

Hierarchical Average Precision Training for Pertinent Image Retrieval

Elias Ramzi^{1,2} , Nicolas Audebert¹ , Nicolas Thome^{1,3} , Clément Rambour¹ 
and Xavier Bitot²

¹ CEDRIC, Conservatoire National des Arts et Métiers, Paris, France
{elias.ramzi,nicolas.audebert,nicolas.thome,clement.rambour}@cnam.fr

² Coexya, Paris, France
xavier.bitot@coexya.eu

³ Sorbonne Université, CNRS, ISIR, F-75005 Paris, France

Abstract. Image Retrieval is commonly evaluated with Average Precision (AP) or Recall@k. Yet, those metrics, are limited to binary labels and do not take into account errors’ severity. This paper introduces a new hierarchical AP training method for pertinent image retrieval (HAPPIER). HAPPIER is based on a new \mathcal{H} -AP metric, which leverages a concept hierarchy to refine AP by integrating errors’ importance and better evaluate rankings. To train deep models with \mathcal{H} -AP, we carefully study the problem’s structure and design a smooth lower bound surrogate combined with a clustering loss that ensures consistent ordering. Extensive experiments on 6 datasets show that HAPPIER significantly outperforms state-of-the-art methods for hierarchical retrieval, while being on par with the latest approaches when evaluating fine-grained ranking performances. Finally, we show that HAPPIER leads to better organization of the embedding space, and prevents most severe failure cases of non-hierarchical methods. Our code is publicly available at <https://github.com/elias-ramzi/HAPPIER>.

Keywords: Hierarchical Image Retrieval, Hierarchical Average Precision, Ranking

1 Introduction

Image Retrieval (IR) consists in ranking images with respect to a query by decreasing order of visual similarity. IR methods are commonly evaluated using Recall@k (R@k) or Average Precision (AP). Because those metrics are non-differentiable, a rich literature exists on finding adequate surrogate loss functions to optimize them with deep learning, with tuple-wise losses [26,32,41,39,40], proxy based losses [42,38,11,35] and direct AP optimization methods [6,29,23,31,2,28].

These metrics are only defined for binary (\oplus/\ominus) labels, which we denote as *fine-grained labels*: an image is negative as soon as it is not strictly similar to the query. Binary metrics are by design unable to take into account the severity of the mistakes in a ranking. On Fig. 1, some negative instances are “less negative” than others, *e.g.* given the “Brown Bear” query, “Polar bear” is more relevant than

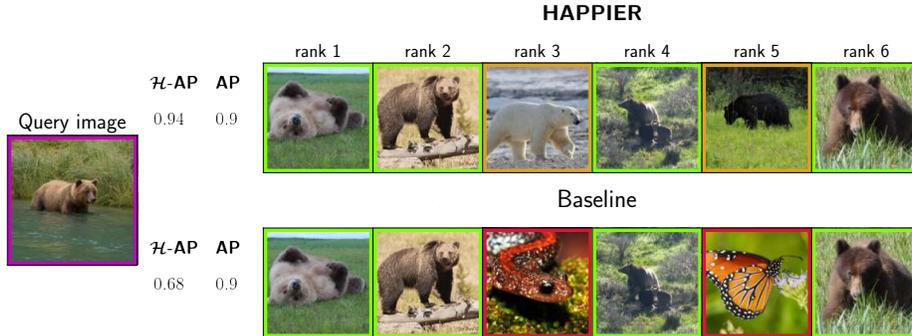


Fig. 1: **Proposed HAPPIER framework for pertinent image retrieval.** Standard ranking metrics based on binary labels, *e.g.* Average Precision (AP), assign the same score to the bottom and top row rankings (0.9). We introduce the \mathcal{H} -AP metric based on non-binary labels, that takes into account mistakes’ severity. \mathcal{H} -AP assigns a smaller score to the bottom row (0.68) than the top one (0.94). HAPPIER maximizes \mathcal{H} -AP during training and thus explicitly supports to learn rankings similar to the top one, in contrast to binary ranking losses.

“Butterfly”. However, AP is 0.9 for both the top and bottom rankings. Consequently, training on binary metrics (*e.g.* AP or R@k) develops no incentive to produce ranking such as the top row, and often produces rankings similar to the bottom one. To address this problem, we introduce the HAPPIER method dedicated to Hierarchical Average Precision training for Pertinent Image Retrieval. HAPPIER provides a smooth training objective, amenable to gradient descent, which explicitly takes into account the severity of mistakes when evaluating rankings.

Our first contribution is to define a new Hierarchical AP metric (\mathcal{H} -AP) that leverages the hierarchical tree between concepts and enables a fine weighting between errors in rankings. As shown in Fig. 1, \mathcal{H} -AP assigns a larger score (0.94) to the top ranking than to the bottom one (0.68). We show that \mathcal{H} -AP provides a consistent generalization of AP for the non-binary setting. We also introduce our HAPPIER_F variant, giving more weights to fine-grained levels of the hierarchy. Since \mathcal{H} -AP, like AP, is a non-differentiable metric, our second contribution is to use HAPPIER to directly optimize \mathcal{H} -AP by gradient descent. We carefully design a smooth surrogate loss for \mathcal{H} -AP that has strong theoretical guarantees and is an upper bound of the true loss. We then define an additional clustering loss to support having a consistency between partial and global rankings.

We validate HAPPIER on six IR datasets, including three standard datasets (Stanford Online Products [22] and iNaturalist-base/full [37]), and three recent hierarchical datasets (DyML [33]). We show that, when evaluating on hierarchical metrics (*e.g.* \mathcal{H} -AP), HAPPIER outperforms state-of-the-art methods for fine-grained ranking [41,42,35,28], the baselines and the latest hierarchical method of [33], and only slightly under-performs *vs.* state-of-the-art IR methods at the fine-grained level (*e.g.* AP, R@1). HAPPIER_F performs on par on fine-grained metrics while still outperforming fine-grained methods on hierarchical metrics.

2 Related work

2.1 Image Retrieval and ranking

The Image Retrieval community has designed several families of methods to optimize metrics such as AP and R@k. Methods that relies on tuple-wise losses, like pair losses [15,27], triplet losses [41], or larger tuples [32,20,39] learn comparison relations between instances. Methods using proxies have been introduced to lower the computational complexity of tuple based training [21,42,38,11,35]: they learn jointly a deep model and weight matrix that represent proxies using a cross-entropy based loss. Proxies are approximations of the original data points that should belong to their neighbourhood. Finally, there also has been large amounts of work dedicated to the direct optimization of the AP during training by introducing differentiable surrogates [6,29,23,31,2,28], so that models are optimized on the same metric they are evaluated on. However, nearly all of these methods only consider binary labels: two instances are either the same (positive) or different (negative), leading to poor performance when multiple levels of hierarchy are considered.

2.2 Hierarchical predictions and metrics

There has been a recent regain of interest in Hierarchical Classification [12,1,8] with the introductions of methods based either on a hierarchical softmax function or on multiple classifiers. It is considered that learning from hierarchical relations between labels leads to more robust models that make “better mistakes” [1]. Yet, hierarchical classification means that labels are known in advance and are identical in the train and test sets. This is called a *closed set* setting. However, Hierarchical Image Retrieval does not fall into this framework. Standard IR protocols consider the *open set* paradigm to better evaluate the generalization abilities of learned models: the retrieval task at test time pertains to labels that were not present in the train set, making classification poorly suited to IR.

Meanwhile, the broader Information Retrieval community has been using datasets where documents can be more or less relevant depending on the query and the user making the request [16,19]. Instead of the mere positive/negative dichotomy, each instance has a continuous score quantifying its relevance to the query. To quantify the quality of their retrieval engine, Information Retrieval researchers have long used ranking based metrics, such as the NDCG [17,10], that penalize mistakes differently based on whether they occur at the top or the bottom of the ranking and whether wrong documents still have some marginal relevance or not. Average Precision is also used as a retrieval metric [18] and has even been given probabilistic interpretations based on how users interact with the system [13]. Several works have investigated how to optimize those metrics during the training of neural networks, *e.g.* using pairwise losses [4] and later using smooth surrogates of the NDCG in LambdaRank [5], SoftRank [34], ApproxNDCG [25] and Learning-To-Rank [3]. These works however focused on NDCG, the most popular metric for information retrieval, and are without any theoretical guarantees: the surrogates

are approximations of the NDCG but not *lower bounds*, *i.e.* their maximization does not imply improved performances during inference.

An additional drawback of this literature is that NDCG does not relate easily to average precision [14], which is the most common metric in image retrieval. Fortunately, there have been some works done to extend AP in a graded setting where relevance between instances is not binary [30,13]. The graded Average Precision from [30] is the closest to our goal as it leverages SoftRank for direct optimization on non-binary relevance judgements, although there are significant shortcomings. There is no guarantee that the SoftRank surrogate actually minimizes the graded AP, it requires to annotate datasets with pairwise relevances which is unpractical for large scale settings and was only applied to small-scale corpora of a few thousands documents, compared to the hundred thousands of images in IR.

Recently, the authors of [33] introduced three new hierarchical benchmarks datasets for image retrieval, in addition to a novel hierarchical loss CSL. CSL extends proxy-based triplet losses to the hierarchical setting and tries to structure the embedding space in a hierarchical manner. However, this method faces the same limitation as the usual triplet losses: minimizing CSL does not explicitly optimize a well-behaved hierarchical evaluation metric, *e.g.* \mathcal{H} -AP. We show experimentally that our method HAPPIER significantly outperforms CSL [33] both on hierarchical metrics and AP-level evaluations.

3 HAPPIER Model

We detail HAPPIER our Hierarchical Average Precision training method for Pertinent Image Retrieval. We first introduce the Hierarchical Average Precision, \mathcal{H} -AP in Sec. 3.1, that leverages a hierarchical tree (Fig. 2a) of labels. It is based on the hierarchical rank, \mathcal{H} -rank, and evaluates rankings so that more relevant instances are ranked before less relevant ones (Fig. 2b). We then show how to directly optimize \mathcal{H} -AP by stochastic gradient descent (SGD) using HAPPIER in Sec. 3.2. Our training objective combines a carefully designed smooth upper bound surrogate loss for $\mathcal{L}_{\mathcal{H}\text{-AP}} = 1 - \mathcal{H}\text{-AP}$ and a clustering loss $\mathcal{L}_{\text{clust.}}$ that supports consistent rankings.

Context Let us consider a retrieval set $\Omega = \{\mathbf{x}_j\}_{j \in [1;N]}$ composed of N instances. For a query¹ $\mathbf{q} \in \Omega$, we aim to order all $x_j \in \Omega$ so that more relevant (*i.e.* similar) instances are ranked before less relevant instances.

In our hierarchical setting, the relevance of an instance \mathbf{x}_j is non-binary. We assume that we have access to a hierarchical tree defining semantic similarities between concepts as in Fig. 2a. For a query \mathbf{q} , we leverage this knowledge to partition the set of retrieved instances into $L+1$ disjoint subsets $\{\Omega^{(l)}\}_{l \in [0;L]}$. $\Omega^{(L)}$ is the subset of the most similar instances to the query (*i.e.* fine-grained level): for $L=3$ and a “Lada #2” query, $\Omega^{(3)}$ are the images of the same “Lada #2” (green), see Fig. 2. The set $\Omega^{(l)}$ for $l < L$ contains instances with smaller

¹ For the sake of readability, our notations are given for a single query. During training, HAPPIER optimizes our hierarchical retrieval objective by averaging several queries.

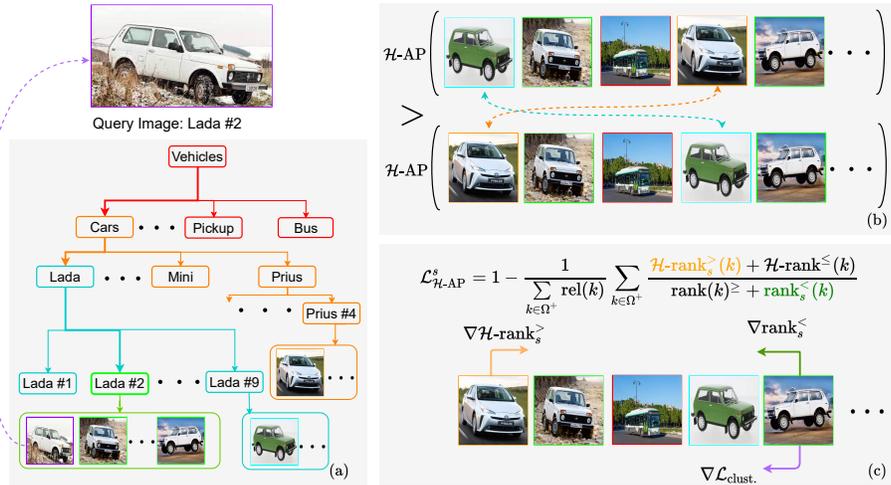


Fig. 2: HAPPIER leverages a hierarchical tree representing the semantic similarities between concepts in (a) to introduce a new hierarchical metric, $\mathcal{H}\text{-AP}$ in Eq. (3), see (b). $\mathcal{H}\text{-AP}$ exploits the hierarchy to weight rankings’ inversion: given the query image of a “Lada #2”, $\mathcal{H}\text{-AP}$ penalizes an inversion with a “Lada #9” less than with a “Prius #4”. To directly train models with $\mathcal{H}\text{-AP}$, we carefully study the structure of the problem and introduce the $\mathcal{L}_{\mathcal{H}\text{-AP}}^s$ loss in Eq. (5), which provides a smooth upper bound of $\mathcal{L}_{\mathcal{H}\text{-AP}}$, see (c). We also train HAPPIER with the $\mathcal{L}_{\text{clust.}}$ loss in Eq. (6) to enforce the partial ordering in stochastic optimization to match the global ones.

relevance with respect to the query: $\Omega^{(2)}$ in Fig. 2 is the set of “Lada” that are not “Lada #2” (blue) and $\Omega^{(1)}$ is the set of “Cars” that are not “Lada” (orange). We also define $\Omega^- := \Omega^{(0)}$ as the set of negative instances, *i.e.* the set of vehicles that are not “Cars” (in red) in Fig. 2 and $\Omega^+ = \bigcup_{l=1}^L \Omega^{(l)}$. Each instance k of $\Omega^{(l)}$ is thus associated a value through the *relevance function* denoted as $\text{rel}(k)$ [16].

To rank the instances $x_j \in \Omega$ with respect to the query q , we compute cosine similarities in an embedding space. More precisely, we extract embedding vectors using a deep neural network f parameterized by θ , $v_j = f_\theta(x_j)$, and compute the cosine similarity between the query and every image $s_j = f_\theta(q)^T v_j$. Images are then ranked by decreasing cosine similarity score. We learn the parameters θ of the network with HAPPIER, our framework to directly minimize $\mathcal{L}_{\mathcal{H}\text{-AP}}(\theta) = 1 - \mathcal{H}\text{-AP}(\theta)$. This enforces a ranking where the instances with the highest cosine similarity scores belong to $\Omega^{(L)}$, then $\Omega^{(L-1)}$ *etc.* and the items with the lowest cosine similarity belong to Ω^- .

3.1 Hierarchical Average Precision

Average Precision (AP) is the most common metric in Image Retrieval. AP evaluates a ranking in a binary setting: for a given query, each instance is either

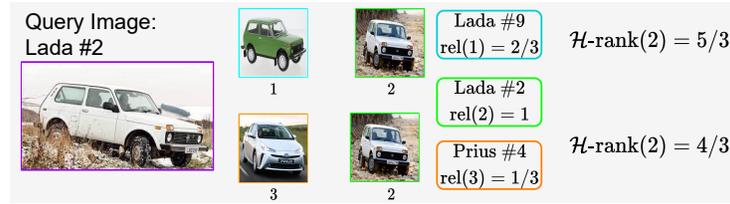


Fig. 3: Given a “Lada #2” query, the top inversion is less severe than the bottom one. Indeed on the top row instance 1 is semantically closer to the query – as it is a “Lada” – than instance 3 on the bottom row. Indeed instance 3’s closest common ancestor with the query, “Cars”, is farther in the hierarchical tree (see Fig. 2a). Because of that $\mathcal{H}\text{-rank}(2)$ is greater on the top row ($5/3$) than on the bottom row ($4/3$), leading to a greater $\mathcal{H}\text{-AP}$ in Fig. 2b for the top row.

positive or negative. It is computed as the average of precision at each rank n over the positive set $\text{AP} = \frac{1}{|\Omega^+|} \sum_{n=1}^N \text{Prec}(n)$. Previous works have written the AP using the ranking operator [2] as in Eq. (1). The rank for an instance k is written as a sum of Heaviside (step) function H [25]: this counts the number of instances j ranked before k , *i.e.* that have a higher cosine similarity ($s_j > s_k$). rank^+ is the rank among the positive instances, *i.e.* restricted to Ω^+ .

$$\text{AP} = \frac{1}{|\Omega^+|} \sum_{k \in \Omega^+} \frac{\text{rank}^+(k)}{\text{rank}(k)}, \text{ with } \begin{cases} \text{rank}(k) = 1 + \sum_{j \in \Omega} H(s_j - s_k) \\ \text{rank}^+(k) = 1 + \sum_{j \in \Omega^+} H(s_j - s_k) \end{cases} \quad (1)$$

Extending AP to hierarchical image retrieval We propose an extension of AP that leverages non-binary labels. To do so, we extend the concept of rank^+ to the hierarchical case with the concept of hierarchical rank, $\mathcal{H}\text{-rank}$:

$$\mathcal{H}\text{-rank}(k) = \text{rel}(k) + \sum_{j \in \Omega^+} \min(\text{rel}(k), \text{rel}(j)) \cdot H(s_j - s_k). \quad (2)$$

Intuitively, $\min(\text{rel}(k), \text{rel}(j))$ corresponds to seeking the closest ancestor shared by instance k and j with the query in the hierarchical tree. As illustrated in Fig. 3, $\mathcal{H}\text{-rank}$ induces a smoother penalization for instances that do not share the same fine-grained label as the query but still share some coarser semantics, which is not the case for rank^+ .

From $\mathcal{H}\text{-rank}$ in Eq. (2) we define the Hierarchical Average Precision, $\mathcal{H}\text{-AP}$:

$$\mathcal{H}\text{-AP} = \frac{1}{\sum_{k \in \Omega^+} \text{rel}(k)} \sum_{k \in \Omega^+} \frac{\mathcal{H}\text{-rank}(k)}{\text{rank}(k)} \quad (3)$$

Eq. (3) extends the AP to non-binary labels. We replace rank^+ by our hierarchical rank $\mathcal{H}\text{-rank}$ and the normalization term $|\Omega^+|$ is replaced by $\sum_{k \in \Omega^+} \text{rel}(k)$, which both represent the “sum of positives”, see more details in supplementary A.2.

\mathcal{H} -AP extends the desirable properties of the AP. It evaluates the quality of a ranking by: i) penalizing inversions of instances that are not ranked in decreasing order of relevances with respect to the query, ii) giving stronger emphasis to inversions that occur at the top of the ranking. Finally, we can observe that, by this definition, \mathcal{H} -AP is equal to the AP in the binary setting ($L=1$). This makes \mathcal{H} -AP a *consistent generalization* of AP (details in supplementary A.2).

Relevance function design The relevance $\text{rel}(k)$ defines how “similar” an instance $k \in \Omega^{(l)}$ is to the query q . While $\text{rel}(k)$ might be given as input in Information Retrieval datasets [24,9], we need to define it based on the hierarchical tree in our case. We want to enforce the constraint that the relevance decreases when going up the tree, *i.e.* $\text{rel}(k) > \text{rel}(k')$ for $k \in \Omega^{(l)}$, $k' \in \Omega^{(l')}$ and $l > l'$. To do so, we assign a total weight of $(l/L)^\alpha$ to each semantic level l , where $\alpha \in \mathbb{R}^+$ controls the decrease rate of similarity in the tree. For example for $L=3$ and $\alpha=1$, the total weights for each level are $1, \frac{2}{3}, \frac{1}{3}$ and 0. The instance relevance $\text{rel}(k)$ is normalized by the cardinal of $\Omega^{(l)}$:

$$\text{rel}(k) = \frac{(l/L)^\alpha}{|\Omega^{(l)}|} \text{ if } k \in \Omega^{(l)} \quad (4)$$

Other definitions fulfilling the decreasing similarity behaviour in the tree are possible. An interesting option for the relevance enables to recover a weighted sum of AP, denoted as $\sum w\text{AP} := \sum_{l=1}^L w_l \cdot \text{AP}^{(l)}$ (supplementary A.2), *i.e.* the weighted sum of AP is a particular case of \mathcal{H} -AP.

We set $\alpha=1$ in Eq. (4) for the \mathcal{H} -AP metric and in our main experiments. Setting α to larger values supports better performances on fine-grained levels as their relevances will relatively increase. This variant is denoted HAPPIER_F and discussed in Sec. 4.

3.2 Direct optimization of \mathcal{H} -AP

\mathcal{H} -AP in Eq. (3) involves the computation of \mathcal{H} -rank and rank, which are non-differentiable due to the summing of Heaviside step functions. We thus introduce a smooth approximation of \mathcal{H} -AP to obtain a surrogate loss amenable to gradient descent, which fulfils theoretical guarantees for proper optimization.

Re-writing \mathcal{H} -AP In order to design our surrogate loss for $\mathcal{L}_{\mathcal{H}\text{-AP}} = 1 - \mathcal{H}\text{-AP}$, we decompose \mathcal{H} -rank and rank into two quantities. Denoting $\mathcal{H}\text{-rank}^>(k)$ (resp. $\mathcal{H}\text{-rank}^\leq(k)$) as the restriction of \mathcal{H} -rank to instances of strictly higher relevances (resp. lower or equal), we can see that $\mathcal{H}\text{-rank}(k) = \mathcal{H}\text{-rank}^>(k) + \mathcal{H}\text{-rank}^\leq(k)$. The rank can be decomposed in a similar fashion: $\text{rank}(k) = \text{rank}^\geq(k) + \text{rank}^<(k)$ where $<$ (resp. \geq) denotes the restriction to instances of strictly lower relevances (resp. higher or equal). The $\mathcal{L}_{\mathcal{H}\text{-AP}}$ can be rewritten as follow:

$$\mathcal{L}_{\mathcal{H}\text{-AP}} = 1 - \frac{1}{\sum_{k \in \Omega^+} \text{rel}(k)} \sum_{k \in \Omega^+} \frac{\mathcal{H}\text{-rank}^>(k) + \mathcal{H}\text{-rank}^\leq(k)}{\text{rank}^\geq(k) + \text{rank}^<(k)}. \quad (5)$$

We choose to optimize over $\mathcal{H}\text{-rank}^>$ and $\text{rank}^<$ in Eq. (5). We maximize $\mathcal{H}\text{-rank}^>$ to enforce that the k^{th} instance must decrease in cosine similarity score if it is ranked before another instance of higher relevance ($\nabla\mathcal{H}\text{-rank}^>$ in Fig. 2 enforces the blue instance to be ranked after the green one as it is less relevant to the query). We minimize $\text{rank}^<$ to encourage the k^{th} instance to increase in cosine similarity score if it is ranked after one or more instances of lower relevance ($\nabla\text{rank}^<$ in Fig. 2 enforces that the last green instance moves before less relevant instances). Optimizing both those terms leads to a decrease in $\mathcal{L}_{\mathcal{H}\text{-AP}}$. On the other hand, we purposely do not optimize the two remaining $\mathcal{H}\text{-rank}^<(k)$ and $\text{rank}^>(k)$ terms, since this could harm training performances as explained in supplementary A.3.

Upper bound of $\mathcal{L}_{\mathcal{H}\text{-AP}}$ Based on the previous analysis, we now design our surrogate loss $\mathcal{L}_{\mathcal{H}\text{-AP}}^s$ by introducing a smooth approximation of $\text{rank}^<$ and $\mathcal{H}\text{-rank}^>(k)$. An important sought property of $\mathcal{L}_{\mathcal{H}\text{-AP}}^s$ is that it is an upper bound of $\mathcal{L}_{\mathcal{H}\text{-AP}}$. To this end, we approximate $\mathcal{H}\text{-rank}^>(k)$ with a piece-wise linear function that is a lower bound of the Heaviside function. $\text{rank}^<$ is approximated with a smooth upper bound of the Heaviside that combines a piece-wise sigmoid function and an affine function, which has been shown to make the training more robust thanks to the induced implicit margins between positives and negatives [31,2,28]. More details are given in supplementary A.3 on those surrogates.

Clustering constraint in HAPPIER Positives only need to have a greater cosine similarity with the query than negatives in order to be correctly ranked. Yet, we cannot optimize the ranking on the entire datasets – and thus the true $\mathcal{L}_{\mathcal{H}\text{-AP}}$ – because of the batch-wise estimation performed in stochastic gradient descent. To mitigate this issue, we take inspiration from clustering methods [42,35] to define the following objective in order to group closely the embeddings of instances that share the same fine-grained label:

$$\mathcal{L}_{\text{clust.}}(\theta) = -\log\left(\frac{\exp(\frac{v_y^T p_y}{\sigma})}{\sum_{p_z \in \mathcal{Z}} \exp(\frac{v_y^T p_z}{\sigma})}\right), \quad (6)$$

where p_y is the normalized proxy corresponding to the fine-grained class of the embedding v_y , \mathcal{Z} is the set of proxies, and σ is a temperature scaling parameter. In Fig. 2, $\nabla\mathcal{L}_{\text{clust.}}$ further clusters “Lada #2” instances. $\mathcal{L}_{\text{clust.}}$ induces a reference shared across batches and thus enforces that the partial ordering in-between batches is consistent with the global ordering over the entire retrieval set.

Our resulting final objective is a linear combination of both our losses, with a weight factor $\lambda \in [0,1]$ that balances the two terms:

$$\mathcal{L}_{\text{HAPPIER}}(\theta) = (1-\lambda) \cdot \mathcal{L}_{\mathcal{H}\text{-AP}}^s(\theta) + \lambda \cdot \mathcal{L}_{\text{clust.}}(\theta) .$$

4 Experiments

4.1 Experimental setup

Datasets We use the standard benchmark Stanford Online Products [22] (SOP) with two levels of hierarchy ($L=2$), and iNaturalist-2018 [37] with the standard splits from [2] in two settings: i) iNat-base with two levels of hierarchy ($L=2$) ii) iNat-full with the full biological taxonomy composed of 7 levels ($L=7$). We also evaluate on the recent dynamic metric learning (DyML) datasets (DyML-V, DyML-A, DyML-P) introduced in [33] for the task of hierarchical image retrieval, each with 3 semantic levels ($L=3$).

Implementation details Our base model is a ResNet-50 pretrained on ImageNet for SOP and iNat-base/full, and a ResNet-34 randomly initialized on DyML-V&A and pretrained on ImageNet on DyML-P, following [33]. Unless specified otherwise, all reported results are obtained with $\alpha=1$ in Eq. (4) and $\lambda=0.1$ for $\mathcal{L}_{\text{HAPPIER}}$. We study the impact of these parameters in Sec. 4.3.

Metrics For SOP and iNat, we evaluate the models based on three hierarchical metrics: \mathcal{H} -AP – which we introduced in Eq. (3) – the Average Set Intersection (ASI) and the Normalized Discounted Cumulative Gain (NDCG), defined in supplementary B.3. We also report the AP for each semantic level. For DyML, we follow the evaluation protocols of [33] and compute AP, ASI and R@1 on each semantic scale before averaging them. We cannot compute \mathcal{H} -AP or NDCG on those datasets as the hierarchical tree is not available on the test set.

Baselines We compare HAPPIER to several recent image retrieval methods optimized at the fine-grained level, which represent strong baselines for IR when training with binary labels: Triplet SH (TL_{SH}) [41], NormSoftMax (NSM) [42], ProxyNCA++ (NCA++) [35] and ROADMAP [28]. We also benchmark against hierarchical methods obtained by summing these fine-grained losses at different levels (denoted by Σ), and with respect to the recent hierarchical CSL loss [33]. Details on the experimental setup are given in supplementary B.

4.2 Main Results

Hierarchical results We first evaluate HAPPIER on global hierarchical metrics. On Tab. 1, we notice that HAPPIER significantly outperforms methods trained on the fine-grained level only, with a gain on \mathcal{H} -AP over the best performing methods of +16.1pt on SOP, +13pt on iNat-base and 12.7pt on iNat-full. HAPPIER also exhibits significant gains compared to hierarchical methods. On \mathcal{H} -AP, HAPPIER has important gains on all datasets (*e.g.* +6.3pt on SOP, +4.2pt on iNat-base over the best competitor), but also on ASI and NDCG. This shows the strong generalization of the method on standard metrics. Compared to the recent CSL loss [33], we observe a consistent gain over all metrics and datasets, *e.g.* +6pt on \mathcal{H} -AP, +8pt on ASI and +2.6pts on NDCG on SOP. This shows the benefits of optimizing a well-behaved hierarchical metric compared to an ad-hoc proxy method.

Table 1: Comparison of HAPPIER on SOP and iNat-base/full when using hierarchical metrics. Best results in **bold**, second best underlined.

Method	SOP			iNat-base			iNat-full			
	\mathcal{H} -AP	ASI	NDCG	\mathcal{H} -AP	ASI	NDCG	\mathcal{H} -AP	ASI	NDCG	
Fine	Triplet SH [41]	42.2	22.4	78.8	39.5	63.7	91.5	36.1	59.2	89.8
	NSM [42]	42.8	21.1	78.3	38.0	51.6	88.9	33.3	51.7	88.2
	NCA++ [35]	43.0	21.5	78.4	39.5	57.0	90.1	35.3	55.7	89.0
	Smooth-AP [2]	42.9	20.6	78.2	41.3	64.2	91.9	37.2	60.1	90.1
	ROADMAP [28]	43.3	19.1	77.9	40.3	61.0	91.2	34.7	59.6	89.5
Hier.	Σ TL _{SH} [41]	<u>53.1</u>	53.3	<u>89.2</u>	44.0	87.4	96.4	39.9	<u>85.5</u>	92.0
	Σ NSM [42]	50.4	49.7	87.0	47.9	75.8	94.4	<u>46.9</u>	74.2	93.8
	Σ NCA++ [35]	49.5	52.8	87.8	48.9	78.7	95.0	44.7	74.3	92.6
	CSL [33]	52.8	<u>57.9</u>	88.1	<u>50.1</u>	89.3	<u>96.7</u>	45.1	84.9	93.0
HAPPIER	59.4	65.9	91.5	54.3	89.3	96.9	47.9	87.2	93.8	

On Tab. 2, we evaluate HAPPIER on the recent DyML benchmarks. HAPPIER again shows significant gains in mAP and ASI compared to methods only trained on fine-grained labels, *e.g.* +9pt in mAP and +10pt in ASI on DyML-V. HAPPIER also outperforms other hierarchical baselines: +4.8pt mAP on DyML-V, +0.9 on DyML-A and +1.8 on DyML-P. In R@1, HAPPIER performs on par with other methods on DyML-V and outperforms other hierarchical baselines by a large margin on DyML-P: 63.7 *vs.* 60.8 for Σ NSM. Interestingly, HAPPIER also consistently outperforms CSL [33] on its own datasets².

Table 2: Performance comparison on Dynamic Metric Learning benchmarks [33].

Method	DyML-Vehicle			DyML-Animal			DyML-Product			
	mAP	ASI	R@1	mAP	ASI	R@1	mAP	ASI	R@1	
Fine	TL _{SH} [41]	26.1	38.6	84.0	37.5	46.3	66.3	36.32	46.1	59.6
	NSM [42]	27.7	40.3	88.7	38.8	48.4	<u>69.6</u>	35.6	46.0	57.4
	Smooth-AP [2]	27.1	39.5	83.8	37.7	45.4	63.6	36.1	45.5	55.0
	ROADMAP [28]	27.1	39.6	84.5	34.4	42.6	62.8	34.6	44.6	<u>62.5</u>
Hier.	Σ TL _{SH} [41]	25.5	38.1	81.0	38.9	47.2	65.9	<u>36.9</u>	46.3	58.5
	Σ NSM [42]	<u>32.0</u>	<u>45.7</u>	89.4	<u>42.6</u>	<u>50.6</u>	70.0	36.8	<u>46.9</u>	60.8
	CSL [33]	30.0	43.6	87.1	40.8	46.3	60.9	31.1	40.7	52.7
HAPPIER	37.0	49.8	<u>89.1</u>	43.8	50.8	68.9	38.0	47.9	63.7	

² CSL’s score on Tab. 2 are above those reported in [33]; personal discussions with the authors [33] validate that our results are valid for CSL, see supplementary B.5.

Detailed evaluation Tabs. 3 and 4 shows the different methods’ performances on all semantic hierarchy levels. We evaluate HAPPIER and also HAPPIER_F ($\alpha > 1$ for Eq. (4) in Sec. 3.1), with $\alpha = 5$ on SOP and $\alpha = 3$ on iNat-base/full. HAPPIER optimizes the overall hierarchical performances, while HAPPIER_F is meant to be optimal at the fine-grained level while still optimizing coarser levels.

Table 3: Comparison of HAPPIER *vs.* methods trained only on fine-grained labels on SOP and iNat-base. Metrics are reported for both semantic levels.

		SOP			iNat-base		
		Fine		Coarse	Fine		Coarse
	Method	R@1	AP	AP	R@1	AP	AP
Fine	TL _{SH} [41]	79.8	59.6	14.5	66.3	33.3	51.5
	NSM [42]	81.3	61.3	13.4	70.2	<u>37.6</u>	38.8
	NCA++ [35]	81.4	61.7	13.6	67.3	37.0	44.5
	Smooth-AP [2]	81.3	61.7	13.4	67.3	35.2	53.1
	ROADMAP [28]	82.2	62.5	12.9	69.3	35.1	50.4
Hier.	CSL [33]	79.4	58.0	<u>45.0</u>	62.9	30.2	<u>88.5</u>
	HAPPIER	81.0	60.4	58.4	<u>70.7</u>	36.7	88.6
	HAPPIER_F	<u>81.8</u>	<u>62.2</u>	36.0	71.0	37.8	78.8

On Tab. 3, we observe that HAPPIER gives the best performances at the coarse level, with a significant boost compared to fine-grained methods, *e.g.* +43.9pt AP compared to the best non-hierarchical TL_{SH} [41] on SOP. HAPPIER even outperforms the best fine-grained methods in R@1 on iNat-base, but is slightly below on SOP. HAPPIER_F performs on par with the best methods at the finest level on SOP, while further improving performances on iNat-base, and still significantly outperforms fine-grained methods at the coarse level.

The satisfactory behaviour and the two optimal regimes of HAPPIER and HAPPIER_F are confirmed and even more pronounced on iNat-full (Tab. 4): HAPPIER gives the best results on coarser levels (from “Order”), while being very close to the best results on finer ones. HAPPIER_F gives the best results at the finest levels, even outperforming very competitive fine-grained baselines.

Again, note that HAPPIER outperforms CSL [33] on all semantic levels and datasets on Tabs. 3 and 4, *e.g.* +5pt on the fine-grained AP (“Species”) and +3pt on the coarsest AP (“Kingdom”) on Tab. 4.

4.3 HAPPIER analysis

Ablation study In Tab. 5, we study the impact of our different choices regarding the direct optimization of \mathcal{H} -AP. The baseline method uses a sigmoid to optimize \mathcal{H} -AP as in [25, 2]. Switching to our surrogate loss $\mathcal{L}_{\mathcal{H}\text{-AP}}^s$ Sec. 3.2 yields a +0.8pt increase in \mathcal{H} -AP. Finally, the combination with $\mathcal{L}_{\text{clust.}}$ in HAPPIER results in an additional 1.3pt improvement in \mathcal{H} -AP.

Table 4: Comparison of HAPPIER *vs.* methods trained only on fine-grained labels on iNat-Full. Metrics are reported for all 7 semantic levels.

Method	Species		Genus	Family	Order	Class	Phylum	Kingdom
	R@1	AP						
TL _{SH} [41]	66.3	33.3	34.2	32.3	35.4	48.5	54.6	68.4
NSM [42]	<u>70.2</u>	37.6	<u>38.0</u>	31.4	28.6	36.6	43.9	63.0
NCA++ [35]	67.3	37.0	37.9	33.0	32.3	41.9	48.4	66.1
Smooth-AP [2]	67.3	35.2	36.3	33.5	35.0	49.3	55.8	69.9
ROADMAP [28]	69.3	35.1	35.4	29.3	29.6	46.4	54.7	69.5
<hr/>								
Hier. CSL [33]	59.9	30.4	32.4	36.2	50.7	<u>81.0</u>	<u>87.4</u>	<u>91.3</u>
<hr/>								
Hier. HAPPIER	70.2	36.0	37.0	38.0	51.9	81.3	89.1	94.4
HAPPIER_F	70.8	37.6	38.2	38.8	<u>50.9</u>	76.1	82.2	83.1

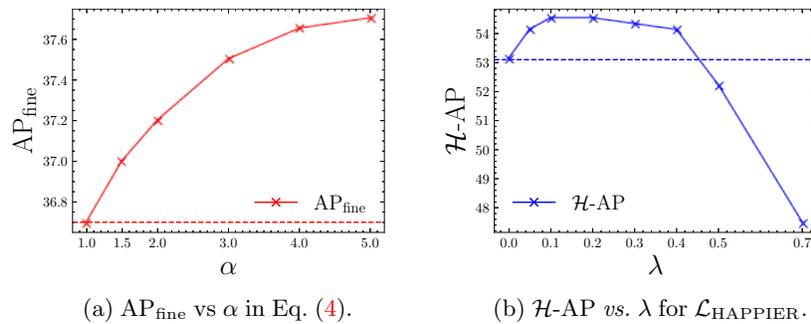
Table 5: Impact of optimization choices for \mathcal{H} -AP (cf. Sec. 3.2) on iNat-base.

$\mathcal{L}_{\mathcal{H}\text{-AP}}^s$	$\mathcal{L}_{\text{clust.}}$	\mathcal{H} -AP
\times	\times	52.3
\checkmark	\times	53.1
\checkmark	\checkmark	54.3

Table 6: Comparison of \mathcal{H} -AP (Eq. (4)) and $\sum w\text{AP}$ from supplementary A.2.

test \rightarrow train \downarrow	\mathcal{H} -AP	$\sum w\text{AP}$	NDCG
\mathcal{H} -AP	53.1	39.8	97.0
$\sum w\text{AP}$	52.0	40.5	96.4

Impact of the relevance function Tab. 6 compares models that are trained with the relevance function of Eq. (4), *i.e.* \mathcal{H} -AP, and $\sum w\text{AP}$ (relevance given in supplementary A.2). We report results for \mathcal{H} -AP, $\sum w\text{AP}$ and NDCG. Both \mathcal{H} -AP, $\sum w\text{AP}$ perform better when trained with their own metric: +1.1pt \mathcal{H} -AP for the model trained to optimize it and +0.7pt $\sum w\text{AP}$ for the model trained to optimize it. Both models show similar performances in NDCG (96.4 *vs.* 97.0).

Fig. 4: Impact on Inat-base of α in Eq. (4) for setting the relevance of \mathcal{H} -AP (a) and of the λ hyper-parameter on HAPPIER results (b).

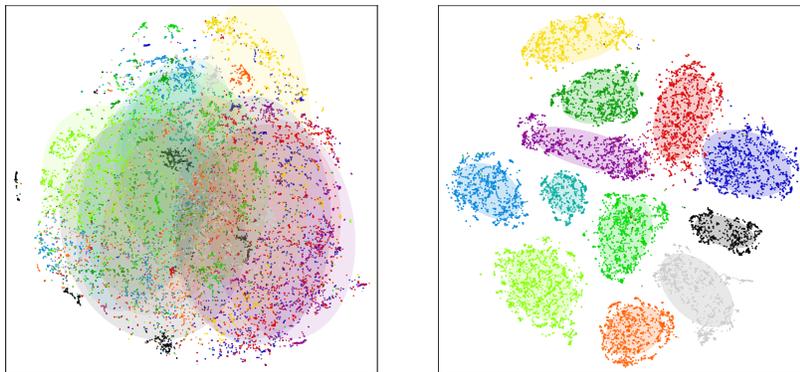
Hyper-parameters Fig. 4a studies the impact of α for setting the relevance in Eq. (4): increasing α improves the performances of the AP at the fine-grained level on iNat-base, as expected. We also show in Fig. 4b the impact of λ weighting $\mathcal{L}_{\mathcal{H}\text{-AP}}^s$ and $\mathcal{L}_{\text{clust.}}$ in HAPPIER performances: we observe a stable increase in $\mathcal{H}\text{-AP}$ within $0 < \lambda < 0.5$ compared to optimizing only $\mathcal{L}_{\mathcal{H}\text{-AP}}^s$, while a drop in performance is observed for $\lambda > 0.5$. This shows the complementarity of $\mathcal{L}_{\mathcal{H}\text{-AP}}^s$ and $\mathcal{L}_{\text{clust.}}$, and how, when combined, HAPPIER reaches its best performance.

4.4 Qualitative study

We provide here qualitative assessments of HAPPIER, including embedding space analysis and visualization of HAPPIER’s retrievals.

t-SNE: organization of the embedding space In Fig. 5, we plot using t-SNE [36,7] how HAPPIER learns an embedding space on SOP ($L=2$) that is well-organized. We plot the mean vector of each fine-grained class and we assign the color based on the coarse level. We show on Fig. 5a the t-SNE visualisation obtained using a baseline method trained on the fine-grained labels, and in Fig. 5b we plot the t-SNE of the embedding space of a model trained with HAPPIER. We cannot observe any clear clusters for the coarse level on Fig. 5a, whereas we can appreciate the the quality of the hierarchical clusters formed on Fig. 5b.

Controlled errors Finally, we showcase in Fig. 6 errors of HAPPIER *vs.* a fine-grained baseline. On Fig. 6a, we illustrate how a model trained with HAPPIER makes mistakes that are less severe than a baseline model trained only on the fine-grained level. On Fig. 6b, we show an example where both models fail to retrieve the correct fine-grained instances, however the model trained with HAPPIER retrieves images of bikes that are visually more similar to the query.



(a) t-SNE visualization of a model trained only on the fine-grained labels. (b) t-SNE visualization of a model trained with **HAPPIER**.

Fig. 5: t-SNE visualisation of the embedding space of two models trained on SOP. Each point is the average embedding of each fine-grained label (object instance) and the colors represent coarse labels (object category, *e.g.* bike, coffee maker).



(a) HAPPIER can help make less severe mistakes. The inversion on the bottom row are with negative instances (in red), whereas with HAPPIER (top row) inversions are with instances sharing the same coarse label “bike” (in orange).



(b) In this example, the models fail to retrieve the correct fine-grained images. However HAPPIER still retrieves images of very similar bikes (in orange) whereas the baseline retrieves images that are dissimilar semantically to the query (in red).

Fig. 6: Qualitative examples of failure cases from a standard fine-grained model corrected by training with HAPPIER.

5 Conclusion

In this work, we introduce HAPPIER, a new training method that leverages hierarchical relations between concepts to learn robust rankings. HAPPIER is based on a new metric \mathcal{H} -AP that evaluates hierarchical rankings and uses a combination of a smooth upper bound surrogate with theoretical guarantees and a clustering loss to directly optimize it. Extensive experiments show that HAPPIER performs on par to state-of-the-art image retrieval methods on fine-grained metrics and exhibits large improvements *vs.* recent hierarchical methods on hierarchical metrics. Learning more robust rankings reduces the severity of ranking errors, and is qualitatively related to a better organization of the embedding space with HAPPIER. Future works include the adaptation of HAPPIER to the unsupervised setting, *e.g.* for providing a relevant self-training criterion.

Acknowledgement This work was done under a grant from the the AHEAD ANR program (ANR-20-THIA-0002). It was granted access to the HPC resources of IDRIS under the allocation 2021-AD011012645 made by GENCI.

References

1. Bertinetto, L., Mueller, R., Tertikas, K., Samangooei, S., Lord, N.A.: Making better mistakes: Leveraging class hierarchies with deep networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12506–12515 (2020) **3**
2. Brown, A., Xie, W., Kalogeiton, V., Zisserman, A.: Smooth-ap: Smoothing the path towards large-scale image retrieval. In: European Conference on Computer Vision. pp. 677–694. Springer (2020) **1, 3, 6, 8, 9, 10, 11, 12**
3. Bruch, S., Zoghi, M., Bendersky, M., Najork, M.: Revisiting approximate metric optimization in the age of deep neural networks. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 1241–1244 (2019) **3**
4. Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., Hullender, G.: Learning to rank using gradient descent. In: Proceedings of the 22nd International Conference on Machine Learning. p. 89–96. ICML '05, Association for Computing Machinery, New York, NY, USA (2005). <https://doi.org/10.1145/1102351.1102363>, <https://doi.org/10.1145/1102351.1102363> **3**
5. Burges, C., Ragno, R., Le, Q.: Learning to rank with nonsmooth cost functions. In: Schölkopf, B., Platt, J., Hoffman, T. (eds.) Advances in Neural Information Processing Systems. vol. 19. MIT Press (2006), <https://proceedings.neurips.cc/paper/2006/file/af44c4c56f385c43f2529f9b1b018f6a-Paper.pdf> **3**
6. Cakir, F., He, K., Xia, X., Kulis, B., Sclaroff, S.: Deep metric learning to rank. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1861–1870 (2019) **1, 3**
7. Chan, D.M., Rao, R., Huang, F., Canny, J.F.: Gpu accelerated t-distributed stochastic neighbor embedding. *Journal of Parallel and Distributed Computing* **131**, 1–13 (2019) **13**
8. Chang, D., Pang, K., Zheng, Y., Ma, Z., Song, Y.Z., Guo, J.: Your” flamingo” is my” bird”: Fine-grained, or not. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11476–11485 (2021) **3**
9. Chapelle, O., Chang, Y.: Yahoo! learning to rank challenge overview. In: Proceedings of the learning to rank challenge. pp. 1–24. PMLR (2011) **7**
10. Croft, W.B., Metzler, D., Strohman, T.: Search engines: Information retrieval in practice, vol. 520. Addison-Wesley Reading (2010) **3**
11. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4690–4699 (2019) **1, 3**
12. Dhall, A., Makarova, A., Ganea, O., Pavlo, D., Greeff, M., Krause, A.: Hierarchical image classification using entailment cone embeddings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 836–837 (2020) **3**
13. Dupret, G., Piwowarski, B.: A user behavior model for average precision and its generalization to graded judgments. In: Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 531–538. SIGIR '10, Association for Computing Machinery, New York, NY, USA (2010). <https://doi.org/10.1145/1835449.1835538>, <https://doi.org/10.1145/1835449.1835538> **3, 4**
14. Dupret, G., Piwowarski, B.: Model based comparison of discounted cumulative gain and average precision. *Journal of Discrete Algorithms* **18**,

- 49–62 (2013). <https://doi.org/https://doi.org/10.1016/j.jda.2012.10.002>, <https://www.sciencedirect.com/science/article/pii/S1570866712001372>, selected papers from the 18th International Symposium on String Processing and Information Retrieval (SPIRE 2011) 4
15. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06). vol. 2, pp. 1735–1742. IEEE (2006) 3
 16. Hjørland, B.: The foundation of the concept of relevance. *Journal of the American Society for Information Science and Technology* **61**(2), 217–237 (2010) 3, 5
 17. Järvelin, K., Kekäläinen, J.: Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)* **20**(4), 422–446 (2002) 3
 18. Järvelin, K., Kekäläinen, J.: Ir evaluation methods for retrieving highly relevant documents. In: *ACM SIGIR Forum*. vol. 51, pp. 243–250. ACM New York, NY, USA (2017) 3
 19. Kekäläinen, J., Järvelin, K.: Using graded relevance assessments in ir evaluation. *Journal of the American Society for Information Science and Technology* **53**(13), 1120–1129 (2002). <https://doi.org/https://doi.org/10.1002/asi.10137>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.10137> 3
 20. Law, M.T., Thome, N., Cord, M.: Learning a distance metric from relative comparisons between quadruplets of images. *International Journal of Computer Vision* **121**(1), 65–94 (2017) 3
 21. Movshovitz-Attias, Y., Toshev, A., Leung, T.K., Ioffe, S., Singh, S.: No fuss distance metric learning using proxies. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 360–368 (2017) 3
 22. Oh Song, H., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4004–4012 (2016) 2, 9
 23. P., M.V., Paulus, A., Musil, V., Martius, G., Rolínek, M.: Differentiation of blackbox combinatorial solvers. In: *ICLR* (2020) 1, 3
 24. Qin, T., Liu, T.: Introducing LETOR 4.0 datasets. *CoRR* **abs/1306.2597** (2013), <http://arxiv.org/abs/1306.2597> 7
 25. Qin, T., Liu, T.Y., Li, H.: A general approximation framework for direct optimization of information retrieval measures. *Information Retrieval* **13**, 375–397 (2009) 3, 6, 11
 26. Radenovic, F., Tolias, G., Chum, O.: CNN image retrieval learns from bow: Unsupervised fine-tuning with hard examples. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 9905, pp. 3–20. Springer (2016). https://doi.org/10.1007/978-3-319-46448-0_1, https://doi.org/10.1007/978-3-319-46448-0_1 1
 27. Radenović, F., Tolias, G., Chum, O.: Cnn image retrieval learns from bow: Unsupervised fine-tuning with hard examples. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *Computer Vision – ECCV 2016*. pp. 3–20. Springer International Publishing, Cham (2016) 3
 28. Ramzi, E., Thome, N., Rambour, C., Audebert, N., Bitot, X.: Robust and decomposable average precision for image retrieval. *Advances in Neural Information Processing Systems* **34** (2021) 1, 2, 3, 8, 9, 10, 11, 12
 29. Revaud, J., Almazán, J., Rezende, R.S., Souza, C.R.d.: Learning with average precision: Training image retrieval with a listwise loss. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5107–5116 (2019) 1, 3

30. Robertson, S.E., Kanoulas, E., Yilmaz, E.: Extending average precision to graded relevance judgments. In: Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval. pp. 603–610 (2010) [4](#)
31. Rolínek, M., Musil, V., Paulus, A., Vlastelica, M., Michaelis, C., Martius, G.: Optimizing rank-based metrics with blackbox differentiation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7620–7630 (2020) [1](#), [3](#), [8](#)
32. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 29. Curran Associates, Inc. (2016), <https://proceedings.neurips.cc/paper/2016/file/6b180037abbea991d8b1232f8a8ca9-Paper.pdf> [1](#), [3](#)
33. Sun, Y., Zhu, Y., Zhang, Y., Zheng, P., Qiu, X., Zhang, C., Wei, Y.: Dynamic metric learning: Towards a scalable metric space to accommodate multiple semantic scales. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5393–5402 (2021) [2](#), [4](#), [9](#), [10](#), [11](#), [12](#)
34. Taylor, M., Guiver, J., Robertson, S., Minka, T.: Softrank: Optimizing non-smooth rank metrics. In: Proceedings of the 2008 International Conference on Web Search and Data Mining. p. 77–86. WSDM '08, Association for Computing Machinery, New York, NY, USA (2008). <https://doi.org/10.1145/1341531.1341544>, <https://doi.org/10.1145/1341531.1341544> [3](#)
35. Teh, E.W., DeVries, T., Taylor, G.W.: Proxynca++: Revisiting and revitalizing proxy neighborhood component analysis. In: European Conference on Computer Vision. pp. 448–464. Springer (2020) [1](#), [2](#), [3](#), [8](#), [9](#), [10](#), [11](#), [12](#)
36. van der Maaten, L., Hinton, G.: Visualizing high-dimensional data using t-sne. Journal of Machine Learning Research **9**, 2579–2605 (2008) [13](#)
37. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8769–8778 (2018) [2](#), [9](#)
38. Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W.: Cosface: Large margin cosine loss for deep face recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5265–5274 (2018) [1](#), [3](#)
39. Wang, X., Han, X., Huang, W., Dong, D., Scott, M.R.: Multi-similarity loss with general pair weighting for deep metric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5022–5030 (2019) [1](#), [3](#)
40. Wang, X., Zhang, H., Huang, W., Scott, M.R.: Cross-batch memory for embedding learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6388–6397 (2020) [1](#)
41. Wu, C.Y., Manmatha, R., Smola, A.J., Krahenbuhl, P.: Sampling matters in deep embedding learning. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2840–2848 (2017) [1](#), [2](#), [3](#), [9](#), [10](#), [11](#), [12](#)
42. Zhai, A., Wu, H.Y.: Classification is a strong baseline for deep metric learning. arXiv preprint arXiv:1811.12649 (2018) [1](#), [2](#), [3](#), [8](#), [9](#), [10](#), [11](#), [12](#)