An Investigation Into the Application of Anomaly Detection and the Meijering Filter in the eKYC Process to Detect Recaptured Identity Documents

John Magee[1,†], Stephen Sheridan[1] and Christina Thorpe[1]

[1]*School of Informatics and Cybersecurity, Technological University Dublin, Dublin, Ireland*

Abstract

As financial institutions move to offer more digital based services, the need for secure and accurate eKYC services increases. Identity documents submitted during the eKYC process are fundamental to establish the identity of customers. The ability of bad actors to modify identity documents using digital imaging software opens these eKYC services to new threats, resulting in identity theft and reputational damage. In this exploratory research we apply auto-encoder/decoder anomaly detection to the domain of recaptured identity document detection, using the Meijering filter as the feature extraction mechanism.

Keywords

image processing, anomaly detection, auto-encoder/decoder, Meijering filter, document fraud, recaptured identity documents,

1. Introduction

Traditional Know Your Customer (KYC) solutions require banking customers to present themselves at a physical building, a slow and costly process [1]. eKYC solutions allow customers to sign up for new services remotely, greatly benefiting banking institutions. It also opens up new attack vectors from bad actors looking to commit fraud or identity theft. Regulations [1], such as Anti-Money Laundering (AML, including Politically Exposed Persons (PEPs) screening [2]) and Countering the Financing of Terrorism (CFT) regulations impose strict identity proofing standards. The ability to prove the authenticity of an identity document during a remote account opening is fundamental in establishing the true identity of a new customer and to ensure compliance with such regulations.

The advanced nature of computer graphic software means it is easy to create/modify images of identity documents. Recent research by Magee et al. [2, 3, 4] demonstrated that biomedical imaging filters (the Meijering filter, the Sato filter and Frangi filter) can be used to help identify recaptured identity documents. Their best results were obtained using the Meijering filter and a Random Forest classifier (APCER=7.5%, BPCER=6.6%)[4]. However, it is not possible to know what new attack vectors bad actors will use when attempting to by-pass security mechanisms designed to detect fraudulent documents, thus exposing a limitation on the use of a classification based systems in this domain.

Machine learning techniques are divided into two categories, classification and regression. Classification techniques are designed around the premise that a dataset is divided into discrete instances, referred to as 'classes'. Training techniques ensure the model learns features to distinguish between classes. Regression techniques predict a value across a continuous range [5]. These techniques require datasets that are labeled. It is not always possible to have such datasets available when researching sensitive domains, such as identity document fraud. A paradigm known as anomaly detection [6, 7] can be useful in such circumstances. This technique takes advantage of an understanding of expected data and

AICS'24: 32nd Irish Conference on Artificial Intelligence and Cognitive Science, December 09–10, 2024, Dublin, Ireland

*Corresponding author.

†Primary author.

B00149241@mytudublin.ie (J. Magee); stephen.sheridan@mytudublin.ie (S. Sheridan); christina.thorpe@mytudublin.ie (C. Thorpe)

[1] https://www.centralbank.ie/regulation/anti-money-laundering-and-countering-the-financing-of-terrorism

[2] https://www.dataprotection.ie/en/faqs/banking-insurance-sector/what-politically-exposed-persons-pep-screening

attempts to identify anomalous patterns [7]. Hilton & Salakhutdinov [8] showed that auto-encoders can reduce the dimensions of the input data into a more efficient latent space. The decoder regenerates the latent space into the input format. The difference between the input and reconstruction is leveraged to detect anomalous events [7]. The auto-encoder/decoder architecture is also fundamental to the modern development of generative AI [9]. Based on our previous research, we pose the following research question: "Can anomaly detection using neural network based auto-encoder/decoder architecture be used to detect recaptured identity documents processed by the Meijering filter?".

The structure of this paper is as follows: Section 2 presents a brief overview of the related work; Section 3 introduces the experiment design; Section 4 presents our results and discussion; Section 5 presents our conclusion and future work.

2. Related Research

This section introduces a review of the identity document datasets, techniques used to date for document recapture/forgery detection and the usage of anomaly detection across industries.

2.1. Identity Document Data Sets

The Mobile Identity Document Video (MIDV) is a data set released by Smart Engines [3]. It consists of captured synthetic identity documents, the purpose is to further advancement in Optical Character Recognition (OCR) and Machine Readable Zone (MRZ) processing. Three versions of MIDV exist, MIDV-500[10], MIDV-2019[11] and MIDV-2020[12]. MIDV-500 consists of 500 video clips for 50 identity document types as well as annotated images in TIFF format. MIDV-2019 extended the original dataset by including high resolution images captured under different lighting conditions. MIDV-2020 further expanded the data set, including JPEG images of the documents. As this dataset does not include recaptured identity documents, it is not applicable for use in this experiment.

Kumar et al. [13] describe an image generation system for synthetic passports, driver's licenses and Visa stickers. They have published a data set consisting of 15,000 images, available on Kaggle [4]. Special character sets and fonts exist in official documents. The goal of this dataset is to further the study in this field as one of the leading indicators of forged documents is typically font related i.e. the forger used the wrong fonts.

More recently, Benalcazar et al. [14] published another mechanism for synthetic data generation using three different approaches, including computer vision algorithms and Generative Adversarial Networks (GANs). Their dataset includes print/scan recapture and screen recapture images of documents. Using the Chilean ID Card as a template, they used a range of random facial images (from the FERET database), signatures, names and dates to generate random instances of ID Cards. They developed the ability to transfer different types of noise into the images to simulate actual capture from a device e.g. screen recapture. Their data set is not publicly available as access is controlled, otherwise this would considered for use in this experiment.

Soares et al. [15] released the Brazilian Identity Document (BID) data set, a synthetic identity document dataset generated using the Brazilian ID Card as a template. The purpose of this data set is to advance the domain of document segmentation and Optical Character Recognition (OCR) research. This is the source dataset used to build the dataset for our experiments as it contains both portrait and landscape captures and when printed are approximately the same size as the information page on a passport book. This is published on Kaggle [5].

Chen et al. [16] published a data augmentation method to assist with the training of algorithms for forged document detection. Their method was shown to increase the accuracy of deeply trained models by 6.8%. While this work doesn't specifically focus on identity documents, some of the techniques are

[3] https://smartengines.com/

[4] https://www.kaggle.com/turabbajeer

[5] https://www.kaggle.com/datasets/johnmageetud/recaptured-identity-documents

applicable to this domain as special documents e.g. certificates, contain complex features specifically designed to ensure authenticity, and by design to make the documents harder to forge. It is not clear if this approach can be used in the context of anomaly detection.

2.2. Classification for Document Recapture/Forgery Detection

Berenguel et al. [17] developed a document classification based on texture analysis. Their system uses a Naïve Bayes classifier based on multivariate Bernoulli distribution to detect if a document was genuine or counterfeit. The Naïve Bayes classifier inputs are constructed by a sequence of components that include Principal Component Analysis (PCA) and Linear Support Vector Machines (SVM). They construct a dataset of crowd sourced recaptures of the Spanish ID Card, capturing both sides of the ID Card. This is a binary classification network and they report a classification F-score of approximately 98% across each side of the ID Card. Further work by Berenguel et al. [18] was the development of the Counterfeit Recurrent Comparator (CRC) to identify counterfeit documents. This design takes into consideration existing research of the human perception system [19]. This network is trained using an existing dataset of counterfeit bank notes from their own research [20] and they report a AUC of 0.984.

Yang et al. [21] developed a Convolutional Neural Network (CNN) to detect recaptured images. Their network is a binary classification network, classifying if an image is genuine or if it is recaptured. This research is not specific to identity documents, but it can be applied to identity document images if a sufficient dataset was available. Their network is trained using lower quality images, only 512 by 512 pixels in size. Their research reports a classification accuracy of 99.74%. Testing with even lower sized images reduced this accuracy slightly. However, this approach shows that processing images in patches is possible alternative to processing whole images which is something we use in the experiment documented in this paper.

Chen et al. [22] developed a Siamese neural network to detect recaptured documents. A Siamese network is a complicated neural network design containing of two identical components that find similarities between inputs. Despite its complexity, in this case it was ultimately configured to be a binary classification network. They train and test their network using synthetic document data and achieve 6.92% Attack Presentation Classification Error Rate (APCER) and 8.51% Bonafide Presentation Classification Error Rate (BPCER).

Research from Magee et al. [2, 3, 4]. has also focused on a classification approach to recaptured identity document detection. This research utilises the Meijering filter as a feature extraction process and the Random Forest classification algorithm. They achieved APCER 7.53% and BPCER 6.59% using an optimised Meijering filter configuration discovered using a grid search approach. These accuracy results are comprable to the results obtained by Chen et al. using their complex Siamese neural network architecture.

The papers presented above [18, 20, 21, 22] show a clear trend toward the use of neural network architectures for classification models. Our own work [2, 3, 4] used traditional machine learning models for the purposes of classification. All these algorithms depend on the existence of structured data for training, which is a limitation we have identified. There are no examples of the use of anomaly detection.

2.3. Anomaly Detection Algorithms

A comparative analysis by Kharitonov et al.[23] shows the most common anomaly detection algorithms include neural network based auto-encoders, Isolation Forest, K-Nearest Neighbour(KNN). The success of neural networks in recent years has propelled the use of auto-encoders in this domain.

Autonomous vehicles are a recent literature source for anomaly detection. Rezaei et al. [24] used a GAN based auto-encoder/decoder architecture for sensor fault detection as well as to protect vehicle systems from cyber attacks. Their network follows the auto-encoder/decoder architecture where the encoder is under-complete [25] (the number of neurons in each layer decreases, forcing the network to learn a smaller dimension latent space representation of the input). The decoder network then uses the over-complete network architecture to reconstruct the reduced dimensional representation back to that of the original input. They report their results in terms of accuracy, sensitivity and specificity derived from the range of anomaly scores. Han et al. [26] developed an anomaly detection system based on the transformer architecture for GPS spoofing threats. Their system combines image and positional data to detect anomalies, the images broken down into patches. The image and positional data are merged and encoded by the transformer. This encoded feature set is the input to a one-class classifier Multi-Layer Perceptron (MLP), computing the probability that the input is anomalous. Han et al. report results using the F1 and Area Under the Curve (AUC) metrics but no threshold is defined. Di Biase et al. [27] present a computer vision based pixel-wise anomaly detection framework. This is specifically designed to detect anomalous objects in the path of a vehicle (e.g. a person, a dog or another vehicle) using an image re-synthesis approach. Their encoder is an under-complete architecture based on the VGG-16 [28] network architecture. Their decoder component is an over-complete architecture but adds the novelty by concatenating a generated semantic image into each decoder network layer, augmenting the role of the decoder. Their metrics are based on the false positive rate as a fixed threshold of 95% true positive rate.

An overview of the application of anomaly detection in industry is presented by Kharitonov et al. [23]. They show traditional machine learning methods, e.g. KNN and Isolation Forest, have a wide application. They investigate the robustness of each algorithm to predict machine breakdowns based on historical data. Their experiments show that traditional machine learning algorithms can still outperform neural network based auto-encoders. Gruber & Heselmann [29] developed a Frangi filter [30] based detection system to detect defects in transparent materials. The Frangi filter is a biomedical image filter, designed to enhance vessels in images. Their system uses feature extraction from filtered images and the KNN algorithm for fault detection. Chen et al. [31] present a system for fibre inspection in an industrial textile setting. They use an under-complete LeNet-5 [32] CNN based encoder component and a Support Vector Machine to predict fibre classification. Input images are split into patches for processing. They measure their solution effectiveness using the precision metric when varying the amount of data used in the training process. When their system is trained with 80% of the available training data it is able to outperform existing state of the art.

Medical applications of anomaly detection is using autoencoder/decoder architectures is common in electrocardiogram (ECG) event monitoring. Verardo et al. [33], Lomoio et al. [34] and Shan et al. [35] all follow this approach, using the MSE or MAE loss functions.

2.4. Gap Analysis

It is clear that anomaly detection is a widely accepted method in industry and we take inspiration from the diverse examples of its application across different domains. We note that there are no examples of anomaly detection being used in the domain of recaptured identity document detection, classification models[17, 18, 20, 21, 22, 2, 3, 4] are used exclusively. Therefore, we will apply anomaly detection to the domain of recaptured identity document detection in this experiment.

3. Methodology

3.1. Data set

This research reuses our dataset used in [4] that consists of recaptured images from the BID dataset [15]. The scope of this research extends the base dataset to include captures from an Android device, a Motorola G62.

The recaptured dataset consists of two types of recaptured identity documents, 306 screen recaptured images and 918 printed recaptured documents, captured using two iPhone models (8 & 12) and a Motorola G62 Android device. Printed recaptured documents are subdivided into printed paper recaptures and plastic covered printed paper recaptures. The final breakdown of the dataset used in this research is represented in Table 1. A sample capture of the raw documents and the resultant images after being processed by the Meijering filter are shown in Figure 1.

Table 1
Type and count of recaptured documents per device.

Source	Printer	iPhone 8	iPhone 12	M G62
Paper	Inkjet	102	102	102
	Laser	102	102	102
Plastic covered	Inkjet	102	102	102
	Laser	102	102	102
Screen Recapture	N/A	102	102	102

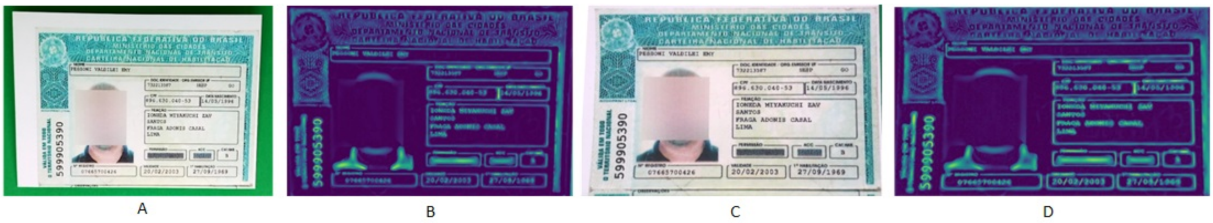

Figure 1: Shows (A) the raw captured image of an inkjet printout from the Motorola G62 and (B) the resultant image after being processed using the Meijering filter. (C) the raw captured image of an inkjet printout from the iPhone8 and (D) the resultant image after being processed using the Meijering filter.

3.2. Network Architecture

The network architecture used in this research is described in Table 2. The encoder component represents a traditional under-complete network, where the network is forced to learn a lower dimension representation of the input [25]. The decoder component is the traditional over-complete network that learns to reconstruct the original input. The network is not symmetrical as the decoder component contains one more layer than the encoder component. It is well-known that auto-encoder/decoders always produce a lower resolution reconstruction of the input data, therefore this is included to provide the decoder with additional weights (degrees of freedom) to better reconstruct the input. As this research is a proof of concept, a minimal network architecture is used consisting only of dense layers.

Table 2
Describes the individual layers in the encoder-decoder architecture.

Component	Layer	Input	Output	Activation
Encoder	Input	9176	9176	
	1	9176	1248	ReLU
	2	1248	512	ReLU
Decoder	Input	512		ReLU
	1	512	768	ReLU
	2	768	1248	ReLU
	Output	1248	9176	ReLU

3.3. Loss Functions

The most common loss functions used by neural network based auto encoders are the Mean Absolute Error (MAE), Mean Square Error (MSE) and Cosine Similarity (CS). The MSE [36] is defined in Equation 1 and computes the squared error between the value, y, and the predicted value $f(x)$, and computes the mean of these values across the number of samples n.

$$MSE = \frac{1}{n}\sum_{i=1}^{n}\left(y_{\mathrm{i}} - f\left(x_{\mathrm{i}}\right)\right)^{2} \tag{1}$$

The MAE [37] loss function is defined in Equation 2. A variation on the MSE, the absolute error value is computed, and the mean is taken across the number of samples *n*.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}\left|y_{\mathrm{i}} - f\left(x_{\mathrm{i}}\right)\right| \tag{2}$$

The CS [38] loss function is defined in Equation 3. This measures the angular distance between the given and predicted values. This has been shown to be better than metrics like MSE and MAE that attempt to measure a linear distance. This metric has resulted in exceptional results in the domain of natural language processing and is integral to the functioning of the transformer architecture [39].

$$CS = \frac{\sum_{i=1}^{n}\left(y_{\mathrm{i}} \times f\left(x_{\mathrm{i}}\right)\right)}{\sqrt{\sum_{i=1}^{n}\left(y_{\mathrm{i}}\right)^{2}} \times \sqrt{\sum_{i=1}^{n} f\left(x_{\mathrm{i}}\right)^{2}}} \tag{3}$$

3.4. Data Processing

The image processing pipeline used in this research consists of the following steps:

1. *Image filtering*. Apply the Meijering filter to each of the input images. Input image dimensions are 1344x848 pixels. The filtered output are gray scale images of dimensions 1244x748 pixels, saved using the viridis colour space.
2. *Image cropping*. The images are cropped to a final dimensions of 1240x740 by removing the outer most pixels.
3. *Generating image patches*. Each image is divided evenly into 100 patches of dimensions 124x74 pixels.

The final step is to aggregate the individual files containing the intensity pixel values so they can be used in the training process. This step simply copies individual patch data into CSV files, generating the training and validation dataset files. The output of this process provides us with:

1. *Training dataset*. Consists of 27,540 patches of filtered screen recaptured documents.
2. *Test dataset*. Consists of 3,060 patches, reserved from the training dataset, that is used to test all models. This test dataset remains static across all models and all training files.
3. *Validation dataset*. This consists of 122,400 patches of the filtered printed recaptured documents. This is considered anomalous data for the purposes of this experiment.

3.5. Training the network

The NVIDIA GeForce GTX 1650 Ti GPU used in this research can only accommodate approximately 2GB of data, therefore the data is distributed across 10 separate files. The training procedure uses each file sequentially, enabling the ability to train in batch mode. The Keras [6] framework, version 2.10.0, is used to train the models, with the ADAM optimizer. The three popular loss functions in the field of anomaly detection are used, MSE, MAE and CS. The following safe guards are employed in order to avoid over fitting the training data: a) we randomise the input data so the algorithm cannot learn from a predictable input sequence, b) we enable shuffling of training data after each epoch. The models are trained using four different epoch levels, 50, 100, 150 and 200. We then compare the performance of models trained for the same number of epochs. No significant increase in performance was observed past 200 epoches.

[6] https://keras.io/

3.6. Testing the Model

The performance of machine learning classifiers typically report results using the confusion matrix and its constituent metrics i.e., True Positive Rate (TPR), True Negative Rate (TNR), False Positive Rate (FPR) and the False Negative Rate (FNR) [37]. When measuring the performance of an anomaly detection model, these metrics cease to be as useful because the model is not aware of classes. The results in this research are presented using the range based Equal Error Rate (EER) metric [40]. This approach determines a threshold value at which both error rates (FPR and FNR) are equal. Measuring error rates is industry standard for biometric and document classification systems as error rates are a measure of security of the system. This is formalized in the ISO Presentation Attack Detection standard (30107) [7].

3.6.1. Algorithm to compute EER

The output of the testing procedure is a range of MAE scores, one for each patch, from the test and validation datasets. These are processed by the algorithm described in Algorithm 1.

Algorithm 1 Algorithm to compute logistics curves to determine EER

```
Require: $len(test\_scores) > 0$
Require: $len(validation\_scores) > 0$
 1: procedure ComputeEER($test\_scores, validation\_scores$)
 2:     $min\_threshold \leftarrow min(test\_scores)$
 3:     $max\_threshold \leftarrow max(validation\_scores)$
 4:     $threshold\_range \leftarrow seq(min\_threshold, max\_threshold, 0.1)$
 5:     $index \leftarrow 0$
 6:     $index \leftarrow$next_$value$
 7:     while $index \leq len($threshold_range$)$ do
 8:         $threshold\_value \leftarrow$threshold_range$[index]$
 9:         $fa \leftarrow count($test_scores > $threshold\_value)$
10:         $fn \leftarrow count($validation_scores <= $threshold\_value)$
11:         save( $threshold\_value, fa, fn)$
12:         $index \leftarrow$index + 1
13:     end while
14: end procedure
```

The algorithm is based on a computation of the false classifications at a specific threshold value. The algorithm is initialised on lines 2 and 3 by obtaining the minimum score from the test dataset and the maximum score from the validation dataset. These values are used in line 4 to compute a sequence of threshold values, in increments of 0.1, between the minimum and maximum values using the *seq* function in **R**. The procedure iterates over the threshold values and computes the number of misclassifications at each threshold (lines 8 & 9). The threshold and the misclassifcation rates are then saved to file (line 11). A chart plotting the misclassification rates is shown in Figure 2, where the reader can see the intersection point of the two logistic curves, representing the threshold to provides the equal error rate.

4. Results and Discussion

The results are shown in Table 3. The model trained using Cosine Similarity produced lower EER values than the MAE and MSE models. The lowest EER obtained using MAE=9.03% compared to Cosine Similarity=8.08%, while MSE=9.12%. The lowest mean EER is achieved by the models trained using Cosine Similarity=39.58% compared to MAE=44.94% and MSE=45.45%. The highest EERs are produced by dark patches where no real distinguishing information is available.

[7] https://www.iso.org/standard/83828.html

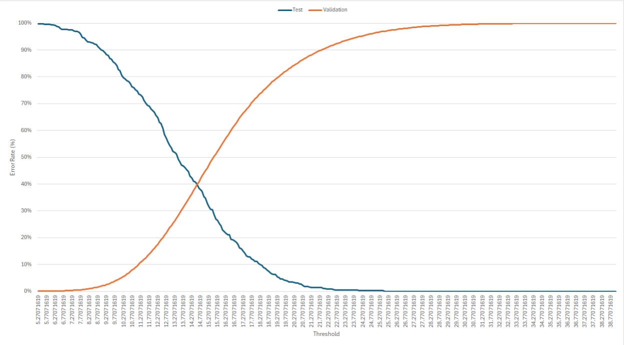

Figure 2: Chart showing the EER plotted to show the intercept point of the two curves.

Table 3
EER statistics for models trained using loss functions MAE, MSE and Cosine Similarity

	MAE			**MSE**			**Cosine**		
Epochs	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max
50	11.39%	47.10%	93.97%	9.16%	46.36%	93.86%	8.57%	41.33%	**81.86%**
100	9.13%	46.06%	93.35%	9.57%	45.78%	93.49%	**8.08%**	40.73%	82.48%
150	**9.03%**	45.55%	93.37%	9.55%	45.64%	93.25%	8.14%	39.71%	83.33%
200	9.13%	**44.94%**	**92.95%**	**9.12%**	**45.45%**	**93.07%**	**8.08%**	**39.58%**	82.76%

All models show a large EER range, due to the relevant performance of different patches. We see that each loss function responds differently to the input patch. Models trained using the MAE and MSE loss function produce models that result in significantly lower EER for vertical document features while models trained using the Cosine Similarity produce lower EER on horizontal document features. Figures 3 and 4 show patches with vertical and horizontal features, overlaid with the mean patch EER, highlighting for the reader the difference in performance. Further research is required to understand if this is a true response to features intrinsic to the documents or if this is just co-incidence. Visual inspection shows that patches tend to be reconstructed to a higher degree of accuracy for test samples (see Figure 5) compared to validation samples (see Figure 6). Validation patches reconstructed by the models also shows signs of hallucination, as shown in Figures 7 and 8.

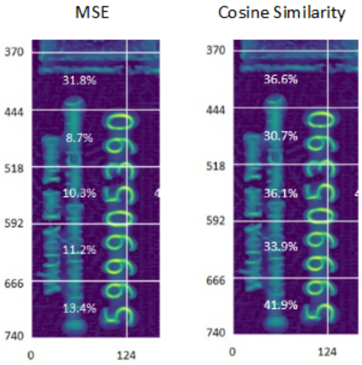

Figure 3: Comparing the performance of MSE and Cosine Similarity models based on vertical features. See the patch EER for highlighting the difference in performance.

The loss function used to train a model is critical. While MSE and MAE are extensively used, as shown in our literature review in Section 2, they do not have the sensitivity to accurately represent the error in image reconstruction scenarios like the one we undertook in this research. Similarly, the CS loss function was capable of training a model that produced lower mean EER, but still far too high for practical use.

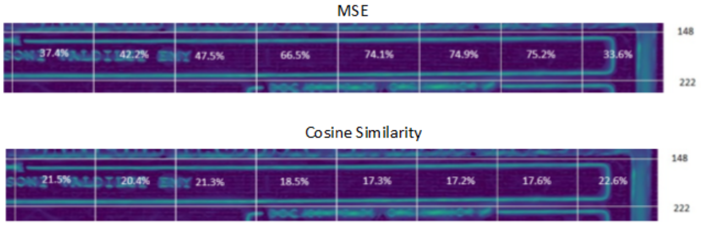

Figure 4: Comparing the performance of MSE and Cosine Similarity models based on horizontal features. See the patch EER for highlighting the difference in performance.

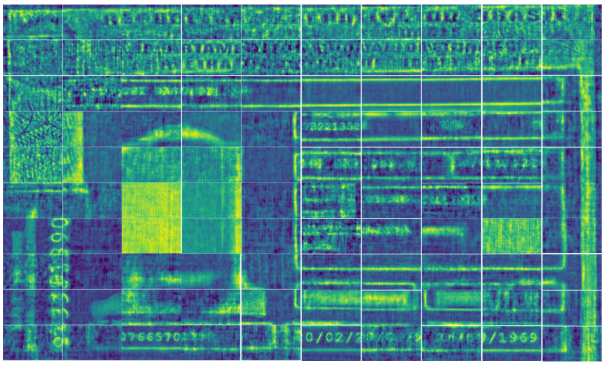

Figure 5: Manual reconstruction of patches from a screen recaptured document, image recaptured using an iPhone8.

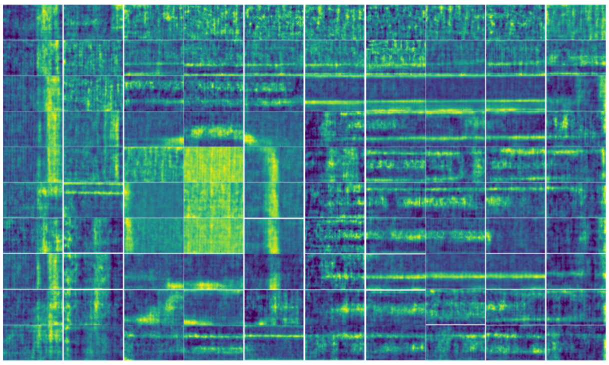

Figure 6: Manual reconstruction of patches from a paper recaptured document from an inkjet printed, image recaptured using an iPhone8.

5. Conclusion

We posed the following research question: "Can anomaly detection using neural network based autoencoder-decoder architecture be used to detect recaptured identity documents processed by the Meijering filter?". Based on our results, the answer is a qualified "Yes". We see the MSE and CS loss learn different patterns, CS performed better with horizontal features while MSE performed better with vertical features. These features correspond to parts of the document that produce visually different impulse responses from the Meijering filter, thus validating this approach.

Care is needed to ensure we don't infer too much about how real identity document images will behave under the same experimental conditions, but we have demonstrated the potential of this approach. Another aspect that requires caution is that this research evaluated the EER at a patch level, not considering the document as a whole entity with a single metric. Despite its limitations, this proof-of-concept shows the potential for the use of anomaly detection when detecting recaptured identity documents. While this paper specifically targeted the eKYC process for financial institutions, the practical use of this technology can extent to any sector where establishment of identity is important,

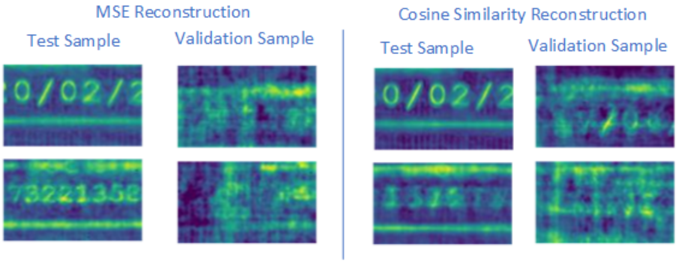

Figure 7: Visualising the reconstruction of sample patches for models trained using MSE and Cosine Similarity loss functions.

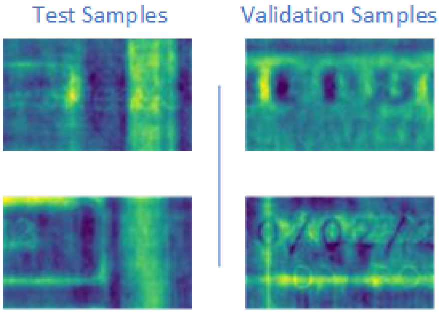

Figure 8: Examples of generated patches that demonstrate hallucinations by the network. Samples from the test dataset recreate the actual patch whereas the same patches from the validation set appear to make up features that are not present in the actual input.

such as the government and health sectors.

5.1. Future Work

Based on the limitations of the loss functions, investigations into different loss functions is a priority. Research by Yan et al. made use of a loss function to combine MSE and Cosine Similarity. Other loss functions typically associated with generative models, MinMax, Wasserstein and Diffusion loss functions, are obvious next steps. Experimenting with deep model architectures like ResNet-50 and Diffusion models are on our list of next steps. In this research we trained our own model from the beginning as this was simply a proof of concept, however, it is well established that transfer learning is a valid strategy to improve the training of neural networks by leveraging the weights from models that are already trained. Expanding our dataset is critical to ensure the results of this research generalise across other mobile devices, and to capture realistic input data to represent genuine identity document captured images. Investigating the use of an error metric for the entire document rather than individual patches is also in the scope for future work.

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), IEEE, 2018, pp. 1129–1136. doi:10.1109/Cybermatics_2018.2018.00205.

[2] J. Magee, S. Sheridan, C. Thorpe, An investigation into the application of the meijering filter for document recapture detection, in: International Conference on Cloud and Big Data Computing (ICCBDC2023), 2023. doi:10.12720/jait.15.1.132-137.

[3] J. Magee, S. Sheridan, C. Thorpe, Classifying recaptured identity documents using the biomedical meijering and sato algorithms, in: APWG.EU Technical Summit and Researchers Sync-Up, APWG, 2023. URL: https://ceur-ws.org/Vol-3631/paper2.pdf.

[4] J. Magee, S. Sheridan, C. Thorpe, Optimization of biomedical imaging filters for use in recaptured identity document classification, in: Cybersecurity Ireland (ICCBDC2023), 2023. doi:10.1109/Cyber-RCI59474.2023.10671521.

[5] C. M. Bishop, Pattern Recognition and Machine Learning, Springer, New York, 2009.

[6] C. C. Aggarwal, Neural Networks and Deep Learning, Springer, New York, 2018.

[7] S. K. Adari, S. Alla, Beginning Anomaly Detection Using Python-Based Deep Learning, Springer, 2024.

[8] G. E. Hinton, R. R. Salakhutdinov, Reducing the dimensionality of data with neural networks (2006) 504–507. doi:10.1126/science.1127647.

[9] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models (2020). doi:10.48550/ARXIV.2006.11239.

[10] V. V. Arlazarov, K. Bulatov, T. Chernov, V. L. Arlazarov, MIDV-500: A dataset for identity documents analysis and recognition on mobile devices in video stream (2019). doi:10.18287/2412-6179-2019-43-5-818-824.

[11] K. Bulatov, D. Matalov, V. V. Arlazarov, MIDV-2019: Challenges of the modern mobile-based document OCR (2020) 64. doi:10.1117/12.2558438. arXiv:1910.04009.

[12] K. Bulatov, E. Emelianova, D. Tropin, N. Skoryukina, Y. Chernyshova, A. Sheshkus, S. Usilin, Z. Ming, J.-C. Burie, M. M. Luqman, V. V. Arlazarov, MIDV-2020: A comprehensive benchmark dataset for identity document analysis (2022). doi:10.18287/2412-6179-CO-1006.

[13] T. Kumar, M. Turab, S. Talpur, R. Brennan, M. Bendechache, Forged character detection datasets: Passports, driving licences and visa stickers (2022). doi:10.5121/ijaia.2022.13202.

[14] D. Benalcazar, J. E. Tapia, S. Gonzalez, C. Busch, Synthetic id card image generation for improving presentation attack detection, IEEE Transactions on Information Forensics and Security (2023) 1814–1824. doi:10.1109/TIFS.2023.3255585.

[15] 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.

[16] C. Chen, B. Li, R. Cai, J. Zeng, J. Huang, Distortion model-based spectral augmentation for generalized recaptured document detection (2024). doi:10.1109/TIFS.2023.3333548.

[17] A. Berenguel, O. R. Terrades, J. Llados, C. Canero, E-counterfeit: A mobile-server platform for document counterfeit detection, in: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR), IEEE, 2017, pp. 15–20. doi:10.1109/ICDAR.2017.390.

[18] A. Berenguel, O. Ramos Terrades, J. Llados Canet, C. Canero Morales, Recurrent comparator with attention models to detect counterfeit documents, in: 2019 International Conference on Document Analysis and Recognition (ICDAR), IEEE, 2019, pp. 1332–1337. doi:10.1109/ICDAR.2019.00215.

[19] 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.

[20] 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, pp. 1237–1242. doi:10.1109/ICDAR.2017.204.

[21] P. Yang, R. Ni, Y. Zhao, Recapture image forensics based on laplacian convolutional neural networks, in: Digital Forensics and Watermarking, volume 10082, Springer International Publishing, 2017, pp. 119–128. doi:10.1007/978-3-319-53465-7_9.

[22] C. Chen, S. Zhang, F. Lan, J. Huang, Domain-agnostic document authentication against practical recapturing attacks (2022) 2890–2905. doi:10.1109/TIFS.2022.3197054.

[23] A. Kharitonov, A. Nahhas, M. Pohl, K. Turowski, Comparative analysis of machine learning models for anomaly detection in manufacturing (2022). doi:10.1016/j.procs.2022.01.330.

[24] S. Rezaei, N. Masoud, A. Khojandi, GAAD: GAN-enabled autoencoder for real-time sensor anomaly detection and recovery in autonomous driving (2024). doi:10.1109/JSEN.2024.3361460.

[25] Y. B. Ian Goodfellow, A. Courville, Deep Learning, MIT Press, 2016.

[26] M. Du, Z. Chen, C. Liu, R. Oak, D. Song, Lifelong anomaly detection through unlearning, in: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, ACM, 2019, pp. 1283–1297. doi:10.1145/3319535.3363226.

[27] G. Di Biase, H. Blum, R. Siegwart, C. Cadena, Pixel-wise anomaly detection in complex driving scenes, in: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, 2021, pp. 16913–16922. doi:10.1109/CVPR46437.2021.01664.

[28] S. Liu, W. Deng, Very deep convolutional neural network based image classification using small training sample size, in: 2015 3rd IAPR Asian Conference on Pattern Recognition (ACPR), 2015, pp. 730–734. doi:10.1109/ACPR.2015.7486599.

[29] D. P. Gruber, M. Haselmann, Inspection of transparent objects with varying light scattering using a frangi filter (2021) 27. doi:10.3390/jimaging7020027.

[30] A. F. Frangi, W. J. Niessen, K. L. Vincken, M. A. Viergever, Multiscale vessel enhancement filtering, in: W. M. Wells, A. Colchester, S. Delp (Eds.), Medical Image Computing and Computer-Assisted Intervention — MICCAI'98, volume 1496, Springer Berlin Heidelberg, 1998, pp. 130–137.

[31] Z. Chen, Y. Xiu, Y. Zheng, X. Wang, Q. Wang, D. Guo, Y. Wan, A weakly supervised learning pipeline for profiled fibre inspection (2024) 772–784. doi:10.1049/ipr2.12984.

[32] Y. Lecun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, Proceedings of the IEEE 86 (1998) 2278–2324. doi:10.1109/5.726791.

[33] G. Verardo, M. Boman, S. Bruchfeld, M. Chiesa, S. Koch, G. Q. Maguire, D. Kostic, FMM-head: Enhancing autoencoder-based ECG anomaly detection with prior knowledge (2023). doi:10.48550/ARXIV.2310.05848.

[34] U. Lomoio, P. Vizza, R. Giancotti, G. Tradigo, S. Petrolo, S. Flesca, P. Hiram Guzzi, P. Veltri, AUTAN-ECG: An AUToencoder bAsed system for anomaly detectioN in ECG signals, 2023. doi:10.36227/techrxiv.24638856.v1.

[35] L. Shan, Y. Li, H. Jiang, P. Zhou, J. Niu, R. Liu, Y. Wei, J. Peng, H. Yu, X. Sha, S. Chang, Abnormal ECG detection based on an adversarial autoencoder (2022) 961724. doi:10.3389/fphys.2022.961724.

[36] G. James, D. Witten, T. Hastie, R. Tibshirani, An Introduction to Statistical Learning: with Applications in R, Springer Texts in Statistics, Springer US, 2021. doi:10.1007/978-1-0716-1418-1.

[37] J. D. Kelleher, B. MacNamee, A. D'Arcy, Fundamentals of Machine Leanring for Predictive Data Analytics, MIT Press US, 2015.

[38] S. Yan, H. Shao, Y. Xiao, B. Liu, J. Wan, Hybrid robust convolutional autoencoder for unsupervised anomaly detection of machine tools under noises (2023) 102441. doi:10.1016/j.rcim.2022.102441.

[39] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need (2017). arXiv:1706.03762.

[40] L. Zhang, X. Wang, E. Cooper, N. Evans, J. Yamagishi, Range-Based Equal Error Rate for Spoof Localization, in: Proc. INTERSPEECH 2023, 2023, pp. 3212–3216. doi:10.21437/Interspeech.2023-1214.