

A Web Application for Watermark Recognition

Oumayma Bounou^{1, 2, 3}, Tom Monnier^{2, 3}, Ilaria Pastrolin², Xi Shen³, Christine Bénévent², Marie-Françoise Limon-Bonnet^{2, 5}, François Bougard⁴, Mathieu Aubry³, Marc Smith², Olivier Poncet², Pierre-Guillaume Raverdy¹

¹Inria, France

²École Nationale des Chartes-PSL (Centre Jean-Mabillon, EA 3624), France

³LIGM, IMAGINE, Ecole des Ponts, Univ Gustave Eiffel, CNRS, Marne-la-Vallée, France

⁴CNRS-IRHT, France

⁵Archives Nationales, France

Abstract

The study of watermarks is a key step for archivists and historians as it enables them to reveal the origin of paper. Although highly practical, automatic watermark recognition comes with many difficulties and is still considered an unsolved challenge. Nonetheless, Shen et al. [2019] recently introduced a new approach for this specific task which showed promising results. Building upon this approach, this work proposes a new public web application dedicated to automatic watermark recognition entitled *Filigranes pour tous*. The application not only hosts a detailed catalog of more than 17k watermarks manually collected from the French National Archives (Minutier central) or extracted from existing online resources (Briquet database), but it also enables non-specialists to identify a watermark from a simple photograph in a few seconds. Moreover, additional watermarks can easily be added by the users making the enrichment of the existing catalog possible through crowdsourcing. Our Web application is available at <http://filigranes.inria.fr/>.

Keywords

cross-domain recognition; deep learning; watermark recognition; web application, paper analysis

I INTRODUCTION

Paper appears in China at the beginning of the second century B.C. The principle of manufacture is based on a paste made with water in which various vegetable fibres are incorporated. A floating mould made of bamboo mesh or wood collects sheets of paste that are dried standing. The technological transfer to the West took place via the Arabs from the middle of the eighth century. In Central Italy, in Fabriano, it is accompanied in the 13th century by an important innovation: floating moulds are replaced by metal grids called shapes that work in identical pairs for manufacturing procedures. From the 13th to the 19th century, most European paper-makers identified their production by adding a small wire, formed in a specific shape, into the paper mold. The paper was thinner where the wire was present, and the motif can thus be seen when looking at light transmitted through the paper. The wires were however not very stable, and had to be replaced regularly with new ones, typically using a related but different shape. The watermark is specific to each manufacturer and in an original way (compared to many other medieval manufactured products) constitutes a traditional identification of our trademarks. As said by a remarkable historian of the watermarks, Charles-Moïse Briquet: ‘Every sheet of watermarked paper carries its birth certificate’.

This was therefore the first trademark of the ancient world that allows us, knowing the place, the approximate dating and sometimes the name of the production, to understand the traffic and the economy of this special material and to date documents (texts or drawings) which do not contain any explicit chronological indication. Identifying the watermark in a paper and pairing it with other metadata such as paper dimensions or thickness thus enables to identify the paper maker and date the paper production quite precisely by searching existing watermark design catalogs, making it a key tool for archivists, historians, curators and all those who are interested in the trade of written documents (booksellers, antiquarians, etc.). We don't know how many watermarks have existed in the West since the 13th century. Probably tens of thousands until 1500, hundreds of thousands until the French Revolution: the development of printing, the progress of the government in writing and the growing literacy of the population has indeed led to an increase and multiplication of producers in all countries.

In this work, we introduce a new web application for automatic large-scale watermark recognition based on the approach developed by Shen et al. [2019]. This application enables archivists and historians not only to search for a particular watermark in our large database of watermarks, but also to identify a watermark from a simple photograph. Our database is composed of a new photograph dataset manually collected from the Minutier central des notaires de Paris (more than 2,250 photos of almost 400 watermarks) as well as the well-known Briquet design catalog (16,112 watermarks). Besides, the existing database is designed to be easily enriched by users with their own watermark pictures.

II MOTIVATIONS

Since the end of the 19th century, many researchers have sought to establish the most comprehensive watermarks catalogs possible on the basis of the paper stored in archives and libraries in Western Europe. There has been many attempts to build a general database. While they have always remained unfulfilled, the first of them was a masterpiece. Charles-Moïse Briquet (1839-1918, the greatest man in the history of Switzerland after Guillaume Tell and Roger Federer, according to Neil Harris, professor of Bibliography at the University of Udine in Italy and passionate filigranist) visited 235 archives and collections of manuscripts and printed books, worked on 30,840 volumes and 1,432 unrelated documents (particularly letters).

He made 45,930 tracings, which he reproduced on a more robust paper now in the Briquet archives in Geneva. In 1907, he published *Les filigranes, dictionnaire historique des marques de papier dès leur apparition vers 1282 jusqu'en 1600*. The book (4 volumes) contains 16,112 watermark designs reproduced in real size. Other researchers carried out new surveys a few decades later. The most prolific of them for the German space, Gerhard Piccard, produced *Die Wasserzeichenkartei Piccard im Hauptstaatsarchiv Stuttgart* in 25 volumes published between 1961 and 1997: it contains 92,000 reproductions, almost six times more than Briquet catalog.

More recently, online repositories were built to browse these catalogs. A project called Briquet Online¹ led by Ezio Ornato within the LAMOP is hosted on the website of the Austrian Academy of Sciences and contains just over 16,000 watermarks collected by Briquet. Another project² launched in 2010 takes over and amplifies the Piccard catalog and corresponds to a Deutsche Forschungsgemeinschaft company: it gathers 103,000 watermark images preserved in Germany with a more intuitive and visual search according to a progressive prioritization of requests (for example hand-hand with hand-fingers with three fingers, etc.). The Bernstein

¹http://www.ksbm.oew.ac.at/_scripts/php/BR.php

²<https://www.wasserzeichen-online.de/wzis/index.php>

project³ brings them together with another Austrian edition (Wasserzeichen des Mittelalters) in one searchable database on the website of the Austrian Academy of Sciences. Much of the ingenuity developed by these online databases lies in the search for the best possible semantic description of watermarks to offer as quickly as possible the answer to the questions posed by the discovery of one design or another. The drafting of thesauri and description standards followed different paths. In 2002, the International Association of Paper Historians (IHP) indexed the descriptive terms of watermarks in six international Paper languages. In 2012, the Austrian project Bernstein published on its website a glossary of *Watermark-Terms: Vocabulary for Watermark Description*, in seven languages. The structure is highly hierarchical to gradually lead from general ratings to very precise ratings.

However, identifying a watermark remains very hard because on one side the catalogs do not cover all the existing watermarks and on the other side, they contain hundreds of thousands of designs and are already too large for non-experts to search. Indeed, identifying the semantic content of a watermark, which can be used to search the catalogs, is both difficult and subjective, requiring expert knowledge of each catalog, collection and period. Leveraging simple cell-phone photographs to search existing watermark design catalogs or build and search crowd sourced photograph catalogs could thus be a real game changer for archivists and historians. Previous work has tackled this problem in the context of retrieving similar watermark designs given an input watermark image. Rauber et al. [1997] for instance have used histogram-based methods to compare the similarity between a query watermark image to their database of binarized digitized watermarks. But their approach relies on the ability to separate watermarks from other lines in the paper, which is itself a complex task, and thus requires many pre-processing steps. Others like Picard et al. [2016] have used Machine Learning based methods such as dictionary learning, where an image is seen as a collection of regions where each region is encoded on a dictionary; similar regions are then encoded by similar atoms of the dictionary. These approaches differ from ours as their objective is not to find the exact matching design of a given watermark photograph, but to find a similar watermark design given an input design. A more similar approach to the one of Shen et al. [2019], which we based our work on, Pondenkandath et al. [2018] used convolutional neural networks to classify a non-public dataset of more than 100k watermarks into 12 categories. But their approach differs from our work as our objective is not to classify into a fixed number of coarse categories. Instead, we are trying to find the exact class for each watermark, making each one of the 16,112 watermarks in the Briquet database a single class.

While we believe the association of the photograph of a watermark on a paper to its corresponding design from the Briquet catalog is of great help to anyone looking to identify a watermark, other prior information the user might have about the original paper would be complementary to fully identify and date it, and future work could consider how to properly integrate it to the proposed matching algorithm.

III DATA COLLECTION

Developing the automatic recognition system detailed in Shen et al. [2019] requires specific watermark photographs, taken in a standardized way. In this section, we detail how we collected these images.

³http://www.memoryofpaper.eu/BernsteinPortal/appl_start_disp

3.1 Primary source and database

We first collected watermark photographs in the Minutier central des notaires de Paris, the department of the French National Archives where all the archives of the notaries practising in Paris (and in the old municipalities depending from the capital in 1860) from the end of the 15th century to the beginning of the 20th century are preserved. Notaries produced familiar and economic deeds or any other private deed with authentic value, so that the Minutier central des notaires de Paris holds today around 22 million deeds coming from 122 notary's offices. Notaries drew up tens of documents every month and conserved the minutes on paper in their archives, bound in registers or collected in sets of loose sheets. The abundance of daily and weekly deeds drawn up by the same office implies a significant consumption of paper, which is perfectly suited to the efficient collection of many samples of the same watermark required to train the recognition algorithm.

A database was thus created from original photos from the documents of the Minutier central des notaires de Paris, where the voluminous registers and the numerous loose sheets are made by consistent series of the same paper (that is the same watermark), and where the richness and variety of watermarks made it possible to collect easily and quickly significant series of the same pattern. As it is in the common sense for paper historians, we decided to define a class of watermark by all the watermarks produced using the same watermark mould, including the couple of inevitable twins. Indeed, in the regular process of hand-making paper, two paper moulds (where the twin watermarks were sewn on) were simultaneously and alternately used at the same vat by a couple of workers. In this way, one ream of paper contains the twin watermarks. A watermark must thus be seen as a couple of watermarks showing the same pattern, used at the same time to produce the same stock of paper. The data collection systematically concerned the minutes of the years 1500, 1525, 1550, 1575, 1600, 1625, 1650, 1675 and 1700, which resulted in 361 different watermarks each photographed multiple times.

To ensure the most efficient training for the algorithm, the photos have been taken following a specific protocol: the watermarks should fit inside a 2:3 rectangle and appear in a wide variety of conditions, for example on written or blank sheets of paper, with a lot of luminosity or in the darkness.

3.2 Enrichment with existing resources

Since we wanted to address watermark photographs recognition from existing design collections, we also required paired data of photographs and designs to train and evaluate the recognition algorithm. We focused on the well-known repertory *Les filigranes* (1907), by the Swiss scholar Charles-Moïse Briquet (1839-1918). Apart from its authority and accuracy in the dating and location of paper and watermarks, Briquet's database was an interesting extension because it reproduces many different patterns, some of them already found at the Minutier central des notaires de Paris in a later version. We thus created another training and evaluation database for automatic recognition using the watermarks listed by Briquet. On one hand, all the 16,112 PDF versions of the original published tracings were retrieved from the Briquet Online website (built by the LAMOP in Paris and currently hold in Vienna by the OAW) and we semi-automatically removed information irrelevant for the visual classification algorithm. On the other hand, we tracked down 298 real watermarks traced by Briquet in Paris, which are still visible in the same documents he consulted, presently preserved at the DMAAR, the department of the National Archives for the Medieval Age and the Ancien Régime.

IV WEB APPLICATION

Our main goal is to introduce a mobile application for watermark recognition. Indeed, we believe the wide availability of smartphones can allow the general public, which has little knowledge about watermarks, to get relevant information right away when sorting through documents. The mobile application was also seen as the proper vector for crowd-sourcing knowledge about documents at diverse institutions (Terras [2016]). We adopted a set of recent mobile Web standards, also known as *Progressive Web Apps*, for developing and maintaining our native mobile application.

In this section, we first introduce the concept of Progressive Web Apps. We then detail the technical design of the application and present its search, catalog and contribution workflows. The full code for the Web application is publicly available⁴, and scripts as well as Dockerfiles are provided to help and encourage future development.

4.1 Progressive Web Apps

Building mobile applications has been so far dominated by native development (Android’s Java, iOS’s Objective-C/Swift) or by hybrid development frameworks (Ionic, React Native) that generate native applications from a common source. Building native applications enables to leverage the phone’s full capabilities, either online or offline, and benefits from App stores’ distribution channels. However, despite the domination of the mobile ecosystem by two platforms, Android (75%) and iOS (20%), mobile app development is still hindered by device fragmentation due to the various underlying OS versions, the different form-factors and the vendor-specific modifications (for Android).

As an alternative to native and hybrid applications, Progressive Web Apps (PWA) are application bundles delivered through the Web that leverage common Web technologies (e.g., HTML, CSS, Javascript) and advanced Web standards supported in modern Web browsers (e.g., service workers, manifests, local storage, notifications, geo-localisation and sensors). As the support for these advanced standards has dramatically increased in both desktop and mobile browsers (Chrome, Firefox, Safari, Opera, . . .), developing native-like mobile applications from a single code base and deploying through a unique compilation and delivery process is now possible.

Major services (e.g., Uber, Tinder, Trivago) now leverage PWA for their mobile application as it, comparatively to native applications, successfully addresses heterogeneity (through well-supported standards), produces smaller bundles (which improves download/update times) and enables to start a relationship with Web visitors (which addresses the issue of overcrowded App stores).

4.2 Architecture

To build the application, we leveraged the Angular framework for the Web frontend, NodeJS for the Web backend and the Python Flask framework for the Search backend application, whose goal is to perform the actual matching of a given source watermark image to all the watermarks in the database. Both NodeJS and Flask backend applications expose a REST API to be used by respectively the Angular and NodeJS applications. These technological choices were first and foremost guided by their wide usage and support history in order to ease the maintenance. The full architecture is illustrated in Figure 1.

The Web and the Search backend applications access a Mongo database currently storing the

⁴<https://gitlab.inria.fr/dhai/filigrane>




Figure 1: Web application architecture.

watermark images and associated metadata for both the Minutier and the Briquet repositories. The database may be extended either with individual contributions from registered users (i.e., crowdsourcing), or through the direct integration of additional repositories.

For the Web frontend application, we leveraged a range of Web APIs⁵. In addition to the Service Worker API that enables the execution of the code in the Web page as a stand-alone application on the smartphone, we also leverage the APIs for local storage and database, media access, push notifications and background tasks. The compiled application bundles are downloaded by the client's Web browser when visiting the project's website from a smartphone, and a request to install it as a mobile app on the phone's home screen is displayed. Once installed, the application can be launched directly from the home screen, even offline, and in a fullscreen mode (hiding the fact that the application is executed inside the phone's Web browser).

A definitive advantage of PWA is therefore the seamless installation while visiting the project's website on a smartphone. Moreover, it benefits from the efficient application updates which notify the user directly from the application and only update parts of the application that have actually been updated.

4.3 Workflows

The first three main purposes of our application are 1) to enable non-specialists to identify a watermark using a simple photograph by searching for a potential matching watermark in our large-scale database (*search workflow*), 2) to visualize the database of watermark images and their corresponding details (*catalog workflow*) and 3) to enable the user to enrich the existing database (*contribution workflow*).

4.3.1 Search

The search workflow is illustrated in Figure 2. First, a document of watermark is scanned (2(a)) and square resized for better matching capacity (2(b)). Then, the matching algorithm is launched and once finished, the first 10 matching watermarks, both from the photographs and designs datasets, are transferred to the Web backend which automatically displays the results (2(c)). Finally, matching entries can be selected to display their details (2(d)) and explored further through image zooming (2(e)) or external links (2(f)).

More specifically, as a mobile user takes a snapshot of a document with a watermark (or selects an existing snapshot), the photograph is sent to the NodeJS Web backend application, which in turns makes the relevant calls to the search backend. Indeed, the search backend exposes multiple API calls, each accessing a search algorithm dedicated to a specific dataset: the photograph

⁵<https://developer.mozilla.org/en-US/docs/Web/API>

dataset (currently corresponding to the Minutier central des notaires de Paris dataset for now) and the design dataset (currently corresponding to Briquet for now). Thus, the Web backend calls matching algorithms against the different datasets in the database. Once the results for all the matching calls are returned, the Web backend returns them to the frontend application, each set of results is displayed separately and the user may explore and visualize them independently.

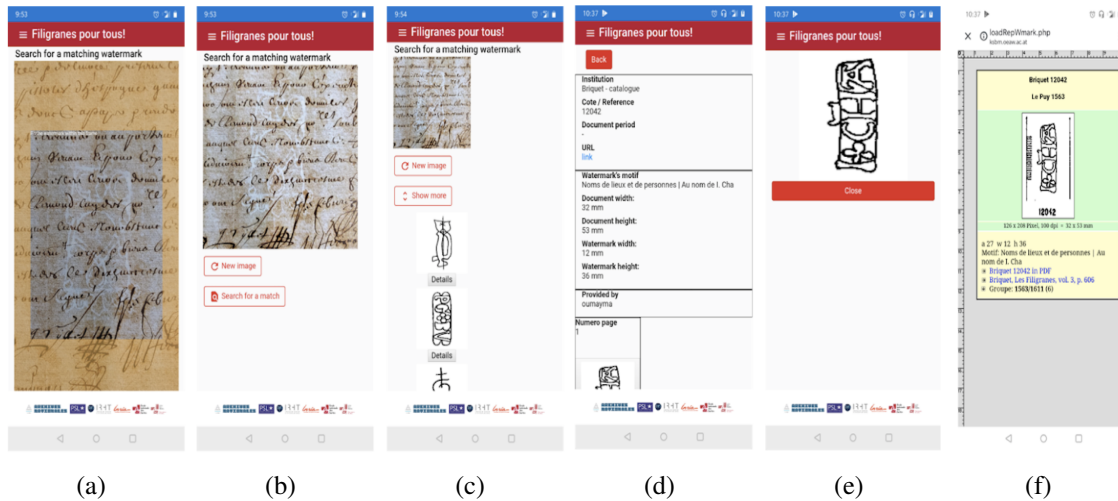


Figure 2: Example of the search workflow (see text for details).

4.3.2 Catalog

The catalog workflow enables the user to visualize our database as illustrated in Figure 3. A user may navigate the full database of registered watermarks (3(a) and 3(b)) and then explore the details of individual watermarks (3(c) for a design from Briquet and 3(d) for a photograph from Minutier).

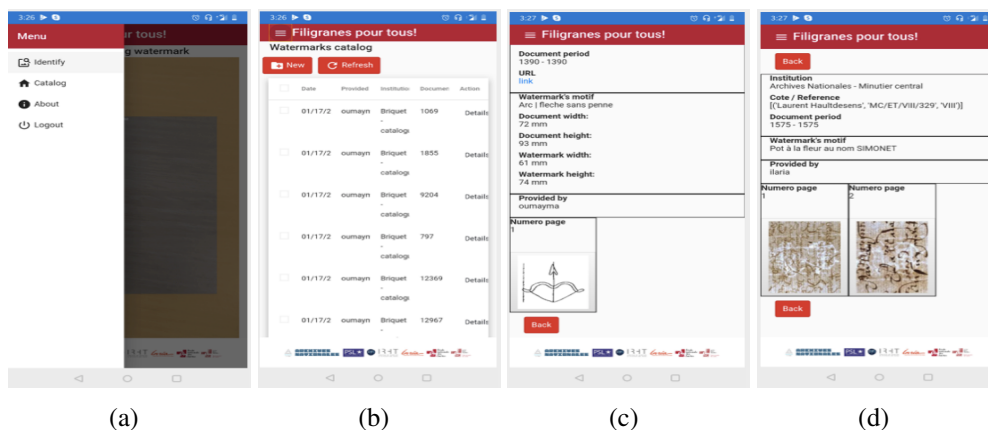


Figure 3: Example of the catalog workflow (see text for details).

4.3.3 Contribution

The contribution workflow in Figure 4 illustrates how a user can scan (4(a)) and annotate (4(b) and 4(c)) a watermark before it is added to the watermarks database on the Web backend side. Crowdsourcing is foreseen to extend the sets of watermarks that the matching algorithms search

against (i.e., to identify more watermarks), and also to gather descriptive metadata related to the document of the watermarks being added. Leveraging the storage and background tasks Web APIs, these steps may be undertaken offline (e.g., when inside a building with no Internet access), waiting for later connectivity to upload the data. Once uploaded, the contributed watermark is not yet added to the database used for the matching algorithm. Indeed, the contributed data needs to be curated and validated. How to properly integrate institutional and crowd-contributed content remains an open challenge. At this stage, the contribution feature of the application is therefore experimental and limited to internal contributors.




Figure 4: Example of the contribution workflow (see text for details).

V WATERMARK RECOGNITION ALGORITHM

As mentioned above, we follow the approach described in Shen et al. [2019]. Here, we specify the problem, summarize the main elements of the recognition algorithms and give technical details about our implementation in the web application.

5.1 Task and challenges

Given a photograph of a watermark on an archive document, our goal is identify the watermark in order to find information about the document such as its origin or the time when it was written. In order to do so, we would like to compare this photograph to the largest possible database to find the corresponding watermark. As discussed above, we chose to use the database collected by Briquet which is available online and contains 16,112 engraved watermarks along with the one collected from the Minutier central des notaires de Paris which contains 2,250 photographs of almost 400 watermarks. The recognition task is challenging for several reasons. First, some photographs are very cluttered (Figure 5). This could be because the state of conservation of the archive document is bad, or simply because a lot of writing has overshadowed the watermark in the paper. Second, even if the watermark was perfectly visible in the document, it would still be hard to identify in some cases because it could be very similar to dozens of other watermarks, as can be seen in Figure 6. Additionally, the matching we are doing is large-scale since given one photograph we are looking for the corresponding design among more than 16,000 candidate matches. Finally, from a computer vision perspective, it is a difficult task to solve because we want to match images from different modalities, where one is a photograph and the other is a design.




Figure 5: Examples of cluttered archive documents photographs where the watermark is barely visible.




Figure 6: Example of different but very similar watermarks.

5.2 Approach

5.2.1 Convolutional Neural Networks

To perform watermark recognition, the key is to obtain a good representation of the images of watermarks, i.e. a representation that is discriminative enough for the computer to distinguish the different types of watermarks and at the same time is not sensitive to other factors such as the presence of writings or the illumination conditions. Such a representation is called a feature or descriptor. One solution is to try to design a feature that will suit these requirements, typically by extracting edges from the image. However, this proves to be extremely difficult. Another approach is to learn a good feature from data. This can be done using several approaches, and in particular using deep neural networks, which recently showed impressive performance for many visual tasks, and can be seen as a class of computational models with multiple levels of abstraction (see LeCun et al. [2015]). They are defined by a set of parameters that can be automatically learned to perform a specific task on a set of examples called training data. Convolutional Neural Networks (CNNs) are a class of neural networks that use a specific type of operation, convolutions, and have proved very successful for image analysis. We thus use CNNs to extract relevant features from the images. The next section presents how the parameters for the CNNs we used were obtained.

5.2.2 Network Training

We used two different CNNs for recognition, the first one designed to match a query photograph to a photograph database, the second one to match a query photograph to a design database.

The first CNN was a standard ResNet-18 network (He et al. [2015]) whose parameters were

optimized to classify well a subset of our watermarks. More precisely, we selected 60 watermarks with 100 different photographs for each and optimized the network parameters so that the network could classify each photograph into one of the 60 classes. Our model reached 98% accuracy on new watermark photographs from these classes. An intermediate representation learned by this deep network can directly be used to recognize new watermarks, as detailed in the next section.

To match photographs to designs, we additionally fine-tuned our model on the joint photograph and design database. As the matching is done locally, the fine-tuning was done on intermediary image representations rather than on the whole image level. In this fine-tuning, we optimized similarity between features of corresponding patches using a standard metric learning loss, the triplet loss introduced by Chechik et al. [2010]. Details of the training procedure for this second step can be found in Shen et al. [2019].

5.2.3 Matching procedure

Given a pair of images, for example a photograph and a design, as visualized in Figure 7, we can compute their features using the network trained according to the procedure described in the previous section. We then need to compare the two features to decide if they correspond to the same watermark. We again followed the approach developed by Shen et al. [2019] and compare the content of the two images in two steps. The first step is to compute a global similarity score between the global features of each of the images. The second step aims at refining the matching by emphasizing spatial information and uses a local similarity score that can roughly be interpreted as the number of small patches that are similar in the two images.



Figure 7: Cross-domain matching. Given a source photograph of a watermark, we want to find its corresponding design in a database of 16,112 designs.

In the example of the Briquet catalog, we would like to find the matching design of the photograph of a watermark among 16,112 designs. We first compute the global feature similarity between the query photograph and the 16,112 designs (a few seconds). We then compute the more computationally expensive local similarity score only on the most similar designs. This allows to refine the ranking of the most similar designs in a reasonable time. Indeed, computing this local similarity on the full database takes 27 minutes for each query, while the two-steps approach leads to similar results in a matter of seconds. Following Shen et al. [2019], this matching is performed at multiple scales and takes rotations and flips into account as illustrated in Figure 8.

5.3 Results

5.3.1 Results and execution time

Here, we present results obtained when matching a photograph of a watermark to designs from the Briquet database. This is much harder than matching a watermark photograph to another, as is done in the web application on data from the Minutier central des notaires de Paris. We expect much higher performances for this photograph-to-photograph retrieval task, but it was difficult to evaluate performances in a realistic setup. Our experiments indicate around 90% accuracy



Figure 8: Matching is performed at five scales of the original photograph, each including 90° rotation and horizontal and vertical flipping.

with the photographs we obtained from the Minutier Central, but these only include a limited amount of variations and are probably over-optimistic. However, obtaining images from very different sources is very challenging. We thus focus on the more difficult photograph-to-design experiments.

The original matching method described in Shen et al. [2019] achieves state-of-the-art results in cross-domain watermark recognition. Computing the described local similarity using different scales on the full 16,112 designs leads to an accurate recognition in 55% of the cases, meaning that in 55% of the cases, the exact design amongst the full 16,112 possibilities is recovered, which we think is an impressive performance considering the watermark is barely visible in many of the photographs and that many designs are extremely similar. The counterpart of this result is that it takes 27 minutes to match a single photograph. Since we want to integrate the matching algorithm to the web application described in Section IV, we can not expect users to be willing to wait this long for one result. We thus used the approach described in Section 5.2.3 and pre-filtered the targets using the faster but less accurate global similarity. We also experimented matching the watermarks at a single scale instead of using five scales. The results presented in Table 1 show that using pre-filtering degrades slightly the results ($\sim 5\%$) but leads to much faster matching with more than 100x speedup. Performing local matching at a single scale further reduced computation time ($\sim 5x$) but degraded the results by another 5%. We thus decided to keep the multiscale matching for our application, which leads to a processing time of 13 seconds per photograph and 49% top-1 accuracy.

In practice, performance could easily be boosted since the user could take a second photograph if the result with the first shoot was not successful. We give hints of the boost that could be obtained with such an approach in Table 2). While the first line reports results from a single shot in the setup we selected for the application (multiscale matching and 100 pre-filtering), the second line reports the best result using two photographs. In 70% of the cases, the top retrieved design was the good one for one of the two shots, and in almost 80% of the cases, the good design was amongst the top 10 retrieved for one of the two photographs. We believe these results make our application highly practical.

		Top-1 acc	Top-5 acc	Top-10 acc	Runtime
One scale	100 pre-filtering	44	56	58	3s
	no pre-filtering	46	61	63	432s
Multi-scale	100 pre-filtering	49	60	65	13s
	no pre-filtering	55	68	72	1584s

Table 1: Comparison of the impact of pre-filtering on 100 candidate designs, and of multi-scale matching on top-K accuracies in % and execution times.

	Top-1 accuracy	Top-5 accuracy	Top-10 accuracy	Top-100 accuracy
First shot	49	60	65	69
Second shot	70	77	79	81

Table 2: Top-K accuracy in %. First shot corresponds to a case where a user only takes one photograph and looks for a match. Second shot is when a user takes a second photographs if the first one fails.

5.3.2 Qualitative Results

Figures 9 and 10 show cases of challenging queries for the matching algorithm where there are multiple similar designs in the search database. Figure 9 highlights challenging cases where the algorithm was still successful, because enough details were present in the watermark and could be identified in the photograph. On the contrary, Figure 10 shows failure cases, but also highlights the extreme difficulty of our data which includes a wide variety of very similar designs where identifying the correct watermark in the photograph is very challenging.



Figure 9: Top 5 results returned by our algorithm. Challenging examples where the algorithm obtains the correct match as first results.



Figure 10: Top 5 results returned by our algorithm. Challenging examples where the algorithm fails to get the correct match as first result.

VI CONCLUSION

We introduced a new public and user-friendly web application aimed at making the recognition and study of watermarks easier and more efficient. First, the application hosts a large-scale annotated watermark database. Second, it incorporates an adaptation of a state-of-the-art automatic watermark recognition algorithm, enabling users to identify a watermark from a simple photograph in a few seconds. Third, although the feature is not public yet, the database can directly be enriched through user contributions.

In the future, we plan to improve the matching algorithm to get even better recognition performance, and to extend our database through existing resources or individual contributions, which we ambition to collect by making our web application a go-to platform for watermark crowdsourcing. Indeed, we believe an automatic recognition tool such as the one we introduce could be key to integrate data from multiple sources in a single database without requiring expert knowledge to annotate every sample. Even though this is challenging due to questions related to data property, it would be especially valuable as it would help bring together the available collections of watermark holders.

Acknowledgements

This work has received support under the program "Investissements d'Avenir" launched by the French Government and implemented by ANR with the references ANR-10-IDEX-0001-02 PSL. It was also supported by the ANR project EnHerit (ANR-17-CE23-00080), project RAPID Tabasco and gifts from Adobe to Ecole des Ponts.

References

- Gal Chechik, Varun Sharma, Uri Shalit, and Samy Bengio. Large scale online learning of image similarity through ranking. *J. Mach. Learn. Res.*, 11:1109–1135, March 2010. ISSN 1532-4435.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. 2015.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. doi: 10.1038/nature14539. URL <https://doi.org/10.1038/nature14539>.
- D. Picard, T. Henn, and G. Dietz. Non-negative dictionary learning for paper watermark similarity. In *2016 50th Asilomar Conference on Signals, Systems and Computers*, pages 130–133, 2016.

- Vinaychandran Pondenkandath, Michele Alberti, Nicole Eichenberger, Rolf Ingold, and Marcus Liwicki. Identifying cross-depicted historical motifs, 2018.
- Christian Rauber, Thierry Pun, and Peter F. Tschudin. Retrieval of images from a library of watermarks for ancient paper identification. 1997.
- Xi Shen, Iliaria Pastrolin, Oumayma Bounou, Spyros Gidaris, Marc Smith, Olivier Poncet, and Mathieu Aubry. Large-scale historical watermark recognition: dataset and a new consistency-based approach. 2019.
- Melissa Terras. *Crowdsourcing in the Digital Humanities*, page 420 – 439. 01 2016. doi: 10.1002/9781118680605.ch29.