

Benchmarking unsupervised near-duplicate image detection

Original

Benchmarking unsupervised near-duplicate image detection / Morra, Lia; Lamberti, Fabrizio. - In: EXPERT SYSTEMS WITH APPLICATIONS. - ISSN 0957-4174. - STAMPA. - 135:(2019), pp. 313-326. [10.1016/j.eswa.2019.05.002]

Availability:

This version is available at: 11583/2733022 since: 2021-11-24T17:39:32Z

Publisher:

Elsevier

Published

DOI:10.1016/j.eswa.2019.05.002

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Elsevier postprint/Author's Accepted Manuscript

© 2019. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>. The final authenticated version is available online at:
<http://dx.doi.org/10.1016/j.eswa.2019.05.002>

(Article begins on next page)

Accepted Manuscript

Benchmarking unsupervised near-duplicate image detection


Lia Morra, Fabrizio Lamberti

PII: S0957-4174(19)30315-X
DOI: <https://doi.org/10.1016/j.eswa.2019.05.002>
Reference: ESWA 12655

To appear in: *Expert Systems With Applications*

Received date: 6 February 2019
Revised date: 15 April 2019
Accepted date: 7 May 2019

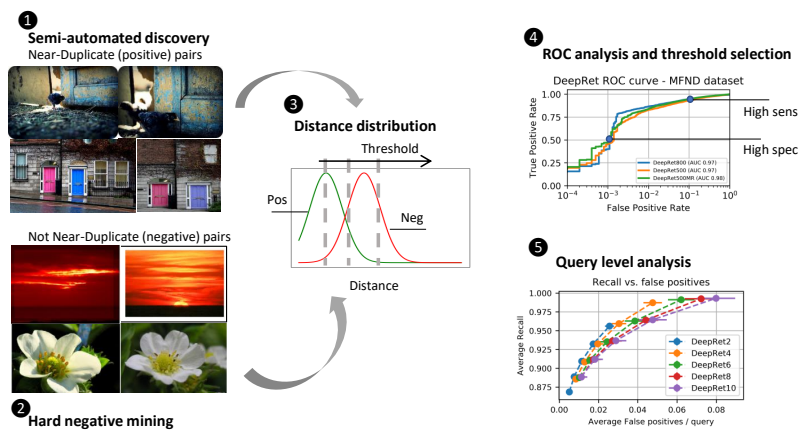
Please cite this article as: Lia Morra, Fabrizio Lamberti, Benchmarking unsupervised near-duplicate image detection, *Expert Systems With Applications* (2019), doi: <https://doi.org/10.1016/j.eswa.2019.05.002>



This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Highlights

- Unsupervised near-duplicate image detection requires high specificity up to $10^{-6} - 10^{-9}$
- Empirical comparison of CNN-based descriptors for near-duplicate image detection
- Validated, principled methodology to estimate sensitivity and estimate false alarms
- Fine-tuning CNNs for retrieval is beneficial but may suffer in specificity
- New set of annotations released for near-duplicate detection benchmarking



Benchmarking unsupervised near-duplicate image detection

Lia Morra^{a,*}, Fabrizio Lamberti^a

^a*Department of Control and Computer Engineering, Politecnico di Torino, Torino, Italy*

Abstract

Unsupervised near-duplicate detection has many practical applications ranging from social media analysis and web-scale retrieval, to digital image forensics. It entails running a threshold-limited query on a set of descriptors extracted from the images, with the goal of identifying all possible near-duplicates, while limiting the false positives due to visually similar images. Since the rate of false alarms grows with the dataset size, a very high specificity is thus required, up to $1 - 10^{-9}$ for realistic use cases; this important requirement, however, is often overlooked in literature. In recent years, descriptors based on deep convolutional neural networks have matched or surpassed traditional feature extraction methods in content-based image retrieval tasks. To the best of our knowledge, ours is the first attempt to establish the performance range of deep learning-based descriptors for unsupervised near-duplicate detection on a range of datasets, encompassing a broad spectrum of near-duplicate definitions. We leverage both established and new benchmarks, such as the Mir-Flick Near-Duplicate (MFND) dataset, in which a known ground truth is provided for all possible pairs over a general, large scale image collection. To compare the specificity of different descriptors, we reduce the problem of unsupervised detection in arbitrarily sized datasets to that of binary classification of near-duplicate vs. not-near-duplicate images. The latter can be conveniently characterized using Receiver Operating Curve (ROC). Our findings in general favor the choice of fine-tuning deep convolutional networks, as opposed to using off-the-shelf features, but differences at high specificity settings depend on the specific dataset and are often small. The best performance was observed on the MFND benchmark, achieving 96% sensitivity at a false positive rate of 1.43×10^{-6} .

Keywords: Near-duplicate detection, convolutional neural networks, instance-level retrieval, unsupervised detection, performance analysis, image forensics

*I am corresponding author

Email addresses: lia.morra@polito.it (Lia Morra), fabrizio.lamberti@polito.it (Fabrizio Lamberti)

1. Introduction

Near-duplicate (ND) image detection or discovery entails finding altered or alternative versions of the same image or scene in a large scale collection. This technique has plenty of practical applications, ranging from social media analysis and web-scale retrieval, to digital image forensics. Our work was motivated in particular by applications in the latter domain, as detecting the re-use of photographic material is a key component of several passive image forensics techniques. Examples of such applications include detection of copyright infringements (Zhou et al., 2017c; Chiu et al., 2012; Ke et al., 2004), digital forgery attacks such as cut-and-paste, copy-move and splicing (Chennamma et al., 2009; Hirano et al., 2006), analysis of media devices seized during criminal investigations (Connor & Cardillo, 2016; Battiato et al., 2014), tracing the online origin of sequestered content (de Oliveira et al., 2016; Amerini et al., 2017), and fraud detection (Li et al., 2018; Cicconet et al., 2018).

In all the above-mentioned applications, we cannot resort to standard hashing techniques, given that even minimal alterations would make different copies untraceable. Similarly, it is not possible to rely on associated text, tags or taxonomies for retrieval, as done for instance in (Gonçalves et al., 2018), since they would likely change in different sites or devices where content is used. Images may be subject to digital forgery, with parts of one or more existing images combined to create fake ones. Therefore, it is imperative to resort to content-based image retrieval techniques for the task of locating near-duplicates.

Let us consider, as a motivating example, the case of fraud detection. Many companies, like insurance ones, are relying on user-supplied photographic evidence to support business processes (Li et al., 2018). Photos of the same object or scene may be re-used multiple times to obtain an unfair advantage: such frauds are unlikely to be detected unless a largely automatic image analysis system is in place.

It should be noticed that we are adopting a very broad definition of ND, encompassing all images of the same object or scene, whereas many papers in literature restrict the definition to copies of the same digital sources that have been digitally manipulated (Connor et al., 2015; Foo & Sinha, 2007). Naturally occurring NDs, such as images of the same scene or object acquired at different times or from different viewpoints, are often more challenging to detect. However, in emerging applications such as fraud detection, which motivate our work, we do not wish to restrict ourselves to either definition: as a matter of fact, we have no reason to assume that, when constructing fraudulent claims, digital content manipulation is more likely than simply acquiring different shots of the same scene. This broad definition brings ND detection closer to the task of instance-level image retrieval, which is abundantly studied in the literature, but with a crucial difference: while the latter is usually formulated as a human-guided supervised search, the former needs as little human supervision as possible. To achieve this goal, we need to re-frame the problem from a supervised K -nearest neighbors search to an unsupervised threshold-limited search, where the distance is used as a classification function to distinguish ND from

non-ND pairs.

Realistic datasets in image forensics and fraud detection range between 10^5 – 10^7 images (Connor & Cardillo, 2016). Since the number of possible pairs grows quadratically with the dataset size, a very low false positive rate (or conversely, a high specificity) is needed to obtain a tractable number of false alarms and therefore be acceptable by the end user. For a dataset of 10^6 images, a false positive rate of 10^{-9} , which would be considered exceptionally low in many applications, would still translate to 500 false alarms.

In recent years, deep convolutional neural networks (CNNs) have shown unprecedented performance in many computer vision tasks, and content-based image retrieval is no exception. To the best of our knowledge, very few papers have exploited CNNs-based descriptor for ND detection, but if we look at the closely related task of instance-level retrieval, a consistent body of research has emerged in recent years favoring the adoption of CNN-based representation over traditional SIFT-based methods (Zheng et al., 2017). Experimental results on several benchmark datasets show that they achieve better performance, use more compact representations and are faster to compute (Zheng et al., 2017). However, given the need to re-frame the problem as an *unsupervised* threshold-limited search (where the overall performance is dominated by specificity rather than sensitivity), it is not straightforward to evaluate whether unsupervised near-duplicate search lies within the grasp of the current state-of-the-art. To the best of our knowledge, only Connor & Cardillo (2016) have previously addressed the issue of quantifying the performance of unsupervised near-duplicate detection (Connor & Cardillo, 2016), and ours is the first contribution to specifically characterize deep learning descriptors on a wide range of ND categories.

One of the underlying reasons for is certainly the lack of suitably annotated benchmark datasets, as well as of an established methodology to measure a descriptor’s performance. It is crucial that benchmarks for ND detection include a sufficiently large number of *negative queries*, i.e., images for which the absence of NDs has been established, in order to assess both specificity and sensitivity. In some cases, we can resort to digital transformations to simulate NDs, but this is not applicable to all transformations.

Instance-level retrieval benchmarks, such as the Oxford5k, Ukbench and Holidays datasets, comprise a variety of naturally occurring and challenging NDs, but are rather small scale and include only clusters of related images (Zheng et al., 2017). Recently, a new benchmark has become available to address the specific needs of ND detection: the Mir-Flickr Near Duplicate (MFND) dataset, based on the pre-existing MIR-Flickr collection (Connor et al., 2015). In this benchmark, a large number of NDs were mined using a semi-automatic procedure, so that the remaining images can be assumed to be negative queries; however, in their initial search Connor and colleagues focused on specific subclasses of NDs that limit the representativeness of this benchmark for applications such as fraud detection. Connor & Cardillo (2016) showed that the problem of unsupervised detection could thus be characterized as a binary classification problem, and we build upon their contribution for our experimental methodology.

The overarching objective of this study is to evaluate the performance of

state-of-the-art deep learning descriptors and establish a baseline against which future research can be compared. A thorough experimental comparison includes a wide range of established and emerging public benchmarks, as well as data from a real-life fraud detection case study. Our contributions can be summarized as follows:

- we compare the performance of CNN-based descriptors on the task of unsupervised near-duplicate detection, and show empirically on a variety of datasets that specificity has a large impact on the relative ranking of different descriptors;
- we extend considerably the available annotations for the MFND benchmark to obtain a large-scale benchmark which supports a wide range of ND definitions and use cases;
- finally, we extend previous work by Connor & Cardillo (2016) towards a principled evaluation methodology that captures the performance requirements of unsupervised ND discovery; **we show analytically and experimentally that by using hard negative mining, we can approximate the Area under the ROC curve (AUC) that can be used to rank the performance different descriptors.**

The rest of the paper is organized as follows: in Section 2, related work on instance retrieval and ND detection is reviewed. Section 3 introduces the datasets that are considered in the experiments. The evaluation methodology is presented in Section 4, **whereas** the experimental setup is described in Section 5. Results are presented and discussed in Sections 6 and 7, respectively.

2. Related work

2.1. Content-based image retrieval and instance-level retrieval

Content-based image retrieval systems (CBIR) are designed to retrieve semantically similar images within a database based on a specific query (e.g., by providing another image). This problem can be decomposed in two main challenges: describing image content in terms of visual features, and conducting an exact or approximate nearest neighbor search based on a distance measure (Zheng et al., 2017; Bay et al., 2008). Such features can be hand-crafted, or learned from data by using deep CNNs. In this section, we will review feature extraction techniques, and refer to existing surveys for the challenges related to feature aggregation, quantization, indexing and distance measures (Zheng et al., 2017; Zhou et al., 2017b).

2.1.1. Hand-crafted features

Global features based on the characteristics of the entire image (color, shape, texture, histogram, etc..) were extensively used in early CBIR systems. In the early 2000s, *local feature extraction* emerged as a more effective alternative, which generally involves two key steps: key interest point detection and local

region description. In the first step, key salient features in the image are identified with high repeatability, using dense sampling or ~~but~~ more commonly by detecting local extrema in the scale-space domain (e.g. Difference of Gaussians, Hessian matrix, etc.). One or multiple descriptors are then extracted from the local region centered at each interest point, usually designed to be invariant to rotation changes and robust to affine distortions, addition of noise, illumination changes, etc. The most popular local feature descriptors are SIFT and SURF (Zheng et al., 2017). SIFT-based approaches generally yield very large feature sets, in the order of the thousands per image. The Bag of Visual Word (BoVW) is the most common approach for feature reduction and quantization in CBIR and instance retrieval.

2.1.2. Deep learning approaches

Since 2015, deep learning has become the state of the art approach to CBIR (Wan et al., 2014; Gordo et al., 2016; Balntas et al., 2016; Babenko & Lempitsky, 2015; Zagoruyko & Komodakis, 2015). Deep CNNs have the distinct advantage of learning hierarchical, high-level abstractions close to the human cognition processes. Similarly to SIFT, CNNs can be trained to extract features from local regions of interests (patches), after detecting key interest points, which are then quantized e.g., using the BoVW (Zagoruyko & Komodakis, 2015; Balntas et al., 2016). Alternatively, it is possible to extract semantic-aware features from the activations of top convolutional layers in an image: it can be shown that such feature vectors can be interpreted as an approximate many-to-many region matching, without the need to explicitly extract key points, and with the advantage of obtaining faster and more compact representations. To this aim, two fundamental approaches are available. In the first approach, feature extraction is based on pre-trained models, like the VGG network trained for object recognition, alone or in combination with traditional visual features (Babenko & Lempitsky, 2015; Wan et al., 2014). In the second one, a CNN can be trained to learn a ranking function in an end-to-end fashion, mapping the input space to a target latent space such that the Euclidean distance in latent space approximates visual similarity (Gordo et al., 2016; Wang et al., 2014). In order to optimize a ranking loss, a special architecture called a Siamese network is used (Wang et al., 2014; Gordo et al., 2016, 2017). Usually, descriptors are pre-trained on ImageNet to learn image semantics, and then fine-tuned on a second training set with relevance information (Wang et al., 2014; Gordo et al., 2016).

2.2. Near-duplicate image detection

Several works have focused on near-duplicate image detection as a distinct application from content-based image retrieval (Chennamma et al., 2009; Foo & Sinha, 2007; Chum et al., 2008; Li et al., 2015; Xie et al., 2014; Hu et al., 2009; Liu et al., 2015; Xu et al., 2010; Kim et al., 2015; Battiato et al., 2014; Zhou et al., 2017c; Cicconet et al., 2018; Chen et al., 2017; Connor & Cardillo, 2016).

In order to frame our contribution with respect to previous literature, a more precise working definition of near-duplicate is needed. Given the range

of potential applications, it comes as no surprise that the definition of near-duplicate image is indeed quite varied. Starting from the work by Foo & Sinha (2007), two main sources of near-duplicates have been identified in the literature, namely *identical* and *non-identical near duplicates* (Foo & Sinha, 2007; Connor et al., 2015; Jinda-Apiraksa et al., 2013; Chen et al., 2017). Identical near-duplicates (INDs) are derived from the same digital source after applying some transformations, including cropping and rescaling, changes in image format, thumbnail resizing, *insertion* of logos or watermarks, or cosmetic changes (black & white conversion, image enhancement and so forth).

Non-identical near-duplicates (NINDs), on the contrary, are defined as images that share the same content (i.e., they depict the same scene or object), but with different illuminations, subject movement, viewpoint changes, occlusions, etc. (Foo & Sinha, 2007; Jinda-Apiraksa et al., 2013). Detecting NINDs is deemed more challenging than INDs, and their definition is more subjective (Jinda-Apiraksa et al., 2013); for these reasons, many authors have mostly targeted INDs.

Depending on the type of ND targeted and the level of transformation involved, most papers in literature have either focused on global features or on SIFT features combined with BoWV quantization. Global features have been mostly used for IND detection (Connor & Cardillo, 2016; Chen et al., 2017; Li et al., 2015). Local descriptors, such as SIFT features combined with BoWV quantization, allow detecting more aggressive alterations (including NINDs), sub-image retrieval or image forgery (e.g. copy-move attacks). Local descriptors are prone to false positive matches, as they do not take into account spatial coherence; to reduce false alarms, some authors have proposed pruning techniques to improve specificity and scalability (Foo & Sinha, 2007; Liu et al., 2015), whereas other authors have focused on post-query verification (Zhou et al., 2017c; Hu et al., 2009; Xu et al., 2010).

From an evaluation point of view, many papers framed the problem of ND detection as a supervised K -nearest neighbor search, and few papers have addressed the issue of quantifying the specificity of descriptors when performing unsupervised, threshold-limited near-duplicate discovery (Connor & Cardillo, 2016; Chen et al., 2017; Kim et al., 2015). The most relevant prior work is that by Connor et al, who proposed a method to evaluate the specificity of ND detectors and choosing the optimal distance threshold, based on Receiver Operating Curve (ROC) analysis (Connor & Cardillo, 2016). A more in-depth analysis of this methodology, and the extensions that we propose, is available in Section 4. Other authors have used small test sets to establish the optimal threshold, that was subsequently applied to a larger dataset (Chen et al., 2017). For instance, Chen et al. used bloom filtering and range queries to detect duplicate images under scaling, watermarking and format change transformation (Chen et al., 2017); for evaluation purposes, they estimated the percentage of false and correct rejections, as well as precision and recall curves, on a smaller *ease* dataset.




Figure 1: Examples of near duplicates pairs of varying complexity from the four datasets included in the comparison. For the CLAIMS dataset, difficulty was evaluated subjectively by one rater, whereas for the California-ND, it was established based on the agreement between 10 independent raters.

Dataset	Size	IND clusters (pairs)	NIND clusters (pairs)
CLAIMS	201,961	NA	1037 (1,475)
MFND	1,000,000	3,825 (4,672)	10,454 (18,299)
California-ND	701	NA	107 (4,609)
Holidays	1,491	NA	500 (2,072)

Table 1: Comparison of the benchmark datasets. Related IND or NIND pairs were grouped into clusters; the average number of images per cluster ranges from 2.07 to 6.55. IND and NIND pairs are counted separately, where applicable.

3. Datasets

The experimental results presented in this paper were based on four image collections, including a private dataset (Section 3.1) and three publicly available benchmarks (Sections 3.2, 3.3 and 3.4). The CLAIMS dataset was collected for insurance purposes, and therefore constitutes a realistic case study for fraud detection applications. The MFND dataset (Connor et al., 2015), based on the MIR-Flickr image retrieval benchmark (Huiskes & Lew, 2008), contains a variety of both INDs and NINDs. Both the California-ND and Holidays datasets contain personal holiday photos and, while much smaller in size, include several challenging NIND examples (Jegou et al., 2008; Jinda-Apiraksa et al., 2013). Examples of ND pairs of different complexity from the various databases are given in Fig. 1, whereas a summary of the datasets characteristics is reported in Table 1.

3.1. CLAIMS dataset

The CLAIMS dataset includes a variety of indoor and outdoor scenes, mostly from residential and commercial buildings. It contains a total of 201,961 images

235 coming from 22,327 claims. Image subsets that are associated with a given claim generally include images of the same scenes or objects, representing a source of relatively rapidly identifiable NINDs. More infrequently, images from different claims may also represent the same scene or object. This dataset contains both IND (e.g. insertion of small captions or logos, changes in aspect ratio, format
240 change, compression, etc.) and NIND clusters (e.g. sequential snapshots of the same scene, viewpoint changes, etc.).

The collection was annotated to generate both positive queries (i.e., with known NDs) and negative queries (i.e., for which absence of NDs was confirmed). NDs were annotated following a semi-manual procedure, in which a set of claims
245 was randomly selected. For each claim, all potential image pairs were generated and the ND pairs were manually selected. Connected pairs of NDs from the same claim were grouped to form clusters.

Non-near duplicate (NND) pairs were randomly extracted following a hard negative mining strategy (see Section 4.1 for details). The results were visually
250 inspected obtaining additional 103 near-duplicate pairs. The final annotated set included 1,475 ND pairs, forming 1,037 distinct clusters; the average number of images per cluster is 2.2.

3.2. *MirFlickr Near Duplicate*

The MIR-Flickr Near Duplicate (MFND) collection is a recent revisitation
255 of the MIR-Flickr image retrieval benchmark (Huiskes & Lew, 2008). Connor and colleagues observed a significant number of NDs in this one million image collection, which were semi-automatically retrieved using different ND finders (Connor et al., 2015). We have expanded their annotations by adopting a broader definition of ND, as well as using different descriptors.

260 The first MFND annotation was generated using a set of four global descriptors (based on MPEG-7 and perceptual Hashing global features) and five distance measures, which were combined to form different similarity functions (Connor et al., 2015). A threshold-limited nearest-neighbor search was conducted using approximated metric search techniques, yielding a few thousand
265 potential ND pairs for every function. We have expanded this annotation by using the three CNN-based descriptors included in this study, and the Euclidean distance. Several threshold-limited, K -nearest neighbor searches were performed (with $K=5$ and $K=1$), yielding a few hundred thousands potential ND pairs which were visually inspected. Exact duplicates were eliminated based
270 on the MD5 hash.

Each of the resulting image pairs was manually assigned to one of three categories, IND, NIND or other, following the categorization illustrated in Section 2.2 (Connor et al., 2015). The strength of this methodology is that it minimizes biases with respect to the images in the collection, as well as to the method
275 with which the near-duplicates have been detected. We assumed, as in previous work (Connor et al., 2015), that both IND and NIND relations are transitive, allowing the identification of clusters of images that share the same content. The resulting clusters were also visually inspected for consistency.

As for the CLAIMS collection, NND pairs were generated through a hard
 280 negative mining procedure; results were visually inspected identifying 120 additional NIND pairs.

The available annotations were thus substantially extended from 1,958 to
 3,825 IND clusters (4,672 vs. 2,407 pairs) and from 379 to 10,454 NIND clusters
 (18,299 pairs). Many new IND pairs detected were subject to digital content
 285 manipulations, cropping or color alterations; we found that CNN-based descriptors
 were particularly robust to colorization techniques. A total of 30,925 images
 were found to have at least one IND or NIND in the collection, with a mean
 cluster size of 2.2.

3.3. California-ND

290 The California-ND collection comprises 701 photos taken from a real user's
 personal photo collection (Jinda-Apiraksa et al., 2013). It includes many chal-
 lenging NIND cases, without resorting to artificial image transformations. To
 account for the intrinsic ambiguity of NIND definition, the collection was man-
 ually annotated by 10 different observers, including the photographer himself.
 295 Instructions such as "If any two (or more) images look similar in visual ap-
 pearance, or convey similar concepts to you, label them as near-duplicates."
 were given to the raters. Out of 245,350 unique possible combinations, 4,609
 image pairs were identified as ND by at least one subject; notably, in 82% of
 the cases raters disagreed to some extent on whether or not a pair of images
 300 should be considered ND. The image pairs form 107 clusters of NIND images,
 where each cluster contains on average 6.55 images; the ND pairs were grouped
 assuming that the ND relationship is transitive (which is not generally the case,
 but seemed reasonable in this particular situation).

3.4. Holidays

305 The INRIA Holidays dataset (Jegou et al., 2008), a popular benchmark for
 instance retrieval, is mainly composed by the authors' personal holidays photos.
 The images, all high resolution, include a large variety of scene types (natural,
 man-made, water and fire effects, etc). Images were grouped by the authors in
 500 disjoint image clusters, each representing a distinct scene or object, for a
 310 total of 1,491 images and an average cluster size of 2.98 images. From the 500
 clusters, a total of 2,072 ND pairs, mostly NIND, can be identified.

4. Performance evaluation

315 In this section, we illustrate the performance metrics and protocol used to
 evaluate the specificity and sensitivity of unsupervised ND detection. In the
 first battery of test, we extended the work of Connor and colleagues (Connor
 & Cardillo, 2016), reducing the problem of unsupervised discovery to that of
 binary classification; ROC analysis can be used to measure the ability of a ND
 detector to distinguish ND pairs from visually similar examples, as detailed in
 Section 4.1. The second battery of test involves estimating the average false

positive rates generated by a negative query, and is explained in Section 4.3; the relationship between these two performance measures is also explored. An overview of the methodology is presented in Fig. 2.

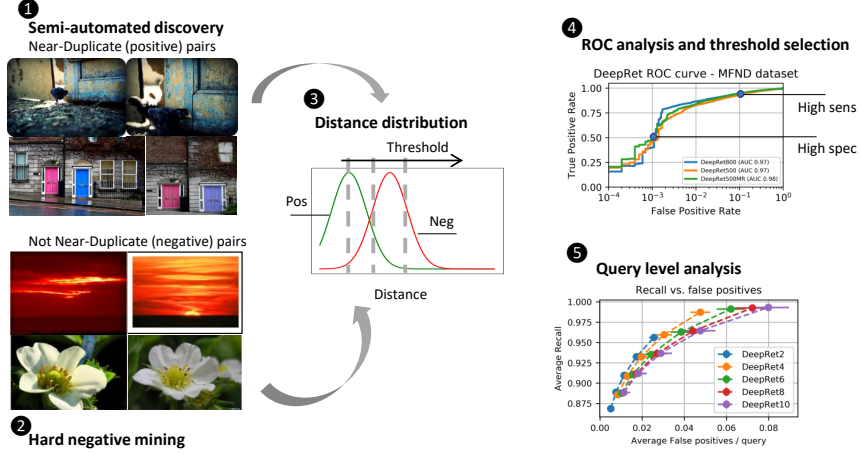


Figure 2: Methodology employed to calculate the performance on the MFND benchmark. First, all near-duplicate pairs are discovered through a semi-supervised search technique (step 1). On the remainder of the collection, hard negative mining is used to identify hard samples of visually similar, but not near-duplicate pairs (step 2); crucially, this step needs to be repeated for each descriptor. Using the distance as a classification function (step 3), ROC analysis can be used to characterize the ability of the descriptor to distinguish near-duplicates from not-near duplicate pairs. From ROC analysis, suitable thresholds on the distance can be selected based on the application requirements. Performance at query time can be thus be reliably estimated (step 5).

4.1. ROC analysis

As suggested by Connor & Cardillo (2016), a near-duplicate finder can be modeled as a positive numeric function D over any two image descriptors x and y , where normally D will be a proper distance metric. To run an unsupervised search, it is necessary to use D as a classification function over images pairs, which without loss of generality can be achieved by choosing a threshold t :

$$D_t(x, y) = D(x, y) < t \quad (1)$$

The problem of unsupervised discovery can be characterized as finding the near-duplicate intersection of two image sets $X \cap_{ND} Y$, that is the set of pairs of images from sets X and Y that satisfy the conceptual near-duplicate relation ND (Connor & Cardillo, 2016). If $Sens(t)$ and $Spec(t)$ are the sensitivity and specificity of $D_t(x, y)$, the number of true positive (TP) matches will be

$$TP(t) = Sens(t) |X \cap_{ND} Y| \quad (2)$$

and the number of false positives (FP) will be

$$FP(t) = (1 - Spec(t)) |X| |Y| \quad (3)$$

assuming that $|X \cap_{ND} Y| \ll |Y|$. In our setting, $|X| = K$ is the number of query images and $|Y| = M$ is the size of the collection. Another useful figure to define is the number of false positives / query image, which can be computed as

$$FP_q(t) = (1 - Spec(t)) M \quad (4)$$

Given a set of ND and NND pairs, the sensitivity (or recall) and the specificity of a ND detector can be estimated as follows:

$$Sens(t) = \frac{\text{No. of correctly identified ND pairs}}{\text{Total no. of ND pairs}} \quad (5)$$

$$Spec(t) = \frac{\text{No. of correctly identified NND pairs}}{\text{Total no. of NND pairs}} \quad (6)$$

Both quantities are function of the threshold t , and the overall performance can be characterized by ROC analysis.

4.2. Hard negative mining

In a realistic dataset the pool of NND images is very large, compared to the number of ND pairs - it is not **feasible** to evaluate all possible pairs. Hard negative mining extracts a compact set of NND from a large image collection, starting from a subset of randomly selected query images, for which we can assume that a near-duplicate match does not exist in the collection. For each query image, the pairwise distances between the query images and all the other images in the collection are calculated, and the most "difficult" examples are selected.

Starting from a random sample of query images, two hard negative mining strategies were considered:

- the nearest neighbor for each query is selected (*hn1*);
- the K -nearest neighbors for all **queries** are retrieved and sorted; the most difficult pairs (i.e., those with the smallest distances) are then selected (*hn2*).

Notably, the hard negative mining procedure depends on the relative ranking of the images, and hence has to be repeated for each descriptor and for each distance formulation.

The distances of the hard negatives are among the smallest of the $K \times M$ distances measured, where K is the number of query images and M is the dimension of the dataset: if a distance threshold exists such that all the "difficult" pairs are successfully identified, **then** we can assume that all potential NND pairs in the collection will be identified as well. For instance, for the CLAIMS dataset $K = 4400$ and $M = 80,000$, yielding a specificity of

350 $1 - 1/353,000,000 = 1 - 2.83 \times 10^{-9}$, which is the smallest possible sensitivity that can be measured in this setting.

The specificity measured on the hard negative samples can be used to approximate the true specificity that would be observed at large. For instance, a .9 specificity (.1 FP rate), would allow to successfully discard $0.9M$ NND pairs; however, it would also fail to discard at least $0.1K$ NND pairs, and hence would correspond to a specificity of at most $1 - 1.25 \times 10^{-6}$. In this way, it is possible to estimate a lower bound on the amount of FPs generated on datasets of arbitrary size.

4.2.1. Area under the ROC curve

360 The AUC is a common summary metrics that quantifies the global performance of a classifier (Bradley, 1997). We estimated the AUC for each descriptor and 95% confidence intervals were calculated under the normal assumption according to Hanley & McNeil (1982).

365 Since the ROC is calculated on a subset of all possible negative pairs, the resulting AUC will be an approximation of the true AUC if all pairs were taken into account. We will refer in the following to AUC_{hn1} and AUC_{hn2} to denote the AUC calculated on pairs extracted using hard negative mining strategies $hn1$ and $hn2$, respectively.

Let p_1, \dots, p_{N^+} be ND pairs (i.e., positive samples) and n_1, \dots, n_{N^-} be all the NND pairs (i.e., negative samples), where in our case $N^- = K \times M$. The AUC can be expressed as a sum of indicator functions (Hanley & McNeil, 1982):

$$AUC = \frac{1}{N^+N^-} \sum_{i=1}^{N^+} \sum_{j=1}^{N^-} \mathbb{1}_{f(p_i) > (n_j)} \quad (7)$$

370 where $f(\cdot)$ is a scoring function which, in our case, is the distance between the descriptors of the two images in the pair ¹. For simplicity, we omit $f(\cdot)$ from the notation in the rest of the paper.

Let n_l, \dots, n_{H^-} be the hard negative NND pairs, where $H^- \ll N^-$. The estimated AUC can be expressed as follows:

$$AUC_{hn} = \frac{1}{N^+H^-} \sum_{i=1}^{N^+} \sum_{l=1}^{H^-} \mathbb{1}_{p_i > n_l} \quad (8)$$

Since in general

$$\mathbb{1}_{p_i > n_j} \leq \mathbb{1}_{p_i > n_l} \quad \forall n_j \in N^- - H^- \quad \forall n_l \in H^- \quad (9)$$

it can be demonstrated that for both hard negative mining strategies AUC_{hn} is an upper bound for the true AUC. It follows that the most appropriate choice

¹We follow here the notation normally used in ROC literature where the positive samples are expected to be scored higher than negative samples, whereas in our case the scoring function is a distance and pairs with lower distance would be scored higher

would be to use the strategy that provides the tighter bound. Analytical proof
 375 is provided in Appendix A.

4.3. Range query performance

ROC analysis does not directly represent the observed system performance, which also depends on the distribution of the type of images, the size of the clusters, and so forth. We analyzed an alternative performance measure, obtained
 380 by simulating the case of a single query image x compared against a collection of images Y , which is a special case of the general problem of near-duplicate detection described in Section 4.1. An unsupervised, threshold-limited range search is conducted to retrieve a list of potential near-duplicates, and used to estimate the number of FPs/query or $FP_q(t)$. In practice, it is convenient to restrict the
 385 search to the K -nearest neighbors in order to cap the number of FPs/query to a reasonable number. The proposed experimental setup executes a number of positive queries (i.e., images with one or more known ND), and negative queries (i.e., images that have no expected NDs), over a dataset constructed as follows:

- positive queries were derived from the clusters of ND images, where the
 390 first image are used as queries and the rest are inserted in the database, as normally done for Holidays and other image retrieval benchmarks;
- negative query images were selected from the NND pairs, and a set of distractors are used to evaluate specificity; in practice, we use for convenience
 the same image pool used for hard negative mining.

For varying values of the threshold $t = T_i$ on the distance measure, we
 395 compared the *average recall*, calculated over all positive queries, and the *average number of FPs/query*, calculated over all negative queries. Note that average recall is different from pair-wise sensitivity used in ROC analysis, as each query may contain multiple pairs of varying “difficulty”. The FPs/query depend on
 400 the size of the dataset and the specificity as predicted by Equation 3.

5. Experimental setup

In this section, a detailed analysis of the experimental setup is given concerning the descriptors selection, their implementation, and the hard negative mining parameters.

5.1. Descriptors

Two sets of descriptors were compared in this work: global descriptors, and CNN-based descriptors; for the latter, we compared examples of the two main approaches (aggregation of raw deep convolutional features without embedding and Siamese networks) described in Section 2.

410 Among global descriptors, *GIST* (Oliva & Torralba, 2001) was selected based on previous results on the MFND collection (Connor & Cardillo, 2016). The *GIST*, or spatial envelope, is a bio-inspired feature that simulates human visual

perception to extract rough but concise context information (Oliva & Torralba, 2001). The input image is decomposed using spatial pyramid into N blocks, filtered by a number of multi-scale, multi-orientation Gabor filters (4 scales, 8 orientations per scale), and then summarized by a feature extractor that captures the "gist" of the image, handling translational, angular, scale and illumination changes. We experimented with perceptual Hashing, however the results are not reported as they were generally very poor.

The *SPoC* (*Sum-Pooled Convolutional*) descriptor was initially proposed by Babenko & Lempitsky (2015). The features are extracted from the top convolutional layer of a pre-trained neural network and spatially aggregated using sum pooling. The length of the feature vector will thus be equal to the depth of the final convolutional layer (usually in the order of the hundreds). Best results were obtained extracting features after ReLU activation, confirming previous findings (Babenko & Lempitsky, 2015). PCA whitening and compression is applied, and the vectors are normalized to unit length (L2 normalization).

The *R-MAC* architecture, proposed by Tolias et al. (2015) builds a compact feature vectors by encoding several image regions in a single pass. First, sub-regions are defined using a fixed grid over a range of progressively finer scales l ranging from 1 to L ; then, max-pooling is used to extract features from each individual region. Each region feature vector is post-processed with PCA-whitening and L2 normalization. Finally, the regional feature vectors are summed into a single image vector, which is again L2 normalized.

The *DeepRetrieval* architecture, proposed by Gordo et al. (2016), employs a Siamese network to learn a ranking function based on the triplet loss function. Let I_q be a query image with descriptor q , I_+ be a relevant image with descriptor d_+ , and I_- be a non-relevant image with descriptor d_- . The ranking triplet loss is defined as

$$L(I_q, I_+, I_-) = \frac{1}{2} \max(0, m + \|q - d_+\|^2 - \|q - d_-\|^2) \quad (10)$$

where m is a scalar that controls the margin. At test time, the features are extracted from the top convolutional layer and aggregated using sum-pooling and normalization. The Deep Retrieval architecture includes an additional proposal network, similar to the R-MAC grid network, so that the features are calculated on several potential regions of interest, as opposed to the entire image. The Deep Retrieval network is pre-trained on the Landmarks dataset (Babenko et al., 2014).

5.2. Implementation

The tested descriptors, and related parameters, are summarized in Table 2. For SPoC and GIST, images were resized to 512×512 , whereas for Deep Retrieval images were rescaled so that the longest side is equal to S .

A Python re-implementation of the original Matlab code by Olive and Torralba was used for GIST, after converting images to grayscale². The SPoC

²<http://people.csail.mit.edu/torralba/code/spatialenvelope/>

descriptor was computed from pre-trained networks architectures such as VGG (Simonyan & Zisserman, 2014) and Residual Networks (He et al., 2016). We included both models pre-trained on ImageNet (Deng et al., 2009), as well as on the Places205 or Places365 datasets (Zhou et al., 2014, 2017a), and on a hybrid dataset including images from both ImageNet and Places. **The R-MAC descriptor was re-implemented in Python based on the original Matlab implementation by the authors³; R-MAC was calculated only for the ResNet101 architecture.** All networks were available in Caffe; pre-trained models were downloaded from the Caffe Model Zoo⁴, or were made available by the authors for the Places dataset^{5,6}.

We trained the PCA parameters, without reducing the number of features, on a representative set of the collection (95,230 for CLAIMS and 100,000 for MFND), which was not used in testing. The parameters trained on the MFND collection were also used for the Holidays and California-ND collection; although the image characteristics in the two datasets is different, this is consistent with previous works in literature (Jegou et al., 2008).

The FAISS library (Johnson et al., 2017), specifically the flat index with L2 exact search, was used to index the collection and perform range queries. All experiments were run on a system with a i7-7700 CPU @3.60GHz and GTX1080Ti nVIDIA GPU.

Descriptor	Label	Size	Parameters
GIST (Oliva & Torralba, 2001)	GIST4	512	number of blocks = 4
GIST (Oliva & Torralba, 2001)	GIST8	512	number of blocks = 8
Deep Retrieval (Gordo et al., 2016)	DeepRet800	2048	ResNet101, Fine-tuned on Landmarks dataset, S = 800, no multiresolution
Deep Retrieval (Gordo et al., 2016)	DeepRet500	2048	ResNet101, Fine-tuned on Landmarks dataset, S = 500, no multiresolution
Deep Retrieval (Gordo et al., 2016)	DeepRet500MR	2048	ResNet101, Fine-tuned on Landmarks dataset, S = 500, multiresolution (2)
SPOC (Babenko & Lempitsky, 2015)	SP_VGG19IN	512	VGG19, Trained on ImageNet
SPOC (Babenko & Lempitsky, 2015)	SP_VGG16PL	512	VGG16, Trained on Places205
aSPOC (Babenko & Lempitsky, 2015)	SP_VGG16HY	512	VGG16, Trained on Hybrid (Places205 & ImageNet) dataset
SPOC (Babenko & Lempitsky, 2015)	SP_ResNet101IM	2048	ResNet101, Trained on ImageNet dataset
SPOC (Babenko & Lempitsky, 2015)	SP_ResNet152IM	2048	ResNet152, Trained on ImageNet dataset
SPOC (Babenko & Lempitsky, 2015)	SP_ResNet152HY	2048	ResNet152, Trained on Hybrid (Places365 & ImageNet) dataset
R-MAC Tolias et al. (2015)	RMAC	2048	ResNet101, Trained on ImageNet dataset, L=2

Table 2: Synthetic description of the descriptors used.

5.3. Hard negative mining

For each dataset (CLAIMS and MFND), we randomly selected a set of negative query images (4,500 and 5,000, respectively), and a larger pool for mining (80,000 and 70,000 images, respectively), after excluding IND or NIND pairs and images used to train the PCA parameters. The two hard negative mining strategies introduced in Section 4.2, were compared: in *hn2*, the 10 nearest neighbors were found for each query images, and then the 10,000 most difficult pairs were selected. The number of samples was increased for *hn2* to account for

³<https://github.com/gtolias/rmac>

⁴<https://github.com/BVLC/caffe/wiki/Model-Zoo>

⁵<https://github.com/CSAILVision/places365>

⁶<http://www.europe.naverlabs.com/Research/Computer-Vision/Learning-Visual-Representations/Deep-Image-Retrieval>

images that belong to more than one pair, which we observed experimentally, and ensure sufficient diversity. For GIST and Deep Retrieval, the hard negative mining procedure was calculated for one parameter set, to reduce the computational cost. For the Holidays and California-ND datasets, we used the NND pairs mined for MFND; this is consistent with previous works that have used the same collection as distractors for large scale retrieval testing (Jegou et al., 2008). For both datasets, the hard negative mining procedure was repeated for a large number of descriptors (8 for MFND, and 13 for the CLAIMS dataset), and the results were visually inspected for the presence of near duplicates. A total of 101 (2.2%) and 121 (2.4%) ND pairs were found for the CLAIMS and MFND datasets, respectively, and their labels were changed accordingly.

6. Results

In this section, results of the ROC analysis are compared across different descriptors (Section 6.1) and different datasets (Section 6.2). Finally, the performance obtained from random queries is analyzed and compared with that predicted by ROC analysis in Section 6.3.

Descriptor	CLAIMS	MFND		California-ND	Holidays
	AUROC	AUROC-IND	AUROC-all	AUROC	AUROC
GIST4	0.397 (0.381 – 0.413)	0.808 (0.799 – 0.817)	0.51 (0.491 – 0.509)	0.598 (0.587 – 0.610)	0.365 (0.351 – 0.378)
GIST8	0.45 (0.433 – 0.467)	0.854 (0.847 – 0.862)	0.561 (0.552 – 0.569)	0.696 (0.685 – 0.706)	0.493 (0.478 – 0.508)
DeepRet800	0.891 (0.88 – 0.903)	0.994 (0.993 – 0.996)	0.983 (0.981 – 0.984)	0.929 (0.924 – 0.934)	0.744 (0.731 – 0.758)
DeepRet500	0.88 (0.868 – 0.892)	0.992 (0.991 – 0.994)	0.979 (0.978 – 0.981)	0.917 (0.911 – 0.923)	0.748 (0.734 – 0.761)
DeepRet500MR	0.891 (0.88 – 0.903)	0.992 (0.99 – 0.994)	0.983 (0.982 – 0.984)	0.938 (0.933 – 0.943)	0.789 (0.777 – 0.802)
SP_VGG19_IM	0.391 (0.375 – 0.407)	0.934 (0.929 – 0.940)	0.904 (0.901 – 0.908)	0.88 (0.873 – 0.887)	0.671 (0.656 – 0.685)
SP_VGG16_PL	0.446 (0.429 – 0.463)	0.907 (0.901 – 0.914)	0.881 (0.877 – 0.886)	0.909 (0.903 – 0.915)	0.675 (0.66 – 0.689)
SP_VGG16_HY	0.418 (0.401 – 0.434)	0.929 (0.923 – 0.934)	0.906 (0.903 – 0.910)	0.896 (0.89 – 0.903)	0.697 (0.683 – 0.711)
SP_ResNet101IM	0.518 (0.501 – 0.536)	0.961 (0.957 – 0.965)	0.941 (0.938 – 0.943)	0.93 (0.925 – 0.936)	0.776 (0.763 – 0.789)
SP_ResNet512IM	0.522 (0.505 – 0.540)	0.963 (0.959 – 0.967)	0.944 (0.941 – 0.946)	0.921 (0.916 – 0.927)	0.78 (0.767 – 0.793)
SP_ResNet512HY	0.459 (0.442 – 0.476)	0.924 (0.918 – 0.930)	0.916 (0.913 – 0.920)	0.866 (0.859 – 0.874)	0.737 (0.723 – 0.751)
RMAC	0.323 (0.308 – 0.338)	0.99 (0.988 – 0.992)	0.945 (0.942 – 0.947)	0.88 (0.873 – 0.888)	0.737 (0.723 – 0.751)
Average	0.549	0.937	0.870	0.863	0.684

Table 3: Area under the ROC curve (AUC) with 95% confidence intervals. The NND pairs are extracted using the first hard negative mining strategy (*hn1*). For the MFND dataset, the AUC is calculated separately for both IND and NIND pairs (AUROC-all), and for IND pairs vs. NND pairs; in the latter case, NIND are not counted as either FP or TP. The average AUC across all descriptors provides a semi-quantitative estimate of the "difficulty" of each dataset.

6.1. ROC analysis

The Area under the ROC curve (AUC) for all descriptors, along with 95% confidence intervals, is reported in Tables 3 and 4. There is a large difference in estimated performance depending on the hard negative mining technique employed, with *hn1* yielding optimistically biased estimates. This is most evident for the CLAIMS dataset, which contains a larger number of visually similar, but not duplicate images.

For IND detection, the difference between global features and deep learning based features is less pronounced. The results are lower than previously reported in literature, because the IND dataset has been significantly expanded, and the

Descriptor	CLAIMS	MFND			California-ND	Holidays
	AUROC	AUROC-IND	AUROC-all	AUROC	AUROC	AUROC
GIST4	0.108 (0.101 – 0.114)	0.706 (0.696 – 0.715)	0.317 (0.31 – 0.323)	0.398 (0.389 – 0.408)	0.178 (0.17 – 0.186)	
GIST8	0.121 (0.114 – 0.128)	0.756 (0.747 – 0.765)	0.346 (0.339 – 0.352)	0.462 (0.452 – 0.472)	0.243 (0.234 – 0.253)	
DeepRet800	0.381 (0.366 – 0.395)	0.99 (0.988 – 0.992)	0.969 (0.967 – 0.970)	0.929 (0.924 – 0.934)	0.628 (0.614 – 0.642)	
DeepRet500	0.428 (0.412 – 0.443)	0.987 (0.985 – 0.989)	0.962 (0.96 – 0.964)	0.917 (0.911 – 0.923)	0.614 (0.6 – 0.628)	
DeepRet500MR	0.46 (0.444 – 0.476)	0.987 (0.985 – 0.989)	0.97 (0.968 – 0.971)	0.938 (0.933 – 0.943)	0.676 (0.663 – 0.690)	
SP_VGG19_IM	0.24 (0.229 – 0.251)	0.882 (0.876 – 0.889)	0.82 (0.816 – 0.825)	0.787 (0.779 – 0.796)	0.52 (0.507 – 0.534)	
SP_VGG16_PL	0.288 (0.274 – 0.302)	0.866 (0.859 – 0.873)	0.821 (0.817 – 0.826)	0.859 (0.851 – 0.866)	0.577 (0.563 – 0.591)	
SP_VGG16_HY	0.267 (0.255 – 0.279)	0.881 (0.874 – 0.887)	0.834 (0.83 – 0.838)	0.814 (0.806 – 0.822)	0.563 (0.55 – 0.577)	
SP_ResNet101IM	0.397 (0.382 – 0.412)	0.943 (0.939 – 0.948)	0.911 (0.908 – 0.914)	0.892 (0.885 – 0.898)	0.685 (0.671 – 0.698)	
SP_ResNet512IM	0.396 (0.381 – 0.411)	0.947 (0.943 – 0.952)	0.917 (0.914 – 0.920)	0.885 (0.878 – 0.891)	0.694 (0.68 – 0.707)	
SP_ResNet512HY	0.327 (0.313 – 0.340)	0.881 (0.874 – 0.887)	0.866 (0.862 – 0.870)	0.784 (0.776 – 0.793)	0.627 (0.613 – 0.641)	
RMAC	0.336 (0.322 – 0.350)	0.985 (0.983 – 0.988)	0.917 (0.914 – 0.920)	0.825 (0.817 – 0.833)	0.641 (0.627 – 0.655)	
Average	0.312	0.901	0.804	0.791	0.554	

Table 4: Area under the ROC curve (AUC) with 95% confidence intervals. The NND pairs are extracted using the **red** hard negative mining strategy (*hn2*). For the MFND dataset, the AUC is calculated separately for both IND and NIND pairs (AUROC-all), and for IND pairs vs. NND pairs; in the latter case, NIND are not counted as either FP or TP. The average AUC across all descriptors provides a semi-quantitative estimate of the “difficulty” of each dataset.

new pairs include transformations to which previous descriptors were less robust.

The DeepRetrieval architecture generally outperforms SPoC for all datasets, despite being trained on a different dataset (Landmarks) with no fine-tuning.

The actual gap in performance is very low for the MFND dataset, and increases for other datasets, with CLAIMS exhibiting the highest gap. It is worth noting, however, that DeepRetrieval is also more prone to FPs due to visually similar images, and the performance estimates are more sensitive than for SPoC to the hard negative mining strategy. **The R-MAC descriptor performs slightly better than SPoC for the MFND dataset, and slightly worse for the other datasets.**

The performance of the SPoC descriptor strongly depends on the network architecture, with Residual Networks consistently outperforming VGG on all datasets. The dataset on which the network was trained has instead a limited impact, possibly due to the effect of PCA whitening. ROC curves for selected descriptors and datasets are reported in Fig. 3. The remaining ROC curves are available as supplementary material. On the MFND collection, the best performance is obtained by the DeepRet descriptor, retrieving 96% of the true positives at a FP rate of 1.43×10^{-6} .

6.2. Dataset comparison

In order to better highlight differences between the datasets, we computed the false FP and TP rate w.r.t. the distance threshold for each dataset and for each of the two best performing descriptors, as detailed in Fig. 4.

Not surprisingly, INDs are more easily detected than NINDs. The CLAIMS dataset contains the most challenging near duplicates, closely followed by the Holidays dataset. Given the annotation procedure followed for the MFND benchmark, it is possible that the NIND examples are skewed towards examples that are more easily detected using the present descriptors, and future experiments will likely find new examples. Examples of ND pairs that were poorly scored are reported in Fig. 5; empirically, large changes in viewpoint appear among the most challenging differences.

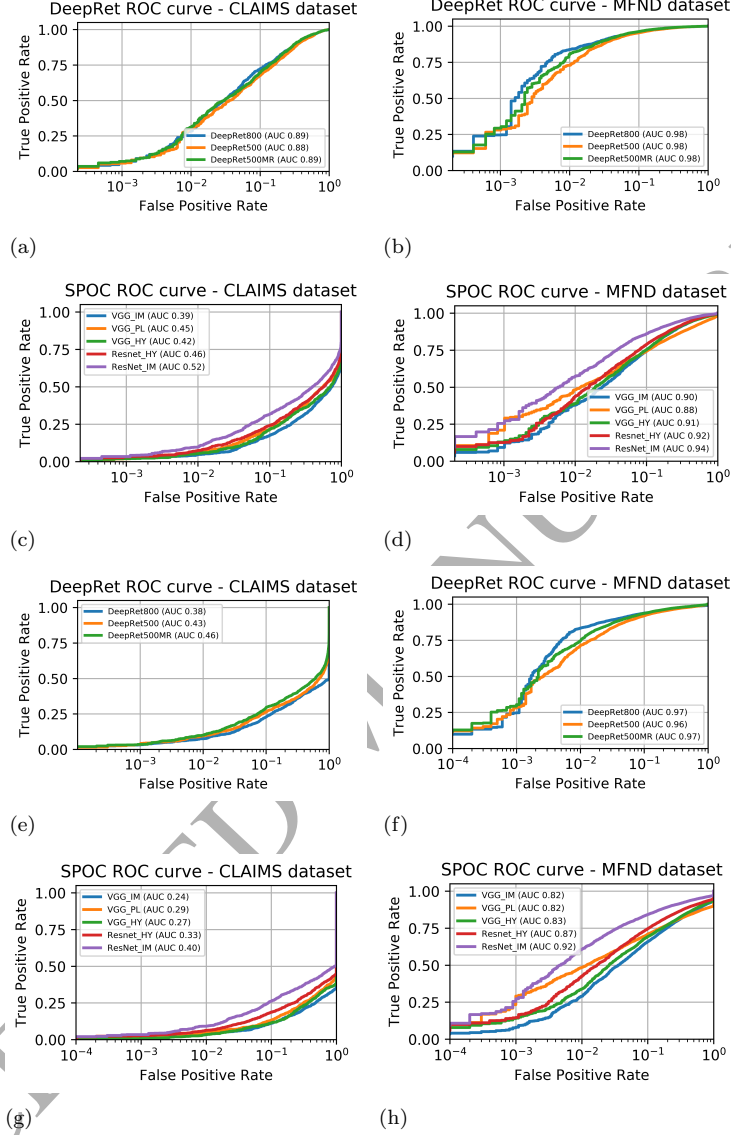


Figure 3: ROC curves for the Deep Retrieval descriptors for hard negative mining strategy *hn1* (a-d) and *hn2* (e-h) respectively. A logarithmic scale was used for the FP rate axis to highlight low values in the 0.01 – 0.1 range. Since the NND pairs were extracted using a hard negative mining strategy, a 0.1 FP rate corresponds to a projected minimum FP of 1.25×10^{-6} and 1.43×10^{-6} for the CLAIMS and MFND datasets, respectively.

We also compared FP rates on the MFND and CLAIMS datasets with the two hard negative mining strategies. For *hn1*, MFND appears to be more dif-

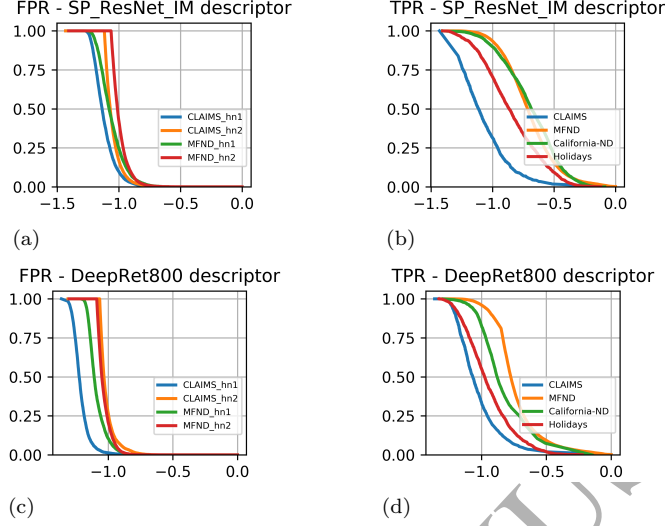


Figure 4: Comparison of TP rate and FP rates across different datasets for selected descriptors: SP_ResNet101_ImageNet (a-b), and DeepRet800 (c-d). For FPs, pairs selected with both hard negative mining strategies *hn1* and *hn2* are separately plotted.

535 difficult than CLAIMS, whereas for *hn2* the two datasets are quite comparable for both descriptors. Given a random query image, it is more likely to find a similar image for MFND than CLAIMS, but CLAIMS contains larger clusters of images that are both semantically and visually similar, as is likely going to be the case for any dataset that comes from a focused domain. Examples of hard negatives (*hn2*) for both datasets are reported in Fig. 5.

6.3. Query performance analysis

540 In this section, the two best performing descriptors at ROC analysis were compared: DeepRet800 and SP_ResNet152IN, using the experimental setup detailed in Section 4.3.

545 We performed threshold-limited queries at thresholds T corresponding to a FP rate in the $[0.01 - 0.1]$ range, and a maximum number of results/query K between 2 and 10.

The results are plotted in Fig. 6. Since we are using the same dataset for both hard negative mining and estimating query performance, it can be easily shown from Eq. 4 that a FP rate of 0.1 should correspond to an average number of FPs/query of 0.1 as well. Estimates based on *hn1* have larger deviations from expected values, especially for the DeepRet descriptor on the CLAIMS dataset, which is a 20x larger than expected. For *hn2*, actual measured FPs/query are usually slightly lower than predicted. Since we limit the maximum number of images retrieved by each query, this factor may explain the discrepancy, which is higher for the CLAIMS dataset where images are more tightly clustered in



Figure 5: Examples of challenging image pairs for unsupervised near-duplicate detection. Examples (a-e) are challenging negative examples which, despite high semantic and visual similarity, are not near duplicates. Examples derived from CLAIMS (a-c) are related to image types that are particularly common this collection, whereas examples from MFND (d-e) are mostly of subjects which are particularly popular on Internet, such as sunsets and cats. Examples (g-h) are challenging near-duplicates from the CLAIMS dataset which were given low similarity scores by all descriptors; common patterns that are difficult to detect include drastic changes in viewpoint, or one of the two images in the pair represents a detail of a larger scene.

feature space. It should also be noticed that in Eq. 4 the specificity depends only on the threshold, and not on the query image; our experiments, however, suggest that this does not hold true in practice, and that certain types of images are more prone to false positives.

7. Discussion

7.1. Dataset and methodology

Our contributions are a crucial step towards a principled evaluation methodology through which estimating the specificity of unsupervised detection in arbitrarily sized datasets is reduced to the simpler problem of binary classification of ND vs. NND pairs; a tractable number of NND pairs can be extracted through hard negative mining strategies. In the simplest implementation, hard negatives can be mined by finding the nearest neighbors in the dataset, using exact or approximate search depending on the size.

We established the first benchmark for unsupervised NIND detection, an extension of the MFND benchmark comprising more than 20,000 pairs of INDs or NINDs. We followed a semi-automatic procedure that potentially could locate almost all pairs of NDs in the dataset (Connor et al., 2015). In our experiments,

occasionally hard negatives mined may still contain a small percentage (1-2%) of NDs: hence, annotation of the MFND benchmark should be regarded as an ongoing process, that will grow as new descriptors will be tested. For comparison, in an initial experiment performed before extending the dataset (Connor et al., 2015), roughly 8% of the hard negative mined were either NIND or IND pairs. Our experimental comparison on state-of-the-art descriptors suggests that, when compared with a real-life dataset representative of a fraud detection application, MFND is a surprisingly realistic benchmark for estimating the specificity. On the contrary, NIND samples in MFND are on average slightly easier to detect than other datasets, albeit the difference is much reduced compared to IND samples. The presented methodology builds upon previous results from Connor & Cardillo (2016) on IND detection; we proved that the accuracy of the estimated specificity crucially depends on choosing a proper hard negative mining strategy. We provide two additional contributions that strengthen the adoption of this methodology: first, we show analytically that the AUC of the ROC obtained on the hard negative subset is an upper bound of the true AUC. Secondly, we show experimentally that, starting from the experimental ROC, we are able to predict quite accurately the false positive rate per query, which is an indirect proof that the ROC is indeed a good approximation of the true curve. For this experimental comparison, we used the same dataset for hard negative mining and performance evaluation, but in principle, it would be more convenient to perform the hard negative mining on a smaller dataset. Future work is needed to determine whether the false positive rate can be extrapolated to a larger dataset.

An alternative, more intuitive, figure of merit would be the average recall and FPs/query as a function of the distance threshold t . This curve is less practical to use as it depends on the size of the dataset and, being unbounded, defining summary performance measures such as the Area under the ROC curve is not straightforward. It closely resembles the Free-Response Receiver Operating Characteristics (FROC), an extension of ROC analysis used for many diagnostic tasks where the observer (human or machine) can identify the location of an arbitrary number of potential abnormalities, as opposed to the binary prediction task of determining whether an abnormality is present or not (Petrick et al., 2013). In that context, alternatives to the AUC have been proposed and could be extended to our use case.

7.2. Performance comparison

To the best of our knowledge, this is the first attempt to evaluate deep learning descriptors on unsupervised discovery of non-identical near duplicates.

Connor & Cardillo (2016) argued that global descriptors are sufficient for IND detection. Our experience on the GIST descriptor, which obtained the highest performance in the previous comparison, suggests that CNN-based descriptors offer significant advantages also in this case, and compare favorably in terms of execution time.

We have included in our comparison three widely used architectures: SPoC, R-MAC and DeepRetrieval. Note that the DeepRetrieval architecture includes

region pooling (like R-MAC), but unlike other descriptors the features are fine-tuned on the Landmarks dataset for the retrieval task using a Siamese network. Confirming previous results on instance-level image retrieval benchmarks, nicely summarized by Zheng et al. (2017), our experimental results overall favor the choice of fine-tuning the representation for retrieval, as opposed to using off-the-shelf features trained using classification loss (Gordo et al., 2016). The actual performance gap, however, strongly depends on factors related to both the network architecture, the chosen trade-off between specificity and sensitivity, and the underlying dataset structure.

The Holidays dataset has been extensively used to benchmark instance-level retrieval tasks, and all descriptors analyzed in this paper were also previously tested on this dataset, albeit using a different approach for performance assessment. The performance (mean Average Precision) is reported in previous literature as follows: 75.9 (SPoC), 85.2 (R-MAC) and 86.7 (DeepRetrieval) (Gordo et al., 2016). For the Holidays near-duplicate detection task, the best results for the three descriptors are 0.641 (R-MAC), 0.694 (SPoC) and 0.676 (DeepRetrieval), suggesting that SPoC may outperform architectures that are significantly more complex to train and deploy. We should note that none of the descriptors were trained on the Holidays dataset, but the PCA for SPoC and R-MAC was trained on the MFND dataset, which is used as distractors for the near-duplicate detection task.

First, the task is different, not only because the performance measure is different, but also because in our experimental setting, images from the MFND collection are used as negative samples; this is needed to evaluate specificity, which is difficult to do directly on Holidays due to the small size of the dataset and the absence of distractor images. We found experimentally that in many cases the increase in sensitivity is counterbalanced by a corresponding increase in the false positive rate. This is especially evident for the R-MAC descriptor, for which the overall performance decreases in all datasets except MFND. Secondly, each descriptor has many parameters, and the best combination is dataset dependent. While exploring all possible combinations is a daunting task, our experiments provide some useful insights. We found that the backbone depth and architecture were the single most important factor affecting performance. The original SPoC paper, and many subsequent comparisons (Babenko & Lempitsky, 2015; Gordo et al., 2016; Zheng et al., 2017), employed the VGG architecture as backbone, but we found a major boost in performance by using Residual Networks; the DeepRetrieval architecture, on the contrary, uses ResNet101 as backbone (Gordo et al., 2017). In our experiments, the depth of the architecture appears a more relevant factor than the specific feature training, and this an important consideration that should be kept in mind by practitioners.

When compared on the same backbone architecture (Resnet101), the DeepRetrieval outperformed SPoC on CLAIMS and MFND, but not on Holidays. The Holidays dataset contains a lot of outdoors and natural scenes imagery, which may not sufficiently covered by the Landmarks dataset. We expected that SPoC features extracted from networks trained on a scene recognition task, for instance on the Places dataset, or a mixture of Places and ImageNet,

could perform better for near-duplicate detection, since many near-duplicates include complex scenes. However, we did not find consistent advantages, especially when using Residual Networks as the backbone architecture. In a high specificity setting, the difference between pre-trained and fine-tuned networks is further reduced, as visually similar images tend to generate many false positives. Future work will be dedicated to training a specific descriptor for unsupervised near-duplicate detection, incorporating specificity requirements at training time as well as test time. In literature, feature weighting schemes have also been proposed (Mohedano et al., 2018; Kalantidis et al., 2016); such descriptors could be trained in an unsupervised fashion, or do not require any training at all. The performance of such schemes from the point of view of specificity is another direction worth exploring.

In this work, we have used the same descriptor and distance function for all images, regarding of their content. Notably, images are not uniformly distributed in the embedded feature space, and the specificity is largely affected by the presence of clusters of images that are very similar from a semantic and visual point of view. This behaviour is observed in both CLAIMS and MFND datasets, despite their different origin. Exploiting this underlying structure to improve the performance of ND discovery is an important avenue for future research.

8. Conclusions

Unsupervised discovery of near-duplicate detection is an important problem in digital forensics and fraud detection. As the number of false alarms grows quadratically with the size of the input dataset, practical applications require a very high specificity, or conversely low false positive rate, often in the range of 10^{-7} – 10^{-10} . Hard negative mining can be used to select a subset of the dataset, on which ROC analysis can be used to evaluate the performance.

We have evaluated a selection of descriptors based on Convolutional Neural Networks following the proposed methodology. While the task of NIND detection is conceptually similar to instance-level image retrieval, we experimentally found that the same descriptors may be ranked differently, as the Area under the ROC curve depends more strongly on specificity than the mean Average Precision. This strengthens the need for a dedicated benchmark, targeting applications where unsupervised search is required. Our findings in general favor the choice of fine-tuning deep convolutional networks, as opposed to using off-the-shelf features, but differences at high specificity settings strongly depend on the specific dataset and are often small. On the MFND collection, promising performance is obtained by the DeepRet descriptor, retrieving 96% of the true positives at a FP rate of 1.43×10^{-6} . However, further improvement in specificity would benefit many applications, especially in the forensics domain.

Appendix A. Hard negative mining provides an upper bound for the AUC

705 In this section, proof that the AUC calculated using either hard negative mining strategies is an upper bound for the true AUC is provided.

Proposition 1. *When using hard negative mining strategy hn2, the resulting AUC_{hn2} is an upper bound for the true AUC.*

Proof. Hard negative mining strategy hn2 ensures that the selected n_l pairs are the most difficult pairs within the set N^- ; it follows that:

$$f(n_j) \leq f(n_l) \quad \forall n_j \in N^- - H^- \quad \forall n_l \in H^- \quad (A.1)$$

and consequently:

$$\mathbb{1}_{p_i > n_j} \leq \mathbb{1}_{p_i > n_l} \quad \forall n_j \in N^- - H^- \quad \forall n_l \in H^- \quad (A.2)$$

The AUC can be decomposed in two terms

$$AUC = \frac{1}{N^+ N^-} \left[\sum_{i=1}^{N^+} \sum_{l=1}^{H^-} \mathbb{1}_{p_i > n_l} + \sum_{i=1}^{N^+} \sum_{j=1}^{N^- - H^-} \mathbb{1}_{p_i > n_j} \right] \quad (A.3)$$

where the first term is known and is proportional to AUC_{hn2} from Eq. 8, and the second term is the contribution of the negative samples that are not observed. However, we can substitute the second term by replicating the hard negative samples $\lceil \frac{N^- - H^-}{H^-} \rceil$ times, and combining with Eq. A.2 we conclude that:

$$\begin{aligned} AUC &\leq \frac{1}{N^+ N^-} \left[\sum_{i=1}^{N^+} \sum_{l=1}^{H^-} \mathbb{1}_{p_i > n_l} + \sum_{k=1}^{\frac{N^- - H^-}{H^-}} \sum_{i=1}^{N^+} \sum_{l=1}^{H^-} \mathbb{1}_{p_i > n_l} \right] = \\ &= \frac{1}{N^+ N^-} \left[N^+ H^- AUC_{hn2} + \left(\frac{N^- - H^-}{H^-} N^+ H^- \right) AUC_{hn2} \right] = AUC_{hn2} \end{aligned}$$

□

710 **Proposition 2.** *When using hard negative mining strategy hn1, the resulting AUC_{hn1} is an upper bound for the true AUC.*

Proof. Again, let us decompose the AUC as the sum of two terms, where the first term is known and is proportional to AUC_{hn1} , and the second term is the contribution of the negative samples that are not observed, as detailed in Eq.

715 A.3.

Each sample n_j consists of a pair of images (x_k, y_m) , where $x_k \in X$ and $y_m \in Y$; in other terms, $N^- = \{(x_k, y_m), \quad k = 1, \dots, K, \quad m = 1, \dots, M\}$. Then according to the definition of hn1,

$$H^- = \{(x_k, y_{m^*}) \mid m^* = \arg \max_m f(x_k, y_m)\} \quad (A.4)$$

The sum over $l = 1, \dots, H^-$ and $j = 1, \dots, N^-$ in Eq. A.3 can be decomposed in terms of $k = 1, \dots, K$ and $m = 1, \dots, M$ as follows:

$$AUC = \frac{1}{N^+N^-} \left[\sum_{i=1}^{N^+} \sum_{k=1}^K \mathbb{1}_{p_i > (x_k, y_{m^*})} + \sum_{i=1}^{N^+} \sum_{k=1}^K \sum_{m=1, m \neq m^*}^M \mathbb{1}_{p_i > (x_k, y_m)} \right] \quad (\text{A.5})$$

where $N^- = KM$, and according to the definition of *hn1* there are exactly K hard negative pairs.

By definition, $f(x_k, y_m) \leq f(x_k, y_{m^*})$ and thus

$$\mathbb{1}_{p_i > (x_k, y_m)} \leq \mathbb{1}_{p_i > (x_k, y_{m^*})} \quad \forall k = 1, \dots, K \quad \forall m \neq m^* \quad (\text{A.6})$$

Combining Eqs. A.5 and A.6, we conclude that:

$$\begin{aligned} AUC &\leq \frac{1}{N^+N^-} \left[N^+H^- AUC_{hn1} + \sum_{i=1}^{N^+} \sum_{k=1}^K \sum_{m=1, m \neq m^*}^M \mathbb{1}_{p_i > (x_k, y_{m^*})} \right] = \\ &\frac{1}{N^+N^-} \left[N^+K AUC_{hn1} + (M-1)N^+K AUC_{hn1} \right] = AUC_{hn1} \end{aligned}$$

□

Acknowledgment

The authors gratefully acknowledge the financial support of Reale Mutua Assicurazioni. We are grateful to Lucia Sabatino, Lucia Romano and Giulia Gemesio for help in data cleaning and annotation.

Conflict of Interest The research was funded by a private grant from Reale Mutua Assicurazioni (Torino, Italy). The authors report no conflict of interest, and have no direct or indirect commercial financial incentive associated with publishing the manuscript.

References

References

- Amerini, I., Uricchio, T., & Caldelli, R. (2017). Tracing images back to their social network of origin: A cnn-based approach. In *Information Forensics and Security (WIFS), 2017 IEEE Workshop on* (pp. 1–6). IEEE.
- Babenko, A., & Lempitsky, V. (2015). Aggregating local deep features for image retrieval. In *Proceedings of the IEEE international conference on computer vision* (pp. 1269–1277).
- Babenko, A., Slesarev, A., Chigorin, A., & Lempitsky, V. (2014). Neural codes for image retrieval. In *European conference on computer vision* (pp. 584–599). Springer.

- 740 Balntas, V., Riba, E., Ponsa, D., & Mikolajczyk, K. (2016). Learning local feature descriptors with triplets and shallow convolutional neural networks. In *BMVC* (p. 3). volume 1.
- Battiato, S., Farinella, G. M., Puglisi, G., & Ravì, D. (2014). Aligning codebooks for near duplicate image detection. *Multimedia Tools and Applications*, 72, 1483–1506.
- 745 Bay, H., Ess, A., Tuytelaars, T., & Van Gool, L. (2008). Speeded-up robust features (surf). *Computer vision and image understanding*, 110, 346–359.
- Bradley, A. P. (1997). The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30, 1145–1159.
- 750 Chen, M., Li, Y., Zhang, Z., Hsu, C.-H., & Wang, S. (2017). Real-time, large-scale duplicate image detection method based on multi-feature fusion. *Journal of Real-Time Image Processing*, 13, 557–570.
- Chennamma, H., Rangarajan, L., & Rao, M. (2009). Robust near duplicate image matching for digital image forensics. *International Journal of Digital Crime and Forensics (IJDcf)*, 1, 62–79.
- 755 Chiu, C.-Y., Li, S.-Y., & Hsieh, C.-Y. (2012). Video query reformulation for near-duplicate detection. *IEEE Transactions on Information Forensics and Security*, 7, 1594–1603.
- Chum, O., Philbin, J., Zisserman, A. et al. (2008). Near duplicate image detection: min-hash and tf-idf weighting. In *BMVC* (pp. 812–815). volume 810.
- 760 Cicconet, M., Elliott, H., Richmond, D. L., Wainstock, D., & Walsh, M. (2018). Image forensics: Detecting duplication of scientific images with manipulation-invariant image similarity. *arXiv preprint arXiv:1802.06515*, .
- Connor, R., Cardillo, F., MacKenzie-Leigh, S., & Moss, R. (2015). Identification of mir-flickr near-duplicate images. In *10th International Conference on Computer Vision Theory and Applications*.
- 765 Connor, R., & Cardillo, F. A. (2016). Quantifying the specificity of near-duplicate image classification functions. In *11th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*.
- 770 Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on* (pp. 248–255). Ieee.
- 775 Foo, J. J., & Sinha, R. (2007). Pruning sift for scalable near-duplicate image matching. In *Proceedings of the eighteenth conference on Australasian database-Volume 63* (pp. 63–71). Australian Computer Society, Inc.

- Gonçalves, F. M. F., Guilherme, I. R., & Pedronette, D. C. G. (2018). Semantic guided interactive image retrieval for plant identification. *Expert Systems with Applications*, 91, 12–26.
- 780 Gordo, A., Almazán, J., Revaud, J., & Larlus, D. (2016). Deep image retrieval: Learning global representations for image search. In *European Conference on Computer Vision* (pp. 241–257). Springer.
- Gordo, A., Almazan, J., Revaud, J., & Larlus, D. (2017). End-to-end learning of deep visual representations for image retrieval. *International Journal of Computer Vision*, 124, 237–254.
- 785 Hanley, J. A., & McNeil, B. J. (1982). The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143, 29–36.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
- 790 Hirano, Y., Garcia, C., Sukthankar, R., & Hoogs, A. (2006). Industry and object recognition: Applications, applied research and challenges. In *Toward Category-Level Object Recognition* (pp. 49–64). Springer.
- Hu, Y., Cheng, X., Chia, L.-T., Xie, X., Rajan, D., & Tan, A.-H. (2009). Coherent phrase model for efficient image near-duplicate retrieval. *IEEE Trans. Multimedia*, 11, 1434–1445.
- 795 Huiskes, M. J., & Lew, M. S. (2008). The mir flickr retrieval evaluation. In *Proceedings of the 1st ACM international conference on Multimedia information retrieval* (pp. 39–43). ACM.
- Jegou, H., Douze, M., & Schmid, C. (2008). Hamming embedding and weak geometric consistency for large scale image search. In *European conference on computer vision* (pp. 304–317). Springer.
- 800 Jinda-Apiraksa, A., Vonikakis, V., & Winkler, S. (2013). California-nd: An annotated dataset for near-duplicate detection in personal photo collections. In *Quality of Multimedia Experience (QoMEX), 2013 Fifth International Workshop on* (pp. 142–147). IEEE.
- 805 Johnson, J., Douze, M., & Jégou, H. (2017). Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, .
- Kalantidis, Y., Mellina, C., & Osindero, S. (2016). Cross-dimensional weighting for aggregated deep convolutional features. In *European conference on computer vision* (pp. 685–701). Springer.
- 810 Ke, Y., Sukthankar, R., & Huston, L. (2004). An efficient parts-based near-duplicate and sub-image retrieval system. In *Proceedings of the 12th annual ACM international conference on Multimedia* (pp. 869–876). ACM.

- Kim, S., Wang, X.-J., Zhang, L., & Choi, S. (2015). Near duplicate image
815 discovery on one billion images. In *Applications of Computer Vision (WACV), 2015 IEEE Winter Conference on* (pp. 943–950). IEEE.
- Li, J., Qian, X., Li, Q., Zhao, Y., Wang, L., & Tang, Y. Y. (2015). Mining near
duplicate image groups. *Multimedia Tools and Applications*, 74, 655–669.
- Li, P., Shen, B., & Dong, W. (2018). An anti-fraud system for car insurance
820 claim based on visual evidence. *arXiv preprint arXiv:1804.11207*, .
- Liu, L., Lu, Y., & Suen, C. Y. (2015). Variable-length signature for near-
duplicate image matching. *IEEE Transactions on Image Processing*, 24,
1282–1296.
- Mohedano, E., McGuinness, K., Giró-i Nieto, X., & O'Connor, N. E. (2018).
825 Saliency weighted convolutional features for instance search. In *2018 International Conference on Content-Based Multimedia Indexing (CBMI)* (pp. 1–6).
IEEE.
- Oliva, A., & Torralba, A. (2001). Modeling the shape of the scene: A holistic
830 representation of the spatial envelope. *International journal of computer vision*, 42, 145–175.
- de Oliveira, A. A., Ferrara, P., De Rosa, A., Piva, A., Barni, M., Goldenstein,
S., Dias, Z., & Rocha, A. (2016). Multiple parenting phylogeny relationships
in digital images. *IEEE Transactions on Information Forensics and Security*,
11, 328–343.
- 835 Petrick, N., Sahiner, B., Armato, S. G., Bert, A., Correale, L., Delsanto, S.,
Freedman, M. T., Fryd, D., Gur, D., Hadjiiski, L. et al. (2013). Evaluation
of computer-aided detection and diagnosis systems. *Medical physics*, 40.
- Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for
large-scale image recognition. *arXiv preprint arXiv:1409.1556*, .
- 840 Tolias, G., Sicre, R., & Jégou, H. (2015). Particular object retrieval with integral
max-pooling of cnn activations. *arXiv preprint arXiv:1511.05879*, .
- Wan, J., Wang, D., Hoi, S. C. H., Wu, P., Zhu, J., Zhang, Y., & Li, J. (2014).
Deep learning for content-based image retrieval: A comprehensive study. In
845 *Proceedings of the 22nd ACM international conference on Multimedia* (pp.
157–166). ACM.
- Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen,
B., & Wu, Y. (2014). Learning fine-grained image similarity with deep rank-
ing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern
Recognition* (pp. 1386–1393).
- 850 Xie, L., Tian, Q., Zhou, W., & Zhang, B. (2014). Fast and accurate near-
duplicate image search with affinity propagation on the imageweb. *Computer
Vision and Image Understanding*, 124, 31–41.

- Xu, D., Cham, T. J., Yan, S., Duan, L., & Chang, S.-F. (2010). Near duplicate identification with spatially aligned pyramid matching. *IEEE Transactions on Circuits and Systems for Video Technology*, 20, 1068–1079.
- Zagoruyko, S., & Komodakis, N. (2015). Learning to compare image patches via convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4353–4361).
- Zheng, L., Yang, Y., & Tian, Q. (2017). Sift meets cnn: A decade survey of instance retrieval. *IEEE transactions on pattern analysis and machine intelligence*, .
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., & Torralba, A. (2017a). Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, .
- Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., & Oliva, A. (2014). Learning deep features for scene recognition using places database. In *Advances in neural information processing systems* (pp. 487–495).
- Zhou, W., Li, H., & Tian, Q. (2017b). Recent advance in content-based image retrieval: A literature survey. *arXiv preprint arXiv:1706.06064*, .
- Zhou, Z., Wang, Y., Wu, Q. J., Yang, C.-N., & Sun, X. (2017c). Effective and efficient global context verification for image copy detection. *IEEE Transactions on Information Forensics and Security*, 12, 48–63.

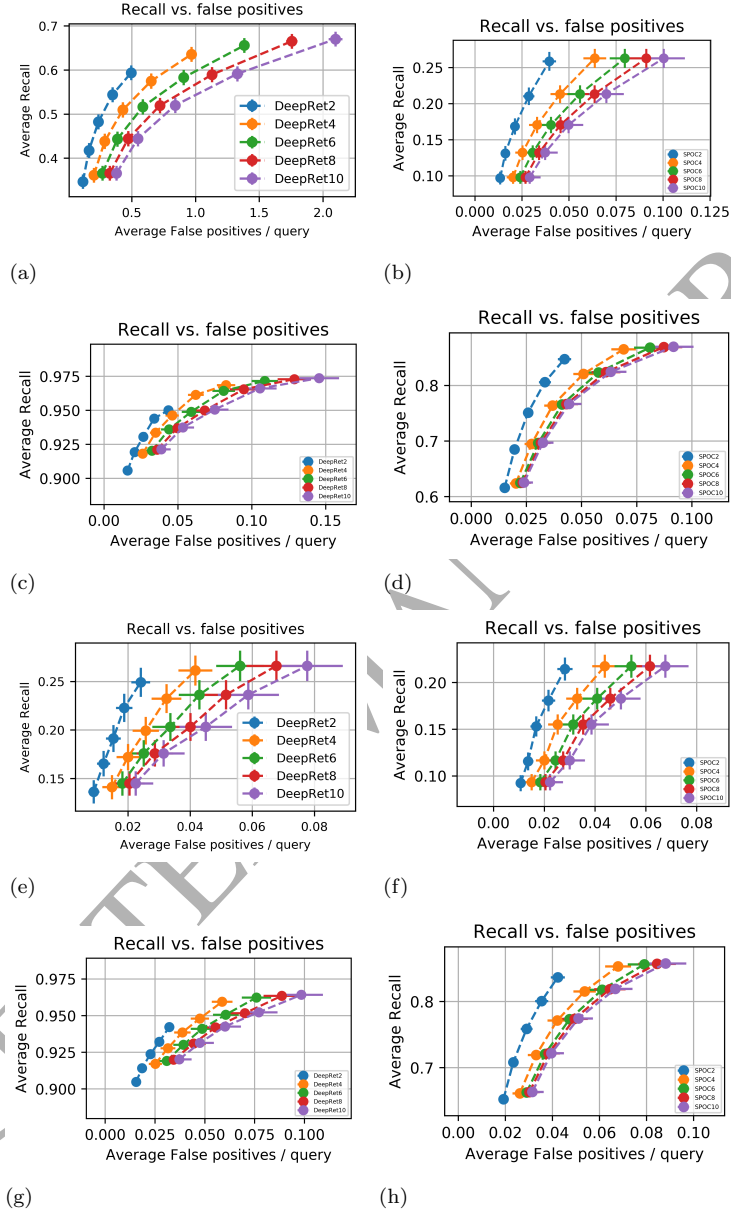


Figure 6: Average recall vs. FPs/query for the CLAIMS (a-b, e-f) and MFND dataset (c-d, g-h), with thresholds calculated using hard negative mining *hn1* (top row, a-d) and *hn2* (bottom row, e-h). Performance is measured at fixed thresholds (dots in the above curves), bars indicate the standard error. The maximum number of images retrieved by each query is limited to $K = 2, 4, 6, 8, 10$, results are plotted as separate curves.

Lia Morra: Conceptualization, Methodology, Software, Investigation, Writing-
Original Draft **Fabrizio Lamberti:** Conceptualization, Writing-Review & Edit-
875 ing, Funding Acquisition

ACCEPTED MANUSCRIPT