



HAL
open science

Spatially-consistent Feature Matching and Learning for Heritage Image Analysis

Xi Shen, Robin Champenois, Shiry Ginosar, Ilaria Pastrolin, Morgane Rousselot, Oumayma Bounou, Tom Monnier, Spyros Gidaris, François Bougard, Pierre-Guillaume Raverdy, et al.

► To cite this version:

Xi Shen, Robin Champenois, Shiry Ginosar, Ilaria Pastrolin, Morgane Rousselot, et al.. Spatially-consistent Feature Matching and Learning for Heritage Image Analysis. *International Journal of Computer Vision*, 2022, 10.1007/s11263-022-01576-x. hal-03620996v2

HAL Id: hal-03620996

<https://hal.science/hal-03620996v2>

Submitted on 30 Mar 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Spatially-consistent Feature Matching and Learning for Heritage Image Analysis

Xi Shen, Robin Champenois, Shiry Ginosar, Ilaria Pastrolin, Morgane Rousselot, Oumayma Bounou, Tom Monnier, Spyros Gidaris, François Bougard, Pierre-Guillaume Raverdy, Marie-Françoise Limon, Christine Bénévent, Marc Smith, Olivier Poncet, K. Bender, Béatrice Joyeux-Prunel, Elizabeth Honig, Alexei A. Efros, Mathieu Aubry

Received: date / Accepted: date

Abstract Progress in the digitization of cultural assets leads to online databases that become too large for a human to analyze. Moreover, some analyses might be challenging, even for experts. In this paper, we explore two applications of computer vision to analyze historical data: watermark recognition and one-shot repeated pattern detection in artwork collections. Both problems present computer vision challenges which we believe to be representative of the ones encountered in cultural heritage applications: limited supervision is available, the tasks are fine-grained recognition, and the data comes in several different modalities. Both applications are also highly practical, as recognizing watermarks makes it possible to date and locate documents, while detecting repeated patterns allows exploring visual links between artworks. We demonstrate on both tasks the benefits of relying on deep mid-level features. More precisely, we define an image similarity score based on geometric verification of mid-level fea-

tures and show how spatial consistency can be used to fine-tune out-of-the-box features for the target dataset with weak or no supervision. This paper relates and extends our previous works [62,63]. Our code and data are available at <http://imagine.enpc.fr/~shenx/HisImgAnalysis/>.


Keywords Feature learning · Self-supervised learning · Artwork analysis · Watermark recognition

1 Introduction


Learning global image features has been successful in many image recognition tasks, such as image classification [31], landmark retrieval [29,55], metric learning [40,20,72] and few-shot classification [74,67]. In this paper, we highlight that matching and learning mid-level features and taking into account spatial information is better adapted to two cultural heritage applications required to recognize exactly repeated patterns across wide appearance changes, namely, historical watermark recognition and one-shot repeated pattern detection in artwork collections.

For watermark recognition (Figure 1(a)), we aim at retrieving the exact drawing (green) corresponding to the query photograph (blue) of a watermark from a very large database, which includes many similar watermarks each corresponding to a different category. In artwork collections, we tackle one-shot repeated pattern detection (Figure 1(b)), i.e. given a query detail (blue) we retrieve it in artworks with various styles (green). Both applications share similar Computer Vision challenges: (i) since annotations are very costly and require expert knowledge, there are only few category-level annotations available for training for watermark recog-

X. Shen, T. Monnier and M. Aubry
LIGM (UMR 8049), École des Ponts, Univ. Gustave Eiffel,
CNRS, Marne-la-Vallée, France
E-mail: firstname.lastname@enpc.fr
S. Ginosar, E. Honig and A. Efros
University of California, Berkeley, US
I. Pastrolin, M. Rousselot, C. Bénévent, M. Smith, O. Poncet
École Nationale des Chartres, Paris, France
M.-F. Limon
Archives Nationales, Paris, France
B. Joyeux-Prunel
Université de Genève, Switzerland
K. Bender
independent researcher, Gent, Belgium
F. Bougard
IRHT, Paris, France
O. Bounou and P.-G. Raverdy
INRIA, Paris, France
S. Gidaris
Valeo AI, Paris, France



(a) Historical watermark recognition



(b) One-shot repeated pattern detection

Fig. 1: We present a unified approach to two problems: (a) historical watermark recognition, and (b) repeated-pattern detection in artwork collections.

nition and no training annotations for repeated detail detection in art collections; (ii) the artwork details and the watermarks are represented in different visual domains; (iii) both tasks require fine-grained separation between similar but distinct classes. We believe these challenges are common to many cultural heritage applications.

The two problems we tackle are also highly practical for historians and archivists. For watermark recognition, the drawings we recognize come from large catalogs [10, 1] where they are associated with information such as the date and location of the paper fabrication, which are important clues to analyze historical documents. Similarly, repeated patterns in art collections are important to art historians and allow them to analyze influences, find provenance, and even establish authorship [11].

For each application, we introduced datasets for training and evaluation adapted to modern deep learning approaches [62, 63] and extend them in this paper. For historical watermark recognition, we first collected more than 6k new photographs for hundreds of classes, then searched and photographed systematically the exact original watermarks corresponding to hundreds of line drawings printed in Briquet’s classic catalog of more

than 16k [10] watermarks. This allows us to tackle at scale for the first time the scenarios of practical interest for scholars: one-shot instance recognition and cross-domain one-shot instance recognition among more than 16k fine-grained classes. For one-shot detection of the artwork’s details, we used a database of 1,587 artworks from Brueghel’s workshop [2, 36], selected 10 of the most commonly repeated details, and annotated their 273 occurrences in the full dataset for evaluation. Note that these visual patterns are near-duplicated but with various styles as artworks are in different media (*e.g.* oil, ink, chalk, watercolor) and on different materials (*e.g.* paper, panel, copper). In this paper, we provide additional results on a collection of 25,681 depictions of Venus [7]. It is more than an order of magnitude larger and more diverse than the Brueghel dataset, with works spreading over centuries and continents.

We demonstrate that both problems can be successfully addressed by leveraging the geometric consistency of mid-level feature matches both for training features and for scoring candidate correspondences. Our idea can be seen as revisiting the geometric verification of classical feature matches [66] for mid-level deep feature training and matching. More precisely, we propose a local matching score to compare pairs of images, which combines spatial consistency and local feature similarity. We also present a feature fine-tuning strategy that can be used without any annotations or with weak category-level annotations: we mine positive and negative pairs using geometric verification and use them to optimize a standard triplet loss.

We demonstrate experimentally that the proposed local matching score provides important performance gains for both tasks compared to global features, and that our feature fine-tuning – without any annotations on the Brueghel dataset and with weak category-level supervision on the watermark dataset – further improves the results. For example, on the challenging cross-domain one-shot recognition over more than 16k fine-grained watermark categories, for which average pooling fails (0% accuracy), our local matching score using classification pretrained features achieves 45% accuracy, and our fine-tuning further improves it to 55%.

Our code and data are available at our project page <http://imagine.enpc.fr/~shenx/HisImgAnalyses/>.

2 Related work

In this section, we first discuss the use of mid-level features for image recognition, then works related to our two applications, artwork analysis and watermark recognition.

Mid-level features for image recognition. Recent work analyzing the performance of CNNs [9, 24] suggests that they might ignore a large part of the spatial information present in the image, and rather work in a way similar to classical order-less bags-of-features methods [75, 30, 80]. This might not be suitable for recognizing repeated patterns such as watermarks and near-duplicated details in artworks, where the actual shape is key. To build on the local features learned by CNNs but consider spatial information, we follow an approach closely related to the classic spatial verification step introduced in Video Google [66] with SIFT features [43]. Rather than using SIFTs, which we found were not adapted to non-photographic depictions, we use intermediary deep features, that can be thought of as mid-level image features. Mid-level features [65, 16, 17] have been used in cross-domain matching [5], video instance segmentation [76] and coarse parametric transformation estimation [61]. Our feature fine-tuning also leverages the spatial structure of images: thus, it is related both to self-supervised feature learning methods that use spatial information to learn deep visual features in images [18, 47] or videos [76, 37], and to [58] that uses neighborhood consensus to learn correspondences from the correlation map. We adapted [58] to use as baseline.

Computer vision and art. There is a long-standing and fruitful collaboration between computer vision and art. On the synthesis side, promising results have been obtained for transferring artistic style to a photograph [33, 23, 81], or even trying to create art [21, 32]. On the analysis side, there are several efforts on the collection and annotation of large-scale art datasets [39, 46, 49, 78, 68, 44, 48], and using them for genre and authorship classification [39, 71, 68] or fragment alignment [48]. Others focus on applying and generalizing visual correspondence and object detection methods to paintings using both classical [64, 26, 13, 5, 60], as well as deep [12, 14, 77, 28] approaches. In particular, Yin et al. [79] used the same Brueghel data [2, 36] as us and annotate it to train detectors for five object categories (carts, cows, windmills, rowboats, and sailboats). Beyond classification, several projects have built search engines for artworks [3, 4]. An original and related approach was also proposed by [38] to relate artwork by estimating human poses rather than simply comparing abstract features. Our work is also related to [73] which performs pattern spotting in historical documents but for patterns with little appearance change. To the best of our knowledge, our work is the first to demonstrate a method capable of detecting small copied details across different depiction styles: this fine-grained analysis is the most relevant for

art historians, as it may allow them to discover influences, find provenance and establish authorship.

Historical watermark imaging and recognition. Manual tracing and back-lit photography are two standard ways to reproduce watermarks. Manual tracing [51, 10] consists in copying the watermark pattern on tracing paper and was typically used to create the classical catalogs of watermarks, which are nowadays aggregated in online databases such as Briquet Online [1] which includes specifically the drawings from [10] we worked with. Back-lit photography is the most common technique to acquire an actual photograph of a watermark. While a watermark is often barely visible, placing it in front of a light source and looking at the transmitted light reveals it more or less clearly.

This cross-modality between the drawings and the photographs one would like to identify, is one of the main challenges of watermark recognition. Several studies have focused on localizing and extracting the pattern of a watermark from a photograph [35, 34, 59], which could potentially be helpful to match a photograph to a database of drawings [56]. However, these approaches are complex, not end-to-end and have not yet been demonstrated on a large scale. Some works on watermark recognition focus on drawings. [56] uses histogram-based descriptors in a spirit similar to shape context [6], while more recent works use machine learning approaches, such as dictionary learning [50] or neural networks [52]. The study most similar to us is probably [52], that used a non-publicly available database [22] to train a convolutional neural network classifying 106,000 watermark into 12 coarse categories.

3 Method

In this section, we start by introducing our similarity score named “local matching score” which combines spatial consistency and local feature similarity in Section 3.1. Then, in Section 3.2, we detail our feature learning strategies for both the weakly supervised and the unsupervised scenario, which correspond to the two applications focused in this work. Precisely, for watermark recognition, a small number of image-level annotations are accessible and watermarks are approximately centered. For artworks, no annotations are available. The overview of our feature learning strategies is shown in Figure 2. Finally, we provide some implementation details in Section 3.3.

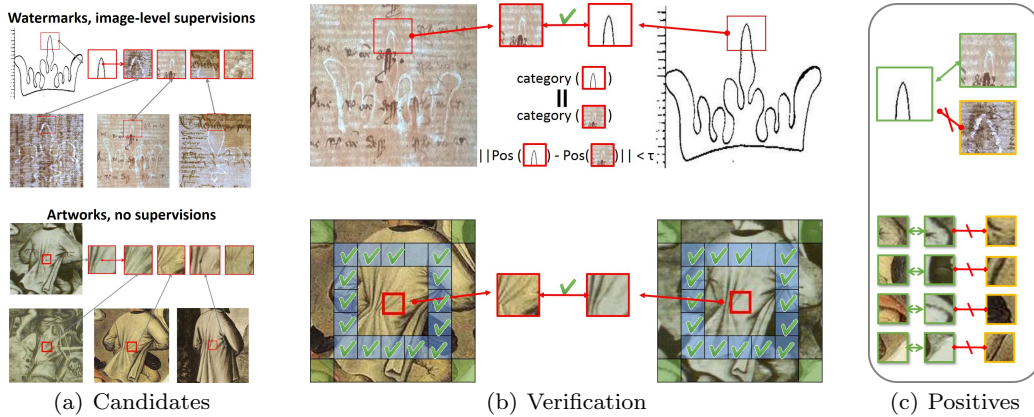


Fig. 2: Mid-level feature learning *with weak supervisions* or *without supervisions*. (a) Our approach relies on candidate correspondences obtained by matching the features of a proposal region (in red) to the full database. (b) In case of category-level supervisions (top), we verify whether matched pairs belong to the same category and their spatial difference is small. If no supervisions is available (bottom), the candidate correspondences are verified by matching the features of the verification region (in blue, note that, to make the figure visible, the blue regions are visualized as 5×5 squares while we use 10×10 squares in our algorithm.) of the query in the candidate images and checking for consistency. (c) Finally, we extract features from the positive (in green) and negative (in yellow) regions and use them to improve the features using a metric learning loss.

3.1 Local matching score

Let us assume we have access to a trained convolutional neural network that we want to use to compare two images I_1 and I_2 . The simplest solution is to use the similarity between aggregated features, typically the cosine similarity between average-pooled features for a ResNet network. However, such features will completely discard location information, which is likely not well adapted for fine-grained differentiation between similar classes. Another natural solution is then to consider convolutional features and use the similarity between the resulting descriptors. However, this gives the same importance to all local features, while some might be discriminative and others might not be, and this amounts to comparing only convolutional features at the same spatial location, i.e., assuming that the images are exactly aligned. This is a strong assumption and is not realistic in the cases we study: precisely centering and scaling the watermarks is difficult and the copied details in artworks might have undergone some amount of local deformation.

Thus, we propose a similarity metric to compare a pair of images I_1 and I_2 which relies on matching densely mid-level CNN features (e.g., *conv4* in ResNet [31]) computed at different scales. To select only discriminative features, we check that they correspond to reciprocal matches and discard them otherwise. More formally, for every local feature f_1^i in I_1 , we first find its best match in I_2 , f_2^i . We then check that it is a reciprocal

match, i.e. that the best match i.e. the best match in I_1 of f_2^i is f_1^i and we write \mathcal{M} the set of all reciprocal matches. We also write x_1^i the position of f_1^i in I_1 , and x_2^i the position f_2^i in I_2 .

Assuming that the two images are approximately aligned, which is true for example in the watermark dataset, we define the local matching score S by combining a *Spatial Consistency* score (SC) measuring the similarity between the positions x_1^i and x_2^i , and a *Feature Similarity* (FS) measuring the distance between f_1^i and f_2^i :

$$S(I_1, I_2) = \sum_{i \in \mathcal{M}} e^{-\underbrace{\frac{\|x_1^i - x_2^i\|^2}{2\sigma^2}}_{SC}} \underbrace{s(f_1^i, f_2^i)}_{FS} \quad (1)$$

where s is a feature level similarity – for which we use cosine similarity in all of our experiments –, and σ is a tolerance parameter.

If the images are not coarsely aligned, which is the case for repeated details detection in artworks, we need to further model the alignment between the two images. We do so by estimating a parametric transformation \mathcal{T} between the images, using RANSAC on the set of reciprocal matches \mathcal{M} . We then compute the score:

$$S(I_1, I_2) = \sum_{i \in \mathcal{M}} e^{-\underbrace{\frac{\|\mathcal{T}(x_1^i) - x_2^i\|^2}{2\sigma^2}}_{SC}} \underbrace{s(f_1^i, f_2^i)}_{FS} \quad (2)$$

\mathcal{T} is an Affine transformation and therefore has 6 parameters. Classical RANSAC is employed to estimate \mathcal{T}

: for each iteration, we sample 3 matches to compute a candidate transformation ; a match is considered as an inlier if its difference with the candidate transformation is smaller than 2 in the feature space; the final transformation is the one with maximum number of inliers after 1k iterations.

3.2 Spatially-consistent feature learning

The overview of our feature learning strategies is shown in Figure 2, which can be separated into three parts: 1. searching candidates of training samples (Figure 2(a)); 2. verification of candidates with weak category-level supervision or no supervision (Figure 2(b)); 3. training on selected samples (Figure 2(c)). We first explain the objective function we use to improve our features in Section 3.2.1, then present how we sample candidate matches in Section 3.2.2. Finally, we detail how to filter training samples in the candidate matches in Section 3.2.3.

3.2.1 Hypothesis and objective function

Our goal is to improve local features for matching. Starting with a standard pretrained deep feature extractor, we extract matching regions from the dataset that we can then circle back and use to improve the features. Our two key hypotheses are that: (i) our dataset includes large parts of images that contain repeated patterns, and (ii) the initial feature descriptor is good enough to extract some positive matches. We follow a metric learning approach. Assuming we have a set of positive pairs \mathcal{P} and a set of negative pairs \mathcal{N} , we learn our feature extractor f by minimizing a standard triplet loss:

$$\mathcal{L}(f) = \sum_{(n_1, n_2) \in \mathcal{N}} \max(1 - \lambda, s(f(n_1), f(n_2))) - \sum_{(p_1, p_2) \in \mathcal{P}} \min(\lambda, s(f(p_1), f(p_2))) \quad (3)$$

where the similarity measure s is the cosine similarity and λ is a hyper-parameter.

The main challenge for learning such features is defining the sets of positive and negative pairs. In the following sections, we show how to find these pairs in both weakly supervised and unsupervised cases.

3.2.2 Candidates

We randomly sample several query features from each image in the dataset and match them densely at every




Fig. 3: Different region configurations. red: query regions, blue: verification regions, green: positive regions.

scale to all the images in the dataset using cosine similarity. This can be done efficiently and in parallel for many queries using the normalized query features as the weights of a convolution applied to the normalized features. For each query, we select its top K matches in the entire dataset as candidate correspondences (Figure 2(a)).

3.2.3 Verification

We now explain how we filter the positive and negative pairs from these candidate local feature pairs both in the case with and without image-level supervision.

Image-level supervisions In the case where image level category labels are available, they can be used as a first cue to select positive matches: they can only exist between images belonging to the same category. Furthermore, if we assume that the images are coarsely aligned, as is the case for our watermark recognition problem, we can additionally filter only matches that are spatially consistent up to a given spatial threshold τ . This threshold allows to take into account misalignment between the source and target images and we study its influence in our experiments. Once a positive pair has been identified, we look for hard negatives by matching one of the positive features with all the photographs of watermarks from other categories, and select the the most similar feature as negative (Figure 2(b) top).

Unsupervised setting In the more challenging case where no image-level label is available and spatial alignment cannot be assumed, as in the case of repeated details in art collections, we rely on spatial consistency to verify the quality of candidates: a match will be considered valid if its neighbors agree with it. More precisely, let us assume we have a candidate match between features from the proposal region p_1 in image I_1 and a corresponding region p_2 in image I_2 , visualized in red in Figure 2(b) bottom. We define a *verification region* around p_1 , visualized in blue. Every feature in this region is individually matched to image I_2 , and votes for the candidate match if it matched consistently with p_2 . Summing the votes of all features in the verification region allows us to rank the candidate matches. A fixed

percentage of the candidates are then considered verified.

The choice of the verification region is important to the verification step. The key aspect is that the features in the verification region should be, as much as possible, independent of the features in the proposal region. On the other hand, having them too far apart would reduce the chances of the region being completely matched. For our experiments, we used the 10x10 feature square centred around the query region.

Finally, given a set of verified correspondences, we have to decide which features to use as positive training pairs (Figure 2(c) bottom). One possibility would be to directly use features from the proposal region, since they have been verified. However, since the proposal region has already been “used” once (to verify the matches), it does not bring enough independent signal to make quality hard positives. Instead, we propose to sample positives from a different *positive region*. We evaluated different configurations for the positive region, as visualized in Figure 3 (in green). We choose to keep only 4 positive pairs per verified proposal, positioned at the corners of a square and denote the different setups as P2 to P14, the number corresponding to the size of the square. We will show in the experiments (Section 5.2) that P12 and P14 perform better than the alternatives. The features from the positive regions (Figure 2(b) bottom in green) are then used as hard positives for feature fine-tuning (Figure 2(c) bottom).

For a given positive pair, we extract top-N closest features from the same image as negatives. This selects hard negatives and avoids any difference in the distribution of the depiction styles in our positive and negative samples. We chose a relatively high number ($N = 20$) of negatives to account for the fact that some of them might actually correspond to matching regions, for example, in the case of repeated elements, or to locations near the optimal match.

3.3 Implementation details

In all of our experiments, we used *conv4* features of the ResNet-18 [31] architecture and train the models with the Adam [41] optimizer with learning rate $1e-5$ and $\beta = [0.9, 0.99]$. For both applications, we use $K = 10$ candidate matches for each query.

For fine-tuning on watermark datasets, the initial features are pretrained on our classification dataset (see Section 4.1). The hyper-parameter λ is set to 1. Since the watermarks might be flipped and rotated with respect to the references, we consider matches with four rotated ($0^\circ, 90^\circ, 180^\circ, 270^\circ$) reference images and their

flipped versions. Each source image was resized to 352×352 so that it was represented by 22×22 features. To be robust to scale discrepancies, we matched the source features to features extracted from the target image resized at five scales, corresponding to square features of widths 16, 19, 22, 25 and 28. Using a single Geforce GTX 1080 Ti GPU, training converged in approximately 2 hours

For learning features on art collections, the initial features are pretrained on ImageNet [15]. The hyper-parameter λ is set to 0.8. Since repeated patterns might have large scale differences in different artworks, we used 7 different scales ranging from 20 to 80 features in the largest dimension, regularly sampled on two octaves with 3 scales per octave. From these candidates, the top 10% with the most votes from neighbours were considered verified. Note that these parameters might need to be adjusted depending on the diversity and size of the dataset, but we found that they performed well for the Brueghel data [2]. Using a single GPU Geforce GTX 1080 Ti, the training converged in approximately 10 hours.

4 Experiments on watermark recognition

In this section, we provide experimental results on watermarks recognitions. We first introduce the datasets we used, then present detailed results and analysis of the proposed local matching score and the proposed feature learning strategies.

4.1 Datasets

In this section, we explain how we created datasets to evaluate **(A) one-shot** and **(B) cross-domain** watermark recognition in the case where the watermarks are coarsely aligned. In the supplementary material available on our project page, we explain in detail the definition of classes and preprocessing procedures. More samples can be found in the project page.

Dataset A The goal of this dataset is to train and evaluate methods for one-shot fine-grained watermark recognition from photographs. It thus only includes photographs and is split into two parts: a first part that can be used for feature training or meta-training, with many examples of each watermark, and a second part, with few examples, to evaluate one-shot recognition.

We first created **dataset A-classification** by collecting 50 training and 10 validation images for 100 watermarks, which we found was large enough to perform pre-training / meta-training of CNNs. Examples of images from the same class are given in Figure 4(a).




Fig. 4: Images in watermark datasets: (a) Samples from the same class in A-classification; (b) three categories in A-one-shot. (c) - (e): Photograph, Drawing and Synthetic in B-cross-domain.

We then collected **dataset A-one-shot** with 100 other test classes, each made of 3 photographs: one ‘clean’ image without any writing and two standard test images. We used ‘clean’ images as references for fine-grained one-shot recognition, as they are representative of what archivists typically collect as reference images and allow us to ensure that recognition is not related to the writing style of the document. We show examples from three different categories in Figure 4(b) of references and associated test images. Note that some challenges can be directly seen from the examples: (i) many of the training and testing examples are cluttered and in poor condition, and (ii) some testing categories are highly similar.

Dataset B An important challenge is to use the drawings from existing watermark catalogs to perform recognition. Indeed, collecting watermarks and information

for such catalogs is a very tedious and expensive work. We focused on the Briquet catalog [10, 1]. The available images include additional information, such as IDs of watermarks, paper line positions, or complementary marks that can be found at another position on the paper sheet. Such information is valuable to experts, but cannot easily be used in an automatic recognition system based on a single photograph. We thus extracted the main part of the watermarks whenever it was clear and ended up with 16,753 drawings that could be used as references for photograph recognition. Examples of these drawings can be seen on the project page.

The challenge that we want to evaluate on this dataset is to recognize photographs in a drawing database. We thus collected **dataset B** by searching the original archives that provided the material for the Briquet catalog [10], in a specific city (Paris) and collected photographs for 240 classes (see examples in Figure 4(d)). We use 140 classes with a total of 463 photographs as training and keep the remaining 100 classes with 2 photographs in each class for testing. Because comparing photographs (Fig. 4(c)) directly with a line drawing (Fig. 4(d)) as a reference is very challenging, we also report results using a synthetic image generated from the drawing simply by using the average watermark color as background and making the drawing pattern lighter (Fig. 4(e)). In our experiments, we first compare the methods for 100-class one-shot cross-domain classification and then give the results for the even more challenging 16,753-class classification. In the following sections, we refer the 100-class one-shot cross-domain classification as the task on the **dataset B**. The 16,753-class classification is the task on **Briquet**.

4.2 One-shot cross-domain recognition

In this section, we present results on the one-shot recognition on the dataset A-one-shot and one-shot cross-domain recognition on the dataset B. In all experiments, we use a network pretrained on A-classification (100 classes, 50 training and 10 validation images for each class). The best performances were obtained with a strong dropout (0.7 ratio), which is not surprising given the relatively small size of our dataset. The validation accuracy is 98.8%, the mis-classified images being only very difficult or ambiguous cases, which shows that our 6k images dataset was large enough to train a good network for fine-grained watermark classification.

Comparison to one-shot recognition methods On dataset A-one-shot, which does not include any domain shift, we compare our method with some one-shot recognition methods:

Method \ Features Resolution	AvgPool		Concat		Local. Sim.	
	256	352	256	352	256	352
Baseline	69	78	74	86	75	81
Cosine Classifier [53,25]	84	83	82	80	84	82
Matching Networks [74]	74	76	76	78	82	84
Weights Prediction [25]	86	83	84	-	85	-
Ours Resolution 256			85			
Ours Resolution 352			90			

Table 1: Comparison with one-shot recognition approaches on dataset A-one-shot (200 images to classify in 100 categories unseen during training and described by a single ‘clean’ image). Accuracy in %. Our score based on local matches clearly outperforms all baselines.

- *Baseline*: directly using the features learned by training a network on the classification task and computing the dot-product as the distance between classification weights and features.
- *Cosine Classifier*: [25,53] have shown that the performance of the baseline can be improved if the dot-product operation (between classification weights and features) in the last linear layer of the network is replaced with the cosine similarity operation.
- *Matching Networks*: we tried the metric-learning approach of Matching Networks [74], performing meta-training to solve one-shot recognition tasks using a differentiable nearest-neighbor-like classifier.
- *Weights Prediction*: the one-shot recognition approach of Gidaris and Komodakis [25]. The key insight is that the novel class weights have two components: i) average features of annotated samples; ii) linear combinations of the base class weights, which are meta-learned during the training. It uses a feature extractor learned with a cosine-similarity-based classifier which remains frozen during the meta-training procedure.

For each feature, we report three different similarities:

- *AvgPool*: cosine similarity using the average pooled features.
- *Concat*: cosine similarity on the descriptor formed by the concatenation of all the spatial features.
- *Local Sim.*: computing the cosine similarity over each local feature individually, then averaging.

For all baselines, we report the best performance over *conv4/conv5* features. The optimal parameters and training strategy for each baseline and a qualitative analysis of our results are reported on the project page.

The results are in Table 1, our local matching score leads to 85% accuracy for 256×256 images, which is close to the best one-shot approach, Weights Prediction [25], but without any specific feature learning. This demonstrates the interest of our local matching score

Method	A	B-4(d)	B-4(e)	Time (s / Query)
Exact features comparison				
AvgPool	78	4	12	1
Concat	86	55	61	2
Local Sim.	81	56	65	2
Our local matching score				
Ours	90	66	72	15

Table 2: Comparison of our local matching score (eq. 1) with alternative feature similarities. Accuracy in % for one-shot recognition on dataset A-one-shot (A column) and one-shot cross-domain recognition on dataset B using either the drawing (B-4(d)) or our synthetic image (B-4(e)).

for one-shot fine-grained watermark recognition. The performance can further be boosted to 90% by resizing images to a larger resolution, 352×352 pixels, which was not possible for Weights Prediction for very large features due to computational cost. This is clearly above the performances of the other methods.

Comparison of feature similarities for one-shot cross-domain recognition In Table 2, we compare our local matching score (eq. 1) with alternative feature similarities on our two datasets. On dataset B, we use either the drawing (B-4(d)) or our synthetic image (B-4(e)) as reference. We always used the features trained for classification on dataset A-classification, and compare on each dataset to the similarities described in the previous paragraph (AvgPool, Concat, and LocalSim.). Our local matching score consistently outperforms all these baseline similarities. In our naive implementation, our approach is slower than direct feature comparison, but both can be mixed to obtain fast results on very large datasets.

4.3 Learning features for cross-domain recognition

We now focus on cross-domain recognition. We first compare our approach with different feature training strategies on our dataset B. We scale our watermark recognition method to the full Briquet catalog, showing we can perform classification with more than 16k classes.

Feature fine-tuning In Table 3, we compare the results from our fine-tuning strategy to different baselines:

- *Unsupervised (translation)*: in a spirit similar to [70], we translated the features in our target domain so that, on our training set, they have the same mean as the features from the target domain. We then use our score to perform nearest-neighbor classification.

Method	B-4(d)	B-4(e)
w/o Fine-tuning	66	72
Unsupervised (Translation) [70]	63	70
Supervised (Affine) [45,54]	64	72
Randomization [69,42]	53	75
Triplet-loss	64	65
NC-Net [58]	61	65
Ours (unsupervised)	60	71
Our weakly supervised fine-tuning		
$\tau = 0$	65	72
$\tau = 1/22$	75	81
$\tau = 3/22$	75	83
$\tau = 5/22$	73	77
$\tau = \text{inf}$	61	74

Table 3: Accuracy (in %) on one-shot cross-domain recognition with different feature fine-tuning approaches. τ is the hyper-parameter to constraint the misalignment of matches (Sec. 3.2.3). $\tau = 3/22$ indicates we tolerates a difference of 3 features while the entire feature map is 22×22 . We experiment with different reference images (“B-4(e)” : synthetic image ; “B-4(d)” : drawing). Note that all the approaches except NC-Net [58] use local matching score

- *Supervised (affine)*: since we have aligned images from both domains, we can adapt features from the source and target domains in a supervised way, similarly to [45,54]. We found that a simple affine adaptation gave the best results, likely because of the small size of our dataset.
- *Randomization*: we also report results using randomized generated synthetic images from the drawings. The randomized synthetic image S are generated by computing $S = B + R \times (G * E)$, where B is a background sampled from photographs of paper without watermarks, G is a Gaussian filter, R is a random image and E is the binary watermark pattern extracted from the drawing. Generated images can be found in the supplementary material. We trained a standard classifier on the generated images. Such an approach has been shown to be very successful for example for 3D pose estimation [69,42].
- *Triplet-loss*: similarly to our method, we tried to improve the features using a triplet loss on local features, but using as positive all aligned features in the images from the same category.
- *NC-Net* [58]: while it was not initially designed for domain adaptation, we trained NC-Net on our database because of the intuition that, similarly to our method, it is able to learn to leverage spatial information. We use our pre-trained ResNet for the feature extractor and freeze it during the training. The other parts are kept the same as the category level matching model proposed in [58]. The positive pairs are composed

Method	Briquet-4(e)		Briquet-4(e)+Fine-tuning	
	acc.@1	acc.@1000	acc.@1	acc.@1000
AvgPool	0	16	0	21
Concat	27	77	29	82
Local Sim.	28	80	28	83
Ours N=1000	45	80	54	83
Ours N=inf	44	86	55	91

Table 4: Top-1 and top-1000 accuracy on our Briquet dataset with different models (“Briquet-4(e)” : model trained on classification on dataset A-classification ; “Briquet-4(e)+Fine-tuning” : fine-tuned model). Our approach first applies Local Similarity to obtain N top ranked references and then uses our score to re-rank the N references.

with one image from each domain, which results in 463 pairs in the training set. The training converges in 20 epochs. We then consider the sum of the scores over all correspondences as the score between a pair of images.

We detail these approaches and the training details in the supplementary material available in our project website. All results except for NC-Net are reported for matching done with our local matching score – using other metrics leads to worse performance. The effectiveness of this metric might be the reason why standard domain adaptation approaches only marginally improve performances over the baseline. Another possible reason is the small size of our training set (463 photographs, 140 references). On the contrary, our fine-tuning strategy boosts performances by a clear margin.

Large-scale recognition on Briquet We finally evaluate one-shot cross-domain recognition using the test photographs of our dataset B and our full curated version of the Briquet dataset as reference. This comparison with 16,753 fine-grained classes is very challenging, but also corresponds to a realistic scenario for watermark recognition. We use our synthetic images (Figure 4(e)) to represent the drawings. Table 4 shows the top-1 and top-1000 accuracy using the different baseline similarities described in Section 4.2. Since our local matching score is computationally more expensive than the baselines, we evaluate a two-step procedure for recognition: for each test photograph, we first select the top- N candidate classes using Local Sim., then re-rank them using our local matching score. Since the local similarity – *i.e.* averaging the cosine distance between the local features over the images – yields the best results, we use it to perform the first step of selection. Using $N = 1000$ leads to performances similar to directly applying our method to the full 16,753 classes.

Ranking the reference drawings with Local Similarity takes approximately 3s and reranking the top-1000



Fig. 5: Top-3 matches retrieved with our approach with synthetic reference (Figure 4(e)) on Briquet dataset (16,753 categories). For all the four queries, our approach is able to recognize the correct categories.

with our local matching score takes 37s on a single GPU Geforce GTX 1080 T. This is acceptable for practical use, and we thus believe the application of our algorithm will be a game-changer and widen considerably the potential use of watermark analysis, which until now has been limited to a small number of experts.

Web application The ultimate goal of watermark recognition is to develop an application to simplify the search of watermarks. We developed a first version of such a web application [8] which can be accessed at <https://filigranes.inria.fr/>.

Qualitative results Typical examples of top-3 matches with synthetic references (Figure 4(e)) are shown in Figure 5. Although the task is challenging, as some queries are cluttered and some other categories are highly similar to the ground-truth, our approach allows to recognize the correct categories. Note that more visual results can be found in the project page.

5 Experiments on one-shot art pattern detection

In this section, we present results on one-shot art pattern detection. We first introduce the datasets we used,

then present detailed results on one-shot detection on Brueghel [2] dataset. We finally show qualitative results on the Venus dataset.

5.1 Datasets

Brueghel The Brueghel [2] dataset contains 1,587 artworks done in different media (*e.g.* oil, ink, chalk, watercolour) and on different materials (*e.g.* paper, panel, copper), describing a wide variety of scenes (*e.g.* landscape, religious, still life). This dataset is especially adapted for our task since it assembles paintings from artists related to the same workshop, who thus had many interactions with each other, and includes many copies, preparatory drawings, and borrowed details.

With the help of our art history collaborators, we selected 10 of the most commonly repeated details in the dataset and annotated the visual patterns in the full dataset using the VGG Image Annotator tool [19]. The 10 annotated patterns can be seen in Figure 6 as queries (blue boxes), and our full annotations can be found in the project page. We were careful to select diverse patterns, and for each of them to annotate only duplicates, and not full object classes. Note for example, that for the horses and lion classes, we annotated separately two variants of the details (front and back facing lion, front and right facing horse). This resulted in 273 annotated instances, with a minimum of 11 and a maximum of 57 annotations per pattern.

These annotations allow us to evaluate one-shot duplicate detection results. In our evaluation, we use an IoU threshold of 0.3 for positives, because precise annotations of the bounding boxes in different environment is difficult and approximate detection would be sufficient for most applications. In practice, our detected bounding boxes, visualised in Figure 6 (green boxes) often appear more consistent than the annotations. We compute the Average Precision for each query, average them per class and report class level mean Average Precision (mAP).

Venus To demonstrate the generality of our approach, we show results on a very different database, containing 25,681 depictions of Venus [7]. This database is much larger than the Brueghel database and more diverse since it includes depictions from many different periods and very diverse styles. Because it is thematic, it however includes many cases of revisited artwork and artistic citations. It is thus much better suited to our purpose than standard artwork databases. The size and diversity of the database makes annotation unpractical, but we report qualitative results which demonstrate matches of interest beyond the case of copied details.



Fig. 6: Detection example with our trained features on the Brueghel dataset. We show the top 5 matches (in transparency) as well as the annotations (green bounding box) for one example of query from each of our 10 annotated categories. Notice how the style of the matches can be different from the one of the query.

Feature \ Method	Cosine similarity	Local matching
ImageNet pre-taining	58.0	54.8
Context Prediction [18]	58.8	64.29
Ours fine-tuning	75.3	76.4

Table 5: Experimental results on Brueghel, IoU > 0.3.

5.2 One-shot detection

We evaluated our feature learning strategy using one-shot detection mainly on the Brueghel dataset which was small enough to be manually annotated. We performed one-shot detection simply by computing dense features on the dataset and computing their cosine similarity with the features corresponding to the query. The query was resized so its largest dimension in the feature map would be 8. Note that unlike standard deep detection approaches [27, 57], we do not use region proposals because we want to be able to match regions which do not correspond to objects.

Comparison and analysis. We compare one-shot detection performance with different features using both cosine similarity and the score described in equation 1. In Figure 8 we show an example of the top-6 matches for the same query using different approaches. In the first line, the matches obtained with the cosine similarity between features trained on ImageNet are all from styles similar to the style of the query, and while three of them include horses, they are not in the same pose as the query. On the contrary, the matches with our trained features, shown on the second line, mainly associate horses in the exact same position, including in depictions with different styles. Finally, the results with the same trained features but using the local matching




Fig. 7: Evolution of the mean Average Precision for one-shot matching on the Brueghel dataset during training. Performance decreases after a few iterations for settings where we extract positive regions correlated with the proposal region.

score, shown on the third line, show a slight improvement.

The corresponding quantitative results are presented in Table 5 and confirm these observations. Indeed, learning our features improves the score by approximately 30%. The discovery procedure and score provide an additional boost.

Positive region configuration. We now focus on evaluating the different positive region settings described in Section 3.2.3 and Figure 3. For each of them, we analyse the performance of the features on one-shot learning on the Brueghel dataset and its evolution during training. The results can be seen in Figure 7. Interestingly, the performance always initially improves over ImageNet features. However, when the positive region is close to the proposal region, the performance decreases after some iterations of our training procedure and ends up



Fig. 8: From a single query, shown on the left, we show the detection results obtained with cosine similarity with ImageNet feature (a) and our trained features (b) as well as the ones (c) obtained with our features and the local matching score presented in Section 3.1.




Fig. 9: Visualizing matched correspondences between a pair of images during training.

with worse performance than the initial features. However, if the positive region is far enough from the query (P12 and P14), the performance improves much more and does not subsequently deteriorate. We thus use P12 for all our other experiments.

Training visualization. To visualise the influence of the feature training for our discovery task, we selected a pair of matching images from the Brueghel dataset and visualised the the inlier set \mathcal{I} (defined in equation 1) at different steps of training in Figure 9. During training, larger and larger parts of the images can be matched in a consistent way. This shows both the efficiency of our feature training and its relevance for our task.

Qualitative results. Qualitative results using our approach for each of the 10 details we annotated on the Brueghel dataset are shown in Figure 6. It gives a sense of the difficulty of the task we target and the quality of the results we obtain. Note for example how the matches are of different styles, and how the two types of lions (top row) and the two types of horses (bottom

row) are differentiated. In the following, we compare these results with baselines and analyze the differences.

In Figure 10, we show qualitative results on the Venus dataset. In this case, the retrieved images might not be exact copies but sets of works that revisit at different periods the same model, sometime borrowing only part of the works and keeping details, such as Flora from the *Primavera* from Botticelli in the last row, sometime focusing on the composition such as for *La Venus del espejo* from Velázquez in the 7th row. Note how our method is able to retrieve relevant results in all of these cases, despite strong variations. These results showcase the interest of our method for Art Historians to explore such collections where one can hardly expect exhaustive annotations and meta-data.

Limitations Our method has several limitations. First, the approach need to compare image pairs, which is not possible for fast retrieval applications. Second, in both applications, our feature learning strategy requires a good feature initialization. For watermark recognition, we start from features learned from A-classification. For artwork detection, we start from ImageNet [15] pre-trained features.

6 Conclusion

In this paper, we have introduced two cultural heritage applications and associated datasets: watermark recognition and one-shot repeated art pattern detection. We have presented a unified mid-level feature training and matching method for both applications. We showed that the proposed local matching score clearly improves over standard feature similarities and that our mid-level feature fine-tuning approach can further boost performances on both tasks.

Acknowledgements This work was supported in part by ANR project EnHerit ANR-17-CE23-0008, PSL Filigrane pour tous project, project Rapid Tabasco, gifts from Adobe to Ecole des Ponts. We thank Minsu Cho, Pascal Monasse and Renaud Marlet for fruitful discussions, and Kenny Oh and Davienne Shields for thier help on defining the task and building the Brueghel dataset.

References

1. Briquet online. http://www.ksbm.oew.ac.at/_scripts/php/BR.php 2, 3, 7
2. Brueghel family: Jan brueghel the elder.” the brueghel family database. university of california, berkeley. <http://www.janbrueghel.net/>. Accessed: 2018-10-16 2, 3, 6, 10
3. imgs.ai. <https://imgs.ai/> 3




Fig. 10: Each row shows the top 6 matches on the Venus database given queries outlined in blue on the left images.

4. timemachine. diamond.timemachine.eu 3
5. Aubry, M., Russell, B.C., Sivic, J.: Painting-to-3d model alignment via discriminative visual elements. *ACM Transactions on Graphics (ToG)* (2014) 3
6. Belongie, S., Malik, J., Puzicha, J.: Shape context: A new descriptor for shape matching and object recognition. In: *NeurIPS* (2001) 3
7. Bender, K.: Distant viewing in art history. a case study of artistic productivity. *International Journal for Digital Art History* (1) (2015) 2, 10
8. Bounou, O., Monnier, T., Pastrolin, I., Shen, X., Ben-event, C., Limon-Bonnet, M.F., Bougard, F., Aubry, M., Smith, M., Poncet, O., et al.: A web application for watermark recognition. *Journal of Data Mining and Digital Humanities* (2020) 10
9. Brendel, W., Bethge, M.: Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. In: *ICLR* (2019) 3
10. Briquet, C.M.: *Les filigranes* (1907) 2, 3, 7
11. Castellano, G., Lella, E., Vessio, G.: Visual link retrieval and knowledge discovery in painting datasets. *Multimedia Tools and Applications* (2021) 2
12. Crowley, E.J., Parkhi, O.M., Zisserman, A.: Face painting: querying art with photos. In: *BMVC* (2015) 3
13. Crowley, E.J., Zisserman, A.: Of gods and goats: Weakly supervised learning of figurative art. In: *BMVC* (2013) 3
14. Crowley, E.J., Zisserman, A.: The art of detection. In: *ECCV* (2016) 3
15. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *CVPR* (2009) 6, 12
16. Doersch, C., Gupta, A., Efros, A.A.: Mid-level visual element discovery as discriminative mode seeking. In: *NeurIPS* (2013) 3
17. Doersch, C., Gupta, A., Efros, A.A.: Context as supervisory signal: Discovering objects with predictable context. In: *ECCV* (2014) 3
18. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: *ICCV* (2015) 3, 11
19. Dutta, A., Gupta, A., Zisserman, A.: VGG image annotator (VIA). <http://www.robots.ox.ac.uk/~vgg/software/via/> (2016) 10
20. Elezi, I., Vascon, S., Torcinovich, A., Pelillo, M., Leal-Taixé, L.: The group loss for deep metric learning. In: *ECCV* (2020) 1
21. Elgammal, A., Liu, B., Elhoseiny, M., Mazzone, M.: Can: Creative adversarial networks, generating” art” by learning about styles and deviating from style norms. *arXiv* (2017) 3
22. Frauenknecht, E., Stiegler, M.: *Wzis – wasserzeichen-informationssystem: Verwaltung und präsentation von wasserzeichen und ihrer metadaten. Kodikologie und Paläographie im Digitalen Zeitalter 3: Codicology and Palaeography in the Digital Age 3* (2015) 3
23. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: *CVPR* (2016) 3
24. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In: *ICLR* (2019) 3
25. Gidaris, S., Komodakis, N.: Dynamic few-shot visual learning without forgetting. In: *CVPR* (2018) 8
26. Ginosar, S., Haas, D., Brown, T., Malik, J.: Detecting people in cubist art. In: *Workshop at ECCV* (2014) 3
27. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *CVPR* (2014) 11
28. Gonthier, N., Gousseau, Y., Ladjal, S., Bonfait, O.: Weakly supervised object detection in artworks. *arXiv* (2018) 3
29. Gordo, A., Almazan, J., Revaud, J., Larlus, D.: End-to-end learning of deep visual representations for image retrieval. *IJCV* (2017) 1
30. Grauman, K., Darrell, T.: Pyramid match kernels: Discriminative classification with sets of image features. In: *ICCV* (2005) 3
31. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR* (2016) 1, 4, 6
32. Hertzmann, A.: Can computers create art? In: *Arts. Multidisciplinary Digital Publishing Institute* (2018) 3
33. Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B., Salesin, D.H.: Image analogies. In: *SIGGRAPH* (2001) 3
34. Hiary, H.: Paper-based watermark extraction with image processing. Ph.D. thesis (2008) 3
35. Hiary, H., Ng, K.: A system for segmenting and extracting paper-based watermark designs. *International Journal on Digital Libraries* (2007) 3
36. Honig, E.: *Jan Brueghel and the Senses of Scale*. Pennsylvania State University Press (2016) 2, 3
37. Jabri, A., Owens, A., Efros, A.A.: Space-time correspondence as a contrastive random walk. In: *NeurIPS* (2020) 3
38. Jenicek, T., Chum, O.: Linking art through human poses. In: *ICDAR* (2019) 3
39. Karayev, S., Trentacoste, M., Han, H., Agarwala, A., Darrell, T., Hertzmann, A., Winnemoeller, H.: Recognizing image style. *arXiv* (2013) 3
40. Kim, S., Kim, D., Cho, M., Kwak, S.: Proxy anchor loss for deep metric learning. In: *CVPR* (2020) 1
41. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv* (2014) 6
42. Loing, V., Marlet, R., Aubry, M.: Virtual training for a real application: Accurate object-robot relative localization without calibration. *IJCV* (2018) 9
43. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *IJCV* (2004) 3
44. Mao, H., Cheung, M., She, J.: Deepart: Learning joint representations of visual arts. In: *ACM Multimedia* (2017) 3
45. Massa, F., Russell, B.C., Aubry, M.: Deep exemplar 2d-3d detection by adapting from real to rendered views. In: *CVPR* (2016) 9
46. Mensink, T., Van Gemert, J.: The rijksmuseum challenge: Museum-centered visual recognition. In: *ICMR* (2014) 3
47. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: *ECCV* (2016) 3
48. Paumard, M.M., Picard, D., Tabia, H.: Jigsaw puzzle solving using local feature co-occurrences in deep neural networks. In: *ICIP* (2018) 3
49. Picard, D., Gosselin, P.H., Gaspard, M.C.: Challenges in content-based image indexing of cultural heritage collections. *IEEE Signal Processing Magazine* (2015) 3
50. Picard, D., Henn, T., Dietz, G.: Non-negative dictionary learning for paper watermark similarity. In: *ACSSC* (2016) 3
51. Piccard, G.: *Die Wasserzeichenkartei Piccard im Hauptstaatsarchiv Stuttgart: Wasserzeichen Buchstabe P*. (1977) 3
52. Pondenkandath, V., Alberti, M., Eichenberger, N., In-gold, R., Liwicki, M.: Identifying cross-depicted historical motifs. *arXiv* (2018) 3

53. Qi, H., Brown, M., Lowe, D.G.: Low-shot learning with imprinted weights. In: CVPR (2018) [8](#)
54. Rad, M., Oberweger, M., Lepetit, V.: Feature mapping for learning fast and accurate 3d pose inference from synthetic images. In: CVPR (2018) [9](#)
55. Radenović, F., Tolias, G., Chum, O.: Fine-tuning cnn image retrieval with no human annotation. TPAMI (2018) [1](#)
56. Rauber, C., Tschudin, P., Pun, T.: Retrieval of images from a library of watermarks for ancient paper identification. In: Electronic Visualisation and the Arts (1997) [3](#)
57. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NeurIPS (2015) [11](#)
58. Rocco, I., Cimpoi, M., Arandjelović, R., Torii, A., Papdja, T., Sivic, J.: Neighbourhood consensus networks. In: NeurIPS (2018) [3](#), [9](#)
59. Said, J., Hiary, H.: Watermark location via back-lighting modelling and verso registration. Multimedia Tools and Applications (2016) [3](#)
60. Seguin, B., diLenardo, I., Kaplan, F.: Tracking transmission of details in paintings. In: DH (2017) [3](#)
61. Shen, X., Darmon, F., Efros, A.A., Aubry, M.: Ransac-flow: generic two-stage image alignment. In: ECCV (2020) [3](#)
62. Shen, X., Efros, A.A., Aubry, M.: Discovering visual patterns in art collections with spatially-consistent feature learning. In: CVPR (2019) [1](#), [2](#)
63. Shen, X., Pastrolin, I., Bounou, O., Gidaris, S., Smith, M., Poncet, O., Aubry, M.: Large-scale historical watermark recognition: dataset and a new consistency-based approach. In: ICPR (2020) [1](#), [2](#)
64. Shrivastava, A., Malisiewicz, T., Gupta, A., Efros, A.A.: Data-driven visual similarity for cross-domain image matching. SIGGRAPH ASIA (2011) [3](#)
65. Singh, S., Gupta, A., Efros, A.A.: Unsupervised discovery of mid-level discriminative patches. In: ECCV (2012) [3](#)
66. Sivic, J., Zisserman, A.: Video google: A text retrieval approach to object matching in videos. In: ICCV (2003) [2](#), [3](#)
67. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: NeurIPS (2017) [1](#)
68. Strezoski, G., Worring, M.: Omniart: multi-task deep learning for artistic data analysis. arXiv (2017) [3](#)
69. Su, H., Qi, C.R., Li, Y., Guibas, L.J.: Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views. In: ICCV (2015) [9](#)
70. Sun, B., Feng, J., Saenko, K.: Return of frustratingly easy domain adaptation. In: AAAI (2016) [8](#), [9](#)
71. Tan, W.R., Chan, C.S., Aguirre, H.E., Tanaka, K.: Ceci n'est pas une pipe: A deep convolutional network for fine-art paintings classification. In: ICIP (2016) [3](#)
72. Teh, E.W., DeVries, T., Taylor, G.W.: Proxynca++: Revisiting and revitalizing proxy neighborhood component analysis. In: ECCV (2020) [1](#)
73. Úbeda, I., Saavedra, J.M., Nicolas, S., Petitjean, C., Heutte, L.: Pattern spotting in historical documents using convolutional models. In: Proceedings of the 5th International Workshop on Historical Document Imaging and Processing, pp. 60–65 (2019) [3](#)
74. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: NeurIPS (2016) [1](#), [8](#)
75. Wallraven, C., Caputo, B., Graf, A.: Recognition with local features: the kernel recipe. In: ICCV (2003) [3](#)
76. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: CVPR (2019) [3](#)
77. Westlake, N., Cai, H., Hall, P.: Detecting people in artwork with cnns. In: ECCV (2016) [3](#)
78. Wilber, M.J., Fang, C., Jin, H., Hertzmann, A., ColloMosse, J., Belongie, S.J.: Bam! the behance artistic media dataset for recognition beyond photography. In: ICCV (2017) [3](#)
79. Yin, R., Monson, E., Honig, E., Daubechies, I., Maggioni, M.: Object recognition in art drawings: Transfer of a neural network. In: ICASSP (2016) [3](#)
80. Zhang, J., Marszałek, M., Lazebnik, S., Schmid, C.: Local features and kernels for classification of texture and object categories: A comprehensive study. ICCV (2007) [3](#)
81. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. ICCV (2017) [3](#)