

# SIMULTANEOUS FORGERY IDENTIFICATION AND LOCALIZATION IN PAINTINGS USING ADVANCED CORRELATION FILTERS

Paul Buchana\*, Irina Cazan\*, Manuel Diaz-Granados\*, Felix Juefei-Xu and Marios Savvides

ECE Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA

## ABSTRACT

With the availability of high resolution digital technology, there has been increased interest in developing statistical and image processing techniques that can enhance the existing capabilities of analyzing works of art for authenticity. This work explores the merits of using advanced correlation filters in supplementing art experts efforts in identifying forgeries among disputed paintings. We show that by training the optimal trade-off synthetic discriminant function (OTSDF) filter on each section of a coarsely parceled image of an original painting, we are not only able to distinguish between a low-quality digitized representation of a painting and its forgery, but also specifically indicate where the differences occur and where the replica is particularly faithful to the original. This method is also valuable in determining whether an original painting has undergone any modifications, given that a representation of the initial version is available.

**Index Terms**— Forgery Detection, Forgery Localization, Advanced Correlation Filters

## 1. INTRODUCTION

Whether motivated by material gain or the desire to express admiration for another artist, art forgery has been a lucrative and active business for centuries. Traditionally, forgery detection has been based on the discerning abilities of “connoisseurs”, relying on their ability to deduce authenticity from an artist’s known work, life, and influences. However, over time, these traditional methods have greatly been enhanced by quantitative methods, from X-ray analysis to isotope content and most recently, mathematical tools of describing an artists style applied to high-resolution digitized versions of the paintings. A lot of progress has been made in using digital techniques of feature extraction to describe an artists style. Based on these results, studies have managed to classify test paintings as originals or forgeries. However, some of these methods have shown weaknesses when faced with varying quality of the digital scans or photographs of the paintings.

In this work, we describe the application of advanced correlation filters to (1) detecting forgeries in paintings under varying qualities of the paintings’ digital reproduction and (2)



**Fig. 1:** A: Originals. B: Copies. C: Image processing steps: a regular printout scan is shown. Notice on the final image a slight padding on the right side (highlighted).

localizing where alterations have been made to the original. This project differs from the literature in two ways. First, it explores the merits of using advanced correlation filters in determining the authenticity of disputed paintings, an approach not attempted previously in this problem space. Second, it does not attempt to use an artists style signature for authentication; given a known (original) work of art that was lost or stolen, we want to determine whether the version that resurfaced is indeed the original or an imitation, not to attribute never-before seen paintings to one artist or another.

**Previous Work:** Computational image analysis and machine learning techniques have been shown to be a promising way of assisting art experts in the authentication of unknown or disputed paintings. One of the most widely used techniques to date have involved multi-resolution analysis to extract salient features from digitized version of the paintings investigated in searching to analyze different artists’ styles. Among the feature extraction tools utilized there are wavelets ([1], [2], [3]), curvelets ([4]), craqueleure, and contourlet transforms ([5]) with the purpose of detecting subtle differences in brushstrokes and painting degradation that may point to whether the painting is an original or a forgery.

However, these methods are used mostly for classification, and to support art experts in the evaluation of the authenticity of a painting. On the other hand, correlation filters have been widely used for automatic target recognition [6], object alignment [7], biometrics recognition [8], and many other applications in which true identification is crucial, and a low margin or no false positives are allowed. To our knowl-

\*These authors contributed equally to this work.

edge, correlation filters have not been used previously for the purpose of judging the authenticity of paintings.

## 2. PROPOSED METHOD

Our proposed method of addressing the problem of forgery detection and alteration localization is centered around advanced correlations filters. In particular, we train the OTSDF [9] filter on localized regions of the original painting and test for global forgery or local alterations with a range of different digital representations of the same painting.

We transform each painting in our dataset from RGB to gray-scale and down-sample by a factor of 10 (10x). We then extract HOG features because they capture the artist’s brushstroke direction at a granular level, a characteristic of the painting that is very hard, if not impossible to imitate exactly. In order to test locally, we segment each image into a predefined number of patches. The smaller the window of the patch, the more accurate the detection of differences from the original painting. These patches are first computed for the training image (original), generating multiple OTSDF filters; the same process is carried out on the testing images.

To atone for the distortions that the testing images may have experienced while the painting was scanned or photographed, we generate multiple shifts in all directions for each patch of the testing image. In the testing process, each patch of the testing image is correlated with its corresponding patch-OTSDF filter. We evaluate the performance of the filter in recognizing original images by calculating the peak-to-sidelobe (PSR) ratio. When the test image is an original (or a low-resolution rendering of the original), the correlation output should exhibit a sharp peak, which should not be observed when the filter is applied to forgeries. Once the peak is located in the correlation output, the mean and standard deviation are calculated for a square sidelobe region surrounding it (except for a small mask region right around the peak).

The reasoning behind using the OTSDF filter is that we require a noise-tolerant filter, so that external factors, such as the collection medium (scanner, digital camera), do not cause false negatives. The OTSDF filters achieve the desired trade-off between average correlation energy and output noise variance through the following closed-form solution:

$$H = P^{-1}X(X^*P^{-1}X)^{-1}c, P = \alpha C + \beta D, \alpha, \beta \in [0, 1]$$

$X$  is the training image data matrix, where the size is  $d \times N$  ( $N$  is the number of training images).  $D$  is  $\frac{1}{N} \sum D_i$ , where  $D_i$  is a diagonal matrix of size  $d \times d$  whose diagonal elements are the magnitude square of the associated element in  $X_i$ . Here,  $C$  is a diagonal matrix containing the elements of the input noise power spectral density along its diagonal.  $\alpha^2 + \beta^2 = 1$ , represents the trade off between sharp peaks and tolerance to noise. In this application, we are using white noise, which assumes  $C = I$ , the identity matrix.

## 3. EXPERIMENTS

The dataset [10] contains seven original paintings and their known copies (forgeries), which have been digitized under uniform acquisition conditions. Both originals and copies serve as ground truth since their authenticity is completely known. The series of paintings were produced by Charlotte Caspers, and initially used in [2]. This dataset serves as an ideal start for designing, training, and testing a correlation filter since we can compare results with reality. To supplement this dataset and further test the proposed method, we have generated additional digital representations of these paintings using both high and low quality collection methods. First, a green border was added to each image so that the actual image borders are easily detected after printing and scanning. The high resolution images were printed on glossy photograph paper and on regular printer paper. Both photographs and printouts were then scanned at decreasing resolutions (600dpi-100dpi). Data samples are included in Fig. 1.

### Experiment I: Benchmarking

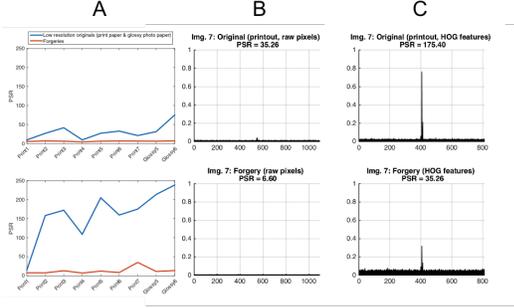
For the first experiment, we follow a similar protocol as used in [2]. Each painting is partitioned into  $1024 \times 1024$  patches and the features used are either the raw pixels of the gray-scale image or HOG features, which capture local brushstroke direction. The partitioning allows for an artificial augmentation of the dataset, as each patch can be considered as a stand-alone painting, since it contains a slightly different composition. For each patch of each painting, we train the OTSDF filter with the original patch and test it using both the original and copy corresponding patch. The two tests in Table 1 contain the results of using raw pixels as features (Test1) and HOG features (Test2). The results are remarkably good given access to very high resolution testing images, but that may not be the case in practice. In the following experiments, we show the performance of the proposed method under real-world noise and degradation conditions.

### Experiment II: Global and Patch Detection

*Global-OTSDF Filter:* For each image pair in the dataset, the OTSDF filter is trained using the original image. Testing is performed using lower resolution originals, the known copy, and a few distorted versions of the original. In each case, the OTSDF filter is trained with either raw pixels or HOG features. The low resolution original images captured using different methods are recognized with varying degrees of confidence when raw pixels are used as features, and the forgeries are rejected in each case. Significantly better performance is observed when HOG features are used, as captured by the PSR measure (Fig. 2(A)). Using HOG features is a good alternative for capturing authenticity of an image, since they are able to capture the directions of the painter’s brushstrokes and micro-level features that are close to impossible to imitate. However, filtering using HOG features is prone to sometimes returning false positives, as can be seen in Fig. 2(B, C). To avoid this situation and enhance the recog-

**Table 1:** Experiment 1 Results.

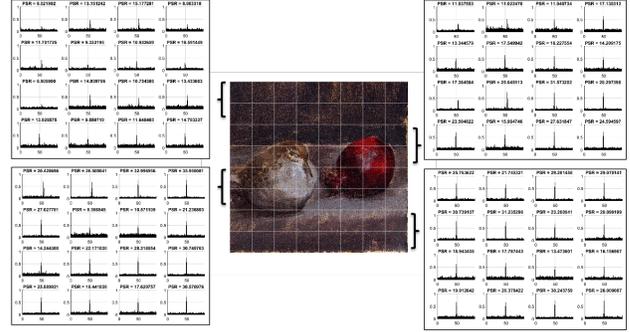
Pair	Ground	Paint	Brushes	Style	Test1			Test2		
					Total	Copy	Original	Total	Copy	Original
1	Smooth CP Board	Oils	S&H		96%	94%	100%	100%	100%	100%
2	CP Canvas	Oils	S&H		98%	96%	100%	100%	100%	100%
3	CP Canvas	Acrylics	S&H		99%	98%	100%	99%	98%	100%
4	Bare Linen Canvas	Oils	S	TI	94%	88%	100%	100%	100%	100%
5	Chalk and Glue	Oils	S	TI	95%	90%	100%	97.5%	95%	100%
6	CP Canvas	Acrylics	S	TI	98%	96%	100%	100%	100%	100%
7	Smooth CP Board	Oils	S	Sm,Bl	95%	90%	100%	99%	98%	100%



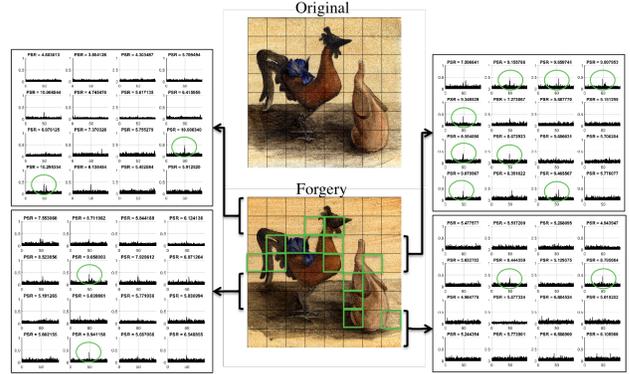
**Fig. 2:** A: PSR values for a subset of tested images using raw pixels (top) and HOG features (bottom). Low PSR values signal that the test image is not an original. B,C: A forgery is incorrectly recognized as original using the OTSDF filter with HOG features.

nition performance, we experimented with the patch-OTSDF filter, the proposed method.

**Patch-OTSDF Filter:** The patch-OTSDF filter is trained on 64 patches of the original image and for each patch, tested with 169 shifts. From the  $64 \times 169$  correlation outputs, the highest PSR value outputs are selected per patch. Each of the 64 outputs represents a specific location in the image, as can be seen in Fig. 3. This method of patching and shifting is effective, since we can pinpoint exactly which part of the testing image has discrepancies from the original, something that the global OTSDF cannot accomplish. With the patch-OTSDF filter method we can determine if any changes have been made to the original painting. For example, if the painting is lost and then recovered, using the patching-OTSDF analysis we can find (1) if the painting is the original, and (2) if any changes were made and specifically where, which may be unnoticeable to the human eye. This is exemplified in Fig. 5, where the image appears identical to the original, but has minor, imperceptible alterations on the top-right corner; our method is able to recognize which part of the original painting was altered. Patches corresponding to the top right corner have very low PSR values, coinciding with the regions where we altered the painting. Fig. 6 shows another mildly modified image (added red brushstrokes to the green object on the right); unfortunately, the 64 patch-OTSDF filter is unable to identify the differences, primarily because the individual patches are too large and the filter matches the alteration background. To address this issue, the patch size is decreased, using a total of 2500, allowing the identification of



**Fig. 3:** The correlation outputs corresponding to the highest PSR values for each patch of the test image (a low-resolution reproduction of the original). The PSR peaks for each of the 64 patches indicate a match.



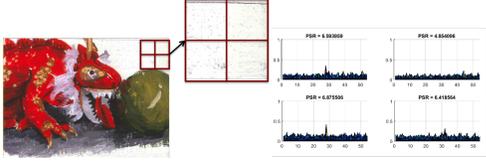
**Fig. 4:** Similar (highlighted) and different parts of a copy are identified.

these minute changes.

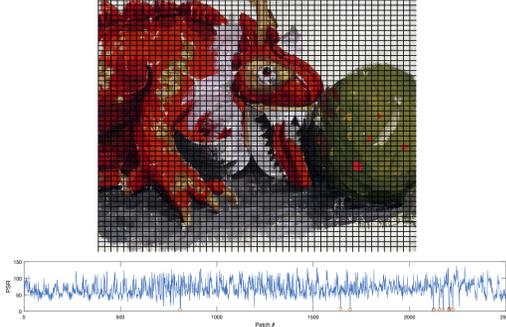
### Experiment III: Robustness

The proposed method is tested for robustness when either the training image, the testing image, or both have been corrupted in some way due to digital processing. Specifically, we are interested in assessing the accuracy of our method in detecting whether the testing image is an original or a forgery under increasing levels of noise and missing pixels.

First, we assess the detection accuracy when training with a high fidelity digital representation of the original painting, but testing with an image (original) to which noise was globally manually added. The results may be seen in Fig. 7, where a white Gaussian noise (WGN) level of 25 dB still leads to very good detection accuracy; however, we observe a sharp



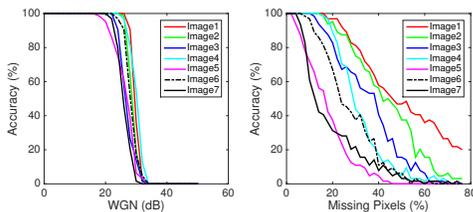
**Fig. 5:** The distorted section of the original image is correctly identified as a non-match.



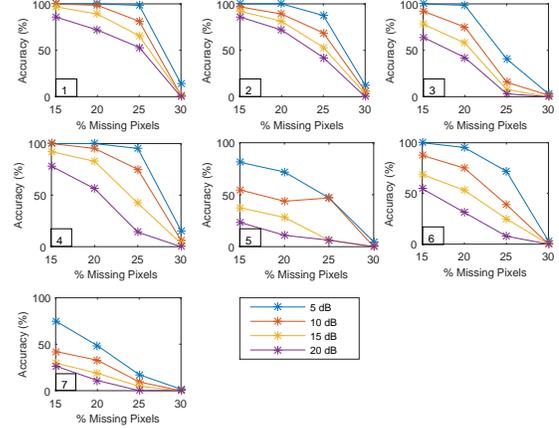
**Fig. 6:** Using 2500 patches, the 12 areas of the original image containing distortions are correctly identified.

accuracy drop-off between 25 and 30 dB of added WGN.

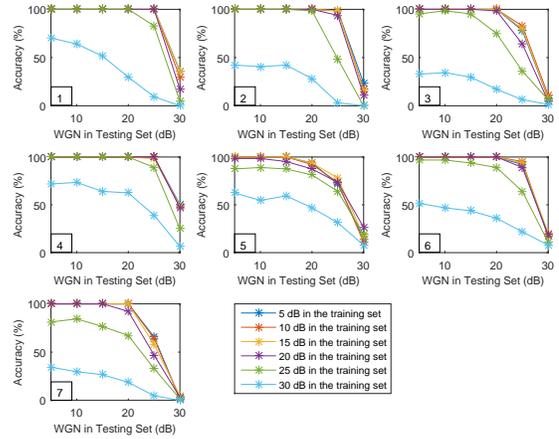
Second, we test the robustness of the proposed method to testing with images (originals) that are missing pixels uniformly at random. Results are included in Fig. 7. The experiment indicates that the accuracy of our method is highly dependent on the specific painting examined, but maintains a relatively good (up to 50%) accuracy when  $<12\%$  of the pixels are missing. Thus, even when a realistic random percentage of the image pixels have been compromised due to malfunctions or defects in the capture medium (photography, scanning, etc.), the proposed method is sufficiently robust in detecting originals from forgeries. Next, we have merged the distortions from the first and second experiments, where the effect of adding more than 25 dB of WGN creating a sharp decrease of accuracy, and randomly removing pixels (which is prone to affect the detection differently from one image to another) is combined in one graph Fig. 8. For five out of seven images, we get more than 50% accuracy by adding 15 dB WGN and removing 20% of the pixels in the testing images. This indicates an overall robustness to these two combined distortions (more than 50% of patches recognized means the



**Fig. 7:** Robustness to noise addition (left) and random pixel removal (right).



**Fig. 8:** Robustness to both noise and random pixel removal.



**Fig. 9:** Robustness to noise in both training and test sets.

detection of an original painting). We have also experimented with varying levels of noise in both the training and testing sets. The results are included in Fig. 9.

#### 4. CONCLUSION

In this work, we have demonstrated the merits of using correlation filters to detect painting forgeries when the a digital representation of the original is available or can be obtained. We have shown that by training the patch-OTSDF filter on each section of a gridded image of an original painting, we are not only able to distinguish between a low-quality digitized representation of a painting and its forgery, but also indicate with required precision where the differences occur and where the replica is particularly faithful to the original. This method is also valuable in determining whether an original painting has undergone any modifications, given that a representation of the initial version is available. The method has good generalization power via its tunable granularity and resilience to noise in the case of low-quality digital representations of the original.

## 5. REFERENCES

- [1] C Richard Johnson Jr, Ella Hendriks, Igor J Bereznoy, Eugene Brevdo, Shannon M Hughes, Ingrid Daubechies, Jia Li, Eric Postma, and James Z Wang, “Image processing for artist identification,” *Signal Processing Magazine, IEEE*, vol. 25, no. 4, pp. 37–48, 2008.
- [2] GÜngör Polatkan, Sina Jafarpour, Andrei Brasoveanu, Shannon Hughes, and Ingrid Daubechies, “Detection of forgery in paintings using supervised learning,” in *Image Processing (ICIP), 2009 16th IEEE International Conference on*. IEEE, 2009, pp. 2921–2924.
- [3] Siwei Lyu, Daniel Rockmore, and Hany Farid, “A digital technique for art authentication,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 101, no. 49, pp. 17006–17010, 2004.
- [4] James M Hughes, Daniel J Graham, and Daniel N Rockmore, “Quantification of artistic style through sparse coding analysis in the drawings of pieter bruegel the elder,” *Proceedings of the National Academy of Sciences*, vol. 107, no. 4, pp. 1279–1283, 2010.
- [5] Christian Robert Jacobsen and Morten Nielsen, “Authentication of paintings using hidden markov modelling of contourlet coefficients,” Tech. Rep., Department of Mathematical Sciences, Aalborg University, 2011.
- [6] Ryan Kerekes, BVK Kumar, et al., “Correlation filters with controlled scale response,” *Image Processing, IEEE Transactions on*, vol. 15, no. 7, pp. 1794–1802, 2006.
- [7] Vishnu Naresh Boddeti, Takeo Kanade, and BVK Kumar, “Correlation filters for object alignment,” in *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*. IEEE, 2013, pp. 2291–2298.
- [8] Vijayakumar Bhagavatula and Marios Savvides, “Correlation pattern recognition for biometrics,” *The international research of optical Engineering*, 2006.
- [9] BVK Kumar, Abhijit Mahalanobis, and Daniel W Carlson, “Optimal trade-off synthetic discriminant function filters for arbitrary devices,” *Optics Letters*, vol. 19, no. 19, pp. 1556–1558, 1994.
- [10] C Caspers, “Caspers Data Set,” <http://web.math.princeton.edu/ipai/datasets.html>.