Learning Banknote Fitness for Sorting

Jan-Mark Geusebroek

Informatics Institute

University of Amsterdam, The Netherlands

Email: jm.geusebroek@gmail.com

Peter Markus and Peter Balke

De Nederlandsche Bank N.V.

Amsterdam, The Netherlands

Email: P.Markus@dnb.nl, P.Balke@dnb.nl

Abstract—**In this work, a machine learning method is proposed for banknote soiling determination. We apply proven techniques from computer vision to come up with a robust and effective method for automatic sorting of banknotes. The proposed method is evaluated with respect to various invariance classes. The method shows excellent performance on a large validation set of over 8,000 banknotes from the Eurosystem, while being learned on only 300 banknotes per denomination.**

I. Introduction

A typical problem central banks are facing is whether used banknotes are suitable for recirculation, or rather should be shredded and replaced by new ones. Obviously, more frequent recirculation reduces the printing costs and environmental burden. Given the huge amounts of banknotes in circulation for even small countries, determining the fitness of banknotes is not only of importance in cost control, but also poses a serious technical challenge in terms of processing speed and accuracy. Besides this, the dirtiness of a banknote cannot easily be captured in objective rules.

The Dutch central bank, De Nederlandsche Bank (DNB) has shown interest in an efficient use of banknotes since 1953. This interest is not only from a cost point of view, as cleaner notes are also more secure and more attractive. It appears that soiling is the main reason for classifying banknotes unfit for circulation [1]. Other aspects of banknote fitness are stains and limpness, which show a high correlation with the level of soiling. Further fitness criteria, especially the so called mechanical defects are relatively easy to detect. Therefore, we will concentrate on the soiling aspect in this paper.

In general the performance of existing soil detection is poor [2]. The majority of methods to recognize soil on banknotes are based on optical measurements. These measurements may be based on either intensity or intensity contrast or a combination of the two. In 2009 DNB concluded that the main soiling mechanism of euro banknotes is that fingerprint deposits will cumulate and eventually will form a yellow/brownish layer of aged sebum [3], see Fig. 1. As a consequence, soil is predominantly visible in the blue part of the spectrum. In addition it is the (gentle) touch of the human finger causing soil (particularly sebum) adhesion on the elevated parts, the crumble- or fold lines, of the banknote revealing a structural yet inhomogeneous appearance.

The challenge in soil detection is to have all parameters correct, meaning measuring the right contrast or intensity areas of the banknote, for the right (combination of) colors.

Fig. 1.   Images of euro banknotes at four different deterioration stages.

In this paper, we consider a learning based method on the combination of intensity and contrast differences on color images. The evidence from above outlined soiling studies show that such a method could in principal work. Yet up till now, the algorithms used in automatic sorting could not meet the requested requirements to separate well unfit from fit [2].

This paper is organized as follows. First, we describe the

image processing steps taken to extract the banknote data (Section II). In Section III, we describe the banknote features, feature selection and classification. Results of the approach are detailed in Section IV. We wrap up with the conclusions (Section V).

II. EXTRACTING THE BANKNOTE IMAGE DATA

In order to apply machine learning algorithms, the banknotes have to be segmented and aligned such that image based machine learning can be applied. We consider a two stage process. In the first stage, we segment the banknote and extract the banknote image from the conveyor belt image. Therefore, we detect the sides of the banknote paper, and de-skew the banknote to an image of fixed size and fixed banknote orientation. In a second stage, the alignment of banknotes to the background offset layers of the printing process is considered. Both processes are evaluated for their impact on classification performance.

A. Segmenting the banknote region

Banknotes have to be accurately segmented and aligned from the raw sensor data before we can reliably and reproducibly extract features from the banknote area. The banknote paper area in the input image is obtained by intensity thresholding above noise level. To reduce noise, the red, green and blue channels are added together to form the intensity image $I = R+G+B$. In this way, signal-to-noise ratio is optimized, and a threshold yields most of the paper region. Remaining single pixel noise, both in foreground and background, is removed by a binary 3x3 opening followed by a 3x3 closing operation. This proved a reliable method when applied to the training set of 300 Euro 10 banknotes. The binary image resulting from this step is referred to as the banknote mask in the sequel. As an alternative method, one could automatically determine the threshold level by clustering the input values, using a 2-means clustering approach. Even then, removal of binary single pixel noise is required.

In a next step, the banknote region needs to be de-skewed, allowing consistent feature extraction over similar regions of the banknote. As such, a box needs to be fitted around the banknote, from which the skew parameters can be estimated and a linear transformation can map the pixels to a rectangular and fixed sized banknote image. This will be the image used in consecutive machine learning and classification. We aim at line fitting over the four sides of the banknote paper. This is approached by constructing the contour chain of the banknote mask image using a simple contour following algorithm. The algorithm starts at the first object pixel detected when scanning the binary mask image from top-left to bottom-right. After that, it keeps tracing the contour by counterclockwise searching for a neighbor object pixel. When found, it adds that pixel-location to the contour chain, and subsequently looks for the next contour pixel at its new location.

Having the contour pixel locations of the banknote, finding its four sides boils down to line fitting. A RANSAC based method [4, Chapter 15] is applied, which initially "guesses"

Fig. 2. The camera image with segment lines fitted through the contour (top) and the de-skewed banknote area (bottom).

a line through the contour data, counts how many inliers there are, and re-calculates the exact line parameters by least square line fitting through the inlier points. Four lines with highest inlier count are accepted, as they represent the edges of the banknote. The number of trials –i.e. how often a line is "guessed" as initialization for RANSAC– depends on the "guessing" strategy. One could apply brute force by trying all combinations of two points a certain distance apart on the contour. Alternatively, one could randomly sample a few iterations with points being at least half the width or length of the banknote apart. Furthermore, one could sample all combinations of points where the contour curvature, measured over a certain distance along the contour, is above some threshold. As tuning this last option is not completely trivial under the presence of dogears, we applied a brute force strategy with a large sample distance. This limits the amount of trials, while being perfectly reproducible. The outcome of this procedure are the parameters of the four lines which together confine the banknote paper region. The area is de-skewed into a fixed sized banknote image, the result shown in Fig. 2.

B. Aligning banknote offset prints

The method described above aligns the banknote paper-wise; that is, the banknote paper is cut out and de-skewed to a rectangular image of fixed size for all banknotes. However, a significant tolerance is allowed in the exact positioning of the banknote offset prints relative to the paper boundary. Hence, to reduce the variation between regions introduced by the allowed tolerances in the printing process, our aim here is to make a more precise alignment of the banknote offset prints. From a selection of superfit banknotes, the one banknote inducing the least amount of variation when "overlaid" on the other superfits is determined. Therefore, the banknote with minimum summed absolute difference between the pixels colors of all other banknotes is taken, using the paper-wise de-skewed

images as input. Here, color difference between two pixels is considered to be the sum of the absolute differences between the three respective color channels. The resulting image yields the typical (or modal) positioning of the banknote offset layers within a set of new (or almost new) banknotes, considered the reference for alignment of all other banknotes.

Alignment of a given banknote image now proceeds as follows. The banknote is de-skewed as outlined previously. After that, the image data is shifted in the x- and y-direction within an $n \times n$ neighborhood, and matched against the reference image for all possible shifts within the neighborhood. The shift with minimum summed color difference to the reference image yields the best alignment between the current banknote and the reference. Feature extraction and machine classification is than performed on these aligned images. In this way, banknotes are aligned to the major content of the printing layers, rather than to the banknote paper area.

III. Machine learning of banknote fitness

Machine categorization of banknotes being fit or unfit is certainly not trivial. First of all, dirtiness of banknotes is not something easily captured in rules. Current technology seem to establish the amount of reflection of a white part of the banknote, for example the watermark region. New banknotes tend to reflect more light in such a region, as they are more white than dirty banknotes. However, paper and printing quality varies, over time and per manufacturer, causing already large intensity differences over new, unused banknotes. Given the inherently complex manufacturing process of banknotes as a major defence to counterfeiting, further constrains on the printing specifications are hard to implement. Rather than explicitly designing rules to distinguish fit from unfit notes, a machine learning approach is attractive, using a limited set of example fits and unfits for every type of banknote that should be categorized.

Another issue is the huge bulk of banknotes returned to central banks everyday, putting harsh demands on processing speed. Computer vision has dealt with such processing issues and some techniques for fast visual detection or categorization are available nowadays. We combine the idea of measuring reflectance, contrast, and color content of regions. We design our algorithm to profit from the available knowledge on sebum deposit on banknotes as the major source of soiling by explicitly encoding yellow-blue information. Furthermore, machine learning is applied to automatically select relevant regions for classification. As such, our design yields a fast and accurate banknote categorization method.

A. Image features

Given that certain regions of the image will be better suited to detect dirt than other regions, the aim is to automatically determine these regions and their contribution to fitness categorization. Based on earlier work [5], each banknote image is subdivided into a pre-defined set of overlapping rectangles of various sizes and aspect ratios, thereby densely sampling the image. For each rectangle, the average intensity is measured,

together with the standard deviation of intensity - the latter yielding a measure of local contrast. Hence, per region, two features are calculated. Note that these numbers can be calculated very efficient using integral images, as put forward in [6]. To counteract variations in overall illumination, due to accumulated dust on the image sensor during operations and variations in overall printing quality, both these numbers are normalized by the average intensity of the whole banknote region.

As we do not only have access to intensity information, but also color information, the same procedure is repeated for the R, G, and B color channel, respectively, and for their opponent color combinations: $YB = R + G - 2B$, $RG = R - 2G + B$. So, we are left with two numbers (ave, stdev) per rectangle for six color channels (including intensity), yielding a total of 12 features per rectangle. Furthermore, the features of the front and back side of the banknote are combined, resulting in a large set of features per banknote for feature selection and classification. Obviously, after training a feature selector, only a small part of all these features need to be calculated for the classification into fit or unfit.

As an alternative to exploiting only grey or full color images, we investigate the performance of the blue-red ratio as an alternative feature. In this case, only the ratio of the blue channel over the red channel is used as input image. The average ratio and its standard deviation over all regions is calculated, similar to the grey-value method described above. As sedum deposit is mainly apparent in the blue channel, with normalization by the red channel, we expect this to improve classification compared to intensity information alone, with only a slight increase of computing power. Exploiting only the blue and red channel allows for much smaller bandwidth between image capture device and memory, simplifying the overall design of fitness sensors.

B. Feature selection and classification

A well known and proven method for feature selection is from the area of face detection [6]. We adopt this approach and classical AdaBoost [7] is applied to select the N most discriminating features (and hence banknote regions) for classification. AdaBoost with simple linear weak classifiers is used, the weak classifiers being combined into a strong linear classifier. After training the classifier, a set of regions is obtained, each being taken from a single color channel, with a single feature (average or stdev) and its threshold most relevant to distinguish fit from unfit banknotes. See Fig. 3 for an illustration of the result. Based on training set classification results for 300 euro 10 banknotes, we have set the parameter $N = 40$. In this way, the training set could be well classified, without obvious over-fitting to the training data.

IV. Results

The dataset used consists of euro 5, 10, 20, and 50 banknotes, scanned in color, two-sided. The banknotes are extracted and de-skewed to a standard image size of 332×635 pixels. Table I shows the number of images available for

Fig. 3. The by AdaBoost selected regions which discriminate best between the fit and unfit euro 10 banknotes.

TABLE I
DATA SET USED IN THE EXPERIMENTS.

Denomination	Train		Validation	
	Fit	Unfit	Fit	Unfit
Euro 5	150	150	1348	950
Euro 10	150	150	1456	1000
Euro 20	150	150	450	633
Euro 50	150	150	1397	595

training and validation of the proposed method. Randomly 150 fits and 150 unfit banknotes have been selected from the available data, which have been used for training the classifier. The remaining images are used as test set to validate the approach. For all 300 training images per denomination, feature values are determined per region, and AdaBoosting is applied to the features. The 40 features most discriminant between the fit and unfit training samples, with their polarities and thresholds, and the weak classifier relative weight, together constitute the strong classifier.

Where all features had to be calculated over the whole training images, for classification only the features relevant are calculated over the areas selected by AdaBoost. These features are fed into the now trained classifier and output confidence is ranked in order to assess classifier performance. Therefore, we assess the classifier Receiver Operating Characteristic (ROC). The resulting ROC graphs are given in Fig. 4. The following relevant cases are considered: *color* - the classifier has all six available (opponent) color channels at its disposition for training and classification, but normalized for overall intensity; *grey* - the classifier had only intensity information available, again normalized for overall intensity; *B/R-ratio* - the classifier had the ratio between the blue and red channel available, which is by construction normalized for intensity. Here, the

weakly aligned case is considered, where the banknote paper is used to de-skew the images to a standard image rectangle. Furthermore, we consider the *aligned* cases, where the extra processing step is included to align the offset printing layers, the variations indicated in the ROC graphs below. Although the tested algorithm variations base there classification on input data with varying richness in information content - including three or two color channels or only assessing intensity information, the AdaBoost algorithm does in all cases select 40 regions or features for classification. Hence, the dimensionality for each method is identical, allowing fair comparison of the classification result.

When assessing the graphs in Fig. 4, most striking is the performance of *color* and *color-aligned* for euro 20 notes. The performance is characterized by the selection of a single feature by AdaBoost. The method was able to discriminate perfectly between the fit and unfit notes in the training set based on this one region. Hence, there is only one linear classifier, with a single threshold, and so only a single rectangular area in the ROC curve for the validation set. Apparently, separating the euro 20 notes in our collection turns out to be a simple task when having color information - quite different from the same task for intensity images only. However, note that the number of available euro 20 examples of fit notes is significantly less than for other denominations.

Regarding the euro 5 and euro 10 notes, color information improves considerably over intensity information. Performance of *RB-ratio* improves over intensity, and for Euro 5 even reaches similar performance as *color-aligned*. The effect of aligning the banknote offset printing layers remains disputed - although improvement seems theoretically obvious, in practice the proposed alignment method depends on the quality of the banknotes. More crumpled paper introduces larger errors in the alignment method, which is reflected in the results. The overall quality, including the quality of what is still being considered fit or not, is lower for euro 5 banknotes than for euro 50 banknotes. Hence, the effect for alignment is poor for euro 5, while improving over the denominations up to euro 50. Euro 50 banknotes are simply cleaner and of higher quality to start with, as the lower quality banknotes already have been taken out at an earlier stage. Given the computational effort involved in alignment, one might rather concentrate on the weakly (banknote paper) aligned methods.

Besides ROC curves, the classification rate when assigning fit/unfit labels based on the probability indicated by the classifier (fit if $p > 0.5$) have also been assessed, see Table II below. Just by straightforwardly applying the learned strong classifier, the method already reaches a false fit performance close to acceptable for many central banks, while providing an excellent false unfit rate. When accepting a 5% of false fits into circulation, the results of false unfits indicated by the ROC analysis are shown in Table III.

When considering which features have been selected by AdaBoost for the various denominations, indeed the red-blue opponency plays an important role. Many of the selected features consider threshold on the blue presence in the yellow-
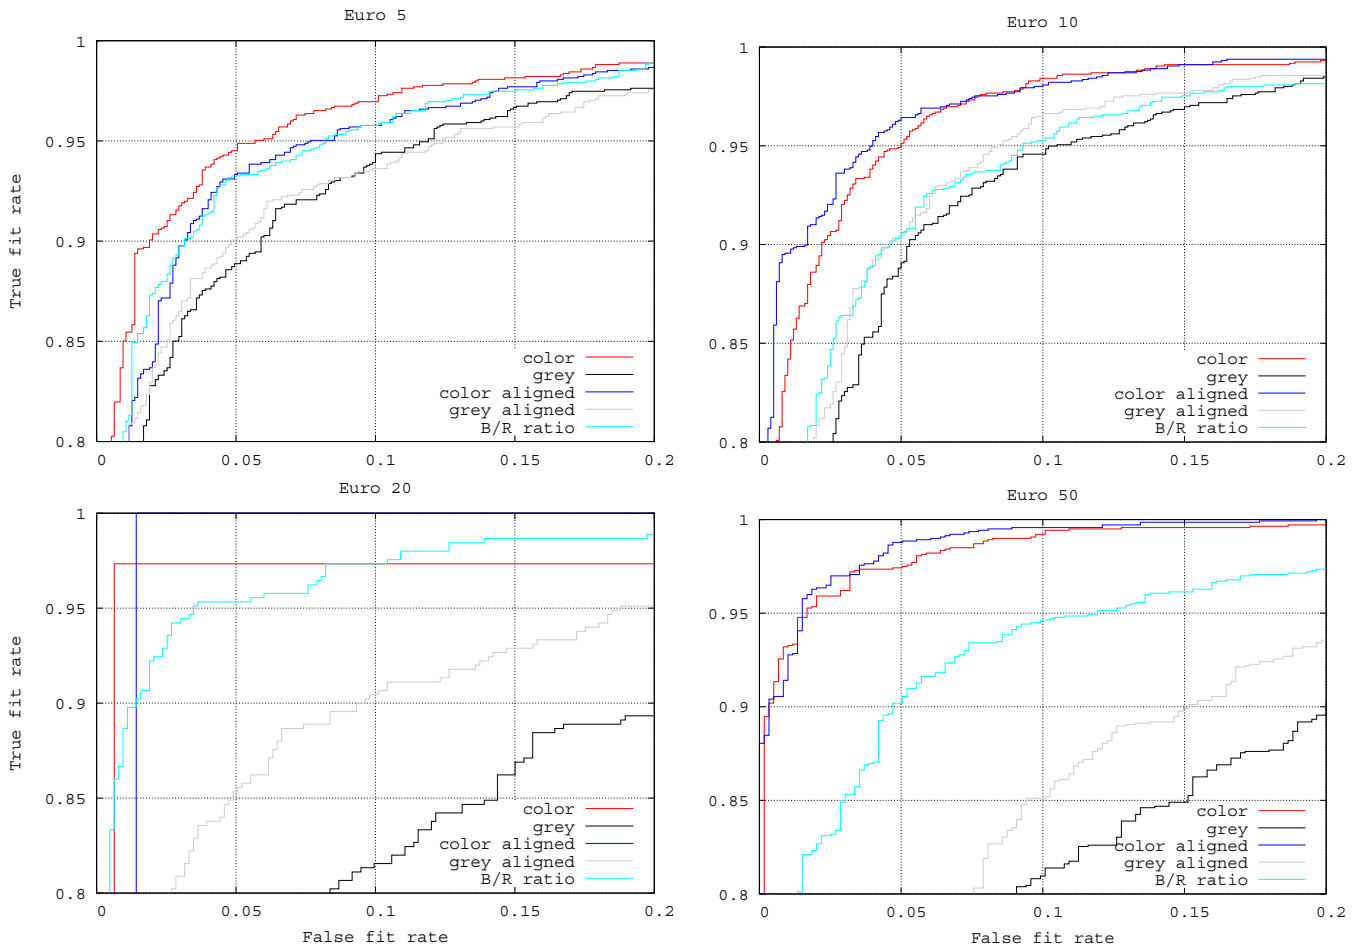
Fig. 4.   ROC curves for 5, 10, 20, and 50 euro banknotes, respectively.

TABLE II
RESULTS WHEN APPLYING THE LEARNED STRONG CLASSIFIER.

Denom-ination	Color		Grey		RB-ratio	
	False fit	False unfit	False fit	False unfit	False fit	False unfit
Euro 5	5.68%	5.12%	8.32%	7.12%	6.11%	6.38%
Euro 10	4.90%	5.08%	9.10%	6.04%	7.10%	6.73%
Euro 20	0.63%	2.67%	12.95%	15.78%	4.90%	4.67%
Euro 50	3.19%	3.58%	13.95%	15.32%	6.55%	7.66%

TABLE III
FALSE UNFIT RATE AT 5% ACCEPTANCE OF UNFITS IN CIRCULATION.

	Color	Grey	RB-ratio
Euro 5	5.49%	11.13%	6.75%
Euro 10	5.08%	11.20%	9.62%
Euro 20	2.67%	29.11%	4.67%
Euro 50	2.58%	27.34%	9.81%

blue opponent features, or compare the red-blue content to green. The normalized intensity is of some importance, and few of these features are selected. Furthermore, it is interesting to see that regions all over the banknote area are selected by AdaBoost, rather than only the regions with dominantly white paper being visible.

V. CONCLUSIONS

In this paper, a machine learning approach to banknote soiling fitness is demonstrated. The approach is based on simple integral features over the color channels and opponent color combinations over a plurality of banknote regions. Feature selection is performed by AdaBoost, which requires a small training set of fit and unfit banknotes per denomination. Performance of the proposed method is demonstrated on a large set of approximately 8,000 banknotes from the Eurosystem. The use of color is demonstrated to significantly outperform grey value features, especially for the euro 20 and euro 50 denominations. More precisely aligning the banknote printing layers did not significantly improve performance.

The proposed machine learning method for banknote soiling exceeded the expectations. It showed that the use of color information increases significantly the performance with respect to intensity information only. Where in general soil detectors focus on the watermark region, our results demonstrate the use of the entire banknote surface for analysis has added

value. A combination of mean and standard deviation features normalized to intensity using color information showed the most promising results for all denominations. This method can be applied to raw images, meaning the de-skewing processing step can even be skipped. This will contribute in a positive sense to the processing speed. We expect that a processing speed of 2,000 banknotes per minute will be viable.

ACKNOWLEDGMENT

REFERENCES

[1] H. de Heij, "Durable banknotes: an overview," in *BPC General Meeting*, 2002.

[2] T. Buitelaar, "The colour of soil," in *DNB Cash Seminar*, 2008.

[3] P. Balke, "New soiling test method: Anti-dirty fingers," in *Banknote*, 2009.

[4] D. A. Forsyth and J. Ponce, *Computer Vision: A Modern Approach*, ser. Prentice Hall series in artificial intelligence. New Jersey: Prentice Hall, 2003.

[5] J. C. van Gemert, C. J. Veenman, A. W. M. Smeulders, and J. M. Geusebroek, "Visual word ambiguity," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 7, pp. 1271–1283, 2010.

[6] P. A. Viola and M. J. Jones, "Robust real-time face detection," *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.

[7] Y. Freund and R. E. Shapire, "A decision-theoretic generalization of on-line learning and an application to boosting," in *Computational Learning Theory: Eurocolt 95*, 1995, pp. 23–37.