of MCD \textcircled{R} and Random \textcircled{R} , while Fig. 9c) displays the average K of different pipelines. It can be seen that Random \textcircled{R} is still valid in controlling the risk, even with a poorly calibrated uncertainty. This is because the validity of the Theorem 1 is not dependent on the uncertainty estimator. However, using a poorly calibrated uncertainty estimator leads to a larger retrieval set size due to suboptimal computation of κ . This observation is consistent with the observation found by (Angelopoulos and Bates 2021).

Qualitative visualization The distribution of retrieval size on the Pittsburgh test set is presented in Fig. 11 (see the supplementary materials for similar results on the other datasets). It is clear that the retrieval size varies with the risk level α : a smaller α results in a larger retrieval size, and vice versa. This helps end users save time on easy queries and focus on more difficult ones.




Figure 10: The risks on the Pittsburgh test set with different error rates by BTL \textcircled{R} : smaller δ leads to more conservative retrievals (e.g., larger retrieval size).

Fig. 12 shows the qualitative visualization of retrievals on the CAR-196 dataset by different methods. In the first row, it is evident that when a relatively easy query is provided, all methods successfully retrieve the correct results. Moving on to the second row, where the query is more challenging,




Figure 11: The retrieval set size, K , on the Pittsburgh test set by BTL \textcircled{R} : retrieval set size adapts to the risk level α ($\delta = 0.1$).

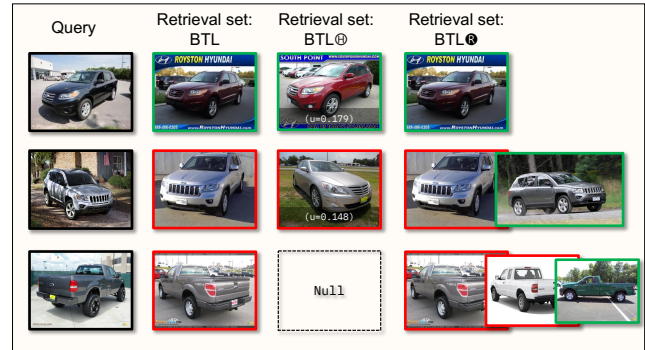


Figure 12: The qualitative visualization of retrieval sets on the CAR-196 test set by different methods (Red border \square denotes a wrong candidate, while green border \square correct): BTL \textcircled{R} adjusts the size of retrievals to reflect the uncertainty while ensuring coverage ($\alpha = 0.2$ and $\delta = 0.1$).

the fixed-size retrieval method fails without any indication, and BTL \textcircled{R} fails as well with a low heuristic number. However, BTL \textcircled{R} takes into account the difficulty of the query and retrieves an additional candidate, resulting in the correct answer. The third row represents a more demanding query. BTL \textcircled{R} generates no candidate, while BTL \textcircled{R} retrieves more candidates to meet the coverage requirement.

5 Conclusion

This paper proposes RCIR as a risk-controlled improvement to the current image retrieval systems, for the scenarios when the reliability of image retrieval systems is essential. RCIR offers a coverage guarantee, regardless of data distribution or model selection. Nevertheless, RCIR still has room for improvement. For example, while the choice of heuristic uncertainty estimator does not affect the validity of RCIR, a poorly calibrated uncertainty estimator may increase the average retrieval set size. It would be interesting to figure out a method that can not only quantify but also reduce the risks.

Acknowledgements

Xiaowei Huang’s contribution is partially supported by the UK EPSRC through the End-to-End Conceptual Guarding of Neural Architectures [EP/T026995/1]. Xingyu Zhao’s contribution is supported by the UK EPSRC New Investigator Award through Harnessing Synthetic Data Fidelity for Assured Perception of Autonomous Vehicles. Special

thanks to <https://instagram.com/littlemandyart> for the cute dog sticker pack design that is featured in Fig. 1.

References

- Angelopoulos, A. N.; and Bates, S. 2021. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. *arXiv*, abs/2107.07511.
- Angelopoulos, A. N.; Kohli, A. P.; Bates, S.; Jordan, M.; Malik, J.; Alshaabi, T.; Upadhyayula, S.; and Romano, Y. 2022a. Image-to-image regression with distribution-free uncertainty quantification and applications in imaging. In *International Conference on Machine Learning*, 717–730. PMLR.
- Angelopoulos, A. N.; Kohli, A. P. S.; Bates, S.; Jordan, M. I.; Malik, J.; Alshaabi, T.; Upadhyayula, S.; and Romano, Y. 2022b. Image-to-Image Regression with Distribution-Free Uncertainty Quantification and Applications in Imaging. In *International Conference on Machine Learning (ICML)*, 717–730.
- Arandjelovic, R.; Gronat, P.; Torii, A.; Pajdla, T.; and Sivic, J. 2016. NetVLAD: CNN architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5297–5307.
- Babenko, A.; Slesarev, A.; Chigorin, A.; and Lempitsky, V. 2014. Neural codes for image retrieval. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*, 584–599. Springer.
- Bates, S.; Angelopoulos, A.; Lei, L.; Malik, J.; and Jordan, M. 2021. Distribution-free, risk-controlling prediction sets. *Journal of the ACM (JACM)*, 68(6): 1–34.
- Chang, J.; Lan, Z.; Cheng, C.; and Wei, Y. 2020. Data uncertainty learning in face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5710–5719.
- Cheng, Y.; and Wang, H. 2019. A modified contrastive loss method for face recognition. *Pattern Recognition Letters*, 125: 785–790.
- Dubey, S. R. 2022. A Decade Survey of Content Based Image Retrieval Using Deep Learning. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(5): 2687–2704.
- Fort, S.; Hu, H.; and Lakshminarayanan, B. 2019. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*.
- Gal, Y.; and Ghahramani, Z. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, 1050–1059. PMLR.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On Calibration of Modern Neural Networks. In *International Conference on Machine Learning (ICML)*, 1321–1330.
- Hadsell, R.; Chopra, S.; and LeCun, Y. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, 1735–1742. IEEE.
- Kendall, A.; and Gal, Y. 2017. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, 554–561.
- Lian, J.; Liu, J.; Zhang, S.; Gao, K.; Liu, X.; Zhang, D.; and Yu, Y. 2021. A structure-aware relation network for thoracic diseases detection and segmentation. *IEEE Transactions on Medical Imaging*, 40(8): 2042–2052.
- Lowe, D. G. 1999. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, 1150–1157. Ieee.
- Musgrave, K.; Belongie, S.; and Lim, S.-N. 2020. A metric learning reality check. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, 681–699. Springer.
- Romano, Y.; Patterson, E.; and Candes, E. 2019. Formalized quantile regression. *Advances in neural information processing systems*, 32.
- Schroff, F.; Kalenichenko, D.; and Philbin, w. c. d. t. r. . h. o. o. w. a. p. o. a. m. d. m., JaUsing (11). 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 815–823.
- Sohn, K. 2016. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems*, 29.
- Taha, A.; Chen, Y.-T.; Yang, X.; Misu, T.; and Davis, L. 2019. Exploring uncertainty in conditional multi-modal retrieval systems. *arXiv preprint arXiv:1901.07702*.
- Torii, A.; Sivic, J.; Pajdla, T.; and Okutomi, M. 2013. Visual place recognition with repetitive structures. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 883–890.
- Vovk, V.; Gammerman, A.; and Shafer, G. 2005. *Algorithmic learning in a random world*, volume 29. Springer.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset.
- Warburg, F.; Jørgensen, M.; Civera, J.; and Hauberg, S. 2021. Bayesian triplet loss: Uncertainty quantification in image retrieval. In *Proceedings of the IEEE/CVF International conference on Computer Vision*, 12158–12168.
- Yao, H.; Zhang, S.; Hong, R.; Zhang, Y.; Xu, C.; and Tian, Q. 2019. Deep representation learning with part loss for person re-identification. *IEEE Transactions on Image Processing*, 28(6): 2860–2871.
- Zhang, Y.; Wang, C.; and Deng, W. 2021. Relative uncertainty learning for facial expression recognition. *Advances in Neural Information Processing Systems*, 34: 17616–17627.