

A Self-Supervised Descriptor for Image Copy Detection

Ed Pizzi Sreya Dutta Roy Sugosh Nagavara Ravindra Priya Goyal Matthijs Douze
Meta AI

Abstract

Image copy detection is an important task for content moderation. We introduce *SSCD*, a model that builds on a recent self-supervised contrastive training objective. We adapt this method to the copy detection task by changing the architecture and training objective, including a pooling operator from the instance matching literature, and adapting contrastive learning to augmentations that combine images.

Our approach relies on an entropy regularization term, promoting consistent separation between descriptor vectors, and we demonstrate that this significantly improves copy detection accuracy. Our method produces a compact descriptor vector, suitable for real-world web scale applications. Statistical information from a background image distribution can be incorporated into the descriptor.

On the recent DISC2021 benchmark, *SSCD* is shown to outperform both baseline copy detection models and self-supervised architectures designed for image classification by huge margins, in all settings. For example, *SSCD* outperforms SimCLR descriptors by 48% absolute.

Code is available at <https://github.com/facebookresearch/sscd-copy-detection>.

1. Introduction

All online photo sharing platforms use content moderation to block or limit the propagation of images that are considered harmful: terrorist propaganda, misinformation, harassment, pornography, *etc.* Some content moderation can be performed automatically, for unambiguous data like pornographic pictures, but this is much harder for complex data like memes [31] or misinformation [2]. In these cases, content is moderated manually. For viral images, where copies of same image may be uploaded thousands of times, manual moderation of each copy is tedious and unnecessary. Instead, each image for which a manual moderation decision is taken can be recorded in a database, so that it can be re-identified later and handled automatically.

This paper is concerned with this basic task of re-identification. This is non trivial because copied images are

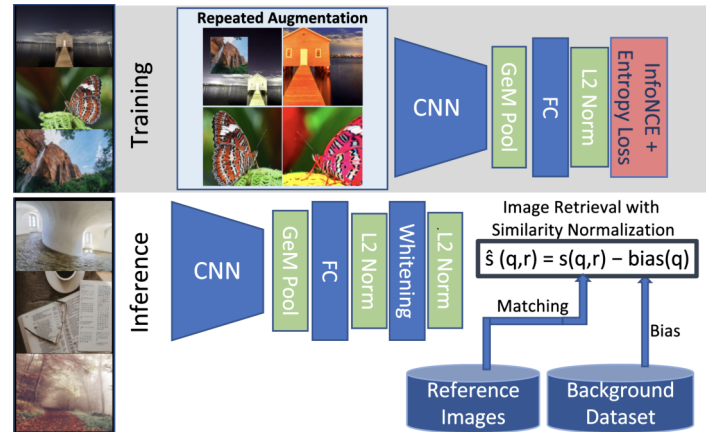


Figure 1. The SSCD architecture for image copy detection. It is based on SimCLR, with the following additions: the entropy regularization, cutmix/mixup-aware InfoNCE, and inference-time score normalization.

often altered, for technical reasons (*e.g.* a user shares a mobile phone screenshot that captures additional content), or users may make adversarial edits to evade moderation.

Image re-identification is an image matching problem, with two additional challenges. The first is the enormous scale at which copy detection systems are deployed. At this scale, the only feasible approach is to represent images as short descriptor vectors, that can be searched efficiently with approximate nearest neighbor search methods [23, 30]. Copy detection systems typically proceed in 2 stages: a *retrieval* stage that produces a shortlist of candidate matches and a *verification* stage, often based on local descriptor matching that operates on the candidates. In this work, we are concerned with the first stage. Figure 1 shows the overall architecture of our Self Supervised Copy Detection (SSCD) approach.

The second challenge is that there is a hard match/non-match decision to take, and positive image pairs are rare. We wish to limit verification candidates using a threshold, which is a harder constraint than the typical image retrieval setting, where only the order of results matter.

SSCD uses differential entropy regularization [42] to

promote a uniform embedding distribution, which has three effects: (1) it makes distances from different embedding regions more comparable; (2) it avoids the embedding collapse described in [29], making full use of the embedding space; (3) it also improves ranking metrics that do not require consistent thresholds across queries.

Score normalization is important for ranking systems. An advanced score normalization relies on matching the query images with a set of background images. In this work, we show how this normalization can be incorporated in the image descriptor itself. We anticipate that this work will set a strong single-model baseline for image copy detection. We plan to release code and models for our method.

Section 2 discusses works related to this paper. Section 3 motivates the use of an entropy loss term in a simplified setting. Section 4 carefully describes SSCD. Section 5 presents results and ablations of our method. Section 6 points out a few observations about the copy detection task.

2. Related work

Content tracing approaches. Content tracing on a user-generated photo sharing platform aims at re-identifying images when they circulate out and back into the platform. There are three broad families of tracing methods: metadata-based [1, 3], watermarking [13, 32, 51, 63] and content-based. This work belongs to this last class.

Classical image datasets for content tracing, like Casia [16, 36] focus on image alterations like splicing, removal and copy-move transformations [16, 44, 54] that alter only a small fraction of the image surface, so the re-identification is done reliably with simple interest-point based techniques. The challenge is to detect the tampered surface, which is typically approached with deep models inspired by image segmentation [34, 62]. A related line of research is image phylogeny: the objective is to identify the series of edits that were applied to an image between an initial and a final state [14, 15, 33]. The Nimble/Media forensics series of competitions organized by NIST aim at benchmarking these tasks [40, 57]. In this work we focus on the identification itself, with strong transformations and near duplicates that need to be distinguished (see Figure 2).

Semantic and perceptual image comparison Several definitions of near-duplicate image matching, form a continuum between pixel-wise copy and instance matching [18, 28]. The definition we use in this work is: images are considered copies iff they come from the same 2D image source. More relaxed definitions allow, for example, to match nearby frames in a video.

There is a large body of literature about solving instance matching [7, 11, 26, 35, 37, 46–48] *i.e.*, recognizing images of the same 3D object with viewpoint/camera changes. In this



Figure 2. Example retrieval results from the DISC2021 dataset. Each row is an example. From left to right: query image, first result returned by SSCD, first result returned by the SimCLR baseline.

work, we build on this literature because it addresses complex image matching, and to our knowledge, recent works and benchmarks for strict copy detection are rare [17, 53].

Instance matching. Classical instance matching relies on 3D matching tools, like interest points [26, 35, 43]. CNN-based approaches use backbones from image classification, either pre-trained [4, 20, 49] or trained end-to-end [21, 38], with two adaptations: (1) the pooling layer that converts the last CNN activation map to a vector is a max-pooling [49], or more generally GeM pooling [39], a form of L_p normalization where p is adapted to the image resolution [7]; (2) careful normalization of the vectors. In addition to simple L2-normalization [4], “whitening” is often used to compare descriptors [25, 49]. An additional normalization technique contrasts the distances w.r.t. a background distribution of images [18, 27]. In this work, we apply these pooling and normalization techniques to copy detection.

Contrastive self-supervised learning. A recent line of self-supervised learning research uses contrastive objectives that learn image representations that bring transformed im-

ages together. These methods either discriminate image features [10, 22, 24] or the cluster assignments of these image features [8]. These methods either rely on memory banks [24, 56] or large batch sizes [10]. In particular, SimCLR [10] uses matching transformed image copies as a surrogate task to learn a general image representation that transfer well to other tasks, such as image classification. A contrastive InfoNCE loss [52] is used to map copies of the same source image nearby in the embedding space.

Differential entropy regularization. Increasing the entropy of media descriptors forces them to spread over the representation space. Sablayrolles *et al.* [42] observed that the entropy can be estimated locally with the Kozachenko-Leononenko differential entropy estimator [6], that can be incorporated directly into the loss to maximize descriptor entropy. The work of El-Nouby *et al.* [19] is closest to our approach. It adds the entropy term to a contrastive loss at fine-tuning time to improve the accuracy for category and instance retrieval. Our approach is similar, applied to a self-supervised objective and image copy detection.

3. Motivation

In this section, we start from the SimCLR [10] method, then perform a simple experiment where we combine it with the entropy loss from [42] and witness how it impacts classification and copy detection tasks.

3.1. Preliminaries: SimCLR

SimCLR training is best described at the mini-batch level. For batches of N images, it creates two augmented copies of each image (repeated augmentations), yielding $2N$ transformed images. The positive pairs of matching images are $P = \{(i, i + N), (i + N, i)\}_{i=1..N}$. We denote positive matches for image i as $P_i = \{j \mid (i, j) \in P\}$. Each image is transformed by a CNN backbone network. The final activation map of the CNN is average pooled, then projected using a two-layer MLP into a L2-normalized descriptor $z_i \in \mathbb{R}^d$. Descriptors are compared with a cosine similarity: $\text{sim}(z_i, z_j) = z_i^T z_j$. A contrastive InfoNCE loss maximizes the similarity between copies relative to the similarity of non-copies. For inference (*e.g.* to transfer to image classification), SimCLR discards the training-time MLP, using globally pooled features from the CNN trunk directly.

The InfoNCE loss. SimCLR’s InfoNCE loss is a softmax cross-entropy with temperature, that matches descriptors to other descriptors. Let $s_{i,j}$ be the temperature-adjusted cosine similarity $s_{i,j} = \text{sim}(z_i, z_j)/\tau$. The InfoNCE loss is defined as a mean of $\ell_{i,j}$ terms for positive pairs $(i, j) \in P$:

$$\ell_{i,j} = -\log \frac{\exp(s_{i,j})}{\sum_{k \neq i} \exp(s_{i,k})} \quad (1)$$

$$\mathcal{L}_{\text{InfoNCE}} = \frac{1}{|P|} \sum_{i,j \in P} \ell_{i,j}. \quad (2)$$

3.2. Entropy regularization

We use the differential entropy loss proposed in [42], based on the Kozachenko-Leononenko estimator. We adapt it to the repeated augmentation setting by only regularizing neighbors from different source images:

$$\mathcal{L}_{\text{KoLeo}} = -\frac{1}{N} \sum_{i=1}^N \log \left(\min_{j \notin \hat{P}_i} \|z_i - z_j\| \right), \quad (3)$$

where $\hat{P}_i = P_i \cup \{i\}$. Since this entropy loss is a log of the distance to the nearest neighbor, its impact is very high for nearby vectors but dampens quickly when the descriptors are far apart. The effect is to “push” apart nearby vectors.

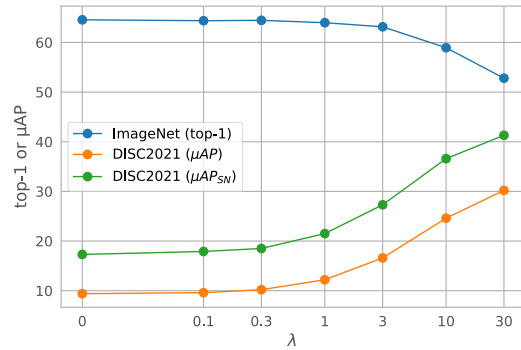


Figure 3. Preliminary experiment: we train SimCLR models on ImageNet with varying differential entropy regularization strength, (regular SimCLR: $\lambda = 0$). We measure: ImageNet linear classification accuracy and DISC2021 micro average precision (μAP), with optional score normalization (μAP_{SN}). The ImageNet and DISC2021 measures are not comparable, but trends within each curve are significant.

3.3. Experiment: SimCLR and entropy

For this experiment, we combine our contrastive loss with the entropy loss, using a weighting factor λ , similar to [19, 42]:

$$\mathcal{L}_{\text{basic}} = \mathcal{L}_{\text{InfoNCE}} + \lambda \mathcal{L}_{\text{KoLeo}}. \quad (4)$$

We then evaluate the impact of the combined loss on an image classification setting and a copy detection setting, see Section 5.1 for more details about the setup.

Figure 3 shows how varying entropy loss weight λ impacts both tasks. As the entropy loss weight increases, ImageNet linear classification accuracy decreases: this loss

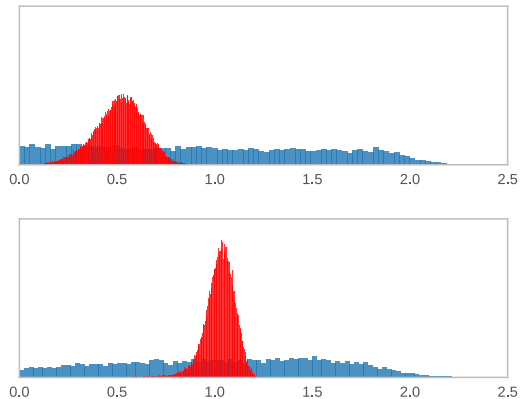


Figure 4. Preliminary experiment: histogram of squared distances for DISC2021 **matching images** and **non-matching nearest neighbors**. Above: baseline SimCLR. Below: SimCLR combined with entropy regularization (weight $\lambda = 30$), without whitening or similarity normalization.

term is not helpful for classification. Conversely, for copy detection the accuracy increases significantly.

Figure 4 shows the distribution of distances between matching images (positive pairs) and the nearest non-matching neighbors (negative pairs). Applying the entropy loss increases all distances and makes the negative distance distribution more narrow. The result is that there is a larger contrast between positive pairs and the mode of the negative distribution, *i.e.* they are more clearly separated.

4. Method

Having seen how the entropy loss improves copy detection accuracy, in this section we expand it into a robust image copy detection approach: SSCD. This entails adapting the architecture, the data augmentation, the pooling and adding a normalization stage, as shown in Figure 1.

4.1. Architecture

SSCD uses a ResNet-50 convolutional trunk to extract image features. We standardize on this architecture because it is widely used, well optimized and still very competitive for image classification [55], but any CNN or transformer backbone could be used (see Section 5).

Pooling. For classification, the last CNN activation map is converted to a vector by mean pooling. We use generalized mean (GeM) pooling instead, which was shown [7, 39] to improve the discriminative ability of descriptors. This is desirable for instance retrieval and our copy detection case alike. GeM introduces a parameter p , equivalent to average pooling when $p = 1$ and max-pooling when $p \rightarrow \infty$. SSCD uses $p = 3$, following common practice for image retrieval models [7, 39, 49].

type	details
SimCLR	horizontal flip, random crop, color jitter, grayscale, Gaussian blur
Strong blur	50% large-radius Gaussian blur ($\sigma \in [1, 5]$)
Advanced	10% rotation, 10% text, 20% emoji, 20% JPEG compression
Adv. + mixup	2.5% mixup, 2.5% cutmix

Table 1. List of data augmentations used for SSCD. The presentations is incremental: each set of augmentations includes the ones from all rows before. The percentages are probabilities to apply each augmentation.

While GeM pooling at inference time systematically improves accuracy, we observe that it is beneficial at training time only in combination with the differential entropy regularization, *i.e.* with a vanilla InfoNCE it is better to train with average pooling. We conjecture that GeM pooling may reduce the difficulty of the training task without the additional objective of maximally separating embedding points. We observe that learning the scalar p , as proposed in [39], fails for contrastive learning: the pooling parameter grows unbounded until training becomes numerically unstable.

Descriptor projection. SimCLR uses a 2-layer MLP projection at training time. For inference, the MLP is discarded and CNN trunk features are used directly. The MLP is partly motivated to retain transformation-covariant features in the base network, which may be useful for downstream tasks, despite a training task that requires a transformation-invariant descriptor. Jing *et al.* [29] also find that the MLP insulates the trunk model from an embedding collapse into a lower-dimensional space caused by the InfoNCE loss.

For SSCD, the training and inference tasks are the same, obviating the need for transformation-covariant features, and differential entropy regularization prevents the dimensional collapse. We replace the MLP with a simple linear projection to the target descriptor size, and retain this projection for inference.

4.2. Data Augmentation

Self-supervised contrastive objectives learn to match images across image transforms. These methods are sensitive to the augmentations seen at training time [10], since invariance to these transforms is the only supervisory signal.

Table 1 lists the SSCD augmentations used in our experiments. Note that since our main evaluation dataset (DISC2021) is built in part with data augmentation, there is a risk of overfitting to the augmentations of that dataset. This is mitigated by (1) DISC2021’s set of augmentations is not known precisely and (2) we present strong results trained using a simple blur augmentation. Our starting baseline is the default set of SimCLR augmentations.

Strong blur. Empirically, copy detection benefits from a stronger blur than is typically used for contrastive learning. We strengthen the blur augmentation compared to SimCLR. We suggest that invariance to blur confers a low-frequency bias, reducing the model’s sensitivity to high-frequency noise common to real world copies. We use this setting for most ablation steps, because it is easy to reproduce, and provides a good baseline setting for comparing methods. This augmentation was initially tuned on a proprietary dataset, and is unlikely to overfit to DISC2021.

Advanced augmentations. We evaluate our method with additional augmentations, to demonstrate how SSCD extends as augmentations are added. Half of rotations rotate by multiples of 90 degrees and half are unconstrained. The text has a random font, text, opacity, font size, and color. We add emoji of random size. We apply JPEG compression with randomly sampled compression quality. These augmentations are somewhat inspired by DISC2021 but are still fairly generic for image copy detection problems.

Mixed images. We use two augmentations that combine content from two images within a training batch. In a copy detection context, these augmentations model partial copies, where part of an image is included in a composite image. Mixup [61] is a pixelwise weighted average of two images (a and b) with parameter $\gamma \in [0, 1]$: $\gamma \cdot a + (1 - \gamma) \cdot b$. CutMix [59] moves rectangular regions from one image into another. See Appendix D for implementation details. Mixed images match multiple images in the batch, requiring changes to our losses, outlined below.

4.3. Loss Functions

SSCD uses a weighted combination of the contrastive InfoNCE and the entropy loss, as in Equation (4). However, we need to adapt both losses for the mixed-image augmentation case, where P_i may contain multiple matching images.

InfoNCE with MixUp/CutMix augmentations. We adapt the InfoNCE loss (see Section 3.1) to accommodate augmentations that mix features from multiple images. Given an image i with full or partial matches $j \in P_i$, we modify the pairwise loss term from Equation (1) as:

$$\hat{\ell}_{i,j} = -\log \frac{\exp(s_{i,j})}{\exp(s_{i,j}) + \sum_{k \notin \hat{P}_i} \exp(s_{i,k})}, \quad (5)$$

where $\hat{P}_i = P_i \cup \{i\}$. We then combine these terms by taking a mean per image, so that each image contributes similarly to the overall loss, and average per-image losses. Note that this is equivalent to InfoNCE for non-mixed images.

$$\mathcal{L}_{\text{InfoNCE-mix}} = \frac{1}{2N} \sum_{i=1}^{2N} \frac{1}{|P_i|} \sum_{j \in P_i} \hat{\ell}_{i,j}. \quad (6)$$

Entropy loss. Our formulation of the entropy loss in Equation (3) remains the same, with \hat{P}_i updated to include multiple matching images.

Combination. The losses are combined with entropy weight parameter λ :

$$\mathcal{L} = \mathcal{L}_{\text{InfoNCE-mix}} + \lambda \mathcal{L}_{\text{KoLeo}} \quad (7)$$

Multi-GPU implementation. The contrastive matching task benefits from a large batch size, since this provides stronger negatives. Losses are evaluated over the global batch, after aggregating image descriptors across GPUs. Descriptors from all GPUs are included in the negatives InfoNCE matches against, and we choose nearest neighbors for entropy regularization from the global batch. Batch normalization statistics are synchronized across GPUs to avoid leaking information within a batch. We use the LARS [58] optimizer for stable training at large batch size.

4.4. Inference and retrieval

For inference, the loss terms are discarded. Features are extracted from the images using the convolutional trunk followed by GeM pooling, the linear projection head, and L2 normalization. Then we apply whitening to the descriptors. The whitening matrix is learned on the DISC2021 training set. The descriptors are compared with cosine similarity or equivalently with simple L2 distance.

4.5. Similarity normalization

We follow [18] using similarity normalization [12,27] as one of our evaluation settings. It uses a background dataset of images as a noise distribution, and produces high similarity scores only for queries whose reference similarity is greater than their similarity to nearest neighbors in the background dataset. Given a query image q and a reference image r with similarity $s(q,r) = \text{sim}(z_q, z_r)$, the adjusted similarity is $\hat{s}_0(q,r) = s(q,r) - \beta s(q, b_n)$ where b_n is the n^{th} nearest neighbor from the background dataset, and $\beta \geq 0$ is a weight.

We generalize this by aggregating an average similarity across multiple neighbors (n to n_{end}) from the background dataset:

$$\hat{s}(q,r) = s(q,r) - \underbrace{\frac{\beta}{1 + n_{\text{end}} - n} \sum_{i=n}^{n_{\text{end}}} s(q, b_i)}_{\text{bias}(q)}. \quad (8)$$

Integrated bias. Carrying around a bias term makes indexing of descriptors more complex. Therefore, we include the bias into the descriptors as an additional dimension:

$$\hat{z}_q = [z_q \quad -\text{bias}(q)] \quad \hat{z}_r = [z_r \quad 1] \quad (9)$$

Then we are back to $\hat{s}(q, r) = \text{sim}(\hat{z}_q, \hat{z}_r)$. The descriptors are not normalized, *i.e.* the dot product similarity is not equivalent to L2 distance. If L2 distance is preferred for indexing, it is possible to convert the max dot product search task into L2 search using the approach from [5].

Similarity normalization consistently improves metrics. However it adds operational complexity, and may make it difficult to detect content similar to the background distribution. Therefore, we report results both with and without this normalization.

5. Experiments

In this section we evaluate SSCD for image copy detection. Despite its relative simplicity, it depends on various settings that we evaluate in an extensive ablation study.

5.1. Datasets

DISC2021. Most evaluations are on the validation dataset of the Image Similarity Challenge, DISC2021 [18]. DISC2021 contains both automated image transforms and manual edits. There are 1 million reference images and 50,000 query images, of which 10,000 are true copies. A disjoint 1 million image training set is used for model training and as background dataset for score normalization. The training set contains no copies or labels, but is representative of the image distribution of the dataset. The performance is evaluated with micro average precision (μAP) that measures the precision-recall tradeoff with a uniform distance threshold.

ImageNet. For some experiments we train models on the ImageNet [41] training set (ignoring the class labels). We use ImageNet linear classification to measure how our copy detection methods affect semantic representation learning.

Copydays [17] is a small copy detection dataset. Following common practice [7, 9], we augment it with 10k distractors from YFCC100M [45], a setting known as CD10K, and evaluate the retrieval performance with mean average precision (mAP) on the “strong” subset of robustly transformed copies. In addition to this standard measure, we evaluate the μAP on the overall dataset.

5.2. Training implementation

We use the training schedule and hyperparameters from SimCLR [10]: batch size $N = 4096$, resolution 224×224 , learning rate of $0.3 \times N/256$, and a weight decay of 10^{-6} . We train models for 100 epochs on either ImageNet or the DISC training set, using a cosine learning rate schedule without restarts and with a linear ramp-up. We use the LARS optimizer for stable training at large batch size. We train at spatial resolution 224×224 .

We use a lower temperature than SimCLR, $\tau = 0.05$ versus 0.1, following an observation in [10] that this setting

yields better accuracy on the training task, while reducing accuracy of downstream classification tasks.

5.3. Evaluation protocol

Inference. We resize the small edge of an image to size 288 preserving aspect ratio for fully convolutional models. We use a larger inference size than seen at training to avoid train-test discrepancy [50]. We use different preprocessing for the DINO [9] ViT baseline, following their copy detection method. See Appendix D for details.

Descriptor postprocessing. Image retrieval benefits from PCA whitening. SSCD descriptors are whitened, then L2 normalized. For baseline methods that use CNN trunk features, we L2 normalize both before and after whitening. SimCLR projection features often occupy a low-dimensional subspace, making whitening at full descriptor size unstable, and many representations perform better when whitened with low-variance dimensions excluded. For baseline methods, we try dimensionalities $\{d, \frac{3}{4}d, \frac{d}{2}, \frac{d}{4}, \dots\}$ and choose the one that maximizes the final accuracy. For SSCD, we whiten at full descriptor size.

We use the FAISS [30] library to apply embedding postprocessing and perform exhaustive k-nearest neighbor search. We train PCA on the DISC2021 training dataset, following standard protocol for this dataset.

5.4. Results

method	trained on	transforms	dims	μAP	μAP_{SN}
Multigrain [7, 18]	ImageNet*		1500	16.5	36.5
HOW [18, 48]	SfM-120k*			17.3	37.2
Multigrain [7]	ImageNet*		2048	20.5	41.7
DINO [9] †	ImageNet		1500	32.2	53.8
SimCLR [10] trunk	ImageNet	SimCLR	2048	13.1	33.9
SimCLR [10] proj	ImageNet	SimCLR	128	9.4	17.3
SimCLR _{CD} trunk	ImageNet	strong blur	2048	39.8	56.8
SSCD	ImageNet	strong blur	512	50.4	64.5
SSCD	ImageNet	advanced	512	55.5	71.0
SSCD	ImageNet	adv.+mixup	512	56.8	72.2
SSCD	DISC	strong blur	512	54.8	63.6
SSCD	DISC	advanced	512	60.4	71.1
SSCD	DISC	adv.+mixup	512	61.5	72.5
SSCD _{large} †	DISC	adv.+mixup	1024	63.7	75.3

Table 2. Copy detection performance in μAP on the DISC2021 dataset. *: methods that use supervised labels. †: trunk larger than ResNet50. DINO baseline uses ViT-B/16.

DISC results. Table 2 reports DISC2021 results from the baseline methods published in [18] and SSCD. Our evaluation protocol obtains somewhat stronger results for the Multigrain baseline (3rd row). The first observation is that SSCD improves the baseline accuracy by $2\times$ to $5\times$ before

score normalization, demonstrating that copy detection benefits from specific architectural and training adaptations.

We present results on a few different SSCD models trained on ImageNet or DISC2021, using the three augmentation settings we propose. The intermediate model SimCLR_{CD} has all of our proposed changes except the entropy loss. SSCD_{large} model uses a larger descriptor size and a ResNeXt-101 trunk.

We evaluate SimCLR using both trunk and projected features, and find trunk features ($\mu AP = 13.1$) to outperform features from the projection head ($\mu AP = 9.4$) with and without score normalization. Further experiments (Appendix A) show the reverse when training with entropy loss: projected features have similar accuracy to trunk features, despite a much more compact representation.

The gain of SimCLR_{CD} ($\mu AP = 39.8$ without score normalization) over SimCLR (13.1) is decomposed in Section 5.5. Introducing the entropy loss in SSCD contributes an additional 10% absolute of μAP , which is further increased by stronger augmentations (+6.2%) and training on a dataset with less domain shift (+4.7%). These findings are confirmed after score normalization.

Copydays results. Table 3 reports results for baseline methods using publicly released models, but omit Multi-grain settings that we were unable to reproduce. We used published preprocessing settings for baselines and whitening. Our DINO results outperform published results.

model	trunk	dims	size	mAP	μAP
Multigrain [7]	ResNet50	1500	long 800	82.3	77.3
DINO [9]	ViT-B/16	1536	224 ²	82.8	92.3
DINO [9]	ViT-B/8	1536	320 ²	86.1	88.4
SSCD	ResNet50	512	short 288	86.6	98.1
SSCD _{large}	ResNeXt101	1024	long 800	93.6	97.1

Table 3. Copydays (CD10K) accuracy measured in mAP on the “strong” subset, and μAP on the full dataset.

The first SSCD result is with all settings from our DISC2021 experiments, where we resize the short side of each image to 288 pixels. With no tuning on this dataset, our method outperforms published results. We also show results for SSCD_{large} using a ResNeXt101 trunk and 1024 descriptor dimensions, at larger inference size. We report more results on CD10K in Appendix B.

In addition to state-of-the-art accuracy using the customary mAP ranking metric, our method provides a significant improvement in the global μAP metric, indicating better distance calibration. On high-resolution images that are common for image retrieval, we observe improved mAP but degraded μAP . SSCD descriptors are more compact than baselines.

5.5. Ablations

Comparison with SimCLR. We provide a stepwise comparison between SimCLR and our method in Table 4. SimCLR projection features are not particularly strong for this task until we apply several of our adaptations. SimCLR is unable to exploit a \mathbb{R}^{512} descriptor, only slightly outperforming its \mathbb{R}^{128} setting. SimCLR_{CD} represents our architectural and hyper-parameter changes before adding differential entropy representation. Differential entropy regularization alone adds +17.4% μAP and +12.9% μAP_{SN} , more than any other step.

name	Score normalization: method	No			Yes	
		dims	μAP	256d	μAP_{SN}	256d
SimCLR	trunk features	2048	13.1	7.3	33.9	26.8
	+ GeM pooling	2048	21.5	12.1	45.3	35.8
SimCLR	projection	128	9.4	9.4	17.3	17.3
	+ GeM pooling	128	11.1	11.1	18.8	18.8
	+ strong blur	128	14.1	14.1	26.0	26.0
	+ low temp	128	26.0	26.0	41.5	41.5
	+ 512d proj	512	27.5	27.5	43.5	43.5
SimCLR _{CD}	+ linear proj	512	33.0	32.4	51.6	50.5
SSCD	+ entropy loss	512	50.4	44.0	64.5	57.8
SSCD	+ adv. augs	512	55.5	49.7	71.0	65.8
SSCD	+ mixup	512	56.8	51.1	72.2	67.1

Table 4. Ablation from SimCLR to our method, showing DISC2021 μAP performance for models trained on ImageNet. To compare descriptors of different sizes, we also show metrics after PCA reduction to 256 dimensions.

Entropy weight. Table 5 compares how varying entropy loss weight (λ) affects copy detection accuracy, using SimCLR_{CD} as a baseline. Models for this experiment are trained using the strong blur augmentation setting.

model	μAP	μAP_{SN}	recall@1	MRR
SimCLR _{CD}	33.0	51.6	58.6	60.5
$\lambda = 1$	33.1	51.9	58.7	60.9
$\lambda = 3$	38.0	56.1	62.9	65.1
$\lambda = 10$	45.3	61.5	67.7	69.5
$\lambda = 30$	50.4	64.5	69.8	71.4

Table 5. DISC2021 accuracy metrics with varying entropy weight λ for models trained on ImageNet.

As the entropy weight increases, we see a corresponding increase in global accuracy metrics. We also see a similar increase in per-query ranking metrics, such as recall at 1 and mean reciprocal rank (MRR). The increase in ranking metrics demonstrates that differential entropy regularization improves copy detection accuracy in general, beyond creating a more uniform notion of distance.

In contrast to metric learning contexts where entropy regularization has been used, copy detection benefits from higher λ values. Our standard setting is $\lambda = 30$, while [19] reports reduced accuracy with $\lambda > 1$, and [42] uses values < 0.1 . At $\lambda > 40$, training becomes unstable, and tends to minimize the entropy loss at the expense of the InfoNCE loss: embeddings are uniformly distributed, but meaningless because image copies are not near anymore.

Additional ablations. We explore how batch size, training schedule, descriptor dimensions, and score normalization affect accuracy in Appendix A.

6. Discussion

Dimensional collapse. We find, similar to [29, 60], that SimCLR collapses to a subspace of approximately 256 dimensions when trained in 512 dimensions. Table 4 shows that SimCLR’s accuracy does not improve much when the descriptor size increases from 128 to 512 dimensions. SSCD’s entropy regularization resolves this collapse, and allows the model to use the full descriptor space.

Entropy regularization and whitening. SSCD is much more accurate than baselines when compared *without* whitening or similarity normalization: 47.8 μAP for $\lambda = 30$ when trained on ImageNet, versus 26.8 for an equivalent $\lambda = 0$ model. Both the entropy loss and post-training PCA whitening aim at creating a more uniform descriptor distribution. However PCA whitening can distort the descriptor space learned during training, particularly when many dimensions have trivial variance. Differential entropy regularization promotes an approximately uniform space, allowing the model to adapt to an approximately whitened descriptor during training, reducing the distortion whitening induces.

Uniform distribution as a perceptual prior. For most experiments in this work we focus on the μAP metric that requires a separation between matches and non-matches at a fixed threshold. However Table 5 shows that ranking metrics *also* improve with increased the entropy loss weight, *i.e.* better calibration across queries does not fully explain the benefit of entropy regularization.

Differential entropy regularization acts as a kind of prior, selecting for an embedding space that is uniformly distributed. We argue that, when applied to contrastive learning, this regularization is a **perceptual prior**, selecting for stronger copy detection representations. An ideal copy detection descriptor would map copies of the same image together, while keeping even semantically similar (same “class”) images far apart *i.e.* the descriptor distribution is uniform. This differs from the ideal properties of a representation for transfer learning to classification, where images depicting the same class should be nearby (a dense region) and well separated other classes (a sparse region between classes).

Visual results. Figure 2 shows a few retrieval results, where SSCD outperforms the vanilla SimCLR. The two first examples demonstrate the impact of more appropriate data augmentation at training time: SSCD ignores text overlays and blur/color balance. The two last examples show that SimCLR falls back on low-level texture matching (grass) when SSCD correctly recovers the source image.

Limitations. Our method is explicitly text-insensitive when training with text augmentation, and we find that it is somewhat text-insensitive even when trained without text augmentation. For this reason, SSCD is not precise when matching images composed entirely of text. Different photos of the same scene (*e.g.* of landmarks) may be identified as copies, even if the photos are distinct. Sometimes, images are combined to create a composite image or collage, where the copied content may occupy only a small region of the composite image. “Partial” copies of this kind are hard to detect with global descriptor models like SSCD, and local descriptor methods may be necessary in this case. Finally, matching at high precision often requires an additional verification step.

Ethical considerations. We focus our investigation on the DISC2021 dataset, which is thoughtful in its approach to images of people, using only identifiable photos of paid actors who gave consent for their images to be used for research. Copy detection for content moderation is adversarial. There is a risk that publishing research for this problem will better inform actors aiming to evade detection. We believe that this is offset by the improvements that open research will bring.

This technology allows scaling manual moderation, which helps protect users from harmful content. However, it can also be used for *e.g.* political censorship. We still believe that advancing this technology is a net benefit.

7. Conclusion

We presented a method to train effective image copy detection models. We have demonstrated architecture and objective changes to adapt contrastive learning to copy detection. We show that the differential entropy regularization dramatically improves copy detection accuracy, promoting consistent separation of image descriptors.

Our method demonstrates strong results on DISC2021, significantly surpassing baselines, and transfers to Copydays, yielding state-of-the-art results. Our method is efficient because it relies on a standard trunk, uses smaller inference sizes than are typical for image retrieval, and produces a compact descriptor. Additionally, its calibrated distance metric limits candidates for verification. We believe that these results demonstrate a unique compatibility between uniform embedding distributions and the task of copy detection.

References

- [1] Social media sites photo metadata test results 2019. <https://iptc.org/standards/photo-metadata/social-media-sites-photo-metadata-test-results-2019/>. Accessed: 2020-10-20. **2**
- [2] Hunt Allcott, Matthew Gentzkow, and Chuan Yu. Trends in the diffusion of misinformation on social media. *Research & Politics*, 6(2):2053168019848554, 2019. **1**
- [3] J Aythora, R Burke-Agüero, A Chamayou, S Clebsch, M Costa, J Deutscher, N Earnshaw, L Ellis, P England, C Fournet, et al. Multi-stakeholder media provenance management to counter synthetic media risks in news publishing. In *Proc. Intl. Broadcasting Convention (IBC)*, volume 1, page 8, 2020. **2**
- [4] Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. Neural codes for image retrieval. In *Proc. ECCV*, pages 584–599. Springer, 2014. **2**
- [5] Yoram Bachrach, Yehuda Finkelstein, Ran Gilad-Bachrach, Liran Katzir, Noam Koenigstein, Nir Nice, and Ulrich Paquet. Speeding up the xbox recommender system using a euclidean transformation for inner-product spaces. In *Proceedings of the 8th ACM Conference on Recommender systems*, pages 257–264, 2014. **6**
- [6] Jan Beirlant, E J. Dudewicz, L Gyor, and E.C. Meulen. Non-parametric entropy estimation: An overview. *International Journal of Mathematical and Statistical Sciences*, 6, 1997. **3**
- [7] Maxim Berman, Hervé Jégou, Andrea Vedaldi, Iasonas Kokkinos, and Matthijs Douze. Multigrain: a unified image embedding for classes and instances. *arXiv preprint arXiv:1902.05509*, 2019. **2, 4, 6, 7, 13**
- [8] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*, 2020. **3**
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. **6, 7, 13**
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proc. ICML*, pages 1597–1607. PMLR, 2020. **3, 4, 6**
- [11] Ondrej Chum, James Philbin, Josef Sivic, Michael Isard, and Andrew Zisserman. Total recall: Automatic query expansion with a generative feature model for object retrieval. In *Proc. ICCV*, 2007. **2**
- [12] Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*, 2017. **5**
- [13] Ingemar Cox, Matthew Miller, Jeffrey Bloom, Jessica Fridrich, and Ton Kalker. *Digital watermarking and steganography*. Morgan kaufmann, 2007. **2**
- [14] Zaroni Dias, Siome Goldenstein, and Anderson Rocha. Large-scale image phylogeny: Tracing image ancestral relationships. *Ieee Multimedia*, 20(3):58–70, 2013. **2**
- [15] Zaroni Dias, Anderson Rocha, and Siome Goldenstein. Image phylogeny by minimal spanning trees. *IEEE Transactions on Information Forensics and Security*, 7(2):774–788, 2011. **2**
- [16] Jing Dong, Wei Wang, and Tieniu Tan. Casia image tampering detection evaluation database. In *2013 IEEE China Summit and International Conference on Signal and Information Processing*, pages 422–426. IEEE, 2013. **2**
- [17] Matthijs Douze, Hervé Jégou, Harsimrat Sandhawalia, Laurent Amsaleg, and Cordelia Schmid. Evaluation of gist descriptors for web-scale image search. In *Proc. CIVR*, 2009. **2, 6**
- [18] Matthijs Douze, Giorgos Toliás, Ed Pizzi, Zoë Papanikopos, Lowik Chanussot, Filip Radenovic, Tomas Jenicek, Maxim Maximov, Laura Leal-Taixé, Ismail Elezi, et al. The 2021 image similarity dataset and challenge. *arXiv preprint arXiv:2106.09672*, 2021. **2, 5, 6**
- [19] Alaaeldin El-Nouby, Natalia Neverova, Ivan Laptev, and Hervé Jégou. Training vision transformers for image retrieval. *arXiv:2102.05644*, 2021. **3, 8**
- [20] Yunchao Gong, Liwei Wang, Ruiqi Guo, and Svetlana Lazebnik. Multi-scale orderless pooling of deep convolutional activation features. In *Proc. ECCV*, 2014. **2**
- [21] Albert Gordo, Jon Almazán, Jérôme Revaud, and Diane Larlus. Deep image retrieval: Learning global representations for image search. In *Proc. ECCV*, 2016. **2**
- [22] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020. **3**
- [23] Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. Accelerating large-scale inference with anisotropic vector quantization. In *International Conference on Machine Learning*, pages 3887–3896. PMLR, 2020. **1**
- [24] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proc. CVPR*, pages 9729–9738, 2020. **3**
- [25] Hervé Jégou and Ondřej Chum. Negative evidences and co-occurrences in image retrieval: The benefit of pca and whitening. In *Proc. ECCV*, pages 774–787. Springer, 2012. **2**
- [26] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *Proc. ECCV*. Springer, 2008. **2**
- [27] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. Exploiting descriptor distances for precise image search. Technical report, INRIA, 2011. **2, 5**
- [28] Amornched Jinda-Apiraksa, Vassilios Vonikakis, and Stefan Winkler. California-nd: An annotated dataset for near-duplicate detection in personal photo collections. In *2013 Fifth International Workshop on Quality of Multimedia Experience (QoMEX)*, pages 142–147. IEEE, 2013. **2**
- [29] Li Jing, Pascal Vincent, Yann LeCun, and Yuandong Tian. Understanding dimensional collapse in contrastive self-supervised learning. *arXiv:2110.09348*, 2021. **2, 4, 8, 13**

- [30] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *arXiv*, 2017. 1, 6
- [31] Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. The hateful memes challenge: Detecting hate speech in multimodal memes. *arXiv preprint arXiv:2005.04790*, 2020. 1
- [32] Xiyang Luo, Ruohan Zhan, Huiwen Chang, Feng Yang, and Peyman Milanfar. Distortion agnostic deep watermarking. In *CVPR*, 2020. 2
- [33] Daniel Moreira, Aparna Bharati, Joel Brogan, Allan Pinto, Michael Parowski, Kevin W Bowyer, Patrick J Flynn, Anderson Rocha, and Walter J Scheirer. Image provenance analysis at scale. *IEEE Transactions on Image Processing*, 27(12):6109–6123, 2018. 2
- [34] Eric Nguyen, Tu Bui, Viswanathan Swaminathan, and John Collomosse. Oscar-net: Object-centric scene graph attention for image attribution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14499–14508, 2021. 2
- [35] David Nister and Henrik Stewenius. Scalable recognition with a vocabulary tree. In *Proc. CVPR*, 2006. 2
- [36] Nam Thanh Pham, Jong-Weon Lee, Goo-Rak Kwon, and Chun-Su Park. Hybrid image-retrieval method for image-splicing validation. *Symmetry*, 11(1):83, 2019. 2
- [37] Filip Radenovic, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *Proc. CVPR*, 2018. 2
- [38] Filip Radenovic, Giorgos Tolias, and Ondrej Chum. CNN image retrieval learns from BoW: Unsupervised fine-tuning with hard examples. In *Proc. ECCV*, 2016. 2
- [39] Filip Radenović, Giorgos Tolias, and Ondřej Chum. Fine-tuning cnn image retrieval with no human annotation. *IEEE transactions on pattern analysis and machine intelligence*, 41(7):1655–1668, 2018. 2, 4
- [40] Eric Robertson, Haiying Guan, Mark Kozak, Yooyoung Lee, Amy N Yates, Andrew Delgado, Daniel Zhou, Timothee Kheyrkhan, Jeff Smith, and Jonathan Fiscus. Manipulation data collection and annotation tool for media forensics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 29–37, 2019. 2
- [41] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *IJCV*, 2015. 6
- [42] Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, and Hervé Jégou. Spreading vectors for similarity search. 2019. 1, 3, 8
- [43] Josef Sivic and Andrew Zisserman. Video google: A text retrieval approach to object matching in videos. In *null*, page 1470. IEEE, 2003. 2
- [44] NIST MediFor Team. Nimble challenge 2017 evaluation plan, 2017. 2
- [45] Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: the new data in multimedia research. *Commun. ACM*, 59:64–73, 2016. 6
- [46] Giorgos Tolias, Yannis Avrithis, and Hervé Jégou. To aggregate or not to aggregate: Selective match kernels for image search. In *Proc. ICCV*, 2013. 2
- [47] Giorgos Tolias, Yannis Avrithis, and Hervé Jégou. Image search with selective match kernels: aggregation across single and multiple images. *IJCV*, 116(3):247–261, 2016. 2
- [48] Giorgos Tolias, Tomas Jenicek, and Ondřej Chum. Learning and aggregating deep local descriptors for instance-level recognition. In *Proc. ECCV*, pages 460–477. Springer, 2020. 2, 6
- [49] Giorgos Tolias, Ronan Sifre, and Hervé Jégou. Particular object retrieval with integral max-pooling of cnn activations. In *Proc. ICLR*, pages 1–12, 2016. 2, 4
- [50] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy. *arXiv preprint arXiv:1906.06423*, 2019. 6
- [51] Matthieu Urvoy, Dalila Goudia, and Florent Atrousseau. Perceptual dft watermarking with improved detection and robustness to geometrical distortions. *IEEE Transactions on Information Forensics and Security*, 2014. 2
- [52] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018. 3
- [53] Shuang Wang and Shuqiang Jiang. Instre: a new benchmark for instance-level object retrieval and recognition. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 11(3):37, 2015. 2
- [54] Bihan Wen, Ye Zhu, Ramanathan Subramanian, Tian-Tsong Ng, Xuanjing Shen, and Stefan Winkler. Coverage—a novel database for copy-move forgery detection. In *2016 IEEE international conference on image processing (ICIP)*, pages 161–165. IEEE, 2016. 2
- [55] Ross Wightman, Hugo Touvron, and Hervé Jégou. Resnet strikes back: An improved training procedure in timm, 2021. 4
- [56] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. 3
- [57] Amy N Yates, Haiying Guan, Yooyoung Lee, Andrew P Delgado, Daniel F Zhou, Jonathan G Fiscus, et al. Nimble challenge 2017 evaluation data and tool. 2017. 2
- [58] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv:1708.03888*, 2017. 5
- [59] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. *arXiv preprint arXiv:1905.04899*, 2019. 5
- [60] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *Proc. ICML*, volume 139, pages 12310–12320. PMLR, 2021. 8, 13

- [61] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *Proc. ICLR*, 2018. 5
- [62] Peng Zhou, Xintong Han, Vlad I Morariu, and Larry S Davis. Learning rich features for image manipulation detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1053–1061, 2018. 2
- [63] Jiren Zhu, Russell Kaplan, Justin Johnson, and Li Fei-Fei. Hidden: Hiding data with deep networks. In *ECCV*, 2018. 2