Classifying recaptured identity documents using the biomedical Meijering and Sato algorithms

John Magee¹, Stephen Sheridan¹ and Christina Thorpe¹

¹School of Informatics and Cybersecurity, Technological University Dublin, Dublin, Ireland

Abstract

Recent research into the domain of identity document recapture detection demonstrated the capability of the Meijering filter, a biomedical image processing algorithm, to detect features present in recaptured documents. Manipulation of identity documents using image processing software is a low-cost, high-risk threat to modern financial systems, opening these institutions to fraud through crimes related to identity theft. In this paper we extend the research into the application of biomedical image processing algorithms, including the Meijering filter and the Sato filter. We build support vector machine and decision tree classifiers based on histograms of images generated from these filters and apply some rudimentary feature reduction techniques. The results show that both filters can be applied to this domain, with the Meijering filter slightly outperforming the Sato filter in most tests.

Keywords

Identity documents, document recapture detection, Meijering filter, Sato filter

1. Introduction

Traditional Know Your Customer (KYC) channels for financial institutions are slow, inefficient, and costly [1]. Remote customer onboarding services and electronic Know Your Customer (eKYC) services are viable alternative to traditional methods and help reduce costs and friction experienced by customers signing up for services. Retail banking institutions are increasingly including eKYC services within mobile banking apps, including the ability to open a new account remotely. This can present a hole in the security architecture as bad actors can easily use modern digital imaging software to manipulate identity documents, exposing financial institutions to fraud [2] through simple document recapture attacks. The consequences of fraud are significant; in 2022 the UK National Crime Agency reported that money laundering cost the UK economy in the region of "hundreds of billion pounds per year" ¹. A recaptured identity document is one where a copy is made of a legitimate identity document, possibly altered, and then printed on to paper as a hard copy. Document recapture attacks are a low-cost, high risk to eKYC services and such presentation attacks should be detected and immediately rejected.

Recent work by Magee et al. [3] demonstrated poten-

tial that the biomedical algorithm known as the Meijering filter can be used in the domain of identity document recapture detection. The Meijering filter [4] is a technique designed to assist the analysis of neurite growth of fluorescence images captured using microscopes. As a form of texture detection, it was used successfully to detect paper texture by Magee et al. Another biomedical filter, known as the Sato filter [5] is used to detect and enhance tubular and linear structures in medical images. Both the Meijering and Sato filters are built on techniques that use the eigenvalues of the Hessian matrix for image enhancement.

The goal of this research is to build on the recent research by Magee et al. that demonstrated the capability of the Meijering filter as a feature extraction process to help detect recaptured identity document images. This research improves on the original work by 1) attempting to improve the data set used by Magee et al., 2) reproduce the original work with the new data set and compare our results with the reported results 3) use a range of input features to test how these influence the classification accuracy, 4) test the application of the Sato filter 5) train a decision tree classifier using the Meijering and Sato filtered data and 6) compare the classification performance of the decision tree and SVM algorithms using the Meijering and Sato filtered data across the range of input features.

Workshop ISSN 1613-0073

APWG.EU Technical Summit and Researchers Sync-Up 2023 (Tech 2023), June 21 22, 2023, Dublin, Ireland

B00149241@mytudublin.ie (J. Magee); ss@mytudublin.ie

⁽S. Sheridan); ct@mytudublin.ie (C. Thorpe) © 0 * 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Authorities of Johrenational (Co RY 4.0)

Attribution 4.0 International (CC BY 4.0). CEUR Workshop Proceedings (CEUR-WS.org)

¹https://www.parliament.uk/business/lords/media-centre/houseof-lords-media-notices/2022/november-2022/the-governmentmust-take-the-fight-to-the-fraudsters-by-slowing-down-fasterpayments-and-prosecuting-corporates-for-failure-to-preventfraud/

2. Related Work

2.1. Introduction

This research reports metrics using the ISO Presentation Attack Detection standard, ISO 30107². This publication introduced specific terminology to address error rates for classification algorithms. These terms are Attack Presentation Classification Error Rate (APCER) and Bona Fide Presentation Classification Error Rate (BPCER). Readers familiar with the confusion matrix may not recognise these terms, but they are just synonyms for False Negative Rate (FNR) and False Positive Rate (FPR).

2.2. The Sato Filter

The Sato filter, introduced by Sato et al. [5], is a type of image filter used for the detection and enhancement of tubular and linear structures in medical images and is an extension of the widely used vessel enhancement filter known as the Frangi filter [6]. It is a multi-scale filter that uses the eigenvalues of the Hessian matrix to enhance linear structures in an image while reducing background noise. While this filter is typically used to detect blood vessels, in this exploratory research we are using it to detect surface texture of recaptured document images.

2.3. The Meijering Filter

Introduced by Meijering et al. [4], this technique was designed to assist the analysis of neurite growth of fluorescence images captured using microscopes and is also an extension of the Frangi filter [6]. Their algorithm is implemented in 2 phases, it first assigns each pixel in the image a probability of that pixel belonging to a neurite (known as the detection phase), it then it links together the center lines of the neurites to form the neurite tracing (this is known as the tracing phase). Conceptually, this can be considered a form of texture detection where the neurite is the textured object being detected against a fluorescent background. Texture detection is a technique previously used for document recapture detection [7, 8].

2.4. Document Recapture and Forgery Detection

Berenguel et al. [9] developed a system using image acquisition from a mobile device for counterfeit detection. Their research focused on Spanish identity documents but was also used to detect counterfeit bank notes. Their data set was generated through crowd sourcing using a mobile app, individuals were encouraged to submit identity documents and classify them (genuine/not genuine) in the mobile app. The disadvantage of this solution is that it requires model generation per document type, which increases the complexity when the number of document types and issuing countries/institutions increase.

Yang et al. [10] developed a Convolutional Neural Network (CNN) based solution to detect recaptured images. They reused images from existing data sets [11, 12] of general photographs from public sources, not identity documents. Their data set consists of 10,000 genuine and recaptured images. All images are 512 x 512 pixels in size, considerably low resolution compared to the capabilities of modern mobile cameras. Their contribution is to add a Laplician filter into the CNN as an enhancement layer. They report an classification accuracy of 99.74% for images 512 x 512 pixels, while smaller images resulted in a slightly lower accuracy.

Berenguel et al. [9] proposed a Counterfeit Recurrent Comparator (CRC) network design to identify counterfeit documents. The design is based on research on the human perception system [13]. The researchers reused a data set from their own work [8] that consists of identity documents and counterfeit bank notes. The document is split into patches and the CRC network iterates over all patches until the complete document is assessed. Their network performance is compared to PeleeNet[14], outperforming it with a mean AUC score of 0.984.

Chen et al. [15] proposed a Siamese network design to address the detection of recaptured documents. A Siamese network is a neural network design that contains two or more identical components used to find similarities between inputs. Such network designs require samples of genuine documents, as well as recaptured documents, to train the network. To address this problem, they created a database of 320 captured and 2627 recaptured document images based on generated synthetic data. Chen et al. achieved 6.92% APCER and 8.51% BPCER by their proposed network.

Magee et al. [3] investigated the application of a biomedical imaging algorithm, the Meijering filter, to the domain of document recapture detection. The Meijering filter was applied to recaptured document images and histograms were generated from these filtered images as a form of rudimentary feature extraction. A support vector machine was trained using these features as a classification model to distinguish between document images recaptured from a screen and document images recaptured from printed hard copies. Without applying any data cleaning techniques, the results were promising with a mean APCER of 15.45% for an iPhone8 and 29.35% for an iPhone12 mobile device, while both models had a BPCER of approximately 24%. These results are not state-of-the-art, but they do show the potential to use biomedical imaging algorithms in the domain of identity document recaptured detection.

²https://www.iso.org/standard/53227.html

3. Procedure

3.1. The Data Set

Magee et al. [3] concluded that their data set required cleaning as the source of their recaptured images did not truly represent that of an identity document data set. The source data set was the BID data set [16], a synthetic data set generated based on images of real Brazilian driver licenses. The purpose of BID data set was to assist with document segmentation and OCR, meaning the authors intentionally added variation (different background colours, adding noise or bold text) that is not present in genuine identity documents, consistency being a security feature of identity documents. Based on this observation, we undertook an attempt to clean the data set used by Magee et al. and we removed 10 source images from the data set due to noise and replaced them with 8 new images that contained a consistent look and feel in an attempt to reduce the amount of variance in the data set. We plan to recapture more images from the source BID data set to augment our data set in future work. Magee et al. used screen recaptured images as a ground truth data set, something to measure the printed recaptured documents against. As no data set of genuine recaptured identity documents is available, these screen recaptured images act as a proxy for genuine recaptured documents. We also replaced some of the recaptured screen images that contained obvious screen artifacts. This resulted in 22 screen recaptured images being replaced by the iPhone8 and 23 screen recaptured images being replaced by the iPhone12. This effort was only partially successful as it was not possible to obtain recaptures without some visible screen artifacts. A breakdown of the data set after the cleaning exercise is represented in Table 1.

3.2. Feature Extraction and Reduction

The feature extraction process used by Magee et al. was limited to the histogram intensity values of the Meijering filtered images with a bin width value of 1, meaning the greyscale value with 256 grey values produced 256 input features. This represents the highest resolution histogram possible from a greyscale image. As part of this research, we are applying some rudimentary data reduction by using different numbers of bins during the histogram generation process. The bins used in this work are 8, 10, 16, 32, 48, 50, 64, 128 and 256. We are also introducing the Sato filter as a new feature extraction process. The same feature extraction process described above is applied to the Sato filtered images. After applying the filters, the same screen artifacts described by Magee et al. were observed in the screen recaptured images. These are unique artifacts close to the edges of the image, indicating the area where the identity document image

transitioned to the screen background on the monitor. These artifacts are not what we expect to see in genuine recaptured documents, therefore all filtered images are cropped to remove these artifacts by removing the bordering 50 pixels from each side of the image. Python scripts were used to process all of the recaptured images and the data processing pipeline is represented in Figure 1. Image processing scripts are written in Python, version 3.8.5, Scikit-Image version 1.2.2 and OpenCV 4.4.0.42. All default parameters in the Scikit-Image library are used when generating the Meijering and Sato images.

3.3. Support Vector Machine Classifier

Magee et al. used the Support Vector Machine (SVM) classification algorithm to train a classifier to distinguish between screen and printed recaptured images. Previous research has shown the SVM classifier to be useful in this domain [17, 18, 7]. SVM models were generated for each bin number as outlined in Section 3.2. The SVM is used as a binary classifier, therefore a class value of 0 represents a label for screen recaptured images and a class value of 1 represents the label for printed recaptured images. To generate sufficient accuracy metrics, we used the same training and testing procedure as outlined by Magee et al., twenty different seeded tests were run, each using stratified 10-fold cross validation, resulting in 200 accuracy metrics for each test. The same seeds used by Magee et al. were used in this work. The APCER and BPCER metrics are computed for each test and the average metric across all tests for each device is reported. The SVM was trained using Python 3.8.5 and scikit-learn 1.2.2, all default parameters are used.

3.4. Decision Tree Classifier

The application of the No Free Lunch theorem [19] tells us that each machine learning algorithm is biased in its own way. In an effort to avoid being misled by any inherent bias in the SVM algorithm, we use the decision tree algorithm as a measure of comparison. Decision tree algorithms are easy to explain and can be visualised very easily but may also over-fit training data [20]. The exact same process reported in Section 3.3 is used to train a decision tree classification algorithm. We use three different split criteria in this research, they are Gini, Entropy and logloss. The APCER and BPCER metrics are computed for each test and the average metric across all tests for each device is reported. The decision tree was trained using Python 3.8.5 and Scikit-learn 1.2.2. Except for the split criterion, all default parameters are used.



Figure 1: This figure shows the image processing pipelines applied to all the images in the data set. Part A represents the initial data processing to apply the Meijering and Sato filters while part B represents the histogram data generation using the 9 different bin numbers (8, 10, 16, 32, 48, 50, 64, 128 and 256).

Table 1

Type and count of captured documents per device.

	Source	Printer	iPhone 8	iPhone 12
Printed		Inkjet	102	102
		Laser	102	102
Plastic		Inkjet	102	102
		Laser	102	102
Screen Re	ecapture	N/A	102	102

3.5. Results

The results of the SVM classification models are displayed in Table 2 and the results of the decision tree models are shown in Table 3.

The results of both classification algorithms indicate that models trained on the Meijering filtered data result in more accurate models than those trained on Sato data, but this is not the case across the board. For example, the SVM results in Table 2 show the BPCER for the iPhone12 using only 8 bins is 10% higher for Meijering data compared to Sato, which is a considerable difference. The decision tree results also show that models trained on the Meijering filtered data result in more accurate models than those trained on Sato data, regardless of the split criteria used. Very little variation was observed across the 3 different split criteria, Gini, Entropy and logloss, indicating that each split criteria was invariant to any bias in the data. It is also remarkable how invariant the decision tree accuracy metrics are relative to the number of bins used, the variance shown by the SVM models appears much higher. For example, for the Meijering data, the iPhone8 SVM model has an APCER of 11.84% for 8 bins and 17.64% for 256 bins whereas the decision tree APCER is 14.55% and 14.83% respectively for the same data. Decision tree models trained on the Meijering filtered data are approximately 5% more accurate that those trained on Sato filtered data. However, all these test results need to be interpreted in the context of the limited number of samples in the test data sets, meaning a single misclassification will result in a minimum error rate of 9%. The difference between the Meijering and Sato classification metrics can be down to only 1 more misclassified sample.

4. Conclusion and Future Work

The objective of this research was to build on the work of Magee et al. and continue the investigation into the use of biomedical imaging algorithms into the domain of identity document recapture detection. In this paper we introduced a rudimentary feature reduction technique by selecting different bin numbers in the histogram generation process and measuring how the features influenced the model accuracy. We introduced a second biomedical algorithm, the Sato filter, and applied all the same model generation and testing to images produced by this filter. Finally, we added a new machine learning technique, the decision tree, as a new method to compare performance against that of the SVM. We have shown that the decision tree algorithm typically out performs the SVM model when comparing the APCER and BPCER metrics and they also appear more invariant to changes in the number of input features. The results also show that the Meijering filter typically results in models that provide higher accuracy than those trained using the Sato filter, but this is not across the board and the differences relatively minor in the context of the small data sets used for testing (small classification differences result in large percentage differences). We strictly controlled the procedure used in this work, meaning there is no variation in the process used to create and test models based on data from the Meijering and Sato filters. The exact same data set is used, the same random seeds are used to control the kFold cross validation process and the same machine learning algorithm implementations are used to train and evaluate the models. As a result of this, we are confident

Table 2

This table shows the performance statistics for the SVM model trained from histogram data generated using the Sato and Meijering filtered images. The statistics are shown for each model trained using different numbers of histogram bins.

SVM Performance Statistics											
							Bins				
			8	10	16	32	48	50	64	128	256
iPhone8											
	APCER	Meijering	11.84	10.05	8.76	12.95	14.16	14.75	14.23	15.55	17.64
		Sato	11.88	19.46	10.21	17.82	17.34	15.70	17.30	15.94	18.58
	BPCER	Meijering	31.58	32.14	28.51	27.36	25.31	25.65	24.74	26.61	26.71
		Sato	38.76	30.90	39.19	25.70	24.93	23.81	23.92	25.78	25.58
iPhone12											
	APCER	Meijering	18.19	19.90	18.53	19.44	23.93	23.14	22.25	23.68	27.40
		Sato	20.26	20.84	21.55	21.26	23.05	21.31	22.98	24.39	26.68
	BPCER	Meijering	36.71	32.58	29.43	22.96	21.66	23.30	22.03	24.14	26.28
		Sato	26.83	25.16	30.60	26.75	24.29	22.46	25.68	24.20	23.32

Table 3

This table shows the performance statistics for the SVM model trained from histogram data generated using the Sato and Meijering filtered images. The statistics are shown for each model trained using different numbers of histogram bins.

Decision Tree Performance Statistics											
Metric	Filter	Split					Bins				
			8	10	16	32	48	50	64	128	256
iPhone8											
APCER	Meijering	Gini	14.55	15.66	15.59	15.60	15.33	14.31	14.15	15.91	14.83
		Entropy	14.20	14.81	14.66	14.82	15.13	13.98	14.97	14.82	15.31
		log loss	14.05	14.88	14.78	14.67	14.98	13.85	14.98	15.00	15.28
	Sato	Gini	20.70	19.16	19.42	20.88	21.86	20.62	22.20	23.89	21.20
		Entropy	19.33	19.82	19.17	20.25	20.80	20.51	20.35	23.22	21.06
		log loss	19.36	19.55	18.86	20.09	20.87	20.60	20.62	23.01	21.03
BPCER	Meijering	Gini	16.97	16.41	16.32	14.81	13.72	14.57	15.08	16.34	13.76
		Entropy	15.69	15.37	14.47	14.55	15.12	13.55	14.47	17.03	15.21
		log loss	16.08	15.40	14.31	14.73	15.29	13.05	14.12	16.76	15.23
	Sato	Gini	21.42	18.50	18.88	21.53	22.22	19.78	20.60	23.30	21.35
		Entropy	20.90	19.90	18.10	20.87	21.42	19.55	20.98	23.39	20.46
		log loss	20.80	19.70	18.23	20.77	21.22	18.37	20.93	23.06	19.78
iPhone12											
APCER	Meijering	Gini	17.83	17.15	16.11	16.80	16.57	14.44	16.59	18.75	19.24
		Entropy	18.20	15.93	15.52	15.40	14.82	15.12	16.21	18.00	18.00
		log loss	18.05	16.13	15.85	15.61	14.84	15.63	15.92	18.44	17.86
	Sato	Gini	21.13	21.79	20.94	21.94	23.12	19.69	21.02	24.66	23.50
		Entropy	20.35	20.65	21.40	21.36	21.19	18.96	21.75	25.46	24.24
		log loss	20.51	20.72	21.39	21.25	21.20	18.97	21.98	25.49	23.71
BPCER	Meijering	Gini	17.58	17.55	15.99	18.32	15.82	17.42	17.47	19.61	19.31
		Entropy	18.03	17.24	15.47	16.50	15.03	16.35	15.41	18.53	17.68
		log loss	17.81	17.52	15.71	16.70	15.03	16.46	15.26	18.88	18.16
	Sato	Gini	19.52	22.10	22.01	23.19	23.62	20.60	22.88	25.72	24.00
		Entropy	19.01	21.67	21.27	22.79	20.52	18.99	21.20	26.05	25.14
		log loss	19.06	21.43	21.55	22.47	20.59	19.29	21.69	25.63	25.40

that any variation in performance between the two filters is directly related to the algorithms.

Results reported by Magee et al., only using 256 bins, are APCER 15.45% and BPCER 24.40% for the iPhone8 compared to APCER of 17.64% and BPCER of 26.71%

for this work. They report APCER 29.35% and BPCER 24.05% for the iPhone12 compared to APCER of 27.40% and BPCER of 26.28% of this work. The results of this work show degraded classification performance in comparison to Magee et al. in 3 of the 4 results, when using

256 bins, although the difference is small. It should be noted that the best classification accuracy metrics obtained in this work are for models trained with less than 256 bins. The difference in results is due to the data cleaning exercise that we undertook. Future work will include more data capture to augment the results obtained using the Meijering and Sato filters.

References

- [1] R. Soltani, U. Trang Nguyen, A. An, A new approach to client onboarding using self-sovereign identity and distributed ledger, in: 2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (Smart-Data), IEEE, 2018, pp. 1129–1136.
- [2] H. B. Macit, A. Koyun, Tamper detection and recovery on RGB images, in: D. J. Hemanth, U. Kose (Eds.), Artificial Intelligence and Applied Mathematics in Engineering Problems, Springer International Publishing, 2020, pp. 972–981.
- [3] J. Magee, S. Sheridan, C. Thorpe, An investigation into the application of the meijering filter for document recapture detection. under review for the 12th international conference on intelligent information processing (iciip 2023), not yet published., 2023.
- [4] E. Meijering, M. Jacob, J.-C. Sarria, P. Steiner, H. Hirling, M. Unser, Design and validation of a tool for neurite tracing and analysis in fluorescence microscopy images, Cytometry Part A 58A (2004) 167–176.
- [5] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, R. Kikinis, Threedimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images 2 (1998) 143–168.
- [6] A. F. Frangi, W. J. Niessen, K. L. Vincken, M. A. Viergever, Multiscale vessel enhancement filtering, in: Medical Image Computing and Computer-Assisted Intervention – MICCAI'98, volume 1496, Springer Berlin Heidelberg, 1998, pp. 130–137.
- [7] X. Hou, T. Zhang, G. Xiong, Y. Zhang, X. Ping, Image resampling detection based on texture classification 72 (2014-09) 1681–1708.
- [8] A. B. Centeno, O. R. Terrades, J. L. i. Canet, C. C. Morales, Evaluation of texture descriptors for validation of counterfeit documents, in: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR), IEEE, 2017-11, pp. 1237–1242.
- [9] A. Berenguel, O. R. Terrades, J. Llados, C. Canero, E-counterfeit: A mobile-server platform for doc-

ument counterfeit detection, in: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR), IEEE, 2017, pp. 15–20.

- [10] P. Yang, R. Ni, Y. Zhao, Recapture image forensics based on laplacian convolutional neural networks, in: Y. Q. Shi, H. J. Kim, F. Perez-Gonzalez, F. Liu (Eds.), Digital Forensics and Watermarking, volume 10082, Springer International Publishing, 2017, pp. 119–128. Series Title: Lecture Notes in Computer Science.
- [11] H. Cao, A. C. Kot, Identification of recaptured photographs on LCD screens, in: 2010 IEEE International Conference on Acoustics, Speech and Signal Processing, IEEE, 2010, pp. 1790–1793.
- [12] R. Li, R. Ni, Y. Zhao, An effective detection method based on physical traits of recaptured images on LCD screens, in: Y.-Q. Shi, H. J. Kim, F. Pérez-González, I. Echizen (Eds.), Digital-Forensics and Watermarking, volume 9569, Springer International Publishing, 2016, pp. 107–116.
- [13] P. Shyam, S. Gupta, A. Dukkipati, Attentive recurrent comparators, in: D. Precup, Y. W. Teh (Eds.), Proceedings of the 34th International Conference on Machine Learning, volume 70 of *Proceedings of Machine Learning Research*, PMLR, 2017, pp. 3173– 3181.
- [14] R. J. Wang, X. Li, C. X. Ling, Pelee: A real-time object detection system on mobile devices, in: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 31, Curran Associates, Inc., 2018.
- [15] C. Chen, S. Zhang, F. Lan, J. Huang, Domainagnostic document authentication against practical recapturing attacks 17 (2022) 2890–2905.
- [16] D. S. Soares, R. B. Das Neves Junior, B. L. D. Bezerra, BID dataset: a challenge dataset for document processing tasks, in: Anais Estendidos da Conference on Graphics, Patterns and Images (SIBRAPI Estendido 2020), Sociedade Brasileira de Computação, 2020, pp. 143–146.
- [17] C.-Y. Yeh, W.-P. Su, S.-J. Lee, Employing multiplekernel support vector machines for counterfeit banknote recognition 11 (2011) 1439–1447.
- [18] V. Lohweg, J. L. Hoffmann, H. Dörksen, R. Hildebrand, E. Gillich, J. Hofmann, J. Schaede, Banknote authentication with mobile devices, 2013, p. 866507.
- [19] H. Wolpert, D, G. Macready, W., No free lunch theorems for optimization, in: IEEE Transactions on Evolutionary Computation, 1997.
- [20] G. James, D. Witten, T. Hastie, R. Tibshirani, An Introduction to Statistical Learning: with Applications in R, Springer Texts in Statistics, Springer US, 2021.