

A Benchmark and Asymmetrical-Similarity Learning for Practical Image Copy Detection

Wenhao Wang^{1,2*}, Yifan Sun², Yi Yang³

¹ ReLER, University of Technology Sydney

² Baidu Research

³ Zhejiang University

wangwenhao0716@gmail.com, sunyifan01@baidu.com, yangyics@zju.edu.cn

Abstract

Image copy detection (ICD) aims to determine whether a query image is an *edited copy* of any image from a reference set. Currently, there are very limited public benchmarks for ICD, and all overlook a critical challenge in real-world applications, *i.e.*, the distraction from hard negative queries. Specifically, some queries are not edited copies but are *inherently similar* to some reference images. These hard negative queries are easily false recognized as edited copies, therefore significantly compromising the ICD accuracy. This observation motivates us to build the first ICD benchmark featuring this characteristic. Based on existing ICD datasets, this paper constructs a new dataset by additionally adding 100,000 and 24,252 hard negative pairs into the training and test set, respectively. Moreover, this paper further reveals a unique difficulty for solving the hard negative problem in ICD, *i.e.*, there is a fundamental conflict between current metric learning and ICD. This conflict is: the metric learning adopts symmetric distance while the edited copy is an asymmetric (unidirectional) process, *e.g.*, a partial crop is close to its holistic reference image and is an edited copy, while the latter cannot be the edited copy of the former (in spite of the distance is equally small). This insight results in an Asymmetrical-Similarity Learning (ASL) method, which allows the similarity in two directions (the query \leftrightarrow the reference image) to be different from each other. Experimental results show that ASL outperforms state-of-the-art methods by a clear margin, confirming that solving the symmetric-asymmetric conflict is critical for ICD. The NDEC dataset and code are available at <https://github.com/WangWenhao0716/ASL>.

Introduction

Image copy detection (ICD) is to detect whether a query image is an edited copy of any image from a reference set (Fig. 1 (a)). It has extensive applications, such as preserving information legitimacy and integrity. However, current resources available for ICD research are not sufficient: the public datasets are rare and all overlook a critical challenge, *i.e.*, the distraction from hard negative samples. Concretely, in realistic ICD systems, some queries are not edited copies but look inherently similar to the reference images, as shown in Fig. 1 (b). These hard negative samples are prone to being



Figure 1: Respective examples for edited copy and hard negative sample. In (a), the right-side image is an *edited copy* of the left-side image. In (b), the right-side image is NOT the edited copy of the left-side image, in spite that they are *inherently similar*. Therefore, the right-side image is a hard negative sample which easily confuses the ICD algorithm. This paper builds the first Negative Distractor for Edited Copy (NDEC) benchmark to emphasize the distraction from hard negative samples.

falsely recognized as edited copies and significantly compromise the ICD accuracy. Therefore, while the hard positive problem (*i.e.*, some edited copies look quite dissimilar to their reference images) has already been well investigated, the hard negative problem is another critical challenge that deserves investigation in the ICD task.

To draw attention to the hard negative problem in ICD, this paper builds a Negative Distractors for Edited Copy (NDEC), the first benchmark featuring this characteristic from two aspects. First, NDEC adds abundant hard negative queries into the testing data, significantly increasing ICD difficulty. Among the 49,252 queries, there are 24,252 hard negative queries, which are quite similar to some reference images but are not edited copies. Second, the NDEC adds abundant hard negative pairs into the training data. Besides of the one million normal training images from (Douze et al. 2021), the NDEC training set contains 100,000 additional hard negative pairs. We believe adding these hard negative pairs for training is potential to improve the ICD accuracy. In a word, *the NDEC arouses the attention on hard negative distractors in ICD and meanwhile provides the prerequisite data for exploring the solutions.*

We benchmark NDEC with state-of-the-art ICD methods and find that this new dataset is very challenging. Generally, existing ICD methods in their core rely on deep metric learning. They aim to learn a deep feature space where each edited copy and its reference image are close to each other,

*Work done during an internship at Baidu Research.



Figure 2: The “reference \rightarrow edited copy” is an unidirectional process. From left to right, the latter image is partially cropped from the former image and contains less content/information. Consequentially, the latter image is the edited copy of the former, while the former image can NOT be the edited copy of the latter. Contrary to this asymmetric process, the distance by its definition is symmetric: the distance from the former to the latter is equal to the distance from the latter to the former.

while the non-edited copies are far away from the reference images. Although these methods are already equipped with modern hard mining strategies, We observe that they still have difficulties to distinguish those hard negative distractors against true edited copies. This important observation indicates that in ICD, the hard negative problem is especially difficult (compared to many other metric learning tasks).

An important reason causing the above-described difficulty is: there is a fundamental conflict between current metric learning and ICD. The metric learning generally adopts symmetric distance while the edited copy is an asymmetric (unidirectional) process. For example, a partial crop is close to its holistic reference image and is an edited copy, as shown in Fig. 2. However, in spite that the reference image is equally close to the partial crop, it cannot be the edited copy of the partial crop. More generally, the “reference \rightarrow edited copy” process usually discards some content/information and is thus unidirectional. Therefore, the symmetric distance is not sufficient for identifying an edited copy, especially when two images are close to each other.

In response to the symmetric-asymmetric conflict, this paper proposes a novel Asymmetric-Similarity Learning (ASL) method for ICD. ASL desires the deep feature of the image containing more content/information to have larger norm, which is learnable. Correspondingly, the norm ratio between the reference and its edited copy is larger than 1, while the norm ratio between the edited copy and the reference is smaller than 1. Therefore, ASL is asymmetric and well fits the unidirectional property of edited copy. Experimental results show that ASL substantially improves ICD accuracy on NDEC. It confirms that solving the symmetric-asymmetric conflict is beneficial for realistic ICD.

To sum up, this paper makes the following contributions:

1. We contribute a new ICD dataset, *i.e.*, Negative-Distractor for Edited Copy (NDEC), with emphasis on the seldom-noticed hard negative problem (while preserving the popular hard positive problem).
2. We benchmark NDEC with state-of-the-art methods and correspondingly reveals a fundamental conflict between the commonly-adopted symmetric distance and the asymmetric “reference \rightarrow edited copy” process.

3. We propose a novel Asymmetric-Similarity Learning (ASL) for ICD. ASL uses the norm ratio as an asymmetric similarity metric to distinguish edited copy against hard negative samples and substantially improves ICD.

Related Works

Existing Image Copy Detection Methods

Currently, public ICD methods are rare because organizations intend to keep the techniques as obscure as possible (Douze et al. 2021). Before deep learning era, ICD relies on hand-crafted methods, *e.g.*, global descriptor (Kim 2003; Wan et al. 2008), local descriptor, Bag-of-Words models and spatial constraints (Zhou et al. 2010, 2016). Deep learning has provided ICD with substantial improvement. In a recent large-scale ICD competition DISC2021 (Douze et al. 2021) held by Meta AI Research, all the winners use deep learning models. Basically, they follow distance-based deep metric learning approach. For example, D^2LV (Wang et al. 2021a) and BoT (Wang et al. 2021b) form each image in the training set as its own class and then use triplet loss (Hermans, Beyer, and Leibe 2017) and classification loss to train models. EfNet (Papadakis and Addicam 2021) uses ArcFace (Deng et al. 2019) to gain better representations. In this paper, we find that the new NDEC dataset significantly challenges these distance-based metric learning methods with a unique conflict between the unidirectional edited copy and the symmetric distance.

Self-supervised Learning

So far as we know, existing ICD methods rely on self-supervised learning (He et al. 2020; Grill et al. 2020; Chen et al. 2020b; Pizzi et al. 2022) due to a common sense, *i.e.*, the edited copies for training can be automatically generated based on unlabeled images. Specifically, the popular pipeline uses automatic transformations to generate positive training pairs and then apply deep metric learning.

In contrast to this common sense, we find that automatically-generated positive pairs and self-supervised learning do not suffice to learn discriminative ICD algorithms because they do not provide hard negative pairs. Therefore, we provide 100,000 manually-annotated hard negative pairs in addition to the unlabeled samples in NDEC and propose ASL to utilize these manually-annotated images. ASL combines both self-supervised learning (on the unlabeled data) and supervised learning (on the paired samples) and improves ICD.

Datasets

This section first briefly reviews the publicly available ICD datasets and then elaborately presents the proposed NDEC dataset.

Available Datasets

Currently, there are only two publicly available ICD datasets, *i.e.*, CopyDays (Douze et al. 2009) and DISC21 (Douze et al. 2021).

The CopyDays dataset (Douze et al. 2009) was proposed in 2009. It is a small-scaled dataset with only 3,157 test




Figure 3: The query images in the NDEC dataset. To simultaneously feature the hard negative and hard positive problems, we design some queries as hard negative distractors, and at the same time, there are also some edited copies with true matches in the reference dataset.

samples (157 query images + 3,000 reference images) and does not provide training data. Correspondingly, CopyDays does not accommodate deep learning methods. Moreover, the query images are easy to recognize, because the transformations for generating the edited copy are relatively simple, *e.g.*, resizing, cropping, blurring and changing the contrast. Therefore, CopyDays does not well present the challenge in realistic ICD systems.

Another existing dataset DISC21 (Douze et al. 2021) was proposed by Meta AI Research in NeurIPS 2021. Compared with CopyDays, DISC21 has three advantages. **1)** it is a large-scale dataset with one million (unlabeled) training images. **2)** it provides a challenging test set that better reflects the reality (than CopyDays does). Specifically, it has one million reference images and 50,000 queries produced through sophisticated transformations. Many query images look quite dissimilar to their corresponding reference images and are thus hard positive samples. **3)** it mimics the realistic needle-in-haystack setting, *i.e.*, the query dataset contains many distractors which are not edited copies.

Although DISC21 already notices the significance of distractors, its distractors are not confusing, deviating from the realistic scenario. Specifically, in DISC21, the query set does not contain hard negative samples which look very similar to the reference images. Therefore, using DISC21 for evaluation is prone to over-optimistic results.

The Proposed NDEC Dataset

To promote the research on realistic ICD, this paper proposes the NDEC (Negative Distractor for Edit Copy) dataset.

The NDEC dataset is based on two sources: the DISC21 dataset (Douze et al. 2021) and the OpenImage dataset (Kuznetsova et al. 2020; Krasin et al. 2017). Section already introduces DISC21. OpenImage is collected by Google and contains about nine million images annotated with image-level labels and other meta information. The NDEC does not use the annotations and merely uses OpenImage as a source for collecting samples.

Test set. ICD needs to inspect whether a given query image is the edited copy of any reference image. To simulate

this procedure, the test set basically consists of a query set and a reference set. In NDEC, the query and reference set contain 49,252 and 1,000,000 images, respectively. Specifically, among the 49,252 query images, only 5,009 images are edited copy and have true matches in the reference set. The remaining 44,243 query images are not edited copies but negative queries. Some of queries are shown in Fig. 3.

Compared with the negative queries in DISC21, the negative queries in our NDEC are much more difficult: there are 24,252 hard negative images selected from the OpenImage dataset. The selection procedure is semi-automatic: we use MultiGrain (Berman et al. 2019) to automatically evaluate the similarity score between OpenImage samples and the NDEC reference images. MultiGrain (Berman et al. 2019) returns 50,000 candidate images that are very similar to reference images. We note that some of these 50,000 candidate images are true edited copies. Therefore, we filter out those edited copies from the candidates and preserve the hard negative samples through manual inspection. The manual inspection over the 50,000 candidate images costs 2 annotators about 5 days’ work.

Another difference from DISC21 is that the NDEC prohibits training on the query and reference sets. We note that some state-of-the-art methods (Yokoo 2021; Sun et al. 2021) actually use 25,000 queries from DISC21 for training and achieves extra benefits. However, in realistic ICD, using queries for training is not quite feasible. Therefore, to ensure fair evaluation on NDEC, we remove these 25,000 queries from DISC21 when re-using its partial queries.

Training set. While the test set increases the ICD difficulty, the training set also provides extra hard negative pairs to promote the potential solutions. Specifically, the training set contains 1,100,000 images, which are comprised of 900,000 basic images and 100,000 hard negative pairs (*i.e.*, $100,000 \times 2$ paired images). Note that, in each pair, there is a manually labeled direction, *i.e.* the latter is a hard negative image to the former, and the former may be an edited copy of the latter (See Fig. 4). The basic images are from DISC21, while the hard negative images are from OpenImage. To collect these hard negative pairs, we use the same semi-automatic procedure for preparing the hard negative queries. Concretely, we use MultiGrain (Berman et al. 2019) to automatically select 150,000 candidates images (which are similar to some basic training images) and then manually filter out the edited copies. The manual filtering costs three annotators about 10 days for initial annotation, and another three annotators 5 days for a double-check. Fig. 4 visualizes some hard negative pairs in the training set.

Method

We first revisit the popular deep metric learning baseline and reveal the conflict between symmetric distance and the unidirectional reference \rightarrow edited copy process. Then we propose our Asymmetrical-Similarity Learning (ASL) under an intuitive edited copy scenario (*i.e.*, crop-to-copy) and generalize ASL to learn from general hard negative pairs. Finally, we emphasize our ASL with ICD baselines can handle other (basic) edited copies.

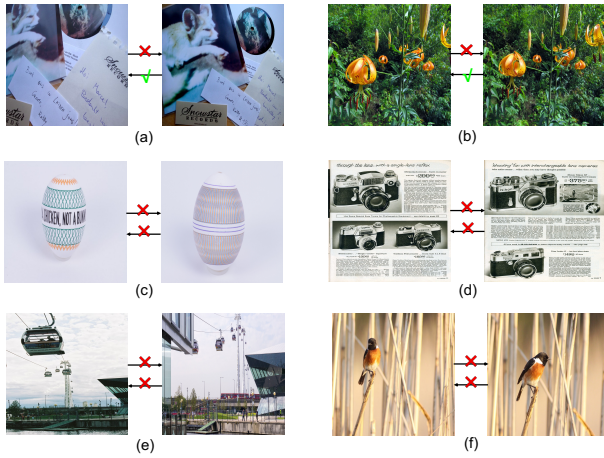


Figure 4: The hard negative pairs in the training set of the NDEC dataset. In each pair, we emphasize the right-side image is not an edited copy (cropped region) of the left-side image, though the left-side image may be an edited copy of the right-side image as in (a) and (b). In (c) and (d), these two images are visually similar. Two images from (e) are captured from different camera angles. In (f), two images are with the same location but different shooting time.

The Limitation of Deep Metric Learning Baseline

The ICD baseline. Existing ICD methods in their core rely on deep metric learning. The keynote is to learn a deep feature space where the reference image and its edited copy are close to each other, while the reference image and its non-edited copy are far away. To this end, during training, they treat a reference image and its edited copy as a positive pair, and use deep metric learning to minimize the distance within the positive pair. Concretely, the training procedure can be summarized into a two-step pipeline as below:

- *Generating positive pairs.* Since the ICD training set does not provide “reference-copy” pairs but only the unpaired training samples, the training pipeline relies on multiple transformations to automatically generate the edited copies, *e.g.*, cropping, blurring, and rotation. Consequentially, the auto-generated edited copy and its original image can be packed into a positive pair (or an individual class) for deep metric learning.

- *Deep metric learning.* Given the generated pairs or classes, deep metric learning usually uses them for pairwise training or classification training or pairwise plus classification training. The pairwise training directly optimizes the distances between samples, *i.e.*, an anchor image’s distance to a positive sample should be smaller than that its distance to a negative sample, through triplet loss (Hermans, Beyer, and Leibe 2017), N-Pair loss (Sohn 2016), *etc.* The classification training optimizes the distances between samples and a set of class proxies, *i.e.*, a sample’s distance to its target proxy should be smaller than its distance to the non-target proxies, through Softmax loss (Liu et al. 2016), CosFace (Wang et al. 2018), Circle loss (Sun et al. 2020), *etc.*

The symmetric-asymmetric conflict. A basic property

of the distance is the symmetry. Let us consider a reference image R and its edited copy E . The distance from R to E consistently equals the distance from E to R , *i.e.*, $D(R \rightarrow E) \equiv D(E \rightarrow R)$. Through deep metric learning, the distance from E to R is optimized to a small value, *i.e.*, $D(E \rightarrow R) < \epsilon$ (ϵ is a threshold), indicating E as an edited copy of R . Due to the symmetric property, we have $D(R \rightarrow E) < \epsilon$ as well, indicating that R is an edited copy of E . The latter indication is incorrect, because R cannot be the edited copy of E . As we explained in Section and Fig. 2, the process of reference \rightarrow edited copy is unidirectional. Therefore, R is close to E and is a hard negative sample of E . Existing deep metric learning is not capable to distinguish such hard negative samples, because it inevitably reduces $D(R \rightarrow E)$ when it tries to reduce $D(E \rightarrow R)$.

Asymmetrical-Similarity Learning (ASL)

ASL under Crop-to-Copy Scenario. We propose Asymmetrical-Similarity Learning (ASL) to avoid the above-described symmetric-asymmetric conflict. For intuitive understanding, we first focus on the popular crop-to-copy scenario*, where the edited copy is generated through cropping (Fig. 2). Another reason for we specifically focus on the crop-to-copy scenario is that it can be learned through self-supervision and thus achieves accuracy gains without any annotation cost.

Given two deep features \mathbf{x}_i and \mathbf{x}_j from image l_i and l_j , respectively, ASL defines the ratio between their norm value as:

$$R(\mathbf{x}_i \rightarrow \mathbf{x}_j) = |\mathbf{x}_i| / |\mathbf{x}_j| \quad (1)$$

Obviously, the norm ratio is asymmetric, because $R(\mathbf{x}_i \rightarrow \mathbf{x}_j) \neq R(\mathbf{x}_j \rightarrow \mathbf{x}_i)$. Given this asymmetric property, ASL uses norm ratio as the metric to identify crop-to-copy. Specifically, if l_j is a cropped copy of l_i (*i.e.*, l_i is the reference image of l_j), ASL desires $R(\mathbf{x}_i \rightarrow \mathbf{x}_j) > 1$. To this end, ASL generates abundant cropped copies and learns their norm ratios in a self-supervised manner, as illustrated in Fig. 5. Given a reference image, ASL crops a series of sub-images with decreasing sizes. During training, ASL selects a larger crop l_i and a smaller crop l_j as the pseudo reference image and the edited copy, respectively.

The training objective is

$$\mathcal{L} = \exp(1 - R(\mathbf{x}_i \rightarrow \mathbf{x}_j)) + \lambda \cdot \mathcal{L}_{mtr}, \quad (2)$$

where \mathcal{L}_{mtr} is a popular distance-based metric learning loss and λ is a balance parameter. \mathcal{L}_{mtr} can be implemented as CosFace, where the angles between different classes are learned to be larger than a margin. During training, decreasing the loss value \mathcal{L} increases $R(\mathbf{x}_i \rightarrow \mathbf{x}_j)$, therefore making the norm value of reference image l_i larger than that of the edited copy l_j . Therefore, ASL has two simultaneous learning effects, *i.e.*, it pulls inherently similar images close to each other and makes the image with more content/information have a larger feature norm, as illustrated in Fig. 5.

*CopyDays (Douze et al. 2009) points out most practical copy-right violations occurs with cropping.




Figure 5: The illustration of Asymmetrical-Similarity Learning (ASL). In ASL, distance-based metric learning loss pulls the features of inherently similar images close to each other. Meanwhile, the norm ratio based loss makes images with more content/information have a larger norm. For better visualization, on the right side, larger image sizes represent images with more content/information (though when training, all the images are resized to the same size).

During testing, given a query image l_Q and a reference image l_E , we first calculate their distance to identify whether they are close to each other. If the distance between them is small, we further compare their norm ratio $R(x_Q \rightarrow x_E)$. If $R(x_Q \rightarrow x_E) > 1$, it indicates that the query is actually larger than the reference image and thus is unlikely to be an edited copy.

ASL Using Hard Negative Pairs. The ASL in before section focuses on the crop-to-edit copy scenario in a self-supervised training manner. Given the manually annotated hard negative training pairs in NDEC, we generalize ASL to learn from more scenarios. We recall that the hard negative pairs in NDEC training set are annotated in single directions, *i.e.*, within each negative pair, the latter image l_L is similar to, but not the edited copy of the former image l_F , as illustrated in Fig. 4.

The principle of learning from these hard negative pairs is: ASL desires $R(x_L \rightarrow x_F) > 1$ because l_L is not the edited copy of l_F (the first item in Eq. 2), while considering l_L is positive to l_F in \mathcal{L}_{mtr} (the second item in Eq. 2) because these two images are inherently similar, as detailed below:

- **The norm ratio item in ASL treats l_L as negative to l_F .** Specifically, for the norm ratio item in Eq. 2, ASL considers that l_L contains no less content/information than l_F because l_L is not the edited copy of l_F . Therefore, ASL optimizes the norm ratio $R(x_L \rightarrow x_F)$ to be larger than 1.

- **The distance-based metric learning item treats l_L as positive to l_F .** Specifically, \mathcal{L}_{mtr} in ASL actually pulls l_L and l_F close to each other. This is because l_L and l_F are manually identified as being inherently similar to each other (in spite that there is no edited copy). If \mathcal{L}_{mtr} treats them as negative pair and pushes them far apart, it will make the true positive queries hard to recognize, therefore compromising ICD (*w.r.t.* the recall, in particular).

We note that although the above two items seem to stand against each other (one treats l_L as a negative to l_F while

the other treats l_L as a positive to l_F), they actually brings joint benefits for ICD. We validate this design through ablation study in Table 2. This confirms that ASL effectively alleviates the symmetric-asymmetric conflict in ICD task.

Combining ASL with the baseline. We note that ASL by itself only focuses on distinguishing the hard negative samples. Therefore, it is important to combine ASL with the ICD baseline, so as to recognize those basic edited copy patterns, *e.g.* color-jittered and padded images. Consequently, the baselines recognize the basic edited copies and ASL further filters out the hard negative ones.

Experiments

Protocols and Baselines

Evaluation protocols. We adopt two protocols, *i.e.* μAP and precision. Following DISC21 (Douze et al. 2021), μAP reflects the objectives of ICD. Besides μAP , we also use precision, defined as $\text{precision} = \text{TP} / (\text{TP} + \text{FP})$ where “TP” and “FP” denote the number of true positive and false positive matches in the top N pairs, respectively. Specifically, the N is set as 5,009 (the number of queries with true matches), which causes that $\text{precision} = 1$ ideally.

Baselines. We use two baselines, *i.e.*, a strong baseline used in D²LV (Wang et al. 2021a) and a simple (and relatively weak) baseline used in BoT (Wang et al. 2021b). The difference between these two baselines is that the former one uses additional model ensemble. The strong baseline allows ASL to compete against state-of-the-art methods, while the simple baseline is more convenient and clearer for ablation studies. To be computationally efficient, we replace the 8192-dim features (Wang et al. 2021a,b) with commonly-used 2048-dim features. Moreover, instead of combining cross-entropy loss and triplet loss (Hermans, Beyer, and Leibe 2017), we use a single CosFace (Wang et al. 2018). On DISC21, our implementation achieves consistent re-




Figure 6: The performance change of different proportions of hard negative distractors.

sults (90.13% and 72.42% μAP) as the baselines in D²LV (90.10% μAP) (Wang et al. 2021a) and BoT (72.73% μAP) (Wang et al. 2021b), respectively.

The Challenge from NDEC

Increasing hard negative queries decreases the accuracy.

Based on the simple baseline, we analyze how the distraction from hard negative queries impacts the ICD accuracy. We recall that the test set of NDEC contains 24,252 hard negative queries from OpenImage and 25,000 (5,009 positive and 19,991 relatively easy negative) queries from DISC21 dataset. We maintain the latter 25,000 queries unchanged in all experiments and gradually increase the hard negative queries. The corresponding μAP scores of the baseline are summarized in Fig. 6, from which we draw two observations. First, we clearly observe that as the hard negative queries increase from 0 to 24,252 (100%), the μAP scores gradually decrease from 72.42% to 47.00% (−25.42%). It indicates that the hard negative queries are very difficult to distinguish and significantly compromise the ICD accuracy. Second, when the volume of hard negative queries approaches 100%, the decreasing speed becomes very slow (e.g., −2.47% μAP drop for the last 25% increase of hard negative queries), indicating that the 24,252 hard negative queries in NDEC are roughly sufficient for presenting the challenge.

NDEC significantly challenges the state-of-the-art methods. We further benchmark NDEC with state-of-the-art methods (Wang et al. 2021a; Sun et al. 2021; Yokoo 2021; Papadakis and Addicam 2021; Wang et al. 2021b) (top performing methods in DISC2021 (Douze et al. 2021)) in Fig. 7. It clearly shows that all the methods undergo a substantial accuracy decrease when the evaluation dataset is changed from DISC21 to our NDEC. Specifically, the μAP drops are ranged from 23.13% (CNNCL) to 39.09% (EsViTp), indicating that NDEC is much more challenging than DISC21. Given that all the other testing conditions on DISC21 and NDEC (except for the hard negative distraction) are roughly the same, we infer the hard negative distraction in NDEC, which reflects the realistic challenge, is the major reason for this accuracy decrease.




Figure 7: Evaluation on DISC21 and our NDEC. All the methods undergo a substantial accuracy decrease when the evaluation dataset is changed from DISC21 to NDEC.

The Effectiveness of ASL

ASL achieves new state of the art. Fig. 7 further compares the proposed ASL against state-of-the-art methods, from which we draw two observations. First, on the challenging NDEC, ASL (with strong baseline) surpasses all the competing methods by a large margin. For example, ASL is higher than D²LV, the strongest competitor by +5.22% μAP . It indicates that ASL has the strongest capacity for solving the hard negative problem, which is the characteristic of NDEC as well as an important challenge in realistic ICD.

Second, we observe that on DISC21, ASL (with strong baseline) is not the best (−0.97% lower than D²LV). In other words, although ASL has significant superiority on NDEC, this superiority does not benefit its performance on DISC21. This is reasonable because DISC21 has no hard negative queries, and D²LV already well distinguishes those relatively easy negative distractors. It is another evidence that DISC21 does not well reflect the hard negative challenge in realistic ICD, and NDEC is valuable.

ASL consistently improves various baselines. Besides of our strong and simple baselines, we use more state-of-the-art methods as additional baselines to show that ASL brings general improvement, as illustrated in Table 1. Please note that for these additional baselines, we do not re-implement them and then integrate them into ASL. Instead, we use their models in an off-the-shelf manner and use an independently-trained ASL as a post processing. Therefore, “BoT + ASL” and “D²LV + ASL” achieve different (inferior) results compared to our “Simple + ASL” and “Strong + ASL”, respectively, in spite that the underlying baseline methods are the same.

Table 1 shows that ASL consistently improves most the indicators (i.e., μAP , false positive and precision) over all the baselines. It indicates that ASL is effective, even if we simply combine it with an off-the-shelf baseline. We also note that for most baselines, ASL marginally compromises the true positive. It is because when ASL uses the predicted norm ratio to filter the “inherently similar” negative queries, it inevitably false recognizes some true edited copies as negative queries. This observation should draw two-fold attention. On the one hand, given the significant improvement on “false positive” and precision, we think the slight degradation on true positive is worthy and the overall benefit is valuable. On the other hand, the hard negative problem in NDEC

Method	$\mu AP \uparrow$	True Positive \uparrow	False Positive \downarrow	Precision \uparrow
EsViTp	45.55%	2,346	2,663	46.84%
EsViTp+ASL	48.31%	1,265	501	71.63%
CNNCL	54.30%	2,620	2,389	52.31%
CNNCL+ASL	56.95%	2,110	545	79.47%
EfNet	51.97%	2,579	2,430	51.49%
EfNet+ASL	53.81%	1,997	562	78.04%
BoT	42.06%	2,075	2,934	41.43%
BoT+ASL	45.07%	1,620	838	65.91%
D ² LV	58.87%	2,836	2,173	56.62%
D ² LV+ASL	61.28%	2,227	583	79.25%
Simple Basel.	47.00%	2,269	2,740	45.30%
Simple + ASL	49.30%	1,829	648	73.84%
Strong Basel.	61.03%	2,968	2,041	59.25%
Strong + ASL	64.09%	2,331	567	80.43%

Table 1: ASL brings general improvement over various baselines. In the first group, the baselines are employed in an off-the-shelf manner.

Method	$\mu AP \uparrow$	True Positive \uparrow	False Positive \downarrow	Precision \uparrow
Simple Basel.	47.00%	2,269	2,740	45.30%
ASL-Crop	49.10%	2,098	1,233	62.98%
ASL-Negative	48.14%	1,877	845	68.96%
ASL-Positive	48.17%	1,932	794	70.87%
Triplet	45.37%	1,774	1,200	59.65%
ASL	49.30%	1,829	648	73.84%

Table 2: The ablation studies based on our simple baseline.

is still far from being solved and calls for more efforts from the research community.

Ablation Studies

We investigate three aspects of ASL, *i.e.*, ASL under crop-to-copy scenario, ASL using hard negative pairs and the superiority of norm ratio against triplet loss in Table 2. Specifically, for the ASL using hard negative pairs, there is an important design, *i.e.*, given a hard negative training pair, the distance-based metric learning item in Eq. 2 treats it as a positive pair. We compare this design (“ASL-Positive”) against its counterpart (“ASL-Negative”), in which the metric learning item treats the negative pair as it is. From Table 2, we draw four observations as below:

First, comparing “ASL-Crop” against “Simple Basel.”, we observe ASL under crop-to-copy scenario already brings significant improvement. Specifically, it reduces the false positive matches from 2,740 to 1,233 and increases the precision from 45.30% to 62.98%.

Second, comparing “ASL-Negative” against “ASL-Positive”, we find that when learning from additional hard negative samples, “ASL-Positive” treating these samples as positive pair is better. It is because these samples, though containing no edited copies, are inherently similar. Therefore, ASL **1**) treats them as positive pairs for distance-based metric learning to make them close to each other in the deep feature space and **2**) leaves the function of distinguishing them from true edited copies with the norm ratio item.

Third, replacing the norm ratio item with a canonical triplet loss is inferior. Although using a triplet loss brings

some improvement *w.r.t.* the false positive and precision over the baseline, it is much worse than ASL with norm ratio. It validates that the symmetric-asymmetric conflict confuses distance-based metric learning and using norm ratio is critical for ASL to alleviate this conflict.

Fourth, comparing “ASL” (ASL combining crop and hard negative pairs) against other ASL editions, we observe another round of substantial improvement, *e.g.*, +10.86% precision over “ASL-Crop” and +2.97% precision over “ASL-Positive”. It indicates that both the crop-to-copy learning and the hard negative learning are beneficial and are complementary to each other.

Conclusion

This paper introduces the hard negative problem into the Image Copy Detection task and makes three contributions. First, we contribute the NDEC dataset, which highlights the hard negative distractor challenge. NDEC adds abundant hard negative queries to draw the attention to this critical challenge and meanwhile provides hard negative training pairs to promote exploring the solutions. Second, we reveal a fundamental conflict between the distance-based metric and the edited copy process, which makes the hard negative problems very difficult. Third, we propose a novel Asymmetrical-Similarity Learning to alleviate the conflict. ASL learns an asymmetric similarity metric based on norm ratio and substantially improves ICD. We also notice that ASL is still far from solving the hard negative challenge and call for more research efforts from the ICD community.

References

- Berman, M.; Jégou, H.; Vedaldi, A.; Kokkinos, I.; and Douze, M. 2019. Multigrain: a unified image embedding for classes and instances. *arXiv preprint arXiv:1902.05509*.
- Chen, D.; Zhang, S.; Yang, J.; and Schiele, B. 2020a. Norm-aware embedding for efficient person search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12615–12624.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020b. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4690–4699.
- Douze, M.; Jégou, H.; Sandhawalia, H.; Amsaleg, L.; and Schmid, C. 2009. Evaluation of gist descriptors for web-scale image search. In *Proceedings of the ACM International Conference on Image and Video Retrieval*, 1–8.
- Douze, M.; Tolias, G.; Pizzi, E.; Papakipos, Z.; Chanussot, L.; Radenovic, F.; Jenicek, T.; Maximov, M.; Leal-Taixé, L.; Elezi, I.; et al. 2021. The 2021 Image Similarity Dataset and Challenge. *arXiv preprint arXiv:2106.09672*.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33: 21271–21284.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- Hermans, A.; Beyer, L.; and Leibe, B. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- Kim, C. 2003. Content-based image copy detection. *Signal Processing: Image Communication*, 18(3): 169–184.
- Krasin, I.; Duerig, T.; Alldrin, N.; Ferrari, V.; Abu-El-Haija, S.; Kuznetsova, A.; Rom, H.; Uijlings, J.; Popov, S.; Kamali, S.; Mallocci, M.; Pont-Tuset, J.; Veit, A.; Belongie, S.; Gomes, V.; Gupta, A.; Sun, C.; Chechik, G.; Cai, D.; Feng, Z.; Narayanan, D.; and Murphy, K. 2017. Open-Images: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://storage.googleapis.com/openimages/web/index.html>*.
- Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Mallocci, M.; Kolesnikov, A.; Duerig, T.; and Ferrari, V. 2020. The Open Images Dataset V4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*.
- Liu, W.; Wen, Y.; Yu, Z.; and Yang, M. 2016. Large-margin softmax loss for convolutional neural networks. In *ICML*, volume 2, 7.
- Meng, Q.; Zhao, S.; Huang, Z.; and Zhou, F. 2021. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14225–14234.
- Papadakis, S. M.; and Addicam, S. 2021. Producing augmentation-invariant embeddings from real-life imagery. *arXiv preprint arXiv:2112.03415*.
- Pizzi, E.; Roy, S. D.; Ravindra, S. N.; Goyal, P.; and Douze, M. 2022. A Self-Supervised Descriptor for Image Copy Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14532–14542.
- Sohn, K. 2016. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems*, 29.
- Sun, X.; Qin, Y.; Xu, X.; Gong, G.; Fang, Y.; and Wang, Y. 2021. 3rd Place: A Global and Local Dual Retrieval Solution to Facebook AI Image Similarity Challenge. *arXiv preprint arXiv:2112.02373*.
- Sun, Y.; Cheng, C.; Zhang, Y.; Zhang, C.; Zheng, L.; Wang, Z.; and Wei, Y. 2020. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6398–6407.
- Wan, Y.; Yuan, Q.; Ji, S.; He, L.; and Wang, Y. 2008. A survey of the image copy detection. In *2008 IEEE Conference on Cybernetics and Intelligent Systems*, 738–743. IEEE.
- Wang, H.; Wang, Y.; Zhou, Z.; Ji, X.; Gong, D.; Zhou, J.; Li, Z.; and Liu, W. 2018. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5265–5274.
- Wang, W.; Sun, Y.; Zhang, W.; and Yang, Y. 2021a. D² 2LV: A Data-Driven and Local-Verification Approach for Image Copy Detection. *arXiv preprint arXiv:2111.07090*.
- Wang, W.; Zhang, W.; Sun, Y.; and Yang, Y. 2021b. Bag of Tricks and A Strong baseline for Image Copy Detection. *arXiv preprint arXiv:2111.08004*.
- Yokoo, S. 2021. Contrastive Learning with Large Memory Bank and Negative Embedding Subtraction for Accurate Copy Detection. *arXiv preprint arXiv:2112.04323*.
- Zhou, W.; Lu, Y.; Li, H.; Song, Y.; and Tian, Q. 2010. Spatial coding for large scale partial-duplicate web image search. In *Proceedings of the 18th ACM international conference on Multimedia*, 511–520.
- Zhou, Z.; Wang, Y.; Wu, Q. J.; Yang, C.-N.; and Sun, X. 2016. Effective and efficient global context verification for image copy detection. *IEEE Transactions on Information Forensics and Security*, 12(1): 48–63.